

Hitachi Content Software for File

3.9.4

CLI Reference

Hitachi Content Software for File is a high performance storage solution for AI, ML, analytics, and other GPU-accelerated workloads. It provides the speed of a distributed file system (DFS) with the capacity and hybrid cloud capabilities of an object store. The unique integration of file and object storage offers a tightly coupled, single solution for an appliance-like experience designed for ultra-high performance and capacity applications

© 2021 Hitachi Vantara. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at <https://support.hitachivantara.com/en-us/contact-us.html>.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

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

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.

Contents

| | |
|--|-----------|
| Preface | 7 |
| Intended audience..... | 7 |
| Release notes..... | 7 |
| Document conventions..... | 7 |
| Accessing product documentation..... | 9 |
| Getting help..... | 9 |
| Comments..... | 9 |
| Chapter 1: Managing Content Software for File | 10 |
| Overview..... | 10 |
| Connecting to another host..... | 14 |
| Using an auto-completion script..... | 14 |
| Displaying the cluster status..... | 14 |
| Chapter 2: Managing object stores, filesystem groups, and filesystems | 16 |
| Managing object stores..... | 16 |
| Viewing object stores using the CLI..... | 16 |
| Adding an object store using the CLI..... | 16 |
| Editing an object store using the CLI..... | 22 |
| Deleting an object store using the CLI..... | 24 |
| Attaching or detaching object stores using the CLI..... | 25 |
| Managing filesystem groups..... | 25 |
| Viewing filesystem groups using the CLI..... | 26 |
| Adding a filesystem group using the CLI..... | 26 |
| Editing a filesystem group using the CLI..... | 26 |
| Deleting a filesystem group using the CLI..... | 27 |
| Managing filesystems..... | 28 |
| Viewing filesystems using the CLI..... | 28 |
| Adding a filesystem using the CLI..... | 28 |
| Editing a filesystem using the CLI..... | 30 |
| Deleting a filesystem using the CLI..... | 31 |
| Chapter 3: Managing KMS | 33 |
| Adding or updating a KMS using the CLI..... | 33 |

| | |
|---|-----------|
| Viewing the KMS using the CLI..... | 35 |
| Removing the KMS using the CLI..... | 35 |
| Re-wrapping filesystem keys using the CLI..... | 35 |
| Chapter 4: Managing manual fetch and release of data..... | 37 |
| Pre-fetching API for data lifecycle management using the CLI..... | 37 |
| Fetching files from an object store using CLI..... | 37 |
| Fetching a directory containing many files using the CLI..... | 38 |
| Release API for data lifecycle management..... | 38 |
| Releasing files from SSD to an object store using the CLI..... | 38 |
| Releasing a directory containing many files..... | 39 |
| Chapter 5: Managing clients..... | 40 |
| Joining the cluster using the CLI..... | 40 |
| Configuring the host as client using the CLI..... | 40 |
| Configuring client networking using the CLI..... | 41 |
| Applying the host configuration using the CLI..... | 42 |
| Chapter 6: Managing snapshots..... | 44 |
| Viewing snapshots using the CLI..... | 44 |
| Creating a snapshot using the CLI..... | 44 |
| Deleting a snapshot using the CLI..... | 45 |
| Restoring a filesystem or snapshot from another snapshot using the CLI..... | 46 |
| Updating a snapshot using the CLI..... | 46 |
| Uploading a snapshot using the CLI..... | 47 |
| Creating a filesystem from a snapshot using the CLI..... | 48 |
| Chapter 7: Managing NFS..... | 50 |
| Managing NFS networking configuration (interface groups)..... | 50 |
| Defining interface groups using the CLI..... | 50 |
| Setting interface group ports using the CLI..... | 50 |
| Setting interface group IPs using the CLI..... | 51 |
| Managing NFS Access Control (client access groups)..... | 52 |
| Defining client access groups using the CLI..... | 52 |
| Managing client access groups using the CLI..... | 52 |
| Adding or deleting an IP using the CLI..... | 53 |
| Managing NFS client permissions..... | 53 |
| Managing NFS client permissions using the CLI..... | 53 |
| Chapter 8: Managing SMB..... | 56 |
| Showing an SMB cluster using the CLI..... | 56 |
| Showing an SMB domain configuration using the CLI..... | 56 |
| Creating an SMB cluster using the CLI..... | 56 |

| | |
|---|-----------|
| Checking status of SMB host readiness using the CLI..... | 60 |
| Joining an SMB cluster to an Active Directory using the CLI..... | 60 |
| Deleting an SMB cluster using the CLI..... | 61 |
| Configuring trusted domains using the CLI..... | 61 |
| Listing trusted domains using the CLI..... | 61 |
| Adding trusted domains using the CLI..... | 61 |
| Removing trusted domains using the CLI..... | 62 |
| Listing SMB shares using the CLI..... | 62 |
| Adding SMB shares using the CLI..... | 63 |
| Updating SMB shares using the CLI..... | 66 |
| Controlling SMB shares users lists using the CLI..... | 66 |
| Showing SMB share user lists using the CLI..... | 66 |
| Adding SMB share user to list using the CLI..... | 67 |
| Removing SMB share user from lists using the CLI..... | 68 |
| Resetting SMB share user lists using the CLI..... | 69 |
| Removing SMB shares using the CLI..... | 70 |
| Chapter 9: Managing alerts..... | 72 |
| Listing alert types using the CLI..... | 72 |
| Describing alerts using the CLI..... | 72 |
| Viewing alerts using the CLI..... | 72 |
| Muting alerts using the CLI..... | 73 |
| Unmuting alerts using the CLI..... | 73 |
| Chapter 10: Managing events..... | 75 |
| Viewing events using the CLI..... | 75 |
| Listing local events using the CLI..... | 77 |
| Triggering a custom event using the CLI..... | 78 |
| Chapter 11: Managing statistics..... | 79 |
| Listing statistic types using the CLI..... | 79 |
| Viewing statistics realtime using the CLI..... | 79 |
| Reading statistic information using the CLI..... | 80 |
| Setting statistic retention using the CLI..... | 82 |
| Chapter 12: Managing users..... | 84 |
| Creating users using the CLI..... | 84 |
| Changing user passwords using the CLI..... | 85 |
| Deleting users using the CLI..... | 86 |
| Configuring an LDAP user directory using the CLI..... | 86 |
| Configuring an LDAP server using the CLI..... | 86 |
| Viewing a configured LDAP user directory using the CLI..... | 89 |
| Disabling or enabling a configured LDAP user directory using the CLI..... | 90 |

| | |
|---|------------|
| Chapter 13: Managing organizations..... | 91 |
| Managing organizations..... | 91 |
| Creating an organization using the CLI..... | 91 |
| Viewing organizations using the CLI..... | 92 |
| Renaming organizations using the CLI..... | 92 |
| Updating an organization's quotas using the CLI..... | 93 |
| Deleting an organization using the CLI..... | 93 |
| Revoking user access using the CLI..... | 94 |
| Chapter 14: Expansion of specific resources..... | 95 |
| Dynamic modifications using the CLI..... | 95 |
| Memory modifications..... | 96 |
| Network modifications..... | 96 |
| Host IPs modifications..... | 96 |
| Local resources editing commands using the CLI..... | 97 |
| Chapter 15: Managing clusters..... | 98 |
| Shrinking a cluster using the CLI..... | 98 |
| Options for shrinking a cluster..... | 98 |
| Listing drives and their states using the CLI..... | 98 |
| Deactivating a drive using the CLI..... | 98 |
| Removing a drive using the CLI..... | 99 |
| Deactivating an entire host using the CLI..... | 100 |
| Removing a host using the CLI..... | 100 |
| Chapter 16: Managing background tasks..... | 101 |
| Viewing background tasks using the CLI..... | 101 |
| Limiting background tasks using the CLI..... | 101 |
| Chapter 17: Running cluster diagnostics..... | 102 |
| Managing diagnostics using the CLI..... | 102 |

Preface

This guide provides information and instructions for managing the Hitachi Content Software for File system from the Command Line (CLI).

Please read this document carefully to understand how to use this product, and maintain a copy for your reference.

Intended audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who install, configure, and operate the Hitachi Content Software for File system.

Readers of this document should be familiar with the following:

- Storage system and performance concepts, including clustering and networking.
- Storage array and tiering concepts.
- Object stores, including S3, Hitachi Content Platform, and Hitachi Content Platform for cloud scale.
- Data lifecycle management concepts.

Release notes


Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>.




Document conventions

This document uses the following typographic conventions:

| Convention | Description |
|---------------------|--|
| Bold | <ul style="list-style-type: none"> Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. Indicates emphasized words in list items. |
| <i>Italic</i> | <ul style="list-style-type: none"> Indicates a document title or emphasized words in text. Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p> |
| Monospace | Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code> |
| < > angle brackets | Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings. |
| [] square brackets | Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing. |
| { } braces | Indicates required or expected values. Example: { a b } indicates that you must choose either a or b. |
| vertical bar | Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b. |

This document uses the following icons to draw attention to information:

| Icon | Label | Description |
|---|-------|---|
|  | Note | Calls attention to important or additional information. |

| Icon | Label | Description |
|---|---------|--|
|  | Tip | Provides helpful information, guidelines, or suggestions for performing tasks more effectively. |
|  | Caution | Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash). |
|  | WARNING | Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury. |

Accessing product documentation

Product user documentation is available on the Hitachi Vantara Support Website: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

[Hitachi Vantara Support Connect](#) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](#) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Managing Content Software for File

The commands used to manage the Hitachi Content Software for File software system are described along with instructions for using the command-line interface (CLI).

Overview

The Command Line Interface (CLI) is installed on each host and is available through the `weka` command. Running this command will display a list of all available top-level commands.

Available top-level commands

```

$ weka
Usage:
    weka [--help] [--build] [--version] [--legal]

Description:
    The base command for all weka related CLIs

Subcommands:
    agent      Commands that control the weka agent (outside the weka
containers)
    alerts     List alerts in the Weka cluster
    cloud      Cloud commands. List the cluster's cloud status, if no
subcommand supplied.
    cluster    Commands that manage the cluster
    debug      Commands used to debug a weka cluster
    diags      Diagnostics commands to help understand the status of the
cluster and its environment
    events     List all events that conform to the filter criteria
    fs         List filesystems defined in this Weka cluster
    local      Commands that control weka and its containers on the local
machine
    mount      Mounts a wekafs filesystem. This is the helper utility
installed at /sbin/mount.wekafs.
    nfs        Commands that manage client-groups, permissions and
interface-groups
    org        List organizations defined in the Weka cluster
    security   Security commands.
    smb        Commands that manage Weka's SMB container
    stats      List all statistics that conform to the filter criteria
    status     Get an overall status of the Weka cluster
    umount     Unmounts wekafs filesystems. This is the helper utility
installed at /sbin/umount.wekafs.
    user       List users defined in the Weka cluster
    version    When run without arguments, lists the versions available on
this machine. Subcommands allow for
              downloading of versions, setting the current version and
other actions to manage versions.

Options:
    --agent      Start the agent service
    -h, --help   Show help message
    --build      Prints the CLI build number and exits
    -v, --version Prints the CLI version and exits
    --legal      Prints software license information and exits

```

A number of options are common across many commands:

- `-J|--json` flag prints the raw JSON value returned by the cluster.
- `-H|--hostname` flag directs the CLI to communicate with the cluster through the given hostname.
- `--raw-units` flag causes units such as capacity and bytes to be printed in their raw format, as returned by the cluster.
- `--UTC` flag causes timestamps to be printed in the UTC timezone, rather than in the local time of the machine running the CLI command.
- `-f|--format` flag specifies the format to output the result (view, csv, markdown or JSON).
- `-o|--output` flag specifies the columns of the output to be included.
- `-s|--sort` flag specifies the order to sort the output. May include a '+' or '-' before the column name to sort by ascending or descending order.
- `-F|--filter` flag specifies the filter values for a member (without forcing it to be in the output).
- `--no-header` flag indicates that the column header should not be shown when printing the output.
- `-C|--CONNECT-TIMEOUT` flag can be used to change the default timeout used for connecting to the system using the JRPC protocol.
- `-T|--TIMEOUT` flag can be used to change the default timeout for which the commands waits for a response before giving up.

Command hierarchy

Most top-level commands are the default list command for their own collection. Additional sub-commands may be available under them.

For example, the `weka fs` command displays a list of all filesystems and is also the top-level command for all filesystems, filesystem groups and snapshot-related operations. It is possible to use the `-h/--help` flags or the `help` command to display a list of available commands at each level, as shown below:

```
$ weka fs
| FileSystem | Name      | Group   | SSD Bu | Total  | Is re | Is creat | Is
remov
| ID        |          |         | dget   | Budget | ady   | ing      |
ing
+-----+-----+-----+-----+-----+-----+-----+-----+
| FSId: 0   | default | default | 57 GiB | 57 GiB | True  | False   |
False $ weka fs -h
Usage:
  weka fs [--name name]
          [--HOST HOST]
          [--PORT PORT]
          [--CONNECT-TIMEOUT CONNECT-TIMEOUT]
          [--TIMEOUT TIMEOUT]
```

```

[--format format]
[--output output]...
[--sort sort]...
[--filter filter]...
[--help]
[--raw-units]
[--UTC]
[--no-header]
[--verbose]
[--json]

```

Description:

List filesystems defined in this Weka cluster

Subcommands:

```

create      Create a filesystem
download    Download a filesystem from object store
update      Update a filesystem
delete      Delete a filesystem
restore     Restore filesystem content from a snapshot
group       List filesystem groups
snapshot    List snapshots
tier        Show object storage connectivity for each node in the cluster

```

Options:

```

--name      Filesystem name
-H, --HOST  Specify the host. Alternatively, use the
WEKA_HOST env variable
-P, --PORT  Specify the port. Alternatively, use the
WEKA_PORT env variable
-C, --CONNECT-TIMEOUT  Timeout for connecting to cluster, default: 10
secs (format: 3s, 2h, 4m, 1d, 1d5h, 1w)
-T, --TIMEOUT  Timeout to wait for response, default: 1 minute
(format: 3s, 2h, 4m, 1d, 1d5h, 1w)
-f, --format  Specify in what format to output the result.

```

Available options are:

```

view|csv|markdown|json|oldview (format: 'view',
'csv', 'markdown', 'json' or 'oldview')
-o, --output  Specify which columns to output. May include
any of the following:
id,name,group,usedSSDD,usedSSDM,usedSSD,freeSSD,
availableSSDM,availableSSD,usedTotalD,usedTotal,freeTotal,availableTotal,
maxFiles,status,encrypted,stores,auth
-s, --sort    Specify which column(s) to take into account
when sorting the output. May include a '+' or
 '-' before the column name to sort in ascending
or descending order respectively. Usage:
[+|-]column1[, [+|-]column2[,...]]
-F, --filter  Specify what values to filter by in a specific
column. Usage:
column1=val1[,column2=val2[,...]]

```

```

-h, --help          Show help message
-R, --raw-units     Print values in raw units (bytes, seconds,
etc.). When not set, sizes are printed in
                    human-readable format, e.g 1KiB 234MiB 2GiB.
-U, --UTC          Print times in UTC. When not set, times are
converted to the local time of this host.
--no-header        Don't show column headers when printing the
output
-v, --verbose      Show all columns in output

```

Connecting to another host

Most system commands deliver the same result on all cluster hosts. However, it is sometimes necessary to execute a command on a specific host. This is performed using the `-H/--hostname` option and specifying the hostname or IP address of the target host.

Using an auto-completion script

Using `bash` you can use auto-completion for CLI commands and parameters. The auto-completion script is automatically installed.

- To disable the auto-completion script, run `weka agent autocomplete uninstall`
- To reinstall the script on a host, run `weka agent autocomplete install` and re-enter your shell session.

You can also use `weka agent autocomplete export` to get the bash completions script and write it to any desired location.

Displaying the cluster status

The `weka status` command displays the overall status of the Content Software for File system.

For example, if the cluster is healthy, a result similar to the following should be displayed:

```

$ weka status
WekaIO v3.7.2 (CLI build 3.7.2)

    cluster: WekaProd (00569cef-5679-4e1d-afe5-7e82748887de)
    status: OK (8 backends UP, 48 drives UP)
  protection: 4+2
    hot spare: 2 failure domains
drive storage: 82.94 TiB total, 82.94 TiB unprovisioned

```

For example, if the cluster has one failed host, a result similar to the following should be displayed:

```
$ weka status
WekaIO v3.7.2 (CLI build 3.7.2)

    cluster: WekaProd (00569cef-5679-4e1d-afe5-7e82748887de)
      status: DEGRADED (7 backends UP, 42 drives UP)
            Rebuild in progress (3%)
    protection: 4+2
      hot spare: 2 failure domains
    drive storage: 82.94 TiB total, 82.94 TiB unprovisioned
```

Chapter 2: Managing object stores, filesystem groups, and filesystems

The management of object stores, filesystem groups and filesystems is an integral part of the operation and performance of the Content Software for File system and overall data lifecycle management.

Managing object stores

Manage object stores using the CLI.

Viewing object stores using the CLI

Command

```
weka fs tier s3
```

This command is used to view information on all the object stores configured to the system.

Adding an object store using the CLI

Before you begin

- For Hitachi Content Platform for cloud scale object store, make sure that the storage component contains an already configured bucket. For details on object store configuration, see [Bucket management](#).
- For HCP for cloud scale only, after logging in to the HCP for cloud scale GUI to generate new S3 credentials for the `Access Keys` and `Secret Keys`, which are required to configure the object store in Content Software for File, use the option to go the built-in S3 Console. For HCP for cloud scale, in the S3 Console, perform these applicable steps:
 1. Select the user icon in the upper right corner of the screen.
 2. Select Generate credentials. A warning screen displays.
 3. Select Generate. New values for `Access Keys` and `Secret Keys` display.
 4. Select Copy to copy the values.
 5. Select Done.

For more information about HCP for Cloud Scale, see [HCP for Cloud Scale](#).

- For Hitachi Content Platform object store, configure a namespace, a tenant, set the pruning value to zero (0), enable versioning, and create credentials, as described in [Creating a namespace](#). Creating the `Access Keys` and `Secret Keys` is based on the user that has access to the bucket or namespace. Generating new credentials removes the previously generated S3 credentials..
- Additionally, the generated values for `Access Keys` and `Secret Keys` are not viewable again, so maintain them for your records.

For more information about Content Platform, see [Content Platform](#).


Command

```
weka fs tier s3 add
```

Use the following command line to add an object store:

```
weka fs tier s3 add <name> [--hostname=<host>] [--port=<port>] [--bucket=<bucket>] [--auth-method=<auth-method>] [--region=<region>] [--access-key-id=<access-key-id>] [--secret-key=<secret-key>] [--protocol=<protocol>] [--bandwidth=<bandwidth>] [--errors-timeout=<errors-timeout>] [--prefetch-mib=<prefetch-mib>] [--max-concurrent-downloads=<max-concurrent-downloads>] [--max-concurrent-uploads=<max-concurrent-uploads>] [--max-concurrent-removals=<max-concurrent-removals>]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|---|----------------------|-----------|---------|
| name | String | Enter the display name of the object store being created, up to 1,024 alphanumeric characters.  Note: If you leave this field blank, the storage component is listed without a name. | Must be a valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|---|-------------------------|-----------|---------|
| hostname | String | Enter the object store host identifier, valid name/IP. Either DNS or IP Address. Type as many as 255 URI unreserved characters A-Z, a-z, 0-9, hyphen (-), period (.), underscore (_), and tilde (~) only. | Must be a valid name/IP | Yes | |
| port | String | Default (80). Field is not mandatory. This is an AWS S3 term. | Must be a valid name | No | 80 |

| Name | Type | Value | Limitations | Mandatory | Default |
|--------|--------|--|----------------------|-----------|---------|
| bucket | String | <ul style="list-style-type: none"> <li data-bbox="706 262 873 1291">▪ For HCP for cloud scale, enter the object store bucket name found on the storage component. The name can be from 3 to 63 characters long and must contain only lowercase characters (a-z), numbers (0-9), periods (.), or hyphens (-). <li data-bbox="706 1312 873 1543">▪ For Content Platform, enter the Tenant information. | Must be a valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|--------|---|---------------------------------------|-----------|---------|
| auth-method | String | Choose from: None, AWSSignature2, AWSSignature4. Content Platform supports AWSSignature2 and AWSSignature4. HCP for cloud scale supports AWSSignature4. | None, AWSSignature2, or AWSSignature4 | Yes | |
| region | String | Enter region name of up to 1,024 characters. For example, the region <code>us-west-2</code> is a possible valid region. | | Yes | |
| access-key-id | String | Enter the object store access key ID. | | Yes | |
| secret-key | String | Enter object store secret key ID. | | Yes | |
| bandwidth | Number | Default (100,000 Mb/s). Bandwidth limitation per core (Mbps). Entry not mandatory. | | No | |

| Name | Type | Value | Limitations | Mandatory | Default |
|--------------------------|--------|--|---------------------------------------|-----------|---------|
| errors-timeout | Number | If the OBS link is down for longer than this timeout period, all IOs that need data return with an error | 1-15 minutes, for example, 5m or 300s | No | 300 |
| prefetch-mib | Number | How many MiB of data to prefetch when reading a whole MiB on object store | | No | 0 |
| max-concurrent-downloads | Number | Maximum number of downloads concurrently performed on this object store in a single IO node | 1-64 | No | 64 |
| max-concurrent-uploads | Number | Maximum number of uploads concurrently performed on this object store in a single IO node | 1-64 | No | 64 |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------------------|--------|--|-------------|-----------|---------|
| max-concurrent-removals | Number | Maximum number of removals concurrently performed on this object store in a single IO node | 1-64 | No | 64 |



Note: By default, when using the CLI, a misconfigured object store will not be created. To create an object store even when it is misconfigured, use the `--skip-verification` option.

Up to 2 different object stores buckets can be configured per filesystem in the system.

Editing an object store using the CLI

Command

weka fs tier s3 update

Use the following command line to edit an object store:

```
weka fs tier s3 update <name> [--new-name=<new-name>] [--hostname=<hostname>] [--port=<port>] [--bucket=<bucket>] [--auth-method=<auth-method>] [--region=<region>] [--access-key-id=<access-key-id>] [--secret-key=<secret-key>] [--protocol=<protocol>] [--bandwidth=<bandwidth>] [--errors-timeout=<errorstimeout>] [--prefetch-mib=<prefetch-mib>] [--max-concurrent-downloads=<maxconcurrent-downloads>] [--max-concurrent-uploads=<max-concurrent-uploads>] [--max-concurrent-removals=<max-concurrent-removals>]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|---------------------------------------|----------------------|-----------|---------|
| name | String | Name of the object store being edited | Must be a valid name | Yes | |
| new-name | String | New name for the object store | Must be a valid name | No | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|--------|--|---------------------------------------|-----------|---------|
| hostname | String | Object store host identifier | Must be a valid name/IP | Yes | |
| port | String | Object store port | Must be a valid name | Yes | |
| bucket | String | Object store bucket name | Must be a valid name | Yes | |
| auth-method | String | Authentication method | None | Yes | |
| region | String | Region name | | Yes | |
| access-key-id | String | Object store access key ID | | Yes | |
| secret-key | String | Object store secret key | | Yes | |
| bandwidth | Number | Bandwidth limitation per core (Mbps) | | No | |
| errors-timeout | Number | If the OBS link is down for longer than this timeout period, all IOs that need data return with an error | 1-15 minutes, for example, 5m or 300s | No | |
| prefetch-mib | Number | How many MiB of data to prefetch when reading a whole MiB on object store | | No | |

| Name | Type | Value | Limitations | Mandatory | Default |
|--------------------------|--------|---|-------------|-----------|---------|
| max-concurrent-downloads | Number | Maximum number of downloads concurrently performed on this object store in a single IO node | 1-64 | No | |
| max-concurrent-uploads | Number | Maximum number of uploads concurrently performed on this object store in a single IO node | 1-64 | No | |
| max-concurrent-removals | Number | Maximum number of removals concurrently performed on this object store in a single IO node | 1-64 | No | |

Deleting an object store using the CLI

Command

```
weka fs tier s3 delete
```

Use the following command line to delete an object store:

```
weka fs tier s3 delete <name>
```


Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|--|----------------------|-----------|---------|
| name | String | Name of the object store being deleted | Must be a valid name | Yes | |

Attaching or detaching object stores using the CLI**Command**

```
weka fs tier s3 attach
```

To attach an object store to a filesystem, use the following command:

```
weka fs tier s3 attach <fs-name> <obs-name>
```

Command

```
weka fs tier s3 detach
```

To detach an object store from a filesystem, use the following command:

```
weka fs tier s3 detach <fs-name> <obs-name>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|---|----------------------|-----------|---------|
| fs-name | String | Name of the filesystem to be attached to / detached from the object store | Must be a valid name | Yes | |
| obs-name | String | Name of the object store to be attached or detached | Must be a valid name | Yes | |

Managing filesystem groups

Manage filesystem groups using the CLI.

Viewing filesystem groups using the CLI

Command

```
weka fs group
```

Use this command to view information about the filesystem groups in the system.

Adding a filesystem group using the CLI

Command

```
weka fs group create
```

Use the following command to add a filesystem group:

```
weka fs group create <name> [--target-ssd-retention=<retention>] [--start-demote=<demote>]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------------|--------|---|------------------------|-----------|------------------|
| name | String | Name of the filesystem group being created | Must be a valid name | Yes | |
| target-ssd-retention | Number | Target retention period (in seconds) before tiering to the object store | Must be a valid number | No | 86400 (24 hours) |
| start-demote | Number | Target tiering cue (in seconds) before tiering to the object store | Must be a valid number | No | 10 |

Editing a filesystem group using the CLI

Command

```
weka fs group update
```

Use the following command to edit a filesystem group:

```
weka fs group update <name> [--new-name=<new-name>] [--target-ssd-
retention=<retention>] [--start-demote=<demote>]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------------|--------|---|------------------------|-----------|---------|
| name | String | Name of the filesystem group being edited | Must be a valid name | Yes | |
| new-name | String | New name for the filesystem group | Must be a valid name | Yes | |
| target-ssd-retention | Number | New target retention period (in seconds) before tiering to the object store | Must be a valid number | No | |
| start-demote | Number | New target tiering cue (in seconds) before tiering to the object store | Must be a valid number | No | |

Deleting a filesystem group using the CLI

Command

```
weka fs group delete
```

Use the following command line to delete a filesystem group:

```
weka fs group delete <name>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|--|----------------------|-----------|---------|
| name | String | Name of the filesystem group to be deleted | Must be a valid name | Yes | |

Managing filesystems

Manage filesystems using the CLI.

Viewing filesystems using the CLI

Command

weka fs

Use this command to view information about the filesystems in the system.

Adding a filesystem using the CLI

Command

weka fs create

Use the following command line to add a filesystem:

```
weka fs create <name> <group-name> <total-capacity> [--ssd-capacity <ssd>]
[--max-files <max-files>] [--filesystem-id <id>] [--encrypted] [--obs-name
<obs-name>] [--auth-required auth]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|--------------------------------------|----------------------|-----------|---------|
| name | String | Name of the filesystem being created | Must be a valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|---------|---|------------------------|----------------------------------|--|
| group-name | String | Name of the filesystem group to which the new filesystem is to be connected | Must be a valid name | Yes | |
| total-capacity | Number | Total capacity of the new filesystem | Minimum of 1GiB | Yes | |
| ssd-capacity | Number | For tiered filesystems, this is the SSD capacity. If not specified, the filesystem is pinned to SSD | Minimum of 1GiB | No | SSD capacity will be set to total capacity |
| max-files | Number | Metadata allocation for this filesystem | Must be a valid number | No | Automatically calculated by the system based on the SSD capacity |
| encrypted | Boolean | Encryption of filesystem | | No | No |
| obs-name | String | Object store name for tiering | Must be a valid name | Mandatory for tiered filesystems | |
| auth-required | String | Determines if mounting the filesystem requires to be authenticated to Content Software for File | yes or no | No | no |



Note: When creating an encrypted filesystem a KMS must be defined. To define an encrypted filesystem without a KMS, it is possible to use the `--allow-no-kms` parameter in the command. This can be useful when running POCs but should not be used in production, since the security chain is compromised when a KMS is not used. If filesystem keys exist when adding a KMS, they are automatically re-encrypted by the KMS for any future use.

Editing a filesystem using the CLI

Command

weka fs update

Use the following command line to edit an existing filesystem:

```
weka fs update <name> [--new-name=<new-name>] [--total-capacity=<totalcapacity>] [--ssd-capacity=<ssd-capacity>] [--max-files=<max-files>] [--authrequired=<auth-required>]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|--------|---|------------------------|-----------|----------------|
| name | String | Name of the filesystem being edited | Must be a valid name | Yes | |
| new-name | String | New name for the filesystem | Must be a valid name | Optional | Keep unchanged |
| total-capacity | Number | Total capacity of the edited filesystem | Must be a valid number | Optional | Keep unchanged |
| ssd-capacity | Number | SSD capacity of the edited filesystem | Must be a valid number | Optional | Keep unchanged |
| max-files | Number | Metadata limit for the filesystem | Must be a valid number | Optional | Keep unchanged |

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|--------|---|-------------|-----------|---------|
| auth-required | String | Determines if mounting the filesystem requires to be authenticated to Content Software for File | yes or no | No | no |

Deleting a filesystem using the CLI

Command

weka fs delete

Use the following command line to delete a filesystem:

```
weka fs delete <name> [--purge-from-obs]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|---------|---|----------------------|-----------|---------|
| name | String | Name of the filesystem to be deleted | Must be a valid name | Yes | |
| purge-from-obs | Boolean | For a tiered filesystem, if set, all filesystem data is deleted from the object store bucket. | | No | False |



Note: Using `purge-from-obs` will remove all data from the object store. This includes any backup data or snapshots created from this filesystem (if this filesystem has been downloaded from a snapshot of a different filesystem, it will leave the original snapshot data intact). If any of the removed snapshots have been (or are) downloaded and used by a different filesystem, that filesystem will stop functioning correctly, data might be unavailable, and errors might occur when accessing the data.



Note: It is possible to either un-tier or migrate such a filesystem to a different object store bucket before deleting the snapshots it has downloaded.

Chapter 3: Managing KMS

Manage Key Management System (KMS) within the Content Software for File system using the CLI.

Adding or updating a KMS using the CLI

Command

```
weka security kms set
```

Use the following command line to add or update the Vault KMS configuration in the Content Software for File system:

```
weka security kms set <type> <address> <key-identifier> [--token token] [--client-cert client-cert] [--client-key client-key] [--ca-cert ca-cert]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|--------|---|---|---|---------|
| type | String | Type of the KMS | Either <code>vault</code> or <code>kmip</code> | Yes | |
| address | String | KMS server address | URL for Vault, <code>hostname:port</code> for KMIP | Yes | |
| key-identifier | String | Key to be used for encryption as a service in the KMS | Key name (for Vault) or a key UID (for KMIP) | Yes | |
| token | String | API token to access Vault KMS | Must have: Read permissions to <code>transit/keys/<master-key-name></code> | Must be supplied for Vault and must not be supplied for <code>kmip</code> | |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|---|--|--|---------|
| | | | write permissions to transit/encrypt/<master-keyname> and transit/decrypt/<masterkeyname> permissions to /transit/rewrap and auth/token/lookup | | |
| client-cert | String | Path to the client certificate PEM file | Must permit encrypt and decrypt permissions | Must be supplied for kmip and must not be supplied for Vault | |
| client-key | String | Path to the client key PEM file | | Must be supplied for kmip and must not be supplied for vault | |
| ca-cert | String | Path to the CA certificate PEM file | | Optional for kmip and must not be supplied for vault | |



Note: For the add or update command to succeed, the KMS should be preconfigured and available with the key and a valid token.

For example: Setting the Content Software for File system with a Vault KMS:

```
weka security kms set vault https://vault-dns:8200 weka-key --token
s.nRucA9Gtb3yNVmLUK221234
```

Setting the Content Software for File system with a KMIP compliant KMS (for example, SmartKey):

```
weka security kms set kmip amer.smartkey.io:5696 b2f81234-c0f6-4d63-b5b3-
84a82e231234 --client-cert smartkey_cert.pem --client-key smartkey_key.pem
```

Viewing the KMS using the CLI

Command

```
weka security kms
```

Use this command to show the details of the configured KMS.

Removing the KMS using the CLI

Command

```
weka security kms unset
```

Use this command to remove the KMS from the Content Software for File system. It is only possible to remove a KMS configuration if no encrypted filesystems exist.



Note: To force remove a KMS even if encrypted filesystems exist, use the `--allow-downgrade` attribute. In such cases, the encrypted filesystem keys are re-encrypted with local encryption and may be compromised.

Re-wrapping filesystem keys using the CLI

Command

```
weka security kms rewrap
```

If the KMS key is compromised or requires rotation, the KMS admin can rotate the key in the KMS. In such cases, this command is used to re-encrypt the encrypted filesystem keys with the new KMS master key.

```
weka security kms rewrap [--new-key-uid new-key-uid]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|--|-------------|--|---------|
| new-key-uid | String | Unique identifier for the new key to be used to wrap filesystem keys | | Must be supplied for kmip and must not be supplied for Vault | |



Note: Existing filesystem keys that are part of the Snap-To-Object feature will not be automatically re-encrypted with the new KMS key.



Caution: Unlike in Vault KMS, re-wrapping a KMIP-based KMS requires generating a new key in the KMS, rather than rotating the same key. Hence, the old key should be preserved in the KMS in order to be able to decrypt old Snap2Obj snapshots.

Chapter 4: Managing manual fetch and release of data

How to manage manually force-fetching tiering data back to NVMe SSDs and force-releasing SSD data to the object store.

Pre-fetching API for data lifecycle management using the CLI

Pre-fetching API for data lifecycle management using the CLI.

Tiered files are always accessible and should generally be treated like regular files. Moreover, while files may be tiered, their metadata is always maintained on the NVMe SSDs. This allows traversing files and directories without worrying about how such operations may affect performance.

Sometimes, it may be necessary to access previously-tiered files quickly. In such situations, it is possible to request the Content Software for File system to fetch the file back to the SSD without accessing them directly. This is performed using the prefetch `weka fs tier fetch` command, which can be issued using the command.

Fetching files from an object store using CLI

Command

```
weka fs tier fetch
```

Use the following command to release files:

```
weka fs tier fetch <path> [-v]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|----------------------------------|--|-------------|-----------|---------|
| path | A comma-separated list of string | List of file paths | | Yes | |
| -v, --verbose | Boolean | Showing fetch requests as they are submitted | | No | Off |

Fetching a directory containing many files using the CLI

In order to fetch a directory that contains a large number of files, it is recommended to use the `xargs` command in a similar manner, as follows:

```
find -L <directory path> -type f | xargs -r -n512 -P64 weka fs tier fetch -v
```



Note: The pre-fetching of files does not guarantee that they will reside on the SSD until they are accessed.

In order to ensure that the fetch is effective, the following must be taken into account:

- **Free SSD Capacity:** There has to be sufficient free SSD capacity to retain the filesystems that are to be fetched.
- **Tiering Policy:** The tiering policy may release some of the files back to the object store after they have been fetched, or even during the fetch if it takes longer than expected. The policy should be long enough to allow for the fetch to complete and the data to be accessed before it is released again.

Release API for data lifecycle management

How to release API for data lifecycle management.

Releasing files from SSD to an object store using the CLI

Using the manual release command, it is possible to clear SSD space in advance (for example, for shrinking one filesystem SSD capacity for a different filesystem without releasing important data, or for a job that needs more SSDs space from different files). The metadata will still remain on SSD for fast traversal over files and directories but the data will be marked for release and will be released to the object-store as soon as possible, and before any other files are scheduled to release due to other lifecycle policies.

Command

```
weka fs tier release [-v]
```

Use the following command to release files:

```
weka fs tier release <path>
```

Table 1 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|----------------------------------|--|-------------|-----------|---------|
| path | A comma-separated list of string | List of file paths | | Yes | |
| -v, --verbose | Boolean | Showing release requests as they are submitted | | No | Off |

Releasing a directory containing many files

In order to release a directory that contains a large number of files, it is recommended to use the **xargs** command in a similar manner, as follows:

```
# directory
find -L <directory path> -type f | xargs -r -n512 -P64 weka fs tier release

# similarly, a file containing a list of paths can be used
cat file-list | xargs -P32 -n200 weka fs tier release
```

Chapter 5: Managing clients

How to manage clients using the CLI.

Joining the cluster using the CLI

Command

```
weka cluster host add
```

Once the client host is in the stem mode (this is the mode defined immediately after running the `install.sh` command), use the following command line on the client host to add it to the cluster:

```
weka -H <backend-hostname> cluster host add <client-hostname>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------------|--------|---|-----------------------------|-----------|---------|
| backend-hostname | String | IP/hostname of one of the existing backend instances in the cluster | Valid hostname (FQDN or IP) | Yes | |
| client-hostname | String | IP/hostname of the client currently being added | Valid hostname (FQDN or IP) | Yes | |



Note: On completion of this stage, the host-ID of the newly added host will be received. Make a note of it for the next steps.

Configuring the host as client using the CLI

Command

```
weka cluster host cores
```


To configure the new host as a client, run the following command:

```
weka cluster host cores <host-id> <cores> --frontend-dedicated-cores=<fe_cores>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|--------------------------|--------|--|---|---|---------|
| host-id | String | Identifier of the host to be added to the cluster | Must be a valid host identifier | Yes | |
| cores | Number | Number of physical cores to be allocated to the Content Software for File client | Maximum 19 cores | Yes | |
| frontend-dedicated-cores | Number | Number of physical cores to be dedicated to FrontEnd processes | For clients, the number of total cores and frontendedicated-cores must be equal | Yes, in order to configure a host as a client | |

Configuring client networking using the CLI

Command

```
weka cluster host net add
```



Note: If the new client is to communicate with the Content Software for File system cluster over the kernel UDP stack, it is not necessary to run this command. If a high-performance client is required and the appropriate network NIC is available, use the following command to configure the networking interface used by the client to communicate with the Content Software for File system cluster hosts:

```
weka cluster host net add <host-id> <device> --ips=<ip-address> --netmask=<netmask> --gateway=<gateway>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|---------|------------|--|--|-----------|---------|
| host-id | String | Identifier of the host to be added to the cluster | Must be a valid host identifier | Yes | |
| device | String | Network interface device name for example, <code>eth1</code> | Must be a valid network device name | Yes | |
| ips | IP address | The IP address of the new interface | Must be a valid IP address | Yes | |
| gateway | IP address | The IP address of the default routing gateway | The gateway must reside within the same IP network of ip-address (as described by netmask). Not relevant for IB / L2 non-routable networks | No | |
| netmask | Number | Number of bits in the net mask, for example, the net mask of 255.255.0.0 has 16 netmask bits | Describes the number of bits that identify a network ID (also known as CIDR) | No | |



Note: When configuring an InfiniBand client, do not pass the `--ips`, `--netmask`, and `--gateway` parameters.



Note: InfiniBand clients can only join a cluster with InfiniBand backends. It is not possible to mix InfiniBand and Ethernet clients/backends.

Applying the host configuration using the CLI

Command

```
weka cluster host apply
```

After successfully configuring the host and its network device, run the following command to finalize the configuration by activating the host:

```
weka cluster host apply <host-id> [--force]
```

Parameter

| Name | Type | Value | Limitation | Mandatory | Default |
|-------------|------------------------|---|---------------------------------|------------------|----------------|
| hostid | Comma-separated string | Identifier of host to be added to the cluster | Must be a valid host identifier | Yes | |
| force | Boolean | Do not prompt for confirmation | | No | Off |

Chapter 6: Managing snapshots

How to manage snapshots using the CLI.

Viewing snapshots using the CLI

Command

```
weka fs snapshot
```

This command is used to display all snapshots of all filesystems in a single table.

Creating a snapshot using the CLI

Command

```
weka fs snapshot create
```

Use the following command line to add a snapshot:

```
weka fs snapshot create <file-system> <name> [<access-point>] [--source-snap=<source>] [--is-writable]
```

Table 2 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|-------------------------------------|----------------------|-----------|---------|
| file-system | String | A valid filesystem identifier | Must be a valid name | Yes | |
| name | String | Unique name for filesystem snapshot | Must be a valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|--------------|---------|--|----------------------|-----------|-------------------------------------|
| access-point | String | Name of newly-created directory for filesystem level snapshots, which will serve as the access point for the snapshots | Must be a valid name | Yes | |
| source | String | Must be an existing snapshot | Must be a valid name | No | Filesystem snapshot the file system |
| is_writable | Boolean | Sets the created snapshot to be writable | | No | False |

Deleting a snapshot using the CLI

Command

```
weka fs snapshot delete
```

Use the following command line to delete a snapshot:

```
weka fs snapshot delete <file-system> <name>
```

Table 3 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|-------------------------------------|----------------------|-----------|---------|
| file-system | String | A valid filesystem identifier | Must be a valid name | Yes | |
| name | String | Unique name for filesystem snapshot | Must be a valid name | Yes | |

Restoring a filesystem or snapshot from another snapshot using the CLI

You can restore a filesystem or snapshot from another snapshot.

Command

```
weka fs restore
```

Use the following command line to restore a snapshot:

```
weka fs restore <file-system> <source-name>
```

Command

```
weka fs snapshot copy
```

Use the following command line to restore a snapshot to another snapshot:

```
weka fs snapshot copy <file-system> <source-name> <destination-name>
```

Table 4 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------------|--------|--|------------------------------|-----------|---------|
| file-system | String | A valid filesystem identifier | Must be a valid name | Yes | |
| source-name | String | Unique name for the source of the snapshot | Must be a valid name | Yes | |
| destination-name | String | Name of the destination to which the snapshot should be copied | Must be an existing snapshot | Yes | |



Caution: When restoring a filesystem from a snapshot (or copying over an existing snapshot), the filesystem data and metadata are changed. Make sure IOs to the filesystem are stopped during this time.

Updating a snapshot using the CLI

Command

```
weka fs snapshot update
```

This command changes the snapshot attributes. Use the following command line to update an existing snapshot:

```
weka fs snapshot update <file-system> <name> [--new-name=<new-name>] [--iswritable] [--access-point=<access-point>]
```

Parameters

| Name | String | Value | Limitations | Mandatory | Default |
|--------------|---------|--|-----------------------|-----------|---------|
| file-system | String | A valid filesystem identifier. | Must be a valid name | Yes | |
| name | String | Unique name for the updated snapshot | Must be a valid name. | Yes | |
| new-name | String | New name for the updated snapshot | Must be a valid name. | No | |
| is-writable | Boolean | Sets the snapshot to be writable | | No | |
| access-point | String | Name of directory for snapshot, which will serve as the access point for the snapshot. | Must be a valid name. | No | |

Uploading a snapshot using the CLI

Command

weka fs snapshot upload

Use the following command line to upload an existing snapshot:

```
weka fs snapshot upload <file-system> <snapshot>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|------------------------|-------------|-----------|---------|
| filesystem | String | Name of the filesystem | | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|----------------------------|--|-----------|---------|
| snapshot | String | Name of snapshot to upload | Must be a snapshot of the <file-system> filesystem | Yes | |

Creating a filesystem from a snapshot using the CLI

Command

weka fs download

Use the following command line to create a filesystem from an existing snapshot:

```
weka fs download <name> <group-name> <total-capacity> <ssd-capacity> <obs> <locator>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|----------|--|-------------|-----------|---------|
| name | String | Name of the filesystem to be created | | Yes | |
| group-name | String | Name of the filesystem group in which the new filesystem will be placed | | Yes | |
| total-capacity | Capacity | Total capacity of the downloaded filesystem | | Yes | |
| ssd-capacity | Capacity | SSD capacity of the downloaded filesystem | | Yes | |
| obs | String | Object store name for tiering | | Yes | |
| locator | String | Object store locator obtained from a previously successful snapshot upload | | Yes | |

The locator is either a locator saved previously for disaster scenarios, or can be obtained using the **weka fs snapshot** command on a system with a live filesystem with snapshots.



Note: For encrypted filesystem, when downloading the same KMS master-key should be used to decrypt the snapshot data. For more information about encryption, see KMS management in the *Content Software for File User Guide*.

Chapter 7: Managing NFS

How to manage NFS networking configuration (interface groups) and access control (client access groups) using the CLI.

Managing NFS networking configuration (interface groups)

Defining interface groups using the CLI

Command

```
weka nfs interface-group add
```

Use the following command line to add an interface group:

```
weka nfs interface-group add <name> <type> [--subnet=<subnet>] [--gateway=<gw>]
```

Parameters

| Name | Type | Value | Limitation | Mandatory | Default |
|--------|--------|---|-----------------|-----------|-----------------|
| name | String | Unique interface group name | None | Yes | |
| type | String | Group type | Can only be NFS | Yes | NFS |
| subnet | String | The subnet mask in the 255.255.0.0 format | Valid netmask | No | 255.255.255.255 |
| gw | String | Gateway IP | Valid IP | No | 255.255.255.255 |

Setting interface group ports using the CLI

Command

```
weka nfs interface-group port add
```

Command

```
weka nfs interface-group port delete
```

Use the following command lines to add or delete an interface group port:

```
weka nfs interface-group port add <name> <host-id> <port>
weka nfs interface-group port delete <name> <host-id> <port>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|--|--------|------------------------------------|---------------|-----------|---------|
| name | String | Interface group name | None | Yes | |
| host-id | String | Host ID on which the port resides* | Valid host ID | Yes | |
| port | String | Port's device, for example, eth1 | Valid device | Yes | |
| * It is possible to obtain host IDs using the following command: | | | | | |
| weka cluster host -H=<hostname> | | | | | |

Setting interface group IPs using the CLI

Command

```
weka nfs interface-group ip-range add
```

Command

```
weka nfs interface-group ip-range delete
```

Use the following command lines to add or delete an interface group IP:

```
weka nfs interface-group ip-range add <name> <ips>
weka nfs interface-group ip-range delete <name> <ips>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|----------------------|-------------|-----------|---------|
| name | String | Interface group name | None | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|----------|----------------|-----------|---------|
| ips | String | IP range | Valid IP range | Yes | |

Managing NFS Access Control (client access groups)

Defining client access groups using the CLI

Command

weka nfs client-group

Use the following command lines to add or delete a client access group:

```
weka nfs client-group add <name>
weka nfs client-group delete <name>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|------------|-------------|-----------|---------|
| name | String | Group name | Valid name | Yes | |

Managing client access groups using the CLI

Command

weka nfs rules

Use the following command lines to add or delete a client group DNS:

```
weka nfs rules add dns <name> <dns>
weka nfs rules delete dns <name> <dns>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|------------|-------------|-----------|---------|
| name | String | Group name | Valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|-----------------------------------|-------------|-----------|---------|
| dns | String | DNS rule with *?[] wildcard rules | | Yes | |

Adding or deleting an IP using the CLI

Command

weka nfs rules

Use the following command lines to add or delete a client group IP:

```
weka nfs rules add ip <name> <ip>
weka nfs rules delete ip <name> <ip>
```

Parameters

| Name | Type | Value | Limitation | Mandatory | Default |
|------|--------|---|------------|-----------|---------|
| name | String | Group name | Valid name | Yes | |
| ip | String | IP with netmask rule, in the 1.1.1.1/255.255.0.0 format | Valid IP | Yes | |

Managing NFS client permissions

Managing NFS client permissions using the CLI

Command

weka nfs permission

Use the following command lines to add, update, or delete NFS permissions:

```
weka nfs permission add <filesystem> <group> [--path path] [--permission-type permission-type] [--root-squashing root-squashing] [--anon-uid anon-uid] [--anon-gid anon-gid] [--obs_direct]
```

```
weka nfs permission update <filesystem> <group> [--path path] [--permission-type permission-type] [--root-squashing root-squashing] [--anon-uid anon-uid] [--anon-gid anon-gid] weka nfs permission delete <filesystem> <group> [--path path]
```

```
weka nfs permission delete <filesystem> <group> [--path path]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-----------------|--------|--|---------------------------------|------------------------------------|---------|
| filesystem | String | Filesystem name | Existing filesystem | Yes | |
| group | String | Client group name | Existing client group | Yes | |
| path | String | The root of the share | Valid path | No | / |
| permission-type | String | Permission type RO: readonly RW: readwrite | RO: readonly RW: readwrite | No | RW |
| root-squashing | String | Root squashing | on/off | No | On |
| anon-uid | Number | Anonymous user ID (relevant only for root squashing) | Valid UID (between 0 and 65535) | Yes (if root squashing is enabled) | 65534 |
| anon-gid | Number | Anonymous user group ID (relevant only for root squashing) | Valid GID (between 0 and 65535) | Yes (if root squashing is enabled) | 65534 |

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|------|---------|--|-----------|---------|
| obs-direct | | Boolean | For more information, see object-store direct mount option, in the <i>Content Software for File User Guide</i> . | No | No |

Chapter 8: Managing SMB

Server message block (SMB) management, such as setting up an SMB cluster over Content Software for File filesystems, is described including managing the cluster itself using the CLI.

Showing an SMB cluster using the CLI

Command

```
weka smb cluster
```

Use this command to view information about the SMB cluster managed by the Content Software for File system.

Showing an SMB domain configuration using the CLI

Command

```
weka smb domain
```

Use this command to view information about the SMB domain configuration.

Creating an SMB cluster using the CLI

Command

```
weka smb cluster create
```

Use the following command line to create a new SMB cluster to be managed by the Content Software for File system:

```
weka smb cluster create <name> <domain> [--samba-hosts samba-hosts]... [--smb-ips-pool smb-ips-pool]... [-- smb-ips-range smb-ips-range] [--domainnetbios-name domain-netbios-name] [--idmap-backend idmap-backend] [--joined-domainmapping-from-id joined-domain-mapping-from-id] [--joined-domainmapping-to-id joined-domain-mapping-to-id] [--default-domain-mapping-from-id default-domain-mapping-from-id] [--default-domain-mapping-to-id defaultdomain-mapping-to-id] [--encryption encryption]
```


Table 5 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------------------------|------------------------------|---|--|-----------|--------------------------------|
| name | String | NetBIOS name for the SMB cluster | Must be a valid name (ASCII) | Yes | |
| domain | String | The domain which the SMB cluster is to join | Must be a valid name (ASCII) | Yes | |
| samba-hosts | Comma-separated strings | List of 3-8 Content Software for File system hosts to participate in the SMB cluster, based on the host IDs in Content Software for File | Must be valid host IDs | Yes | |
| smb-ips-pool | Comma-separated IP addresses | The public IPs used as floating IPs for the SMB cluster to serve the SMB over and thereby provide HA; should not be assigned to any host on the network | Must be valid IP addresses | No | |
| smb-ips-range | IP address range | The public IPs used as floating IPs for the SMB cluster to serve the SMB over and thereby provide HA; should not be assigned to any host on the network | Format: A.B.C.D-E for example, 10.10.0.1-100 | No | |
| domain-netbios-name | String | Domain NetBIOS name | Must be a valid name (ASCII) | No | First part of domain parameter |
| idmap-backend | String | The Id mapping method to use | <code>rfc2307</code> or <code>rid</code> | No | <code>rfc2307</code> |
| joined-domain-mapping-from-id | Number | The first ID of the range for the main AD ID mapping | | No | 0 |

| Name | Type | Value | Limitations | Mandatory | Default |
|--------------------------------|--------|--|-------------|-----------|------------|
| joined-domain-mapping-to-id | Number | The last ID of the range for the main AD ID mapping | | No | 0 |
| default-domain-mapping-from-id | Number | The first ID of the range for the default AD ID mapping (for trusted domains that have no range defined) | | No | 4290000001 |
| default-domain-mapping-to-id | Number | The last ID of the range for the default AD ID mapping (for trusted domains that have no range defined) | | No | 4291000000 |

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|--|--|-----------|---------|
| encryption | String | The global encryption policy to use. <code>enabled</code> - enables encryption negotiation but doesn't turn it on automatically for supported sessions and share connections. <code>disabled</code> - doesn't support encrypted connections. <code>desired</code> - enables encryption negotiation and turns on data enabled, disabled, desired or required No enabled encryption on supported sessions and share connections. <code>required</code> - enforces data encryption on sessions and share connections. Clients that do not support encryption will be denied access to the server. | enabled, disabled, desired or required | No | enabled |



Note: All IPs must reside on the same subnet, in order to enable HA through IP takeover.

- The IPs must be configured but MUST NOT be in use by any other application/host in the subnet, including Content Software for File system management nodes, system IO nodes, or system NFS floating IPs.
- The `--smb-ips` parameter is supposed to accept the public IPs that the SMB cluster will expose. To mount the SMB cluster in an HA manner, clients should be mounted using one of the exposed public IPs, thereby ensuring that they will automatically reconnect if one of the SMB hosts fails.
- If it is necessary to set global options to the SMB library, contact customer support.

For example:

```
weka smb cluster create wekaSMB mydomain --samba-hosts 0,1,2,3,4 --smb-ips-pool 1.1.1.1,1.1.1.2 --smb-ips-range 1.1.1.3-5
```

In this example of a full command, an SMB cluster is configured over the Content Software for File system hosts 0-4. The SMB cluster is called `wekaSMB`, the domain name is called `mydomain`, and is directed to use public IPs 1.1.1.1 to 1.1.1.5.

Checking status of SMB host readiness using the CLI

Command

```
weka smb cluster status
```

Use this command to check the status of the hosts which are part of the SMB cluster. Once all host are prepared and ready, it is possible to join an SMB cluster to an Active Directory.

Joining an SMB cluster to an Active Directory using the CLI

Command

```
weka smb domain join
```

Use the following command line to join an SMB domain to an Active Directory:

```
weka smb domain join <username> <password>
```

Parameters

| Name | Type | Value | Limitation | Mandatory | Default |
|----------|--------|--|----------------------------------|-----------|---------|
| username | String | Name of a user with permissions to add a machine to the domain | Must be a valid name (ASCII) | Yes | |
| password | String | The password of the user | Must be a valid password (ASCII) | Yes | |

In order to join another Active Directory to the current SMB cluster configuration, it is necessary to leave the current Active Directory. This is performed using the following command line:

```
weka smb domain leave <username> <password>
```

On completion of this operation, it is possible to join another Active Directory to the SMB cluster.



Note: To configure a new SMB cluster, the current SMB cluster has to be deleted.

Deleting an SMB cluster using the CLI

Command

```
weka smb cluster destroy
```

Use this command to destroy an SMB cluster managed by the Content Software for File system. Deleting an existing SMB cluster managed by the system does not delete the backend Content Software for File filesystems, but removes the SMB share exposures of these filesystems.



Note: Editing an existing cluster is not supported. Consequently, to change an SMB cluster configuration, the cluster has to be deleted and recreated.

Configuring trusted domains using the CLI

How to configure trusted domains using the CLI.

Listing trusted domains using the CLI

Command

```
weka smb cluster trusted-domains
```

Use this command to list all the configured trusted domains and their ID ranges.

Adding trusted domains using the CLI

Command

```
weka smb cluster trusted-domains add
```

Use the following command line to add an SMB trusted domain:

```
weka smb cluster trusted-domains add <domain-name> <from-id> <to-id>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|---|---|-----------|---------|
| domain-name | String | The name of the domain being added | Must be a valid name (ASCII) | Yes | |
| from-id | Number | The first ID of the range for the domain ID mapping | The range cannot overlap with other domains | Yes | |
| to-id | Number | The last ID of the range for the domain ID mapping | The range cannot overlap with other domains | Yes | |

Removing trusted domains using the CLI**Command**

```
weka smb cluster trusted-domains remove
```

Use the following command line to remove an SMB trusted domain:

```
weka smb cluster trusted-domains remove <domain-id>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-----------|--------|---|-------------|-----------|---------|
| domain-id | Number | The internal ID of the domain to remove | | Yes | |

Listing SMB shares using the CLI**Command**

```
weka smb share
```

Use this command to list all existing SMB shares.

Adding SMB shares using the CLI

Command

```
weka smb share add
```

Use the following command line to add a new share to be exposed to SMB:

```
weka smb share add <share-name> <fs-name> [--description description] [--internalpath internal-path] [--file-create-mask mask] [--directory-create-mask mask] [--obs_direct] [--encryption encryption] [--read-only] [--user-list-type list-type] [--users users]...
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-----------------------|---------|--|------------------------------|-----------|---------|
| share-name | String | Name of the share being added | Must be a valid name (ASCII) | Yes | |
| fs-name | String | Name of the filesystem to share | Must be a valid name | Yes | |
| description | String | Description of what the share will receive when viewed remotely | Must be a valid string | No | |
| internal-path | String | The internal path within the filesystem (relative to its root) which will be exposed | Must be a valid path | No | |
| file-create-mask | String | POSIX permissions for the file created through the SMB share | Numeric (octal) notation | No | 0744 |
| directory-create-mask | String | POSIX permissions for directories created through the SMB share | Numeric (octal) notation | No | 0755 |
| acl | Boolean | Enable Windows ACLs on the share (which will be translated to POSIX) | Up to 16 ACEs per file | No | No |

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|---------|--|-------------------------------------|-----------|---------|
| obs-direct | Boolean | For more information about obs-direct, see Time-based Policies for the Control of Data Storage Location in the Content Software for File User Guide. | Direct Mount section | No | No |
| encryption | String | The share encryption policy. cluster_default - the share encryption policy will follow the global SMB cluster setting desired - turns on data encryption for this share for clients that support encryption if negotiation has been enabled globally. required - enforces encryption for the shares. Clients that do not support encryption will be denied access to the share. If the global option is set to disabled access will be denied to these shares for all clients | cluster_default desired or required | No | Cluster |
| read-only | Boolean | Sets the share as read-only. Users cannot create or modify files in this share | | | No |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|-----------------------------------|---|--|-----------|---------|
| user-list-type | String | The type of initial permissions list for users | The type of initial permissions list for users <code>read_only</code> - list of users that will not be given write access to the share, regardless of the read-only setting. <code>read_write</code> - list of users that will be given write access to the share, regardless of the read-only setting. <code>valid</code> - list of users that are allowed to log-in to this share SMB service (empty list - all users are allowed) <code>invalid</code> - list of users that are not allowed to log-in to this share SMB service | No | |
| users | A comma-separated list of strings | A list of users to use with the user-list-type list. Can use the <code>@</code> notation to allow groups of users, for example, <code>root</code> , <code>Jack</code> , <code>@domain\admins</code> | Up to 8 users/groups for all lists combined per share | No | Empty |



Note: If it is necessary to set share specific options to the SMB library, contact customer support.

For example: The following is an example for adding users to a share mounted on a filesystem named "default":

```
weka smb share add rootShare default
weka smb share add internalShare default --internal-path some/dir --
description "Exposed share"
```

In this example, the first SMB share added has the Content Software for File system share for default. The second SMB share has internal for default.

Updating SMB shares using the CLI

Command

```
weka smb share update
```

Use the following command line to update an existing share:

```
weka smb share update <share-id> [--encryption encryption]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|--|-------------------------------------|-----------|---------|
| share-id | Number | The ID of the share to be updated | Must be a valid share ID | Yes | |
| encryption | String | The share encryption policy. desired - turns on data encryption for this share for clients that support encryption if negotiation has been enabled globally. required - enforces encryption for the shares. Clients that do not support encryption will be denied access to the share. If the global option is set to disabled access will be denied to these shares for all clients | cluster_default desired or required | No | |

Controlling SMB shares users lists using the CLI

How to control SMB shares users lists using the CLI.

Showing SMB share user lists using the CLI

Command

```
weka smb share lists show
```

Use this command to view the various user-list settings:

Adding SMB share user to list using the CLI

Command

```
weka smb share lists add
```

Use the following command line to add users to a share user-list:

```
weka smb share lists add <share-id> <user-list-type> <--users users>...
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|--------|--|---|-----------|---------|
| share-id | Number | The ID of the share to be updated | Must be a valid share ID | Yes | |
| user-list-type | String | The type of permissions list for users | <p><code>read_only</code> - list of users that will not be given write access to the share, regardless of the read-only setting.</p> <p><code>read_write</code> - list of users that will be given write access to the share, regardless of the read-only setting.</p> <p><code>valid</code> - list of users that are allowed to log-in to this share SMB service (empty list - all users are allowed)</p> <p><code>invalid</code> - list of users that are not allowed to log-in to this share SMB service</p> | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------|-----------------------------------|---|---|-----------|---------|
| users | A comma-separated list of strings | A list of users to add to the <code>user-list-type</code> list. Can use the <code>@</code> notation to allow groups of users, for example <code>root, Jack, @domain\admins</code> | Up to 8 users/groups for all lists combined per share | Yes | |

Removing SMB share user from lists using the CLI

Command

```
weka smb share lists remove
```

Use the following command line to remove users from a share user-list:

```
weka smb share lists remove <share-id> <user-list-type> <--users users>...
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|-----------------------------------|--------------------------|-----------|---------|
| share-id | Number | The ID of the share to be updated | Must be a valid share ID | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|-----------------------------------|--|--|-----------|---------|
| user-list-type | String | The type of permissions list for users | <p><code>read_only</code> - list of users that will not be given write access to the share, regardless of the read-only setting.</p> <p><code>read_write</code> - list of users that will be given write access to the share, regardless of the read-only setting.</p> <p><code>valid</code> - list of users that are allowed to log-in to this share SMB service (empty list - all users are allowed)</p> <p><code>invalid</code> - list of users that are not allowed to login to this share SMB service</p> | Yes | |
| users | A commas eparated list of Strings | A list of users to remove from the user-list-type list. Can use the @ notation to allow groups of users, for example, <code>root, Jack, @domain \admins</code> | Up to 8 users/groups for all lists combined per share | Yes | |

Resetting SMB share user lists using the CLI

Command

```
weka smb share lists reset
```

Use the following command line to remove all users from a share user-list:

```
weka smb share lists reset <share-id> <user-list-type>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------------|--------|---------------------------------------|---|-----------|---------|
| share-id | Number | The ID of the share to be updated | Must be a valid share ID | Yes | |
| user-list-type | String | The type of permissions list to reset | <p><code>read_only</code> - list of users that will not be given write access to the share, regardless of the <code>read-only</code> setting.</p> <p><code>read_write</code> - list of users that will be given write access to the share, regardless of the <code>read-only</code> setting.</p> <p><code>valid</code> - list of users that are allowed to log-in to this share SMB service (empty list - all users are allowed)</p> <p><code>invalid</code> - list of users that are not allowed to log-in to this share SMB service</p> | Yes | |

Removing SMB shares using the CLI**Command**

```
weka smb share remove
```

Use the following command line to remove a share exposed to SMB:

```
weka smb share remove <share-id>
```

Parameters

| Name | Type | Value | Limitation | Mandatory | Default |
|----------|--------|-----------------------------------|--------------------------|-----------|---------|
| share-id | String | The ID of the share to be removed | Must be a valid share ID | Yes | |

For example: The following is an example for removing an SMB share defined as ID 1:

```
weka smb share remove 1
```

Chapter 9: Managing alerts

How to manage alerts using the CLI. For a list of alerts, see List of alerts in the *Content Software for File User Guide*.

Listing alert types using the CLI

Command

```
weka alerts types
```

Use this command to lists all possible types of alerts that can be returned from the Content Software for File cluster.

Describing alerts using the CLI

Command

```
weka alerts describe
```

Use this command to describe all the alert types that might be returned from the Content Software for File cluster along with possible action items for each alert.

Viewing alerts using the CLI

Command

```
weka alerts
```

Use the following command line to list all alerts (muted and unmuted) in the Content Software for File cluster:

```
weka alerts [--muted]
```


Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------|---------|--|-------------|-----------|---------|
| muted | Boolean | List muted alerts alongside the unmuted ones | | No | False |

Muting alerts using the CLI**Command**

```
weka alerts mute
```

Use the following command line to mute an alert-type:

```
weka alerts mute <alert-type> <duration>
```

Muted alerts will not be prompted when listing active alerts. Alerts cannot be suppressed indefinitely, so a duration for the muted period must be provided. After expiry of the muted period, the alert-type is automatically unmuted.

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|---|----------------------------------|-----------|---------|
| alert-type | String | An alert-type to mute, use <code>weka alerts types</code> to list types | | Yes | |
| duration | String | How long to mute this alert type for | Format: 3s, 2h, 4m, 1d, 1d5h, 1w | Yes | |

Unmuting alerts using the CLI**Command**

```
weka alerts unmute
```

Use the following command line to unmute a previously-muted alert-type:

```
weka alerts unmute <alert-type>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|--|-------------|-----------|---------|
| alert-type | String | An alert-type to unmute, use <code>weka alerts types</code> to list types. For a list of alerts, see List of alerts in the <i>Content Software for File User Guide</i> . | Yes | Yes | |

Chapter 10: Managing events

How to manage events using the CLI. For a list of events, see List of events in the *Content Software for File User Guide*.

Viewing events using the CLI

Command

weka events

Use the following command line to list events in the Content Software for File cluster:

```
weka events [--num-results num-results] [--start-time <start>] [--end-time <end>] [--severity severity] [--direction direction] [-- fetch-order fetchorder] [--type-list type-list] [--exclude-type-list exclude-type-list] [--category-list category-list] [--by-digested- time] [--show-internal] [--rawunits] [--UTC]
```

Parameters

| Name | Type | Value | Limitations | Mandato ry | Default |
|-----------------|---------|---|--|---------------|--------------------------------|
| num- results | Integer | Maximum number of events to display | Positive integer. 0 shows all events | No | 50 |
| start-time | String | Include events that occurred at this start time and later | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | No | -365 days |
| end-time | String | Include events that occurred up to this time | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | No | Set to a time represents 'now' |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------------|---------|---|---|-----------|--|
| severity | String | Include events with this level of severity and higher | 'info', 'warning', 'minor', 'major' or 'critical' | No | INFO |
| direction | String | Sort events by ascending or descending time | 'asc' or 'dsc' | No | asc |
| fetch-order | String | Fetch from end-time backwards or from start-time forwards | 'fw' or 'bw' | No | bw |
| type-list | String | Filter events by type (can be used multiple times) | Use <code>weka events list-types</code> to see available types | No | None |
| exclude-type-list | String | Filter-out events by type (can be used multiple times) | Use <code>weka events list-types</code> to see available types | | |
| category-list | String | Include only events matching the defined category | Categories can be Alerts, Cloud, Clustering, Drive, Events, Filesystem, IO, InterfaceGroup, Licensing, NFS, Network, Node, ObjectStorage, Raid, Statistics, System, Upgrade, and User | No | All |
| digested-time | Boolean | Query and sort results by digested time | | No | False |
| show-internal | Boolean | Also displays internal events | | No | False |
| raw-units | Boolean | Print values in raw units (bytes, seconds, among others) | | No | Human-readable format, for example, 1KiB 234MiB 2GiB |

| Name | Type | Value | Limitations | Mandatory | Default |
|------|---------|--------------------|-------------|-----------|-------------------|
| UTC | Boolean | Print times in UTC | | No | Host's local time |

Listing local events using the CLI

Command

weka events list-local

Use the following command line to list recent events on the specific host running the command from (can be useful for cases there is no connectivity to support cloud, no connectivity from a specific host, or for hosts which are not part of the cluster):

```
weka events list-local [--start-time <start>] [--end-time <end>] [--next next] [--stem-mode] [--show-internal] [--raw-units] [--UTC]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------|--------|---|--|-----------|--------------------------------|
| start-time | String | Include events that occurred at this start time and later | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | no | -365 days |
| end-time | String | Include events that occurred up to this time | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | No | Set to a time represents 'now' |
| next | String | Identifier to the next page of events | As returned in the previous call to <code>weka events list-local</code> | No | |

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|---------|--|-------------|-----------|--|
| stem-mode | Boolean | Displays events when the host has not been attached to the cluster | | No | False |
| show-internal | Boolean | Also displays internal events | | No | False |
| raw-units | Boolean | Print values in raw units (bytes, seconds, among others.) | | No | Human-readable format, for example, 1KiB 234MiB 2GiB |
| UTC | Boolean | Print times in UTC | | No | Host's local time |

Triggering a custom event using the CLI

Command

```
weka events trigger-event
```

It can be useful to mark specific activities, maintenance work, or important changes or new usage of the system, and see that as part of the system events timeline. To trigger a custom event use:

```
weka events trigger-event <text>
```

Chapter 11: Managing statistics

How to manage statistics available in the Content Software for File system using the CLI. For a list of statistics, see List of statistics in the *Content Software for File User Guide*.

Listing statistic types using the CLI

Command

```
weka stats list-types
```

Use the following command line to obtain statistics definition information:

```
weka stats list-types [<name-or-category>] [--show-internal]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------------|---------|-----------------------------------|------------------------|-----------|---------|
| name-or-category | String | Name or category to filter by | Valid name or category | No | |
| show-internal | Boolean | Also displays internal statistics | | No | False |

Viewing statistics realtime using the CLI

Command

```
weka stats realtime
```

Use the following command line to obtain the current performance-related statistics of the hosts, in a one-second interval:

```
weka stats realtime [<node-ids>] [--raw-units] [--UTC]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-----------|-------------------------|--|-------------|-----------|---|
| node-ids | Comma-separated strings | Only show realtime stats of these nodes | | No | |
| raw-units | Boolean | Print values in raw units (bytes, seconds, among others) | | No | Humanreadable format, for example, 1KiB 234MiB 2GiB |
| UTC | Boolean | Print times in UTC | | No | Host's local time |

Reading statistic information using the CLI**Command****weka stats**

The collected statistics are helpful to analyze system performance and determine the source of any problems as the Content Software for File system runs, according to several categories. When each category is selected, a list of the possible statistics that can be selected is displayed.



Note: All statistics are averaged over one-second intervals. Consequently, "total" or other aggregates relate to a specific minute.

Use the following command line to manage filters and read statistics:

```
weka stats [--start-time <start>] [--end-time <end>] [--interval interval]
[--resolution-secs <secs>] [--category category] [--stat stat]
[--node-ids odeids] [--param param] [--accumulated] [--per-node] [--no-zeros]
[--showinternal] [--json] [--raw-units] [--UTC]
```


Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-----------------|--------|---|--|-----------|--------------|
| start-time | String | Start time of the reported period | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | Yes | |
| end-time | String | End time of the reported period | Format: 5m, -5m, -1d, -1w, 1:00, 01:00, 18:30, 18:30:07, 2018-12-31 10:00, 2018/12/31 10:00, 2018-12-31T10:00, 9:15Z, 10:00+2:00 | No | Current time |
| interval | String | Period of time to be reported | Valid interval in seconds (positive integer number) | Yes | |
| resolution-secs | String | Length of each interval in the reported period | Must be multiples of 60 seconds | No | 60 |
| category | String | Specific categories for retrieval of appropriate statistics | Valid existing categories: CPU, Object Store, Operations, Operations (NFS), Operations (Driver), and SSD | No | All |
| stat | String | Statistics names | Valid statistics names | No | All |
| node-ids | String | Node id | Valid node-id | No | All |

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------|---------|--|-----------------|-----------|--|
| param | String | For parameterized statistics, retrieve only the instantiations where the specified parameter is of the specified value. Multiple values can be supplied for the same key, for example, '-- param method:putBlocks -- param method:initBlock' | Format: key:val | No | |
| accumulated | Boolean | Display accumulated statistics, not rate statistics | | No | False |
| per-node | Boolean | Does not aggregate statistics across nodes | | No | False |
| no-zeros | Boolean | Filters results where the value is 0 | | No | False |
| show-internal | Boolean | Also displays internal statistics | | No | False |
| raw-units | Boolean | Print values in raw units (bytes, seconds, among others.) | | No | Human readable format, for example 1KiB 234MiB 2GiB |
| UTC | Boolean | Print times in UTC | | No | Host's local time |

Setting statistic retention using the CLI

Command

```
weka stats retention
```

Use the following command line to set the statistics retention period:

```
weka stats retention <--days days> [--dry-run]
```

Parameters

| Name | Type | Value | Limitation | Mandatory | Default |
|---------|---------|---|--|-----------|---------|
| days | Number | The number of days to keep the statistics. | Should have enough free disk space per server. | Yes | |
| dry-run | Boolean | Only test the required capacity per the retention period. | | No | |

Use `weka stats retention status` to view the current retention and `weka stats retention restore-default` to restore the default retention settings.

Chapter 12: Managing users

How to manage users for Content Software for File using the CLI.

Creating users using the CLI

Command

```
weka user add
```

Use the following command line to create a user:

```
weka user add <username> <role> [password]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|---|---|-----------|---|
| username | String | Name of the user to change the password for | Must be a valid local user | Yes | |
| role | String | Role of the new created user | regular, readonly, orgadmin or clusteradmin | Yes | |
| password | String | New password | | No | If not supplied, command will prompt to supply the password |

For example,

```
$ weka user add my_new_user S3cret regular
```

This command line creates a user with a username of `my_new_user`, a password of `S3cret` and a role of Regular user. It is then possible to display a list of users and verify that the user was created:

```
1 $ weka user
2 Username | Source | Role
3 -----+-----+-----
4 my_new_user | Internal | Regular
5 admin | Internal | Admin
```

Using the `weka user whoami` command, it is possible to receive information about the current user running the command. To use the new user credentials, use the `WEKA_USERNAME` and `WEKA_PASSWORD` environment variables:

```
$ WEKA_USERNAME=my_new_user WEKA_PASSWORD=S3cret 1 weka user whoami
2 Username | Source | Role
3 -----+-----+-----
4 my_new_user | Internal | Regular
```

Changing user passwords using the CLI

Command

`weka user passwd`

Use the following command line to change a local user password:

```
weka user passwd <password> [--username username]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|---|----------------------------|-----------|------------------------|
| password | String | New password | | Yes | |
| username | String | Name of the user to change the password for | Must be a valid local user | No | Current logged-in user |



Note: If necessary, provide or set `WEKA_USERNAME` or `WEKA_PASSWORD`.

Deleting users using the CLI

Command

```
weka user delete
```

To delete a user, use the following command line:

```
weka user delete <username>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|-----------------------------|-----------------------------|-----------|---------|
| username | String | Name of the user to delete. | Must be a valid local user. | Yes | |

For example:

```
$ weka user add my_new_user
```

Then run the **weka user** command to verify that the user was deleted:

```
1 $ weka user
2 Username | Source | Role
3 -----+-----+-----
4 admin | Internal | Admin
```

Configuring an LDAP user directory using the CLI

How to configure an LDAP user directory using the CLI.

Configuring an LDAP server using the CLI

Command

```
weka user ldap setup
```

Command

```
weka user ldap setup-ad
```

Two CLI commands are used to configure an LDAP user directory for user authentication. The first is for configuring a general LDAP server. The second is for configuring an Active Directory server.

To configure an LDAP server, use the following command line:

```

1 weka user ldap 1 setup <server-uri>
2 <base-dn>
3 <user-object-class>
4 <user-id-attribute>
5 <group-object-class>
6 <group-membership-attribute>
7 <group-id-attribute>
8 <reader-username>
9 <reader-password>
10 <cluster-admin-group>
11 <org-admin-group>
12 <regular-group>
13 <readonly-group>
14 [--server-timeout-secs server-timeout-secs]
15 [--protocol-version protocol-version]
16 [--user-revocation-attribute user-revocation-attrib

```

To configure an Active Directory server, use the following command line:

```

1 weka user ldap setup-ad <server-uri>
2 <domain>
3 <reader-username>
4 <reader-password>
5 <cluster-admin-group>
6 <org-admin-group>
7 <regular-group>
8 <readonly-group>
9 [--server-timeout-secs server-timeout-secs]

```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------------|--------|--|---|-----------|---------|
| server-uri | String | Either the LDAP server host name/IP or a URI | URI must be in format ldap://hostname:port or ldaps://hostname:port | Yes | |
| base-dn | String | Base DN under which users are stored | Must be valid name | Yes | |
| user-id-attribute | String | Attribute storing user IDs | Must be valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------------------------------|--------|--|---|-----------|---------|
| user-object-class | String | Object class of users | Must be valid name | Yes | |
| group-object-class | String | Object class of groups | Must be valid name | Yes | |
| group-membership-attribute | String | Attribute of group containing the DN of a user membership in the group | Must be valid name | Yes | |
| group-id-attribute | String | Attribute storing the group name | Name has to match names used in the <admin-group>, <regular group> and <readonly group> | Yes | |
| reader-username and reader-password | String | Credentials of a user with read access to the directory | Password is kept in the Content Software for File cluster configuration in plain text, as it is used to authenticate against the directory during user authentication | Yes | |
| cluster-admin-group | String | Name of group containing users defined with cluster admin role | Must be valid name | Yes | |
| org-admin-group | String | Name of group containing users defined with organization admin role | Must be valid name | Yes | |
| regular-group | String | Name of group containing users defined with regular privileges | Must be valid name | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|---------------------------|--------|--|---|-----------|---------|
| readonly-group | String | Name of group containing users defined with read only privileges | Must be valid name | Yes | |
| server-timeout-secs | Number | Server connection timeout | Seconds | No | |
| protocol-version | String | Selection of LDAP version | LDAP v2 or v3 | No | LDAP v3 |
| user-revocation-attribute | String | The LDAP attribute; when its value changes in the LDAP directory, user access and mount tokens are revoked | User must re-login after a change is detected | No | |

Viewing a configured LDAP user directory using the CLI

Command

```
weka user ldap
```

This command is used for viewing the current LDAP configuration used for authenticating users. The following command line is used:

```
weka user ldap [--HOST HOST]{--PORT PORT}
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|--------|-----------|----------------------|-----------|---------|
| HOST | String | Host name | Only valid host name | Yes | |
| PORT | String | Port name | Only valid port name | Yes | |

Disabling or enabling a configured LDAP user directory using the CLI

Command

```
weka user ldap disable
```

Command

```
weka user ldap enable
```

These commands are used for disabling or enabling user authentication through a configured LDAP user directory.



Note: It is not possible to delete an LDAP configuration; only disable it.

Chapter 13: Managing organizations

How to manage organizations using the CLI.

Managing organizations

Only users defined as Cluster Admins can manage organizations. When no organization is created, the root organization is the default organization and all operations are regular, that is, it is not necessary to authenticate the mounts or supply an organization name when logging in using the GUI or CLI.

Once a new organization is created, the organization name must be provided in every login command, using the `--org` attribute in the `weka user login` command.

Creating an organization using the CLI

Command

`weka org create`

Use the following command line to create an organization:

```
weka org create <name> <username> <password> [--ssd-quota ssd-quota] [--total-quota total-quota]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|--|----------------------|-----------|---------|
| name | String | Organization name | Must be a valid name | Yes | |
| username | String | Username of the created Organization Admin | Must be a valid name | Yes | |
| password | String | Password of the created Organization Admin | | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|--------|---|------------------------|-----------|-----------------|
| ssd-quota | Number | Allowed quota out of the system SSDs to be used by the organization | Must be a valid number | No | 0 (not limited) |
| total-quota | Number | Total allowed quota for the organization (SSD and object store) | Must be a valid number | No | 0 (not limited) |

Viewing organizations using the CLI

Command

weka org

```

1 # weka org
2
3 ID | Name | Allocated SSD | SSD Quota | Allocated Total | Total Q
4 ---+-----+-----+-----+-----+-----
5 0 | Root | 0 B | 0 B | 0 B | 0 B
6 1 | Local IT | 500.00 GB | 500.00 GB | 500.00 GB | 0 B
7 2 | CUSTOMER_1 | 100.00 GB | 300.00 GB | 200.00 GB | 900.00

```

Renaming organizations using the CLI

Command

weka org rename

Use the following command line to rename an organization:

```
weka org rename <org> <new-name>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|----------------|---------------------------------|-------------|-----------|---------|
| org | String/Integer | Current organization name or ID | | Yes | |

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|--------|-----------------------|-------------|-----------|---------|
| new-name | String | New organization name | | Yes | |

Updating an organization's quotas using the CLI

Command

```
weka org set-quota
```

Use the following command line to update an organization's quota:

```
weka org set-quota <org> [--ssd-quota ssd-quota] [--total-quota total-quota]
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------------|----------------|---|--|-----------|---------|
| org | String/Integer | Organization name or ID | The root organization (org ID = 0 cannot be limited) | Yes | |
| ssd-quota | Number | Allowed quota out of the system SSDs to be used by the organization | Must be a valid number | No | |
| total-quota | Number | Total allowed quota for the organization (SSD and object store) | Must be a valid number | No | |

Deleting an organization using the CLI

Command

```
weka org delete
```

Use the following command line to delete an organization:

```
weka org delete <org>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------|----------------|-------------------------|-------------|-----------|---------|
| org | String/Integer | Organization name or ID | | Yes | |



Caution: Deleting an organization is irreversible and will also remove all entities related to the organization, such as filesystems, object stores, and users.

Revoking user access using the CLI**Command**

```
weka user revoke-tokens
```

Use the following command to revoke internal user access to the system and mounting filesystems:

```
weka user revoke-tokens <username>
```

For LDAP users, access can be revoked by changing the `user-revocation-attribute` defined in the LDAP server configuration.

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|----------|----------------|--|-------------|-----------|---------|
| username | String/Integer | Valid user in the organization of the Organization Admin running the command | | Yes | |



Note: NFS and SMB are different protocols from WekaFS which require additional security considerations when used, for example, NFS permissions are granted per host, so permissions for accessing these hosts for NFS export should be handled carefully.

Chapter 14: Expansion of specific resources

Guidelines provided for the expansion processes that only involve the addition of specific resources.

Dynamic modifications using the CLI

Most modifications to host configurations can be performed dynamically, without deactivating the host. Such configurations include the addition or removal of memory and network resources, changing IPs, extending network subnets and limiting the Content Software for File system bandwidth on the host.

All these changes can be performed using the relevant **weka cluster host** command. Once this command is used with a specific `host-id` selected, it will be staged for update on the cluster. To view the un-applied configuration use:

Command

weka cluster host

To apply the changes or view un-applied configurations, use:

```
weka cluster host resources <host-id>
```

To view this configuration, use:

```
weka cluster host resources <host-id> --stable
```

The last local configuration (of a host that successfully joined a cluster) is saved. If a failure or problem occurs with the new configuration, the host will automatically revert to the last known good configuration. You can also apply these changes locally using:

Command

weka local resources apply

To apply the changes, use next the command.

```
weka cluster host apply <host-ids>
```

To add another network device to host-id 0 , run the following commands:

```
weka cluster host net add 0 --device=eth2  
weka cluster host apply 0
```

Memory modifications

To dynamically change the memory configuration, use the steps described for the configuration of memory on an active host, followed by the `weka cluster host apply` command.

For example: to change host-id 0 memory to 1.5 GB, run the following commands:

```
weka cluster host memory 0 1.5GB
weka cluster host apply 0
```



Note: It is possible to accumulate several changes on a host and apply only once on completion.

Network modifications

To dynamically change the network configuration, use the steps described for the Configuration of networking section in the *Content Software for File User Guide* on an active host, followed by the `weka cluster host apply` command.

For example: To add another network device to host-id 0, run the following commands:

```
weka cluster host net add 0 --device=eth2
weka cluster host apply 0
```



Note: It is possible to accumulate several changes on a host and apply only once on completion.

For additional information, contact customer support.

Host IPs modifications

To dynamically change the host's management IPs, you can use the `management-ips resource editing` command.

For example: To change the management IPs on host-id 0, run the following commands:

```
weka cluster host management-ips 192.168.1.10 192.168.1.20
weka cluster host apply 0
```



Note: The number of management IPs determines whether the host will use Highly Available Networking mode (HA), causing each IO process to use both hosts NICs. A host with 2 IPs will use HA mode and a host with only 1 IP will not use HA mode.

Local resources editing commands using the CLI

It is also possible to run modification commands locally on the host by connecting to the desired host and running a `local resources` command equivalent to its `weka cluster host` counterpart. These local commands have the same semantics of their remote counterparts only that they don't receive the host-id as the first parameter and operate instead on the local host. Commands that can be performed dynamically on an Active host:

```
weka local resources [--stable]
weka local resources apply
weka local resources net
weka local resources net add
weka local resources net remove
weka local resources memory
weka local resources bandwidth
weka local resources management-ips
weka local resources dedicate
```

The following commands cannot be performed on an Active host and require deactivating the host first using `weka cluster host deactivate`:

```
weka local resources failure-domain
weka local resources cores
```

Chapter 15: Managing clusters

How to expand and shrink a cluster in a homogeneous Content Software for File system configuration.



Note: The cluster expansion process described here is only applicable to a homogeneous Content Software for File system configuration, which is highly recommended. For non-homogeneous system configurations, contact your Hitachi representative.

Shrinking a cluster using the CLI

The procedures are described that are involved in the shrinking of a cluster and may be required when it is necessary to reallocate cluster hardware.

Options for shrinking a cluster

Cluster shrinking can involve either the removal of some of the assigned SSDs or the removal of hosts from the system. The following operations are available:

1. Listing all the drives and their states, in order to receive a view of currently-allocated resources and their status.
2. Deactivating drives as the first step before removing a host.
3. Removing (a subset of) SSD drives allocated for the cluster.
4. Deactivating hosts, which can be used after deactivating drives in preparation for the removal of the host.
5. Removing hosts in order to complete the cluster shrinking.

Listing drives and their states using the CLI

Command

```
weka cluster drive
```

Use this command to display a list of all the drives in the cluster and their status.

Deactivating a drive using the CLI

Command

```
weka cluster drive deactivate
```

Running this command will redistribute the stored data on the remaining drives and can be performed on multiple drives.



Note: After running this command, the deactivated drives will still appear in the list.



Note: It is not possible to deactivate a drive if it will lead to an unstable state, that is, if the system capacity after drive deactivation is insufficient for the SSD capacity of currently-provisioned filesystems.

Drive deactivation starts an asynchronous process known as phasing out, which is a gradual redistribution of the data between the remaining drives in the system. On completion, the phased-out drive is in an inactive state, that is, not in use by the Content Software for File system, but still appearing in the list of drives.



Note: Running the `weka cluster drive` command will display whether the redistribution is still being performed.

To deactivate a drive, run the following command:

```
weka cluster drive deactivate <uuids>
```

Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------|-------------------------|-----------------------------------|-------------|-----------|---------|
| uuids | Comma-separated strings | Comma-separated drive identifiers | | Yes | |

Removing a drive using the CLI

Command

```
weka cluster drive remove
```

This command is used to completely remove a drive from the cluster. After removal, the drive will not be recoverable. To remove a drive, run the following command:

```
weka cluster drive remove <uuids>
```

Table 6 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|-------|-------------------------|-------------------------|-------------|-----------|---------|
| uuids | Comma-separated strings | Comma-separated strings | | Yes | |

Deactivating an entire host using the CLI

Command

```
weka cluster host deactivate
```

This command is used as the first step when seeking to shrink a cluster. Running this command will automatically deactivate all the host's drives. To deactivate an entire host, run the following command:

```
weka cluster host deactivate <host-ids> [--allow-unavailable-host]
```

Table 7 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|------------------------|--------------------------|---|--|-----------|---------|
| host-ids | Space-separated integers | Space-separated host identifiers | | Yes | |
| allow-unavailable-host | Boolean | Allow deactivation of an unavailable host | If the host returns, it will join the cluster in an active state | No | No |

Removing a host using the CLI

Command

```
weka cluster host remove
```

Running this command will eliminate the host from the cluster, that is, the host will switch to the `stem` mode after the removal, at which point it can be reallocated either to another cluster or purpose. To remove the host from the cluster, run the following command:

```
weka cluster host remove <host-id>
```

Table 8 Parameters

| Name | Type | Value | Limitations | Mandatory | Default |
|---------|-------------------------|----------------------------------|-------------|-----------|---------|
| host-id | Comma-separated strings | Comma-separated host identifiers | | Yes | |

Chapter 16: Managing background tasks

The management of background tasks running on Content Software for File clusters is described.

Viewing background tasks using the CLI

Command

```
weka cluster tasks
```

It is possible to view currently-running background tasks, including their status and progress. This command is used for viewing all background tasks. For each task, a range of data can be displayed, as shown in the following example:

```
1 # weka cluster tasks
2 Type | State | Progress | Description
3 -----+-----+-----+-----
4 OBS_DETACH | RUNNING | 94 | Detaching Object Storage `obs_1` from
```

Limiting background tasks using the CLI

Command

```
weka cluster tasks limits
```

This command is used to view the currently-defined limits.

Command

```
weka cluster tasks limits set <cpu-limit limit>
```

This command is used to update the CPU limit.

Chapter 17: Running cluster diagnostics

The details for running cluster diagnostic commands are provided.

Managing diagnostics using the CLI

Command

```
weka diags
```

Command

```
weka local diags
```

The command `weka diags` is used for cluster-wide diagnostics from any host in the cluster. The `weka local diags` command creates diagnostics information about the Content Software for File software and saves it for further analysis by customer support.

The commands can be run with the following options:

```
weka [local] diags <--collect|--upload> [--pack-to dir]
```

When `weka local diags` receives a directory using the `-o` option, the diagnostics dump of the host is moved to that directory on completion of the collection process.



Note: In the following situations the local option should be used: when no functioning manager in the originating host or the hosts being addressed or when there is no connectivity between the manager and the cluster leader, the cluster has no leader, the local container is down, the host cannot reach the leader or a remote host fails to respond to the `weka diags remote` command.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact