AssuredSAN Pro 5000 Series
CLI Reference Guide

Adobe PostScript™

Contents

4   Contents

Tables

About this guide

This guide provides information about managing an AssuredSAN™ Pro 5000 Series storage system by using its command-line interface (CLI).

Intended audience

This guide is intended for storage system administrators.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel, Serial Attached SCSI (SAS), and Ethernet protocols

Related documentation

For information about	See
Enhancements, known issues, and late-breaking information not included in product documentation	Release Notes
Overview of product shipkit contents and setup tasks	Getting Started*
Regulatory compliance and safety and disposal information	AssuredSAN Product Regulatory Compliance and Safety*
Installing and using optional host-based software components (CAPI Proxy, VDS Provider, VSS Provider, SES Driver)	AssuredSAN Installing Optional Software for Microsoft Windows® Server
Using a rackmount bracket kit to install an enclosure into a rack	AssuredSAN Rackmount Bracket Kit Installation* *or* AssuredSAN 2-Post Rackmount Bracket Kit Installation*
Product hardware setup and related troubleshooting	AssuredSAN Pro 5000 Series Setup Guide
Obtaining and installing a license to use licensed features	AssuredSAN 5000 Series Obtaining and Installing a License
Using the web interface to configure and manage the product	AssuredSAN Pro 5000 Series Storage Management Guide
Using the command-line interface (CLI) to configure and manage the product	AssuredSAN Pro 5000 Series CLI Reference Guide
Event codes and recommended actions	AssuredSAN Pro 5000 Series Event Descriptions Reference Guide
Identifying and installing or replacing field-replaceable units (FRUs)	AssuredSAN Pro 5000 Series FRU Installation and Replacement Guide

* Printed document included in product shipkit.

For additional information, see Dot Hill's Customer Resource Center web site: http://crc.dothill.com.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text	Cross-reference links
Blue, underlined text	Email addresses
Blue, underlined text	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
Italic text	Text emphasis
Monospace text	• File and directory names • System output • Code • Commands, their arguments, and argument values
Monospace, italic text	• Code variables • Command variables
Monospace, bold text	Emphasized of file and directory names, system output, code, and text typed at the command line

△ **CAUTION:**   Indicates that failure to follow directions could result in damage to equipment or data.

📝 **IMPORTANT:**   Provides clarifying information or specific instructions.

📝 **NOTE:**   Provides additional information.

☼ **TIP:**   Provides helpful hints and shortcuts.

1 Using the CLI

This chapter introduces the command-line interface (CLI).

Accessing the CLI

The CLI software embedded in the controller modules enables you to manage and monitor storage-system operation. You can access the CLI in two ways:

- By using HTTP, HTTPS, Telnet, an SSH application, or a terminal emulator on a management host that is remotely connected through a LAN to a controller module's network port. See your product's Setup Guide for information about setting management port IP addresses using the CLI.
- By using a terminal emulator on a management host that is directly connected to a controller module's serial CLI port.

NOTE: Using HyperTerminal on a Microsoft Windows host with the CLI:

- On a host connected to a controller module's mini-USB CLI port, incorrect command syntax in a HyperTerminal session can cause the CLI to hang. To avoid this problem, use correct syntax, use a different terminal emulator, or connect to the CLI using telnet rather than the mini-USB cable.
- Close the HyperTerminal session before shutting down the controller or restarting its Management Controller; otherwise the host's CPU cycles may rise unacceptably.

Table 2 Default usernames and passwords

Username	Password	Roles
monitor	!monitor	Monitor (view only)
manage	!manage	Monitor, Manage (view and change)

CLI output formats

The CLI has two output formats:

- Console format, which is the human-to-computer interface (HCI).
- XML API format, which is the computer-to-computer interface (CCI).

Console format enables users to interact with the CLI and obtain easily readable information. This format automatically sizes fields according to content and adjusts content to window resizes. These capabilities would present problems for a CCI in the form of scripts or other client software.

XML API format enables any external application to interact with the storage system. The XML format is constructed to permit new fields to be added without impacting existing clients if they follow standard XML parsing conventions. The XML format also contains alternate fields for values which may be localized, such as a state value. These values are guaranteed not to change with different languages and can be used for scripting.

Scripting is not supported using console format because labels, field sizes, and order of fields may change in future firmware releases. To properly script CLI commands use XML API format, which is expected to remain consistent from release to release; field names will be consistent and new functionality will be added as new fields. These types of changes in XML output will not impact a conventional XML parsing engine.

You can change the CLI output format by using the set cli-parameters command; see Example on page 129.

Using CLI interactively

By default the CLI is an interactive application. When you are logged into the CLI, the CLI waits for a command to be entered and then responds to it.

The following example shows interactively starting a Telnet session, logging into the CLI, executing a command to show information about the system, and exiting the CLI:

```
$: telnet 172.22.5.55
172.22.5.55 login: monitor
Password:

product
System Name: Test
System Location: Lab
Version: version
# show system
System Information
------------------
System Name: Test
System Contact: J.Smith
...
Success: Command completed successfully. (2012-06-15 14:42:57)

# exit
```

Using a script to access the CLI

The following example shows how to construct a Perl script to communicate with the CLI via Telnet. cLogin is called at the start of the script to log a user into the CLI. The script uses the command-line parameters specified as the IP address, username, and password. After the user has been logged in, other commands can be sent to the CLI.

```
use Net::Telnet;
sub cLogin {
    $telnet->open($_[0]);
    $telnet->waitfor(/(login|username)[: ]*$/i);
    $telnet->print("$_[1]");
    $telnet->waitfor(/password[: ]*$/i);
    $telnet->print("$_[2]");
    # either got a login or a prompt
    @ok = $telnet->waitfor(/(#|login:*) /i);
    if ($debug_commands == 1) { print "-"; print @ok; print "-\n"; }
    if ($ok[1] =~ m/login/gi)
    {
        return 0;
    }
    else
    {
        return 1;
    }
}
$ipAddr = $ARGV[0];
$username = $ARGV[1];
$password = $ARGV[2];
$telnet = new Net::Telnet ( Timeout=>10,
Errmode=>'die',
Prompt => '/\# $/i');
if ( !cLogin($ipAddr, $username, $password) == 1 )
{
    print("Error: $username user failed to log in. Exiting.\n");
```

```
            $telnet->close;
            exit(0);
        }
```

For scripting support, the following command can be used to access the XML API and to disable the paging mode that pauses for each full screen of command output.

```
        $telnet->cmd("set cli-parameters api-embed pager disabled");
```

The following code segment shows how to get the entire configuration information from the CLI and print the output. The output can easily be redirected to a file for archiving.

```
        @sV = $telnet->cmd("show configuration");
        for ($i=0; $i<scalar(@sV); $i++)
        {
            print ("@sV[ $i ]");
        }
```

Because basic command-line semantics provide prompts for user input and response time is indeterminate, scripts must use an "expect"-type mechanism to scan output for prompts. It is more efficient to use the HTTP interface to access the XML API. The following example shows how to construct a Perl script to communicate with the XML API via HTTP.

```
        use LWP::UserAgent;
        use Digest::MD5 qw(md5_hex);
        use XML::LibXML;

        # generate MD5 hash using default username/password
        my $md5_data = "manage_!manage";
        my $md5_hash = md5_hex( $md5_data );
        print "$md5_hash\n";

        # create the URL and send an http GET request

        $ua = LWP::UserAgent->new;
        $url = 'http://10.0.0.2/api/login/' . $md5_hash;

        print ("Sending to $url\n");
        $req = HTTP::Request->new(GET => $url);

        $res = $ua->request($req);

        # Parse the XML content using LibXML to obtain the session key

        print $res->content;

        my $parser = XML::LibXML->new();
        my $doc = $parser->parse_string( $res->content );

        my $root = $doc->getDocumentElement;
        my @objects = $root->getElementsByTagName('OBJECT');
        my @props = $objects[0]->getElementsByTagName('PROPERTY');

        my $sessionKey;

        foreach my $prop ( @props ) {
            my $name = $prop->getAttribute('name');

            print "Property = " . $name . "\n";

            if( $name eq 'response' ) {
                $sessionKey = $prop->textContent;
            }
```

```
    }

    print "Session Key = $sessionKey\n";

    # Run a sample command to obtain the disks in the system.

    $url = 'http://10.0.0.2/api/show/disks';
    $req = HTTP::Request->new(GET => $url);
    $req->header('sessionKey' => $sessionKey );
    $req->header('dataType' => 'ipa' );

    $res = $ua->request($req);

    $url2 = 'http://10.0.0.2/api/exit';
    $req2 = HTTP::Request->new(GET => $url3);
    $req2->header('sessionKey' => $sessionKey );
    $req2->header('dataType' => 'api' );

    $res2 = $ua->request($req2);

    print $res->content;
```
The next section provides more information about using the XML API.

Using the XML API

The Management Controller provides access for monitoring and management via the Telnet and SSH protocols for command-line interface semantics, or via the HTTP and HTTPS protocols for XML API request/response semantics.

You can use an XML parser, such as `XML::Parser` in Perl, to process the XML output and store this information as objects. The XML parser should use the Document Type Definition (DTD) version that corresponds to the firmware level to ensure that the XML is validated. The DTD provides the structure of all content returned by the CLI when XML API format is enabled. By obtaining the latest DTD for validation, the parser will be forward compatible. To obtain the DTD, go to crc.dothill.com.

The output of each CLI command is composed of valid XML data until the CLI prompt (typically #) is encountered. The output contains a valid XML header followed by the XML elements described in the following table.

Table 3 XML API elements

Element	Description and attributes
RESPONSE	The RESPONSE element is the top-level element, which contains all data output for the CLI command that was issued. The response includes: • A number of OBJECT elements, which varies by command. • A status object that provides a message and return code. A return code of 0 indicates that the command succeeded. Any other return code is an error code. There is only one RESPONSE element per issued command.
OBJECT	In general, an OBJECT element describes a storage-system component such as a disk or a volume. An object has these attributes: • `basetype`. This attribute allows output in brief mode to be correlated with metadata to reduce the overhead of each command, as described in XML API optimization. This is also a good field to use to detect the type of the object (e.g., a disk, a volume, etc.). • `name`. The name of the object. • `oid`. The unique identifier for the object in the scope of the response. The OBJECT element can contain PROPERTY elements.

Table 3 XML API elements (continued)

Element	Description and attributes
PROPERTY	A PROPERTY element provides detail about the attributes of an OBJECT. A property has these attributes: • name. The unique name for the property within the object. • key. Indicates whether this property is a key value to identify this object. • type. The type of data represented by the element data. • size. Typically the maximum size of the output. Usually only important if the console output is displayed in rows. • draw. Whether to show or hide this data in console format. • sort. The type of sorting that can be applied to this property. • display-name. The label for this data to show in user interfaces.
COMP	A COMP (composition) element associates nested objects, such as a task object within a schedule object. A composition element has these attributes: • P. The oid of the part component. • G. The oid of the group component. An alternative to using COMP elements is described in XML API optimization.
ASC	The association element provides a simple association description between two objects in the response. • A. First object. • B. Second object.

Scripting guidelines

When scripting command input, use CLI syntax as defined in this guide. For use with Telnet or SSH, use a space character between command names, parameters, and their values (as shown throughout this guide). For use with the HTTP interface, use a '/' character instead of a space character between command names, parameters, and their values.

When writing scripts to parse XML API output, use an XML library to parse the data. For parsing, a script should *not* rely on ordering, spacing, or column position. To find a specific property, a script should compare property names as it searches through the data. This allows the script to be compatible with future versions that could potentially add new fields to the output.

The output of show commands is intended for monitoring or obtaining the current configuration. Other commands provide configuration data and display one or more status objects that specify the status of command processing. The last status object specifies the overall status of the command; other status objects indicate intermediate processing status.

The following example shows the XML API status object:

```
<OBJECT basetype="status" name="status" oid="1">
    <PROPERTY name="response-type" type="string" size="12" draw="false"
sort="nosort" display-name="Response Type">Success</PROPERTY>
    <PROPERTY name="response-type-numeric" type="uint32" size="12" draw="false"
sort="nosort" display-name="Response Type">0</PROPERTY>
    <PROPERTY name="response" type="string" size="180" draw="true" sort="nosort"
display-name="Response">Command completed successfully. (2012-11-13
14:17:52)</PROPERTY>
    <PROPERTY name="return-code" type="sint32" size="15" draw="false"
sort="nosort" display-name="Return Code">0</PROPERTY>
    <PROPERTY name="component-id" type="string" size="80" draw="false"
sort="nosort" display-name="Component ID"></PROPERTY>
    <PROPERTY name="time-stamp" type="string" size="25" draw="false"
sort="datetime" display-name="Time">2012-11-13 14:17:52</PROPERTY>
    <PROPERTY name="time-stamp-numeric" type="uint32" size="25" draw="false"
sort="datetime" display-name="Time">1352816272</PROPERTY>
</OBJECT>
```

In a script, each command should check the previous command's status before proceeding. If the value of the `status` object's `return-code` property is `0`, the command succeeded; any other value means that the command failed.

XML API examples

The following example shows a command, formatted for use with the command-line interface or for use with the HTTP interface, and its XML API output.

- Command-line interface format: `set system name Test2 contact J.Smith`
- HTTP interface format: `set/system/name/Test2/contact/J.Smith`

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<RESPONSE VERSION="L100">
  <OBJECT basetype="status" name="status" oid="1">
    <PROPERTY name="response-type" type="string" size="12" draw="false"
sort="nosort" display-name="Response Type">Success</PROPERTY>
    <PROPERTY name="response-type-numeric" type="uint32" size="12" draw="false"
sort="nosort" display-name="Response Type">0</PROPERTY>
    <PROPERTY name="response" type="string" size="180" draw="true" sort="nosort"
display-name="Response">Command completed successfully. (2012-11-13
14:21:59)</PROPERTY>
    <PROPERTY name="return-code" type="sint32" size="15" draw="false"
sort="nosort" display-name="Return Code">0</PROPERTY>
    <PROPERTY name="component-id" type="string" size="80" draw="false"
sort="nosort" display-name="Component ID"></PROPERTY>
    <PROPERTY name="time-stamp" type="string" size="25" draw="false"
sort="datetime" display-name="Time">2012-11-13 14:21:59</PROPERTY>
    <PROPERTY name="time-stamp-numeric" type="uint32" size="25" draw="false"
sort="datetime" display-name="Time">1352816519</PROPERTY>
  </OBJECT>
</RESPONSE>
```

XML API optimization

The following are two ways to optimize XML API performance:

- Use embedded objects. This allows one object to contain not only properties but also other objects. In general, parsing a structure such as this is easier as the association between objects is simpler. This is an alternative to using COMP elements.
- Use brief mode. In brief mode, which is disabled by default, returns a subset of attributes of object properties. The name and type attributes are always returned. Other properties can be obtained by using the meta command with the basetype of the object. This optimization reduces the number of bytes transmitted for each request and allows caching of CLI metadata. Brief mode can be enabled or disabled by using the set cli-parameters command.

In the following example, which uses brief mode, embedded objects contain media-specific detail for ports:

```
# show ports
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<RESPONSE VERSION="L100">
  <OBJECT basetype="port" name="ports" oid="1" format="rows">
    <PROPERTY name="durable-id" type="string">hostport_A0</PROPERTY>
    <PROPERTY name="controller" key="true" type="string">A</PROPERTY>
    <PROPERTY name="controller-numeric" key="true" type="string">1</PROPERTY>
    <PROPERTY name="port" key="true" type="string">A0</PROPERTY>
    <PROPERTY name="port-type" type="string">FC</PROPERTY>
    <PROPERTY name="port-type-numeric" type="string">6</PROPERTY>
    <PROPERTY name="media" type="string">FC(P)</PROPERTY>
    <PROPERTY name="target-id" type="string">target-ID</PROPERTY>
    <PROPERTY name="status" type="string">Up</PROPERTY>
    <PROPERTY name="status-numeric" type="string">0</PROPERTY>
    <PROPERTY name="actual-speed" type="string">8Gb</PROPERTY>
    <PROPERTY name="actual-speed-numeric" type="string">7</PROPERTY>
```

```
            <PROPERTY name="configured-speed" type="string">Auto</PROPERTY>
            <PROPERTY name="configured-speed-numeric" type="string">3</PROPERTY>
            <PROPERTY name="health" type="string">OK</PROPERTY>
            <PROPERTY name="health-numeric" type="string">0</PROPERTY>
            <PROPERTY name="health-reason" type="string"></PROPERTY>
            <PROPERTY name="health-recommendation" type="string"></PROPERTY>
            <OBJECT basetype="fc-port" name="port-details" oid="2" format="rows">
               <PROPERTY name="configured-topology" type="string">PTP</PROPERTY>
               <PROPERTY name="primary-loop-id" type="string">255</PROPERTY>
            </OBJECT>
         </OBJECT>
         <OBJECT basetype="port" name="name" oid="3" format="rows">
            <PROPERTY name="durable-id">hostport_A1</PROPERTY>
            ...
         </OBJECT>
         <OBJECT basetype="status" name="status" oid="17">
            <PROPERTY name="response-type" type="string">Success</PROPERTY>
            <PROPERTY name="response-type-numeric" type="string">0</PROPERTY>
            <PROPERTY name="response" type="string">Command completed successfully.
      (2012-11-13 14:42:43)</PROPERTY>
            <PROPERTY name="return-code" type="sint32">0</PROPERTY>
            <PROPERTY name="component-id" type="string"></PROPERTY>
            <PROPERTY name="time-stamp" type="string">2012-11-13 14:42:43</PROPERTY>
            <PROPERTY name="time-stamp-numeric" type="string">1352817763</PROPERTY>
         </OBJECT>
      </RESPONSE>
```

Command syntax

General rules for specifying commands

Command and parameter names are not case sensitive.

Parameters enclosed in square brackets ([]) are optional. Do not type the bracket characters.

Parameter values separated by '|' characters are options. Enter only one of the values. These values are not case sensitive.

Parameter values in italics are variables. Substitute text that is appropriate for the task you want to perform. Unless specified otherwise, variable values are case sensitive and have a maximum length in bytes. ASCII characters are 1 byte; most Latin (Western European) characters with diacritics are 2 bytes; most Asian characters are 3 bytes.

Unless otherwise specified, a parameter value can include spaces and printable UTF-8 characters except angle brackets (<>), backslash (\), comma (,), and double quote (").

A parameter value that includes a space must be enclosed in double quotes.

Parameters can be entered in any order. However, if the value of a parameter with no keyword is the same as the keyword of an optional parameter, the optional parameter must precede the value.

Specifying disks

Disks are specified by enclosure ID and slot number. Enclosure IDs increment from 0. Disk IDs increment from 0 in each enclosure. You can specify:

- A disk. Example: `0.4`
- A hyphenated range of disks. Example: `0.4-7`
- A comma-separated list of individual disks, ranges, or both (with no spaces). Example: `0.4,0.6-9`

Specifying storage pools and storage-pool components

You can specify:

- A storage pool by its name or serial number.
- A storage-pool component by its name or serial number.
- A list of storage-pool-component names or serial numbers separated by commas (with no spaces). Not all commands support lists. Example: `A_01_01,A_02_01`

Specifying volumes and volume groups

You can specify:

- A volume by its name or serial number.
- A list of volume names or serial numbers separated by commas (with no spaces). Not all commands support lists. Example: `Vol0003,"Vol #1"`
- A volume group by its name in the format *volume-group-name.**, where * represents all volumes in the group. Example: `TestVolumes.*`

Specifying ports

Controller host ports are specified by port number only (to use the same port in each controller, which is recommended) or by controller ID and port number (to specify an individual port). Controller IDs are A for the upper controller and B for the lower controller, and are not case sensitive. Port IDs increment from 0 in each controller module. You can specify:

- A port ID in one controller. Example: `a1`
- A port ID in both controllers. Example: `1`
- A hyphenated range of IDs in both controllers or in the same controller. Do not mix controller IDs in a range. Example: `0-2` or `a0-a2`
- A comma-separated list of IDs, ranges, or both (with no spaces). Example: `a0,1-2`

Specifying initiators, hosts, and host groups

You can specify:

- An FC or SAS initiator by its nickname or 16-hex-digit WWPN.
- A host by name in the format *host-name.**, where * represents all initiators in the host. Example: `FC-Server.*`
- A host group by name in the format *host-group-name.*.**, where the first * represents all hosts in the group and the second * represents all initiators in those hosts. Example: `TestLab.*.*`

Command completion, editing, and history

When entering commands interactively you can abbreviate their names and keywords. For example, you can enter **sho cl** to run the `show cli-parameters` command. If you press **Tab** or **Ctrl+i** after typing sufficient characters to uniquely identify the command or keyword, the remainder of the command or keyword is displayed so you can confirm your intent. If you enter too few letters to uniquely identify a keyword, pressing **Tab** or **Ctrl+i** will list commands or keywords that match the entered string and redisplays the string so you can complete it.

When scripting commands, type commands in full to aid readability.

The history contains commands entered in the active CLI session. You can recall a command from the history, edit it, and run it.

Table 4 Keyboard shortcuts for command completion, editing, and history

To	Press
Complete a partially entered keyword	**Tab** or **Ctrl+i**
Show command history	F6
Get previous command from history	**Up Arrow**
Get next command from history	**Down Arrow**
Move cursor left	**Left Arrow**
Move cursor right	**Right Arrow**
Delete previous character	**Backspace**

Viewing help

To view brief or full help for commands, use the help command.

Size representations

Parameters such as names of users and volumes have a maximum length in bytes. ASCII characters are 1 byte; most Latin (Western European) characters with diacritics are 2 bytes; most Asian characters are 3 bytes.

Operating systems usually show volume size in base 2. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2.

In the CLI, the base for entry and display of storage-space sizes can be set per user or per session; see create user and set cli-parameters. When entering storage-space sizes only, either base-2 or base-10 units can be specified.

Table 5 Size representations in base 2 and base 10

Base 2		Base 10	
Unit	Size in bytes	Unit	Size in bytes
KiB (kibibyte)	1,024	KB (kilobyte)	1,000
MiB (mebibyte)	$1,024^2$	MB (megabyte)	$1,000^2$
GiB (gibibyte)	$1,024^3$	GB (gigabyte)	$1,000^3$
TiB (tebibyte)	$1,024^4$	TB (terabyte)	$1,000^4$
PiB (pebibyte)	$1,024^5$	PB (petabyte)	$1,000^5$
EiB (exbibyte)	$1,024^6$	EB (exabyte)	$1,000^6$

The locale setting determines the character used for the decimal (radix) point, as shown below.

Table 6 Decimal (radix) point character by locale

Language	Character	Examples
English, Chinese, Japanese	Period (.)	146.81 GB 3.0 Gb/s
Spanish	Comma (,)	146,81 GB 3,0 Gb/s

Event log

A controller enclosure's event log records all events that have occurred in or been detected by the controller modules and encompasses all field-replaceable units (FRUs) in the storage system.

Each event has one of the following levels, in decreasing severity:

- Critical. A failure occurred that may cause a controller to shut down. Correct the problem *immediately*.
- Error. A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
- Warning. A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct it if necessary.
- Informational. A configuration or state change occurred, or a problem occurred that the system corrected. No action is required.

For information about viewing events, see the show events command.

2 Categorical list of commands

This chapter helps you find a command within a category of functionally related commands. A command might appear in more than one category.

Table 7 Commands by category

Category	Commands
CLI and users	create user, delete user, exit, help, set cli-parameters, set password, set prompt, set user, show cli-parameters, show users
Storage pools, storage-pool components, disks, and spares	abort scrub, abort verify, add storage, clear disk-metadata, dequarantine, remove storage, rescan, scrub disks, scrub storage-pool-components, set led, set storage-pool, show disks, trust, verify storage-pool-components
Volumes, initiators, hosts, and mapping	add host-group-members, add host-members, create host, create host-group, create volume, create volume-group, create volume-set, delete all-volumes, delete initiator-nickname, delete host-groups, delete hosts, delete volume-groups, delete volumes, expand volume, map volume, release volume, remove host-group-members, remove host-members, remove volume-group-members, set cache-parameters, set host, set host-group, set initiator, set volume, set volume-group, show cache-parameters, show host-groups, show host-maps, show hosts, show initiators, show maps, show unwritable-cache, show volume-maps, show volume-names, show volume-reservations, show volumes, unmap volume
Snap pools and snapshots	create snapshots, delete all-snapshots, delete snapshot, delete snapshot-write-data, reset snap-pool, reset snapshot, rollback volume, set priorities, set snap-pool-parameters, set snapshot-policy, show priorities, show snap-pools, show snapshots
Volume copy	abort volumecopy, copy volume, show volume copy-status
Scheduled tasks	create schedule, create task, delete schedule, delete task, set schedule, set task, show schedules, show tasks
Event notification	set email-parameters, set snmp-parameters, show email-parameters, show events, show snmp-parameters, test
System configuration and utilities	clear cache, load license, ping, reset host-link, restart, set advanced-settings, set auto-write-through-trigger, set controller-date, set disk-parameters, set enclosure, set host-parameters, set led, set network-parameters, set ntp-parameters, set protocols, set system, show advanced-settings, show auto-write-through-trigger, show configuration, show controller-date, show controllers, show disk-parameters, show enclosures, show enclosure-status, show fans, show frus, show inquiry, show license, show network-parameters, show ntp-status, show ports, show power-supplies, show protocols, show sas-link-health, show sensor-status, show system, show system-parameters, show versions, shutdown
Service utilities	clear events, clear expander-status, fail, restore defaults, set debug-log-parameters, set expander-fault-isolation, set expander-phy, show debug-log-parameters, show expander-status
Remote systems	create remote-system, delete remote-system, remote, set remote-system, show remote-systems, verify links, verify remote-link
Remote replication	abort replication, add replication-volume, create replication-set, delete replication-set, export snapshot, remove replication-volume, replicate snapshot, replicate volume, resume replication, set replication-primary-volume, set replication-volume-parameters, show replication-images, show replication-sets, show replication-volumes, suspend replication
Statistics	reset all-statistics, reset controller-statistics, reset disk-error-statistics, reset disk-statistics, reset host-port-statistics, reset storage-statistics, reset volume-statistics, show controller-statistics, show disk-statistics, show host-port-statistics, show volume-statistics

Table 7 Commands by category (continued)

Category	Commands
API specific	meta, show refresh-counters
Diagnostic	abort scrub, abort verify, add spares, clear disk-metadata, delete spares, dequarantine, fail, reset host-link, scrub storage-pool-components, set debug-log-parameters, set snap-pool-parameters, show storage-pool-components, trust, verify storage-pool-components

3 Alphabetical list of commands

This chapter is organized to help you find a command by name. Each command topic includes one or more of the following sections:

Description	The command's purpose and notes about its usage
Level	The minimum user level/role to access the command.
Syntax	The command's syntax
Parameters	Descriptions of the command's parameters
Output	Descriptions of fields shown in console format
Example	One or more examples of the command's usage in console format
See also	References to commands that are used with the command

abort replication

Description Aborts the current replication operation for the specified secondary volume. The current replication may be running or suspended. This command must be issued on the system that owns the secondary volume.

Level Manage

Syntax `abort replication`
    `[set` *`replication-set`*`]`
    `replication-volume`

Parameters `set` *`replication-set`*
Optional. The name or serial number of the replication set.

`replication-volume`
The name or serial number of the secondary volume. If the name is not unique across replication sets, specify the set parameter.

Example Abort replication of primary volume `V1` to secondary volume `rV1`:

`# abort replication rV1`

See also • resume replication
• suspend replication
• show replication-sets
• show replication-volumes

abort scrub

Description Aborts a scrub operation for specified storage-pool components or disks.

Level Diagnostic

Syntax `abort scrub`
`    [storage-pool-components` *`storage-pool-components`*`]`
`    [disks` *`disks`*`]`

Parameters `storage-pool-components` *`storage-pool-components`*
Optional. A comma-separated list of the names or serial numbers of the storage-pool components to stop scrubbing. For storage-pool component syntax, see Command syntax on page 18.

`disks` *`disks`*
Optional. The IDs of the disks to stop scrubbing. For disk syntax, see Command syntax on page 18.

Example Abort scrubbing storage-pool component `A_01_01`:

`# abort scrub storage-pool-component A_01_01`

Abort scrubbing disk 1.15:

`# abort scrub disks 1.15`

See also • scrub disks
• scrub storage-pool-components
• show disks
• show storage-pool-components

abort verify

Description Aborts the verify operation for specified storage-pool components.

Level Diagnostic

Syntax `abort verify storage-pool-components` *`storage-pool-components`*

Parameters `storage-pool-components` *`storage-pool-components`*
A comma-separated list of the names or serial numbers of the storage-pool components to stop
verifying. For storage-pool component syntax, see Command syntax on page 18.

Example Abort verifying storage-pool component `A_01_01`:

`# abort verify storage-pool-component A_01_01`

See also • show storage-pool-components
• verify storage-pool-components

abort volumecopy

Description Aborts copying a volume. When the abort is complete, the destination volume is deleted.

Level Manage

Syntax `abort volumecopy volume`

Parameters *volume*
The name or serial number of the source volume, the destination volume, or if a snapshot is being copied, its source volume. For volume syntax, see Command syntax on page 18.

Example Abort creating destination volume `v1_copy`:

`# abort volumecopy v1_copy`

See also
- copy volume
- show volume copy-status
- show volumes

add host-group-members

Description Adds hosts to a host group.

Level Manage

Syntax `add host-group-members`
`    hosts host-names`
`    host-group-name`

Parameters `hosts host-names`
A comma-separated list of the names of hosts to add to the specified host group. From 1 to 256 hosts can be specified.

`host-group-name`
The name of an existing host group.

Example Add existing hosts `Host3` and `Host4` to existing host group `HostGroup1`:

`# add host-group-members hosts Host3,Host4 HostGroup1`

See also • remove host-group-members
• show host-groups
• show hosts

add host-members

Description Adds initiators to a host.

Level Manage

Syntax
```
add host-members
      initiators initiators
      host-name
```

Parameters initiators *initiators*
A comma-separated list of the name or IDs of initiators to add to the specified host. From 1 to 128 initiators can be specified.

host-name
The name of an existing host.

Example Add existing initiators `FC-init3` and `FC-init4` to existing host `FC-Host`:

```
# add host-members initiators FC-init3,FC-init4 FC-Host
```

See also • remove host-members
• show hosts
• show initiators

add replication-volume

Description Adds an existing secondary volume to a replication set. The volume can be in the same system as the primary volume or in a remote system.

A secondary volume is a standard volume created specifically for use as the destination for a replication by using the create volume command's `prepare-replication-volume` parameter.

Secondary volumes serve as the destination for replicated data from the primary volume. When a replicate snapshot or replicate volume command is issued, data is replicated from the primary volume to the associated secondary volume in the replication set.

This command must be run on the primary system.

Level Manage

Syntax
```
add replication-volume
    link-type FC
    [max-queue #]
    [nowait]
    primary-volume volume
    [priority low|medium|high]
    [remote-system system]
    [secondary-address ip=IPs|wwnn=WWNNs|wwpn=WWPNs]
    [set replication-set]
    replication-volume
```

Parameters `link-type FC`
Specifies the type of ports being used for the inter-system link:

- `FC`: FC ports

`max-queue #`
Optional. The number of replication images to consider when determining the next image to replicate: 1–64. Used only if the `on-collision` parameter is set to `oldest`.

`nowait`
Optional. Adding a volume to a replication set can take the Storage Controller several minutes to complete. This parameter allows that processing to continue in the background so the Management Controller can process other commands.

`primary-volume volume`
The name or serial number of the replication volume to be the primary volume for the replication set.

`priority low|medium|high`
Optional. The priority of the replication process for the replication volume: low, medium, or high.

`remote-system system`
Optional for a local volume; required for a remote volume if the `secondary-address` parameter is not specified. The name or network-port IP address of the remote system.

`secondary-address ip=IPs|wwnn=WWNNs|wwpn=WWPNs`
Optional for a local volume; required for a remote volume if the `remote-system` parameter is not specified. Specifies host ports in the remote system by IP address, WWNN, or WWPN. An IP address value can include a port number; for example, 10.134.2.1:3260. Multiple values must be separated by commas and no spaces; for example: ip=10.134.2.1,10.134.2.2.

`set replication-set`
Optional. The name or serial number of the replication set.

`replication-volume`
The name or serial number of the secondary volume to add. If the name is not unique across replication sets, specify the set parameter.

Example Add secondary volume `MV2` to the replication set whose primary volume is `MV1`; set the replication priority to high; and allow a maximum of 2 queued images to be considered for replication:

```
# add replication-volume link-type FC secondary-address
wwpn=207000c0ffd52c31,217000c0ff52c31 primary-volume MV1 priority high max-queue
2 MV2
```

See also • show replication-sets
• show replication-volumes

add spares

Description Designates specified available disks to be spares. A spare can replace a failed disk of the same type (enterprise SAS or midline SAS) and the same or lower capacity that is in a storage-pool component. An SSD cannot be a spare.

Normally, a spare must be in a disk slot that is reserved for spares (slots 20–23 in a 24-disk enclosure or slots 10–11 in a 12-disk enclosure); however, this command allows any enterprise or midline SAS disk to be designated as a spare. This command does not affect any existing spares.

Level Diagnostic

Syntax `set spares disks` *`disks`*

Parameters `disks` *`disks`*
The IDs of the disks to designate as spares. For disk syntax, see Command syntax on page 18.

Example Designate disks 1.23 and 1.24 as spares:

```
# add spares disks 1.23,1.24
```

See also • delete spares
• show disks

add storage

Description Provisions disks into tiers in storage pools, according to rules defined by the storage-system manufacturer. In a new system, this command quickly provisions disks in enclosures so that you can proceed with creating and mapping volumes. In an existing system, this command quickly provisions disks in new enclosures that are added to expand storage.

Depending on a disk's type and the slot in which is located, a disk will be provisioned into a storage-pool component in the Performance, Standard, or Archive tier, as a spare disk for a certain storage-pool tier, or for use as storage-pool read cache. The system alternates allocation of disk resources between the two storage pools.

Some disk configurations will cause an entire enclosure to be associated with one storage pool, leaving the system unbalanced in terms of pool allocation.

- For a 24-disk enclosure the system will automatically balance disk resources between the two storage pools.
- For a 12-disk enclosure, the system will automatically assign all disk resources to one storage pool. If a second 12-disk enclosure is added, its resources are assigned to the partner storage pool to balance the configuration. To specify which pool to add the disks to, use the `storage-pool` parameter.

📝 **NOTE:** If you replace a disk that is configured as a spare, the replacement disk's `Usage` value will be `AVAIL` (available). You must use the `add storage` command to configure the replacement disk as a spare (`Usage` value `SPARE`) so that it can be incorporated into a storage pool upon failure of an equivalent disk.

Level Manage

Syntax add storage
```
[enclosures enclosure-IDs]
[full-enclosure-list all-enclosure-IDs]
[storage-pool A|B]
[configuration-preference performance|capacity]
[noprompt]
```

Parameters enclosures *enclosure-IDs*
Optional. Adds disks in specified enclosures only. Multiple enclosure IDs must be separated by commas (with no spaces). If this parameter is omitted, the command will add disks from all enclosures.

full-enclosure-list *all-enclosure-IDs*
Optional. When adding storage to a configured system, use this parameter to specify the range of enclosure IDs in the system. To configure storage, the system must contain at least 48 disks.

storage-pool A|B
Optional. Specifies the storage pool to which to assign disks that will result in an unbalanced configuration; the value is case sensitive. If the disks being added will result in a balanced configuration, this parameter will be ignored. If this parameter is omitted, the system will add the disks to the next storage pool in the alternation sequence.

configuration-preference performance|capacity
Optional. When initially configuring storage, use this parameter to optimize the creation of storage-pool components in the Standard tier for better I/O performance or for greater capacity. If this parameter is omitted, the capacity option will be used. After storage is provisioned the first time, this setting cannot be changed; if specified, it will be ignored.

noprompt
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

Example When initially configuring storage, optimize the Standard tier for I/O performance:

```
# add storage configuration-preference performance
```

Add storage from all enclosures:

```
# add storage
```

Add storage from enclosure 2, which is a 12-disk enclosure, and assign its disks to storage pool B:

```
# add storage enclosures 2 storage-pool B
```

See also • show storage

add volume-group-members

Description   Adds volumes to a volume group.

Level   Manage

Syntax   `add volume-group-members`
      `volumes` *`volume-IDs`*
     `volume-group-name`

Parameters   `volumes` *`volume-IDs`*
A comma-separated list of the names or serial numbers of volumes to add to the specified volume group. From 1 to 20 volumes can be specified.

`volume-group-name`
The name of an existing volume group.

Example   Add existing volumes `Vol0002` and `Vol0003` to existing volume group `VolumeGroup1`:

```
# add volume-group-members volumes Vol0002,Vol0003 VolumeGroup1
```

See also   •   remove volume-group-members
        •   show volume-groups
        •   show volumes

clear cache

Description Clears unwritable cache data from both controllers. This data cannot be written to disk because it is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host and disk. Unwritable cache is also called orphan data.

You can clear unwritable cache data for a specified volume or for all volumes.

When there are several volumes' worth of unwritable data, running the `clear cache` command will clear only the first volume's unwritable data. Therefore, it might become necessary to run the "clear cache" command multiple times to remove all cache data. Use the following process to ensure all cache data has been cleared.

1. Run the `show unwritable-cache` command to view the percent of unwritable cache in each controller.
2. Run the `clear cache` command.
3. Run the `show unwritable-cache` command to view the percent of unwritable cache in each controller. If the value is not 0, repeat steps 1–3 until the value is 0.

Level Manage

Syntax `clear cache [volume volume]`

Parameters `volume volume`
Optional. The name or serial number of the volume whose cache data should be cleared. For volume syntax, see Command syntax on page 18. If this parameter is omitted, the command clears any unwritable cache data for volumes that are no longer online or that no longer exist.

Example Clear unwritable cache data for volume `V1` from both controllers:

`# clear cache volume v1`

See also • show unwritable-cache
• show volumes

clear disk-metadata

Description Clears metadata from leftover disks. **For use by or with direction from technical support.**

⚠ **CAUTION:**
- Only use this command when all storage-pool components are online and leftover disks exist. Improper use of this command may result in data loss.
- Do not use this command when a storage-pool component is offline and one or more leftover disks exist.

If you are uncertain whether to use this command, contact technical support for further assistance.

Each disk in a storage-pool component has metadata that identifies the owning storage-pool component, the other disks in the storage-pool component, and the last time data was written to the storage-pool component. The following situations cause a disk to become a *leftover*:

- The disks' timestamps do not match so the system designates members having an older timestamp as leftovers.
- A disk is not detected during a rescan, then is subsequently detected.

When a disk becomes a leftover, the following changes occur:

- The disk's health becomes `Degraded` and its `Usage` value becomes `LEFTOVR`.
- The disk is automatically excluded from the storage-pool component, causing the storage-pool component's health to become `Degraded` or `Fault`, depending on the RAID level.
- The disk's fault LED is illuminated amber.

If a spare is available, and the health of the storage-pool component is `Degraded`, the component will use that spare to start reconstruction. When reconstruction is complete, you can clear the leftover disk's metadata. Clearing the metadata will change the disk's health to `OK` and its `Usage` value to `AVAIL`. The disk may then become the target of a copyback operation from the original spare, or the disk may become available for use in a new storage-pool component.

If a spare is not available to begin reconstruction, or reconstruction has not completed, keep the leftover disk so that you'll have an opportunity to recover its data.

This command clears metadata from leftover disks only. If you specify disks that are not leftovers, the disks are not changed.

Level Diagnostic

Syntax `clear disk-metadata` *disks*

Parameters *disks*
IDs of the leftover disks to clear metadata from. For disk syntax, see

Example Show disk usage:

```
# show disks
Location ... Usage    ...
---------------------...
1.1      ... LEFTOVR  ...
1.2      ... STORAGE  ...
...
```

Clear metadata from a leftover disk:

```
# clear disk-metadata 1.1
```

See also show disks

clear events

Description Clears the event log for controller A, B, or both. **For use by or with direction from technical support.**

Level Manage

Syntax `clear events [a|b|both]`

Parameters `a|b|both`
Optional. The controller event log to clear. If this parameter is omitted, both event logs are cleared.

Example Clear the event log for controller A:

`# clear events a`

See also • show events

clear expander-status

Description Clears the counters and status for SAS expander lanes. **For use by or with direction from technical support.**

Counters and status can be reset to a good state for all enclosures, or for a specific enclosure whose status is `Error` as shown by the show expander-status command.

📝 **NOTE:** If a rescan is in progress, the clear operation will fail with an error message saying that an EMP does exist. Wait for the rescan to complete and then retry the clear operation.

Level Manage

Syntax `clear expander-status [enclosure ID]`

Parameters `enclosure ID`
Optional. The enclosure number.

Example Clear the expander status for the first enclosure:

`# clear expander-status enclosure 0`

See also • show expander-status

copy volume

Description   Copies a standard or snapshot volume to a new standard volume in the same storage pool. If the source volume is a snapshot, you can choose whether to include its modified data (data written to the snapshot since it was created). The new volume is completely independent of the source volume.

For a standard volume, the volume copy creates a transient snapshot, copies the data from the snapshot to the new volume, and deletes the snapshot when the copy is complete. For a snapshot, the copy operation is performed directly from the source; this source data may change if modified data is to be included in the copy and the snapshot is mounted and I/O is occurring to it.

To ensure the integrity of a copy of a standard volume, unmount the volume or at minimum perform a system cache flush and refrain from writing to the volume. Since the system cache flush is not natively supported on all operating systems, it is recommended to unmount temporarily. The copy operation is for all data on the disk at the time of the request, so if there is data in the OS cache, that data will not be copied over. Unmounting the volume forces the cache flush from the OS. After the copy operation has started, it is safe to remount the volume and/or resume I/O.

To ensure the integrity of a copy of a snapshot with modified data, unmount the snapshot or perform a system cache flush. The snapshot will not be available for read or write access until the copy operation is complete, at which time you can remount the snapshot. If modified write data is not to be included in the copy, then you may safely leave the snapshot mounted. During a copy operation using snapshot modified data, the system takes the snapshot offline.

While the copy operation is in progress, the new volume's type is shown as `standard*`; when complete, it changes to `standard`. To see copy progress, use the show volume copy-status command.

Level   Manage

Syntax
```
copy volume
    source-volume source-volume
    [modified-snapshot yes|no]
    [prompt yes|no]
    destination-volume
```

Parameters   `source-volume` *source-volume*
The name or serial number of the volume or snapshot to copy. For volume syntax, see Command syntax on page 18.

`modified-snapshot yes|no`
Optional. Specifies whether to include or exclude modified write data from the snapshot in the copy. This parameter applies only when the source volume is a snapshot.

- `yes`: Include modified snapshot data.
- `no`: Exclude modified snapshot data.

If this parameter is omitted for a snapshot, modified snapshot data is excluded.

`prompt yes|no`
Optional. For scripting, this specifies an automatic response to the prompt to unmount the source volume before proceeding:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

destination-volume
A name for the volume to create. For volume syntax, see Command syntax on page 18.

Example   Copy standard volume `V1` to new volume `V1copy`:

```
# copy volume source-volume V1 V1copy
The source volume MUST be unmounted from all hosts prior to starting a volume
copy.
Are you ready to continue? (y/n) yes
```

See also   • abort volumecopy
   • show volume copy-status
   • show volumes

create certificate

Description   Creates or removes a custom security certificate.

The storage system supports use of unique certificates for secure data communications, to authenticate that the expected storage systems are being managed. Use of authentication certificates applies to the HTTPS protocol, which is used by the web server in each controller module.

After using this command you must restart the system's Management Controllers to have the change take effect.

Level   Manage

Syntax
```
create certificate
    [unique]
    [contents content-string]
    [restore]
    [noprompt]
```
One of the first three optional parameters must be supplied.

Parameters   `unique`
Optional. A security certificate is generated based on the system's serial number and other standard values. This certificate is installed, and the original certificate is archived.

`contents content-string`
Optional. A security certificate is generated based on the supplied content. The content becomes the subject of the certificate creation request and must be formatted as `/type0=value0/type1=value1/type2=...`, where types include `CO` for country, `ST` for state or province, `L` for location, `CN` for common name, and `O` for organization. The content string cannot exceed 100 characters and cannot contain space or semicolon characters. An example is `/CO=US/ST=CO/O=MyOrganization/CN=www.mysite.com`. This certificate is installed, and the original certificate is archived.

`restore`
Optional. The archived original certificate is restored and the custom certificate is discarded. The custom certificate may have been created with this CLI command or uploaded using FTP.

`noprompt`
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a response. Specifying this parameter allows the command to proceed without user interaction.

Example   Create a custom certificate based on the system's serial number:

```
# create certificate unique
To make the certificate change take effect you must restart both Management
Controllers by entering "restart mc both". Press Enter to acknowledge this
message.
```

Restore the original certificate that existed before the custom one was generated:

```
# create certificate restore
To make the certificate change take effect you must restart both Management
Controllers by entering "restart mc both". Press Enter to acknowledge this
message.
```

See also   • restart mc both
• show certificate

create host

Description Creates a host that includes specified initiators, and optionally adds the host to a host group.

Level Manage

Syntax
```
create host
     initiators initiators
     host-name
     [host-group host-group-name]
```

Parameters   initiators *initiators*
A comma-separated list of initiator nicknames, IDs, or both, with no spaces. For FC or SAS the ID is a WWPN. A WWPN can include a colon between each byte but the colons will be discarded.

host-name
A name for the host. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

host-group *host-group-name*
Optional. The name of an existing host group to which to add the new host.

Example Create an entry named Host1 for an FC host whose WWPN is 207000C0FF001122:

```
# create host initiators 207000C0FF001122 Host1
```

Create an entry for an FC host by pasting a WWPN that includes colons:

```
# create host initiators 20:70:00:c0:ff:d7:4c:07 Host2
```

See also   • set host
• show hosts
• show initiators

create host-group

Description Creates a host group that includes specified hosts.

Level Manage

Syntax
```
create host-group
    hosts hosts
    host-group-name
```

Parameters hosts *hosts*
A comma-separated list of the names of hosts to include in the host group. From 1 to 256 hosts can be specified.

host-group-name
A name for the host group. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

Example Create a host group named `HostGroup1` that includes hosts `Host1` and `Host2`:

```
# create host-group hosts Host1,Host2 HostGroup1
```

See also
- set host-group
- show host-groups
- show hosts

create remote-system

Description Creates a persistent association with a remote storage system. This allows a local system to track remote systems by their network-port IP addresses and cache their login credentials. The IP address you specify is used to connect to the remote system and obtain information such as the system name and both controllers' IP addresses. You can then use the system name or an IP address in commands that need to interact with the remote system.

Level Manage

Syntax
```
create remote-system
     username username
     password password
     IP-address
```

Parameters username *username*
The name of a user in the remote system. This must be a manage-level user to remotely configure or provision that system.

password *password*
The password for that user.

IP-address
The network-port IP address of the remote system.

Example Create a remote system:

```
# create remote-system username manage password !manage 10.122.1.21
```

See also
- delete remote-system
- remote
- set remote-system
- show remote-systems

create replication-set

Description Creates a replication set from a specified standard volume. The specified volume becomes the replication set's primary volume.

You can specify a remote volume to be the replication set's secondary volume, or one will be created. When a replicate snapshot or replicate volume command is issued, data in the primary volume is replicated to the associated secondary volume.

You can associate the primary volume with a remote system in two ways:

- If the local and remote MCs can communicate, you can specify a storage pool in a remote system. A standard volume is created in the remote storage pool.
- If the local and remote MCs cannot communicate or you want to use an existing replication-prepared volume in the remote system, you can specify the volume to use as the secondary volume.

You can also start the initial replication by specifying the `snapshot` parameter.

If you create a replication set without specifying a replication destination, a partial set is created and you must use the add replication-volume command to complete the set before you can perform replication.

If the create transaction fails, a prompt asks if you want to revert the transaction, which undoes any changes made in attempting to create the replication set. To revert, enter `yes`; otherwise, enter `no`.

Level Manage

Syntax
```
create replication-set
     [link-type FC]
     [max-queue #]
     [noprompt]
     [nowait]
     [primary-address ip=IPs|wwnn=WWNNs|wwpn=WWPNs]
     [priority low|medium|high]
     [remote-system system]
     [remote-storage-pool storage-pool-ID]
     [remote-volume volume]
     [secondary-address ip=IPs|wwnn=WWNNs|wwpn=WWPNs]
     [set name]
     [snapshot snapshot]
     primary-volume
```

Parameters `link-type FC`
Optional; required if the `primary-address` parameter is *not* specified. Specifies the type of ports being used for the inter-system link:

- `FC`: FC ports.

`max-queue #`
Optional. The number of replication images to consider when determining the next image to replicate: 1–64. Used only if the `on-collision` parameter is set to `oldest`.

`noprompt`
Optional in console format; required for XML API format. Suppresses the prompt to revert the transaction if the command fails, which requires a yes or no response. If this parameter is specified and the command fails, the transaction is automatically reverted.

`nowait`
Optional. Adding a volume to a replication set can take the Storage Controller several minutes to complete. This parameter allows that processing to continue in the background so the Management Controller can process other commands.

primary-address ip=*address*|iqn=*IQN*|wwn=*WWN*
Optional. Specifies host ports in the local system by IP address, WWNN, or WWPN. An IP address value can include a port number; for example, 10.134.2.1:3260. Multiple values must be separated by commas and no spaces; for example: ip=10.134.2.1,10.134.2.2.

priority low|medium|high
Optional. The priority of the replication process for the replication volumes: low, medium, or high.

remote-system *system*
Optional; required if the `remote-storage-pool` or `remote-volume` parameter specifies a destination in a remote system. The name or network-port IP address of the remote system.

remote-storage-pool *storage-pool-ID*
Optional. The name or serial number of the storage pool in which the remote volume should be created. If the storage pool is in a remote system:

- You must specify the `remote-system` parameter.
- If the local and remote MCs can communicate, you can specify a name or serial number; otherwise, you must specify a serial number.

remote-volume *volume*
Optional; do not use with the `remote-storage-pool` parameter. The name or serial number of a replication-prepared volume to use as a secondary volume. If the volume is in a remote system and the local and remote MCs can communicate, you can specify a name or serial number; otherwise, you must specify a serial number.

secondary-address ip=*IPs*|wwnn=*WWNNs*|wwpn=*WWPNs*
Optional if the MCs can communicate and the `remote-system` parameter is specified; required if the MC cannot communicate and a replication-prepared volume is specified. Specifies host ports in the remote system by IP address, WWNN, or WWPN. An IP address value can include a port number; for example, 10.134.2.1:3260. Multiple values must be separated by commas and no spaces; for example: ip=10.134.2.1,10.134.2.2.

set *name*
Optional. A name for the new replication set. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes. A name that includes a space must be enclosed in double quotes. If this parameter is omitted, the set is named rs*primary-volume*.

snapshot *snapshot*
Optional. A name for the replication snapshot that will be created during the replication process. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes. A name that includes a space must be enclosed in double quotes. Specifying this parameter will start the initial replication.

primary-volume
The name or serial number of a standard volume to use as the primary volume.

Example For two systems connected with FC links, set up replication from local standard volume `Data` to remote storage pool `A` in remote system `System2`:

```
# create replication-set link-type FC remote-system System2 remote-storage-pool
A Data
```

For two systems connected with FC links, set up replication from local standard volume `Source` to remote replication-prepared volume `Dest`:

```
# create replication-set link-type FC remote-system System2 remote-volume Dest
Source
```

For two systems whose MCs aren't currently connected, set up replication from local standard volume `Vol0001` to a remote replication-prepared volume:

```
# create replication-set primary-address wwpn=207000c0ffd52c31,217000c0ffd52c31
remote-volume 00c0ffda421f000089b16b4d01000000 secondary-address
wwpn=207000c0ffd52c31,217000c0ffd52 Vol0001
```

Create a local replication set using existing standard volume `Personnel` as the primary volume and a new volume in storage pool B for the secondary volume:

```
# create replication-set link-type FC remote-storage-pool B Personnel
```

Create a local replication set using existing standard volume `Test` as the primary volume and existing replication-prepared volume `Test-backup` as the secondary volume:

```
# create replication-set link-type FC remote-volume Test-backup Test
```

See also
- add replication-volume
- delete replication-set
- remove replication-volume
- replicate snapshot
- show remote-systems
- show replication-images
- show replication-sets
- show replication-volumes

create schedule

Description Schedules a task to run automatically.

When scheduling `ReplicateVolume` tasks, a best practice is to schedule no more than three volumes to start replicating at the same time, and for those replications to recur no less than 60 minutes apart. If you schedule more replications to start at the same time, or schedule replications to start more frequently, some scheduled replications may not have time to complete.

Level Manage

Syntax
```
create schedule
    schedule-specification "specification"
    task-name task-name
    schedule-name
```

Parameters schedule-specification "*specification*"
Defines when the task will first run, and optionally when it will recur and expire. You can use a comma to separate optional conditions. Dates cannot be in the past. For times, if neither AM nor PM is specified, a 24-hour clock is used.

- start *yyyy-mm-dd hh*:*mm* [AM|PM]
  Specifies a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.

- [every # minutes|hours|days|weeks|months|years]
  Specifies the interval at which the task will run.
  For better performance when scheduling a TakeSnapshot task that will run under heavy I/O conditions or on more than three volumes, the retention count and the schedule interval should be set to similar values; for example if the retention count is 10 then the interval should be set to 10 minutes.
  For a ReplicateVolume task, the default and minimum interval is 30 minutes.

- [between *hh*:*mm* [AM|PM] and *hh*:*mm* [AM|PM]]
  Constrains the time range during which the task is permitted to run. Ensure that the start time is within the specified time range.

- [only any|first|second|third|fourth|fifth|last|#st|#nd|#rd|#th weekday|weekendday|Sunday|Monday|Tuesday|Wednesday|Thursday|Friday |Saturday of year|month|January|February|March|April|May|June|July |August|September|October|November|December]
  Constrains the days or months when the task is permitted to run. Ensure that this constraint includes the start date.

- [count #]
  Constrains the number of times the task is permitted to run.

- [expires *yyyy-mm-dd hh*:*mm* [AM|PM]]
  Specifies when the schedule expires, after which the task will no longer run.

task-name *task-name*
The task to run. The name is case sensitive.

schedule-name
A name for the new schedule. The name is case sensitive; cannot include angle brackets, comma, double quote, or backslash; and can have a maximum of 32 bytes. A name that includes a space must be enclosed in double quotes.

Example Create schedule Sched1 that runs Task1 for the first time on March 1, 2012; runs daily between midnight and 1:00 AM; and runs for the last time in the morning of January 1, 2013:

```
# create schedule schedule-specification "start 2012-03-01 00:01, every 1 days,
between 12:00 AM and 1:00 AM, expires 2013-01-01 1:00 AM" task-name Task1 Sched1
```

Create schedule `Sched2` that runs `Task2` for the first time on March 1, 2012, and on the first weekday of each month, with no expiration:

```
# create schedule schedule-specification "start 2012-03-01 00:01 only first
weekday of month" task-name Task2 Sched2
```

See also
- delete schedule
- set schedule
- show schedules
- show tasks

create snapshots

Description   Creates a snapshot of each specified standard volume.

Level   Manage

Syntax   `create snapshots`
  `volumes` *`volumes`*
  `snap-names`

Parameters   volumes *`volumes`*
A comma-separated list of up to 16 standard volumes of which to take snapshots. For volume syntax, see Command syntax on page 18.

`snap-names`
A comma-separated list of names for the resulting snapshots. A snapshot name is case sensitive; can include spaces and printable UTF-8 characters except angle brackets, backslash, comma, and double quote; and can have a maximum of 20 bytes. A name that includes a space must be enclosed in double quotes.

Example   Create snapshots of standard volumes `V1` and `V2`:

`# create snapshots volumes V1,V2 V1snap,V2snap`

See also   • show snapshots
  • show volumes

create task

Description Creates a task that can be scheduled. You can create a task to take a snapshot of a standard volume; to copy a standard or snapshot volume to a new standard volume; to replicate a replication set's primary volume to a remote system; to reset a snapshot; or to enable or disable drive spin down.

⚠ **CAUTION:** Before scheduling a `ResetSnapshot` task, consider that if the snapshot is mounted to a host, the snapshot must be unmounted before the reset is performed; leaving it mounted can cause data corruption. You should create a scheduled job on the host to unmount the snapshot prior to resetting it.

Level Manage

Syntax To create a task to take a snapshot:

```
create task
    type TakeSnapshot
    source-volume volume
    snapshot-prefix prefix
    retention-count #
    task-name
```

To create a task to reset a snapshot:

```
create task
    type ResetSnapshot
    snapshot-volume volume
    task-name
```

To create a task to copy a volume:

```
create task
    type VolumeCopy
    source-volume volume
    dest-prefix prefix
    [modified-snapshot yes|no]
    task-name
```

To create a task to replicate a volume:

```
create task
    type ReplicateVolume
    source-volume volume
    snapshot-prefix prefix
    retention-count #
    [replication-mode new-snapshot|last-snapshot]
    task-name
```

To create a task to enable spin down for all disks:

```
create task
    type EnableDSD
    task-name
```

To create a task to disable spin down for all disks:

```
create task
    type DisableDSD
    task-name
```

Parameters `type TakeSnapshot|ResetSnapshot|VolumeCopy|ReplicateVolume|EnableDSD`
`|DisableDSD`
The task type:

- `TakeSnapshot`: Takes a snapshot of a standard volume.
- `ResetSnapshot`: Deletes the data in the snapshot and resets it to the current data in the source volume. The snapshot's name and other volume characteristics are not changed.
- `VolumeCopy`: Copies a standard or snapshot volume to a new standard volume. The command creates the destination volume in the storage pool that contains the source volume.
- `ReplicateVolume`: Replicates a replication set's primary volume to a remote system.
- `EnableDSD`: Enables spin down for all disks. You can use this to enable or resume spin down during hours of infrequent activity.
- `DisableDSD`: Disables spin down for all disks. You can use this to disable or suspend spin down during hours of frequent activity.

`source-volume` *volume*
For a `TakeSnapshot` task, the name or serial number of the standard volume of which to take a snapshot. For a `VolumeCopy` task, the name or serial number of the standard or snapshot volume to copy. For a `ReplicateVolume` task, the name or serial number of the primary volume to replicate. For volume syntax, see Command syntax on page 18.

`snapshot-prefix` *prefix*
A label to identify snapshots created by this task. Snapshot names have the format *prefix*_s#, where # starts at `001`.

`retention-count` *#*
For a `TakeSnapshot` task this parameter specifies the number of snapshots created by this task to retain, from 1 to the licensed limit. When a new snapshot exceeds the limit, the oldest snapshot with the same prefix is deleted.

For a `ReplicateVolume` task this parameter specifies the number of replication images created by this task to retain, from 2 to 32. When a new image exceeds this limit, the oldest image with the same prefix is deleted. This parameter applies to the primary volume only; for the secondary volume, images will accumulate until either the secondary storage pool's space limit is reached or the maximum number of images is reached, after which the oldest image will be deleted as new images are created.

`snapshot-volume` *volume*
Name or serial number of the snapshot to reset. For volume syntax, see Command syntax on page 18.

`dest-prefix` *prefix*
A label to identify the volume copy created by this task. Copy names have the format *prefix*_c#, where # starts at `001`.

`modified-snapshot yes|no`
Optional. Specifies whether to include or exclude modified write data from the snapshot in the copy. This parameter applies only when the source volume is a snapshot; it is ignored if the source volume is a standard volume.

- `yes`: Include modified snapshot data.
- `no`: Exclude modified snapshot data.

If this parameter is omitted for a snapshot, modified snapshot data is excluded.

`replication-mode new-snapshot|last-snapshot`
Optional. Specifies whether to replicate a new snapshot of the volume to the remote system, or to replicate the last (most recent existing) snapshot of the volume to the remote system.

- `new-snapshot`: Replicate a new snapshot.
- `last-snapshot`: Replicate the most recent existing snapshot.

If this parameter is omitted, a new snapshot is replicated.

`task-`*name*
A name for the new task. The name is case sensitive; cannot include angle brackets, comma, double quote, or backslash; and can have a maximum of 32 bytes. A name that includes a space must be enclosed in double quotes.

Example Create task `Snap` that takes a snapshot of standard volume `V1` and retains only the latest four snapshots with the prefix `V1` (e.g., `V1_S0001`):

```
# create task type TakeSnapshot source-volume V1 snapshot-prefix V1
retention-count 4 Snap
```

Create task `Reset` that resets snapshot `V1_S0001`:

```
# create task type ResetSnapshot snapshot-volume V1_S0001 Reset
```

Create task `Copy` that copies volume `V1` to a new volume named `C_V0001` in the same storage pool:

```
# create task type VolumeCopy source-volume V1 dest-prefix C modified-snapshot
yes Copy
```

Create task `Replicate` that replicates primary volume `V2`:

```
# create task type ReplicateVolume source-volume V2 snapshot-prefix V2
retention-count 4 Replicate
```

Create a task to enable or resume spin down for all disks:

```
# create task type EnableDSD DSDresume
```

Create a task to disable or suspend spin down for all disks:

```
# create task type DisableDSD DSDsuspend
```

See also • create schedule
- delete task
- set task
- show tasks
- show volumes

create temp-license

Description Creates and installs a temporary license to try licensed features of the storage system. You can create a temporary license one time. A temporary license will expire 60 days from the time it is created. If you do not install a permanent license before the temporary license expires, you cannot create new items using these features; however, you can continue to use existing items. After a temporary license is created or a permanent license is installed, the command cannot be re-run.

Level Manage

Syntax `create temp-license`

Example Enable licensed features for 60 days in order to try them:

`# create temp-license`

See also • show license

create user

Description Creates a user account. The system supports 12 user accounts. You can create a general user that can access the WBI, CLI, or FTP interface, or an SNMPv3 user that can access the MIB or receive trap notifications. SNMPv3 user accounts support SNMPv3 security features such as authentication and encryption.

Level Manage

Syntax
```
create user
      [authentication-type MD5|SHA|none]
      [base 2|10]
      [interfaces interfaces]
      [level monitor|manage|diagnostic]
      [roles roles]
      [locale English|en|Spanish|es|Japanese|ja|Chinese-simplified|zh-s]
      [password password]
      [precision #]
      [privacy-password encryption-password]
      [privacy-type DES|AES|none]
      [storage-size-base 2|10]
      [storage-size-precision #]
      [storage-size-units auto|MB|GB|TB]
      [temperature-scale celsius|c|fahrenheit|f]
      [timeout #]
      [trap-host IP-address]
      [type novice|standard|advanced|diagnostic]
      [units auto|MB|GB|TB]
      name
```

Parameters authentication-type MD5|SHA|none
Optional. For an SNMPv3 user, this specifies whether to use a security authentication protocol. Authentication uses the user password.

- MD5: MD5 authentication. This is the default.
- SHA: SHA (Secure Hash Algorithm) authentication.
- none: No authentication.

base 2|10
Optional. Sets the base for entry and display of storage-space sizes:

- 2: Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
- 10: Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is the default.

Operating systems usually show volume size in base 2. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2.

`interfaces` *interfaces*
Optional. Specifies the interfaces that the user can access. Multiple values must be separated by commas and no spaces. The defaults are `cli` and `wbi`. A command that specifies `snmpuser` or `snmptarget` cannot also specify a non-SNMP interface.

- `cli`: Command-line interface.
- `wbi`: Web-browser interface (Storage Management Console).
- `ftp`: File transfer protocol interface.
- `smis`: Storage Management Initiative Specification (SMI-S) interface.
- `snmpuser`: Allows an SNMPv3 user to view the SNMP MIB.
- `snmptarget`: Allows an SNMPv3 user to receive SNMP trap notifications. This option requires the `trap-host` parameter.
- `none`: No interfaces.

`level monitor|manage|diagnostic`
Optional.

- `monitor`: User can view but not change system settings. This is the default.
- `manage`: User can view and change system settings.
- `diagnostic`: User can access additional, diagnostic commands.

You cannot specify both this parameter and the `roles` parameter.

`roles` *roles*
Optional. Specifies the user's roles as one or more of the following values:

- `monitor`: User can view but not change system settings. This is the default.
- `manage`: User can view and change system settings.
- `diagnostic`: User can access additional, diagnostic commands.

Multiple values must be separated with a comma (with no spaces). If multiple values are specified, the user's access to commands will be determined by the highest-level role specified. You cannot specify both this parameter and the `level` parameter.

`locale English|en|Spanish|es|Japanese|ja|Chinese-simplified|zh-s`
Optional. The display language. The default is English.

`password` *password*
Optional in console format; required for XML API format. Sets a new password for the user. The value is case sensitive; can include a maximum of 32 bytes using printable UTF-8 characters except angle brackets, backslash, comma, or double quote. A value that includes a space must be enclosed in double quotes. If this parameter is omitted, the command prompts you to enter and re-enter a value, which is displayed encrypted. For an SNMPv3 user whose `authentication-type` parameter is set to use authentication, this specifies the authentication password and must include at least 8 characters.

`precision #`
Optional. Sets the number of decimal places (1–10) for display of storage-space sizes. Default is 1.

`privacy-password` *encryption-password*
Optional. For an SNMPv3 user whose `privacy-type` parameter is set to use encryption, this specifies the encryption password. The value is case sensitive; can include a maximum of 32 bytes using printable UTF-8 characters except angle brackets, backslash, comma, or double quote; must include at least 8 characters.

`privacy-type DES|AES|none`
Optional. For an SNMPv3 user, this specifies whether to use a security encryption protocol. This parameter requires the `privacy-password` parameter and the `authentication-type` parameter.

- `DES`: Data Encryption Standard.
- `AES`: Advanced Encryption Standard.
- `none`: No encryption. This is the default.

`storage-size-base 2|10`
Optional. Alias for `base`.

`storage-size-precision #`
Optional. Alias for `precision`.

`storage-size-units auto|MB|GB|TB`
Optional. Alias for `units`.

`temperature-scale celsius|c|fahrenheit|f`
Optional. Sets the scale for display of temperature values:

- `fahrenheit` or `f`: Temperatures are shown in degrees Fahrenheit.
- `celsius` or `c`: Temperatures are shown in degrees Celsius. This is the default.

`timeout #`
Optional. Sets the timeout value in seconds for the login session. Valid values are 120–43200 seconds (2–720 minutes). The default is 1800 seconds (30 minutes).

`trap-host IP-address`
Optional. For an SNMPv3 user whose `interface` parameter is set to `snmptarget`, this specifies the IP address of the host that will receive SNMP traps.

`type novice|standard|advanced|diagnostic`
Optional. Identifies the user's experience level. The default is `standard`. This parameter does not affect access to commands.

`units auto|MB|GB|TB`
Optional. Sets the unit for display of storage-space sizes:

- `auto`: Sizes are shown in units determined by the system. This is the default.
- `MB`: Sizes are shown in megabytes.
- `GB`: Sizes are shown in gigabytes.
- `TB`: Sizes are shown in terabytes.

Based on the `precision` setting, if a size is too small to meaningfully display in the selected unit, the system uses a smaller unit for that size. For example, if `units` is set to TB, `precision` is set to 1, and `base` is set to 10, the size 0.11709 TB is instead shown as 117.1 GB.

name
A name for the new user, which cannot already exist in the system. The name is case sensitive; cannot include a comma, double quote, backslash, or space; and can have a maximum of 29 bytes.

📝 **NOTE:** The user names `admin` and `api` are reserved for internal use.

Example Create user `John` who will view system information using base 2 in Storage Management Console:

```
# create user base 2 interfaces wbi level monitor John
Enter new password: ****
Re-enter new password: ****
```

Create user `MIB` that can view the SNMP MIB, using authentication and encryption:

```
# create user interfaces snmpuser password whoAreYou authentication-type SHA
privacy-type AES privacy-password whoAreYou MIB
```

Create user `Traps` that can receive SNMP trap notifications, using authentication without encryption:

```
# create user interfaces snmptarget authentication-type MD5 trap-host
172.22.4.171 Traps
Enter new password: *********
Re-enter new password: *********
```

See also
- delete user
- set snmp-parameters
- set user
- show users

create volume

Description   Creates a standard volume in a storage pool. The volume is snapshot-enabled and is associated with the snap pool for the storage pool. You can create the volume unmapped or set its default mapping. Default mapping settings apply to all initiators, unless overridden by an explicit mapping between specific initiators and the volume. You can later change the mapping by using the map volume and unmap volume commands.

By default, this command will not map the created volume.

If you intend to use the volume as the secondary volume in a replication set, you can create a replication-prepared volume. A secondary volume cannot be mapped.

To create multiple volumes at once, use the create volume-set command.

Level   Manage

Syntax
```
create volume
    storage-pool storage-pool-ID
    size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]
    [access read-write|rw|read-only|ro|no-access]
    [lun LUN]
    [ports ports]
    [prepare-replication-volume]
    [volume-group volume-group-name]
    [preference archive|standard|performance]
    name
```

Parameters   storage-pool storage-pool-ID
The name or serial number of the storage pool in which to create the volume.

size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]
Sets the volume size using the current base, as shown by the show cli-parameters command. The unit can be specified as follows:

- If base 2 is in use: B (bytes), KiB (kibibytes), MiB (mebibytes), or GiB (gibibytes)
- If base 10 is in use: B (bytes), KB (kilobytes), MB (megabytes), or GB (gigabytes)
- If no unit is specified, the unit is 512-byte blocks.

Although smaller values are accepted, the minimum volume size is 1 MB (1 MiB). If overcommit is enabled, the size can exceed the physical capacity of the storage pool. The maximum volume size is 140 TB (128 TiB). To see whether overcommit is enabled, use the show system-parameters command.

access read-write|rw|read-only|ro|no-access
Optional. The access permission for initiators to the volume: read-write (rw), read-only (ro), or no-access. Read-only access applies to all ports. If no-access is specified, the volume is not mapped. The default is read-write.

lun LUN
Optional if the access parameter is set to no-access. Specifies the LUN to assign to the mapping on all ports.

ports ports
Optional. The controller host ports to use for the mapping. For example, if read-write access is specified you can specify the ports parameter to restrict the ports through which this access is available. For port syntax, see Command syntax on page 18. If this parameter is omitted, all ports are selected.

prepare-replication-volume
Optional. Specifies to create a secondary volume for use in a replication set. This parameter precludes use of the lun and ports parameters because a secondary volume cannot be mapped.

`volume-group` *volume-group-name*
Optional. The name of a volume group to which to add the volume. A name that includes a space must be enclosed in double quotes. If the group does not exist, it will be created.

`preference archive|standard|performance`
Optional. Specifies how to tune the tier-migration algorithm for the volume:

- `archive`: This setting specifies never to move this volume to the SSD tier, regardless of how often the data is accessed.
- `standard`: This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving the volume to the appropriate tier. This is the default.
- `performance`: This setting specifies to keep this volume in the SSD tier as much as possible.

name
A name for the new volume. The name is case sensitive; can include spaces and printable UTF-8 characters except angle brackets, backslash, comma, and double quote; and can have a maximum of 20 bytes. A name that includes a space must be enclosed in double quotes.

Example Create a 100 GB standard volume named `MyVolume` in pool A, map it to use LUN 5 with read-write access through port 1 in each controller, add it to volume group `MyGroup`, tune tier-migration for performance:

```
# create volume MyVolume storage-pool A size 100GB access rw lun 5 ports 1
volume-group MyGroup preference performance
```

Create a 20 GB standard volume named `Secrets` in storage pool A, and map it to use LUN 333 with read-only access through all ports:

```
# create volume Secrets storage-pool A size 20GB access read-only
```

Create a 1 TB secondary volume named `NewYork` in storage pool B:

```
# create volume NewYork storage-pool B size 1TB prepare-replication-volume
```

See also
- create volume-set
- delete volumes
- set volume
- show ports
- show storage-pools
- show volume-groups
- show volumes

create volume-group

Description   Creates a volume group that includes specified volumes.

Level   Manage

Syntax   `create volume-group`
`    volumes volumes`
`     volume-group-name`

Parameters   `volumes volumes`
A comma-separated list of the names of volumes to include in the volume group. From 1 to 20 volumes can be specified.

`volume-group-name`
A name for the volume group. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

Example   Create a volume group named `VGroup1` that includes volumes `Vol0001` and `Vol0002`:

`# create volume-group volumes Vol0001,Vol0002 VGroup1`

See also   • set volume-group
• show volume-groups
• show volumes

create volume-set

Description Creates a specified number of standard volumes in a storage pool. The volumes are snapshot-enabled and are associated with the snap pool for the storage pool. You can create the volumes unmapped or set their default mapping. Default mapping settings apply to all initiators, unless overridden by an explicit mapping between specific initiators and the volume. You can later change mappings by using the map volume and unmap volume commands.

By default, this command will not map the created volumes.

To create a single volume, use the create volume command.

Level Manage

Syntax
```
create volume-set
    storage-pool storage-pool-ID
    basename base-name
    count #
    size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]
    [baselun base-LUN]
    [access read-write|rw|read-only|ro|no-access]
    [ports ports]
    [volume-group volume-group-name]
    [preference archive|standard|performance]
```

Parameters `storage-pool storage-pool-ID`
The name or serial number of the storage pool in which to create the volume.

`basename base-name`
A name to which a number will be appended to generate a different name for each volume. A name is case sensitive; can include spaces and printable UTF-8 characters except angle brackets, backslash, comma, and double quotes; and can have a maximum of 16 bytes. A name that includes a space must be enclosed in double quotes.

Resulting volumes are numbered sequentially starting with 0000. If volumes with the specified basename already exist, names of new volumes start with the first available name in the sequence. For example: for basename `Data`, if `Data0000` and `Data0002` exist, the next volumes created will be `Data0001` and `Data0003`.

`count #`
The number of volumes to create, from 1 to 512. Volumes will be created up to the maximum number supported per storage pool.

`size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]`
Sets the size for each volume using the current base, as shown by the show cli-parameters command. The unit can be specified as follows:

- If base 2 is in use: `B` (bytes), `KiB` (kibibytes), `MiB` (mebibytes), or `GiB` (gibibytes)
- If base 10 is in use: `B` (bytes), `KB` (kilobytes), `MB` (megabytes), or `GB` (gigabytes)
- If no unit is specified, the unit is 512-byte blocks.

Although smaller values are accepted, the minimum volume size you can create is 1 MB (1 MiB). If overcommit is enabled, the volume size can exceed the physical capacity of the storage pool. The maximum volume size is 140 TB (128 TiB). To see whether overcommit is enabled, use the show system-parameters command. If overcommit is disabled and the combined size of the volumes will exceed the capacity of the storage pool, an error message is displayed and no volumes are created.

`baselun base-LUN`
Optional. The first in a sequence of LUNs to assign to map the volumes through ports specified by the `ports` parameter. If the `baselun` and `ports` parameters are omitted, the volumes are not mapped. If a LUN to be assigned to a volume is already in use, an error message is displayed and that volume and any subsequent volumes are not mapped.

`access read-write|rw|read-only|ro|no-access`
Optional. The access permission for initiators to the volumes: read-write (`rw`), read-only (`ro`), or no-access. Read-only access applies to all ports. If `no-access` is specified, the volume is not mapped. The default is `read-write`.

`ports` *ports*
Optional. The controller host ports to use for the mapping. For example, if read-write access is specified you can specify the ports parameter to restrict the ports through which this access is available. For port syntax, see Command syntax on page 18. If not all ports are specified, the unspecified ports are not mapped. If the `ports` and `baselun` parameters are omitted, the volumes are not mapped.

`volume-group` *volume-group-name*
Optional. The name of a volume group to which to add the volumes. A name that includes a space must be enclosed in double quotes. If the group does not exist, it will be created.

`preference archive|standard|performance`
Optional. Specifies how to tune the tier-migration algorithm for the volumes:

- `archive`: This setting specifies never to move the volumes to the SSD tier, regardless of how often the data is accessed.
- `standard`: This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving the volumes to the appropriate tier. This is the default.
- `performance`: This setting specifies to keep the volumes in the SSD tier as much as possible.

Example   Create three 100 GB standard volumes with base name `MyVolume` in pool A, map them to use LUNs 5–7 with read-write access on port 1 in each controller, add them to volume group `MyGroup`, and tune tier-migration for performance:

```
# create volume-set basename MyVolume storage-pool A size 100GB count 3
access rw baselun 5 ports 1 volume-group MyGroup preference performance
```

See also   • create volume
- delete volumes
- map volume
- set volume
- show maps
- show storage-pools
- show volume-groups
- show volumes

delete all-snapshots

Description Deletes all snapshots associated with a specified volume. All data associated with the snapshots is deleted and their space in the snap pool is freed for use. The specified volume can be a standard volume or a snap pool. The snapshots' schedules are also deleted.

Syntax `delete all-snapshots`
    `volume `*`volume`*
    `[delete-type all-standard-snapshots|all-replication-snapshots`
    `|all-snapshot-types]`

Parameters `volume `*`volume`*
The name or serial number of the standard volume or snap pool. For volume syntax, see Command syntax on page 18.

`delete-type all-standard-snapshots|all-replication-snapshots`
`|all-snapshot-types`
Optional; for replication-enabled systems only. Specifies the type of snapshots that can be deleted. If this parameter is omitted, the default is `all-standard-snapshots`, which means that the command will try to delete only standard snapshots and will not succeed if replication snapshots exist.

Example Delete all snapshots associated with standard volume `V1`:

`# delete all-snapshots volume V1`

Delete all replication snapshots associated with standard volume `V2`, which is the primary volume in a replication set:

`# delete all-snapshots volume V2 delete-type all-replication-snapshots`

See also • show snapshots
    • show volumes

delete all-volumes

Description Deletes all volumes in both storage pools. Associated mappings, schedules, and snapshots are also deleted.

⚠ **CAUTION:**   This command will delete all data in the system.

To delete specific volumes, use the delete volumes command.

Level Manage

Syntax `delete all-volumes`

Example Delete all volumes in both storage pools:

`# delete all-volumes`

See also • show volumes

delete host-groups

Description Deletes specified host groups and optionally all hosts in those groups.

Before using the option to delete all the hosts in the groups, ensure that the hosts are unmapped. If any host is mapped, the command will fail and no changes will be made.

Level Manage

Syntax
```
delete host-groups
       [delete-hosts]
       host-groups|all
```

Parameters delete-hosts
Optional. Specifies to delete all hosts in the groups. If this parameter is omitted, the host groups will be deleted but their hosts will not be deleted.

host-groups
Specifies a comma-separated list of the names of host groups to delete, or to delete all host groups.

Example Delete host groups HGroup1 and HGroup2 but not the hosts in those groups:

```
# delete host-groups HGroup1,HGroup2
```

Delete all host groups and the hosts in those groups:

```
# delete host-groups delete-hosts all
```

See also
- show host-groups
- show host-maps

delete hosts

Description   Deletes specified hosts that are not in a host group. Mapped and unmapped hosts can be deleted. Deleting a host does not delete its initiators.

Level   Manage

Syntax   `delete hosts hosts|all`

Parameters   `hosts|all`
Specifies a comma-separated list of the names of hosts to delete, or to delete all hosts.

Example   Delete hosts `Host1` and `Host2`:

`# delete hosts Host1,Host2`

Delete all hosts:

`# delete hosts all`

See also   • create host
• set host
• show host-maps
• show hosts

delete initiator-nickname

Description Deletes the nickname of an initiator. If the initiator was discovered or is mapped, the initiator's ID (WWN or IQN) will remain in the system. If you unmap the initiator, its ID will be removed from the system.

To rename an initiator, use the set initiator command.

Level Manage

Syntax `delete initiator-nickname nickname|initiator|all`

Parameters `nickname|initiator|all`
Specifies the nickname to delete, or the initiator ID for which to delete the nickname, or to delete all nicknames. For an FC or SAS initiator, the ID is a WWPN. A WWPN can include a colon between each byte but the colons will be discarded.

Example Delete nickname `FC-port1`:

`# delete initiator-nickname FC-port1`

Delete the nicknames of all initiators:

`# delete initiator-nickname all`

See also • set initiator
• show initiators

delete remote-system

Description Deletes the persistent association with a remote system.

After establishing replication to a remote system, if you choose to delete the remote system you can safely do so without affecting replications. However, because the remote system's name and IP address will no longer appear in user interfaces, record this information before deleting the remote system so that you can access it at a later time, such as to delete old replication images or for disaster recovery.

Level Manage

Syntax `delete remote-system` *`system`*

Parameters *`system`*
Name or network-port IP address of the remote system.

Example Delete remote system `System2`:

```
# delete remote-system System2
```

See also • create remote-system
• remote
• set remote-system
• show remote-systems

delete replication-set

Description Dissolves a replication set. The replication volumes associated with the replication set are converted to standard volumes and any replication snapshots associated with the replication volumes are converted to standard snapshots. Snapshots are converted regardless of the number of snapshots allowed by the system's license. This command must be run on the primary system.

Level Manage

Syntax `delete replication-set` *ID*

Parameters *ID*
Name or serial number of either the replication set or its primary volume.

Example Delete replication set `RS1`:

`# delete replication-set RS1`

Delete the replication set with primary volume `V1`:

`# delete replication-set V1`

See also • show replication-sets
  • show replication-volumes

delete schedule

Description   Deletes a task schedule.

If the schedule uses a task that is not used by any other schedule, a confirmation prompt will ask whether you want to delete the schedule and the task. Reply yes to delete both, or no to delete only the schedule.

When a volume is deleted, its schedules and tasks are also deleted.

Level   Manage

Syntax   `delete schedule`
           `[prompt yes|no]`
           `schedule`

Parameters   `prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

`schedule`
The name of schedule to delete.

Example   Delete schedule `Sched1`, whose task is used by another schedule:

```
# delete schedule Sched1
```

Delete schedule `Sched2`, whose task is not used by another schedule:

```
# delete schedule Sched2
This schedule uses a task (Task1) that is not used by any other schedule.
Do you want to delete this task, also? (y/n) no
```

See also
- create schedule
- show schedules

delete snapshot

Description Deletes specified snapshots. All data uniquely associated with the snapshot is deleted and associated space in the snap pool is freed for use. The snapshot's schedules are also deleted.

Level Manage

Syntax
```
delete snapshot
    [delete-priority standard-snapshot|volume-copy-snapshot|replication-snapshot
    |replicating-snapshot|common-sync-point-snapshot|only-sync-point-snapshot
    |queued-snapshot]
    [force]
    snapshots
```

Parameters
```
delete-priority standard-snapshot|volume-copy-snapshot
    |replication-snapshot|replicating-snapshot|common-sync-point-snapshot
    |only-sync-point-snapshot|queued-snapshot
```
Optional. Priority of snapshots that can be deleted. If the specified priority is less than the snapshot's priority, deletion is prevented. This is intended to protect against accidentally deleting high-priority snapshots. You must specify this parameter or the `force` parameter, but not both.

```
force
```
Optional. Overrides priority protection and forces the specified snapshot to be deleted. You must specify this parameter or the `delete-priority` parameter, but not both.

```
snapshots
```
Names or serial numbers of the snapshots to delete. For volume syntax, see Command syntax on page 18.

Example Delete snapshot `SS1`, which is being used in a replication operation:

```
# delete snapshot delete-priority replicating-snapshot SS1
```

Force deletion of snapshot `SS2`:

```
# delete snapshot force SS2
```

Delete three standard snapshots:

```
# delete snapshot s1,s2,s3
```

See also
- delete all-snapshots
- delete snapshot-write-data
- show snapshots

delete snapshot-write-data

Description Deletes data written to a standard snapshot after it was created. Deleting this modified data reverts the snapshot to the state when it was first taken. This command is not allowed for a replication snapshot.

Level Manage

Syntax `delete snapshot-write-data` *snapshot*

Parameters *snapshot*
Name or serial number of the snapshot from which to delete modified data. For volume syntax, see Command syntax on page 18.

Example Delete only modified data from snapshot `SS1`:

`# delete snapshot-write-data SS1`

See also • delete snapshot
• show snapshots

delete spares

Description Changes specified spares back into available disks.

Level Diagnostic

Syntax `delete spares disks disks`

Parameters `disks disks`
The IDs of the spares to delete. For disk syntax, see Command syntax on page 18.

Example Delete spares 1.23 and 1.24:

`# delete spares disks 1.23,1.24`

See also
- add spares
- show disks

delete task

Description   Deletes a task. If the task is scheduled, a confirmation prompt will ask whether you want to delete the task and its schedules. Reply yes to delete both, or no to cancel the command.

Level   Manage

Syntax   `delete task`
      `[prompt yes|no]`
       `task`

Parameters   `prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

`task`
The name of the task to delete.

Example   Delete unscheduled task `Task1`:

```
# delete task Task1
```

Delete scheduled task `Task2`:

```
# delete task Task2
Task Task2 has 2 schedules associated with it that will also be deleted.
Do you want to continue? (y/n) no
```

See also
- delete schedule
- show schedules
- show tasks

delete user

Description Deletes a user account. You can delete any user that has been created and the default users `monitor` and `ftp`.

Level Manage

Syntax delete user
      [noprompt]
       name

Parameters noprompt
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

name
The user to delete. Names are case sensitive.

Example Delete user `jsmith`:

```
# delete user jsmith
Are you sure you want to delete user jsmith? (y/n) yes
```

Delete user `Kim` and suppress the confirmation prompt:

```
# delete user noprompt Kim
```

See also • create user
• show users

delete volume-groups

Description Deletes specified volume groups and optionally all volumes in those groups.

Before using the option to delete all the volumes in the groups, ensure that the volumes are unmapped. If any volume is mapped, the command will fail and no changes will be made.

Level Manage

Syntax `delete volume-groups`
`    [delete-volumes]`
`    volume-groups|all`

Parameters `delete-volumes`
Optional. Specifies to delete all volumes in the groups. If the volumes have scheduled tasks, the schedules and tasks will be deleted. If this parameter is omitted, the volume groups will be deleted but their volumes will not be deleted.

`volume-groups|all`
Specifies a comma-separated list of the names of volume groups to delete, or to delete all volume groups.

Example Delete volume groups `VGroup1` and `VGroup2` but not the volumes in those groups:

`# delete volume-groups VGroup1,VGroup2`

Delete all volume groups and the volumes in those groups:

`# delete volume-groups delete-volumes all`

See also • show maps
• show volume-groups

delete volumes

Description Deletes specified volumes. If the volumes have scheduled tasks, the schedules and tasks will be deleted.

⚠ **CAUTION:** Deleting a volume will delete all data in that volume.

Level Manage

Syntax `delete volumes` *volumes*

Parameters *volumes*
Names or serial numbers of the volumes to delete. For volume syntax, see Command syntax on page 18.

Example Delete volumes `V1` and `V2`:

`# delete volumes V1,V2`

See also • create volume
• show volumes

dequarantine

Description Removes a storage-pool component from quarantine. **For use by or with direction from technical support.**

⚠ **CAUTION:**   If a component is removed from quarantine and does not have enough disks to continue operation, its status will change to `OFFL` and its data cannot be recovered. To continue operation, a RAID-6 component can have only one or two inaccessible disks

---

The system will automatically quarantine a component having a fault-tolerant RAID level if one or more of its disks becomes inaccessible, or to prevent invalid ("stale") data that may exist in the controller from being written to the component. Quarantine will not occur if a known failed disk becomes inaccessible or if a disk becomes inaccessible after failover or recovery. If quarantine occurs because of an inaccessible disk, event 172 is logged. If quarantine occurs to prevent writing invalid data, event 485 is logged.

Examples of when quarantine can occur are:

- At system power-up, a component has fewer disks online than at the previous power-up. This may happen because a disk is slow to spin up or because an enclosure is not powered up. The component will be automatically dequarantined if the inaccessible disks come online and the component status becomes `FTOL` (fault tolerant and online), or if after 60 seconds the component status is `QTCR` or `QTDN`.
- During system operation, a component loses redundancy plus one more disk; for example, three disks are inaccessible in a RAID-6 component or two disks are inaccessible in a RAID-1 component. The component will be automatically dequarantined if after 60 seconds its status is `FTOL`, `FTDN`, or `CRIT`.

Quarantine isolates the component from host access and prevents the system from changing the component status to `OFFL` (offline). The number of inaccessible disks determines the quarantine status; from least to most severe:

- `QTDN` (quarantined with a down disk): The RAID-6 component has one inaccessible disk. The component is fault tolerant but degraded.
- `QTCR` (quarantined critical): The component is critical with at least one inaccessible disk. For example, two disks are inaccessible in a RAID-6 storage-pool component or one disk is inaccessible for other fault-tolerant RAID levels.
- `QTOF` (quarantined offline): The component is offline with multiple inaccessible disks causing user data to be incomplete.

When a component is quarantined, its disks become write-locked, all volumes in its storage pool become inaccessible, and they are not available to hosts until the component is dequarantined. Depending on the operation, the length of the outage, and the settings associated with the operation, the operation may automatically resume when the component is dequarantined or may require manual intervention. A component can remain quarantined indefinitely without risk of data loss.

A component is dequarantined when it is brought back online, which can occur in three ways:

- If the inaccessible disks come online, making the component FTOL, the component is automatically dequarantined.
- If after 60 seconds from being quarantined the component is QTCR or QTDN, the component is automatically dequarantined. The inaccessible disks are marked as failed and the component status changes to CRIT (critical) or FTDN (fault tolerant with a down disk). If the inaccessible disks later come online, they are marked as LEFTOVR (leftover).
- The dequarantine command is used to manually dequarantine the component. If the inaccessible disks later come online, they are marked as LEFTOVR (leftover). If event 485 was logged, use the dequarantine command only as specified by the event's recommended-action text to avoid data corruption or loss.

A quarantined component can be fully recovered if the inaccessible disks are restored. Make sure that all disks are properly seated, that no disks have been inadvertently removed, and that no cables have been unplugged. Sometimes not all disks in the component power up. Check that all enclosures have restarted after a power failure. If these problems are found and then fixed, the component recovers and no data is lost.

If the inaccessible disks cannot be restored (for example, they failed), and the component's status is FTDN or CRIT, and compatible spares are available to replace the inaccessible disks, reconstruction will automatically begin.

If a replacement disk (reconstruct target) is inaccessible at power up, the component becomes quarantined; when the disk is found, the component is dequarantined and reconstruction starts. If reconstruction was in process, it continues where it left off.

📝 **NOTE:** The only command you may on a quarantined component is dequarantine.

Level Diagnostic

Syntax dequarantine storage-pool-component *storage-pool-component*

Parameters storage-pool-component *storage-pool-component*
The name or serial number of the storage-pool component to remove from quarantine. For storage-pool-component syntax, see Command syntax on page 18.

Example After determining that component A_01_01 is quarantined, remove it from quarantine and re-check its status:

```
# show storage-pool-components
Name      ... Status ...
----------------------
A_01_01  ... QTDN   ... (Quarantined with a down disk)
----------------------


# dequarantine storage-pool-component A_01_01

# show storage-pool-components
Name      ... Status ...
----------------------
A_01_01  ... FTDN   ... (Fault tolerant with a down disk)
----------------------
```

See also • show storage-pool-components

exit

Description Log off and exit the CLI session.

Level Monitor

Syntax `exit`

expand volume

Description Expands a standard volume by a specific size. A volume that has snapshots cannot be expanded. Before you can expand a volume that has snapshots, you must delete all of the snapshots.

Level Manage

Syntax
```
expand volume
      size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]
      volume
```

Parameters `size size[B|KB|MB|GB|TB|KiB|MiB|GiB|TiB]`
Specifies the amount of space to add to the volume, using the current base. To see the current base, use the show cli-parameters command. The unit can be specified as follows:

- If base 2 is in use: `B` (bytes), `KiB` (kibibytes), `MiB` (mebibytes), or `GiB` (gibibytes)
- If base 10 is in use: `B` (bytes), `KB` (kilobytes), `MB` (megabytes), or `GB` (gigabytes)
- If no unit is specified, the unit is 512-byte blocks.

If overcommit is disabled, expansion is restricted to the space available in the storage pool that contains the volume. If overcommit is enabled, the volume size can exceed the physical capacity of the storage pool. The maximum volume size is 140 TB (128 TiB). To see whether overcommit is enabled, use the show system-parameters command.

`volume`
The name or serial number of the volume to expand. For volume syntax, see Command syntax on page 18.

Example Expand standard volume `Vol0004`, which has no snapshots, by 100 GB:

```
# expand volume size 100GB Vol0004
```

See also • show volumes

export snapshot

Description Exports a replication snapshot in the local system to a new standard snapshot. The standard snapshot will reside in the same snap pool, take a snapshot license, and be independent of the replication snapshot. The standard snapshot can used like any other standard snapshot, and changes to it will not affect the replication snapshot.

The standard snapshot is subject to the snap pool's deletion policies. If the snap pool reaches its critical threshold, the snapshot may be deleted, even if it is mapped. If you want to preserve the snapshot's data, you can create a standard volume from the snapshot; see the copy volume command.

The export command will not succeed if the replication snapshot is in a remote system or if the resulting snapshot would exceed license limits.

Level Manage

Syntax
```
export snapshot
    name standard-snapshot
    [drm-snapshot yes|no]
    [set replication-set]
    replication-snapshot
```

Parameters name *standard-snapshot*
A name for the resulting snapshot.

drm-snapshot yes|no
Optional. Specifies whether to export a snapshot that was taken for remote replication.

set *replication-set*
Optional. The name or serial number of the replication set.

replication-snapshot
The name or serial number of the replication snapshot to export. If the name is not unique across replication sets, specify the set parameter.

Example Export local replication snapshot RepSnap1 to standard snapshot Snap1:

```
# export snapshot name Snap1 RepSnap1
```

See also • show replication-sets
• show snapshots

fail

Description	Forces the partner controller module to crash for a non-maskable interrupt. **For use by or with direction from technical support.**
	The command causes the crash by issuing an internal kill command to the Storage Controller in the partner controller module. This might be helpful to gather debug information that is only available via a crash dump.
Level	Diagnostic
Syntax	`fail controller a|b`
Parameters	`controller a|b` Specifies whether to kill controller A or B. You cannot kill the controller on which the command is issued.
Example	From controller A, fail controller B:
	`# fail controller b`
See also	• unfail

help

Description Shows brief help for all available commands or full help for a specific command. For information about using command shortcuts, see Command completion, editing, and history on page 19.

Level Monitor

Syntax To view brief descriptions of all commands that are available to the user level you logged in as, enter:

```
help
```

To view full help for a specific command, enter:

```
help [command-name]
```

To view information about the syntax for specifying parameters, disks, and so forth, enter:

```
help syntax
```

Example Show brief help for all available commands:

```
# help
```

Show full help for the show cli-parameters command:

```
# help show cli-parameters
```

load license

Description   Used by the WBI to install a license file to enable licensed features of the storage system.

See also   • show license

map volume

Description Maps specified volumes using settings that override the volumes' default mapping.

When a volume is created, if no mapping settings are specified the volume is not mapped; otherwise, those settings become its default mapping, which specifies the controller host ports and access level that all connected initiators have to the volume, and the LUN presented to all initiators to identify the volume. The default mapping's LUN is known as the volume's *default LUN*.

The `map volume` command creates mappings with different settings for different initiators. Optionally, you can specify the LUN, ports, and access level for a mapping. A mapping can make a volume accessible to initiators, or inaccessible to initiators (known as *masking*). For example, assume a volume's default mapping allows read-only access using LUN 5. You can give one initiator read-write access using LUN 6, and you can give a second initiator no access to the volume.

> 📝 **NOTE:** You cannot map a replication set's secondary volume.

> 📝 **NOTE:** When mapping a volume to an initiator that is using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount the volume and will report an error such as "unknown partition table."

Level Manage

Syntax `map volume`
```
[access read-write|rw|read-only|ro|no-access]
[mapping ports.LUN]
[lun LUN]
[ports ports]
[host hosts]
[initiator initiators|hosts|host-groups]
volumes|volume-groups
```

Parameters `access read-write|rw|read-only|ro|no-access`
Optional. The access permission for initiators to the volumes: read-write (rw), read-only (ro), or no-access. If the access parameter is specified as read-write or read-only, either the `mapping` parameter or the `lun` parameter must be specified. For an explicit mapping, `no-access` causes the volume to be masked from specified initiators. If the `access` parameter is omitted, access is set to `read-write`.

`mapping ports.LUN`
Optional. The ports and LUN to use for the mappings; any unspecified ports become unmapped. The `mapping` parameter is ignored if `access` is set to `no-access`.

`lun LUN`
Optional. The LUN to use for the mapping. If a single volume and multiple initiators are specified, the same LUN is used for each initiator. If multiple volumes and a single initiator are specified, the LUN will increment for the second and subsequent volumes. If multiple volumes and initiators are specified, each initiator will have the same LUN for the first volume, the next LUN for the second volume, and so on. The `lun` parameter is ignored if `access` is set to `no-access`. If the `lun` parameter is omitted, the default LUN is presented. Do not use the `lun` parameter with the `mapping` parameter.

`ports ports`
Optional. The controller host ports to use for the mapping; any unspecified ports become unmapped. If the `ports` parameter is specified, the `lun` parameter must also be specified. The `ports` parameter is ignored if `access` is set to `no-access`. If the `ports` parameter is omitted, all ports are mapped. For port syntax, see Command syntax on page 18. Do not use the `ports` parameter with the `mapping` parameter.

`host` *hosts*
Optional. This parameter is superseded by the initiator parameter. For backward compatibility with legacy scripts, the keyword `host` is aliased to the keyword `initiator`.

`initiator` *initiators|hosts|host-groups*
Optional. Specifies initiators, hosts, or host groups to which to map the volumes. If the initiator parameter is specified, both the lun and ports parameters must be specified. If the initiator parameter is omitted, the mapping applies to all initiators that are not explicitly mapped. For initiator, host, and host-group syntax, see Command syntax on page 18.

volumes|volume-groups
The names or serial numbers of the volumes or volume groups to map. For volume or volume-group syntax, see Command syntax on page 18.

Example Map volume `v2` with read-only access for initiator `Init1`, using port 1 in each controller and LUN 100:

```
# map volume access ro ports 1 lun 100 initiator Init1 v2
```

Map volumes `v2` and `v3` with read-write access for host group `HostGroup2`, using port 1 in each controller and LUN 101:

```
# map volume access rw ports 1 lun 101 initiator HostGroup2.*.* v2,v3
```

Mask volume `v4` from hosts `Host1` and `Host3`:

```
# map volume v4 access no-access initiator Host1.*,Host3.*
```

Map volumes `v1` and `v2` to initiators `Init1` and `Init2`, using port 1 in each controller and starting with LUN 6:

```
# map volume ports 1 lun 6 initiator Init1,Init2 v1,v2
```

See also
- show host-groups
- show hosts
- show initiators
- show maps
- show ports
- show volume-groups
- show volumes
- unmap volume

meta

Description   In XML API format only, shows all property metadata for objects. This includes data not shown in `brief` mode.

Level   Monitor

Syntax   `meta` *basetypes*

Parameters   *basetypes*
A basetype or a list of basetypes separated by commas (with no spaces) to specify the objects for which to show metadata. The basetypes are:

```
advanced-settings-table            remote-links
attribute-priorities               remote-system
auto-write-through-trigger         replicate-volume-tasks
cache-parameter                    replication-image
cache-settings                     replication-image-params
cli-parameters                     replication-set
compact-flash                      replication-volume
configured-storage                 replication-volume-summary
controller-cache-parameters        reset-snapshot-tasks
controllers                        sas-port
controller-statistics              sas-status-controller-a
copy-volume-status                 sas-status-controller-b
debug-log-parameters               schedules
disk-hist-statistics               security-communications-protocols
disk-statistics                    sensors
drive-parameters                   ses
drives                             shutdown-status
drive summary                      snap-pools
email-parameters                   snap-shot-parameters
enclosure-components               snapshots
enclosure-fru                      snapshot-with-retention-tasks
enclosure-list                     snap-tasks
enclosures                         snmp-parameters
enclosure-sku                      snmp-users
events                             spares
expander-ports                     status
fan                                storage-pool-components
fc-port                            storage-pools
host                               storage-pool-status
host-group                         storage-status
host-group-view                    storage-tier-status
host-group-view-mappings           system
host-port-statistics               system-parameters-table
hosts                              tasks
hosts-view                         tier-hist-statistics
host-view                          tier-summary
host-view-mappings                 time-settings-table
initiator                          unconfigured-disks
initiator-view                     unconfigured-storage
initiator-view-mappings            unhealthy-component
inquiry                            unwritable-cache
license                            users
log-header-table                   versions
managed-logs                       virtual-disks
master-volumes                     virtual-disk-summary
network-parameters                 volume-copy-status
ntp-status                         volume-copy-tasks
paged-storage-statistics           volume-groups
policy-threshold                   volume-group-view
pool-hist-statistics               volume-group-view-mappings
pool-summary                       volume-group-view-mappings-luns
port                               volume-names
power-supplies                     volume-reservations
readcache-hist-statistics          volumes
redundancy                         volume-statistics
refresh_counters                   volume-view
remote-addresses                   volume-view-mappings
```

Example Show all metadata for objects returned by the show disks command:

```
# meta drives
```

See also • set cli-parameters

ping

Description Tests communication with a remote host. The remote host is specified by IP address. Ping sends ICMP echo response packets and waits for replies.

Level Monitor

Syntax `ping`
    host-address
    [*count*]

Parameters *host-address*
The remote host's IP address in dotted decimal form.

count
Optional. The number of packets to send. The default is 4 packets. Use a small count because the command cannot be interrupted.

Example Send two packets to the remote computer at 10.134.50.6:

`# ping 10.134.50.6 2`

release volume

Description Clears host registrations and releases persistent reservations for all or specified volumes. Normally, reservations placed on volumes by hosts accessing those volumes can be released by host software. This command should be used only when the system is in an abnormal state, perhaps due to a configuration problem, and you need to remove all reservations for specified volumes and return them to a "clean" state.

⚠ **CAUTION:** Releasing reservations for volumes may allow unintended access to those volumes by other hosts, which may result in data corruption. Before issuing this command, quiesce all hosts that have visibility to the volumes whose reservations will be released.

Level Manage

Syntax `release volume all|`*`volumes`*

Parameters `all|`*`volumes`*
Specifies all volumes, or the names or serial numbers of specific volumes. For volume syntax, Command syntax on page 18.

Example Release reservations for a specific volume:

`# release volume v2`

See also • show volume-reservations
• show volumes

remote

Description Runs a command on a remote system that is associated with the local system. If the command cannot connect to remote controller A, it tries to connect to remote controller B; if unsuccessful, the remote command is not run. Output is displayed in console or XML API format depending on the local system's setting.

Level Monitor

Syntax `remote`
    system
    command

Parameters *system*
The name or network-port IP address of the remote system.

command
The full name of any CLI command that is valid for the remote user's access level.

Example Run the `show system` command on remote system `System2`:

`# remote System2 show system`

See also • show remote-systems

remove host-group-members

Description Removes specified hosts from a host group. You cannot remove all hosts from a group; at least one must remain. The hosts are ungrouped but not deleted.

Level Manage

Syntax `remove host-group-members`
      `hosts` *hosts*
    host-group-name

Parameters `hosts` *hosts*
A comma-separated list of the names of hosts to remove from the host group.

host-group-name
The name of the host group.

Example Remove two hosts from a group that contains three hosts:

`# remove host-group-members hosts Host2,Host3 HostGroup1`

See also • delete host-groups
    • show host-groups
    • show hosts

remove host-members

Description Removes specified initiators from a host. You cannot remove all initiators from a host; at least one must remain. The initiators are ungrouped but not deleted.

Level Manage

Syntax
```
remove host-members
     initiators initiators
     host-name
```

Parameters initiators *initiators*
A comma-separated list of the names or serial numbers of initiators to remove from the host.

host-name
The name of the host.

Example Remove two initiators from a group that contains three initiators:

```
# remove host-members initiators FC-init2,FC-init3 FC-host1
```

See also • delete hosts
• show hosts
• show initiators

remove replication-volume

Description Removes a secondary volume from a replication set. Any replication snapshots associated with that volume are converted to standard snapshots, regardless of the number of snapshots allowed by the system's license.

You must run this command on the primary system. You cannot remove the primary volume.

Level Manage

Syntax
```
remove replication-volume
     [nowait]
     [primary-volume volume]
     [set replication-set]
     replication-volume
```

Parameters `nowait`
Optional. Removing a volume from a replication set can take the Storage Controller several minutes to complete. This parameter allows that processing to continue in the background so the Management Controller can process other commands.

`primary-volume volume`
Optional; use only if the replication set has a primary-volume conflict. Name or serial number of the primary volume.

`set replication-set`
Optional. The name or serial number of the replication set.

`replication-volume`
The name or serial number of the secondary volume to remove. If the name is not unique within the replication set, the volume that is not the primary volume is removed. If the name is not unique across replication sets, specify the `set` parameter.

Example Remove secondary volume `rData` from a replication set:

```
# remove replication-volume rData
```

See also
• show replication-sets
• show replication-volumes

remove storage

Description Removes storage elements from the system and frees their disks for removal or reuse. You can remove all storage, a specific enclosure, a specific tier in an enclosure, or all spares in a specific enclosure. Each option has a confirmation prompt.

When using the option to remove all storage, storage pools and all volumes and the data that they contain will be deleted.

When using the options to remove a specific enclosure or a specific tier in an enclosure, the volumes and data in the corresponding storage-pool components will be migrated to other storage-pool components in the same pool. If there is insufficient space available elsewhere in the pool, the command will fail and before it can proceed you must either add storage or free existing storage by deleting snapshots and volumes. Assuming space is available for the migration, the process can take hours depending on the amount of data in those disks; you will be prompted to monitor the disks to see when they become available. There are several ways to know when the data has been successfully migrated:

- You can periodically run the show events command and look for event 470, which marks the successful removal of a storage-pool component.
- You can periodically run the show storage command to see when the applicable components are no longer shown in its output.
- You can run the show storage-status command to see the progress (percent complete) of the REMV job for each applicable component when the removal process is in progress.
- You can periodically run the show disks command to see when the Usage value for each applicable disk becomes AVAIL (available).

Level Manage

Syntax
```
remove storage
    [all]
    [enclosure enclosure-ID]
    [tier performance|standard|archive|readcache]
    [spares]
    [noprompt]
```

Parameters all
Optional. Removes both storage pools and all disks (including spares) from available storage.

△ **CAUTION:** This parameter will delete all data in all volumes in the system.

---

enclosure enclosure-ID
Optional. Removes the specified enclosure's disks from available storage.

tier performance|standard|archive|readcache
Optional. Removes the specified tier in an enclosure specified by the enclosure parameter from available storage. The tier parameter will remove the tier from both storage pools (if applicable) and will not affect other disks in that enclosure.

spares
Optional. Removes all spares in an enclosure specified by the enclosure parameter, or in all enclosures. The spares parameter affects only spares and will not affect any disks that are currently operating in a storage pool.

noprompt
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

Example Remove all storage:

```
# remove storage all
Volumes will be automatically deleted and data will be lost. Do you want to
continue? (yes/no) yes
```

Remove enclosure 2 from available storage:

```
# remove storage enclosure 2
Volume data will be automatically migrated elsewhere in the storage pool.
Depending on the amount of storage consumed, this process can be very
time-consuming and all disks should remain in place until the migration is
complete. Do you want to continue? (y/n) yes
```

Remove the standard tier in enclosure 3 from available storage:

```
# remove storage enclosure 3 tier standard
Volume data will be automatically migrated elsewhere in the storage pool.
Depending on the amount of storage consumed, this process can be very
time-consuming and all disks should remain in place until the migration is
complete. Do you want to continue? (y/n) yes
```

Remove spares in enclosure 1:

```
# remove storage enclosure 1 spares
Removing spares from the system is not recommended and can reduce the overall
availability and performance of the storage system. Do you want to continue?
(y/n) yes
```

See also • show enclosures
             • show storage-pools

remove volume-group-members

Description Removes volumes from a volume group. You cannot remove all volumes from a group; at least one must remain. The volumes are ungrouped but not deleted.

Level Manage

Syntax
```
remove volume-group-members
      volumes volume-IDs
    volume-group-name
```

Parameters volumes *volume-IDs*
A comma-separated list of the names or serial numbers of volumes to remove from the volume group.

volume-group-name
The name of the volume group.

Example Remove volumes `Vol0002` and `Vol0003` from volume group `VolumeGroup1`:

```
# remove volume-group-members volumes Vol0002,Vol0003 VolumeGroup1
```

See also
- delete volume-groups
- show volume-groups
- show volumes

replicate snapshot

Description Initiates a replication operation using an existing snapshot as the data source. This command replicates the specified external standard snapshot.

If the specified snapshot has not already been replicated on the replication volume, each replication volume in the replication set is requested to replicate the snapshot data. Only snapshot preserved data is replicated; snapshot modified data is not replicated.

If you instead want to create and replicate a snapshot as a single task, use the replicate volume command.

Level Manage

Syntax
```
replicate snapshot
     [name replication-snapshot]
     [set replication-set]
     external-snapshot
```

Parameters name *replication-snapshot*
Optional. A name for the resulting replication snapshot.

set *replication-set*
Optional. The name or serial number of the replication set.

external-snapshot
The name or serial number of the external snapshot to use as the data source. If the name is not unique across replication sets, specify the set parameter.

Example Replicate external snapshot Snap1 and name the resulting replication snapshot RepSnap1:

```
# replicate snapshot name RepSnap1 Snap1
```

See also • show replication-sets
• show snapshots

replicate volume

Description Creates a replication snapshot of the specified volume and initiates a replication operation. This command can follow a create replication-set command that did not specify to initiate replication.

If the snapshot you want to replicate already exists, use the replicate snapshot command instead.

📝 **NOTE:** If replication requests are sent to a secondary system whose temporary replication license has expired, the requests are queued but are not processed, and the secondary system reports event 472. If this condition occurs, check for this event in the event log, event-notification emails, and SNMP traps. To continue using replication, purchase a permanent replication license.

Level Manage

Syntax
```
replicate volume
    [set replication-set]
    snapshot replication-snapshot
    [preserve-snapshot external-snapshot]
    volume
```

Parameters set *replication-set*
Optional. The name or serial number of the replication set.

snapshot *replication-snapshot*
The name for the new replication snapshot.

preserve-snapshot *external-snapshot*
Optional. The name of the external snapshot to use as the source for this replication. This preserves the snapshot that is being used to replicate the volume as an external snapshot; otherwise, the snapshot is converted to a replication snapshot.

volume
The name or serial number of the primary volume to replicate. If the name is not unique across replication sets, specify the set parameter.

Example Replicate volume V1 and name the replication snapshot RepSnap1:

```
# replicate volume snapshot RepSnap1 V1
```

See also • show replication-sets
• show replication-volumes

rescan

Description This command forces rediscovery of connected disks and enclosures. If both Storage Controllers are online and able to communicate with both expansion modules in each connected enclosure, this command also reassigns enclosure IDs based on controller A's enclosure cabling order. A manual rescan may be needed after system power-up to display enclosures in the proper order.

A manual rescan is not required to detect when disks are inserted or removed; the controllers do this automatically. When disks are inserted they are detected after a short delay, which allows the disks to spin up.

When you perform a manual rescan, it temporarily pauses all I/O processes, then resumes normal operation.

Level Manage

Syntax `rescan`

Example Scan for device changes and re-evaluate enclosure IDs:

`# rescan`

reset all-statistics

Description Resets performance statistics for both controllers. You can specify either to reset all live statistics to zero, or to reset (clear) all historical performance statistics for all disks. If you reset historical statistics, an event will be logged and new data samples will continue to be stored every quarter hour.

Level Manage

Syntax
```
reset all-statistics
      [historical]
      [prompt yes|no]
```

Parameters `historical`
Optional. Specifies to reset historical statistics instead of live statistics. If this parameter is omitted, the command will reset live statistics instead of historical statistics.

`prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt that will appear if the historical parameter is specified:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If the historical parameter is specified and the prompt parameter is omitted, you must manually reply to the prompt. If the historical parameter is omitted, the prompt parameter has no effect. There is no confirmation prompt for live statistics.

Example Reset all live statistics for both controllers:

```
# reset all-statistics
```

Reset all historical disk-performance statistics for both controllers:

```
# reset all-statistics historical
This command will reset all historical disk statistics.
Do you want to reset? (y/n) yes
```

See also
- reset controller-statistics
- reset disk-error-statistics
- reset disk-statistics
- reset host-port-statistics
- reset storage-statistics
- reset volume-statistics
- show controller-statistics

reset controller-statistics

Description Resets performance statistics for controllers.

This command resets all controller statistics except Power On Time. To reset this, restart or power cycle a controller.

Level Manage

Syntax `reset controller-statistics [a|b|both]`

Parameters `a|b|both`
Optional. Specifies whether to reset statistics for controller A, B, or both. If this parameter is omitted, statistics are reset for both controllers.

Example Reset statistics for both controllers:

`# reset controller-statistics`

See also
- reset all-statistics
- reset disk-error-statistics
- reset disk-statistics
- reset host-port-statistics
- reset storage-statistics
- reset volume-statistics
- show controller-statistics

reset disk-error-statistics

Description Resets error statistics for all or specified disks. Statistics that are reset include:

- Number of SMART events recorded
- Number of I/O timeouts accessing the disk
- Number of times the disk did not respond
- Number of attempts by the controllers to spin up the disk
- Number of media errors (errors generated by the disk as specified by its manufacturer)
- Number of non-media errors (errors generated by the controllers or by the disk and not categorized as media errors)
- Number of block reassignments
- Number of bad blocks found

To reset other disk statistics, use the reset disk-statistics command.

Level Manage

Syntax `reset disk-error-statistics [disks]`

Parameters *disks*
Optional. IDs of the disks for which to reset statistics. For disk syntax, see Command syntax on page 18. If this parameter is omitted, statistics are reset for all disks.

Example Reset error statistics for disks 1.1 and 2.1:

`# reset disk-error-statistics 1.1,2.1`

See also
- reset all-statistics
- reset controller-statistics
- reset disk-statistics
- reset host-port-statistics
- reset storage-statistics
- reset volume-statistics
- show disk-statistics
- show disks

reset disk-statistics

Description Resets performance statistics for disks.

This command resets basic disk statistics but not disk error statistics. To reset these, use the reset disk-error-statistics command.

Level Manage

Syntax `reset disk-statistics`

Example Reset statistics for all disks:

```
# reset disk-statistics
```

See also
- reset all-statistics
- reset controller-statistics
- reset disk-error-statistics
- reset host-port-statistics
- reset storage-statistics
- reset volume-statistics
- show disk-statistics

reset host-link

Description Resets specified controller host ports (channels). **For use by or with direction from technical support.**

For an FC host port configured to use FC-AL (loop) topology, a loop initialization primitive (LIP) is issued. For SAS, resetting a host port issues a `COMINIT/COMRESET` sequence and might reset other ports.

Level Diagnostic

Syntax
```
reset host-link
    ports ports
```

Parameters port *ports*
A controller host port ID, a comma-separated list of IDs, a hyphenated range of IDs, or a combination of these. A port ID is a controller ID and port number, and is not case sensitive. Do not mix controller IDs in a range.

Example Reset the host link on port A1:

```
# reset host-link ports A1
```

See also • show ports

reset host-port-statistics

Description Resets performance statistics for controller host ports.

Level Manage

Syntax `reset host-port-statistics [ports ports]`

Parameters ports *ports*
Optional. The controller ID and port number of ports for which to reset statistics. For port syntax, see Command syntax on page 18. If this parameter is omitted, statistics are reset for all controller host ports.

Example Reset statistics for all controller host ports:

```
# reset host-port-statistics
```

See also • reset all-statistics
• reset controller-statistics
• reset disk-error-statistics
• reset disk-statistics
• reset storage-statistics
• reset volume-statistics
• show host-port-statistics
• show ports

reset snap-pool

Description Clears a specified snap pool and resets its size, thresholds, and policies to default values.

You can use this command to decrease the size of a snap pool that is too large. The command deletes all snapshots in the snap pool, sets the size of the snap pool to the minimum (about 5.3 GB), and sets all thresholds and policies to their default values.

⚠ **CAUTION:** Deleting snapshots will delete all data in those snapshots.

Level Manage

Syntax
```
reset snap-pool
     snap-pool-ID|storage-pool-ID
     [prompt yes|no]
```

Parameters *snap-pool-ID*|*storage-pool-ID*
The name or serial number of either the snap pool or its storage pool. Names are case sensitive.

`prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must reply to the prompt.

Example Reset the snap pool in storage pool A:

```
# reset snap-pool A
Resetting snap pool SNAP_A will delete all snapshots in the snap pool.
Do you want to continue? (y/n) yes
```

See also
- show storage-pools
- show volumes

reset snapshot

Description Deletes the data in a standard snapshot and resets it to the current data in the source volume. The snapshot's volume characteristics are not changed. The command prompts you to unmount the snapshot from hosts before performing the reset; leaving it mounted can cause data corruption. This command is not allowed for a replication snapshot.

⚠ **CAUTION:** All data represented by the snapshot as it exists prior to issuing this command is lost.

Level Manage

Syntax `reset snapshot`
     `[prompt yes|no]`
     volume

Parameters `prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

volume
The name or serial number of the snapshot to reset. For volume syntax, see Command syntax on page 18.

Example Reset snapshot `SS1`:

```
# reset snapshot SS1
You MUST unmount the snapshot from all hosts before resetting it.
Are you ready to continue? (y/n) yes
```

See also • show snapshots

reset storage-statistics

Description Resets live (not historical) performance statistics for storage pools.

This command resets basic disk statistics but not disk error statistics. To reset these, use the reset disk-error-statistics command.

Level Manage

Syntax `reset storage-statistics [`*`storage-pool-ID`*`]`

Parameters *`storage-pool-ID`*
Optional. The name or serial number of the storage pool for which to reset statistics. If this parameter is omitted, statistics are reset for all storage pools.

Example Reset live statistics for storage pool B:

```
# reset storage-statistics B
```

See also
- reset all-statistics
- reset controller-statistics
- reset disk-error-statistics
- reset host-port-statistics
- reset volume-statistics
- show disk-statistics

reset volume-statistics

Description Resets performance statistics for all or specified volumes.

Level Manage

Syntax `reset volume-statistics [`*`volumes`*`]`

Parameters *volumes*
Optional. The names or serial numbers of the volumes for which to reset statistics. For volume syntax, see Command syntax on page 18. If this parameter is omitted, statistics are reset for all volumes.

Example Reset statistics for volume `v1`:

```
# reset volume-statistics v1
```

See also
- reset all-statistics
- reset controller-statistics
- reset disk-error-statistics
- reset disk-statistics
- reset host-port-statistics
- reset storage-statistics
- show volume-statistics
- show volumes

restart

Description Restarts the Storage Controller or Management Controller in a controller module.

If you restart a Storage Controller, it attempts to shut down with a proper failover sequence, which includes stopping all I/O operations and flushing the write cache to disk, and then the controller restarts. The Management Controller is not restarted so it can provide status information to external interfaces.

If you restart a Management Controller, communication with it is lost until it successfully restarts. If the restart fails, the partner Management Controller remains active with full ownership of operations and configuration information.

⚠ **CAUTION:** If you restart both controller modules, all users will lose access to the system and its data until the restart is complete.

📝 **NOTE:** When a Storage Controller is restarted, live performance statistics that it recorded will be reset; historical performance statistics are not affected. In a dual-controller system, disk statistics may be reduced but will not be reset to zero, because disk statistics are summed between the two controllers. For more information, see help for commands that show statistics.

Level Manage

Syntax restart
    sc|mc
    [a|b|both]
    [noprompt]

Parameters sc|mc
The controller to restart:

- sc: Storage Controller
- mc: Management Controller

a|b|both
Optional. The controller module containing the controller to restart. If this parameter is omitted, the command affects the controller being accessed.

noprompt
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

Example Restart the Management Controller in controller A, which you are logged in to:

```
# restart mc a
During the restart process you will briefly lose communication with the
specified Management Controller(s).
Do you want to continue? (y/n) yes
```

From controller A, restart the Storage Controller in controller B:

```
# restart sc b
```

Restart both Storage Controllers:

```
# restart sc both
Restarting both controllers can cause a temporary loss of data availability.
Do you want to continue? (y/n) yes
```

See also • shutdown

restore defaults

Description Restores the default configuration to the controllers. **For use by or with direction from technical support.**

For details about which settings are restored, see Settings changed by restore defaults on page 291.

△ **CAUTION:** This command changes how the system operates and might require some reconfiguration to restore host access to volumes.

Level Manage

Syntax
```
restore defaults
     [noprompt]
     [prompt yes|no]
```

Parameters noprompt
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

prompt yes|no
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- yes: Allow the command to proceed.
- no: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

Example Restore the controllers' default configuration:

```
# restore defaults
WARNING: The configuration of the controller will be reset to default settings.
The Management Controller will restart once this is completed. Do you want to
continue? (y/n) yes
```

See also • restart

resume replication

Description Resumes a suspended replication operation for the specified secondary volume. This command must be issued on the system that owns the secondary volume.

Level Manage

Syntax
```
resume replication
    [set replication-set]
    replication-volume
```

Parameters `set replication-set`
Optional. The name or serial number of the replication set.

`replication-volume`
The name or serial number of the secondary volume. If the name is not unique across replication sets, specify the `set` parameter.

Example Resume replication of primary volume `V1` to secondary volume `rV1`:

```
# resume replication rV1
```

See also
- abort replication
- show replication-sets
- show replication-volumes

rollback volume

Description Rolls back the data in a standard volume to the data that exists in a specified snapshot of that volume. You can choose whether to include modified write data from the snapshot in the rollback. You must unmount the source volume from hosts before using this command. The command will prompt you to ensure the source volume is unmounted before proceeding.

⚠ **CAUTION:** All data that differs between the source volume and the snapshot is lost. Create a snapshot of the source volume as it currently exists before performing a rollback.

Level Manage

Syntax
```
rollback volume
    snapshot snapshot
    [modifiedsnapshot yes|no]
    [prompt yes|no]
    volume
```

Parameters `snapshot` *snapshot*
The name or serial number of the snapshot containing the data to roll back to. For volume syntax, see Command syntax on page 18.

`modifiedsnapshot yes|no`
Optional. Specifies whether to include or exclude modified write data from the snapshot in the rollback.

- `yes`: Include modified snapshot.
- `no`: Exclude modified snapshot data.

If this parameter is omitted, modified snapshot data is excluded.

`prompt yes|no`
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

volume
The name or serial number of the standard volume to roll back. For volume syntax, see Command syntax on page 18.

Example Roll back standard volume `V1` to snapshot `SS1`:

```
# rollback volume snapshot SS1 V1
You MUST unmount the volume and the snapshot from all hosts before starting a
rollback operation.
Are you ready to continue? (y/n) yes
```

See also
- show snapshots
- show volumes

scrub disks

Description   Analyzes specified disks to find and fix disk errors. This command acts on disks that are in a storage-pool component; it does not act on disks that are spares or have `LEFTOVR` status. This command will fix parity mismatches for RAID 6 and mirror mismatches for RAID 1.

A disk can be used while it is being scrubbed. To check the progress of a disk scrub (`DRSC`) job, use the show disks command.

When scrub is complete, event 207 is logged and specifies whether errors were found and whether user action is required.

Level   Manage

Syntax   `scrub disks disks`

Parameters   *disks*
The IDs of the disks to scrub. For disk syntax, see Command syntax on page 18.

Example   Start scrubbing disks 1.1 and 1.3:

```
# scrub disks 1.1,1.3
```

See also   • abort scrub
   • set advanced-settings
   • show disks

scrub storage-pool-components

Description Analyzes specified storage-pool components to find and fix disk errors. **For use by or with direction from technical support.**

This command acts on disks that are associated with a storage-pool component and are neither dedicated spares nor leftovers. This command will fix parity mismatches for RAID 6, mirror mismatches for RAID 1, and media errors for all RAID levels.

A scrub can last over an hour, depending on storage-pool component size, utility priority, and amount of I/O activity. However, a "foreground" scrub performed with this command is typically faster than a background scrub enabled with the set advanced-settings command. A storage-pool component can be used while it is being scrubbed. To check the progress of a scrub (VRSC) job, use the show storage-pool-components command.

When scrub is complete, event 207 is logged and specifies whether errors were found and whether user action is required.

Level Diagnostic

Syntax `scrub storage-pool-components` *component-names*

Parameters `storage-pool-component` *component-names*
The names of the storage-pool components to scrub.

Example Start scrubbing storage-pool component `A_01_01`:

`# scrub storage-pool-components A_01_01`

See also • abort scrub
  • set advanced-settings

set advanced-settings

Description Sets advanced system parameters.

Level Manage

Syntax
```
set advanced-settings
     [auto-write-back enabled|disabled|on|off]
     [background-disk-scrub enabled|disabled|on|off]
     [background-scrub enabled|disabled|on|off]
     [background-scrub-interval interval]
     [compact-flash-failure enabled|disabled|on|off]
     [controller-failure enabled|disabled|on|off]
     [emp-poll-rate rate]
     [fan-failure enabled|disabled|on|off]
     [host-cache-control enabled|disabled|on|off]
     [managed-logs enabled|disabled|on|off]
     [missing-lun-response notready|illegal]
     [partner-firmware-upgrade enabled|disabled|on|off]
     [partner-notify enabled|disabled|on|off]
     [power-supply-failure enabled|disabled|on|off]
     [smart enabled|disabled|on|off|detect-only]
     [super-cap-failure enabled|disabled|on|off]
     [sync-cache-mode immediate|flush]
     [temperature-exceeded enabled|disabled|on|off]
     [utility-priority low|medium|high]
     [spin-down enabled|disabled|on|off]
     [spin-down-delay delay]
```

Parameters `auto-write-back enabled|disabled|on|off`
Optional. Sets whether the cache mode will change from write-through to write-back when the trigger condition is cleared.

- `disabled` or `off`: Auto-write-back is disabled.
- `enabled` or `on`: Auto-write-back is enabled. This is the default.

`background-disk-scrub enabled|disabled|on|off`
Optional. Sets whether disks that are not in storage-pool components are automatically checked for disk defects to ensure system health. The interval between background disk scrub finishing and starting again is 24 hours.

- `disabled` or `off`: Background disk scrub is disabled. This is the default.
- `enabled` or `on`: Background disk scrub is enabled.

`background-scrub enabled|disabled|on|off`
Optional. Sets whether disks in storage-pool components are automatically checked for disk defects to ensure system health. The interval between background component scrub finishing and starting again is specified by the `background-scrub-interval` parameter.

- `disabled` or `off`: Background component scrub is disabled. This is the default.
- `enabled` or `on`: Background component scrub is enabled.

`background-scrub-interval interval`
Optional. Sets the interval in hours between background component scrub finishing and starting again, from 1–360 hours. The default is 24 hours.

`compact-flash-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when CompactFlash memory is not detected during POST (Power-On Self-Test), fails during POST, or fails during controller operation.

- `disabled` or `off`: The CompactFlash failure trigger is disabled.
- `enabled` or `on`: The CompactFlash failure trigger is enabled. This is the default.

`controller-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when a controller fails.

- `disabled` or `off`: The controller failure trigger is disabled. This is the default.
- `enabled` or `on`: The controller failure trigger is enabled.

`emp-poll-rate` *rate*
Optional. Sets the interval at which the storage system will poll each enclosure's Enclosure Management Processor (EMP) for status changes, from 5–3600 seconds. The default is 5 seconds.

`fan-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when a fan fails.

- `disabled` or `off`: The fan failure trigger is disabled. This is the default.
- `enabled` or `on`: The fan failure trigger is enabled.

`host-cache-control enabled|disabled|on|off`
Optional. Sets whether hosts are allowed to use the SCSI MODE SELECT command to change the storage system's write-back cache setting.

- `disabled` or `off`: Host control of caching is disabled. This is the default.
- `enabled` or `on`: Host control of caching is enabled.

`managed-logs enabled|disabled|on|off`
Optional. Enables or disables the managed logs feature, which allows log files to be transferred from the storage system to a log collection system to avoid losing diagnostic data.

- `disabled` or `off`: The managed logs feature is disabled. This is the default.
- `enabled` or `on`: The managed logs feature is enabled.

`missing-lun-response notready|illegal`
Optional. Sets whether host drivers may probe for LUNs until the host drivers reach the LUN to which they have access.

- `notready`: Sends a reply that there is a LUN where a gap has been created but that it's "not ready." Sense data returned is sensekey = 2, code = 4, qualifier = 3. This option is the default.
- `illegal`: Sends a reply that there is a LUN but that the request is "illegal." Sense data returned is sensekey = 5, code = 25h, qualifier = 0. If the system is used in a VMware environment, use this option.

`partner-firmware-upgrade enabled|disabled|on|off`
Optional. Sets whether partner firmware update (PFU) will monitor component firmware versions and automatically update them in the partner controller.

- `disabled` or `off`: PFU is disabled.
- `enabled` or `on`: PFU is enabled. This is the default.

`partner-notify enabled|disabled|on|off`
Optional. Sets whether to notify the partner controller that a trigger condition occurred. Enable this option to have the partner also change to write-through mode for better data protection. Disable this option to allow the partner continue using its current caching mode for better performance. The default is `disabled`.

- `disabled` or `off`: Notification is disabled. This is the default.
- `enabled` or `on`: Notification is enabled.

`power-supply-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy automatically changes to write-through when a power supply fails.

- `disabled` or `off`: The power-supply failure trigger is disabled. This is the default.
- `enabled` or `on`: The power-supply failure trigger is enabled.

`smart enabled|disabled|on|off|detect-only`
Optional. Enables or disables SMART (Self-Monitoring Analysis and Reporting Technology) monitoring for all disks in the storage system.

- `disabled` or `off`: Disables SMART for all disks in the system and for all disks added to the system.
- `enabled` or `on`: Enables SMART for all disks in the system and for all disks added to the system. This is the default.
- `detect-only`: Detects but does not change the SMART setting of each disk in the system, and for each new disk added to the system.

`super-cap-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when the supercapacitor that provides backup power for cache is not fully charged or fails.

- `disabled` or `off`: The super-capacitor failure trigger is disabled.
- `enabled` or `on`: The super-capacitor failure trigger is enabled. This is the default.

`sync-cache-mode immediate|flush`
Optional. Sets how the `SCSI SYNCHRONIZE CACHE` command is handled.

- `immediate`: Good status is returned immediately and cache content is unchanged. This option is the default.
- `flush`: Good status is returned only after all write-back data for the specified volume is flushed to disk.

`temperature-exceeded enabled|disabled|on|off`
Optional. Sets whether the system will shut down a controller when its temperature exceeds the critical operating range.

- `disabled` or `off`: The over-temperature trigger is disabled. This is the default.
- `enabled` or `on`: The over-temperature trigger is enabled.

`utility-priority low|medium|high`
Optional. Sets the priority at which data-redundancy utilities, such as verify and reconstruct for storage-pool components, run with respect to I/O operations competing for the system's processors. (This does not affect background scrub of storage-pool components, which always runs at "background" priority.)

- `high`: Utilities have higher priority than host I/O. Use when your highest priority is to return the system to a fully fault-tolerant state. This can cause heavy I/O to be slower than normal. This is the default.
- `medium`: Utility performance is balanced with host I/O performance.
- `low`: Utilities run at a slower rate with minimal effect on host I/O. Use when streaming data without interruption, such as for a web server, is more important than data redundancy.

```
spin-down enabled|disabled|on|off
```
Optional. Sets whether available disks and spares will spin down after a period of inactivity shown by the `spin-down-delay` parameter.

- `disabled` or `off`: Drive spin down for available disks and spares is disabled. This is the default.
- `enabled` or `on`: Drive spin down for available disks and spares is enabled.

```
spin-down-delay delay
```
Optional. Sets the period of inactivity after which available disks and spares will spin down. The default is 15 minutes. Setting the delay to 1–360 minutes will enable spin down; setting the delay to 0 will disable spin down.

Example Enable partner firmware upgrade:

```
# set advanced-settings partner-firmware-upgrade enabled
```

Enable managed logs:

```
# set advanced-settings managed-logs enabled
```

See also • scrub storage-pool-components
 • show advanced-settings

set auto-write-through-trigger

Description Sets the trigger conditions that cause the controller to change the cache policy from write-back to write-through. You can set multiple triggers. By default `super-cap-failure` and `auto-write-back` are enabled.

When the cache mode is changed, an event is logged.

Level Manage

Syntax
```
set auto-write-through-trigger
    [controller-failure enabled|disabled|on|off]
    [super-cap-failure enabled|disabled|on|off]
    [compact-flash-failure enabled|disabled|on|off]
    [power-supply-failure enabled|disabled|on|off]
    [fan-failure enabled|disabled|on|off]
    [temperature-exceeded enabled|disabled|on|off]
    [partner-notify enabled|disabled|on|off]
    [auto-write-back enabled|disabled|on|off]
```

Parameters `controller-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when a controller fails.

- `disabled` or `off`: The controller failure trigger is disabled. This is the default.
- `enabled` or `on`: The controller failure trigger is enabled.

`super-cap-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when the supercapacitor that provides backup power for cache is not fully charged or fails.

- `disabled` or `off`: The super-capacitor failure trigger is disabled.
- `enabled` or `on`: The super-capacitor failure trigger is enabled. This is the default.

`compact-flash-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when CompactFlash memory is not detected during POST (Power-On Self-Test), fails during POST, or fails during controller operation.

- `disabled` or `off`: The CompactFlash failure trigger is disabled.
- `enabled` or `on`: The CompactFlash failure trigger is enabled. This is the default.

`power-supply-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy automatically changes to write-through when a power supply fails.

- `disabled` or `off`: The power-supply failure trigger is disabled. This is the default.
- `enabled` or `on`: The power-supply failure trigger is enabled.

`fan-failure enabled|disabled|on|off`
Optional. Sets whether the cache policy will change from write-back to write-through when a fan fails.

- `disabled` or `off`: The fan failure trigger is disabled. This is the default.
- `enabled` or `on`: The fan failure trigger is enabled.

`temperature-exceeded enabled|disabled|on|off`
Optional. Sets whether the system will shut down a controller when its temperature exceeds the critical operating range.

- `disabled` or `off`: The over-temperature trigger is disabled. This is the default.
- `enabled` or `on`: The over-temperature trigger is enabled.

`partner-notify enabled|disabled|on|off`
Optional. Sets whether to notify the partner controller that a trigger condition occurred. Enable this option to have the partner also change to write-through mode for better data protection. Disable this option to allow the partner to continue using its current caching mode for better performance.

- `disabled` or `off`: Notification is disabled. This is the default.
- `enabled` or `on`: Notification is enabled.

`auto-write-back enabled|disabled|on|off`
Optional. Sets whether the cache mode will change from write-through to write-back when the trigger condition is cleared.

- `disabled` or `off`: Auto-write-back is disabled.
- `enabled` or `on`: Auto-write-back is enabled. This is the default.

Example Enable the controller-failure trigger and disable the partner-notification trigger:

```
# set auto-write-through-trigger controller-failure enabled partner-notify
disabled
```

See also • show auto-write-through-trigger
- show events

set cache-parameters

Description Sets a cache parameters for a specified volume.

> **NOTE:** Only change the read-ahead cache settings if you fully understand how the host operating system, application, and adapter move data so that you can adjust the settings accordingly. Be prepared to monitor system performance and adjust read-ahead size until you find the optimal size for your application.

Level Manage

Syntax
```
set cache-parameters
    [write-policy write-back|write-through|wb|wt]
    [optimization standard|no-mirror]
    [read-ahead-size disabled|default|maximum|64KB|128KB|256KB|512KB|1MB|2MB|4MB
    |8MB|16MB|32MB]
    volume
```

Parameters `write-policy write-back|write-through|wb|wt`
Optional. Sets the cache write policy, which determines when cached data is written to the disks. The ability to hold data in cache while it is being written to disk can increase storage device speed during sequential reads.

- `write-back` or `wb`: Write-back caching does not wait for data to be completely written to disk before signaling the host that the write is complete. This is the default and preferred setting for a fault-tolerant environment because it improves the performance of write operations and throughput.
- `write-through` or `wt`: Write-through caching significantly impacts performance by waiting for data to be completely written to disk before signaling the host that the write is complete. Use this setting only when operating in an environment with low or no fault tolerance.

You can configure the write policy to automatically change from write-back to write-through when certain environmental events occur, such as a fan failure. For details, see help for the set auto-write-through-trigger command.

`optimization standard|no-mirror`
Optional. Sets the cache optimization mode:

- `standard`: Optimizes cache for both sequential and random reads. This is the default.
- `no-mirror`
  When this mode is enabled, each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. ULP behavior is not affected, with the exception that during failover any write data in cache will be lost.

`read-ahead-size`

Optional. Controls the use and size of read-ahead cache. You can optimize a volume for sequential reads or streaming data by changing the amount of data read in advance after two back-to-back reads are made. Read ahead is triggered by two back-to-back accesses to consecutive logical block address (LBA) ranges. Read ahead can be forward (that is, increasing LBAs) or reverse (that is, decreasing LBAs). Increasing the read-ahead size can greatly improve performance for multiple sequential read streams. However, increasing read-ahead size will likely decrease random read performance.

- `disable`: Turns off read-ahead cache.
- `default`: Sets one chunk for the first access in a sequential read and one stripe for all subsequent accesses. The controllers treat RAID-1 storage-pool components internally as if they have a stripe size of 64 Kbyte, even though they are not striped. This setting works well for most applications.
- `maximum`: Allows the controller to dynamically calculate the maximum read-ahead cache size for the volume.
- `64 KB`, `128 KB`, `256 KB`, `512 KB`, `1 MB`, `2 MB`, `4 MB`, `8 MB`, `16 MB`, or `32 MB`: Sets a specific cache size.

volume

The name or serial number of the volume to change. For volume syntax, see Command syntax on page 18.

Example   Set the cache policy and cache optimization mode for volume `V1`:

```
# set cache-parameters optimization standard read-ahead-size maximum V1
```

See also • show cache-parameters
- show volumes

set cli-parameters

Description Sets options that control CLI behavior. If you are accessing the CLI through the network port, settings apply to the current CLI session only. If you are accessing the CLI through the enclosure's CLI port, settings persist across sessions.

The base, locale, precision, temperature scale, timeout, and units settings are read from the user's account, and can be overridden by using this command.

Level Manage

Syntax
```
set cli-parameters
    [base 2|10]
    [console|api|api-embed]
    [brief enabled|disabled|on|off]
    [locale English|en|Spanish|es|Japanese|ja|Chinese-simplified|zh-s]
    [pager enabled|disabled|on|off]
    [precision #]
    [storage-size-base 2|10]
    [storage-size-precision #]
    [storage-size-units auto|MB|GB|TB]
    [temperature-scale celsius|c|fahrenheit|f]
    [timeout #]
    [units auto|MB|GB|TB]
```

Parameters `base 2|10`
Optional. Sets the base for entry and display of storage-space sizes:

- `2`: Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
- `10`: Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is the default.

Operating systems usually show volume size in base 2. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2.

`console|api|api-embed`
Optional. Sets the output format:

- `console`: Supports interactive use of the CLI by displaying command output in easily readable format. This format automatically sizes fields according to content and adjusts content to window resizes. This is the default.
- `api`: Supports scripting by displaying command output in XML. All objects are displayed at the same level, related by COMP elements.
- `api-embed`: Alternate form of XML output which displays "child" objects embedded (indented) under "parent" objects.
- `ipa`: Alternate form of XML output for internal use only.
- `json`: Alternate data-interchange format for internal use only.

`brief enabled|disabled|on|off`
Optional.

- `enabled` or `on`: In XML output, shows a subset of attributes of object properties. The `name` and `type` attributes are always shown.
- `disabled` or `off`: In XML output, shows all attributes of object properties. This is the default.

`locale English|en|Spanish|es|Japanese|ja|Chinese-simplified|zh-s`
Optional. The display language. The default is English.

```
pager enabled|on|disabled|off
```
Optional.

- `enabled` or `on`: Halts output after each full screen to wait for keyboard input. This is the default.
- `disabled` or `off`: Output is not halted. When displaying output in XML API format, which is intended for scripting, disable paging.

```
precision #
```
Optional. Sets the number of decimal places (1–10) for display of storage-space sizes. The default is 1.

```
storage-size-base 2|10
```
Optional. Alias for `base`.

```
storage-size-precision #
```
Optional. Alias for `precision`.

```
storage-size-units auto|MB|GB|TB
```
Optional. Alias for `units`.

```
temperature-scale celsius|c|fahrenheit|f
```
Optional. Sets the scale for display of temperature values:

- `fahrenheit` or `f`: Temperatures are shown in degrees Fahrenheit.
- `celsius` or `c`: Temperatures are shown in degrees Celsius. This is the default.

```
timeout #
```
Optional. Sets the timeout value in seconds for the login session. Valid values are 120–43200 seconds (2–720 minutes). The default is 1800 seconds (30 minutes).

```
units auto|MB|GB|TB
```
Optional. Sets the unit for display of storage-space sizes:

- `auto`: Sizes are shown in units determined by the system. This is the default.
- `MB`: Sizes are shown in megabytes.
- `GB`: Sizes are shown in gigabytes.
- `TB`: Sizes are shown in terabytes.

Based on the `precision` setting, if a size is too small to meaningfully display in the selected unit, the system uses a smaller unit for that size. For example, if `units` is set to TB, `precision` is set to 1, and `base` is set to 10, the size 0.11709 TB is instead shown as 117.1 GB.

Example Set CLI parameters:

```
# set cli-parameters timeout 600 precision 2 units GB temperature-scale f
```

For scripting, display XML output in `api-embed` format and disable paging:

```
# set cli-parameters api-embed pager off
```

For scripting, display brief XML output in `api-embed` format and disable paging:

```
# set cli-parameters api-embed pager off brief on
```

Set CLI to show output in console format:

```
# set cli-parameters console
```

See also • show cli-parameters

set controller-date

Description Sets the date and time parameters for the system. You can set the date and time manually or configure the system to communicate with a Network Time Protocol (NTP) server. Alternatively, you can configure NTP by using the set ntp-parameters command.

📝 **NOTE:** If you specify valid NTP parameters and manual date/time parameters in the same command, the NTP parameters will take precedence. If the NTP server cannot be contacted, the date and time will not be changed and no error message will be displayed. If you specify the `timestamp` parameter and other manual date/time parameters in the same command, the `timestamp` parameter will take precedence.

Level Manage

Syntax To set the date and time manually:

```
set controller-date
    jan|feb|mar|apr|may|jun|jul|aug|sep|oct|nov|dec
    day
    hh:mm:ss
    year
```

To set the date and time manually by specifying a timestamp:

```
set controller-date
    timestamp timestamp
    timezone +|-hh[:mm]
```

To configure use of NTP:

```
set controller-date
    ntp enabled|disabled|on|off
    ntpaddress IP-address
    timezone +|-hh[:mm]
```

Parameters `jan|feb|mar|apr|may|jun|jul|aug|sep|oct|nov|dec`
The month.

`day`
The day number (1–31).

`hh:mm:ss`
The hour (0–23), the minutes (0–59), and the seconds (0–59).

`year`
The year as a four-digit number.

`ntp enabled|disabled|on|off`
Enables or disables use of NTP. When NTP is enabled and the specified NTP server is available, each controller's time is synchronized with the server.

`ntpaddress IP-address`
The network address of an available NTP server.

`timezone +|-hh[:mm]`
The system's time zone as an offset in hours (-1 to -12, +1 to +13) and minutes (0–59) from Coordinated Universal Time (UTC). For example, the Pacific Time Zone offset is -8 during Pacific Standard Time or -7 during Pacific Daylight Time.

`timestamp timestamp`
The date and time represented as the number of seconds (not counting leap seconds) that have elapsed since 1970-01-01 00:00:00 UTC. The resulting time will be in UTC, unless you also specify the `timezone` parameter.

Example Manually set the system time and date to 1:45 PM on September 22, 2012:

```
# set controller-date sep 22 13:45:0 2012
```

Manually set the system date and time to 4:30:50 PM on November 2, 2011 by specifying a timestamp and an offset for the Central Time zone:

```
# set controller-date timestamp 1320273050 timezone -6
```

Set the system to use NTP with an offset for the Mountain Time zone:

```
# set controller-date ntp enabled ntpaddress 69.10.36.3 timezone -7
```

Set the system to use NTP with an offset for the Bangalore, India, time zone:

```
# set controller-date ntp enabled ntpaddress 69.10.36.3 timezone +5:30
```

See also
- set ntp-parameters
- show controller-date
- show ntp-status

set debug-log-parameters

Description Sets the types of debug messages to include in the Storage Controller debug log. **For use by or with direction from technical support.**

Level Diagnostic

Syntax `set debug-log-parameters` *message-type*+|- [...]

Parameters *message-type*+|-
One of the following message types, followed by a plus (+) to enable or a minus (-) to disable inclusion in the log:

- `awt`: Auto-write-through cache triggers debug messages. Disabled by default.
- `bkcfg`: Internal configuration debug messages. Enabled by default.
- `cache`: Cache debug messages. Enabled by default.
- `capi`: Internal Configuration API debug messages. Enabled by default.
- `capi2`: Internal Configuration API tracing debug messages. Disabled by default.
- `disk`: Disk interface debug messages. Enabled by default.
- `dms`: Snapshot feature debug messages. Enabled by default.
- `emp`: Enclosure Management Processor debug messages. Enabled by default.
- `fo`: Failover and recovery debug messages. Enabled by default.
- `fruid`: FRU ID debug messages. Enabled by default.
- `hb`: Not used.
- `host`: Host interface debug messages. Enabled by default.
- `init`: Not used.
- `ioa`: I/O interface driver debug messages (standard). Enabled by default.
- `iob`: I/O interface driver debug messages (resource counts). Disabled by default.
- `ioc`: I/O interface driver debug messages (upper layer, verbose). Disabled by default.
- `iod`: I/O interface driver debug messages (lower layer, verbose). Disabled by default.
- `mem`: Internal memory debug messages. Disabled by default.
- `misc`: Internal debug messages. Enabled by default.
- `msg`: Inter-controller message debug messages. Enabled by default.
- `mui`: Internal service interface debug messages. Enabled by default.
- `ps`: Paged storage debug messages. Disabled by default.
- `raid`: RAID debug messages. Enabled by default.
- `rcm`: Removable-component manager debug messages. Disabled by default.
- `res2`: Internal debug messages. Disabled by default.
- `resmgr`: Reservation Manager debug messages. Disabled by default.

Example Include RAID and cache messages, exclude EMP messages, and leave other message types unchanged:

```
# set debug-log-parameters raid+ cache+ emp-
```

See also • show debug-log-parameters

set disk-parameters

Description Sets parameters that affect disk operation. Two features controlled by these parameters are disk Self-Monitoring Analysis and Reporting Technology (SMART) and drive spin down.

- Disks equipped with SMART technology can alert the controller of impending disk failure. When SMART is enabled, the system checks for SMART events one minute after a restart and every five minutes thereafter. SMART events are recorded in the event log. Changes to the SMART setting take effect after a rescan or a controller restart.

- The drive spin down feature monitors disk activity within system enclosures and spins down inactive disks, based on user-specified settings. This command sets spin-down parameters for available disks and spares.

> **NOTE:** Drive spin down affects disk operations as follows:
> - Spun-down disks are not polled for SMART events.
> - Operations requiring access to disks may be delayed while the disks are spinning back up.

Level Manage

Syntax
```
set disk-parameters
    [smart enabled|disabled|on|off|detect-only]
    [spin-down enabled|disabled|on|off]
    [spin-down-delay delay]
```

Parameters `smart enabled|disabled|on|off|detect-only`
Optional. Sets whether SMART is enabled or disabled for disks:

- `disabled` or `off`: Disables SMART for all disks in the system and for all disks added to the system.
- `enabled` or `on`: Enables SMART for all disks in the system and for all disks added to the system. This is the default.
- `detect-only`: Detects but does not change the SMART setting of each disk in the system, and for each new disk added to the system.

`spin-down enabled|disabled|on|off`
Optional. Sets whether available disks and spares will spin down after a period of inactivity shown by the `spin-down-delay` parameter.

- `disabled` or `off`: Drive spin down for available disks and spares is disabled. This is the default.
- `enabled` or `on`: Drive spin down for available disks and spares is enabled.

`spin-down-delay delay`
Optional. Sets the period of inactivity after which available disks and spares will spin down. The default is 15 minutes. Setting the delay to 1–360 minutes will enable spin down; setting the delay to 0 will disable spin down.

Example Enable SMART and drive spin down, and set the spin-down delay to 10 minutes:

```
# set disk-parameters smart on spin-down on spin-down-delay 10
```

See also • show disk-parameters

set drive-parameters

Alias for the set disk-parameters command.

set email-parameters

Description Sets SMTP notification parameters for events and managed logs.

Level Manage

Syntax
```
set email-parameters
    domain domain
    email-list email-addresses
    [include-logs enabled|disabled|on|off]
    notification-level crit|error|warn|info|none
    [sender sender]
    server server
```

Parameters domain *domain*
The domain name that is joined with an @ symbol to the sender name to form the "from" address for remote notification. The domain name can have a maximum of 255 bytes. Because this name is used as part of an email address, do not include spaces. For example: `MyDomain.com`. If the domain name is not valid, some email servers will not process the mail.

email-list *email-addresses*
Enter up to four comma-separated email addresses for recipients of event notifications. Each email address can have a maximum of 320 bytes. If the managed logs feature is enabled, the fourth email-address is used to specify the address of the log collection system. For example:
`IT-team@MyDomain.com,,,LogCollector@MyDomain.com`

include-logs enabled|disabled|on|off
Optional. When the managed logs feature is enabled, this option activates the "push" mode, automatically attaching system log files to managed logs email notifications that are sent to the log collection system. This option is disabled by default.

notification-level crit|error|warn|info|none
The minimum severity for which the system should send notifications:

- `crit`: Sends notifications for Critical events only.
- `error`: Sends notifications for Error and Critical events.
- `warn`: Sends notifications for Warning, Error, and Critical events.
- `info`: Sends notifications for all events.
- `none`: Disables email notification. If this option is specified, no other parameters are required and their current values are not changed.

sender *sender*
Optional. The sender name that is joined with an @ symbol to the domain name to form the "from" address for remote notification. This name provides a way to identify the system that is sending the notification. The sender name can have a maximum of 64 bytes. Because this name is used as part of an email address, do not include spaces. For example: `Storage-1`. If no sender name is set, a default name is created.

server *server*
The IP address of the SMTP mail server to use for the email messages.

Example Set the system to send an email from `RAIDsystem@mydomain.com` to both `sysadmin@mydomain.com` and `JSmith@domain2.com` when a non-Informational event occurs, and to send an email with attached logs to `logcollector@mydomain.com` when logs need to be transferred:

```
# set email-parameters server 10.1.1.10 sender RAIDsystem domain mydomain.com
notification-level warn include-logs enabled email-list
sysadmin@mydomain.com,JSmith@domain2.com,,logcollector@mydomain.com
```

See also
- show email-parameters
- test email

set enclosure

Description Sets an enclosure's name, location, rack number, and rack position. Set these parameters to values that help you identify and locate the enclosure. A value that contains a space must be enclosed in double quotes.

These values are used when user interfaces show enclosure-related data; for example, in output of the show enclosures command and in event-log entries related to enclosures.

Level Manage

Syntax `set enclosure`
```
[name new-name]
[location location]
[rack-number rack-number]
[rack-position rack-position]
enclosure-number
```

Parameters `name new-name`
A new name for the enclosure. The name can include a maximum of 20 bytes, using printable UTF-8 characters except angle brackets, double quote, or backslash. A value that contains a space must be enclosed in double quotes.

`location location`
The location of the enclosure. The name can include a maximum of 20 bytes, using printable characters except angle brackets, double quote, or backslash.

`rack-number rack-number`
The number of the rack containing the enclosure, from 0–255.

`rack-position rack-position`
The enclosure's position in the rack, from 0–255.

`enclosure-number`
The enclosure ID.

Example Set parameters for enclosure 1:

```
# set enclosure 1 name Storage-5 location Lab rack-number 9 rack-position 3
```

See also • show enclosures

set expander-fault-isolation

Description Temporarily disables PHY fault isolation for a specific Expander Controller. **For use by or with direction from technical support.**

By default, the Expander Controller in each I/O module performs fault-isolation analysis of SAS expander PHY statistics. When one or more error counters for a specific PHY exceed the built-in thresholds, the PHY is disabled to maintain storage system operation.

While troubleshooting a storage system problem, a service technician can use this command to temporarily disable fault isolation for a specific Expander Controller in a specific enclosure.

> **NOTE:** If fault isolation is disabled, be sure to re-enable it before placing the system back into service. Serious problems can result if fault isolation is disabled and a PHY failure occurs.

Level Manage

Syntax
```
set expander-fault-isolation
    [encl enclosure-ID]
    [wwn enclosure-WWN]
    [controller a|b|both]
    enabled|disabled|on|off
```

Parameters `encl enclosure-ID`
Optional. The enclosure ID of the enclosure containing the PHY. Specify either this parameter or the wwn parameter.

`wwn enclosure-WWN`
Optional. The WWN of the enclosure containing the PHY. Specify either this parameter or the `encl` parameter.

`controller a|b|both`
Optional. The I/O module containing the Expander Controller whose setting you want to change: A, B, or both. If this parameter is omitted, the setting is changed in both I/O modules.

`enabled|disabled|on|off`
Whether to enable or disable PHY fault isolation.

Example Disable PHY fault isolation for Expander Controller A in an enclosure:

```
# set expander-fault-isolation encl 1 controller a disabled
```

Re-enable PHY fault isolation for Expander Controller A in the same enclosure:

```
# set expander-fault-isolation encl 1 controller a enabled
```

See also • set expander-phy
• show enclosures
• show expander-status

set expander-phy

Description   Disables or enables a specific PHY. **For use by or with direction from technical support.**

Level   Manage

Syntax   ```
set expander-phy
      [encl enclosure-ID]
      [wwn enclosure-WWN]
      [controller a|b|both]type drive|inter-exp|sc-0|sc-1|ingress|egress
      phy phy-ID
      enabled|disabled|on|off
```

**Parameters**   encl *enclosure-ID*
Optional. The enclosure ID of the enclosure containing the PHY. Specify either this parameter or the wwn parameter.

wwn *enclosure-WWN*
Optional. The WWN of the enclosure containing the PHY. Specify either this parameter or the encl parameter.

controller a|b|both
Optional. The I/O module containing the PHY to enable or disable: A, B, or both. If this parameter is omitted, the setting is changed in both I/O modules.

type drive|inter-exp|sc-0|sc-1|ingress|egress
The PHY type:

- drive: PHY connected to a disk drive.
- inter-exp: PHY in an expansion module that communicates between its expander and the expander in the partner expansion module.
- sc-0: PHY in the ingress bus to the local Storage Controller.
- sc-1: PHY in the ingress bus to the partner Storage Controller.
- ingress: PHY in an ingress port.
- egress: PHY in an egress port.

phy *phy-ID*
The logical PHY number.

enabled|disabled|on|off
Whether to enable or disable the specified PHY.

**Example**   Disable the first egress PHY in controller A:

```
# set expander-phy encl 0 controller a type egress phy 0 disabled
```

Enable the PHY for disk 5 in controller B:

```
# set expander-phy encl 0 controller b type drive phy 5 enabled
```

See also   • set expander-fault-isolation
- show enclosures
- show expander-status

set host

Description Sets the name of a host.

Level Manage

Syntax
```
set host
     [name new-name]
     host-name
```

Parameters name *new-name*
Optional. Sets the name of the host to the specified name. The name is case sensitive; cannot include angle brackets, comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

host-name
The current name of the host.

Example Change the name of `Host1` to `MyHost`

```
# set host name MyHost Host1
```

See also • show hosts

set host-group

Description Sets the name of a host group.

Level Manage

Syntax `set host-group`
`    name `*`new-name`*
`    `*`host-group-name`*

Parameters `name `*`new-name`*
A new name for the host group. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

`host-group-name`
The current name of the host group.

Example Change the name of `HostGroup1` to `MyHostGroup`:

`# set host-name name MyHostGroup HostGroup1`

See also • show host-groups

set host-parameters

Description Sets controller host-port parameters for communication with connected hosts. You can set parameters for FC ports; there are no parameters for SAS ports.

Attempting to change FC loop IDs requires restarting the controllers; otherwise, new settings take effect without restarting the controllers.

Level Manage

Syntax To set FC-port parameters:

```
set host-parameters
    [controller a|b|both]
    [fibre-connection-mode loop|point-to-point|auto]
    [fibre-loop-id values]
    [noprompt]
    [ports ports|all]
    [prompt yes|no|expert]
    [speed 2g|4g|8g|auto]
```

Parameters `controller a|b|both`
Deprecated; use the `ports` parameter.

`fibre-connection-mode loop|point-to-point|auto`
Optional. For FC, sets the topology for the specified ports to:

- `loop`: Fibre Channel-Arbitrated Loop (public or private).
- `point-to-point`: Fibre Channel point-to-point.
- `auto`: Loop preferred, otherwise point-to-point, based on the detected connection type. This is the default.

You must also specify the `ports` parameter.

`fibre-loop-id values`
Optional. For FC, specifies comma-separated loop ID values to request for host ports when controllers arbitrate during a LIP. Use this option if you want ports to have specific addresses, if your system checks addresses in reverse order (lowest address first), or if an application requires that specific IDs be assigned to recognize the controller. If the loop ID is changed for one port, the same ID is used for other ports in the same controller. If the ports parameter is specified, loop IDs are set based on the controllers that the ports are in. You cannot specify the same value for ports in different controllers.

- `soft` or 255: Soft target addressing enables the LIP to determine the loop ID. Use this setting if the loop ID is permitted to change after a LIP or power cycle.
- 0–125: Specify a hard target address if you do not want the loop ID to change after a LIP or power cycle. If the port cannot acquire the specified ID, it is assigned a soft target address.

You must restart affected controllers to make loop ID changes take effect.

`noprompt`
Optional in console format; required for XML API format. Suppresses the confirmation prompt, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

`ports ports|all`
Optional. Specific host port numbers or all ports. For port syntax, see Command syntax on page 18.

```
prompt yes|no|expert
```
Optional. For scripting, this specifies an automatic response to the confirmation prompt:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.
- `expert`: Allow the command to proceed.

If this parameter is omitted, you must manually reply to the prompt.

```
speed 2g|4g|8g|auto
```
Optional. For FC, sets a forced link speed in Gbit/s or lets the speed be auto-negotiated (`auto`). Because a speed mismatch prevents communication between the port and host, set a speed only if you need to force the port to use a known speed for testing, or you need to specify a mutually supported speed for more than two FC devices connected in an arbitrated loop.

Example On a system with FC ports, set the link speed to 8 Gbit/s for ports A1 and B1, and set controller A ports to request loop ID 5 and controller B ports to request loop ID 32:

```
# set host-parameters speed 8g ports a1,b1 fibre-loop-id 5,32
WARNING: This change will take effect immediately. Changes may affect access to
data. Do you want to continue? (y/n) yes
```

On a system with FC ports, set controller A ports to request loop ID 14 and controller B ports to use soft target addressing, and suppress the confirmation prompt:

```
# set host-parameters fibre-loop-id 14,soft noprompt
```

See also
- restart
- show ports

set initiator

Description Sets the nickname of an initiator.

Level Manage

Syntax
```
set initiator
    [nickname nickname]
    id initiator
```

Parameters `nickname nickname`
Optional. Sets the nickname of the initiator to the specified name. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

`id initiator`
The ID of the initiator. For FC or SAS the ID is a WWPN. A WWPN can include a colon between each byte but the colons will be discarded.

Example For FC initiator 21000024ff3dfed1, set the nickname to `FC-port1`:

```
# set initiator nickname FC-port1 id 21000024ff3dfed1
```

See also • show initiators

set led

Description Changes the state of the identification LED on a specified disk or enclosure. For a disk this affects the fault LED. For an enclosure this affects the unit locator LED. LEDs are described in the Setup Guide.

Level Manage

Syntax To set a disk LED:

```
set led
    disk ID
    enable|disable|on|off
```

To set an enclosure LED:

```
set led
    enclosure ID
    enable|disable|on|off
```

Parameters disk *ID*
The disk to locate. For disk syntax, see Command syntax on page 18.

enclosure *ID*
The enclosure to locate.

enable|disable|on|off
Specifies to set or unset the LED.

Example Identify disk 5 in enclosure 1:

```
# set led disk 1.5 on
```

Stop identifying enclosure 1:

```
# set led enclosure 1 off
```

set network-parameters

Description Sets parameters for controller module network ports.

You can manually set static IP values for each controller, or you can specify that IP values should be set automatically for both controllers through communication with a Dynamic Host Configuration Protocol (DHCP) server.

Each controller has the following factory-default IP settings:

* DHCP: disabled
* Controller A IP address: 10.0.0.2
* Controller B IP address: 10.0.0.3
* IP subnet mask: 255.255.255.0
* Gateway IP address: 10.0.0.1

When DHCP is enabled, the following initial values are set and remain set until the system is able to contact a DHCP server for new addresses.

* Controller IP addresses: 169.254.x.x (where the value of $x.x$ is the lowest 16 bits of the controller serial number)
* IP subnet mask: 255.255.0.0
* Gateway IP address: 0.0.0.0

$169.254.x.x$ addresses (including gateway 169.254.0.1) are on a private subnet that is reserved for unconfigured systems and the addresses are not routable. This prevents the DHCP server from reassigning the addresses and possibly causing a conflict where two controllers have the same IP address. As soon as possible, change these IP values to proper values for your network.

To switch a controller from DHCP addressing to static addressing, you must set the IP address, netmask, and gateway values.

You can also set link speed, auto-negotiation, and duplex-mode parameters to enable Ethernet switches to communicate with the storage system. These settings will persist if the Ethernet cable is disconnected and reconnected, but will revert to default values if the controller is removed from the enclosure.

Level Manage

Syntax
```
set network-parameters
     [dhcp]
     [ipv 4|6]
     [ip address]
     [netmask netmask]
     [gateway gateway]
     [controller a|b]
     [ping-broadcast enabled|disabled|on|off]
     [link-speed 10mbps|100mbps]
     [duplex-mode full|half]
     [auto-negotiation enabled|disabled|on|off]
     [prompt yes|no]
     [noprompt]
```

Parameters dhcp
Optional. Specifies to use DHCP to set both controllers' IP values.

`ipv 4|6`
Optional. Specifies whether to use IP version 4 (IPv4) or 6 (IPv6) for addressing the network ports.

- `4`: Lets you specify addresses in dot-decimal format, where the four octets of the address use decimal values without leading zeroes and the octets are separated by a period; for example, `10.132.2.205`. This option is the default.
- `6`: Lets you specify addresses using eight groups of four hexadecimal digits, where the groups are separated by a colon. All groups must be specified. For example, 0000:0000:0000:0000:0000:0000:0A90:3442.

`ip address`
Optional. An IP address for the port.

`netmask netmask`
Optional. An IP subnet mask for the port.

`gateway gateway`
Optional. A gateway IP address for the port.

`controller a|b`
Optional. For IP-related parameters, this specifies whether to apply the settings to controller A or B. If this parameter is omitted, settings are applied to the controller being accessed. This parameter does not apply to Ethernet switch-related parameters, whose settings are always applied to the controller being accessed.

`ping-broadcast enabled|disabled|on|off`
Optional. Enables the storage system to respond when a ping to a broadcast address is issued on the system's subnet. The default is `Disabled`.

`link-speed 10mbps|100mbps`
Optional. Specifies the maximum link speed for auto-negotiation.

- `10mbps`: Sets the link speed to 10 Mb/s.
- `100mbps`: Sets the link speed to 100 Mb/s. If this parameter is omitted and a port's link speed is not set, it will be set to 100 Mb/s.

`duplex-mode full|half`
Optional. Specifies the duplex mode.

- `full`: Sets the port to use full duplex. If this parameter is omitted and a port's duplex mode is not set, it will be set to full duplex.
- `half`: Sets the port to use half duplex.

`auto-negotiation enabled|disabled|on|off`
Optional. Specifies whether to auto-negotiate the link speed.

- `enabled` or `on`: Enables auto-negotiation. If you enable auto-negotiation you must also specify the `link-speed` parameter to set the maximum supported speed for auto-negotiation. Because the controller will choose the proper duplex mode for the negotiated speed, you cannot specify the `duplex-mode` parameter.
- `disabled` or `off`: Disables auto-negotiation. If you disable auto-negotiation and don't specify the `link-speed` and `duplex-mode` parameters, the controller will use the current settings of the network port.

If this command is specified and the `auto-negotiation` parameter is omitted, auto-negotiation will be enabled.

```
prompt yes|no
```
Optional. For scripting, this specifies an automatic response to the confirmation prompt that appears when an Ethernet switch parameter is changed:

- `yes`: Allow the command to proceed.
- `no`: Cancel the command.

If this parameter is omitted, you must manually reply to the prompt.

```
noprompt
```
Optional in console format; required for XML API format. Suppresses the confirmation prompt that appears when an Ethernet switch parameter is changed, which requires a yes or no response. Specifying this parameter allows the command to proceed without user interaction.

Example Use DHCP to set network port IP values:

```
# set network-parameters dhcp
```

Manually set network port IP values for controller A (disabling DHCP for both controllers, if it was enabled) using IPv4 addressing:

```
# set network-parameters ip 192.168.0.10 netmask 255.255.255.0 gateway
192.168.0.1 controller a
```

Manually set network port IP values for controller A (disabling DHCP for both controllers, if it was enabled) using IPv6 addressing:

```
# set network-parameters ipv6 ip 2107:cdba::a9fe:88b0 netmask
FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:0000 gateway 2107:cdba::a9fe:1 controller a
```

For the controller you are accessing, set the network port to auto-negotiate with a maximum link speed of 100 Mb/s:

```
# set network-parameters auto-negotiation on link-speed 100mbps
WARNING: You are attempting to change the Ethernet settings for this Management
Controller, which will affect the link speed and duplex mode settings. Are you
sure? yes
```

For the controller you are accessing, for which auto-negotiation is disabled, set the network port to use a maximum link speed of 10 Mb/s and half-duplex mode:

```
# set network-parameters link-speed 10mbps duplex-mode half
WARNING: You are attempting to change the Ethernet settings for this Management
Controller, which will affect the link speed and duplex mode settings. Are you
sure? yes
```

See also • show network-parameters

set ntp-parameters

Description Sets Network Time Protocol (NTP) parameters for the system. You can manually set system date and time parameters by using the set controller-date command.

Level Manage

Syntax
```
set ntp-parameters
     [ntp enabled|disabled|on|off]
     [ntpaddress IP-address]
     [timezone +|-hh[:mm]]
```

Parameters `ntp enabled|disabled|on|off`
Optional. Enables or disables use of NTP. When NTP is enabled and the specified NTP server is available, each controller's time is synchronized with the server.

`ntpaddress IP-address`
Optional. The network address of an available NTP server.

`timezone +|-hh[:mm]`
Optional. The system's time zone as an offset in hours (-1 to -12, +1 to +13) and minutes (0–59) from Coordinated Universal Time (UTC). For example, the Pacific Time Zone offset is -8 during Pacific Standard Time or -7 during Pacific Daylight Time.

Example Set the system to use NTP with an offset for the Mountain Time zone:

```
# set ntp-parameters ntp enabled ntpaddress 69.10.36.3 timezone -7
```

Set the system to use NTP with an offset for the Bangalore, India, time zone:

```
# set ntp-parameters ntp enabled ntpaddress 69.10.36.3 timezone +5:30
```

See also • set controller-date
• show controller-date
• show ntp-status

set password

Description Sets a user's password for system interfaces (such as the CLI). A password can be entered as part of the command, or the command prompts you to enter and re-enter the new password.

Level Manage

Syntax `set password`
`    [password `*`password`*`]`
`    [`*`user`*`]`

Parameters `password `*`password`*
Optional. Sets a new password for the user. The value is case sensitive; can include a maximum of 32 bytes using printable UTF-8 characters except angle brackets, backslash, comma, or double quote. A value that includes a space must be enclosed in double quotes. If this parameter is omitted, the command prompts you to enter and re-enter a value, which is shown as a string of asterisks. For an SNMPv3 user whose `authentication-type` parameter is set to use authentication, this specifies the authentication password and must include at least 8 characters.

user
Optional. The user name for which to set the password. If this parameter is omitted, this command affects the logged-in user's password.

Example Change the password for the default user, `manage`:

```
# set password manage
Enter new password: ****
Re-enter new password: ****
```

Change the password for user `JDoe`:

```
# set password JDoe password Abc%123
```

See also • show users

set priorities

Description Sets snapshot-retention priorities for a specified snap pool. Snap-pool priorities, in conjunction with snapshot priorities, determine which snapshots are retained if system resource limitations require some snapshots to be automatically deleted.

Lower-priority snapshots will be deleted before higher-priority snapshots. Priority values are 1–65535. To use a default priority, specify the value 0.

Level Manage

Syntax `set priorities`
```
    [standard-snap #]
    [volume-copy #]
    [replication-snap #]
    [replicating #]
    [common-sync-point #]
    [only-sync-point #]
    [queued #]
    [drm-snapshot #]
    snap-pool
```

Parameters `standard-snap #`
Optional. The retention priority for a standard snapshot. The default is 0x6000 (24576).

`volume-copy #`
Optional. The retention priority for a snapshot that is being used to copy data from a source volume to a destination volume. This attribute is temporary for the duration of the volume-copy process. The default is 0xa000 (40960).

`replication-snap #`
Optional. The retention priority for a replication snapshot. The default is 0x4000 (16384).

`replicating #`
Optional. The retention priority for a snapshot that is being replicated to a secondary volume. This snapshot is required in order to resume the replication. The attribute is temporary for the duration of the replication process. The default is 0xc000 (49152).

`common-sync-point #`
Optional. The retention priority for the latest snapshot that is copy complete in all secondary volumes. It identifies a common point in time that is known by all destinations. The default is 0x8000 (32768).

`only-sync-point #`
Optional. The retention priority for the only sync point that is available in at least one secondary volume. If this snapshot is removed, then the next replication requires a full sync to be performed. The default is 0xe000 (57344).

`queued #`
Optional. The retention priority for a snapshot that was taken for remote replication but is queued waiting for the previous replications to complete. The default is 0x2000 (8192).

`drm-snapshot #`
Optional. The retention priority for a snapshot that was taken for remote replication.

`snap-pool`
The name or serial number of the snap pool.

Example Set attribute priorities for snap pool `SP1`, raising the priority for standard snapshots and leaving other priorities unchanged:

```
# set priorities only-sync-point 65535 SP1
```

See also • show priorities
  • show snap-pools

set prompt

Description Sets the prompt for the current CLI session.

Level Monitor

Syntax `set prompt` *prompt*

Parameters *prompt*
The new prompt, which can include any printable UTF-8 characters except angle brackets, backslash, and double quote, and can have a maximum of 16 bytes. A prompt that includes a space must be enclosed in double quotes.

Example Change the prompt from "# " to "CLI$ " and start entering another command:

```
# set prompt "CLI$ "
Success: Command completed successfully. (2012-08-14 14:47:23)
CLI> show ...
```

set protocols

Description Enables or disables management services and protocols.

Level Manage

Syntax `set protocols`
`      [debug enabled|disabled|on|off]`
`      [ftp enabled|disabled|on|off]`
`      [http enabled|disabled|on|off]`
`      [https enabled|disabled|on|off]`
`      [ses enabled|disabled|on|off]`
`      [smis enabled|disabled|on|off]`
`      [snmp enabled|disabled|on|off]`
`      [ssh enabled|disabled|on|off]`
`      [telnet enabled|disabled|on|off]`
`      [usmis enabled|disabled|on|off]`
`      [activity enabled|disabled|on|off]`

Parameters `debug enabled|disabled|on|off`
Optional. Enables or disables debug capabilities, including Telnet debug ports and privileged diagnostic user IDs. This is disabled by default.

`ftp enabled|disabled|on|off`
Optional. Enables or disables the expert interface for updating firmware. This is enabled by default.

`http enabled|disabled|on|off`
Optional. Enables or disables the standard Storage Management Console web server. This is enabled by default.

`https enabled|disabled|on|off`
Optional. Enables or disables the secure Storage Management Console web server. This is enabled by default.

`ses enabled|disabled|on|off`
Optional. Enables or disables the in-band SCSI Enclosure Management Services interface. This is enabled by default.

`smis enabled|disabled|on|off`
Optional. Enables or disables the secure Storage Management Initiative Specification interface (SMI-S) interface. This option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTPS port 5989. HTTPS port 5989 and HTTP port 5988 cannot be enabled at the same time, so enabling this option will disable port 5988. This is enabled by default.

`snmp enabled|disabled|on|off`
Optional. Enables or disables the Simple Network Management Protocol interface. Disabling this option disables all SNMP requests to the MIB and disables SNMP traps. To configure SNMP traps use the set snmp-parameters command. This is enabled by default.

`ssh enabled|disabled|on|off`
Optional. Enables or disables the secure shell CLI. This is enabled by default.

`telnet enabled|disabled|on|off`
Optional. Enables or disables the standard CLI. This is enabled by default.

`usmis enabled|disabled|on|off`
Optional. Enables or disables the unsecure Storage Management Initiative Specification (SMI-S) interface. This option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTP port 5988. HTTP port 5988 and HTTPS port 5989 cannot be enabled at the same time, so enabling this option will disable port 5989. This is disabled by default.

`activity enabled|disabled|on|off`
Optional. Allows the user to moitor progress of activities (like partner firmware upgrade) on the patner controller.

Example Disable unsecure HTTP connections and enable FTP:

`# set protocols http disabled ftp enabled`

See also • show protocols

set remote-system

Description Changes remote-system credentials stored in the local system. Do this when the user name or password to access a remote system has been changed in that system.

Level Manage

Syntax `set remote-system`
  `[password` *password*`]`
  `[username` *username*`]`
  system

Parameters `password` *password*
Optional. The new password to access the remote system. The value is displayed in clear text.

`username` *username*
Optional. The new user name to access the remote system.

`system`
The name or network-port IP address of the remote system.

Example Change the password that is stored to access a remote system:

`# set remote-system password Abc_123 System2`

See also • create remote-system
  • delete remote-system
  • remote
  • show remote-systems

set replication-primary-volume

Description Changes the primary volume for a replication set. You must issue this command to each volume in the replication set. If the volumes in a replication set have different primary-volume settings — for example, if the primary volume was changed while one volume was offline — this results in a primary-volume conflict. Change the primary volume on the secondary system, then, if possible, change the primary volume on the primary system.

As part of this command, you can initiate a rollback to a specified snapshot to synchronize the new primary volume with a known data image.

When the secondary volume becomes the primary volume, it only retains the replication images that the primary volume had and deletes any images that the primary volume did not have. Because the secondary volume may not have successfully replicated all the images associated with the primary volume, the secondary volume might have a subset of the primary volume's images.

Level Manage

Syntax
```
set replication-primary-volume
    [nowait]
    primary-volume replication-volume
    [set replication-set]
    [snapshot snapshot]
    volume replication-volume
```

Parameters `nowait`
Optional. Changing the primary volume can take the Storage Controller several minutes to complete. This parameter allows that processing to continue in the background so the Management Controller can process other commands.

`primary-volume` *replication-volume*
The name or serial number of the replication volume to designate as the new primary volume for the replication set. If the name is not unique, you must specify the serial number.

`set` *replication-set*
Optional. The name or serial number of the replication set.

`snapshot` *snapshot*
Optional. Name or serial number of a snapshot to roll the replication volume data back to.

`volume` *replication-volume*
The name or serial number of a volume in the replication set whose primary volume you want to change. If one of the volumes in the replication set is offline, this must specify the volume that remains online. If the name is not unique across replication sets, specify the `set` parameter.

Example For this example, assume that:

- Replication set `RS` has the primary volume `Data` and the secondary volume `rData`.
- Primary volume `Data` resides in the primary system, `System1`.
- Secondary volume `rData` resides in the secondary system, `System2`.

On `System1`, view the status of replication set `RS`:

```
# show replication-sets RS
Replication Set [Name (RS) Serial Number (SN) ] Primary Volume:
   Name   ...   Status ... Location Primary-Volume ... Primary-Volume-Status
   --------------------------------------------------------------------
   Data   ...   Online ... Local    Data            ... Online
   rData  ...   Online ... Remote   Data            ... Online
```

On `System1`, unmap the primary volume from hosts.

On `System2`, set secondary volume `rData` to be the primary volume:

```
# set replication-primary-volume volume rData primary-volume rData
```

On `System2`, view the set's status and notice that the primary volume has changed and that a primary-volume conflict exists:

```
# show replication-sets RS
Replication Set [Name (RS) Serial Number (SN) ] Primary Volume:
   Name  ...   Status ... Location Primary-Volume ... Primary-Volume-Status
   ------------------------------------------------------------------------
   Data  ...   Online ... Remote   Data          ... Conflict
   rData ...   Online ... Local    rData         ... Conflict
```

On `System1`, view the set's status and notice that it does not reflect the primary-volume change, thereby causing the conflict:

```
# show replication-sets RS
Replication Set [Name (RS) Serial Number (SN) ] Primary Volume:
   Name  ...   Status ... Location Primary-Volume ... Primary-Volume-Status
   ------------------------------------------------------------------------
   Data  ...   Online ... Local    Data          ... Conflict
   rData ...   Online ... Remote   Data          ... Conflict
```

On `System1`, as already done on `System2`, set `rData` to be the primary volume:

```
# set replication-primary-volume volume Data primary-volume rData
```

On `System1` (the new secondary system), view the set's status and notice that the system is set to use the new primary volume and the conflict is resolved:

```
# show replication-sets RS
Replication Set [Name (RS) Serial Number (SN) ] Primary Volume:
   Name  ...   Status ... Location Primary-Volume ... Primary-Volume-Status
   ------------------------------------------------------------------------
   Data  ...   Online ... Local    rData         ... Online
   rData ...   Online ... Remote   rData         ... Online
```

Wait a couple of minutes for processing to complete. Then, on `System2` (the new primary system), view the set's status and notice that the system is set to use the new primary volume and that the conflict is resolved:

```
# show replication-sets
Replication Set [Name (RS) Serial Number (SN) ] Primary Volume:
   Name  ...   Status ... Location Primary-Volume ... Primary-Volume-Status
   ------------------------------------------------------------------------
   Data  ...   Online ... Local    rData         ... Online
   rData ...   Online ... Remote   rData         ... Online
```

Map the new primary volume to hosts.

See also
- show replication-sets
- show replication-volumes
- show snapshots
- unmap volume

set replication-volume-parameters

Description Sets parameters for a specified replication volume. This command must be run separately on each system where the volume resides; changes to these parameters are not automatically synchronized across systems.

Level Manage

Syntax
```
set replication-volume-parameters
    [link-type FC]
    [max-queue #]
    [max-retry-time #]
    [monitor-interval #]
    [on-collision newest|oldest]
    [on-error retry|suspend]
    [priority low|medium|high]
    [remote-address ip=IPs|wwnn=WWNNs|wwpn=WWPNs]
    [set replication-set]
    replication-volume
```

Parameters `link-type FC`
Optional. Specifies the type of ports being used for the inter-system link:

- `FC`: FC ports.

`max-queue #`
Optional. The number of replication images to consider when determining the next image to replicate: 1–64. Used only if the `on-collision` parameter is set to `oldest`.

`max-retry-time #`
Optional; valid only if the `on-error` parameter is set to `retry`. The maximum time in seconds to retry a single replication if an error occurs. Allowed values are 0-64000; the default is 1800 (30 minutes). A value of 0 means do not time out retries; that is, retry forever. By default, a retry will occur 5 minutes after an error occurs. If another error occurs and the difference in time between when the error occurs and the initial retry time is greater than the max-retry-time value, the replication will be suspended.

In order to prevent a replication set from suspending when multiple independent, discontinuous errors occur during a single replication, set `max-retry-time` for the secondary volume either to 0 (retry forever) or to 60 minutes for each 10GB of volume size.

`monitor-interval #`
Optional. The interval in seconds at which the primary volume should query the secondary volume. Values less than 300 (5 minutes) or greater than 1800 (30 minutes) are not allowed. The default is 300.

`on-collision newest|oldest`
Optional. The collision policy to use when the image queue depth is met:

- `newest`: Only the latest replication image should be considered for the next replication operation.
- `oldest`: Only the latest N replication images should be considered for the next replication operation, where N is defined by the `max-queue` parameter and the oldest of these images should be considered first.

`on-error retry|suspend`
Optional. The error policy to use when errors occur during the replication process:

- `retry`: Retry the operation for the time specified in the `max-retry-time` parameter.
- `suspend`: Suspend the replication operation.

`priority low|medium|high`
Optional. The priority of the replication process on the replication volume: low, medium, or high.

`remote-address ip=`*`IPs`*`|wwnn=`*`WWNNs`*`|wwpn=`*`WWPNs`*
Optional. The remote addresses associated with the replication volume. Specifies host ports in the remote system by IP address, WWNN, or WWPN. An IP address value can include a port number; for example, `10.134.11.10:3260`. Multiple values must be separated by commas and no spaces; for example: `ip=10.134.2.1,10.134.2.2`.

`set `*`replication-set`*
Optional. The name or serial number of the replication set.

`replication-volume`
The name or serial number of the replication volume. If the name is not unique within the replication set, the local volume is assumed. If the name   is not unique across replication sets, specify the set parameter.

Example Set a new remote address and error policy for replication volume `V2` in replication set `RS1`:

```
# set replication-volume-parameters remote-address ip=10.1.66.55 on-error
suspend set RS1 V2
```

See also • show replication-sets
• show replication-volumes

set schedule

Description Changes parameters for a specified schedule. If you want to change the schedule name, create a new schedule to replace the existing one.

When scheduling `ReplicateVolume` tasks, a best practice is to schedule no more than three volumes to start replicating at the same time, and for those replications to recur no less than 60 minutes apart. If you schedule more replications to start at the same time, or schedule replications to start more frequently, some scheduled replications may not have time to complete.

Level Manage

Syntax
```
set schedule
    schedule-specification "specification"
    task-name task-name
    schedule-name
```

Parameters `schedule-specification "specification"`
Defines when the task will first run, and optionally when it will recur and expire. You can use a comma to separate optional conditions. Dates cannot be in the past. For times, if neither `AM` nor `PM` is specified, a 24-hour clock is used.

- `start yyyy-mm-dd hh:mm [AM|PM]`
  Specifies a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.

- `[every # minutes|hours|days|weeks|months|years]`
  Specifies the interval at which the task will run.
  For better performance when scheduling a `TakeSnapshot` task that will run under heavy I/O conditions or on more than three volumes, the retention count and the schedule interval should be set to similar values; for example if the retention count is 10 then the interval should be set to 10 minutes.
  For a `ReplicateVolume` task, the minimum interval is 30 minutes.

- `[between hh:mm [AM|PM] and hh:mm [AM|PM]]`
  Constrains the time range during which the task is permitted to run. Ensure that the start time is within the specified time range.

- `[only any|first|second|third|fourth|fifth|last|#st|#nd|#rd|#th weekday|weekendday|Sunday|Monday|Tuesday|Wednesday|Thursday|Friday |Saturday of year|month|January|February|March|April|May|June|July |August|September|October|November|December]`
  Constrains the days or months when the task is permitted to run. Ensure that this constraint includes the start date.

- `[count #]`
  Constrains the number of times the task is permitted to run.

- `[expires yyyy-mm-dd hh:mm [AM|PM]]`
  Specifies when the schedule expires, after which the task will no longer run.

`task-name task-name`
The name of the task to run. The name is case sensitive.

`schedule-name`
The name of the schedule to change.

Example Change parameters, including the associated task, for schedule `Sched1`:

```
# set schedule schedule-specification "start 2012-01-01 00:01 every 1 days
expires 2012-12-31 00:01" task-name Snap Sched1
```

See also • show schedules
- show tasks

set snap-pool-parameters

Description   Sets consumption thresholds and associated recovery policies for both snap pools. **For use by or with direction from technical support.**

Each storage pool has a single snap pool. Each snap pool has the following thresholds:

- Warning: The snap pool is moderately full. When this threshold is reached, an event is logged.
- Error: The snap pool is nearly full and unless corrective action is taken, snapshot data loss is probable. When this threshold is reached, an Informational event is logged and the error policy is invoked.
- Critical: The snap pool is essentially full and data loss is imminent unless action is taken. When this threshold is reached, an event is logged and the critical policy is invoked.

> 📝 **NOTE:**   The `deleteoldestsnapshot` policy does not apply business logic to the delete decision and may delete snapshots that are mounted or modified. You may set retention priorities for a snap pool to indicate that some snapshots are more important than others, but these priorities do not ensure any specific snapshot is protected. For more information see the set priorities and show priorities commands.

Level   Diagnostic

Syntax
```
set snap-pool-parameters
    [storage-pool storage-pool-ID]
    [warning-threshold #%]
    [error-threshold #%]
    [critical-threshold #%]
    [error-policy nochange|haltwrites|notifyonly|deleteoldestsnapshot]
    [critical-policy nochange|haltwrites|notifyonly|deleteoldestsnapshot]
```

Parameters   `storage-pool storage-pool-ID`
Optional. The name or serial number of the storage pool. Names are case sensitive. If this parameter is omitted, specified parameter changes will be applied to the snap pool in each storage pool.

`warning-threshold #%`
Optional. The percent of snap-pool consumption that triggers the warning policy. This value must be less than the error-threshold value. The default is 75%.

`error-threshold #%`
Optional. The percent of snap-pool consumption that triggers the error policy. This value must be less than the critical-threshold value. The default is 90%.

`critical-threshold #%`
Optional. The percent of snap-pool consumption that triggers the critical policy. This value must be less than or equal to 98%. The default is 98%.

`error-policy nochange|haltwrites|notifyonly|deleteoldestsnapshot`
Optional. The action to take when the error threshold is reached.

- `nochange`: Retains the current setting.
- `haltwrites`: Halts writes to all volumes and snapshots associated with the snap pool to prevent existing snapshots from getting out of sync when subsequent volume writes require additional snap-pool resources, which are unavailable.
- `notifyonly`: Logs an event to notify the administrator. This is the default.
- `deleteoldestsnapshot`: Deletes the oldest snapshot that is not the current sync point of a replication set.

`critical-policy nochange|haltwrites|notifyonly|deleteoldestsnapshot`
Optional. The action to take when the critical threshold is reached.

- `nochange`: Retains the current setting.
- `haltwrites`: Halts writes to all volumes and snapshots associated with the snap pool to prevent existing snapshots from getting out of sync when subsequent volume writes require additional snap-pool resources, which are unavailable.
- `notifyonly`: Logs an event to notify the administrator. This is the default.
- `deleteoldestsnapshot`: Deletes the oldest snapshot that is not the current sync point of a replication set.

Example   Change the thresholds and policies for the snap pool in storage pool A:

```
# set snap-pool-parameters storage-pool A warning-threshold 70% error-threshold
85% critical-threshold 90% error-policy deleteoldestsnapshot critical-policy
haltwrites
```

See also
- set snapshot-policy
- set storage-pool
- show snap-pools
- show storage-pools
- reset snap-pool

set snapshot-policy

Description Sets the action to take when a snap pool reaches its error threshold. By default, this threshold indicates that 90% of the snap-pool capacity has been consumed. This command applies to both snap pools.

Level Manage

Syntax `set snapshot-policy autodelete enabled|disabled|on|off`

Parameters `autodelete enabled|disabled|on|off`

- `enabled` or `on`: Deletes the oldest snapshot that is not the current sync point of a replication set.
- `disabled` or `off`: Logs an event to notify the administrator that the threshold has been reached. This is the default.

Example Enable automatic snapshot deletion for both snap pools:

`# set snapshot-policy autodelete on`

See also • show snap-pools

set snmp-parameters

Description   Sets SNMP parameters for event notification. To enable or disable SNMP requests to the MIB use the set protocols command.

Level   Manage

Syntax
```
set snmp-parameters
     [enable crit|error|warn|info|none]
     [add-trap-host address]
     [del-trap-host address]
     [trap-host-list trap-host-list]
     [read-community string]
     [write-community string]
```

Parameters   enable crit|error|warn|info|none
Optional. Sets the level of trap notification:

- `crit`: Sends notifications for Critical events only.
- `error`: Sends notifications for Error and Critical events.
- `warn`: Sends notifications for Warning, Error, and Critical events.
- `info`: Sends notifications for all events.
- `none`: All events are excluded from trap notification and traps are disabled.

add-trap-host address
Optional. Specifies the IP address of a destination host that will receive traps. Three trap hosts can be set.

del-trap-host address
Optional. Deletes a trap destination host.

trap-host-list trap-host-list
Optional. Replaces the current list.

read-community string
Optional. Sets a community string for read-only access. This value is also included in traps that are sent. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes.

write-community string
Optional. Sets a community string for write access. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes.

Example   Enable Critical events only, specify a trap host, and set the community string for read-only access:

```
# set snmp-parameters enable crit add-trap-host 172.22.4.171 read-community
public
```

See also
- set protocols
- show snmp-parameters
- test

set storage

Description   Sets the system to use thin provisioning or full provisioning when creating volumes.

Level   Manage

Syntax   `set storage overcommit enabled|disabled|on|off`

Parameters   `overcommit enabled|disabled|on|off`

- `enabled or on`: The system will use thin provisioning, which means that more capacity can be allocated to volumes than physically exists in the system. When stored data approaches the limit of physical capacity, the administrator can add more enclosures to the system.
- `disabled or off`: The system will use full provisioning, which means that the capacity allocated to volumes when they are created cannot exceed the physical capacity of the system. This is the default.

NOTE: If you try to disable overcommit and the total space allocated to thin-provisioned volumes exceeds the physical capacity of their storage pool, an error will say that there is insufficient free disk space to complete the operation and overcommit will remain enabled.

Example   Enable overcommit for storage:

`# set storage overcommit on`

See also
- show storage
- show system-parameters

set storage-pool

Description Sets thresholds and maximum snap-pool size for a specified storage pool.

Level Manage

Syntax
```
set storage-pool
    [low-threshold #%]
    [middle-threshold #%]
    [snap-pool-size #%]
    storage-pool-name
```

Parameters `low-threshold #%`
The low threshold for page allocation as a percentage of storage-pool capacity. When this threshold is reached, event 462 is logged with Informational severity. This value must be less than the `middle-threshold` value. The default is 25%.

`middle-threshold #%`
The middle threshold for page allocation as a percentage of storage-pool capacity. When this threshold is reached, event 462 is logged to notify the administrator to add capacity to the pool. This value must be between the `low-threshold` value and the high-threshold value, which is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space. The default `middle-threshold` value is 50%. If overcommit is enabled, the event has Informational severity; if overcommit is disabled, the event has Warning severity. (There is also a high threshold, which when reached logs event 462 with Critical severity to alert the administrator that it is critical to add capacity to the pool.)

`snap-pool-size #%`
Optional. The size of the snap-pool as a percentage of storage-pool capacity. The actual size of the snap pool will be the greater of the minimum snap-pool size (approximately 5 GB) or the specified percentage of the storage pool (rounded up to the nearest storage-pool page size). Physical storage will be set aside or reserved for exclusive use by the snap pool and will not be available for use by user volumes. When overcommit is disabled, this reduces the available storage pool capacity by the size of the snap pool.

📝 **IMPORTANT:**

- The default snap-pool size (0%) is not adequate for *any* snapshot activity.
- A snap pool that is too small can become consumed quickly with heavy snapshot activity, leading to threshold events for the snap pool; in the worst case, writes to user volumes will be halted. Therefore, determine and set `snap-pool-size` before using snapshot features.
- The snap-pool size can be increased but it *cannot* be decreased. Therefore, be careful not to make the snap pool too large. A snap pool can be reset (by using the reset snap-pool command) but doing so will delete all current snapshots for that snap pool.

`storage-pool-name`
The name of the storage pool for which to change settings.

Example For storage pool A, set the low threshold to 30% and the maximum snap-pool size to 10%:

```
# set storage-pool low-threshold 30% snap-pool-size 10% A
```

See also
- set storage
- show storage-pools

set system

Description Sets the system's name, contact person, location, and description. Each value can include a maximum of 79 bytes, using printable UTF-8 characters except angle bracket, double quote, or backslash. A value that contains a space must be enclosed in double quotes.

The name, location, and contact are included in event notifications. All four values are included in system debug logs for reference by service personnel. When using the WBI, the system name appears in the browser title bar or tab.

Level Manage

Syntax
```
set system
    [name value]
    [contact value]
    [location value]
    [info value]
```

Parameters
name *value*
Optional. A name to identify the system.

contact *value*
Optional. The name of the person who administers the system.

location *value*
Optional. The location of the system.

info *value*
Optional. A brief description of what the system is used for or how it is configured.

Example Set the system name to `Test` and the contact to `J. Doe`:

```
# set system name Test contact "J. Doe"
```

See also • show system

set task

Description Changes parameters for a TakeSnapshot, ReplicateVolume, or VolumeCopy task. For these types of tasks, you can change parameters other than name, type, or associated volumes. If you change the parameters for a running task, the changes will take effect the next time the task runs.

If you want to change parameters for a ResetSnapshot task or the name, type, or associated volumes for another type of task, create a new task to replace the existing one.

Level Manage

Syntax
```
set task
    [snapshot-prefix prefix]
    [retention-count #]
    [dest-prefix prefix]
    [modified-snapshot yes|no]
    [replication-mode new-snapshot|last-snapshot]
    name
```

Parameters snapshot-prefix *prefix*
Optional. A label to identify snapshots created by this task. Snapshot names have the format *prefix*_s#, where # starts at 001.

retention-count #
Optional. For a TakeSnapshot task, the number of snapshots created by this task to retain, from 1 to the licensed limit. When a new snapshot exceeds the limit, the oldest snapshot with the same prefix is deleted. If you reduce the retention count for a task, excess snapshots will be removed the next time the task runs.

For a ReplicateVolume task, the number of replication images created by this task to retain, from 2 to 32. When a new image exceeds this limit, the oldest image with the same prefix is deleted. This parameter applies to the primary volume only; for the secondary volume, images will accumulate until either the secondary storage pool's space limit is reached or the maximum number of images is reached, after which the oldest image will be deleted as new images are created.

dest-prefix *prefix*
Optional. A label to identify the volume copy created by this task. Copy names have the format *prefix*_c#, where # starts at 001.

modified-snapshot yes|no
Optional. For a VolumeCopy task, specifies whether to include or exclude modified write data from the snapshot in the copy. This parameter applies only when the source volume is a snapshot; it is ignored if the source volume is a standard volume.

- yes: Include modified snapshot data.
- no: Exclude modified snapshot data.

If this parameter is omitted for a snapshot, modified snapshot data is excluded.

replication-mode new-snapshot|last-snapshot
Optional. Specifies whether to replicate a new snapshot of the volume to the remote system, or to replicate the last (most recent existing) snapshot of the volume to the remote system.

- new-snapshot: Replicate a new snapshot.
- last-snapshot: Replicate the most recent existing snapshot.

If this parameter is omitted, a new snapshot is replicated.

name
The name of the task to change.

Example Change parameters for a TakeSnapshot-type task named Snap:

```
# set task snapshot-prefix v1 retention-count 2 Snap
```

Change parameters for a VolumeCopy-type task named `Copy`:

```
# set task dest-prefix v1 modified-snapshot no Copy
```

Change parameters for a ReplicateVolume-type task named `Replicate`:

```
# set task snapshot-prefix v2 replication-mode last-snapshot Replicate
```

See also
- create task
- delete task
- set schedule
- show schedules
- show tasks

set user

Description Changes user preferences for the session or permanently. You cannot change the access level of user `manage`.

📝 **NOTE:** User changes take effect when the user next logs in.

Level Manage

Syntax
```
set user
     [authentication-type MD5|SHA|none]
     [base 2|10]
     [interfaces interfaces]
     [roles roles]
     [level monitor|manage|diagnostic]
     [locale English|en|Spanish|es|French|fr|German|de|Italian|it|Japanese|ja
     |Dutch|nl|Chinese-simplified|zh-s]
     [password password]
     [precision #]
     [privacy-password encryption-password]
     [privacy-type DES|AES|none]
     [session-preferences]
     [storage-size-base 2|10]
     [storage-size-precision #]
     [storage-size-units auto|MB|GB|TB]
     [temperature-scale celsius|c|fahrenheit|f]
     [timeout #]
     [trap-host IP-address]
     [type novice|standard|advanced|diagnostic]
     [units auto|MB|GB|TB]
     name
```

Parameters `authentication-type MD5|SHA|none`
Optional. For an SNMPv3 user, this specifies whether to use a security authentication protocol. This parameter requires the password parameter and, for the `snmptarget` interface, the `trap-host` parameter.

- `MD5`: MD5 authentication. This is the default.
- `SHA`: SHA (Secure Hash Algorithm) authentication.
- `none`: No authentication.

`base 2|10`
Optional. Sets the base for entry and display of storage-space sizes:

- `2`: Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
- `10`: Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is the default.

Operating systems usually show volume size in base 2. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2.

`interfaces` *interfaces*
Optional. Specifies the interfaces that the user can access. Multiple values must be separated by commas and no spaces. The defaults are `cli` and `wbi`.

- `cli`: Command-line interface.
- `wbi`: Web-browser interface (Storage Management Console).
- `ftp`: File transfer protocol interface.
- `smis`: Storage Management Initiative Specification (SMI-S) interface.
- `snmpuser`: Allows an SNMPv3 user to view the SNMP MIB.
- `snmptarget`: Allows an SNMPv3 user to receive SNMP trap notifications. This option requires the `trap-host` parameter.
- `none`: No interfaces.

`level monitor|manage|diagnostic`
Optional.

- `monitor`: User can view but not change system settings. This is the default.
- `manage`: User can view and change system settings.
- `diagnostic`: User can access additional, diagnostic commands.

You cannot specify both this parameter and the `roles` parameter.

`roles` *roles*
Optional. Specifies the user's roles as one or more of the following values:

- `monitor`: User can view but not change system settings. This is the default.
- `manage`: User can view and change system settings.
- `diagnostic`: User can access additional, diagnostic commands.

Multiple values must be separated with a comma (with no spaces). If multiple values are specified, the user's access to commands will be determined by the highest-level role specified. You cannot specify both this parameter and the `level` parameter.

`locale English|en|Spanish|es|Japanese|ja|Chinese-simplified|zh-s`
Optional. The display language. The default is English.

`password` *password*
Optional in console format; required for XML API format. Sets a new password for the user. The value is case sensitive; can include a maximum of 32 bytes using printable UTF-8 characters except a backslash, comma, or double quote. A value that includes a space must be enclosed in double quotes. For an SNMPv3 user whose `authentication-type` parameter is set to use authentication, this specifies the authentication password and must include at least 8 characters.

`precision #`
Optional. Sets the number of decimal places (1–10) for display of storage-space sizes. The default is 1.

`privacy-password` *encryption-password*
Optional. For an SNMPv3 user whose `privacy-type` parameter is set to use encryption, this specifies the encryption password. The value is case sensitive; can include a maximum of 32 bytes using printable UTF-8 characters except a backslash, comma, or double quote; and must include at least 8 characters.

`privacy-type DES|AES|none`
Optional. For an SNMPv3 user, this specifies whether to use a security encryption protocol. This parameter requires the `privacy-password` parameter and the `authentication-type` parameter.

- `DES`: Data Encryption Standard.
- `AES`: Advanced Encryption Standard.
- `none`: No encryption. This is the default.

`session-preferences`
Optional. Specifies that the current CLI settings will become permanent settings for the user. This parameter cannot be combined with any other parameter.

`storage-size-base 2|10`
Optional. Alias for `base`.

`storage-size-precision #`
Optional. Alias for `precision`.

`storage-size-units auto|MB|GB|TB`
Optional. Alias for `units`.

`temperature-scale celsius|c|fahrenheit|f`
Optional. Sets the scale for display of temperature values:

- `fahrenheit` or `f`: Temperatures are shown in degrees Fahrenheit.
- `celsius` or `c`: Temperatures are shown in degrees Celsius. This is the default.

`timeout #`
Optional. Sets the timeout value in seconds for the login session. Valid values are 120–43200 seconds (2–720 minutes). The default is 1800 seconds (30 minutes).

`trap-host IP-address`
Optional. For an SNMPv3 user whose `interface` parameter is set to `snmptarget`, this specifies the IP address of the host that will receive SNMP traps.

`type novice|standard|advanced|diagnostic`
Optional. Identifies the user's experience level. The default is `standard`. This parameter does not affect access to commands.

`units auto|MB|GB|TB`
Optional. Sets the unit for display of storage-space sizes:

- `auto`: Sizes are shown in units determined by the system. This is the default.
- `MB`: Sizes are shown in megabytes.
- `GB`: Sizes are shown in gigabytes.
- `TB`: Sizes are shown in terabytes.

Based on the `precision` setting, if a size is too small to meaningfully display in the selected unit, the system uses a smaller unit for that size. For example, if `units` is set to TB, `precision` is set to 1, and `base` is set to 10, the size 0.11709 TB is instead shown as 117.1 GB.

`name`
Specifies the user account to change. Names are case sensitive.

Example Change the temperature scale and accessible interfaces for user `jsmith`:

`# set user jsmith temperature-scale f interfaces wbi,cli`

Change the password for user `JDoe`:

`# set user JDoe password Abc%123`

Change the authentication type for SNMPv3 user `Traps`:

`# set user Traps authentication-type MD5 password snmptraps`

See also
- set password
- show users

set volume

Description Changes a volume's name and identifying information.

Level Manage

Syntax
```
set volume
    [access read-write|rw|read-only|ro]
    [identifying-information description]
    [name new-name]
    [preference archive|standard|performance]
    volume
```

Parameters `access read-write|rw|read-only|ro`
Optional. The access permission for initiators the volume: read-write (`rw`), read-only (`ro`), or no-access. Read-only access applies to all ports. If `no-access` is specified, the volume is not mapped. The default is `read-write`.

`identifying-information description`
Optional. A description of the volume to help a host-side user identify it. The value can have a maximum of 127 bytes, using printable UTF-8 characters except angle brackets, double quote, or backslash. A value that contains a space must be enclosed in double quotes.

`name new-name`
Optional. A new name for the volume. The name is case sensitive; can include spaces and printable UTF-8 characters except angle brackets, backslash, comma, and double quote; and can have a maximum of 20 bytes. A name that includes a space must be enclosed in double quotes.

`preference archive|standard|performance`
Optional. Specifies how to tune the tier-migration algorithm for the volume:

- `archive`: This setting specifies never to move this volume to the SSD tier, regardless of how often the data is accessed.
- `standard`: This setting attempts to balance the frequency of data access, disk cost, and disk availability by moving the volume to the appropriate tier. This is the default.
- `performance`: This setting specifies to keep this volume in the SSD tier as much as possible.

`volume`
The name or serial number of the volume to change. For volume syntax, see Command syntax on page 18.

Example Rename volume `V1` to `V2`:

```
# set volume name V2 V1
```

Set identifying information for `V3`:

```
# set volume identifying-information "Project X data" V3
```

Restrict volume `OldFiles` to the Archive tier with read-only access:

```
# set volume preference archive access ro OldFiles
```

See also
- show maps
- show volumes

set volume-group

Description Sets the name of a volume group.

Level Manage

Syntax `set volume-group`
`    name new-name`
`    volume-group-name`

Parameters `name new-name`
A new name for the volume group. The name is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes. A name that includes a space must be enclosed in double quotes.

`volume-group-name`
The current name of the volume group.

Example Change the name of `VGroup1` to `MyVGroup`:

`# set volume-group name MyVGroup VGroup1`

See also • show volume-groups

show advanced-settings

Description Shows the settings for advanced system-configuration parameters.

Level Monitor

Syntax `show advanced-settings`

Output `Storage-Pool-Component Background Scrub`
Shows whether disks in storage-pool components are automatically checked for disk defects to ensure system health. The interval between background component scrub finishing and starting again is specified by the `Storage-Pool-Component Background Scrub Interval` field.

- `Disabled`: Background component scrub is disabled. This is the default.
- `Enabled`: Background component scrub is enabled.

`Storage-Pool-Component Background Scrub Interval`
Shows the interval between background component scrub finishing and starting again, from 1–360 hours. The default is 24 hours.

`Partner Firmware Upgrade`
Shows whether component firmware versions are monitored and will be automatically updated on the partner controller.

- `Disabled`: Partner firmware upgrade is disabled.
- `Enabled`: Partner firmware upgrade is enabled. This is the default.

`Utility Priority`
Priority at which data-redundancy utilities, such as verify and reconstruct for storage-pool components, run with respect to I/O operations competing for the system's processors. (This does not affect background scrub of storage-pool components, which always runs at "background" priority.)

- `High`: Utilities have higher priority than host I/O. This can cause heavy I/O to be slower than normal. This is the default.
- `Medium`: Utility performance is balanced with host I/O performance.
- `Low`: Utilities run at a slower rate with minimal effect on host I/O.

`SMART`
Shows whether SMART (Self-Monitoring Analysis and Reporting Technology) is enabled or disabled for disks.

- `Detect-Only`: Each disk in the system retains its individual SMART setting, as will new disks added to the system.
- `Enabled`: SMART is enabled for all disks in the system and will be enabled for new disks added to the system. This is the default.
- `Disabled`: SMART is disabled for all disks in the system and will be disabled for new disks added to the system.

`Enclosure Polling Rate`
Shows the interval in seconds at which the storage system will poll each enclosure's Enclosure Management Processor (EMP) for status changes, from 5–3600 seconds. The default is 5 seconds.

`Host Control of Caching`
Shows whether hosts are allowed to use the `SCSI MODE SELECT` command to change the storage system's write-back cache setting.

- `Disabled`: Host control of caching is disabled. This is the default.
- `Enabled`: Host control of caching is enabled.

Sync Cache Mode
Shows how the `SCSI SYNCHRONIZE CACHE` command is handled:

- `Immediate`: Good status is returned immediately and cache content is unchanged.
- `Flush To Disk`: Good status is returned only after all write-back data for the specified volume is flushed to disk.

Missing LUN Response
Shows whether host drivers may probe for LUNs until the host drivers reach the LUN to which they have access.

- `Not Ready`: Sends a reply that there is a LUN where a gap has been created but that it's "not ready." Sense data returned is sensekey = 2, code = 4, qualifier = 3. This is the default.
- `Illegal Request`: Sends a reply that there is a LUN but that the request is "illegal." Sense data returned is sensekey = 5, code = 25h, qualifier = 0.

Controller Failure
Shows whether the cache policy will change from write-back to write-through when a controller fails.

- `Disabled`: The controller failure trigger is disabled. This is the default.
- `Enabled`: The controller failure trigger is enabled.

Supercap Failure
Shows whether the cache policy will change from write-back to write-through when the supercapacitor that provides backup power for cache is not fully charged or fails.

- `Disabled`: The super-capacitor failure trigger is disabled.
- `Enabled`: The super-capacitor failure trigger is enabled. This is the default.

CompactFlash Failure
Shows whether the cache policy will change from write-back to write-through when CompactFlash memory is not detected during POST (Power-On Self-Test), fails during POST, or fails during controller operation.

- `Disabled`: The CompactFlash failure trigger is disabled.
- `Enabled`: The CompactFlash failure trigger is enabled. This is the default.

Power Supply Failure
Shows whether the cache policy will change from write-back to write-through when a power supply fails.

- `Disabled`: The power-supply failure trigger is disabled. This is the default.
- `Enabled`: The power-supply failure trigger is enabled.

Fan Failure
Shows whether the cache policy will change from write-back to write-through when a fan fails.

- `Disabled`: The fan failure trigger is disabled. This is the default.
- `Enabled`: The fan failure trigger is enabled.

Temperature Exceeded
Shows whether the system will shut down a controller when its temperature exceeds the critical operating range.

- `Disabled`: The over-temperature failure trigger is disabled. This is the default.
- `Enabled`: The over-temperature failure trigger is enabled.

`Partner Notify`
Shows whether the partner controller will be notified when a trigger condition occurs.

- `Disabled`: Notification is disabled; the partner controller will continue using its current caching mode. This is the default.
- `Enabled`: Notification is enabled; the partner controller will change to write-through mode for better data protection.

`Auto Write Back`
Shows whether the cache mode will change from write-through to write-back when the trigger condition is cleared.

- `Disabled`: Auto-write-back is disabled.
- `Enabled`: Auto-write-back is enabled. This is the default.

`Idle Drive Spin Down`
Shows whether available disks and spares will spin down after a period of inactivity shown by the `Idle Drive Spin Down Delay` field.

- `Disabled`: Drive spin down for available disks and spares is disabled. This is the default.
- `Enabled`: Drive spin down for available disks and spares is enabled.

`Idle Drive Spin Down Delay`
Shows the period of inactivity in minutes after which available disks and spares will spin down, from 1–360 minutes. The default is 15 minutes. The value 0 means spin down is disabled.

`Disk Background Scrub`
Shows whether disks that are not in storage-pool components are automatically checked for disk defects to ensure system health. The interval between background disk scrub finishing and starting again is 24 hours.

- `Disabled`: Background disk scrub is disabled. This is the default.
- `Enabled`: Background disk scrub is enabled.

`Managed Logs`
Shows whether the managed logs feature is enabled, which allows log files to be transferred from the storage system to a log collection system to avoid losing diagnostic data as logs fill.

- `Disabled`: The managed logs feature is disabled. This is the default.
- `Enabled`: The managed logs feature is enabled.

Example Show advanced system-configuration settings:

```
# show advanced-settings
```

See also • set advanced-settings

show auto-write-through-trigger

Description Shows the system's write-through trigger settings. When a trigger condition occurs and the trigger is enabled, the RAID controller cache mode changes from write-back to write-through.

Level Monitor

Syntax `show auto-write-through-trigger`

Output `Controller Failure`
Shows whether the cache policy will change from write-back to write-through when a controller fails.

- `Disabled`: The controller failure trigger is disabled. This is the default.
- `Enabled`: The controller failure trigger is enabled.

`Supercap Failure`
Shows whether the cache policy will change from write-back to write-through when the supercapacitor that provides backup power for cache is not fully charged or fails.

- `Disabled`: The super-capacitor failure trigger is disabled.
- `Enabled`: The super-capacitor failure trigger is enabled. This is the default.

`CompactFlash Failure`
Shows whether the cache policy will change from write-back to write-through when CompactFlash memory is not detected during POST (Power-On Self-Test), fails during POST, or fails during controller operation.

- `Disabled`: The CompactFlash failure trigger is disabled.
- `Enabled`: The CompactFlash failure trigger is enabled. This is the default.

`Power Supply Failure`
Shows whether the cache policy will change from write-back to write-through when a power supply fails.

- `Disabled`: The power-supply failure trigger is disabled. This is the default.
- `Enabled`: The power-supply failure trigger is enabled.

`Fan Failure`
Shows whether the cache policy will change from write-back to write-through when a fan fails.

- `Disabled`: The fan failure trigger is disabled. This is the default.
- `Enabled`: The fan failure trigger is enabled.

`Temperature Exceeded`
Shows whether the system will shut down a controller when its temperature exceeds the critical operating range.

- `Disabled`: The over-temperature failure trigger is disabled. This is the default.
- `Enabled`: The over-temperature failure trigger is enabled.

`Partner Notify`
Shows whether the partner controller will be notified when a trigger condition occurs.

- `Disabled`: Notification is disabled; the partner controller will continue using its current caching mode. This is the default.
- `Enabled`: Notification is enabled; the partner controller will change to write-through mode for better data protection.

`Auto Write Back`
Shows whether the cache mode will change from write-through to write-back when the trigger condition is cleared.

- `Disabled`: Auto-write-back is disabled.
- `Enabled`: Auto-write-back is enabled. This is the default.

Example Show the system's auto-write-through trigger settings:

```
# show auto-write-through-trigger
```

See also • set auto-write-through-trigger

show cache-parameters

Description   Shows cache settings and status for the system and optionally for a volume.

Level   Monitor

Syntax   `show cache-parameters [`*`volume`*`]`

Parameters   *volume*
Optional. Name or serial number of the volume for which to show settings. For volume syntax, see
Command syntax on page 18. If this parameter is not specified, only system-wide settings are shown.

Output   **System/controller cache parameters:**

`Operation Mode`
Shows the system's operating mode, also called the cache redundancy mode:

- `Active-Active ULP`: Both controllers are active using ULP (Unified LUN Presentation). Data
  for volumes configured to use write-back cache is automatically mirrored between the two
  controllers to provide fault tolerance.
- `Fail Over`: Operation has failed over to one controller because its partner is not operational.
  The system has lost redundancy.
- `Down`: Both controllers are not operational.

`Write Back Status`
Shows the current, system-wide cache policy as determined by auto-write-through logic. This value is
not settable by users. If an auto-write-through trigger condition (such as a CompactFlash failure) is
met, the cache policy for all volumes changes to write-through, overriding the volume-specific
settings. When the problem is corrected, the cache policy reverts to the value configured for each
individual volume.

- `Enabled`: Write-back. This is the normal state.
- `Disabled`: Write-through.
- `Not up`: The controller is not up.

`CompactFlash Status`

- `Not Installed`: The CompactFlash card is not installed.
- `Installed`: The CompactFlash card is installed.
- `Unknown`: The CompactFlash card's status is unknown.

`CompactFlash Health`

- `OK`
- `Fault`

`Cache Flush`

- `Enabled`: If the controller loses power, it will automatically write cache data to the
  CompactFlash card. Cache flush is normally enabled, but is temporarily disabled during
  controller shut down.
- `Disabled`: Cache flush is disabled.

Volume cache parameters:

`Serial Number`
If a volume is specified, its serial number.

`Name`
If a volume is specified, its name.

```
Cache Write Policy
```
If a volume is specified, its cache write policy:

- `write-back`: Write-back caching does not wait for data to be completely written to disk before signaling the host that the write is complete. This is the default and preferred setting for a fault-tolerant environment because it improves the performance of write operations and throughput.
- `write-through`: Write-through caching significantly impacts performance by waiting for data to be completely written to disk before signaling the host that the write is complete. Use this setting only when operating in an environment with low or no fault tolerance.

```
Cache Optimization
```
If a volume is specified, its cache optimization mode:

- `standard`: Optimizes cache for both sequential and random reads. Appropriate for applications that read and write small files in random order, such as transaction-based and database update applications. This is the default.
- `no-mirror`: When this mode is enabled, each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. ULP behavior is not affected, with the exception that during failover any write data in cache will be lost.

```
Read Ahead Size
```
If a volume is specified, its read-ahead cache setting:

- `Disabled`: Read-ahead caching is disabled.
- `Default`: One chunk for the first access in a sequential read and one stripe for all subsequent accesses.
- `Maximum`: Maximum read-ahead size calculated by the controller.
- `64 KB`, `128 KB`, `256 KB`, `512 KB`, `1 MB`, `2 MB`, `4 MB`, `8 MB`, `16 MB`, or `32 MB`: Size selected by a user.

Example Show the cache parameters for the system and for volume `V1`:

```
# show cache-parameters V1
```

See also
- set cache-parameters
- show volumes

show certificate

Description Shows the status of the system's security certificate.

Level Manage

Syntax `show certificate`

Output `Certificate Status`

- `Default`: No customer-supplied or alternate certificate has been installed.
- `Customer-supplied`: A custom or alternate certificate has been installed.

`Time Created`
Date and time in the format $year\text{-}month\text{-}day\ hour\text{:}minutes\text{:}seconds$ when the custom certificate was created.

Example Show certificate status for the system:

```
# show certificate
```

See also • create certificate

show channels

Alias for the show ports command.

show cli-parameters

Description Shows the current CLI session preferences.

Level Monitor

Syntax `show cli-parameters`

Output `Timeout`
Time in seconds that the session can be idle before it automatically ends. Valid values are 120–43200 seconds (2–720 minutes). The default is 1800 seconds (30 minutes).

`Output Format`

- `Console`: Supports interactive use of the CLI by displaying command output in easily readable format. This format automatically sizes fields according to content and adjusts content to window resizes. This is the default.
- `api`: Supports scripting by displaying command output in XML. All objects are displayed at the same level, related by `COMP` elements.
- `api-embed`: Alternate form of XML output which displays "child" objects embedded (indented) under "parent" objects.
- `ipa`: Alternate form of XML output for internal use only.
- `json`: Alternate data-interchange format for internal use only.

`Brief Mode`

- `Enabled`: In XML output, shows a subset of attributes of object properties. The `name` and `type` attributes are always shown.
- `Disabled`: In XML output, shows all attributes of object properties. This is the default.

`Base`
Base for entry and display of storage-space sizes:

- `2`: Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
- `10`: Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is the default.

Operating systems usually show volume size in base 2. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2.

`Pager`

- `Enabled`: Halts output after each full screen to wait for keyboard input. This is the default.
- `Disabled`: Output is not halted. When displaying output in XML API format, which is intended for scripting, disable paging.

`Locale`
Display language. The default is English.

`Precision`
Number of decimal places (1–10) shown for display of storage-space sizes. The default is 1.

`Units`
Unit for display of storage-space sizes:

- `Auto`: Sizes are shown in units determined by the system. This is the default.
- `MB`: Sizes are shown in megabytes.
- `GB`: Sizes are shown in gigabytes.
- `TB`: Sizes are shown in terabytes.

Based on the precision setting, if a size is too small to meaningfully display in the selected unit, the system uses a smaller unit for that size. For example, if `Units` is set to TB, `Precision` is set to 1, and `Base` is set to 10, the size 0.11709 TB is instead shown as 117.1 GB.

```
Temperature Scale
```
- Fahrenheit: Temperatures are shown in degrees Fahrenheit.
- Celsius: Temperatures are shown in degrees Celsius. This is the default.

Example Show current CLI settings:

```
# show cli-parameters
```

See also • set cli-parameters

show configuration

Description Shows system configuration information.

Level Monitor

Syntax `show configuration`

Output
- System information from `show system` command.
- Controller information from the `show controllers` command.
- Controller firmware and hardware version information from the `show versions` command with the `details` parameter.
- Controller host port information from the `show ports` command.
- Disk information from the `show disks` command.
- Disk information by enclosure from the `show disks` command with the `encl` parameter.
- Storage-pool-component information from the `show storage` command.
- Storage-pool information from the `show storage-pools` command.
- Storage-pool-component status information from the `show storage-status` command.
- Enclosure status information from the `show enclosures` command.
- Field-replaceable unit (FRU) information from the `show frus` command.

Example Show system configuration information:

```
# show configuration
```

show controller-date

Description   Shows the system's current date and time.

Level   Monitor

Syntax   `show controller-date`

Output   `Controller Date`
Date and time in the format *year-month-day hour:minutes:seconds.*

`Time-Zone Offset`
The system's time zone as an offset in hours and minutes from Coordinated Universal Time (UTC). This is shown only if NTP is enabled.

Example   Show the date and time:

`# show controller-date`

See also   •   set controller-date
            •   show ntp-status

show controllers

Description Shows information about each controller module in the storage system.

Level Monitor

Syntax `show controllers`

Output `Controller ID`

- `A`: Controller A.
- `B`: Controller B.

`Serial Number`

- Serial number of the controller module.
- `Not Available`: The controller module is down or not installed.

`Hardware Version`
Controller module hardware version.

`CPLD Version`
Complex Programmable Logic Device firmware version.

`MAC Address`
Controller network port MAC address.

`WWNN`
Storage system World Wide Node Name (WWNN).

`IP Address`
Controller network port IP address.

`IP Subnet Mask`
Controller network port IP subnet mask.

`IP Gateway`
Controller network port gateway IP address.

`Disks`
Number of disks in the storage system.

`Storage Pools`
Number of storage pools in the system.

`Cache Memory Size (MB)`
Controller module cache memory size (MB).

`Host Ports`
Number of host ports in the controller module.

`Disk Channels`
Number of expansion ports in the controller enclosure.

`Disk Bus Type`
Type of interface between the controller module and disks:

- `SAS`

`Status`

- `Operational`
- `Down`
- `Not Installed`
- `Unknown`

`Failed Over to This Controller`
Indicates whether the partner controller has failed over to this controller:

- `No`: The partner controller has not failed over to this controller.
- `Yes`: The partner controller has either failed or been shut down and its responsibilities have been taken over by this controller. There will be a delay between the time that the value of `Status` becomes `Down` for one controller and the time that the value of `Failed Over to This Controller` becomes `Yes` for the other controller. This time period is the time that it takes for a controller to take over the responsibilities of its partner.

`Fail Over Reason`
If `Failed Over` is `Yes`, a reason for the failover appears; otherwise, `Not applicable` appears.

`Health`

- `OK`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

`Position`
Position of the controller in the enclosure:

- `Top`: The controller is in the top slot.
- `Bottom`: The controller is in the bottom slot.

`Phy Isolation`
Shows whether the automatic disabling of SAS expander PHYs having high error counts is enabled or disabled for this controller.

- `Enabled`: PHY fault isolation is enabled. This is the default.
- `Disabled`: PHY fault isolation is disabled.

`Controller Redundancy Mode`
Shows the system's operating mode, also called the cache redundancy mode:

- `Active-Active ULP`: Both controllers are active using ULP (Unified LUN Presentation). Data for volumes configured to use write-back cache is automatically mirrored between the two controllers to provide fault tolerance.
- `Fail Over`: Operation has failed over to one controller because its partner is not operational. The system has lost redundancy.
- `Down`: Both controllers are not operational.

`Controller Redundancy Status`

- `Redundant`: Both controllers are operational.
- `Operational but not redundant`: In active-active mode, one controller is operational and the other is offline.
- `Down`: This controller is not operational.
- `Unknown`: Status information is not available.

Example Show controller information:

```
# show controllers
```

See also
- show configuration
- show frus

show controller-statistics

Description Shows live performance statistics for controller A, B, or both.

Level Monitor

Syntax `show controller-statistics [a|b|both]`

Parameters `a|b|both`
Optional. Specifies whether to show information for controller A, B, or both. If this parameter is omitted, information is shown for both controllers.

Output `Durable ID`
Controller ID in the form `controller_ID`.

`CPU Load`
Percentage of time the CPU is busy, from 0–100.

`Power On Time (Secs)`
Number of seconds since the controller was restarted.

`Bytes per second`
Data transfer rate calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`IOPS`
Input/output operations per second, calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`Number of Reads`
Number of read operations since these statistics were last reset or since the controller was restarted.

`Number of Writes`
Number of write operations since these statistics were last reset or since the controller was restarted.

`Data Read`
Amount of data read since these statistics were last reset or since the controller was restarted.

`Data Written`
Amount of data written since these statistics were last reset or since the controller was restarted.

`Num Forwarded Cmds`
The current count of commands that are being forwarded or are queued to be forwarded to the partner controller for processing. This value will be zero if no commands are being forwarded or are queued to be forwarded.

`Reset Time`
Date and time, in the format `year-month-day hour:minutes:seconds`, when these statistics were last reset, either by a user or by a controller restart.

`Total Power On Hours`
The total amount of hours the controller has been powered on in its life time.

Example Show statistics for controller A:

`# show controller-statistics A`

See also
- reset all-statistics
- reset controller-statistics
- show disk-statistics
- show host-port-statistics
- show volume-statistics

show debug-log-parameters

Description Shows which debug message types are enabled (On) or disabled (Off) for inclusion in the Storage Controller debug log. **For use by or with direction from technical support.**

Level Monitor

Syntax `show debug-log-parameters`

Output
- `host`: Host interface debug messages. Enabled by default.
- `disk`: Disk interface debug messages. Enabled by default.
- `mem`: Internal memory debug messages. Disabled by default.
- `fo`: Failover and recovery debug messages. Enabled by default.
- `msg`: Inter-controller message debug messages. Enabled by default.
- `ioa`: I/O interface driver debug messages (standard). Enabled by default.
- `iob`: I/O interface driver debug messages (resource counts). Disabled by default.
- `ioc`: I/O interface driver debug messages (upper layer, verbose). Disabled by default.
- `iod`: I/O interface driver debug messages (lower layer, verbose). Disabled by default.
- `misc`: Internal debug messages. Enabled by default.
- `rcm`: Removable-component manager debug messages. Disabled by default.
- `raid`: RAID debug messages. Enabled by default.
- `cache`: Cache debug messages. Enabled by default.
- `emp`: Enclosure Management Processor debug messages. Enabled by default.
- `capi`: Internal Configuration API debug messages. Enabled by default.
- `mui`: Internal service interface debug messages. Enabled by default.
- `bkcfg`: Internal configuration debug messages. Enabled by default.
- `awt`: Auto-write-through cache triggers debug messages. Disabled by default.
- `res2`: Internal debug messages. Disabled by default.
- `capi2`: Internal Configuration API tracing debug messages. Disabled by default.
- `dms`: Snapshot feature debug messages. Enabled by default.
- `fruid`: FRU ID debug messages. Enabled by default.
- `resmgr`: Reservation Manager debug messages. Disabled by default.
- `init`: Not used.
- `ps`: Paged storage debug messages. Disabled by default.
- `hb`: Not used.

Example Show debug log parameters:

```
# show debug-log-parameters
```

See also
- set debug-log-parameters

show disk-parameters

Description   Shows disk settings.

Level   Monitor

Syntax   `show disk-parameters`

Output   `SMART`
Shows whether SMART (Self-Monitoring Analysis and Reporting Technology) is enabled or disabled for disks.

- `Detect-Only`: Each disk in the system retains its individual SMART setting, as will new disks added to the system.
- `Enabled`: SMART is enabled for all disks in the system and will be enabled for new disks added to the system. This is the default.
- `Disabled`: SMART is disabled for all disks in the system and will be disabled for new disks added to the system.

`Drive Write Back Cache`

- `Disabled`: Disk write-back cache is disabled for all disks in the system and will be disabled for new disks added to the system. This parameter cannot be changed.

`Timeout Retry Maximum`
Maximum number of times a timed-out I/O operation can be retried before the operation is failed.

`Attempt Timeout`
Number of seconds before an I/O operation is aborted and possibly retried.

Overall Timeout
Total time in seconds before an I/O operation is failed regardless of the `Attempt Timeout` and `Number of Retries` settings.

`Idle Drive Spin Down`
Shows whether available disks and spares will spin down after a period of inactivity shown by the `Idle Drive Spin Down Delay` field.

- `Disabled`: Drive spin down for available disks and spares is disabled. This is the default.
- `Enabled`: Drive spin down for available disks and spares is enabled.

`Idle Drive Spin Down Delay`
Shows the period of inactivity in minutes after which available disks and spares will spin down, from 1–360 minutes. The default is 15 minutes. The value 0 means spin down is disabled.

Example   Show disk settings:

```
# show disk-parameters
```

See also   • set disk-parameters

show disks

Description Shows information about disks or disk slots in the storage system. If no parameter is specified, the command shows information for all installed disks.

> **NOTE:** In console format, to aid reading, disks are sorted to display in order by enclosure and disk number. In API formats, output is not sorted because it is expected to be manipulated by a host application.

Level Monitor

Syntax
```
show disks
    [storage-pool storage-pool-ID]
    [disks]
    [free]
    [all]
    [perf]
    [encl]
```

Parameters storage-pool *storage-pool-ID*
Optional. The name or serial number of the storage pool that contains the disks about which to show information. Names are case sensitive.

disks
Optional. The IDs or serial numbers of disks about which to show information. For disk syntax, see Command syntax on page 18.

free
Optional. Shows information about all available disks. Do not specify this parameter with other parameters.

all
Optional. Shows information about all installed disks. Do not specify this parameter with other parameters.

perf
Optional. For all or specified disks, this parameter shows performance statistics from the latest historical sample for each disk. Statistics shown include total I/Os (reads and writes), total amount of data transferred, and average I/O response time.

encl
Optional. Shows information about each disk slot, whether it contains a disk or not. Do not specify this parameter with other parameters; they will be ignored.

Output **Without the `encl` parameter:**

The disk location in the format *enclosure-number.disk-slot-number*.

Serial Number
The disk serial number

Vendor
The disk vendor.

Rev
The firmware revision number.

Type

- SAS: Dual-port, spinning, enterprise SAS disk.
- SAS MDL: Dual-port, spinning, midline SAS disk.
- sSAS: Dual-port, solid-state SAS disk (SSD).

Speed (kr/min)
The speed of a spinning disk, in thousands of revolutions per minute, as specified by the disk vendor. For an SSD, 0 is shown.

Size
The disk capacity, formatted to use the current base, precision, and units.

Rate*(Gb/s)
The data transfer rate in Gbit/s. A footnote indicates that it is normal behavior for the rate to vary.

Some 6-Gbit/s disks might not consistently support a 6-Gbit/s transfer rate. If this happens, the controller automatically adjusts transfers to those disks to 3 Gbit/s, increasing reliability and reducing error messages with little impact on system performance. This rate adjustment persists until the controller is restarted or power-cycled.

Usage

- `AVAIL`: The disk is available.
- `SPARE`: The disk is configured as a spare.
- `STORAGE`: The disk is part of a storage-pool component.
- `FAILED`: The disk is unusable and must be replaced. Reasons for this status include: excessive media errors; SMART error; disk hardware failure; unsupported disk.
- `LEFTOVR`: The disk is part of a storage-pool component that is not found in the system.

Recon State

- `From`: This disk is being used as the source of a reconstruct operation.
- `To`: This disk is being used as the target of a reconstruct operation.
- `N/A`: This disk is not being used in a reconstruct operation.

Copyback State

- `From`: This disk is being used as the source of a copyback operation.
- `To`: This disk is being used as the target of a copyback operation.
- `N/A`: This disk is not being used in a copyback operation.

Tier

- `Performance`: The disk is in the highest storage tier, which uses SSDs (high speed, low capacity).
- `Standard`: The disk is in the storage tier that uses enterprise-class spinning SAS disks (lower speed, higher capacity).
- `Archive`: The disk is in the lowest storage tier, which uses midline spinning SAS disks (low speed, high capacity).
- `Read Cache`: The disk is an SSD providing high-speed read cache for a storage pool.

Pool
The name of the storage pool that contains the disk.

Operation

- `CPYBK`: The disk is being used in a copyback operation.
- `DRSC`: The disk is being scrubbed.
- `IDLE`: The disk has no active operation.
- `INIT`: The disk's storage-pool component is being initialized.
- `POOL`: The disk is part of a storage-pool component.
- `RCON`: The disk's storage-pool component is being reconstructed.
- `VRFY`: The disk's storage-pool component is being verified.

Life Remaining%

- 100%–0%: The percentage of life remaining for the disk, which is an SSD. As wear increases, this value will decrease from 100% to 0%, at which time the disk will be unusable.
- `N/A`: The disk is not an SSD.

Total I/Os
If the `perf` parameter is specified, this field shows the total number of I/O operations (reads and writes).

Data Transferred
If the `perf` parameter is specified, this field shows the total number of bytes transferred.

I/O Resp Time
If the `perf` parameter is specified, this field shows the average time in microseconds to complete an I/O operation.

Health

- `OK`
- `Degraded`
- `Fault`
- `N/A`
- `Unknown`

Health Reason
If Health is not `OK`, this field shows the reason for the health state.

Health Recommendation
If Health is not `OK`, this field shows recommended actions to take to resolve the health problem.

With the `encl` parameter:

Status

- `Up`: The disk is present and is properly communicating with the expander.
- `Spun Down`: The disk is present and has been spun down by the drive spin down feature.
- `Warning`: The disk is present but the system is having communication problems with the disk LED processor. For disk and midplane types where this processor also controls power to the disk, power-on failure will result in `Error` status.
- `Error`: The disk is present but is not detected by the expander.
- `Unknown`: Initial status when the disk is first detected or powered on.
- `Not Present`: The disk slot indicates that no disk is present.

Encl
The number of the enclosure that contains the disk.

Slot
The number of the disk slot that contains the disk.

Vendor
The disk vendor.

Model
The disk model number.

Serial Number
The disk serial number.

Size
The disk capacity, formatted to use the current base, precision, and units.

Example Show information about all disks:

```
# show disks
```

Show information about all disk slots:

```
# show disks encl
```

Show performance statistics for all disks in storage pool A:

```
# show disks storage-pool A perf
```

Show information about disks 1.1, 1.2, 1.3, and 1.5:

```
# show disks 1.1-3,1.5
```

See also • show storage
 • show storage-status

show disk-statistics

Description Shows live or historical performance statistics for disks. You can view live statistics for all or specified disks, or historical statistics for a specified disk. The system samples disk-performance statistics every quarter hour and retains performance data for 6 months.

The historical option allows you to specify a time range or a number (count) of data samples to include. It is not recommended to specify both the `time-range` and `count` parameters; if both parameters are specified, and more samples exist for the specified time range, the samples' values will be aggregated to show the required number of samples.

Level Monitor

Syntax To show live statistics:

```
show disk-statistics [disks]
```

To show historical statistics:

```
show disk-statistics
    [disk]
    [historical]
    [time-range "date/time-range"]
    [count number-of-data-samples]
    [all]
```

Parameters *disk*
Optional. Identifies one or more disks for which to show live statistics. If this parameter is omitted, statistics will be shown for all disks. For disk syntax, see Command syntax on page 18.

`historical`
Optional. Specifies to show historical statistics. If this parameter is omitted, live statistics will be shown.

`time-range "date/time-range"`
Optional. Specifies the date/time range of historical statistics to show, in the format `"start yyyy-mm-dd hh:mm [AM|PM] end yyyy-mm-dd hh:mm [AM|PM]"`. If the start date/time is specified but no end date/time is specified, the current date/time will be used as the end date/time. The system will return the oldest sample taken after the start time and the latest sample taken before the end time. If the specified start date/time is earlier than the oldest sample, that sample will be used as the start date/time. If you specify this parameter, do not specify the `count` parameter. If this parameter is omitted, the most recent 100 data samples will be displayed.

`count number-of-data-samples`
Optional. Specifies the number of data samples to display, from 1–100. Each sample will be shown as a separate row in the command output. If this parameter is omitted, 100 samples will be shown. If you specify this parameter, do not specify the `time-range` parameter.

`all`
Optional. Specifies to show the full set of performance metrics. If this parameter is omitted, the default set of performance metrics will be shown.

Output
(Live) `Durable ID`
Disk ID in the form disk_*enclosure-number.disk-number*.

`Serial Number`
Disk serial number.

`Bytes per second`
Data transfer rate calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

IOPS
Input/output operations per second, calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

Number of Reads
Number of read operations since these statistics were last reset or since the controller was restarted.

Number of Writes
Number of write operations since these statistics were last reset or since the controller was restarted.

Data Read
Amount of data read since these statistics were last reset or since the controller was restarted.

Data Written
Amount of data written since these statistics were last reset or since the controller was restarted.

Reset Time
Date and time, in the format *year-month-day hour*:*minutes*:*seconds*, when these statistics were last reset, either by a user or by a controller restart.

Output (Historical)

Durable ID
Disk ID in the form disk_*enclosure-number*.*disk-number*.

Serial Number
Disk serial number.

Total I/Os
Total number of read and write operations since the last sampling time.

Number of Reads
Shown by the all parameter. Number of read operations since the last sampling time.

Number of Writes
Shown by the all parameter. Number of write operations since the last sampling time.

Data Transferred
Total amount of data read and written since the last sampling time.

Data Read
Shown by the all parameter. Amount of data read since the last sampling time.

Data Written
Shown by the all parameter. Amount of data written since the last sampling time.

Total IOPS
Total number of read and write operations per second since the last sampling time.

Read IOPS
Shown by the all parameter. Number of read operations per second since the last sampling time.

Write IOPS
Shown by the all parameter. Number of write operations per second since the last sampling time.

Total Bps
Total data transfer rate, in bytes per second, since the last sampling time.

Read Bps
Shown by the all parameter. Data transfer rate, in bytes per second, for read operations since the last sampling time.

Write Bps
Shown by the all parameter. Data transfer rate, in bytes per second, for write operations since the last sampling time.

`Queue Depth`
Shown by the `all` parameter. Average number of pending read and write operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity.

`I/O Resp Time`
Average response time, in microseconds, for read and write operations since the last sampling time.

`Read Resp Time`
Shown by the `all` parameter. Average response time, in microseconds, for read operations since the last sampling time.

`Write Resp Time`
Shown by the `all` parameter. Average response time, in microseconds, for write operations since the last sampling time.

`Average I/O Size`
Shown by the `all` parameter. Average data size of read and write operations since the last sampling time.

`Average Read I/O Size`
Shown by the `all` parameter. Average data size of read operations since the last sampling time.

`Average Write I/O Size`
Shown by the `all` parameter. Average data size of write operations since the last sampling time.

`Number of Disk Errors`
Shown by the `all` parameter. Total number of disk errors detected since the last sampling time. Error types include: number of SMART events; number of timeouts accessing the disk; number of times the disk did not respond; number of attempts by the storage system to spin-up the disk; media errors generated by the disk as specified by its manufacturer; non-media errors (generated by the storage system, or by the disk and not categorized as media errors); number of bad-block reassignments.

`Sample Time`
Date and time, in the format *year-month-day hour*:*minutes*:*seconds*, when the data sample was taken.

Example   Show live statistics for disks 1.1 and 2.1:

`# show disk-statistics 1.1,2.1`

Show historical statistics from a specified date and time range for disk 1.5:

`# show disk-statistics 1.5 historical time-range "start 2012-12-05 4:40 PM end 2012-12-05 5:00 PM"`

Show all samples of historical statistics for disk 1.5:

`# show disk-statistics 1.5 historical all`

See also   • reset all-statistics
             • reset disk-error-statistics
             • reset disk-statistics
             • show controller-statistics
             • show disks
             • show host-port-statistics
             • show volume-statistics

show email-parameters

Description Shows email (SMTP) notification parameters for events and managed logs.

Level Monitor

Syntax `show email-parameters`

Output `Email Notification`

- `Disabled`: Email notification is disabled. This is the default.
- `Enabled`: Email notification is enabled.

`Email Notify Filter`

Shows the minimum severity for which the system should send event notifications:

- `crit`: Sends notifications for Critical events only.
- `error`: Sends notifications for Error and Critical events.
- `warn`: Sends notifications for Warning, Error, and Critical events.
- `info`: Sends notifications for all events.
- `none`: Disables email notification and clears the settings. This is the default.

This parameter does not apply to managed-logs notifications.

`Email Address (1-3)`
Shows up to three email addresses for recipients of event notifications.

`Log Destination`
Shows the email address for the log collection system used by the managed logs feature.

`Email Server`
The IP address of the SMTP mail server to use for the email messages.

`Email Domain`
The domain name that, with the sender name, forms the "from" address for remote notification.

`Email Sender`
The sender name that, with the domain name, forms the "from" address for remote notification.

`Include Logs`
Shows whether system log files will automatically be attached to email notification messages generated by the managed logs feature. This is the "push" mode for managed logs. This option is disabled by default.

Example Show settings for email notification:

`# show email-parameters`

See also • set email-parameters

show enclosures

Description Shows information about the enclosures in the storage system. Full detail available in XML API output only.

Level Monitor

Syntax `show enclosures`

Output `Encl`
Enclosure ID

`Encl WWN`
Enclosure WWN

`Name`
Enclosure name

`Location`
Enclosure location; blank if not set

`Rack`
Number of the rack containing the enclosure

`Pos`
Position of the enclosure in the rack

`Vendor`
Enclosure vendor

`Model`
Enclosure model

`EMP # BUS:ID Rev`
Address and firmware revision of the Enclosure Management Processor in each controller's Expander Controller

`Midplane Type`

- `2U24-6Gv2`: Midplane for 2U, reduced-depth, 24-disk enclosure with 6-Gbit/s maximum data rate to disks
- `2U12-6Gv2`: Midplane for 2U, reduced-depth, 12-disk enclosure with 6-Gbit/s maximum data rate to disks
- `N/A`: Other type of midplane

`Health`

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show information about enclosures in the system:

```
# show enclosures
```

See also
- set enclosure
- show sensor-status

show enclosure-status

Description Shows the status of system enclosures and their components. For each attached enclosure, the command shows general SCSI Enclosure Services (SES) information followed by component-specific information.

Level Monitor

Syntax `show enclosure-status`

Output **General SES fields:**

`Chassis`
Chassis serial number.

`Vendor`
Enclosure vendor name.

`Product ID`
Product model identifier.

`CPLD`
Complex Programmable Logic Device version.

`EMP # BUS:ID Rev`
Address and firmware revision of the Enclosure Management Processor in each controller's Expander Controller.

`WWPN`
World wide port name of the SES device reporting the enclosure status.

`Status`
Overall status of the enclosure:

- `Absent`: The enclosure is not present.
- `Error`: The enclosure has a fault.
- `OK`: The enclosure is operating normally.

`Not Available`: Status is not available.

Enclosure Component Status fields:

`Type`
Component type:

- `FAN`: Cooling fan unit
- `PSU`: Power supply unit
- `Temp`: Temperature sensor
- `Voltage`: Voltage sensor
- `Disk`: Disk drive module

`#`
Component ID

`Status`
Component status:

- `Absent`: The component is not present.
- Error: The component or at least one subcomponent has failed.
- Warning: The component or at least one subcomponent is not working normally.
- `OK`: The component and any subcomponents are working normally. Temperature status `OK` indicates that the sensor is working properly, not that the temperature is within an acceptable range.
- `N/A`: Status is not available.

`FRU P/N`
Part number of the field-replaceable unit (FRU) that contains the component.

`FRU S/N`
Serial number of the FRU that contains the component.

`Add'l Data`
Additional data, if applicable:

- `addr=`: For a disk, the slot address
- `temp=`: For a temperature sensor, the temperature
- `voltage=`: For a voltage sensor, the voltage
- `--`: No data

Example Show enclosure status:

```
# show enclosure-status
```

See also • show enclosures
- show frus
- show sensor-status

show events

Description Shows events logged by each controller in the storage system. A separate set of event numbers is maintained for each controller. Each event number is prefixed with a letter identifying the controller that logged the event.

Events are listed from newest to oldest, based on a timestamp with one-second granularity; therefore the event log sequence matches the actual event sequence within about one second.

For further information about diagnosing and resolving problems, see:

- The troubleshooting chapter and the LED descriptions appendix in your product's Setup User Guide
- The topics about verifying component failure in your product's FRU Installation and Replacement Guide
- The full list of event codes, descriptions, and recommended actions in the Event Descriptions Reference Guide

Level Monitor

Syntax
```
show events
    [detail]
    [last #]
    [from timestamp]
    [to timestamp]
    [from-event event-ID]
    [to-event event-ID]
    [a|b|both|error]
    [logs yes|no]
```

Parameters detail
Optional. Shows additional information and recommended actions for displayed events. This information is also in the Event Descriptions Reference Guide.

last #
Optional. Shows the latest specified number of events. If this parameter is omitted, all events are shown.

from timestamp
Optional. Shows events including and after a timestamp specified with the format *MMDDYYhhmmss*. For example, 043012235900 represents April 30 2012 at 11:59:00 p.m. This parameter can be used with the to parameter or the to-event parameter.

to timestamp
Optional. Shows events before and including a timestamp specified with the format *MMDDYYhhmmss*. For example, 043012235900 represents April 30 2012 at 11:59:00 p.m. This parameter can be used with the from parameter or the from-event parameter.

from-event event-ID
Optional. Shows events including and after the specified event ID. If this number is smaller than the ID of the oldest event, events are shown from the oldest available event. Events are shown only for the controller that the event ID specifies (A or B). This parameter can be used with the to parameter or the to-event parameter.

to-event event-ID
Optional. Shows events before and including the specified event ID. If this number is larger than the ID of the oldest event, events are shown up to the latest event. Events are shown only for the controller that the event ID specifies (A or B). This parameter can be used with the from parameter or the from-event parameter.

```
a|b|both|error
```
Optional. Specifies to filter the event listing:

- `a`: Shows events from controller A only. Do not use this parameter with the `from-event` parameter or the `to-event` parameter.
- `b`: Shows events from controller B only. Do not use this parameter with the `from-event` parameter or the `to-event` parameter.
- `both`: Shows events from both controllers. Do not use this parameter with the `from-event` parameter or the `to-event` parameter.
- `error`: Shows Warning, Error, and Critical events.

```
logs yes|no
```
Optional.

- `no`: Lists events as described in the Output section, below. This is the default.
- `yes`: Shows events in tabular format, with columns for event ID, date and time, event code, severity, and message.

Output
- Date and time when the event was logged
- Event code identifying the type of event to help diagnose problems; for example, `[181]`
- Event ID prefixed by `A` or `B`, indicating which controller logged the event; for example, `#A123`
- Model, serial number, and ID of the controller module that logged the event
- Severity:
  - `CRITICAL`: A failure occurred that may cause a controller to shut down. Correct the problem *immediately*.
  - `ERROR`: A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
  - `WARNING`: A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct it if necessary.
  - `INFORMATIONAL`: A configuration or state change occurred, or a problem occurred that the system corrected. No action is required.
- Event-specific message giving details about the event

Example Show the last two events:

```
# show events last 2
```
Show the last three non-Informational events:

```
# show events last 3 error
```
Show all events from April 30 2012 at 11:59:00 p.m. through May 2 2012 at 11:59:00 a.m.:

```
# show events from 043012235900 to 050212115900
```
Show a range of events logged by controller A:

```
# show events from-event a100 to-event a123
```
Show detailed output for a specific event:

```
# show events from-event A2264 to-event A2264 detail
```

See also
- clear events
- set snmp-parameters
- show snmp-parameters

show expander-status

Description Shows diagnostic information relating to SAS Expander Controller physical channels, known as PHY lanes. **For use by or with direction from technical support.**

For each enclosure, this command shows status information for PHYs in I/O module A and then I/O module B.

Level Monitor

Syntax `show expander-status`

Output `Encl`
Enclosure that contains the SAS expander.

`Ctlr`
I/O module that contains the SAS expander.

`Phy`
Identifies a PHY's logical location within a group based on the PHY type. Logical IDs are 0–23 for drive PHYs; 0–1 for SC PHYs; and 0–3 for other PHYs. If the PHY's controller module or expansion module is not installed, this field shows "--".

`Type`
- `Drive`: 1-lane PHY that communicates between the expander and a disk drive.
- `Egress`: 4-lane PHY that communicates between the expander and an expansion port or SAS Out port.
- `SC-1`: (Controller module only) 2-lane PHY that communicates between the expander and the partner's expander.
- `SC-0`: (Controller module only) 4-lane PHY that communicates between the expander and the SC.
- `Ingress`: (Expansion module only) 4-lane PHY that communicates between the expander and an expansion port.
- `Inter-Exp`: (Expansion module only) Communicates between the expander and the partner's expander.
- `Undefined`: No status information is available.
- `Unused`: The PHY exists in the expander but is not connected, by design.

`Status`
- `Enabled - Healthy`: The PHY is enabled and healthy.
- `Enabled - Degraded`: The PHY is enabled but degraded.
- `Disabled`: The PHY has been disabled by a user or by the system.

`Elem Status`
A standard SES status for the element:

- `Disabled`: Critical condition is detected.
- `Error`: Unrecoverable condition is detected. Appears only if there is a firmware problem related to PHY definition data.
- `Non-critical`: Non-critical condition is detected.
- `Not Used`: Element is not installed in enclosure.
- `OK`: Element is installed and no error conditions are known.
- `Unknown`: Either:
  - Sensor has failed or element status is not available. Appears only if an I/O module indicates it has fewer PHYs than the reporting I/O module, in which case all additional PHYs are reported as unknown.
  - Element is installed with no known errors, but the element has not been turned on or set into operation.

```
Disabled
```

- `Enabled`: PHY is enabled.
- `Disabled`: PHY is disabled.

```
Reason
```

- Blank if `Elem Status` is `OK`.
- `Error count interrupts`: PHY disabled because of error-count interrupts.
- `Phy control`: PHY disabled by a SES control page as a result of action by a Storage Controller or user.
- `Not ready`: PHY is enabled but not ready. Appears for SC-1 PHYs when the partner I/O module is not installed. Appears for Drive, SC-1, or Ingress PHYs when a connection problem exists such as a broken connector.
- `Drive removed`: PHY disabled because drive slot is empty.
- `Unused - disabled by default`: PHY is disabled by default because it is not used.
- `Excessive Phy changes`: PHY is disabled because of excessive PHY change counts.

Example Show expander status:

```
# show expander-status
```

See also • clear expander-status
 • set expander-fault-isolation
 • set expander-phy

show fans

Description Shows information about each fan in the storage system.

Level Monitor

Syntax `show fans`

Output `Name`
Fan name in the form `Fan loc:`*`position`*`-PSU `*`power-supply-ID`*. The position is as viewed from the back of the enclosure.

`Location`
Fan location in the form `Enclosure `*`enclosure-ID - position`*. The position is as viewed from the back of the enclosure.

`Status`

- `Up`
- `Warning`
- `Error`
- `Not Present`
- `Unknown`

`Speed`
Fan speed (RPM).

`Position`
Fan position, as viewed from the back of the enclosure:

- Left
- Right

`Serial Number`

- (blank): Not applicable.

`Firmware Version`

- (blank): Not applicable.

`Hardware Version`

- (blank): Not applicable.

`Health`

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show fan information:

`# show fans`

See also • show frus

- show power-supplies

show frus

Description Shows FRU (field-replaceable unit) information for the storage system. Some information is for use by service technicians.

Level Monitor

Syntax `show frus`

Output **FRU fields:**

`Name`

- `CHASSIS_MIDPLANE`: 2U chassis and midplane circuit board
- `RAID_IOM`: Controller module
- `BOD_IOM`: Expansion module
- `POWER_SUPPLY`: Power supply module

`Description`
FRU description

`Part Number`
FRU part number

`Serial Number`
FRU serial number

`Revision`
Hardware revision level

`Dash Level`
FRU template revision number

`FRU Shortname`
Short description

`Manufacturing Date`
Date and time in the format *year-month-day hour:minutes:seconds* when a PCBA was programmed or a power supply module was manufactured

`Manufacturing Location`
City, state/province, and country where the FRU was manufactured

`Manufacturing Vendor ID`
JEDEC ID of the manufacturer

`FRU Location`
Location of the FRU in the enclosure:

- `MID-PLANE SLOT`: Chassis midplane
- `UPPER IOM SLOT`: Controller module or expansion module A
- `LOWER IOM SLOT`: Controller module or expansion module B
- `LEFT PSU SLOT`: Power supply module on the left, as viewed from the back
- `RIGHT PSU SLOT`: Power supply module on the right, as viewed from the back

`Configuration SN`
Configuration serial number

`FRU Status`

- `Absent`: Component is not present
- `Fault`: One or more subcomponents has a fault
- `OK`: All subcomponents are operating normally
- `Not Available`: Status is not available

Original SN
For a power supply module, the original manufacturer serial number; otherwise, N/A.

Original PN
For a power supply module, the original manufacturer part number; otherwise, N/A.

Original Rev
For a power supply module, the original manufacturer hardware revision; otherwise, N/A.

Enclosure ID
Enclosure ID of the enclosure that contains the FRU.

Example Show information about all FRUs in the system:

```
# show frus
```

show host-groups

Description Shows information about specified host groups or all host groups. For each host group, its member hosts and their initiators are also shown.

Level Monitor

Syntax `show host-groups [groups host-groups]`

Parameters `groups host-groups`
Optional. A comma-separated list of the names of host groups for which to show information. If this parameter is omitted, information is shown for all host groups.

Output **Host group information**

`Name`
The host group name.

`Number of Members`
The number of hosts in the host group.

Host information

`Name`
The host name.

`Number of Members`
The number of initiators in the host.

Initiator information

`Nickname`
The nickname of the initiator.

`Discovered`

- `Yes`: The initiator was discovered and its entry was automatically created.
- `No`: The initiator was manually created.

`Mapped`
Shows whether the initiator is explicitly mapped to any volumes:

- `Yes`: At least one volume is explicitly mapped to the initiator.
- `No`: No volumes are explicitly mapped to the initiator.

`ID`

- For an FC or SAS initiator, its WWPN.

Example Show information about all host groups:

`# show host-groups`

Show information about host groups `HGroup1` and `HGroup3` only:

`# show host-groups groups HGroup1,HGroup3`

See also
- create host-group
- delete host-groups
- set host-group

show host-maps

Description Shows mapping information for host groups, hosts, and initiators.

Level Monitor

Syntax `show host-maps [IDs]`

Parameters *IDs*
Optional. A comma-separated list of the names or serial numbers of initiators, hosts, and host groups for which to show mappings. If this parameter is omitted, all mappings are shown. For name syntax, see Command syntax on page 18.

Output **Host group information**

`Serial Number`
The serial number of the host group.

`Group Name`
The host group name in the format *host-group-name.*.**, where the first * represents all hosts in the group, and the second * represents all initiators in those hosts.

Host information

`Serial Number`
The serial number of the host.

`Group Name`
The host group name in the format *host-name.**, where the represents all initiators in those hosts.

Initiator information

`ID`

• For an FC or SAS initiator, its WWPN.

`Name`
The initiator name.

Mapping information

`Volume`

• For a volume, its name.
• For a volume group, its name in the format *volume-group-name.**, where the * represents all volumes in the group.

`Serial Number`
The serial number of the volume.

`LUN`
The LUN used to access the volume.

`Access`
The type of host access to the volume:

• `read-write`: Read and write. This is the default.
• `read-only`: Read only.
• `no-access`: No access (masked).

`Ports`
The controller host ports to which the mapping applies.

Example Show mappings for all hosts:

`# show host-maps`

Show mappings for host group `HGroup1`:

```
# show host-maps HGroup1.*
```

See also
- show hosts
- show maps
- show volume-maps
- show volumes

show host-port-statistics

Description Shows live performance statistics for each controller host port. For each host port these statistics quantify I/O operations through the port between a host and a volume. For example, each time a host writes to a volume's cache, the host port's statistics are adjusted.

Level Monitor

Syntax `show host-port-statistics [ports ports]`

Parameters `ports ports`
Optional. A comma-separated list of port IDs for which to show information. For port syntax, see Command syntax on page 18. If this parameter is omitted, statistics are shown for all host ports.

Output `Durable ID`
Host port ID in the form `hostport_controller-ID-and-port-number.`

`Bytes per second`
Data transfer rate calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`IOPS`
Input/output operations per second, calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`Number of Reads`
Number of read operations since these statistics were last reset or since the controller was restarted.

`Number of Writes`
Number of write operations since these statistics were last reset or since the controller was restarted.

`Data Read`
Amount of data read since these statistics were last reset or since the controller was restarted.

`Data Written`
Amount of data written since these statistics were last reset or since the controller was restarted.

`Queue Depth`
Number of pending I/O operations being serviced.

`I/O Resp Time`
Average response time in microseconds for read and write operations, calculated over the interval since these statistics were last requested or reset.

`Read Resp Time`
Average response time in microseconds for all read operations, calculated over the interval since these statistics were last requested or reset.

`Write Resp Time`
Average response time in microseconds for all write operations, calculated over the interval since these statistics were last requested or reset.

`Reset Time`
Date and time, in the format `year-month-day hour:minutes:seconds`, when these statistics were last reset, either by a user or by a controller restart.

Example Show host-port statistics:

```
# show host-port-statistics
```

Show host-port statistics for port A1 only:

```
# show host-port-statistics ports a1
```

See also
- reset all-statistics
- reset host-port-statistics
- show controller-statistics
- show disk-statistics
- show volume-statistics
- show host-port-statistics
- show ports

show hosts

Description Shows information about specified hosts or all hosts.

Level Monitor

Syntax
```
show hosts
     [groups host-groups]
     [hosts]
```

Parameters groups *host-groups*
Optional. A comma-separated list of host-group names containing hosts for which to show information. If this parameter is omitted, information is shown for all hosts.

hosts
Optional. A comma-separated list of host names for which to show information. If this parameter is omitted, information is shown for all hosts.

Output

Host information

Name
The host name.

Number of Members
The number of initiators in the host.

Initiator information

Nickname
The nickname of the initiator.

Discovered

- Yes: The initiator was discovered and its entry was automatically created.
- No: The initiator was manually created.

Mapped
Shows whether the initiator is explicitly mapped to any volumes:

- Yes: At least one volume is explicitly mapped to the initiator.
- No: No volumes are explicitly mapped to the initiator.

ID

- For an FC or SAS initiator, its WWPN.

Example Show information about all hosts:

```
# show hosts
```

Show information about host Host3 only:

```
# show hosts Host3
```

Show information about hosts in host groups HGroup1 and HGroup3 only:

```
# show hosts groups HGroup1,HGroup3
```

See also
- create host
- delete hosts
- set host

show initiators

Description Shows information about initiators in specified hosts or all initiators.

Initiator entries are automatically created for host initiators that have sent an `inquiry` command or a `report luns` command to the system. This typically happens when the physical host containing an initiator boots up or scans for devices. When the command is received, the system saves the host port information; however, the information is retained after a restart only if you have set a nickname for the initiator.

Level Monitor

Syntax
```
show initiators
     [initiator-id]
     [hosts host-names]
```

Parameters *initiator-id*
Optional.

hosts *host-names*
Optional. A comma-separated list of the names of hosts containing initiators for which to show information. If this parameter is omitted, information is shown for all initiators.

Output `Nickname`
The nickname of the initiator.

`Discovered`

- `Yes`: The initiator was discovered and its entry was automatically created.
- `No`: The initiator was manually created.

`Mapped`
Shows whether the initiator is explicitly mapped to any volumes:

- `Yes`: At least one volume is explicitly mapped to the initiator.
- `No`: No volumes are explicitly mapped to the initiator.

`ID`

- For an FC or SAS initiator, its WWPN.

Example Show information about all initiators:

```
# show initiators
```

Show information about initiators in hosts `Host1` and `Host3` only:

```
# show initiators hosts Host1,Host3
```

See also • set initiator
   • show hosts

show inquiry

Description Shows inquiry data for each controller module.

Level Monitor

Syntax `show inquiry`

Output
- Product vendor name, product ID, and vendor ID
- Management Controller firmware version and loader version
- Storage Controller firmware version and loader version
- Controller module serial number
- Media Access Control (MAC) address
- Network port IP address

Example Show inquiry data for both controllers:

```
# show inquiry
```

See also
- show versions

show license

Description Shows the status of licensed features in the storage system.

Level Monitor

Syntax `show license`

Output `License Key`

- The license key, if a license is installed and valid.
- `not installed`, if a license is invalid or is not installed.

`Maximum Licensable Snapshots`
Number of snapshots that the highest-level license allows.

`Base Maximum Snapshots`
Number of snapshots allowed without an installed license.

`Licensed Snapshots`
Number of snapshots allowed by the installed license.

`In-Use Snapshots`
Number of existing licensed snapshots.

`Snapshots Expire`

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

`Volume Copy`
Shows whether Volume Copy functions are enabled or disabled.

`Volume Copy Expires`

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

`Replication`
Shows whether AssuredRemote functions are enabled or disabled.

`Replication Expires`

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

`VDS`
Shows whether the VDS (Virtual Disk Service) Hardware Provider is enabled or disabled.

```
VDS Expires
```

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

```
VSS
```
Shows whether the VSS (Volume Shadow Copy Service) Hardware Provider is enabled or disabled.

```
VSS Expires
```

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

```
SRA
```
Shows whether Storage Replication Adapter (SRA) support is enabled or disabled.

```
SRA Expires
```

- `Never`. License is purchased and doesn't expire.
- Number of days remaining for a temporary license.
- `Expired`. Temporary license has expired and cannot be renewed.
- `Expired/Renewable`. Temporary license has expired and can be renewed.
- `N/A`. No license installed.

Example   Show information about the installed license:

```
# show license
```

See also   • create temp-license
   • load license

show maps

Description Shows information about mappings. By default, the command shows mappings whose access type is `read-write`, `read-only`, or `no-access` (but not `not-mapped`) for all volumes.

Level Monitor

Syntax
```
show maps
    [initiator]
    [all]
    [IDs]
```

Parameters **initiator**
Optional. Specifies to show mappings whose access type is `read-write`, `read-only`, or `no-access` (but not `not-mapped`) for all initiators.

all
Optional. Specifies to include `not-mapped` entries in mapping output.

IDs
Optional. A comma-separated list of the names or serial numbers of host-type items (initiators, hosts, and host groups) or volume-type items (volumes and volume groups) for which to show mappings. Do not include both host-type and volume-type items in a list. For name syntax, see Command syntax on page 18.

Output **Without the `initiator` parameter:**

Volume group information:

`Serial Number`
The serial number of the volume group.

`Group Name`
The volume group name in the format *volume-group-name*`.*`, where the * represents all volumes in the group.

`Ports`

- The controller host ports to which the mapping applies.
- Blank if not mapped or mapped as `no-access`.

`Access`
The type of host access to the volume:

- `read-write`: Read and write. This is the default.
- `read-only`: Read only.
- `no-access`: No access (masked).
- `not-mapped`: Not mapped.

`Initiator-Identifier`

- For an FC or SAS initiator, its WWPN.
- `all other initiators`: The volume's default mapping.
- For a host or host group, its serial number.

`Nickname`

- The host name in the format *host-name*`.*`, where the * represents all initiators in the host.
- The host-group name in the format *host-group-name*`.*.*`, where the first * represents all hosts in the host group and the second * represents all initiators in those hosts.
- Blank if not set or for `all other initiators`.

`Name`
The name of the volume.

```
Serial Number
```
The serial number of the volume.

```
LUN
```

- The LUN that identifies the volume to a host.
- Blank if not mapped or mapped as `no-access`.

Volume information:

```
Serial Number
```
The serial number of the volume.

```
Name
```
The name of the volume.

```
Ports
```

- The controller host ports to which the mapping applies.
- Blank if not mapped or mapped as `no-access`.

```
LUN
```

- The LUN that identifies the volume to a host.
- Blank if not mapped or mapped as `no-access`.

```
Access
```
Type of host access to the volume:

- `read-write`: Read and write. This is the default.
- `read-only`: Read only.
- `no-access`: No access (masked).
- `not-mapped`: Not mapped.

```
Identifier
```

- For an FC or SAS initiator, its WWPN.
- `all other initiators`: The volume's default mapping.

```
Nickname
```

- The host name in the format *host-name.**, where the * represents all initiators in the host.
- Blank if not set or for `all other initiators`.

Output **With the `initiator` parameter:**

Host group information:

```
Serial Number
```
The serial number of the host group.

```
Group Name
```
The host group name in the format *host-group-name.**.*, where the first * represents all hosts in the group, and the second * represents all initiators in those hosts.

Host information:

```
Serial Number
```
The serial number of the host.

```
Host Name
```
The host name in the format *host-name.**, where the * represents all initiators in those hosts.

Initiator information:

```
ID
```

- For an FC or SAS initiator, its WWPN.

`Name`
The initiator name.

Mapping information:

`Volume`

- For a volume, its name.
- For a volume group, its name in the format `volume-group-name.*`, where the `*` represents all volumes in the group.

`Serial Number`
The serial number of the volume.

`LUN`
The LUN used to access the volume.

`Access`
Type of host access to the volume:

- `read-write`: Read and write. This is the default.
- `read-only`: Read only.
- `no-access`: No access (masked).

`Ports`
The controller host ports to which the mapping applies.

Example   Show mapping information for all volumes:

`# show maps`

Show mapping information for all initiators:

`# show maps initiator`

Show mapping information for volume group `VGroup1` and ungrouped volume `Vol0004`:

`# show volume-maps VGroup1.*,Vol0004`

See also   • show host-groups
- show host-maps
- show hosts
- show initiators
- show volume-groups
- show volume-maps
- show volumes

show network-parameters

Description Shows the settings and health of each controller module's network port.

Level Monitor

Syntax `show network-parameters`

Output `IP Address`
Network port IP address.

`Gateway`
Network port gateway IP address.

`Subnet Mask`
Network port IP subnet mask.

`MAC Address`
Controller's unique Media Access Control address.

`Addressing Mode`

- `Manual`: Network settings set manually (statically).
- `DHCP`: DHCP used to set network parameters.

`Link Speed`

- `Unknown`: Either the link speed has not been set, or it has been unset because the controller module was removed from its enclosure.
- `10mbps`: The network port link speed is set to 10 Mb/s.
- `100mbps`: The network port link speed is set to 100 Mb/s.

`Duplex Mode`

- `Undefined`: Either the duplex mode has not been set, or it has been unset because the controller module was removed from its enclosure.
- `Half`: The network port duplex mode is set to half duplex.
- `Full`: The network port duplex mode is set to full duplex.

`Auto Negotiation`

- `Disabled`: Either the network port has not been set, or it has been unset because the controller module was removed from its enclosure, or the port is connected to a switch and is set to use the link speed and duplex mode shown by the `Link Speed` and `Duplex Mode` fields.
- `Enabled`: The network port is set to auto-negotiate a link speed (up to the maximum speed shown by the `Link Speed` field) and duplex mode with a connected Ethernet switch.

`Health`
The health of the network connection.

- `OK`
- `Degraded`
- `N/A`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show network parameters:

```
# show network-parameters
```

See also • set network-parameters

show ntp-status

Description   Shows the status of the use of Network Time Protocol (NTP) in the system.

Level   Monitor

Syntax   `show ntp-status`

Output   NTP Status

- `activated`: NTP is enabled.
- `deactivated`: NTP is disabled.

NTP Server Address
NTP server IP address, or 0.0.0.0 if not set.

Last Server Contact
Date and time in the format `year-month-day hour:minutes:seconds` of the last message received from the NTP server; or none.

Example   Show NTP status for the system:

```
# show ntp-status
```

See also   • set controller-date

show ports

Description Shows information about host ports in both controllers.

Level Monitor

Syntax `show ports`

Output `Ports`
Controller ID and port number

`Media`

- `FC(P):` Fibre Channel Point-to-Point (public or private)
- `FC(L):` Fibre Channel-Arbitrated Loop
- `FC(-):` Not applicable, as when the port is disconnected
- `SAS:` Serial Attached SCSI

`Target ID`
Port WWN

`Status`

- `Up:` The port is cabled and has an I/O link.
- `Disconnected:` Either no I/O link is detected or the port is not cabled.

`Speed (A)`

- Actual host-port link speed in Gbit/s
- Blank if not applicable

`Speed (C)`
Configured host-port link speed:

- FC: `Auto`, `8Gb`, `4Gb`, or `2Gb` (Gbit/s)
- Blank if not applicable

`Health`

- `OK`
- `Degraded`
- `Fault`
- `N/A`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

`Topo (C)`
FC and SAS only. Configured topology

`Width`
SAS only. Number of PHY lanes in the SAS port.

`PID`
FC and SAS only. Primary ID, or blank if not applicable.

Example Show port information for both controllers:

`# show ports`

See also
- set host-parameters
- show sas-link-health

show power-supplies

Description Shows information about each power supply in the storage system.

Level Monitor

Syntax `show power-supplies`

Output `Encl`
The ID of the enclosure that contains the power supply.

`Serial Number`
The serial number of the power supply.

`Name`
The power-supply identifier and location.

`Health`

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show power-supply information:

`# show power-supplies`

See also
- show fans
- show frus

show priorities

Description Shows snapshot-retention priorities for a specified snap pool. Snap-pool priorities, in conjunction with snapshot priorities, determine which snapshots are retained if system resource limitations require some snapshots to be automatically deleted.

Lower-priority snapshots will be deleted before higher-priority snapshots. Priority values are 0–65535.

Level Monitor

Syntax `show priorities snap-pool`

Parameters `snap-pool`
The name or serial number of the snap pool, as shown by the show snap-pools command.

Output `Attribute Name`

- `Standard Snapshot`
- `Volume Copy Snapshot`: A snapshot that is being used to copy data from a source volume to a destination volume. This attribute is temporary for the duration of the volume-copy process.
- `Replication Snapshot`
- `Replicating Snapshot`: A snapshot that is being replicated to a secondary volume. This snapshot is required in order to resume the replication. The attribute is temporary for the duration of the replication process.
- `Common Sync Point Snapshot`: The latest snapshot that is copy-complete in all secondary volumes. It identifies a common point in time that is known by all destinations.
- `Only Sync Point Snapshot`: The only sync point that is available in at least one secondary volume. If this snapshot is removed, then the next replication requires a full sync to be performed.
- `Queued Snapshot`: A snapshot that was taken for remote replication but is queued waiting for the previous replications to complete.
- `DRM Snapshot`: A temporary standard snapshot created from a replication snapshot for the purpose of doing a test failover for disaster recovery management (DRM).

`Priority`
Retention priority for the corresponding attribute.

Example Show priorities for snap-pool `SNAP_A`:

```
# show priorities SNAP_A
```

See also • set priorities

show protocols

Description Shows which management services and protocols are enabled or disabled.

Level Monitor

Syntax `show protocols`

Output `Web Browser Interface (HTTP)`
Shows whether the standard Storage Management Console web server is enabled or disabled. The default is Enabled.

`Secure Web Browser Interface (HTTPS)`
Shows whether the secure Storage Management Console web server is enabled or disabled. The default is Enabled.

`Command Line Interface (Telnet)`
Shows whether the standard CLI is enabled or disabled. The default is Enabled.

`Secure Command Line Interface (SSH)`
Shows whether the secure shell CLI is enabled or disabled. The default is Enabled.

`Storage Management Initiative Specification (SMI-S)`
Shows whether the secure SMI-S interface is enabled or disabled. When enabled, this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTP port 5989. The default is Enabled.

`Unsecure Storage Management Initiative Specification (SMI-S 5988)`
Shows whether the secure SMI-S interface is enabled or disabled. When enabled, this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTP port 5988. The default is Disabled.

`File Transfer Protocol (FTP)`
Shows whether the expert interface for performing actions such as updating firmware is enabled or disabled. The default is Enabled.

`Simple Network Management Protocol (SNMP)`
Shows whether the SNMP interface is enabled or disabled. When this is disabled, all SNMP requests to the MIB are disabled and SNMP traps are disabled. The default is Enabled.

`Service Debug`
Shows whether the Telnet debug port is enabled or disabled. The default is Disabled.

`In-band SES Management (SES)`
Shows whether the in-band SES interface is enabled or disabled. The default is Enabled.

Example Show the status of service and security protocols:

```
# show protocols
```

See also • set protocols

show refresh-counters

Description In XML API format only, shows when the data represented by the basetype was last updated.

The value 0 means the data has never been updated and is not cached. A value other than 0 is a timestamp indicating that the data has been updated. If the value has changed since the last time you called this command then the data has changed.

Level Monitor

Syntax `show refresh-counters`

See also • set cli-parameters

show remote-systems

Description Shows information about remote systems associated with the local system. Information shown is retrieved from the remote systems.

Level Monitor

Syntax `show remote-systems [system]`

Parameters *system*
Optional. The name or network-port IP address of the remote system about which to show information.

Output `System Name`
The name of the remote system.

`System Contact`
The name of the person who administers the remote system.

`System Location`
The location of the remote system.

`System Information`
A brief description of the remote system.

`Vendor Name`
The vendor name of the remote system.

`Product ID`
The product model identifier of the remote system.

`Product Brand`
The brand name of the remote system.

`IP Address Controller A`
The IP address of the network port in controller A in the remote system.

`IP Address Controller B`
The IP address of the network port in controller B in the remote system.

`Username`
The name of a Manage-level user in the remote system.

`Status`

- `Uninitialized`: This system hasn't communicated with the remote system.
- `Ready`: This system has contacted the remote system and it is ready to use.
- `Connected`: This system is transferring data to the remote system.
- `Not Connected`: The system is not connected to the remote system.

`Last Connected`
Date and time, in the format *year-month-day hour*:*minutes*:*seconds* (UTC), when successful communication was last established between the MC in the local system and the MC in the remote system. This value does not indicate when connection status was last determined, and will not be updated if the remote MC is not accessible or if the connection status is `Not Connected`.

Example Show information about remote system `System2`:

`# show remote-systems System2`

See also
- create remote-system
- delete remote-system
- remote
- set remote-system

show replication-images

Description Shows information about replication images for a specified replication volume.

Level Monitor

Syntax
```
show replication-images
     [set replication-set]
     [replication-volume]
```

Parameters set *replication-set*
Optional. Name or serial number of the replication set.

replication-volume
Optional. Name or serial number of the replication volume. If the name is not unique within the replication set, the local volume is assumed. If the name is not unique across replication sets, specify the set parameter.

Output **Replication volume summary information:**

Name
Replication volume name.

Serial Number
Replication volume serial number.

Type
Replication volume type:

- `Primary Volume`: The volume is the primary volume in a replication set.
- `Secondary Volume`: The volume is the secondary volume in a replication set.

Replication volume image information:

Image Serial Number
Replication image serial number.

Image Name
User-defined name assigned to the primary image.

Snapshot Serial
Replication snapshot serial number associated with the image. The replication snapshot is associated with the replication volume specified in the request.

Snapshot Name
Replication snapshot name associated with the image. For a secondary image, this value is not filled in until the replication is completed.

Creation Date/Time
Date and time when the replication image was created in the replication volume.

Information shown for secondary images, not primary images:

Status
Status of the replication image:

- `N/A`: The image information is not valid.
- `Queued`: The image is known to exist in the primary-view volume but replication has not started.
- `Replicating`: The image is being replicated.
- `Suspended`: The image is being replicated but replication is suspended.
- `Complete`: The image is created, fully replicated, and available.
- `Create-Snapshot`: The image is fully replicated but a snapshot of the image is being created.
- `Offline`: The image has been replicated but is unusable due to an error.

`Progress`
Percentage complete if the image is being replicated. Applies only to secondary volumes.

`Start Date/Time`
Date and time when replication started on the replication volume.

`Last Update`
Date and time when the image was last updated (either due to an ongoing replication operation or the replication being completed).

`Suspended`
Date and time when the image was suspended or resumed.

`Est Complete`
Estimated time when replication is expected to complete.

`Time [HH:MM:SS]`
Total time of replication (in hours, minutes, and seconds) including any suspension time.

Example Show information about replication images for replication set `rsvol1`:

```
# show replication-images rsvol1
```

See also • show replication-sets
• show replication-volumes

show replication-sets

Description   Shows information about replication sets in the local system. You can view information about all replication sets or a specific replication set.

Level   Monitor

Syntax   `show replication-sets`
        `[`*`replication-set`*`]`
        `[shallow]`

Parameters   `replication-set`
Optional. The name or serial number of a replication set or volume for which to display replication-set-level information. If this parameter is omitted, information is shown for all replication sets.

`shallow`
Optional. In the output, displays only the replication volumes table, not the following table that lists the replication volumes in the replication sets.

Output   **Replication set information:**

`Name`
Replication set name.

`Serial Number`
Replication set serial number.

Replication volume information:

`Name`
Replication volume name.

`Serial Number`
Replication volume serial number.

`Status`
Replication volume status: `Replicating`, `Suspended`, `Initializing`, `Inconsistent`, `Offline`, `Online`, or `Establishing proxy`. After a secondary volume is reattached it has `Establishing proxy` status while establishing a proxy connection with the remote (primary) system in preparation for replication.

`Status-Reason`
More information about the status value, or `N/A` for `Online` status.

`Monitor`
Replication volume monitoring status:

- `OK`: Communication to the remote volume is successfully occurring on the FC network.
- `Failed`: Communication to the remote volume has failed because of an FC network issue or because the remote volume has gone offline.

`Location`
Replication volume location: `Local` or `Remote`.

`Primary-Volume`
Primary volume name. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.

`Primary-Volume-Serial`
Primary volume serial number. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.

`Primary-Volume-Status`
Primary volume status: `Online`, `Offline`, `Conflict`, or `N/A`.

MaxQueue
Number of replication images to consider when determining the next image to replicate. Used only if the `On Collision` parameter is set to `Oldest`.

MaxRetryTime
Amount of time in seconds that the replication volume should retry a replication operation on any specific image when errors occur. Used only if the `On Error` parameter is set to `Retry`.

On Error
Error policy to invoke when errors occur during the replication process: `Retry` or `Suspend`.

Link Type
Type of ports used to link the primary and secondary volumes:

- `FC`: FC ports.

On Collision
Collision policy used to determine the next image to replicate when multiple replication images are queued: `Newest` or `Oldest`.

Monitor Interval
Interval in seconds at which the primary volume should query the secondary volume.

Priority
Priority of the replication process on the replication volume: `Low`, `Medium`, or `High`.

Connection Status

- `Not Attempted`: Communication has not been attempted to the remote volume.
- `Online`: The volumes in the replication set have a valid connection but communication is not currently active.
- `Active`: Communication is currently active to the remote volume.
- `Offline`: No connection is available to the remote system.

Connection Time
Date and time of the last communication with the remote volume, or `N/A`.

Remote link information:

Connected Ports

- For a remote primary or secondary volume, this field shows the IDs of up to two hosts ports in the local system that are connected to the remote system. If two ports are connected but only one is shown, this indicates that a problem is preventing half the available bandwidth from being used.
- For a local primary or secondary volume, this field shows N/A.

Remote Address
The address of each host port in the remote system through which the volume is accessible.

Example Show information about all replication sets:

```
# show replication-sets
```

Show information about replication set `RS1`:

```
# show replication-sets RS1
```

See also • show replication-images
• show replication-volumes

show replication-volumes

Description Shows information about volumes in replication sets. You can view information about all replication volumes, volumes in a specified replication set, or a specified replication volume.

Level Monitor

Syntax `show replication-volumes`
`[set `*`replication-set`*`]`
`[`*`replication-volume`*`]`
`[shallow]`

Parameters `set `*`replication-set`*
Optional. The name or serial number of the replication set.

`replication-volume`
Optional. The name or serial number of the replication volume. If the name is not unique within the replication set, the local volume is assumed. If the name is not unique across replication sets, specify the `set` parameter.

`shallow`
Optional. In the output, displays only the replication volumes table, not the following table that lists replication images in the replication volumes.

Output **Replication volume information:**

`Name`
Replication volume name.

`Serial Number`
Replication volume serial number.

`Status`
Replication volume status:

- `Initializing`: The initial (full) replication to the volume is in progress.
- `Online`: The volume is online and is consistent with the last replicated image.
- `Inconsistent`: The volume is online but is in an inconsistent state. A full replication is required to initialize it.
- `Replicating`: The volume is online and replication is in progress.
- `Replicate-delay`: The volume is online but the in-progress replication has been temporarily delayed; a retry is occurring.
- `Suspended`: The volume is online but the in-progress replication has been suspended.
- `Offline`: The volume can be accessed but is unusable due to an error.
- `Establishing proxy`: The volume is establishing a proxy connection to a remote volume. This will occur when a detached secondary volume is reattached and is re-establishing a connection with the primary system in preparation for replication.
- `Detached`: The volume is detached for removal.

`Status-Reason`
More information about the status value, or `N/A` for `Online` status.

`Monitor`
Replication volume monitoring status:

- `OK`: Communication to the remote volume is successfully occurring on the network.
- `Failed`: Communication to the remote volume has failed because of a communication issue or because the remote volume has gone offline.

`Location`
Replication volume location: `Local` or `Remote`.

`Primary-Volume`
Primary volume name. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.

`Primary-Volume-Serial`
Primary volume serial number. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.

`Primary-Volume-Status`
Primary volume status: `Online`, `Offline`, `Conflict`, or `N/A`.

`MaxQueue`
Maximum number of replication images to consider when determining the next image to replicate. Used only if the `On Collision` parameter is set to `Oldest`. The default is 32.

`MaxRetryTime`
Maximum amount of time in seconds that the replication volume should retry a replication operation on any specific image when errors occur. Used only if the `On Error` parameter is set to `Retry`. The default is 1800.

`On Error`
Error policy to invoke when errors occur during the replication process:

- `Retry`: Retry the replication for the time specified by the `MaxRetryTime` value. This is the default.
- `Suspend`: Suspend the replication until the error is resolved automatically or through user intervention.

`Link Type`
Type of ports used to link the primary and secondary volumes:

- `FC`: FC ports.

`On Collision`
Collision policy used to determine the next image to replicate when multiple replication images are queued:

- `Newest`: Only the latest replication image should be considered for the next replication operation.
- `Oldest`: Only the latest *n* replication images should be considered for the next replication operation, where *n* is defined by the `MaxQueue` value and the oldest of these images should be considered first. This is the default.

`Monitor Interval`
Interval in seconds at which the primary volume should query the secondary volume. The default is 300.

`Priority`
Priority of the replication process on the replication volume:

- `High`: Replication has higher priority than host I/O. This can cause heavy I/O to be slower than normal. This is the default.
- `Medium`: Replication performance is balanced with host I/O performance.
- `Low`: Replication runs at a slower rate with minimal effect on host I/O. Use when streaming data without interruption is more important than data redundancy.

`Connection Status`

- `Not Attempted`: Communication has not been attempted to the remote volume.
- `Online`: The volumes in the replication set have a valid connection but communication is not currently active.
- `Active`: Communication is currently active to the remote volume.
- `Offline`: No connection is available to the remote system.

```
Connection Time
```
Date and time of the last communication with the remote volume, or `N/A`.

Remote link information:

```
Connected Ports
```

- For a remote primary or secondary volume, this field shows the IDs of up to two hosts ports in the local system that are connected to the remote system. If two ports are connected but only one is shown, this indicates that a problem is preventing half the available bandwidth from being used.
- For a local primary or secondary volume, this field shows `N/A`.

```
Remote Address
```
The address of each host port in the remote system through which the volume is accessible.

Example Show information about all replication volumes:

```
# show replication-volumes
```

Show information about replication volume `MyData` in replication set `RS1`:

```
# show replication-volumes set RS1 MyData
```

See also • set replication-volume-parameters
- show replication-sets

show sas-link-health

Description Shows the condition of SAS expansion-port connections.

Level Monitor

Syntax `show sas-link-health`

Output `Encl`
Enclosure ID.

`Ctlr`
ID of the controller module or expansion module.

`Name`

- `Out Port`: Egress (expansion) port in a controller module or an expansion module. Can be connected to an ingress port in an expansion module.
- `In Port`: Ingress port in an expansion module. Can be connected to an egress (expansion) port in a controller module or an expansion module.

`Status`

- `Up`
- `Disconnected`
- `Not Present`
- `Warning`
- `Error`
- `Unknown`

`Health`

- `OK`
- `Degraded`
- `Fault`
- `N/A`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows the recommended actions to take to resolve the health issue.

Example Show the condition of SAS expansion-port connections:

`# show sas-link-health`

See also
- show expander-status
- show ports

show schedules

Description Shows information about task schedules.

Level Monitor

Syntax `show schedules [`*`schedule-name`*`]`

Parameters *`schedule-name`*
Optional. Shows information about the specified schedule only. If this parameter is omitted, information is shown for all schedules.

Output `Schedule Name`
The schedule name.

`Schedule Specification`
The schedule settings for running the associated task.

`Status`

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run at the next scheduled time.
- `Suspended`: The task had an error and is holding in its current state.
- `Expired`: The task has exceeded a constraint and will not run again.
- `Invalid`: The task is invalid.

`Next Time`
Date and time, in the format *year-month-day hour*.`minutes`:`seconds` (UTC), when the schedule will next run.

`Task To Run`
The name of the task that the schedule runs.

`Error Message`
If an error occurred while running the schedule, the error message; otherwise, blank.

Task-specific information, as shown by the show tasks command.

Example Show information about all task schedules:

`# show schedules`

Show information about schedule `Sched2` only:

`# show schedules Sched2`

See also
- create schedule
- delete schedule
- show tasks

show sensor-status

Description Shows the status of each environmental sensor in each enclosure.

Information shown only for a controller enclosure: on-board temperature, disk controller temperature, memory controller temperature, super-capacitor voltage and charge, overall unit (enclosure) status.

Information shown for all enclosures: temperature, voltage, and current for each IOM (controller module or expansion module); temperature, voltage, and current for each PSU (power supply).

Normal and error ranges for temperature and voltage are specified in the Setup Guide.

Level Monitor

Syntax `show sensor-status`

Output `Encl`
Enclosure ID.

`Sensor Name`
Sensor name and location.

`Value`

- For a sensor, its value.
- For overall unit status, one of the status values below.

`Status`

- `OK`: The sensor is present and detects no error condition.
- `Warning`: The sensor detected a non-critical error condition. Temperature, voltage, or current is between the warning and critical thresholds.
- `Error`: The sensor detected a critical error condition. Temperature, voltage, or current exceeds the critical threshold.
- `Unavailable`: The sensor is present with no known errors, but has not been turned on or set into operation because it is initializing. This typically occurs during controller startup.
- `Unrecoverable`: The enclosure management processor (EMP) cannot communicate with the sensor.
- `Unknown`: The sensor is present but status is not available.
- `Not Installed`: The sensor is not present.
- `Unsupported`: Status detection is not implemented.

Example Show sensor status for all enclosures:

`# show sensor-status`

show snap-pools

Description Shows information about snap pools.

Level Monitor

Syntax `show snap-pools [controller a|b|both]`

Parameters `controller a|b|both`
Optional. Shows information about the snap pool owned by controller A, the snap pool owned by controller B, or both snap pools. If this parameter is omitted, information is shown for all snap pools.

Output `Pool`
The name of the storage pool that contains the snap pool.

`Serial Number`
The serial number of the snap pool.

`Name`
The name of the snap pool.

`Size`
The total capacity of the snap pool.

`Reserved Size`
The amount of storage-pool capacity that is reserved for use by the snap pool.

`Free`
The amount of free space available in the snap pool.

`Volumes`
The number of standard volumes using this snap pool.

`Snapshots`
The number of snapshots using this snap pool.

`Status`

- `Available`: The snap pool is available for use.
- `Offline`: The snap pool is not available for use, as in the case where its disks are not present.
- `Corrupt`: The snap pool's data integrity has been compromised; the snap pool can no longer be used.

`Threshold`

- `Warning`: The snap pool is moderately full. When this threshold is reached, an event is logged. The default is 75%.
- `Error`: The snap pool is nearly full and unless corrective action is taken, snapshot data loss is probable. When this threshold is reached, an Informational event is logged and the error policy is invoked. The default is 90%.
- `Critical`: The snap pool is essentially full and data loss is imminent unless action is taken. When this threshold is reached, an event is logged and the critical policy is invoked. The default is 98%.

`%Usage`
The percent of snap-pool consumption that invokes the policy for the threshold.

`Policy`
The action to take when a threshold is reached:

- `Auto Expand`: Try to expand the snap pool. If the snap pool's space usage reaches the percentage specified by its error threshold, the system will log Warning event 230 and will try to expand the snap pool.
    - If the snap pool is successfully expanded, the system will log Informational event 444.
    - If the snap pool cannot be expanded because there is not enough available space in its storage pool, the system will log Warning event 444 and will automatically delete the oldest snapshot that is not a current sync point.

    Each time the snap-pool's error threshold is reached and the system cannot auto-expand the storage pool, the oldest remaining snapshot that is not a current sync point will be deleted. This behavior occurs for each snap pool independently, based on its space usage.
- `Delete Oldest Snapshot`: Delete the oldest snapshot.
- `Delete Snapshots`: Delete all snapshots.
- `Halt Writes`: Halt writes to the snap pool.
- `Notify Only`: Log an event to notify the administrator.

Example Show information for both snap pools:

```
# show snap-pools
```

Show information for the snap pool owned by controller A:

```
# show snap-pools controller a
```

See also • set snap-pool-parameters

show snapshots

Description Shows information about snapshots for a specified controller, volume, or snap pool. If no parameters are specified, information about all snapshots is shown.

Level Monitor

Syntax `show snapshots`
    `[controller a|b|both]`
    `[standard-volume volume]`
    `[snap-pool volume]`
    `[type standard|replication|all]`

Parameters `controller a|b|both`
Optional. Shows snapshots owned by controller A only, by controller B only, or by either controller (`both`). If this parameter is omitted, snapshots owned by either controller are shown.

`standard-volume volume`
Optional. Shows snapshots associated with the specified source volume name or serial number. For volume syntax, see Command syntax on page 18. If this parameter is omitted, snapshots of all volumes are shown.

`snap-pool volume`
Optional. Shows snapshots associated with the specified snap pool name or serial number. For volume syntax, see Command syntax on page 18. If this parameter is omitted, snapshots in both snap pools are shown.

`type standard|replication|all`
Optional. Shows only standard (non-replication) snapshots, only replication snapshots, or snapshots of all types. If this parameter is omitted, snapshots of all types are shown.

Output `Serial Number`
Serial number of the snapshot.

`Name`
Name of the snapshot.

`Creation Date/Time`
The date and time when the snapshot was prepared or committed.

`Status`

- `Available`
- `Unavailable`: See the `Status-Reason` value.

`Status-Reason`
Shows `N/A` for `Available` status, or one of the following reasons for `Unavailable` status:

- `MV Not Accessible`: Source volume is not accessible
- `MV Not Found`: Source volume is not found
- `SP Not Accessible`: Snap pool is not accessible
- `SP Not Found`: Snap pool is not found
- `SS Pending`: Snapshot is pending
- `VC-MD In Progress`: Volume-copy with modified data is in progress
- `RB-MD In Progress`: Rollback with modified data is in progress

`Source Volume`
The name of the volume of which the snapshot was taken.

`Snap Data`
The total amount of preserved and write data associated with the snapshot.

`Unique Data`
The amount of preserved and write data that is unique to the snapshot.

```
Shared Data
```
The amount of preserved and write data that is shared between this snapshot and other snapshots.

```
Priority
```
The retention priority for the snapshot, based on the snapshot attributes and the user-defined retention priority for the snapshot type.

```
User Priority
```
The user-defined retention priority for the snapshot type.

```
Type
```

- `Standard snapshot`: Snapshot of a standard volume.
- `Standard snapshot(DRM)`: A temporary standard snapshot created from a replication snapshot for the purpose of doing a test failover for disaster recovery management (DRM).
- `Replication snapshot`: For a primary or secondary volume, a snapshot that was created by a replication operation but is not a sync point.
- `Replication snapshot(Replicating)`: For a primary volume, a snapshot that is being replicated to a secondary system.
- `Replication snapshot(Current sync point)`: For a primary or secondary volume, the latest snapshot that is copy-complete on any secondary system in the replication set.
- `Replication snapshot(Common sync point)`: For a primary or secondary volume, the latest snapshot that is copy-complete on all secondary systems in the replication set.
- `Replication snapshot(Old Common sync point)`: For a primary or secondary volume, a common sync point that has been superseded by a new common sync point.
- `Replication snapshot(Only sync point)`: For a primary or secondary volume, the only snapshot that is copy-complete on any secondary system in the replication set.
- `Replication snapshot(Queued)`: For a primary volume, a snapshot associated with a replication operation that is waiting for a previous replication operation to complete.
- `Replication snapshot(Awaiting replicate)`: For a primary volume, a snapshot that is waiting to be replicated to a secondary system.

Example Show information about snapshots associated with snap pool `SNAP_B`:

```
# show snapshots snap-pool SNAP_B
```

See also
- show volumes
- show snap-pools

show snmp-parameters

Description Shows SNMP settings for event notification.

Level Monitor

Syntax `show snmp-parameters`

Output `SNMP Trap Notification Level`

- `crit`: Only Critical events are sent as traps.
- `error`: Error and Critical events are sent as traps.
- `warn`: Warning, Error, and Critical events are sent as traps.
- `info`: All events are sent as traps.
- `none`: No events are sent as traps and traps are disabled.

`SNMP Trap Host IP#`
The IP address of each trap host.

`SNMP read community`
The community string for read-only access. The value is represented by six asterisks to monitor-level users and is shown in clear text to manage-level and diagnostic-level users.

`SNMP write community`
The community string for write access. The value is represented by six asterisks to monitor-level users and is shown in clear text to manage-level and diagnostic-level users.

Example Show SNMP notification settings:

`# show snmp-parameters`

See also
- set snmp-parameters
- set protocols
- show protocols

show storage

Description   Shows information about the storage configuration of the system.

Level   Monitor

Syntax   `show storage [unconfigured]`

Parameters   `unconfigured`
Optional. Determines what storage is available to be pooled and how it would be configured. This parameter always attempts to balance components between pools. If this parameter is omitted, the current storage configuration is shown.

Output   `Pool`
The name of the storage pool.

`Component`
The name of the storage-pool component.

`Type`

- `SAS`: Dual-port, spinning, enterprise SAS disk.
- `SAS MDL`: Dual-port, spinning, midline SAS disk.
- `sSAS`: Dual-port, solid-state SAS disk (SSD).

`Size`
The amount of space allocated to the storage-pool component.

`Configuration`

- `RAID1`
- `RAID6`
- `RFC` (read flash cache)
- `Spare`

`Tier`

- `Performance`: The disk is in the highest storage tier, which uses SSDs (high speed, low capacity).
- `Standard`: The disk is in the storage tier that uses enterprise-class spinning SAS disks (lower speed, higher capacity).
- `Archive`: The disk is in the lowest storage tier, which uses midline spinning SAS disks (low speed, high capacity).
- `Read Cache`: The disk is an SSD providing high-speed read cache for a storage pool.

`Encl`
The IDs of enclosures that contain disks used by the storage-pool component.

`Disks`
The IDs of disks used by the storage-pool component. A series of disks in the same enclosure is shown in the format `enclosure-ID.first-disk-last-disk`.

Example   Show information about the current storage configuration:

```
# show storage
```

See also
- add storage
- show storage-status

show storage-pool-components

Description   Shows information about storage-pool components.

Level   Diagnostic

Syntax   `show storage-pool-components [`*`storage-pool-ID`*`]`

Parameters   *`storage-pool-ID`*
Optional. The name or serial number of the storage pool for which to show information. If this parameter is omitted, information is shown for both storage pools.

Output   `Name`
The name of the storage-pool component.

`Serial Number`
The serial number of the storage-pool component.

`Storage Pool Serial Number`
The serial number of the storage pool.

`Storage Pool`
The name of the storage pool.

`Tier`

- `Performance`: The storage-pool component is in the highest storage tier, which uses SSDs (high speed, low capacity).
- `Standard`: The storage-pool component is in the storage tier that uses enterprise-class spinning SAS disks (lower speed, higher capacity).
- `Archive`: The storage-pool component is in the lowest storage tier, which uses midline spinning SAS disks (low speed, high capacity).
- `Read Cache`: The storage-pool component is an SSD providing high-speed read cache for a storage pool.

`Total Pages`
The total pages.

`Allocated Pages`
The number of pages allocated to the storage-pool component.

`Available Pages`
The number of available pages in the storage-pool component.

`% of Pool`
The percentage of the pool size that the storage-pool component occupies.

`Utility Running`

- `DRSC`: The disk is being scrubbed.
- `EXPD`: The storage-pool component is being expanded.
- `INIT`: The storage-pool component is being initialized.
- `RCON`: The storage-pool component is being reconstructed.
- `VRFY`: The storage-pool component is being verified.
- `VRSC`: The storage-pool component is being scrubbed.
- Blank if no job is running.

`Utility Progress`

- `0–99`: The progress, in percent, of a utility that is running on the storage-pool component.
- `100`: The utility has finished running.

```
Status
```
- `CRIT`: Critical. The storage-pool component is online but isn't fault tolerant because some of its disks are down.
- `FTDN`: Fault tolerant with a down disk. The storage-pool component is online and fault tolerant, but some of its disks are down.
- `FTOL`: Fault tolerant and online.
- `OFFL`: Offline. Either the storage-pool component is using offline initialization, or its disks are down and data may be lost.
- `QTCR`: Quarantined critical. The storage-pool component is offline and quarantined because at least one disk is missing; however, the storage-pool component could be accessed.
- `QTDN`: Quarantined with a down disk. The storage-pool component is offline and quarantined because at least one disk is missing; however, the storage-pool component could be accessed and would be fault tolerant. For instance, one disk is missing from a RAID-6 component.
- `QTOF`: Quarantined offline. The storage-pool component is offline and quarantined because multiple disks are missing and user data is incomplete.
- `UNKN`: Unknown.
- `UP`: Up. The storage-pool component is online and does not have fault-tolerant attributes.

```
Health
```
- `OK`
- `Degraded`
- `Fault`
- `Unknown`
- `N/A`

```
Health Reason
```
If Health is not `OK`, this field shows the reason for the health state.

```
Health Recommendation
```
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show all storage-pool components:

```
# show storage-pool-components
```

Show components for storage pool B only:

```
# show storage-pool-components B
```

See also
- show storage
- show storage-status

show storage-pools

Description Shows information about storage pools.

Level Monitor

Syntax `show storage-pools [`*`storage-pool-ID`*`]`

Parameters *`storage-pool-ID`*
Optional. The name or serial number of the storage pool for which to show information. If this parameter is omitted, information is shown for both storage pools.

Output `Name`
The name of the storage pool.

`Serial Number`
The serial number of the storage pool.

`Total Size`
The total capacity of the storage pool.

`Avail`
The available capacity in the storage pool.

`Components`
The number of storage-pool components in the storage pool.

`Volumes`
The number of volumes in the storage pool.

`Low Thresh`
The low threshold for page allocation as a percentage of storage-pool capacity. When this threshold is reached, event 462 is logged with Informational severity. The default is 25%.

`Mid Thresh`
The middle threshold for page allocation as a percentage of storage-pool capacity. When this threshold is reached, event 462 is logged to notify the administrator to add capacity to the pool. If overcommit is enabled, the event has Informational severity; if overcommit is disabled, the event has Warning severity. The default is 50%.

`High Thresh`
The high threshold for page allocation as a percentage of storage-pool capacity. When this threshold is reached, event 462 is logged with Critical severity to alert the administrator that it is critical to add capacity to the pool. The threshold value is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space.

`Snap-Pool% (Size)`
The size of the snap pool as a percentage of storage-pool capacity and in bytes.

`Snap-Pool Avail`
The available capacity in the snap pool.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show information for both storage pools:

`# show storage-pools`

Show information for storage pool A:

`# show storage-pools A`

See also • set storage-pool
• show storage-pool-components

show storage-status

Description Shows information and statistics for storage-pool components, storage pools, or tiers.

You can view interactive statistics for storage-pool components, or historical statistics for pools and tiers. The system samples disk-performance statistics every quarter hour and retains performance data for 6 months.

The historical option allows you to specify a time range or a number (count) of data samples to include. It is not recommended to specify both the `time-range` and `count` parameters; if both parameters are specified, and more samples exist for the specified time range, the samples' values will be aggregated to show the required number of samples.

Level Monitor

Syntax
```
show storage-status
    [storage-pool storage-pool-ID]
    [tier tier-name]
    [historical]
    [time-range "date/time-range"]
    [count number-of-data-samples]
    [all]
    [filename filename]
```

Parameters `storage-pool` *storage-pool-ID*
Optional. The name or serial number of the storage pool about which to show information. If this parameter and the `tier` parameter are omitted, information is shown for all storage-pool components.

`tier` *tier-name*
Optional. The name of the tier about which to show information. If this parameter and the `storage-pool` parameter are omitted, information is shown for all storage-pool components.

`historical`
Optional. Specifies to show historical statistics. If this parameter and the `tier` parameter are specified, you must also specify the `pool` parameter. If the `historical` parameter is omitted, interactive statistics will be shown.

`time-range` "*date/time-range*"
Optional. Specifies the date/time range of historical statistics to show, in the format "start *yyyy-mm-dd hh:mm* [AM|PM] end *yyyy-mm-dd hh:mm* [AM|PM]". If the start date/time is specified but no end date/time is specified, the current date/time will be used as the end date/time. The system will return the oldest sample taken after the start time and the latest sample taken before the end time. If the specified start date/time is earlier than the oldest sample, that sample will be used as the start date/time. If you specify this parameter, do not specify the `count` parameter. If this parameter is omitted, the most recent 100 data samples will be displayed.

`count` *number-of-data-samples*
Optional. Specifies the number of data samples to display, from 1–100. Each sample will be shown as a separate row in the command output. If this parameter is omitted, 100 samples will be shown. If you specify this parameter, do not specify the `time-range` parameter.

`all`
Optional. Specifies to show the full set of performance metrics. If this parameter is omitted, the default set of performance metrics will be shown.

`filename` *filename*
Optional. Only applies when the `historical` parameter is specified. Creates a file of CSV historical statistics which can be accessed using FTP.

Output	Storage-pool component status

```
Pool
Component
Tier
Job
Job %
Total Size
Allocated Size
Pages Allocated Per Min
Pages Deallocated Per Min
Available RFC Size
Hot Pages
Pages Reclaimed
Status
Health
Health Reason
Health Recommendation
```

Storage-pool status:

```
Pool
Total Size
Allocated Size
Available Size
Over-Committed
Pages Allocated Per Min
Pages Deallocated Per Min
Total RFC Size
Available RFC Size
Reserved Size
Health
Health Reason
Health Recommendation
```

Storage-tier status:

```
Pool
Tier
% of Pool
Disks
Total Size
Allocated Size
Available Size
Pages Allocated Per Min
Pages Deallocated Per Min
Hot Pages
Pages Reclaimed
```

Interactive statistics (for components, pools, or tiers):

```
Time (ms) Since Reset
Number of Reads
Number of Writes
Data Read
Data Written
Bytes per second
IOPS
```

Historical statistics (for pools or tiers):

Pool information:

```
Pool
Total I/Os
Data Transferred
Total IOPS
Total B/s
Number Of Allocated Pages
Number Of Hot Pages
Sample Time
```

Tier information:

```
Pool
Tier
Total I/Os
Data Transferred
Total IOPS
Total B/s
Number Of Allocated Pages
Number Of Hot Pages
Number of Page Moves In
Number of Page Moves Out
Total I/Os
Number Of Initial Allocations
Number Of Unmaps
Number Of Rfc Copies
Sample Time
```

Example Show interactive statistics for all storage-pool components:

```
# show storage-status
```

Show interactive statistics for storage pool A:

```
# show storage-status storage-pool A
```

Show historical statistics from a specified date and time range:

```
# show storage-status historical time-range "start 2012-06-24 4:40 PM
end 2012-06-18 5:00 PM"
```

Show historical statistics for the Standard tier in storage pool A:

```
# show storage-status storage-pool A tier Standard historical
```

See also • reset all-statistics
- show controller-statistics
- show disk-statistics
- show host-port-statistics
- show volume-statistics

show system

Description Shows information about the storage system. If the system's health is not `OK`, each unhealthy component is listed with information to help you resolve the health problem.

Level Monitor

Syntax `show system [detail]`

Parameters `detail`
Optional. Shows additional information about the system.

Output `System Name`
The name of the system. The default is `Uninitialized Name`.

`System Contact`
The name of the person who administers the system. The default is `Uninitialized Contact`.

`System Location`
The location of the system. The default is `Uninitialized Location`.

`System Information`
A brief description of what the system is used for or how it's configured. The default is `Uninitialized Info`.

`Midplane Serial Number`
The serial number of the controller enclosure midplane.

`Vendor Name`
The vendor name.

`Product ID`
The product model identifier.

`Product Brand`
The product brand name.

`SCSI Vendor ID`
The vendor name returned by the SCSI INQUIRY command.

`SCSI Product ID`
The product identifier returned by the SCSI INQUIRY command.

`Enclosure Count`
The number of enclosures in the system.

`Health`

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Other MC Status`

The operational status of the partner MC.

- `Operational`
- `Not operational`

```
PFU Status
```
Shows whether partner firmware update is running on the system, or is idle.

- `Running`
- `Idle`

```
Supported Locales
```
Supported display languages: English, Spanish, Japanese, Chinese-Simplified.

Example Show system information:

```
# show system
```

See also
- set system
- show system-parameters

show system-parameters

Description Shows the current settings for the storage system.

Level Monitor

Syntax `show system-parameters`

Output `ULP Enabled`
Shows that the system is using Unified LUN Presentation, which can expose all LUNs through all host ports on both controllers. The interconnect information is managed in the controller firmware. ULP appears to the host as an active-active storage system where the host can choose any available path to access a LUN regardless of volume ownership. When ULP is in use, the system's operating/cache-redundancy mode is shown as Active-Active ULP. ULP uses the T10 Technical Committee of INCITS Asymmetric Logical Unit Access (ALUA) extensions, in SPC-3, to negotiate paths with aware host systems. Unaware host systems see all paths as being equal.

`Storage Overcommit Enabled`

- `true`: The system will use thin provisioning, which means that more capacity can be allocated to volumes than physically exists in the system. When stored data approaches the limit of physical capacity, the administrator can add more enclosures to the system.

- `false`: The system will use full provisioning, which means that the capacity allocated to volumes when they are created cannot exceed the physical capacity of the system. This is the default.

`Number of Host Ports`
The number of host-interface ports in the controller enclosure.

`Maximum Disks`
The number of disks that the system supports.

`Maximum Volumes`
The number of volumes that the system supports.

`Maximum LUNs`
The number of LUNs that the system supports.

`Maximum Storage Pools Per Controller`
The number of storage pools that each controller supports.

`Maximum Host Groups Per System`
The number of host groups that the system supports.

`Maximum Hosts per Host Group`
The maximum number of hosts that a host group can contain.

`Maximum Initiators Per Host`
The maximum number of initiators that a host can contain.

`Maximum Volume Groups Per Storage Pool`
The maximum number of volume groups that a storage pool can contain.

`Maximum Volumes Per Volume Group`
The maximum number of volumes that a volume group can contain.

`Local Controller`
The ID of the controller you are accessing.

`Serial Number`
The last six digits of the midplane serial number.

Example Show settings and configuration limits for the storage system:

```
# show system-parameters
```

See also • show license
• show system

show tasks

Description Shows information about all tasks.

Level Monitor

Syntax `show tasks [`*`task-name`*`]`

Parameters *task-name*
Optional. Shows information about the specified task only. If this parameter is omitted, information is shown for all tasks.

Output **For a TakeSnapshot task:**

`Task Name`
The name of the task.

`Task Type`
`TakeSnapshot`

`Status`

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

`Task State`
The current step of the task:

- `Start`
- `VerifyVolume`
- `ValidateLicensingLimit`
- `CreateName`
- `CreateSnap`
- `VerifySnap`
- `InspectRetention`
- `FindOldestSnap`
- `UnmapSnap`
- `ResetSnap`
- `RenameSnap`

`Error Message`
If an error occurred while processing the task, the error message; otherwise, blank.

`Source Volume`
The name of the standard volume.

`Source Volume Serial`
The serial number of the standard volume.

`Prefix`
The label that identifies snapshots created by this task. Snapshot names have the format *prefix*`_s#`, where # starts at `001`.

`Count`
The number of snapshots to retain with this prefix. When a new snapshot exceeds this limit, the oldest snapshot with the same prefix is deleted.

`Last Created`
The name of the last snapshot created by the task, or blank.

```
Snapshot Name
```
The name of each snapshot taken, or blank.

```
Snapshot Serial
```
The serial number of each snapshot taken, or blank.

For a ResetSnapshot task:

```
Task Name
```
The name of the task.

```
Task Type
ResetSnapshot
```

```
Status
```

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

```
Task State
```
The current step of the task:

- `Start`
- `VerifySnap`
- `UnmapSnap`
- `ResetSnap`

```
Error Message
```
If an error occurred while processing the task, the error message; otherwise, blank.

```
Snapshot Name
```
The name of the snapshot to reset.

```
Snapshot Serial Number
```
The serial number of the snapshot to reset.

For a VolumeCopy task:

```
Task Name
```
The name of the task.

```
Task Type
VolumeCopy
```

```
Status
```

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

```
Task State
```
The current step of the task:

- `Start`
- `VerifyVolume`
- `CreateName`
- `ObtainMap`
- `UnmapVolume`
- `CreateVolume`
- `RemapVolume`
- `VerifyCreatedVolume`

```
Error Message
```
If an error occurred while processing the task, the error message; otherwise, blank.

```
Source Volume
```
The name of the volume to be copied.

```
Source Volume Serial
```
The serial number of the volume to be copied.

```
Destination Storage Pool
```
The name of the storage pool in which the new volume will be created.

```
Destination Storage-Pool Serial
```
The serial number of the destination storage pool.

```
Prefix
```
The label that identifies copies created by this task. Volume names have the format *prefix*`_c#`, where # starts at `001`.

```
Modified Data
```

- `modified`: The copy includes modified snapshot data.
- `preserved`: The copy excludes modified snapshot data.

```
Last Created
```
The name of the last volume created by the task, or blank.

For a ReplicateVolume task:

```
Task Name
```
The name of the task.

```
Task Type
ReplicateVolume
```

```
Status
```

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

`Task State`
The current step of the task:

- `Start`
- `VerifyVolume`
- `CreateName`
- `RepVolume`
- `VerifySnap`

`Error Message`
If an error occurred while processing the task, the error message; otherwise, blank.

`Primary Volume`
The name of the volume to replicate.

`Primary Volume Serial Number`
The serial number of the volume to replicate.

`Prefix`
The label that identifies snapshots created by this task. Volume names have the format *prefix*`_c#`, where # starts at `001`.

`Mode`

- `new-snapshot`: Replicate a new snapshot of the primary volume.
- `last-snapshot`: Replicate the most recent existing snapshot of the primary volume.

`Last Created`
The name of the last volume created by the task, or blank.

`Last Used Snapshot`
For a task whose replication mode is `last-snapshot`, the name of the last snapshot used for replication; otherwise, `N/A`.

For an EnableDSD task:

`Task Name`
The name of the task.

`Task Type`
`EnableDSD`

`Status`

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

`Task State`
The current step of the task; always `Start`.

`Error Message`
If an error occurred while processing the task, the error message; otherwise, blank.

For a DisableDSD task:

```
Task Name
```
The name of the task.

```
Task Type
DisableDSD
```

```
Status
```

- `Uninitialized`: The task is not yet ready to run.
- `Ready`: The task is ready to run.
- `Active`: The task is running.
- `Error`: The task has an error.
- `Invalid`: The task is invalid.

```
Task State
```
The current step of the task; always `Start`.

```
Error Message
```
If an error occurred while processing the task, the error message; otherwise, blank.

Example   Show information about all tasks:

```
# show tasks
```

Show information about task `Task1` only:

```
# show tasks Task1
```

See also
- create schedule
- create task
- delete task
- set task
- show schedules

show unwritable-cache

Description Shows the percentage of unwritable data in the system. This data cannot be written to disk because it is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host and disk.

Level Monitor

Syntax `show unwritable-cache`

Example Show the percentage of unwritable cache data in each controller:

```
# show unwritable-cache
```

See also • clear cache

show users

Description Shows settings for configured users.

Level Monitor

Syntax `show users [user]`

Parameters *user*
Optional. Shows settings for the specified user only. If this parameter is omitted, settings are shown for all users.

Output `Username`
The user name.

`Roles`

- `monitor`: View-only access to selected user interfaces.
- `manage`: Modify access to selected user interfaces.
- `diagnostic`: User can access additional, diagnostic commands.

`User Type`
The user's experience level: `Novice`, `Standard`, `Advanced`, or `Diagnostic`. All user types have access to all commands.

`User Locale`
The display language for this user. The default is English.

`WBI`

- `x`: The user can access the web-browser interface. This is the default.
- (blank): The user cannot access this interface.

`CLI`

- `x`: The user can access the command-line interface. This is the default.
- (blank): The user cannot access this interface.

`FTP`

- `x`: The user can access the file transfer protocol interface.
- (blank): The user cannot access this interface. This is the default.

`SMI-S`

- `x`: The user can access the Storage Management Initiative Specification (SMI-S) interface.
- (blank): The user cannot access this interface. This is the default.

`SNMP`

- `U`: The user can access the SNMPv3 interface and view the MIB.
- `T`: The user can access the SNMPv3 interface and receive trap notifications.
- (blank): The user cannot access this interface. This is the default.

`Authentication Type`

- `MD5`: MD5 authentication. This is the default.
- `SHA`: SHA (Secure Hash Algorithm) authentication.
- `none`: No authentication.

`Privacy Type`

- `DES`: Data Encryption Standard.
- `AES`: Advanced Encryption Standard.
- `none`: No encryption. This is the default.

`Password`
The user password. For a standard user the password is represented by eight asterisks. For an SNMPv3 user this is the authentication password; it is shown in clear text for reference when configuring users in the corresponding management application.

`Privacy Password`
The encryption password for an SNMPv3 user whose privacy type is set to `DES` or `AES`. The password is shown in clear text for reference when configuring users in the corresponding management application.

`Trap Host`
The SNMP trap destination for an SNMPv3 user that can receive trap notifications.

Example Show settings for all configured users:

```
# show users
```

Show settings for user `JSmith`:

```
# show users JSmith
```

See also
- create user
- delete user
- set user

show versions

Description Shows firmware and hardware version information for each controller module.

Level Monitor

Syntax `show versions [detail]`

Parameters `detail`
Optional. Shows information about the versions of firmware and hardware in each controller module. If this parameter is omitted, only firmware-bundle information is shown.

Example Show firmware-bundle version information for a system in which controller B is not installed:

```
# show versions
Controller A Versions
---------------------
Bundle Version: value
Build Date: value

Controller B Versions
---------------------
Bundle Version: Not Present
Build Date: Not Present
```

Show detailed version information for each controller:

```
# show versions detail
Controller A Versions
---------------------
Storage Controller CPU Type: value
Bundle Version: value
Build Date: value
Storage Controller Code Version: value
Storage Controller Code Baselevel: value
Memory Controller FPGA Code Version: value
Storage Controller Loader Code Version: value
CAPI Version: value
Management Controller Code Version: value
Management Controller Loader Code Version: value
Expander Controller Code Version: value
CPLD Code Version: value
Hardware Version: value
Host Interface Module Version: value
Host Interface Module Model: value
Backplane Type: value
Host Interface Hardware (Chip) Version: value
Disk Interface Hardware (Chip) Version: value
SC Boot Memory Reference Code Version: value

Controller B Versions
---------------------
...
```

See also • show inquiry

show volume copy-status

Description Shows information about in-progress volume copy operations. While a volume copy is in progress, the destination volume cannot be accessed.

Level Monitor

Syntax `show volume copy-status [controller a|b|both]`

Parameters `controller a|b|both`
Optional. Shows volume copy operations for volumes owned by controller A only, by controller B only, or by either controller (`both`). If this parameter is omitted, all volume copy operations are shown.

Output `VC Volume Name`
The name of the destination volume.

`Serial Number`
The serial number of the destination volume.

`Pool`
The name of the destination storage pool.

`Source Volume`
The name of the standard volume.

`Progress`
The percent complete of the volume copy.

`Status`

- `Unavailable`: A volume copy is in progress to the destination volume.
- `Suspended`: The source volume went offline while a volume copy was in progress. When the source volume comes back online, the copy process resumes from the point where it stopped.

`Status-Reason`
More information about the `Status` value.

Example Show information about in-progress copy operations for controller A:

`# show volume copy-status controller a`

See also • abort volumecopy
• copy volume

show volume-groups

Description Shows information about specified volume groups or all volume groups.

Level Monitor

Syntax `show volume-groups [`*`volume-groups`*`]`

Parameters *`volume-groups`*
Optional. A comma-separated list of the names of volume groups for which to show information. If this parameter is omitted, information is shown for all volume groups.

Output **Volume group information:**

`Group Name`
The name of the volume group.

`Serial Number`
The serial number of the volume group.

`Type`
The group type, which is `Volume`.

`Number of Members`
The number of volumes in the volume group.

Volume information:

`Pool Name`
The name of the storage pool that contains the volume.

`Total Size`
The total size of the volume.

`Allocated Size`
The amount of space currently allocated to the volume.

`Serial Number`
The serial number of the volume.

`WR Policy`
Cache write policy:

- `write-back`: Write-back caching does not wait for data to be completely written to disk before signaling the host that the write is complete. This is the default and preferred setting for a fault-tolerant environment because it improves the performance of write operations and throughput.
- `write-through`: Write-through caching significantly impacts performance by waiting for data to be completely written to disk before signaling the host that the write is complete. Use this setting only when operating in an environment with low or no fault tolerance.

`Cache Opt`
Cache optimization mode:

- `standard`: Optimizes cache for both sequential and random reads. Appropriate for applications that read and write small files in random order, such as transaction-based and database update applications. This is the default.
- `no-mirror`: When this mode is enabled, each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. ULP behavior is not affected, with the exception that during failover any write data in cache will be lost.

`Read Ahead Size`
Read-ahead cache setting:

- `Disabled`: Read-ahead caching is disabled.
- `Default`: One chunk for the first access in a sequential read and one stripe for all subsequent accesses.
- `Maximum`: Maximum read-ahead size calculated by the controller.
- `64 KB`, `128 KB`, `256 KB`, `512 KB`, `1 MB`, `2 MB`, `4 MB`, `8 MB`, `16 MB`, or `32 MB`: Size selected by a user.

`Preference`

- `Archive`: The volume will stay in the SSD tier, regardless of how often the data is accessed.
- `Standard`: The volume will move to the appropriate tier, depending on how frequently its data is accessed. This is the default.
- `Performance`: The volume will stay in the SSD tier as much as possible.

`Snapshot`

- `No`: The volume is not a snapshot.
- `Yes`: The volume is a snapshot.

`Qualifier`

- `N/A`: Non-replication-specific volume such as a standard volume, snapshot, or snap pool.
- `RSR`: Replication-specific volume, such as a primary volume, secondary volume, replication snapshot, or replication image.
- `RSR (DRM Promoted Secondary)`: During an actual site failover (not a test failover), the replication set's primary and secondary volumes are shown as primary volumes with a primary-volume conflict. This qualifier enables host-based Storage Replication Adapter (SRA) software to determine which volume is the failed-over secondary volume for disaster recovery management (DRM). If both systems are online and the communication link between them is up, both systems will show similar information. When the SRA completes a restore-replication or reverse-replication operation, this volume will be shown as a secondary volume with the `RSR` qualifier.

`Volume Description`

- Blank if not set.

`WWN`
World Wide Name of the volume, used by host-based Storage Replication Adapter (SRA) software to identify the volume.

`Health`

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

`Health Reason`
If Health is not `OK`, this field shows the reason for the health state.

`Health Recommendation`
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show information about all volume groups:

```
# show volume-groups
```

Show information about volume groups `VGroup1` and `VGroup2` only:

```
# show volume-groups VGroup1,VGroup2
```

See also
- create volume-group
- delete volume-groups
- set volume-group

show volume-maps

Description Shows mapping information for volume groups and volumes.

Level Monitor

Syntax
```
show volume-maps
    [IDs]
    [all]
```

Parameters *IDs*
Optional. A comma-separated list of the names or serial numbers of volumes and volume groups for which to show mappings. If this parameter is omitted, all mappings are shown. For name syntax, see Command syntax on page 18.

`all`
Optional. Includes entries for volumes without explicit maps.

Output **Volume group information:**

`Serial Number`
The serial number of the volume group.

`Group Name`
The volume group name in the format *volume-group-name*.*, where the * represents all volumes in the group.

`Ports`

- The controller host ports to which the mapping applies.
- Blank if not mapped or mapped as `no-access`.

`Access`
The type of host access to the volume:

- `read-write`: Read and write. This is the default.
- `read-only`: Read only.
- `no-access`: No access (masked).
- `not-mapped`: Not mapped.

`Initiator-Identifier`

- For an FC or SAS initiator, its WWPN.
- `all other initiators`: The volume's default mapping.
- For a host or host group, its serial number.

`Nickname`

- The host name in the format *host-name*.*, where the * represents all initiators in the host.
- The host-group name in the format *host-group-name*.*.*, where the first * represents all hosts in the host group and the second * represents all initiators in those hosts.
- Blank if not set or for `all other initiators`.

`Name`
The name of the volume.

`Serial Number`
The serial number of the volume.

`LUN`

- The LUN that identifies the volume to a host.
- Blank if not mapped or mapped as `no-access`.

Volume information:

`Serial Number`
The serial number of the volume.

`Name`
The name of the volume.

`Ports`

- The controller host ports to which the mapping applies.
- Blank if not mapped or mapped as `no-access`.

`LUN`

- The LUN that identifies the volume to a host.
- Blank if not mapped or mapped as `no-access`.

`Access`
Type of host access to the volume:

- `read-write`: Read and write. This is the default.
- `read-only`: Read only.
- `no-access`: No access (masked).
- `not-mapped`: Not mapped.

`Identifier`

- For an FC or SAS initiator, its WWPN.0
- `all other initiators`: The volume's default mapping.

`Nickname`

- The host name in the format *host-name*.*, where the * represents all initiators in the host.
- Blank if not set or for `all other initiators`.

Example   Show mappings for all mapped volumes:

```
# show volume-maps
```

Show mappings for all volumes:

```
# show volume-maps all
```

Show mappings for volume `Vol0005`:

```
# show volume-maps Vol0005
```

See also
- show host-maps
- show hosts
- show maps
- show volumes

show volume-names

Description Shows volume names and serial numbers. This reduced form of the show volumes command is optimized for seeking basic volume information from a remote system.

Level Monitor

Syntax `show volume-names [volumes]`

Parameters *volumes*
Optional. A comma-separated list of the names or serial numbers of the volumes for which to show information. For volume syntax, see Command syntax on page 18. If this parameter is omitted, information is shown for all volumes.

Output `Name`
The name of the volume.

`Serial Number`
The serial number of the volume.

Example Show volume names and serial numbers:

`# show volume-names`

See also • show volume-maps
• show volumes

show volume-reservations

Description Shows persistent reservations for all or specified volumes. The persistent group reservations (PGR) mechanism enables application clients on multiple hosts to control access to a storage volume, and limits access by other hosts.

Each host must be registered with the storage system in order to establish a persistent reservation for a volume, thereby becoming a reservation holder.

If the system gets into an abnormal state and you need to remove all registrations and reservations for specified volumes to return them to a "clean" state, you can use the release volume command. This command must be used with care, as described in its help.

For more information about persistent reservations, see the SPC-3 specification at http://www.t10.org.

Level Monitor

Syntax `show volume-reservations all|`*`volumes`*

Parameters `all|`*`volumes`*
Specifies all volumes, or the names or serial numbers of specific volumes. For volume syntax, Command syntax on page 18.

Output `Name`
The name of the volume.

`Serial Number`
The serial number of the volume.

`Volume Reserved`

- `Free`: The volume is not reserved.
- `Reserved`: The volume has been reserved by a host.

`Host ID`
The host WWPN.

`Ports`
The controller host-port identifiers.

`Reservation Type`

- `Write Exclusive`: Write commands are only allowed for a single reservation holder.
- `Exclusive Access`: Certain access (read, write) commands are only allowed for a single reservation holder.
- `Write Exclusive - Registrants Only`: Write commands are only allowed for registered hosts. There is a single reservation holder.
- `Exclusive Access - Registrants Only`: Certain access (read, write) commands are only allowed for registered hosts. There is a single reservation holder.
- `Write Exclusive - All Registrants`: Write commands are only allowed for registered hosts. There is a single reservation holder.
- `Exclusive Access - All Registrants`: Certain access (read, write) commands are only allowed for registered hosts. There is a single reservation holder.
- `Undefined`: The volume has no persistent reservations.

Example Show reservations for all volumes:

```
# show volume-reservations all
```

See also
- release volume
- show volumes

show volumes

Description Shows information about volumes and snapshots.

Level Monitor

Syntax ```
show volumes
    [storage-pool storage-pool-ID]
    [type snapshot|standard|primary-volume|secondary-volume|replication-volume
    |all]
    [volumes|volume-groups]
```

Parameters `storage-pool` *storage-pool-ID*
Optional. The name or serial number of the storage pool that contains the volumes for which to show information. Names are case sensitive.

`type snapshot|standard|primary-volume|secondary-volume`
`|replication-volume|all`
Optional. Specifies the type of volumes to show. The `replication-volume` option shows primary and secondary volumes.

volumes|volume-groups
Optional. A comma-separated list of the names or serial numbers of volumes or volume groups for which to show information. For volume syntax, see Command syntax on page 18.

Output `Pool`
The name of the storage pool that contains the volume.

`Name`
The name of the volume.

`Total Size`
The total size of the volume.

`Allocated Size`
The amount of space currently allocated to the volume.

`Serial Number`
The serial number of the volume.

`WR Policy`
Cache write policy:

- `write-back`: Write-back caching does not wait for data to be completely written to disk before signaling the host that the write is complete. This is the default and preferred setting for a fault-tolerant environment because it improves the performance of write operations and throughput.
- `write-through`: Write-through caching significantly impacts performance by waiting for data to be completely written to disk before signaling the host that the write is complete. Use this setting only when operating in an environment with low or no fault tolerance.

`Cache Opt`
Cache optimization mode:

- `standard`: Optimizes cache for both sequential and random reads. Appropriate for applications that read and write small files in random order, such as transaction-based and database update applications. This is the default.
- `no-mirror`: When this mode is enabled, each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. ULP behavior is not affected, with the exception that during failover any write data in cache will be lost.

Read Ahead Size
Read-ahead cache setting:

- `Disabled`: Read-ahead caching is disabled.
- `Default`: One chunk for the first access in a sequential read and one stripe for all subsequent accesses.
- `Maximum`: Maximum read-ahead size calculated by the controller.
- `64 KB`, `128 KB`, `256 KB`, `512 KB`, `1 MB`, `2 MB`, `4 MB`, `8 MB`, `16 MB`, or `32 MB`: Size selected by a user.

Preference

- `Archive`: The volume will stay in the SSD tier, regardless of how often the data is accessed.
- `Standard`: The volume will move to the appropriate tier, depending on how frequently its data is accessed. This is the default.
- `Performance`: The volume will stay in the SSD tier as much as possible.

Snapshot

- `No`: The volume is not a snapshot.
- `Yes`: The volume is a snapshot.

Qualifier

- `N/A`: Non-replication-specific volume such as a standard volume, snapshot, or snap pool.
- `RSR`: Replication-specific volume, such as a primary volume, secondary volume, replication snapshot, or replication image.
- `RSR (DRM Promoted Secondary)`: During an actual site failover (not a test failover), the replication set's primary and secondary volumes are shown as primary volumes with a primary-volume conflict. This qualifier enables host-based Storage Replication Adapter (SRA) software to determine which volume is the failed-over secondary volume for disaster recovery management (DRM). If both systems are online and the communication link between them is up, both systems will show similar information. When the SRA completes a restore-replication or reverse-replication operation, this volume will be shown as a secondary volume with the `RSR` qualifier.

Volume Description

- Blank if not set.

WWN
World Wide Name of the volume, used by host-based Storage Replication Adapter (SRA) software to identify the volume.

Health

- `OK`
- `Degraded`
- `Fault`
- `Unknown`

Health Reason
If Health is not `OK`, this field shows the reason for the health state.

Health Recommendation
If Health is not `OK`, this field shows recommended actions to take to resolve the health issue.

Example Show volume information for standard volumes only:

```
# show volumes type standard
```

Show volume information for storage pool A only:

```
# show volumes storage-pool A
```

See also
- create volume
- delete volumes
- expand volume
- set volume
- show maps

show volume-statistics

Description Shows live performance statistics for all or specified volumes. For each volume these statistics quantify I/O operations between hosts and the volume. For example, each time a host writes to a volume's cache, the volume's statistics are adjusted.

Level Monitor

Syntax `show volume-statistics [volumes]`

Parameters *volumes*
Optional. A comma-separated list of the names or serial numbers of the volumes for which to show statistics. For volume syntax, see Command syntax on page 18. If this parameter is omitted, statistics are shown for all volumes.

Output `Name`
The name of the volume.

`Serial Number`
The serial number of the volume.

`Bytes per second`
The data transfer rate calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`IOPS`
Input/output operations per second, calculated over the interval since these statistics were last requested or reset. This value will be zero if it has not been requested or reset since a controller restart.

`Number of Reads`
The number of read operations since these statistics were last reset or since the controller was restarted.

`Number of Writes`
The number of write operations since these statistics were last reset or since the controller was restarted.

`Data Read`
The amount of data read since these statistics were last reset or since the controller was restarted.

`Data Written`
The amount of data written since these statistics were last reset or since the controller was restarted.

`Allocated Pages`
The number of pages allocated to the volume.

`% Performance`
The percentage of volume capacity occupied by data in the Performance tier.

`% Standard`
The percentage of volume capacity occupied by data in the Standard tier.

`% Archive`
The percentage of volume capacity occupied by data in the Archive tier.

`% RFC`
The percentage of read-cache capacity that is occupied.

`Reset Time`
Date and time, in the format *year-month-day hour*:*minutes*:*seconds*, when these statistics were last reset, either by a user or by a controller restart.

Example Show statistics for volume `Vol0001`:

```
# show volume-statistics Vol0001
```

See also
- reset all-statistics
- reset volume-statistics
- show controller-statistics
- show disk-statistics
- show host-port-statistics
- show volumes

shutdown

Description Shuts down the Storage Controller in a controller module. This ensures that a proper failover sequence is used, which includes stopping all I/O operations and writing any data in write cache to disk. If the Storage Controller in each controller module is shut down, hosts cannot access the system's data. Perform a shut down before removing a controller module or powering down the system.

△ **CAUTION:** You can continue to use the CLI when either or both Storage Controllers are shut down, but information shown might be invalid.

Level Manage

Syntax `shutdown [a|b|both]`

Parameters `a|b|both`
Optional. Specifies to shut down the Storage Controller in controller A, B, or both. If this parameter is omitted, the command affects the controller being accessed.

Example Shut down the Storage Controller in controller A:

```
# shutdown a
```

See also • restart
• show controllers

suspend replication

Description Suspends the current replication operation on the specified secondary volume. This command must be issued on the system that owns the secondary volume. Once suspended, the replication must be resumed or aborted to allow the volume to resume normal operation.

Level Manage

Syntax
```
suspend replication
    [set replication-set]
    replication-volume
```

Parameters set *replication-set*
Optional. The name or serial number of the replication set

replication-volume
The name or serial number of the secondary volume. If the name is not unique across replication sets, specify the set parameter.

Example Suspend replication of primary volume V1 to secondary volume rV1:

```
# suspend replication rV1
```

See also • abort replication
• resume replication
• show replication-sets
• show replication-volumes

test

Description Sends a test message to configured destinations for event notification and managed logs. After issuing this command, verify that the test message reached its destinations.

Level Manage

Syntax `test`
`email|managedlogs|managedlogswarn|managedlogswrap|notification|snmp`
`[region crash1|crash2|crash3|crash4|ecdebug|mc|scdebug]`

Parameters `email|managedlogs|managedlogswarn|managedlogswrap|notification|snmp`

- `email`: This option behaves the same as the notification option and remains for backward compatibility only.
- `managedlogs`: Specify this option to test receipt of the managed-logs notification that logs need to be transferred. (Event 400)
- `managedlogswarn`: Specify this option to test receipt of the managed-logs notification that logs are nearly full and must be transferred to avoid losing older entries. (Event 401)
- `managedlogswrap`: Specify this option to test receipt of the managed-logs notification that logs have wrapped and older entries may be lost. (Event 402)
- `notification`: Specify this option to test receipt of event-notification messages by every interface that is configured to receive them, such as email, SNMP, and SMI-S. (Event 312)
- `snmp`: This option behaves the same as the notification option and remains for backward compatibility only

`region crash1|crash2|crash3|crash4|ecdebug|mc|scdebug`
Optional. For use with the managed logs feature, this parameter specifies the log type (debug-data region) for which to send notifications.

- `crash1`, `crash2`, `crash3`, or `crash4`: Specify one of these options to send notification for one of the Storage Controller's four crash logs.
- `ecdebug`: Specify this option to send notification for the Expander Controller log.
- `mc`: Specify this option to send notification for the Management Controller log.
- `scdebug`: Specify this option to send notification for the Storage Controller log, which includes the event log.

If this parameter is omitted, the command sends four representative log types: `crash1`, `ecdebug`, `scdebug`, and `mc`.

Example Test receipt of event notifications by every interface that is configured to receive them:

```
# test notification
```

Test receipt of the managed-logs notification that the SC log needs to be transferred:

```
# test managedlogs region scdebug
```

See also
- set email-parameters
- set snmp-parameters

trust

Description Enables an offline storage-pool component to be brought online for emergency data recovery. This command must be enabled before each use. If used improperly this command can cause unstable operation and data loss; before use, carefully read the cautions and procedures below.

The trust command resynchronizes the time and date stamp and any other metadata in a bad disk. This makes the disk an active member of the storage-pool component again. You might need to do this when a storage-pool component is offline because a disk is failing, you have no data backup, and you want to try to recover the data from the storage-pool component. In this case, trust may work, but only as long as the failing disk continues to operate.

When the "trusted" storage-pool component is back online, back up its data and audit the data to make sure that it is intact. Then delete that storage-pool component, create a new storage-pool component, and restore data from the backup to the new storage-pool component. Using a trusted storage-pool component is only a disaster-recovery measure; the storage-pool component has no tolerance for any additional failures.

△ **CAUTION:**
1. Do not use the `trust` command when the storage system is unstable; for example, if there are many power or topology-change events.
2. The `trust` command cannot be run on a quarantined storage-pool component.
3. Never update controller-module, expansion-module, or disk firmware when the storage-pool component is offline.
4. Never clear unwritten data cache when a storage-pool component is offline.
5. Do not use the `trust` command on a storage-pool component that went offline while being expanded.
6. Do not use the `trust` command on a storage-pool component with status `CRIT`. Instead, add spares and let the system reconstruct the storage-pool component.

Steps for running the `trust` command

1. Disable background scrub.
2. Identify the cause for the storage-pool component going offline.
3. If an external issue (power, cabling, and so forth) caused the storage-pool component to go offline, fix the external issue before continuing to the next step.
4. Disable host access to the offline storage-pool component.:
   a. Determine the owning controller of the offline storage-pool component.
   b. For all online storage-pool components owned by that controller, change ownership to the partner controller.
   c. Remove the host-port cables of the owning controller of the offline storage-pool component.
5. Note the order in which the disks failed.
6. If the disks went `LEFTOVR`/failed at different times, before running the `trust` command, physically remove all disks that were members of the storage-pool component that were not in use or available when the storage-pool component was last in the critical state. This includes disks added for reconstruction of the storage-pool component. For a RAID-6 storage-pool component, remove the first two failed disks of the offline storage-pool component, according to the logs.
7. If the storage-pool component went offline in the middle of reconstruction, remove the disk being used as the reconstruction target.
8. Unseat the spare disks associated with the storage-pool component to prevent reconstruction.*
9. Reseat the remaining affected disks.
10. Enable the `trust` command.
11. Run the `trust` command on the storage-pool component.

After running the trust command

1. Perform a complete backup of the storage-pool component.
2. Delete the storage-pool component.
3. Replace disks, if necessary.
4. Re-create the storage-pool component.
5. Restore the data from the backup performed in step 1.
6. Run `verify storage-pool-components` *component-name* `fix no` to verify that the correct disk was used in the `trust` command.
7. Restore original storage-pool component ownership and reinsert host-port cables.
8. Re-enable background scrub.

* It is recommended to avoid reconstruction after using the trust command. Reconstruction causes heavy usage of disks that were already reporting errors. This usage could cause the disks to fail during reconstruction, which can cause data to be unrecoverable.

Level Diagnostic

Syntax To enable the `trust` command:

```
trust enable
```

To trust a storage-pool component:

```
trust storage-pool-component component-ID
```

Parameters `enable`
Enables the `trust` command before use.

`storage-pool component` *component-ID*
The name or serial number of the storage-pool component to trust.

Example Enable the `trust` command and then trust storage-pool component `A_01_01`:

```
# trust enable
Success: Command completed successfully.

# trust storage-pool-component A_01_01
Success: Command completed successfully.
```

See also
- show storage-pool-components
- verify storage-pool-components

unfail

Description	Allows the partner controller module to recover from a simulated failure performed with the diagnostic-level `fail` command. If you attempt to unfail a controller that is operating, the command will have no effect.
Level	Monitor
Syntax	`unfail controller a|b`
Parameters	`controller a|b` Specifies whether to unfail controller A or B. You cannot unfail the controller on which the command is issued.
Example	From controller A, unfail controller B: `# unfail controller b`
See also	• fail • show controllers

unmap volume

Description Deletes explicit mappings or the default mapping for specified volumes. When an explicit mapping is deleted, access by that host to the volume is controlled by the volume's default mapping (described in help for create volume). When a default mapping is deleted, access by hosts to the volume is controlled by any explicit mappings of those hosts to the volume.

If you want to mask access for a specific host to a specific volume, use the map volume command and set the access parameter to `no-access`.

Level Manage

Syntax
```
unmap volume
     [initiator initiators|hosts|host-groups]
     volumes|volume-groups
```

Parameters `initiator initiators|hosts|host-groups`
Optional. A comma-separated list of initiators, hosts, or host groups from which to unmap volumes. For an FC or SAS initiator, specify the nickname or 16-hex-digit WWPN. For a host, specify the name as *host-name.**. For a host group, specify the name as *host-group-name.*.**. If this parameter is omitted, mapping changes apply to all initiators not explicitly mapped (that is, to the default mapping).

volumes|volume-groups
A comma-separated list of volumes or volume groups to unmap. For a volume, specify its name or serial number. For a volume group, specify the name as *volume-group-name.**. For volume syntax, see Command syntax on page 18.

Example Delete explicit mappings for host `Host1` to volumes `V1` and `V3` (leaving the default mappings unchanged):

```
# unmap volume initiator Host1.* V1,V3
```

Delete volume `V2`'s default mapping (leaving explicit mappings unchanged):

```
# unmap volume V2
```

Delete explicit mappings for initiator `FC-port1` to volume group `MyVolumes` (leaving the default mappings unchanged):

```
# unmap volume initiator FC-port1 MyVolumes.*.*
```

See also
- map volume
- show host-maps
- show hosts
- show volume-maps
- show volumes

verify links

Description Verifies host-port links. If a remote system is specified, all link paths between the local system and the remote system are tested; otherwise, link paths between controller A and controller B in the local system are tested. The remote system must already have been added by using the create remote-system command.

For replication purposes, this tests the links to be used for replication from one system to another system. To verify bidirectional communication, run this command from the primary system to the secondary system, and then from the secondary system to the primary system.

Level Manage

Syntax
```
verify link
     [remote-system system]
     [link-type FC]
```

Parameters `remote-system` *system*
Optional. The remote system's name or the IP address of one of its controller network ports. If this parameter is omitted, links between the local controllers are verified.

`link-type FC`
Optional. Specifies the type of host-port links to verify:

- `FC`: Verify FC-to-FC links only.

If this parameter is omitted, all links are verified.

Output `Port`
The port ID in the local system.

`Type`

- `FC`: FC port.

`Links`
The IDs of linked ports in the target system.

Example Verify all links between controllers A and B in the local system:

```
# verify links
```

Verify all links between the local system and remote system `System2`:

```
# verify links remote-system System2
```

See also
- show remote-systems
- verify remote-link

verify remote-link

Description Verifies host-port links between the local system and a specified remote system. All link paths, or only paths having a specified link type, between the two systems are tested. The remote system must already have been added by using the create remote-system command.

For replication purposes, this tests the links to be used for replication from one system to another system. To verify bidirectional communication, run this command from the primary system to the secondary system, and then from the secondary system to the primary system.

Level Manage

Syntax ```
verify remote-link
     remote-system system
     [link-type FC]
```

**Parameters** remote-system *system*
The remote system's name or the IP address of one of its controller network ports.

link-type FC
Optional. Specifies the type of host-port links to verify:

• FC: Verify FC-to-FC links only.

If this parameter is omitted, all links are verified.

**Output** Port
The port ID on the local system.

Type

• FC: FC port.

Links
The IDs of linked ports on the remote system.

**Example** Verify all links between two systems:

```
# verify remote-link remote-system System2
```

See also • show remote-systems
• verify links

verify storage-pool-components

Description Analyzes fault-tolerant storage-pool components to find and fix inconsistencies between their redundancy data and their user data. This command acts on disks that are associated with a storage-pool component and are neither dedicated spares nor leftovers. This command will fix parity mismatches for RAID 6 and mirror mismatches for RAID 1. This command can be performed only for a storage-pool component whose status is FTOL (fault tolerant and online).

Verification can last over an hour, depending on storage-pool component size, utility priority, and amount of I/O activity. You can use a storage-pool component while it is being verified. To view the progress of a verify (VRFY) job, use the show storage-pool-components command.

When verification is complete, event 21 is logged and specifies the number of inconsistencies found. Such inconsistencies can indicate that a disk in the storage-pool component is going bad.

TIP: Unless you want to verify without fixing errors, it is better to use the `scrub storage-pool-components` command, which operates similarly to `verify storage-pool-components` but also can find and fix media errors for any RAID level.

Level Diagnostic

Syntax
```
verify storage-pool components
    component-IDs
    [fix yes|no]
```

Parameters *component-IDs*
A comma-separated list of the names or serial numbers of the storage-pool components to verify.

`fix yes|no`
Optional. Specifies whether or not to automatically fix parity mismatches by making parity match the data in all cases. The default is `yes`. If you specify `no`, event 21 will report any errors found and they will not be fixed.

Example Start verifying storage-pool component `A_01_01`:

```
# verify storage-pool-component A_01_01
```

See also • abort verify
• scrub storage-pool-components
• show storage-pool-components

A   Settings changed by restore defaults

This appendix summarizes the system settings that result from using the `restore defaults` command.

Table 8 Settings changed by restore defaults

Setting	Value
System information settings: • System name • System contact • System location • System information	 • `Uninitialized Name` • `Uninitialized Contact` • `Uninitialized Location` • `Uninitialized Info`
Management protocols settings: • CLI/Telnet • CLI/SSH • FTP • SNMP • WBI/HTTP • WBI/HTTPS • SMI-S • Unsecure SMI-S • Debug • In-band SES	 • `Enabled` • `Enabled` • `Enabled` • `Enabled` • `Enabled` • `Enabled` • `Enabled` • Disabled • `Disabled` • `Enabled`
Users	All configured users are deleted and replaced with default user definitions and default settings: • User: `manage`; password: `!manage` • User: `monitor`; password: `!monitor` • User: `ftp`; password: `!ftp` • Temperature scale: `Celsius` • Timeout: 30 minutes
CLI/Telnet timeout	30 minutes
Tasks and schedules	(preserved)[1]
Remote system definitions	(preserved)[2]
MC debug logs	(preserved)[3]
SC event logs	(preserved)
Time/date and NTP settings	(preserved)[4]
Network settings	(preserved)[5]
SNMP settings: • SNMP • SNMP trap notification level • SNMP trap host IPs • SNMP read community • SNMP write community	 • `Disabled` • `none` • `0.0.0.0` • `public` • `private`

Table 8 Settings changed by restore defaults (continued)

Setting	Value
SMTP settings: • Email notification • Email notify filter • Email addresses • Email server • Email domain • Email sender • Log destination • Include logs	 • `Disabled` • (none) • (none) • (none) • (none) • (none) • (none) • `Disabled`
SSL/SSH certificates	(preserved)[6]
Licenses	(preserved)
Host port settings: • FC link speed • FC topology (3XX0 models) • FC topology (3XX3 models)	 • `Auto` • `Auto` • `Point-to-Point`
Host names	(preserved)[7]
Drive spin down	`Disabled`
Advanced settings: • Partner firmware upgrade • Utility priority • SMART • Dynamic spare configuration • Enclosure polling rate • Supercap failure • CompactFlash failure • Auto write back • Disk background scrub	 • `Disabled` • `High` • (preserved)[8] • `Disabled` • 5 seconds • `Enabled` • `Enabled` • `Enabled` • `Disabled`
Native Command Queuing (NCQ)	Enabled
Managed Logs	`Disabled`

1. Factory default: no tasks or schedules.
2. Factory default: no remote system definitions.
3. Factory default: MC logs are cleared.
4. Factory default: NTP is disabled; NTP host IP address is 0.0.0.0; NTP offset is 0.
5. Factory default: IP address mode is manual; controller A IP is 10.0.0.2; controller B IP is 10.0.0.3; subnet mask is 255.255.255.0; gateway IP is 10.0.0.1.
6. Factory default: certificates are initialized.
7. Factory default: host names are cleared.
8. Factory default: no action for all disks.

Glossary

AES	Advanced Encryption Standard.
allocated page	A page of storage-pool space that has been allocated to a volume to store data.
ALUA	Asymmetric Logical Unit Access.
array	*See* storage system.
ATS	Automated tiered storage. A paged-storage feature that automatically uses the appropriate tier of disks to store data based on how frequently the data is accessed. This enables higher-cost, higher-speed disks to be used only for frequently needed data, while infrequently needed data can reside in lower-cost, lower-speed disks.
automated tiered storage	See ATS.
CAPI	Configuration Application Programming Interface. The proprietary protocol used for communication between the Storage Controller and the Management Controller in a controller module. CAPI is always enabled.
chassis	The sheetmetal housing of an enclosure.
controller A (or B)	A short way of referring to controller module A (or B).
controller enclosure	An enclosure that contains one or two controller modules.
controller module	A FRU that contains the following subsystems and devices: a Storage Controller processor; a Management Controller processor; a SAS expander and Expander Controller processor; management interfaces; cache protected by a supercapacitor pack and nonvolatile memory (CompactFlash); host, expansion, network, and service ports; and midplane connectivity. In a controller enclosure, the upper controller module is designated *A* and the lower one is designated *B*.
chunk size	The amount of contiguous data that is written to a disk in a storage-pool component before moving to the next disk in that component.
compatible disk	A disk that has enough capacity to replace a failed disk and is the same type (SAS or SATA).
default mapping	Host-access settings that are configured when a volume is created, and that apply to all initiators that are not explicitly mapped to that volume using different settings. See also explicit mapping and masking.
DES	Data Encryption Standard.
drive enclosure	An enclosure that contains one or two expansion modules. Drive enclosures can be connected to a controller enclosure to provide additional storage capacity.
drive spin down	See DSD.
DRM	Disaster recovery management. Storage-system firmware features that, when the Site Replication Adapter (SRA) feature is enabled, support the use of VMware's Site Recovery Manager to automate disaster-recovery failover and failback tasks. See also SRA.
DSD	A power-saving feature that monitors disk activity in the storage system and spins down inactive SAS and SATA disks, based on user-selectable policies.
dual-port disk	A disk that is connected to both controllers so its data path is fault-tolerant.
EC	Expander Controller. The processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. See also EMP.
EMP	Enclosure management processor. An EC subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks.
enclosure	A physical storage device that contains disk drives and other FRUs. If the enclosure contain.
Expander Controller	See EC.

expansion enclosure	See drive enclosure.
expansion module	A FRU that contains the following subsystems and devices: a SAS expander and Expander Controller processor; host, expansion, and service ports; and midplane connectivity. In a drive enclosure, the upper expansion module is designated *A* and the lower one is designated *B*.
explicit mapping	Access settings for an initiator to a volume that override the volume's default mapping. See also default mapping and masking.
failback	See recovery.
failover	In an active-active configuration, failover is the act of temporarily transferring ownership of controller resources from an offline controller to its partner controller, which remains operational. The resources include storage pools, volumes, cache data, host ID information, and LUNs and WWNs. See also recovery.
FC	Fibre Channel interface protocol.
FRU	Field-replaceable unit.
host	A user-defined group of initiators that represents a server or switch. In this product, *host* is also used to refer generically to an initiator, host, or host group.
host group	A user-defined group of hosts for ease of management.
image ID	A globally unique serial number that identifies the point-in-time image source for a volume. All volumes that have identical image IDs have identical data content, whether they be snapshots or stand-alone volumes.
initiator	An external port that the storage system is connected to. The external port may be a port in an I/O adapter in a server, or a port in a network switch.
I/O module	See IOM.
IOM	Input/output module, which can be either a controller module or an expansion module.
IOPS	I/O operations per second.
JBOD	"Just a bunch of disks." See drive enclosure.
leftover	The state of a disk that has been automatically excluded from a storage pool, and is no longer needed by the storage pool after it is reconstructed.
LFF	Large form factor (disk drive).
loop	Fibre Channel Arbitrated Loop (FC-AL) topology.
Management Controller	See MC.
mapping (or map)	Settings that specify whether a volume is presented as a storage device to a host, and how the host can access the volume. Mapping settings include an access type (read-write, read-only, or no access), controller host ports through which initiators may access the volume, and a LUN that identifies the volume to the host. See also default mapping and explicit mapping.
masking	Volume-mapping settings that specify no access to that volume by hosts. See also default mapping and explicit mapping.
MC	Management Controller. The processor (located in a controller module) that is responsible for human-computer interface and computer-computer interface functions, and interacts with the SC.
metadata	Data in the first sectors of a disk drive that stores all disk, storage-pool component, and volume specific information including storage-pool component membership or spare identification, storage-pool component ownership, volumes and snapshots in the storage-pool component, host mapping of volumes, and results of the last media scrub.
MIB	Management Information Base.
mount	To enable access to a volume from a host OS.
network port	The Ethernet port on a controller module through which its Management Controller is connected to the network.

over-committed	The amount of storage capacity that is allocated to volumes exceeds the physical capacity of the storage system.
page	A range of contiguous LBAs in a storage-pool component.
paged storage	A method of mapping logical host requests to physical storage that maps the requests to virtualized "pages" of storage that are in turn mapped to physical storage. This provides more flexibility for expanding capacity and automatically moving data than the traditional, linear method in which requests are directly mapped to storage devices.
partner firmware update	See PFU.
PFU	Partner firmware update. A feature that synchronizes the firmware in a pair of controller modules. If PFU is enabled and a firmware update is performed in one controller, that firmware revision will become the active revision and will be synchronized to the partner controller.
point-to-point	Fibre Channel Point-to-Point topology.
pool A (or B)	A short way of referring to storage pool A (or B).
primary volume	The volume that is the source of data in a replication set and that can be mapped to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume.
proxy volume	A virtual volume in the local system that represents a volume in a remote system. Proxy volumes are used internally by the controllers to perform actions such as transferring replication data.
quick rebuild	A feature that reduces the time to restore fault tolerance to a RAID-6 storage-pool component that has experienced disk failure. The quick-rebuild process rebuilds only data stripes that contain user data; data stripes that have not been allocated to user data are rebuilt in the background.
RAID 1	This RAID level uses a pair of disks, where each disk contains a complete copy of data to protect against the failure of one disk. RAID 1 is used for storage-pool components in the Performance tier and for read cache.
RAID 6	This RAID level uses block-level data striping with double distributed parity to protect against failure of two disks. RAID 6 is used for storage-pool components in the Standard and Archive tiers. Each RAID-6 component in the system contains 10 disks (8 data disks and 2 parity disks).
RAID head	See controller enclosure.
read cache	A tiered-storage feature that uses SSDs as read cache only while keeping a separate copy of the data on spinning disks. Read cache is also referred to as read flash cache.
read flash cache	See read cache.
recovery	In an active-active configuration, recovery is the act of returning ownership of controller resources to a controller (which was offline) from its partner controller. The resources include storage pools, volumes, cache data, host ID information, and LUNs and WWNs. See also failover.
RFC	Read flash cache. See read cache.
remote replication	Asynchronous (batch) replication of block-level data from a volume in a primary system to a volume in one or more secondary systems by creating a replication snapshot of the primary volume and copying the snapshot data to the secondary systems via Fibre Channel links. The capability to perform remote replication is a licensed feature (AssuredRemote).
replication image	A conceptual term for replication snapshots that have the same image ID in primary and secondary systems. These synchronized snapshots contain identical data and can be used for disaster recovery.
replication-prepared volume	A volume created for the purpose of being the secondary volume in a replication set. Replication-prepared volumes are automatically created by the Storage Management Console's Replication Setup Wizard, or they can be created manually in the CLI.

replication set	Associated primary and secondary volumes that are enabled for replication and that typically reside in two physically or geographically separate storage systems. See primary volume and secondary volume.
replication snapshot	A special type of snapshot, created by the remote replication feature, that preserves the state of data of a primary volume as it existed when the snapshot was created. For a primary volume, the replication process creates a replication snapshot on both the primary system and—when the replication of primary-volume data to the secondary volume is complete—on the secondary system. Replication snapshots cannot be mapped and are not counted toward a license limit, although they are counted toward the system's maximum number of volumes. A replication snapshot can be exported to a regular, licensed snapshot. See also replication sync point.
replication sync point	The state of a replication snapshot whose corresponding primary or secondary snapshot exists and contains identical data. For a replication set, four types of sync point are identified: the only replication snapshot that is copy-complete on any secondary system is the "only sync point"; the latest replication snapshot that is copy-complete on any secondary system is the "current sync point"; the latest replication snapshot that is copy-complete on all secondary systems is the "common sync point"; a common sync point that has been superseded by a new common sync point is an "old common sync point."
SAS	Serial Attached SCSI interface protocol or disk-drive architecture.
SC	Storage Controller. The processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller.
secondary volume	The volume that is the destination for data in a replication set and that is not accessible to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The contents of a secondary volume are in a constant state of flux and are not in a consistent state while a replication is in process. Only snapshots that are associated with a secondary volume are data consistent.
SES	SCSI Enclosure Services.
SFF	Small form factor (disk drive).
SHA	Secure Hash Algorithm.
shelf	See enclosure.
snap pool	A type of volume that stores data that is specific to snapshots of volumes in a storage pool, including copy-on-write data and data written explicitly to the snapshots. Each storage pool has one snap pool. A snap pool cannot be mapped.
snapshot	A "virtual" volume that preserves the state of a standard volume's data as it existed when the snapshot was created. Snapshots are created with a copy-on-write mechanism. A snapshot is initially created as a sparse copy of the standard (source) volume, and as new data blocks are written to the source volume, the old data blocks are written to the snapshot. A snapshot can be mapped and written to by hosts. The capability to create snapshots is a licensed feature (AssuredSnap). Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.
SRA	Storage Replication Adapter. A host-based software component that allows VMware's Site Recovery Manager to manage the storage-system firmware's disaster recovery management (DRM) features, automating disaster-recovery failover and failback tasks. The SRA uses the CLI XML API to control the storage system. See also DRM.
SSD	Solid-state drive.
standard volume	A volume that can be mapped to initiators and presented as a storage device to a host, and that is enabled for snapshots. In user interfaces, a standard volume is often referred to simply as a volume.
Storage Controller	See SC.
Storage Management Console	The web application that is embedded in each controller module and is the primary management interface for the storage system.

storage pool	One or more storage-pool components that, as a group, serve up storage pages to volumes.
storage-pool component	A RAID set in a storage pool. Storage-pool components that serve up storage pages to volumes can use RAID 1 or RAID 6. Each RAID-1 component uses 2 disks. Each RAID-6 component uses 10 disks: 8 data disks and 2 parity disks.
storage system	A controller enclosure with at least one connected drive enclosures. Product documentation and interfaces use the terms storage system and system interchangeably.
thin provisioning	A feature that allows actual storage for a volume to be assigned as data is written, rather than storage being assigned immediately for the eventual size of the volume.
tier	A class of physical storage, based on disk type, in a hierarchy of performance. The predefined tiers are: Performance, which uses SAS SSDs (high speed, low capacity); Standard, which uses enterprise-class spinning SAS disks (lower speed, higher capacity); and Archive, which uses midline spinning SAS disks (low speed, high capacity).
tier migration	The process of moving data to the appropriate performance tier based on how frequently the data is accessed. Tier migration can occur automatically in accordance with configured thresholds and policies, or can be done manually.
ULP	Unified LUN Presentation. A RAID controller feature that enables a host to access mapped volumes through any controller host port. ULP incorporates Asymmetric Logical Unit Access (ALUA) extensions.
under-provisioned	The amount of storage capacity that is allocated to volumes is less than the physical capacity of the storage system.
unwritable cache data	Cache data that has not been written to disk and is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host and disk. Unwritable cache is also called orphan data.
UTC	Universal Coordinated Time.
VDS	Volume Disk Service.
volume	A portion of a storage pool that can be used to store user data. See also standard volume, primary volume, replication-prepared volume, secondary volume, snapshot, and replication snapshot.
volume copy	An independent copy of the data in a volume. The capability to create volume copies is a licensed feature (AssuredCopy) that makes use of snapshot functionality.
volume group	A user-defined group of volumes.
VSS	Volume Shadow Copy Service.
WBI	Web-browser interface. See Storage Management Console.
WWN	World Wide Name. A globally unique 64-bit number that identifies a node process or node port.
WWNN	World Wide Node Name. A globally unique 64-bit number that identifies a node process.
WWPN	World Wide Port Name. A globally unique 64-bit number that identifies a node port.

Index