



system controller commands

ONTAP 9.15.1 commands

NetApp
May 17, 2024

Table of Contents

- system controller commands 1
 - system controller show 1
 - system controller bezel-led modify 2
 - system controller bezel-led show 3
 - system controller bootmedia show-serial-number 4
 - system controller bootmedia show 6
 - system controller clus-flap-threshold show 8
 - system controller config show-errors 9
 - system controller config show 10
 - system controller config pci show-add-on-devices 12
 - system controller config pci show-hierarchy 15
 - system controller coredump-device show-serial-number 18
 - system controller coredump-device show 20
 - system controller environment show 23
 - system controller flash-cache show 25
 - system controller flash-cache secure-erase run 28
 - system controller flash-cache secure-erase show 28
 - system controller fru show-manufacturing-info 29
 - system controller fru show 32
 - system controller fru led disable-all 35
 - system controller fru led enable-all 35
 - system controller fru led modify 36
 - system controller fru led show 38
 - system controller ioxm show 39
 - system controller location-led modify 41
 - system controller location-led show 42
 - system controller memory dimm show 43
 - system controller nvram-bb-threshold show 46
 - system controller pci show 46
 - system controller pcicerr threshold modify 48
 - system controller pcicerr threshold show 49
 - system controller platform-capability show 49
 - system controller replace cancel 50
 - system controller replace pause 51
 - system controller replace resume 51
 - system controller replace show-details 52
 - system controller replace show 54
 - system controller replace start 55
 - system controller service-event delete 58
 - system controller service-event show 58
 - system controller slot module insert 60
 - system controller slot module remove 61
 - system controller slot module replace 61

system controller slot module show	62
system controller sp config show	63
system controller sp upgrade show	65

system controller commands

system controller show

Display the controller information

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller show` command displays information about all the controllers in the cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about all the controllers in the cluster:

- Controller name
- System ID
- System serial number
- Controller model name
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information about all the controllers in the cluster.

[-node {<nodename>|local}] - Node

Selects information about the specified controller.

[-system-id <text>] - System ID

Selects information about the controller with the specified System ID.

[-model <text>] - Model Name

Selects information about the controllers with the specified model name.

[-part-number <text>] - Part Number

Selects information about the controllers with the specified part number.

[-revision <text>] - Revision

Selects information about the controllers with the specified revision.

[-serial-number <text>] - Serial Number

Selects information about the controller with the specified system serial number.

[`-controller-type <text>`] - Controller Type

Selects information about the controllers with the specified controller type.

[`-status {ok|ok-with-suppressed|degraded|unreachable|unknown}`] - Status

Selects information about the controllers with the specified health monitor status.

[`-chassis-id <text>`] - Chassis ID

Selects information about the controllers with the specified chassis ID.

Examples

The below example displays information about all controllers in the cluster.

```
cluster1::> system controller show
      Controller Name           System ID           Serial Number       Model
      Status
      -----
node1           140733730268652    700001456939       FAS2520
ok
node2           140733730268667    700001456941       FAS2520
ok
      2 entries were displayed.
```

The example below displays detailed information about specified controller in the cluster.

```
cluster1::> system controller show -instance -node node1
      Node: node1
      System ID: 140733730268652
      Model Name: FAS2520
      Part Number: 111-01316
      Revision: 21
      Serial Number: 700001456939
      Controller Type: none
      Status: ok
      Chassis ID: 4591227214
```

system controller bezel-led modify

Modify the state of the bezel LED log (lit or unlit)

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller bezel-led modify` command controls the illumination state of the LED logo on the chassis bezel (lit or unlit) of a node. On platforms where two controllers in a high-availability (HA) configuration share the same chassis, any changes to the LED logo's illumination state are kept synchronized on both nodes.

Parameters

-node {<nodename>|local} - Node

The node whose bezel LED logo illumination state is being modified.

[-state {off|on|disconnected|unknown}] - Bezel LED Logo State

The bezel LED logo illumination state of the node. The possible values are: on (lit), off (unlit), disconnected. The default state is on .

Examples

The following example turns on the bezel LED logo:

```
cluster1::> system controller bezel-led modify -node node1 -state on
```

The following example turns off the bezel LED logo:

```
cluster1::> system controller bezel-led modify -node node1 -state off
```

system controller bezel-led show

Display the state of the bezel LED logo (lit or unlit)

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller bezel-led show` command shows the current illumination state of the bezel LED logo (lit or unlit).

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-node` {<nodename>|local}] - Node

The node whose bezel LED logo illumination state is being shown.

[`-state` {off|on|disconnected|unknown}] - Bezel LED Logo State

The bezel LED logo illumination state of the node.

Examples

The following example lists the current illumination state of the bezel LED logo:

```
cluster1::> system controller bezel-led show
Node           Bezel LED Logo State
-----
node1          off
node2          off
```

system controller bootmedia show-serial-number

Display the Boot Media Device serial number

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller bootmedia show-serial-number` command displays the Boot Media Device serial number. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about the bootmedia:

- Node name
- Display name
- Serial Number
- Size
- Bootmedia state
- Status

To display more details, use the `-instance` parameter.

Parameters

{ [`-fields` <fieldname>,...]

Selects the fields that you specify.

| [`-instance`] }

Displays detailed information for all the bootmedia devices.

[-node {<nodename>|local}] - Node

Selects the bootmedia device that is present on the specified node.

[-serial-num <text>] - Serial Number

Selects the bootmedia devices with the specified serial number.

[-vendor-id <Hex Integer>] - Vendor ID

Selects the bootmedia devices with the specified vendor ID.

[-device-id <Hex Integer>] - Device ID

Selects the bootmedia devices with the specified device ID.

[-display-name <text>] - Display Name

Selects the bootmedia devices with the specified display name.

[-unique-name <text>] - Unique Name

Selects the bootmedia device with the specified unique name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects the bootmedia devices with the specified health monitor.

[-usbmon-status {present|not-present}] - Bootmedia Health Monitor

Selects the bootmedia devices with the specified USBMON status.

[-device-state {good|warn|bad}] - Bootmedia State

Selects the bootmedia devices with the specified device state.

[-size <integer>] - Max Memory Size (MB)

Selects the bootmedia devices with the specified memory size.

[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects the bootmedia devices with the specified health monitor status.

Examples

The following example displays the information of the bootmedia devices present in all the nodes in a cluster:


```

cluster1::> system controller bootmedia show-serial-number

Node           Display Name           Serial Number           (MB) State
Status
-----
or-12-01
ok             BootMedia/SAMSUNG     S2J4NXAGA08186        122104 good
              MZVLV128HCGR-00000
              BootMedia-2/SAMSUNG  S2J4NXAGA08198        122104 good
ok             MZVLV128HCGR-00000
2 entries were displayed.

```

The following example displays the detailed information about the bootmedia present in a node:

```

cluster1::> system controller bootmedia show-serial-number -instance -node
node1
Node: node1
      Vendor ID: 8086
      Device ID: 8d02
      Display Name: TOSHIBA THNSNJ060GMCU
      Unique Name: /dev/ad4s1 (TOSHIBA THNSNJ060GMCU)
      Health Monitor Name: controller
Bootmedia Health Monitor: present
      Bootmedia State: good
      Max memory size(in MB): 16367
      Status: ok
      Serial number: Y4IS104FTNEW

```

system controller bootmedia show

Display the Boot Media Device Health Status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller bootmedia show` command displays details of the bootmedia devices present in all the nodes in a cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about the bootmedia:

- Node name
- Display name

- Vendor ID
- Device ID
- Memory size
- Bootmedia state
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information for all the bootmedia devices.

[-node {<nodename>|local}] - Node

Selects the bootmedia device that is present on the specified node.

[-serial-num <text>] - Serial Number

Selects the bootmedia devices with the specified serial number.

[-vendor-id <Hex Integer>] - Vendor ID

Selects the bootmedia devices with the specified vendor ID.

[-device-id <Hex Integer>] - Device ID

Selects the bootmedia devices with the specified device ID.

[-display-name <text>] - Display Name

Selects the bootmedia devices with the specified display name.

[-unique-name <text>] - Unique Name

Selects the bootmedia device with the specified unique name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects the bootmedia devices with the specified health monitor.

[-usbmon-status {present|not-present}] - Bootmedia Health Monitor

Selects the bootmedia devices with the specified USBMON status.

[-device-state {good|warn|bad}] - Bootmedia State

Selects the bootmedia devices with the specified device state.

[-size <integer>] - Max Memory Size (MB)

Selects the bootmedia devices with the specified memory size.

[`-health {ok|ok-with-suppressed|degraded|unreachable|unknown}`] - Status

Selects the bootmedia devices with the specified health monitor status.

Examples

The following example displays the information of the bootmedia devices present in all the nodes in a cluster:

```
cluster1::> system controller bootmedia show
Size Bootmedia
Node          Display Name          Vendor ID Device ID      (MB) State
Status
-----
node1         Micron Technology     634       655      1929 good
ok
              0x655
node2         Micron Technology     634       655      1929 good
ok
              0x655
```

The example below displays the detailed information about the bootmedia present in a node.

```
cluster1::> system controller bootmedia show -instance -node node1
Node: node1
      Vendor ID: 634
      Device ID: 655
      Display Name: Micron Technology 0x655
      Unique Name: Micron Technology 0x655 (ad.0)
      Health Monitor Name: controller
      USBMON Health Monitor: present
      Bootmedia State: good
      Max memory size(in MB): 1929
      Status: ok
```

system controller clus-flap-threshold show

Display the controller cluster port flap threshold

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller clus-flap-threshold show` command allows the display of the threshold for link flapping counts for all nodes. This threshold would be the number of times the cluster port links for a given node can flap (go down) within a polling period before triggering an alert.

system controller config show-errors

Display configuration errors

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller config show-errors` displays configuration errors.

- node
- description

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information for all the PCI devices.

[-node {<nodename>|local}] - Node

Displays configuration errors on the specified node.

[-verbose <true>] - Verbose Output?

The `-verbose` parameter enables verbose mode, resulting in the display of more detailed output.

[-description <text>] - Error Description

Displays the node with the specified configuration error.

Examples

The example below displays configuration errors on all the nodes in the cluster.

```

cluster1::> system controller config show-errors

Configuration Info and Errors for Node: cluster1-01
-----
----
Chelsio T320E 2x10G NIC card (PN X1008A) in slot 1 is not supported on
model FAS3210

Configuration Info and Errors for Node: cluster1-02
-----
----
PCI-E Dual 10/100/1000 Ethernet G20 card (PN X1039A) in slot 2 is not
supported on model FAS3210

cluster1::>

cluster1::> system controller config show-errors -verbose

Configuration Info and Errors for Node: cluster1-01
-----
----
sysconfig: Card in slot 2 (7-1275-0008-46848) is not supported.
sysconfig: slot 12 OK: X2067: Proprietary embedded SAS HBA

cluster1::>

```

system controller config show

Display System Configuration Information

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller config show` command displays system configuration information for the devices present in the controller. To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

Selects the nodes that match this parameter value.

[-device <text>] - Device

Selects the configuration information that matches the specified device.

[-subslot <integer>] - Subslot Number

Selects the configuration information that matches the specified subslot.

[-info <text>] - Device Info

Selects the configuration information that matches the specified device information.

[-slot <text>] - Slot

Selects the configuration information that matches the specified slot.

Examples

The following example displays configuration information for slot 1 of the controller:

```
cluster1::> system controller config show -slot 1

Node: node1
Sub- Device/
Slot slot Information
---- ----
-----
 1    - NVRAM10 HSL
      Device Name:      Interconnect HBA: Generic OFED Provider
      Port Name:        ib1a
      Default GUID:     fe80:0000:0000:0000:0000:0000:0000:0104
      Base LID:         0x104
      Active MTU:       8192
      Data Rate:        0 Gb/s (8X)
      Link State:       DOWN
      QSFP Vendor:      Amphenol
      QSFP Part Number: 112-00436+A0
      QSFP Type:        Passive Copper 1m ID:00
      QSFP Serial Number: APF16130066875
      QSFP Vendor:      Amphenol
      QSFP Part Number: 112-00436+A0
      QSFP Type:        Passive Copper 1m ID:00
      QSFP Serial Number: APF16130066857

cluster1::>
```

system controller config pci show-add-on-devices

Display PCI devices in expansion slots

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller config pci show-add-on-devices` command displays information about the PCIe devices in I/O expansion slots. The command displays the following information about the PCIe devices:

- Node
- Model
- Type
- Slot
- Device
- Vendor
- Sub-device ID

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information about PCI devices.

[-node {<nodename>|local}] - Node

Selects the PCI devices that are present in the specified node.

[-model <text>] - Model String

Selects the PCI devices that are present on the system with the specified model name.

[-type <integer>] - Device Type

Selects the PCI devices with the specified device type.

[-slot-and-sub <text>] - PCI Slot Number

Selects the PCI devices present in the specified slot or slot-subslot combination.

[-device <text>] - Device

Selects the PCI devices with the specified device ID.

[-vendor <text>] - Vendor Number

Selects the PCI devices with the specified vendor ID.

`[-sub-device-id <integer>]` - Sub Device ID

Selects the PCI devices with the specified sub-device ID.

Examples

The example below displays information about PCI devices found in I/O expansion slots of all the nodes in the cluster.


```

cluster1::> system controller config pci show-add-on-devices
Node           Model           Slot Type Device  Vendor Sub-Device
ID
-----
cluster1-01    FAS6240
10             6      7  0x2532 0x1077
0              5      1  0x1527 0x8086
0              2      7  0x6732 0x15B3
0              3      1  0x8030 0x1077
0              1      2  0x8001 0x11F8
0              15     1  0x10FB 0x8086
0              13     1  0x150E 0x8086
1              7      1  0x1528 0x8086
0
cluster1-02    FAS6240
10             6      7  0x2532 0x1077
0              5      1  0x1527 0x8086
0              2      7  0x6732 0x15B3
0              3      1  0x8030 0x1077
0              1      2  0x8001 0x11F8
0              15     1  0x10FB 0x8086
0              13     1  0x150E 0x8086
1              7      1  0x1528 0x8086
0
16 entries were displayed.

cluster1::>

```

system controller config pci show-hierarchy

Display PCI hierarchy

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller config pci show-hierarchy` command displays the PCI Hierarchy of all PCI devices found in a controller. The command displays the following information about the PCI devices:

- Node
- Level
- Device
- Link Capability
- Link Status

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information for PCI devices.

[-node {<nodename>|local}] - Node

Displays the PCI hierarchy of the specified node.

[-level <integer>] - PCI Device Level

Displays the PCI devices that match the specified level within the PCI hierarchy.

[-pci-device <text>] - PCI Device

Displays the PCI devices that match the specified device description.

[-link-cap <text>] - Link Capability

Displays the PCI devices that match the specified link capability.

[-link-status <text>] - Link Status

Displays the PCI devices that match the specified link status.

Examples

The example below displays the PCI hierarchy for all of the nodes in the cluster.

```
cluster1:> system controller config pci show-hierarchy
PCI Hierarchy
```

Node: cluster1-01

```
Level Device      Link
-----
1      Br[3721](0,3,0): PCI Device 8086:3721 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(2),LkWd(4),DLAct),
2      Dv[8001](1,0,0): PMC SAS adapter on Controller
LinkCap(MaxLkSp(2),MaxLkWd(8),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(2),LkWd(4),SCLk),
1      Br[3722](0,4,0): PCI Device 8086:3722 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),DLAct),
2      Dv[6274](2,0,0): PCI Device 15b3:6274 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(8),ASPM(1),L0(7),L1(7),Port(68))
      LinkStatus(LkSp(1),LkWd(4)),
1      Br[3723](0,5,0): PCI Device 8086:3723 on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(0)),
1      Br[3b42](0,28,0): PCI Device 8086:3b42 on Controller
LinkCap(MaxLkSp(1),MaxLkWd(4),ASPM(3),L0(4),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk,DLAct),
2      Dv[150e](4,0,0): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk),
2      Dv[150e](4,0,1): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk),
2      Dv[150e](4,0,2): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk),
2      Dv[150e](4,0,3): Intel 1G NIC on Controller
LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(6),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk),
1      Br[3b4a](0,28,4): PCI Device 8086:3b4a on Controller
```

```

LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(1),SCLk,DLAct),
2      Dv[10d3](5,0,0): Intel 1G NIC on Controller

LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(1),SCLk),
1      Br[3b4e](0,28,6): PCI Device 8086:3b4e on Controller

LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(4),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(1),SCLk,DLAct),
2      Dv[10d3](7,0,0): Intel 1G NIC on Controller

LinkCap(MaxLkSp(1),MaxLkWd(1),ASPM(3),L0(1),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(1),SCLk),

Node: cluster1-02

Level Device      Link
-----
1      Br[3721](0,3,0): PCI Device 8086:3721 on Controller

LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(2),LkWd(4),DLAct),
2      Dv[8001](1,0,0): PMC SAS adapter on Controller

LinkCap(MaxLkSp(2),MaxLkWd(8),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(2),LkWd(4),SCLk),
1      Br[3722](0,4,0): PCI Device 8086:3722 on Controller

LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),DLAct),
2      Dv[6274](2,0,0): PCI Device 15b3:6274 on Controller

LinkCap(MaxLkSp(1),MaxLkWd(8),ASPM(1),L0(7),L1(7),Port(68))
      LinkStatus(LkSp(1),LkWd(4)),
1      Br[3723](0,5,0): PCI Device 8086:3723 on Controller

LinkCap(MaxLkSp(2),MaxLkWd(4),ASPM(3),L0(3),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(0)),
1      Br[3b42](0,28,0): PCI Device 8086:3b42 on Controller

LinkCap(MaxLkSp(1),MaxLkWd(4),ASPM(3),L0(4),L1(6),Port(68))
      LinkStatus(LkSp(1),LkWd(4),SCLk,DLAct),
2      Dv[150e](4,0,0): Intel 1G NIC on Controller

```

```

LinkCap (MaxLkSp (2) ,MaxLkWd (4) ,ASPM (3) ,L0 (6) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (4) ,SClk) ,
2      Dv [150e] (4,0,1) : Intel 1G NIC on Controller

LinkCap (MaxLkSp (2) ,MaxLkWd (4) ,ASPM (3) ,L0 (6) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (4) ,SClk) ,
2      Dv [150e] (4,0,2) : Intel 1G NIC on Controller

LinkCap (MaxLkSp (2) ,MaxLkWd (4) ,ASPM (3) ,L0 (6) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (4) ,SClk) ,
2      Dv [150e] (4,0,3) : Intel 1G NIC on Controller

LinkCap (MaxLkSp (2) ,MaxLkWd (4) ,ASPM (3) ,L0 (6) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (4) ,SClk) ,
1      Br [3b4a] (0,28,4) : PCI Device 8086:3b4a on Controller

LinkCap (MaxLkSp (1) ,MaxLkWd (1) ,ASPM (3) ,L0 (4) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (1) ,SClk,DLAct) ,
2      Dv [10d3] (5,0,0) : Intel 1G NIC on Controller

LinkCap (MaxLkSp (1) ,MaxLkWd (1) ,ASPM (3) ,L0 (1) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (1) ,SClk) ,
1      Br [3b4e] (0,28,6) : PCI Device 8086:3b4e on Controller

LinkCap (MaxLkSp (1) ,MaxLkWd (1) ,ASPM (3) ,L0 (4) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (1) ,SClk,DLAct) ,
2      Dv [10d3] (7,0,0) : Intel 1G NIC on Controller

LinkCap (MaxLkSp (1) ,MaxLkWd (1) ,ASPM (3) ,L0 (1) ,L1 (6) ,Port (68) )
      LinkStatus (LkSp (1) ,LkWd (1) ,SClk) ,
28 entries were displayed.
cluster::>

```

system controller coredump-device show-serial-number

Display the coredump device serial number

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller coredump-device show-serial-number` command displays the serial number of coredump devices. This command is only applicable to AFF A700 systems. By default, the command displays the following information about the coredump device:

- Node name

- Display name
- Serial number
- Size
- Device state
- Status

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information for the coredump device.

[-node {<nodename>|local}] - Node

Selects the coredump device that is present on the specified node.

[-serial-num <text>] - Serial Number

Selects the coredump device with the specified serial number.

[-vendor-id <Hex Integer>] - Vendor ID

Selects the coredump device with the specified vendor ID.

[-device-id <Hex Integer>] - Device ID

Selects the coredump device with the specified device ID.

[-display-name <text>] - Display Name

Selects the coredump device with the specified display name.

[-unique-name <text>] - Unique Name

Selects the coredump device with the specified unique name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects the coredump device with the specified health monitor.

[-health-status {present|not-present}] - Health Status

Selects the coredump device with the specified health status.

[-device-state {good|warn|bad}] - Device State

Selects the coredump device with the specified device state.

[-size <integer>] - Max Memory Size (GB)

Selects the coredump device with the specified memory size.

[`-health {ok|ok-with-suppressed|degraded|unreachable|unknown}`] - Status

Selects the coredump device with the specified health monitor status.

[`-shelf-status {present|not-present}`] - Shelf Connected Status

Selects the coredump device with the specified shelf connected status.

Examples

The following example displays the information of the coredump device present in all the nodes in a cluster:

```
cluster1::> system controller coredump-device show-serial-number
Size Device
Node          FRU/Display Name      Serial Number          (GB) State
Status
-----
cluster1-01
ok            Coredump/X9170A      A22P5061550000644     1024 good
cluster1-02
ok            Coredump/X9170A      A22P5061550000137     1024 good
2 entries were displayed.
```

The following example displays the detailed information about the coredump device present in a node:

```
cluster1::> system controller coredump-device show-serial-number -instance
-node cluster1-01
Node: cluster1-01
  Serial Number: A22P5061550000644
  Vendor ID: 1b85
  Device ID: 4018
  Display Name: Coredump/X9170A
  Unique Name: 0X917011900207410CZ000Z63000001T0003500
  Health Monitor Name: controller
  Health Status: present
  Device State: good
  Max Memory Size (GB): 1024
  Status: ok
  Shelf Connected Status: present
```

system controller coredump-device show

Display the coredump device health status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller coredump-device show` command displays details of the coredump devices present in all the nodes in a cluster. This command is only applicable to AFF A700 systems. By default, the command displays the following information about the coredump device:

- Node name
- Display name
- Vendor ID
- Device ID
- Memory size
- Device state
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information for the coredump device.

[-node {<nodename>|local}] - Node

Selects the coredump device that is present on the specified node.

[-serial-num <text>] - Serial Number

Selects the coredump device with the specified serial number.

[-vendor-id <Hex Integer>] - Vendor ID

Selects the coredump device with the specified vendor ID.

[-device-id <Hex Integer>] - Device ID

Selects the coredump device with the specified device ID.

[-display-name <text>] - Display Name

Selects the coredump device with the specified display name.

[-unique-name <text>] - Unique Name

Selects the coredump device with the specified unique name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects the coredump device with the specified health monitor.

[-health-status {present|not-present}] - Health Status

Selects the coredump device with the specified health status.

[-device-state {good|warn|bad}] - Device State

Selects the coredump device with the specified device state.

[-size <integer>] - Max Memory Size (GB)

Selects the coredump device with the specified memory size.

[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects the coredump device with the specified health monitor status.

[-shelf-status {present|not-present}] - Shelf Connected Status

Selects the coredump device with the specific shelf connected status.

Examples

The following example displays the information of the coredump devices present in all the nodes in a cluster:

```
cluster1::> system controller coredump-device show
Size Device
Node      FRU/Display Name      Vendor ID Device ID      (GB) State
Status
-----
cluster1-01
          Coredump/X9170A      1b85      4018      1024 good
ok
cluster1-02
          Coredump/X9170A      1b85      4018      1024 good
ok
2 entries were displayed.
```

The example below displays the detailed information about the coredump device present in a node.

```

cluster1::> system controller coredump-device show -instance -node
cluster1-01
Node: cluster1-01
    Serial Number: A22P5061550000644
        Vendor ID: 1b85
        Device ID: 4018
    Display Name: Coredump/X9170A
    Unique Name: 0X91701190020741OCZ000Z63000001T0003500
Health Monitor Name: controller
    Health Status: present
        Device State: good
Max Memory Size (GB): 1024
    Status: ok
Shelf Connected Status: present

```

system controller environment show

Display the FRUs in the controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller environment show` displays information about all environment FRUs in the cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about the environment FRUs in the cluster:

- Node
- FRU name
- FRU state

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information about the environment FRUs.

[-node {<nodename>|local}] - Node

Selects information about all the environment FRUs that the specified node owns.

[-serial-number <text>] - FRU Serial Number

Selects information about all the environment FRUs with the specified serial number.

[-fru-name <text>] - FRU Name

Selects information about the environment FRU with the specified FRU name.

[-type {controller|psu|fan|dimm|bootmedia|ioxm|nvram|nvdim}] - FRU Type

Selects information about all the environment FRUs with the specified FRU type.

[-name <text>] - Name

Selects information about all the environment FRUs with the specified unique name.

[-state <text>] - FRU State

Selects information about all the environment FRUs with the specified FRU state.

[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects information about all the environment FRUs with the specified health monitor status.

[-display-name <text>] - Display Name for the FRU

Selects information about all the environment FRUs with the specified display name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Monitor Name

Selects information about all the environment FRUs with the specified monitor.

[-model <text>] - Model Type

Selects information about all the environment FRUs with the specified FRU model.

[-shared {shared|not_shared}] - Shared Resource

Selects information about all the environment FRUs with the specified sharing type.

[-chassis-id <text>] - Chassis ID

Selects information about all the environment FRUs in the specified chassis.

[-additional-info <text>] - Additional Information About the FRU

Selects information about all the environment FRU with specified additional information.

[-seq-state-cnt <integer>] - Count of Same State

Selects information about all the environment FRU with specified sequential state count.

Examples

The following example displays information about all major environment FRUs in the cluster:

```
cluster1::> system controller environment show
```

Node	FRU Name	State
node1	PSU1 FRU	GOOD
node1	PSU2 FRU	GOOD
node2	PSU1 FRU	GOOD
node2	PSU2 FRU	GOOD

The following example displays detailed information about a specific environment FRU:

```
cluster1::> system controller environment show -node node1 -fru-name "PSU1  
FRU" -instance
```

```
Node: node1
```

```
FRU Serial Number: XXT122737891
```

```
FRU Name: PSU1 FRU
```

```
FRU Type: psu
```

```
Name: XXT122737891
```

```
FRU State: GOOD
```

```
Status: ok
```

```
Display Name for the FRU: PSU1 FRU
```

```
Monitor Name: controller
```

```
Model Type: none
```

```
Shared Resource: shared
```

```
Chassis ID: 4591227214
```

```
Additional Information About the FRU: Part Number: 114-00065+A0
```

```
Revision: 020F
```

```
Manufacturer: NetApp
```

```
FRU Name: PSU
```

system controller flash-cache show

Display the Flash Cache device status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller flash-cache show` command displays the current Flash Cache device information.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

If this parameter is specified, only status information for the matching node is displayed.

[-device-id <integer>] - Device ID

If this parameter is specified, only status information for the matching device-id is displayed.

[-slot <integer>] - Slot

If this parameter is specified, only status information for the matching slot is displayed. Slot can have a format of slot, slot-subslot, or domain:bus:slot.

[-subslot <integer>] - Subslot

If this parameter is specified, only status information for the matching subslot is displayed.

[-domain <integer>] - Domain

If this parameter is specified, only status information for the matching PCI domain is displayed.

[-bus <integer>] - Bus

If this parameter is specified, only status information for the matching PCI bus is displayed.

[-slot-string <text>] - Slot String

If this parameter is specified, only status information for the matching slot is displayed. Format can be slot or slot-subslot.

[-device-state

{ok|erasing|erased|failed|removed|online|offline_failed|degraded|offline_threshold}] - Device State

If this parameter is specified, only status information for the matching device-state is displayed.

[-model-number <text>] - Model Number

If this parameter is specified, only status information for the matching model-number is displayed.

[-part-number <text>] - Part Number

If this parameter is specified, only status information for the matching part-number is displayed.

[-serial-number <text>] - Serial Number

If this parameter is specified, only status information for the matching serial-number is displayed.

[-firmware-version <text>] - Firmware Version

If this parameter is specified, only status information for the matching firmware-version is displayed.

[-firmware-file <text>] - Firmware Filename

If this parameter is specified, firmware-file is the base name of the firmware file located in disk_fw to update

the device firmware.

[~~-hardware-revision~~ <text>] - Hardware Revision

If this parameter is specified, only status information for the matching hardware-revision is displayed.

[~~-capacity~~ <integer>] - Capacity

If this parameter is specified, only status information for the matching capacity is displayed.

[~~-last-change-time~~ <integer>] - Time Last State Change

If this parameter is specified, only status information for the matching last-change-time is displayed.

[~~-service-time~~ <integer>] - Service Time

If this parameter is specified, only status information for the matching service-time is displayed.

[~~-percent-online~~ <integer>] - Percent Online

If this parameter is specified, only status information for the matching percent-online is displayed.

[~~-average-erase-cycle-count~~ <integer>] - Avg Erase Cycle Count

If this parameter is specified, only status information for the matching average-erase-cycle-count is displayed.

[~~-threshold-profile~~ <text>] - Threshold Profile

If this parameter is specified, only status information for the matching threshold-profile is displayed.

Examples

The following example displays the current state of all Flash Cache devices:

```
cluster1::> system controller flash-cache show
          Device      Model Part      Serial      Firmware
Capacity Device
Node     ID      Slot Number Number      Number      Version
(GB) State
-----
node1
          0 6-1   X9172A 119-00209 A22P7061550000004  NA00
4096 ok
          1 6-2   X9170A 119-00207 A22P5061550000135  NA00
1024 ok
node2
          0 6-1   X9172A 119-00209 A22P7061550000007  NA00
4096 ok
          1 6-2   X9170A 119-00207 A22P5061550000091  NA00
1024 ok
4 entries were displayed.
```

system controller flash-cache secure-erase run

Perform a secure-erase operation on the targeted devices

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller flash-cache secure-erase run` command securely erases the given Flash Cache device.

Parameters

-node {<nodename>|local} - Node

Selects the node of the specified Flash Cache devices.

-device-id <integer> - Device Id

Selects the device-id of the specified Flash Cache devices.

Examples

The following example securely erases the selected Flash Cache device:

```
cluster1::> system controller flash-cache secure-erase -node nodel -device  
-id 0
```

system controller flash-cache secure-erase show

Display the Flash Cache card status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller flash-cache secure-erase show` command displays the current Flash Cache device secure-erase status.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-node` {<nodename>|local}] - Node

If this parameter is specified, only status information for the matching node is displayed.

[`-device-id` <integer>] - Device Id

If this parameter is specified, only status information for the matching device-id is displayed.

[`-device-state` {ok|erasing|erased|failed|removed}] - Device State

If this parameter is specified, only status information for the matching device-state is displayed.

Examples

The following example displays the current state of all the Flash Cache devices:

```
cluster1::> system controller flash-cache secure-erase show
Node          Device ID Slot Device State
-----
node1
              0 6-1  ok
              1 6-2  erasing
node2
              0 6-1  erased
              1 6-2  ok
4 entries were displayed.
```

system controller fru show-manufacturing-info

Display manufacturing information of FRUs

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller fru show-manufacturing-info` command displays manufacturing information for field replaceable units (FRUs) installed in the system. The information includes FRU-description, serial number, part number, and revision number. To display more details, use the `-instance` parameter.

Parameters

{ [`-fields` <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [`-instance`] }

Displays detailed information about the installed FRUs in the system.

[`-node {<nodename>|local}`] - Node

Selects a specific node's installed FRUs.

[`-system-sn <text>`] - System Serial Number

Selects information about installed FRUs with the specified system serial number.

[`-model-name <text>`] - Model Name

Selects information about installed FRUs with the specified model name.

[`-system-id <text>`] - System ID

Selects information about installed FRUs with the specified system ID.

[`-kernel-version <text>`] - Kernel Version

Selects information about installed FRUs with the specified kernel version.

[`-firmware-release <text>`] - Firmware Release

Selects information about installed FRUs with the specified firmware release.

[`-description <text>`] - FRU Description

Selects information about installed FRUs with the specified FRU description.

[`-vendor-id <text>`] - FRU Vendor ID

Selects information about the FRU with the vendor identification.

[`-serial-number <text>`] - FRU Serial Number

Selects information about the FRU with the specified serial number.

[`-part-number <text>`] - FRU Part Number

Selects information about the FRU with the specified part number.

[`-revision <text>`] - FRU Revision of Part Number

Selects information about the FRU with the specified revision.

[`-manufacturer <text>`] - FRU Manufacturer

Selects information about the FRU with the specified manufacturer.

[`-manufacture-date <text>`] - FRU Manufacturing Date

Selects information about the FRU with the specified manufacture date.

[`-product-id <text>`] - FRU Product Identifier

Selects information about the FRU with the specified product ID.

[`-firmware-version <text>`] - FRU Firmware Version

Selects information about the FRU with the specified firmware version.

Examples

The following example displays all installed FRUs in the system:



```
cluster1::> system controller fru show-manufacturing-info
```

```
Node: platsw-lodi-1-01
```

```
System Serial Number: 791541000047
```

```
Model Name: FAS9040
```

```
System ID: 0537024373
```

```
Firmware release: 10.0X18
```

```
Kernel Version: NetApp Release sysMman_3887886_1608151712: Mon Aug  
15
```

```
15:54:00 PDT 2016
```

FRU Description Rev.	FRU Serial Number	FRU Part Number	FRU
-----	-----	-----	
Mother Board	031537000390	111-02419	40
Chassis	031536000252	111-02392	40
DIMM-1	CE-01-1510-02A8DC73	SHB722G4LML23P2-SB	-
DIMM-3	CE-01-1510-02A8DCCC	SHB722G4LML23P2-SB	-
DIMM-8	CE-01-1510-02A8DE54	SHB722G4LML23P2-SB	-
DIMM-9	CE-01-1510-02A8DE1C	SHB722G4LML23P2-SB	-
DIMM-11	CE-01-1510-02A8DF42	SHB722G4LML23P2-SB	-
DIMM-16	CE-01-1510-02A8DD9B	SHB722G4LML23P2-SB	-
FAN1	031534001263	441-00058	40
FAN2	031534001292	441-00058	40
FAN3	031534001213	441-00058	40
PSU1	PSD092153200591	114-00146	40
PSU3	PSD092153200700	114-00146	40
mSATA boot0	1439100B02C3	-	MU03
1/10 Gigabit Ethernet Controller IX4-T	031538000121	111-02399	40
QLogic 8324 10-Gigabit Ethernet Controller	031535000664	111-02397	40
NVRAM10	031537000846	111-02394	40
NVRAM10 BATT	31534000932	NetApp, Inc.	111- 02591
NVRAM10 DIMM	CE-01-1510-02A8DC03	SHB722G4LML23P2-SB	-
PMC-Sierra PM8072 (111-02396)	031537000246	111-02396	41
PMC-Sierra PM8072 (111-02396)	031537000246	111-02396	41
PMC-Sierra PM8072 (111-02396)	031537000246	111-02396	41
PMC-Sierra PM8072 (111-02396)	031537000246	111-02396	41
PMC-Sierra PM8072 (111-02396)	031537000179	111-02396	41
Disk Serial Number	PNHH1J0B	X421_HCOBD450A10	-
Disk Serial Number	PNHH2BKB	X421_HCOBD450A10	-
Disk Serial Number	PNHJPZ8B	X421_HCOBD450A10	-
Disk Serial Number	PNHG6SKB	X421_HCOBD450A10	-
Disk Serial Number	PNHKJYTB	X421_HCOBD450A10	-
Disk Serial Number	PNHSVMEY	X421_HCOBD450A10	-
Disk Serial Number	PNHT8KWY	X421_HCOBD450A10	-
Disk Serial Number	PNHSVLOY	X421_HCOBD450A10	-

```

Disk Serial Number      PNHG5RHB                X421_HCOBD450A10      -
Disk Serial Number      PNHEXLWB                X421_HCOBD450A10      -
Disk Serial Number      PNHT8LJY                X421_HCOBD450A10      -
Disk Serial Number      PNHSVKZY                X421_HCOBD450A10      -
PMC-Sierra PM8072 (111-02396) 031537000179          111-02396              41
PMC-Sierra PM8072 (111-02396) 031537000179          111-02396              41
PMC-Sierra PM8072 (111-02396) 031537000179          111-02396              41
DS2246                  6000113106              0190                    -
DS2246-Pwr-Supply      XXT111825308            114-00065+A0           9C
DS2246-Pwr-Supply      XXT111825314            114-00065+A0           9C
DS2246-MODULE          8000675532              111-00690+A3           23
DS2246-MODULE          8000751790              111-00690+A3           23
DS2246-CABLE           512130075               112-00430+A0           -
DS2246-CABLE           -                         -                         -
DS2246-CABLE           512130118               112-00430+A0           -
DS2246-CABLE           -                         -                         -
49 entries were displayed.

```

system controller fru show

Display Information About the FRUs in the Controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller fru show` command displays information about all the controller specific Field Replaceable Units (FRUs) in the cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about all the FRUs in the cluster:

- Node
- FRU name
- Health monitor subsystem
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

```
{ [-fields <fieldname>,...]
```

Selects the fields that you specify.

```
| [-instance ] }
```

Displays detailed information about the controller specific FRUs in the cluster.

[-node {<nodename>|local}] - Node

Selects information about the FRUs in the specified node.

[-subsystem <Subsystem>] - Subsystem

Selects information about the FRUs of the specified subsystem.

[-serial-number <text>] - FRU Serial Number

Selects information about the FRU with the specified serial number.

[-fru-name <text>] - Name of the FRU

Selects information about the FRU with the specified FRU name.

[-type {controller|psu|fan|dimm|bootmedia|ioxm|nvram|nvdimm}] - FRU Type

Selects information about the FRU with the specified FRU type.

[-name <text>] - FRU Name

Selects information about the FRU with the specified unique name.

[-state <text>] - FRU State

Selects information about the FRU with the specified state.

[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects information about the FRU with the specified health monitor status.

[-display-name <text>] - Display Name for the Fru

Selects information about the FRU with the specified display name.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Monitor Name

Selects information about the FRU with the specified health monitor type.

[-model <text>] - Model Type

Selects information about the FRU with the specified model.

[-chassis-id <text>] - Chassis ID

Selects information about the FRU with the specified chassis ID.

[-location <text>] - Location of the FRU

Selects information about the FRU with the specified FRU location.

[-additional-info <text>] - Additional Information About the FRU

Selects information about the FRU with the specified additional information.

Examples

The example below displays information about all controller specific FRUs in the cluster.

```

cluster1::> system controller fru show
      Node              FRU Name              Subsystem
Status -----
-----
ok    node1              PSU1 FRU              Environment
ok    node1              PSU2 FRU              Environment
ok    node1              DIMM-NV1              Memory
ok    node1              DIMM-1                Memory
ok    node1              Micron Technology 0x655 (ad.0) Motherboard
ok    node2              PSU1 FRU              Environment      ok
ok    node2              PSU2 FRU              Environment
ok    node2              DIMM-NV1              Memory
ok    node2              DIMM-1                Memory
ok    node2              Micron Technology 0x655 (ad.0) Motherboard
ok
10 entries were displayed.

```

The example below displays information about the specific FRU.

```

cluster1::> system controller fru show -instance -serial-number AD-01-
1306-2EA01E9A
      Node: node1
      Subsystem: Memory
      FRU Serial Number: AD-01-1306-2EA01E9A
      Name of the FRU: DIMM-1
      FRU Type: dimm
      FRU Name: DIMM-1
      FRU State: ok
      Status: ok
      Display Name for the Fru: DIMM-1
      Monitor Name: controller
      Model Type: none
      Chassis ID: 4591227214
      Location of the FRU: Memory Slot: 1
      Additional Information About the FRU: Part No: HMT82GV7MMR4A-H9

```

system controller fru led disable-all

Turn off all the LEDs Data Ontap has lit

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller fru led disable-all` command turns off all the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which there are corresponding fault LEDs. In addition, there is a summary FRU fault LED on the external face-plate of both the controller and IOXM; labeled with a "!". A summary fault LED will be on when any of the internal FRU fault LEDs are on. Only the controller and IOXM internal FRU fault LEDs can be controlled by the end-user. The summary fault LEDs are turned on and off based on the simple policy described above. If you want to turn off the summary fault LED, you must turn off all internal FRU fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the system are FRU fault LEDs. Externally visible fault LEDs are labeled with a "!" and internal FRU fault LEDs remain on, even when the controller or IOXM is removed from the chassis. In addition, internal FRU fault LEDs will remain on until explicitly turned off by the end-user, even after a FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs, cards in PCI slots, boot media devices, NV batteries and coin cell batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is the number of FRUs of that particular type that exist in the controller or IOXM. Both controller and IOXM have a FRU map label for use in physically locating internal FRUs. The FRU ID/slot tuple used by the [system controller fru led show](#) command matches that specified on the FRU map label.

Examples

Turn off all FRU fault LEDs.

```
cluster1::*> system controller fru led disable-all
14 entries were modified.
```

Related Links

- [system controller fru led show](#)

system controller fru led enable-all

Light all the LEDs

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller fru led enable-all` command turns on all the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which there are corresponding fault LEDs. In addition, there is a summary FRU fault LED on the external face-plate of both the controller and IOXM; labeled with a "!". A summary fault LED will be on when any of the internal FRU fault LEDs are on. Only the controller and IOXM internal FRU fault LEDs can be controlled by the end-user. The summary fault LEDs are turned on and off based on the simple policy described above. If you want to turn off the summary fault LED, you must turn off all internal FRU fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the system are FRU fault LEDs. Externally visible fault LEDs are labeled with a "!" and internal FRU fault LEDs remain on, even when the controller or IOXM is removed from the chassis. In addition, internal FRU fault LEDs will remain on until explicitly turned off by the end-user, even after a FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs, cards in PCI slots, boot media devices, NV batteries and coin cell batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is the number of FRUs of that particular type that exist in the controller or IOXM. Both controller and IOXM have a FRU map label for use in physically locating internal FRUs. The FRU ID/slot tuple used by the [system controller fru led show](#) command matches that specified on the FRU map label.

Examples

Turn on all FRU fault LEDs.

```
cluster1:*> system controller fru led enable-all
14 entries were modified.
```

Related Links

- [system controller fru led show](#)

system controller fru led modify

Modify the status of FRU LEDs

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller fru led modify` command modifies the current state of the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which there are corresponding fault LEDs. In addition, there is a summary FRU fault LED on the external face-plate of both the controller and IOXM; labeled with a "!". A summary fault LED will be on when any of the internal FRU fault LEDs are on. Only the controller and IOXM internal FRU fault LEDs can be controlled by the end-user. The summary fault LEDs are turned on and off based on the simple policy described above. If you want to turn off the summary fault LED, you must turn off all internal FRU fault LEDs.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the system are FRU fault LEDs. Externally visible fault LEDs are labeled with a "!" and internal FRU fault LEDs remain on, even when the controller or IOXM is removed from the chassis. In addition, internal FRU fault LEDs will remain on until explicitly turned off by the end-user, even after a FRU has been replaced.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs, cards in PCI slots, boot media devices, NV batteries and coin cell batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is the number of FRUs of that particular type that exist in the controller or IOXM. Both controller and IOXM have a FRU map label for use in physically locating internal FRUs. The FRU ID/slot tuple used by the [system controller fru led show](#) command matches that specified on the FRU map label.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)

Selects FRU fault LEDs on the specified nodes.

-fru-id <FRU LED key> - FRU ID (privilege: advanced)

Selects the FRU fault LEDs that match the specified FRU type.

-fru-slot <integer> - FRU Slot (privilege: advanced)

Selects the FRU fault LEDs that match the specified slot.

[-fru-state {on|off|unknown}] - FRU State (privilege: advanced)

Specifies the target state for the FRU fault LED.

Examples

Turn off DIMM 3's FRU fault LED.

```
cluster1::*> system controller fru led modify -node node1 -fru-id dimm
-fru-slot 3 -fru-state off
```

The example below turns on all PCI FRU fault LEDs.

```
cluster1::*> system controller fru led modify -node node1 -fru-id pci -fru
-slot * -fru-state on
```

Related Links

- [system controller fru led show](#)

system controller fru led show

Display the status of FRU LEDs

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller fru led show` command displays information about the current state of the controller and IOXM FRU fault LEDs.

A FRU (Field Replaceable Unit) is any piece of the system that is designed to be easily and safely replaced by a field technician.

Both the controller and IOXM FRUs have a number of internal FRUs for which there are corresponding fault LEDs. In addition, there is a summary FRU fault LED on the external face-plate of both the controller and IOXM; labeled with a "!". A summary fault LED will be on when any of the internal FRU fault LEDs are on.

All FRU fault LEDs are amber in color. However, not all amber LEDs in the system are FRU fault LEDs. Externally visible fault LEDs are labeled with a "!" and internal FRU fault LEDs remain on, even when the controller or IOXM is removed from the chassis.

FRUs are identified by a FRU ID and slot tuple. FRU IDs include: DIMMs, cards in PCI slots, boot media devices, NV batteries and coin cell batteries. For each FRU ID, the FRUs are numbered 1 through N, where N is the number of FRUs of that particular type that exist in the controller or IOXM. Both controller and IOXM have a FRU map label for use in physically locating internal FRUs. The FRU ID/slot tuple used by the `system controller fru led show` command matches that specified on the FRU map label.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects FRU fault LEDs on the specified nodes.

[-fru-id <FRU LED key>] - FRU ID (privilege: advanced)

Selects the FRU fault LEDs that match the specified FRU type.

[-fru-slot <integer>] - FRU Slot (privilege: advanced)

Selects the FRU fault LEDs that match the specified slot.

[-fru-bay <text>] - FRU Bay (privilege: advanced)

Selects the FRU fault LEDs that match the specified bay.

[*-fru-state* {on|off|unknown}] - FRU State (privilege: advanced)

Selects the FRU fault LEDs that match the specified status.

[*-lit-by* <text>] - Lit By (privilege: advanced)

Selects the FRU fault LEDs that were lit by the specified source.

Examples

List the current state of all FRU fault LEDs.

```
cluster1::*> system controller fru led show
Node           FRU Type      Bay Slot State  Lit By
-----
host1
               controller  A 1    on    SP
               ioxm       B 1    off   -
               pci        - 1    off   -
               pci        - 2    off   -
               pci        - 3    off   -
               pci        - 4    off   -
               pci        - 5    off   -
               pci        - 6    off   -
               dimm-nv    - 1    off   -
               dimm-nv    - 2    off   -
               dimm      - 1    off   -
               dimm      - 2    off   -
               dimm      - 3    off   -
               dimm      - 4    off   -
               identify  - 1    off   -

5 entries were displayed.
```

The example below displays the status of only a specific FRU.

```
cluster1::*> system controller fru led show -node host1 -fru-id controller
-fru-slot 1
Node           FRU Type      Bay Slot State  Lit By
-----
host1
               controller  A 1    off   -
```

system controller ioxm show

Displays IOXM Device Health Status

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller ioxm show` command displays the details of the IO expansion modules (IOXMs) that are connected to the nodes in a cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about the IOXMs:

- Node name
- Display name
- Is IOXM present?
- Power status
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

`{ [-fields <fieldname>, ...]`

Selects the fields that you specify.

`| [-instance] }`

Displays detailed information for all the IOXMs.

`[-node {<nodename>|local}] - Node`

Selects the IOXM that is connected to the specified node.

`[-chassis-config {c-i|c-c|c-b}] - Controller-IOXM or Controller-Controller or Controller-Blank`

Selects the IOXMs with the specified chassis configuration.

`[-is-present {present|not-present}] - IOXM Presence`

Selects the IOXMs that are connected and detected (`present`) or connected but not detected (`not-present`).

`[-power {good|bad}] - Power to IOXM`

Selects the IOXMs with the specified power state.

`[-display-name <text>] - Display Name`

Selects the IOXMs with the specified display name.

`[-unique-name <text>] - Unique Name`

Selects the IOXM with the specified unique name.

`[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name`

Selects the IOXMs with the specified health monitor.

`[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - IOXM Health`

Selects the IOXMs with the specified health monitor status.

Examples

The example below displays the information of all the IOXMs that are connected to the nodes in a cluster.

```
cluster1::> system controller ioxm show
```

Node	Display Name	Is-Present?	Power	Status
node1	IOXM	present	good	ok
node2	IOXM	present	good	ok

The example below displays detailed information of an IOXM that is connected to a node.

```
cluster1::> system controller ioxm show -instance -node node1
Node: node1
Controller-IOXM or Controller-Controller or Controller-Blank: c-i
  IOXM Presence: present
  Power to IOXM: good
  Display Name: node1/IOXM
  Unique Name: 8006459930
Health Monitor Name: controller
  IOXM Health: ok
```

system controller location-led modify

Modify the location LED state of a controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller location-led modify` command modifies the current state of the location LED. When lit, the location LED can help you find the controller in the data center.

There is a blue location LED on every controller and on the front of the chassis. When you turn on the location LED for either controller, the chassis location LED automatically turns on. When both controller location LEDs are off, the chassis location LED automatically turns off.

After the location LED is turned on, it stays illuminated for 30 minutes and then automatically shuts off.

Parameters

-node {<nodename>|local} - Node

Selects the location LED on the specified filers.

[`-state {on|off}`] - LED State

Modifies the state of the location LED on the filer.

Examples

The following example turns on the location LED:

```
cluster1::*> system controller location-led modify -node node1 -state on
```

Turn off Location LED.

```
cluster1::*> system controller location-led modify -node node1 -state off
```

system controller location-led show

Display the location LED state on controllers

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller location-led show` command shows the current state of the location LED. When lit, the location LED can help you find the controller in the data center.

There is a blue location LED on every controller and on the front of the chassis. When you turn on the location LED for either controller, the chassis location LED automatically turns on. When both controller location LEDs are off, the chassis location LED automatically turns off.

After the location LED is turned on, it stays illuminated for 30 minutes and then automatically shuts off.

Parameters

{ [`-fields <fieldname>`,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [`-instance`] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-node {<nodename>|local}`] - Node

Selects the location LED on the specified filers.

[`-state {on|off}`] - LED State

Displays the location LED's status.

Examples

The following example lists the current state of the location LED:

```
cluster1::*> system controller location-led show
Node                Location LED State
-----
node1               Off
node2               Off
```

system controller memory dimm show

Display the Memory DIMM Table

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller memory dimm show` command displays information about the DIMMs in all the nodes in the cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about all the DIMMs in the cluster:

- Node
- DIMM name
- Uncorrectable ECC error count
- Correctable ECC error count
- CECC Alert Method
- CPU socket
- Channel
- Slot number
- Health monitor status
- Failure reason

To display more details, use the `-instance` parameter.

Parameters

{ [-fields <fieldname>,...]

Selects the fields that you specify.

| [-instance] }

Displays detailed information about the DIMMs in all the controllers in the cluster.

[-node {<nodename>|local}] - Node

Selects information about the DIMMs in the specified node.

[-pds-id <integer>] - DIMM ID

Selects information about the DIMMs with the specified DIMM ID.

[-slotname <text>] - Slot Name

Selects information about the DIMMs with the specified slot name.

[-socket <integer>] - CPU Socket

Selects information about the DIMMs with the specified socket ID.

[-channel <integer>] - Channel

Selects information about the DIMMs with the specified channel number.

[-slot-no <integer>] - Slot Number on a Channel

Selects information about the DIMMs with the specified slot number.

[-serial <text>] - Serial Number

Selects information about the DIMMs with the specified serial number.

[-part-no <text>] - Part Number

Selects information about the DIMMs with the specified part number.

[-cecc-count <integer>] - Correctable ECC Error Count

Selects information about the DIMMs with the specified correctable ECC error count.

[-uecc-count <integer>] - Uncorrectable ECC Error Count

Selects information about the DIMMs with the specified uncorrectable ECC error count.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects information about the DIMMs with the specified health monitor.

[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects information about the DIMMs with the specified health monitor status.

[-name <text>] - Unique Name of DIMM

Selects information about the DIMMs with the specified unique name.

[-display-name <text>] - Display Name for the DIMM

Selects information about the DIMMs with the specified display name.

[-alt-cecc-method <text>] - CECC Alert Method

Selects information about the DIMMs with the specified CECC error alert method.

[-alt-cecc-dimm {true|false}] - Replace DIMM

Selects information about the DIMMs with the specified replace DIMM value.

[-failure-reason <text>] - Failure Reason

Selects information about the DIMMs with the specified failure reason.

Examples

The example below displays information about the DIMMs in all the nodes in the cluster.

```
cluster1::> system controller memory dimm show
          DIMM      UECC      CECC      Alert      CPU
Slot      Failure
Node      Name      Count      Count      Method Socket Channel
Number  Status      Reason
-----  -
          node1      DIMM-1      0          0 bucket      0          0
0 unknown disabled
          node1      DIMM-NV1     0          0 bucket      0          1
1 ok
          node2      DIMM-1      1          0 bucket      0          0
0 ok
          node3      DIMM-NV1     0          0 bucket      0          1
1 ok
4 entries were displayed.
```

The example below displays detailed information about a specific DIMM in a specific controller.

```
cluster1::> system controller memory dimm show -instance -node node1 -pds
-id 1
          Node: node1
          DIMM ID: 1
          Slot Name: DIMM-1
          CPU Socket: 0
          Channel: 0
          Slot Number on a Channel: 0
          Serial Number: AD-01-1306-2EA01E9A
          Part Number: HMT82GV7MMR4A-H9
          Correctable ECC Error Count: 0
          Uncorrectable ECC Error Count: 0
          Health Monitor Name: controller
          Status: unknown
          Unique Name of DIMM: DIMM-1
          Display Name for the DIMM: DIMM-1
          CECC Alert Method: bucket
          Replace DIMM: false
          Failure Reason: disabled
```


system controller nvram-bb-threshold show

Display the controller NVRAM bad block threshold

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller nvram-bb-threshold show` command displays the threshold for the NVRAM bad block counts for a node.

system controller pci show

Display the PCI Device Table

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller pci show` command displays details of the PCI devices present in all of the nodes in a cluster. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported. By default, the command displays the following information about the PCI devices:

- Node name
- Display name
- Correctable error count
- Functional link width
- Functional link speed
- Health monitor status

To display more details, use the `-instance` parameter.

Parameters

`{ [-fields <fieldname>,...]`

Selects the fields that you specify.

`| [-instance] }`

Displays detailed information for all of the PCI devices.

`[-node {<nodename>|local}] - Node`

Selects the PCI devices that are present in the specified node.

`[-bus-number <integer>] - Bus Number`

Selects the PCI devices with the specified bus number.

[-device-number <integer>] - Device Number

Selects the PCI devices with the specified device number.

[-function-number <integer>] - Function Number

Selects the PCI devices with the specified function number.

[-slot-number <integer>] - Slot Info

Selects the PCI devices with the specified slot number.

[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name

Selects the PCI devices monitored by the specified health monitor.

[-vendor-id <Hex Integer>] - Vendor ID

Selects the PCI devices with the specified vendor ID.

[-device-id <Hex Integer>] - Device ID

Selects the PCI devices with the specified device ID.

[-physical-link-width <integer>] - Physical Link Width

Selects the PCI devices with the specified physical link width.

[-functional-link-width <integer>] - Functional Link Width

Selects the PCI devices with the specified functional link width.

[-physical-link-speed <text>] - Physical Link Speed(GT/s)

Selects the PCI devices with the specified physical link speed.

[-functional-link-speed <text>] - Functional Link Speed(GT/s)

Selects the PCI devices with the specified functional link speed.

[-unique-name <text>] - Unique Name

Selects the PCI devices with the specified unique name.

[-corr-err-count <integer>] - Correctable Error Count

Selects the PCI devices with the specified correctable error count.

[-health {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects the PCI devices with the specified health monitor status.

[-display-name <text>] - Display Name

Selects the PCI devices with the specified display name.

[-cerr-diff <integer>] - Correctable Error Difference

Selects the PCI devices with the specified difference in correctable error count.

Examples

The example below displays the information about the PCIe devices present in all of the nodes in the cluster.

```

cluster1::> system controller pci show
          Display          Correctable Functional Functional
Node      Name            Error Count Link Width Link Speed
Status
-----
cluster1-01  Ontap PCI Device 0          0           4           5GT/s
ok
cluster1-02  Ontap PCI Device 4          0           4           5GT/s
ok

```

The example below displays detailed information about a PCIe device in a node.

```

cluster1::> system controller pcie show -instance -node cluster1-01 -bus
-number 1
Node: cluster1-01
          Bus Number: 1
          Device Number: 0
          Function Number: 0
          Slot Info: 0
          Health Monitor Name: controller
          Vendor ID: 11f8
          Device ID: 8001
          Physical Link Width: 4
          Functional Link Width: 4
          Physical Link Speed(GT/s): 5GT/s
          Functional Link Speed(GT/s): 5GT/s
          Unique Name: ontap0@pci0:1:0:0
          Correctable Error Count: 0
          Status: ok
          Display Name: Ontap PCI Device 0
          Correctable Error Difference: 0

```

system controller pcicerr threshold modify

Modify the Node PCIe error alert threshold

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller pcicerr threshold modify` command modifies node-wide PCIe correctable error threshold counts in the cluster.

Parameters

`[-pcie-cerr-threshold <integer>]` - Corr. Error Limit

The PCIe error threshold count that would trigger an alert if exceeded.

`[-nvram-bb-threshold <integer>]` - NVRAM Bad Block limit

The NVRAM bad block threshold count that would trigger an alert if exceeded.

Examples

The example below displays the information about setting node-wide PCIe error threshold count in the cluster:

```
cluster1::> system controller threshold modify -pcie-cerr-threshold 100
```

system controller pcicerr threshold show

Display the Node PCIe error alert threshold

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller pcicerr threshold show` command displays information about node-wide PCIe correctable error threshold counts in the cluster.

Examples

The example below displays the information about node-wide PCIe error threshold count in the cluster:

```
cluster1::> system controller pcicerr threshold show

PCIe Error Threshold
-----
                        200
```

system controller platform-capability show

Display platform capabilities

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller platform-capability show` command displays information about all platform capabilities for each controller in the cluster. By default, the command displays the following information about all controllers in the cluster:

- Controller Name
- Capability ID
- Capability Supported?
- Capability Name

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

Displays detailed information about all controllers in the cluster.

[-node {<nodename>|local}] - Node

Selects information about the specified controller.

[-capability-id <integer>] - Capability ID

Selects the desired capability ID.

[-supported <text>] - Supported?

Selects the desired capability support state (true or false).

[-name <text>] - Capability Name

Selects the desired capability name.

Examples

The following example displays platform capability information for the controller:

```
cluster1::> system controller platform-capability show
Node                Capability ID Supported? Capability Name
-----
or-099-diag-01
                   0                false    CAP_CMCI_ENABLED
                   1                false    CAP_HA_CONFIG_ONLY
                   2                true     CAP_SUPPORT_CARD_FRU
                   3                true     CAP_SCORPIO_EN
                   4                false    CAP_NVD_EN
                   5                false    CAP_ENABLE_HPET
                   6                false    CAP_VERIFY_ACPI_TABLE
7 entries were displayed.
```

system controller replace cancel

Cancel ongoing controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace cancel` command is used to cancel a controller replacement that is in a paused state (paused-on-request, paused-on-error or paused-for-intervention). The update cannot be canceled if it is not in a paused state.

Examples

The following example displays a cancel operation:

```
cluster1::> system controller replace cancel
Warning: The controller replacement will be canceled and any changes will
have to be reverted manually.
Do you want to continue? {y|n}: y
Controller replacement canceled successfully.
```

system controller replace pause

Pause ongoing controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace pause` command is used to pause a currently running replacement. The operation pauses at the next predefined update stage (for example, after finishing the current task it will pause the next restartable task) which might take some time to reach. When the update reaches the pause point, it transitions into the pause-on-request state.

Examples

The following example displays pause operation:

```
cluster1::> system controller replace pause

A pause requested for Controller Replacement operation.
The current task will continue and the next restartable task will be
paused.
```

system controller replace resume

Resume paused controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace resume` command is used to resume an update that is currently in one of paused-on-request, paused-on-error or paused-for-intervention states. If the update is not paused then an error is returned.

Examples

The following example shows a resume operation:

```
cluster1::> system controller replace resume

Controller replacement resumed successfully.
```

system controller replace show-details

Display detailed status of controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace show-details` command displays detailed information about the currently running and previously run non-disruptive controller replacement operations. The command displays the following information:

- Phase
- Node
- Task name
- Task status
- Error message

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-operation-identifier {None|Controller-replacement}] - Operation Identifier (privilege: advanced)

Specifies the NDO operation identifier.

[-task-identifier <integer>] - Task Identifier (privilege: advanced)

Specifies the identification number of the task.

[-node <nodename>] - Node That Performs Operation (privilege: advanced)

Specifies the node that is to be replaced.

[-task-phase {None|Initialization|Precheck|Commit|Collect-info|Preop|Resource-release|Verification|Resource-regain|Postop|Uncommit|Postcheck|Cleanup}] - Task Phase (privilege: advanced)

Specifies the phase of the operation.

[-task-name <text>] - Name of the Task (privilege: advanced)

Specifies the name of the task.

[-task-status {None|Ready-to-start|In-progress|Waiting|Paused-on-error|Paused-for-intervention|Paused-on-request|Completed-on-first-node|Completed|Failed|Pause_req|Canceled}] - Status of the Task (privilege: advanced)

Specifies the status of the task.

[-task-error <text>] - Error During the Task Execution (privilege: advanced)

Specifies the error occurred.

[-task-recovery-action <text>] - Action to Recover from Error (privilege: advanced)

Specifies the action to be taken in case of error.

Examples

The following example displays detailed information about the non-disruptive replacement operation:


```

cluster1::*> system controller replace show-details
  Task Phase          Node          Task Name          Operation-State
  -----
Precheck            node1          Cluster Health Check  Completed
                  MCC Cluster Check  Completed
                  Aggr Relocation    Completed
                  Status Check
                  Model Name Check   Completed
                  Cluster Quorum Check Completed
                  Image Version Check Completed
                  HA Status Check   Completed
                  Aggregate Status  Completed
                  Check
                  Disk Status Check   Completed
                  Data LIF Status Check Completed
                  Cluster LIF Status  Completed
                  Check
                  ASUP Status Check   Completed
                  CPU Utilization Check Completed
                  Aggr Reconstruction  Completed
                  Check
                  Node Affinity Job    Completed
                  Check
Collect-info        node1          Verify Details      Paused-for-
intervention
                  node2          Verify Details      Paused-for-
intervention
17 entries were displayed.

```

system controller replace show

Display status of controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace show` command displays overall information about the currently running, or previously run controller replacement operation. The command displays the following information:

- Operation Status
- Error message
- Recommended action

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>`, ... parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

Examples

The following example displays information about automated nondisruptive operation:

```
cluster1::*> system controller replace show
Node                Status                Error-Action
-----
node1                Paused-for-intervention  Follow the instructions given
in
                                                                Step Details
node2                None
Step Details:
-----
Controller replacement operation has been paused for user intervention.
Collect the following info from the current node:
 1. vserver services name-service dns show
 2. service-processor show -node * -instance
 3. network port ifgrp show
 4. network port vlan show
 5. network interface failover-groups show
 6. storage array config show -switch switchname
 7. storage encryption disk show
2 entries were displayed.
```

system controller replace start

Start controller replacement

Availability: This command is available to *cluster* administrators at the *advanced* privilege level.

Description

The `system controller replace start` command is used to initiate a controller-replacement operation. The update is preceded by a validation of the HA pair to ensure that any issues that might affect the update are identified.

There are predefined points in the update when the update can be paused (either requested by the user or by

the operation in case of an error or for manual intervention).

Parameters

[-nodes <nodename>, ...] - Nodes for Controller Replacement (privilege: advanced)

Specifies the nodes that are to be replaced.

[-simulate <true>] - Simulate Controller Replacement (privilege: advanced)

Dry run the operation. Checks for all validations.

[-skip-metrocluster-check {true|false}] - Skip Metrocluster Check before Replacement (privilege: advanced)

Skips the DR mirror offline check when performing Metrocluster validation. In 4-node Metrocluster configuration, if controller replacement is already complete one one site, then the partner site should replace its controllers by setting this parameter to `true` . The default value is `false`.

[-nso {true|false}] - Select NSO Procedure (privilege: advanced)

Follow the negotiated switchover switchback based controller replacement procedure for 4 node MCC FC system. To select ARL based procedure for 4 node MCC FC , this parameter needs to be set to `false` .This parameter is only honored in MCC configuration. Hence, setting or unsetting of this parameter has no effect in HA configuration. The default value is `true` for MCC FC.

Examples

The following example shows the replacement operation:

```
cluster1::> system controller replace start -nodes node1,node2 -simulate true
```

Warning: 1. Current version of node is 9.4.0

Before starting controller replacement, please ensure that the new controllers are in the version 9.4.0

2. Verify that NVMEM or NVRAM batteries of the new nodes are charged, and charge them if they are not. You need to physically check the new nodes to see if

the NVMEM or NVRAM batteries are charged. You can check the battery status either by connecting to a console or using SSH, logging into the Service Processor

(SP) for your system, and use the system sensors to see if the battery has a sufficient charge.

Attention: Do not try to clear the NVRAM contents. If there is a need to clear the contents of NVRAM, contact NetApp technical support.

3. If you are replacing the controllers with an used one, please ensure to run wipeconfig before controller replacement

4. Current model name is FAS8040

Before starting the operation, please ensure that the new controller model is supported for controller replacement.

Do you want to continue? {y|n}: y

Controller replacement: Prechecks in progress.....

Controller replacement has been paused for user intervention.

Please collect the following info from the current node:

```
vserver services name-service dns show
```

```
network interface show -curr-node node -role cluster,intercluster,node-  
mgmt,cluster-mgmt
```

```
network port show -node node -type physical
```

```
service-processor show -node * -instance
```

```
network fcp adapter show -node node
```

```
network port ifgrp show
```

```
network port vlan show
```

```
system node show -instance -node node
```

```
run -node node sysconfig
```

```
storage aggregate show -node node
```

```
volume show -node node
```

```
network interface failover-groups show
```

```
storage array config show -switch switchname
```

```
system license show -owner node
```

```
storage encryption disk show
```

system controller service-event delete

Manually clear a selected service event

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller service-event delete` command removes the service event from the list and extinguishes all related FRU attention LEDs.

In some cases, where the underlying fault condition remains, the service event might be reported again, causing it to reappear in the list. In such cases, it is necessary to remedy the underlying fault condition in order to clear the service event.

Parameters

-node {<nodename>|local} - Node

Selects service events on the specified nodes.

-event-id <integer> - Service Event ID

Selects the service events that match the specified event identifier. Together with the node, this field uniquely identifies the row to delete. Use the [system controller service-event show](#) command to find the event identifier for the service event to delete.

Examples

The following example lists the currently active service events. Then, using the listed Service Event ID, the service event is deleted:

```
cluster1::> system controller service-event show

Node                ID  Event Location                Event Description
-----
-----
plata4-1a           1   DIMM in slot 1 on Controller A  Uncorrectable ECC

cluster1::> system controller service-event delete -event-id 1
```

Related Links

- [system controller service-event show](#)

system controller service-event show

Display the active service events causing attention LEDs to be lit

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller service-event show` command displays one or more events that have been detected by the system for which a physical service action might be required. Physical service actions sometimes involve replacing or re-seating misbehaving FRUs. In such cases FRU attention LEDs will be illuminated to assist in physically locating the FRU in need of attention. When the FRU in question is contained within another FRU, both the inner and outer FRU attention LEDs will be lit. It creates a path of LEDs that starts at the chassis level and leads to the FRU in question. For example, if a DIMM is missing from the controller motherboard, the storage OS will detect this and log a service event whose location is the DIMM slot on the controller. The DIMM slot LED, controller LED and chassis LED will all be lit to create a path of LEDs to follow.

FRU Attention LEDs that are not visible from outside of the system (e.g. those on the controller motherboard such as DIMMs, boot device etc.) will remain on for a few minutes, even after power is removed from the containing FRU. As such, when the controller is removed from the chassis, a DIMM slot FRU attention LED will remain on, helping to locate the FRU in need of attention.

Generally, service events are cleared automatically when the issue is resolved. The corresponding FRU attention LEDs are extinguished accordingly. In cases where the service event request is caused by an environmental issue, it might be necessary to manually remove the service event from the list. This can be done using the [system controller service-event delete](#) command.

Parameters

{ [-fields <fieldname>,...]

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

| [-instance] }

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[-node {<nodename>|local}] - Node

Selects service events on the specified nodes.

[-event-id <integer>] - Service Event ID

Selects the service events that match the specified event identifier. Together with the node, this field uniquely identifies the row for use with the [system controller service-event delete](#) command

[-event-loc <text>] - Location

Selects the service events that match the specified event location.

[-event-desc <text>] - Description

Selects the service events that match the specified event description.

[-event-timestamp <text>] - Timestamp

The time that the event occurred, recorded by the Service Processor

Examples

The following example lists the currently active service events.

```
cluster1::> system controller service-event show
```

Node	ID	Event Location	Event Description
plata4-1a	1	DIMM in slot 1 on Controller A	Uncorrectable ECC

Related Links

- [system controller service-event delete](#)

system controller slot module insert

Add a module on the controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller slot module insert` command adds a module on the controller.

Parameters

-node {<nodename>|local} - Node

Selects the PCIe modules that are present in the specified node.

-slot <text> - Slot Number

Selects the PCIe modules present in the specified slot or slot-subslot combination.

Examples

The following example adds a module in the local node:

```
p2i030::> system controller slot module insert -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be
powered
      on and initialized.
Do you want to continue? {y|n}:y
The module has been successfully powered on, initialized and placed into
service.

p2i030::>
```

system controller slot module remove

Remove a module on the controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller slot module remove` command removes a module on the controller.

Parameters

-node {<nodename>|local} - Node

Selects the PCIe modules that are present in the specified node.

-slot <text> - Slot Number

Selects the PCIe modules present in the specified slot or slot-subslot combination.

Examples

The following example removes a module in the local node:

```
p2i030::> system controller slot module remove -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be
         powered off for removal.
Do you want to continue? {y|n}: y
The module has been successfully removed from service and powered off. It
can now be safely removed.

p2i030::>
```

system controller slot module replace

Power off a module on the controller for replacement

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller slot module replace` command powers off a module on the controller for replacement.

Parameters

-node {<nodename>|local} - Node

Selects the PCIe modules that are present in the specified node.

-slot <text> - Slot Number

Selects the PCIe modules present in the specified slot or slot-subslot combination.

Examples

The following example powers off a module in the local node:

```
p2i030::> system controller slot module replace -node local -slot 1

Warning: IO_CARRIER_NIANTIC_NIC module in slot 1 of node p2i030 will be
powered
      off for replacement.
Do you want to continue? {y|n}: y
The module has been successfully powered off. It can now be safely
replaced. After the replacement module is inserted, use the "system
controller slot module insert" command to place the module into service.

p2i030::>
```

system controller slot module show

Display hotplug status of a module on the controller

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller slot module show` command displays hotplug status of a module on the controller. The command displays the following information about the PCIe modules:

- Node
- Slot
- Module
- Status

To display more details, use the `-instance` parameter.

Parameters

```
{ [-fields <fieldname>,...]
```

If you specify the `-fields <fieldname>, ...` parameter, the command output also includes the specified field or fields. You can use `'-fields ?'` to display the fields to specify.

```
| [-instance ] }
```

If you specify the `-instance` parameter, the command displays detailed information about all fields.

[`-node` {<nodename>|local}] - Node

Selects the PCIe modules that are present in the specified node.

[`-slot` <text>] - Slot Number

Selects the PCIe modules present in the specified slot or slot-subslot combination.

[`-status` <text>] - Module Status

Selects hotplug status for PCIe modules.

[`-card` <text>] - Module Name

Selects module name for PCIe modules.

Examples

The following example displays hotplug status of PCI modules found in the local node:

```
::> system controller slot module show -node local
Node          Slot  Module                               Status
-----
localhost     1     IO_CARRIER_NIANTIC_NIC             powered-on
localhost     2     IO_4X_10GBT_INTL_NIC                 powered-on
localhost     3     IO_4X_12Gb_PMC_SAS                   powered-on
localhost     4     IO_4X_10GBE_16GFC_QLGC_CNA          powered-on
localhost     5     IO_4X_12Gb_PMC_SAS                   powered-on
localhost     6     NVRAM10                              hotplug-not-
supported
localhost     6-1   empty                                empty
localhost     6-2   empty                                empty
localhost     7     IO_4X_12Gb_PMC_SAS                   powered-on
localhost     8     IO_4X_10GBT_INTL_NIC                 powered-on
localhost     9     IO_4X_12Gb_PMC_SAS                   powered-on
localhost    10     IO_4X_12Gb_PMC_SAS                   powered-on
localhost    11     IO_4X_12Gb_PMC_SAS                   powered-on
13 entries were displayed.

::>
```

system controller sp config show

Display the Service Processor Config Table

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller sp config show` command displays the following configuration information of the service processor for all nodes in the cluster:

- Node name
- Service processor status
- Service processor firmware version
- Booted firmware version
- Service processor configuration status
- Physical Ethernet link status of service processor
- Health monitor status

To display more details, use the `-instance` parameter. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported.

Parameters

`{ [-fields <fieldname>,...]`

Selects the field that you specify.

`| [-instance] }`

Displays detailed configuration information of the service processor.

`[-node {<nodename>|local}] - Node`

Use this parameter to list the service processor configuration of the specific node.

`[-version <text>] - Firmware Version`

Selects the service processor configuration with the specified firmware version.

`[-boot-version {primary|backup}] - Booted Version`

Selects the service processor configuration with the specified version of the currently booted partition.

`[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name`

Selects the service processor configuration with the specified monitor name.

`[-sp-status {online|offline|sp-daemon-offline|node-offline|degraded|rebooting|unknown|updating}] - SP Status`

Selects the service processor configuration with the specified status of service processor.

`[-sp-config {true|false}] - Auto Update Configured`

Selects information about the service processor with the specified configuration status of the service processor.

`[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status`

Selects the service processor configuration information with the specified service processor status.

[`-link-status {up|down|disabled|unknown}`] - Public Link Status

Selects the service processor configuration with the specified physical ethernet link status.

[`-name <text>`] - Display Name

Selects the service processor configuration with the specified unique name.

Examples

The example below displays configuration of the service processor in all the nodes in the cluster:

```
cluster1::> system controller sp config show
      Firmware   Booted   Auto Update   SP       Link
Node  Version     Version   Configured   Status   Status   Status
----  -
node1  2.2.2       primary   true         online   up        ok
node2  2.2.2       primary   true         online   up        ok
```

The example below displays configuration of the service processor of a particular node in detail:

```
cluster1::> system controller sp config show -instance -node node1
      Node: node1
      Firmware Version: 2.2.2
      Booted Version: primary
      Health Monitor Name: controller
      SP Status: online
      Auto Update Configured: true
      Status: ok
      Public Link Status: up
      Display Name: SP Config
```

system controller sp upgrade show

Display the Service Processor Upgrade Table

Availability: This command is available to *cluster* administrators at the *admin* privilege level.

Description

The `system controller sp upgrade show` command displays the following information about the service processor firmware of all the nodes in the cluster:

- Node name
- Is new firmware available?
- Is autoupdate enabled?
- Status of autoupdate

- Health monitor status

To display more details, use the `-instance` parameter. These commands are available for 80xx, 25xx and later systems. Earlier models are not supported.

Parameters

`[-fields <fieldname>,...]`

Selects the fields that you specify.

`[-instance]`

Displays detailed upgrade information of the service processor.

`[-node {<nodename>|local}] - Node`

Use this parameter to list the upgrade information of the service processor on the specified node.

`[-new-fw-avail {true|false}] - New Firmware Available`

Selects the information of the service processors which have new firmware available.

`[-new-fw-version <text>] - New Firmware Version`

Selects the information about service processors with the specified firmware version.

`[-auto-update {true|false}] - Auto Update`

Selects the information about service processors with the specified state.

`[-auto-update-stat {installed|corrupt|updating|auto-updating|none}] - Auto Update Status`

Selects the information about service processors with the specified auto update status.

`[-auto-update-sttime <MM/DD/YYYY HH:MM:SS>] - Auto Update Start Time`

Selects the information about service processors with the specified start time.

`[-auto-update-entime <MM/DD/YYYY HH:MM:SS>] - Auto Update End Time`

Selects the information about service processors with the specified end time.

`[-auto-update-per <integer>] - Auto Update Percent Done`

Selects the information about service processors with the specified auto update percentage completed.

`[-auto-update-maxret <integer>] - Auto Update Maximum Retries`

Selects the information about service processors with the specified maximum number of retries.

`[-auto-update-curret <integer>] - Auto Update Current Retries`

Selects the information about service processors with the specified number of current retries.

`[-auto-update-prevstat {failed|passed}] - Previous AutoUpdate Status`

Selects the information about service processors with the specified automatic update status.

`[-monitor {node-connect|system-connect|system|controller|chassis|cluster-switch|example|ethernet-switch}] - Health Monitor Name`

Selects the information about service processors with the specified monitor name.

[-status {ok|ok-with-suppressed|degraded|unreachable|unknown}] - Status

Selects the information about service processors with the specified health monitor status.

[-name <text>] - Display Name

Selects the information about service processors with the specified display name.

Examples

The example below displays service processor upgrade information for all nodes in the cluster:

```
cluster1::> system controller sp upgrade show
New    Firmware      Auto Update  Auto Update
Node  Available          Feature      Status      Status
----  -
node1 false             true         installed   ok
node2 false             true         installed   ok
2 entries were displayed.
```

The example below displays the detailed service processor upgrade information for a specific node:

```
cluster1::> system controller sp upgrade show -instance -node node1
Node: node1
  New Firmware Available: false
  New Firmware Version: Not Applicable
  Auto Update: true
  Auto Update Status: installed
  Auto Update Start Time: Thu Oct 20 20:06:03 2012 Etc/UTC
  Auto Update End Time: Thu Oct 20 20:09:19 2012 Etc/UTC
  Auto Update Percent Done: 0
  Auto Update Maximum Retries: 5
  Auto Update Current Retries: 0
  Previous AutoUpdate Status: passed
  Health Monitor Name: controller
  Status: ok
  Display Name: SP Upgrade
```

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.