Storage Inventory

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 MountsServers *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 VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior.The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink.

Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. 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 on the various platforms. 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
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.	The failure pattern is displayed until: • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 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. Default = Option 1
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.	Default = Enabled Can be: 1. Disabled (only on one drive) 2. Enabled (on all drives)

Status LED	Behavior	Options
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.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 2: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller **BLACKLIST**	`ledmon` will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing **BLINK_ON_INIT**	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
Set `ledmon` scan interval **INVERVAL**	Defines the time interval between `ledmon sysfs` scans. The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding **REBUILD_BLINK_ON_ALL**	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
RAID volume is migrating **BLINK_ON_MIGR**	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 ledmon debug level **LOG_LEVEL**	Corresponds with –log-level flag from `ledmon`.	Acceptable values are: quiet, error, warning, info, debug, all - 0 means ‘quiet’ and 5 means ‘all Default = 2

Status LED	Behavior	Options
Set manage one RAID member or All RAID **RAID_MEMBRES_ONLY**	If the flag is set to `ledmon` true, 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 **WHITELIST**	`ledmon` 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.

Note On B200 and B480 M5 blade servers, NVMe drives cannot be used directly with SAS controllers. Use an LSTOR-PT pass-through controller instead.

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.

Note Virtual Controller for direct CPU-attached NVME drives will be shown in UCSM only when it detects the NVME drives and such NVME switch details are retained between drive insertions. These NVME switch entries will be removed only when the server is decommissioned.

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.

Name	Description
OOB Interface Support	Whether the NVMe PCIe SSD storage device support out-of-band management .
PCIe Address	The NVMe PCIe SSD storage device on the virtual interface card (VIC). **Note** PCIe Address does not come upon hot insertion of the NVMe card. To view this info, re-acknowledge the server.
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.

Enabling Volume Management Device on UCS Storage

Enabling Intel® Volume Management Device

Volume Management Device (VMD) Setup

The Intel® Volume Management Device (VMD) is a tool that provides NVMe drivers to manage PCIe Solid State Drives attached to VMD-enabled domains. This includes Surprise hot-plug of PCIe drives and configuring blinking patterns to report status. PCIe Solid State Drive (SSD) storage lacks a standardized method to blink LEDs to represent the status of the device. With VMD, you can control LED indicators on both direct attached and switch attached PCIe storage using a simple command-line tool.

To use VMD, you must first enable VMD through a UCS Manager BIOS policy and set the UEFI boot options. Enabling VMD provides Surprise hot plug and optional LED status management for PCIe SSD storage that is attached to the root port. VMD Passthrough mode provides the ability to manage drives on guest VMs.

Enabling VMD also allows configuration of Intel® Virtual RAID on CPU (VRoC), a hybrid RAID architecture on Intel® Xeon® Scalable Processors. Documentation on the use and configuration of VRoC can be found at the Intel website.

IMPORTANT: VMD must be enabled in the UCS Manager BIOS settings before Operating System install. If enabled after OS installation, the server will fail to boot. This restriction applies to both standard VMD and VMD Passthrough. Likewise, once enabled, you cannot disable VMD without a loss of system function.

Enabling VMD on UCS Manager

To configure a BIOS and local boot Policy for VMD in UCS Manager, use the following procedure. The VMD platform default is disabled.

Note VMD must be enabled before OS installation.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 3 Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.

Step 4 Go to **Policies** > **Root** > **BIOS Policies** and select the new policy.

Step 5 Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.

Step 6 Scroll down to **VMD Enable** and select **Enable**.

Step 7 Click **Save Changes** to enable VMD functions.

Step 8 In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode** and **Add NVMe** from the **Local Devices** menu. Click **Save Changes** to create the policy.

Enabling Volume Management Device (VMD) in Passthrough Mode

Volume Management Device (VMD) Passthrough Mode

The Intel® Volume Management Device (VMD) driver release package for Direct Device Assignment contains the Intel VMD UEFI Driver version for Direct Assign (PCIe PassThru) in VMware ESXi Hypervisor. The Intel VMD NVMe driver assists in the management of CPU-attached Intel PCIe NVMe SSDs.

The Intel VMD driver is required to enable the Direct Assign and discovery of the VMD physical addresses from a supported guest VM. Drivers are only provided for Passthrough mode for ESXi support of Red Hat

Linux or Ubuntu. VMD Passthrough is enabled by configuring a UCS Manager BIOS policy before loading the Operating System. Once the Operating System has been loaded, you cannot enable or disable the VMD Passthrough option.

Note Passthrough mode is enabled by default, but you should always confirm that it is enabled before proceding.

Configuring VMD Passthrough

Passthrough mode is only supported on ESXi drivers for Red Hat Linux or Ubuntu guest operating systems.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 3 Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.

Step 4 Go to **Policies** > **Root** > **BIOS Policies** and select the new policy.

Step 5 Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.

Step 6 Scroll down to **VMD Enable** and select **Enable**.

Step 7 Click **Save Changes** to enable VMD functions.

Step 8 To finish enabling VMD Passthrough mode, select **Advanced** and **Intel Directed IO** from the submenus and scroll down to **Intel VT Directed IO**. Verify that the dropdown is set to **Enabled**. If not, set it.

Step 9 Click **Save Changes** to enable the VMD Passthrough policy.

Step 10 In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode**. Click **OK** to create the policy.

Downloading VMD Drivers

Intel® Volume Management Device Drivers

Intel® Volume Management Device (VMD) for NVMe enables drive management options using hardware logic inside the Intel Xeon processor. Specific drivers are available for the following operating systems:

- Linux
- Windows 2016, 2019
- VMWare

Note The latest VMWare drivers are available directly from the VMWare site. Following links in the VMWare driver download on the Cisco download site will take you directly to the VMWare login page.

For guest Operating Systems on ESXi, use VMD Passthrough mode. Supported Operating Systems for VMD Passthrough are:

- Red Hat Linux
- Ubuntu

To use the features of Intel VMD, you must:

- Enable VMD by creating a BIOS policy in the UCS Manager.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

- Install the appropriate VMD NVMe driver.
- Install the appropriate management tools for the driver package.
- Boot from UEFI.

Intel® Virtual RAID on CPU (VRoC) with VMD

Intel® Virtual RAID on CPU (VRoC) allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel NVMe SSD drives using hardware logic inside the Intel Xeon processor. More information on Intel VRoC can be found at: https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html.

The User Guides for Intel VRoC can be accessed at the direct link at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en

The Windows and Linux user documentation also contains information on how to configure Intel VRoC in the pre-boot environment. Creation of RAID volumes in VRoC is through the HII interface. The Windows documentation provides information on using the BIOS HII option to set up and configure RAID volumes in VRoC.

To use Intel VRoC, you must:

- Enable VMD in the BIOS settings
- Use UEFI boot mode
- Have sufficient drive resources to create the volume
- Use the BIOS HII option to set up and configure VRoC.

The Cisco implementation of Intel VRoC supports RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity) and RAID 10 (combined mirroring and striping).

Downloading the Linux VMD Drivers

Complete these steps to download and install the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 3 Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.

Step 4 Click on the latest release in the left panel.

Note The ISO image for VMD on blade servers is available from the 4.0(4f) release onward.

Step 5 Click on **ISO image of UCS-related linux drivers only** and download the driver bundle.

Step 6 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD** > **RHEL***x.x*.

Step 7 Click on the version of Red Hat Linux that you wish to install.

Step 8 Extract the contents of the folder. The folder contains both the driver package and associated documentation. Follow the installation procedure packaged with the drivers.

What to do next

The Intel® Virtual RAID on CPU (VRoC) Linux Software User Guide can be found with the user documentation at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en. It provides information on performing BIOS HII VRoC setup in the pre-boot environment, as well as how to install and use the programmable LED utility.

Downloading the Windows VMD Drivers

Complete these steps to download the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 3 Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.

Step 4 Click on the latest release in the left panel.

The ISO image for VMD is available from the 4.0(4f) release onward.

Step 5 Click on **ISO image of UCS-related windows drivers only** and download the driver bundle.

Step 6 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD** > **KIT_x_x_x_xxxx**.

Step 7 Extract the contents of the folder.

Step 8 Click on the entry for the kit and **KIT** > **Install**.

Step 9 The folder contains both the driver package and associated documentation. Expand the zip file for **VROC_x_x_x_xxxxInstall**.

Step 10 Follow the installation procedure packaged with the drivers.

What to do next

For setting up Intel® Virtual RAID on CPU (VRoC), refer to the online instructions at https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html.

Information on VRoC RAID features and management can be found in the *Windows Intel Virtual RAID on CPU Software User's Guide* at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/ssd-software/Windows_VROC_User_Guide.pdf.

Downloading the VMD Passthrough Drivers

Complete these steps to download and install the driver bundle for VMD Passthrough mode:

Note The VMD Passthrough driver bundle includes packages for both ESXi and Ubuntu.

Before you begin

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on**Servers - Unified Computing.**

Step 3 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 4 Choose the UCS utilities from the Software Type selections: **Unified Computing System (UCS) Utilities**.

Step 5 Click on the latest release in the left panel.

Note The ISO image for VMD is available from UCSM 4.0(4f) release onward.

Step 6 Click on **ISO image of UCS-related vmware utilities only** and download the utilities bundle.

Step 7 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD**.

The bundle provides both the driver installation package for the desired version of ESXi or VMD Direct Asssign with Ubuntu, passthrough mode, and the Signed LED Offline bundle. Also included is a pdf that provides steps to configure an Ubuntu Virtual Machine in ESXi.

Step 8 Click on either the version of ESXi that you wish to install or the zip file for Ubuntu.

For ESXi versions, Click on **ESXi**_*x* > **Direct Assign** and chose the desired zip file.

Step 9 Extract the contents of the folder. Follow the installation procedure packaged with the driver software.

What to do next

Extract the contents of the LED management tools zip file. Install the management tools according to the instructions included with the driver package.

Before using the command line tools, the ESXi command line shell should be enabled from either the vSphere client or from the direct console of the ESXi host system.

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior.The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink. Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. 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 on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 4: 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
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.	The failure pattern is displayed until: • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 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. Default = Option 1
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.	Default = Enabled Can be: 1. Disabled (only on one drive) 2. Enabled (on all drives)
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.

Status LED	Behavior	Options
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.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 5: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller **BLACKLIST**	`ledmon` will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing **BLINK_ON_INIT**	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
Set `ledmon` scan interval **INVERVAL**	Defines the time interval between `ledmon sysfs` scans. The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding **REBUILD_BLINK_ON_ALL**	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
RAID volume is migrating **BLINK_ON_MIGR**	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 ledmon debug level **LOG_LEVEL**	Corresponds with –log-level flag from `ledmon`.	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 **RAID_MEMBRES_ONLY**	If the flag is set to `ledmon` true, will limit monitoring only to drives that are RAID members.	1. False / (all RAID member and PT) 2. True / (RAID member only) Default = False

Status LED	Behavior	Options
Limited scans only to specific controllers **WHITELIST**	`ledmon` limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 6: 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.

Enabling Intel® Volume Management Device

Volume Management Device (VMD) Setup

The Intel® Volume Management Device (VMD) is a tool that provides NVMe drivers to manage PCIe Solid State Drives attached to VMD-enabled domains. This includes Surprise hot-plug of PCIe drives and configuring blinking patterns to report status. PCIe Solid State Drive (SSD) storage lacks a standardized method to blink LEDs to represent the status of the device. With VMD, you can control LED indicators on both direct attached and switch attached PCIe storage using a simple command-line tool.

To use VMD, you must first enable VMD through a UCS Manager BIOS policy and set the UEFI boot options. Enabling VMD provides Surprise hot plug and optional LED status management for PCIe SSD storage that is attached to the root port. VMD Passthrough mode provides the ability to manage drives on guest VMs.

Enabling VMD also allows configuration of Intel® Virtual RAID on CPU (VRoC), a hybrid RAID architecture on Intel® Xeon® Scalable Processors. Documentation on the use and configuration of VRoC can be found at the Intel website.

IMPORTANT: VMD must be enabled in the UCS Manager BIOS settings before Operating System install. If enabled after OS installation, the server will fail to boot. This restriction applies to both standard VMD and VMD Passthrough. Likewise, once enabled, you cannot disable VMD without a loss of system function.

Enabling VMD on UCS Manager

To configure a BIOS and local boot Policy for VMD in UCS Manager, use the following procedure. The VMD platform default is disabled.

Note VMD must be enabled before OS installation.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 3 Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.

Step 4 Go to **Policies** > **Root** > **BIOS Policies** and select the new policy.

Step 5 Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.

Step 6 Scroll down to **VMD Enable** and select **Enable**.

Step 7 Click **Save Changes** to enable VMD functions.

Step 8 In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode** and **Add NVMe** from the **Local Devices** menu. Click **Save Changes** to create the policy.

Enabling Volume Management Device (VMD) in Passthrough Mode

Volume Management Device (VMD) Passthrough Mode

The Intel® Volume Management Device (VMD) driver release package for Direct Device Assignment contains the Intel VMD UEFI Driver version for Direct Assign (PCIe PassThru) in VMware ESXi Hypervisor. The Intel VMD NVMe driver assists in the management of CPU-attached Intel PCIe NVMe SSDs.

The Intel VMD driver is required to enable the Direct Assign and discovery of the VMD physical addresses from a supported guest VM. Drivers are only provided for Passthrough mode for ESXi support of Red Hat Linux or Ubuntu. VMD Passthrough is enabled by configuring a UCS Manager BIOS policy before loading the Operating System. Once the Operating System has been loaded, you cannot enable or disable the VMD Passthrough option.

Note Passthrough mode is enabled by default, but you should always confirm that it is enabled before proceding.

Configuring VMD Passthrough

Passthrough mode is only supported on ESXi drivers for Red Hat Linux or Ubuntu guest operating systems.

Procedure

Step 1 In the **Navigation** pane, click **Servers**.

Step 2 Expand the node for the organization where you want to create the policy.

If the system does not include multi tenancy, expand the **root** node.

Step 3 Configure the BIOS policy for VMD: select a service profile and go to the **Policies** tab. In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.

Step 4 Go to **Policies** > **Root** > **BIOS Policies** and select the new policy.

Step 5 Expand **BIOS Policies** and select **Advanced** and **LOM and PCIe Slots** from the submenus.

Step 6 Scroll down to **VMD Enable** and select **Enable**.

Step 7 Click **Save Changes** to enable VMD functions.

Step 8 To finish enabling VMD Passthrough mode, select **Advanced** and **Intel Directed IO** from the submenus and scroll down to **Intel VT Directed IO**. Verify that the dropdown is set to **Enabled**. If not, set it.

Step 9 Click **Save Changes** to enable the VMD Passthrough policy.

Step 10 In the **Boot Policy** tab, create a local boot policy. Select **Uefi** for the **Boot Mode**. Click **OK** to create the policy.

Downloading VMD Drivers

Intel® Volume Management Device Drivers

Intel® Volume Management Device (VMD) for NVMe enables drive management options using hardware logic inside the Intel Xeon processor. Specific drivers are available for the following operating systems:

- Linux
- Windows 2016, 2019
- VMWare

Note The latest VMWare drivers are available directly from the VMWare site. Following links in the VMWare driver download on the Cisco download site will take you directly to the VMWare login page.

For guest Operating Systems on ESXi, use VMD Passthrough mode. Supported Operating Systems for VMD Passthrough are:

- Red Hat Linux
- Ubuntu

To use the features of Intel VMD, you must:

- Enable VMD by creating a BIOS policy in the UCS Manager.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

- Install the appropriate VMD NVMe driver.
- Install the appropriate management tools for the driver package.
- Boot from UEFI.

Intel® Virtual RAID on CPU (VRoC) with VMD

Intel® Virtual RAID on CPU (VRoC) allows you to create and manage RAID volumes within the BIOS of VMD-enabled Intel NVMe SSD drives using hardware logic inside the Intel Xeon processor. More information on Intel VRoC can be found at: https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html.

The User Guides for Intel VRoC can be accessed at the direct link at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en

The Windows and Linux user documentation also contains information on how to configure Intel VRoC in the pre-boot environment. Creation of RAID volumes in VRoC is through the HII interface. The Windows documentation provides information on using the BIOS HII option to set up and configure RAID volumes in VRoC.

To use Intel VRoC, you must:

- Enable VMD in the BIOS settings
- Use UEFI boot mode
- Have sufficient drive resources to create the volume
- Use the BIOS HII option to set up and configure VRoC.

The Cisco implementation of Intel VRoC supports RAID 0 (striping), RAID 1 (mirroring), RAID 5 (striping with parity) and RAID 10 (combined mirroring and striping).

Downloading the Linux VMD Drivers

Complete these steps to download and install the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 3 Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.

Step 4 Click on the latest release in the left panel.

Note The ISO image for VMD on blade servers is available from the 4.0(4f) release onward.

Step 5 Click on **ISO image of UCS-related linux drivers only** and download the driver bundle.

Step 6 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD** > **RHEL***x.x*.

Step 7 Click on the version of Red Hat Linux that you wish to install.

Step 8 Extract the contents of the folder. The folder contains both the driver package and associated documentation. Follow the installation procedure packaged with the drivers.

What to do next

The Intel® Virtual RAID on CPU (VRoC) Linux Software User Guide can be found with the user documentation at: https://www.intel.com/content/www/us/en/support/articles/000030445/memory-and-storage/ssd-software.html?productId=122484&localeCode=us_en. It provides information on performing BIOS HII VRoC setup in the pre-boot environment, as well as how to install and use the programmable LED utility.

Downloading the Windows VMD Drivers

Complete these steps to download the driver bundle:

Before you begin

Make sure that VMD is enabled in the BIOS settings.

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 3 Choose the UCS drivers from the Software Type selections: **Unified Computing System (UCS) Drivers**.

Step 4 Click on the latest release in the left panel.

The ISO image for VMD is available from the 4.0(4f) release onward.

Step 5 Click on **ISO image of UCS-related windows drivers only** and download the driver bundle.

Step 6 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD** > **KIT_x_x_x_xxxx**.

Step 7 Extract the contents of the folder.

Step 8 Click on the entry for the kit and **KIT** > **Install**.

Step 9 The folder contains both the driver package and associated documentation. Expand the zip file for **VROC_x_x_x_xxxxInstall**.

Step 10 Follow the installation procedure packaged with the drivers.

What to do next

For setting up Intel® Virtual RAID on CPU (VRoC), refer to the online instructions at https://www.intel.com/content/www/us/en/support/products/122484/memory-and-storage/ssd-software/intel-virtual-raid-on-cpu-intel-vroc.html.

Information on VRoC RAID features and management can be found in the *Windows Intel Virtual RAID on CPU Software User's Guide* at https://www.intel.com/content/dam/support/us/en/documents/memory-and-storage/ssd-software/Windows_VROC_User_Guide.pdf.

Downloading the VMD Passthrough Drivers

Complete these steps to download and install the driver bundle for VMD Passthrough mode:

Note The VMD Passthrough driver bundle includes packages for both ESXi and Ubuntu.

Before you begin

Note The system will fail to boot if VMD is enabled or disabled after OS installation. Do not change the BIOS setting after OS installation.

Procedure

Step 1 In a web browser, navigate to https://software.cisco.com/download/home.

Step 2 Search on**Servers - Unified Computing.**

Step 3 Search on **UCS B-Series Blade Server Software** or **UCS C-Series Rack-Mount UCS-Managed Server Software**, depending on your platform.

Step 4 Choose the UCS utilities from the Software Type selections: **Unified Computing System (UCS) Utilities**.

Step 5 Click on the latest release in the left panel.

Note The ISO image for VMD is available from UCSM 4.0(4f) release onward.

Step 6 Click on **ISO image of UCS-related vmware utilities only** and download the utilities bundle.

Step 7 When the driver bundle is downloaded, open it and select **Storage** > **Intel** > **VMD**.

The bundle provides both the driver installation package for the desired version of ESXi or VMD Direct Asssign with Ubuntu, passthrough mode, and the Signed LED Offline bundle. Also included is a pdf that provides steps to configure an Ubuntu Virtual Machine in ESXi.

Step 8 Click on either the version of ESXi that you wish to install or the zip file for Ubuntu.

For ESXi versions, Click on **ESXi**_*x* > **Direct Assign** and chose the desired zip file.

Step 9 Extract the contents of the folder. Follow the installation procedure packaged with the driver software.

What to do next

Extract the contents of the LED management tools zip file. Install the management tools according to the instructions included with the driver package.

Before using the command line tools, the ESXi command line shell should be enabled from either the vSphere client or from the direct console of the ESXi host system.

Custom LED Status with VMD on NVMe

Once you have set up VMD, you can customize LED blinking patterns on PCIe NVMe drives. Information on LED customization can be found in the User Guides included in the driver packages.

LED Blinking

PCIe SSD drives lack a standard way to manage the LEDs that indicate drive status and health. Without this, there is a risk of removing the wrong drive, resulting in data loss. SSD drives have two indicators, the first being a green activity LED whose signals come directly from the SSD, and the second being a status LED whose signals come from the backplane. VMD manages only the status LEDs, not the activity LEDs.

LED Management only applies to NVMe and/or SATA drives. It does not support drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard .

LED Activity During Drive Hot-plug

VMD with NVMe supports Surprise hot-plugging. When a disk is hot-removed, then re-inserted into the same slot, the fault LED blinks for 10 seconds. This is expected behavior.The fail state is imposed on a slot's LEDs when the drive is removed, but the backplanes require the drive to be present in the slot for a LED to blink. Thus, the fail state exists once the drive is removed, but a LED blinks only when the new drive is inserted and discovered. The LED will return to normal once hot-plug event is handled.

Custom Blinking Patterns

VRoC with VMD allows you to perform basic LED management configuration of the status LEDs on compatible backplanes. Once the VMD NVMe driver is installed, you can install the VMD LED Management Tool, which lets you manage the LED through a command line interface. 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 on the various platforms. As individualized patterns are programmable, these tables provide only representative guidelines.

Table 7: 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
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.	The failure pattern is displayed until: • 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. • 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. Default = Option 1
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.	Default = Enabled Can be: 1. Disabled (only on one drive) 2. Enabled (on all drives)
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.	Default = Enabled Can be: 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs)
Rebuild	Only the migrating drive blinks.	Default = Disabled

Table 8: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller **BLACKLIST**	`ledmon` will exclude scanning controllers listed on the blacklist. When the whitelist is also set in the config file, the blacklist is ignored.	Exclude controllers on the blacklist. Default = Support all controllers
RAID volume is initializing, verifying, or verifying and fixing **BLINK_ON_INIT**	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
Set `ledmon` scan interval **INVERVAL**	Defines the time interval between `ledmon sysfs` scans. The value is given in seconds.	10s (5s maximum) Default = 10s
RAID volume is rebuilding **REBUILD_BLINK_ON_ALL**	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
RAID volume is migrating **BLINK_ON_MIGR**	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 ledmon debug level **LOG_LEVEL**	Corresponds with –log-level flag from `ledmon`.	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 **RAID_MEMBRES_ONLY**	If the flag is set to `ledmon` true, 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 **WHITELIST**	`ledmon` limits changing the LED state to controllers listed on whitelist.	Limit changing LED state in whitelist controller. Default = No limit.

Table 9: 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.

Status LED	Behavior	Options
"Off"	The ability to turn off the "Identify" LED once a specific device in an enclosure has been located.	None. Default is Off.