PanCLI Reference Guide

The VDURA Data Platform command-line interface

Documentation for PanFS 12.0.0

© 2002–2026 VDURA Inc. All rights reserved. VDURA, the VDURA logo, VDURA Data Platform, VeLO, VPOD, DirectFlow, PanView, PanMove, and Multi-Level Erasure Coding (MLEC) are trademarks or registered trademarks of VDURA, Inc. in the U.S. and/or other countries. All other trademarks, registered trademarks, trade names, company names, and service marks are the respective properties of their holders.

Parts of this product are derived from Berkeley Database software, licensed from Sleepycat Software.

Parts of this product are derived from Berkeley BSD systems, licensed from the University of California. The BDB software is also BSD licensed, and carries the following copyrights:

- Copyright © 1990–2009 Oracle. All rights reserved.
- Copyright © 1990, 1993, 1994, 1995 The Regents of the University of California. All rights reserved.
- Copyright © 1995, 1996 The President and Fellows of Harvard University. All rights reserved.
- Copyright © 2000–2005 INRIA, France Telecom

Parts of this product are derived from gd Version 1.8.4, licensed by Boutell.Com, Inc.

Copyright © 1982, 1986, 1991, 1993 The Regents of the University of California. All rights reserved.

© UNIX System Laboratories, Inc.

All or some portions of this file are derived from material licensed to the University of California by American Telephone and Telegraph Co. or Unix System Laboratories, Inc. and are reproduced herein with the permission of UNIX System Laboratories, Inc.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Disclaimer

VDURA has carefully checked the information in this document and believes it to be accurate. However, VDURA assumes no responsibility for any inaccuracies that this document may contain. In no event will VDURA be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

Any alterations or inconsistent replications of this document cannot be considered accurate and are void. VDURA may not be able to support issues related to altered documents that are not approved by subject matter experts.

Document version 12.0.0-2026-11-12-11:47:44

1. Introduction 1
1.1 About This Guide 1
1.2 Document Conventions 1
1.3 Additional Resources 2
1.4 Contact VDURA 3

2. PanCLI Usage Basics 4
2.1 Connect to PanCLI 4
2.2 Display Online Help 4
2.3 General Navigation 4
2.4 Edit Lines 5
2.5 Meta Character Syntax 6
2.6 Support for Pipe 6
2.7 Output Redirect Options 7
2.8 Glob Pattern Matching 8

3. PanCLI Commands 9
3.1 Command Summary 9
3.2 Command Details and Usage 21

1. Introduction

Welcome to the VDURA ® Data Platform CLI Reference Guide. System administrators who prefer command-line management to graphic user interface (GUI) management can use the VDURA command-line interface (**PanCLI**) to perform the same monitoring and management tasks that they can perform in the PanActive Manager GUI.

1.1 About This Guide

In this guide you will learn how to use the commands available in the latest version of PanCLI, the VDURA Data Platform command-line interface. For information and instructions on legacy systems, refer to older versions of documentation available at support.VDURA.com.

This guide is written for system and network administrators who are familiar with Network Attached Storage (NAS) systems and Storage Area Networks (SANs). Users should also be familiar with administering and managing file systems using Network File Systems (NFS) and Common Internet File Systems (CIFS).

A working knowledge of network protocols, network administration and management, and the Linux®, UNIX®, or Windows® operating systems is also required.

1.2 Document Conventions

This manual uses the following standard set of typographical conventions to introduce new terms, illustrate screen displays, and describe command syntax.

Convention	Meaning
boldface	Identifiers (names of classes, objects, constants, events, functions, program variables, forms, labels, and reports), environment variables, filenames, table names, column names, icons, menu items, command names, and other similar terms appear in boldface.

Convention	Meaning
`courier`	Information that the product displays and information that you are instructed to enter appear in a courier typeface.

1.3 Additional Resources

This section covers Online Help, Product Documentation, Release Notes, and Knowledge Base.

Online Help

Online help is available in PanActive Manager in two ways:

- Entire online help contents: From the Welcome page, click the Help link to view the entire online help contents for PanActive Manager.
- Page-level help: From a particular page in PanActive Manager, click the Help link to view help for that page.

Product Documentation

The full VDURA documentation set is available in PDF format from the Welcome page of PanActive Manager or from the Content tab of support.VDURA.com .

Consult the following VDURA documentation for additional information regarding VDURA products:

- VDURA Data Platform User Guide: Describes how to use PanActive Manager to monitor and manage your VDURA system.
- VDURA V5000 Certified Platform Installation and Service Guide: Describes how to install and configure your VDURA system.

Release Notes

Release notes supplement the information in this guide and are available at support.VDURA.com or from PanActive Manager Home page.

Be sure to examine the release notes before you begin the installation and setup process, as they contain important information that can improve your success with the VDURA system, such as fixed issues, version compatibility information, and known issues/workarounds from previous versions of PanFS that may still occur in the current release.

Knowledge Base

VDURA customers have access to the articles in the Knowledge Base tab on support.VDURA.com.

1.4 Contact VDURA

There are several ways to contact VDURA, depending on your needs.

Internet	**The support.VDURA.com website provides access to submit a service request and track orders online. Support.VDURA.com also offers access to custom software downloads and a robust knowledge base to assist with technical issues.**
Technical Support	Get technical support and file service requests using any of these methods: - Telephone–North America: 1-888-726-2727 - Telephone–International: 1.412.323.6439 - Email: technical-support@vdura.com
File a Service Request	support.VDURA.com

2. PanCLI Usage Basics

This chapter covers the following topics.

2.1 Connect to PanCLI

You can connect to PanCLI by connecting to the Director host via ssh. You must use the login and password for the VDURA administrator. If configured, you may also log in as guest.

The PanCLI interface is case sensitive. After you type a command at the system prompt, press Enter to execute the command. The PanCLI interface supports VT-100 command-line editing.

2.2 Display Online Help

You can display help for PanCLI commands by entering a question mark (?) after the command or typing help and the command. For example:

```
arp ?
```

or

```
help arp
```

2.3 General Navigation

General navigation features within PanCLI include the following:

Keystroke	Description
`?`	Display a list of all commands.
`[command-name] ?`	Display help for a specific command.
`[command-prefix] TAB`	Complete the partial command, starting with the given prefix.
`[command-prefix] TAB TAB`	Display the list of commands, starting with the given prefix.

The PanCLI interface is based on Tcl shell and uses double quotes or curly braces to group arguments. Single quotes are not special characters in this shell. The following special characters must be escaped (preceded by a backslash) to be used in this shell:

```
` " ; | * ? & < > () [ ] { } $ \
```

2.4 Edit Lines

The following table describes the PanCLI interface editing features.

Keystroke	Description
Ctrl-P or the up arrow key	Move the cursor to the previous line in the command.
Ctrl-N or the down arrow key	Move the cursor to the next line in the command.
Ctrl-B or the left arrow key	Move the cursor back one character.
Ctrl-F or the right arrow key	Move the cursor forward one character.
Ctrl-A	Move the cursor to the beginning of the line.
Ctrl-E	Move the cursor to the end of the command line.
Ctrl-W	Cut a region of text.
Ctrl-Y	Paste a region of text.
Ctrl-K	Delete from the cursor to the end of the line.
Ctrl-D	Delete forward one letter.
Ctrl-U	Delete to beginning of the line.
Meta-B followed by Ctrl-K	Delete to beginning of the word.
Ctrl-C	Quit the currently running command.

2.5 Meta Character Syntax

The following conventions are used to describe PanCLI command syntax:

Convention	Description
`[X]`	Required parameter
`?X?`	Optional parameter
`X|Y`	Alternative parameters
`...`	A list of parameters

Service mode commands and the pan_* tools follow standard UNIX usage notation instead, in which square brackets `[ ]` mark an optional parameter.

2.6 Support for Pipe

You can pipe to the following commands in PanCLI:

```
more, zmore, less, zless, grep, egrep, zgrep, head, tail, awk, cut, sort, uniq, wc, xargs
```

All other pipe requests will be passed as arguments to the initial command.

Operator	Description
`|`	Separates distinct commands in the pipeline. The standard output of the preceding command will be piped into the standard input of the next command.
`|&`	Separates distinct commands in the pipeline. Both standard output and standard error of the preceding command will be piped into the standard input of the next command. This form of redirection overrides forms such as 2> and >&.

2.7 Output Redirect Options

PanCLI supports the following file output redirect options:

Option	Description
`>& fileName`	Both standard output from the last command and standard error from all commands are redirected to the file named fileName, overwriting its previous contents.
`>> fileName`	Standard output from the last command is redirected to the file named fileName, appending to it rather than overwriting it.
`2>> fileName`	Standard error from all commands in the pipeline is redirected to the file named fileName, appending to it rather than overwriting it.
`>>& fileName`	Both standard output from the last command and standard errors from all commands are redirected to the file named fileName, appending to it rather than overwriting it.
`>@ fileId`	FileId must be the identifier for an open file, such as the return value from a previous call to open. Standard output from the last command is redirected to fileId's file, which must have been opened for writing.
`2>@ fileId`	FileId must be the identifier for an open file, such as the return value from a previous call to open. Standard error from all commands in the pipeline is redirected to fileId's file. The file must have been opened for writing.

Option	Description
`>&@ fileId`	FileId must be the identifier for an open file, such as the return value from a previous call to open. Both standard output from the last command and standard error from all commands are redirected to fileId's file. The file must have been opened for writing.

2.8 Glob Pattern Matching

PanCLI supports a number of standard UNIX commands. For these commands, filename substitution is supported. For any arguments to the UNIX commands, if the word contains any of the special glob characters, it is a candidate for filename substitution, also known as globbing. This word is then regarded as a pattern (glob-pattern), and replaced with an alphabetically sorted list of file names which match the pattern. To avoid globbing, precede any special glob characters with two backslashes. Special characters include:

Character	Description
`?`	Match any single character.
`*`	Match any sequence of zero or more characters.
`\[chars\]`	Match any single character in chars. If chars contains a sequence of the form a-b then any character between a and b (inclusive) will match.
`\x`	Match the character x.
`{a,b,...}`	Match any of the strings a, b, etc.
`~`	Words beginning with ~ refer to the home directory.

3. PanCLI Commands

IMPORTANT: As of PanFS 11.5.0, the use of bladeset in PanCLI command names and options has been deprecated. As reflected in the software and this guide, the correct terminology is now storageset. For now, users can still use bladeset, however the system will return a message indicating that the command has been deprecated: "The bladeset command is deprecated; please use storageset instead."

3.1 Command Summary

The following PanCLI commands are described in this chapter.

Command	Summary
about	Shows customer ID, current system software revision, and copyright information.
ad	Manages the VDURA Active Directory client configuration. AD feature is unavailable in this release.
admin	Displays and sets the admin user identity and group access.
arp	Displays the IP-Ethernet address translation table.
awk	Pattern-directed scanning and processing language.
cacert	Displays or downloads the CA certificate file.
cat	Concatenates and prints file.
cd	Changes the working directory.
cifs	Manages CIFS connectivity to the system. CIFS feature is unavailable in this release.
clear	Clears the screen.

Command	Summary
collect_logs	Collects logs, corefiles, and system databases from nodes.
config	Saves or loads system configuration on another host.
core	Lists and views corefiles and stack traces (service mode command).
cp	Copies files.
cut	Cuts out selected portions of each line of a file.
date	Displays and sets the current time and date.
df	Displays free disk space for this node.
diag	Runs system diagnostics.
diag_exec	Executes the given command on the specified nodes (service mode command).
directflow	Displays and copies DirectFlow software release files.
dns	Displays and sets DNS domain names, servers, and delegation.
drain	Manages draining of nodes.
du	Displays file system block usage for each file argument and for each directory in the file hierarchy rooted in each directory argument.
enclosure_info	Displays diagnostic information for the enclosure. This command is only eligible to run on enclosure master nodes.
eventlog	Displays the historical log of system events.
exit	Exits the command line interface.

Command	Summary
firmwareupdate	Manages system-wide disk firmware update for enclosure-based nodes.
fsckutil	Verifies and repairs file system on one or more VPODs (service mode command).
fsdb_osd	Allows metadata editing for OSDfs filesystem (service mode command).
fsdb_osdv4	Interactive editor for osdv4 file system (service mode command).
getdevinfo	Displays or sets local node configuration values (service mode command).
grep	Searches files for a pattern.
guest	Manages access to the read-only guest account.
gwctl	Displays and sets gateway tuning parameters (service mode command).
gzip	Compresses a file.
halt	Halts this node (service mode command).
head	Displays the first part of a file.
help	Lists commands. Type a question mark (?) after a command for help with a specific command.
history	Saves and runs history of commands.
ifconfig	Displays network interface parameters.
install	Runs the system Install Wizard.
ip	Displays and configures network interface parameters.

Command	Summary
ipmi_setup	Configures baseboard management controllers (BMCs) for remote IPMI access (service mode command).
ipmitool	Controls IPMI-enabled devices (service mode command).
iplist	Displays the list of IP addresses for the system.
kill	Terminates or signals a running process (service mode command).
killall	Terminates a running process by name (service mode command).
last	Displays last logins of users and ttys.
ldap	Manages VDURA LDAP client configuration.
ldapsearch	Opens a connection to an LDAP server, binds, and performs a search using specified parameters.
ldapwhoami	Opens a connection to an LDAP server, binds, and performs a 'who am i' operation.
ls	Lists directory contents.
lw-lsa	Manages Likewise security and authentication service (service mode command).
lwnet	Manages CIFS/SMB server (service mode command).
lwregshell	Manages the Likewise registry (service mode command).
lwsm	Controls Likewise services (service mode command).

Command	Summary
md5	Calculates an MD5 message-digest fingerprint (checksum) for a file.
mkdir	Creates the directories named as operands, in the order specified, using mode rwxrwxrwx (0777) as modified by the current umask.
more	Interactively views a text file.
nameresorder	Manages priority ordering of username resolution services.
netconfig	Displays or sets system network configuration.
netdiag	Runs network diagnostics and displays results.
netperf	Turns the netperf daemon on and off on all nodes in the system (service mode command).
netstat	Shows network status.
netswitch	Displays switch information and status.
nfs	Manages NFS connectivity to the system. NFS feature is unavailable in this release.
nis	Configures Network Information Services (NIS) domain and servers.
notify	Displays and updates error notification preferences.
ntp	Configures the Network Time Protocol (NTP) server.
objutil	Lists and manipulates filesystem objects (service mode command).
osdv4_admin	Modifies the configuration of a running iopathv4 (only for Ultra nodes - only run under supervision of VDURA Technical Support).

Command	Summary
osdv4_compress_metadata_db.py compress_db	Compacts VPOD metadata databases. This command should only be run under the supervision of VDURA Technical Support (service mode command).
osdv4_rebalance	Reconfigures OSDv4 filesystem to store copy-on-write metadata of objects on flash drive and triggers rebalancing (only for Ultra nodes - only run under supervision of VDURA Technical Support).
pan_atime	Obtains Storage Node estimate of access time on a file or directory.
pan_attr	Allows setting and querying attributes of a PanFS object.
pan_collect_logs	Copies log files and collects various information that is useful in debugging problems with the DirectFlow Client.
pan_debug_enum_datasets	Lists available datasets (service mode command).
pan_debug_print_layout_table	Prints the SM internal device table (service mode command).
pan_debug_print_one_stat	Prints a named debug_dataset from a single target (service mode command).
pan_debug_trace_archive_collect	Collects traces from a circular buffer to a file, with archiving (service mode command).
pan_debug_trace_collect	Collects traces from a circular file (service mode command).
pan_debug_trace_destroy	Deletes the specified trace buffer from memory and closes the file if spilling traces.
pan_debug_trace_info	Gives details about the specified trace buffer.

Command	Summary
pan_debug_trace_list	Lists all the traces in use.
pan_debug_trace_modify	Modifies trace filters for an existing trace buffer (service mode command).
pan_debug_trace_reset	Resets trace configuration to system default.
pan_debug_trace_spill	Spills remaining traces in trace buffer to file.
pan_debug_trace_start	Starts tracing.
pan_debug_trace_stop	Stops tracing.
pan_df	Displays information about the file system on which each file resides or all file systems by default.
pan_du	Summarizes disk usage of each file, recursively for directories.
pan_filecount	Counts the number of files and directories in VDURA volumes.
pan_fm_config_get	Views internal details for an FM ID (service mode command).
pan_fm_id_translate	Reports internal mappings between ID domains (service mode command).
pan_fm_vol_cache_examine	Examines the internal FM volume cache (service mode command).
pan_fsrc_get_progress	Displays the current status of FSRC and progress details.
pan_get_master_mserver	Displays the IP address of the Director running the MConsole service.
pan_iperf	Pan_iperf is iperf renamed to avoid name conflicts.

Command	Summary
pan_journal_config_maxlevelstore	Queries and updates journald MaxLevelStore parameter (service mode command).
pan_lm_lock_list	Views lock manager internals (service mode command).
pan_net_enclosure_nettest	Performs basic network connectivity and throughput tests between nodes and between nodes and external hosts (service mode command).
pan_nettest	Runs iperf TCP throughput tests to characterize the network between a set of N clients and a set of M servers.
pan_op_log_tool	Formats File Manager, Storage Manager, and Lock Manager transaction logs (service mode command).
pan_rm_blobdb_util	Views and edits the RM blobdb database (service mode command).
pan_rm_brick_exhume	Removes a reconstruction record and optionally returns a buried brick to service (service mode command).
pan_rm_configdb_util	Views and edits the RM configdb database (service mode command).
pan_rm_get_repset	Views replication set (repset) details specific to the RM (service mode command).
pan_rm_move_backup	Changes the backup location of a volume service (service mode command).
pan_rm_move_service	Moves a service to another node (service mode command).
pan_rm_ptp_tool	Updates the live system PTP database (service mode command).

Command	Summary
pan_rm_rehome_realm_root	Moves the root volume to a new StorageSet (service mode command).
pan_rm_volume_info	Lists internal volume details specific to the RM database (service mode command).
pan_sam_util_verify_parity	Checks or corrects parity in a file (service mode command).
pan_sitesync	The manager portion of SiteSync, the utility which replicates an existing directory tree in a new location.
pan_sm_attr_check	Repairs storage attributes of an object (service mode command).
pan_sm_cap_query	Displays caps that are outstanding in Storage Manager (service mode command).
pan_sm_cleanup_bg_scrubber_dir	Cleans up bg_scrubber files (service mode command).
pan_sm_print_map	Prints the layout of an object across Storage Nodes.
pan_sm_repair_query	Prints the repair data from specified storage manager (service mode command).
pan_switch_software_update	Upgrades switch firmware (service mode command).
pan_tcp_capture_client	Client utility to capture tcpdump data and related network or NFS statistics and configuration data.
pan_trace	Collects diagnostic trace buffers.
pan_trace_control	Controls the following trace actions: reset, prego (service mode command).

Command	Summary
panfs_trace	Collects traces from PanFS clients and gateways.
panic	Panics this node (service mode command).
panstack	Decodes stack traces (service mode command).
paritywarnings	Enables/disables parity warnings (service mode command).
pasxml	Gets pasxml URL output.
pax	Reads, writes, and lists the members of an archive file, and will copy directory hierarchies.
phonehome	Modifies the phone home schedule (service mode command).
ping	Sends ICMP ECHO_REQUEST packets to network hosts.
process	Views status for, start, and stop local processes (service mode command).
prompt	Displays or sets the PanCLI prompt for all nodes.
property	Manages property bits (service mode command).
ps	Shows the status of system processes.
pwd	Returns working directory name.
raidutil	Displays RAID information (service mode command).
realm	Configures the realm DirectFlow exports file.
reboot	Interactive command to reboot this node (service mode command).
register	Completes product registration.

Command	Summary
rm	Attempts to remove the non-directory type files specified on the command line.
rmdir	Removes the directory entry specified by each directory argument, provided it is empty.
rpancli	Executes PanCLI commands on multiple nodes.
rsyslog	Views or changes remote syslog collectors and facility.
s3	Implements support of the Amazon AWS S3 API as one of the interfaces for accessing the data on VDURA realms. Deprecated as of 12.0.0; use PanCTL to configure S3.
samba	Manages Samba server (service mode command).
scp	Copies files between hosts on a network.
secure	Displays and updates security status.
sha1	Calculates a SHA-1 message-digest fingerprint (checksum) for a file.
smtp	Displays, sets, or tests the SMTP email server and from address.
snmp	Configures SNMP support.
sort	Sorts lines of text files.
storageset	Configures StorageSets.
support	Configures communication with my.VDURA.com.
syscontrol	Sets system and hardware controls.
sysctl	Manages sysctl values (service mode command).

Command	Summary
sysmap	Displays the system hardware map.
sysstat	Displays system status information.
tail	Displays the contents of a file or, by default, its standard input, to the standard output.
tar	Creates, adds files to, or extracts files from an archive file in tar format, called a tarfile.
tcpdump	Allows capturing network traffic into or out of the node, with arbitrary filters (service mode command).
terminal	Sets or shows terminal characteristics.
timezone	Sets or shows time zone.
top	Displays or updates information about the top CPU processes.
traceroute	Prints the route packets to the given network host.
trpost	Creates post-process binary trace output for easy viewing (service mode command).
uniq	Reads the specified input_file comparing adjacent lines, and writes a copy of each unique input line to the output_file.
upgrade	Manages system software revision.
uptime	Displays this node's uptime and recent load averages.
usermap	Displays and configures usermap configuration.
userquota	Displays and configures user and group quota configuration.
vi	Edits a text file.

Command	Summary
volservice	Manages backup of volume metadata.
volume	Manages volumes.
wbinfo	Queries idmapd_panasas.
wc	Displays the number of lines, words, and bytes contained in each input file or standard input (if no file is specified) to the standard output.
whoami	Prints the username associated with the current effective userid.
xargs	Constructs and executes argument lists.
ypcat	Prints the values of all keys from the NIS database specified by map name, which may be a map name or a map nickname.
ypmatch	Prints the values of one or more keys from an NIS map.
zcat	Decompresses files then concatenates and prints in the manner of the cat command.
zgrep	Prints lines matching a pattern.
zless	Interactively views a zipped file.
zmore	Allows examination of compressed or plain text files one screenful at a time on a soft-copy terminal.

3.2 Command Details and Usage

This section provides detailed descriptions of available PanCLI commands. Please read and be sure you understand the behavior of any command before running it. If you have any questions, please contact VDURA Technical Support.

about

Description

Displays the software version number, build, customer ID, and copyright information.

ad

Description

The CIFS feature on which this command depends is currently unsupported. Use of this command may be refused.

Options

Option	Description
debug	Display the current debug level for name resolution services. To set the debug level, use the **nameresorder** command. **Note:** Any debug level above 0 will impact performance.
enabled ?yes|no?	Display or enable/disable Active Directory for UNIX user, group, and netgroup name resolution. The AD feature allows the system to perform username to UID mappings, group name to GID mappings, and netgroup lookups via Active Directory. To enable AD, the following CIFS values must be set: - Windows Domain Name - Windows Account Name - Windows Account Password These values can be updated using the cifs PanCLI command. **Note:** You must also add AD to the name resolution search ordering before the system will perform name resolution via Active Directory. Use the **nameresorder** command to manage name resolution search order.

admin

Description

Displays and sets the admin user identity and group access.

Syntax

```
admin ?[subcommand] ?[args]??
```

Example

Set the admin's email to bob@example.com:

```
admin email bob@example.com
```

Options

Option	Description
`email ?clear|[addr]?`	Display, clear, or set the administrator email address.
`group unix|windows ?[group]?`	Display or set the UNIX and WINS groups with administrator access. Use an empty string ("") to clear the value.
`name`	Display the name of the administrator for the system.
`password`	Give a prompt to change the admin password.
`timeout ?[minutes]?`	Display or set the current timeout setting (the number of minutes a Web login session remains valid).

arp

Description

Displays the Internet-to-Ethernet address translation tables used by the Address Resolution Protocol (ARP).

Consult the UNIX manual page for details about the **arp** command.

awk

Description

Pattern-directed scanning and processing language.

Awk scans each input file for lines that match any of a set of patterns specified literally in prog or in one or more files specified as -f progfile. With each pattern there can be an associated action that will be performed when a line of a file matches the pattern. Each line is matched against the pattern portion of every pattern-action statement; the associated action is performed for each matched pattern.

- The file name **-** means the standard input.
- Any file of the form **var=value** is treated as an assignment, not a filename, and is executed at the time it would have been opened if it were a filename.
- The option **-v** followed by **var=value** is an assignment to be done before **prog** is executed; any number of **-v** options may be present.
- The **-F fs** option defines the input field separator to be the regular expression **fs**.

An input line is normally made up of fields separated by white space or by regular expression **fs**. The fields are denoted **$1**, **$2**, etc., while **$0** refers to the entire line. If **fs** is null, the input line is split into one field per character.

For more information about statements and variables, see the PanCLI online help command **help awk** .

Syntax

```
awk [-F fs] [-v var=value] ['prog' | -f progfile] [file ...]
```

cacert

Description

This command displays or downloads the CA certificate file. Without the optional argument, it displays the CA certificate content. If the optional **[file]** argument is given, the command copies the CA certificate file to the specified file on a remote host via **scp**. The **[file]** argument must be formatted as **[user]@[host]:[file]**. The command prompts for a password for **scp** to copy the file.

Syntax

```
cacert ?[file]?
```

Example

Display the CA certificate:

```
cacert
```

Download the CA certificate file using the [user]@[host]:[file] argument:

```
cacert b_smith@10.72.7.11:/var/tmp/admin/cacert.pem
Enter password:
```

cat

Description

The **cat** utility reads files sequentially, writing them to the standard output. The file operands are processed in command-line order. If file is a single dash (-) or absent, **cat** reads from the standard input. If file is a UNIX domain socket, **cat** connects to it and then reads it until EOF. This complements the UNIX domain binding capability available in **inetd(8)**.

Syntax

```
cat [-benstuv] [file ...]
```

Options

Option	Description
`-b`	Number the non-blank output lines, starting at 1.
`-e`	Display non-printing characters (see the **-v** option), and display a dollar sign ($) at the end of each line.
`-n`	Number the output lines, starting at 1.
`-s`	Squeeze multiple adjacent empty lines, causing the output to be single spaced.
`-t`	Display non-printing characters (see the **-v** option), and display tab characters as **^I**.
`-u`	Guarantee that the output is unbuffered.
`-v`	Display non-printing characters so they are visible. Control characters print as **^X** for control-X; the delete character (octal 0177) prints as **^?**. Non-ASCII characters (with the high bit set) are printed as **M-** (for meta) followed by the character for the low 7 bits.

cd

Description

Change working directory.

Syntax

```
cd [directory]
```

cifs

Description

The CIFS gateway feature is currently unsupported. Use of this command and its subcommands may be refused.

Syntax

```
cifs ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Specify the Organizational Unit (OU) where computer accounts are created:

```
cifs computerou Computers/Servers/Unix
```

Add a group:

```
cifs localgroups add Engg 231
```

Add users to a group:

```
cifs localgroups addusers Engg Tom,Mark
```

Delete matching groups:

```
cifs localgroups delete OperationsTmp SalesTmp2
```

Delete all the matching users:

```
cifs localgroups deleteusers Engg Mark
```

Deploy groups file:

```
cifs localgroups deploy [file]
```

File format:

```
# "Group name" GID
"Marketing Group" 107
IT 109
Engg 200
```

Display the list of groups:

```
cifs localgroups list
```

Display groups and members:

```
cifs localgroups members Engg
```

Deploy members file:

```
cifs localgroups membersdeploy [file]
```

File format:

```
# "Group Name" Members(comma separated list of users)
IT Mark,Joe,Tom
"Marketing Group" Tom,"Jon S",Tim
```

Add a user:

```
cifs localusers add Tom uid 205 group Engg enable 1
```

Delete a user:

```
cifs localusers delete Tom
```

Deploy a users file:

```
cifs localusers deploy [file]
```

File format:

```
# "User name" UID "Group name" Enabled Password
Administrator 0 Administrators 0 ""
Tom 105 Users 0 ""
"Joe Senior" 106 Users 1 81D6CEE0D16AE931B73C59D7E0C089C0
Guest 0 Guests 0 ""
```

List users:

```
cifs localusers list
```

Change password:

```
cifs localusers password Mark
```

Print the status of a user:

```
cifs localusers status Mark
```

Disable the user:

```
cifs localusers status Mark 0
```

Change the primary group of the user:

```
cifs localusers update Mark Engg
```

Options

Option	Description
`account ?[name]?`	Display or set the name of the CIFS administrator user. Use an empty string (“”) to clear the value.
`computerou [computer-account-OU]`	Display or set the CIFS Computer account Organizational Unit (OU). The OU string reads from top to bottom without RDNs and is delimited by /. If left empty, this will default to the standard Active Directory location of Computers. Use an empty string (“”) to clear the value.
`debug ?[level]?`	A service mode command to manage CIFS debug detail. The level can be any number between 0 and 10, with indicating no debugging and 10 being the most verbose debug level. Debugging CIFS can impact system performance.
`domainname [domain-name]`	Display or set the CIFS domain name. Use an empty string (“”) to clear the value.

Option	Description
`enabled ?no|[yes ?workgroup?]?`	To enable CIFS, a number of other CIFS values are required to be set. - Select no to disable CIFS. This is the default. - Select yes to enable CIFS in domain mode. Before doing this, run cifs subcommands netbios, domainname, account, and password to set required configuration values. - Select yes workgroup to enable CIFS in workgroup mode. Before doing this, run the cifs netbios command to set the CIFS NetBIOS name.
`idletimeout ?[minutes]?`	Display or set the number of minutes to timeout idle connections to CIFS shares.
`legacymode ?yes|no?`	Display or enable/disable CIFS Legacy mode. - If no option is given, the current mode will be displayed. - yes: disable client caching, but run CIFS across all Director Nodes. - no: enable client caching, but run CIFS only on a single DirectorNode.

Option	Description

`localgroups add [groupname] ?[GID]?`
`localgroups addusers [groupname] [username] ?[,username] ...?`
`localgroups delete [groupname-pattern1] ? [groupname-pattern2] ...?`
`localgroups deleteusers [groupname] [username-pattern1] ? [username-pattern2] ...?`
`localgroups deploy [file]`
`localgroups ?list ?[groupname-pattern1] ? [groupname-pattern2] ...??`
`localgroups members ?[groupname-pattern1] ? [groupname-pattern2] ...?`
`localgroups membersdeploy [file]`
	Manage localgroups in the realm:
	- The add option adds a group name with its optional GID. Group name can contain only the A-Za-z 0-9_.#-" characters. If a GID has not already been assigned by another name resolution provider, such as NIS, you can assign one here by entering a number between 100-99999. This number cannot match the GID for any other localgroup you have created. To assign no GID, either specify no argument or enter 0.
	- The addusers option adds the given list of users to the given group. Enter a comma-separated list of existing usernames to join the group.
	- The delete option deletes all groups matching the given pattern(s), except built-in groups. Only groups with no members can be deleted.
	- The deleteusers option deletes all the matching users from the given group name. You cannot delete a user from its primary group.
	- The deploy option deploys the localgroups configuration file. If [file] looks like `?[user]@?[host]:[file],` the command prompts for a password and uses scp to fetch the file; otherwise it looks for the file in PanFS. If no matching PanFS file is found, it checks the local filesystem of the node that hosts PanActive Manager.
	- The list option prints groups. If patterns are given, the command only lists groups whose names match the patterns. If no option is given, the command prints all groups.
	- The members option prints all groups and their members. If patterns are given, it only lists members of the groups whose names match the patterns.
	- The membersdeploy option deploys the local

Option	Description
	members configuration file. If [file] looks like `?[user]@?[host]:[file], the command` prompts for a password and uses scp to fetch the file; otherwise it looks for the file in PanFS. If no matching PanFS file is found, it checks the local filesystem of the node that hosts PanActive Manager.

Option	Description

`localusers add [username] ?[UID]? ?[primary group name]? ?[enable [0|1]?`
`localusers delete [username-pattern1] ?[username-pattern2] ...?`
`localusers deploy [file]`
`localusers ?list? ?[username-pattern1] ?[username-pattern2] ...?`
`localusers password [username]`
`localusers status [username] ?[enable]?`
`localusers update [username] [groupname]`
	Manage localusers in the realm.
	- The add option adds a username with its UID, groupname, enable status, and prompts for a password if enable is selected. All parameters except username are optional. The value of enable can be either 0 to disable or 1 to enable. Username can contain only the A-Za-z 0-9_.#-" characters. If a UID has not already been assigned by another name resolution provider such as NIS, you can assign one here by entering a number between 100-99999. This number cannot match the UID for any other localuser you have created. To assign no UID, either specify no argument or enter 0. The default values of the UID, group, and enable are 0, Users, and 0 respectively.
	- The delete option deletes all users matching the given pattern(s) except the system users Administrator and Guest.
	- The deploy option deploys the localusers configuration file. If [file] looks like `?[user]@?[host]:[file]` , the command prompts for a password and uses scp to fetch the file; otherwise it looks for the file in PanFS. If no matching PanFS file is found, it checks the local filesystem of the node that hosts PanActive Manager.
	- The list option prints users. If patterns are given, the command only lists users whose names match the patterns. If no option is given, the command prints all users.
	- The password option prompts for a new password and confirmation. Empty passwords are allowed.
	- The status option prints the status of a given user. The enable argument changes the status of the user. The value of enable can be either 0 to

Option	Description
	disable or 1 to enable. Disabling an account prevents logging in by clients using that account. - The update option updates the primary group of the user.
`cifs master ?[director]?`	Display CIFS master server IP address or move CIFS master server to another Director.
`netbios ?[name]?`	Display or set the system's NETBIOS name. Use an empty string ("") to clear the value. NetBIOS name cannot be longer than 15 characters.
`password`	Interactive command to store the CIFS password.
`shares`	Lists all CIFS shares; same as **cifs shares list** .

Option	Description
`shares add [sharename] [path] ?[perm] ?[maxusers] ?[comment]???` `shares edit [sharename] [new-sharename] ?[path] ?[perm] ?[comment] ?[forceUpdate]???`	Add or edit a CIFS share. - sharename : the name of the share. - path : The root volume is referred to as `/` . Other volumes appear as top-level directories beneath `/` . For example, a volume named “home” appears as /home. The correct path must exist and be online when the share is created. - perm : Permissions specifies share-level access for non-admin users. You can specify all, readonly, or none. The default value is all, which means all users have full control and only file-level permissions are used. - comment : A descriptive comment contained within double-quotes. - *forceUpdate:* If the share’s status indicates that the actual share permission is different from the database entry, forceUpdate with value 1 will update the share’s permissions to the database value. Adding a share or changing the path to a share automatically grants members of the Domain Admins group full control rights over the shared directory. Use a Windows client to set additional permissions or remove this access.
`shares delete [sharename-pattern] ?[sharename-pattern2] ...?`	Delete the CIFS shares whose names match the glob-style patterns.
`shares ?list ?sharename-pattern ?sort name|path|perm|maxusers???`	List CIFS shares matching the given glob pattern. The default is to list all shares. Choose between sorting the results by share name, path, permissions, or the maximum number of users. Shares are sorted by share name by default.

Option	Description

`cifs vip ?list ?primary|secondary??`
`cifs vip add (primary|secondary) [address] [netmask] ?[subnet]?`
`cifs vip delete (primary|secondary) [address] [netmask]`
	Manage CIFS Virtual IP addresses.
	CIFS Virtual IP addresses provide CIFS clients a way to connect to the CIFS Master service, which can be assigned to any Director Node. If the CIFS Master node dies or restarts, the CIFS Master role will be reassigned to another node and CIFS Virtual IPs will be serviced by this node providing IP takeover functionality. Unlike IP aliases, CIFS Virtual IP addresses are accessible from the CIFS Master node only; they are not distributed across several nodes in the cluster. CIFS Virtual IP address are not available in CIFS Legacy mode; all configured VIPs will be inactive if CIFS Legacy mode is on. See the help cifs legacymode command.

clear

Description

Clears the screen.

collect_logs

Description

Collects logs commonly used to diagnose issues from specified nodes. VDURA recommends to always collect logs from all replication set (repset) nodes in addition to any other nodes of interest. In the event of a PanActive Manager outage, the repset option will continue to find the correct nodes to collect from. The host option followed by an IP address should also always work.

Syntax

```
collect_logs ?[option] [option] ...?
```

Example

Gather logs for all nodes in VCH-2 and VCH-5, as well as all nodes in the replication set (repset):

```
collect_logs enclosure VCH-2 enclosure VCH-5 repset
```

Gather logs for the enclosure master nodes in VCH-2 and VCH-5, as well as all nodes in the replication set (repset):

```
collect_logs enclosuremaster enclosure VCH-2 enclosure VCH-5 repset
```

Gather logs for all Storage Nodes in the StorageSet Set 2, as well as all nodes in the replication set (repset):

```
collect_logs storageset "Set 2" repset
```

Gather logs excluding corefiles for all Directors in VCH-4, as well as all nodes in the replication set (repset):

```
collect_logs nocores directors enclosure Shelf-4 repset
```

Gather logs for both primary and backup Directors for the /home and /build volumes, as well as all Directors in the replication set (repset):

```
collect_logs volumeowner home volumeowner build repset
```

Gather logs for the Storage Nodes that cover the file /home/myFile, both the primary and backup Directors for the /home volume, and all Directors in the replication set (repset):

```
collect_logs fileowner /home/myFile repset
```

Gather all core/stack files that are less than two days old for the given node:

```
collect_logs host VCH-1,3 match /var/corefiles/* age 2 days
```

By default, collect_logs runs with a timeout value of 900 seconds, or 15 minutes. This timeout is often not adequate for logs from many hosts or for large logs with corefiles. To allow the command to run longer, use a long timeout value such as 3600 (one hour) or 7200 (two hours) for large logs collections.

You can also specify timeout 0 to specify running with no timeout. If timeout 0 is used, collect_logs will run until it either completes or the user ends it with CTRL+C. If you're collecting logs from a host that you think may not respond, consider not using the timeout 0 parameter.

The example below collects files needed to read a corefile; if the nocores option is used, these files will not be collected:

```
[pancli] collect_logs timeout 3600 host {node} waitsavecore age 2 days
```

Note the longer value to prevent a timeout while the command is running.

The example below collects logs for a core where the core sweeper should have already had time to run:

```
[pancli] collect_logs timeout 3600 host {node} age 2 days
```

The age 2 days option instructs the command to only gather the previous two days of data. This covers the date of the crash so you get cores from that date based on error messages telling you what day the crash occurred without extra data from prior to then.

In the following examples, host {node} is used to specify a particular node or nodes by either name or IP address:

```
[pancli] collect_logs host {node} age 12 hours
[pancli] collect_logs host {node} host {node} age 2 days
```

In the following examples, the parameter directly after collect_logs is used to specify a particular group of nodes:

```
[pancli] collect_logs repset age 2 days
[pancli] collect_logs directors age 12 hours
[pancli] collect_logs storage age 2 days
[pancli] collect_logs storageset "Set 1" age 2 days
[pancli] collect_logs enclosure {enclosure} age 2 days
```

In the following examples, the parameter directly after collect_logs is used to specify volume owner, file owner, or enclosuremaster nodes.

```
[pancli] collect_logs volumeowner {volume_name} age 2 days
[pancli] collect_logs volumeowner {volume_name} age 12 hours
[pancli] collect_logs fileowner {/volume_name/path/to/file or object_ID} age 12 hours
[pancli] collect_logs enclosuremaster age 4 days
[pancli] collect_logs enclosuremaster enclosure {enclosure} age 4 days
```

> **Note:** The fileowner option does not work on fenced objects.

For a failed or stuck reconstruction, use the age argument to go back to a time shortly before when reconstruction failed:

```
[pancli] collect_logs directors match /var/log/all* age 12 hours
```

If the volume has snapshots, also get the realm database:

```
[pancli] support servicemode yes
[pancli_service] diag_exec -timeout 1000 "pan_dump_realm" > /tmp/rmdb_{SR-XXXXXXX}
```

Use the match option to collect files that match a specified pattern:

```
[pancli] collect_logs all match /var/corefiles/*stack match /var/crash/*stack age 2 days
[pancli] collect_logs host {node} match /var/log/all* age 8 months
[pancli] collect_logs match /var/log/diag/driveS*.log match /var/log/all.log directors
```

Linux node health check only:

```
[pancli] collect_logs match /pan/var/log/diag/driveS*.log match /pan/var/log/all.log
storagemodules
```

Log collections containing match report an error if the node does not have matching files; this is normal.

If logs fail to collect, ping the node. If the ping succeeds, try to ssh into the host and collect logs locally.

Options

Option	Description
`age [count] [units]`	Do not collect /**var/log/*.[0-9]+** files and corefiles that are more than count units old, where valid units are minute(s), hour(s), day(s), week(s), month(s). Default is one week. If the **match** option is specified, it defaults to unlimited age. If the count is **0**, age is not measured.
`all`	Collect logs from all nodes in the system. This option is only available when **collect_logs** is run on a Director in the replication set (repset).
`container`	Collect logs from all VDURA Data Platforms in the system. This option is only available when **collect_logs** is run on a Director in the replication set (repset).
`directors`	Collect logs from all Directors in the system. This option is available only when running **collect_logs** on a Director in the replication set (repset).

Option	Description
`enclosure [name]`	Collect logs from all nodes in the named enclosure. You can use this option multiple times to specify multiple enclosures. This option is available only when running **collect_logs** on a Director in the replication set (repset).
`enclosuremaster`	Collect logs from nodes with the enclosuremaster role. One enclosuremaster node in each enclosure controls the enclosure hardware. This option is only available when the command is run on a Director in the replication set (repset).
`fileowner [PanFS file]|[Object ID]`	Collect logs from the nodes that store a portion of the given PanFS file or Object ID, as well as the primary and backup Directors for the file or object's volume. Owners of multiple files can be specified by using this option more than once.
`host [name]`	Collect logs from the node with this name, IP address, or serial number. You can use this option multiple times to specify multiple nodes.
`match [file path pattern]`	Collect only those files whose names match the specified glob-style pattern. Use a glob pattern that captures the full absolute path to the desired files (e.g. `/var/log/*http*` for files in **/var/log** that have the substring **http** somewhere in their name). Multiple patterns can be specified by using this option more than once.
`nocores`	Do not collect any corefiles or executables. Stack files are always gathered.
`nodbs`	Do not collect any system databases.
`nodiags`	Do not run any diagnostic tools before collecting logs.

Option	Description
`panfs [dir]`	Copy the output file to the given PanFS directory . The directory will be created if it does not already exist.
`range [ip]-[ip]`	Collect logs from nodes in this range of IP addresses. You can use this option multiple times to specify multiple ranges. This option is available only when running **collect_logs** on a Director in the replication set (repset).
`repset`	Collect logs from all Directors in the replication set (repset). This option is available only when running **collect_logs** on a Director in the replication set (repset).
`scp [user] [host] [dir]`	**Scp** the output file to the given user, host, and directory arguments. You will be prompted to enter a password. The target directory must exist on that host already.
`storage`	Collect logs from all Storage Nodes in the system. This option is only available when the command is run on a Director in the replication set (repset).
`storagemodules`	Collect logs from all storage modules in the system. This option is only available when **collect_logs** is run on a Director in the replication set (repset).
`storageset [name]`	Collect logs from all nodes in the named StorageSet. You can use this option multiple times to specify multiple StorageSets. This option is available only when running **collect_logs** on a Director in the replication set (repset).
`timeout [seconds]`	The timeout for **scp** in seconds. The default is 900 seconds (15 minutes). Use **0** for no timeout.

Option	Description
`upload [user]`	Upload the output file to **upload.vdura.com** as the given user. You will be prompted to enter a password.
`volumeowner [volume name]`	Collect logs from the Director that owns the named volume. If Volume Services Failover is enabled, also collect logs from the backup Director for the given volume. You can use this option multiple times to specify multiple volume owners. This option is available only when running **collect_logs** on a Director in the replication set (repset).
`waitsavecore`	Wait until **savecore** is no longer running before collecting logs.

config

Description

Saves or loads system configuration on another host.

An automatic hourly backup of config data is available on secondary Directors (multi-Director systems) or on Storage Nodes (single-Director systems). This archive can be used for emergency restoration of the system. The backup file can be found at **/var/tmp/realm_config.tgz.latest.hourly**.

Syntax

```
config [subcommand]
```

Options

Option	Description
`load`	Load system configuration from a file on a remote host.
`save`	Save system configuration to a file on a remote host.

core

Description

Service mode command to list and view corefiles and stack traces.

Syntax

```
core ?[subcommand] ?[arg]??
```

Example

Run the contents of the script /home/textdump_script:

```
core textscript /home/textdump_script
```

Example script:

```
textdump set; capture on; show allpcpu; bt;ps; alltrace; show alllock; call doadump; reset
```

Options

Option	Description
`(none)`	List corefiles and corresponding executable files that exist on this node.
`dumptype ?text|mini?`	Specify the type of kernel core dump for all 64-bit hardware in the system. - **mini**: Minidumps only dump in-use kernel pages, not unmapped or userspace pages. - **text**: Textdumps write a TAR file of text-format information (rather than a memory image) to the swap partition.
`gcore [execfile] [pid]`	Run **gcore -s** on the given executable and process.

Option	Description
`gdb [execfile]|-k [corefile]`	Run **gdb** on the specified executable and corefile. No executable file is needed if the core is a kernel corefile; just give **-k** as a placeholder.
`stack [corefile]`	Display the stack for the corefile.
`textscript ?[file]?`	The textscript is used as a preprocessor for **ddb** to create text dump output. This script is only applicable on 64-bit hardware in systems where the **dumptype** is configured as **text**. The script runs when the kernel debugger is entered as the result of a panic. To deploy an updated **textscript**, first create a PanFS file containing the desired script. The file must contain the commands to run on **ddb**. Note that the system will not validate the contents of the script.

cp

Description

Copies files.

If the destination file does not exist, the mode of the source file is used as modified by the file mode creation mask (**umask**, see **csh(1)**).

If the source file has its set user ID bit on, that bit is removed unless both the source file and the destination file are owned by the same user. If the source file has its set group ID bit on, that bit is removed unless both the source file and the destination file are in the same group and the user is a member of that group. If both the set user ID and set group ID bits are set, all of the above conditions must be fulfilled or both bits are removed.

Appropriate permissions are required for file creation or overwriting. Symbolic links are always followed unless the **-R** flag is set, in which case symbolic links are not followed, by default. The **-H** and **-L** flags (in conjunction with the **-R** flag) cause symbolic links to be followed as described below. The **-H**, **-L**, and **-P** options are ignored unless the **-R** option is specified. In addition, these options override each other and the command's actions are determined by the last one specified.

Syntax

```
cp ?[options]? [source] [target]
```

Example

Copy the contents of the source_file to the target_file:

```
cp [-R [-H | -L | -P]] [-f | -i] [-pv] source_file target_file
```

For each destination file that already exists, its contents are overwritten if permissions allow. Its mode, user ID, and group ID are unchanged unless the -p option was specified.

Copy the contents of each named source_file to the destination target_directory:

```
cp [-R [-H | -L | -P]] [-f | -i] [-pv] source_file ... target_directory
```

Target_directory must exist unless there is only one named source_file which is a directory, and the -R flag is specified.

The names of the files themselves are not changed. If cp detects an attempt to copy a file to itself, the copy will fail.

Options

Option	Description
`-H`	If the **-R** option is additionally specified, symbolic links on the command line are followed. Symbolic links encountered in the tree traversal are not followed.
`-L`	If the **-R** option is additionally specified, all symbolic links are followed.
`-P`	If the **-R** option is additionally specified, no symbolic links are followed. This is the default.
`-R`	If **source_file** designates a directory, **cp** copies the directory and the entire subtree connected at that point. This option also causes symbolic links to be copied rather than indirected through, and for **cp** to create special files rather than copying them as normal files. Created directories have the same mode as the corresponding source directory, unmodified by the process' umask. Note that **cp** copies hard linked files as separate files. If you need to preserve hard links, consider using **tar(1)**, **cpio(1)**, or **pax(1)** instead.
`-f`	For each existing destination pathname, remove it and create a new file without prompting for confirmation regardless of its permissions. The **-f** option overrides any previous **-i** options.

Option	Description
`-i`	Cause **cp** to write a prompt to the standard error output before copying a file that would overwrite an existing file. If the response from the standard input begins with the characters **y** or **Y**, the file copy is attempted. The **-i** option overrides any previous **-f** options.
`-p`	Cause **cp** to preserve in the copy as many of the modification time, access time, file flags, file mode, user ID, and group ID values as allowed by permissions. If the user ID and group ID cannot be preserved, no error message is displayed and the exit value is not altered. If the source file has its set user ID bit on and the user ID cannot be preserved, the set user ID bit is not preserved in the copy's permissions. If the source file has its set group ID bit on and the group ID cannot be preserved, the set group ID bit is not preserved in the copy's permissions. If the source file has both its set user ID and set group ID bits on, and either the user ID or group ID cannot be preserved, neither the set user ID nor set group ID bits are preserved in the copy's permissions.
`-v`	Cause **cp** to be verbose, showing files as they are copied.

cut

Description

Cuts out selected portions of each line of a file.

The **cut** utility cuts out selected portions of each line (as specified by list) from each file and writes them to the standard output. If no file arguments are specified, or a file argument is a single dash (-), **cut** reads from the standard input. The items specified by list can be in terms of column position or in terms of fields delimited by a special character. Column numbering starts from 1.

The **list** option argument is a comma or white space separated set of increasing numbers and/or number ranges. Number ranges consist of a number, a dash (-), and a second number and select the fields or columns from the first number to the second, inclusive. Numbers or number ranges may be preceded by a dash, which selects all fields or columns from 1 to the last number. Numbers or number ranges may be followed by a dash, which selects all fields or columns from the last number to the end of the line. Numbers and number ranges may be repeated, overlapping, and in any order. It is not an error to select fields or columns not present in the input line.

Syntax

```
cut -b list [-n] [file ...]
cut -c list [file ...]
cut -f list [-d delim] [-s] [file ...]
```

Options

Option	Description
`-b list`	The list specifies byte positions.
`-c list`	The list specifies character positions.
`-d delim`	Use *delim* as the field delimiter character instead of the tab character.
`-f list`	The list specifies fields, separated in the input by the field delimiter character (see the **-d** option.) Output fields are separated by a single occurrence of the field delimiter character.

Option	Description
`-n`	Do not split multi-byte characters. Characters will only be output if at least one byte is selected and, after a prefix of zero or more unselected bytes, the rest of the bytes that form the character are selected.
`-s`	Suppress lines with no field delimiter characters. Unless specified, lines with no delimiters are passed through unmodified.

date

Description

Displays or sets the current date and time.

> **Note:** The **date** command cannot be used to change date/time on external NTP servers.

Syntax

```
date ?[mm/dd/ccyy] [hh:mm]?:[ss]??
```

Example

Show the current date and time:

```
date
```

Set the date to November 13, 2019, 4:40 PM:

```
date 11/13/2019 16:40
```

Options

Option	Description
`(none)`	Display the current date and time of the local host.

Option	Description
`?[mm/dd/ccyy] [hh:mm]?:[ss]??`	Change the current date and/or time (sets the system clock). Changing the system clock using this command will force a system reboot. This action can only be taken on hosts in the replication set (repset). Changing the system clock is only allowed if the system is not configured to use external NTP servers for time synchronization. Use the **ntp** command to check NTP configuration. Date/time formats include: - cc – Century, 19 or 20, prepended to yy . - yy – Year, last two digits. - mm – Month, 01 through 12. - dd – Day, 01 through 31. - hh – Hour, 0 through 23. - mm – Minutes, 0 through 59 - ss – Seconds, 0 through 61 (59 plus a maximum of two leap seconds)

df

Description

Displays free disk space for the node.

Syntax

```
df [-b | -g | -h | -H | -k | -m | -P] [-ailn] [-t type]
```

Options

Option	Description
`-a`	Show all mount points, including those that were mounted with the MNT_IGNORE flag.
`-b`	Use 512-byte blocks rather than the default. This overrides the BLOCKSIZE specification from the environment.
`-g`	Use 1073741824-byte (1-Gbyte) blocks rather than the default. Note that this overrides the BLOCKSIZE specification from the environment.
`-H`	Human-readable output. Use unit suffixes: Byte, Kilobyte, Megabyte, Gigabyte, Terabyte, and Petabyte to reduce the number of digits to three or less using base 10 for sizes.
`-h`	Human-readable output. Use unit suffixes: Byte, Kilobyte, Megabyte, Gigabyte, Terabyte, and Petabyte to reduce the number of digits to three or less using base 2 for sizes.
`-i`	Include statistics on the number of free inodes.
`-k`	Use 1024-byte (1-Kbyte) blocks rather than the default. This overrides the BLOCKSIZE specification from the environment.

Option	Description
`-l`	Only display information about locally-mounted filesystems.
`-m`	Use 1048576-byte (1-Mbyte) blocks rather than the default. This overrides the BLOCKSIZE specification from the environment.
`-n`	Print the previously obtained statistics from the filesystems. This option should be used if it is possible that one or more filesystems are in a state such that they will not be able to provide statistics without a long delay. When this option is specified, **df** will not request new statistics from the filesystems, but will respond with the possibly stale statistics that were previously obtained.
`-P`	Use POSIX compliant output of 512-byte blocks rather than the default. This overrides the BLOCKSIZE specification from the environment.
`-t`	Only print statistics for filesystems of the specified types. More than one type may be specified in a comma separated list. The list of filesystem types can be prefixed with **no** to specify the filesystem types for which action should not be taken. For example, the **df** command `df -t nonfs,mfs` lists all filesystems except those of type NFS and MFS. The **lsvfs(1)** command can be used to find out the types of filesystems that are available on the system.

diag

Description

Runs system diagnostics.

Syntax

```
diag ?-args "[test options]"? check|start|list|result ?[options]?
```

Example

Check the status of netstat command executed with -i and -n options on all nodes:

```
diag check -args "-i -n" netstat
```

Run the NTP diagnostic on all Directors:

```
diag start ntp -directorblades
```

Run the enclosure diagnostic on all enclosures:

```
diag start -shelfmaster
```

Run the enclosure diagnostic for a specific enclosure:

```
diag start -shelfmaster -shelf [name|serial]
```

Options

Option	Description
`check [testPattern] ?[node options]?`	Check the status of the executed diagnostics matching the test pattern. Node options identify which node(s) to check the diagnostic results on. Values can be IP addresses, the keyword “all”, or any other options accepted by the **iplist** command. If no node options are specified, the results will be displayed for all applicable nodes. Use the **-args** option to check the result of the tests run with specific arguments.
`list ?[testPattern]?`	List the available diagnostic tests. If a pattern is given, only list tests whose names match the pattern.
`result [testPattern] ?[node options]? [file]`	Copy the results of a diagnostic test to a PanFS file. Node options identify which node(s) to check the diagnostic results on. Values can be IP addresses, the keyword “all”, or any other options accepted by the **iplist** command. If no node options are specified, the results will be gathered for all applicable nodes. Files must be specified using a full path starting with **/volume-name**.
`start ?-args "[test options]"? [testName] ? [node options]?`	Start a diagnostic test on any number of nodes. Node options identify which node(s) to run the diagnostic test on. Values can be IP addresses, the keyword “all”, or any other options accepted by the **iplist** command. If no node options are specified, the diagnostic test will be run on all applicable nodes. Use the **-args** option to provide arguments to the test.

diag_exec

Description

Executes the given command on the specified nodes.

Either the **-interactive** option or a command must be entered. The command can be any shell command that can be executed on a node. Wildcards will not be substituted. Note that commands will be run serially on the requested nodes. Do not use this command to launch interactive commands remotely.

Note that this command can only be run in service mode and should only be run under the direction of VDURA Technical Support.

Syntax

```
diag_exec ?[options]? interactive|["command"] ?[node-options]?
```

Options

Node Options

If running interactively, one or more node options are required. Node options include:

Option	Description
`host [ip]`	Use this option for each IP that you want to contact.
`all`	Use this option to run the command on all nodes.

Categories available in the **iplist** PanCLI command can also be used:

Option	Description
`containers`	All VDURA Data Platforms in the system.
`directors`	All Directors in the system.
`enclosure [name]`	Nodes in the specified enclosure. Multiple enclosures can be specified by using this option more than once.

Option	Description
`fileowner [PanFS file]|[Object ID]`	Storage Nodes that store a portion of the given PanFS file or Object ID as well as the primary and backup Directors for the file or object's volume. The path specified in **[PanFS file]** must not include **/panfs/<realm-name>/** at the start of the path. Owners of multiple files can be specified by using this option more than once.
`host [name]`	Node with this name, IP address, or serial. Multiple hosts can be specified by using this option more than once. If using an IP address, you can omit the "host" term and just enter the IP address.
`range [ip]-[ip]`	The specified range of IP addresses. Multiple hosts can be specified by using this option more than once.
`repset`	Directors in the replication set (repset).
`storage`	All Storage Nodes in the system.
`storageset [name]`	All nodes in a StorageSet. Multiple StorageSets can be specified by using this option more than once.
`volume`	Volume with the specified name.
`volumeowner [volume name]`	Director that owns the given volume. If Volume **Services Failover** is enabled, also get the IP address of the backup Director for the given volume. Multiple volume owners can be specified by using this option more than once.

Other Options

Option	Description
`help`	Show usage message.
`timeout [sec]`	The number of seconds to wait for a node to respond before quitting. A 0 indicates no timeout. Default timeout is 20 seconds.

Example

Send a command to all Directors in the realm:

```
[pancli_service] diag_exec ?[PanCLI command]? directors
```

Send a command to a specific IP address:

```
[pancli_service] diag_exec ?[PanCLI command]? 10.1.2.3
```

Send a command to a node with this name, IP address, or serial. Multiple hosts can be specified by using this option more than once.

```
[pancli_service] diag_exec ?[PanCLI command]? host VCH-1,1
```

Send a command to all Directors in the replication set (repset):

```
[pancli_service] diag_exec ?[PanCLI command]? repset
```

Example command issuing ipmitool fru list to 10.101.1.64:

```
[pancli_service] diag_exec "ipmitool fru list" 10.101.1.64
Contacting 10.101.1.64:
Tue Oct 21 19:07:16 GMT 2025: Executing "ipmitool fru list"...
Tue Oct 21 19:07:17 GMT 2025: Executed "ipmitool fru list" successfully:
FRU Device Description : Builtin FRU Device (ID 0)
Chassis Type : Other
Chassis Extra : 532200003f4005001
Board Mfg Date : Tue Jan 14 14:37:00 2025 GMT
Board Mfg : AIC
Board Product : CASSIOPEIA
Board Serial : BR80HC010200242
Product Name : CASSIOPEIA
Product Extra : 6a8cc2bebfde11d301b79c6b007aaedf
```

Selecting an appropriate category will run a command on all the nodes matching that category. Example

command issuing ipmitool fru list to all Directors:

```
[pancli_service] diag_exec "ipmitool fru list" -directors
Contacting 10.101.1.64:
Tue Oct 21 19:18:38 GMT 2025: Executing "ipmitool fru list"...
Tue Oct 21 19:18:39 GMT 2025: Executed "ipmitool fru list" successfully:
FRU Device Description : Builtin FRU Device (ID 0)
Chassis Type : Other
Chassis Extra : 532200003f4005001
Board Mfg Date : Tue Jan 14 14:37:00 2025 GMT
Board Mfg : AIC
Board Product : CASSIOPEIA
Board Serial : BR80HC010200242
Product Name : CASSIOPEIA
Product Extra : 6a8cc2bebfde11d301b79c6b007aaedf
Contacting 10.101.1.65:
Tue Oct 21 19:18:39 GMT 2025: Executing "ipmitool fru list"...
Tue Oct 21 19:18:39 GMT 2025: Executed "ipmitool fru list" successfully:
FRU Device Description : Builtin FRU Device (ID 0)
Chassis Type : Other
Chassis Part Number : XPS102CAYP00D03
Chassis Serial : 515-25032500600002
Chassis Extra : 52f200013f4005001
Board Mfg Date : Thu Jan 16 20:24:00 2025 GMT
Board Mfg : AIC
Board Product : CASSIOPEIA
Board Serial : BR80HC010200020
Product Manufacturer : AIC
Product Name : SB102-CA
Product Part Number : XPS102CAYP00D03
Product Version : 01
Product Serial : 515-25032500600002
Product Asset Tag : 52f200013f4005001
Product Extra : 6a92ab8ebfde11d301b79c6b007aafbc
Contacting 10.101.1.15:
Tue Oct 21 19:18:39 GMT 2025: Executing "ipmitool fru list"...
Tue Oct 21 19:18:39 GMT 2025: Executed "ipmitool fru list" successfully:
FRU Device Description : Builtin FRU Device (ID 0)
Chassis Type : Other
```

```
Chassis Part Number : XPS102CAYP00D03
Chassis Serial : 515-25032500600004
Chassis Extra : 52f200003f4005001
Board Mfg Date : Wed Jan 15 04:50:00 2025 GMT
Board Mfg : AIC
Board Product : CASSIOPEIA
Board Serial : BR80HC010200083
Product Manufacturer : AIC
Product Name : SB102-CA
Product Part Number : XPS102CAYP00D03
Product Version : 01
Product Serial : 515-25032500600004
Product Asset Tag : 52f200003f4005001
Product Extra : 6a89b4e8bfde11d301b79c6b007aaf1a
```

directflow

Description

Displays or copies DirectFlow® client software.

Syntax

```
directflow ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Use the name of the file given by calling the directflow versions command. To copy the DirectFlow image file to the /home/bob/client_software directory:

```
directflow copy [file] /home/bob/client_software
```

Options

Option	Description
`cancel`	Cancel the copy in progress.
`copy [file] [dir]`	Copy the given DirectFlow software image file to the specified PanFS directory.
`progress`	Display the progress of the current copy request.

Option	Description
`versions`	Display the release information for all available versions of DirectFlow client software. Currently not available: PanActive Link cannot be used to download DirectFlow client software or to show the latest available versions. Download the client from support.VDURA.com instead. PanActive Link must be enabled (with web access) for downloads and version information to be available via PanActive Manager or PanCLI. In addition, if the system is behind a firewall, the customer may need to specify a web proxy host and port.

dns

Description

Displays and sets DNS domain names, servers, and delegation.

Syntax

```
dns ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`delegation ?yes|no?`	Display whether DNS delegation is enabled or disabled and turns on or off DNS delegation.
`domains`	Display the DNS domain names.
`domains add [domain_name ]`	Add a DNS domain name.
`domains delete [domain_name]`	Delete a DNS domain name.
`domains clear`	Clear all DNS domain names.
`servers`	Display the DNS server IP addresses.
`servers add [ip]`	Add a DNS IP server address. A maximum of three total DNS servers are supported.
`servers delete [ip]`	Delete a DNS server IP address.
`servers clear`	Clear all DNS server IP addresses.
`services`	Display the DNS services.

Option	Description
`services add [name] (nfs|cifs) (all|dhcp|aliases|preferaliases| [ipsAndSubnets])`	The **add** option adds a new DNS service with the given DNS component name for either NFS or CIFS. The last argument specifies which nodes the service applies to: - **all**: The DNS name will resolve to all IP addresses of all the Directors. This includes both DHCP IP addresses and IP aliases. - **dhcp**: The DNS name will resolve only to DHCP assigned IP addresses of all the Directors. - **aliases**: The DNS name will resolve to only alias assigned addresses of all the Directors. This could resolve to an empty set of addresses if no aliases are configured. - **preferaliases**: The DNS name will resolve to only alias assigned addresses of all the Directors. If there are no aliases assigned, then resolution will fall back to DHCP assigned addresses of all the Directors. - **ipsAndSubnets**: A list of space separated IP addresses (including DHCP and/or alias IPs) and/or subnet/masks (e.g. 10.1.81.0/255.255.255.0). The address of a Director will be included if it matches an IP address in this list or if it falls in a subnet/netmask in the list.
`services delete [name]`	Delete the DNS service with the given name.
`zone ?[name]?`	Display or set the zone used for DNS delegation.

drain

Description

Manages draining of nodes.

Syntax

```
drain ?status|start|stop? [args...]
```

Options

Option	Description
`start [component]`	Start draining nodes for the given component. An enclosure or node component can be specified by name or serial. If an enclosure is given, start draining all eligible nodes in that enclosure.
`status ?[component]? ?-p?`	Display the drain status for all nodes in the system. Specify a component to limit results to a specific node or to nodes in a specific enclosure. An enclosure or node component can be specified by name or serial. Use the **-p** option to keep checking status every few seconds until none of the requested nodes are draining. Drain status also includes a listing of matching nodes in the drained and failed states.
`stop ?[component]?`	Cancel drain midway for all draining nodes in the system. Specify a component to stop drain on a specific node or on nodes in a specific enclosure. An enclosure or node component can be specified by name or serial.

du

Description

Displays file system block usage for each file argument and for each directory in the file hierarchy rooted in each directory argument. If no file is specified, usage of the hierarchy rooted in the current directory is displayed.

This command counts the storage used by symbolic links and not the files they reference unless the **-H** or **-L** option is specified. If either the **-H** or **-L** option is specified, storage used by any symbolic links that are followed is not counted or displayed.

Files having multiple hard links are counted and displayed a single time per **du** execution.

ENVIRONMENT BLOCKSIZE: If the environment variable BLOCKSIZE is set and the **-k** option is not specified, the block counts will be displayed in units of that size block. If BLOCKSIZE is not set and the **-k** option is not specified, the block counts will be displayed in 512-byte blocks.

Syntax

```
du [-H | -L | -P] [-I mask] [-a | -s | -d depth] [-c] [-h | -k] [-r] [-x] [file ...]
```

Options

Option	Description
`-H`	Symbolic links on the command line are followed; symbolic links in file hierarchies are not followed.
`-L`	Symbolic links on the command line and in file hierarchies are followed.
`-I mask`	Ignore files and directories matching the specified mask.
`-P`	No symbolic links are followed. This is the default.
`-a`	Display an entry for each file in a file hierarchy.

Option	Description
`-h`	Human-readable output. Use unit suffixes Byte, Kilobyte, Megabyte, Gigabyte, Terabyte, and Petabyte.
`-r`	Generate messages about directories that cannot be read, files that cannot be opened, and so on. This is the default. This option exists solely for conformance with X/Open Portability Guide Issue 4 (XPG4).
`-s`	Display an entry for each specified file. Equivalent to **-d 0**.
`-d depth`	Display an entry for all files and directories **depth** directories deep.
`-c`	Display a grand total.
`-k`	Display block counts in 1024-byte (1-Kbyte) blocks.
`-x`	Filesystem mount points are not traversed.

enclosure_info

Description

Displays diagnostic information for the enclosure. This command is only eligible to run on enclosure master nodes.

Syntax

```
enclosure_info
```

Example

Sample output:

```
Enclosure: 46190002402001001 - ActiveStor Enclosure
================================================================================
Enclosure Fans
--------------
Enclosure Fan 1 4900 RPM okay (500 - 25400)
Enclosure Fan 2 5000 RPM okay (500 - 25400)
================================================================================
Server: 46190005302001001 - ASD-200
================================================================================
Server Temperatures
-------------------
Server CPU 1 Temp. 41 C okay (5 - 100)
Server Socket1 VRM Temp. 43 C okay (5 - 100)
Server Socket1 DIMM A1 Temp. 31 C okay (5 - 85)
Server Socket1 DIMM B1 Temp. 30 C okay (5 - 85)
Server Socket1 DIMM C1 Temp. (---- - ----)
Server Socket1 DIMM C2 Temp. (---- - ----)
Server Socket1 DIMM D1 Temp. 28 C okay (5 - 85)
Server Socket1 DIMM E1 Temp. 28 C okay (5 - 85)
Server Socket1 DIMM F1 Temp. 28 C okay (5 - 85)
Server Socket1/A-C VRM Temp. 38 C okay (5 - 100)
Server Socket1/D-F VRM Temp. 38 C okay (5 - 100)
Server Incoming Air Temp. 24 C okay (5 - 50)
Server Motherboard Temp. 56 C okay (5 - 90)
Server System Temp. 22 C okay (5 - 85)
Server SIOM Slot Temp. 48 C okay (5 - 85)
Server Drive nda0 Temp. 42 C okay (0 - 68)
Server Core 1 Temp. 37 C okay (5 - 99)
Server Core 2 Temp. 40 C okay (5 - 99)
Server Core 3 Temp. 35 C okay (5 - 99)
Server Core 4 Temp. 36 C okay (5 - 99)
Server Core 5 Temp. 37 C okay (5 - 99)
Server Core 6 Temp. 39 C okay (5 - 99)
Server Core 7 Temp. 39 C okay (5 - 99)
Server Core 8 Temp. 38 C okay (5 - 99)
Server Voltages
---------------
```

```
Server PVnn 1.01 V okay (0.86 - 1.14)
Server 1.05V 1.03 V okay (0.90 - 1.20)
Server Socket1/A-C 1.19 V okay (1.03 - 1.37)
Server Socket1/D-F 1.19 V okay (1.03 - 1.37)
Server CPU 1 1.82 V okay (1.26 - 2.30)
Server 1.8V 1.80 V okay (1.46 - 2.13)
Server 3.3V (Vcc) 3.27 V okay (2.70 - 3.96)
Server 3.3V (Vsb) 3.36 V okay (2.70 - 3.97)
Server 5V (Vcc) 5.10 V okay (4.14 - 6.09)
Server 5V (Vsb) 4.99 V okay (4.06 - 5.95)
Server 12V 12.18 V okay (10.54 - 13.28)
Server Fans
-----------
================================================================================
PS1: P2K22CL13EB0230 - ActiveStor Director Power Supply
================================================================================
PS1 Temperatures
----------------
Power Supply 1 Temp. 1 36 C okay (5 - 75)
Power Supply 1 Temp. 2 47 C okay (5 - 75)
PS1 Voltages
------------
Power Supply 1 AC Input 199.75 V okay (80.00 - 275.00)
Power Supply 1 12V 12.19 V okay (11.16 - 12.84)
PS1 Fans
--------
Power Supply 1 Fan 1 4160 RPM okay (3000 - 9200)
Power Supply 1 Fan 2 5024 RPM okay (3000 - 12000)
================================================================================
PS2: P2K22CL13EB0229 - ActiveStor Director Power Supply
================================================================================
PS2 Temperatures
----------------
Power Supply 2 Temp. 1 36 C okay (5 - 75)
Power Supply 2 Temp. 2 46 C okay (5 - 75)
PS2 Voltages
------------
Power Supply 2 AC Input 199.75 V okay (80.00 - 275.00)
Power Supply 2 12V 12.18 V okay (11.16 - 12.84)
```

```
PS2 Fans
--------
Power Supply 2 Fan 1 3840 RPM okay (3000 - 9200)
Power Supply 2 Fan 2 5120 RPM okay (3000 - 12000)
================================================================================
/var/log/diag/46190002402001001.2022.05.11.16.54.21.pancli_diag_enclosure_info.log
0
```

eventlog

Description

Displays the historical log of system events.

Syntax

```
eventlog ?[options ...]?
```

Example

Copy the event log database to another machine:

```
eventlog -count max -output csv > eventLog.csv
scp eventLog.csv user@host:/path
rm eventLog.csv
```

Find events from the 24 to 48 hours ago for enclosure VCH-2:

```
eventlog -from "48 hours ago" -to "24 hours ago" -output tab -enclosure VCH-2
```

List the most recent 1000 events sorted in descending order:

```
eventlog -count 1000
```

List the most recent 10 Critical or Warning events:

```
eventlog -category critical -category warning
```

List all events for StorageSet “Set 1”:

```
eventlog -count max -storageset set1
```

List all events whose message matches *offline*

```
eventlog -count max -message offline
```

Options

Option	Description
`(none)`	Display the last 10 system events of all categories, sorted by ascending time.
`node [name|ip|serial]`	Limit results to events for the given node.
`category [category]`	Limit results to events with the given category. Valid categories include: critical, warning, syschanges, and adminaction. Multiple category options may be used to show multiple categories.
`code [code]`	Limit results to events with the given event code. Multiple code options may be used to show multiple codes.
`columns [summary|all]`	Display summary or all details of events. This option applies only if the output format is plaintext. Default value for plaintext output is **summary**. For **tab** and **csv** outputs, all columns are always displayed.
`count ([number]|max)`	Results will include up to the given number of events. Use **0** or **max** to get the system maximum of 20,000. If no **to**, **from**, or **count** option is given, the default count is 10.

Option	Description
`from [time]`	Limit results to events that occurred at or after the given time. Times can be specified in a variety of formats: - 10/31/06 1:56 pm - October 31, 2006 13:56:00 - 31 October 2006 01:56:00 PM - Tue, 31 Oct 06 13:56 - 061031 1:56 PM - 2006-10-31 13:56:00 - 31-OCT-06 01:56:00 pm In addition, relative times, such as **yesterday** or **3 days ago** may be used.
`hardware [hardware-serial]`	Limit results to events that occurred on the given hardware. The serial can identify a node, enclosure, battery, fan, switch, etc.
`message [glob-pattern]`	Limit results to events whose message matches the given glob pattern.
`output (plaintext|tab|csv)`	Control the output format. Choose between plaintext, tab separated, and csv formats. The default value is **plaintext**. **Tab** and **csv** formats produce one line of output for each log entry with tabs or commas separating the fields.
`enclosure [name|serial]`	Limit results to events for the given enclosure.
`sort (ascending|descending)`	Control the sort order. By default, results are sorted by time in ascending order.
`storageset [name|id]`	Limit results to events for the given StorageSet.
`to [time]`	Limit results to events that occurred at or before the given time. See **from** option above for acceptable time formats.

Option	Description
`type [component-type]`	Limit results to events for the given component type. Valid component types include: **Battery**, **StorageSet**, **Director**, **Fan**, **OSD**, **Power Supply**, **Enclosure**, **Snapshot**, **Switch**, **Volume**, and **Volume Service**.
`user [username]`	Limit results to events attributed to the given user.
`volservice [volservice-id]`	Limit results to events attributed to the given volume service.
`volume [name|id]`	Limit results to events for the given volume.

exit

Description

Exits the **PanCLI** utility.

```
exit ?code?
```

Options

Option	Description
`(none)`	Exit command line interface.
`code`	If an integer code is given, the PanCLI shell will exit with that code.

firmwareupdate

Description

Manages system-wide disk firmware update for enclosure-based nodes.

Syntax

```
firmwareupdate ?[options]?
```

Options

Option	Description
`cancel`	Cancel system firmware update. Once a system-wide disk firmware update is cancelled, updates will not be started on any more nodes. Any update that is already underway will continue, and its progress/success will not be monitored.
`start (storage|director)`	Start system-wide disk firmware update. The administrator can choose whether to run the disk firmware update on all Storage Nodes or all Directors. Each of the requested nodes will be checked to see which ones, if any, need an update. Nodes that need an update will be updated and rebooted sequentially, ordered by rack/enclosure name and slot. If any node does not return from reboot, then the system-wide update is aborted and no further nodes will be updated.

Option	Description
`?status ?-poll? ?[numEvents]??`	Check the status of system-wide disk firmware update. This command displays the status of the currently running system-wide disk firmware update. If the **-poll** option is used, the command will continue to post status updates periodically until the currently running disk firmware update completes on all of the requested nodes. Status includes display of the last **[numEvents]** disk firmware update related events, starting with the beginning of the most recent disk firmware update request; **[numEvents]** defaults to 5. Use the keyword **all** to specify no maximum.

fsckutil

Description

Service mode command to verify and repair the file system on one or more VPODs.

Syntax

```
fsckutil check|start ?[options]? ?[node options]?
```

Example

Check the status of fsck executed on the local VPOD:

```
fsckutil check
```

Check the status of fsck executed on all VPODs:

```
fsckutil check osds
```

Run fsck on local VPOD:

```
fsckutil start
```

Run fsck on all VPODs:

```
fsckutil start osds
```

Options

Option	Description
`check ?[node options]?`	Check the status of the executed **fsck** on specified nodes. Node options identify which Storage Node(s) or VPOD(s) to check. Values can be any of the following options: - **no node options specified**: the command will be run on the local Storage Node only. - **storageset [name]**: Include all VPODs in this StorageSet. - **host [name]**: Include the nodes with this name, IP address, or serial. - **ip**: Include the node with this IP address. - **ip-ip**: Include the nodes in this IP range. - r**ange ip-ip**: Include the nodes in this IP range. - **storageblades**: Include all nodes. - **osds**: Include all VPODs.

Option	Description
`start ?[node options]?`	Start **fsck** on specified VPODs. Node options identify which node(s) or VPOD(s) to run **fsck** on. Values can be any of the following options: - **no node options specified**: the command will be run on the local Storage Node only. - **storageset [name]**: Include all VPODs in this StorageSet. - **host [name]**: Include the nodes with this name, IP address, or serial. - **ip**: Include the node with this IP address. - **ip-ip**: Include the nodes in this IP range. - r**ange ip-ip**: Include the nodes in this IP range. - **storageblades**: Include all nodes. - **osds**: Include all VPODs. **Note:** This command can be run on a Storage Node to invoke **fsck** locally. To invoke **fsck** on multiple VPODs or a remote VPOD, this command must be run on a Director.

fsdb_osd

Description

Service mode interactive metadata editor for OSDfs filesystem.

The **-r** option keeps the editor in read-only mode. The read-only mode is useful for looking at the on-disk state of a mounted filesystem. It is also useful for just looking at metadata without the risk of accidentally changing it.

Without the **-r** option, the editor can only be run when the filesystem is unmounted. If the filesystem is mounted, this command will give a "device busy" message. Use of this tool requires extreme caution and should only be performed under the direction of VDURA Technical Support.

Launching **fsdb_osd** gives you an interactive **fsdb** prompt. Type **help** to see command help or **quit** to exit from the interactive tool.

Syntax

```
fsdb_osd ?-r? [device]
```

Example

Start the editor in read-only mode:

```
fsdb_osd -r /dev/ccd0c
```

Start the editor in read-write mode (first be sure the filesystem is unmounted):

```
fsdb_osd /dev/ccd0c
```

fsdb_osdv4

Description

Service mode interactive metadata editor for OSDv4 filesystem.

Launching **fsdb_osdv4** gives you an interactive **fsdb** prompt. Type **help** to see command help or **quit** to exit from the interactive tool.

Use of this tool requires extreme caution and should only be performed under the direction of VDURA Technical Support.

```
fsdb_osdv4 ?-D [datastore path]? ?-I [intentlog path]? ?-L [bdb log path]? ?-M [mds path]?
```

getdevinfo

Description

Service mode command to display or set local node configuration values.

- Run without options, **getdevinfo** displays local node configuration values.
- Specify the local configuration variable name to display that variable.
- Specify the local configuration variable name and value to set the variable. Manually setting a local configuration value may make the node inoperable, so this usage is not recommended without assistance from VDURA Technical Support.
- The **-q** option displays the configuration values without displaying the variable name(s).

Syntax

```
getdevinfo ?[options]?
```

Options

Option	Description
`-q`	Display the configuration values without displaying the variable name(s).

grep

Description

Searches files for a pattern.

Use the PanCLI **help grep** command for a full description of **grep** usage, options, and expressions.

guest

Description

Manages access to the read-only guest account.

Enables or disables access to the read-only guest account. Enabling access will prompt you to set a password. If no arguments are given, the status of the guest account is displayed.

Syntax

```
guest ?enable|disable?
```

gwctl

Description

Service mode command to display and set gateway tuning parameters.

Syntax

```
gwctl [subcommand] ?[arg1] [arg2] ...?
```

Options

Option	Description
`atime ?enable|disable?`	Display, enable, and disable **atime** support. By default, access time updates are not enabled. Enabling **atime** updates with this command enables updates for files accessed through the Directors via NFS or CIFS. To track **atime** for files accessed via DirectFlow, you must add the **osd-timestamp-bypass=0** mount option to each DirectFlow client's **/etc/fstab** entry. Access times can be retrieved via the **pan_atime** utility from the **panfs-apps** RPM.

Option	Description

`readahead`
`readahead on|off|medium|high`
`readahead custom [trigger-offset] [max-size] [window-size]`
	If no arguments are given, display the current **readahead** settings.
	- **on**: Enables readahead with the default policy, which is currently the same as **medium**.
	- **off**: Disables **readahead** entirely, and may improve performance for applications that primarily do random I/O, such as databases.
	- **medium**: Enables a moderate amount of read-ahead, and is intended to be suitable for a broad range of applications.
	- **high**: Provides more aggressive read-ahead and may improve performance for applications that primarily read whole files in order.
	The **custom** argument allows you to change the command algorithm parameters to specific settings.
	- **trigger-offset**: The minimum offset in a file that will trigger readahead, in bytes.
	- **max-size**: The largest readahead I/O that will be launched, in bytes.
	- **window-size**: determines the maximum number of concurrent readahead I/O operations of size **[max-size]** that the system will issue as part of the read-ahead mechanism.
	Custom **readahead** parameters should be set only in consultation with VDURA Technical Support. Setting incorrect custom parameters may significantly degrade both random and sequential read performance.

Option	Description
`stripemultiple ?read|write [value]?`	The read and write stripe multiple parameters control the maximum size I/O that the gateway client will launch for reads and writes, respectively. The parameters are in units of RAID stripes. Stripe multiple parameters should be changed only in consultation with VDURA Technical Support.

gzip

Description

Compresses a file.

For a complete description of usage and options, see the PanCLI **help gzip** command.

halt

Description

Interactive service mode command to halt the node.

Halting individual nodes can cause system-wide outages and should only be done under the direction of VDURA Technical Support.

Syntax

```
halt -p
```

head

Description

Displays the first part of a file.

This filter displays the first count lines or bytes of each of the specified files or of the standard input if no files are specified. If **count** is omitted, it defaults to 10.

Syntax

```
head [-n count | -c bytes] [file] ...
```

help

Description

Displays general or command/subcommand-specific help.

Type **?** after a command for help with a specific command. The format of the usage statements in this help system is as follows:

- Arguments enclosed in square brackets indicate a value that can vary.
- Arguments surrounded by question marks are optional.

Syntax

```
help [command] ?[subcommand]?
```

Example

View a listing of available commands type:

```
help
```

View help for the netconfig command:

```
help netconfig
```

View help for the netconfig command's edit subcommand:

```
help netconfig edit
```

Options

Option	Description
`(none)`	Display general help information, and a list of **PanCLI** commands.
`docs`	Display documentation for all commands.

Option	Description
`[command] ?[subcommand]?`	Display help for a specific command or subcommand.

history

Description

Saves and runs a history of commands.

Syntax

```
history ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`(none)`	Display command history.
`add 'command' [exec]`	Add an event to the history and optionally execute it.
`change 'command' [event]`	Change a given event (0,-1,-2...) to a new command.
`clear`	Clear the command history.
`event ?event?`	Look up the command for a given event name or number.
`info [count]`	Show current events in history and optionally limit to count.
`keep ?[count]?`	Display or change number of events that are kept in history. Default is 20.
`nextid`	Display the ID number of the next item to be added to the history.
`redo ?event?`	Repeats the last command or the specified history event.

ifconfig

Description

Displays and configures network interface parameters.

The **ifconfig** command called with no arguments can be used at any time to check the network interface status of the local host. Calling this command with additional arguments requires service mode and should only be done at the request of VDURA Technical Support, as some ifconfig requests can disable networking and cause system outages.

This command cannot be used on ActiveStor Ultra nodes; the **ip** command should be used instead.

Consult the UNIX manual page for details about the **ifconfig** command.

Syntax

```
ifconfig -C
ifconfig [-L] interface address_family [address [dest_address]] [parameters]
ifconfig interface create
ifconfig -a [-L] [-d] [-m] [-u] [address_family]
ifconfig -l [-d] [-u] [address_family]
ifconfig [-L] [-d] [-m] [-u]
```

install

Description

Use this command to perform the initial setup of the system or to clear the setup and start over. The installation wizard will wait for the background management server process before starting the guided setup procedure.

Setup prompts associated with unsupported/unlicensed features (CIFS and NFS) will be skipped.

Options

Option	Description
`network`	Set the initial system IP, netmask, and default router. Netmasks must be entered in the form 255.255.255.0.

ip

Description

The **ip** command called with no arguments can be used at any time to check the network interface status of the local host. Calling this command with additional arguments requires service mode and should only be done at the request of VDURA Technical Support, as some **ip** requests can disable networking and cause system outages. This command can only be used on ActiveStor Ultra nodes; use the **ifconfig** command on other nodes.

Syntax

```
ip addr
Service mode:
ip [ OPTIONS ] OBJECT { COMMAND | help }
ip [ -force ] -batch filename where OBJECT := { link | address | addrlabel | route | rule |
neigh | ntable | tunnel | tuntap | maddress | mroute | mrule | monitor | xfrm |
netns | l2tp | fou | macsec | tcp_metrics | token | netconf | ila | vrf | sr }
```

Options

Option	Description
`-V`	version
`-s`	statistics
`-d`	details
`-r`	resolve
`-h`	human readable
`-iec`	print human-readable rates in IEC units
`-f`	family { inet | inet6 | ipx | dnet | mpls | bridge | link }
`-4 | -6 | -I | -D | -B | -0 |`	shortcuts for -family inet, inet6, ipx, decnet, bridge, link

Option	Description
`-l`	loops { maximum-addr-flush-attempts }
`-br`	brief
`-o`	oneline
`-t`	timestamp
`-ts`	short
`-b`	batch [filename]
`-rc`	vbuf [size]
`-n`	netns name
`-a`	all
`-c`	color

ipmi_setup

Description

The **ipmi_setup** command is a service mode command to configure baseboard management controllers (BMCs) for remote IPMI access. This script can either be run manually for each BMC you wish to configure or in batch mode to configure multiple units at once. For all valid commands, the BMC is reset to factory defaults before the desired configuration is applied.

> **Note:** IPMI passwords must meet the following criteria: Length of 8 to 20 characters; Cannot be reverse of username; Includes characters from at least three of these character classes: A-Z, a-z, 0-9, and special characters

Syntax

```
ipmi_setup [<ip> <netmask> <gateway> | batch <cfg> | dhcp | factory]
```

Example

Static IP, command-line

Configure the BMC to use IP address 1.2.3.4, netmask 255.255.255.0, and gateway (router) 1.2.3.255. The script will also prompt for a password to apply to the BMC.

```
ipmi_setup 1.2.3.4 255.255.255.0 1.2.3.255
```

DHCP, command-line

Configure the BMC to obtain IP address, netmask, and gateway (router) via DHCP. This configuration requires an appropriately-configured DHCP server on the broadcast domain. As a final step, the script will prompt for a password to apply to the BMC.

```
ipmi_setup dhcp
```

Factory reset, command-line

Reset the BMC to its factory default configuration.

```
ipmi_setup factory
```

Batch

Configure three nodes using three different modes:

- Node 1: Configure the BMC for static IP using IP address 10.11.12.13 as the node-id.
- Node 2: Configure the BMC for DHCP using data serial number 4a7a2042a1ae36 as the node-id.
- Node 3: Reset the BMC to factory defaults using hardware serial number 42af0001302001001 as the node-id.

Create a batch file containing one line for each node; this example will use /tmp/ipmi_setup.cfg. Note that any type of node-id may be used, and that the desired BMC password is included for each node if the command requires it.

```
10.11.12.13 1.2.3.4 255.255.255.0 1.2.3.255 password1
4a7a2042a1ae36 dhcp password2
42af0001302001001 factory
```

Use scp to copy /tmp/ipmi_setup.cfg to all three nodes:

```
diag_exec " scp root@[director IP address]:/tmp/ipmi_setup.cfg /tmp" all
```

Note that Director IP address above is the address of the node where you created the config file:

```
diag_exec " scp root@192.168.52.10:/tmp/ipmi_setup.cfg /tmp" all
```

After all nodes have the configuration file, run ipmi_setup on all nodes to be configured. To run the command in the background simultaneously on the nodes, use an ampersand (&):

```
diag_exec -timeout 600 "ipmi_setup batch /tmp/ipmi_setup.cfg &" all
```

When run, ipmi_setup will read the batch file, find the specified configuration for each node, and apply it.

> **Note:** A node may go into Non-reconciled state while being configured.

Once finished, you should verify that all nodes have been configured successfully. If any node(s) failed, either re-run ipmi_setup with just the failed node(s) or manually run ipmi_setup on the affected node(s).

Options

Modes: The **ipmi_setup** script has three configuration modes for each BMC.

Option	Description
Static IP	Use **ipmi_setup** followed by the BMC IP address, netmask, and gateway IP address. The script will then prompt for a new BMC password.
DHCP	Use **ipmi_setup dhcp** to automatically configure the BMC for DHCP. The script will then prompt for a new BMC password.
Factory reset	Use **ipmi_setup factory** to reset the BMC to its factory default configuration, including the BMC password.

Argument Styles: The **ipmi_setup** script can be run in two different ways.

Option	Description
Command-line	Enter arguments via the command line. If needed, the script asks for a password to assign to the BMC.

Option	Description

Batch	Enter the command-line argument **batch** followed by the path to a batch file you have previously created and copied to all nodes you are modifying. This file must contain one configuration line for each node on which the script is to be run. Each line must start with a **node-id** that designates what node the line applies to. The **node-id** can be either the node's IP address, data serial number, or hardware serial number of the node. After the **node-id**, the line contains a list of arguments equivalent to the command-line version. For options that require an IPMI password, the password is included in the configuration file in plain text. Since it contains cleartext passwords, the batch file will be deleted after use.
	To configure a static IPMI SOL IP address in batch mode, each line must contain the following:
	`[Node Identifier] [IPMI IP address] [IPMI netmask] [IPMI gateway IP address] [IPMI password]`
	To configure IPMI SOL to obtain an IP address via DHCP in batch mode, each line must contain the following:
	`[Node Identifier] dhcp [IPMI password]`
	To reset the IPMI SOL configuration to the factory default in batch mode, each line must contain the following:
	`[Node Identifier] factory`

ipmitool

Description

Service mode command to manage Intelligent Platform Management Interface (IPMI) functions of individual nodes or of the realm. These functions include reporting field replaceable unit (FRU) information and sensor readings, reporting and changing LAN and serial-over-LAN configuration, and remote power control.

Syntax

```
ipmitool ?[options]? command subcommand ?[suboptions]?
```

Options

Option	Description
`channel`	Configure Management Controller channels.
`chassis`	Get chassis status and set power state.
`lan`	Configure LAN Channels.
`mc`	Management Controller status and global enables.
`power`	Shortcut to chassis power commands.
`sol`	Configure and connect IPMIv2.0 Serial-over-LAN.
`user`	Configure Management Controller users.

iplist

Description

Gets the list of IP addresses for the system.

Syntax

```
iplist ?[option] [option] ...?
```

Options

Option	Description
`(none)`	Return a list of all IP ranges currently in use by the system.
`containers`	See the IP addresses of all VDURA Data Platforms in the system.
`directors`	See the IP addresses of all Directors in the system.
`enclosure [name]`	Return only those IP ranges in use by nodes in the specified enclosure. Multiple enclosures can be specified by using this option more than once.
`fileowner [PanFS file]|[Object ID]`	See the IP addresses of Storage Nodes that store a portion of the given PanFS file or Object ID as well as the primary and backup Directors for the file or object's volume. The path specified in **[PanFS file]** must not include **/panfs/<realm-name>/** at the start of the path. Owners of multiple files can be specified by using this option more than once.
`host [name]`	See the IP address of the node with this name, IP address, or serial. Multiple hosts can be specified by using this option more than once.

Option	Description
`range [ip]-[ip]`	Return only those IP ranges in use in the specified range of IP addresses. Multiple hosts can be specified by using this option more than once.
`repset`	Return only those IP ranges in use by Directors in the replication set (repset).
`storage`	See the IP addresses of all Storage Nodes in the system.
`storageset [name]`	See the IP addresses of all nodes in a StorageSet. Multiple StorageSets can be specified by using this option more than once.
`volume`	See the IP address of the volume with the specified name.
`volumeowner [volume name]`	Get the IP address of the Director that owns the given volume. If Volume **Services Failover** is enabled, also get the IP address of the backup Director for the given volume. Multiple volume owners can be specified by using this option more than once.

kill

Description

Service mode command to terminate or signal a running process. The kill utility sends a signal to the processes specified by the **pid** operand(s).

Only the superuser may send signals to other users' processes.

Syntax

```
kill [pid]
kill [-s signal_name] pid ...
kill -l [exit_status]
kill -signal_name pid ...
kill -signal_number pid ...
```

Options

Option	Description
`-s signal_name pid`	A symbolic signal name specifying the signal to be sent instead of the default TERM.
`-l [exit_status]`	If no operand is given, list the signal names; otherwise, write the signal name corresponding to **exit_status**.
`-signal_name pid`	A symbolic signal name specifying the signal to be sent instead of the default TERM.
`-signal_number pid`	A non-negative decimal integer, specifying the signal to be sent instead of the default TERM.

The **-1 pid** has special meaning: If superuser, broadcast the signal to all processes; otherwise broadcast to all processes belonging to the user.

Some of the more commonly used signals:

1	**HUP (hang up)**
2	INT (interrupt)
3	QUIT (quit)
6	ABRT (abort)
9	KILL (non-catchable, non-ignorable kill)
14	ALRM (alarm clock)
15	TERM (software termination signal)

killall

Description

Service mode command to terminate a running process by name.

This command kills processes selected by name, as opposed to the selection by pid as done by **kill**. By default, it will send a TERM signal to all processes with a real UID identical to the caller of **killall** that matches the name procname. The superuser is allowed to kill any process.

Syntax

```
killall [procname]
```

Options

Option	Description
`-d | -v`	Be more verbose about what will be done. For a single **-d** option, either a list of the processes that will be sent the signal or a message indicating that no matching processes have been found will be printed.
`-h | -? -help`	Show usage message.
`-l`	List the names of the available signals and exit, like in **kill**.
`-m`	Match the argument **procname** as a case insensitive regular expression against the names of processes found. This can be dangerous, as a single dot will match any process running under the real UID of the caller.
`-s`	Show only what would be done, but do not send any signal.

Option	Description
`-SIGNAL`	Send a different signal instead of the default TERM. The signal may be specified either as a name (with or without a leading SIG) or numerically.
`-u user`	Limit potentially matching processes to those belonging to the specified user.
`-t tty`	Limit potentially matching processes to those running on the specified tty.
`-c procname`	When used with the **-u** or **-t** flags, limit potentially matching processes to those matching the specified progname.

last

Description

Displays last logins of users and ttys.

Syntax

```
last [-d [[CC]YY][MMDD]hhmm[.SS]] [-f file] [-h host] [-n] [-s] [-t tty] [-w] [user ...] [-y]
```

Options

Option	Description
`-d date`	Specify the snapshot date and time. All users logged in at the snapshot date and time will be reported. This may be used with the **-f** option to derive the results from stored **wtmp** files. When this argument is provided, all other options except for **-f** and **-n** are ignored. The argument should be in the form **[[CC]YY] [MMDD]hhmm[.SS]** where each pair of letters represents the following: - **CC**: The first two digits of the year (i.e., the century - 19, 20, etc.). - **YY**: The second two digits of the year. If **YY** is specified, but **CC** is not, a value for **YY** between 69 and 99 results in a **CC** value of 19. Otherwise, a **CC** value of 20 is used. - **MM**: Month of the year, from 1 to 12. - **DD**: Day of the month, from 1 to 31. - **hh**: Hour of the day, from 0 to 23. - **mm**: Minute of the hour, from 0 to 59. - **SS**: Second of the minute, from 0 to 61.
`-f file`	Last reads the file entered instead of the default **/var/log/wtmp**.
`-h host`	Host names may be names or internet numbers.

Option	Description
`-n`	Limit the report to *n* lines.
`-s`	Report the duration of the login session in seconds, instead of the default days, hours, and minutes.
`-t tty`	Specify the tty. Names may be given fully or abbreviated, for example, **last -t 03** is equivalent to **last -t tty03**.
`-w`	Widen the duration field to show seconds, as well as the default days, hours, and minutes.
`-y`	Report the year in the session start time.

ldap

Description

Manages VDURA LDAP client configuration.

Syntax

```
ldap ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`account ?[userid]?`	If your LDAP server does not allow anonymous binds, you must enter the username or distinguished name (DN) to use for binding. If you enter a username instead of a DN, the system will construct the bind DN as **cn=[username],[baseDN]**. This user must have permission to search the directory and retrieve user, group, and netgroup entries. Use “” to clear the value.
`certificate ?[file]|""?`	A CA certificate for the CA that issued the LDAP server’s certificate is required to use secure communication (TLS). Certificates must follow the standard PEM format. If no file is given, display the contents of the certificate that is saved as part of the LDAP configuration. If a file is specified, load that file. The filename can be a path in PanFS (relative to the system root directory) or a path on the local filesystem of the node that hosts PanActive Manager. File paths must be absolute to be found on the local filesystem. The contents of the file replace the current certificate, if any. Use “” to clear the certificate.

Option	Description
`debug`	This is the debug level used by name resolution services. To set this debug level, use the **nameresorder** command. Any debug level above 0 will impact performance.
`domain ?[domain-name]?`	The base distinguished name for all LDAP queries will be derived by the system from this domain name. The domain name should be a fully qualified DNS name. For example, if you specify **eng.company.com**, the base DN will be set to **dc=eng,dc=company,dc=com**. If you need to specify a base DN that does not match this format, you can also specify the actual base DN to use. Use the following command to clear the domain: `ldap domain ""`
`enabled ?yes|no?`	The LDAP feature allows the system to perform username to UID mappings, group name to GID mappings, and netgroup lookups via LDAP. To enable LDAP, the domain name and server must be valid. If your LDAP server disallows anonymous binds, you must also specify the account name and password. Note that you must also add LDAP to the name resolution search ordering before the system will perform LDAP resolution. Use the **nameresorder** command to manage name resolution search order.
`password`	Interactive command to change the LDAP bind password.

Option	Description
`queryparams ?[attribute] [value]?`	If no attributes and values are given, display the current settings. Specify the search attributes to be used in LDAP queries. The valid search attributes are: - **usrobjectclass**: Object class for user records - **usrlogin**: User login name - **usrpassword**: User password - **usrid**: Numeric user ID - **usrgid**: User's numeric group ID - **usrhome**: User's home directory - **usrshell**: User's login shell - **usrgecos**: GECOS name for a user - **grpobjectclass**: Object class for group records - **grpname**: Group name - **grpgid**: Group numeric ID - **grpmembers**: Group members - **netgrpobjectclass**: Object class for netgroup records - **ngrpname**: Netgroup name - **ngrpmembers**: Netgroup members - **ngrpgroups**: Netgroup subgroup To clear a value, type "".
`secure ?yes|no?`	When mandatory security is disabled, the system will attempt to establish a secure TLS session with the LDAP server, but will fall back to an insecure connection if TLS negotiation fails. When mandatory security is enabled, the system will refuse to connect unless it can establish a TLS session and the server presents a valid certificate. When establishing the TLS session, the system will first try StartTLS and fall back to LDAPS if StartTLS is not supported. You must also enter a CA certificate via the **ldap certificate** command to use TLS.

Option	Description
`server ?[ip|hostname]?`	Use this command to get or set the LDAP server. Either IP address or hostname can be given. Use the following command to clear the server: `ldap server ""`
`timeout [seconds]`	Display or set the LDAP timeout value in seconds. Allowed timeout value is 0 to 15 seconds. A 0 value means no timeout.

ldapsearch

Description

This command opens a connection to an LDAP server, binds, and performs a search using specified parameters.

Syntax

```
ldapsearch [options] [filter [attributes...]]
```

Arguments

Argument	Description
`filter`	RFC-2254 compliant LDAP search filter.
`attributes`	Whitespace-separated list of attribute descriptions. An attribute description can be **1.1** (no attributes), ***** (all user attributes), or **+** (all operational attributes).

Search Options

Option	Description
`-a deref`	One of **never** (default), **always**, **search**, or **find**.
`-A`	Retrieve attribute names only (no values).
`-b basedn`	Base DN for the search.

Option	Description
`-E [!][ext][=[extparam]]`	Search extensions (**!** indicates criticality). `[!]domainScope` (domain scope) `[!]mv=[filter]` (matched values filter) `[!]pr=[size][/prompt|noprompt]` (paged results/prompt) `[!]subentries[=true|false]` (subentries) `[!]sync=ro[/[cookie]]` (LDAP Sync refreshOnly) `rp[/[cookie]][/[slimit]]` (LDAP Sync refreshAndPersist)
`-F prefix`	URL prefix for files (default: `file:///var/tmp//`).
`-l limit`	Time limit (in seconds, or **none** or **max**) for the search.
`-L`	Print responses in LDIFv1 format.
`-LL`	Print responses in LDIF format without comments.
`-LLL`	Print responses in LDIF format without comments and version.
`-s scope`	One of **base**, **one**, **sub**, or **children** (search scope).
`-S attr`	Sort the results by the attribute **attr**.
`-t`	Write binary values to files in the temporary directory.
`-tt`	Write all values to files in the temporary directory.
`-T path`	Write files to the directory specified by **path** (default: `/var/tmp/`).
`-u`	Include User Friendly entry names in the output.

Option	Description
`-z limit`	Size limit (in entries, or **none** or **max**) for the search.

Common Options

Option	Description
`-c`	Continuous operation mode (do not stop on errors).
`-C`	Chase referrals (anonymously).
`-d level`	Set the LDAP debugging level to **level**.
`-D binddn`	Bind DN.

`-e [!][ext][=[extparam]]`	General extensions (**!** indicates criticality).
	`[!]assert=[filter]` (an RFC 2254 filter)
	`[!]authzid=[authzid]` ("dn:[dn]" or "u:[user]")
	`[!]manageDSAit`
	`[!]noop`
	`ppolicy`
	`[!]postread[=[attrs]]` (a comma-separated attribute list)
	`[!]preread[=[attrs]]` (a comma-separated attribute list)
	`abandon`, `cancel` (SIGINT sends abandon/cancel; not really controls)

`-f file`	Read operations from **file**.
`-h host`	LDAP server.
`-H URI`	LDAP Uniform Resource Identifier(s).
`-I`	Use SASL Interactive mode.
`-k`	Use Kerberos authentication.
`-K`	Like **-k**, but do only step 1 of the Kerberos bind.

Option	Description
`-M`	Enable the Manage DSA IT control (**-MM** to make it critical).
`-n`	Show what would be done but do not actually do it.
`-O props`	SASL security properties.
`-p port`	Port on the LDAP server.
`-P version`	Protocol version (default: 3).
`-Q`	Use SASL Quiet mode.
`-R realm`	SASL realm.
`-U authcid`	SASL authentication identity.
`-v`	Run in verbose mode (diagnostics to standard output).
`-V`	Print version information (**-VV** only).
`-w passwd`	Bind password (for simple authentication).
`-W`	Prompt for the bind password.
`-x`	Simple authentication.
`-X authzid`	SASL authorization identity (“dn:[dn]” or “u:[user]”).
`-y file`	Read the password from **file**.
`-Y mech`	SASL mechanism.
`-Z`	Start TLS request (**-ZZ** to require a successful response).

See Also

ldap, **ldapwhoami**

ldapwhoami

Description

This command opens a connection to an LDAP server, binds, and issues a “Who am I?” operation to request the user’s authorization identity (authzid).

Syntax

```
ldapwhoami [options]
```

Options

Option	Description
`-d level`	Set the LDAP debugging level to **level**.
`-D binddn`	Bind DN.

`-e [!][ext][=[extparam]]`	General extensions (**!** indicates criticality).
	`[!]assert=[filter]` (an RFC 2254 filter)
	`[!]authzid=[authzid]` (“dn:[dn]” or “u:[user]”)
	`[!]manageDSAit`
	`[!]noop`
	`ppolicy`
	`[!]postread[=[attrs]]` (a comma-separated attribute list)
	`[!]preread[=[attrs]]` (a comma-separated attribute list)
	`abandon`, `cancel` (SIGINT sends abandon/cancel; not really controls)

`-f file`	Read operations from **file**.
`-h host`	LDAP server.
`-H URI`	LDAP Uniform Resource Identifier(s).
`-I`	Use SASL Interactive mode.

Option	Description
`-n`	Show what would be done but do not actually do it.
`-O props`	SASL security properties.
`-p port`	Port on the LDAP server.
`-Q`	Use SASL Quiet mode.
`-R realm`	SASL realm.
`-U authcid`	SASL authentication identity.
`-v`	Run in verbose mode (diagnostics to standard output).
`-V`	Print version information (**-VV** only).
`-w passwd`	Bind password (for simple authentication).
`-W`	Prompt for the bind password.
`-x`	Simple authentication.
`-X authzid`	SASL authorization identity (“dn:[dn]” or “u:[user]”).
`-y file`	Read the password from **file**.
`-Y mech`	SASL mechanism.
`-Z`	Start TLS request (**-ZZ** to require a successful response).

See Also

ldap, **ldapsearch**

ls

Description

Lists directory contents.

- For each operand that names a file of a type other than directory, **ls** displays its name as well as any requested associated information.
- For each operand that names a file of type directory, ls displays the names of files contained within that directory, as well as any requested associated information.
- If no operands are given, the contents of the current directory are displayed.
- If more than one operand is given, non-directory operands are displayed first; directory and non-directory operands are sorted separately and in lexicographical order.

Syntax

```
ls [-ABCFGHLPRTWabcdfghiklnoqrstuw1] [file ...]
```

Options

Option	Description
`-A`	List all entries except for `.` and `..`. Always set for the superuser.
`-B`	Force printing of non-printable characters (as defined by **ctype(3)** and current locale settings) in file names as `\xxx`, where **xxx** is the numeric value of the character in octal.
`-C`	Force multi-column output. This is the default when output is to a terminal.
`-F`	Display a slash (`/`) immediately after each pathname that is a directory, an asterisk (`*`) after each that is executable, an at sign (`@`) after each symbolic link, an equals sign (`=`) after each socket, a percent sign (`%`) after each whiteout, and a vertical bar (`|`) after each that is a FIFO.

Option	Description
`-G`	Enable colorized output.
`-H`	Follow symbolic links on the command line. This option is assumed if none of the **-F**, **-d**, or **-l** options are specified.
`-L`	If the argument is a symbolic link, list the file or directory the link references rather than the link itself. This option cancels the **-P** option.
`-P`	If the argument is a symbolic link, list the link itself rather than the object the link references. This option cancels the **-H** and **-L** options.
`-R`	Recursively list subdirectories encountered.
`-T`	When used with the **-l** (lowercase letter “ell”) option, display complete time information for the file, including month, day, hour, minute, second, and year.
`-W`	Display whiteouts when scanning directories.
`-a`	Include directory entries whose names begin with a dot (`.`).
`-b`	As **-B**, but use C escape codes whenever possible.
`-c`	Use the time when the file status was last changed for sorting or printing.
`-d`	List directories as plain files (not searched recursively).
`-f`	Do not sort the output.
`-g`	This option is deprecated and is only available for compatibility with 4.3BSD. It was used to display the group name in the long (**-l**) format output.

Option	Description
`-h`	When used with the **-l** option, use unit suffixes (Byte, Kilobyte, Megabyte, Gigabyte, Terabyte, and Petabyte) to reduce the number of digits to three or less, using base 2 for sizes.
`-i`	For each file, print the file serial number (inode number).
`-k`	If the **-s** option is specified, print the file size allocation in kilobytes, not blocks.
`-l`	(The lowercase letter “ell”.) List in long format. If the output is to a terminal, a total sum for all the file sizes is output on a line before the long listing.
`-n`	Display user and group IDs numerically rather than converting them to a user or group name in a long (**-l**) output.
`-o`	Include the file flags in a long (**-l**) output.
`-q`	Force printing of non-graphic characters in file names as the character `?`. This is the default when output is to a terminal.
`-r`	Reverse the order of the sort to get reverse lexicographical order or the oldest entries first.
`-s`	Display the number of file system blocks actually used by each file, in units of 512 bytes, where partial units are rounded up to the next integer value. If the output is to a terminal, a total sum for all the file sizes is output on a line before the listing.
`-t`	Sort by time modified (most recently modified first) before sorting the operands by lexicographical order.

Option	Description
`-u`	Use the time of last access, instead of the last modification of the file, for sorting (**-t**) or printing (**-l**).
`-w`	Force raw printing of non-printable characters. This is the default when output is not to a terminal.
`-1`	(The numeric digit “one”.) Force output to be one entry per line. This is the default when output is not to a terminal.

The **-1**, **-C**, and **-l** options all override each other. The last one specified determines the format used.

The **-c** and **-u** options override each other. The last one specified determines the file time used.

The **-B**, **-b**, **-w**, and **-q** options all override each other. The last one specified determines the format used for non-printable characters.

The **-H**, **-L**, and **-P** options all override each other (either partially or fully). They are applied in the order specified.

By default, **ls** lists one entry per line to standard output. The exceptions are to terminals or when the **-C** option is specified.

File information is displayed with one or more blanks separating the information associated with the **-i**, **-s**, and **-l** options.

The Long Format

If the **-l** option is given, the following information is displayed for each file.

- File mode
- Number of links
- Owner name
- Group name
- Number of bytes in the file
- Abbreviated month the file was last modified
- Day of the month the file was last modified
- Hour the file was last modified
- Minute the file was last modified
- Pathname

In addition, for each directory whose contents are displayed, the total number of 512-byte blocks used by the files in the directory is displayed on a line by itself immediately before the information for the files in the directory.

If the modification time of the file is more than 6 months in the past or future, the year of the last modification is displayed in place of the hour and minute fields.

If the owner or group names are not a known user or group name, or the **-n** option is given, the numeric IDs are displayed.

If the file is a character special or block special file, the major and minor device numbers for the file are displayed in the size field. If the file is a symbolic link, the pathname of the linked-to file is preceded by `->`.

The file mode printed under the **-l** option consists of the entry type, owner permissions, and group permissions. The entry type character describes the type of file, as follows.

Character	Meaning
`b`	Block special file.
`c`	Character special file.
`d`	Directory.
`l`	Symbolic link.
`s`	Socket link.
`p`	FIFO.
`-`	Regular file.

The next three fields are three characters each, for owner permissions, group permissions, and other permissions. Each field has three character positions.

1. If **r**, the file is readable. If **-**, it is not readable.
2. If **w**, the file is writable. If **-**, it is not writable.
3. The first of the following that applies.

Character	Meaning
`S`	If in the owner permissions, the file is not executable and set-user-ID mode is set. If in the group permissions, the file is not executable and set-group-ID mode is set.
`s`	If in the owner permissions, the file is executable and set-user-ID mode is set. If in the group permissions, the file is executable and set-group-ID mode is set.
`x`	The file is executable or the directory is searchable.
`-`	The file is neither readable, writable, executable, nor set-user-ID nor set-group-ID mode, nor sticky.

The next two characters apply only to the third character in the last group (other permissions).

Character	Meaning
`T`	The sticky bit is set (mode 1000), but not execute or search permission.
`t`	The sticky bit is set (mode 1000), and is searchable or executable.

lw-lsa

Description

Manages the Likewise security and authentication service. This command must be run in service mode.

Active Directory (AD) is not supported in this release, so the **ad-cache**, **ad-get-machine**, and **ad-reload-join-state** options do not apply.

Syntax

```
lw-lsa [command] [args...]
```

Options

Option	Description
`ad-cache --delete-all ?--domain domain?`	Delete everything from the cache.
`ad-cache --delete-user ?--domain domain? [--name user-name | --uid uid]`	Delete a user from the cache with the given name or UID.
`ad-cache --delete-group ?--domain domain? [--name group-name | --gid gid]`	Delete a group from the cache with the given name or GID.
`ad-cache --enum-users ?--domain domain? ?[--batchsize 1...1000]?`	Enumerate the users in the cache returning *[batchsize]* users at a time.
`ad-cache --enum-groups ?--domain domain? ?[--batchsize 1...1000]?`	Enumerate the groups in the cache returning *[batchsize]* groups at a time.
`ad-get-machine account password ?[domain]?`	Get machine account information for the specified domain.
`ad-reload-join-state`	Instructs lsass provider to reload the domain join information.
`get-status ?--all? ?[--provider provider]? ?domain?`	Get status for the specified provider for the specified (or all) domains.

Option	Description
`find-objects [--user | --group] ?[--by-[alias|dn|name|nt4|sid|upn|unix-id] key]? ?--nss? ?--provider [name]?`	Find objects of the specified object.
`query-member-of [--user | --group] ?[--by-[alias|dn|name|nt4|sid|upn|unix-id] key]? ?--nss? ?--provider [name]?`	Find which groups the objects specified by keys are members of.

lwnet

Description

Manages the CIFS/SMB server. This command must be run in service mode.

Syntax

```
lwnet [command] [args...]
```

Options

Option	Description
`file [id [close]]`	View or close the specified file ID.
`session [\\computername [delete | del]]`	View or close open sessions from [\ *computername]* .
`share enum-info [enum-options]`	Get information about all shares on the server.
`share get-info name [info-options]`	Get information about a single share name.
`share set-info name set-info-options`	Set various share options. - –c**omment comment**: Set share comment - –**[enable|disable]-abe**: Enable/disable access based enumeration - –**allow [nt4 name] [read-only | read-write]**: Allow user/group the specified access - –**deny [nt4 name]**: Deny user/group access to the share - –c**lear-allow**: Clear allowed list - –**clear-deny**: Clear denied list

lwregshell

Description

Views and manages the Likewise registry. This command must be run in service mode.

Syntax

```
lwregshell [command] [args...]
```

Options

Option	Description
`export [filename]`	Export the complete registry to filename.

lwsm

Description

Controls Likewise services. This command must be run in service mode.

Syntax

```
lwsm [command] [args...]
```

Options

Option	Description
`list`	List all services and their current statuses.
`info [service]`	Get information about a service.
`stop [service]`	Stop a service after stopping all its transitive dependencies.
`start [service]`	Start a service after starting all its transitive dependencies.
`refresh [service]`	Service (or lwsmd) reloads its configuration.
`get-log [service]`	Get current logging state for the service.
`set-log-level [service] [facility] [level]`	Set log level for a given service.
`set-log-target [service] [facility] [type] [target]`	Set log target (syslogd or file) for service.
`tap-log [service] [facility] [level]`	Temporarily change the log level for service and redirect it to stdout.

md5

Description

Calculates an MD5 message-digest fingerprint (checksum) for a file. The utility takes as input a message of arbitrary length and produces as output a “fingerprint” or “message digest” of the input. The hexadecimal checksum of each file listed on the command line is printed. Also see the **sha1** command.

Syntax

```
md5 file [...]
```

mkdir

Description

The **mkdir** utility creates the directories named as operands, in the order specified, using **mode rwxrwxrwx (0777)** as modified by the current umask. The user must have write permission in the parent directory.

Syntax

```
mkdir [-pv] [-m mode] directory_name ...
```

Options

Option	Description
`-m`	Set the file permission bits of the final created directory to the specified mode. The mode argument can be in any of the formats specified to the **chmod(1)** command. If a symbolic mode is specified, the operation characters **+** and **-** are interpreted relative to an initial mode of **a=rwx**.
`-p`	Create intermediate directories as required. If this option is not specified, the full path prefix of each operand must already exist. On the other hand, with this option specified, no error will be reported if a directory given as an operand already exists. Intermediate directories are created with permission bits of **rwxrwxrwx (0777)** as modified by the current umask, plus write and search permission for the owner.
`-v`	Be verbose when creating directories, listing them as they are created.

more

Description

Interactively views a text file.

More is a program that allows backward movement in the file as well as forward movement.

For information about the program and a list/description of its commands and options, see the PanCLI **help more** command.

nameresorder

Description

Manages priority ordering of username resolution services.

Active Directory (AD) is not supported in this release.

Syntax

```
nameresorder ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`debug ?[level]?`	Manage the name resolution services debug level. This debug level affects name resolution services. Note that a non-zero value can impact system performance. The level can be any number between 0 and 10, inclusive. A value of 0 allows no debugging, while 10 is the most verbose debug level. Setting the debug level is only available in service mode and should only be done under the direction of VDURA Technical Support.

`netgroup`
`netgroup [ldap|ad|nis] ?[ldap|ad|nis]? ?[ldap|ad|nis] ?`
`netgroup clear`
	Specify the order of name resolution services to use for UNIX netgroup name resolution. Supported services are LDAP, AD, and NIS.
	If the first argument is **clear**, then the order list will be cleared, and no name resolution services will be used for UNIX netgroup name resolution. Including a service in this list will only be effective if that service is enabled. Use the **nis**, **ad**, and **ldap** commands to enable the various services.

Option	Description

`usergroup`
`usergroup [ldap|ad|nis] ?[ldap|ad|nis]? ?[ldap|ad|nis] ?`
`usergroup clear`
	Specify the order of name resolution services to use for UNIX user/group name resolution. Supported services are LDAP, AD, and NIS.
	If the first argument is **clear**, then the order list will be cleared, and no name resolution services will be used for UNIX user/group name resolution.
	Including a service in this list will only be effective if that service is enabled. Use the **nis**, **ad**, and **ldap** commands to enable the various services.

netconfig

Description

Displays and sets system network related values.

Syntax

```
netconfig ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
none	If no options are given, display the system IP, DNS name, netmask, default routers, and maximum ethernet framesize.

Option	Description
`admin-nic all [enable|disable]`	This command applies only to ActiveStor Directors. The administrative interface is used to access a Director node to do the initial out-of-the-box (OOBE) realm configuration or to access the node if it otherwise becomes inaccessible over the network. The **netconfig admin-nic all** commands are realm-wide. The **enable** command enables all Director node administrative interfaces, while the **disable** command disables all Director node administrative interfaces. Enabling the administrative interface takes the interface out of the disabled state, assigns it a static IP address of 169.254.12.34/16, and puts the interface into the Administratively UP state. Disabling the administrative interface completely erases its configuration and puts the interface into the Administratively DOWN state. To avoid the potential for duplicate IP address conflicts, physically disconnect the interface from the network after initial (OOBE) realm configuration is complete.
`alias`	Manage IP aliases and NFS failover. IP aliases can be used for various purposes, most notably to enable NFS failover. ActiveStor Directors can support both port bonding and secondary aliases. For this configuration, the 10 GE NICs must be connected to a different VLAN than the NET card uplinks. To do this, put the two 1 GE NICs into lagg0 and use it as the primary interface, then put the two 10 GE NICs into lagg1 and use it as the secondary interface.

Option	Description
`alias add (primary|secondary) [start] [end] [netmask] ?[subnet]?`	Add a new IP alias range to the primary or secondary interface. If your shelves span multiple subnets, when you add an IP alias, you will need to enter the subnet in which to apply the alias. Output of the **netconfig subnet** command lists the subnets in the form required by this command.
`alias delete (primary|secondary) [start] [end] [netmask]`	Delete an IP alias range on the primary or secondary interface.
`alias edit (primary|secondary) [old-start] [new-start] [old-end] [new-end] [old-netmask] [new-netmask]`	Edit an IP alias range on the primary or secondary interface.
`alias split (primary|secondary) [old-start] [new-end] [new-start] [old-end] [old-netmask] [new-netmask]`	Split an IP alias range on the primary or secondary interface.
`alias failover ?enable|disable?`	Enable or disable IP alias failover. For additional IP alias troubleshooting, it may become necessary to freeze the IP alias assignment. The **netconfig alias failover disable** command allows this service-mode action. The consequence of taking this action is that no new aliases will be assigned, and no aliases will failover to other nodes in the event of a node failure. This command should only be used under the direction of VDURA Technical Support. To cancel this special mode, use the **netconfig alias failover enable** command. To see whether IP alias failover has been disabled, use the command **netconfig alias failover**. To disable IP alias assignment on specific Directors, use the **syscontrol node [id] alias** command.

Option	Description
`edit ?[primary|secondary]?`	This interactive command allows you to edit the primary or secondary network settings. If neither the primary nor the secondary option is specified, the default is primary. Use the **primary** command to edit the system name, IP address, netmask, default routers, IP address ranges, and primary network maximum ethernet framesize. Netmasks must be entered in the form **255.255.255.0**. Changing the primary settings forces a system reboot or shutdown to apply the changes. Use the **secondary** command to edit the secondary network maximum ethernet framesize. Changing the secondary settings does NOT require a system reboot or shutdown to apply the changes.
`framesize ?update?`	If the word update is given after the **framesize** subcommand (`framesize update`), the interactive netconfig **edit networking changes wizard** is invoked to prompt the admin for new networking configuration values. Completion of this wizard results in an automatic system reboot. The maximum Ethernet framesize is the same as the Maximum Transfer Unit (MTU). While the default value of 1500 can be used for compatibility with all Ethernet networks, VDURA recommends using 9000 (jumbo frame) after verification that your network environment supports it (i.e., all clients are using the same MTU of 9000 and ports on switches/routers are set to use MTU larger than 9000) to achieve maximum performance. If no additional value is given after the **framesize** subcommand, just display the current maximum ethernet framesize.

Option	Description
`ip`	If the system is configured with automatic Director failover, then the IP address of each replication node will be listed.
`ipranges`	Display the current IP address ranges.
`ipranges add [start] ?[end]?`	Add an IP address range to the list. If **[end]** is not specified, the **[start]** and **[end]** addresses are the same.
`ipranges assigned ?[start]?`	Display the ranges of IP addresses currently assigned in the system. If a **[start]** IP is specified, display the IP address, MAC address, and lease expiration time for each of the nodes whose IP address falls in the specified range.
`ipranges delete [start]`	Remove an IP address range from the list with the given **[start]** address.
`ipranges edit [oldstart] [oldend] [newstart] [newend]`	Edit the range start and end addresses.
`ipranges ipchange`	An interactive service mode command will be started to change the IP address of a node that has already joined the system. Note that this change will force a reboot of that node. The command cannot be used to modify the IP address of any Director that is in the replication set (repset).
`ipranges split [oldstart] [newend] [newstart] [oldend]`	Split the range start and end addresses into two new ranges.

Option	Description
`leases ?release [ip]?`	Display all MAC address and DHCP lease information. The lease information includes the IP, MAC address, lease status, and expiration for each lease in the system. Release of an IP lease can only be done in service mode. This should only be performed under the direction of VDURA Technical Support.
`name`	Display the system DNS name.
`netgroup ?list?`	If the **list** option or no option is given, prints the contents of the netgroup file.
`netgroup copy [file]`	Copy the currently deployed netgroup file to the given PanFS file path. Files must be specified using a full path starting with **/volume-name**.
`netgroup deploy [file]`	Deploy and save the netgroup file that exists in the given PanFS file location. Files must be specified using a full path starting with **/volume-name**. Netgroup files must be 53999 bytes or smaller to successfully deploy. File sizes of 54000 bytes or larger will result in an error and fail to deploy.
`netmask`	Display the netmask address.

Option	Description
`nic`	Display or configure the system network configuration. If no options are given, display the system networking configuration. **Command Option:** Use the **-f** option to execute the command with no confirmation prompt and minimal explanatory text. Example: `netconfig nic -f node 418e0019302001001 primary link-count 2 mode default` **Scope Options:** Some commands allow you to specify the scope as an entire realm (**all**), an entire enclosure (**enclosure**), or an individual Director (**node**). - **all**: The setting is applied to all the appropriate Director Nodes in all enclosures. - **enclosure**: The setting is applied to all the appropriate Director Nodes in the specified enclosure. - **node**: The setting is applied only to the specified Director Node. **Primary and Secondary Options:** - **primary**: Specifies the primary interface. - **secondary**: Specifies the secondary interface. **Link Count Option:** The **link-count** specifies the number of physical links (4|2|1|0). Not all **link-count** values are valid all the time, depending on current hardware and software settings. A **link-count** of 0 is never valid for a primary interface. **Mode Options:** Not all modes are valid for every platform. - **single**: Specifies single-link mode. - **failover**: Specifies failover mode. - **static**: Specifies staticlinkagg mode. - **lacp**: Specifies LACP mode. - **default**: Specifies the platform-specific default networking mode. The default primary

Option	Description
	mode is lacp; the default secondary mode is failover.
`nic [all|enclosure|node] secondary [enable|disable]`	This command explicitly configures the associated Director Node secondary interface settings. - **netconfig nic ... secondary enable**: Configures the specified secondary interfaces with the default dual-network settings. - **netconfig nic ... secondary disable**: Deletes all the configuration settings from the specified secondary interfaces.
`nic [all|enclosure|node] [primary|secondary] link-count [4|2|1|0] mode [single|failover|static|lacp|default] ? timeout [short|long]?`	**netconfig nic ... [primary|secondary] ...**: Explicitly configures the various Director primary and secondary interface settings. **netconfig nic ... secondary ...**: Explicitly configures the associated Director Node secondary interface settings. The optional **timeout** argument applies only to **lacp** mode and specifies the LACP protocol timeout. A short timeout equals 3 seconds; a long timeout equals 90 seconds. The default value is short.
`nic node [Node-ID] warn [enable|disable] [nic ID]`	This command configures the warning settings for a node's physical NICs. Valid NIC ID values, depending on the node type, are: `NIC4|NIC5|NIC6|NIC7|NIC8|NIC9|1ge1|1ge2|10ge1|10ge2|10ge3|10ge4|25ge1|25ge2|40ge1|40ge2|bge0|bge1|cxl0|cxl1|cxl2|cxl3|igb0|igb1|ix0|ix1|enp179s0f0|enp179s0f1|mxge0|mxge1|all` The NIC ID value of **all** will enable or disable warnings for all NICs in the node. NIC warnings are an important safeguard, so use caution if you choose to disable NIC warnings.

Option	Description
`repair`	This interactive command allows you to repair networking settings that were not entered correctly during a prior call to **netconfig edit**. This tool must be run on each replication set (repset) node and, once run on all repset nodes, all repset nodes must be rebooted. You will be forced to choose a single node to be the sole repset member. You will be prompted for a new system IP address, netmask, default router, and IP address range. Netmasks must be entered in the form **255.255.255.0**.
`routers`	Display the default routers.
`routers add [ip1] ?[ip2] ...?`	Add default routers.
`routers delete [ip1] ?[ip2] ...?`	Remove default routers.
`sshkey ?list?`	Display the list of ssh-authorized key files.
`sshkey copy [user] [PanFS-file]`	Output the current SSH keys for the given user to the given PanFS file. Files must be specified using a full path starting with **/volume-name**.
`sshkey clear [user]`	Delete the current SSH key for the given user.
`sshkey deploy [user] [PanFS-file]`	Reads the contents of the given PanFS file and saves its contents for the given user. The new setting is automatically deployed to all nodes in the system. Files must be specified using a full path starting with **/volume-name**.

Option	Description
`sshkey publickey ?keygen?`	Print the system's public SSH key. The key is automatically generated once by the system and is kept permanently or until the admin invokes the **keygen** option to generate a new system key. Once a new public key is generated, the old one is no longer effective. The public key allows all nodes to make password-free **ssh** and **scp** contact with each other, as well as with remote hosts configured with this key. Configuration of this public key is needed for use with features such as replication. The admin can simply add this key to the **/root/.ssh/authorized_keys** file on an external host.
`staticroute`	Display configured static routes. Static routes may be used to route traffic for an IP subnet via a router other than the default router. The router may be reachable via the primary or secondary interface.
`staticroute add [host|network] [netmask] [gateway1] ?[gateway2] ...?`	Add a host or network route via the gateway. You must give a host IP, netmask, and at least one gateway IP address.
`staticroute delete [host|network] [netmask] [gateway]`	Delete a static route. You must give the host IP, netmask, and one of the gateway IP addresses for an existing static route.
`staticroute edit [host|network] [netmask] [old-gateway] [new-gateway1] ?[new-gateway2] ...?`	Edit the gateways for a static route. You must give the host IP, netmask, and one of the gateway IP addresses for an existing static route, as well as one or more new gateway IP addresses. The gateway address is the next hop to route packets through. Netmasks must be entered in the form **255.255.255.0**.
`subnet`	List all the subnets in the system.

Option	Description
`subnet add [netmask] [default router] [start] ?[end]?`	Add a subnet. You must supply a netmask, default router IP address, and an IP address range. Netmasks must be entered in the form **255.255.255.0**. Additional default routers and IP ranges can be added later using the **netconfig** routers **add** and **netconfig ipranges add** commands.
`subnet delete [subnet]`	Delete a subnet. You can specify an unused subnet in the form **IP/prefix**. Please note all the IP ranges configured in the subnet will be deleted from the system. If any IP aliases have been configured in the subnet, you must delete the IP aliases before you can delete the subnet.

netdiag

Description

Runs network diagnostic tests and displays results.

Syntax

```
netdiag
```

netperf

Description

Turns the **netperf** daemon on/off on all nodes in the system.

Netperf is a benchmark that can be used to measure various aspects of networking performance. Its focus is on bulk data transfer and request/response performance using either TCP or UDP and the Berkeley Sockets interface.

When enabled, the **netperf** server will run on port 4005.

Syntax

```
netperf ?enable|disable?
```

netstat

Description

Displays network status, including the state of all sockets and numeric network addresses.

Syntax

```
netstat ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`[-AaLnSW] [-f protocol_family | -p protocol] [-M core] [-N system]`	Display a list of active sockets (protocol control blocks) for each network protocol, for a particular protocol_family or for a single protocol. - -**A**: Show the address of a protocol control block (PCB) associated with a socket. Used for debugging. - -**a**: Show the state of all sockets; normally sockets used by server processes are not shown. - -**L**: Show the size of the various listen queues. The first count shows the number of unaccepted connections, the second count shows the amount of unaccepted incomplete connections, and the third count is the maximum number of queued connections. - -**S**: Show network addresses as numbers (as with -n) but show ports symbolically.

Option	Description
`-i | -I interface [-abdnt] [-f address_family] [-M core] [-N system]`	Show the state of all network interfaces or a single interface which have been auto-configured. Interfaces statically configured into a system, but not located at boot time, are not shown. An asterisk (*) after an interface name indicates that the interface is down. - -**a**: Show multicast addresses currently in use for each Ethernet interface and for each IP interface address. Multicast addresses are shown on separate lines following the interface address with which they are associated. - -**b**: Show the number of bytes in and out. - -**d**: Show the number of dropped packets. - -**t**: Show the contents of watchdog timers.
`-w wait [-I interface] [-d] [-M core] [-N system]`	At intervals of **wait** seconds, display the information regarding packet traffic on all configured network interfaces or a single interface. If **-d** is also present, show the number of dropped packets.
`-s [-s] [-z] [-f protocol_family | -p protocol] [-M core] [-N system]`	Display system-wide statistics for each network protocol, for a particular protocol_family, or for a single protocol. - -**s**: If repeated, counters with a value of zero are suppressed. - -**z**: Reset statistic counters after displaying them.
`-i | -I interface -s [-f protocol_family | -p protocol] [-M core][-N system]`	Display per-interface statistics for each network protocol, for a particular protocol_family, or for a single protocol.
`-m [-M core] [-N system]`	Show statistics recorded by the memory management routines. The network manages a private pool of memory buffers.

Option	Description
`-r [-AanW] [-f address_family] [-M core] [-N system]`	Display the contents of all routing tables or a routing table for a particular address_family. - -**A**: Show the contents of the internal Patricia tree structures; used for debugging. - -**a**: Show protocol-cloned routes (routes generated by an RTF_PRCLONING parent route). When -**W** is also present, show the path MTU for each route.
`-rs [-s] [-M core] [-N system]`	Display routing statistics. If -**s** is repeated, counters with a value of zero are suppressed.
`-g [-W] [-f address_family] [-M core] [-N system]`	Show information related to multicast (group address) routing. By default, show the IP multicast virtual-interface and routing tables.
`-gs [-s] [-f address_family] [-M core] [-N system]`	Show multicast routing statistics. If -**s** is repeated, counters with a value of zero are suppressed.

For additional general options, see the PanCLI **help netstat** command.

netswitch

Description

This command displays switch hardware and firmware information, as well as switch status for all types of switches. Switch reset action is available for all switch types except Direct Connect 10 switches. There is additional support for port information, switch statistics, and other switch management actions for 10 GE switches.

Syntax

```
netswitch [subcommand] [arg1] ?[arg2] [arg3]?
```

Options

Option	Description
`cli [switch-number]`	This command is only supported for 10 GE switches. This command launches the interactive **remote_cli** of the local switch identified by the switch number. Use **Ctrl-c** to exit the **remote_cli** shell. - **[switch-number]**: Must be either 1 or 2, indicating the NET1 or NET2 slot for the switch to operate on.
`fdb [show]|[flush] [switch-number]`	This command is only supported for 10 GE switches. - **[switch-number]**: Must be either 1 or 2, indicating the NET1 or NET2 slot for the switch to operate on. - *show* : Display the learned MAC address table of the given local switch. - *flush* : Flush the learned MAC address table of the given local switch. The user will be prompted interactively for confirmation before any action is taken.

Option	Description

`info`
`info [enclosurename|enclosure-serial] ? [switch-number]?`
`info [switch-serial]`
	Display hardware, firmware, and port information for one or more switches. Port information is only available for 10 GE switches.
	If no enclosure or switch is specified, show information for all switches in the system.
	If a switch serial or switch number is given, this command will give information about the single specified switch. The switch number, if given, must be 1 or 2, indicating the NET1 or NET2 slot. Switch serial is the unique 17-digit alphanumeric hardware serial number of the switch.
	If an enclosure name or enclosure serial is given without specifying a switch number, then this command will give information for both switches in the enclosure.

`reset [enclosurename|enclosure-serial] [switch-number]` `reset [switch-serial]`	Reset the switch. The switch-number, if given, must be either 1 or 2, indicating the NET1 or NET2 slot. Any switch can be reset using this command. Switch serial is the unique 17-digit alphanumeric hardware serial number of the switch. The user will be prompted interactively for confirmation before any action is taken.

Option	Description

`statistics`
`statistics [enclosurename|enclosure-serial] ?[switch-number]?`
`statistics [switch-serial]`
	Display switch status and statistics for one or more switches. Statistics are only available for 10 GE switches.
	If no enclosure or switch is specified, show information for all switches.
	If a switch serial or switch number is given, then this command will only give statistics for the single specified switch. The switch number, if given, must be 1 or 2, indicating the NET1 or NET2 slot.
	If an enclosure name or enclosure serial is given without specifying a switch number, then this command will give statistics for both switches in the enclosure.

nfs

Description

The NFS gateway feature is currently unsupported. Use of this command and its subcommands may be refused.

> **Note:** See the PanFS **netconfig alias** command to enable NFS failover.

Syntax

```
nfs ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`deletemount [client]?:[path]? ?[IP address]?`	Use this command to clear outdated entries for the given client (IP or hostname) from the NFS mount list. If a PanFS file path is given, only mount table entries for that path will be deleted. If an IP address is given, only mounts to that IP address will be deleted.
`enabled ?yes|no?`	Display or change whether or not NFS access to the system is enabled.

`exports ?list v3?`
`?copy [file] v3?`
`?upload [file] v3?`
	View and configure NFS v3 exports.
	- **list**: Display the NFS v3 configuration.
	- **copy**: Copy the NFS v3 exports file contents to a given location.
	- **upload**: Deploy and save the given NFS v3 exports PanFS file.

Option	Description
`mountport ?[port-number]?`	Service mode command to display or set the port used by NFS **mountd**. Setting this parameter will force NFS mountd to bind to the specified port. - The port number should be a number between 0 and 65535 (inclusive). If the port number is set to 0, NFS mountd will use a system assigned port number. - If no argument is given, the current port number is displayed. Changing this parameter will stop and restart the NFS mountd running on each Director, which will abort running sessions.
`showmount ?ip [IP address]? ?client [client]? ?path?`	Display the NFS mount list. - **ip**: limit the results to mounts against a specific IP address. - **client**: Limit the results to mounts made by a specific client (IP or hostname). - **path**: Results will include the PanFS file path for each mount.

nis

Description

Manages NIS settings.

Syntax

```
nis ?domain|hosts|servers|unix|debug ?[arg1]??
```

Example

Configure with 3 NIS servers:

```
nis servers "nis1.mycompany.com nis2.mycompany.com nis3.mycompany.com"
```

Clear all NIS servers:

```
nis servers ""
```

Options

Option	Description
`debug`	This is the debug level used by name resolution services. To set this debug level, use the nameresorder command. Any debug level above 0 will impact performance.
`domain ? [domainname]?`	Display or set the NIS domain name. Use an empty string (“”) to clear the value.
`hosts ?yes|no?`	Display whether or not NIS hostname resolution is enabled. Use **yes** to use yp/NIS for hostname lookups.
`servers ?[hosts]?`	Display or set the list of NIS servers. Use an empty string to clear.

Option	Description
`unix ?yes|no?`	The NIS UNIX name resolution feature allows the system to perform username to UID mappings, group name to GID mappings, and netgroup lookups via NIS. The NIS domain and servers must be valid to enable this feature. You must also add NIS to the name resolution search ordering before the system will perform NIS resolution. Use the nameresorder command to manage name resolution search order.

notify

Description

Displays and updates error notification preferences.

Syntax

```
notify ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Start sending status emails to bob@email.com:

```
notify config statusemail address add bob@email.com
```

Start sending critical emails to pager address 8675309@pager.com:

```
notify config critical pager address add 8675309@pager.com
```

Send a summary test message to bob@email.com and sysadmin@yourco.com:

```
notify test summary bob@email.com sysadmin@yourco.com
```

Options

Option	Description
none	Show current notification settings.
`config (critical|warning) (email|pager) address (add|delete) [address] ?[address1] ...?`	Add or remove one or more email or pager addresses to those receiving critical or warning messages.
`config (critical|warning) (email|pager) address clear`	Clear all critical or warning email or pager addresses.

Option	Description
`config (critical|warning) (email|pager) frequency (incident|day|hour|minute|never)`	Set the frequency for sending critical or warning alerts.
`config statusemail address (add|delete) [address] ?[address1]...?`	Add or remove one or more email addresses to those receiving status email messages.
`config statusemail address clear`	Clear all status email addresses.
`config statusemail frequency (week|day|hour|never)`	Set the frequency for sending status emails.
`test ?[address1] ...?`	Send a status message to the given addresses. If no addresses are specified, the current status email settings are used.
`warnclientawol ?yes|no?`	Control whether or not to raise warning alerts when clients fail to respond. The system automatically detects failure of clients to respond and reports such failures in the event log. This control tells the system whether to additionally raise a system-wide warning alert that triggers email and pager messages.

ntp

Description

Displays and manages NTP servers.

Specify three NTP servers for redundancy. If a node reports clock drift, remove any server that responds inconsistently, or use one reliable server. If no NTP servers are specified, one Director is automatically chosen to operate as a stratum 0 time source for the system. In this case, the **date** command can be used to set the time for all nodes in the system.

Syntax

```
ntp ?[subcommand] ?[args]??
```

Options

Option	Description
none	Displays current list of NTP servers.
`add ?-noreboot? [server-name]|[ip]`	NTP servers will be used in the order in which they are listed. Configuration of external NTP servers in a system that uses a local NTP server triggers a system reboot. Interactive confirmation is required in this case. -**noreboot** : Use under the direction of VDURA Technical Support to avoid rebooting the system.
`clear`	Clear the list of NTP servers. Configuration of external NTP servers in a system that uses a local NTP server triggers a system reboot. Clearing all NTP servers from the configuration requires interactive confirmation, since the change cannot be reversed without a system reboot.

Option	Description
`delete [server-name]|[ip]`	Remove the given NTP server name or IP address from the list. Configuration of external NTP servers in a system that uses a local NTP server triggers a system reboot. Deletion of the last NTP server from the configuration requires interactive confirmation, since the change cannot be reversed without a system reboot.

objutil

Description

Service mode command to list and manipulate filesystem objects.

Syntax

```
objutil [subcommand] [arg]
```

Example

An administrator can view and modify ACLs using the `objutil listacl` and `objutil setacl` commands, respectively. The `objutil setacl` command may add, remove, or modify an ACE if the ACL contains a ACE with a matching type, identity, and inheritance flags. If no matching ACE is found, a new ACE will be added to the ACL. If a match is found, the permissions will be updated unless :none: is specified, as :none: will delete the ACE from the ACL.

In the example below, an ACE is removed for everyone, an ACE is added for UID 640, and UID 640 permissions are modified to have all permissions.

Listing the default setting:

```
[pancli_service] objutil listacl /home
```

ACL for /home is:

```
1 +uid:0,:all:
2 +gid:0,rxnRP
3 +pan_gid:1,rxnRP
```

Remove the ACE for everyone:

```
[pancli_service] objutil setacl +pan_gid:1,:none: /home
[pancli_service] objutil listacl /home
```

ACL for /home is:

```
1 +uid:0,:all:
2 +gid:0,rxnRP
```

Add an ACE for UID 640:

```
[pancli_service] objutil setacl +uid:640,rxnRP /home
[pancli_service] objutil listacl /home
```

ACL for /home is:

```
3 +uid:640,rxnRP
1 +uid:0,:all:
2 +gid:0,rxnRP
```

Modify UID 640 permissions to have all permissions:

```
[pancli_service] objutil setacl +uid:640,:all: /home
[pancli_service] objutil listacl /home
```

ACL for /home is:

```
3 +uid:640,:all:
1 +uid:0,:all:
2 +gid:0,rxnRP
```

Options

Option	Description
`attrs [objectID|path]`	Display attributes for the specified pathname or object ID. Files must be specified using a full path starting with **/volume-name**. Object IDs are of the form **l-xDnnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnnn**, where each *n* is a hex digit.
`compattrs [objectID|path]`	Display all component object attributes for specified object or file. Files must be specified using a full path starting with **/volume-name**. Object IDs are of the form **l-xDnnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnnn**, where each *n* is a hex digit.

Option	Description
`dircompact [objectID1|path1] ? [objectID2|path2] ... [objectIDN|pathN]?`	Reclaim space used by directories after data is removed. Directories may continue to occupy excessive space after entries are removed. This tool reclaims that space by removing holes left after entries are removed from the directory. The **objutil ls** command may be used to list the directory and determine if it would benefit from compaction. Compacting the directory will not damage the filesystem, but may lead to inconsistent results during the compaction if filesystem clients are in the process of reading the directory. Directories must be specified using a full path starting with **/volume-name**. Object IDs are of the form **I-xDnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnn**, where each *n* is a hex digit.

`id [path]`
`id snapshot:[snapshot-name]:[path]`
	Display the object ID of the given pathname.
	Files must be specified using a full path starting with **/volume-name**.
	Example:
	`objutil id /home/myfile`
	Use the **snapshot:[snapshot-name]:** prefix to the file path to get the ID of a file or directory object that resides in a snapshot.
	Example:
	`objutil id snapshot:2009.04.21.16.32.17.hourly:/home/myfile`

Option	Description
`listacl [objectID|path]`	Display ACL for specified path or objectID. Files must be specified using a full path starting with **/volume-name.** Object IDs are of the form **I-xDnnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnnn**, where each *n* is a hex digit.
`ls [objectID|path]`	List the contents of the specified directory. Files must be specified using a full path starting with **/volume-name.** Object IDs are of the form **I-xDnnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnnn**, where each *n* is a hex digit.
`paths [objectID|path]`	Display path(s) to the specified object ID or file. Files must be specified using a full path starting with **/volume-name.** Object IDs are of the form **I-xDnnnnnnnnnnnnnnnn-xGnnnnnnnn-xUnnnnnnnnnnnnnnnn**, where each *n* is a hex digit.

Option	Description
`setacl [ACE] [objectID]`	Set an ACE (ACL entry) for specified object. ACL entries (ACEs) are printed in the form **(+/-)identity,perms[,I:inheritance flags]**. - **+** indicates a positive ACE that grants the specified permissions. - **-** indicates a negative ACE that denies the specified permissions. - Identity specifies the target user or group and is one of the following: Identity, Type, Description: `uid:NNN` , UNIX/NIS, UID NNN `gid:NNN` , UNIX/NIS, GID NNN `user_sid:NNN` , Windows, user SID NNN `group_sid:NNN` , Windows, group SID NNN `pan_gid:NNN` , Internal, PanFS group NNN The following internal PanFS groups are recognized: Group, Description: `pan_gid:1` , everyone group - The perms portion of the ACE is a string composed of the following characters: Permission, Abv, **Effect on** File, **Effect on** Dir: read/list dir, `r` , read data contents, list dir write, `w` , modify data contents, n/a execute, `x` , read data contents (note 1), list dir append, `a` , append data to object, n/a create, `c` , n/a, create files delete, `d` , delete files, delete dir lock, `k` , take advisory or exclusive locks, n/a change ACL, `p` , modify ACL, modify ACL take ownership, `o` , change owner, change owner sandbox, `s` , n/a, (note 2) create subdirs, `C` , n/a, create subdirs read named attributes, `n` , read named attributes, n/a write named attributes, `N` , write named attributes, n/a

Option	Description
	delete child, `D`, n/a, delete files or subdirectories read attributes, `R`, read attributes, read attributes write attributes, `W`, write attributes, write attributes read ACL, `P`, read ACL, read ACL synchronize, `S`, n/a, n/a write retention, `h`, n/a, n/a write retention hold, `H`, n/a, n/a For the perms you can also specify **:all:** to specify all permissions or **:none:** to specify no permissions. An ACE that specifies no permissions will remove an existing ACE that matches the type (**+** or -), identity, and inheritance flags. Notes: 1. From the perspective of the file server, there is no way to control what a client does with the contents of a file after reading it on Windows. Therefore, the **execute** permission is treated as an advisory hint to the Windows client as to whether the contents of the file are executable, rather than being an enforced prohibition on executing the file. On Unix, a binary file cannot be executed without the **execute** bit on. 2. A directory with the **sandbox** bit set operates similarly to a directory with the Unix **sticky** bit set. Users can only delete files and subdirectories that they own, not those owned by other users. The inheritance flags is a string composed of one or more of the following 2 letter flags: Flag, Meaning: `OI`, ACE will be inherited to objects (files) `CI`, ACE will be inherited to containers (directories)

Option	Description
	`IO` , ACE will be inherited but does not affect permissions of this object `NP` , ACE will be inherited one level down only `ID` , ACE was inherited from the parent directory VDURA ACLs are similar, but not identical, to Windows ACLs. ACEs have the following precedence, from most important to least important: 1. explicit deny ACE 2. explicit allow ACE 3. inherited deny ACE 4. inherited allow ACE When a file is accessed by a NFS or DirectFlow client, the ACL is translated to a set of POSIX mode bits (read/write/execute for user/group/other).

osdv4_admin

Description

Modifies the configuration of a running **iopathv4**. This command should only be run under the supervision of VDURA Technical Support.

Syntax

```
osdv4_admin [commands] ?[options]?
```

Options

Commands

Option	Description
`add-migration-policy`	Add a migration policy to the DSM (requires –**migration**-**policy**).
`dump-config`	Dump the current configuration used by iopathv4.
`dump-tasks`	Dump the iopathv4 task queues to the journal.
`get-assignment-policy`	Get the current DSM assignment policy.
`get-migration-bandwidth`	Get the bandwidth allotted to object migration.
`get-migration-search-period`	Get the period at which migration entries are identified.
`get-ssd-bucket-ceiling`	Get the SSD bucket ceiling.
`get-ssd-fullness-target`	Get the SSD fullness target percentage.
`list-all-assignment-policies`	List all assignment policies available to the DSM.
`list-all-migration-policies`	List all migration policies available to the DSM.
`list-migration-policies`	List the migration policies in the DSM.

Option	Description
`locate-object`	Locate the managing DS for an object (requires –**object**).
`remove-migration-policy`	Remove a migration policy from the DSM (requires –**migration**-**policy**).
`set-assignment-policy`	Set the DSM assignment policy (requires –**assignment**-**policy**).
`set-migration-bandwidth`	Set the bandwidth allotted to object migration (requires –**bandwidth**).
`set-migration-rate`	Set the migration rate used by the DSM (requires –**rate**).
`set-migration-search-period`	Set the period at which migration entries are identified (requires –**search**-**period**).
`set-ssd-bucket-ceiling`	Set the SSD bucket ceiling (requires –**bucket**).
`set-ssd-fullness-target`	Set the SSD fullness target percent (requires –**percent**, only valid up to 95%).

Options

Option	Description
`--assignment-policy=uint32`	DSM assignment policy
`--bandwidth=uint32`	Bandwidth cap for object migration (MiB/s)
`--bucket=uint8`	Bucket number
`--migration-policy=uint32`	DSM migration policy
`--migration-rate=uint64`	Migration rate used by the DSM (MiB/s)
`--object=stor_id`	Object ID
`--percent=uint8`	Percentage value (0 - 100)
`--search-period=sec:nse`	Time period between migration searches

osdv4_compress_metadata_db.py compress_db

Description

Command to compact VPOD metadata databases. This command performs compaction operations based on the size of the database, potentially utilizing HDDs for additional space. This command should only be run under the supervision of VDURA Technical Support.

Syntax

```
diag_exec "osdv4_compress_metadata_db.py compress_db ?[options]?" storagemodules
```

Example

Compress VPOD metadata and:

- Use the /pandata/dump_here directory.
- Wait 9600 seconds before timing out.
- Clean the backup directory before compression.

```
diag_exec "osdv4_compress_metadata_db.py compress_db --dump-dir /pandata/dump_here -t 9600 --force-clean-workdir" storagemodules
```

Options

Option	Description
`-h, --help`	Show usage message.
`--threshold THRESHOLD`	Minimum percentage of free space required to stop the compression. Default is 10.0.
`--compress-all`	Indicate that all databases should be compressed. Default is **false**.
`--dump-dir DUMP_DIR`	Path to the directory where the uncompressed database will be dumped. Default is **/pandata/1**.

Option	Description
`-t TIMEOUT, --timeout TIMEOUT`	Maximum amount of time allowed for the operation in seconds. Default is 3600.
`--force-clean-workdir`	Indicate that the backup directory should be forcefully cleaned before compression. Default is **false**.
`--try-restore-databases-on-fail`	Indicate that the script should attempt to restore the databases from backup if compression fails. Default is **false**.

osdv4_rebalance

Description

Command to rebalance metadata for object storage. This command should only be run under the supervision of VDURA Technical Support.

Syntax

```
osdv4_rebalance.py [-h] {migrate,balance} ...
```

Options

Option	Description
`migrate`	Reconfigure OSDv4 filesystem to store Copy-On-Write metadata of objects on flash drive and trigger rebalancing.
`balance`	Trigger rebalancing of OSDv4 metadata.

pan_atime

Description

Obtains Storage Node estimate of access time on a file or directory.

Pan_atime provides a mechanism to retrieve an estimate of the access time on a file by directly fetching attribute values from Storage Nodes. The utility analyzes object creation time, per-component data modification, and access time to determine the correct access time for the file system object. The algorithm does not take into account object components that were moved as a result of reconstruction.

Pan_atime is not a POSIX-compliant access time, and it is not possible to set the value of this timestamp directly. A realm must be mounted in order for the utility to work. It functions over both DirectFlow and NFS mounts but is substantially more efficient when run over DirectFlow. The utility prints its output in a format approximating **ls -lu**.

Access time is updated only by read operations that reach Storage Nodes. If a client read is satisfied entirely from a client or Director Node cache, no **atime** update will occur.

A DirectFlow Client must use a set of mount options to control **atime** update behavior:

osd-timestamp-bypass: If disabled (set to 0), atime updates on each data access, regardless of pan-atime or pan-diratime values. If enabled (set to 1,) atime remains unchanged across reads by the client, unless pan-atime or pan-diratime is enabled.

pan-atime: If enabled (set to 1), file **atime** will be updated once per **pan-atime-update-interval** seconds on each attribute storing component of the object.

pan-diratime: The same functionality as **pan_atime**, but for directories only.

pan-atime-update-interval: Define time period in seconds which should pass since last data access to update **atime**. Applies only to **pan-atime** or **pan-diratime**.

NFS clients require no special mount options, but **atime** updates must be enabled on the Director Nodes in the realm. This is done using the service mode PanCLI command **gwctl atime enable**. Without this option, reads by NFS clients will not update **atime**.

CONTACTING A REALM

Pan_atime will make direct contact with the realm from user level, even when run over a DirectFlow mount. If the client where **pan_atime** is running has a global DirectFlow mount (**/panfs**) and the file being accessed is in the default realm, no additional options are required. Otherwise, you must specify the realm name and the realm IP address. This can be done in one of three ways:

- Command line arguments: Invoke the program using both **–bindcache-realm-name** and **–bindcache-rm-addr**.
- Environment variables: Set the **PAN_BINDCACHE_RM_ADDR** environment variable in the shell that runs the utility.
- Configuration file: Create a configuration file in the following format:

```
<realm>
name = "REALM NAME HERE";
<realm_mgr>
<mgr_addr>
type = "ipv4";
ip_addr = "REALM MANAGER IP ADDRESS HERE";
ip_port = 3095;
</mgr_addr>
</realm_mgr>
</realm>
```

The default location of this file is **/var/pan/pan_rm_client_config.cfg**. This file can be overridden via the **–bindcache-config-dir** and/or **–bindcache-config-name** command-line options or by setting the **PAN_BINDCACHE_CONFIG_DIR** and/or **PAN_BINDCACHE_CONFIG_NAME** environment variables in the shell that runs the utility.

CAVEATS

The timestamps returned are measured against the Storage Node's internal clock and converted to the client's time zone using the standard Unix system calls for this purpose. Therefore, the client should be synchronized to the same NTP server as the realm to fetch results meaningful to the client.

If you run the utility without using any of the three methods for locating a realm, and there is no **/panfs** global mount, the system will print a message asking you to create a configuration file and then immediately exit.

Syntax

```
pan_atime [options] file [file ...]
```

Options

Option	Description
`--bindcache-config-dir=local-dir`	Set the directory to search for the config file. See “Contacting a Realm”.
`--bindcache-config-name=filename`	Set the file name for the config file. See “Contacting a Realm”.
`--bindcache-rm-addr=ip-address`	Set the realm IP address of the realm to contact. This should be the realm’s primary IP address as configured in the GUI and should be in dotted-decimal format. See “Contacting a Realm”.
`-d, --debug`	Display debug-oriented information to **stderr**; useful only for debugging problems with the utility itself.
`--defaults`	Display defaults.
`--env-equiv`	Display environment var equivalents.
`-f, --format=string`	Set the format of the printed timestamp. This utility supports all the formatting primitives defined in **strftime(3)**, plus **%P**, which displays as the number of days since the last **atime** update.
`-h, --help`	Display help message.
`-m, --max-retries=uint32`	Maximum number of retries on system calls receiving EAGAIN.

Example

```
% setenv PAN_BINDCACHE_RM_ADDR 10.65.126.11
% pan_atime /panfs/panfs.panasas.com/home/dir0/file0
-rw-r--r-- 1 root root 3072 Jun 27, 2025 08:26:47 /panfs/panfs.panasas.com/home/dir0/file0
% pan_atime -f '%s' /panfs/panfs.panasas.com/home/dir0/file0
-rw-r--r-- 1 root root 3072 1119875207 /panfs/panfs.panasas.com/home/dir0/file0
```

pan_attr

Description

Allows setting and querying attributes of a PanFS object.

SUPPORTED ATTRIBUTES

cw_open: Legal values are 0 and 1. When set to 1, DirectFlow Clients will always attempt to open the file in concurrent write (CW) mode. This flag has no effect for files opened via NFS or CIFS. Note that CW mode relaxes distributed cache consistency and is not suitable for all applications.

def_layout: Set the default RAID layout for files created in this directory. Only files created after the attribute is set are affected; existing files use their existing RAID layout. The value of this attribute is also inherited statically, at create time, by new subdirectories of this directory. The value should be a valid string description of a PanFS RAID layout, See the "Layout Descriptors" section below.

LAYOUT DESCRIPTORS

PanFS RAID layouts can be described in long form or short form. A long form layout descriptor consists of a token representing the RAID type, followed by one or more key/value pairs separated by an equals sign (=). Note that since the layout descriptors include spaces, they must typically be quoted when specified on the command line. All parameters of the long form layout must be provided, otherwise the layout will not be parsed correctly.

A value of 0 indicates that the system default should be used.

The recognized long form layout types and description formats are:

- RAID-10 max_comps=mcomps stripe_unit=unit policy=policy
- RAID-10 (striped mirrors).
- mcomps: The number of components to stripe across. Must be a multiple of 2.
- unit: The amount of data (in bytes) to place in one mirror pair before moving to the next mirror pair. Must be a multiple of 4 KB.
- policy: The striping policy. The only supported policy is 1 (round-robin).
- Specifying 0 for any parameter will select the system default.
- RAID-1/5-GROUPS max_comps=mcomps stripe_unit=unit group_width=gwidth group_depth=gdepth policy=policy
- RAID groups (parity stripes).
- mcomps: The total number of components to stripe across all RAID groups. Must be a multiple of **gwidth**.

- unit: The amount of data (in bytes) to place on one component before moving to the next component in the layout. Must be a multiple of 4 KB.
- gwidth: The number of components in each RAID group.
- gdepth: The number of stripes to place in a given RAID group before moving to the next RAID group.
- policy: The striping policy. The only supported policy is 1 (round-robin).
- Specifying 0 for any parameter will select the system default.
- RAID-N-GROUPS max_comps=mcomps stripe_unit=unit group_width=gwidth group_depth=gdepth max_faults=mfaults encoding=encoding: Similar to RAID-1/5-GROUPS above.
- mcomps, **unit**, **gwidth** and **gdepth** have the same meaning as above.
- mfaults: The number of faults tolerated by this RAID layout. The only supported value of **mfaults** is 2, making this essentially a RAID6 layout.
- The only supported encoding value is 1, which utilizes Reed-Solomon encoding.
- Specifying 0 for any parameter will select the system default.

The recognized short form layout description formats are:

- RAID-10
- RAID-1/5-GROUPS
- RAID-N-GROUPS

EXIT STATUS

- 0: Success.
- 1: Abnormal exit.
- 2: Syntax error.
- 3: **ioctl** not supported on this device.
- 4: **ioctl** version error.
- 5: Permission denied.
- 6: Target not a directory.
- 7: Error parsing layout.

Syntax

```
pan_attr [options] filename
```

Options

Option	Description
`--defaults`	Display defaults.

Option	Description
`--env-equiv`	Display environment var equivalents.
`-e, --examples`	Display examples.
`-g, --get`	Get value of an attribute, e.g., **-g cw_open**.
`-h, --help`	Display help message.
`-q, --quiet`	Do not display error messages.
`-s, --set`	Set value of an attribute, e.g., **-s cw_open:1** or **-s cw_open:0**.
`--version`	Display program version.

Example

The pathname parameter is the path to a specific object, e.g., /panfs/REALM/home/myfile.c or /panfs/REALM/home/mydir.

Getting attributes

Get cw_open attribute from a file:

```
pan_attr -g cw_open /panfs/REALM/home/myfile.c
```

Setting attributes

Set cw_open attribute to a file:

```
pan_attr -s cw_open:1 /panfs/REALM/home/myfile.c
```

Set RAID-1/5-GROUPS layout with a maximum of 20 components and RAID group width 10:

```
pan_attr -s "def_layout:RAID-1/5-GROUPS max_comps=20 stripe_unit=0 group_width=10 group_depth=0
policy=0" /panfs/REALM/home/mydir
```

Set RAID-1/5-GROUPS layout using the short-form description:

```
pan_attr -s "def_layout:RAID-1/5-GROUPS" /panfs/REALM/home/mydir
```

Clearing attributes

Clear cw_open attribute to a file:

```
pan_attr -s cw_open:0 /panfs/REALM/home/myfile.c
```

pan_collect_logs

Description

Copies log files (e.g., **/var/log/messages**) and collects various information that is useful in debugging problems with the DirectFlow Client.

The output is organized under a directory that includes the host-name and timestamp, which lets you collect logs from various clients at various times all in one place.

For best results, this command should be run as root so the required log files can be read and all utilities can be executed.

RETURN VALUES

Pan_collect_logs exits with 0 on success or >0 if an error occurred.

Syntax

```
pan_collect_logs [OPTION]...
```

Options

Option	Description
`--defaults`	Display program defaults.
`-d,--dir DIRECTORY`	This is a mandatory argument that specifies the directory under which to place the collected files. An additional sub-directory is created under directory with the hostname of the client and a timestamp. These directories are created automatically by **pan_collect_logs**.
`-h, --help`	Display help message.

pan_debug_enum_datasets

Description

Service mode command to list available datasets.

Syntax

```
pan_debug_enum_datasets ?[options]? (-n [host-ip] -p [port])|(-m [mgr-id])
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Don't consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_debug_print_layout_table

Description

Service mode command to print the SM internal device table.

Syntax

```
pan_debug_print_layout_table ?[options]? (-n [host-ip] -p [port])|(-m [mgr-id])
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.
`-s, --sm-mgr-id=mgr_id`	SM mgr_id to fetch layout table from (needed if using -**n**/-**p**).

pan_debug_print_one_stat

Description

Service mode command to print a named debug_dataset from a single target.

Syntax

```
pan_debug_print_one_stat ?[options]? (-n [host-ip] -p [port])|(-m [mgr-id]) [data_name]
```

The command must always include the data_name (string) argument, which is the dataset name.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID(dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Don't consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.
`-z, --zero-after-fetch`	Zero stats after fetching.

pan_debug_trace_archive_collect

Description

Service mode command to collect traces from a circular buffer to a file, with archiving.

This command is a wrapper for **pan_debug_trace_collect** with additional archiving and compression functionality. Options such as –**mgr-id** or –**host-ip** and –**port** are passed uninterpreted to **pan_debug_trace_collect**.

Once started, a **pan_debug_trace_archive_collect** task generally runs forever in the background, collecting trace data to **[collect-file]** and managing the archive files. It will exit when trace collection fails consistently for a period of time, adjustable with –**collect-retry**. The default retry duration is 90 seconds. Hence, it will tolerate a simple restart of the target process, but not a long network outage or reboot.

Errors and other diagnostic messages go to **[collect-file].out** while the task runs in the background.

Syntax

```
pan_debug_trace_archive_collect (--archive-num [n]) (--archive-size [kbytes]) (--archive-zip)
(--collect-retry [seconds]) ?collect-options? [id] [collect-file]
pan_debug_trace_archive_collect --cancel [collect-file]
```

The recommended way to stop this task and signal the task associated with [collect-file] to stop is:

```
pan_debug_trace_archive_collect --cancel [collect-file]
```

The command must always include the following arguments:

Argument	Description
`id (uint32)`	ID of trace buffer to collect.
`collect-file (string)`	Collect file.

Options

Option	Description
`--cancel`	Cancel any currently running **pan_debug_trace_archive_collect** associated with **[collect-file]**.
`--archive-num [n]`	Number of archive files to keep; default is: 49.
`--archive-size [kbytes]`	Size threshold for archiving **[collect-file];** default is 100MB or 500 MB with –**archive**-**zip**.
`--archive-zip`	Compress archive files with **gzip**.
`--collect-retry [seconds]`	Seconds to retry when collect is failing; default is 90 or use 0 for unlimited.

Collect Options

Option	Description
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.
`-w, --wait`	Wait to collect until buffer is 3/4 full.

pan_debug_trace_collect

Description

Service mode command to collect traces from a circular file.

Syntax

```
pan_debug_trace_collect ?options? [id] [collect-file]
```

The command must always include the following arguments:

Argument	Description
`id (uint32)`	ID of trace buffer to collect.
`collect-file (string)`	Collect file.

Options

Option	Description
`-a, --append`	Append data to file.
`-c, --cancel`	Cancel a collect at the server.
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host id (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-1, --one-shot`	Make at most one collect RPC.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

Option	Description
`-s, --size=byte size`	Size of collect buffer in bytes; this argument is a **uint64** with an optional **[bBkKmMgGtT]** suffix. For instance, 1 means 1 byte, and 1 G means 1 gigabyte.
`-w, --wait`	Wait to collect until buffer is 3/4 full.

pan_debug_trace_destroy

Description

This command deletes the specified trace buffer from memory and closes the file if spilling traces. Note that **pan_debug_trace_stop** does not call destroy.

Syntax

```
pan_debug_trace_destroy ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_debug_trace_info

Description

This command gives details about the specified trace buffer.

Syntax

```
pan_debug_trace_info ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.
`-a, --show-all`	Show all modules regardless of filter.

pan_debug_trace_list

Description

This command lists all the traces in use.

Syntax

```
pan_debug_trace_list [options]
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_debug_trace_modify

Description

Available only in service mode. This command modifies trace filters for an existing trace buffer.

Syntax

```
pan_debug_trace_modify ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-f, --filter`	Trace module filter: **module_name:sets**.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.
`-s, --size=byte size`	Size of collect buffer in bytes; this argument is a **uint64** with an optional **[bBkKmMgGtT]** suffix. For instance, 1 means 1 byte, and 1 G means 1 gigabyte.

pan_debug_trace_reset

Description

This command resets the trace configuration to the system default.

Syntax

```
pan_debug_trace_reset ?options?
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_debug_trace_spill

Description

This command spills remaining traces in trace buffer to file.

Syntax

```
pan_debug_trace_spill ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_debug_trace_start

Description

This command starts tracing.

Syntax

```
pan_debug_trace_start ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--append=on/off`	Append to trace spill file; do not truncate.
`--clear=on/off`	Clear trace buffer before collecting.
`--collect-first`	Collect first data; stop when buffer full.
`--collect-last`	Collect last data; overwrite earlier data.
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-f, --filter`	Trace module filter: **module_name:sets**.
`--filter-all=string`	Use this filter set on all modules.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

Option	Description
`-s, --size=byte size`	Size of collect buffer in bytes; this argument is a **uint64** with an optional **[bBkKmMgGtT]** suffix. For instance, 1 means 1 byte, and 1 G means 1 gigabyte.
`-F, --spill-file=string`	Spill file.
`--spill-non-buf`	Non-buffered spill style.
`--spill-write-sync`	Spill writes to the disk are synchronous.
`--ttl=uint32`	**Time To Live** of this trace; 0=forever.

pan_debug_trace_stop

Description

This command stops tracing. This command does not call **pan_debug_trace_destroy**.

Syntax

```
pan_debug_trace_stop ?options? [id]
```

The command must always include the id (uint32) argument, which is the trace buffer ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host-ip=string`	Host ID (dotted IP address) to contact.
`-m, --mgr-id=mgr_id`	Talk to manager with this ID.
`--no-try-dp-mgr`	Do not consider datapath **pan_mgr**.
`-p, --port=string`	Host port to contact.
`-r, --realm-name=string`	Specify the system name.

pan_df

Description

Displays information about the file system on which each file resides or all file systems by default.

Pan_df displays the amount of disk space available on the file system containing each file name argument. If no file name is given, the space available on all currently mounted file systems is displayed. Disk space is shown in 1 K blocks by default, unless the environment variable **POSIXLY_CORRECT** is set, in which case 512-byte blocks are used.

If an argument is the absolute file name of a disk device node containing a mounted files system, **pan_df** displays the space available on that file system rather than on the file system containing the device node (which is always the root file system). This version of **pan_df** cannot show the space available on unmounted file systems because on most kinds of systems doing so requires nonportable intimate knowledge of file system structures.

If an argument is the absolute file name of a file or directory in a mounted VDURA DirectFlow file system, **pan_df** will attempt to use the **STATFS_EXTENDED ioctl** to fetch space usage information for the volume and/or StorageSet that contains the named file or directory. If the volume has no hard quota, the information displayed for that file will reflect the total space and space available in the StorageSet. If a hard quota is set for the volume, the total and available space will be adjusted to reflect the volume's quota.

Syntax

```
pan_df [OPTION]...[FILE]...
```

Options

Mandatory arguments for long options are also mandatory for short options.

Option	Description
`-a, --all`	Include filesystems having 0 blocks.
`-B, --block-size=SIZE`	Use SIZE-byte blocks. SIZE may be (or may be an integer optionally followed by) one of following: kB 1000, K 1024, MB 1000*1000, M 1024*1024, and so on for G, T, P, E, Z, Y.

Option	Description
`-h, --human-readable`	Display sizes in human-readable format (e.g., 1K, 234 M, 2G).
`-H, --si`	Display sizes in human readable format but use powers of 1000, not 1024.
`--help`	Display help message.
`-i, --inodes`	List inode information instead of block usage.
`-k`	Use 1 K size blocks.
`-l, --local`	Limit listing to local filesystems.
`--no-sync`	Do not invoke sync before getting usage info. This is default behavior.
`-P, --portability`	Use the POSIX output format.
`--sync`	Invoke sync before getting usage info.
`-t, --type=TYPE`	Limit listing to filesystems of type TYPE.
`-T, --print-type`	Display filesystem type.
`-x, --exclude-type=TYPE`	Limit listing to filesystems not of type TYPE.
`-v`	Ignored.
`--version`	Display version information.

pan_du

Description

Summarizes disk usage of each file, recursively for directories. Directories are traversed in parallel by multiple threads.

Syntax

```
pan_du [options] path [path ...]
```

Options

Mandatory arguments for long options are also mandatory for short options.

Option	Description
-a, –print-age-info	Display space summary by file age. The files are sorted into buckets by **mtime**. The buckets are < 30 days, 30-60 days, 60-90 days, 90-120 days, and > 120 days.
–apparent-size	Use apparent size (length) instead of disk capacity used. Typically, disk capacity used is larger than the file length, but it can be smaller for sparse files.
–env-equiv	Display environment var equivalents.
-f, –print-filenames	Display a line for each file, including the file's size and last modification time (**mtime**)
-h, –help	Display help message.
-n, –raw-uid	Display numeric UID and GID instead of attempting to translate to user/group names. Only applies when output format is "normal"; "perl" and "script" formats always print numeric UIDs and GIDs.

Option	Description
-o, –output-format=STRING	Output format, either “normal”, “perl”, or “script”.
-r, –raw-times	Display timestamps for -f in “raw” (seconds-since-epoch) format.
-s, –summary	Suppress per-directory space usage in output.
-t, –num-threads=UINT32	Number of worker threads.
-u, –print-user-info	Display space summary by username.
-z, –print-size-info	Display space summary by file length. The files are sorted into buckets by file length (file length is used for the buckets, regardless of the **–apparent-size** argument). The buckets are <1 KiB, 1 KiB to 10 KiB, 10 KiB to 100 KiB, 100 KiB to 1 MiB, 1 MiB to 10 MiB, 10 MiB to 100 MiB, 100 MiB to 1 GiB, 1 GiB to 10 GiB, 10 GiB to 100 GiB, 100 GiB to 1 TiB, and >1 TiB.

Example

The following is an example of normal output:

```
dir /panfs/REALM1: 180 files, 1295319688 KiB
user root: 180 files, 1295319688 KiB
size 1KiB: 0 files, 0 KiB
size 10KiB: 26 files, 416 KiB
size 100KiB: 0 files, 0 KiB
size 1MiB: 0 files, 0 KiB
size 10MiB: 0 files, 0 KiB
size 100MiB: 0 files, 0 KiB
size 1GiB: 0 files, 0 KiB
size 10GiB: 0 files, 0 KiB
size 100GiB: 0 files, 0 KiB
size 1TiB: 0 files, 0 KiB
size MAX: 154 files, 1295319272 KiB
age 30D: 180 files, 1295319688 KiB
age 60D: 0 files, 0 KiB
age 90D: 0 files, 0 KiB
age 120D: 0 files, 0 KiB
age MAX: 0 files, 0 KiB
```

The following is an example of script output:

```
D:/panfs/REALM1:files=180:kibibytes=1295319688
U:0:files=180:kibibytes=1295319688
S:1000:files=0:kibibytes=0
S:10000:files=26:kibibytes=416
S:100000:files=0:kibibytes=0
S:1000000:files=0:kibibytes=0
S:10000000:files=0:kibibytes=0
S:100000000:files=0:kibibytes=0
S:1000000000:files=0:kibibytes=0
S:10000000000:files=0:kibibytes=0
S:100000000000:files=0:kibibytes=0
S:1000000000000:files=0:kibibytes=0
S:MAX:files=154:kibibytes=1295319272
A:2592000:files=180:kibibytes=1295319688
A:5184000:files=0:kibibytes=0
A:7776000:files=0:kibibytes=0
A:10368000:files=0:kibibytes=0
A:MAX:files=0:kibibytes=0
```

The following is an example of perl output:

```
my %dirs;
my %users;
my %size_histo;
my %age_histo;
sub load_data() {
$dirs{'/panfs/REALM1'}->{'files'} = 180;
$dirs{'/panfs/REALM1'}->{'kibibytes'} = 1295319688;
$users{0}->{'files'} = 180;
$users{0}->{'kibibytes'} = 1295319688;
$size_histo{1000}->{'files'} = 0;
$size_histo{1000}->{'kibibytes'} = 0;
$size_histo{10000}->{'files'} = 26;
$size_histo{10000}->{'kibibytes'} = 416;
$size_histo{100000}->{'files'} = 0;
$size_histo{100000}->{'kibibytes'} = 0;
$size_histo{1000000}->{'files'} = 0;
$size_histo{1000000}->{'kibibytes'} = 0;
$size_histo{10000000}->{'files'} = 0;
$size_histo{10000000}->{'kibibytes'} = 0;
$size_histo{100000000}->{'files'} = 0;
$size_histo{100000000}->{'kibibytes'} = 0;
$size_histo{1000000000}->{'files'} = 0;
$size_histo{1000000000}->{'kibibytes'} = 0;
$size_histo{10000000000}->{'files'} = 0;
$size_histo{10000000000}->{'kibibytes'} = 0;
$size_histo{100000000000}->{'files'} = 0;
$size_histo{100000000000}->{'kibibytes'} = 0;
$size_histo{1000000000000}->{'files'} = 0;
$size_histo{1000000000000}->{'kibibytes'} = 0;
$size_histo{MAX}->{'files'} = 154;
$size_histo{MAX}->{'kibibytes'} = 1295319272;
$age_histo{2592000}->{'files'} = 180;
$age_histo{2592000}->{'kibibytes'} = 1295319688;
$age_histo{5184000}->{'files'} = 0;
$age_histo{5184000}->{'kibibytes'} = 0;
$age_histo{7776000}->{'files'} = 0;
$age_histo{7776000}->{'kibibytes'} = 0;
```

```
$age_histo{10368000}->{'files'} = 0;
$age_histo{10368000}->{'kibibytes'} = 0;
$age_histo{MAX}->{'files'} = 0;
$age_histo{MAX}->{'kibibytes'} = 0;
}
```

pan_filecount

Description

Counts the number of files and directories in VDURA volumes. It also provides an estimation mode that counts a portion of the **inode** namespace and uses this to project the total number of files in the volume.

CAVEATS

pan_filecount uses system-level interfaces that are not guaranteed to be compatible between PanFS OS releases. You must use **pan_filecount** from same version of the **panfs-apps** package as the version of the PanFS OS operating system running on your VDURA system. Using a different version of **pan_filecount** may fail to operate or produce incorrect results.

If a global DirectFlow mount is present in **/panfs**, **pan_filecount** will contact the default mounted realm. If there is no global DirectFlow mount or a different realm should be used, you must use the **–bindcache-rm-addr** option to specify the name and IP address of the realm to contact.

Because this utility bypasses normal filesystem permissions, it must be run as root.

Syntax

```
pan_filecount [OPTION]...
```

Options

You must specify either -l or -n to select the volume to count.

Mandatory arguments for long options are also mandatory for short options.

Option	Description
`-a, --countasc`	Count number of attribute-storing components on each Storage Node.
`--bindcache-rm-addr=STRING`	Specify the IP address of the realm manager for the realm to access.
`-d, --debug`	Print debugging info.
`-h, --help`	Display help message.
`-l, --volid=INTEGER`	Specify volume ID to list.

Option	Description
`-m, --maxpct=INTEGER`	Estimate the answer using INTEGER percent of the **inode** space.
`-n, --volname=STRING`	Specify volume name to list.
`-t, --counttype`	Count number of files and directories separately.
`-v, --verbose`	Print verbose status.
`-w, --warn`	Print warning if volume exceeds recommended maximum of 10 million objects.

Example

```
# pan_filecount -n /home -t
volume /home: 88177 objects total (86789 files, 1388 directories)
# pan_filecount -n /test02 -m 5
volume /test02: estimated (5% fraction): 1008682 objects total
```

pan_fm_config_get

Description

Service mode command to view internal details for an FM ID.

Syntax

```
pan_fm_config_get ?[options]? [fm-id]
```

The command must always include the fm-id (mgr_id) argument, which is the ID of the FM to query.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.

pan_fm_id_translate

Description

Service mode command to report internal mappings between ID domains.

Syntax

```
pan_fm_id_translate ?[options]? [id_type] [identity]
```

The command must always include the following arguments:

Argument	Description
`id_type (string)`	Target identity type.
`identity (identity)`	Identity to translate. An identity argument is a string. For instance, uid:0, gid:0, user_sid:S-1-5-10-55-67-34-22, group_sid:S-1-5-10-77-85-66-34, blade:0a0102030f0506, mgr:0x0400000000abcdef, u:nobody, g:nobody, or g:everyone.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-m, --fm-id=mgr_id`	ID of FM to query.
`-h, --help`	Show usage message.

pan_fm_vol_cache_examine

Description

Service mode command to examine the internal FM volume cache.

Syntax

```
pan_fm_vol_cache_examine ?[options]? [vol_id]
```

The command must always include the vol_id (uint32) argument, which is the volume ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-m, --fm-id=mgr_id`	ID of FM to query.
`-h, --help`	Show usage message.

pan_fsrc_get_progress

Description

Displays the current status of **FSRC** and progress details.

While the **volume** details command gives the basic FSRC progress percentage, use **pan_fsrc_get_progress** for additional details about FSRC progress for a given volume.

Syntax

```
pan_fsrc_get_progress [volume-name]
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-e, --eta`	Show ETAs.
`-h, --help`	Show usage message.

pan_get_master_mserver

Description

Displays the IP address of the Director running the **MConsole** service.

Under normal circumstances, **syscontrol realmmgr** shows which Director is running the **MConsole** service. In the rare occurrence of an **MConsole** outage, **pan_get_master_mserver** can be used to find which Director should be running it.

Syntax

```
pan_get_master_mserver [options]
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-q, --quiet`	Only print primary mserver IP.
`-h, --help`	Show usage message.

pan_iperf

Description

Pan_iperf is **iperf** renamed to avoid name conflicts.

Syntax

```
pan_iperf [-s|-c host] [OPTION]...
```

Options

Client/Server

Option	Description
`-f, --format [kmKM]`	Format to report: Kbits, Mbits, KBytes, or MBytes.
`-i, --interval INTEGER`	Number of seconds between periodic bandwidth reports.
`-l, --len INTEGER[KM]`	Length of buffer to read or write. Default is 8 KB. [KM] Indicates options that support a K or M suffix for kilo- or mega-.
`-m, --print_mss`	Print TCP maximum segment size (MTU - TCP/IP header).
`-o, --output <filename>`	Output the report or error message to this specified file.
`-p, --port NUMBER`	Use server port NUMBER to listen on/connect to.
`-q, --fault`	Client touches (faults) the page after writing.
`-u, --udp`	Use UDP rather than TCP.
`-w, --window INTEGER[KM]`	TCP window size (socket buffer size). [KM] Indicates options that support a K or M suffix for kilo- or mega-.

Option	Description
`-z, --randseed INTEGER`	Seed for random number generator.
`-B, --bind <HOST>`	Bind to <HOST>, an interface or multicast address.
`-C, --compatibility`	For use with older versions. Does not send extra messages.
`-M, --mss INTEGER`	Set TCP maximum segment size (MTU - 40 bytes).
`-N, --nodelay`	Set TCP no delay, disabling Nagle's Algorithm.
`-V, --IPv6Version`	Set the domain to IPv6.
`-X, --pause INTEGER[us]`	Maximum pause (in microseconds). Must be >2.
`-Y, --burstsize INTEGER[bytes]`	Burst size in bytes.

Server

Option	Description
`-D, --daemon`	Run the server as a daemon.
`-s, --server`	Run in server mode.
`-U, --single_udp`	Run in single threaded UDP mode.

Client

Option	Description
`-b, --bandwidth INTEGER[KM]`	For UDP, bandwidth to send at in bits/sec. Default is 1 Mbit/sec.; implies **-u**.
`-c, --client <HOST>`	Run in client mode, connecting to <HOST>.
`-d, --dualtest`	Simultaneously run a bidirectional test.

Option	Description
`-n, --num INTEGER[KM]`	Number of bytes to transmit (instead of **-t**). [KM] Indicates options that support a K or M suffix for kilo- or mega-.
`-r, --tradeoff`	Individually run a bidirectional test.
`-t, --time INTEGER`	Time in seconds to transmit for. Default is 10 seconds.
`-F, --fileinput <NAME>`	Input the data to be transmitted from a file.
`-I, --stdin`	Input the data to be transmitted from **stdin**.
`-L, --listenport INTEGER`	Port to receive bidirectional tests.
`-P, --parallel INTEGER`	Number of parallel client threads to run.
`-T, --ttl INTEGER`	Time-to-live for multicast. Default is 1.
`-Z, --linux-congestion <algo>`	Set TCP congestion control algorithm (Linux only).

Miscellaneous

Option	Description
`-h, --help`	Show usage message.
`-v, --version`	Print the version.
`-x, --reportexclude [CDMSV]`	Exclude C (connection), D (data), M (multicast), S (settings), and/or V (server) reports.
`-y, --reportstyle C`	Report as comma-separated values.

> **Note:** The TCP window size option can be set by the environment variable.

pan_journal_config_maxlevelstore

Description

Service mode command to query and update **journald MaxLevelStore** parameter.

Syntax

```
pan_journal_config_maxlevelstore [-h] [-s {emerg,alert,crit, err,warning,notice,info,debug}]
```

Options

Option	Description
`-h, --help`	Show usage message.
`-s {emerg,alert,crit,err,warning,notice,info,debug}, --set {emerg,alert,crit,err,warning,notice,info,debug}`	Set **maxlevelstore**.

pan_lm_lock_list

Description

Service mode command to view lock manager internals.

Syntax

```
pan_lm_lock_list ?[options]?
```

Options

Option	Description
`-a, --all-in-lm`	Print all locks in LM regardless of volume.
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-f, --file=string`	List locks only for this file.
`-F, --flags`	Print the flags.
`-h, --help`	Show usage message.
`-o, --obj-id=stor_id`	List locks only for this **obj-id**.
`-p, --paths`	Print paths to objects.
`-s, --short-format`	Print locks in abbreviated format.
`-t, --type=string`	Lock type (**all**, **byte**, **pending**, **share**).
`-V, --volume=string`	Use the LM responsible for this volume.

pan_net_enclosure_nettest

Description

Service mode command for performing basic network connectivity and throughput tests between nodes and between nodes and external hosts.

Network connectivity tests include small and medium pings. Optional tests include a large **ping** test, **ssh**, and **netperf** throughput tests if **netperf** is enabled on the realm. The large ping test runs by default only if the system MTU is 9000. Test results are output for each node tested and summarized by enclosure.

If **pan_net_enclosure_nettest** is run with any of the following options, the results are summarized by enclosure in the summary log:

```
--all --encl --directors --storage --dms --sms
```

Results for each node are shown in the detail log.

If **pan_net_enclosure_nettest** is run either with the **–repset** option or with IP addresses, the results for those nodes are shown in the summary log.

Pan_net_enclosure_nettest prints the path and names of both logs to the screen when it completes.

The small (56 bytes) and medium (1472 bytes) **ping** tests are always run and may not be skipped. The large (8972 bytes) **ping** test is run if the realm is configured for MTU 9000. The large ping test may optionally be skipped via the **–nolargeping** option.

The **netperf** test requires that **netperf** be enabled on the realm. To enable **netperf** on the realm:

```
[pancli_service] netperf enable
```

Netperf may optionally be disabled once **pan_net_enclosure_nettest** has completed with the following:

```
[pancli_service] netperf disable
```

Netserver must be installed on any hosts whose IP addresses are supplied on the command line if those hosts are not nodes in the realm. **Netserver** must be configured to listen on port 4005.

Netperf results for the local enclosure will be unusually high. The test is executed across the enclosure backplane.

Syntax

```
pan_net_enclosure_nettest [options] ?[space separated IPs]?
```

Options

Option	Description
`--help`	Display this usage message.
`--examples`	Display explanation of **pan_net_enclosure_nettest** use and example command options.
`--all`	Test all nodes in the realm.
`--encl [name]`	Test all nodes in the specified enclosure. Multiple **–encl** arguments may be used.
`--repset`	Test the repset member Directors.
`--directors`	Test all Director Nodes in the realm.
`--storage`	Test all Storage Nodes in the realm.
`--dms`	Test all Director modules in the realm.
`--sms`	Test all Storage modules in the realm.
`--set`	Test all Storage Nodes in the specified StorageSet. Use single quote marks to enclose a StorageSet name that contains a space, as in ‘Set 1’.
`--offline`	Include any nodes that are not marked online in the tests. By default, non-replication set (repset) member nodes that are not online are excluded. Hosts specified on the command line are assumed to be online.

Option	Description
`--quick`	Run shorter **ping** and **netperf** tests. The number of pings is reduced from five to one for each of the small, medium, and large **ping** tests. The **netperf** test is reduced from ten seconds to one second.
`--nolargeping`	Skip the large **ping** test.
`--nossh`	Skip the **ssh** test. This test is automatically skipped for external hosts.
`--noperf`	Skips the **netperf** test.
`--noquery`	Do not query the realm to build a list of nodes to test. A space-separated list of node IP addresses to test must be provided. If **–noquery** is used, all options that specify types of nodes to test are disregarded. The large **ping** and **netperf** tests are run without querying the realm to see if **netperf** is enabled or MTU is set to 9000. The **ssh** test is not run. This option is useful if the replication set (repset) is down.

Example

Run all tests to all nodes in the realm:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --all
```

If **netperf** is not enabled on the realm, the command stops and asks you to enable it. Enter **netperf enable** at the **pancli_service** prompt, or add the **–noperf** option.

The output below is an example of the console output and summary from the command shown above, run on a realm with three Director Nodes, each in its own enclosure. The ping tests pass if the number of packets sent are received. Large **ping** shows N/A because the system MTU is not 9000. The ssh test

passes if the remote command uptime returns in five seconds. The netperf results are 10^6bits/sec. The netperf results for the enclosures are presented as the lowest, highest, and average result for a Storage (S:) or Director (D:) Node in the enclosure.

```
pan_net_enclosure_nettest is using options: --all

pan_net_enclosure_nettest is running on node 10.128.1.0
   Host              Ping Small  Ping Med    Ping Large  SSH     Netperf
   10.128.1.0        OK          OK          N/A         OK      30849.00
   10.128.1.1        OK          OK          N/A         OK      7902.04
   10.128.1.2        OK          OK          N/A         OK      7888.72

By Enclosure:
   Enclosure         Ping Small  Ping Med    Ping Large  SSH     Netperf (Low/High/Avg)
   VCH-1*            1/1         1/1         N/A         1/1     D: 30849.00/30849.00/30849.00
   VCH-4             1/1         1/1         N/A         1/1     D: 7888.72/7888.72/7888.72
   VCH-8             1/1         1/1         N/A         1/1     D: 7902.04/7902.04/7902.04

*Test results to enclosure VCH-1 may be misleadingly high since
 pan_net_enclosure_nettest was run on 10.128.1.0, which is in enclosure VCH-1.

Tue Oct 27 00:25:53 2026 GMT: Elapsed time: 0 hours, 0 minutes, 56 seconds.
pan_net_enclosure_nettest has completed.
Results are logged to /tmp/pan_net_enclosure_nettest-1793060697-summary.log
Detailed results per node are logged to /tmp/pan_net_enclosure_nettest-1793060697-details.log
```

Test results to enclosure VCH-1 may be misleadingly high since pan_net_enclosure_nettest was run on 10.128.1.0, which is in enclosure VCH-1. Repeat pan_net_enclosure_nettest from a node in another enclosure to see a test through the network for 10.128.1.0.

Run all tests to the replication set (repset) Directors:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --repset
```

The output below is an example from the command, run on the same realm. The node running the command is itself a repset member, so its own results are marked with an asterisk and may be misleadingly high. The results to the repset members are shown both per node and by enclosure.

```
pan_net_enclosure_nettest is using options: --repset

pan_net_enclosure_nettest is running on node 10.128.1.0
   Host              Ping Small  Ping Med    Ping Large  SSH     Netperf
   10.128.1.0        OK          OK          N/A         OK      30157.35
   10.128.1.1        OK          OK          N/A         OK      7917.49
   10.128.1.2        OK          OK          N/A         OK      7912.62

Repset:
   Host              Ping Small  Ping Med    Ping Large  SSH     Netperf
   10.128.1.0*       OK          OK          N/A         OK      30157.35
   10.128.1.2        OK          OK          N/A         OK      7912.62
   10.128.1.1        OK          OK          N/A         OK      7917.49

By Enclosure:
   Enclosure         Ping Small  Ping Med    Ping Large  SSH     Netperf (Low/High/Avg)
   VCH-1*            1/1         1/1         N/A         1/1     D: 30157.35/30157.35/30157.35
   VCH-4             1/1         1/1         N/A         1/1     D: 7912.62/7912.62/7912.62
   VCH-8             1/1         1/1         N/A         1/1     D: 7917.49/7917.49/7917.49

*Test results to node 10.128.1.0 may be misleadingly high since pan_net_enclosure_nettest was
run on 10.128.1.0.
*Test results to enclosure VCH-1 may be misleadingly high since
 pan_net_enclosure_nettest was run on 10.128.1.0, which is in enclosure VCH-1.

Tue Oct 27 00:28:25 2026 GMT: Elapsed time: 0 hours, 0 minutes, 55 seconds.
pan_net_enclosure_nettest has completed.
Results are logged to /tmp/pan_net_enclosure_nettest-1793060850-summary.log
Detailed results per node are logged to /tmp/pan_net_enclosure_nettest-1793060850-details.log
```

Run all tests to all Storage Nodes in the realm, in addition to the client hosts 10.70.15.105 and 10.70.15.127:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --storage 10.70.15.105 10.70.15.127
```

Run all tests to all nodes in enclosures VCH-1 and VCH-4 without the netperf test:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --encl VCH-1 --encl VCH-4 --noperf
```

Run all tests to all Storage Nodes in StorageSet Set 2:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --set 'Set 2'
```

Run a quick connectivity test to find out which enclosures are up:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest --all --nossh --noperf --nolargeping --quick
```

Test only node 10.65.23.106. 10.65.23.106 may either be a node in the realm or an external host. No ssh test will be run:

```
[pancli] support servicemode yes
[pancli_service] pan_net_enclosure_nettest 10.65.23.106
```

pan_nettest

Description

Runs **iperf** TCP throughput tests to characterize the network between a set of N clients and a set of M servers.

The **Network Characterization** test runs simultaneous NxM pairwise throughput tests to estimate the full bisection bandwidth between the set of clients and servers. This test is useful to detect if any network element (switch or router) or the network design is incapable of providing the expected bisection bandwidth.

The **Path Characterization** test runs 1-1 pairwise throughput tests to estimate the bandwidth between any pair of client and server. This test is useful to detect if some network link is operating at 100 or 10 Mbps instead of an expected 1000 Mbps.

The **Topology Discovery** test runs 1-1 pairwise traceroute tests to discover the IP topology between any pair of client and server. This test detects if some clients are more hops away or on higher latency paths than other clients, which could result in uneven performance.

Syntax

```
pan_nettest [OPTION]...
```

Options

Option	Description
`--bad-ip`	Skip bad IP addresses.
`-c,--clients`	Space-separated list of names or IP address ranges of clients. IP address ranges are of the form a.b.c.d-l.m.n.o, where a.b.c.d is the starting IP address and l.m.n.o is the ending IP address for the IP address range.
`-C,--expand-client`	Expand client IP address into Storage Node IP addresses.
`-d,--dir`	Name of directory where results will be stored.

Option	Description
`--defaults`	Show program defaults.
`--explain-tests`	Describe the tests that can be run.
`-g,--gui`	Use GUI to view results.
`-h,--help`	Display help message.
`--keep`	Do not cleanup remote scripts and executables. This avoids copying scripts and executables when running the test several times with the same set of servers and clients.
`--nctime`	How many seconds to run Network Characterization test.
`--num-clients`	Number of clients to include in the Path Characterization tests.
`--num-packets`	Number of packets to send.
`-s,--servers`	Space-separated list of names or IP address ranges of servers. IP address ranges are of the form a.b.c.d-l.m.n.o, where a.b.c.d is the starting IP address and l.m.n.o is the ending IP address for the IP address range.
`--num-servers`	Number of servers to include in the Path Characterization tests.
`--num-tests`	Number of times to repeat the tests.
`--num-threads`	Number of threads to use in parallel.
`--pctime`	How many seconds to run Path Characterization test.
`--remove`	Execute an **rm -f** before copying executables to servers/clients. This is used to remove older versions of scripts and executables left over from earlier runs with the **–keep** option.

Option	Description
`--show-notes`	Display notes for help with interpreting the results.
`--show-platforms`	Show supported platforms.
`--size-packets`	Size of packets to send.
`--straggler`	Time to wait in seconds for stragglers to finish in the Network Characterization test.
`-t,--tests`	Specify tests to run (see **–explain-tests**). Must be one of **pc** (for Path Characterization), **nc** (for Network Characterization), **td** (for Topology Discovery), or **all** (for all tests).
`-u,--user-id`	Space separated list of **user-ids** to use instead of **root**.
`-S,--expand-server`	Expand server IP address into Storage Node IP addresses.

pan_op_log_tool

Description

Service mode command to format File Manager, Storage Manager, and Lock Manager transaction logs.

Syntax

```
pan_op_log_tool ?[options]? [log_file]
```

The command must always include the log_file [string] argument, which is the path to the File Manager, Storage Manager, or Lock Manager op-log.

Options

Option	Description
`-c, --create`	Create an empty log file.
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.
`-s, --scrub`	Scrub the log by tossing out entries that are not compatible with AS 3.x or earlier.

pan_rm_blobdb_util

Description

Service mode command to view and edit the RM **blobdb** database.

Syntax

```
pan_rm_blobdb_util ?[options]?
```

Options

Option	Description
`--defaults`	Show defaults.
`-d, --delval=string`	Values to delete (format: **type:name**).
`--env-equiv`	Show environment **var** equivalents.
`-g, --getval=string`	Values to get (format: **type:name**).
`-h, --help`	Show usage message.
`-s, --setval=string`	Values to set (format **type:name:val**).

pan_rm_brick_exhume

Description

Service mode command to remove a reconstruction record and optionally return a buried brick to service.

Syntax

```
pan_rm_brick_exhume data_serial
```

The command must always include the data_serial (string) argument, which is the data serial number of the brick to exhume.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.

pan_rm_configdb_util

Description

Service mode command to view and edit the RM **configdb** database.

Syntax

```
pan_rm_configdb_util ?[options]?
```

Options

Option	Description
`--defaults`	Show defaults.
`--delfile`	**Mcv** file of values to delete.
`-d, --delval`	Values to delete (format: **name**).
`--env-equiv`	Show environment **var** equivalents.
`-i, --getfile`	**Mcv** file of values to get.
`-g, --getval`	Values to get (format: **type:name**).
`-h, --help`	Show usage message.
`-f, --setfile`	**Mcv** file of values to set.
`-s, --setval`	Values to set (format **type:name:val**).

pan_rm_get_repset

Description

Service mode command to view replication set (repset) details specific to the RM.

Syntax

```
pan_rm_get_repset ?[options]?
```

Options

Option	Description
`-d, --dbinit-format`	Print results in dbinit format.
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-n, --host=string`	IP address of president.
`-p, --port=uint16`	Port number of president.

pan_rm_move_backup

Description

Service mode command to change the backup location of a volume service.

> **Note:** This action can also be performed using the **volservice move** command.

Syntax

```
pan_rm_move_backup ?[options]? [id] [serialno]
```

The command must always include the following arguments:

Argument	Description
`id (mgr_id)`	Mgr_id of the service to move (FM,SM,FSRC all move together).
`serialno (string)`	Dataserial number of the brick to move to.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.

pan_rm_move_service

Description

Service mode command to move a service to another node.

Syntax

```
pan_rm_move_service ?[options]? [id] [serialno]
```

The command must always include the following arguments:

Argument	Description
`id (mgr_id)`	Mgr_id of the service to move (FM,SM,FSRC all move together).
`serialno (string)`	Dataserial number of the node to move to.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`--force-dirty-start=on/off`	Force a dirty start after the move.
`-h, --help`	Show usage message.

pan_rm_ptp_tool

Description

Service mode command to update the live system PTP database.

Syntax

```
pan_rm_ptp_tool ?[options]? [ptp-dir]
```

The command must always include the ptp-dir (string) argument, which is the ptp database directory.

Options

Option	Description
`--coup-mgr-id=mgr_id`	Modify the replication set (repset) so that this **mgr_id** is the only RM.
`--defaults`	Show defaults.
`--detect-pending`	Detect if there are pending ptp decrees; exit 0 if there are no pending decrees or != 0 if there are.
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.
`-p, --print`	Print ptp databases.
`-l, --print-log`	Print abbreviated decree log.
`--purge-decree`	Purge db of decree logs.
`--purge-pending`	Purge db of pending (unpassed) decrees.
`--set-thb=secs:nsecs`	Modify the ptp heartbeat interval in the dbs.
`--set-tsu=secs:nsecs`	Modify the ptp election setup time interval in the dbs.

pan_rm_rehome_realm_root

Description

Service mode command to move the root volume to a new StorageSet.

This command moves the root volume to a new StorageSet. It can also be used to move the root volume to a new layout. This service mode command should only be used under the direction of VDURA Technical Support.

> **Note:** This action can also be performed using the **volume moveroot** command.

Syntax

```
pan_rm_rehome_realm_root ?options? [storageset ID]
```

The StorageSet ID can be found using the storageset list id PanCLI command.

The command must always include the storageset ID (uint32) argument.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.
`--layout-type=string`	Layout type.

pan_rm_volume_info

Description

Service mode command to list internal volume details specific to the RM database.

Syntax

```
pan_rm_volume_info ?[options]? (-n [volume-name]|-i [volume-id])
```

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.
`-i, --id=uint32`	ID of the volume to list.
`-n, --name=string`	Name of the volume to list.
`-v, --verbose`	Display list verbosely.

pan_sam_util_verify_parity

Description

Service mode command to check or correct parity in a file.

> **Note:** This command can also be performed using the **raidutil** command.

Syntax

```
pan_sam_util_verify_parity [options]
```

Options

Option	Description
`-d, --comp_dev_id=mgr_id`	Component Device ID 0
`-1, --comp_dev_id_1=mgr_id`	Component Device ID 1
`-2, --comp_dev_id_2=mgr_id`	Component Device ID 2
`-c, --correction_mode=uint32`	Set correction mode (0: no correction; 1: correct specified components; 2: correct parity; 3: auto-correct).
`--defaults`	Show defaults.
`--env-equiv`	Show environment var equivalents.
`-f, --filename=string`	Filename
`-h, --help`	Show usage message.
`-l, --length=uint64`	Length
`-i, --list_id`	List component object IDs when parity is found to be bad.
`--log_path=string`	Parity log file path; default is **cwd**.

Option	Description
`-q, --no_progress_display`	Do not print progress spinner during parity scrub.
`-o, --obj-id=stor_id`	Object ID
`-s, --offset=uint64`	Offset
`-r, --use_rvp=int32`	Verify parity by calling sam RVP read.
`-S, --use_sam=int32`	Verify parity by calling sam instead of VPOD.
`-v, --verbose`	Display list verbosely.

pan_sitesync

Description

This is the main script for starting a data replication operation, which can copy a source tree to a destination, using multiple clients in parallel.

Pan_sitesync is the manager portion of SiteSync, the utility which replicates an existing directory tree in a new location. It does this by running **rsync** in parallel using multiple client hosts.

The primary use of SiteSync is to replicate files and directories between a legacy file system and a VDURA storage file system. The **pan_sitesync** command is used to start and manage a SiteSync replication operation.

> **Note: Rsync** 2.6.8 or later is required on all host clients used for replication.

PARAMETERS

src-path: Fully qualified source path to transfer. The **src-path** must be visible on all the worker clients specified in the **client-list** file.

dst-dir: Fully qualified target directory to replicate to. The **dst-dir** must be visible on all the worker clients specified in the **client-list** file.

PRINCIPLE OF OPERATION

Pan_sitesync determines the directories and/or files to be replicated and sends instructions to worker clients (running **pan_sitesync_worker**) to handle the copying using **rsync**. Unless given an explicit list of items to replicate (with **–files-from** option), **pan_sitesync** begins with the specified source directory and traverses the tree in a depth-first manner. For each directory visited, the manager sends an instruction to a worker client to replicate the directory and its files. The manager traverses the entire source directory tree structure, and the tree content is replicated by the workers.

During a replication operation, **pan_sitesync** periodically saves state and progress in a special temporary directory. By default, this directory is located in the root of the destination; use the **–temp-dir** option to specify an alternate location. If **pan_sitesync** is interrupted, this saved progress may be used to resume the operation where it left off. If pan_sitesync is interrupted with **SIGKILL** (i.e. **kill -9**), then no termination handling can occur and latest progress may not be saved. To insure the latest progress is saved when interrupting **pan_sitesync**, use **Ctrl-C** from the command line (or **kill -SIGINT**).

Pan_sitesync_worker is the **pan_sitesync** worker application and is usually started automatically by the manager. It can also be started manually; see the man page for **pan_sitesync_worker**.

By default, **pan_sitesync** will attempt to start worker processes directly on the same host that is running **pan_sitesync**, using the current user ID and default executables for **pan_sitesync_worker** and **rsync**. To use other hosts, user IDs, or executable paths, use the **–client-list** option. See below for more details about the format of a **client-list** file.

The **panfs-apps** RPM must be installed on all of the worker clients specified in the **client-list** file.

CLIENT LIST FILE

At startup time, **pan_sitesync** may be given a file containing a list of available worker clients. Each line in the **client-list** file has the following form:

```
IP_or_hostname userID <pw> <pan_sitesync_worker> <rsync>
```

Where:

- **userID** is the ID for SSH to use to access the client machine.
- pw is the optional password in open text.
- pan_sitesync_worker is the optional worker executable.
- rsync is the optional **rsync** executable.

Each of the optional items is enclosed in angle brackets, and the angle brackets must always be present. Paths to executables may be included as well (either absolute or relative to the **userID** home directory). For any executable not specified, the default executable names are used, and it is assumed they will be found in the default search path for the SSH **userid**.

Notes:

- A single dot can be used for the **userID** to represent the current user at the time **pan_sitesync** is started (see Example(s) below).
- The **userID** (or current user if a dot is specified) must have at least read access to the source directory, and read+write+execute access to the destination directory.
- Any valid host can be used as a worker and listed in this file, including the host on which the user runs **pan_sitesync** (the SiteSync manager) itself.
- If using key-based SSH, then the password field should be empty.

- Comment lines must begin with a **#** in the first column. Sample client list file contents:

```
# ssh to 169.132.32.10 using the current userid and key-based access, and run
pan_sitesync_worker and rsync from default PATH
169.132.32.10 . <> <> <>
# ssh to the host as the current user, with key-based access,
# and run pan_sitesync_worker from root PATH and use specific rsync
169.132.32.10 . <> <> </test/bin/rsync>
# ssh to the host as root with key-based access
# and run pan_sitesync_worker from root PATH and use specific rsync
169.132.32.10 root <> <> </test/bin/rsync>
# ssh to the host as user 'rnwp' with password 'foopass'
# and use the specified pan_sitesync_worker executable
169.132.32.4 rnwp <foopass></global/utilities/pan_sitesync_worker> <>
```

RETURN VALUES

- 0: Success.
- 2: Improper options or argument specified.
- 255: General failure (see log for more info).

ERRORS

Pan_sitesync reports warnings, errors, and fatal errors to the log. Fatal errors are also always sent to stdout. Files that failed to replicate are reported at the conclusion of replication, and may be retried using the –retry option.

Syntax

```
pan_sitesync [OPTION]... src-path... dst-dir
```

Options

Option	Description
`--addr=<IP>`	Use the specified IP address for worker communications. This option allows selection of the host address (and therefore, the attached subnet) to use for exchanging messages with workers. Generally, this is only needed with multi-homed hosts, when there is no routing between attached subnets. By default, pan_sitesync will use the first routable interface address.
`--client-list=<FILE>`	Specify a file with a list of clients to use for workers. If not specified, the localhost will be used as the client to run workers. The total number of workers will be the number of clients specified, times the value of the workers-per-client option. Changes to this file that occur once processing has started are ignored until the next restart or resume.
`--delete`	Delete files and directories that exist only in the destination.
`--eta`	Before starting replication, perform a dry-run to gather statistics, compute estimates for completion time, and include ETA in progress messages. Implies –progress.
`--exclude-from=<FILE>`	Specify a file that contains rsync exclude patterns (one per line). Blank lines in the file and lines starting with ";" or "#" are ignored. FILE cannot be "-" (i.e. the list cannot be read from standard input as with regular rsync). NOTE: Unlike rsync, SiteSync does not support excluding subdirectories with an exclude file. Only files may be excluded by entries in an exclude file.

Option	Description
`--files-from=<FILE>`	Specify a file containing a list of files and directories to be replicated. Unlike include and exclude files, this file cannot contain wildcards. Each directory or file listed must be specified with a path relative to the source directory. When using this option, directory recursion is disabled.
`-h,--help`	Show usage message.
`-I,--ignore-times`	Do not skip files that match size and time. This forces all files to be updated regardless of the last-changed time or file size.
`--include-from=<FILE>`	Specify a file that contains rsync include patterns (one per line). Blank lines in the file and lines starting with ";" or "#" are ignored. FILE cannot be "-" (i.e. the list cannot be read from standard input as with regular rsync).
`--log-date`	Tell SiteSync to add the current date to log timestamps (dd-MMM-yy HH:MM:SS). Default is HH:MM:SS.
`--log-file=<FILE>`	Specify the path and name of the SiteSync log file to create. Default is "-", which writes log output to stdout.
`--log-msec`	Tell SiteSync to extend log timestamps to millisecond accuracy (HH:MM:SS.mmm). Default is HH:MM:SS.
`--modify-window=<NUM>`	Compare last-modification times with reduced accuracy; allows a difference of up to NUM seconds to be considered equal.
`-n,--dry-run`	Do not copy any directories or files. Instead, determine which files to copy and obtain statistics about the potential replication.

Option	Description
`--no-itemize`	Instruct SiteSync to skip generating ITEM log messages. Normally when any verbose options are specified, the SiteSync manager and workers will produce a log message for every file and directory examined or updated. Use of this option suppresses these itemized messages, while still allowing other verbose log messages to appear. This can greatly reduce log size and processing overhead when these individual messages are not needed.
`--no-sparse`	Instruct SiteSync to omit the -S parameter to rsync; i.e. do not preserve source file sparseness on the destination copy.
`--no-xattrs`	Instruct SiteSync to omit the -X parameter to rsync; i.e. do not copy or update extended attributes on the destination.
`--port-m=<NUM>`	Set the port number to use for the manager's XML-RPC service. Default is 9980.
`--port-w=<NUM>`	Set the base port number to use for the workers' XML-RPC service. Default is dynamic.
`--progress`	Instruct pan_sitesync to periodically show a progress message to stdout.
`-r,--resume`	Instruct pan_sitesync to resume a previous (incomplete or interrupted) operation. If –resume is specified, then a valid temporary directory must exist, and contain the state from a previous operation. Similarly, if –resume is not given, the temporary directory must not already exist.

Option	Description
`--retry`	Restart a replication which completed with one or more file/directory errors. Only the failed items will be retried, and no other options specifying which files to replicate (i.e. include/exclude, files-from, etc.) are allowed.
`--save-period=<MINUTES>`	Specify how long to wait in minutes between periodic updates of the saved replication state in the temporary directory. Increasing this value may cause more reprocessing when resuming SiteSync. Default is 5.
`--size-only`	Skip files that match in size, regardless of whether the source and destination have matching last-modified times.
`--stats-only`	Display stats of an interrupted operation and exit without resuming.
`--sync-size=<SIZE>`	Set the minimum size for processing a single directory with multiple clients. Larger values allow for greater use of single-client processing (which is more efficient) but may reduce the overall performance due to decreased use of parallel processing. Default is 5 G.

Option	Description
`--sync-tax=<SIZE>`	Set the per-file penalty in bytes to use in computing the effective workload for each directory. Larger values will increase multi-client directory processing, which is less efficient but helps distribute workloads more evenly. Default is 1 M. NOTE: The value of sync-size divided by the value of sync-tax sets the maximum number of files to be handled at one time by a single worker. If this ratio is set too high, excessive delays and/or errors during processing may occur.
`--temp-dir`	Set the temporary directory to use for saving state. This directory allows the replication to be interrupted and resumed later without starting over or losing all progress or performance statistics. Default is <dest_dir>/.pps_temp.
`--temp-dir-del=<WHEN>`	Specify whether to delete the temporary directory when pan_sitesync exits. Choices are: - always: Always delete the temporary directory. - never: Never delete the temporary directory. - success: Delete upon successful completion only. Default is success. NOTE: When using –temp-dir-del=always and pan_sitesync is interrupted, the temp dir will only be removed if pan_sitesync is interrupted with Ctrl-C from the command line (or kill -SIGINT). It will be left intact if pan_sitesync is killed with SIGKILL (i.e. kill -9).

Option	Description
`--test-workers`	Do not perform replication. Instead, test the ability of the manager to start the workers, and test each worker's access to the file system directories. NOTE: All other options and arguments are ignored, except –client-list.
`--timeout=<SECONDS>`	Set the rsync I/O timeout. NOTE: To prevent worker processes from running indefinitely waiting for blocked I/O, the timeout default is not zero (infinite) as it usually is for rsync. Therefore, on congested file systems or networks where I/O timeouts may be more likely, this option can be used to raise the timeout duration and reduce or eliminate rsync timeout errors. Default is 300.
`-u,--update`	Skip files that are newer on the receiver.
`--use-rsync=<FILE>`	Instruct manager to use the specified executable as rsync. Default is rsync (in standard search path).
`-v,--verbose`	Increase the amount of information output to the log. By default, SiteSync outputs only warnings and errors. A single -v will add information about what files are being transferred and a brief summary at the end. Two -v options (or -vv) will add additional information about SiteSync operation. Additional -v options are for debugging SiteSync. Note that a single -v will enable individual messages for each file handled. This can add a great deal of log content and processing overhead. Use the –no-itemize option to suppress such messages.
`--version`	Display current version of SiteSync and exit. All other arguments are ignored.

Option	Description
`-w,--workers-per-client=<NUM>`	Specify the number of worker processes to start on each worker client.

Example

Replicate the directory tree within /vol1/src to /vol3/archive3 using the local host as the only worker:

```
pan_sitesync /vol1/src /vol3/archive3
```

If files are found in /vol3/archive3 that do not exist in /vol1/src, they are left untouched.

> **Note:** Pan_sitesync differs from rsync in terms of handling trailing slashes on the source path. Unlike rsync, which will copy the last named directory if a trailing slash is omitted, pan_sitesync will always, regardless of trailing slash or not, copy the contents and not the last named directory from the source into the destination.

Replicate the directory tree from /vol1/src to /vol3/archive3 and display the name of each file replicated. In this case, if any files or directories exist in /vol3/archive3 that are not found in /vol1/src, they are deleted:

```
pan_sitesync -v --delete /vol1/src /vol3/archive3
```

Replicate the directory tree from /vol1/src to /vol3/archive3 using the workers that are defined in the file /usr/foo/sync_clients:

```
pan_sitesync --client-list=/usr/foo/sync_clients /vol1/src /vol3/archive3
```

Replicate the directory tree from /vol1/src to /vol3/archive3 and dump all log output (including periodic progress messages) to /usr/foo/sync.log:

```
pan_sitesync --progress --log-file=/usr/foo/sync.log /vol1/src /vol3/archive3
```

Tell manager to bind to http://123.43.5.6 for worker communications while replicating the directory tree from /vol1/src to /vol3/archive3. If the address does not exist on the host, the operation will immediately fail with an error message:

```
pan_sitesync --addr=123.43.5.6 /vol1/src /vol3/archive3
```

pan_sm_attr_check

Description

Service mode command to repair storage attributes of an object.

Syntax

```
pan_sm_attr_check [obj-id]
```

pan_sm_cap_query

Description

Service mode command to display caps that are outstanding in Storage Manager.

Syntax

```
pan_sm_cap_query [options] sm-id
```

Options

Option	Description
`--volume -v`	Volume name
`--group -g`	Group ID
`--count -c`	Output the number of caps outstanding.
`--scrubs -s`	Output the list of caps with scrub operations outstanding, the scrub ranges, and object IDs.

pan_sm_cleanup_bg_scrubber_dir

Description

Service mode command to clean up **bg_scrubber** files.

When the system encounters an object with a parity inconsistency, it populates the background scrubber directory **/var/log/bg_scrubber** with files that are useful for troubleshooting. Although there is an automated script that clears these files when available space gets low, there may be some occasions where it is useful to clear only files older than the given file on demand rather than waiting for automated cleanup.

This utility removes the given file and all other files older than the given file from the **/var/log/bg_scrubber** directory. The file given must exist relative to **/var/log/bg_scrubber**.

Because this utility removes files that can be used for troubleshooting, it should only be used under the direction of VDURA Technical Support.

Syntax

```
pan_sm_cleanup_bg_scrubber_dir [file]
```

pan_sm_print_map

Description

Prints the layout of an object across Storage Nodes.

Syntax

```
pan_sm_print_map ?[options]? [obj-id]
```

The command must always include the obj-id (stor_id) argument, which is the storage manager ID.

Options

Option	Description
`--defaults`	Show defaults.
`--env-equiv`	Show environment **var** equivalents.
`-h, --help`	Show usage message.

pan_sm_repair_query

Description

Service mode command to print the repair data from a specified Storage Manager.

Syntax

```
pan_sm_repair_query ?[options]? [sm-id]
```

Options

Option	Description
`--count -c`	Output the count of repairs in the repair hash.
`--list -l`	List all objects with repairs outstanding and repair types (maximum 100).
`--attempts -a n`	List objects which have repairs that have been attempted n times or greater.

pan_switch_software_update

Description

Service mode command to upgrade switch firmware.

This command upgrades a 10 Gigabit Ethernet switch. In the rare situation where switch firmware was expected to, but did not upgrade during a system-wide software upgrade, this command can be used to upgrade an individual switch. This is a service mode command and should only be used under the direction of company service personnel to ensure safe usage.

The filename given must be a glob pattern matching an absolute path to a tarred, gzipped switch firmware file. If the **AUTO** option is given instead of a filename, the default firmware file will be used.

In a multi-enclosure system, this command should be run from the enclosure master node on each enclosure whose switch needs to be upgraded.

> **Note:** It takes two to three minutes for firmware to download to the switch. There should be no interruption of switch service during this time. Once the download completes, the upgrade activity begins, causing the switch to be down for 45 to 60 seconds.

To downgrade the switch firmware to a previous version, contact a VDURA service representative.

Syntax

```
pan_switch_software_update ?[options]? AUTO|[upgrade-file]
```

Options

Option	Description
`-l`	Display install log if **-q** option was specified.
`-q`	Run software upgrade quietly.
`-e (NET1|NET2)`	Select a port (both ports are upgraded by default).

pan_tcp_capture_client

Description

Client utility to capture **tcpdump** data and related network or NFS statistics and configuration data.

- Requires **tcpdump** version 4.0.0 or higher unless the **–nodump** option is used.
- Supported platforms: FreeBSD and Linux.

Tcpdump allows only the root user to monitor a network interface. Since **pan_tcp_capture_client** runs **tcpdump**, **pan_tcp_capture_client** must be run by the root user. **Pan_tcp_capture_client** may be run by a non-root user with the **–nodump** option.

Syntax

```
pan_tcp_capture_client [OPTIONS]...
```

Options

Option	Description
`--cycle=INTEGER`	Set how often **pan_tcp_capture_client** will collect relevant statistics and verify that the tcpdump process is still running, in seconds. If used, the value must be between 10 and 300 seconds, and less than or equal to the value set by **–seconds**. The default is 30 seconds.
`--filter=STRING/ INTEGER`	Set a **tcpdump** filter. If the value is an IP address, sets a **host [IP address]** filter. If the value is an integer that is a VDURA port, sets a **port [port number]** filter. **Pan_tcp_capture_client** supports port filters 2049, 3095, 3251, 3260, and 10622. No other **tcpdump** filters are supported. Multiple **–filter** options may be used.
`--help`	Display help message.

Option	Description
`--hours=INTEGER`	If **–roll** is used, set the total number of hours that **pan_tcp_capture_client** should run. Fractional values are allowed. If used, the value must be between 0.25 and 24 hours. Value is not set by default.
`--interface=STRING`	Set the network interface that **tcpdump** will use to capture packets. The default is for **pan_tcp_capture_client** to search for the active interface.
`--minutes=INTEGER`	Sets the maximum number of minutes that **pan_tcp_capture_client** should run. If **–roll** is used, sets the number of minutes that **tcpdump** will run before overwriting the first capture file. If used, the value must be between 2 and 60 minutes. The default is 15 minutes.
`--nfs`	Highlight NFS TCP connections. Off by default.
`--nodump`	Omit the **tcpdump** collection. **Tcpdump** data is collected by default.
`--pssleep=INTEGER`	Set the number of seconds between the action of starting or stopping a process and checking the result. If used, the value must be between 2 and 10 seconds. The default is 2 seconds.
`--roll`	Set **tcpdump** to run for the number of hours specified by the **–hours limit**. When the **–minutes** limit is reached, **pan_tcp_capture_client** will begin to delete the oldest capture files, so that the most recent data for the value of **–minutes** is kept. If used, **–hours** is also required. This option is not set by default.

Option	Description
`--seconds=INTEGER`	Set the number of seconds for **tcpdump** to collect data into an output capture file. After this number of seconds has elapsed, **tcpdump** will create a new capture file. **Tcpdump** embeds the current epoch time into the name of each capture file. If used, the value must be between 10 and 60 seconds. The default is 60 seconds.
`--size=INTEGER`	Set the approximate maximum total size of the output files in GiB. **Pan_tcp_capture_client** checks the total size at the interval specified with the **–cycle** option (30 seconds by default), so if the output files are rapidly growing, the size limit may be exceeded. The default is 30 GiB. If used, the value must be between 1 and 50 GiB.
`--snap=INTEGER`	Sets the **tcpdump snaplen** value. This is the number of bytes from each packet that tcpdump will capture. The **pan_tcp_capture_client** default is 384. If used, the value must be between 64 and 65535.
`--unlock=INTEGER`	Delete the lock symlink that a previous run of **pan_tcp_capture_client** created but failed to delete. Requires the PID of the previous **pan_tcp_capture_client** run that created the symlink. This PID should be provided by the error message that reported that the **pan_tcp_capture_client** lock file exists.

Option	Description
`--workpath=STRING`	Set the directory where **pan_tcp_capture_client** will create its output directory and where it will write the final tarred output file. **Pan_tcp_capture_client** will output the path and name of the final output file to the screen when it completes. The default is set by platform, and is usually **/var/tmp/** or **/tmp/**.

Example

```
$ pan_tcp_capture_client
$ pan_tcp_capture_client --interface eth0
$ pan_tcp_capture_client --minutes 4 --nfs
$ pan_tcp_capture_client --roll --hours 1
$ pan_tcp_capture_client --filter 2049 --filter 38.82.195.140 --filter 10.64.8.179
$ pan_tcp_capture_client --nodump
```

pan_trace

Description

Manages diagnostic tracing and collects trace buffers.

Syntax

```
pan_trace [subcommand] ?[arg1] [arg2] ...?
```

Options

Option	Description
`flags ?basic|extended [flags]?`	Display or set the flags for state collection or diagnostic tracing. The **basic** flags affect state collection. The **extended** flags affect diagnostic tracing. Modification of the flags beyond the default settings should only be done under the direction of VDURA Technical Support. Also consider using local PanCLI commands matching **pan_debug_trace_*** for additional commandline options not supported in the **pan_trace** command.
`getdatasets system|[ip] [file]`	Write the datasets for all Directors in the system or given host to the given PanFS file. Files must be specified to full path starting with **/volume-name**.

Option	Description
`getstate system|[ip] [file]`	Write state information for the entire system or given host to the given PanFS or remote file. To specify a PanFS file, give the full file path starting with the volume name. This command launches an interactive **scp** session if a remote file is specified. To specify a remote file, use the format [**user]@[host]:[file]**, and the system will transfer the file there. If the filename does not end in **.tgz**, .tgz will be appended to the name. Do not specify a directory name.
`query system|[ip]`	Display whether diagnostic trace is enabled for all Directors in the system or a Director specified by IP address.
`start system|[ip]`	Start diagnostic tracing for all Directors in the system or a Director specified by IP address.
`stop system|[ip] [file]`	Stop diagnostic tracing, collect data from the entire system or from only from the given host, and write it to the given PanFS file. Files must be specified using a full path starting with **/volume-name**.
`timeout ?[minutes]?`	Display or set the diagnostic trace timeout. The maximum timeout is 720 minutes. The default is 30 minutes.

pan_trace_control

Description

Service mode command for the following trace actions: **reset**, **prego**.

Syntax

```
pan_trace_control ?[options]? ?[action ...]?
```

Options

Option	Description
`--env-equiv`	Show environment var equivalents.
`-h, --help`	Show usage message.

panfs_trace

Description

Collect traces from PanFS clients and gateways.

Syntax

```
panfs_trace ?options? [target]
```

The command must always include the target (string) argument, which is a path to any file or directory in PanFS.

Options

Option	Description
`-b, --bindcache`	Dump bindcache configuration.
`-c, --callbacks`	Dump callback info table.
`-C, --caps`	Dump capability list for a single object.
`--defaults`	Show defaults.
`-y, --dir_buf_collect`	Directory buffer collection.
`--env-equiv`	Show environment **var** equivalents.
`-e, --exported_opstat_print`	Print client's exported opstats.
`-f, --filter`	Set trace filters (format: **module_name:uint32**. Commonly used module names: **fs_client**, **sam**, **obsd_client**, **ips**, **rpc**).
`-F, --filter_mask=[int3]`	Predefined trace filters (0:def, 1:vis, 2:ext, 3:ext+caps+rpcs, 4:3+pannodes, 5:4+all caches, 6:5+snapshots, 7:6+detailed)
`-h, --help`	Show usage message.
`-X, --hexdump_mod`	Dump information in hex.

Option	Description
`-I, --init-perf-meter`	Initialize client's performance meter module (compiled out by default).
`-j, --ioctl_panfs_flus`	Flush data and attributes.
`-K, --ioctl_panfs_getattr_cached`	Get panfs attributes from cache.
`-J, --ioctl_panfs_getattr_ondisk`	Get panfs attributes from disk.
`-k, --ioctl_panfs_getcbinfo`	Get panfs callback info.
`-L, --ioctl_panfs_statfs`	Get panfs statfs info.
`-Y, --memory_stats_dump`	Dump out memory allocation stats.
`-n, --name_cache`	Dump name cache info.
`-x, --off`	Turn off **pan_log**-based tracing.
`-o, --opstat_print`	Prints client's opstats.
`-O, --opstat_reset`	Resets client's opstats.
`-R, --packed_args`	Dump inflight operations.
`-z, --pannode_extra`	Dump extra **pannode** information.
`-p, --pannodes`	Dump **pannode** table.
`-M, --print_map`	Print map for a particular object.
`-P, --print_perf_meter`	Prints client's ptimers (compiled out by default).
`-d, --proc_trace_dump=[string]`	Dump trace buffers using the proc trace file from **/proc/fs/panfs/tracebuf**.
`-q,--query_mount_settings`	Query the mount point for all available mount settings.
`-m, --realm`	Dump realms table.
`-Z, --reset_perf_meter`	Resets client's ptimers (compiled out by default).

Option	Description
`-S, --silly_rename_translate`	Translate a **sillyrename** filename to owning machine.
`-N, --sock_dump`	Dump socket info.
`-s, --stack_dum`	Dump the thread stacks in the kernel.
`-T, --trace_collect=[string]`	Collect traces to a file.
`-u, --trace_id=[uint32]`	ID of **trace_id** to collect from.
`-U, --trace_resize=[byte size]`	Resize the trace buffers in the client. Note: A **byte size** argument is a **uint64** with an optional **[bBkKmMgGtT]** suffix. For instance, 1 means 1 byte, 1 G means 1 gigabyte, etc.
`-v, --version`	Print the version.

panic

Description

Service mode command to panic this node.

Panic will cause the node to dump kernel core and reboot.

Panic of individual nodes can cause system-wide outages and should only be done under the direction of VDURA Technical Support.

Syntax

```
panic
```

panstack

Description

Service mode command to decode stack traces.

Syntax

```
panstack [-d debugger] [-p] [-c corefile] programfile stack...
panstack [-d debugger] [-p] [-c corefile] programfile < infile > outfile
```

Specifying a programfile of kernel will cause translation to be done in the context of the running kernel, assumed to be /kernel, with the present set of loaded kernel modules. Images for these kernel modules are expected to be available in /boot/kernel/. In this case, -p is disallowed. If the full path is given for the programfile, the -p option will be ignored.

The panstack command decodes *{...}* pan_{printf,log} ~S stack traces given on its command line or found in its standard input. When used as a filter, panstack looks for the stack traces in the standard input and translates them to symbolic names, using the symbols in the provided executable. It copies everything else from the standard input to the standard output unchanged. For best results, run panstack on a machine with the same shared libraries (libc.so, etc.) as the machine where the stack traces were generated, as they are just as important as the programfile executable, and panstack locates and loads them via the normal mechanisms.

Filter processing is generally much faster when the standard input is a seekable file instead of a pipe.

Example

```
panstack pan_mgr "_{28ef6cdc|28eb8ba3|8ccec49|28ebb4a3}_"
panstack pan_mgr "28ef6cdc|28eb8ba3|8ccec49|28ebb4a3"
panstack pan_mgr 8ccec49
panstack kernel < /var/log/all.log > outfile
```

Options

Option	Description
`-d debugger`	Specify an alternate name for **gdb**/**dbx**. The default is **gdb53** on FreeBSD 4.6 systems and **gdb** elsewhere.
`-p`	Search for the given programfile in the environment variable **PATH**.
`-c corefile`	Specify an optional **corefile**. This is needed in instances where shared libraries or kernel modules are not loading at the same address as when the stack data was captured, and a core (or vmcore) is available from the same era as the stack data (and is therefore known to specify matching addresses).

paritywarnings

Description

Service mode command to enable/disable parity warnings. The default is **enable** (i.e., to display parity warnings). If no arguments are given, the current setting is displayed.

Syntax

```
paritywarnings ?enable|disable?
```

pasxml

Description

Gets **pasxml** URL output. If no URL is given, the **index** URL is gathered, listing the supported URLs and their options keys.

Optional keys and values can be given to specify which output is desired. Since version 5.5.0, the following XML tag or attribute names have been changed:

Version 4.1.x to 5.0.1	Version 5.5.x and newer
blade	node
blades	nodes
bladeSerial	nodeSerial
bladeImbalanceGB	osdImbalanceGB
shelf	enclosure
shelves	enclosures
shelfName	enclosureName
shelfMaster	enclosureMaster

A **version** key can be given to get the **pasxml** output that is compatible with the old version, such as 4.1.x and 5.0.x. The value for this key should be a valid version earlier than 5.5.0, e.g. **4.1.0** or **5.0.0**. If no version key is given or the version value is not valid, the command defaults to the latest version.

See the **pasxml** chapter in the *VDURA ActiveStor User Guide* to get the complete documentation for this command.

Syntax

```
pasxml ?[url] ?[key1] [value1] ...??
```

Example

Get the list of accepted pasxml URLs:

```
pasxml
```

Get the node listing, including statistics:

```
pasxml nodes stat 1
```

Get all listings, including alerts and stats:

```
pasxml all stat 1 alert 1
```

Get all listings, including alerts and stats that are compatible with the 4.1.0 version:

```
pasxml all stat 1 alert 1 version 4.1.0
```

pax

Description

The **pax** command will read, write, and list the members of an archive file and copy directory hierarchies. **Pax** operation is independent of the specific archive format and supports a wide variety of different archive formats. A list of supported archive formats can be found under the description of the *-x* option. The presence of the *-r* and *-w* options specify which of the following functional modes **pax** will operate under (list, read, write, or copy).

While processing a damaged archive during a read or list operation, **pax** will attempt to recover from media defects and will search through the archive to locate and process the largest number of archive members possible (see the *-E* option for more details on error handling).

About Operands

Directory operands specify a destination directory pathname. If the directory operand does not exist, or it is not writable by the user, or it is not of type directory, **pax** will exit with a non-zero exit status.

Pattern operands are used to select one or more pathnames of archive members. Archive members are selected using the pattern matching notation described by **fnmatch(3)**. When a pattern operand is not supplied, all members of the archive will be selected. When a pattern matches a directory, the entire file hierarchy rooted at that directory will be selected. When a pattern operand does not select at least one archive member, **pax** will write these pattern operands in a diagnostic message to standard error and then exit with a non-zero exit status.

A file operand specifies the pathname of a file to be copied or archived. When a file operand does not select at least one archive member, **pax** will write these file operand pathnames in a diagnostic message to standard error and then exit with a non-zero exit status.

Syntax

```
pax [-cdnvz] [-f archive] [-s replstr] ... [-U user] ... [-G group] ... [-T [from_date]
[,to_date]] ... [pattern ...]
pax -r [-cdiknuvzDYZ] [-f archive] [-o options] ... [-p string] ... [-s replstr] ... [-E limit]
[-U user] ... [-G group] ... [-T [from_date] [,to_date]] ... [pattern ...]
pax -w [-dituvzHLPX] [-b blocksize] [[-a] [-f archive]] [-x format] [-s replstr] ... [-o
options] ... [-U user] ... [-G group] ... [-B bytes] [-T [from_date] [,to_date] [/[c][m]]] ...
[file ...]
pax -r -w [-diklntuvDHLPXYZ] [-p string] ... [-s replstr] ... [-U user] ... [-G group] ... [-T
[from_date] [,to_date] [/[c][m]]] ... [file ...] directory
```

Options

Option	Description
none	Write a table of contents to standard output of the members of the archive file read from standard input whose pathnames match the specified patterns. The table of contents contains one filename per line and is written using single line buffering.

Option	Description
`-a`	Append files to the end of an archive that was previously written. If an archive format is not specified with a -**x** option, the format currently being used in the archive will be selected. Any attempt to append to an archive in a format different from the format already used in the archive will cause **pax** to exit immediately with a non-zero exit status. The blocking size used in the archive volume where writing starts will continue to be used for the remainder of that archive volume. **Warning:** Many storage devices are not able to support the operations necessary to perform an append operation. Any attempt to append to an archive stored on such a device may damage the archive or have other unpredictable results. Tape drives in particular are more likely to not support an append operation. An archive stored in a regular file system file or on a disk device will usually support an append operation.
`-b blocksize`	When writing an archive, block the output at a positive decimal integer number of bytes per write to the archive file. The blocksize must be a multiple of 512 bytes with a maximum of 64512 bytes. Archives larger than 32256 bytes violate the POSIX standard and will not be portable to all systems. A blocksize can end with **k** or **b** to specify multiplication by 1024 (1K) or 512, respectively. A pair of blocksizes can be separated by **x** to indicate a product. A specific archive device may impose additional restrictions on the size of blocking it will support. When blocking is not specified, the default blocksize is dependent on the specific archive format being used (see the -**x** option).

Option	Description
`-B bytes`	Limit the number of bytes written to a single archive volume to bytes. The bytes limit can end with **m**, **k**, or **b** to specify multiplication by 1048576 (1M), 1024 (1K) or 512, respectively. A pair of bytes limits can be separated by **x** to indicate a product. **Warning:** Only use this option when writing an archive to a device that supports an end-of-file read condition based on last (or largest) write offset, such as a regular file or a tape drive. The use of this option with a floppy or hard disk is not recommended.
`-c`	Match all file or archive members except those specified by the pattern and file operands.
`-d`	Cause files of type directory being copied or archived, or archive members of type directory being extracted, to match only the directory file or archive member and not the file hierarchy rooted at the directory.
`-D`	This option is the same as the -**u** option, except that the file inode change time is checked instead of the file modification time. The file inode change time can be used to select files whose inode information (e.g., UID, GID, etc.) is newer than a copy of the file in the destination directory.

Option	Description
`-E limit`	Limit the number of consecutive read faults while trying to read a flawed archive to limit. With a positive limit, **pax** will attempt to recover from an archive read error and will continue processing, starting with the next file stored in the archive. A limit of 0 will cause **pax** to stop operation after the first read error is detected on an archive volume. A limit of **NONE** will cause **pax** to attempt to recover from read errors forever. The default limit is a small positive number of retries. Use **NONE** with extreme caution, as **pax** may get stuck in an infinite loop on a very badly flawed archive.
`-f archive`	Specify archive as the pathname of the input or output archive, overriding the default standard input (for list and read) or standard output (for write). A single archive may span multiple files and different archive devices. When required, **pax** will prompt for the pathname of the file or device of the next volume in the archive.
`-G group`	Select a file based on its group name, or when starting with a **#**, a numeric GID. A **** can be used to escape the **#**. Multiple **-G** options may be supplied; checking stops with the first match.
`-H`	Follow only command line symbolic links while performing a physical file system traversal.

Option	Description
`-i`	Interactively rename files or archive members. For each archive member matching a pattern operand or each file matching a file operand, **pax** will prompt to **/dev/tty** giving the name of the file, its file mode, and its modification time. **Pax** will then read a line from **/dev/tty**. If this line is blank, the file or archive member is skipped. If this line consists of a single period, the file or archive member is processed with no modification to its name. Otherwise, its name is replaced with the contents of the line. **Pax** will immediately exit with a non-zero exit status if EOF is encountered when reading a response or if **/dev/tty** cannot be opened for reading and writing.
`-k`	Do not overwrite existing files.
`-l`	Link files. In copy mode (-**r** -**w**), hard links are made between the source and destination file hierarchies whenever possible.
`-L`	Follow all symbolic links to perform a logical file system traversal.
`-n`	Select the first archive member that matches each pattern operand. No more than one archive member is matched for each pattern. When members of type directory are matched, the file hierarchy rooted at that directory is also matched unless -**d** is also specified.
`-o options`	Provide information to modify the algorithm for extracting or writing archive files which are specific to the archive format specified by -**x**. In general, options take the form **name**=**value**.
`-P`	Do not follow symbolic links; perform a physical file system traversal. This is the default mode.

Option	Description
`-p string`	Specify one or more file characteristic options (privileges). The string option-argument is a string specifying file characteristics to be retained or discarded on extraction. The string consists of the specification characters **a**, **e**, **m**, **o**, and **p**. Multiple characteristics can be concatenated within the same string and multiple -**p** options can be specified. The meaning of the specification characters are as follows: - **a**: Do not preserve file access times. By default, file access times are preserved whenever possible. - **e**: Preserve everything–user ID, group ID, file mode bits, file access time, and file modification time. This is intended to be used by root, someone with all the appropriate privileges, to preserve all aspects of the files as they are recorded in the archive. The **e** flag is the sum of the **o** and **p** flags. - **m**: Do not preserve file modification times. By default, file modification times are preserved whenever possible. - **o**: Preserve the user ID and group ID. - **p**: Preserve the file mode bits. This intended to be used by a user with regular privileges who wants to preserve all aspects of the file other than ownership. The file times are preserved by default, but two other flags are offered to disable this and use the time of extraction instead. In the preceding list, “preserve” indicates that an attribute stored in the archive is given to the extracted file, subject to the permissions of the invoking process. Otherwise, the attribute of the extracted file is determined as part of the normal file creation action. If neither the **e** nor the **o**

Option	Description
	specification character is specified, or the user ID and group ID are not preserved for any reason, **pax** will not set the **S_ISUID (setuid)** and **S_ISGID (setgid)** bits of the file mode. If the preservation of any of these items fails for any reason, **pax** will write a diagnostic message to standard error. Failure to preserve these items will affect the final exit status, but will not cause the extracted file to be deleted. If the file characteristic letters in any of the string option-arguments are duplicated or conflict with each other, the one(s) given last will take precedence. For example, if **-p eme** is specified, file modification times are still preserved.
`-r`	Read an archive file from standard input and extract the specified files. Archive format and blocking are automatically determined on input. When an extracted file is a directory, the entire file hierarchy rooted at that directory is extracted. All extracted files are created relative to the current file hierarchy. The setting of ownership, access, and modification times, as well as file mode of the extracted files are discussed in more detail under the *-p* option. If any intermediate directories are needed to extract an archive member, these directories will be created as if **mkdir(2)** was called with the bitwise inclusive OR of **S_IRWXU**, **S_IRWXG**, and **S_IRWXO** as the mode argument. When the selected archive format supports the specification of linked files and these files cannot be linked while the archive is being extracted, pax will write a diagnostic message to standard error and exit with a non-zero exit status at the completion of operation.

Option	Description
`-r -w`	Copies the file operands to the destination directory. When no file operands are specified, a list of files to copy with one per line is read from the standard input. When a file operand is also a directory, the entire file hierarchy rooted at that directory will be included. The effect of the copy is as if the copied files were written to an archive file and then subsequently extracted, except that there may be hard links between the original and the copied files (see the -l option). The destination directory must not be one of the file operands or a member of a file hierarchy rooted at one of the file operands. The result of a copy under these conditions is unpredictable.

Option	Description
`-s replstr`	Modify the file or archive member names specified by the pattern or file operands according to the substitution expression replstr, using the syntax of the **ed(1)** utility regular expressions. The format of these regular expressions is **/old/new/[gp]**. As in **ed(1)**, **old** is a basic regular expression, while **new** can contain an ampersand (**&**), **\n** (where **n** is a digit) back-references, or subexpression matching. The **old** string may also contain newline characters. Any non-null character can be used as a delimiter (**/** is shown here). Multiple -**s** expressions can be specified. The expressions are applied in the order they are specified on the command line, terminating with the first successful substitution. The optional trailing **g** continues to apply the substitution expression to the pathname substring which starts with the first character following the end of the last successful substitution. The first unsuccessful substitution stops the operation of the **g** option. The optional trailing **p** will cause the final result of a successful substitution to be written to standard error in the format `[original pathname] >> [new pathname]`. File or archive member names that substitute to the empty string are not selected and will be skipped.
`-t`	Reset the access times of any file or directory read or accessed by **pax** to be the same as they were before being read or accessed by **pax**.

Option	Description
`-T [from_date] [,to_date] [/[c][m]]`	Allow files to be selected based on a file modification or inode change time falling within a specified time range of **from_date** to **to_date**, inclusive. If only a **from_date** is supplied, all files with a modification or inode change time equal to or younger are selected. If only a **to_date** is supplied, all files with a modification or inode change time equal to or older will be selected. When the **from_date** is equal to the **to_date**, only files with a modification or inode change time of exactly that time will be selected. When **pax** is in write or copy mode, the optional trailing field **[c][m]** can be used to determine which file time (inode change, file modification, or both) is used in the comparison. - If neither is specified, the default is to use file modification time only. - Use **m** to specify the comparison of file modification time (the time when the file was last written). - Use **c** to specify the comparison of inode change time (the time when the file inode was last changed, e.g. a change of owner, group, mode, etc). - When **c** and **m** are both specified, the modification and inode change times are both compared. Inode change time comparison is useful in selecting files whose attributes were recently changed or selecting files which were recently created and had their modification time reset to an older time, as when a file is extracted from an archive and the modification time is preserved. Time comparisons using both file times is useful when **pax** is used to create a time based incremental archive (i.e. only files that were

Option	Description
	changed during a specified time range will be archived). A time range is made up of six different fields, and each field must contain two digits. The format is **[yy[mm[dd[hh]]]]mm[.ss]**. - **yy** is the last two digits of the year - the first **mm** is the month (from 01 to 12) - **dd** is the day of the month (from 01 to 31) - **hh** is the hour of the day (from 00 to 23) - the second **mm** is the minute (from 00 to 59) - **ss** is the seconds (from 00 to 59); may be added independently of the other fields The minute field **mm** is required, while the other fields are optional and must be added in the order of **hh**, **dd**, **mm**, **yy**. Time ranges are relative to the current time, so **-T 1234/cm** would select all files with a modification or inode change time of 12:34 PM today or later. Multiple -**T** time ranges can be supplied, but checking stops with the first match.
`-u`	Ignore files that are older (having a less recent file modification time) than a pre-existing file or archive member with the same name. During read, an archive member with the same name as a file in the file system will be extracted if the archive member is newer than the file. During write, a file system member with the same name as an archive member will be written to the archive if it is newer than the archive member. During copy, the file in the destination hierarchy is replaced by the file in the source hierarchy or by a link to the file in the source hierarchy if the file in the source hierarchy is newer.

Option	Description
`-U user`	Select a file based on its username or, when starting with a **#**, a numeric UID. A **** can be used to escape the **#**. Multiple -**U** options may be supplied and checking stops with the first match.
`-v`	During a list operation, produce a verbose table of contents using the format of the **ls(1)** utility with the -**l** option. For pathnames representing a hard link to a previous member of the archive, the output has the format [ls **-l** listing] **==** [link name]. For pathnames representing a symbolic link, the output has the format `[ls -l listing] => [link name]` where [ls **-l** listing] is the output format specified by the **ls(1)** utility when used with the -l option. For all the other operational modes (read, write, and copy), pathnames are written and flushed to standard error without a trailing newline as soon as processing begins on that file or archive member. The trailing newline is not buffered and is written only after the file has been read or written.
`-w`	Write files to the standard output in the specified archive format. When no file operands are specified, standard input is read for a list of pathnames, with one per line without any leading or trailing **[blanks]**. When a file operand is also a directory, the entire file hierarchy rooted at that directory will be included.
`-X`	When traversing the file hierarchy specified by a pathname, do not descend into directories that have a different device ID. See the **st_dev** field as described in **stat(2)** for more information about device ID's.

Option	Description
`-x [format]`	Specify the output archive format, with the default format being **ustar**. **Pax** currently supports the following formats: - **cpio**: The extended **cpio** interchange format specified in the IEEE Std 1003.2 (POSIX.2) standard. The default blocksize for this format is 5120 bytes. Inode and device information about a file, used for detecting file hard links by this format and which may be truncated by this format, is detected by **pax** and is repaired. - **bcpio**: The old binary **cpio** format. The default blocksize for this format is 5120 bytes. This format is not portable and should not be used when other formats are available. Inode and device information about a file, used for detecting file hard links by this format and which may be truncated by this format, is detected by **pax** and is repaired. - **sv4cpio**: The System V release 4 **cpio**. The default blocksize for this format is 5120 bytes. Inode and device information about a file, used for detecting file hard links by this format and which may be truncated by this format, is detected by **pax** and is repaired. - **sv4crc**: The System V release 4 **cpio** with file CRC checksums. The default blocksize for this format is 5120 bytes. Inode and device information about a file, used for detecting file hard links by this format and which may be truncated by this format, is detected by **pax** and is repaired. - **tar**: The old BSD .tar format as found in 4.3BSD. The default blocksize for this format is 10240 bytes. Pathnames stored by this format must be 100 characters or less in length. Only regular files, hard links, soft links, and directories will be archived; other file system types are not

Option	Description
	supported. For backwards compatibility with even older .tar formats, the -**o** option can be used when writing an archive to omit the storage of directories. This option takes the form -**o write_opt=nodir**. - **ustar**: The extended .tar interchange format specified in the IEEE Std 1003.2 (POSIX.2) standard. The default blocksize for this format is 10240 bytes. Pathnames stored by this format must be 250 characters or less in length. **Pax** will detect and report any file that it is unable to store or extract as the result of any specific archive format restrictions. The individual archive formats may impose additional restrictions on use. Typical archive format restrictions include, but are not limited to file pathname length, file size, link pathname length, and the type of the file.
`-Y`	This option is the same as the **-D** option, except that the inode change time is checked using the pathname created after all the file name modifications have completed.
`-Z`	This option is the same as the **-u** option, except that the modification time is checked using the pathname created after all the filename modifications have completed.
`-z`	-**z**: Use **gzip(1)** to compress/decompress the archive while writing/reading. Incompatible with -**a**.

Options that operate on the names of files or archive members (-**c**, -**i**, -**n**, -**s**, -**u**, -**v**, -**D**, -**G**, -**T**, -**U**, -**Y**, and -**Z**) interact as follows:

- When extracting files during a read operation, archive members are selected based only on the user-specified pattern operands as modified by the -**c**, -**n**, -**u**, -**D**, -**G**, -**T**, and -**U** options. Next, any -**s** and -**i** options will modify, in that order, the names of the selected files. The -**Y** and -**Z** options will then be applied based on the final pathname. Finally, the -**v** option will write the names resulting from these modifications.
- When archiving files during a write operation or copying files during a copy operation, archive members are selected based only on the user-specified pathnames as modified by the -**n**, -**u**, -**D**, -**G**, -**T**, and -**U** options (the -**D** option only applies during a copy operation). Next, any -**s** and -**i** options will modify, in that order, the names of these selected files. During a copy operation, the -**Y** and -**Z** options will then be applied based on the final pathname. Finally, the -**v** option will write the names resulting from these modifications. When one or both of the -**u** or -**D** options are specified along with the -**n** option, a file is not considered selected unless it is newer than the file to which it is compared.

phonehome

Description

Service mode command to modify the phone home schedule.

Syntax

```
phonehome [schedule] [get|set] ?[nightly|weekly]?
```

Option	Description
`get`	Get the current phone home schedule.
`set [nightly|weekly]`	Set the phone home schedule to be nightly or weekly.

ping

Description

Sends ICMP ECHO_REQUEST packets to network hosts.

Syntax

```
ping [-AQRadnqrv] [-c count] [-i wait] [-m ttl] [-p pattern] [-P policy] [-s packetsize] [-S src_addr] [-t timeout] [host | [-L] [-I interface] [-T ttl] mcast-group]
```

Options

Option	Description
`-A`	Audible. Output a bell (ASCII 0x07) character when no packet is received before the next packet is transmitted. To cater for round-trip times that are longer than the interval between transmissions, further missing packets cause a bell only if the maximum number of unreceived packets has increased.
`-a`	Audible. Include a bell (ASCII 0x07) character in the output when any packet is received. This option is ignored if other format options are present.
`-c count`	Stop after sending (and receiving) **count** ECHO_RESPONSE packets. This command has a default of 3 counts.
`-d`	Set the `SO_DEBUG` option on the socket being used.
`-i wait`	Wait **wait** seconds between sending each packet. The default is to wait for one second between each packet.

Option	Description
`-I interface`	Source multicast packets with the given interface address. This flag only applies if the ping destination is a multicast address.
`-m ttl`	Set the IP Time To Live for outgoing packets. If not specified, the kernel uses the value of the `net.inet.ip.ttl` MIB variable.
`-L`	Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.
`-n`	Numeric output only. No attempt is made to look up symbolic names for host addresses.
`-p pattern`	Specify up to 16 “pad” bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, **-p ff** causes the sent packet to be filled with all ones.
`-P policy`	The IPsec policy for the ping session. For details, see **ipsec(4)** and **ipsec_set_policy(3)**.
`-Q`	Somewhat quiet output. Do not display ICMP error messages that are in response to the query messages. Originally, the **-v** flag was required to display such errors, but **-v** displays all ICMP error messages. On a busy machine, this output can be overbearing. Without the **-Q** flag, **ping** prints any ICMP error messages caused by its own ECHO_REQUEST messages.
`-q`	Quiet output. Nothing is displayed except the summary lines at startup time and when finished.

Option	Description
`-R`	Record route. Includes the RECORD_ROUTE option in the ECHO_REQUEST packet and displays the route buffer on returned packets. The IP header is only large enough for nine such routes. The **traceroute(8)** command is usually better at determining the route packets take to a particular destination. If more routes come back than should, such as due to an illegal spoofed packet, **ping** prints the route list and then truncates it at the correct spot. Many hosts ignore or discard the RECORD_ROUTE option.
`-r`	Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it (for example, after the interface was dropped by **routed(8)**).
`-S src_addr`	Use the following IP address as the source address in outgoing packets. On hosts with more than one IP address, this option can be used to force the source address to be something other than the IP address of the interface the probe packet is sent on. If the IP address is not one of the interface addresses of this machine, an error is returned and nothing is sent.
`-t timeout`	Specify a timeout, in seconds, before **ping** exits regardless of how many packets have been received.
`-T ttl`	Set the IP Time To Live for multicasted packets. This flag only applies if the ping destination is a multicast address.

Option	Description
`-v`	Verbose output. ICMP packets other than ECHO_RESPONSE that are received are listed.

Example

Send a ping to a host name happy.example.com:

```
ping happy.example.com
```

See Also

netdiag

process

Description

Service mode command to view status for, start, and stop local processes.

Syntax

```
process ?[subcommand] ?[arg]??
```

Options

Option	Description
none	If no options are given, display the current state of the monitor process followed by basic status of all processes that are allowed or required in the current state. Use the **status** option to get more detailed info.
`list ?[pattern]?`	This listing includes all processes regardless of whether they are currently running or what state they can run in. If a glob pattern is given, only show process names that match.
`monitorstate`	Display the state of the monitor process. A node that is configured into an active realm should have a monitor state of **normal**.
`restart [module-name]`	Kill the given process if it is already running and restart it. If the process is not allowed in the current node state, then it will not be started and an error will be returned. **Warning:** Misuse of this command can result in system unavailability and data loss. Users should only restart processes under the explicit direction of VDURA Technical Support.

Option	Description
`start [module-name]`	Start a process. If the process is already running, this request does nothing. If the process is not allowed in the current node state, it will not be started and an error will be returned. **Warning:** Misuse of this command can result in system unavailability and data loss. Users should only start processes under the explicit direction of VDURA Technical Support.
`status ?[module-name-or-glob-pattern]?`	Show the status of all **pan_monitor** modules that match the given pattern. If no pattern is given, show all modules. If multiple modules match the pattern, show which ones are running and whether they are run-once or restartable. If only one module matches the pattern, give more detailed information about that process. Status for a process will include the following: - process name - status: running, running with kill in progress, or not running - permission: run once or restartable - restart count: [number of times module was restarted since boot] - last start time - last exit time - last exit code - PID, command name, state, %cpu, cpu time, %mem, vsz (current virtual size), rss (resident size), and wchan (wait channel) for each descendant of this process (itself included), See **help ps** for more information on the per-PID information.

Option	Description
`stop [module-name]`	Stop a process. If the given process is not already running, this request does nothing. **Warning:** Misuse of this command can result in system unavailability and data loss. Users should only stop processes under the explicit direction of VDURA Technical Support.

prompt

Description

Display or set the PanCLI prompt for all nodes.

Syntax

```
prompt ?[subcommand] ?[arg1]?
```

Example

Set the prompt to the DNS name and login (e.g. [panfs:admin]):

```
prompt set {[%N:%u] }
```

Set the prompt to the IP address (e.g. 10.1.81.11$):

```
prompt set %i\$
```

Set the prompt to the slot and time (e.g. slot 1 17:17:21%):

```
prompt set "slot %s %T%% "
```

Options

Option	Description
none	If no subcommand is given, print the current prompt.
`clear ?-local?`	Restore the prompt to the system default. If the **-local** option is given, only the current session will be affected by the change. Otherwise, the prompt for all PanCLI users on all nodes in the system will be affected.

Option	Description
`get`	Display the prompt spec without substituting the special characters. If no custom prompt has been set, an empty string is returned.
`set ?-local? [new-prompt-spec]`	Set the prompt for all users (admin and guest) on all nodes to the given string. If the **-local** option is given, only the current session will be affected by the change. The following character sequences will be substituted: - **%u**: Login name - **%N**: Full system DNS name - **%n**: First segment of system DNS name - **%i**: Local node IP address - **%s**: Local node slot number - **%S**: Local node enclosure serial - **%h**: Local node hardware serial number - **%d**: Local node data serial number - **%r**: Local node role: Storage, Director, or replication set (repset) - **%t**: Time of day: HH:MM - **%T**: Time of day: HH:MM:SS - **%%**: A single percent character Prompts must end with a pound, percent, close square bracket, or dollar character followed optionally by whitespace. Prompts which include special characters ([ ] and/or $) must be surrounded with curly braces or escaped.

property

Description

Service mode command to manage property bits. To set a flag, give a **+[flag]** argument. To clear a flag, give a **-[flag]** argument.

Misuse of this command can result in data loss. Users should only set and clear property flags under the explicit direction of VDURA Technical Support.

Syntax

```
property ?(+|-)[flag] ...?
```

Options

Flag	Description
none	If no arguments are given, display each flag in the current property values. A **+** will appear in front of all flags that are currently set. A **-** will appear in front of all flags that are cleared.
`config`	System is configured into a storage cluster.
`offline`	System is offline.
`diag`	System is in diagnostic mode.
`failed`	System has failed due to a software error.
`perm-failed`	System has failed due to a hardware fault.
`factory`	First boot after final factory netboot.
`service`	First boot after final service center netboot.
`customer`	First boot after final customer site netboot.
`shelf-master`	System is enclosure master.

Flag	Description
`realm-manager`	System is a member of the replication set (repset).
`object-server`	System is managing one or more VPODs.
`gateway`	System is running NFS/CIFS gateway services.
`pan-mgr`	System is running PanFS metadata services.

ps

Description

Shows the status of system processes.

Syntax

```
ps [-aCcefhjlmrSTuvwx] [-M core] [-N system] [-O fmt] [-o fmt] [-p pid] [-t tty] [-U username]
ps [-L]
```

Options

Option	Description
`-a`	Display information about other users' processes as well as your own.
`-c`	Change the "command" column output to contain just the executable name, rather than the full command line.
`-C`	Change the way the CPU percentage is calculated by using a "raw" CPU calculation that ignores "resident" time (this normally has no effect).
`-e`	Display the environment as well.
`-f`	Show command line and environment information about swapped out processes. This option is honored only if the UID of the user is 0.
`-h`	Repeat the information header as often as necessary to guarantee one header per page of information.
`-j`	Print information associated with the following keywords: **user**, **pid**, **ppid**, **pgid**, **sess**, **jobc**, **state**, **tt**, **time**, and **command**.

Option	Description
`-L`	List the set of available keywords.
`-l`	Display information associated with the following keywords: **uid**, **pid**, **ppid**, **cpu**, **pri**, **nice**, **vsz**, **rss**, **wchan**, **state**, **tt**, **time**, and **command**.
`-M core`	Extract values associated with the name list from the specified core instead of the default `/dev/kmem`.
`-m`	Sort by memory usage, instead of by process ID.
`-N system`	Extract the name list from the specified system instead of the default `/kernel`.
`-O fmt`	Add the information associated with the space-separated or comma-separated list of keywords specified, after the process ID, in the default information display. Keywords can be appended with an equals sign (`=`) and a string. This causes the printed header to use the specified string instead of the standard header.
`-o fmt`	Display information associated with the space-separated or comma-separated list of keywords specified. Keywords can be appended with an equals sign (`=`) and a string. This causes the printed header to use the specified string instead of the standard header.
`-p pid`	Display information associated with the specified process ID.
`-r`	Sort by current CPU usage, instead of by process ID.
`-S`	Change the way the process time is calculated by summing all exited children to their parent process.

Option	Description
`-T`	Display information about processes attached to the device associated with the standard input.
`-t tty`	Display information about processes attached to the specified terminal device.
`-U username`	Display the processes belonging to the specified user name.
`-u`	Display information associated with the following keywords: **user**, **pid**, **%cpu**, **%mem**, **vsz**, **rss**, **tt**, **state**, **start**, **time**, and **command**. The **-u** option implies the **-r** option.
`-v`	Display information associated with the following keywords: **pid**, **state**, **time**, **sl**, **re**, **pagein**, **vsz**, **rss**, **lim**, **tsiz**, **%cpu**, **%mem**, and **command**. The **-v** option implies the **-m** option.
`-w`	Use 132 columns to display information, instead of the default, which is your window size. If the **-w** option is specified more than once, **ps** uses as many columns as necessary without regard for your window size.
`-x`	Display information about processes without controlling terminals.

Keywords

Some of the available keywords are further specified as follows. To list all available keywords, use **ps -L**.

Keyword	Description
`%cpu`	The CPU utilization of the process. This is a decaying average over up to a minute of previous (real) time. Because the time base over which this is computed varies (processes can be very young), the sum of all **%CPU** fields can exceed 100%.

Keyword	Description
`%mem`	The percentage of real memory used by this process.
`lim`	The soft limit on memory used, specified through a call to **setrlimit(2)**.
`lstart`	The exact time the command started, using the `%c` format described in **strftime(3)**.
`nice`	The process scheduling increment (see **setpriority(2)**).
`rss`	The real memory (resident set) size of the process (in 1024-byte units).
`start`	The time the command started. If the command started less than 24 hours ago, the start time is displayed using the `%l:ps.1p` format described in **strftime(3)**. If the command started less than 7 days ago, the start time is displayed using the `%a6.15p` format. Otherwise, the start time is displayed using the `%e%b%y` format.
`state`	The state is given by a sequence of letters, for example, "RWNA". The first letter indicates the run state of the process (see the first table below). Additional characters after the first letter, if any, indicate additional state information (see the second table below).
`tt`	An abbreviation for the pathname of the controlling terminal, if any. The abbreviation consists of the three letters following `/dev/tty`, or, for the console, "con". This is followed by a "-" if the process can no longer reach that controlling terminal (that is, it has been revoked).
`vsz`	Virtual size in Kbytes (alias **vsize**).

Keyword	Description
`wchan`	The event (an address in the system) on which a process waits. When printed numerically, the initial part of the address is trimmed off and the result is printed in hex. For example, 0x80324000 prints as 324000.

The first letter of the **state** keyword is one of the following.

Letter	Meaning
`D`	A process in disk (or other short-term, uninterruptible) wait.
`I`	A process that is idle (sleeping for longer than about 20 seconds).
`J`	A process that is in **jail(2)**. The host name of the prison can be found in `/proc/[pid]/status`.
`R`	A runnable process.
`S`	A process that is sleeping for less than about 20 seconds.
`T`	A stopped process.
`Z`	A dead process (a “zombie”).

Additional characters after the first letter indicate additional state information.

Character	Meaning
`+`	The process is in the foreground process group of its control terminal.
`<`	The process has raised CPU scheduling priority.

Character	Meaning
`>`	The process has specified a soft limit on memory requirements and is currently exceeding that limit. Such a process is (necessarily) not swapped.
`A`	The process has asked for random page replacement (**MADV_RANDOM**, from **madvise(2)**, for example, **lisp(1)** in a garbage collect).
`E`	The process is trying to exit.
`L`	The process has pages locked in core (for example, for raw I/O).
`N`	The process has reduced CPU scheduling priority (see **setpriority(2)**).
`S`	The process has asked for FIFO page replacement (**MADV_SEQUENTIAL**, from **madvise(2)**, for example, a large image processing program using virtual memory to sequentially address voluminous data).
`s`	The process is a session leader.
`V`	The process is suspended during a vfork.
`W`	The process is swapped out.
`X`	The process is being traced or debugged.

Example

Display all users' processes:

```
ps -ax
```

See Also

top

pwd

Description

Returns working directory name.

raidutil

Description

Service mode command to view RAID information.

Syntax

```
raidutil [subcommand] [args]
```

Options

Option	Description
`mapoffset [path] [offset]`	Compute data and parity offsets in a file. The **raidutil mapoffset** command helps determine the component object locations that correspond to a given offset in a file. Given an offset in the file and the path to the file, it prints the component object IDs and offsets in those component objects that contain the data and parity. Files must be specified using a full path starting with **/volume-name**.
`scrub ?[-c]? ?[-o offset -l length]? [path]`	Check or correct parity in a file. The **raidutil scrub** command checks (without the -**c** flag) or corrects (with the -**c** flag) the RAID parity in a file. If the -**o** and -**l** options are not provided, the entire file is checked or scrubbed. Note that since each RAID stripe must be locked while it is scrubbed, this command may reduce performance for clients that are attempting to write to the file. This command may adversely affect system performance while it is running. Files must be specified using a full path starting with **/volume-name**.

realm

Description

Views and configures realm DirectFlow exports.

Syntax

```
realm exports ?[subcommand] ?[arg1]?
```

Options

Option	Description
`copy [file]`	Copy the realm DirectFlow exports file to the given location. The destination file must be specified using a full PanFS path starting with the volume name.
`upload [file]`	Deploy and save the given DirectFlow exports PanFS file. Use the **df** option to deploy the DirectFlow configuration. The destination file must be specified using a full PanFS path starting with the volume name.
`list`	Display the realm DirectFlow exports contents.

reboot

Description

Interactive command to reboot this node.

Rebooting individual nodes can cause system-wide outages and should only be done under the direction of VDURA Technical Support.

Syntax

```
reboot
```

register

Description

Registers your ActiveStor storage cluster with VDURA Support.

rm

Description

The **rm** utility attempts to remove the non-directory type files specified on the command line. If the permissions of the file do not permit writing, and the standard input device is a terminal, the user is prompted for confirmation.

Syntax

```
rm [-dfiPRrvW] file ...
```

Options

Option	Description
`-d`	Attempt to remove directories as well as other types of files.
`-f`	Attempt to remove the files without prompting for confirmation, regardless of the file's permissions. If the file does not exist, do not display a diagnostic message or modify the exit status to reflect an error. The **-f** option overrides any previous **-i** options.
`-i`	Request confirmation before attempting to remove each file, regardless of the file's permissions or whether or not the standard input device is a terminal. The **-i** option overrides any previous **-f** options.
`-P`	Overwrite regular files before deleting them. Files are overwritten three times, first with the byte pattern **0xff**, then **0x00**, and then **0xff** again, before they are deleted.

Option	Description
`-R`	Attempt to remove the file hierarchy rooted in each file argument. The -**R** option implies the -**d** option. If the -**i** option is specified, the user is prompted for confirmation before each directory's contents are processed, as well as before the attempt is made to remove the directory. If the user does not respond affirmatively, the file hierarchy rooted in that directory is skipped.
`-r`	Equivalent to -**R**
`-v`	Be verbose when deleting files, showing them as they are removed.
`-W`	Can only be used to recover files covered by whiteouts.

rmdir

Description

The **rmdir** utility removes the directory entry specified by each directory argument, provided it is empty.

Arguments are processed in the order given. For example, to remove both a parent directory and a subdirectory of that parent, the subdirectory must be specified first so the parent directory is empty when rmdir tries to remove it.

Syntax

```
rmdir [-p] directory ...
```

Options

Option	Description
`-p`	Each directory argument is treated as a pathname for which all components will be removed, if they are empty, starting with the last-most component. See **rm(1)** for fully non-discriminant recursive removal.

rpancli

Description

Interactive command to execute PanCLI commands on multiple nodes.

Either the **-interactive** option or a command must be entered. The command argument can be any PanCLI command. Note that commands will be run serially (not in parallel) on the requested nodes. Do not use this command to launch interactive commands remotely.

Syntax

```
rpancli ?[options]? -interactive|[command] [node-options]
```

Options

Option	Description
`-help`	Print this message.
`-timeout [sec]`	Wait this number of seconds for a node to respond before quitting. Default timeout is 20 seconds.

Node Options

Option	Description
`-host [ip]`	Specify each IP that you want to contact.
`-all`	Run the command on all nodes.

See **iplist** for other supported node options.

rsyslog

Description

Display or change remote syslog collectors and facility.

Syntax

```
rsyslog ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Send all events to katara.localdomain:

```
rsyslog update katara.localdomain
```

Send critical and warning events to 10.1.55.11 on the local1 facility:

```
rsyslog update 10.1.55.11 local1 critical warning
```

Options

Option	Description
none	If no subcommand is given, display the remote syslog collectors along with their filters and the facility.
`clear`	Clear the list of remote syslog collectors.
`remove [server-name]|[ip]`	Remove a host from the list of remote syslog collectors.

Option	Description
`update [server-name]|[ip] ?[facility]? ? critical? ?warning? ?info?`	Updates or adds a remote syslog collector, along with a facility and a specification of the desired event categories that will be sent to the named remote syslog collector. Legal values for the facility are: **auth**, **authpriv**, **console**, **cron**, **daemon**, **ftp**, **kern**, **lpr**, **mail**, **news**, **ntp**, **security**, **syslog**, **user**, **uucp**, **local0**, **local1**, **local2**, **local3**, **local4**, **local5**, **local6**, and **local7**. If no facility is given, the default facility is **local1**. If neither **critical**, **warning**, nor **info** is listed, then that remote syslog collector is unfiltered, and all events in the event log will be sent to that host. If the host is already listed, then the facility and event categories for this host will be updated to those listed or to the defaults if none are listed. If the host is not already listed, it is added, along with the specified facility and event categories.

s3

Description

The s3 command is deprecated as of 12.0.0. Use PanCTL to configure S3 on VDURA systems.

Implements support of the Amazon AWS S3 API as one of the interfaces for accessing the data on VDURA realms, in addition to Directflow, NFS, and SMB. Even though AWS S3 was originally developed for AWS cloud service, it is also now an industry-standard interface for communication between applications and storage systems in hybrid, on-prem, and private cloud information systems.

Syntax

```
s3 ?[subcommand] ?[arg1] [arg2] ...??
```

Example

At a high level, the steps to configure S3 are as follows:

- Create default volume for S3 buckets in the file system.
- s3 prefix set [path]
- s3 port ?port_number? optional
- s3 fsdefaults owner [uid] optional
- s3 fsdefaults group [gid] optional
- s3 fsdefaults dirmode [bits] optional
- s3 tls enabled ?yes|no? optional, but do the following if enabled:
- s3 tls credentials ?certificate_file_path private_key_file_path?
- s3 tls cacert ?ca_certificate_file_path?
- s3 enabled yes
- s3 user add [username] ?owner [owner]? ?group [group]?
- s3 group addusers [groupname] [space separated usernames]
- s3 bucket add [bucketname] ?[panfspath]
- s3 policy add [policyname] [policyjson] optional

For details, see table below.

Options

See table below.

Option	Description

`bucket`	Display, add, and delete buckets.
	Usage:
	`s3 bucket ?list?`
	`s3 bucket add [bucketname] ?[panfspath]`
	`s3 bucket delete [bucketname]`
	If the **list** option or no option is given, displays a list of all groups:
	`s3 bucket ?list?`
	If the **add** option is given, creates a new bucket:
	`s3 bucket add [bucketname]`
	`s3 bucket add [bucketname] ?[panfspath]`
	If the **delete** option is given, deletes specified bucket:
	`s3 bucket delete [bucketname]`
	Examples:
	`# s3 bucket`
	`Bucket Name Path`
	`test-buc /test/`
	`# s3 bucket list`
	`Bucket Name Path`
	`test-buc /test/`
	`# s3 bucket list`
	`Bucket Name Path`
	`test-buc-2 /another_volume/`
	`test-buc3 /test/test-buc3/`
	`# s3 bucket add test-buc`
	`Created bucket "test-buc" successfully.`
	`# s3 bucket add test-buc-2 another_volume`
	`Created bucket "test-buc-2" successfully.`
	`# s3 bucket delete test-buc-4`
	`Deleted bucket "test-buc-4" successfully.`
	`# s3 bucket delete test-buc-3`
	`Failed to delete bucket: No bucket with name "test-buc-3" was found.`

Option	Description

`enabled`	Display, enable, and disable S3 support.
	Usage:
	If no argument is given, display the current setting:
	`s3 enabled`
	If an argument is given, enable or disable the S3 Gateway support:
	`s3 enabled yes`
	`s3 enabled no`
	Examples:
	`# s3 enabled`
	`enabled`
	`# s3 enabled`
	`disabled`
	`# s3 enabled yes`
	`Enabling S3.`
	`# s3 enabled no`
	`Disabling S3.`
	`# s3 enabled no`
	`No change was needed.`
	`# s3 enabled yes`
	`No change was needed.`
	`# s3 enabled yes`
	`S3 can not be enabled if default prefix is not set.`

Option	Description

`fsdefaults`	View and change of default mode bits for owner/group for S3 configuration.
	Usage:
	If no options are given, command shows current default mode bits and owner/group:
	`s3 fsdefaults`
	Command can be used to set default owner. UID must be provided as a non-negative integer (0 to 4294967294) with or without “uid:” prefix:
	`s3 fsdefaults owner [uid]`
	Command can be used to set default group. GID must be provided as a non-negative integer (0 to 4294967294) with or without “gid:” prefix:
	`s3 fsdefaults group [gid]`
	Command can be used to set default directories mode. Mode bits must be provided as a non-negative integer with three digits, e.g. “755”:
	`s3 fsdefaults dirmode [bits]`
	Command can be used to set default objects (files) mode. Mode bits must be provided as a non-negative integer with three digits, e.g. “755”:
	`s3 fsdefaults objmode [bits]`
	Examples:
	`# s3 fsdefaults`
	`Default Owner: uid:65534`
	`Default Group: gid:65534`
	`Default Objects Mode: 644`
	`Default Directories Mode: 755`
	`# s3 fsdefaults owner 111`
	`S3 default owner has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 fsdefaults group 111`
	`S3 default group has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`

Option	Description

	`# s3 fsdefaults dirmode 555`
	`S3 default directories permissions have been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 fsdefaults objmode 555`
	`S3 default files permissions have been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`

Option	Description

`group`	Display, add, and delete groups.
	Usage:
	If the “list” option or no option is given, displays a list of all groups:
	`s3 group ?list?`
	If the “addusers” option is given, adds users to the new or existing group:
	`s3 group addusers [groupname] [space separated usernames]`
	If the “deleteusers” option is given, deletes users from the existing group:
	`s3 group deleteusers [groupname] [space separated usernames]`
	If the “delete” option is given, deletes specified group:
	`s3 group delete [groupname]`
	If the “details” option is given, displays the details of a specific group:
	`s3 group details [groupname]`
	Examples:
	`# s3 group`
	`There are currently no groups.`
	`# s3 group list`
	`There are currently no groups.`
	`# s3 group`
	`nik root`
	`# s3 group adduser nik nik`
	`Member(s) "nik" successfully added to the "nik" group.`
	`# s3 group deleteuser some_group nik`
	`Member(s) "nik" successfully deleted from the "some_group" group.`
	`# s3 group delete some_group`
	`Deleted group "some_group" successfully.`
	`# s3 group delete root`
	`Failed to delete group: Group "root" is not empty.`

Option	Description

	`# s3 group delete some_group`
	`Failed to delete group: No group with name "some_group" was found.`
	`# s3 group details nik`
	`Group Name: nik`
	`Members: nik`
	`Policy:`
	`# s3 group details some_group`
	`Failed to get group details: No group with name "some_group" was found.`

Option	Description

`policy`	Display, add, delete, get, and list policies.
	Set/unset a policy to either a group or a user.
	Usage:
	If the “list” option or no option is given, displays a list of all policies:
	`s3 policy ?list?`
	If the “add” option is given, creates a new policy:
	`s3 policy add [policyname] [policyjson]`
	- **[policyname]**: name of new policy
	- **[policyjson]**: path to JSON file that describes a policy and can have the following formats:
	- Local path: /panfs/policy.json; /tmp/policy.json
	- Remote path: user@host:/temp/policy.json
	If the “delete” option is given, deletes specified policy:
	`s3 policy delete [policyname]`
	If the “details” option is given, displays the details of a specific policy:
	`s3 policy details [policyname]`
	If the “get” option is given, displays the details of a specific policy in JSON format:
	`s3 policy get [policyname]`
	If the “setuser” option is given, applies a policy to a user:
	`s3 policy setuser [space separated policies] [username]`
	If the “setgroup” option is given, applies a policy to a group:
	`s3 policy setgroup [space separated policies] [groupname]`
	If the “unsetuser” option is given, unsets a policy from a user:
	`s3 policy unsetuser [space separated policies] [username]`
	If the “unsetgroup” option is given, unsets a policy from a group:
	`s3 policy unsetgroup [space separated`

Option	Description

	`policies] [groupname]`
	Examples:
	`# s3 policy`
	`consoleAdmin diagnostics readonly readwrite writeonly`
	`# s3 policy list`
	`consoleAdmin diagnostics readonly readwrite writeonly`
	`# s3 policy add test_policy ./test_policy.json`
	`Added policy "test_policy" successfully.`
	`# s3 policy add test_policy ./test_policy.json`
	`Failed to add policy: A policy by name "test_policy" already exists.`
	`# s3 policy add test_policy`
	`Deleted policy "test" successfully.`

Option	Description

`policy (cont.)`	`# s3 policy delete test`
	`Failed to delete policy: No policy with name "test" was found.`
	`# s3 policy details test`
	`Policy Name: test`
	`Created: 2023-06-29T12:14:31.287Z`
	`Updated: 2023-06-29T12:14:31.287Z`
	`# s3 policy details readonly`
	`Policy Name: readonly`
	`Created:`
	`Updated:`
	`# s3 policy get readonly`
	`{`
	`"Version": "2012-10-17",`
	`"Statement": [`
	`{`
	`"Effect": "Allow",`
	`"Action": [`
	`"s3:GetBucketLocation",`
	`"s3:GetObject"`
	`],`
	`"Resource": [`
	`"arn:aws:s3:::*"`
	`]`
	`}`
	`]`
	`}`
	`# s3 policy setuser test nik`
	`Policy "test" is set to the "nik" user.`
	`# s3 policy setuser test readonly nik`
	`Policies "test", "readonly" are set to the "nik" user.`
	`# s3 policy setgroup test test_group`
	`Policy "test" is set to the "test_group" group.`
	`# s3 policy setgroup test readonly test_group`

Option	Description

	`Policies "test", "readonly" are set to the "test_group" group.`
	`# s3 policy setgroup test test_group`
	`Policy "test" is unset to the "test_group" group.`
	`# s3 policy setgroup test readonly test_group`
	`Policies "test", "readonly" are unset to the "test_group" group.`

`port`	Binds S3 API to a specified port.
	Usage:
	If no options are given, shows current port configuration:
	`s3 port`
	If port_number is provided, sets provided port:
	`s3 port ?port_number?`
	If default option is given, sets port to 9000:
	`s3 port default`
	Examples:
	`# s3 port`
	`9000`
	`# s3 port 9000`
	`No change was needed.`
	`# s3 port 9000`
	`S3 port has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 port default`
	`S3 port has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 port default`
	`No change was needed.`

Option	Description

`prefix`	View and change prefix of a path for a buckets created via S3.
	Usage:
	If no options are given, shows current prefix:
	`s3 prefix`
	Command can be used to set default prefix. Prefix should be an existing path in file system under a non-root volume:
	`s3 prefix set [path]`
	Command can be used to reset default prefix. S3 should be disabled; use **s3 enabled no** command:
	`s3 prefix reset`
	Examples:
	`# s3 prefix set test`
	`S3 default prefix has been set successfully.`
	`# s3 prefix set test`
	`No change was needed.`
	`# s3 prefix set another_vol`
	`S3 default prefix has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 bucket add test-buc test`
	`Created bucket "test-buc" successfully.`
	`# s3 prefix set test`
	`Failed to set default prefix for S3: Prefix should not cross with existing buckets.`
	`# s3 prefix reset`
	`S3 default`
	`# s3 prefix reset`
	`No change was needed.`

Option	Description

`tls`	Manages certificates.
	Usage:
	If the “enabled” option is given with no additional parameters, displays current short status of TLS configuration (“enabled” or “disabled”):
	`s3 tls enabled ?yes|no?`
	- If used with “yes” parameter, enables TLS for S3.
	- If used with “no” parameter, disables TLS for S3.
	- When TLS is enabled, HTTPs is enforced and HTTP is no longer accessible.
	If the “details” option is given, displays current detailed status of TLS configuration:
	`s3 tls details`
	If the “default” option is given, resets CA certificate and sets server certificate and private key to those currently used in PanActive Manager GUI:
	`s3 tls default`
	If the “credentials” option is given, sets server certificate and private key to the given ones.
	If certificate and private key files paths are additionally given, sets server certificate and private key to the provided ones:
	`s3 tls credentials ?certificate_file_path private_key_file_path?`
	If the “cacert” option is given, sets CA certificate to the given one.
	If CA certificate file path is additionally given, sets CA certificate to the provided one. This cannot be done unless server certificate is set:
	`s3 tls cacert ?ca_certificate_file_path?`
	If the “reset” option is given, resets server certificate, private key, and CA certificate to null and disables TLS:
	`s3 tls reset`

Option	Description

	Examples:
	`# s3 tls enabled`
	`disabled`
	`# s3 tls enabled`
	`enabled`
	`# s3 tls enabled yes`
	`Enabling S3 TLS.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls enabled no`
	`Disabling S3 TLS.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls enabled yes`
	`Can not enable tls without server certificate and private key or default toggle set to true.`

Option	Description

`tls (cont.)`	`# s3 tls default`
	`Setting default server certificate and key.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls default`
	`No change was needed.`
	`# s3 tls details`
	`S3 TLS configuration:`
	`Enabled: yes`
	`Default settings: yes`
	`# s3 tls details`
	`S3 TLS configuration:`
	`Enabled: yes`
	`Default settings: no`
	`CA Certificate:`
	`File Name: S3CACertificate.pem`
	`SHA256: cca80ba21e2770e7e67e429992d531d7c82d519c6d5283e9321d2f3d0e5613ec`
	`Common Name: VDURA QA CA`
	`Fingerprint: D9:28:4F:32:2F:A2:E0:07:A2:83:0A:DA:06:0F:4D:2A:E3:79:47:BA:31:B2:3F:DE:7A:A1:36:5B:C7:89:8B:A1`
	`Valid From: 2023-06-29T14:35:47.000Z`
	`Valid To: 2024-06-28T14:35:47.000Z`
	`Server Certificate:`
	`File Name: S3ServerCertificate.pem`
	`SHA256: 24c9629c059208d5698f53e972db12d321c71fddcaf09a7f613ad99f11218546`
	`Common Name: pgh-realm-75-81`
	`Fingerprint: D6:F6:07:DB:09:21:BB:A5:5F:CD:94:F9:09:EB:32:54:24:C9:DC:D6:28:23:AD:5A:0B:DC:82:EF:5D:8C:40:56`
	`Valid From: 2023-06-29T14:35:48.000Z`

Option	Description

	`Valid To: 2024-06-28T14:35:48.000Z`
	`DNSs: pgh-realm-75-81`
	`IPs: 10.75.81.11 10.75.81.110 10.75.81.113`
	`Private Key:`
	`File Name: S3ServerKey.pem`
	`SHA256:`
	`c17479a61d483c1dacb4d043ede71eda511aa5d08e29`
	`7a4e6bb89978c4c207b4`

`tls (cont.)`	`# s3 tls credentials`
	`/tmp/s3_certs/S3ServerCertificate.pem`
	`/tmp/s3_certs/S3ServerKey.pem`
	`S3 Server Certificate and Private Key have been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls cacert`
	`/tmp/s3_certs/S3CACertificate.pem`
	`S3 CA Certificate has been set successfully.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls reset`
	`Resetting S3 TLS.`
	`S3 is currently enabled. S3 server will be restarted on all director nodes.`
	`# s3 tls reset`
	`No change was needed.`

Option	Description

`user`	Display, add, edit, and delete users.
	Usage:
	If the “list” option or no option is given, displays a list of all users:
	`s3 user ?list?`
	If the “add” option is given, creates a new user. You can specify system owner and system group with optional parameters “owner” and “group” with or without “uid:” and “gid:” prefixes, respectively:
	`s3 user add [username] ?owner [owner]? ? group [group]?`
	If the “edit password” option is given, provides a prompt to edit an existing user’s password:
	`s3 user edit password [username]`
	If the “edit mapping” option is given, edits UID and GID mappings.
	You can specify system owner and system group with parameters “owner” and “group” with or without “uid:” and “gid:” prefixes, respectively:
	`s3 user edit mapping [username] ?owner [owner]? ?group [group]?`
	- To reset owner or group to default - pass “default” as a value.
	- At least one parameter should be present (owner or group).
	If the “delete” option is given, deletes specified user:
	`s3 user delete [username]`
	If the “details” option is given, displays the details of a specific user:
	`s3 user details [username]`

Option	Description

`user (cont.)`	Examples:
	`# s3 user list`
	`There are currently no users.`
	`(Or will list names of users, if there are any.)`
	`# s3 user add test_user`
	`New test_user password:`
	`Confirm password:`
	`Added user "test_user" successfully.`
	`# s3 user edit test_user password`
	`New test_user password:`
	`Confirm password:`
	`User "test_user" updated successfully.`
	`# s3 user edit mapping test_user owner uid:111`
	`User "test_user" updated successfully.`
	`# s3 user edit mapping test_user group 222`
	`User "test_user" updated successfully.`
	`# s3 user edit mapping test_user group -5`
	`Failed to edit user: You must provide a GID as a positive integer number (0..4294967294) with or without "gid:" prefix.`
	`# s3 user edit mapping test_user owner -5`
	`Failed to edit user: You must provide a UID as a positive integer number (0..4294967294) with or without "uid:" prefix.`
	`# s3 user delete test_user`
	`Deleted user "test_user" successfully.`
	`# s3 user delete test_user`
	`Failed to delete user: No user with name "test_user" was found.`
	`No changes were applied.`
	`# s3 user details test_user`
	`User Name: test_user`
	`Groups: -`
	`Policy: -`
	`System Owner: uid:111`

Option	Description

	`System Group: gid:222`
	`# s3 user details test_user`
	`Failed to get user details: No user with name "test_user" was found.`
	`No changes were applied.`

samba

Description

A service mode command to manage the Samba server. Use these commands when Legacy CIFS Mode is disabled on the CIFS master node. See the VDURA ActiveStor User Guide for more information.

Active Directory (AD) is not supported in this release.

Syntax

```
samba [subcommand] ?[arg1] [arg2] ...?
```

Example

Clean up cached AD objects and identity mapping cache:

```
samba net cache flush
```

Print information about server domain:

```
samba net ads info
```

Print information about server machine account:

```
samba net ads status -U [AD user]
```

Change debug level:

```
samba smbcontrol smbd debug [log level 0-10]
samba smbcontrol winbindd debug [log level 0-10]
```

Reload configuration file:

```
samba smbcontrol smbd reload-config
samba smbcontrol winbindd reload-config
```

Disconnect a share:

```
samba smbcontrol smbd close-share [share name]
```

Disconnect a client:

```
samba smbcontrol smbd kill-client-ip [client ip]
```

Options

Option	Description
`net [subcmd] [options] ...`	Administration of Samba and remote CIFS servers. - **time show/set time** - **lookup**: Look up host names/IP addresses. - **user**: Manage users. - **group**: Manage groups. - **getlocalsid**: Get the SID for the local domain. - **setlocalsid**: Set the SID for the local domain. - **maxrid**: Display the maximum RID currently used. - **file**: Function on remote opened files. - **share**: Function on shares. - **session**: Manage sessions. - **rpc**: Run functions using RPC transport. - **rap**: Run functions using RAP transport. - **ads**: Run functions using ADS transport. - **cache**: Operate on the cache TDB file. - **service**: List/modify running services. - **status**: Print server status. - **help**: Print usage information. Valid targets (choose one; none defaults to localhost): - **-S** or **–server=[server]**: Server name - **-I** or **–ipaddress=[ipaddr]**: Target server address - **-w** or **–workgroup=[wg]**: Target workgroup or domain. Miscellaneous options: - **-p** or **–port=[port]**: Connection port on target - **-W** or **–myworkgroup=[wg]**: Client workgroup - **-d** or **–debuglevel=[level]**: Debug level (0-10) - **-n** or **–myname=[name]**: Client name - **-U** or **–user=[name]**: Username

Option	Description
	- **-s** or **–configfile=[path]**: Pathname of **smb.conf** file - **-l** or **–long**: Print full information. - **-V** or **–version**: Print samba version information. - **-P** or **–machine-pass**: Authenticate as machine account.

Option	Description
`samba smbcontrol ?[OPTION...]? [destination] [message-type] [parameters]`	Send messages to **smbd**, **nmbd**, or **winbindd** processes. - **-t** or **–timeout=TIMEOUT**: Set timeout value in seconds. Help options: - **-?** or **-help** : Print help message. - **–usage** : Print brief usage message. Common Samba options: - **-d** or **–debuglevel=DEBUGLEVEL**: Set debug level. - **-s** or **–configfile=CONFIGFILE** : Use alternate configuration file. - **-l** or **–log-basename=LOGFILEBASE**: Base name for log files - **-V** or **–version**: Print version. - **–option=name=value**: Set **smb.conf** option from command line. [destination] choices: - **nmbd** - **smbd** - **winbindd** - a **process ID** [message-type] choices: - **debug**: Set debug level. - **idmap**: Manipulate idmap cache. - **force-election**: Force a browse election. - **ping**: Elicit a response. - **profile** - **inject**: Inject a fatal signal into a running **smbd**. - **stacktrace**: Print a stack trace of a daemon. - **profilelevel** - **debuglevel**: Print current **debuglevels**. - **printnotify**: Send a print notify message. - **close-share**: Forcibly disconnect a share. - **kill-client-ip**: Forcibly disconnect a client with a specific IP address.

Option	Description
	- **ip-dropped**: Tell **winbind** that an IP got dropped. - **lockretry**: Force a blocking lock retry. - **brl-revalidate**: Revalidate all **brl** entries. - **pool-usage**: Print **talloc** memory usage. - **ringbuf-log**: Print **ringbuf** log. - **dmalloc-mark** - **dmalloc-log-changed** - **shutdown** : Shutdown daemon - **drvupgrade**: Notify a printer driver has changed. - **reload-config**: Force **smbd** or **winbindd** to reload config file. - r**eload-printers**: Force **smbd** to reload printers. - **nodestatus**: Ask **nmbd** to do a node status request. - **online**: Ask **winbind** to go into online state. - **offline**: Ask **winbind** to go into offline state. - **onlinestatus**: Request **winbind** online status. - **dump-event-list**: Dump event list. - **validate-cache**: Validate **winbind** credential cache. - **dump-domain-list**: Dump **winbind** domain list. - **notify-cleanup** - **num-children**: Print number of smbd child processes. - **msg-cleanup** - **noop**: Do nothing.

Option	Description
`smbstatus ?[OPTION...]?`	Display active connections, shares, and locks for CIFS. - **-p** or **–processes**: Show processes only. - **-v** or **–verbose**: Be verbose. - **-L** or **–locks**: Show locks only. - **-S** or **–shares**: Show shares only. - **-N** or **–notify**: Show notifies. - **-b** or **–brief**: Be brief. - **-P** or **–profile**: Do profiling. - **-R** or **–profile-rates** : Show call rates. - **-B** or **–byterange**: Include byte range locks. - **-n** or **–numeric** : numeric uid/gid - **-f** or **–fast skip**: Check if processes still exist. Common Samba options: - **-d** or **–debuglevel=DEBUGLEVEL**: Set debug level. - **-s** or **–configfile=CONFIGFILE** : Use alternate configuration file. - **-l** or **–log-basename=LOGFILEBASE** : base name for log files - **-V** or **–version**: Print version. - **–option=name=value**: Set s**mb.conf** option from command line.

Option	Description
`testparm ?[OPTION...]? ?[config file]?`	Check Samba configuration file for correctness. - **-s** or **–suppress-prompt**: Suppress prompt for enter. - **-v** or **–verbose**: Also display default options. - **-l** or **–skip-logic-checks**: Skip the global checks. - **–show-all-parameters**: Show the parameters, type, possible values. - **–parameter-name=STRING**: Limit **testparm** to a named parameter. - **–section-name=STRING**: Limit **testparm** to a named section. Common Samba options: - **-V** or **–version**: Print version. - **-d** or **–debuglevel=DEBUGLEVEL**: Set debug level. - **–option=name=value**: Set **smb.conf** option from command line.

Option	Description
`wbinfo ?[OPTION..]?`	- **-u** or **–domain-users**: List all domain users. - **-g** or **–domain-groups**: List all domain groups. - **-N** or **–WINS-by-name=NETBIOS-NAME**: Convert NetBIOS name to IP. - **-I** or **–WINS-by-ip=IP**: Convert IP address to NETBIOS name. - **-n** or **–name-to-sid=NAME**: Convert name to sid. - **-d** or **–sid-to-name=SID**: Convert sid to name. - **–sid-to-fullname=SID**: Convert sid to fullname. - **-R** or **–lookup-rids=RIDs**: Convert RIDs to names. - **–lookup-sids=Sid-List**: Convert SIDs to types and names. - **-U** or **–uid-to-sid=UID**: Convert uid to sid. - **-G** or **–gid-to-sid=GID**: Convert gid to sid. - **-S** or **–sid-to-uid=SID**: Convert sid to uid. - **-Y** or **–sid-to-gid=SID**: Convert sid to gid. - **–allocate-uid**: Get a new UID out of idmap. - **–allocate-gid**: Get a new GID out of idmap. - **–set-uid-mapping=UID,SID**: Create or modify uid to sid mapping in idmap. - **–set-gid-mapping=GID,SID**: Create or modify gid to sid mapping in idmap. - **–remove-uid-mapping=UID,SID**: Remove uid to sid mapping in idmap. - **–remove-gid-mapping=GID,SID**: Remove gid to sid mapping in idmap. - **–sids-to-unix-ids=Sid-List**: Translate SIDs to Unix IDs. - **–unix-ids-to-sids=ID-List (u[num] g[num])**: Translate Unix ids to sids. - **-t** or **–check-secret**: Check shared secret. - **-c** or **–change-secret**: Change shared secret.

Option	Description
	- **-P** or **–ping-dc**: Check the NETLOGON connection. - **-m** or **–trusted-domains**: List trusted domains. - **–all-domains**: List all domains (trusted and own domain). - **–own-domain**: List own domain. - **–sequence**: Deprecated command, see –online-status. - **–online-status**: Show whether domains are marked as online or offline. - **-D** or **–domain-info=STRING**: Show most of the info we have about the domain. - **-i** or **–user-info=USER**: Get user info. - **–uid-info=UID**: Get user info from uid. - **–group-info=GROUP**: Get group info. - **–user-sidinfo=SID**: Get user info from sid. - **–gid-info=GID**: Get group info from gid. - **-r** or **–user-groups=USER**: Get user groups. - **–user-domgroups=SID**: Get user domain groups. - **–sid-aliases=SID**: Get sid aliases. - **–user-sids=SID**: Get user group sids for user SID. - **-a** or **–authenticate=user%password**: Authenticate user. - **–pam-logon=user%password**: Execute a pam logon equivalent. - **–logoff**: Log off user. - **–logoff-user=STRING**: Username to log off. - **–logoff-uid=INT**: uid to log off.

Option	Description
`continued: wbinfo ?[OPTION..]?`	- **–set-auth-user=user%password**: Store user and password used by winbindd (root only). - **–ccache-save=user%password**: Store user and password for ccache operation. - **–getdcname=domainname**: Get a DC name for a foreign domain. - **–dsgetdcname=domainname**: Find a DC for a domain. - **–dc-info=domainname**: Find the currently known DCs. - -**-get-auth-user**: Retrieve user and password used by winbindd (root only). - **–p** or **–ping**: Ping **winbindd** to see if it is alive. - **–domain=domain**: Define to the domain to restrict operation. - **–K** or **–krb5auth=user%password**: Authenticate user using Kerberos. - **–krb5ccname=krb5ccname**: Authenticate user using Kerberos and specific credential cache type. - **–separator**: Get the active **winbind** separator. - **–verbose**: Print additional information per command. - **–change-user-password=STRING**: Change the password for a user. - **–ntlmv**: Use NTLMv1 cryptography for user authentication. - **–ntlmv2**: Use NTLMv2 cryptography for user authentication. - **–lanman**: Use lanman cryptography for user authentication.

scp

Description

Secure copy is a remote file copy program that copies files between hosts on a network. It uses **ssh(1)** for data transfer and authentication and provides the same security as **ssh(1)**. Unlike **rcp(1)**, **scp** will ask for passwords or passphrases if they are needed for authentication.

A file name may contain a host and user specification to show that the file is to be copied to/from that host. Copies between two remote hosts are permitted.

Syntax

```
scp [-pqrvBC46] [-F ssh_config] [-S program] [-P port] [-c cipher] [-i identity_file] [-o ssh_option] [[user@]host1:]file1 [...] [[user@]host2:]file2
```

Options

Option	Description
`-c cipher`	Select the cipher to use for encrypting the data transfer. This option directly passed to **ssh(1)**.
`-i identity_file`	Select the file from which the identity (private key) for RSA authentication is read. This option directly passed to **ssh(1)**.
`-p`	Preserve modification times, access times, and modes from the original file.
`-r`	Recursively copy entire directories.
`-v`	Verbose mode. Cause **scp** and **ssh(1)** to print debugging messages about their progress. This is helpful in debugging connection, authentication, and configuration problems.
`-B`	Select batch mode; prevent asking for passwords or passphrases.
`-q`	Disable the progress meter.

Option	Description
`-C`	Pass the -**C** flag to **ssh(1)** to enable compression.
`-F ssh_config`	Specify an alternative per-user configuration file for **ssh**. This option is directly passed to **ssh(1)**.
`-P port`	Specify the port to connect to on the remote host. Note that this option is written with a capital P. The -**p** option is reserved for preserving the times and modes of the file in **rcp(1)**.
`-S program`	Name of the program to use for the encrypted connection. The program must understand **ssh(1)** options.
`-o ssh_option`	Pass options to **ssh** in the format used in **ssh_config(5)**. This is useful for specifying options for which there are no separate **scp** commandline flag.
`-4`	Force **scp** to use IPv4 addresses only.
`-6`	Force **scp** to use IPv6 addresses only.

secure

Description

Displays security status or changes the security setting.

Syntax

```
secure ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
none	Display whether secure access is required.
`gencert`	Generate a new X509 certificate for the system.
`warnfailedlogin ?yes|no?`	Display or set whether or not to raise warning alerts due to excessive login failures. The system automatically detects excessive ssh, HTTP, and HTTPS login failures on a per-host basis and reports the incidents as info items in the event log. This control tells the system whether to additionally raise a system-wide warning alert that triggers email and pager messages.

sha1

Description

Calculates a SHA-1 message-digest fingerprint (checksum) for a file. The utility takes as input a message of arbitrary length and produces as output a “fingerprint” or “message digest” of the input. The hexadecimal checksum of each file listed on the command line is printed. Also see the **md5** command.

Syntax

```
sha1 file [...]
```

smtp

Description

Displays, sets, or tests the SMTP email server and from address.

Syntax

```
smtp ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Send a test message to bob@mail.com using the saved SMTP server:

```
smtp test bob@mail.com
```

Send a test message to bob@mail.com via SMTP server 192.168.1.120:

```
smtp test -server 192.168.1.120 bob@mail.com
```

Options

Option	Description
none	Display the SMTP email server and the from address that all outbound email will use.
`authpass`	Interactive command to store the password. Only needed if the SMTP server requires authentication.
`authuser ?[username]?`	Display or set the username. The username is only needed if the SMTP server requires authentication. Use “” to clear the username.

Option	Description
`from ?[from-address]?`	Display or set the from address for all outbound email. If a custom from address is specified, all outbound email will use this from address. Otherwise, the default from address **PanActiveManager@[domain]** will be used. The domain is the first DNS domain returned by **dns** domains. If none is set, then the IP address of the node from which email is sent is used in place of the DNS domain. Use “” to clear the custom from address and return to using the system default from address.
`port ?[port]?`	Display or set the SMTP server port. Use “” to clear the custom port configuration and use the standard SMTP server port.
`server ?[server]?`	Set the SMTP server or display the current server. Use an empty string (“”) to clear the value.
`test ?-from [from-address]? ?-server [server]? ?-port [port-number]? ?-authuser [username]? ?[address1] ...?`	Send a sample email message to the given email addresses. If no addresses are given, send to the saved admin email address. SMTP server details may be specified with the -**server** , -**port** , and -**authuser** options. If no such options are given, use the saved SMTP server, port, username/password, and from address. If the -**authuser** option is given, the user will be prompted to enter the password interactively.

snmp

Description

Configures SNMP support. VDURA currently supports SNMP v1.

Syntax

```
snmp [subcommand] ?[arg1] [arg2] ...?
```

Example

Set the syscontact field to admin@bigcompany.com in the SNMP MIB:

```
snmp contact admin@bigcompany.com
```

Send a test trap to 192.1.2.3:

```
snmp send_trap 192.1.2.3
```

Configure multiple trap receivers:

```
snmp trap_receivers 192.1.2.3 192.168.1.3
```

Options

Option	Description
`contact ?value?`	Display or set snmp contact field.

Option	Description
`debug ?value?`	Service mode command to enable snmp debug logging. Enter a value to enable or disable snmp debug logging: 0 = disable logging 1 = enable standard logging 2 = detailed logging with transmit and receive information The logs may be collected with the **collect_logs** command.
`download_mibs [user]@[host]:[file]`	This command will download a zipped TGZ archive file which will contain all MIB definition files. In addition to the MIB definition files, a PANASAS-README file which defines the MIB load order will also be included. This command will prompt for a password and use scp to do the remote copy.
`enable ?yes|no?`	Display or change whether the **snmp** agent is enabled.
`list`	Display a list of **snmp** parameters.
`location ?value?`	Display or set **snmp location** field.
`read_community ?value?`	Display or set **snmp read community** field.
`send_trap [ip]`	Send a test trap to the configured receivers.
`trap_community ?value?`	Display or set **snmp trap community** field.
`trap_receivers ?value?`	Display or set snmp trap receivers field. Enter a space separated list of IP addresses and/or fully qualified hostnames to receive traps.

sort

Description

Writes sorted concatenation of all file(s) to standard output.

POS is F[.C][OPTS], where F is the field number and C the character position in the field. OPTS is one or more single-letter ordering options, which override global ordering options for that key. If no key is given, use the entire line as the key.

SIZE may be followed by the following multiplicative suffixes: % 1% of memory, b 1, K 1024 (default), and so on for M, G, T, P, E, Z, Y.

With no FILE or when FILE is -, read standard input.

The locale specified by the environment affects sort order. Set LC_ALL=C to get the traditional sort order that uses native byte values.

Mandatory long option arguments are also mandatory for short options.

Syntax

```
sort [OPTION]... [file ...]
```

Options

Option	Description
`-b, --ignore-leading-blanks`	Ignore leading blanks.
`-d, --dictionary-order`	Consider only blanks and alphanumeric characters.
`-f, --ignore-case`	Fold lower case to upper case characters.
`-g, --general-numeric-sort`	Compare according to general numerical value.
`-i, --ignore-nonprinting`	Consider only printable characters.
`-M, --month-sort`	Compare (unknown) < JAN < ... < DEC.
`-n, --numeric-sort`	Compare according to string numerical value.

Option	Description
`-r, --reverse`	Reverse the result of comparisons.
`-c, --check`	Check whether input is sorted; do not sort.
`-k, --key=POS1[,POS2]`	Start a key at POS1, end it at POS 2 (origin 1).
`-m, --merge`	Merge already sorted files; do not sort.
`-o, --output=FILE`	Write result to FILE instead of standard output.
`-s, --stable`	Stabilize sort by disabling last-resort comparison.
`-S, --buffer-size=SIZE`	Use SIZE for main memory buffer.
`-t, --field-separator=SEP`	Use SEP instead of non-blank to blank transition.
`-T, --temporary-directory=DIR`	Use DIR for temporaries, not $TMPDIR or /tmp; multiple options specify multiple directories.
`-u, --unique`	With -**c**, check for strict ordering; without -**c**, output only the first of an equal run.
`-z, --zero-terminated`	End lines with 0 byte, not newline.
`--help`	Display this help and exit.
`--version`	Provide version information and exit.

storageset

Description

Manages StorageSets.

Syntax

```
storageset ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Pass the name of an enclosure named VCH-3 as an argument to create a StorageSet named Set 2 containing all VPODs in the VCH-3 enclosure:

```
storageset create "Set 2" VCH-3
```

Pass the name of an individual VPOD named VCH-3,2 as an argument to create a StorageSet named Set 2 containing all VPODs in the same enclosure as VCH-3,2:

```
storageset create "Set 2" "VCH-3,2"
```

Pass the name of an individual VDURA Data Platform named VCH-3,1 as an argument to create a StorageSet named Set 2 containing all VPODs in the same enclosure as VCH-3,1:

```
storageset create "Set 2" "VCH-3,1"
```

Pass the name of an enclosure named VCH-3 as an argument to create a StorageSet named Set 2 containing all VPODs in the VCH-3 enclosure with three spare VPODs:

```
storageset create -spare 3 "Set 2" VCH-3
```

Pass the names of two enclosures named VCH-2 and VCH-3 as arguments to create a StorageSet named Set 2 containing all VPODs in the VCH-2 and VCH-3 enclosures:

```
storageset create "Set 2" VCH-2 VCH-3
```

Pass the name of two VPODs named VCH-2,2 and VCH-3,2 as arguments to create a StorageSet named Set 2 containing all VPODs in the same enclosures as VCH-2,2 and VCH-3,2:

```
storageset create "Set 2" "VCH-2,2" "VCH-3,2"
```

Add all unassigned VPODs to StorageSet Set 1, incrementing the StorageSet spare count by 1:

```
storageset addstorage -addspare 1 "Set 1"
```

Add all unassigned VPODs in enclosure VCH-2 to StorageSet Set 1 for purposes of disk migration:

```
storageset addstorage -migration "Set 1" VCH-2
```

Note that no matter if you specify either the enclosure, VDURA Data Platform, or VPOD name, all unassigned VPODs in the enclosure will be added.

Merge StorageSet Set B into Set A and have a combined spare node count of two:

```
storageset merge "Set B" "Set A" spare 2
```

You will then have a single StorageSet named Set A with a spare count of two that contains all volumes from both Set A and Set B.

List all enclosures in a StorageSet named Set 1:

```
storageset enclosure "Set 1"
```

Options

Option	Description
`(none)`	Show information about all StorageSets in the system, same as **storageset list** . **Note:** To use special characters in StorageSet names, put a backslash before the character.

Option	Description
`addstorage -migration| -addspare [num-VPODs]? [storageset_name] ?[component] ...?`	Add unassigned VPODs represented by the given components to the given StorageSet. Components can be enclosures, VDURA Data Platforms, or individual VPODs. If no components are given, all unassigned VPODs will be added to the StorageSet. The admin will be prompted with one or more questions to ensure the outcome is desired. - The **-addspare** option allows the admin to increase the spare VPOD count for the StorageSet while adding VPODs to expand the StorageSet capacity. If no additional spare VPOD count is given, the requested spare space reservation for the StorageSet will not be changed. This option cannot be used with the **-migration** option. - The **-migration** option increases the requested spare VPOD count of the StorageSet by the number of VPODs being added. The aim is to subsequently drain all VPODs in another enclosure that is already part of the same StorageSet. This option cannot be used with the **-addspare** option. - You must supply either the **-addspare** or **-migration** option, but both may not be used together. Use the **-f** option to execute the command with no confirmation prompt and minimal explanatory text.
`capacity-balancing [storageset-name] ? [on|off]?`	Display current capacity balancing setting or turns on or off capacity balancing on the StorageSet.

Option	Description
`capacity-warning ?[option]?`	Display or set StorageSet capacity warning values. Omitting the **option** displays all current settings. The **option** can be any of the following, where string is a percentage integer value from 1 to 99: - **–set-pmm-blade-near-full-thld=** string : Node nearly full threshold - **–set-pmm-blade-near-full-incr=** string : Node nearly full increment - **–set-pmm-storageset-near-full-thld=** string : StorageSet nearly full threshold - **–set-pmm-storageset-near-full-incr=** string : StorageSet nearly full increment - **–set-pmm-blade-ssd-near-full-thld**=string: Node SSD nearly full threshold - **–set-pmm-blade-ssd-near-full-incr**=string: Node SSD nearly full increment - **–set-pmm-bset-ssd-near-full-thld**=string: StorageSet SSD nearly full threshold - **–set-pmm-bset-ssd-near-full-incr**=string: StorageSet SSD nearly full increment The threshold (**thld**) settings are percentages between 1 and 99, where a warning will be sent when the capacity reaches the specified percentage. Values that are not yet set may be reported as “missing”; this is normal and can be changed by setting a value with one of the options above. The increment (*incr*) percentages are the values at which a subsequent warning will be sent. For example, a threshold value of 95 percent and a increment value of 2 percent will result in capacity warnings at 95, 97, and 99 percent. Changing capacity warning thresholds is a service mode command. This should only be performed under the direction of VDURA Technical Support. Note that StorageSet nearly-

Option	Description
	full threshold triggers the automatic deletion of snapshots if the **snapshot diskfull** feature is enabled.
`create ?-spare [num-VPODs]? ?-mlec enable|disable? [storageset-name] [component] ?[component] ...?`	Create a new StorageSet and populate it with unassigned VPODs represented by the given components. Components can be enclosures, VDURA Data Platforms, or individual VPODs, referred to by name or serial. Be aware that when **storageset create** is run the first time immediately after installing realm, you must use either VDURA Data Platform or enclosure names since no VPODs have been created yet. The StorageSet name is limited to 30 characters. Use the **-spare** option to specify the number of VPODs to be reserved for reconstruction in the new StorageSet. The default spare VPOD count is **2**. You can specify either the name or serial number of an enclosure or VDURA Data Platform when creating a StorageSet or adding to a StorageSet. You can also specify either the name or serial number of a single VPOD when creating a StorageSet or adding to a StorageSet. All VPODs from an enclosure or VDURA Data Platform will be added to the specified StorageSet. It is not possible for VPODs from a single enclosure or VDURA Data Platform to belong to more than one StorageSet. Use the **-mlec** option to enable or disable Multi-Level Erasure Coding feature for the given StorageSet. This can only be set when a StorageSet is created and cannot be changed after. MLEC is disabled by default.

Option	Description
`delete [storageset-name]`	Delete a StorageSet, returning all VPODs to the **Not Activated** state. Before deleting a StorageSet, all volumes in the given StorageSet must already be deleted.
`detail [storageset-name]`	Display configuration and space usage information for a StorageSet.
`enclosure [storageset-name]`	List enclosures with storage in the given StorageSet.
`errors ?[storageset-name]? ?show [max]|all?`	Shows what errors are currently active for all StorageSets in the system, corresponding to events in eventlog. With options, displays errors for the specified StorageSet. If *show* is not included to specify the maximum number to list, five errors are displayed.
`list ?name [pattern]? ?show [max]|all? ?sort name|total? ?order increasing|decreasing? ?capacity-balancing? ?id? ?storagecapacity? ?allcolumns?`	List all StorageSets in the system. A maximum of 50 StorageSets is displayed by default. The options are as follows: - **name**: Specify a glob pattern of StorageSet names to match. - **show**: Specify a different maximum. - **sort**: Sort the results by name or by total capacity. - **order**: Indicate whether to sort in increasing or decreasing order. - **capacity-balancing** : Display capacity balancing information for each StorageSet. - **id** : Display the internal numeric ID for each StorageSet. - **storagecapacity**: Display a summary of node counts and capacities. - **allcolumns**: Include all options that give additional data.

Option	Description
`merge [source storageset name] [target storageset name] ?spare [num-nodes]?`	Combine two existing StorageSets into a single StorageSet. The resultant combined StorageSet will contain all of the volumes of both of the preexisting StorageSets. The resultant combined StorageSet will keep the StorageSet name and vertical parity default of the target StorageSet. By default, the resultant StorageSet will have a requested spare count equal to the maximum of the two original StorageSets. This default can be overridden using the spare option. The override value must be greater than or equal to the default; it cannot be set lower. Use the **-f** option to execute the command with no confirmation prompt and minimal explanatory text.
`mlec`	Shows the state of Multi-Level Erasure Coding (MLEC) for a StorageSet. MLEC provides improved fault tolerance for user data. When enabled, each VPOD in a StorageSet can survive a single HDD failure without loss of access to data. This feature reduces usable capacity by one HDD per VPOD, but significantly improves the durability of data. MLEC can only be set during StorageSet creation by adding the **-mlec enable|disable** flag, with one exception described under the update mlec option. See the **storageset create** command for more details.
`update name [old-name] [new-name]`	Update the name of a StorageSet. The StorageSet name is limited to 30 characters.

Option	Description
`update spare [nodes] [storageset-name]`	Update number of nodes to be reserved for reconstruction. The number of spare VPODs can be increased by updating the number of spares and adding new VPODs to the StorageSet. It is not allowed to set the number of spares to zero.
`update mlec enable|disable`	Change a StorageSet's Multi-Level Erasure Coding (MLEC) feature state. This is the one exception to the rule that MLEC can only be set when a StorageSet is created. The **update mlec** option is used to correct mistakes made during the OOBE process. This command can only be used on the default StorageSet immediately after OOBE and only before VDURA Data Platforms are added to the default StorageSet. MLEC is disabled by default.

support

Description

Configures service mode and automated communication with MyVDURA.

Syntax

```
support ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description
`customerid ?id?`	View or edit the customer ID.

`panactivelink`
`panactivelink ?(yes ?-noweb?)|no?`
`panactivelink cc list`
`panactivelink test ?[address1] ...?`
`panactivelink cc add ?[address1] ...?`
`panactivelink cc delete ?[address1] ...?`
`panactivelink cc clear`
	If no option is given, tell whether PanActive Link is enabled.
	If **yes** or **no** is given, then enable or disable PanActive Link accordingly. Use the additional -**noweb** option to enable PanActive Link without web access.
	If the **cc** option is given, view or modify the cc list for outgoing status email.
	If the **test** option is given, generate and send a sample status email to the given email addresses. If none are given, use the saved cc list.

`servicemode ?yes|no?`	Use this command to switch PanCLI to its debugging service mode. The service mode makes debug commands available for use by or with advice from support personnel. Use of the service mode should only be under the direction of VDURA Technical Support to guarantee safe usage.

Option	Description
`user enabled ?yes|no?` `user name`	Enable/disable a temporary support account. The name is chosen by the system. You can view it by using the **name** option.
`webproxy host ?[host]?` `webproxy port ?[port]?`	Display and set the secure web proxy host and port. This proxy server will be used to make outbound HTTPS requests to VDURA servers if you have chosen to enable PanActive Link. Use “” to clear the web proxy host or port.

syscontrol

Description

Sets system and hardware controls.

Syntax

```
syscontrol [subcommand] ?[arg1] [arg2] ...?
```

Options

Option	Description
`activate all|[node1|shelf1] ?[node2|shelf2] ...?`	This command activates Directors. This command is only needed for activation of components in an enclosure with no storage and for stand-alone box components. The system automatically activates any Director that shares an enclosure with storage components that belong to a StorageSet. Activation of a Director means that it will take on volume service activity, NFS/CIFS client activity, and be eligible for membership in the cluster configuration replication set (repset). Use the **all** keyword to activate all inactive Directors. Use the following command to find inactive Directors: `sysmap nodes directors | grep "Not Activated"` Use **storageset addstorage** or **storageset create** to activate storage (Storage Nodes and VPODs).

Option	Description
`capacity-balance-speed ?(default|[1-7])?`	This command controls the system-wide speed of capacity balancing, drain, and data migration activity. Increased speed will impact system performance in favor of completing drain/data migration and balancing capacity more quickly. Decreased speed will slow down the drain/migration and balancing activity in favor of improved system performance. If called with no arguments, this command displays the current speed: 1-7, where 1 is slowest and 7 is highest. Because of the potentially large performance impact on the system, changing the capacity balancing and drain speed requires service mode operation and should only be taken with the assistance of VDURA Technical Support.

`enclosure [serial]|[name] led on|off`
`enclosure [serial]|[name] rename [new-name]`
`enclosure [serial]|[name] remove|reinit`
`enclosure [serial]|[name]`
	Make changes on an enclosure-wide basis.
	- The **led** option requests that all nodes in the enclosure start or stop blinking LEDs.
	- The **rename** option changes the name of the given enclosure.
	- The **remove** option removes all nodes in the enclosure. This action will only have an effect if the all nodes are in the correct state for safe removal.
	- The **reinit** option reinitializes all nodes in the enclosure. This action will only have an effect if all nodes are in the correct state for safe reinitialization. When removing or reinitializing an enclosure, one should consider reducing the spare count for the enclosure's former StorageSet by the number of Storage Nodes in that enclosure.
	Use the **-f** option to execute the command with no confirmation prompt and minimal explanatory text.

Option	Description
`encryption ?enable|disable?`	Turn on/off system-wide end-to-end (E2EE) encryption. The system-wide encryption setting defines default encryption setting for new volumes. Turning encryption on does not affect already created volumes; it is a default setting for new volumes. However, if there are encrypted volumes in the system, it is not possible to turn this setting off. This encryption cannot be turned on if NFS, CIFS, or S3 are enabled. A valid KMIP configuration should be set as a prerequisite to switching encryption on. Enable system-wide volumes encryption: `syscontrol encryption enable` Disable system-wide volumes encryption: `syscontrol encryption disable` View current encryption setting: `syscontrol encryption`
`lock`	Lock the system. New nodes will not join unless the enclosure contains at least one activated node.
`lockmgr ?list ?[volume_name]??` `lockmgr break [volume_name] [handleID]`	Display or break file locks. To list locks for a particular volume, specify a volume name.

Option	Description
`mlec_default ?enable|disable?`	Manage the system-wide Multi-Level Erasure Coding (MLEC) default state. MLEC provides improved fault-tolerance for user data. When enabled, each VPOD in the StorageSet can survive a single HDD failure without loss of access to data. This feature reduces usable capacity by one HDD per VPOD, but significantly improves the durability of data. MLEC is only available on Capacity Expansion Storage Nodes using JBODs. This command displays or sets MLEC default value for new StorageSets.

Option	Description

`node [node] alias ?enable|disable?`
`node [node] drive [drive-id] led ? normal|locate?`
`node [node] led on|off`
`node [node] halt`
`node [node] jettison`
`node [node] rename [new-name]`
`node [node] reboot`
`node [node] reconstruct`
`node [node] remove`
`node [node] reinit ?secure? ?keep_vpods?`
`node [node] revive`
	Reconstruct, reboot, halt, remove, reinitialize, or revive a node, or toggle its LEDs. To specify a node, give [enclosure-name],[slot], IP address, hardware serial, or data serial.
	- The **alias** option specifies whether a Director can have IP aliases assigned to it. Specifying **enable** or **disable** changes the setting for the given node. By default, all Directors can receive aliases. It may be useful to disable alias assignment for a handful of Directors as they are busy with other specific system tasks. Use the **netconfig alias** command to manage aliases on a system-wide basis.
	- The **drive** option allows the LED on a drive to be turned on. The **drive-id** must be a number between 1 and 6, inclusive. The possible drive LED states are locate and normal. Normal blinks the LED only if the drive encounters a problem. Locate blinks the LED quickly for identification. Normal is the default state for all drives.
	- The **led** option starts or stops flashing of LEDs on the given node. This feature is available for all nodes that have their own IP address.
	- The **halt** option will halt the operating system on a node and power it down. The user will have to manually power-cycle the node to bring it back up. Halting a node will make the node unavailable to the system.
	WARNING: For VPODs, the **halt** option turns off the VPOD itself; however, it should be noted that the VPOD Manager has the capability to automatically turn it back on after some time. This is because the VPOD Manager monitors the state of the VPOD and can restart it if necessary to ensure the stable operation of the system. To completely and permanently stop the VPOD without subsequent automatic startup, it is

Option	Description
	recommended to use the special command **vpod_manager_cli**, which should be executed directly on the VDURA Data Platform. This command will allow you to properly stop the VPOD and prevent its automatic restart by the VPOD Manager. - The **jettison** option forces the system to drop a node even if the fault tolerance level has been exceeded and the system may not be able to reconstruct all of the data on it. The user will have to manually rescue the volume data. This option is available in service mode only. - The **rename** option changes the name of the given server. This feature is only available on some types of hardware. - The **reboot** option will reboot a node. This will make the node unavailable for a short while. - The **reconstruct** option tells the system to start reconstructing data for that node based on parity. This will remove the node from the system immediately. The reconstruction will continue on a per-volume basis in the background until all data for that node is restored. This may take several hours to complete. Use the command **volume list storageset [storageset-name]** to see which volumes are currently getting reconstructed. Use the **volume detail** command to monitor reconstruction progress on individual volumes. - The **remove** option tells the system to stop tracking and displaying the given failed node. This action is only allowed on nodes that have no PanFS data and Directors that are neither primary volume service owners nor in the replication set (repset). There are some unusual conditions under which a node is still known to the system but the user interface does not

Option	Description
	display the node. Such nodes can be removed via data serial number. - The **reinit** option reinitializes the given node. This action is only allowed on nodes that have no PanFS data and Directors that are neither primary volume service owners nor in the replication set (repset). - If the **secure** option is given, reinitialize the system using secure disk erase procedure. The **secure** option is applied only locally on the node that needs to be reinited. - The **keep_vpods** option is used to preserve all VPODs and their data on the node and is applicable only to VDURA Data Platforms. - Note that **keep_vpods** and **secure** options are mutually exclusive. The **reinit** subcommand can be invoked locally on any node without specifying a serial or IP address. In this case, following several interactive confirmations the node on which the request is launched will be reinitialized. The local **syscontrol node reinit** command is only available in service mode, as it should only be performed under the direction of VDURA Technical Support. - The **revive** option cancels the reconstruction of a Storage Node and brings the node back online. Once the node returns to service, the system will automatically clean up the stale data on the node which have been reconstructed. This option is available in service mode only.
`offline`	Brings the system offline. PanFS clients will block and retry I/O requests until the system is brought back online.

Option	Description
`offline-unavailable`	This option is only available in service mode. If you choose this option, PanFS clients will receive I/O Errors until the system is brought back online.
`online`	Brings the system online.
`realmmgr ?[list]?` `realmmgr [add|remove] [node]`	Manage the list of Directors in the failover replication set (repset). To specify a node, give [enclosure-name],[slot], IP address, hardware serial, or data serial.
`reboot`	Reboot the entire system.
`reconstruction timeout ?[seconds]|default?`	The reconstruction timeout is the number of seconds between the time when the system discovers a VPOD is not responding and the time when the system automatically removes that VPOD from the StorageSet and begins reconstructing it. If the timeout option is given with no seconds value, the reconstruction timeout that the system is currently configured to use is displayed. If the timeout option is given with an additional seconds value, the timeout value is set to this number of seconds. The **default** value is 360 seconds. A timeout of 0 means immediate reconstruction when a VPOD is detected as missing and is considered to be an unsafe option. Note that changes to this timeout will not affect any VPOD that is already missing and pending reconstruction. This is a service mode command and should only be used under the direction of company service personnel to ensure safe usage.

Option	Description
`reinit ?secure? ?keepnetworking?`	Delete all user data and the system configuration, and reboot the system. This command is available in service mode only. - If the **keepnetworking** option is given, then the system IP, default router, and netmask are preserved and console access is not needed to perform installation. - If the **secure** option is given, reinitialize the system using secure disk erase procedure.
`shutdown`	Shutdown the entire system.
`unlock`	Unlock the system, allowing new nodes to join the system.

sysctl

Description

Service mode command to manage **sysctl** values. This command should only be used under the direction of VDURA Technical Support.

Syntax

```
sysctl ?[args]?
```

Options

Option	Description
none	List the **sysctl** key-value pairs that have been permanently overridden.
`all [node-type] ?[key] [value]?`	Manage permanent system-wide **sysctl** overrides for all nodes. Valid node types include all Directors and Storage Nodes. When called with no key and value, list the **sysctl** key-value pairs that have been permanently overridden on the requested subset of nodes and optional architecture. When a key and value are given, the key-value pair is saved permanently and immediately deployed on all applicable nodes, based on the node type. Note that some **sysctl** keys can only be updated at boot time, so your change might not take noticeable effect until after the next reboot. New nodes added to the system after the setting was saved will automatically get the saved setting. Note that the system will not validate the **sysctl** key. Use the **sysctl clear** command to clear the override.

Option	Description
`clear [node-type] ?[key]?`	Service mode command to clear **sysctl** overrides. Valid node types include all Directors and Storage Nodes. Note that this change takes effect locally on each node upon reboot. Immediate system reboot following a clear request is recommended.
`directorblade ?[key] [value]?`	Manage permanent system-wide **sysctl** overrides for all Director blades. Refer to **sysctl all** for more information.
`director ?[key] [value]?`	Manage permanent system-wide **sysctl** overrides for all Directors. Refer to **sysctl all** for more information.
`local [key] ?[value]?`	Service mode command to manage local **sysctl** values. This command can be called on any node locally and only takes effect on the local node. - When called with no value, display the current value of the given **sysctl** key on the local node. - When a key and value are given, update the value of the given **sysctl** key on the local node. **Warning**: Changes made using this command only take effect on the local node. Once this node is rebooted, the change will no longer apply.
`storageblade ?[key] [value]?`	Manage permanent system-wide **sysctl** overrides for all Storage Nodes. Refer to **sysctl all** for more information.

sysmap

Description

Displays the system hardware map.

Syntax

```
sysmap [subcommand] ?[arg1] [arg2] ...?
```

Options

Option	Description
`assets`	Display a listing of all hardware components in the system.
`containers ?staticlinkagg? ?allcolumns?`	Display a list of all containers. If the **staticlinkagg** option is given, display a column indicating whether static link aggregation is enabled for each container. Use the **allcolumns** option as a shortcut to include all options that give additional data.
`dimms`	Display DIMM info for each node in the system. This includes node name and type, DIMM name, DIMM serial, and DIMM part number.
`drives ?encryption? ?lun?`	Display drive info for each drive in the system. - If no option is given, display drive info for each drive in the system, including node name and type, device, capacity, drive serial, vendor, model, drive type (HDD or SSD) and drive firmware revision. - If the **encryption** option is given, display per drive removable drive encryption information. - If the **lun** option is given, display removable drive LUN per drive.

Option	Description
`enclosure [serial]|[name] ? asset|errors|statistics|allcolumns?`	- If no option is given, display the list of all nodes on the given enclosure. - If the **asset** option is given, display details about all of the hardware components in the given enclosure. - If the **errors** option is given, display the errors related to hardware components on the given enclosure. - The **statistics** option is deprecated. Please use the command **sysstat enclosure [serial]|[name]** to get enclosure statistics. - Use the **allcolumns** option as a shortcut to include all options that give additional data.
`enclosures ?staticlinkagg? ?allcolumns?`	Display a list of all enclosures. - If the **staticlinkagg** option is given, display a column indicating whether static link aggregation is enabled for each enclosure. - Use the **allcolumns** option as a shortcut to include all options that give additional data.
`errors`	Display all errors related to hardware components.
`ips`	Display a listing of all components with IP addresses in the system.
`nics`	Display a listing of all NICs in the system.

Option	Description

`node [node] ? asset|errors|statistics|syslog?`
`node [node] core`
`node [node] core stack [corefileName] ? email?`
`node [node] core copy [corefileName] [dest-file]`
	To specify a node, give [enclosure-name],[slot], IP address, hardware serial, or data serial.
	- If no option is given, display the general information for the given node. This includes information on aliases as well as the services that the node runs. In this case, nodes can also be specified by IP alias.
	- If the **asset** option is given, display hardware details for the given node.
	- If the **errors** option is given, display the errors related to the given node.
	- The **statistics** option is deprecated. Please use the command **sysstat node [node]** to get node statistics.
	- If the **syslog** option is given, display the latest portion of the syslog for the given node. Note: All syslog date/time entries use the GMT timezone. The rest of the system uses the local time, specified by the **timezone** setting.
	- If the **core** option is given, display corefile information for the given node.
	- If the **core stack** option is given, display the stack trace for given corefile.
	- If the **email** option is given, email the stack trace to VDURA.
	- If the **core copy** option is given, copy the requested corefile to the given location in PanFS. To specify the PanFS path, please use an absolute path starting with the volume name.

Option	Description
`nodes ?storage|storagemodules| directors? ?dataserial? ?aliases? ?services? ? shelf|enclosure [serial]| [name-pattern]? ?storageset ?[name-pattern]?? ?capacity? ? verticalparity? ?allcolumns? ?ipmi?`	Display a list of all nodes' status, address, and serial information. Results include all types of nodes by default. - Use the **storage**, **storagemodules**, **directors**, and **directormodules** options to limit the output by hardware type. These options are not mutually exclusive. - Use the **enclosure** or **shelf option** to limit the listing to nodes in the matching enclosures. If a glob pattern is given, enclosure names must match the pattern. A serial number also can be given. - Use the **storageset** option to limit the listing to nodes belonging to the StorageSet whose name matches the glob pattern. Use **unassigned** as the StorageSet name to list enclosures that are not assigned to any StorageSet. If no name-pattern is given, results will include the StorageSet name for each VPOD. - Use the **aliases** option to display all IP aliases for each listed node. - Use the **services** option to display all services that each listed node runs. The services displayed in this command are: RM, RM President, Enclosure Master, DHCP, DNS, CIFS, and PM. - Use the **dataserial** option to display the internal data serial number for each listed node. - Use the **capacity** option to show the capacity of each Storage Node. - Use the **verticalparity** option to show whether each Storage Node has vertical parity enabled or disabled. - Use the **allcolumns** option as a shortcut to include all options that give additional data.

Option	Description
	- Use the **ipmi** option to show the IPMI card configuration information.
`temps`	This command displays the internal, CPU, and incoming air temperatures for all nodes in the system. Results are listed in order of decreasing temperature.

sysstat

Description

Displays current system statistics, including CPU utilization, operations per second, and capacity.

Syntax

```
sysstat ?[subcommand] ?[arg1] [arg2] ...??
```

Options

Option	Description

`all`
`all copy [path]`
`all scp`
	Display a list of system status and configuration information. This information may be requested by service personnel as a first step in troubleshooting system problems.
	- If no option is given, the information is printed to the console.
	- If the **copy** option is supplied, the information is copied to the given PanFS file path. Files must be specified using a full path starting with **/volume-name**.
	- If the **scp** option is given, once the information is gathered, the user will be prompted to enter the information needed to complete the file transfer to a remote location via **scp**.

Option	Description

`clients`
`clients df ?[#clients] ?[min]??`
`clients df storageset [name] ?[#clients] ?[min]??`
`clients df node [id] ?[#clients] ?[min]??`
`clients nfs ?[director] ...? ?[#clients] ?[min]??`
`clients cifs ?[director] ...? ?[#clients] ?[min]??`
`clients cifs reset ?[director] ...?`
	Display a list of clients with the highest resource usage. You must use one of the following arguments to run this command.
	Specify **nfs** for NFS client information, **cifs** for CIFS client information, or **df** for DirectFlow clients.
	NFS and CIFS clients connect to Directors. When listing NFS or CIFS clients, give the [enclosure-name],[slot] name, serial, or IP address of a Director to limit results to clients connected to a single Director.
	The **cifs reset** option resets the CIFS client statistics on the given list of Directors. If no Director is specified, it resets client stats on all Directors. This is a service mode command and should only be used under the direction of VDURA service personnel to ensure safe usage.
	DF clients connect to Storage Nodes. When listing DF clients, give the VPOD set name to limit results to clients connected to nodes in a specific Object RAID VPOD set. Give the Storage Node name, serial, or IP address to limit results to clients accessing a single Storage Node.
	The number of clients displayed can be specified as a number between 1 and 1024. The default is 10.
	The number of minutes to average the load over can be given as a number between 1 and 60. The default is a 5 minute average. Note that for Storage Nodes, data is only collected at five minute intervals, so setting the average minutes to load over as either less than five or values not divisible by five is not best practice.

`director`	Display performance information for each Director.

Option	Description
`drives ?summary?`	Display performance information for each storage drive. If the **summary** keyword is given, then give the performance information on a per-StorageSet basis rather than for each Storage Node.
`enclosure [serial]|[name]`	Display performance information for the requested enclosure.
`errors` `errors clear [id1] ?[id2] ... [idN]?`	Show all or clear specific errors, warnings, and info items in the system.
`hotfiles [node] ?max [max]? ?timeout [timeout]?`	The 10 hottest files are listed by default. Specify **max** and a nonzero integer to list a different number of files. Specify **timeout** and a nonzero integer to use a different timeout (in seconds) than the default, which is hardware specific, but generally between 25 and 180 seconds.
`node [node-designator]`	Display performance information for the requested node. To specify a node, give [enclosure-name], [slot], IP address, hardware serial, or data serial.

Option	Description
`storage ?summary? ?enclosure|shelf [serial]|[name-pattern]? ?storageset [name-pattern]?`	Display performance information for each Storage Node. - If the **summary** keyword is given, then give the performance information on a per-StorageSet basis rather than for each Storage Node. - Use the **enclosure** option to limit the listing to the enclosure with the given serial number or all enclosures whose name matches the glob pattern. - Use the **storageset** option to limit the listing to nodes belonging to the StorageSet whose name matches the glob pattern. Note that the **summary** option cannot be used in conjunction with the **storageset** or **enclosure** options.
`summary`	Display basic system status, also seen during initial login.

tail

Description

The tail utility displays the contents of file or, by default, its standard input, to the standard output.

The display begins at a byte, line, or 512-byte block location in the input. Numbers having a leading plus (+) sign are relative to the beginning of the input, for example, "-c +2" starts the display at the second byte of the input. Numbers having a leading minus (-) sign or no explicit sign are relative to the end of the input, for example, "-n2" displays the last two lines of the input. The default starting location is "-n 10", or the last 10 lines of the input.

Syntax

```
tail [-F | -f | -r] [-b number | -c number | -n number] [file ...]
```

Options

Option	Description
`-b number`	The location is number 512-byte blocks.
`-c number`	The location is number bytes.
`-f`	The -**f** option causes tail to not stop when end of file is reached, but rather to wait for additional data to be appended to the input. The -**f** option is ignored if the standard input is a pipe, but not if it is a FIFO.
`-F`	The -**F** option implies the -**f** option, but tail will also check to see if the file being followed has been renamed or rotated. The file is closed and reopened when tail detects that the filename being read from has a new inode number. The -**F** option is ignored if reading from standard input rather than a file.
`-n number`	The location is number lines.

Option	Description
`-r`	The -**r** option causes the input to be displayed in reverse order, by line. Additionally, this option changes the meaning of the -**b**, -**c** and -**n** options. When the -r option is specified, these options specify the number of bytes, lines, or 512-byte blocks to display, instead of the bytes, lines, or blocks from the beginning or end of the input from which to begin the display. The default for the -**r** option is to display all of the input.

tar

Description

Creates, adds files to, or extracts files from an archive file in tar format, called a tarfile. A tarfile is often a magnetic tape, but can be a floppy diskette or any regular disk file.

The first argument word of the **tar** command line is usually a command word of bundled function and modifier letters, optionally preceded by a dash. It must contain exactly one function letter from the set **A**, **c**, **d**, **r**, **t**, **u**, and **x**, for append, create, difference, replace, table of contents, update, and extract. The command word can also contain other function modifiers, some of which take arguments from the command line in the order they are specified in the command word. Functions and function modifiers can also be specified with the GNU argument convention (preceded by two dashes, one function or modifier per word). Command-line arguments that specify files to add to, extract from, or list from an archive can be given as shell pattern matching strings.

Syntax

```
tar [[-]bundled-options Args] [gnu-style-flags] [filenames | -C directory-name] ...
```

Functions

Specify exactly one of the following functions.

Option	Description

`-A`
`--catenate`
`--concatenate`
	Append the contents of the named file, which must itself be a tar archive, to the end of the archive (erasing the old end-of-archive block). This adds the files contained in the named file to the first archive, rather than adding the second archive as an element of the first. This option requires a rewritable tarfile, and therefore does not work on quarter-inch cartridge tapes.

`-c` `--create`	Create a new archive (or truncate an old one) and write the named files to it.

Option	Description

`-d`
`--diff`
`--compare`
	Find differences between files in the archive and corresponding files in the file system.

`--delete`	Delete named files from the archive. Does not work on quarter-inch tapes.
`-r` `--append`	Append files to the end of an archive. Does not work on quarter-inch tapes.
`-t` `--list`	List the contents of an archive. If filename arguments are given, only those files are listed. Otherwise the entire table of contents is listed.
`-u` `--update`	Append the named files if the on-disk version has a modification date more recent than their copy in the archive (if any). Does not work on quarter-inch tapes.

`-x`
`--extract`
`--get`
	Extract files from an archive. The owner, modification time, and file permissions are restored, if possible. If no file arguments are given, extract all the files in the archive. If a filename argument matches the name of a directory on the tape, that directory and its contents are extracted (as well as all directories under that directory). If the archive contains multiple entries corresponding to the same file (see **–append**), the last one extracted overwrites all earlier versions.

Options

The other options to **tar** can be combined arbitrarily. Single-letter options can be bundled in with the command word. Verbose options that take arguments are followed by the argument. Single-letter options consume successive command-line arguments.

Option	Description
`--help`	Print a message listing and briefly describing all the command options to **tar**.
`--atime-preserve`	Restore the access times on files that are written to tape (this changes the inode-change time).
`-b` `--block-size number`	Set the block size for reading or writing to **number** times 512-byte blocks.
`-B` `--read-full-blocks`	Reassemble short reads into full blocks (for reading 4.2BSD pipes).
`-C directory` `--directory directory`	Change to **directory** before processing the remaining arguments.
`--checkpoint`	Print the number of buffer reads and writes while reading or writing the archive.
`-f [hostname:]file` `--file [hostname:]file`	Read or write the specified file (the default is `/dev/sa0`). If a host name is specified, **tar** uses **rmt(8)** to read or write the specified file on a remote machine. You can use `-` as a file name for reading or writing to or from standard input or output.
`--force-local`	The archive file is local even if it has a colon.

`-F file`
`--info-script file`
`--new-volume-script file`
	Run a script at the end of each archive volume (implies **-M**).

`--fast-read`	Stop after all non-wildcard extraction targets have been found in the archive.
`-G` `--incremental`	Create, list, or extract an old GNU-format incremental backup.

Option	Description
`-g file` `--listed-incremental file`	Create, list, or extract a new GNU-format incremental backup.
`-h` `--dereference`	Do not write symbolic links as symbolic links. Write the data of the files they name.
`-i` `--ignore-zeros`	Ignore blocks of zeros in the archive (usually means end-of-file).
`--ignore-failed-read`	Do not exit with a non-zero status on unreadable files.

`-j`
`-y`
`--bzip`
`--bzip2`
`--bunzip2`
	Filter the archive through **bzip2(1)**.

`-k` `--keep-old-files`	Keep files that already exist on disk. Do not overwrite them from the archive.
`-K file` `--starting-file file`	Begin at **file** in the archive.
`-l` `--one-file-system`	Stay in the local file system when creating an archive (do not cross mount points).
`-L number` `--tape-length number`	Change tapes after writing **number** times 1024 bytes.
`-m` `--modification-time`	Do not extract the file modified time.
`-M` `--multi-volume`	Create, list, or extract a multi-volume archive.
`-n` `--norecurse`	Do not recurse into subdirectories when creating.
`--volno-file file`	The name of the file with the volume number to start with.

Option	Description

`-N date`
`--after-date date`
`--newer date`
	Store only files with a creation time newer than **date**.

`--newer-mtime date`	Store only files with a modification time newer than **date**.

`-o`
`--old-archive`
`--portability`
	Write a V7 format archive, rather than a POSIX format.

`-O` `--to-stdout`	Extract files to standard output.

`-p`
`--same-permissions`
`--preserve-permissions`
	Extract all protection information.

`--preserve`	Has the effect of **-p -s**.
`-P` `--absolute-paths`	Do not strip the leading `/` from file names.
`-R` `--record-number`	Show the record number within the archive with each message.
`--remove-files`	Remove files after adding them to the archive.

`-s`
`--same-order`
`--preserve-order`
	The list of names to extract is sorted to match the archive.

`--show-omitted-dirs`	Show directories that were omitted while processing the archive.
`-S` `--sparse`	Handle “sparse” files efficiently.

`-T file`
`-I file`
`--files-from file`
	Get the names of files to extract or create from **file**, one per line.

Option	Description
`--null`	Modify the behavior of **-T** to expect null-terminated names. Disables **-C**.
`--totals`	Print the total bytes written with **–create**.

`-U`
`--unlink`
`--unlink-first`
	Unlink files before creating them.

`-v` `--verbose`	List files written to the archive with **–create** or extracted with **–extract**. List file protection information along with file names with **–list**.
`-V volume-name` `--label volume-name`	Create the archive with the given **volume-name**.
`--version`	Print the **tar** program version number.

`-w`
`--interactive`
`--confirmation`
	Ask for confirmation for every action.

`-W` `--verify`	Attempt to verify the archive after writing it.
`--exclude pattern`	Exclude files matching the pattern (do not extract them, add them, or list them).
`-X file` `--exclude-from file`	Exclude files listed in **file**.

`-Z`
`--compress`
`--uncompress`
	Filter the archive through **compress(1)**.

`-z`
`--gzip`
`--gunzip`
	Filter the archive through **gzip(1)**.

`--use-compress-program program`	Filter the archive through **program** (which must accept **-d** to mean “decompress”).

Option	Description
`--block-compress`	Block the output of the compression program for tapes or floppies (otherwise writes are of odd length, which device drivers can reject).
`-[0-7][lmh]`	Specify the tape drive and density.

See Also

gzip, **pax**

tcpdump

Description

A service mode command that allows for capturing network traffic into or out of the node, with arbitrary filters. See the **tcpdump(8)** man page for a detailed description of usage.

When output is redirected to a file with the **-w filename** option, **tcpdump** will be stopped after collecting a maximum of 100,000 packets to avoid filling the local filesystem. This limit on packet count can be overwritten by providing the **-c count** option.

When the output file is specified with the **-w filename** option, the user will be prompted to scp the output file to a remote location.

This command may adversely affect system performance while it is running.

Press the ENTER key to stop **tcpdump**.

Syntax

```
tcpdump [-adeflnNOpqStuvxX] [-c count] [ -F file ] [ -i interface ] [-r file ] [ -s snaplen ] [
-T type ] [ -w file ] [expression]
```

terminal

Description

Sets or displays terminal characteristics.

Options

Option	Description
none	Displays the current terminal settings.
`columns ?[width]?`	Displays or sets the terminal column width. By default, terminal width has no effect on line wrapping of output. Use the terminal wrap subcommand to impose a limit.
`separator ?[separator]?`	By default, two blank spaces are used between columns of data in tables. By setting this variable, table cells will temporarily be delimited by the given text. This setting applies only to the current PanCLI session and is not persistent.
`wrap ?enable|disable?`	The default behavior is to wrap output lines that exceed the terminal width. If line wrapping is bothersome, you can disable it, but note that as a result the output of many commands may be missing information due to the requested line truncation.

Syntax

```
terminal ?[options]...?
```

Example

Display the current separator:

```
terminal separator
```

Delimit the table columns using a pipe:

```
terminal separator |
```

Return to the default separator:

```
terminal separator ""
```

timezone

Description

Displays current realm time zone or a list of possible time zones. Sets the realm time zone.

Syntax

```
timezone ?[subcommand] ?[arg]??
```

Options

Options	Description
none	Display the current time zone.
`display`	Display a list of possible time zones.
`set ?[tzIndex]?`	Interactive command to set the time zone or set the time zone specified by **tzIndex**.

top

Description

Displays and updates information about the top CPU processes.

Syntax

```
top [-SbiInqut][-dcount][-stime][-ofield][-U username][number]
```

Options

Option	Description
`-S`	Show system processes in the display. Normally, system processes such as the pager and swapper are not shown. This option makes them visible.
`-b`	Use **batch** mode. In this mode, all input from the terminal is ignored. Interrupt characters (such as **^C** and ^) still have an effect. This is the default on a dumb terminal or when the output is not a terminal.
`-i`	Use **interactive** mode. In this mode, any input is immediately read for processing. After a command is processed, the screen will immediately be updated, even if the command was not understood. This mode is the default when standard output is an intelligent terminal.
`-I`	Do not display idle processes. By default, top displays both active and idle processes.
`-t`	Do not display the top process.
`-n`	Use **non-interactive** mode. This is identical to **batch** mode.

Option	Description
`-q`	Renice top to -20 so that it will run faster. This can be used when the system is being sluggish to improve the possibility of discovering a problem. This option can only be used by root.
`-u`	Do not take the time to map UID numbers to usernames. Normally, **top** will read as much of the file **/etc/passwd** as is necessary to map all the user ID numbers it encounters into login names. This option disables that, while possibly decreasing execution time. The UID numbers are displayed instead of the names.
`-dcount`	Show only count displays, then exit. A display is considered to be one update of the screen. This option allows the user to select the number of displays he wants to see before top automatically exits. For intelligent terminals, no upper limit is set. The default is 1 for dumb terminals.
`-stime`	Set the delay between screen updates to time seconds. The default delay between updates is two seconds.
`-ofield`	Sort the process display area on the specified field. The field name is the name of the column as seen in the output, but in lower case. Likely values are **cpu**, **size**, **res**, and **time**, but may vary on different operating systems. Note that not all operating systems support this option.
`-Uusername`	Show only those processes owned by username. This option currently only accepts usernames and will not understand UID numbers.

traceroute

Description

Prints the route packets take to a network host.

Syntax

```
traceroute [ip]
```

trpost

Description

Service mode command to post-process binary trace output for easy viewing.

Syntax

```
trpost ?[options]? [trace file(s)]
```

Options

Option	Description
`-A, --absolute-times=on/off`	Print all timestamps as absolute values.
`-c, --count`	Do not create an output file, count trace records.
`-C, --cycles-per-second=uint32`	Used when dump header is missing.
`--defaults`	Show defaults.
`-f, --definitions file=string`	Name of file containing definitions.
`--env-equiv`	Show environment var equivalents.
`--error-count-announce=on/off`	Print the number of parsing/printing errors encountered.
`-e, --error-file=string`	Direct trace decoding/print errors to this file.
`-E, --error-file-name-generate`	Generate the error output filename automatically.
`--etime-detailed`	Print out a record for all etime statements.
`-h, --help`	Show usage message.
`-n, --max-recs=uint64`	Maximum # of recs to process.
`--num-verify=uint32`	Number of trace records to verify before decoding.

Option	Description
`-o, --output-file=string`	Direct trace output to this file.
`-O, --output-file-name-generate`	Generate the output filename automatically.
`-H, --print-dump-header`	Print the dump header.
`-R, --reacquire`	Reacquire framing anywhere in a file.
`-r, --readable-timestamps`	Print out timestamps in calendar (ctime) format.
`--size-histo-buckets=uint32`	Used when dump header is missing.
`--size-histo-max-size=uint32`	Maximum size of histogram (min == 0).
`--size-histo-print`	Print the size histograms.
`-s, --skip-bytes=uint64`	Number of bytes (including header) to first data.
`-q, --suppress-contexts`	Suppress context ID display.
`-m, --suppress-memory-details`	Suppress verbose per-trace memory usage stats.
`-d, --timestamp-deltas`	Print deltas with timestamps.
`-v, --verbose`	Verbose mode.
`--version`	Show program version.

uniq

Description

The **uniq** utility reads the specified **input_file** comparing adjacent lines and writes a copy of each unique input line to the **output_file**. If **input_file** is a single dash (-) or absent, the standard input is read. If **output_file** is absent, standard output is used for output. The second and succeeding copies of identical adjacent input lines are not written. Repeated lines in the input will not be detected if they are not adjacent, so it may be necessary to sort the files first.

Options

Option	Description
`-c`	Precede each output line with the count of the number of times the line occurred in the input, followed by a single space.
`-d`	Only output lines that are repeated in the input.
`-f num`	Ignore the first num fields in each input line when doing comparisons. A field is a string of non-blank characters separated from adjacent fields by blanks. Field numbers are one based, i.e., the first field is field one.
`-s chars`	Ignore the first chars characters in each input line when doing comparisons. If specified in conjunction with the **-f** option, the first chars characters after the first num fields will be ignored. Character numbers are one based, i.e., the first character is character one.
`-u`	Only output lines that are not repeated in the input.
`-i`	Case insensitive comparison of lines.

upgrade

Description

Manages the software revision running on your system.

Syntax

```
upgrade ?[subcommand]? ?[arg1] [arg2] ...?
```

Example

Download the version 12.0.0 upgrade image (deprecated; see the `download` option below):

```
upgrade download 12.0.0
```

Upgrade the entire system to version 12.0.0:

```
upgrade install 12.0.0
```

Upgrade node 000f00d2100100001 to match the system's software:

```
upgrade offversion install 000f00d2100100001
```

Lookup available software version from VDURA (currently not available; the system cannot download from VDURA):

```
upgrade versions remote
```

Options

Option	Description
none	Show the status of the most recent or currently running upgrade transaction.

Option	Description
`cancel`	Cancel the currently running upload, distribution, or install.
`delete [version]`	Remove the system's copy of the given version.
`distribute [version] ?[node options]?`	Distribute a copy of the software image with the given version to all the eligible nodes in parallel. The operation is needed before the **upgrade install** command can be invoked. If for any reason a node is missing its copy of the distribution, this command can be retried. Use the **upgrade status -poll** command to watch the distribution progress. Use the **upgrade diststatus** command to troubleshoot failures. Node options can be used to distribute or redistribute the software image to a specific set of nodes. Values can be IP addresses, the keyword **all**, or any other options accepted by the **iplist** command. When node options are given, distribution will be attempted regardless of whether the given node already has a valid copy of the software image.
`diststatus ?-sortby (starttime|endtime|status)?`	View the individual node details for the latest system-wide software image distribution. By default, node detail is sorted by destination node name. Use the -**sortby** options to sort results by distribution start or end time or by distribution status.
`disttimeout ?timeout?`	View or set upgrade distribution timeout. Specify the number of minutes to wait for completion of the scp of the image to a node. The timeout value can be set to **default** or an integer in the range [1-120]. Changing the **disttimeout** value is a service mode command and should only be performed under the direction of VDURA Technical Support.

Option	Description
`download [version]`	Deprecated as of 12.0.0. The VDURA Data Platform can no longer download upgrade images. Download the image from support.VDURA.com and load it with the `upload` option instead.
`history`	Show the listing of versions and dates for each upgrade of this node.
`install ?-ignoredist? [version]`	Install the given version on the entire system. Note that you must distribute the software image to all nodes before installing it. The -**ignoredist** option is a service mode utility that should only be used at the request of VDURA Technical Support. It launches the install without checking the distribution status on each node. If the software distribute is not present on one or more nodes, the system will be taken offline for upgrade, but the upgrade will not be completed on any of the nodes.
`offversion ?install [node serial or ip]?`	Display the list of off-version nodes. If the install option is given, upgrades the given node to match the system version

Option	Description
`?status ?-poll? ?[numEvents]??`	Show the current status of any upload, distribution, or install, including: - The last **[numEvents]** of the upgrade-related events for the most recent upgrade (started up to 1 week ago); **[numEvents]** defaults to 5. Use the keyword **all** to specify no maximum. - No system upgrade in progress (system currently running version **[v]**). - Uploading a remote software image to the system: – start time: – progress: X% complete - Distributing software to all nodes: – start time: – progress: X% complete - Upgrading system from **[old version]** to **[new version]**: – start time: – progress: X% complete - Upgrading an individual off-version node. - Downloading a software image directly from VDURA (currently not available). - The -**poll** option causes the command to continue to post status updates until the currently running action completes.
`upload [file]`	Upload the software image file from a given location. If **[file]** looks like **? user @ ? host** : **file** , then the system prompts for a password and uses **scp** to copy the file. If **[file]** looks like **ftp:// ? user @ ? host : file** , then the system prompts for a password and uses **ftp** to copy the file. Otherwise, look for the file in PanFS. If no matching PanFS file is found, then check the local filesystem of the node that hosts PanActive Manager. Once the file is copied and validated, it is added to the system's version list.

Option	Description
`versions`	Show the listing of all versions that are available for download or have already been uploaded to the system. Listing versions available for download from VDURA is currently not available: the system cannot download upgrade images from VDURA. Download the image from support.VDURA.com and load it with the `upload` option instead.

uptime

Description

Displays the time since the last boot of the node, and 1, 5, and 15 minute load averages.

usermap

Description

Displays and configures the usermap configuration.

Syntax

```
usermap [subcommand] ?[arg] ... [arg]?
```

Example

Map Windows user foo in default domain to UNIX user bar:

```
foo > bar
```

Map Windows user Fred to UNIX user frederick and vice versa:

```
fred = frederick
```

Qualify a user by name or ID:

```
sid:S-1-5-21-680444065-4293180936-402387372-500 = root
MyWindowsUser < uid:123
```

Qualify the Windows domain in case of multiple domains. All mappings below are considered equal:

```
MyDomain\MyUser > MyUnixUser
MyDomain.MyCompany.com\MyUser > MyUnixUser
MyUser@MyUser.MyCompany.com > MyUnixUser
```

Map all root users in all known domains to nobody - not UNIX user root:

```
*\root > nobody
```

Map all users in SomeDomain to a default user:

```
SomeDomain\* > SomeDefaultUser
```

Any other user gets mapped to global default:

```
*\* > default_unix_user
default_win_user < *
```

Options

Option	Description
`deploy [file]`	Save and deploy the usermap configuration file. This command parses and validates the entries from the given usermap configuration file. It displays validation errors if any. If all the entries in the file pass validation, the file is saved in the system and the usermap configuration is deployed. Check the user guide for the format and syntax of the entries in the usermap file. If **[file]** looks like ?[user]@?[host]:[file], command prompts for a password and uses **scp** to fetch the file. Otherwise, command looks for the file in PanFS. If no matching PanFS file is found, command checks the local filesystem of the node that hosts PanActive Manager. This saved usermap file can be displayed or copied using the **list** subcommand. The user map must meet the following requirements: - One entry per line. - UNIX users may be specified by either name or UID. - UIDs will be prefixed with **uid:**. - Windows users may be specified by either name or SID. - Windows usernames may have spaces in them and must be quoted (“). - Windows usernames may specify a domain name, e.g. **domain\user** OR **user@domain**. - Windows domain names might be fully qualified, e.g. **native2k8.example.com\user**. - Windows usernames are case-insensitive, unlike UNIX usernames. - Mappings can be unidirectional (Windows > UNIX or Windows < UNIX) or bidirectional (Windows = UNIX).

Option	Description
	- Wildcard mappings shall be evaluated if no explicit mapping could be found. - Domain wildcard mappings (***\user** OR **user@***) shall be evaluated first by matching against single users in all domains. - User wildcard mappings (**domain*** OR ***@domain**) shall be evaluated second. - The global wildcard mapping (****** OR ***@***) shall be evaluated last. - In case of UNIX there's only a global mapping, as we can only join a single NIS/LDAP domain.
`list ?[file]?`	Display or copy the saved usermap file. Without the optional argument, it displays the usermap content. If the optional **[file]** argument is given, copies the currently deployed usermap file to the specified **[file]**. If **[file]** looks like ?[user]@?[host]:[file], command prompts for a password and uses **scp** to copy the file. Otherwise, command copies the file to PanFS. If no matching directory is found in PanFS, command copies the file to local filesystem of the node that hosts PanActive Manager.

userquota

Description

Displays and configures user and group quota configuration.

Syntax

```
userquota [subcommand] ?[arg] ... [arg]?
```

Options

Option	Description
`deploylimits [file]`	This command parses and validates the entries from the given limits configuration file and displays validation warnings/errors if any. If all the entries in the file pass validation, the file is saved in the system and the limits configuration is deployed. If **[file]** looks like **?[user]@?[host]:[file]**, the command prompts for a password and uses **scp** to fetch the file. Otherwise command looks for the file in PanFS. If no matching PanFS file is found, command checks the local filesystem of the node that hosts PanActive Manager. This saved limits file can be displayed or copied using the **getlimits** subcommand. It is possible, that one or more entries may fail to be deployed after saving the limits file in the system. If this happens, review the event log for the details of errors. Use the **redeploylimits** subcommand to re-deploy the failed entries. For a description of the contents of the limits file and format of the limits configuration entries in the limits file for deployment, see the **getlimits** subcommand.

Option	Description
`policy ?disabled|tracked|enforced?` `policy ?disabled|tracked|enforced |inherited? ?-volume [pattern1] ?[pattern2] ...??`	This command displays and sets the system and volume level user and group quota enforcement policy. Use this command without the -**volume** argument to display and set system-wide policy settings. The system-wide configuration is displayed if no arguments are given. If a policy value is given, this command sets the system-wide policy to specified value. The optional -**volume** argument is used to display and set volume level policy. If no policy value is given, this command displays the settings for volumes matching the patterns given. If a policy value is given, it will be set for the volumes matching the patterns.
`redeploylimits`	Use this subcommand to re-deploy the saved limits file. This is typically used to re-deploy the limits file when the previous deploy operation had failed, after the file was validated and saved in the system.

Option	Description
`usage ?-volume [pattern1] ?[pattern2] ...?? ?-user [pattern1] ?[pattern2] ...?? ?-group [pattern1] ?[pattern2] ...?? ?-output [plaintext|csv|tab]? ?-violation [soft|hard|both]? ?-summary?`	Without any arguments, display detailed usage data for all volumes, users, and groups. **Note:** Using this may take a while to compile the report. Only byte usage data 10 MB or more and file count data 10 or more will be displayed. - Use optional -**volume**, -**user** and/or -**group** arguments to filter usage report by giving patterns of volumes, users and/or groups. These optional args must be followed by one or more patterns. If no pattern is specified for volume, it will default to all volumes. If both user and group patterns are not specified, both user and group quota usage will be reported. If only user or group pattern is specified, only user or group quota usage will be reported. - Use optional -**violation** argument to filter usage report based on quota violation. Default is to show all usage data. - Use the optional -**summary** argument to specify a summarized report. Choosing the -**summary** option will display a summary total of usage for all users/groups per volume and a summary total of usage for each user/group across all volumes. - Use the optional -**output** argument to display the report in plaintext, csv, or tab delimited style. Defaults to plaintext report in a table format.

Option	Description
`getlimits ?[file]?`	Display or copy the saved limits file. - Without the optional argument, displays saved limits file. - With optional **[file]** argument, copies the currently deployed limits file to the specified file. - If **[file]** looks like **?[user]@?[host]:[file]**, the command prompts for a password and uses **scp** to copy the file. Otherwise command copies the file to PanFS. If no matching directory in PanFS is found, command copies the file to local filesystem of the node that hosts PanActive Manager. This file contains one quota entry per line to specify the limits data. Format for each entry: - One entry per line. - Separate entries for user and group quotas. - Separate entries for system and volume level defaults. - The first three fields–type, **id**, and vol–define the key for the entry. All fields except email address fields are required. If email address field is missing, it will default to **none**. - All lines beginning with **#** are comment lines and will be ignored. - All blank lines will be ignored. - The delimiter for fields is whitespace. Leading and trailing white spaces will be ignored. Any field value that contains a space should be quoted with “double quotes”. - For each entry, the following fields and values are expected from left to right (position dependent): - **type**: User or group - **user**: Key word to set quota for user - **group**: Key word to set quota for group - **name**: User or group id/name - If type is **user** or **group**, * indicates all users or

Option	Description
	groups (use to set volume level defaults). - If value is a text string: If prefixed with a domain (ex. domain1\user_1), it will be treated as a windows user/group name. Otherwise, will be treated as UNIX user/group name. - If it starts with **S**-.... it will be considered as Windows SID. - If it starts with **uid:** or **gid:** followed by a number, will be treated as UNIX user/group ID. - If it starts with **user_sid:** or **group_sid:** followed by a SID, will be treated as windows user or group SID. - **vol**: Volume name - * indicates all volumes - i.e. will set the system default. - **name** field must be * if volume is * (i.e., you cannot set default for a user across all volumes; this must be set with an entry for each volume.) - Byte limits - **sbl**: soft byte limit - **hbl**: hard byte limit - **sfl**: soft file limit - **hfl**: hard file limit - For the above limits: - Use - to indicate inherit and 0 to indicate unlimited. - Use up to 2 decimal places for byte limits. - Use only whole numbers for file limits. - Use a minimum of 10M for byte limits. - Use suffix K, M, G, T, or P to indicate mega-, giga-, tera, or peta bytes (power of 10): +ex: 800 K, 200 M, 10G 12.8T etc. - For bytes quota, minimum is 10M and maximum is approximately 9223 PB (use 0 for unlimited, which is approximately 9223 PB). - For files quota, minimum is 1 and maximum is approximately 9223 P (use 0 to indicate

Option	Description
	unlimited which is approximately 9223 petafiles). - **sqemail**: Email address for **soft_quota_violation_notification**. - **hqemail**: Email address for **hard_quota_violation_notification**. Email address format: - Comma separated list of email IDs or may be empty. - Use **$user** – for UNIX email address of user. - Use **$AD** or **$LDAP** – email address to be fetched from AD or LDAP respectively. - Use - to inherit from volume. - If email ID contains spaces, it must be quoted with “double quotes”. - Use a = character by itself as the field value if the email list entered for the other email field is to be used for this also. Do not use = in both. - Use the word **none** explicitly to indicate no email to be sent. - If multiple entries are found for the same **type**/**id**/**vol key** combination, the last entry will be considered, ignoring all earlier ones. A warning to that effect will be displayed when deploying the file. **Note**: Setting user quota enforcement policy for root (/) volume is not allowed, and is always set to **Disabled**. Therefore, any limits set for root (/) volume will have no effect. For examples, see the PanCLI **help userquota getlimits** command.

vi

Description

Edits a text file. In the PanCLI shell, **vi** is automatically run with the secure edit option set, disallowing all access to external programs.

Syntax

```
vi [-FlRr] [-c cmd] [-w size] [file ...]
```

Options

Option	Description
`-c`	Execute cmd immediately after starting the edit session. Particularly useful for initial positioning in the file, however cmd is not limited to positioning commands.
`-F`	Don't copy the entire file when first starting to edit. The default is to make a copy in case someone else modifies the file during your edit session.
`-l`	Start editing with the lisp and showmatch options set.
`-R`	Start editing in read-only mode, as if the command name was view, or the readonly option was set.
`-r`	Recover the specified files, or, if no files are specified, list the files that could be recovered. If no recoverable files by the specified name exist, the file is edited as if the -**r** option had not been specified.
`-w`	Set the initial window size to the specified number of lines.

volservice

Description

Manages backup of volume metadata. A volume service is the metadata for a group of volumes. Volume services reside on a primary Director and may have a backup location on a secondary Director.

Syntax

```
volservice ?(failover|list|move|moveall|remove|txnlogsize) ?[args]??
```

Options

Option	Description
none	Show whether the volume services failover feature is enabled or disabled.

`autocreate`
`autocreate enable all|[director]`
`autocreate disable [director]`
	Manage automatic volume service creation settings for Directors.
	- If no option is given, lists all Directors that have been disabled for automatic volume service creation. These Directors have been marked not eligible to host volume services.
	- To specify a Director, use [enclosure-name], [slot], IP address, hardware serial, or data serial.
	- If option **all** is given, the command will enable all Directors in the system for automatic volume service creation.
	- **Volservice autocreate disable [director]** is a service mode command and should only be performed under the direction of VDURA Technical Support.

Option	Description
`failover ?enable|disable?`	Enabling volume services failover will start redundant backup services for metadata management on a different Director so that volume services will not be interrupted in case of a Director failure. If volume services failover is disabled and a Director failure occurs, the root filesystem service will automatically move to another node, but other volume services hosted by the failed Director will have to be moved manually by using the **volservice move** command. The failed volume services will be marked **FAILED** in the service listing.
`list ?[director]?`	By default this command lists Directors that are assigned to a StorageSet followed by a listing of volume services including the primary and backup Director and volume list. To specify a Director, give [enclosure-name], [slot], IP address, hardware serial, or data serial. If a Director is given, this command lists both primary and backup volume services hosted on the Director, including the list of volumes for each service.

Option	Description
`move ?backup? [serviceID] [director]`	Move a volume service from its current Director to another Director. This moves the metadata management responsibility for all volumes controlled by that volume service. There is no data movement. If the source Director is down, then a File System Recovery (FSRC) will be done on the destination Director before the volumes are brought online. Note that this action can affect Director load balancing. If the **backup** option is chosen, then this command moves the designated backup location for the volume service to the specified Director. This action will not result in any volume downtime, however it can affect Director load balancing following a subsequent Director failover. Note that this action has no effect if volume services failover is disabled. To specify a destination node, give [enclosure-name],[slot], IP address, hardware serial, or data serial.
`moveall [source director] [destination director] ?[alternate destination director]?`	This command moves non-empty volume services from the source Director to the destination Director. When failover is enabled, an alternate destination Director may be needed, as services cannot have the same primary and backup location. This command additionally removes some volume services if they manage no volumes. All empty volume services whose primary owner is the source Director will be removed. Additionally, if the destination Director owns one or more nonempty volume services, any empty volume services owned by the destination Director will also be removed. The same is true of the alternate destination Director, if one is provided.

Option	Description
`remove [serviceID]`	This service mode command allows the removal of a volume service that manages no volumes. This may become necessary after Director replacement, when a Director is the primary owner of two or more volume services and one of which has no volumes in it. If this command is used to remove an empty volume service from a Director that is not the primary owner of any other volume services, the system may automatically recreate a new empty volume service on this Director, which will cancel the effectiveness of this request.

volume

Description

Displays and manages volumes.

> **Note:** To use special characters in volume names, put a backslash before the character.

Syntax

```
volume ?[subcommand] ?[arg1] [arg2] ...??
```

Example

Create a volume named /vol2 in StorageSet Set 1 with RAID 6+:

```
volume create /vol2 storageset "Set 1" soft 450 hard 500
```

List the top 70 volumes that are using up space:

```
volume list show 70 sort used order decreasing
```

Change recovery priorities of volumes /user* to default:

```
volume recoverypriority default /user*
```

Change recovery priorities of volumes /foo and /home:

```
volume recoverypriority 5 /foo 10 /home
```

List volumes having rescue status:

```
volume rescue list
```

List volumes that eligible to rescue:

```
volume rescue list eligible
```

Start volume rescue in the StorageSet Set 1:

```
volume rescue start storageset "Set 1" *
```

Stop volume rescue in the system:

```
volume rescue stop *
```

Options

Option	Description
none	If no subcommand is given, info about all volumes is listed.
`cache [volume-name]`	Service mode command to display volume cache information from File Manager.
`cleanupquery [volume-name] ?detail?`	View the progress of a volume background stale component cleanup. If the **detail** option is given, the command will display detailed information.

Option	Description

`compression [volume-name]`
`compression [volume-name] ?enable|disable?`
	Display or change compression for the specified volume.
	Use this command to display whether compression is currently enabled or disabled and to enable or disable data compression on a specific volume.
	If no option is given, the command displays whether compression is currently enabled or disabled for the specified volume.
	To specify a volume, use its name. Enabling compression may improve storage efficiency but could affect performance, depending on workload and hardware.
	Disabling compression will stop further data from being compressed, but existing compressed data will remain compressed until it is rewritten.
	Example usage:
	`volume compression vol1`
	`volume compression vol1 enable`
	`volume compression vol1 disable`

Option	Description
`create [name] ?[options...]?`	Create a new volume with the specified name for the specified StorageSet. - **description [description]:** If a description is given, this text will be saved with the volume info for administrative convenience. - Volume descriptions are displayed in “volume list description” and “volume description” output. - **efsa (retry|file-unavailable|default):** Specifies the Extended File System Availability mode of the volume. - The default value is retry. - This option is only available for volumes with a RAID 6+, RAID 5+ or RAID 10+ layout. - **encryption [on/off]:** - Use **volume create encryption on** command to create a volume with encryption on if system-wide encryption is off by default. - Use **volume create encryption** off command to create a volume with encryption off if system-wide encryption is on by default - Use **volume list showencryption** to display volume encryption state (on/off). - **layout (raid6+|raid5+|raid10+) / layout (raid5|raid10):** Specifies the raid layout of the volume. The raid setting for a volume cannot be changed after a volume is created. - The **raid6+** layout provides RAID 6 protection for files. RAID 6+ volumes provide a balance between performance, capacity overhead, and RAID protection. - The **raid5+** layout provides RAID 5 protection. RAID 5 volumes are optimized for performance and capacity overhead. The **raid10+** layout provides RAID 10 (striped mirror) protection. - RAID 10 volumes are optimized for small write performance.

Option	Description
	- The **+** on layout types indicates that directory data and metadata have an extra level of redundancy in the volume. For backward compatibility, raid5 and raid10 layouts are supported but only in service mode. - **maxwidth [number-of-VPODs]:** specifies the maximum number of VPODs to stripe over. - Any integer 1 or greater is a valid width. - If no width is specified, the system allows striping across the StorageSet. - The minimum number of VPODs for RAID 5+ is 2; for RAID 6+, the minimum value is 3; for RAID 10+, the minimum value is 2. - **recoverypriority** ([1-100]|default): specifies the recovery priority of the volume. - A lower number designates a higher priority, where 1 is the highest priority. - If the recoverypriority is specified as **default** or no recoverypriority option is given, the system will choose a default value. - **rgdepth [number-of-stripes]:** specifies the number of stripes written to a RAID parity group before advancing to the next parity group. - This option is only available for volumes with RAID 6+ or RAID 5+ layout. - Any integer 1 or greater is a valid depth. - If no depth is specified, the system will choose a default. - **rgwidth [number-of-VPODs]:** specifies the number of VPODs to stripe over in a single RAID group. - Only available for volumes with RAID 6+ or RAID 5+ layout. - Any integer between 3 and 20 (inclusive) is a valid width. - If no width is specified, the system chooses a default.

Option	Description
	- **soft [soft-quota-limit] / hard [hard-quota-limit]:** The quota limits are measured in GB. - Use 0 to specify an unlimited quota. - The default behavior is unlimited quotas. - **storageset [storageset-name]:** If the system has more than one StorageSet, you must specify the name of the StorageSet for your new volume to belong to. - **stripeunit [bytes]:** specifies the number of bytes written to a VPOD before advancing to the next VPOD in the stripe. - Values must be in the form [number]K or [number]M such that the stripe unit is divisible by 16K, with a minimum of 16K and a maximum of 4M. - If no stripe unit is specified, the system will choose a default. - **uperm [user perms] / gperm [group perms] / operm [others perms]:** Permissions must be one of: none, read-only, write-only, execute-only, read-write, read-execute, write-execute, or all. - **user [username or ID] / group [group name or ID]:** User and group can be specified as either UNIX username, UNIX user ID, or Windows SID. - In the case of Windows SID, please specify in the form **sid:S-*** - **volservice [volservice ID]:** If a volume service ID is specified, the new volume will be managed by the given volume service. - This option offers some control over volume management load balancing across Directors. If the volume service is not specified, the system will assign the new volume to an existing volume service automatically.

Option	Description

`damage delete [obj_id]`
`damage deleteall [volume-name]`
`damage history [volume-name]`
`damage list [volume-name]`
`damage unlock [obj_id]`
`damage unlockall [volume-name]`
`damage clearhistory [obj-id]`
`damage clearrecord [obj-id]`
	Manage damaged files in a volume.
	- **delete**: Delete the given damaged filesystem object.
	- **deleteall**: Delete all the damaged filesystem objects in the given volume.
	- **history**: Print a log of actions taken to deal with damaged objects in the given volume.
	- **list**: Print a list of damaged objects in the given volume.
	- **unlock**: Unlock the given filesystem object.
	- **unlockall**: Unlock all the filesystem objects in the given volume.
	- **clearhistory**: Clear all history about the given damaged object.
	- **clearrecord**: Clear the record from the damaged object list.

`delete [volume-name]`	Delete the volume.
`description [volume-name] ?[description text]?`	View or update the description for the given volume. Volumes can be annotated with descriptions for administrative convenience. If no description text argument is given, display the existing description for the given volume. Volume descriptions also appear in **volume list description** output. If description text is given, save the text as the new description for the given volume. Use the following to clear an existing description: `volume description [volume-name] ""`
`details [volume-name]`	Display volume details and status.

Option	Description
`efsa ?[retry|file-unavailable|default] [vol-pattern-1] ?[vol-pattern-2] ...??`	PanFS supports two Extended File System Availability modes: retry and file-unavailable. These modes apply when an object has exceeded its fault-tolerance level. - In the retry mode, PanFS clients will block and retry I/O requests until needed storage devices are back online (default). - In the file-unavailable mode, I/O requests result in I/O errors. If no EFSA mode option is given, this command is the same as **volume list showefsa**. If EFSA mode **default** is given, the system default value will be applied.
`errors ?[volume-name]? ?show [max]|all?`	Display errors related to all volumes in the system. If a volume name is given, errors related to that volume will be displayed. A maximum of 5 errors will be displayed by default. Use the **show** option to specify a different maximum.
`hard-quota-email list|add|remove [volume-name] [email-address-list]`	Display, add, or remove hard quota notification email. The email address list should be a list of valid email addresses separated by a comma or a space. If no email addresses are specified, a quota violation email will go to the general administrator email address.

Option	Description

`list ?name [pattern]?`
`?allcolumns?`
`?storageset[name-pattern]?`
`?director [name]? ?description? ?id ?`
`?nameonly?`
`?order increasing |decreasing?`
`?show [max]|all?`
`?showdirector?`
`?showencryption?`
`?showefsa?`
`?showrecoverypriority?`
`?showvolservice?`
`?sort name|used|storageset|volservice|director|recoverypriority?`
`?snapshot [name]?`
`?volservice [service-id]?`
	Get the list of all volumes in the system. A maximum of 50 volumes is displayed by default.
	- **allcolumns**: Include all options that give additional data.
	- **description**: Display the volume description for each volume with a nonempty description saved.
	- **director**: Only include volumes primarily hosted on the given Director.
	- **id**: Display the internal numeric ID and object ID values for each volume.
	- **nameonly**: List just the names of the volumes. Volume names will be listed in a space-separated format for use in **foreach** loops. For example, you can add a replicator configuration for each volume matching your search by doing the following:
	`[pancli] foreach volume [volume list nameonly ...search criterion...] {+ puts [replicator add $volume ...options...]+ }`
	- **order**: Indicate whether to sort in increasing or decreasing order.
	- **show**: Specify a different maximum.
	- **showdirector**: Display the primary Director.
	- **showefsa**: Display the Extended File System Availability Mode.
	- **showencryption**: Display E2EE encryption status of volumes.
	- **showrecoverypriority**: Display the recovery priority of each volume.
	- **showvolservice**: Display the volume service ID.
	- **sort**: Sort the results by name, by space used, by volume service ID, by Director, by recovery priority, or by StorageSet.
	- **storageset**: Only include volumes whose StorageSet name matches the given glob

Option	Description
	pattern. - **name**: Specify a glob pattern of volume names to match. - **snapshot**: Only include volumes that are covered by the given snapshot. - **volservice**: Only include volumes managed in the given volumeservice.
`moveroot ?layout [layout]? [storageset-name]`	Change the layout of the / volume or move the / volume to the given StorageSet. This command recreates the root volume with given layout in the given StorageSet. The old root volume will be deleted. The layout option specifies the raid layout of the new root volume. Valid values are **raid6+**, **raid5+**, or **raid10+**. If layout option is not given, the default layout of new root volume will be **raid6+**. For backward compatibility, layouts raid5 and raid10 are supported but only in service mode. Moving the root volume automatically deletes all files, symbolic links, and directories that are not parts of paths to other volumes' mountpoints. Other volumes will not be affected. If you have volumes named **/directory/vol0** and **/directory/vol1**, and the root volume is moved via the volume **moveroot** command or within the application, the permission on **/directory** will likely change after the **moveroot** and you may not be able to access the volume. After the **moveroot**, modify the permissions on directories of this type to allow access to the volume.
`reconquery [volume-name] ?detail?`	View the progress of volume reconstruction. If the option **detail** is given, the command will display detailed information.

Option	Description
`recover ?-ignorestate? ?[options]? [volume-name]`	This command starts File System Recovery (FSRC) on the specified volume. The system will take the volume offline while the check proceeds and will repair any errors found in the volume. Use this command with care! The check/repair process cannot be aborted once it starts. It may take many hours to recover volumes with many files or lots of data. The -**ignorestate** flag is a service mode option that will force the system to start FSRC regardless of the volume's state. This option should only be used under the direction of customer support. **ADVANCED OPTIONS:** - **parity-scrub**: Check RAID parity for all files in the volume. Normally, RAID parity is only checked for recently written files. Parity scrubbing proceeds at about 80 MB/sec, so this can extend FSRC runtime significantly. - **no-scrub**: Do not check RAID parity for any file. Only one of the advanced options can be specified at a time.

Option	Description

`recoverypriority`
`recoverypriority default`
`recoverypriority ?[priority-level-1] [volname-pattern-1] ...?`
	Display or edit the recovery priority of one or more volumes.
	If no priority level or the option **default** is given, this command is the same as **volume list showrecoverypriority**. It also accepts any other options accepted by the **volume list** command.
	If the option **default** is given without the volume name pattern, this command will display the system default volume recovery priority value.
	Priority level can be an integer in range [1-100] or **default**. 1 is the highest priority. If **default** is given, the system default value will be applied.
	The recovery priority of the root volume cannot be changed.

`rename [volume-name] [new-name]`	Rename the given volume.

Option	Description

`rescue list ?eligible? ?storageset [storageset]?`
`rescue report ?storageset [storageset]? ?[volname-pattern]?`
`rescue reportclear ?storageset [storageset]? [volname-pattern]`
`rescue start ?storageset [storageset]? [volname-pattern]`
`rescue stop ?storageset [storageset]? [volname-pattern]`
	Service mode command to manage volume rescue.
	When the level of fault-tolerance on a volume has been exceeded, the admin may request a rescue operation manually. The system will perform a best-effort reconstruction of volume data followed by offline recovery (FSRC) of the volume. Any objects in the volume that cannot be reconstructed due to offline VPODs will be deleted.
	- **eligible**: Only list the volumes that are eligible for rescue.
	- **list**: Print a list of volumes having rescue status.
	- **report**: Collect the rescue report logs for the specified volumes. If the **volname-pattern** is not given, report logs for all volumes will be collected.
	- **reportclear**: Delete the rescue report logs for the specified volumes.
	- **start**: Request rescue for the specified volumes.
	- **stop**: Cancel the rescue requests for the specified volumes.
	- **storageset**: Only process the volumes in the given StorageSet.
	A glob pattern of volume names must be given to start or stop volume rescue.

`set read-only|read-write [volume-name]`	Sets the access permission for the volume. When the volume is created, the access permission is set to read-write. It can be changed to read-only.
`set soft-quota|hard-quota [volume-name] [quota-limit]`	Sets the soft or hard quota limit on the specified volume. The quota limit is measured in GB; use 0 to specify an unlimited quota.

Option	Description
`soft-quota-email list|add|remove [volume-name] [email-address-list]`	Display, add, or remove soft quota notification email. The email address list should be a list of valid email addresses separated by a comma or a space. If no email addresses are specified, quota violation email will go to the general administrator email address.
`state [online|offline|offline-unavailable] [volume-name]`	Moves the specified volume to the online or offline state. In the offline state, PanFS clients will block and retry I/O requests until the volume is brought back online. In the offline-unavailable state, I/O requests to this volume result in I/O errors. The offline-unavailable state is only allowed in service mode.

wbinfo

Description

Queries **idmapd_vdura**.

Syntax

```
wbinfo [option...]
```

Options

Option	Description
`-u, --domain-users`	List all domain users.
`-g, --domain-groups`	List all domain groups.
`-p`	Ping **idmapd_vdura** to see if it is alive.
`--get-config-info`	Get **idmapd_vdura** config info.
`--own-domain`	List own domain.
`--trusted-domains`	List trusted domains.
`--lookup-sids=SID-List`	Translate a list of SIDs (comma separated).
`--pan-id-to-name`	Translate a pan_identity to a name.
`--user-to-uid=USER`	Translate a username to a UID.
`--user-to-sid=USER`	Translate a username to a SID.
`--real-uid-to-name-and-sid=UID`	Translate a Unix UID to Windows.
`--sid-to-name-and-real-uid=SID`	Translate a SID to a Unix username and UID.
`--group-to-gid=GROUP`	Translate a groupname to a GID.
`--group-to-sid=GROUP`	Translate a group to a SID.
`--user-in-group=USER%GID`	Validate username in GID.

Option	Description
`--get-user-groups=USER`	Get the list of groups the user belongs to.
`--rid-is-member-of-group-sid=RID%SID`	Validate user rid member of group SID.
`-T, --validate-user-group=USER%GROUP`	Validate username in group name.
`-Q, --mserver-auth-query=STRINGS`	Validate user password unix-group windows-group for **mserver**.
`--get-netgroup=NAME`	Get a netgroup entry.
`--set-log-level=STRING`	Set log level to specified value (error, warning, info, verbose, debug, trace).
`--validate-usermap=filename`	Validate usermap file at [filename].

WC

Description

The **wc** utility displays the number of lines, words, and bytes contained in each input file or standard input (if no file is specified) to the standard output. A line is defined as a string of characters delimited by a newline character. Characters beyond the final newline character will not be included in the line count.

A word is defined as a string of characters delimited by white space characters. White space characters are the set of characters for which the **iswspace(3)** function returns true. If more than one input file is specified, a line of cumulative counts for all the files is displayed on a separate line after the output for the last file.

When an option is specified, **wc** only reports the information requested by that option. The order of output always takes the form of line, word, byte, and file name. The default action is equivalent to specifying the -**c**, -**l**, and -**w** options.

If no files are specified, the standard input is used and no file name is displayed. The prompt will accept input until receiving EOF, or [^D] in most environments.

Syntax

```
wc [-Lclmw] [file ...]
```

Options

Option	Description
`-L`	The number of characters in the longest input line is written to the standard output. When more than one file argument is specified, the longest input line of all files is reported as the value of the final total.
`-c`	The number of bytes in each input file is written to the standard output. This will cancel out any prior usage of the -**m** option.
`-l`	The number of lines in each input file is written to the standard output.

Option	Description
`-m`	The number of characters in each input file is written to the standard output. If the current locale does not support multibyte characters, this is equivalent to the **-c** option. This will cancel out any prior usage of the **-c** option.
`-w`	The number of words in each input file is written to the standard output.

whoami

Description

Prints the username associated with the current effective user ID.

xargs

Description

The **xargs** utility reads space, tab, newline, and end-of-file delimited strings from the standard input and executes the utility with the strings as arguments.

Any arguments specified on the command line are given to the utility upon each invocation, followed by some number of the arguments read from the standard input of xargs. This is repeated until the standard input is exhausted.

Spaces, tabs, and newlines may be embedded in arguments using single or double quotes or backslashes. Single quotes escape all non-single quote characters, excluding newlines, up to the matching single quote. Double quotes escape all non-double quote characters, excluding newlines, up to the matching double quote. Any single character, including newlines, may be escaped by a backslash.

If **utility** is omitted, **echo(1)** is used.

Undefined behavior may occur if utility reads from the standard input.

The **xargs** utility exits immediately without processing any further input if a command line cannot be assembled, **utility** cannot be invoked, an invocation of **utility** is terminated by a signal, or an invocation of **utility** exits with a value of 255.

Syntax

```
xargs [-0oprt] [-E eofstr] [-I replstr [-R replacements] [-S replsize]] [-J replstr] [-L
number] [-n number [-x]] [-P maxprocs] [-s size] [utility [argument ...]]
```

Options

Option	Description
`-0`	Change **xargs** to expect NUL (``\0'') characters as separators, instead of spaces and newlines. This is expected to be used in concert with the -**print0** function in **find(1)**.
`-E eofstr`	Use **eofstr** as a logical EOF marker.

Option	Description
`-I replstr`	Execute utility for each input line, replacing one or more occurrences of replstr in up to replacements (or 5 if no -**R** flag is specified) arguments to utility with the entire line of input. The resulting arguments, after replacement is done, will not be allowed to grow beyond replsize (or 255 if no -**S** flag is specified) bytes; this is implemented by concatenating as much of the argument containing replstr as possible, to the constructed arguments to utility, up to replsize bytes. The size limit does not apply to arguments to utility which do not contain replstr, and furthermore, no replacement will be done on utility itself. Implies -**x**.
`-J replstr`	If this option is specified, **xargs** will use the data read from standard input to replace the first occurrence of replstr instead of appending that data after all other arguments. This option will not affect how many arguments will be read from input (-n) or the size of the command(s) **xargs** will generate (-**s**). The option just moves where those arguments will be placed in the command(s) that are executed. The **replstr** must show up as a distinct argument to **xargs**. It will not be recognized if, for instance, it is in the middle of a quoted string. Furthermore, only the first occurrence of the **replstr** will be replaced. For example, the following command will copy the list of files and directories which start with an uppercase letter in the current directory to **destdir**: `/bin/ls -1d [A-Z]* | xargs -J % cp -rp % destdir`

Option	Description
`-L number`	Call utility for every number lines read. If EOF is reached and fewer lines have been read than number then utility will be called with the available lines.
`-n number`	Set the maximum number of arguments taken from standard input for each invocation of utility. An invocation of utility will use less than number standard input arguments if the number of bytes accumulated (see the **-s** option) exceeds the specified size or there are fewer than number arguments remaining for the last invocation of **utility**. The current default value for number is 5000.
`-o`	Reopen **stdin** as **/dev/tty** in the child process before executing the command. This is useful if you want **xargs** to run an interactive application.
`-P maxprocs`	Parallel mode: run at most **maxprocs** invocations of **utility** at once.
`-p`	Echo each command to be executed and ask the user whether it should be executed. An affirmative response, “y” in the POSIX locale causes the command to be executed; any other response causes it to be skipped. No commands are executed if the process is not attached to a terminal.

Option	Description
`-r`	Compatibility with GNU **xargs**. The GNU version of **xargs** runs the utility argument at least once, even if **xargs** input is empty, and it supports a -**r** option to inhibit this behavior. The FreeBSD version of xargs does not run the utility argument on empty input, but it supports the -**r** option for command-line compatibility with GNU xargs, but the -**r** option does nothing in the FreeBSD version of **xargs**.
`-R replacements`	Specify the maximum number of arguments that -I will do replacement in. If replacements is negative, the number of arguments in which to replace is unbounded.
`-S replsize`	Specify the amount of space (in bytes) that **-I** can use for replacements. The default for replsize is 255.
`-s size`	Set the maximum number of bytes for the command line length provided to utility. The sum of the length of the utility name, the arguments passed to utility (including NULL terminators) and the current environment will be less than or equal to this number. The current default value for size is **ARG_MAX - 4096**.
`-t`	Echo the command to be executed to standard error immediately before it is executed.
`-x`	Force **xargs** to terminate immediately if a command line containing number arguments will not fit in the specified (or default) command line length.

ypcat

Description

Prints the values of all keys from the NIS database specified by mapname, which may be a map name or a map nickname.

Syntax

```
ypcat [-k] [-d domainname] [-h hostname] [-t] mapname
ypcat -x
```

Options

Option	Description
`-d domain`	Specify a domain other than the default domain as returned by **domainname(1)**.
`-h hostname`	Specify a hostname other than the default one as found by **ypbind(8)**.
`-k`	Display map keys. This option is useful with maps in which the values are null or the key is not part of the value.
`-t`	Inhibit map nickname translation.
`-x`	Display the map nickname translation table.

ypmatch

Description

Prints the values of one or more keys from the NIS database specified by mapname, which may be a map name or a map nickname.

Syntax

```
ypmatch [ -kt ] [ -d domain ] key ... mapname
ypmatch -x
```

Options

Option	Description
`-d domain`	Specify a domain other than the default domain as returned by **domainname(1)**.
`-k`	Display map keys. This option is useful with maps in which the values are null or the key is not part of the value.
`-t`	Inhibit map nickname translation.
`-x`	Display the map nickname translation table.

zcat

Description

Decompresses files then concatenates and prints in the manner of the **cat** command.

For more information, see **gzip**.

Syntax

```
zcat [-fhV] file [file [...]]
```

zgrep

Description

Prints lines matching a pattern. For more information, see **grep**.

Syntax

```
zgrep [options] pattern [file...]
```

zless

Description

Interactively views a zipped file. This command is equivalent to **zmore** but uses **less** as a pager instead of **more**.

Syntax

```
zless [ name ... ]
```

zmore

Description

Zmore is a filter which allows examination of compressed or plain text files one screenful at a time on a soft-copy terminal. The command works on files compressed with compress, pack, or gzip, and also on uncompressed files. If a file does not exist, **zmore** looks for a file of the same name with the addition of a .gz, .z or .Z suffix.

Syntax

```
zmore [options]
```