



Hewlett Packard
Enterprise

HPE Superdome Flex Server Administration Guide

Abstract

System administration commands, procedures, security, and management interfaces for HPE Superdome Flex Server.

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Xeon®, and Optane™ are trademarks of Intel Corporation in the U.S. and other countries.

Redfish® is a trademark of Distributed Management Task Force, Inc.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Revision history

Part number	Publication date	Edition	Summary of changes
P06150-401a	February 2021	14	<ul style="list-style-type: none">Updated the RMC commands <code>apropos</code>, <code>clear</code>, <code>create</code>, <code>disable</code>, <code>enable</code>, <code>modify</code>, <code>remove</code>, <code>set</code>, <code>show</code>, and <code>setwith</code> with new attributes.Updated Monitoring Main Memory to display new attribute details for <code>show</code> command.Updated <code>apropos</code> command with a new attribute <code>enable_ler</code> in RMC commands.
P06150-401	October 2020	13	Minor editorial updates
P06150-009	August 2020	12	<ul style="list-style-type: none">Added information on HPE OneView supportMinor editorial updates
P06150-008a	May 2020	11	Minor editorial updates

Table Continued



Part number	Publication date	Edition	Summary of changes
P06150-008	March 2020	10	<ul style="list-style-type: none"> • Added user account password requirements • Added process for firmware downgrades • Revised RMC <code>set</code> command with new features • Minor editorial updates
P06150-007	December 2019	9	<ul style="list-style-type: none"> • Added support for HPE Serviceguard for Linux (SGLX) • Added HPE Persistent Memory information • Updated RMC commands to add new certificate authority (CA) and certificate revocation list (CRL) features • Revised HPE recommendations for RAS features • Minor editorial updates
P06150-006	August 2019	8	<ul style="list-style-type: none"> • Updated HPE Foundation Software network card consistent naming feature • Added online firmware update procedure • Added support for OneView Manage Mode • Added new RMC commands <code>restore</code> and <code>save</code>
P06150-005a	June 2019	7	Minor editorial updates
P06150-005	May 2019	6	<ul style="list-style-type: none"> • Added procedures for configuring auto-power capability, Hyper-Threading Technology, and RAS features • Added information on user roles • Added procedures to automatically name NICs, and activating extended tuning scripts using HPE Foundation Software • Updated IPMI information to include default behavior • Updated RMC commands to add new features including LDAP and user roles
P06150-004	February 2019	5	Added HPE recommendations for enabling RAS features
P06150-003	October 2018	4	Updated title to <i>HPE Superdome Flex Server Administration Guide</i>
P06150-002	August 2018	3	

Table Continued

Part number	Publication date	Edition	Summary of changes
P06150-001	March 2018	2	
007-6491-001	December 2017	1	Initial release as <i>HPE Superdome Flex Server Administrator Commands and Procedures Guide</i>



Contents

HPE Superdome Flex Server manageability features.....	8
Rack management controller (RMC) tasks.....	9
Connecting to the Rack Management Controller (RMC).....	9
Powering on and booting an HPE Superdome Flex Server system from a complete power off.....	9
Powering off an HPE Superdome Flex Server.....	10
Configuring the autopower capability.....	11
Configuring Hyper-Threading Technology.....	12
Configuring RAS.....	13
Configuring RMC user accounts and user roles.....	14
User roles.....	15
Retrieving HPE Superdome Flex Server system configuration information	15
Retrieving hardware configuration information.....	16
Monitoring log files.....	17
Updating firmware.....	18
Installing the latest complex firmware using SUM.....	19
Installing the latest complex firmware using HPE OneView.....	19
Updating firmware manually.....	19
Downgrading firmware.....	22
Creating a crash dump file.....	22
Managing system performance with FW.....	24
Managing system performance with HPE Foundation Software (HFS).....	26
Monitoring main memory.....	26
Retrieving main memory health information.....	28
About page migration.....	28
Monitoring system performance.....	29
Configuring CPU frequency scaling.....	30
Changing the CPU governor setting and frequency setting.....	30
Configuring turbo mode.....	31
Consistently naming network cards.....	31
Default Linux network card naming.....	32
Activating an extended tuning script	33
Additional HPE Foundation Software (HFS) utilities.....	33
HPE Foundation Software (HFS) commands.....	34
cpupower.....	34
gr_systat.....	35
hpe-auto-config.....	37
hpe_irqbalance	39
hubstats.....	40
nodeinfo.....	42
topology.....	43



Rack management controller (RMC) commands.....	48
acquit.....	49
add.....	50
apropos.....	52
baseiolist.....	53
cancel autopower.....	53
clear.....	54
collect debug.....	55
commands.....	56
connect.....	56
create.....	57
deconfig.....	58
disable.....	59
download.....	60
enable.....	61
exit.....	61
generate certificate.....	62
help.....	62
indict.....	63
ipmi.....	63
modify.....	64
ping.....	66
ping6.....	66
power.....	66
reboot.....	68
remove.....	68
restore.....	70
save.....	71
set.....	71
search and searchevents.....	75
show.....	77
test.....	84
update firmware.....	87
upload.....	87
version.....	88
Partition boot options.....	89
Managing with HPE OneView.....	90
Rack manager components.....	90
Add a rack manager	91
Remote management with Redfish.....	93
Out-of-Band management with Redfishtool - general procedure.....	93
Management with ipmitool.....	98
Retrieving ipmi help output.....	98



Running IPMI commands on the rack management controller (RMC) - general procedure.....	99
Out-of-band management with ipmitool.....	99
Running an <code>ipmi</code> command against a baseboard management controller (BMC) - general procedure.....	100
Running an <code>ipmi</code> command against one BMC	100
Running an <code>ipmi</code> command against all BMCs	101
Retrieving RMC or BMC power status information.....	101
Issuing power cycle commands.....	102
Sending a nonmaskable interrupt (NMI) signal to write a crash dump.....	102
Listing information for field replaceable units (FRUs).....	102
Listing the local area network (LAN) configuration.....	104
Listing environmental sensors and values.....	104
Retrieving the system event log (SEL).....	105
Opening a serial console.....	106
Disaster recovery with HPE Serviceguard for Linux (SGLX).....	107
HPE Persistent Memory on HPE Superdome Flex Server.....	108
Websites.....	109
Support and other resources.....	111
Accessing Hewlett Packard Enterprise Support.....	111
Accessing updates.....	111
Remote support.....	112
Warranty information.....	112
Regulatory information.....	112
Documentation feedback.....	113



HPE Superdome Flex Server manageability features

Superdome Flex Server has an advanced manageability system that is always on, constantly monitoring and managing the system components, fabric, and infrastructure for mission-critical high availability.

Major Superdome Flex Server management components include:

Rack Management Controller (RMC)

The main component in the management subsystem. The RMC or eRMC connects to all the system chassis through a physically secure private LAN. The RMC is a standalone 1U rack-mounted chassis.

Embedded Rack Management Controller (eRMC)

The eRMC is an option for one and two chassis systems. It provides RMC functionality using an existing BMC management processor instead of a separate 1U server.

Baseboard Management Controller (BMC)

Chassis-level management. Configures and manages the hardware in that chassis as well as providing virtual media and virtual keyboard, video, or mouse (KVM) features.

Supported manageability software includes:

- HPE OneView
- HPE Insight Remote Support
- HPE Serviceguard for Linux
- Smart Update Manager

For an overview of management features, see www.hpe.com/support/superdomeflex-management-whitepaper.



Rack management controller (RMC) tasks

The RMC includes a command set that you can use to perform administrative functions on the HPE Superdome Flex Server. This chapter shows how to complete several typical tasks.

Connecting to the Rack Management Controller (RMC)

You can connect to the RMC one of the following ways:

- Through a laptop attached to the RMC.

This connection is a direct-attached serial connection using the USB-A to Micro-USB-B cable that was included in the shipping crate.

For example, this type of connection can be useful if you lost the RMC password. With this kind of connection, the RMC does not prompt you for a password. For this reason, Hewlett Packard Enterprise recommends that you restrict physical access to the system as part of your site security procedures.

For information about connecting to the RMC using a laptop, see [HPE Superdome Flex Server User Guide](#).

- Through a network connection, through secure shell.

The following procedure explains how to connect to the RMC over a network connection. See [HPE Superdome Flex Server User Guide](#) for information about setting networking parameters on the RMC.

Procedure

1. Use the secure shell (`ssh`) command to log into the RMC.

The factory configures the default login for the RMC to be `administrator`.

2. Respond to the prompts for the login name and password.

```
remotehost% ssh administrator@flex_rmc
password: #####

HPE Superdome RMC, Rev. 2.0.44   [Bootloader 1.1.4]
RMC:r001i01c cli>
```

For the password, enter the current RMC password for the `administrator` user, and press **Enter**.

If you configured a site-specific RMC password during the installation and configuration process, use your site-specific password. Hewlett Packard Enterprise recommends that you configure a site-specific password.

If you did not specify a site-specific RMC password, use the factory-default RMC password. The factory-default RMC password appears on the password sticker. The password sticker is a bar coded sticker that appears on your system. The sticker is on the rear of the external RMC, and top-right rear corner of the base chassis on an embedded RMC.

For information about the placement of the password sticker, see [HPE Superdome Flex Server User Guide](#).

Powering on and booting an HPE Superdome Flex Server system from a complete power off

The following procedure explains how to power on an HPE Superdome Flex Server system from a complete power off.



Procedure

1. Visually inspect the system and make sure that the power circuit breakers are on.

There is power to the RMC and to the BMCs as long as the RMC and the BMCs are attached to a power source.

2. Connect to the RMC and log in.

3. Check the state of the system.

To view the partition configuration, enter the following commands.

```
RMC cli> show npar
RMC cli> show complex
```

If no errors are indicated or if the system health status is OK, then proceed to power up the system. If there are errors, run `RMC cli> show logs error` and resolve the errors before powering up the system.

4. To power up the system, assuming it is configured with all chassis in one nPartition numbered 0, enter `power on npar pnum=0`. If you have multiple partitions, each partition can be powered on using `pnum=x` where **x** is the partition number you want to power up.

The power-on takes only a few minutes.

```
RMC cli> power on npar pnum=0
```

5. To access the EFI shell commands, enter the following command:

```
RMC cli> connect npar pnum=0
```

It can take 5 to 10 minutes for the `connect npar pnum=0` command to return boot-to-shell progress information.

6. (Conditional) When the `Shell>` prompt appears, enter `fs0:` to access the boot partition.

```
Shell> fs0:
```

7. (Conditional) Boot the system.

- On Oracle Linux Unbreakable Enterprise Kernel (UEK) systems or on Red Hat Enterprise Linux (RHEL) systems, enter the following command:

```
fs0:\> EFI\redhat\grubx64.efi
```

- On SLES systems, enter the following command:

```
fs0:\> EFI\sles\grubx64.efi
```

More information

[Connecting to the Rack Management Controller \(RMC\)](#)

Powering off an HPE Superdome Flex Server

The following procedure explains how to power off an HPE Superdome Flex Server system.

Procedure

1. Log into the HPE Superdome Flex Server using an account with administrator or operator privileges, and enter the following command to stop the operating system and shut down the system:



```
# shutdown -h
```

2. Enter the following command, and verify that the Run State is Off.

```
RMC cli> show npar  
Partitions: 1
```

Par Num	Run State	Health Status	Chassis OK/In	CPUs OK/In	CPU Cores OK/In	DIMMs OK/In	Memory (GiB) RAM/PMem	IO Cards OK/In
p0	Off	OK	1/1	4/4	1536/1536	48/48	1535/0	0/0

```
* OK/In = OK/Installed
```

3. If the power has not shut down after an extended amount of time, log into the RMC using an account with administrator or operator privileges, and enter the following command to power off a single partition.

```
RMC cli> power off npar pnum=PARTITION_ID [force]
```

Where **PARTITION_ID** is the number of the nPartition to be shut down. The `force` option makes the partition perform an OS immediate (nongraceful) shutdown if the OS is still running.



TIP: If you want to monitor the status of the shutdown process, you can open a second session on the RMC CLI. Using two sessions allows you to monitor the OS shutdown and the RMC status simultaneously.

Configuring the autopower capability

The autopower capability allows the HPE Superdome Flex Server to power up automatically when power is applied after a power outage.

This procedure explains the following:

- How to retrieve information about the current autopower setting.
- How to configure the autopower setting you need.

Procedure

1. Log into the RMC as the administrator user.
2. Enter the `show autopower` command.

The command shows the autopower capability that is enabled on your system.

```
RMC cli> show autopower  
==== r001i01c ====  
auto-power on is disabled
```

3. To enable or disable the autopower capability, enter one of the following commands.

- To enable autopower, enter the following command.

```
RMC cli> enable autopower  
==== r001i01c ====  
auto-power on enabled (120 second delay)
```

- To set the delay for autopower, enter the following command.



```
RMC cli> set autopower delay=SECONDS
```

Where *SECONDS* is the number of seconds before the server powers up when power is applied.

- To disable autopower, enter the following command.

```
RMC cli> disable autopower  
==== r001i01c ====  
auto-power on disabled
```

- To cancel autopower, enter the following command.

```
RMC cli> cancel autopower  
==== r001i01c ====  
auto-power on cancelled
```

Use this command to prevent power from being restored automatically to servers following a power outage.

Configuring Hyper-Threading Technology

The following procedure explains how to enable or disable Hyper-Threading Technology on nPartitions in the HPE Superdome Flex Server. Hyper-Threading Technology can be enabled or disabled when an nPartition is created, or on an existing nPartition.

If not specified, Hyper-Threading Technology is enabled by default.

Procedure

1. Log in to the RMC as an administrator user.

2. Enter one of the following commands.

- To create an nPartition with Hyper-Threading Technology enabled, enter the following command:

```
create npar pnum=PARTITION_ID hthread=on chassis={GEOID [GEOID ...],all,unassigned}
```

Where *GEOID* is the chassis or group of chassis you want to include in the new nPartition.

- To create an nPartition with Hyper-Threading Technology disabled, enter the following command.

```
create npar pnum=PARTITION_ID hthread=off chassis={GEOID [GEOID ...],all,unassigned}
```

Where *GEOID* is the chassis or group of chassis you want to include in the new nPartition.

NOTE: The nPartition must be powered off before Hyper-Threading Technology can be enabled or disabled on an existing nPartition.

By default, nPartitions are created with Hyper-Threading: On.

- To enable Hyper-Threading Technology on an existing nPartition, enter the following command.

```
modify npar pnum=PARTITION_ID hthread=on
```

- To disable Hyper-Threading Technology on an existing nPartition, enter the following command.



```
modify npar pnum=PARTITION_ID hthread=off
```

3. To verify the status of Hyper-Threading Technology, enter the `show npar verbose` command.

Configuring RAS

The following procedure explains how to enable RAS or HPC memory modes on nPartitions in the HPE Superdome Flex Server. Memory modes can be set when creating a nPartition or on an existing nPartition.

⚠ CAUTION: Hewlett Packard Enterprise strongly recommends using the `ras=on` specification to enable memory RAS features (ADDDC mode). HPE RAS features provide higher resiliency to DIMM faults versus standard memory error-correcting code (ECC).

The `ras=hpc` specification disables memory RAS features and therefore could result in compromised system resiliency and a potential server outage.

You will experience high failure rates of DIMMs using the `ras=hpc` specification.

To check the memory RAS setting, enter the `show npar verbose RMC` command.

Procedure

1. Log in to the RMC as an administrator user.

2. Enter one of the following commands.

- To create an nPartition with RAS (ADDDC mode) enabled, enter the following command.

```
create npar pnum=PARTITION_ID ras=on chassis={GEOID [GEOID ...],all,unassigned}
```

Where *GEOID* is the chassis or group of chassis you want to include in the new nPartition.

NOTE: nPartitions are created with `ras=on` by default.

- To create an nPartition with HPC mode enabled, enter the following command.

```
create npar pnum=PARTITION_ID ras=hpc chassis={GEOID [GEOID ...],all,unassigned}
```

Where *GEOID* is the chassis or group of chassis you want to include in the new nPartition.

- **NOTE:** The nPartition must be powered off before RAS (ADDDC mode) or HPC mode can be enabled on an existing nPartition.

To enable RAS (ADDDC mode) on an existing nPartition, enter the following command.

```
modify npar pnum=PARTITION_ID ras=on
```

- To enable HPC mode on an existing nPartition, enter the following command.

```
modify npar pnum=PARTITION_ID ras=hpc
```

3. To verify the status of RAS or HPC mode, enter the `show npar verbose` command.

More information

[Monitoring main memory](#)

Configuring RMC user accounts and user roles

User roles can be defined when creating a user account.

Procedure

1. Enter the `add user` command.

```
add user name=USERNAME role={administrator,monitor,operator}
```

USERNAME

Specify the name of a new user that you want to add to the system. The following are rules regarding the *USERNAME* specification:

- You can add a maximum of 30 users.
- *USERNAME* must begin with a letter and is case-sensitive.
- *USERNAME* must be unique to all other user names and group names.
- *USERNAME* can be 2 to 20 characters long and can include all alphanumeric characters, the dash, and the underscore.
- Certain user names are reserved for the RMC and eRMC. The reserved user names are as follows:

User name	Reserved on eRMC	Reserved on RMC
adminstrator	X	X
backup		X
bin		X
daemon	X	X
dbus		X
ftp		X
haldaemon		X
logger	X	X
mail		X
nobody		X
ntp	X	
operator		X
proxy		X
root	X	X
sshd	X	X
sys		X
sysadmin	X	

Table Continued



User name	Reserved on eRMC	Reserved on RMC
sync		X
www-data		X

When you create a user, the system also prompts you to create a password. The password can be from 8 to 40 characters long and can include all printable characters.

2. Enter a password for the user account. The password can be from 8 to 40 characters long and can include all printable characters.
3. To verify the user list, enter the command `show user list`.
 - To see details of the current user, enter the command `show user`.
 - To delete an existing user, enter the command `remove user name=USERNAME`.

More information

[add](#)

[show](#)

[remove](#)

User roles

Superdome Flex supports role-based user accounts that provide each user account with privileges to improve security. Roles can be configured for RMC or Redfish. Up to 30 user accounts can be created.

RMC role name	Redfish role name	Privileges
administrator	Administrator	All privileges including ability to create, delete, and edit other user accounts
operator	Operator	Power control, setting a profile, and BIOS parameters
monitor	ReadOnly	Change own password, access read-only JViewer, and access read-only console
IPMI User	N/A	Log in to IPMI remotely

Retrieving HPE Superdome Flex Server system configuration information

You can use the `show complex` command to view your system configuration. Use this command, for example, to retrieve the baseboard management controller (BMC) identifiers.

Many RMC commands and `ipmi` commands require you to specify one of these IDs when you want the command to operate on a specific BMC. The command syntax permits you to shorten the IDs by omitting the leading 0 characters. In these commands, if you decide to omit any leading 0 characters, omit all. For example, you can specify the BMC ID `r001i01b` as `r1i1b`, by omitting the 0 characters.

```
RMC cli> show complex
```

```
SSN: 5UF6512345
```

```
RMCs: 1
```



```

r001i11c RMC-UV400

BMCs:          2
r001i01b IP147-BASEIO   P000
r001i06b IP147         P000

Partitions:    1
partition000 BMCs:    2
                r001i01b
                r001i06b

```

The output shows the following:

- This system has one RMC, and its node name is r001i11c.
- There are two chassis BMCs, and their IDs are r001i01b and r001i06b.
- There is one partition, p0, and it contains both chassis.

Retrieving hardware configuration information

The following procedure explains how to retrieve hardware configuration information.

Procedure

1. Log into the RMC as the administrator user.
2. Enter the show hwcfg command.

```

RMC cli> show hwcfg
BASEIO_DISABLE=no..... 1/2 BMC(s)
    r001i01b
BASEIO_DISABLE=yes..... 1/2 BMC(s)
    r001i06b
BIOS_CORE_DISABLE_MASK=0x0
    all targeted BMC(s)
BIOS_FILE=/common/bios.fd
    all targeted BMC(s)
BIOS_HT_DISABLE=yes
    all targeted BMC(s)
BIOS_LABMODE=no
    all targeted BMC(s)
BIOS_RAS_DISABLE=no
    all targeted BMC(s)
CHASSIS_DISABLE=no
    all targeted BMC(s)
DEBUG_SW=0x0
    all targeted BMC(s)
HOST_SPI_LOCATION=HOST_SPI_LOC_PCH
    all targeted BMC(s)
IIO_DISABLE=no
    all targeted BMC(s)
MEMRISER_DISABLE=no
    all targeted BMC(s)
NL_ENABLE=yes
    all targeted BMC(s)
PARTITION=1

```



```
all targeted BMC(s)
SOCKET_DISABLE=no
all targeted BMC(s)
```

Monitoring log files

The HPE Superdome Flex Server logs events that occur on the RMC, on the BMCs, and on the server. The `show logs` command returns the value of various log files.

- `show logs audit`
- `show logs console target={BMC_GEOID | PARTITION_ID} [socket={0-3}] [full]`
- `show logs dmesg [target=target]`
- `show logs error [mca | cmc | type=SEVERITY | id=BUNDLE_ID | pnum=PARTITION_ID]`
- `show logs hel [all] [hex] [target=target]`
- `show logs iel [live]`
Same as `show livelogs`
- `show logs sel [target=target]`
- `show logs syshist [verbosity={0,1,2} | raw]`
- `show logs uvdmp [list | id=HEX_ID | file=FILENAME] [verbose]`
- `show logs support [ieldepth=DAYS]`

The following procedure shows how to retrieve hardware event log file and the RMC start-up log file.

Procedure

1. Log into the RMC as the administrator user.
2. Enter the `show logs hel all` command to retrieve hardware event logs.

```
RMC cli> show logs hel all
UV4_NI1_HCERR_SUMMARY           : 0x0000000000000001
UV4_NI1_HCERR_FIRST            : 0x0000000000000001
UV4_NI1_HCERR_TIME_STAMP       : 0x80003DC074E22444
UV4_NI1_NL0_HCERR_SUMMARY      : 0x000000000000004C
UV4_NI1_NL0_HCERR_FIRST        : 0x0000000000000008
UV4_NI1_NL0_HCERR_OVERFLOW     : 0x0000000000000000
UV4_NI0_AOERR0_SUMMARY         : 0x0000000000000002
UV4_NI0_AOERR0_FIRST           : 0x0000000000000002
UV4_NI0_AOERR0_TIME_STAMP      : 0x80003DC074E224C4
UV4_NI0_NL1_AOERR_SUMMARY      : 0x0000000000000008
UV4_NI0_NL1_AOERR_FIRST        : 0x0000000000000008
UV4_NI0_NL1_AOERR_OVERFLOW     : 0x0000000000000000
UV4_NI1_AOERR0_SUMMARY         : 0x0000000000000001
UV4_NI1_AOERR0_FIRST           : 0x0000000000000001
UV4_NI1_AOERR0_TIME_STAMP      : 0x80003DC074E22504
UV4_NI1_NL0_AOERR_SUMMARY      : 0x0000000000000008
UV4_NI1_NL0_AOERR_FIRST        : 0x0000000000000008
UV4_NI1_NL0_AOERR_OVERFLOW     : 0x0000000000000000
```



```

++++ r001i06b 10/25/17 07:36:56
**** RESET ****
++++ r001i06b 10/25/17 08:04:34
**** RESET ****
++++ r001i06b 10/25/17 08:09:59
**** RESET ****
++++ r001i06b 10/25/17 08:24:25
**** RESET ****
++++ r001i06b 10/25/17 08:39:07
**** RESET ****
++++ r001i06b 10/25/17 09:54:57
**** RESET ****
++++ r001i06b 10/25/17 10:30:33
**** RESET ****
++++ r001i06b 10/25/17 10:50:55
**** RESET ****
++++ r001i06b 10/25/17 11:19:45
**** RESET ****
++++ r001i06b 10/25/17 11:41:36
**** RESET ****
++++ r001i06b 10/25/17 18:33:22
**** RESET ****
++++ r001i06b 10/25/17 18:34:46
**** RESET ****
++++ r001i06b 10/25/17 18:38:26
**** RESET ****
++++ r001i06b 10/25/17 18:40:55
**** RESET ****
++++ r001i06b 10/25/17 18:43:20
**** RESET ****
<type h for help> lines 33209-33256/33256 (END)

```

The `show logs hel` all command output can contain information about correctable memory errors, uncorrectable memory errors, machine checks, and other events.

3. Enter the `show logs dmesg` command to retrieve the RMC start-up log file.

```

RMC cli> show logs dmesg
...
[ 3.024725] VFS: Mounted root (cramfs filesystem) readonly on device 1:0.
[ 3.031567] devtmpfs: mounted
[ 3.034931] Freeing unused kernel memory: 204k freed
...

```

More information

[show](#)

Updating firmware

Hewlett Packard Enterprise recommends that all firmware on all devices in your system be updated to the latest version after hardware installation is complete. Hewlett Packard Enterprise also encourages you to check back often for any updates that might have been posted.

There are three methods for updating the complex firmware; using SUM (Smart Update Manager), using HPE OneView, or manually. Each method uses a different firmware bundle.



Installing the latest complex firmware using SUM

The SUM utility enables you to deploy firmware components from either an easy-to-use interface or a command line. SUM has an integrated hardware discovery engine that discovers the installed hardware and the current versions of firmware in use on target servers. SUM also has logic to install updates in the correct order and ensure that all dependencies are met before deployment of a firmware update. SUM also contains logic to prevent version-based dependencies from destroying an installation. This logic ensures that updates are handled in a manner that reduces any downtime required for the update process. SUM does not require an agent for remote installations. SUM is included in the downloadable firmware bundles that include *for SUM installation* in the title.

Detailed installation instructions are provided in the SUM bundle at www.hpe.com/support/superdome-flex-software.

For more information about SUM, see the [HPE Smart Update Manager User Guide](#).

NOTE: You can also update firmware manually. For more information about manually updating the firmware, see the detailed instructions provided in the release notes for the firmware bundle.

More information

[Updating firmware manually](#)

Installing the latest complex firmware using HPE OneView

The HPE OneView utility enables you to deploy firmware updates through the HPE OneView interface.

NOTE: Updating firmware with HPE OneView requires HPE OneView version 5.0 or later and Superdome Flex Server firmware version 3.0.512 or later.

Detailed installation instructions are provided in the HPE OneView bundle at www.hpe.com/support/superdome-flex-software.

For more information about HPE OneView, see *HPE OneView User Guide* at <http://www.hpe.com/info/oneview/docs>.

NOTE: You can also update firmware manually. For more information about manually updating the firmware, see the detailed instructions provided in the release notes for the firmware bundle.

More information

[Updating firmware manually](#)

Updating firmware manually

The following procedures explain how to update all the firmware on an HPE Superdome Flex Server. These procedures upgrade the firmware on the RMC and on the BMCs.

You can perform firmware updates online or offline. Updating firmware online requires system firmware version 3.0.512 or later.

NOTE: The online firmware update is not supported when updating from 3.25.46 (or earlier) to 3.30.140 (or later) due to a PLD version change.



Updating firmware offline

Procedure

1. Obtain a copy of the HPE Superdome Flex Server Firmware Bundle from www.hpe.com/support/superdome-flex-software.
2. Download the new firmware bundle to a location on a computer which is accessible from the RMC over the local network.
3. Log into the RMC as the `administrator` user and provide the password when prompted.

For example:

```
remotehost% ssh administrator@flex-rmc
```

4. To verify that the RMC is configured to use a DNS, enter the following command:

```
RMC cli> show dns
```

If necessary, use the `add dns` command to configure DNS access.

5. Power off the system.

```
RMC cli> power off npar pnum=x
```

6. To update the firmware, enter the following command:

```
RMC cli> update firmware url=path_to_firmware
```

For *path_to_firmware*, specify the location on the local computer that specifies the location of the firmware bundle that you downloaded. Acceptable paths to the location are `https`, `sftp`, and `scp`.

For example:

```
RMC cli> update firmware url=https://myhost.com/bundle
```

The firmware supports specifying the protocol and port in *path_to_firmware*.

For example:

```
update firmware url=https://myhost.com/bundle
update firmware url=https://myhost.com:1234/bundle
```

7. Wait for the RMC to reboot.
8. To verify the new firmware level, enter the following command:

```
RMC cli> show firmware verbose
```

9. To activate new system firmware, enter the following command:

```
RMC cli> power on npar pnum=x
```

Where *x*=partition number. *x*=0 will power up the entire system.

Updating firmware online

Prerequisites

Updating firmware while online requires system firmware 3.0.512 or later.



Procedure

1. Obtain a copy of the HPE Superdome Flex Server Firmware Bundle from www.hpe.com/support/superdome-flex-software.

2. Download the new firmware bundle to a location on a computer which is accessible from the RMC over the local network.

3. Log into the RMC as the `administrator` user and provide the password when prompted.

For example:

```
remotehost% ssh administrator@flex-rmc
```

4. To verify that the RMC is configured to use a DNS, enter the following command:

```
RMC cli> show dns
```

If necessary, use the `add dns` command to configure DNS access.

5. To update the firmware, enter the following command:

```
RMC cli> update firmware url=path_to_firmware
```

For *path_to_firmware*, specify the location on the local computer that specifies the location of the firmware bundle that you downloaded. Acceptable paths to the location are `https`, `sftp`, and `scp`.

For example:

```
RMC cli> update firmware url=https://myhost.com/bundle
```

The firmware supports specifying the protocol and port in *path_to_firmware*.

For example:

```
update firmware url=https://myhost.com/bundle
update firmware url=https://myhost.com:1234/bundle
```

6. Wait for the RMC to reboot.
7. To verify the new firmware level, enter the following command:

```
RMC cli> show firmware verbose
```

- a. Confirm that the configured complex bundle version matches the new firmware version.
- b. Confirm that the configured nPartition bundle version matches the new firmware version.
- c. Confirm that the firmware on all devices matches the new firmware version.

NOTE: Configured nPartitions show the previous firmware version until they are rebooted.

8. To activate new system firmware on an nPartition, enter the following command:

```
RMC cli> reboot npar pnum=x
```

Where `x`=partition number. `x=0` will reboot the entire system.

NOTE: The new system firmware will not be activated on an nPartition until the `reboot npar` command is entered.



Downgrading firmware

CAUTION: Hewlett Packard Enterprise does not recommend downgrading to a prior firmware release. Performing this task may result in loss of data and/or system functionality.

Firmware releases include fixes for known vulnerabilities and/or updates designed to avoid future vulnerabilities. Downgrading firmware may expose you to security vulnerabilities.

Procedure

1. Apply the prior firmware version using online or offline firmware upgrade processes.

NOTE: If you are using SUM, you cannot use the `exclude npar` option for downgrading to a previous firmware version.

2. After applying the prior firmware version, you must restart the system or partition.

- If you are using the online firmware process, reboot the system or partition. The prior firmware version is not applied until after the reboot.

```
RMC cli>reboot npar pnum=Partition #
```

- If you are using the offline firmware process, power on the system or partition.

```
RMC cli>power on
```

More information

[Updating firmware online](#)

[Updating firmware offline](#)

Creating a crash dump file

You can request that the operating system write a crash dump file. The file name resides in a directory that includes a time stamp, and the file location depends on your operating system, as follows:

- For the Oracle Linux UEK 4 update 5 operating system and the RHEL 7.X operating system, the directory is as follows:

```
/var/crash/ipaddr_or_hostname-year.month.day-hour:minutes:seconds
```

- For the SLES 12 SPX operating systems, the directory is as follows:

```
/var/crash/year-month-day-hour:minutes
```

Procedure

1. Log into the HPE Superdome Flex Server as the root user, and enter the following commands to enable the kernel crash dump service, `kdump`.

```
# systemctl enable kdump  
# systemctl start kdump
```

Make sure to log into the server, not the RMC.



The HPE Foundation Software package installation process enables `kdump` by default. The preceding commands ensure that `kdump` is enabled. You can enable `kdump` on your server at any time. `kdump` must be enabled to create a crash dump file.

2. To ensure that the `uv_nmi` default action is `kdump`, enter the following command.

```
# echo kdump > /sys/module/uv_nmi/parameters/action
```

As an alternative to this command, you could also boot the kernel with the following parameter.

```
uv_nmi.action=kdump
```

3. Use one of the following methods to send a nonmaskable interrupt (NMI) signal to start the dump.

Method 1: When logged in directly to the RMC.

- a. Use the `ssh` command to connect securely.

```
# ssh administrator@flex_rmc
```

- b. To send the NMI signal, enter one of the following commands:

- To send the NMI signal to all nPartitions, enter the command:

```
RMC cli> power nmi -o
```

- To send the NMI signal to an individual nPartition, enter the command:

```
RMC cli> power nmi PARTITION_ID
```

Where *PARTITION_ID* indicates the partition where the NMI signal will be sent.

Method 2: From a remote connection.

NOTE: The `ipmitool` command can only be used on systems with a single nPartition.

```
# ipmitool -I lanplus -H HOST -U USER -P PASSWD chassis power diag
```

The `ipmitool` command uses the following specifiers:

HOST

The RMC host name.

USER

The RMC administrator account user name.

PASSWD

The RMC administrator account password.

NOTE: IPMI is disabled by default and must be enabled in the RMC by using the `enable ipmi` command.



Managing system performance with FW

This section provides the list of common performance tunables.

Table 1: nPar attributes / Power and Performance options

npar attributes / Power and performance options	SDFlex
Power Regulator	Native HWP
Minimum Processor Idle Power Core C-State	No C-States
Minimum Processor Idle Power Package C-State	No Package State
Energy/Performance Bias	OS Control
Intel(R) UPI Link Power Management	Disabled
Intel(R) Turbo Boost Technology	Enabled
Energy Efficient Turbo	Enabled
Uncore Frequency Scaling	Auto
Sub-NUMA Clustering	Disabled
NUMA Group Size Optimization	N/A no _MSCT
Advanced Memory Protection	ADDDC ¹
Intel(R) NIC DMA Channels (IOAT)	Enabled
SR-IOV	Enabled
Intel(R) Virtualization Technology (Intel VT)	Enabled
Intel(R) VT-d	Enabled
Processor x2APIC Support	Auto
Intel(R) Hyper-Threading	Enabled
HW Prefetcher	Enabled
Adjacent Sector Prefetch	Enabled
DCU Stream Prefetcher	Enabled
DCU IP Prefetcher	Enabled
LLC Prefetch	Disabled
LLC Dead Line Allocation	Enabled
Local/Remote Threshold	Manual
<ul style="list-style-type: none"> • UPI RRQ Threshold • UPI IRQ Threshold 	<ul style="list-style-type: none"> • 7¹ • 4¹
Snoop Throttle Configuration	N/A
Intel(R) DMI Link Frequency	Auto

Table Continued



npar attributes / Power and performance options	SDFlex
Maximum Memory Bus Frequency	Auto
Memory Refresh Rate	x1¹
Memory Patrol Scrubbing	Enabled
PMem Snoopy Mode for AppDirect	Disabled¹
PMem Performance Mode	BW Optimized (0)¹
XNC Directory Format	1/64 (0)¹
Hub Write Flush Optimization	Enabled¹

¹ This tunable can be modified.

SDFlex tunables

The following examples show how to modify the tunables:

```

SDFlex CLI> modify npar pnum=NPAP_ID [hthread={on,off}]
SDFlex CLI> modify npar pnum=NPAP_ID [memrefreshrate={x1,x2}]
SDFlex CLI> modify npar pnum=NPAP_ID [ras={on,hpc}]
SDFlex UEFI Shell> setvar UpiRrqThresholdNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x7
SDFlex UEFI Shell> setvar UpiIrqThresholdNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x5
SDFlex UEFI Shell> setvar SnoopyModeForADNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x0
SDFlex UEFI Shell> setvar PMMPerformanceNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x0
SDFlex UEFI Shell> setvar DirectoryFormatNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x1
SDFlex UEFI Shell> setvar Uv4HubWriteFlushModeNext -guid FDD70221-655E-48e8-B77D-EA392EC5F60C -bs -rt -nv =0x1

```

Modifications to FW tunables take effect on the next boot. Use the `sysconfig` command at the UEFI Shell to validate changes after reset.



Managing system performance with HPE Foundation Software (HFS)

HPE Foundation Software (HFS) includes automatic boot-time optimization utilities, reliability features, and technical support tools. Designed for high performance computing, these tools help maximize HPE Superdome Flex Server system performance and availability.

NOTE: HFS is required on the server running Linux.

To install HFS, see [HPE Superdome Flex Server OS Installation Guide](#).

HPE Foundation Software can be downloaded at www.hpe.com/support/superdome-flex-software.

Monitoring main memory

The MEMlog utility monitors the overall system health of each dual inline memory module (DIMM) on your system. The MEMlog utility is configured for your system when the HPE Foundation Software (HFS) package is installed. After the HFS installation, memlog starts automatically and uses BIOS settings to determine the operational memory mode. To determine the operational memory mode, enter the following command:

```
RMC cli> show npar verbose
```

```
This system is nPartition capable
```

```
Partitions: 3
```

```
Partition 1:
```

```
Run State           : OS Boot
Health Status       : OK
Chassis OK/In       : 2/2
CPUs OK/In          : 8/8
Cores OK/In         : 224/224
DIMMs OK/In         : 64/64
IO Cards OK/In      : 3/3
Hyper-Threading     : Off
RAS                 : On
Boot Chassis        : r001i01b
Boot Slots           : 3,5
Secure Boot         : Off
Secure Boot Next    : Off
Memory Refresh Rate : Refreshx1
Memory Refresh Rate Next : Refreshx1
PCH SATA Interface Mode : RAID
PCH SATA Interface Mode Next : RAID
Volatile Memory     : 1535 GiB
Persistent Memory    : 2016 GiB
```

```
Partition 11:
```

```
Run State           : OS Boot
Health Status       : OK
Chassis OK/In       : 1/1
CPUs OK/In          : 4/4
Cores OK/In         : 112/112
DIMMs OK/In         : 32/32
IO Cards OK/In      : 0/0
Hyper-Threading     : On
RAS                 : On
```



```

Boot Chassis          : r001i11b
Boot Slots           : 3,5
Secure Boot          : Off
Secure Boot Next     : Off
Memory Refresh Rate  : Refreshx1
Memory Refresh Rate Next : Refreshx1
PCH SATA Interface Mode : RAID
PCH SATA Interface Mode Next : RAID
Volatile Memory      : 767 GiB
Persistent Memory    : 1008 GiB

Partition 16:
Run State            : OS Boot
Health Status       : OK
Chassis OK/In       : 1/1
CPUs OK/In          : 4/4
Cores OK/In         : 112/112
DIMMs OK/In         : 32/32
IO Cards OK/In      : 0/0
Hyper-Threading     : On
RAS                  : On
Boot Chassis        : r001i16b
Boot Slots          : 3,5
Secure Boot         : Off
Secure Boot Next    : Off
Memory Refresh Rate : Refreshx1
Memory Refresh Rate Next : Refreshx1
PCH SATA Interface Mode : RAID
PCH SATA Interface Mode Next : RAID
Volatile Memory     : 767 GiB
Persistent Memory   : 1008 GiB

* OK/In = OK/Installed

```

The preceding example shows the RAS entry set to `on`. On this system, reliability, availability, and serviceability (RAS) mode is enabled, which is the default. If the output shows the RAS entry set to `Hpc`, HPC mode is enabled and RAS is disabled.

The `ras=on` specification is the default. This setting enables ADDDC mode. This mode allows better error recovery but incurs a small performance penalty. When set, all memory error handling occurs in the BIOS, rather than `memlog`.

The `ras=hpc` specification enables maximum memory performance but incurs a small reliability penalty.

⚠ CAUTION: Hewlett Packard Enterprise strongly recommends using the `ras=on` specification to enable memory RAS features (ADDDC mode). HPE RAS features provide higher resiliency to DIMM faults versus standard memory error-correcting code (ECC).

The `ras=hpc` specification disables memory RAS features and therefore could result in compromised system resiliency and a potential server outage.

You will experience high failure rates of DIMMs using the `ras=hpc` specification.

To check the memory RAS setting, enter the `show npar verbose RMC` command.

The memory modes affect `memlog` processing as follows:

- When RAS mode is enabled, `memlog` is disabled and exits.
- When HPC mode is enabled, `memlog` is enabled and monitors system health.

To determine whether `memlog` is running, enter the following command on the server:

```
remotehost% systemctl status memlog
```



For more information about the MEMlog utility, see the *memlog(8)* manpage.

For more information about the `modify` command, use the `help modify` RMC command.

Retrieving main memory health information

Hewlett Packard Enterprise recommends that you check your system periodically to determine whether the MEMlog utility has reported any hardware errors. The MEMlog utility verifies and diagnoses problems with the DIMMs. The utility messages appear in `/var/log/messages`.

The following are ways to use commands to retrieve information about memory problems and memory health:

- Use the `memlogd` command with the `-s` option to retrieve information about DIMMs tagged for repairs. If there are DIMMs to replace, the command returns a **CRITICAL** message along with detailed information for each failing DIMM. For example, the following message indicates a DIMM to be replaced.

```
uv:~ # memlogd -s
CRITICAL - one or more DIMM failling.
r001i11b04 P1-DIMM3A Size 16384MB Width 4 Rank 2 Row 16 Col 10 Bank 16
Serial 39d959c4 Part M393A2G40DB0-CPB signed 2133
  Mon Dec 14 14:40:02 2015 Rank 1 Dram U14B Bank 3 Row 0x5fe5 Col 0x1e8
  single CB DQ47 Temp = 33C hits 26
```

If there are no DIMMs to replace, the output is as follows.

```
Ok - all DIMMs within specification
```

- Scan the system log for entries that contain the string `MEMLOG`. If problems arise with any of the DIMMs on your system, the MEMlog utility writes a message to `/var/log/messages`. To retrieve these messages, enter the following command.

```
# grep MEMLOG /var/log/messages
```

For example, the following messages indicate some DIMMs to replace.

```
# grep MEMLOG /var/log/messages
rli0n0:Dec 9 07:29:45 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM1A Rank 0 DRAM U9 DQ4 Temp = 21C
rli0n0:Dec 9 07:30:00 rli0n0 MEMLOG[4595]: P1-DIMM1A has a failed DRAM and must be replaced soon.
  Exposure to Uncorrected Error is high
rli0n0:Dec 9 07:30:00 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM1A Rank 0 Bank 0 Row 0x0 Col 0x8 Temp = 21C
rli0n0:Dec 9 07:30:00 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM1A Rank 0 DRAM U9 DQ4 Temp = 21C
rli0n0:Dec 9 07:30:12 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM3A Rank 0 Temp = 22C
rli0n0:Dec 9 07:30:12 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM3A Rank 0 DRAM U9 DQ4 Temp = 22C
rli0n0:Dec 9 07:30:25 rli0n0 MEMLOG[4595]: P1-DIMM3A has a failed DRAM and must be replaced soon.
  Exposure to Uncorrected Error is high
rli0n0:Dec 9 07:30:25 rli0n0 MEMLOG[4595]: Read ECC P1-DIMM3A Rank 0 Bank 0 Row 0x0 Col 0x8 Temp = 22C
```

NOTE: Some lines in the preceding output have been wrapped for inclusion in this documentation.

About page migration

If the main memory encounters correctable memory errors, system software analyzes the failure and determines the severity of the problem. The following events occur:

- If MEMlog determines that the correctable error is a repeatable failure, MEMlog completes the following actions:



- MEMlog asks the kernel to migrate the 4k page that contains the corrected error address. If the page is mapped to user space and is not locked, the kernel copies the data to a new page. In addition, MEMlog remaps the old page to the new page.
- The kernel retires the old page. It is no longer available for use.
- If the system software detects that the program data at that memory location is uncorrectable, the following events occur:
 - The kernel stops the application that was affected by the uncorrectable memory error.
 - The application exits.

All other applications running continue to work. The faulty 4K page is retired and is not made available for use.

Example 1. The messages in this example indicate the following kernel actions:

- The kernel received a request to migrate a page.
- The kernel migrated the page.
- The kernel received additional requests to migrate the page.

```
# grep "soft offline" /var/log/messages
[ +0.302854] soft offline: 0x171596ad8: page leaked
[ +0.417714] soft offline: 0x171596ad8: page already poisoned
```

Typically, you see the preceding messages when there is a delay in completing the migration. When a page is reported a second time, the system software checks the poisoned flag, detects the poisoned flag, and generates the `page already poisoned` message in `/var/log/messages`. When the kernel marks a page as poisoned, this action ensures that the page is not reused.

Example 2. The kernel writes messages such as the following to its log file when a system encounters an uncorrectable memory error.

```
[ 139.475068] {1}[Hardware Error]: Hardware error from APEI Generic Hardware Error Source: 2
[ 139.475069] {1}[Hardware Error]: It has been corrected by h/w and requires no further action
[ 139.475070] {1}[Hardware Error]: event severity: corrected
[ 139.475070] {1}[Hardware Error]: Error 0, type: corrected
[ 139.475071] {1}[Hardware Error]: fru_text:
[ 139.475071] {1}[Hardware Error]: section_type: memory error
[ 139.475072] {1}[Hardware Error]: error_status: 0x0000000000040400
[ 139.475072] {1}[Hardware Error]: physical_address: 0x0000000fba0e9000
[ 139.475073] {1}[Hardware Error]: node: 0
[ 139.541602] Disabling lock debugging due to kernel taint
[ 139.547595] mce: [Hardware Error]: Machine check events logged
[ 139.554109] mce: Uncorrected hardware memory error in user-access at fba0e9000
[ 139.562114] MCE 0xfba0e9: Killing einj_tool:4730 due to hardware memory corruption
[ 139.570782] MCE 0xfba0e9: dirty LRU page recovery: Recovered
```

Monitoring system performance

You can use Linux utilities, HPE Foundation Software (HFS) utilities, and open-source utilities to monitor system performance.

The Linux utilities include `w`, `ps`, `top`, `vmstat`, `iostat`, and `sar`.

The HFS utilities are `gr_systat`, `hubstats`, `nodeinfo`, and `topology`.



Configuring CPU frequency scaling

CPU frequency scaling allows the operating system to scale the processor frequency automatically and dynamically. Hewlett Packard Enterprise configures the CPU frequency scaling setting on all HPE Superdome Flex Server systems before they leave the factory. The default setting is assumed to be correct for most implementations. The CPU frequency scaling setting lets your system take advantage of the Intel Turbo Boost technology that is built into each processor.

The Intel Turbo Boost Technology allows processor cores to run faster than the base operating frequency as long as they are operating within the limits set for power, current, and temperature. The CPU frequency scaling setting also affects power consumption and enables you to manage power consumption. For example, theoretically, you can cut power consumption if you clock the processors from 2 GHz down to 1 GHz.

Changing the CPU governor setting and frequency setting

The default CPU frequency governor setting can inhibit system performance. The `hpe-auto-config` utility automatically sets the CPU frequency setting to `performance` mode and sets CPUs to maximum frequency. This procedure explains how to override the automatic CPU frequency setting.

Procedure

1. Enter the following command to view the available CPU governor settings, and study the output to determine which governor setting is appropriate for your site.

```
# cpupower frequency-info -g
```

Hewlett Packard Enterprise recommends that you verify that the CPU governor setting is `performance` and if some other setting is shown, change it to `performance`.

2. Enter the following command to display the available CPU frequencies.

```
# cpupower frequency-info
```

Inspect the `frequency_steps` field, and choose a minimum and/or maximum frequency.

3. Enter one or more of the following commands to change the governor and/or frequency settings.

- Enter the following command to change the governor setting.

```
# cpupower frequency-set -g GOVERNOR
```

For `GOVERNOR`, specify the setting you want.

- Use the `cpupower frequency-set` command to change one of the following values:
 - Both the minimum frequency and the maximum frequency
 - The maximum frequency
 - The minimum frequency

```
# cpupower frequency-set -u MAX -d min
# cpupower frequency-set -u MAX
# cpupower frequency-set -d MIN
```

For `MAX` and `MIN`, specify a value in the following format: `VALUE [UNIT]`

The default `UNIT` is KHz. You can also specify a `UNIT` of Hz, MHz, GHz, or THz.

4. Enter the following command and verify that the `GOVERNOR` setting you specified appears in the `cpupower` command output in the `current_policy` field.



```
# cpupower frequency-info
```

5. Create a configuration file that includes the settings you configured in this procedure.

Your goal is to create a file that includes the command you ran in this procedure. Make sure that the file has execute permission.

```
# echo "cpupower frequency-set -g performance -u 3000MHz -d 2000MHz" > \  
/etc/hpe-auto-config/90_cpu_frequency.sh  
# chmod 744 /etc/hpe-auto-config/90_cpu_frequency.sh
```

When the system boots, the settings in this file override the default `hpe-auto-config` settings to ensure that the settings you configured in this procedure are included after the boot.

More information

[cpupower](#)

Configuring turbo mode

Procedure

1. Make sure that you configured a governor setting.
2. Use the `cat` command to retrieve the list of available frequencies.

```
# cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_available_frequencies  
3301000 3300000 3200000 3100000 3000000 2900000 2800000 2700000 2600000  
2500000 2400000 2300000 2200000 2100000 2000000 1900000 1800000 1700000  
1600000 1500000 1400000 1300000 1200000
```

The preceding output shows the available frequencies. The output lists frequencies in order from the highest, 3,301,000 KHz, to the lowest, 1,200,000 KHz.

The second frequency listed is always the processor nominal frequency. This processor is a 3.3 GHz processor, so 3,300,000 KHz is the nominal frequency.

You can also obtain the nominal frequency by entering the following command and examining the information in the model name field.

```
# cat /proc/cpuinfo
```

3. Use the `cpupower` command to set the frequency to the nominal frequency of 3.3 GHz plus 1 MHz.

That is, specify a frequency of 3,301 MHz.

```
# cpupower frequency-set -u 3301MHz
```

Later, if you want to disable turbo mode, enter the following command to set the maximum frequency back to the nominal frequency.

```
# cpupower frequency-set -u 3300MHz
```

More information

[Changing the CPU governor setting and frequency setting](#)

Consistently naming network cards

You can use HPE Foundation Software to consistently name NICs (network cards) in the system based on their physical chassis location.



NOTE: This feature requires HPE Foundation Software version 1.3 or later.

Hewlett Packard Enterprise recommends enabling this feature to ensure consistent NIC naming if a NIC fails, is moved or replaced, or when firmware is updated.

The naming scheme uses the *GEOID* of the chassis to identify the NIC.

For example, the NIC name *r001i06s02p0* indicates that the NIC is located in rack 001 (*r001*), in the chassis at rack unit 06 (*i06*), is the NIC in slot 02 (*s02*), and is port 0 on that NIC (*p0*).

Procedure

1. Access the system using the RMC.

NOTE: Accessing the system using the RMC is recommended as renaming NICs may result in invalidating existing network configuration.

2. Edit the `/etc/sysconfig/hpe-auto-config` file and change the `UV_GEOID_NIC_NAMES` entry to `yes`.
3. Confirm the changes by entering the `hpe-auto-config` command.
4. Restart the system.
5. Update your network settings as appropriate.

NOTE: Geold NIC names are implemented using `/etc/systemd/network/10-r[0-9][0-9][0-9]*.link` files. If you add or remove NICs in the chassis, remove those files and re-enter the `hpe-auto-config` command.

Default Linux network card naming

The default Linux naming for network cards in Superdome Flex Server differs between firmware versions:

- On a Superdome Flex Server with Intel Xeon 61xx and 81xx (Skylake) processors, the default NIC naming is based on the path naming (udev property `ID_NET_NAME_PATH`).
- On a Superdome Flex Server with Intel Xeon 62xx and 82xx (Cascade Lake) processors, the default NIC naming is based on the slot naming (udev property `ID_NET_NAME_SLOT`).

For example:

- Systems with Intel Xeon 61xx and 81xx (Skylake) processors:

```
ID_NET_NAME_PATH=enP1p129s0f0
```

```
INTERFACE=enP1p129s0f0
```

```
SYSTEMD_ALIAS=/sys/subsystem/net/devices/enP1p129s0f0
```

- Systems with Intel Xeon 62xx and 82xx (Cascade Lake) processors:

```
ID_NET_NAME_PATH=enP1p129s0f0
```

```
ID_NET_NAME_SLOT=enP1s2088f0
```

```
INTERFACE=enP1s2088f0
```

```
SYSTEMD_ALIAS=/sys/subsystem/net/devices/enP1s2088f0
```



NOTE: The preceding default NIC naming applies to systems that use auto-generated NIC names like *enP10p129s0f0* (based on udev property *ID_NET_NAME_PATH*).

To ensure consistent network card naming, Hewlett Packard Enterprise recommends using HPE Foundation Software version 1.3 or later to consistently name network cards based on their physical location.

More information

[Consistently naming network cards](#)

Activating an extended tuning script

You can use HPE Foundation Software to select and activate an extended tuning script for your system. These scripts optimize performance for applications such as SAP HANA OS.

This document describes a general procedure for using tuning scripts. For more detailed information, see *Configuration Guide for HPE Superdome Flex Solutions for SAP HANA with 3PAR All Flash Storage*.

Procedure

1. Verify boot parameters needed for your system configuration, and make changes as necessary.
2. Create and set your parameters in the configuration script file.
 - For SLES, use the configuration file *HPE-Recommended_OS_settings.conf* located in the */etc/saptune/extra* directory. If the file does not exist, create the file.
 - For RHEL, use the configuration file *tuned.conf* located in the */etc/tuned/sap-hpe-hana* directory. If the file does not exist, create the file.
3. Verify that tuning settings are applied by using the `sysctl` command.

Additional HPE Foundation Software (HFS) utilities

There are additional HFS commands and utilities available that typically require no user involvement. Hewlett Packard Enterprise technical support staff members might guide you in the use of these commands when troubleshooting or tuning.

- [hpe-auto-config](#)
- [hpe_irqbalance](#)



HPE Foundation Software (HFS) commands

HPE Foundation Software (HFS) includes automatic boot-time optimization utilities, reliability features, and technical support tools. These tools are initiated and managed through CLI commands to the HFS interface.

Selected details for commonly used HFS commands are included in this guide.

More detailed descriptive and usage information is available in the manpage content for all HFS commands.

cpupower

Syntax

```
cpupower frequency-info [-g]
```

```
cpupower frequency-set [-u MAX] [-d MIN] [-g GOVERNOR]
```

Description

Enables viewing or changing the CPU operating frequency and power governor settings.

Options

frequency-info

Displays the available CPU operating frequencies

frequency-info -g

Displays the available CPU power governor settings

frequency-set

Changes the CPU operating frequency

frequency-set -g

Changes the CPU power governor setting

Specifiers

-u MAX

Specify the maximum CPU frequency in the following format: `value [unit]`. The default unit is KHz. To specify a frequency in a unit other than KHz, add **Hz**, **MHz**, **GHz**, or **THz** after the value.

-d MIN

Specify the minimum CPU frequency in the following format: `value [unit]`. The default unit is KHz. To specify a frequency in a unit other than KHz, add **Hz**, **MHz**, **GHz**, or **THz** after the value.

-g GOVERNOR

Specify the power governor setting. Possible power governor settings are:

- `ondemand` (default) dynamically switches between the available CPUs if at 95% of CPU load. HPE does not recommend this setting.
- `performance`. This governor setting runs the CPUs at the maximum frequency.
- `conservative`. This governor setting dynamically switches between the available CPUs if at 75% of CPU load.



- `powersave`. This governor setting runs the CPUs at the minimum frequency.
- `userspace`. This governor setting runs the CPUs at user-specified frequencies.

Usage

`cpupower frequency-info`

Displays the available CPU frequencies.

`cpupower frequency-info -g`

Displays the available CPU governor power settings.

`cpupower frequency-set -g GOVERNOR`

Changes the power governor setting, where *GOVERNOR* is the required setting.

`cpupower frequency-set -u MAX`

Changes the maximum CPU frequency.

`cpupower frequency-set -d MIN`

Changes the minimum CPU frequency.

`cpupower frequency-set -u MAX -d MIN`

Changes both the maximum and minimum CPU frequencies.

gr_systat

Syntax

`gr_systat`

Description

Generates information that includes the following:

- CPU and memory activity information
- Input/output traffic statistics for each I/O device
- Other system statistics

Displays the following for the system as a whole and for each node:

- The CPU utilization percentage
- I/O wait times
- Interrupt requests (IRQs)
- Memory utilization

The `gr_systat` command is a useful performance monitoring tool that can help you to identify malfunctioning hardware.



Usage

The `gr_systat` command includes several command parameters that let you control the appearance of the `gr_systat` command output. You can save these parameters to a configuration file for future use. For more information, see the `gr_systat` manpage.

Example1

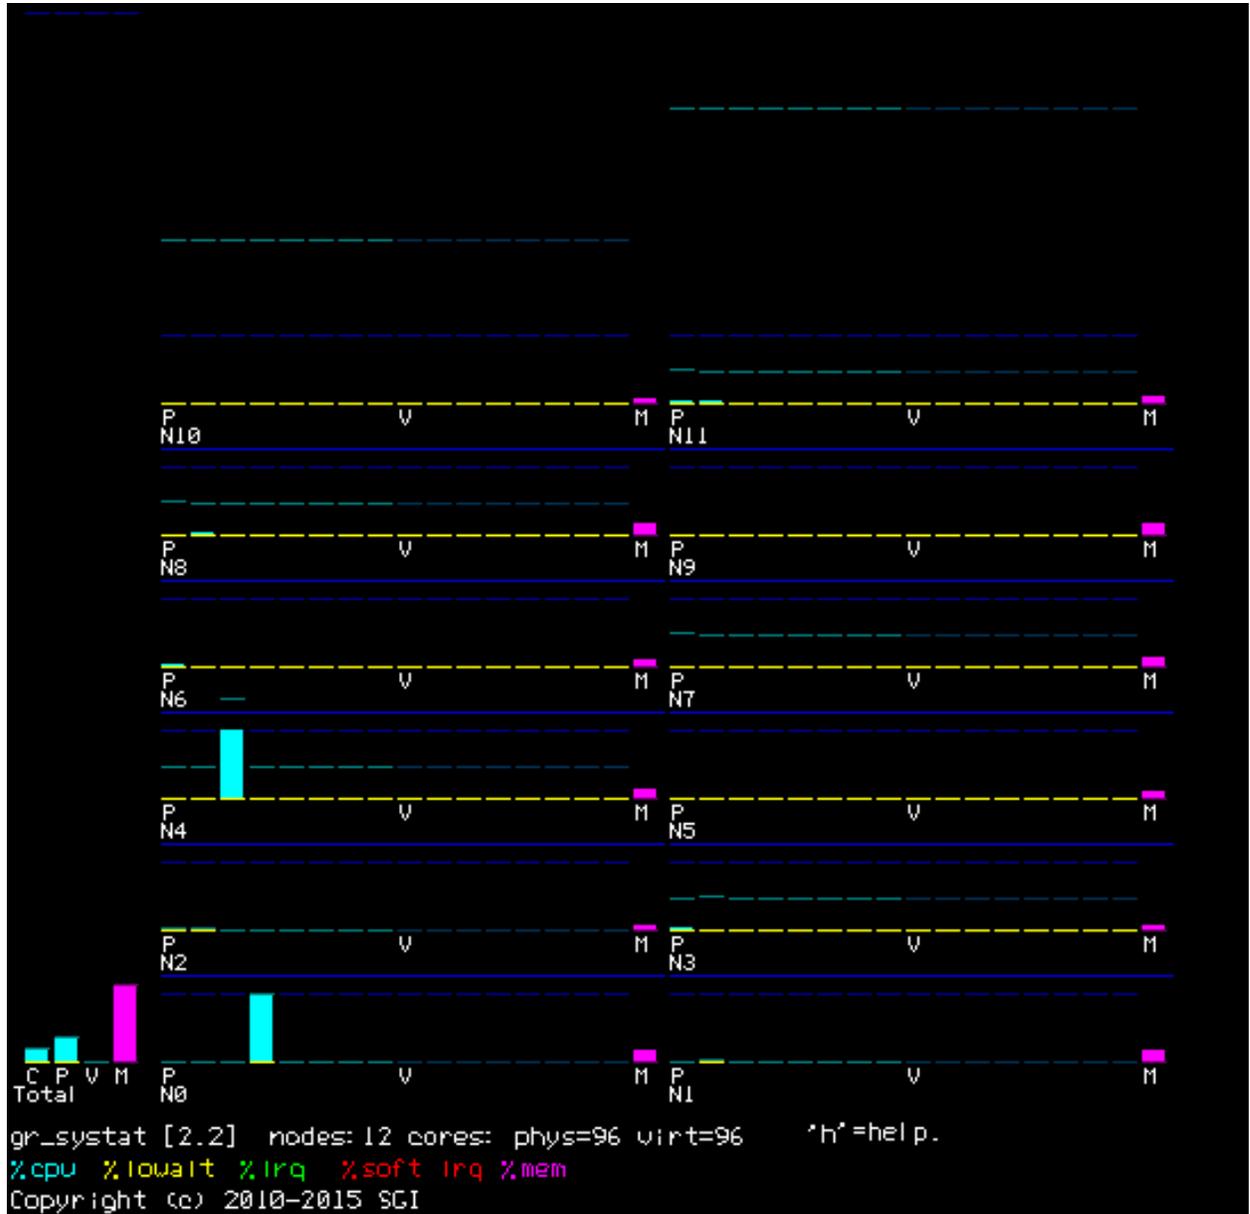


Figure 1: CPU and memory utilization in real time



/etc/hpe-auto-config/*

Executable files in this directory are executed by `hpe-auto-config` in lexical order. A number of scripts are included and custom scripts may be added. A configuration action may be disabled by removing execute permissions from the script that contains it.

/etc/hpe-auto-config/boot_entry_blacklist

List of regular expressions defining which boot entries `hpe-auto-config` are not edited. By default any "rescue" boot entry will not be edited.

/etc/modprobe.d/hpe-auto-config.conf

Destination for kernel module blacklist directives.

/etc/sysconfig/hpe-auto-config

Control behavior of various HPE configuration actions.

/usr/sbin/hpe-auto-config

Main command for running `hpe-auto-config`.

/var/log/hpe-auto-config.log

Each script and shared library function logs its results to this file. Consult this log for a (partial) list of actions taken and for warnings or errors.

`hpe-auto-config` can also use the following options.

-n (--dry-run)

Performs all processing but does not edit `bootloader_config`.

-d (--verbose)

Adds detailed output to `/var/log/hpe-auto-config.log`.

Usage

The `hpe-auto-config` package applies a number of required and suggested edits to the Linux command line on top of default and user-defined parameters. Consult documentation for the installed bootloader program for instructions on how to add or remove Linux boot parameters.

For best results, follow this procedure for changing Linux boot parameters:

1. Follow the documentation for your distribution for adding Linux boot parameters. For grub2 systems, it is recommended to add the boot option to `/etc/default/grub`.
2. To remake the grub boot configuration, restart `hpe-auto-config`.

```
# systemctl restart hpe-auto-config
```

3. Check for conflicts and verify the boot configuration file, `/boot/grub2/grub.cfg` (SLES), or `/boot/efi/EFI/redhat/grub.cfg` (RHEL).

Read the kernel command line in the `grub.cfg` and verify that it contains the option you specified. If any of your boot options conflict with HPE required options, you will see a `syslog` warning message calling out the offending parameters.

```
INFO: boot parameter X in /etc/default/grub conflicts with rule Y
```

To resolve this conflict, remove `X` from `/etc/default/grub` and create a custom `hpe-auto-config` script to override the rule `Y`. See the examples in `hpe-auto-config(1)` for instructions and more information.

For more information on this command, see the `HPE-AUTO-CONFIG(1)` manpage content.



Example 1

Overriding a required boot parameter

In this example, you will add `numa_balancing=enable` to the Linux command line. Add `numa_balancing=enable` to `/etc/default/grub` and restart the `hpe-auto-config`.

```
# echo GRUB_CMDLINE_LINUX+=" numa_balancing=enable" >> /etc/default/grub
# systemctl restart hpe-auto-config
...
hpe-auto-config: 2017/10/05 21:23:42 20_sap_hana.sh: INFO: boot parameter
numa_balancing in /etc/default/grub conflicts with rule replace numa_balancing=disable
```

To override, create an executable file `/etc/hpe-auto-config/21_numa_balancing` with these contents:

```
source /usr/lib64/hpe-auto-config/shlib
boot_option replace numa_balancing=enable
```

Finally, restart `hpe-auto-config` and verify the command line in the `grub.cfg` file. The conflict warning will continue until `numa_balancing=enable` is removed from `/etc/default/grub`.

Enter `nano /etc/default/grub` to update `/etc/default/grub` before restarting `hpe-auto-config`.

```
# nano /etc/default/grub
# systemctl restart hpe-auto-config
```

hpe_irqbalance

Syntax

```
hpe_irqbalance [OPTION]...
```

Description

This utility controls interrupt request (IRQ) affinity. The daemon starts when a system boots. If a device generates IRQs, `hpe_irqbalance` attempts to distribute the interrupts to the CPUs that are on the same chassis (or node) upon which the interrupt originated.

Options

-d

Run interactively in debug mode.

-o

Run in one-shot mode. This mode will distribute the currently existing IRQs once and exit, rather than continually polling for new IRQs.

-t

Specifies the sleeptime in seconds between polling loops, after which time `hpe_irqbalance` will poll for IRQs for which affinity has not been set. This option has no effect if `-o` is specified.

Specifiers

OPTION

hpe_irqbalance_oneshot

Same as the `-o` option.

hpe_irqbalance_sleeptime

Same as the `-t` option.



Usage

By default, this utility starts every two minutes. To change the interval, set `HPE_IRQBALANCE_SLEEPTIME` in `/etc/sysconfig/hpe_irqbalance` to the desired number of seconds and restart `hpe_irqbalance` or reboot your system.

The `hpe-auto-config` package configures the `hpe_irqbalance` utility automatically.

For more information, see the `hpe_irqbalance(8)` manpage.

hubstats

Syntax

```
hubstats [-f ARCHIVE_FILE ] [-t INTERVAL] [-c CONFIG [-1] [-b L-M,N,P,...]]  
cmd [ CMD_ARGS ]
```

```
hubstats -a ARCHIVE_FILE [-b L-M,N,P,... ] -c CONFIG
```

Description

Monitors the system during a command execution and writes a data archive that includes information about the operations that occur when the command runs.

Options

-c

Display *config* monitoring, where *config* is one of: **nlsSummary**, **linkstat**, **gruCmdTraffic**, **gruMemTraffic**, **lhCacheTraffic**, **rhCacheTraffic**, or **dump_metrics**.

-a

Report from a previously recorded `sdstats` archive. When reading from an archive, command is not required, and ignored. May be combined with `-b` and `-c` to give the desired output.

-b L-M,N,P,...

Specify a subset of nodes to monitor. For example to monitor nodes 0-4, 18, and 23:

```
hubstats -b 0-4,18,23
```

-t

Time interval between updates of the data. See *INTERVAL*.

-1

Enable live mode. When command is set and `hubstats` is in collecting mode, a cursor-based GUI displays *config* at every interval tick. For best results, use a terminal emulator that supports colors.

-f

Specify the output file for the data collection.

Specifiers

CONFIG

Specifies the type of statistics. The *CONFIG* value may be one of the following:

NLSUMMARY

Displays Superdome Flex Grid traffic.



LHCACHETRAFFIC

LocalHome (LH) request and reply traffic.

RHCACHETRAFFIC

RemoteHome (RH) request and reply traffic.

GRUCMDTRAFFIC

LocalHome (LH) and RemoteHome (RH) request and reply traffic.

GRUMEMTRAFFIC

Global reference unit (GRU) traffic statistics.

DUMP_METRICS

All. Dumps all `sdstats` metrics to standard output.

LINKSTAT

Interchassis pass-through bandwidth statistics.

CMD_ARGS

Specifies one or more of the arguments to the specified *CMD*.

INTERVAL

Specifies the time interval between updates of the data (`-t`). Units may be specified, for example, microseconds (us) - default, milliseconds (ms), and seconds (s). The live mode update rate (see `-l`) is also set by this interval.

ARCHIVE_FILE

Specify the output file for the data collection. By default an archive named `sdstats_registers.dat` is made in the current working directory. All register values and system topology information are saved in the archive, allowing you to view statistics for any `config` on or off the test system.

```
hubstats -f test1 -c nlSummary sleep 10
```

```
hubstats -a test1 -c gruMemTraffic
```

Usage

- Use the following format to generate and save them in archive:

```
hubstats [ -f ARCHIVE_FILE ] cmd [ CMD_ARGS ]
```

- Use the following format to display information from an existing archive:

```
hubstats -a ARCHIVE_FILE [ -b L-M,N,P,... ] -c CONFIG
```

- The `sdstats` library supplies the `hubstats` utility with statistics from the HPE Superdome Flex Server ASIC board.

-
- **NOTE:** The `sdstats` is not installed by default in HFS. The `sdstats` must be installed to use `hubstats`. Use `yum` (RHEL) or `zypper` (SLES) to install `sdstats`.
-

For more information on the `hubstats` command, see the `HUBSTATS(1)` manpage content.

Example

Run `hubstats` to collect data over the duration of the `sleep` command and display the `nlSummary` statistics.

```
> ht-027:~ # hubstats sleep 60
Creating new archive: sdstats_registers.dat
ht-027:~ # hubstats -a sdstats_registers.dat -c nlSummary
```

```

Reading from archive: sdstats_registers.dat
# Superdome Flex Grid Summary allocated over 60.005779 seconds
#
# (In MiB)          ----- LH Total -----   ----- RH total -----   ----- NL traffic -----
# Hub              Rx          Tx              Rx          Tx              Rx          Tx          Pass
1 r001i11b00h0     0.03     0.01           0.37     0.03           0.41     0.04     0.00
0 r001i11b00h1     2.12     0.67           0.67     2.12           2.79     2.79     0.00
3 r001i11b01h0     0.03     0.05           0.00     0.02           0.05     0.19     0.12
2 r001i11b01h1     0.08     0.07           0.06     0.07           0.14     0.14     0.00

```

nodeinfo

Syntax

nodeinfo [OPTION...]

Description

Monitors per-node Superdome Flex Grid memory statistics.

Options

-l

Print output line by line to stdout, instead of using the full screen display.

-s

Limit the number of data samples that nodeinfo outputs.

-i -t

Wait *i* seconds between printing each data sample.

-o

Sets the starting output options.

k

Show sizes in KB.

m

Show sizes in MB.

p

Show sizes in pages.

H

Show hugepage info.

-h

Show help message.

-v

Show verbose output.

Usage

nodeinfo is a tool for monitoring per-node Superdome Flex Grid memory statistics. The nodeinfo tool reads /sys/devices/system/node/*/meminfo and /sys/devices/system/node/*/numastat on the local system to gather Superdome Flex Grid memory statistics.



For more information on the `nodeinfo` command, see the `NODEINFO(1)` manpage.

From an interactive `nodeinfo` session, enter `h` for a help statement. For example:

```
Display memory statistics by node.
q    quit
+    Increase starting node number. Used only if more nodes than will
    fit in the current window.
-    Decrease starting node number. Used only if more nodes than will
    fit in the current window.
b    Start output with node 0.
e    Show highest node number.
k    show sizes in KB.
m    show sizes in MB.
p    show sizes in pages.
t    Change refresh rate.
A    Show/Hide memory policy stats.
H    Show/Hide hugepage info.
L    Show/Hide LRU Queue stats.

Field definitions:
hit - page was allocated on the preferred node
miss - preferred node was full. Allocation occurred on THIS node
      by a process running on another node that was full

foreign - Preferred node was full. Had to allocate somewhere
         else.

interlv - allocation was for interleaved policy

local - page allocated on THIS node by a process running on THIS node
remote - page allocated on THIS node by a process running on ANOTHER node

(press any key to exit from help screen)
```

Example

Sample memory statistics from the `nodeinfo` command are as follows:

```
# nodeinfo
Memory Statistics  Tue Oct 26 12:01:58 2017
flex-sys
----- Per Node KB -----
node   Total    Free    Used    Dirty    Anon    Slab    hit    Preferred Alloc ----- -- Loc/Rem ---
      miss foreign interlv  local  remote
0      16757488 16277084 480404    52    34284    36288 20724    0    0    0    20720    4
1      16777216 16433988 343228    68    6772    17708 4477    0    0    0    3381    1096
2      16777216 16438568 338648    76    6908    12620 1804    0    0    0    709    1095
3      16760832 16429844 330988    56    2820    16836 1802    0    0    0    708    1094
4      16777216 16444408 332808    88    10124   13588 1517    0    0    0    417    1100
5      16760832 16430300 330532    72    1956    17304 4546    0    0    0    3453    1093
6      16777216 16430788 346428    36    3236    15292 3961    0    0    0    2864    1097
7      16760832 16435532 325300    44    1220    14800 3971    0    0    0    2877    1094
TOT   134148848 131320512 2828336    492   67320   144436 42802    0    0    0    35129   7673
Press "h" for help
```

topology

Syntax

```
topology [OPTION ...]
```

Description

Provides topology information about your system. Application programmers can use the `topology` command to help optimize execution layout for their applications.



Options

If none of the options are specified, then all information is displayed. Any combination of options may be specified to limit the amount of information displayed. Option names may be shortened to their unique values.

--cpus

Show detailed information for the CPUs.

--cops

Show detailed information for attached coprocessors.

--disks

Show detailed information for the disk controllers including SCSI and Serial Attached SCSI controllers.

--gfx

Show detailed information for graphics devices including both VGA (graphics) GPUs and GP (CUDA) GPUs.

--io

Displays information about I/O devices. (**--disks**, **--gfx**, **--networks**)

--networks

Show detailed information for the network devices including Ethernet, Fibre Channel, and InfiniBand controllers.

--nodes

Show detailed information for the system nodes. With `verbose = 1`, it will also show which CPUs are resident on this node (blade). With routed configurations, the blades may not be numbered the same as the physical ordering.

--routers

Show detailed information for the system routers.

-s

Display a system summary (default).

Other operational modifiers are:

--affinity

Show `smp_affinity` for each IRQ (implies arg `--irqs`). With `verbose` set, it will also show the current interrupt count for each IRQ.

-d

Increase the debug level.

--debug=n

Set the debug level to `n`.

--get_node=<device>

Returns the lowest local node associated with where the device is installed. This information is useful for placing an application on the same node as the device, often improving DMA performance. With `verbose` set, it will return all nodes that are local to the device.

--get_cpus=<device>

Returns a list of cpus associated with where the device is installed. Normally returns a cpulist (`cpuJ-cpuK,cpuM-cpuN`). If `verbose`, returns a comma-separated list of individual cpus (`cpuJ,cpuJ+1,...`).

--noheaders

Do not display the column headers. For those cases where multiple devices are on the same node, this option causes the node info to be displayed on every line. This output is useful for parsing the output by other programs. The



format of the data may change and the user is advised to use the `--version` option to maintain compatibility with each new version.

--irqs

Shows which IRQs are assigned to the devices. With `verbose` set, it will also show the current interrupt count for each IRQ. Implies arg `--io` if none of "`--disks`, `--gfx`, or `--networks`" set.

--nox

Do not display the X Server Display column when displaying I/O. This output is implied if `--gfx` is not present or if there is only the BASE I/O VGA installed on the system and it is not configured in the `xorg` configuration file.

-tt

Forces use `/var/run/sgi_uv/topology` if available to obtain system topology information.

--lspci=file

Use file to obtain PCI device information. The file will contain output from the `lspci` command in machine-readable format.

-v

Increase the verbosity.

--verbose=n

Set the verbosity to 'n'.

--version

Display topology version.

--xorgconf=file

Use file to obtain X display information instead of the standard `/etc/X11/xorg.conf` file.

Specifiers

gfxtopology

Same as the `gfx` option to `topology`.

For more information about the `topology` command, see the `TOPOLOGY(1)` manpage content.

Example 1

The `topology` command includes many options. For more information, type `topology --help` on the command line.

The following `topology` command shows the system summary.

```
# topology
System type: Superdome Flex
System name: flex-rmc-009.example.com
Serial number: 5UF7090003
Partition number: 1
    1 Rack
    1 Chassis
    224 CPUs (online: 0-223)
    4 Nodes
487.48 GB Memory Total
125.00 GB Max Memory on any Node
    1 BASE I/O Card
    1 Co-processor
    4 Network Controllers
    1 SATA Storage Controller
```



```

1 USB Controller
1 VGA GPU
1 RAID Controller

```

Example 2

The following topology command explicitly requests the system summary and also shows node and CPU information.

```

# topology --summary --nodes --cpus
System type: Superdome Flex
System name: flex-rcm-009.example.com
Serial number: 5UF7090003
Partition number: 1
    1 Rack
    1 Chassis
    224 CPUs (online: 0-223)
    4 Nodes
487.48 GB Memory Total
125.00 GB Max Memory on any Node
    1 BASE I/O Card
    1 Co-processor
    4 Network Controllers
    1 SATA Storage Controller
    1 USB Controller
    1 VGA GPU
    1 RAID Controller

```

Node	Location	NASID	CPUS	Memory
0	r001i01b00h1	0002	56	126987 MB
1	r001i01b00h0	0000	56	127008 MB
2	r001i01b01h1	0006	56	127008 MB
3	r001i01b01h0	0004	56	127008 MB

CPU	Location	PhysID	CoreID	APIC-ID	Family	Model	Speed	L1 (KiB)	L2 (KiB)	L3 (KiB)
0	r001i01b00h1	00	00	0	6	85	1000	32d/32i	1024	39424
1	r001i01b00h1	00	01	2	6	85	1000	32d/32i	1024	39424
2	r001i01b00h1	00	02	4	6	85	1000	32d/32i	1024	39424
3	r001i01b00h1	00	03	6	6	85	1000	32d/32i	1024	39424
.
.
.

Example 3

The following topology command shows the interrupt requests that are assigned to devices.

```

# topology --irq
Index Location      NASID  PCI Address      IRQ(s)  Device
-----
0 r001i01s00      0002  0000:00:11.5     30      Intel Lewisburg SSATA Controller [AHCI]
. . .             .      0000:00:14.0     29      Intel Lewisburg USB 3.0 xHCI Controller
. . .             .      0000:00:17.0     31      Intel C600/X79 SATA RAID Controller
. . .             .      0000:00:1f.5     0       Intel Lewisburg SPI Controller
. . .             .      0000:01:00.0     19      Emulex x1 PCIe Gen2 Bridge[Pilot4]
. . .             .      0000:02:00.0     19      Matrox MGA G200e [Pilot] ServerEngines (SEPl)
. . .             .      0000:02:01.0     255     Emulex ServerView iRMC HTI
. . .             .      0000:c3:00.0     33-39,42-111 Intel Ethernet Connection X722 for 10GBASE-T
. . .             .      0000:c3:00.1     28,163-291 Intel Ethernet Connection X722 for 10GBASE-T
. . .             .      0000:c3:00.2     28,292-420 Intel Ethernet Connection X722 for 10GBASE-T
. . .             .      0000:c3:00.3     28,421-549 Intel Ethernet Connection X722 for 10GBASE-T

```

Example 4

The following topology command uses the -v option, which includes interrupt count information.



Rack management controller (RMC) commands

This chapter explains the RMC commands that are available to you when you log in as the `administrator` user.

The following RMC commands provide command information from the RMC command line:

- `apropos command`

The `apropos` command retrieves information about the `command` you enter. For `command`, specify the name of any RMC command.

For example, to retrieve command information that pertains to SNMP traps, enter the following:

```
RMC cli> apropos snmp
```

```
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
clear snmp_forwarding
remove snmp forward_address=FORWARD_ADDRESS
set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
show snmp
show snmp forward_address
```

- `help command` or `simply command`

When you enter an incomplete command at the RMC command prompt, the system returns the correct syntax and options for the command. For `command`, specify the name of any RMC command.

For example, to display information about the `enable` command, enter either of the following commands at the prompt.

```
RMC cli> enable
```

```
Usage: enable
       enable autopower
       enable bios_validate
       enable cert_checking
       enable http
       enable https
       enable ipmi
       enable ipv6
       enable ldap
       enable ler vendor_id=VENDOR_ID device_id=DEVICE_ID
       enable script_mode
       enable ssh
```

```
RMC cli> help enable
```

```
Usage: enable
       enable autopower
       enable bios_validate
       enable cert_checking
       enable http
       enable https
       enable ipmi
       enable ipv6
       enable ldap
       enable ler vendor_id=VENDOR_ID device_id=DEVICE_ID
```



```
enable script_mode
enable ssh
```

acquit

Syntax

```
acquit [physloc=HEXADECIMAL_STR | chassis=GEOID | pnum=PARTITION_ID | all]
```

Description

Releases the indicted and deconfigured status of a specified component manually in the following situations:

- After you examine the log files and replace the indicted hardware.
- When you suspect that the hardware is indicted erroneously.

Specifiers

HEXADECIMAL_STR

Specify an 8-byte hexadecimal integer for a physical location. The command does not accept short form strings.

GEOID

Specify the ID number of a chassis.

PARTITION_ID

Specify an nPartition containing one or more chassis.

Usage

- The `acquit physloc=HEXADECIMAL_STR` command acquits all indicted or deconfigured hardware at the specified physical location.
- The `acquit chassis=GEOID` command acquits all indicted or deconfigured hardware in the specified chassis.
- The `acquit pnum=PARTITION_ID` command acquits all indicted or deconfigured hardware in the specified nPartition.
- The `acquit all` command acquits all indicted or deconfigured hardware in the complex.

Example

The following example shows how to retrieve a list of indictments and how to acquit an indictment.

```
RMC cli> show indict
```

```
r001i01b
```

```
FFFF000101FFFF64 rack1/chassis_u1/harp_fan0
```

```
Indicted: Yes
```

```
2017-10-31T17:05:42Z Reason ID: 302
```

```
Summary: Fan reported to be running slowly or stopped
```

```
Cause: A fan is not running at normal speed. There may be something obstructing the airflow, a blocked fan or a fan failure.
```

```
Action: Check the enclosure for obstructions to airflow and
```

ensure the fans are installed properly. If the problem persists, replace the fan. After replacing the faulty fan, perform an acquittal from the RMC command line.

```
RMC cli> acquit physloc=FFFF000101FFFF64
```

add

Syntax

```
add certificate
add certificate ca file=FILENAME
add certificate signed [chassis=GEOID]
add dns [ipaddress=IPADDRESS | search=SEARCH_NAME]
add host name=NAME ipaddress=IPADDRESS
add ldap_group name=GROUP_NAME role={administrator,monitor,operator}
add location [rack=RACK] [bmc_id=BMC_ID] [module={rmc,bmc}] upos=UPOS
add sshkey key="PUBLIC KEY"
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT]
[protocol=PROTOCOL]
add user name=USERNAME role={administrator,monitor,operator}
```

Description

Specifies system characteristics or users to the system.

Specifiers

FILENAME

Specify the name of the file that provides the Certificate Authority.

IPADDRESS

Specify the IP address of a DNS server at your site.

FORWARD_ADDRESS

Specify the IP address of a local computer or device to receive SNMP alerts.

GEOID

Specify the ID number of the chassis to receive the signed certificate.

GROUP_NAME

Specify the LDAP directory group.

PUBLIC KEY

Specify the ssh public key.

USERNAME

Specify the name of a new user that you want to add to the system. The following are rules regarding the *USERNAME* specification:



- You can add a maximum of 30 users.
- *USERNAME* must begin with a letter and is case-sensitive.
- *USERNAME* must be unique to all other user names and group names.
- *USERNAME* can be 2 to 20 characters long and can include all alphanumeric characters, the dash, and the underscore.
- Certain user names are reserved for the RMC and eRMC. The reserved user names are as follows:

User name	Reserved on eRMC	Reserved on RMC
adminstrator	X	X
backup		X
bin		X
daemon	X	X
dbus		X
ftp		X
haldaemon		X
logger	X	X
mail		X
nobody		X
ntp	X	
operator		X
proxy		X
root	X	X
sshd	X	X
sys		X
sysadmin	X	
sync		X
www-data		X

When you create a user, the system also prompts you to create a password. The password can be from 8 to 40 characters long and can include all printable characters.

Usage

- The `add certificate ca file=FILENAME` command adds a specified Certificate Authority file for a signed certificate.
- The `add certificate signed` command adds a default-generated signed security certificate to the RMC.
- The `add certificate signed [chassis=GEOID]` command adds a signed security certificate to the specified chassis.
- The `add dns ipaddress=DNS_IPADDRESS` command specifies your site domain name service to the RMC.



- The `add host name=NAME ipaddress=IPADDRESS` command specifies the host name for the RMC.
- The `add ldap_group name=GROUP_NAME role={administrator,monitor,operator}` command specifies the LDAP group for the RMC and the user role for the RMC.
- The `add location [rack=RACK] [bmc_id=BMC_ID] [module={rmc,bmc}] upos=UPOS` command specifies the physical location of the RMC.
- The `add snmp forward_address=FWD_IPADDRESS` command specifies an IP address of the device to which you want the RMC to send SNMP alerts.
- The `add sshkey key=PUBLIC_KEY` command adds the provided SSH public key to the RMC.
- The `add user name=USERNAME` command adds a user to the RMC.

apropos

Syntax

```
apropos STRING_TO_MATCH [STRING_TO_MATCH ...]
```

Description

Retrieves a list of commands that include the string you specify.

Specifiers

STRING

Specify a topic area, a command name, or another word for which you want the *apropos* command to display information.

Example 1

```
RMC cli> apropos snmp
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
clear snmp_forwarding
remove snmp forward_address=FORWARD_ADDRESS
set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
show snmp
show snmp forward_address
```

Example 2

```
RMC cli> apropos add
add
add certificate
add certificate ca file=FILENAME
add certificate signed [chassis=GEOID]
add dns [ipaddress=IPADDRESS | search=SEARCH_NAME]
add host name=NAME ipaddress=IPADDRESS
add ldap_group name=GROUP_NAME role={administrator,monitor,operator}
add location [rack=RACK] [bmc_id=BMC_ID] [module={rmc,bmc}] upos=UPOS
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT]
[protocol=PROTOCOL]
add sshkey key=PUBLIC_KEY
add user name=USERNAME role={administrator,monitor,operator}
modify npar pnun=PARTITION_ID [hthread={on,off}] [memrefreshrate={x1,x2}]
[ras={on,hpc}] [secure_boot={on,off}] [boot_slots={3,5,8}]
[add_chassis=GEOID [GEOID ...]] [remove_chassis=GEOID [GEOID ...]]
```



```

remove dns [ipaddress=IPADDRESS | search=SEARCH_NAME]
remove host [name=NAME | ipaddress=IPADDRESS]
remove snmp forward_address=FORWARD_ADDRESS
set network [aux] [baseio=GEOID] [addressing={dhcp,static,default}]
[gateway=GATEWAY] [hostname=HOST] [ipaddress=IPADDRESS] [netmask=NETMASK]
set network ipv6 [addressing={dhcp,static}] [ipaddress=IPV6_ADDRESS/
PREFIX_LENGTH] [gateway=GATEWAY]
set network ipv6 aux [addressing={static,default}] [ipaddress=IPV6_ADDRESS/
PREFIX_LENGTH]
set network ipv6 baseio [chassis=GEOID] [addressing={dhcp,static}]
[ipaddress=IPV6_ADDRESS/PREFIX_LENGTH] [gateway=GATEWAY]
set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT]
[protocol=PROTOCOL]
show snmp forward_address

```

baseiolist

Syntax

```
baseiolist
```

Description

Returns the IP address and the MAC address of the BaseIO BMC.

Usage

The IP address is needed as input to the JViewer software. You can use JViewer to install an operating system on an HPE Superdome Flex Server from a remote location.

Example

```
RMC cli> baseiolist
```

```

P011 [r001i11b]: 192.168.1.1 [01:01:01:01:01:01]
                  fddd:6775:1110:93:a00:69ff:fe17:a1e1/64 Scope:Global
                  fe80::a00:69ff:fe17:a1e1/64 Scope:Link
                  fddd:6775:1110:93:1::1c8/64 Scope:Global

```

cancel autpower

Syntax

```
cancel autpower
```

Description

Prevents the HPE Superdome Flex Server from applying power to the server components after a power off.

Usage

When power is restored after a power outage, the power to the RMCs and the BMCs is restored automatically. If the autpower capability is enabled, the system restores power to the servers automatically, too. If you use this command to



cancel the autopower capability, the system does not restore power to the servers automatically after the power is restored.

clear

Syntax

```
clear dns
clear host
clear ldap [server] [port] [login] [binddn] [bindpw] [sid] [basedn]
[userfilter] [userdn] [groupfilter] [groupdn]
clear logs
clear logs console target={BMC_GEOID | PARTITION_ID}
clear logs debug
clear logs error type=SEVERITY [pnum=PARTITION_ID]
clear logs hel [target=TARGET]
clear logs iel [pnum=PARTITION_ID]
clear logs sel [target=TARGET]
clear logs uvdmp [file=FILENAME]
clear ntp
clear remote_log_server
clear rmc [product_data] [verbose]
clear snmp_forwarding
clear sshkey
```

Description

Removes or erases settings, files, or other entities.

Specifiers

FILENAME

Specify the name of the file to remove.

TARGET

Specify the identifier for one of the system RMCs. The `show complex` command returns system information that includes BMC identifiers.

SEVERITY

Specify one of the following severities: MCA, CMC, or ALL.

PARTITION_ID

Specify the numerical identifier for the partition. Specify 0.

GEOID

Specify the identifier of the log.



Usage

- The `clear dns` command removes the domain name service (DNS) setting on the RMC.
- The `clear ldap` command clears LDAP configuration attributes.
- The `clear logs console` command removes the console logs.
- The `clear logs debug` command removes files generated by the `collect debug` command.
- The `clear logs hel` command removes the hardware error logs (HEL).

If you specify the `target=target` argument, the command removes the logs on the specified BMC. For example:

```
RMC cli> clear logs hel target=r001i01b
```

- The `clear logs iel` command removes the integrated event log (IEL).
- The `clear logs sel` command removes the system event log (SEL).

If you specify the `target=TARGET` argument, the command removes the system event log on the specified BMC.

- The `clear logs error type=SEVERITY [pnum=PARTITION_ID]` command removes errors of the specified severity. If a partition number is specified, only errors for the specified partition of the specified severity are removed.

NOTE: HPE does not recommend clearing logs to ensure that failure information is available for diagnostic purposes.

- The `clear logs uvdmp [file=FILENAME]` command removes the specified uvdmp file.
- The `clear ntp` command removes the network time protocol (NTP) setting on the RMC.
- The `clear remote_log_server` command clears the configured remote log server settings and disables sending system log to the remote log server. RMC reboot is required after the clear command to disable logging to the remote server.
- The `clear rmc [product_data] [verbose]` command displays and then clears all product data from the RMC.
- The `clear snmp_forwarding` command removes the simple network management protocol (SNMP) forwarding address setting on the RMC.
- The `clear sshkey` command removes the ssh key.

To clear the key, first use the `show sshkey` command to display the key, and then use the `clear sshkey` command to remove the key.

collect debug

Syntax

```
collect debug
```

Description

Gathers debugging data into a compressed archive file.

Usage

- You can use the `upload debug` command to transfer the file to another system at your site for eventual analysis.
- The approximate time it takes to create the archive file is as follows:
 - Allow up to 90 minutes for the eRMC.
 - Allow up to 25 minutes for the RMC.
 - Allow up to 5 minutes per BMC.

commands

Syntax

commands

Description

Retrieves the list of valid RMC commands and the command formats.

Example output

```
RMC cli> commands
acquit
acquit [physloc=HEXADECIMAL_STR | chassis=GEoid | pnum=PARTITION_ID | all]
add
add certificate
add certificate signed [chassis=GEoid]
add dns [ipaddress=IPADDRESS | search=SEARCH_NAME]
add host name=NAME ipaddress=IPADDRESS
add ldap_group name=GROUP_NAME role={administrator,monitor,operator}
add location [rack=RACK] [bmc_id=BMC_ID] [module={rmc,bmc}] upos=UPOS
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
add sshkey key=PUBLIC KEY
add user name=USERNAME role={administrator,monitor,operator}
apropos
cancel
cancel autopower
clear
clear cae (complex | pnum=PARTITION_NUM)
clear dns
clear host
clear ldap [server] [port] [login] [binddn] [bindpw] [sid] [basedn] [userfilter] [userdn] [groupfilter] [groupdn]
clear logs
.
.
.
```

connect

Syntax

```
connect npar pnum=PARTITION_ID [clear] [full][nocache] [notty] [kill] [spy | steal]
```

Description

Opens a console session to the partition. To exit the console session, type CTRL-] and enter q.

Specifiers

PARTITION_ID

Specify the numerical identifier for the partition.



Usage

- `connect npar pnum=PARTITION_ID` connects you to the partition in a console window.
- `connect npar pnum=PARTITION_ID clear` removes cached console data.
- `connect npar pnum=PARTITION_ID full` dumps out the entire console data cache and then goes to a live console.
- `connect npar pnum=PARTITION_ID nocache` suppresses the display of cached data.
- `connect npar pnum=PARTITION_ID notty` suppresses the interactive console capability.
- `connect npar pnum=PARTITION_ID kill` terminates all existing console connections except for the console at which this command is entered.
- `connect npar pnum=PARTITION_ID spy` enables you to watch another console in read-only mode.
- `connect npar pnum=PARTITION_ID steal` assumes control of the console by force.

You can specify zero or more of the `clear`, `full`, `nocache`, `notty`, or `kill` parameters. You can specify either `spy` or `steal`.

create

Syntax

```
create npar pnum=PARTITION_ID [hthread={on,off}] [ras={on,hpc}]
[memrefreshrate={x1,x2}] [secure_boot={on,off}] [boot_slots={none,3,5,8}]
chassis={GEOID [GEOID ...],all,unassigned}

create npar default
```

Description

Configures partitions.

Specifiers

PARTITION_ID

Specify the numerical identifier for the partition. Valid partition numbers are from 0 to 254.

GEOID

Specify the identifier of the server chassis. The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, *GEOID* `r001i06b` identifies the chassis in rack number 1 with U position 6.

hthread={on,off}

Enables (default) or disables Hyper-Threading Technology for the partition.

ras={on,hpc}

Affects memory error handling.

The `ras=on` specification is the default. This setting enables ADDDC mode. This mode allows better error recovery. When set, all memory error handling occurs in the BIOS, rather than `memlog`.

secure_boot={on,off}

Enables or disables (default) secure boot for the partition.



boot_slots={3, 5, 8}

Specifies the PCI slot numbers that are enabled for boot devices and for option ROM execution. The default is 3 and 5. Only slots 3, 5, and/or 8 support boot capability. If multiple slots are specified, they must be separated by commas.

chassis={*GEOID* [*GEOID* ...], all, unassigned}

Specifies the chassis you want to include in the partition. You can specify one or more *GEOID* specifications, the keyword *all*, or the keyword *unassigned*. This parameter is required.

chassis=GEOID specify the identifier of the server chassis. The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, *GEOID* `r001i06b` identifies the chassis in rack number 1 with U position 6.

chassis=all includes all chassis in the partition. This option is destructive. It removes any existing partitions before attempting to create a single partition. The command does nothing if chassis span fabric zones.

chassis=unassigned includes independent chassis that are currently not configured into other partitions. These chassis must belong to a single fabric zone.

memrefreshrate={x1, x2}

Affects the memory refresh rate for DDR4 DIMMs. The supported options are *x1* or *x2* and configured to *x1* by default.

NOTE: The 2x memory refresh rate might affect the performance and resiliency of the server memory.

Usage

- Use the `create` command to configure partitions. By default, the HPE Superdome Flex Server has one partition, partition 0. You can, however, divide the HPE Superdome Flex Server into multiple partitions. The number of partitions you can create depends on the number and type of chassis included in the system.
- Use the `create npar default` command to automatically create partitions from the available chassis. If you specify the `default` parameter, do not specify a `pnum=` parameter, a `chassis=` parameter, or any optional parameters. This parameter is destructive. It removes any existing partitions before creating the default partitions. Default values are used for all creation options. The number of default partitions created is determined by the number of fabric zones.

deconfig

Syntax

```
deconfig physloc=HEXADECIMAL_STR
```

Description

Manually removes a hardware component from the configuration.

Specifiers

HEXADECIMAL_STR

Can be an 8-byte hexadecimal integer. For example:

- `FF01030106FFFF44`
- `0xFF01030106FFFF44`



The command does not accept short form strings. The string you specify must appear in the output from the `show chassis info` command.

disable

Syntax

```
disable autopower
disable bios_validate
disable cert_checking
disable http
disable https
disable ipmi
disable ipv6
disable ldap
disable ler vendor_id=VENDOR_ID device_id=DEVICE_ID
disable script_mode
disable ssh
```

Usage

- The `disable autopower` command prevents the system from applying power to the server components after a power off.
- The `disable bios_validate` command disables BIOS validation after a power on.
- The `disable cert_checking` command disables certificate checking for TLS connections.
- The `disable http` and `disable https` commands disable HTTP or HTTPS access to the RMC.

NOTE: HTTP is also disabled after performing a factory reset.

- The `disable ipmi` command disables remote access to IPMI through the RMC.

NOTE: IPMI is disabled by default and is also disabled after performing firmware updates or using the `reset factory` command.

- The `disable ipv6` command reconfigures the RMC to disable the external IPv6 interface after the next RMC reboot. IPv4 networking must then be used.
- The `disable ldap` command disables LDAP management of the RMC.
- The `disable ler` command disables PCIe Live Error Recovery for the specified HPE supported I/O card.

NOTE: Disabling LER is only recommended for those cases where FPGA cards are reprogrammed under an OS.

- The `disable script_mode` command returns the command-line interface to interactive mode.
- The `disable ssh` command disables SSH access to the RMC.



download

Syntax

```
download certificate ca url=URL
download certificate crl url=URL
download certificate signed url=URL [chassis=GEOID]
download config (url=URL | file=FILENAME)
download npar config (url=URL | file=FILENAME)
```

Description

Downloads a certificate, configuration or nPartition configuration to the RMC.

Specifiers

URL

Specify the full path to the source location on a local computer. Accepted transfer methods are `https`, `sftp` and `scp`.

The following examples show the supported formats:

- `https://user@localhost/path`
- `sftp://user@localhost/path`
- `scp://user@localhost/path`

The command prompts you for the password to the local host computer.

FILENAME

Specify the name of the file.

GEOID

Specify the identifier of the server chassis where you want to download the certificate. The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, *GEOID* `r001i06b` pertains to the chassis in rack number 1 with U position 6.

Usage

- The `download certificate ca url=URL` command downloads a certificate authority from the specified URL.
- The `download certificate crl url=URL` command downloads a certificate revocation list from the specified URL.
- The `download certificate signed url=URL [chassis=GEOID]` command downloads a signed certificate to the specified chassis. If a chassis is not specified, the command downloads a signed certificate to the RMC.
- The `download config (url=URL | file=FILENAME)` command downloads a configuration file at the specified URL to the RMC.
- The `download npar config (url=URL | file=FILENAME)` downloads the nPartition configuration file at the specified URL to the RMC.



enable

Syntax

```
enable autopower
enable bios_validate
enable cert_checking
enable http
enable https
enable ipmi
enable ipv6
enable ldap
enable ler vendor_id=VENDOR_ID device_id=DEVICE_ID
enable script_mode
enable ssh
```

Usage

- The `enable autopower` command lets the system apply power to the server components after a power off.
- The `enable bios_validate` command enables BIOS validation after a power on.
- The `enable cert_checking` command enables certificate checking for TLS connections.
- The `enable http` and `enable https` commands enable HTTP or HTTPS access to the RMC.

NOTE: HTTP is also disabled after performing a factory reset.

- The `enable ipmi` command enables IPMI through the RMC.

NOTE: IPMI is disabled by default and is also disabled after performing firmware updates or using the `reset factory` command.

- The `enable ipv6` command reconfigures the RMC to enable the external IPv6 interface after the next RMC reboot.
- The `enable ldap` command enables LDAP management of the RMC.
- The `enable ler` command enables PCIe Live Error Recovery for the specified HPE supported I/O card.
- The `enable script_mode` command turns off the console interactive mode. In script mode, when the command line prompts you to enter `y` or `n` in answer to a question, the system assumes a `y` and continues.
- The `enable ssh` command enables SSH access to the RMC.

exit

Syntax

```
exit
```



Description

Closes the RMC CLI session from where the command is entered.

generate certificate

Syntax

```
generate certificate
generate certificate request [chassis=GEOID]
generate certificate selfsigned [chassis=GEOID]
```

Description

Generates a certificate for the RMC/eRMC or the BMC of the specified chassis.

Specifiers

GEOID

Specify the chassis for which to generate the certificate.

Usage

The `generate certificate request` command generates a CSR certificate for the RMC/eRMC.

The `generate certificate request chassis=GEOID` requests a CSR certificate for the BMC of the specified chassis.

The `generate certificate selfsigned` command generates a self-signed certificate for the RMC/eRMC.

The `generate certificate selfsigned chassis=GEOID` command generates a self-signed certificate for the BMC of the specified chassis.

help

Syntax

```
help COMMAND [OPTION]
```

Description

Returns information about an RMC command. Adding a command *OPTION* returns information about that option.

Specifiers

COMMAND

Any of the RMC commands.

OPTION

Any valid option for the RMC command.

Example

The following example uses the `apropos` and `help` commands together.

```
RMC cli> apropos snmp
```

```
add snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
clear snmp_forwarding
```



```
remove snmp forward_address=FORWARD_ADDRESS
set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]
show snmp
show snmp forward_address

RMC cli> help show snmp

Usage: show snmp
       show snmp forward_address

RMC cli> help show snmp forward_address

Shows SNMP forwarding addresses.
Usage: show snmp forward_address
```

indict

Syntax

```
indict physloc=HEXADECIMAL_STR
```

Description

Manually marks a hardware component as failed.

Specifiers

HEXADECIMAL_STR

Can be an 8-byte hexadecimal integer. For example:

- FF01030106FFFF44
- 0xFF01030106FFFF44

The command does not accept short form strings.

Usage

After you replace the failed component, run the `acquit` command to clear the indictment. The string you specify must appear in the output for the affected component from the `show chassis info` command.

ipmi

Syntax

```
ipmi command=CMD_STRING
```

Description

Lets you run industry standard `ipmi` commands on the HPE Superdome Flex Server.



Specifiers

CMD_STRING

Specifies one or more `ipmi` commands. To retrieve a list of possible commands, enter the following at the RMC command prompt:

```
ipmi command=help
```

Enclose the `CMD_STRING` in quotation marks (" ") if you specify more than one command. If the `CMD_STRING` contains spaces, you must enclose it in quotation marks.

Example

```
RMC cli> ipmi command=help
```

Commands:

raw	Send a RAW IPMI request and print response
i2c	Send an I2C Master Write-Read command and print response
spd	Print SPD info from remote I2C device
lan	Configure LAN Channels
chassis	Get chassis status and set power state
power	Shortcut to chassis power commands
event	Send pre-defined events to MC
mc	Management Controller status and global enables
sdr	Print Sensor Data Repository entries and readings

```
.  
.  
.
```

```
RMC cli> ipmi command="chassis power status"
```

```
Chassis Power is off
```

modify

Syntax

```
modify chassis target=GEOID [asset_tag=ASSET_TAG] [serial=SERIAL_NUMBER]  
[type={1-15}] [part_num=PRODUCT_ID] [model=PRODUCT_NAME] [verbose]
```

```
modify ldap_group name=GROUP_NAME role={administrator,monitor,operator}
```

```
modify npar pnum=PARTITION_ID [hthread={on,off}] [memrefreshrate={x1,x2}]  
[ras={on,hpc}] [secure_boot={on,off}] [boot_slots={3,5,8}] [add_chassis=GEOID  
[GEOID ...]] [remove_chassis=GEOID [GEOID ...]]
```

```
modify rmc [asset_tag=ASSET_TAG] [serial=SERIAL_NUMBER]  
[part_num=PRODUCT_ID] [verbose]
```

Description

Lets you change the hardware features on a specified partition.

Specifiers

GROUP_NAME

Specify the LDAP directory group.

PARTITION_ID

Specify the numerical identifier for the partition.



GEOID

Specify the identifier of the server chassis for which you want to disable or enable a unit identification number (UID). The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, *GEOID* `r001i06b` pertains to the chassis in rack number 1 with U position 6.

boot_slots

Specifies the PCI slot numbers that are enabled for boot devices and for option ROM execution. The default is 3 and 5. Only slots 3, 5, and/or 8 support boot capability. If multiple slots are specified, they must be separated by commas.

Usage

- The `modify ldap_group name=GROUP_NAME role={administrator,monitor,operator}` command changes the LDAP directory group or role for the RMC.
- The `modify npar pnum=PARTITION_ID hthread={on,off}` command enables or disables Hyper-Threading Technology.
- The `modify npar pnum=PARTITION_ID [secure_boot={on,off}]` command enables or disables secure boot.
- The `modify npar pnum=PARTITION_ID ras={on,hpc}` command affects memory error handling by enabling or disabling the reliability, availability, and serviceability (RAS) feature.



CAUTION: Hewlett Packard Enterprise strongly recommends using the `ras=on` specification to enable memory RAS features (ADDDC mode). HPE RAS features provide higher resiliency to DIMM faults versus standard memory error-correcting code (ECC).

The `ras=hpc` specification disables memory RAS features and therefore could result in compromised system resiliency and a potential server outage.

You will experience high failure rates of DIMMs using the `ras=hpc` specification.

To check the memory RAS setting, enter the `show npar verbose` RMC command.

- The `modify npar pnum=PARTITION_ID boot_slots=n [,n [,n]]` command adjusts which PCI slot numbers are enabled for boot devices and for option ROM execution.
- The `modify npar pnum=PARTITION_ID add_chassis=GEOID [GEOID ...]` command configures one or more chassis to an existing partition. All specified chassis must belong to the same fabric zone.
- The `modify npar pnum=PARTITION_ID remove_chassis=GEOID [GEOID ...]` command removes one or more chassis from an existing partition.
- The `modify rmc` command changes the RMC's stored asset tag, part number, and serial number. The `modify rmc` command with the exception of changing the asset tag is for use only by service personnel.
- The `modify chassis` command with the exception of the changing chassis asset tag is for use only by service personnel.
- The `modify npar pnum=PARTITION_ID [memrefreshrate={x1,x2}]` command can be used to adjust the memory refresh rate for DDR4 DIMMs, the default memory refresh rate is x1.

NOTE: The 2x memory refresh rate might affect the performance and resiliency of the server memory.

You can specify zero or more of the `hthread`, `memrefreshrate`, `ras`, `boot_slots`, `add_chassis`, or `remove_chassis` parameters.

ping

Syntax

```
ping [-LRUbdnqrvVaAD0] [-c count] [-i interval] [-w deadline]
      [-p pattern] [-s packetsize] [-t ttl] [-I interface]
      [-M pmtudisc-hint] [-m mark] [-S sndbuf]
      [-T tstamp-options] [-Q tos] [hop1 ...] destination
```

Description

Sends an echo request to a server on an IPv4 network to test reachability.

ping6

Syntax

```
ping6 [-LUdfnqrvVaAD0] [-c count] [-i interval] [-w deadline]
       [-p pattern] [-s packetsize] [-t ttl] [-I interface]
       [-M pmtudisc-hint] [-S sndbuf] [-F flowlabel] [-Q tclass]
       [-N nodeinfo-suboptions]
       [hop1 ...] destination
```

Description

Sends an echo request to a server on an IPv6 network to test reachability.

power

Syntax

```
power
power cycle [force]
power cycle bmc
power cycle npar pnum=PARTITION_ID [bootopt=BOOTOPT] [force]
power off [force]
power off npar pnum=PARTITION_ID [force]
power on
power on npar pnum=PARTITION_ID [bootopt=BOOTOPT]
power reset [force]
power reset bmc
power reset npar pnum=PARTITION_ID [bootopt=BOOTOPT] [force]
```

Description

Enables you to control power to the following:

- The HPE Superdome Flex Server system
- Specific components of the HPE Superdome Flex Server system



Specifiers

PARTITION_ID

Specify the numerical identifier for the partition.

BOOTOPT

Specify the boot option.

Usage

- `power` displays the quantity of chassis in each power mode.
- If power is on, the `power cycle npar pnum=PARTITION_ID [bootopt=BOOTOPT] [force]` command performs an OS immediate (non-graceful) shutdown and the momentarily removes power from the specified nPartition. Then the command turns on partition power and starts the nPartition. If `force` is specified, no user confirmation occurs.
- The `power cycle [force]` command functions the same as `power cycle npar` if a single nPartition is configured. This command is invalid if multiple nPartitions are configured. If `force` is specified, no user confirmation occurs.
- The `power cycle bmc` command removes auxiliary power from all BMCs in an nPartition and then turns on auxiliary power.

NOTE: All chassis in the nPartition are powered off when this command is used.

- The `power off npar pnum=PARTITION_ID [force]` command performs a graceful OS shutdown and then removes power from the specified nPartition with no user confirmation. If `force` is specified, the command performs an OS immediate (non-graceful) shutdown instead.
- The `power off [force]` command functions the same as `power off npar` if a single nPartition is configured. This command is invalid if multiple nPartitions are configured.
- The `power on npar pnum=PARTITION_ID [bootopt=BOOTOPT]` command turns on power to the specified nPartition and starts the nPartition power.
- The `power on` command functions the same as `power on npar` if a single nPartition is configured. This command is invalid if multiple nPartitions are configured. If `force` is specified, no user confirmation occurs.
- If power is on, the `power reset npar pnum=PARTITION_ID [bootopt=BOOTOPT] [force]` command performs a graceful OS shutdown then restarts the specified nPartition. If `force` is specified, it performs an OS immediate (non-graceful) shutdown instead. This command is invalid if power is off; you must use `power on` instead.
- The `power reset [force]` command functions the same as `power reset npar` if a single nPartition is configured. This command is invalid if multiple nPartitions are configured.
- The `power reset bmc` command restarts all BMCs in the specified nPartition after user confirmation.
- The RMC and the BMCs are always powered up if the RMC and BMCs are connected to an AC power source. You cannot power off the RMC or the BMCs, however you can power cycle the BMCs.

More information

[Partition boot options](#)



reboot

Syntax

```
reboot
reboot npar pnum=PARTITION_ID [bootopt=BOOTOPT] [force]
reboot rmc
reboot bmc target=GEOID
```

Description

Reboots the RMC or the HPE Superdome Flex Server.

Specifiers

PARTITION_ID

Specify the ID number of the nPartition.

GEOID

Identifies the specific BMC.

BOOTOPT

Specify the boot option.

Usage

- If power is on, the `reboot npar pnum=PARTITION_ID [force]` command performs an OS graceful shutdown, then restarts the specified nPartition. If *force* is specified, the command instead performs an OS immediate (nongraceful) shutdown. This command is invalid if power is off. Use the `power on` command instead.
- The `reboot` command functions the same as the `reboot npar` command when a single nPartition is configured. This command is invalid if multiple nPartitions are configured. No *force* option is supported for this command; use the `reboot npar` command instead.
- The `reboot rmc` command prompts the user that this action may disrupt operations, and then after confirmation restarts the RMC, disconnecting all current user sessions.
- The `reboot bmc target=GEOID` command prompts the user that this action may disrupt operations, and then after confirmation restarts the specified BMC, which may be in any chassis in the complex.

More information

[Partition boot options](#)

remove

Syntax

```
remove certificate ca (file=FILENAME | all)
remove certificate crl (file=FILENAME | all)
remove config (file=FILENAME | all)
remove debug [file=FILENAME]
remove dns [ipaddress=DNS_IPADDRESS | search=SEARCH_NAME]
```



```
remove ermc
remove host [name=NAME | ipaddress=IPADDRESS]
remove ldap_group name=GROUP_NAME
remove ldap_user_data
remove location [module={rmc,bmc}] bmc_id=BMC_ID
remove npar pnum=PARTITION_ID
remove redfish event_subscription (id=ID | all)
remove sshkey key=PUBLIC_KEY
remove snmp forward_address=FORWARD_ADDRESS
remove user name=USERNAME
```

Description

Specifies system characteristics or users to be removed from the HPE Superdome Flex Server.

Specifiers

DNS_IPADDRESS

Specify the IP address of the DNS server at your site.

PARTITION_ID

Specify the numerical ID for the nPartition.

FORWARD_ADDRESS

Specify the IP address of a local computer or device to receive SNMP alerts.

PUBLIC_KEY

Specify the SSH public key. The `show sshkey` command returns the key values.

USERNAME

Specify the name of the user that you want to delete from the system.

FILENAME

Specify the name of the file you want to delete from the system.

Usage

The `remove certificate ca (file=FILENAME | all)` command removes the specified certificate authority file or all certificate authority files.

The `remove certificate crl (file=FILENAME | all)` command removes the specified certificate revocation list file or all certificate revocation list files.

The `remove config file=FILENAME` command removes the specified configuration file.

The `remove config all` command removes all configuration files present in the RMC/eRMC. To see configuration file details, enter the command `show config list`.

The `remove debug [file=FILE]` command removes the specified debug file. To see the debug file details, enter the command `show debug`.

The `remove dns [ipaddress=DNS_IPADDRESS | search=SEARCH_NAME]` command removes the specified DNS server address or the DNS server matching the specified search name.



The `remove ermc` command removes the active eRMC. This command effectively converts the eRMC to a BMC. For use cases and examples, enter the command `help add location`.

The `remove host [name=NAME | ipaddress=IPADDRESS]` command removes the specified host.

The `remove ldap_group name=GROUP_NAME` command removes the specified LDAP group.

The `remove ldap_user_data` command removes all LDAP user data.

The `remove location [module={rmc,bmc}] bmc_id=BMC_ID` command removes a BMC entry or the eRMC from the configuration.

The `remove npar pnum=PARTITION_ID` command removes the specified nPartition. This command can only be used if the specified nPartition is powered off.

The `remove redfish event_subscription (id=ID | all)` command removes the specified or all Redfish event subscription(s).

The `remove sshkey key=PUBLIC_KEY` command removes the specified SSH public key.

The `remove snmp forward_address=FORWARD_ADDRESS` command removes the specified SNMP forwarding IP address.

The `remove user name=USERNAME` command removes the specified user.

NOTE: All existing CLI sessions for the user are logged off when the user is removed.

restore

Syntax

```
restore config (url=URL | file=FILENAME)
```

```
restore npar config (url=URL | file=FILENAME)
```

Description

Download and then execute a previously saved configuration script file. This restores user configurable settings on the complex.

Descriptors

URL

Specify the full path to the source location on a local computer. Accepted transfer methods are `https`, `sftp`, and `scp`.

The following examples show the supported formats:

- `https://user@localhost/path`
- `sftp://user@localhost/path`
- `scp://user@localhost/path`

The command prompts you for the password to the local host computer.

FILENAME

Specify the name of the file.



Usage

- The `restore config (url=URL | file=FILENAME)` command restores RMC configuration using the specified configuration file at the specified URL.
- The `restore npar config (url=URL | file=FILENAME)` restores nPartition configuration from the specified file at the specified URL.

save

Syntax

```
save config (url=URL | file=FILENAME)
```

```
save npar config (url=URL | file=FILENAME)
```

Description

Upload a configuration script file of all user configurable settings to a server or RMC local file.

Descriptors

URL

Specify the full path to the source location on a local computer. Accepted transfer methods are `https`, `sftp`, and `scp`.

The following examples show the supported formats:

- `https://user@localhost/path`
- `sftp://user@localhost/path`
- `scp://user@localhost/path`

The command prompts you for the password to the local host computer.

FILENAME

Specify the name of the file.

Usage

- The `save config (url=URL | file=FILENAME)` command saves the specified RMC configuration file at the specified URL.
- The `save npar config (url=URL | file=FILENAME)` saves the specified nPartition configuration file at the specified URL.

set

Syntax

```
set autopower delay=SECONDS
```

```
set date [yr=YEAR] [mo=MONTH] [day=DAY] [hr=HOURL] [min=MINUTE] |  
[str=DATE_STRING]
```

```
set factory [initialize]
```

```
set failed_login delay=SECONDS
```



```

set hostname [name=HOST_NAME | default]
set ipmi_watchdog (disabled | os_managed)

set ldap [server=SERVER] [port=SSL_PORT] [login=LOGIN_ATTRIBUTE]
[binddn=BIND_DN] [bindpw] [sid=AD_DOMAIN_SID] [basedn=BASE_DN]
[userfilter=USER_FILTER] [userdn=USER_DN] [groupfilter=GROUP_FILTER]
[groupdn=GROUP_DN]

set network [aux] [baseio=GEOID] [addressing={dhcp,static,default}]
[gateway=GATEWAY] [hostname=HOST] [ipaddress=IPADDRESS] [netmask=NETMASK]

set network internal [bmc_subnet=BMC_SUBNET] [cmc_subnet=CMC_SUBNET]
[ermc_subnet=ERMC_SUBNET]

set network internal default

set network ipv6 [addressing={dhcp,static}] [ipaddress=IPV6_ADDRESS/
PREFIX_LENGTH] [gateway=GATEWAY]

set network ipv6 aux [addressing={static,default}] [ipaddress=IPV6_ADDRESS/
PREFIX_LENGTH]

set network ipv6 baseio [chassis=GEOID] [addressing={dhcp,static}]
[ipaddress=IPV6_ADDRESS/PREFIX_LENGTH] [gateway=GATEWAY]

set ntp server=SERVER

set password [username=USERNAME]

set password default

set prompt [string=STRING | default]

set remote_log_server server_address=SERVER_ADDRESS [port=PORT]
[protocol=PROTOCOL]

set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT]
[protocol=PROTOCOL]

set timezone str=TIMEZONE_SPEC

set uid off chassis=GEOID

set uid on chassis=GEOID

```

Description

Specifies system characteristics.

Specifiers

SECONDS

Specify the number of seconds you want to elapse between power being applied to the HPE Superdome Flex Server and the server powering up. When power is reapplied to the system, the RMCs and BMCs power up automatically. You can delay the server power-up by the number of seconds specified.

For example, after a power outage, all servers power up at the same default time interval after power is restored. To stagger each server to power up at a slightly different time, you can specify a different number of *SECONDS* for each server. This practice avoids having all servers power up at the same time, which can lead to a blown circuit breaker in the computer center.

YEAR

Specify the year as a four-digit integer. For example, 2018.

MONTH

Specify the month as a two-digit integer. For example, 06 or 12.

DAY

Specify the day as a two-digit integer. For example, 02 or 31.

HOUR

Specify the hour as a two-digit integer against a 24-hour clock. For example, 08 or 14

MINUTES

Specify the minutes as a two-digit integer.

DATE_STRING

Specify the date in a quoted string. For example:

```
set date str="2020-01-25 13:27"
```

To retrieve the date in a format that you can use as input to the `set date` command, enter the following command on a local Linux or UNIX system:

```
date "+%Y-%m-%d %H:%M"
```

FORWARD_PORT

Specify a network port to receive SNMP notifications.

GATEWAY

Specify a network gateway.

HOST_NAME

Specify a hostname for the RMC. If you specify the keyword of `default`, the RMC displays the default prompt, which is the RMC *geoid*.

SERVER

Specify an appropriate IP address or host name.

NETMASK

Specify the network mask.

PROTOCOL

Specify the network protocol.

STRING

Specify the command prompt you want on the RMC. If *STRING* contains space characters, enclose *STRING* in quotation marks. For example:

```
RMC cli> set prompt string="my rmc"  
RMC cli> set prompt string=flex-rcm
```

NUMBER

Specify the number of the rack in which the RMC resides.

UPOS

Specify the U position of the RMC. In the rack, the U positions are numbered from bottom to top. The bottom-most U position is 1.



TIMEZONE_SPEC

Specify one of the time zone codes. For example, EST or Canada/Central. For the complete time zone code list, log into the RMC and enter the following command:

```
RMC cli> help set timezone
```

GEOID

Specify the identifier of the server chassis for which you want to disable or enable a unit identification number (UID). The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, *GEOID* `r001i06b` pertains to the chassis in rack number 1 with U position 6.

PORT

Specify a network port in the remote log server to receive the system events

PROTOCOL

Specify the network protocol used to transfer the system events to remote log server, only `udp` is supported.

SECONDS

Specify the delay for a failed login in seconds, allowed values are 0 to 240.

Usage

- The `set autopower delay=SECONDS` command sets the delay in seconds before the system powers on after applying power.
- The `set date` command sets the date using the specified parameters.
- The `set factory [initialize]` command resets the system to factory default settings and optionally initializes the system.
- The `set failed_login delay=SECONDS` command sets the login delay value in seconds for a failed login. Login is restricted till the configured delay is reached. This delay value is in seconds. Valid values are 0 to 240 seconds. Default value 0 indicates failed login is disabled. This feature is applicable only for local users but not for remote users such as LDAP users.
- The `set hostname [name=HOST_NAME | default]` command sets the system hostname to the specified name or a default-generated value.
- The `set ipmi_watchdog (disabled | os_managed)` command disables or enables the IPMI watchdog functionality to be managed by the OS. The default is disabled.
- The `set ldap` command sets the LDAP server using the specified parameters.
- The `set network` command sets the primary external networking using the specified parameters or to default generated values.
- The `set network aux` sets the auxiliary external networking using the specified parameters or to default generated values.
- The `set network internal` command sets the internal networking subnet using the specified parameters or to default generated values.
- The `set network ipv6` command sets primary IPv6 networking to the specified parameters.
- The `set network ipv6 aux` command sets auxiliary IPv6 networking to the specified parameters.
- The `set network ipv6 baseio` command sets the IPv6 networking for the BaseIO.
- The `set ntp SERVER=SERVER` command sets the NTP server to the specified IP address or server host name.

- The `set password [username=USERNAME]` command sets the user password for the specified user account. Passwords must meet the following requirements:
 - Passwords may include combinations of these types of characters:
 - Upper case letters
 - Lower case letters
 - Numbers
 - Special characters - including `!@#$%^&*()!@`
 - Password length is dependent on the types of characters used. The minimum length is six characters, with the minimum length increased by two characters for each type not included.
 - Passwords only containing one type of character must be at least 12 characters.
 - Passwords containing two types of character must be at least 10 characters.
 - Passwords containing three types of character must be at least 8 characters.
 - Passwords containing all four types of characters must be at least 6 characters.
- The `set prompt` command sets the CLI prompt to the specified string or to the default value. If you want to use a prompt character such as `>` or `:`, it must be included in the string. The prompt string is stripped of leading and trailing spaces.
- The `set remote_log_server server_address=SERVER_ADDRESS [port=PORT] [protocol=PROTOCOL]` command sets the IP address or hostname of remote log server and enables remote logging. The port number and protocol to be used can optionally be provided for connecting to the remote log server, the supported protocol is UDP.
- The `set snmp forward_address=FORWARD_ADDRESS [port=FORWARD_PORT] [protocol=PROTOCOL]` command sets SNMP forwarding to the specified IP address, optionally using the port and protocol specified.
- The `set timezone str=TIMEZONE_SPEC` command sets the RMC to the specified time zone.
- The `set uid` command turns on or turns off the UID locator light on the specified chassis.

search and searchevents

Syntax

```
search logs iel [begin=BEGIN_TIME] [end=END_TIME] [alert=ALERT_THRESHOLD]
[str=SEARCH_STRING] [excl=EXCLUDE_STRING] [pnum=PARTITION_ID]
```

```
searchevents [-B begin_time] [-E end_time] [-e stringA]...[-e stringB] [-v
stringY]...[-v stringZ] [-A alert_threshold] [-p partition] [-x] [-?] [-h]
  -B : beginning of search time          (default: "last boot" associated
with partition number)
  -E : ending of search time             (default: "now")
  -e : search string (case-insensitive)  (default: no search filters)
  -v : exclude string (case-insensitive) (default: no exclude filters)
  -A : alert threshold                   (default: 0)
  -p : partition number                  (default: unspecified)
  -x : eXhaustive search (entire file)   (default: Not exhaustive.
```

Assumes all events are in chronological order.



-? -h : this help

Description

The `search` command examines log files for one or more characteristics that you specify and returns events that match your search criteria.

The `searchevents -h` command shows additional information and examples on how to specify search criteria.

Specifiers

BEGIN_TIME

Specify a time. Use one of the formats that the `searchevents -h` command returns.

END_TIME

Specify a time. Use one of the formats that the `searchevents -h` command returns.

ALERT_THRESHOLD

Specify one of the following alert levels: 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9.

SEARCH_STRING

Specify a string of your choice. The string cannot include space characters.

EXCLUDE_STRING

Specify a string of your choice. The string cannot include space characters.

Example

```
RMC cli> searchevents -h
```

```
Usage: searchevents [-B begin_time] [-E end_time] [-e stringA]...[-e
stringB] [-v stringY]...[-v stringZ] [-A alert_threshold] [-p partition] [-
x] [-?] [-h]
  -B : beginning of search time          (default: "last boot" associated
with partition number)
  -E : ending of search time             (default: "now")
  -e : search string (case-insensitive) (default: no search filters)
  -v : exclude string (case-insensitive) (default: no exclude filters)
  -A : alert threshold                   (default: 0)
  -p : partition number                  (default: unspecified)
  -x : eXhaustive search (entire file)   (default: Not exhaustive.
Assumes all events are in chronological order.
```

-? -h : this help

All times are assumed to be UTC. Unrecognized/extraneous strings may be silently ignored.

Example formats for Begin/End search times:

```
-B/-E "09/05/2017 20:16:15" -> fully specified date/time
-B/-E "2017-09-05 20:16:15Z" -> ISO standard date/time (trailing Z
optional)
-B/-E "09/05      20:16:15" -> if year is omitted, current year will be
used
-B/-E "9/5       20:16"    -> if seconds are omitted, seconds will be :
00
-B/-E 22:00           -> today at the specified time
-B/-E "XX sec ago"   -> current time minus specified quantity
-B/-E "XX min ago"  -> current time minus specified quantity
-B/-E "XX hour ago" -> current time minus specified quantity
-B/-E "XX day ago"  -> current time minus specified quantity
```

```

-B/-E "XX week ago"      -> current time minus specified quantity
-B/-E "XX month ago"    -> current time minus specified quantity
-B/-E "XX year ago"     -> current time minus specified quantity
-B/-E "XX day ago 17:00" -> combine the above with a specific time
-B/-E yesterday         -> prior date at 00:00:00
-B/-E "yesterday 22:00" -> prior date at specified time
-B/-E today             -> today at 00:00:00
-B/-E "today 7:30"     -> today at the specified time
-B/-E now               -> current time
-B/-E epoch             -> 01/01/1970 00:00:00 (search entire log)
-B/-E IEL_EVENT_KEYWORD -> last IEL_EVENT_KEYWORD (case insensitive)
Additional options for Begin search times:
-B "last poweron"      -> last CHASSIS_POWERED_ON event
-B "last efi reset"    -> last 'EFI reset'
-B "last os reboot"    -> last 'OS reboot'
-B "last boot"         -> last 'poweron' OR 'efi reset' OR 'os reboot'
event

```

show

Syntax

```

show all
show analysis id=BUNDLE_ID [summary]
show autopower
show bios [target=TARGET]
show bios revision [target=TARGET]
show bios state [target=TARGET]
show bios uptime [target=TARGET]
show bios validate
show cae [list] [id=EVENT_ID] [severity=SEVERITY num or text]
[pnum=PARTITION_NUM] [offset=OFFSET] [count=NUMBER] [dump=INDEX]
show certificate [chassis=GEOID]
show certificate request [chassis=GEOID]
show certificate ca [file=FILENAME]
show certificate crl [file=FILENAME]
show cert_checking
show chassis
show chassis info [chassis=GEOID]
show chassis list
show complex
show config [list | file=FILENAME]
show date
show debug
show deconfig [chassis=GEOID | pnum=PARTITION_ID]
show dns

```

```
show fabric [verbose]
show failed_login
show firmware [verbose]
show health [chassis=GEOID | pnum=PARTITION_ID | verbose]
show host
show hostname
show hwcfg
show indict [chassis=GEOID | pnum=PARTITION_ID]
show ipmi_watchdog
show ldap
show ldap_group
show ler [enable] [disable]
show livelogs
show location
show logs
show logs audit
show logs console target={BMC_GEOID | PARTITION_ID} [socket={0-3}] [full]
show logs dmesg [target=TARGET]
show logs error [mca | cmc | type=SEVERITY | id=bundle_id |
pnum=PARTITION_ID]
show logs hel [all] [hex] [target=TARGET]
show logs iel [live]
show logs sel [target=TARGET]
show logs support [ieldepth=DAYS]
show logs syshist [verbosity={0,1,2} | raw]
show logs uvdmp [list | id=HEX_ID | file=FILENAME] [verbose]
show network
show npar [pnum=PARTITION_ID] [verbose]
show ntp
show password_default
show power
show power npar pnum=PARTITION_ID
show privileges [role={administrator,monitor,operator}]
show redfish [uri=URI] [walk]
show redfish event_subscription
show remote_log_server
show sdr [target=TARGET]
show sensor [target=TARGET]
```

```
show snmp
show snmp forward_address
show sshkey
show timezone
show uid chassis=GEOID
show user
show user list
show uvdmp [list | id=HEX_ID | file=FILENAME] [verbose]
```

Specifiers

EVENT_ID

Specify one or more event identifiers. Use a comma to separate each *EVENT_ID*. The `show cae` command returns information for each specified *EVENT_ID*.

FILENAME

Specify the name of the file that you want to view.

OFFSET

Specify an offset value. When you specify an *OFFSET* value, the `show cae` command displays events starting from a particular *INDEX* with *OFFSET* value.

NUMBER

Specify an integer number of events that you want the command to display. You can specify a *NUMBER* and an *offset* to scroll through the entire set of events that are available.

INDEX

Specify an *Index* value from the list.

To retrieve a list of index values, enter `show cae list`.

SEVERITY

Specify one of the following:

- 0 or Unknown
- 1 or Other
- 2 or Informational
- 3 or Degraded or Warning
- 4 or Minor Fault
- 5 or Major Alert
- 6 or Critical
- 7 or Fatal

When you specify a *SEVERITY*, the `show cae` command returns information for events at that severity threshold and above.



BUNDLE_ID

Specify the identifier that is associated with the log file that you want to examine. For example:
0x2110000000000001.

PARTITION_ID

Specify the numerical identifier for the partition.

TARGET

Specify the BMC *geoid*.

GEOID

Specify the identifier of the server chassis for which you want to display a unit identification number (UID). The *GEOID* identifies the rack number and the chassis U position within the rack. To retrieve a *GEOID*, enter `show complex` at the RMC prompt.

For example, a *GEOID* of `r001i06b` pertains to the chassis in rack number 1 with U position 6.

Usage

- The `show all` command runs the following commands in sequence (in a noninteractive mode):
 - `show date`
 - `show timezone`
 - `show autopower`
 - `show bios revision`
 - `show bios state`
 - `show bios uptime`
 - `show bios validate`
 - `show cae list`
 - `show certificate`
 - `show certificate ca`
 - `show certificate crl`
 - `show certificate request`
 - `show cert_checking`
 - `show config`
 - `show config list`
 - `show location`
 - `show complex`
 - `show chassis info`
 - `show debug`
 - `show dns`
 - `show fabric verbose`



- show failed_login
- show firmware verbose
- show hwcfg
- show health
- show health verbose
- show host
- show hostname
- show indict
- show ipmi_watchdog
- show ldap
- show ldap_group
- show ler [enable] [disable]
- show logs audit
- show logs dmesg
- show logs hel
- show logs sel
- show logs uvdmp list
- show deconfig
- show network
- show npar
- show chassis list
- show ntp
- show power
- show redfish event_subscription
- show remote_log_server
- show sdr
- show sensor
- show sshkey
- show snmp
- show uid chassis=r001i01b
- show uid chassis=r001i06b



- show user
- show user list

Example1

The following show indict command returns a list of indictments:

```
RMC cli> show indict
r001i01b
    FFFF000101FFFF64 rack1/chassis_u1/harp_fan0
    Indicted: Yes
      2017-10-31T17:05:42Z Reason ID: 302
      Summary: Fan reported to be running slowly or stopped
      Cause: A fan is not running at normal speed. There may be
something obstructing the airflow, a blocked fan or a fan failure.
      Action: Check the enclosure for obstructions to airflow and
ensure the fans are installed properly. If the problem persists, replace
the fan. After replacing the faulty fan, perform an acquittal from the RMC
command line.
```

This example shows that the fan with resource ID rack1/chassis_u1/harp_fan0 in the chassis with GEOID r001i01b in the physical location FFFF000101FFFF64 has been indicted with an error.

Example 2

The following command displays errors in the machine check architecture (MCA) log files and corrected machine check (CMC) log files:

```
show logs error [mca | cmc | type=SEVERITY | id=BUNDLE_ID]
```

The command returns information to the console. Because the amount of returned information can be large, use your terminal emulation program tools to capture this information. Plan to analyze this program on a host computer other than the RMC.

In a CMC dump, you can, for example, find one or more headers that resemble the following:

```
***** CMC (Corrected Machine Check) Error Log ( 0x2110000000000001 )
=====
CPER Record 0x0000000000000000    Timestamp 2017/09/29 19:45:34
=====
```

The 0x2110000000000001 string is the log identifier that you supply as the *bundle_id* to the show logs error command. For example:

```
RMC cli> show logs error cmc
=====
CPER Record 0x0000000000000000    Timestamp 2017/09/29 19:45:34
=====
Signature Start      0x52455043 (REPC)
Revision             0x0101
Signature End        0xffffffff
Section Count        0x0009 (9 Sections)
Error Severity        0x00000002 (Corrected)
Validation Bits       0x00000002
Record Length         0x00000a58 (2648 Bytes)
Timestamp             0x2017092900194534 (2017/09/29 19:45:34)
Platform ID           00000000-0000-0000-0000000000000000
```



```

(Invalid)
Partition ID      00000000-0000-0000-0000000000000000
                  (Invalid)
Creator ID        00000000-0000-0000-0000000000000000
                  (Invalid)
Notification Type 2dce8bb1-bdd7-450e-b9ad9cf4ebd4f890
                  (CMC)
Record ID         0x00000000000000000000
Flags             0x00000000
Persistence Information 0x00000000000000000000

```

```

=====
CPER Section Descriptor - number 0x0001 (1)
=====

```

```

Section Offset      0x00000308 (776 Bytes)
Section Length      0x00000048 (72 Bytes)
.
.
.

```

Example 3

The following command returns network information:

```

RMC cli> show network
-- Network Configuration --

IPv4 Addressing: dhcp
IPv4 Address: 192.168.1.1
IPv4 Netmask: 255.255.255.0
IPv4 Gateway: 137.38.91.1
Hostname: uv4test14-rcm
IPv6 Addressing: dhcp
IPv6 Address: 1111:1111:1111:11:111:1111:1111:1111/64 Scope:Global
IPv6 Address: 1111:1111:1111:11:1::1111/64 Scope:Global
IPv6 Address: 1111::111:1111:1111:1111/64 Scope:Link

-- AUX Network Configuration --

IPv4 Addressing: static
IPv4 Address: 192.168.1.2
IPv4 Netmask: 255.255.0.0
IPv4 Gateway:
IPv6 Addressing: default

-- Network Information --

eth0      Link encap:Ethernet  HWaddr 00:00:00:00:00:00
          inet addr:192.168.1.3  Bcast:192.168.255.255  Mask:255.255.0.0
          inet6 addr: 1234::567:89ab:cdef:0123/64 Scope:Link
          UP BROADCAST RUNNING ALLMULTI MULTICAST MTU:1500 Metric:1
          RX packets:2209650 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1950104 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:414470219 (395.2 MiB)  TX bytes:233440163 (222.6 MiB)
          Base address:0xc000

eth1      Link encap:Ethernet  HWaddr 01:23:45:67:89:ab
          inet addr:192.168.1.3  Bcast:192.168.1.4  Mask:255.255.255.0
          inet6 addr: abcd:ef0:1234:56:789:abcd:ef0:1234/64 Scope:Global
          inet6 addr: 0123::456:789a:bcde:f012/64 Scope:Link
          inet6 addr: 1234:5678:9abc:de:f::0123/64 Scope:Global
          UP BROADCAST RUNNING ALLMULTI MULTICAST MTU:1500 Metric:1
          RX packets:1610866 errors:0 dropped:43 overruns:0 frame:0

```



```

TX packets:757013 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:249314737 (237.7 MiB) TX bytes:134176764 (127.9 MiB)
Base address:0xe000

-- BaseIO --

P000 [r001i01b]: 192.168.1.5 [ab:cd:ef:01:23:45]
                2345::678:9abc:def0:1234/64 Scope:Link
                bcde:f0123:4567:89:a::bcd/64 Scope:Global
                cdef:0123:4567:89:abc:def0:1234:5678/64 Scope:Global

Network Configuration:
==== r001i01b ====
IPv4 Addressing: dhcp

-- Network protocols --

HTTP:      Enabled
HTTPS:     Enabled
IPMI:      Enabled
IPV6:      Enabled
SSH:       Enabled

-- Internal Network Configuration --

Internal BMC Subnet: 172.30.50.0
Internal CMC Subnet: 172.16.0.0
Internal eRMC Subnet: 172.30.60.0

```

Example 4

The following command returns remote log server information:

```

RMC cli> show remote_log_server

-- Remote log server Information --

Server: 10.11.12.135
Port: 514
Protocol: udp

```

test

Syntax

```

test
test fabric
test ldap [username=USERNAME] [verbose]
test remote_log

```

Description

Test the complex fabric, LDAP configurations, or remote log server settings.



Specifiers

USERNAME

Specify the LDAP username you want to test.

Description

Tests the fabric or LDAP configuration. The `test fabric` command was introduced in version 2.5.246.

Usage

All chassis in the complex must be powered off to run the `test fabric` command.

Example output

```
cli> test fabric
Checking power state
Setting up test
Beginning test

nl8cfg - UV4 NUMALink (NL8) Configuration Verification 1.2.2
-> Read actual cable system configuration from hardware: localhost
-> Open connection to the UV SMN/CMC: host=localhost port=1
-> Read NL8 cable serial numbers from each chassis
-> Build the NL8 cable configuration from the cable definitions
-> Close connection to the UV SMN/CMC: host=localhost port=1
-> Write system configuration to file: /tmp/uv4_sys.scf
Number of partitions = 1
    partition 016 contains 2 chassis, common fabric = YES
Multiple fabric zones = NO
    fabric zone 0 contains [ r001i01b, r001i06b ]
-> Determine the expected configuration for the system containing:
    Selected chassis = 2      (2 total, with 2 base I/Os)
    Number NL cables = 22    (multiple zones = NO)
-> Read expect system cable configuration from file: /usr/local/sysco/data/scan/scf/uv4_001_lr2c.scf
-> Compare system cable configurations: FILE1 and HARDWARE
    EXPECT FILE1 = /usr/local/sysco/data/scan/scf/uv4_001_lr2c.scf
    ACTUAL HARDWARE = localhost
-> Adjusting chassis slot positions in expected configuration file to match actual configuration:
    Expect configuration chassis adjusted from r001i06 to r001i01
    Expect configuration chassis adjusted from r001i11 to r001i06
-> Result: FILE1 and HARDWARE = 0 system configuration differences
==== r001i01b ====
==== r001i06b ====
Cleaning up test
Test is complete
```

```
cli> test ldap
```

```
Beginning test...
```

```
Ping Test
```

```
-----
```

```
# ping -c2 fddd:6777:123:234:aabb:5678:ccdd:cac4
# ping6 -c2 fddd:6777:123:234:aabb:5678:ccdd:cac4
```

```
Ping Test was successful.
```

```
Group Search Test
```

```
-----
```

```
Attempting LDAP search for group3 in cn=Users,dc=ad1,dc=dmn01 DN scope...
```

```
# LDAPTLS_REQCERT=allow
# ldapsearch -LLL -v -x -l 90 -d 4 -H ldaps://[fddd:6777:123:234:aabb:5678:ccdd:cac4]:636 -b "cn=Users,dc=ad1,dc=dmn01" -D "cn=ldapAdmin,cn=Users,dc=ad1,dc=dmn01" -W "(&(objectClass=group))"
```

```

(cn=group3))" cn
Entry for group3 was found.
Attempting LDAP search for group2 in cn=Users,dc=ad1,dc=dmn01 DN scope...
# LDAPTLS_REQCERT=allow
# ldapsearch -LLL -v -x -l 90 -d 4 -H ldaps://[fddd:6777:123:234:aabb:
5678:ccdd:cac4]:636 -b "cn=Users,dc=ad1,dc=dmn01" -D
"cn=ldapAdmin,cn=Users,dc=ad1,dc=dmn01" -W "(&(objectClass=group)
(cn=group2))" cn
Entry for group2 was found.
Attempting LDAP search for group1 in cn=Users,dc=ad1,dc=dmn01 DN scope...
# LDAPTLS_REQCERT=allow
# ldapsearch -LLL -v -x -l 90 -d 4 -H ldaps://[fddd:6777:123:234:aabb:
5678:ccdd:cac4]:636 -b "cn=Users,dc=ad1,dc=dmn01" -D
"cn=ldapAdmin,cn=Users,dc=ad1,dc=dmn01" -W "(&(objectClass=group)
(cn=group1))" cn
Entry for group1 was found.
Group Search Test was successful.

Domain SID Test
-----
Attempting LDAP search for group3 in cn=Users,dc=ad1,dc=dmn01 DN scope...
# LDAPTLS_REQCERT=allow
# ldapsearch -LLL -v -x -l 90 -d 4 -H ldaps://[fddd:6777:123:234:aabb:
5678:ccdd:cac4]:636 -b "cn=Users,dc=ad1,dc=dmn01" -D
"cn=ldapAdmin,cn=Users,dc=ad1,dc=dmn01" -W "(&(objectClass=group)
(cn=group3))" objectsid cn

    decoded Object SID: S-1-5-21-652678-1122334455-1122334455-1111
settings Domain SID: S-1-5-21-652678-1122334455-1122334455

Domain SID Test was successful.

INFO: Optional USERNAME not provided for test. Skipping User Search test.

LDAP Client Test
-----
# nslcd --check
NSLCD is running.

LDAP Client Test was successful.

Test is complete.

cli> test remote_log

ping test:
-----
# ping -c2 10.100.10.100
Ping Test was successful.

```



A test log message has been sent to the system log.
Verify remote server for test message.

update firmware

Syntax

```
update firmware url=PATH_TO_FIRMWARE [dry_run] [reinstall] [exclude_npar_fw]
```

Description

Updates the system firmware.

Specifiers

PATH_TO_FIRMWARE

Specify the path to the firmware. This command assumes the following:

- You downloaded the firmware update bundle from the HPE Support Center.
- The bundle resides on a local computer at your site.

Accepted access methods are `https`, `sftp`, and `scp`.

The firmware also accepts nonstandard ports for `https`, `sftp`, and `scp`.

For example:

```
scp://<user>@<host>:<port>/<path to FW>
```

Options

dry_run

Specify if you want to test the firmware installation.

reinstall

Specify if you want to force an installation when the firmware already matches the complex firmware. The complex firmware includes all firmware on the system except the UEFI drivers for IO. You may update the drivers either through the OS online, or offline from a DVD-booted version using the Smart Update Manager tool.

[exclude_npar_fw]

Specify if you do not want to update any firmware running on an nPartition, such as BIOS.

upload

Syntax

```
upload config (url=URL | file=FILENAME)
```

```
upload debug url=URL [file=FILENAME]
```

```
upload npar config (url=URL | file=FILENAME)
```

```
upload syshist url=URL
```

```
upload uvdmp [id=HEX_ID | file=FILENAME] url=URL
```



Description

Transfers a `tar` file that contains files to a different, local computer at your site. The command compresses the `tar` file automatically.

Specifiers

URL

Specify the full path to the target location on a local computer. Accepted transfer methods are `sftp` and `scp`.

The following examples show the supported formats:

- `sftp://user@localhost/path`
- `scp://user@localhost/path`

The command prompts you for the password to the local host computer.

FILENAME

Specify the name of the file. If you do not specify a `FILENAME` name, the system creates and then uploads a debugging or configuration file.

Usage

The `upload config (url=URL | file=FILENAME)` command uploads the configuration file to the specified URL.

The `upload debug url=URL [file=FILENAME]` command uploads the debugging file to the specified URL.

The `upload npar config (url=URL | file=FILENAME)` command uploads the nPartition configuration file to the specified URL.

The `upload syshist url=URL` command uploads the system history file to the specified URL.

The `upload uvdmp [id=HEX_ID | file=FILENAME] url=URL` command uploads a dump file to the specified URL.

version

Syntax

```
version
```

Description

Displays the eRMC/RMC software version number.

Example output

```
RMC cli> version
HPE Superdome Flex RMC, Rev. 2.30.330 [Bootloader 1.1.4]
```

```
eRMC cli> version
HPE Superdome Flex BMC, Firmware Rev. 2.30.330
```



Partition boot options

The following boot options are available:

- *None* - No boot option specified. Boots from default source.
- *no_oprom* - One time boot option that causes UEFI to boot without loading the default option ROMs.
- *BiosSetup* - Boot to BIOS setup.
- *Cd* - Boot from existing UEFI boot option entries that correspond to CD/DVD drives of any connection type (such as SATA and USB).
- *Hdd* - Boot from existing UEFI boot option entries that correspond to local hard disk drives, excluding USB drives.
- *Pxe* - Boot from existing UEFI boot option entries that correspond to PXE.
- *RemoteDrive* - Boot from existing UEFI boot option entries that correspond to remote (FibreChannel or iSCSI) hard disk drives.
- *SDCard* - Boot from existing UEFI boot option entries that correspond to SD cards.
- *Shell* or *UefiShell* - Boot to UEFI Shell.
- *UefiHttp* - Boot from existing UEFI boot option entries that correspond to HTTP boot.
- *Usb* - Boot from existing UEFI boot option entries that correspond to USB disk drives.



Managing with HPE OneView

HPE OneView 4.2 or later supports monitoring Superdome Flex Server hardware inventory and basic health monitoring. HPE OneView 5.0 or later supports monitoring feature, managing Superdome Flex Server hardware inventory, and basic maintenance functions.

HPE OneView user guides, support matrixes that include supported features, and other manuals are available on the [Enterprise Information Library](#).

To view help in the HPE OneView appliance, click .

Rack manager components

A rack manager platform consists of the following:

- **One or more chassis** — The chassis are individual nodes that are used to build logical components such as systems or nPartitions. Each chassis is a physical container for system components like the compute, storage, or network nodes.

Chassis can be three types:

- Base — Available with every rack manager; includes boot support.
- Expansion — Added to a system with a base chassis to scale up the partition; does not have boot options.
- Partition — Added to a system with a base chassis to scale up the partition or to create additional partitions.

In an HPE Superdome Flex Server and an HPE Superdome Flex 280 Server, the chassis are 5U nodes that represent the compute nodes with CPU, memory, I/O, power supplies, and fans.

NOTE: HPE Superdome Flex 280 Server has only base and expansion chassis.

- **One or more systems or nPartitions** — The system is a logical entity that functions as a server, and runs a single instance of an operating system to host services. In HPE OneView, an nPartition is modeled as a server hardware resource.

In an HPE Superdome Flex Server, the systems are also called nPartitions. An nPartition consists of one or more compute chassis that are interconnected through an ultra-fast fabric. Each nPartition starts with a monarch chassis of the type base chassis or partition chassis, and can be extended with additional chassis to form a larger system.

Every HPE Superdome Flex 280 Server has a base chassis providing BaseIO, management interfaces, and boot support. One expansion chassis can be added to expand the system to eight processor sockets. The required base chassis provides platform management through an embedded Rack Management Controller (eRMC).

- **One or more managers** — The manager component manages the rack manager platform, and hosts the management controller function. HPE OneView communicates with the rack manager through this management controller enabling system administration, control, and platform management of the rack manager platform. Some platforms have multiple manager components to support redundancy.

In an HPE Superdome Flex Server, the manager component is called the Rack Management Controller (RMC). The RMC hosts the DMTF Redfish APIs.

Managers can be one of two types:

- Embedded — The manager is running within the management controller of the base chassis. In an HPE Superdome Flex Server and HPE Superdome Flex 280 Server, the embedded Rack Management Controller (eRMC) can manage one or two chassis (four or eight processor sockets).



NOTE: HPE Superdome Flex 280 Server is managed by eRMC only.

- **Standalone** — The manager is standalone and placed in a unit location within the rack. In an HPE Superdome Flex Server, the standalone RMC can manage the maximum available supported chassis configuration.
- **One or more racks** — The racks physically contain all the rack manager platform components.

Add a rack manager

You can add rack managers as managed or monitored. The HPE Superdome Flex Server is added into HPE OneView as a monitored rack manager.

When you add a rack manager, HPE OneView automatically applies either the HPE OneView Standard license or the HPE OneView Advanced license to the rack manager based on the firmware version installed. HPE OneView discovers and monitors the hardware subcomponents of the rack manager.

HPE OneView supports adding rack managers with IPv4, IPv6, or both IPv4 and IPv6 addresses depending on the network configuration in HPE OneView. You can add the rack manager using the IPv4 address, IPv6 address or with the hostname.

When you add a rack manager, you can select the already uploaded firmware bundle from the firmware baseline option. Once selected, HPE OneView updates the server with the firmware that matches the specified baseline. The baseline is set for each of the rack manager platform components such as chassis, systems (nPartitions), and Rack Management Controller (RMC). You can also choose to manage the firmware update of the rack manager platform components manually.

Prerequisites

Privileges: Infrastructure administrator or Server administrator.

Procedure

1. From the main menu, select **SERVERS > Rack Managers**, and do one of the following:
 - Click **+ Add rack manager** in the master pane.
 - Select **Actions > Add**.
2. Enter the **IP address or host name**.
3. Enter the **Credentials**.

NOTE: Users with **Create** permissions that are restricted by scope must assign resources to one or more restrictive scopes to be able to manage a resource. Users with **Create** permissions that are not restricted by scope are not required to assign resources to a scope to be able to manage a resource. Users with no scope restrictions can assign any resource to any scope.

4. Select the **Firmware baseline** options.
5. Click **Add** to add the rack manager or click **Add+** to add more rack managers.
6. Verify that the rack manager has been added by confirming it is listed in the master pane.

More information



- **About scopes**
- **Enable remote support for a rack manager**
- **Edit remote support**
- **Firmware options when adding rack manager**
- **Adding and monitoring rack managers**
- **License types: Server hardware**
- **Rack manager is unreachable**
- **Rack Managers**
- **Superdome Flex firmware bundle**



Remote management with Redfish

The Redfish standard integrates easily with commonly used tools by specifying a RESTful interface. Redfish uses JSON and OData constructs to facilitate scripting and organization of server management data. Redfish is the main API to manage the HPE Superdome Flex Server through the RMC. Tools like OpenStack and HPE OneView use it to get information about the system, control it, and get Redfish events, and alerts to monitor the system. These tasks can be done as well, using `Redfishtool` or your own scripting methods.

To learn more about Redfish, start at the [Redfish Standard Site](#).

NOTE: HPE OneView version 4.1 or later is supported (and recommended) for monitoring the Superdome Flex Server system.

As an alternative to the procedure in this topic, use the `ssh` command to connect with an RMC. Log in as the administrator user, or any other user you have configured.

Redfish can be used without any particular tool. Simple `curl` commands, or `perl`, `python`, or PowerShell scripting that sends the right JSON data over `https` works well. However, the Distributed Management Task Force (DMTF) has built a specific scripting tool for interacting with systems that support the standard, called `Redfishtool`. `Redfishtool` is able to walk the tree, and complete common tasks. These tasks include gathering the health of objects within the system, and performing actions.

For more information about `Redfishtool`, see <https://github.com/DMTF/Redfishtool>.

You can use the `Redfishtool` commands remotely from a laptop or other computer that is on the same local network as the rack management controller (RMC). The `Redfishtool` commands rely on Python 3.4, which is commonly included in Linux distributions. Python 3.4 can be downloaded for Windows and Linux from <https://www.python.org/>. You can manage a system in the following ways:

The remote management commands use the following terminology:

- The term **HPE Superdome Flex Server system** includes the server, RMC, and the board management controllers (BMCs).
- The term **Superdome Flex Server** refers only to the server component of the system. The server is enclosed within each chassis. The RMC and the BMCs are not part of the server itself.

The RMCs and BMCs have no power on or power off capability. When power is applied to the Superdome Flex Server, the RMC has power, and the BMCs have power. You can issue `Redfishtool` commands to the RMC and to the BMCs as long as power is applied to the system.

NOTE: This documentation shows `Redfishtool` command output. Depending on the revision level of the `Redfishtool` command you use on your remote console, the command output you see can differ from the output shown in this documentation.

Out-of-Band management with Redfishtool - general procedure

The `Redfishtool` is a Redfish client that has to be installed on a computer with network connectivity to the RMC. The tool then can be used to power on or off the HPE Superdome Flex Server and obtain status information.



Procedure

1. Ensure that the computer running the Redfishtool has `https` connectivity to the managed RMC.
2. Ensure that Python 3.4 or later is installed, and working correctly on the workstation running Redfishtool.
3. Follow the instructions for downloading all the files for Redfishtool from <https://github.com/DMTF/Redfishtool>.

NOTE: The implementation of Redfish on the server only uses `https`. One command argument to include, that is not in all the DMTF Redfishtool examples (but is a valid option), is `-S Always` to always use `https`.

Use the following command to get the `root` of the Redfish tree with RMC target `hostname=flex_rmc`, and the default administrator account and password:

```
$> python3 ~/Redfishtool/redfishtool.py -r flex_rmc -u <RMC login> -p
<Password> -S Always root
```

Where `RMC login` is the user name, for example `administrator`, and `Password` is the password for that account. For the rest of this chapter, we will use `administrator` and `DummyPwd` for the user password.

! **IMPORTANT:** Use your own unique password on your RMC for security purposes.

In some network environments, you may need to override a proxy to make things work correctly.

For example, from a Linux workstation the following command overrides once the proxy server and then retrieves the root properties of the remote RMC:

```
$> https_proxy='' python3 ~/Redfishtool/redfishtool.py -r flex_rmc -u
administrator -p DummyPwd -S Always root
```

You can override the proxy server during an entire interactive session with:

```
$> export https_proxy=''
```

Some subsystems have built in help in the `examples` command. In this Bash shell example, the following command displays all the subcommands available to get or set the systems subsystem properties:

```
$> python3 redfishtool.py -r flex_rmc -u administrator -p DummyPwd -S
Always Systems examples
redfishtool -r<ip> Systems                               # shows the systems collection
redfishtool -r<ip> Systems list                          # lists Id, Uri, AssetTag for all systems
redfishtool -r<ip> Systems -I <id>                      # gets the system with Id=<d>
redfishtool -r<ip> Systems -M AssetTag:12345            # gets the system with AssetTag=12345
redfishtool -r<ip> Systems -L <sysUrl>                 # gets the system at URI=<systemUrl>
redfishtool -r<ip> Systems -F                           # get the First system returned (for debug)
redfishtool -r<ip> Systems -l                           # get the first system and verify that there
is only one system
redfishtool -r<ip> Systems -I <id> patch {A: B,C: D,...} # patch the json-formatted {prop: value...}
data to the object
redfishtool -r<ip> Systems -I <id> reset <resetType>    # reset a system. <resetType>=the redfish-defined
values: On, Off, gracefulOff...
redfishtool -r<ip> Systems -I <id> setAssetTag <assetTag> # set the system's asset tag
redfishtool -r<ip> Systems -I <id> setIndicatorLed <state> # set the indicator LED.
<state>=redfish defined values: Off, Lit, Blinking
redfishtool -r<ip> Systems -I <id> setBootOverride <enabledVal> <targetVal> #-- set Boot Override properties.
<enabledVal>=Disabled|Once|Continuous
redfishtool -r<ip> Systems -I<Id> Processors            # get the processors Collection
redfishtool -r<ip> Systems -I<Id> Processors list       # lists Id, Uri, and Socket for all
processors in system with Id=<Id>
redfishtool -r<ip> Systems -I<Id> Processors -i 1      # get the processor with id=1 in
system with Id=<Id>
redfishtool -r<ip> Systems -L <sysUrl> Processors -m Socket:CPU_1 # get processor with property
Socket=CPU_1, on system at url <sysUrl>
```

To list the partitions:

```
$> python3 redfishtool.py -r flex_rmc -u administrator -p DummyPwd -S Always Systems
{
  "Name": "Computer System Collection",
  "@odata.id": "/redfish/v1/Systems",
  "@odata.context": "/redfish/v1/$metadata#ComputerSystemCollection.ComputerSystemCollection",
  "@odata.type": "#ComputerSystemCollection.ComputerSystemCollection",
  "Members": [
    {
      "@odata.id": "/redfish/v1/Systems/Partition0"
    }
  ],
  "Members@odata.count": 1
}
```

To get the top-level model for a partition, in this case number 0:

```
$> python3 redfishtool.py -r flex_rmc -u administrator -p DummyPwd -S Always Systems -I Partition0
{
  "Name": "Npar 0",
  "PowerState": "Off",
  "PCIeFunctions": [],
  "PCIeDevices@odata.count": 2,
  "HostName": null,
  "@odata.context": "/redfish/v1/$metadata#ComputerSystem.ComputerSystem",
  "SecureBoot": {
    "@odata.id": "/redfish/v1/Systems/Partition0/SecureBoot"
  },
  "Bios": {
    "@odata.id": "/redfish/v1/Systems/Partition0/Bios"
  },
  "Processors": {
    "@odata.id": "/redfish/v1/Systems/Partition0/Processors"
  },
  "PCIeFunctions@odata.count": 0,
  "@odata.type": "#ComputerSystem.v1_5_0.ComputerSystem",
  "Oem": {
    "Hpe": {
      "ProductId": "1590PID03020102",
      "@odata.type": "#HpeNpar.v1_0_0.HpeNpar",
      "OSType": null,
      "OSVersion": null,
      "IPv4Addresses": []
    }
  },
  "Links": {
    "ResourceBlocks": [
      {
        "@odata.id": "/redfish/v1/CompositionService/ResourceBlocks/r001i01b"
      },
      {
        "@odata.id": "/redfish/v1/CompositionService/ResourceBlocks/r001i06b"
      }
    ],
    "ResourceBlocks@odata.count": 2,
    "ManagedBy": [
      {
        "@odata.id": "/redfish/v1/Managers/RMC"
      }
    ],
    "Chassis": [
      {
        "@odata.id": "/redfish/v1/Chassis/r001i01b"
      },
      {
        "@odata.id": "/redfish/v1/Chassis/r001i06b"
      }
    ],
    "Chassis@odata.count": 2
  },
}
```



```

"NetworkInterfaces": {
  "@odata.id": "/redfish/v1/Systems/Partition0/NetworkInterfaces"
},
"EthernetInterfaces": {
  "@odata.id": "/redfish/v1/Systems/Partition0/EthernetInterfaces"
},
"Status": {
  "HealthRollup": "OK",
  "State": "Enabled",
  "Health": "OK"
},
"SystemType": "PhysicallyPartitioned",
"@odata.etag": "\"a71e982b056b317ea3ab332038f8de24\"",
"MemorySummary": {
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "TotalSystemMemoryGiB": 6143,
  "TotalSystemPersistentMemoryGiB": 0
},
"@odata.id": "/redfish/v1/Systems/Partition0",
"Storage": {
  "@odata.id": "/redfish/v1/Systems/Partition0/Storage"
},
"ProcessorSummary": {
  "LogicalProcessorCount": 144,
  "Count": 8,
  "Model": "Intel(R) Xeon(R) Gold 6254 Processor",
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  }
},
"Id": "Partition0",
"Memory": {
  "@odata.id": "/redfish/v1/Systems/Partition0/Memory"
},
"UUID": "e3cc8eeb-e4fc-5c4b-af8f-c25d454ed9b9",
"Actions": {
  "#ComputerSystem.Reset": {
    "ResetType@Redfish.AllowableValues": [
      "On",
      "ForceOff",
      "ForceRestart",
      "GracefulRestart",
      "GracefulShutdown",
      "PowerCycle"
    ],
    "target": "/redfish/v1/Systems/Partition0/Actions/ComputerSystem.Reset"
  }
},
"PCIeDevices": [
  {
    "@odata.id": "/redfish/v1/Chassis/r001i01b/PCIeDevice2"
  },
  {
    "@odata.id": "/redfish/v1/Chassis/r001i01b/PCIeDevice3"
  }
],
"Model": "Superdome Flex",
"BiosVersion": "3.30.140",
"Boot": {
  "BootSourceOverrideTarget@Redfish.AllowableValues": [
    "None",
    "BiosSetup",
    "Cd",
    "Hdd",
    "Pxe",
    "RemoteDrive",
    "SDCard",

```



```

        "UefiHttp",
        "UefiShell",
        "Usb"
    ],
    "BootSourceOverrideMode@Redfish.AllowableValues": [
        "UEFI"
    ],
    "BootSourceOverrideTarget": "None",
    "BootSourceOverrideEnabled@Redfish.AllowableValues": [
        "Continuous",
        "Disabled",
        "Once"
    ],
    "BootSourceOverrideEnabled": "Disabled",
    "BootSourceOverrideMode": "UEFI"
}
}
}

```

Or the health of that partition, note it is a property in the preceding list:

```

$> python3 redfishtool.py -r flex_rmc -u administrator -p DummyPwd -S Always Systems -I Partition0 -P Status
{
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  }
}

```

The tool also includes the ability to set up alert monitoring, and can walk the Redfish tree to pull all the health status.

For more information, see:

- **[Redfish Standard Site](#)** – hosted by the Distributed Management Task Force
- **[Redfish - Simple, Modern and Secure Management for Multi-Vendor Cloud and Web-Based Infrastructures](#)** – Technical note
- **[DSP2044: Redfish White Paper](https://www.dmtf.org/standards/redfish)** – more white papers available at: <https://www.dmtf.org/standards/redfish>



Management with `ipmitool`

This section describes HPE Superdome Flex Server support for the Intelligent Platform Management Interface (IPMI) protocol. HPE Superdome Flex Server also supports the Redfish manageability standard for out-of-band management. Redfish is a newer, modern method for interfacing with management processors, using RESTful HTTPS and JSON protocols for greater flexibility and security.

NOTE: IPMI is disabled by default and must be enabled in the rack management controller (RMC) by using the `enable ipmi` command. IPMI is also disabled after performing firmware updates or after using the `reset factory` command.

NOTE: IPMI is only supported for use to remotely interact and control the system through the RMC when all the chassis' in the system are included in one large nPartition. Any other configuration is outside the scope of the IPMI standard.

Therefore if a partition definition exists that does not contain all the chassis in the complex, or if two or more partitions are defined, use Redfish to control and manage it instead of IPMI. Redfish is supported for operation in all server configurations.

For details about the industry standard release of Redfish, see <https://redfish.dmtf.org>. To download Redfishtool, go to <https://github.com/DMTF/Redfishtool>.

Redfish can be used for gathering status, inventory, and health information. Redfish can also be used to subscribe to alerts, controlling power on the server. Redfish is the interface used by software such as HPE OneView.

Redfish can also be used directly without Redfishtool, using simple Python, Perl, or other scripts. Redfish also supports properly formed Linux URL commands and some browser-based JSON web add-ons. Redfishtool is designed for people to get information and manage systems explicitly using Redfish, rather than generically using HTTPS and JSON.

The HPE Superdome Flex Server supports the IPMI. You can access the IPMI v1.5/v2.0 interface through the `ipmi` command. IPMI management software is included in most Linux distributions and in IPMI-compliant management applications. For more information about the `ipmitool` command, see the <http://ipmitool.sourceforge.net/>.

NOTE: IPMI is not recommended for interactive sessions.

You can use the `ipmitool` commands remotely from a laptop or other computer that is on the same local network as the RMC. You can manage a system in the following ways:

The remote management commands use the following terminology:

- The term **HPE Superdome Flex Server system** includes the server, RMC, and the Baseboard Management Controllers.
- The term **Superdome Flex Server** refers only to the server component of the system. The server is enclosed within each chassis. The RMC and the BMCs are not part of the server itself.

More information

[Remote management with Redfish](#)

Retrieving `ipmi` help output

The following procedure explains how to retrieve `ipmi` help output.



Procedure

1. Log into the rack management controller (RMC) as the `administrator` user.
2. Enter the following command:

```
RMC cli> ipmi command=help
```
3. Press `q` to exit the help display.

Running IPMI commands on the rack management controller (RMC) - general procedure

The RMC supports a limited set of IPMI commands that you can use to power on and power off the system, or obtain RMC status information. When you run an IPMI command on the RMC, it returns information about the RMC only.

Use the `ssh` command to connect with an RMC, and log in as the `administrator` user. Enter your `ipmi` commands.

```
RMC cli> bmc exec "ipmi chassis power status" bmc_id
```

Out-of-band management with ipmitool

Procedure

1. Make sure that your local computer is on the same network as the RMC.
2. Enter an `ipmitool` command in the following format.

```
local_computer% ipmitool -I lanplus -H host -U admin -P password cmd
```

The parameters are as follows:

- For `host`, specify the hostname (preferred) or IP address of the RMC you want to target.
- The `admin` user name.
- For `password`, specify the RMC `ipmi` password.

Use the password that is in effect at this time. The initial installation instructions that explained how to attach the system to your site network directed you to change the factory default `ipmi` password. If you changed the `ipmi` password at that time, provide the new `ipmi` password.

If you did not specify a site-specific `ipmi` password, use the factory-default RMC password. The factory-default RMC password appears on the password sticker. The password sticker is a bar coded sticker that appears on your system.

For more information about the placement of the password sticker, see the following:

HPE Superdome Flex Server User Guide

- For `cmd`, specify the specific command you want to run.



Running an `ipmi` command against a baseboard management controller (BMC) - general procedure

Procedure

1. Use the `ssh` command to connect to the rack management controller (RMC).

Log into the RMC as the `administrator` user.

This is necessary if you want to run `ipmi` commands on the system BMCs.

2. Specify the command that targets a specific BMC.

To target a BMC, provide BMC credentials on the `ipmi` command. The user name is `admin`.

Use the password that is in effect at this time. The initial installation instructions that explained how to attach the system to your site network directed you to change the factory default `ipmi` password. If you changed the `ipmi` password at that time, provide the new `ipmi` password.

If you did not specify a site-specific `ipmi` password, use the factory-default RMC password. The factory-default RMC password appears on the password sticker. The password sticker is a bar coded sticker that appears on your system.

For more information about the placement of the password sticker, see the following:

[HPE Superdome Flex Server User Guide](#)

Running an `ipmi` command against one BMC

The following procedure explains how to log into an RMC and then target an `ipmi` command to one BMC.

NOTE: The user name is different for RMC and eRMC. For example, **ADMIN** and **admin** respectively for `ipmitool`. eRMC and RMC have the same set of commands for `ipmitool`.

To run an `ipmi` command on one BMC

Procedure

1. To log into the RMC, use the `ssh` command in the following format.

```
remotehost% ssh administrator@host
```

For `host`, specify the hostname (preferred) or IP address of the RMC.

For example:

```
remotehost% ssh administrator@flex-rmc
```

2. Enter the RMC password when prompted.
3. Use the following command to target the BMC.

```
RMC cli> bmc "ipmi cmd" bmc_id
```

The variables in this command are as follows.



- For *cmd*, specify one of the `ipmitool` commands. You cannot specify the `exec` or `shell` commands, but all other commands are accepted.
- For *bmc_id*, specify the ID of a specific BMC.

Running an `ipmi` command against all BMCs

The following procedure explains how to log into an RMC and then target an `ipmi` command to all BMCs.

To run an `ipmi` command on all BMCs

Procedure

1. Use the `ssh` command in the following format to log into the RMC.

```
remotehost% ssh administrator@host
```

For *host*, specify the hostname (preferred) or IP address of the RMC.

For example

```
remotehost% ssh administrator@flex-rmc
```

2. Enter the RMC password when prompted.
3. Use the following command to target the BMCs.

```
RMC cli> bmc ipmi cmd
```

For *cmd*, specify one of the `ipmi` commands. You cannot specify the `exec` or `shell` commands, but all other commands are accepted.

Retrieving RMC or BMC power status information

Assume the following:

- You are logged into a Linux computer as a regular user, not the root user or the administrator user.
- The hostname of the RMC is `flex-rmc`.
- The chassis BMC IDs are `r001i01b`, `r001i06b`, `r001i11b`, and `r001i02b`.

Example 1: You can use the following command to obtain information about the power status on the RMC.

```
remotehost% ipmitool -I lanplus -H flex-rmc -U ADMIN -P PASSWORD \
power status
Chassis Power is on
```

In this example, assume that the user already added a user with the name `administrator` to IPMI user.

Example 2: You can use the following commands to obtain information about the power status on chassis BMC `r001i01b`.

```
remotehost% ssh administrator@flex-rmc
password:#####
RMC cli> bmc exec "ipmi power status" r1i1b
Chassis Power is on
```



In this example, assume that the user supplied the `administrator` password for the RMC when prompted.

Issuing power cycle commands

The following command sends power control requests to the entire HPE Superdome Flex Server system.

```
remotehost% ipmitool -I lanplus -H host -U ADMIN -P password \  
chassis power action
```

For `host`, specify the hostname (preferred) or IP address of the RMC.

For `action`, specify `status`, `on`, `off`, `cycle`, `reset`, `diag`, or `soft`.

If you want to watch the server boot process, use the `ipmitool sol activate` command to open up a serial console. For information about the `sol activate` command.

NOTE: The `ipmitool power` command in this topic addresses the entire server, not any of the individual chassis. Hewlett Packard Enterprise recommends that you do not attempt to issue power commands to any individual server chassis.

The RMC `power on` command is an alternative to the `ipmitool power` command. For more information about the RMC `power` commands, see:

Power command

The `ipmitool power` command does not work when there are two or more partitions.

More information

[Opening a serial console](#)

Sending a nonmaskable interrupt (NMI) signal to write a crash dump

You can send a nonmaskable interrupt (NMI) signal to the entire HPE Superdome Flex Server system, which responds by writing a system crash dump.

```
remotehost# ipmitool -I lanplus -H flex-rmc -U ADMIN -P password \  
chassis power diag
```

Listing information for field replaceable units (FRUs)

The `fru list` command lists information about each FRU.

NOTE: The remote host will not be running the same version of `ipmitool` as runs on the server RMC, and may display less FRU information. For additional details, such as network card addresses, run `ipmitool` on the RMC.

The following command lists FRUs for an RMC.

```
remotehost% ipmitool -I lanplus -H localhost -U ADMIN -P password fru list  
FRU Device Description : Builtin FRU Device (ID 0)  
Board Mfg Date       : Wed Dec 28 00:37:00 2016  
Board Mfg            : BENCHMARK  
Board Product       : UV3_RMC  
Board Serial        : SAX201  
Board Part Number   : 030_2940_003  
Board Extra         : _A01  
Board Extra         : 00009031  
Board Extra         : 00000002
```



```

Board Extra      : 0000468c
Product Manufacturer : HPE
Product Name     : Rack Management Controller
Product Part Number : Q2N07A
Product Version  : 00
Product Serial    : 5UF7052300
OEM (Hewlett Packard Enterprise) MAC Address Record
  Ethernet 1     : 08:00:69:11:22:33
  Ethernet 2     : 08:00:69:11:22:34

FRU Device Description : UV300 Pwr Bkpln (ID 1)
Product Manufacturer : 3Y POWER
Product Name       : URP1X151AH
Product Part Number : YH-5151E
Product Version    : A04R
Product Serial     : SA080N301617000163
Product Asset Tag  : 10050d
Product Extra      : A08

FRU Device Description : UV300 PS0 (ID 2)
Product Manufacturer : 3Y POWER
Product Name       : URM1A151AM
Product Part Number : YM-2151E
Product Version    : A04R
Product Serial     : SA070N291617000392
Product Asset Tag  : 110104
Product Extra      : A07

FRU Device Description : UV300 PS1 (ID 3)
Product Manufacturer : 3Y POWER
Product Name       : URM1A151AM
Product Part Number : YM-2151E
Product Version    : A04R
Product Serial     : SA070N291617000295
Product Asset Tag  : 110104
Product Extra      : A07

```

Example 2: The following commands list FRUs for a specific chassis BMC.

```

remotehost% ssh administrator@flex-rmc
password: #####
flex-rmc eRMC:r001i01c cli> bmc exec "ipmi fru list" r1i1b

==== r001i01b ====
FRU Device Description : Builtin FRU Device (ID 0)
Board Mfg Date        : Mon Nov 13 03:43:00 2017
Board Mfg             : BENCHMARK
Board Product         : UV4_BASEIO
Board Serial          : SAL054
Board Part Number     : P0003171_005
Board Extra           : _A01
Board Extra           : 00001111
Board Extra           : 00000002
Board Extra           : 00002222
OEM (Hewlett Packard Enterprise) MAC Address Record
  Ethernet 1         : 08:00:69:12:34:51
  Ethernet 2         : 08:00:69:12:34:52
  Ethernet 3         : 08:00:69:12:34:53
  Ethernet 4         : 08:00:69:12:34:54
  Ethernet 5         : 08:00:69:12:34:55
  Ethernet 6         : 08:00:69:12:34:56

FRU Device Description : NODE_EEPROM (ID 1)
Board Mfg Date        : Wed Mar 1 23:41:00 2017

```

```

Board Mfg      : BENCHMARK
Board Product  : IP147
Board Serial   : SBB381
Board Part Number : P0001924_004
Board Extra    : _A01
Board Extra    : 00000111
Board Extra    : 00000002
Board Extra    : 00000222

FRU Device Description : PWR_EEPROM (ID 2)
...

```

Listing the local area network (LAN) configuration

The `lan print 1` command lists the LAN configuration.

Example: The following command lists the LAN configuration for the RMC.

```

remotehost% ipmitool -I lanplus -H flex-rmc -U ADMIN -P password lan print 1
Set in Progress      : Set Complete
Auth Type Support    : NONE MD2 MD5 PASSWORD
Auth Type Enable     : Callback :
                    : User      : MD5
                    : Operator : MD5
                    : Admin    : MD5
                    : OEM      : MD5
IP Address Source    : Static Address
IP Address           : 192.168.1.22
Subnet Mask          : 255.255.255.0
MAC Address          : 00:00:00:00:00:00
SNMP Community String : sgi
IP Header            : TTL=0x40 Flags=0x40 Precedence=0x00 TOS=0x10
BMC ARP Control      : ARP Responses Enabled, Gratuitous ARP Disabled
Gratuitous ARP Intrvl : 2.0 seconds
Default Gateway IP    : 0.0.0.0
Default Gateway MAC   : 00:00:00:00:00:00
Backup Gateway IP     : 0.0.0.0
Backup Gateway MAC    : 00:00:00:00:00:00
RMCP+ Cipher Suites  : 0,1,2,3,6,7,8,11,12
Cipher Suite Priv Max : XuuOXXuuOXXuOXX
                    : X=Cipher Suite Unused
                    : c=CALLBACK
                    : u=USER
                    : o=OPERATOR
                    : a=ADMIN
                    : O=OEM

```

Listing environmental sensors and values

The `sdr list` command retrieves a list of the system's environmental sensors and shows their values.

Example 1: The following command targets the RMC.

```

remotehost% ipmitool -I lanplus -H flex-rmc -U ADMIN -P password sdr list
RMC Inlet      | 23 degrees C      | ok
RMC Exhst     | 24 degrees C      | ok
RMC PSOC Temp  | 23 degrees C      | ok
RMC PSOC 3.3v  | 3.33 Volts        | ok
RMC 12v        | 12.08 Volts       | ok

```



```

RMC 3.3v Aux      | 3.33 Volts      | ok
RMC 2.5v         | 2.47 Volts      | ok
RMC 1.5v         | 1.49 Volts      | ok
RMC 1.2v         | 1.19 Volts      | ok
RMC 1.0v         | 1.00 Volts      | ok
RMC 0.75v        | 0.76 Volts      | ok
Pwr Good         | 0x00            | ok
Pwr Sply Fault   | 0x00            | ok
Pwr Redund OK    | 0x00            | ok
Stacking Pres    | 0x00            | ok
PS0 VOUT         | 11.92 Volts     | ok
PS0 IOUT         | 0.41 Amps       | ok
PS0 Temp1        | 20 degrees C    | ok
PS0 Temp2        | 23 degrees C    | ok
PS0 Fan1         | 10858 RPM       | ok
PS1 VOUT         | 12.05 Volts     | ok
PS1 IOUT         | 0.69 Amps       | ok
PS1 Temp1        | 19 degrees C    | ok
PS1 Temp2        | 23 degrees C    | ok
PS1 Fan1         | 10797 RPM       | ok
Memory Warning   | 0x00            | ok
Memory Critical  | 0x00            | ok
Disk Warning     | 0x00            | ok
Disk Critical    | 0x00            | ok
12V ENA          | 0x00            | ok

```

Example 2: The following commands target a specific BMC.

```

remotehost% ssh administrator@flex-rmc
password: #####
RMC cli> bmc exec "ipmi sdr list" r11b
P0_THROT_OFFSET | -58 degrees C   | ok
P1_THROT_OFFSET | -57 degrees C   | ok
P2_THROT_OFFSET | -59 degrees C   | ok
P3_THROT_OFFSET | -58 degrees C   | ok
CPU0_DIMM_MAX   | 29 degrees C    | ok
CPU1_DIMM_MAX   | 27 degrees C    | ok
CPU2_DIMM_MAX   | 26 degrees C    | ok
CPU3_DIMM_MAX   | 26 degrees C    | ok
CPU0_DIMM_MAX_ID | 0x00           | ok
CPU1_DIMM_MAX_ID | 0x00           | ok
CPU2_DIMM_MAX_ID | 0x00           | ok
CPU3_DIMM_MAX_ID | 0x00           | ok
HARP_INLET_TEMP | 23 degrees C    | ok
HARP_OUTLET_TEMP | 31 degrees C    | ok
...

```

Retrieving the system event log (SEL)

The `sel list` command retrieves the system event log (SEL).

Example 1: The following command targets the RMC.

```

remotehost% ipmitool -I lanplus -H flex-rmc -U administrator -P password sel list
 1 | 07/10/2017 | 12:11:20 | Temperature #0x01 | Upper Non-critical going low | Asserted
 2 | 07/10/2017 | 12:11:20 | Temperature #0x01 | Upper Critical going low | Asserted
 3 | 07/10/2017 | 12:11:20 | Temperature #0x01 | Upper Non-recoverable going low | Asserted
 4 | 07/10/2017 | 12:11:20 | Temperature #0x02 | Upper Non-critical going low | Asserted
 5 | 07/10/2017 | 12:11:20 | Temperature #0x02 | Upper Critical going low | Asserted
 6 | 07/10/2017 | 12:11:20 | Temperature #0x02 | Upper Non-recoverable going low | Asserted
 7 | 07/10/2017 | 12:16:53 | Voltage #0x0a | Upper Non-critical going high | Asserted
 8 | 07/10/2017 | 12:16:53 | Voltage #0x0a | Upper Critical going high | Asserted
 9 | 07/10/2017 | 12:16:53 | Voltage #0x0a | Upper Non-recoverable going high | Asserted
 a | 07/10/2017 | 12:16:56 | Voltage #0x0a | Upper Non-critical going high | Deasserted
 b | 07/10/2017 | 12:16:56 | Voltage #0x0a | Upper Critical going high | Deasserted

```



```
c | 07/10/2017 | 12:16:56 | Voltage #0x0a | Upper Non-recoverable going high | Deasserted
d | 07/10/2017 | 12:23:50 | Voltage #0x0a | Lower Non-critical going low | Asserted
e | 07/10/2017 | 12:23:50 | Voltage #0x0a | Lower Critical going low | Asserted
f | 07/10/2017 | 12:23:50 | Voltage #0x0a | Lower Non-recoverable going low | Asserted
10 | 07/10/2017 | 12:23:53 | Voltage #0x0a | Lower Non-critical going low | Deasserted
11 | 07/10/2017 | 12:23:53 | Voltage #0x0a | Lower Critical going low | Deasserted
12 | 07/10/2017 | 12:23:53 | Voltage #0x0a | Lower Non-recoverable going low | Deasserted
13 | 07/10/2017 | 12:34:50 | Voltage #0x0a | Upper Non-critical going high | Asserted
...
```

Example 2: The following commands target a specific chassis BMC.

```
remotehost% ssh administrator@flex-rmc
password: #####
RMC cli> bmc exec "ipmi sel list" r1i1b
SEL has no entries
```

Opening a serial console

The `sol activate` command opens a serial console to the server. You might want to issue this command, for example, if you want to observe system messages during a boot.

```
remotehost% ipmitool -I lanplus -H flex-rmc -U ADMIN -P password \
sol activate
```



Disaster recovery with HPE Serviceguard for Linux (SGLX)

HPE Serviceguard for Linux provides disaster recovery and failover options for Superdome Flex Server.

HPE Serviceguard for Linux user guides and manuals can be found at <http://www.hpe.com/info/Linux-serviceguard-docs>.



HPE Persistent Memory on HPE Superdome Flex Server

Hewlett Packard Enterprise supports HPE Persistent Memory on Superdome Flex Server with Intel Xeon 62xx and 82xx (Cascade Lake) processors.

For details on supported PMM configurations and related administration and maintenance tasks, see [**HPE Persistent Memory Guide for HPE Superdome Flex Server**](#).



Websites

HPE Superdome Flex Server websites

HPE Superdome Flex Server product page

www.hpe.com/support/superdome-flex-product

HPE Superdome Flex Server customer documentation

www.hpe.com/support/superdome-flex-docs

HPE Superdome Flex Server software

www.hpe.com/support/superdome-flex-software

Server operating systems and virtualization software

www.hpe.com/us/en/servers/server-operating-systems.html

HPE Superdome Flex Server QuickSpecs

www.hpe.com/support/superdome-flex-quickspecs

HPE Superdome Flex Server spare parts list

www.hpe.com/support/superdome-flex-spareparts

HPE Superdome Flex Server release sets (support matrix)

www.hpe.com/support/superdome-flex-release-sets

HPE Superdome Flex Server support documentation

HPE Superdome Flex Server documentation for support specialists is available at www.hpe.com/support/superdome-flex-docs-restricted by signing in to the [Hewlett Packard Enterprise Support Center](#) with an entitled account.

Related product websites

HPE 9361-4i RAID Controller (Q2N11A)

1. Go to the Broadcom MegaRAID SAS 9361-4i product page.

<https://www.broadcom.com/products/storage/raid-controllers/megaraid-sas-9361-4i>

2. Select the **Documentation** tab.

3. Under **User Guide** select:

- *12Gb/s MegaRAID SAS RAID Controllers User Guide*
- *12Gb/s MegaRAID SAS Software User Guide*

HPE 3154-8e RAID Controller (Q6M15A)

- **[HPE Smart Storage Administrator User Guide](#)**
- *Microsemi Adaptec Smart HBA & RAID - Installation And User's Guide*

1. Go to the Microsemi Adaptec SmartRAID 3154-8e product page.

https://storage.microsemi.com/en-us/support/raid/sas_raid/asr-3154-8e/

2. Select the **Documentation** tab and download the guide.



General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

For additional websites, see **Support and other resources**.



Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<https://www.hpe.com/info/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<https://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

Hewlett Packard Enterprise Support Center

<https://www.hpe.com/support/hpesc>

Hewlett Packard Enterprise Support Center: Software downloads

<https://www.hpe.com/support/downloads>

My HPE Software Center

<https://www.hpe.com/software/hpesoftwarecenter>

- To subscribe to eNewsletters and alerts:
<https://www.hpe.com/support/e-updates>
- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:
<https://www.hpe.com/support/AccessToSupportMaterials>





IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

<https://www.hpe.com/services/getconnected>

HPE Proactive Care services

<https://www.hpe.com/services/proactivecare>

HPE Datacenter Care services

<https://www.hpe.com/services/datacentercare>

HPE Proactive Care service: Supported products list

<https://www.hpe.com/services/proactivecaresupportedproducts>

HPE Proactive Care advanced service: Supported products list

<https://www.hpe.com/services/proactivecareadvancedsupportedproducts>

Proactive Care customer information

Proactive Care central

<https://www.hpe.com/services/proactivecarecentral>

Proactive Care service activation

<https://www.hpe.com/services/proactivecarecentralgetstarted>

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

<https://www.hpe.com/support/ProLiantServers-Warranties>

HPE Enterprise and Cloudline Servers

<https://www.hpe.com/support/EnterpriseServers-Warranties>

HPE Storage Products

<https://www.hpe.com/support/Storage-Warranties>

HPE Networking Products

<https://www.hpe.com/support/Networking-Warranties>

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

<https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>



Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

<https://www.hpe.com/info/reach>

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

<https://www.hpe.com/info/ecodata>

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

<https://www.hpe.com/info/environment>

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

