



## Storage Inventory

---

- [Local Disk Locator LED Status, on page 1](#)
- [Toggling the Local Disk Locator LED On and Off, on page 1](#)
- [Custom LED Status with Advanced VMD on NVMe, on page 2](#)
- [NVMe-optimized M5 Servers, on page 5](#)
- [NVMe PCIe SSD Inventory, on page 6](#)
- [Viewing NVMe PCIe SSD Storage Inventory, on page 6](#)

## Local Disk Locator LED Status

The local disk locator LED is located on the slot where you insert the local disk. This LED identifies where a specific disk is inserted in a blade or rack server. The locator LED is useful for maintenance, when you need to remove a disk from among many disks in a server.

You can successfully turn on or off the local disk locator LED when:

- The server is powered on. UCS Manager generates an error if you attempt to turn the locator LED on or off when the server is powered off.
- The CIMC version is UCS Manager 3.1 or higher.
- The RAID controller supports the out-of-band (OOB) storage interface.

When Intel Volume Management Device (VMD) for NVMe is enabled, you can also configure blinking patterns for the LEDs on NVMe-managed devices to show drive status. VMD-enabled drives identified by a failure ID blink pattern can be hot-plugged without a system shutdown.

## Toggling the Local Disk Locator LED On and Off

### Before you begin

#### On and Off

- Ensure the server on which the disk is located is powered on. If the server is off, you are not able to turn on or off the local disk locator LED.

## Procedure

---

- Step 1** In the **Navigation** pane, click the **Equipment** tab.
- Step 2** On the **Equipment** tab, expand **Equipment > Rack Mounts > Servers > Server Number**.
- a) For Rack-mounted servers, go to **Rack Mounts Servers Server Number**.
  - b) For Blade servers, go to **> Sensors > Storage Server Number**.
- Step 3** In the **Work** area, click the **Inventory > Storage > Disks** tabs.  
The Storage Controller inventory appears.
- Step 4** Click a disk.  
The disk details appear.
- Step 5** In the Actions area, click **Turn on Locator LED** or **Turn off Locator LED**.  
The **Locator LED** state appears in the **Properties** area.
- Step 6** Click **Save Changes**.
- 

# Custom LED Status with Advanced VMD on NVMe

## Custom Blinking Patterns

VMD allows you to customize LED blinking patterns on PCIe NVMe drives to better identify failing drives. The tables below provide some brief guidelines for customized blinking. As individualized patterns are programmable, these tables provide only representative guidelines.

*Table 1: LED Blinking Patterns: Windows*

Status LED	Behavior	Options
"Activate LED"	Identifies a specific device in an enclosure by blinking the status LED of that drive in a designated pattern.	1-3600 seconds. Values outside this range default to 12 seconds. Default = 12 seconds

Status LED	Behavior	Options
Drive Failure	Indicates a drive that is in a degraded or failed state by lighting the status LED of that device in a defined failure pattern.	<p>The failure pattern is displayed until:</p> <ul style="list-style-type: none"> <li>• 1. It is physically removed.</li> </ul> <p>or</p> <p>the RAID volume, that contains the failed drive, is either deleted or physically removed.</p> <ul style="list-style-type: none"> <li>• 2. From the time when a non-failed drive that is part of a RAID volume is removed, or the failed drive is identified and removed. It remains in failure state until a new drive is inserted into the same slot or the platform is rebooted.</li> </ul> <p>Default = Option 1</p>
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in Rebuild state, the status LEDs blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume that is being rebuilt.	<p>1. Disabled (only on one drive)</p> <p>2. Enabled (on all drives)</p> <p>Default = Enabled</p>
Managed unplug	During a managed hot unplug, the status LED of the managed drive blinks in the defined Locate pattern until the drive is physically ejected.	None. Enabled by default.
RAID volume is migrating	During RAID volume migration, the status LEDs blink in the defined Rebuild pattern on all drives until the process is complete.	<p>1. Disabled (No Status LED Blinking)</p> <p>2. Enabled (Blinks Status LEDs)</p> <p>Default = Enabled</p>
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 2: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller <b>BLACKLIST</b>	<code>ledmon</code> will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	<p>Exclude controllers on the blacklist.</p> <p>Default = Support all controllers</p>

Status LED	Behavior	Options
RAID volume is initializing, verifying, or verifying and fixing <b>BLINK_ON_INIT</b>	Rebuild pattern on all drives in RAID volume (until initialization, verify, or verify and fix finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
RAID volume is rebuilding <b>REBUILD_BLINK_ON_ALL</b>	Rebuild pattern on a single drive to which RAID volume rebuilds	1. False/Disabled (on one drive) 2. True/Enabled (on all drives) Default = False/Disabled
Set <code>ledmon</code> scan interval <b>INTERVAL</b>	Defines the time interval between <code>ledmon sysfs scans</code> . The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is migrating <b>BLINK_ON_MIGR</b>	Rebuild pattern on all drives in RAID volume (until migration finishes).	1. True/Enabled (on all drives) 2. False/Disabled (no drives) Default = True/Enabled
Set <code>ledmon</code> debug level <b>LOG_LEVEL</b>	Corresponds with <code>-log-level</code> flag from <code>ledmon</code> .	Acceptable values are: quiet, error, warning, info, debug, all - 0 means 'quiet' and 5 means 'all' Default = 2
Set manage one RAID member or All RAID <b>RAID_MEMBERS_ONLY</b>	If the flag is set to <code>ledmon true</code> , will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False
Limited scans only to specific controllers <b>WHITELIST</b>	<code>ledmon</code> limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 3: LED Blinking Patterns: ESXi

Status LED	Behavior	Options
"Identify"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern.	None. Default is Off.
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.

# NVMe-optimized M5 Servers

Beginning with 3.2(3a), Cisco UCS Manager supports the following NVMe-optimized M5 servers:

- **UCSC-C220-M5SN**—The PCIe MSwitch is placed in the dedicated MRAID slot for UCS C220 M5 servers. This setup supports up to 10 NVMe drives. The first two drives are direct-attached through the riser. The remaining eight drives are connected and managed by the MSwitch. This setup does not support any SAS/SATA drive combinations.
- **UCSC-C240-M5SN**—The PCIe MSwitch is placed in the riser-2 at slot-4 for UCS C240 M5 servers. The servers support up to 24 drives. Slots 1-8 are the NVMe drives connected and managed by the MSwitch. The servers also support up to two NVMe drives in the rear and are direct-attached through the riser. This setup supports SAS/SATA combination with the SAS/SATA drives from slots 9-24. These drives are managed by the SAS controller placed in the dedicated MRAID PCIe slot.
- **UCS-C480-M5**—UCS C480 M5 servers support up to three front NVMe drive cages, each supporting up to eight NVMe drives. Each cage has an interposer card, which contains the MSwitch. Each server can support up to 24 NVMe drives (3 NVMe drive cages x 8 NVMe drives). The servers also support a rear PCIe Aux drive cage, which can contain up to eight NVMe drives managed by an MSwitch placed in PCIe slot-10.

This setup does not support:

- a combination of NVMe drive cages and HDD drive cages
- a combination of the Cisco 12G 9460-8i RAID controller and NVMe drive cages, irrespective of the rear Auxiliary drive cage



---

**Note** The UCS C480 M5 PID remains same as in earlier release.

---

The following MSwitch cards are supported in NVMe optimized M5 servers:

- **UCS-C480-M5 HDD Ext NVMe Card (UCSC-C480-8NVME)**—Front NVMe drive cage with an attached interposer card containing the PCIe MSwitch. Each server supports up to three front NVMe drive cages and each cage supports up to 8 NVMe drives. Each server can support up to 24 NVMe drives (3 NVMe drive cages x 8 NVMe drives).
- **UCS-C480-M5 PCIe NVMe Switch Card (UCSC-NVME-SC)**—PCIe MSwitch card to support up to eight NVMe drives in the rear auxiliary drive cage inserted in PCIe slot 10.



---

**Note** Cisco UCS-C480-M5 servers support a maximum of 32 NVMe drives (24 NVMe drives in the front + 8 NVMe drives in the rear auxiliary drive cage)

---

- **UCSC-C220-M5SN** and **UCSC-C240-M5SN** do not have separate MSwitch PIDs. MSwitch cards for these servers are part of the corresponding NVMe optimized server.

## MSwitch Disaster Recovery

You can recover a corrupted MSwitch and roll back to a previous working firmware.



**Note** If you have a setup with Cisco UCS C480 M5 Server, then MSwitch disaster recovery process can be performed only on one MSwitch at a time. If the disaster recovery process is already running for one MSwitch, then wait for it to complete. You can monitor the recovery status from FSM.

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** Expand **Rack-Mounts > Servers**.
  - Step 3** Expand the server for the which contains the MSwitch.
  - Step 4** In the **Work** pane, click **Inventory > Storage > Controller**.
  - Step 5** Select the MSwitch which you want to recover.
  - Step 6** Under the **General** tab, click **Disaster Recovery**.
    - Note** Do not reset the server during the disaster recovery process.
  - Step 7** You can monitor the recovery status from FSM.
- 

## NVMe PCIe SSD Inventory

Cisco UCS Manager GUI discovers, identifies, and displays the inventory of Non-Volatile Memory Express (NVMe) Peripheral Component Interconnect Express (PCIe) SSD storage devices. You can view the health of the storage devices in the server. NVMe with PCIe SSD storage devices reduce latency, increased input/output operations per second (IOPS), and lower power consumption compared to SAS or SATA SSDs.

## Viewing NVMe PCIe SSD Storage Inventory

### Procedure

- 
- Step 1** In the **Navigation** pane, click the **Equipment** tab.
  - Step 2** On the **Equipment** tab, expand **Equipment > Rack Mounts > Servers**.
  - Step 3** Click the **Inventory** tab.
  - Step 4** Do one of the following:
    - a) Click the **Storage** tab.
      - You view the list of NVMe PCIe SSD storage devices named **Storage Controller NVME ID number**.
      - You view the name, size, serial number, operating status, state and other details.

- b) Click the NVMe PCIe SSD storage device.

You see the following inventory details:

Name	Description
ID	The NVMe PCIe SSD storage device configured on the server.
Model	The NVMe PCIe SSD storage device model.
Revision	The NVMe PCIe SSD storage device revision.
RAID Support	Whether the NVMe PCIe SSD storage device is RAID enabled.
OOB Interface Support	Whether the NVMe PCIe SSD storage device support out-of-band management .
PCIe Address	<p>The NVMe PCIe SSD storage device on the virtual interface card (VIC).</p> <p><b>Note</b> PCIe Address does not come upon hot insertion of the NVMe card. To view this info, re-acknowledge the server.</p>
Number of Local Disks	The number of disks contained in the NVMe PCIe SSD storage device.
Rebuild Rate	Not applicable to NVMe PCIe SSD storage devices.
Vendor	The vendor that manufactured the NVMe PCIe SSD storage device.
PID	The NVMe PCIe SSD storage device product ID, also known as product name, model name, product number
Serial	The storage device serial number.

