



Storage-Related Policies

- [About vHBA Templates, on page 1](#)
- [Fibre Channel Adapter Policies, on page 4](#)
- [About the Default vHBA Behavior Policy, on page 13](#)
- [SPDM Security Policy, on page 14](#)
- [SAN Connectivity Policies, on page 17](#)

About vHBA Templates

vHBA Template

This template is a policy that defines how a vHBA on a server connects to the SAN. It is also referred to as a vHBA SAN connectivity template.

You must include this policy in a service profile for it to take effect.

Creating a vHBA Template

Before you begin

This policy requires that one or more of the following resources already exist in the system:

- Named VSAN
- WWNN pool or WWPN pool
- SAN pin group
- Statistics threshold policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** 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 4 Right-click the **vHBA Templates** node and choose **Create vHBA Template**.

Step 5 In the **Create vHBA Template** dialog box, complete the following fields:

Name	Description
Name field	The name of the virtual HBA template. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A user-defined description of the template. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Fabric ID field	The name of the fabric interconnect that vHBAs created with this template are associated with.
Select VSAN drop-down list	The VSAN to associate with vHBAs created from this template.
Create VSAN link	Click this link if you want to create a VSAN.
Template Type field	This can be one of the following: <ul style="list-style-type: none"> • Initial Template—vHBAs created from this template are not updated if the template changes. • Updating Template—vHBAs created from this template are updated if the template changes.
Max Data Field Size field	The maximum size of the Fibre Channel frame payload bytes that the vHBA supports. Enter an integer between 256 and 2112. The default is 2048.
WWPN Pool drop-down list	The WWPN pool that a vHBA created from this template uses to derive its WWPN address.
QoS Policy drop-down list	The QoS policy that is associated with vHBAs created from this template.
Pin Group drop-down list	The SAN pin group that is associated with vHBAs created from this template.
Stats Threshold Policy drop-down list	The statistics collection policy that is associated with vHBAs created from this template.

Step 6 Click **OK**.

What to do next

Include the vHBA template in a service profile.

Binding a vHBA to a vHBA Template

You can bind a vHBA associated with a service profile to a vHBA template. When you bind the vHBA to a vHBA template, Cisco UCS Manager configures the vHBA with the values defined in the vHBA template. If the existing vHBA configuration does not match the vHBA template, Cisco UCS Manager reconfigures the vHBA. You can only change the configuration of a bound vHBA through the associated vHBA template. You cannot bind a vHBA to a vHBA template if the service profile that includes the vHBA is already bound to a service profile template.



Important If the vHBA is reconfigured when you bind it to a template, Cisco UCS Manager reboots the server associated with the service profile.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers** > **Service Profiles**.
 - Step 3** Expand the node for the organization that includes the service profile with the vHBA you want to bind.
If the system does not include multi-tenancy, expand the **root** node.
 - Step 4** Expand *Service_Profile_Name* > **vHBAs**.
 - Step 5** Click the vHBA you want to bind to a template.
 - Step 6** In the **Work** pane, click the **General** tab.
 - Step 7** In the **Actions** area, click **Bind to a Template**.
 - Step 8** In the **Bind to a vHBA Template** dialog box, do the following:
 - a) From the **vHBA Template** drop-down list, choose the template to which you want to bind the vHBA.
 - b) Click **OK**.
 - Step 9** In the warning dialog box, click **Yes** to acknowledge that Cisco UCS Manager may need to reboot the server if the binding causes the vHBA to be reconfigured.
-

Unbinding a vHBA from a vHBA Template

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Service Profiles**.
- Step 3** Expand the node for the organization that includes the service profile with the vHBA you want to unbind.

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

- Step 4** Expand *Service_Profile_Name* > vHBAs.
 - Step 5** Click the vHBA you want to unbind from a template.
 - Step 6** In the **Work** pane, click the **General** tab.
 - Step 7** In the **Actions** area, click **Unbind from a Template**.
 - Step 8** If a confirmation dialog box displays, click **Yes**.
-

Deleting a vHBA Template

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Expand the **vHBA Templates** node.
 - Step 4** Right-click the vHBA template that you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

Fibre Channel Adapter Policies

Ethernet and Fibre Channel Adapter Policies

These policies govern the host-side behavior of the adapter, including how the adapter handles traffic. For example, you can use these policies to change default settings for the following:

- Queues
- Interrupt handling
- Performance enhancement
- RSS hash
- Failover in a cluster configuration with two fabric interconnects



Note For Fibre Channel adapter policies, the values displayed by Cisco UCS Manager may not match those displayed by applications such as QLogic SANsurfer. For example, the following values may result in an apparent mismatch between SANsurfer and Cisco UCS Manager:

- **Max LUNs Per Target**—SANsurfer has a maximum of 256 LUNs and does not display more than that number. Cisco UCS Manager supports a higher maximum number of LUNs. This parameter is applicable only for FC-Initiator.
- **Link Down Timeout**—In SANsurfer, you configure the timeout threshold for link down in seconds. In Cisco UCS Manager, you configure this value in milliseconds. Therefore, a value of 5500 ms in Cisco UCS Manager displays as 5s in SANsurfer.
- **Max Data Field Size**—SANsurfer has allowed values of 512, 1024, and 2048. Cisco UCS Manager allows you to set values of any size. Therefore, a value of 900 in Cisco UCS Manager displays as 512 in SANsurfer.
- **LUN Queue Depth**—The LUN queue depth setting is available for Windows system FC adapter policies. Queue depth is the number of commands that the HBA can send and receive in a single transmission per LUN. Windows Storport driver sets this to a default value of 20 for physical miniports and to 250 for virtual miniports. This setting adjusts the initial queue depth for all LUNs on the adapter. Valid range for this value is 1 - 254. The default LUN queue depth is 20. This feature only works with Cisco UCS Manager version 3.1(2) and higher. This parameter is applicable only for FC-Initiator.
- **IO TimeOut Retry**—When the target device does not respond to an IO request within the specified timeout, the FC adapter cancels the pending command then resends the same IO after the timer expires. The FC adapter valid range for this value is 1 - 59 seconds. The default IO retry timeout is 5 seconds. This feature only works with Cisco UCS Manager version 3.1(2) and higher.

Operating System Specific Adapter Policies

By default, Cisco UCS provides a set of Ethernet adapter policies and Fibre Channel adapter policies. These policies include the recommended settings for each supported server operating system. Operating systems are sensitive to the settings in these policies. Storage vendors typically require non-default adapter settings. You can find the details of these required settings on the support list provided by those vendors.



Important We recommend that you use the values in these policies for the applicable operating system. Do not modify any of the values in the default policies unless directed to do so by Cisco Technical Support.

However, if you are creating an Ethernet adapter policy for an OS (instead of using the default adapter policy), you must use the following formulas to calculate values that work for that OS.

Depending on the UCS firmware, your driver interrupt calculations may be different. Newer UCS firmware uses a calculation that differs from previous versions. Later driver release versions on Linux operating systems now use a different formula to calculate the Interrupt Count. In this formula, the Interrupt Count is the maximum of either the Transmit Queue or the Receive Queue plus 2.

Interrupt Count in Linux Adapter Policies

Drivers on Linux operating systems use differing formulas to calculate the Interrupt Count, depending on the eNIC driver version. The UCS 3.2 release increased the number of Tx and Rx queues for the eNIC driver from 8 to 256 each.

Use one of the following strategies, according to your driver version.

For Linux drivers before the UCS 3.2 firmware release, use the following formula to calculate the Interrupt Count.

$$\begin{aligned} \text{Completion Queues} &= \text{Transmit Queues} + \text{Receive Queues} \\ \text{Interrupt Count} &= (\text{Completion Queues} + 2) \text{ rounded up to nearest power of } 2 \end{aligned}$$

For example, if Transmit Queues = 1 and Receive Queues = 8 then:

$$\begin{aligned} \text{Completion Queues} &= 1 + 8 = 9 \\ \text{Interrupt Count} &= (9 + 2) \text{ rounded up to the nearest power of } 2 = 16 \end{aligned}$$

On drivers for UCS firmware release 3.2 and higher, the Linux eNIC drivers use the following formula to calculate the Interrupt Count.

$$\text{Interrupt Count} = (\# \text{Tx or Rx Queues}) + 2$$

For example:

$$\begin{aligned} \text{Interrupt Count } wq = 32, rq = 32, cq = 64 &- \text{ then Interrupt Count} = \text{Max}(32, 32) + 2 = 34 \\ \text{Interrupt Count } wq = 64, rq = 8, cq = 72 &- \text{ then Interrupt Count} = \text{Max}(64, 8) + 2 = 66 \\ \text{Interrupt Count } wq = 1, rq = 16, cq = 17 &- \text{ then Interrupt count} = \text{Max}(1, 16) + 2 = 18 \end{aligned}$$

Interrupt Count in Windows Adapter Policies

For Windows OS, the recommended adapter policy in UCS Manager for VIC 1400 series and above adapters is Win-HPN and if RDMA is used, the recommended policy is Win-HPN-SMB. For VIC 1400 series and above adapters, the recommended interrupt value setting is 512 and the Windows VIC driver takes care of allocating the required number of Interrupts.

For VIC 1300 and VIC 1200 series adapters, the recommended UCS Manager adapter policy is Windows and the Interrupt would be TX + RX + 2, rounded to closest power of 2. The maximum supported Windows queues is 8 for Rx Queues and 1 for Tx Queues.

Example for VIC 1200 and VIC 1300 series adapters:

$$\text{Tx} = 1, \text{Rx} = 4, \text{CQ} = 5, \text{Interrupt} = 8 (1 + 4 \text{ rounded to nearest power of } 2), \text{Enable RSS}$$

Example for VIC 1400 series and above adapters:

$$\text{Tx} = 1, \text{Rx} = 4, \text{CQ} = 5, \text{Interrupt} = 512, \text{Enable RSS}$$

NVMe over Fabrics using Fibre Channel

The NVM Express (NVMe) interface allows host software to communicate with a non-volatile memory subsystem. This interface is optimized for Enterprise non-volatile storage, which is typically attached as a register level interface to the PCI Express (PCIe) interface.

NVMe over Fabrics using Fibre Channel (FC-NVMe) defines a mapping protocol for applying the NVMe interface to Fibre Channel. This protocol defines how Fibre Channel services and specified Information Units (IUs) are used to perform the services defined by NVMe over a Fibre Channel fabric. NVMe initiators can access and transfer information to NVMe targets over Fibre Channel.

FC-NVMe combines the advantages of Fibre Channel and NVMe. You get the improved performance of NVMe along with the flexibility and the scalability of the shared storage architecture. Cisco UCS Manager Release 4.0(2) supports NVMe over Fabrics using Fibre Channel on UCS VIC 1400 Series adapters.

Starting with UCS Manager release 4.3(2b), NVMeoF using RDMA is supported on Cisco UCS VIC 14000 series adapters.

Starting with UCS Manager release 4.2(2), NVMeoF using Fibre Channel is supported on Cisco UCS VIC 15000 series adapters.

Cisco UCS Manager provides the recommended FC NVME Initiator adapter policies in the list of pre-configured adapter policies. To create a new FC-NVMe adapter policy, follow the steps in the *Creating a Fibre Channel Adapter Policy* section.

NVMe over Fabrics Using RDMA

NVMe over Fabrics (NVMeoF) is a communication protocol that allows one computer to access NVMe namespaces available on another computer. NVMeoF is similar to NVMe, but differs in the network-related steps involved in using the NVMeoF storage devices. The commands for discovering, connecting, and disconnecting a NVMeoF storage device are integrated into the **nvme** utility provided in Linux..

The NVMeoF fabric that Cisco supports is RDMA over Converged Ethernet version 2 (RoCEv2). RoCEv2 is a fabric protocol that runs over UDP. It requires a no-drop policy.

The eNIC RDMA driver works in conjunction with the eNIC driver, which must be loaded first when configuring NVMeoF.

Cisco UCS Manager provides the default Linux-NVMe-RoCE adapter policy for creating NVMe RoCEv2 interfaces. Do not use the default Linux adapter policy. For complete information on configuring RoCEv2 over NVMeoF, refer to the *Cisco UCS Manager Configuration Guide for RDMA over Converged Ethernet (RoCE) v2*.

NVMeoF using RDMA is supported on M5 B-Series or C-Series Servers with Cisco UCS VIC 1400 Series adapters.

Starting with UCS Manager release 4.3(2b), NVMeoF using RDMA is supported on Cisco UCS VIC 14000 series adapters.

Starting with UCS Manager release 4.2(2), NVMeoF using RDMA is supported on Cisco UCS VIC 15000 series adapters.

Creating a Fibre Channel Adapter Policy



Tip If the fields in an area do not display, click the **Expand** icon to the right of the heading.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers** > **Policies**.
- Step 3** 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 4 Right-click **Adapter Policies** and choose **Create Fibre Channel Adapter Policy**.

Step 5 Enter a name and description for the policy in the following fields:

Table 1:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 6 (Optional) In the **Resources** area, adjust the following values:

Name	Description
Transmit Queues field	The number of transmit queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each transmit queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 128. The default is 64.
Receive Queues field	The number of receive queue resources to allocate. This value cannot be changed.
Ring Size field	The number of descriptors in each receive queue. This parameter applies to Extended Link Services (ELS) and Common Transport (CT) fibre channel frames for generic services. It does not affect adapter performance. Enter an integer between 64 and 2048. The default is 64.
I/O Queues field	The number of IO queue resources the system should allocate. Enter an integer between 1 and 64. The default is 1.
Ring Size field	The number of descriptors in each I/O queue. Enter an integer between 64 and 512. The default is 512. Note The number of descriptors can affect the performance of the adapter, so we recommend that you do not change the default value.

Step 7 (Optional) In the **Options** area, adjust the following values:

Name	Description
FCP Error Recovery field	<p>Whether the system uses FCP Sequence Level Error Recovery (FC-TAPE) protocol for sequence level error recovery with tape devices. This enables or disables the Read Exchange Concise (REC) and Sequence Retransmission Request (SRR) functions on the VIC firmware. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—This is the default. • Enabled—You should select this option if your system is connected to one or more tape drive libraries. <p>Note This parameter only applies to a server with a Virtual Interface Card (VIC) adapter.</p>
Flogi Retries field	<p>The number of times that the system tries to log in to the fabric after the first failure.</p> <p>Enter any integer. To specify that the system continue to try indefinitely, enter infinite in this field. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter, or a converged network adapter.</p>
Flogi Timeout (ms) field	<p>The number of milliseconds that the system waits before it tries to log in again.</p> <p>Enter an integer between 1000 and 255000. The default is 4,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter or a converged network adapter.</p> <p>When a Flogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Flogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.</p>
Flogi Retries field	<p>The number of times that the system tries to log into a port after the first failure.</p> <p>Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p>

Name	Description
Plogi Timeout (ms) field	<p>The number of milliseconds that the system waits before it tries to log in again.</p> <p>Enter an integer between 1000 and 255000. The default is 20,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>For an HBA that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 4,000 ms.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p> <p>When a Plogi timeout value of 20 seconds or more is configured for a boot vHBA, it could lead to a SAN boot failure if the adapter does not receive an accept to the initial Plogi. For a boot-enabled vHBA, the recommended timeout values is 5 seconds or less.</p>
Port Down Timeout (ms) field	<p>The number of milliseconds a remote Fibre Channel port should be offline before informing the SCSI upper layer that the port is unavailable. This parameter is important for host multi-pathing drivers and it is one of the key indicators used for error processing.</p> <p>Enter an integer between 0 and 240000. The default is 30,000. For a server with a VIC adapter running ESX, the recommended value is 10,000.</p> <p>For a server with a port that is going to be used to boot a Windows OS from SAN, the recommended value for this field is 5000 milliseconds.</p> <p>We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter.</p>
IO Retry Timeout (seconds)	<p>The number of seconds that the FC adapter waits before aborting the pending command and resending the same IO. This happens when the network device does not responding to an IO request within the specified time.</p> <p>Enter an integer between 0 and 59 seconds. The default IO retry timeout is 5 seconds.</p>
Port Down IO Retry field	<p>The number of times an IO request to a port is returned because the port is busy before the system decides the port is unavailable.</p> <p>Enter an integer between 0 and 255. The default is 8. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter running Windows.</p>

Name	Description
Link Down Timeout (ms) field	<p>The number of milliseconds the uplink port should be offline before it informs the system that the uplink port is down and fabric connectivity has been lost.</p> <p>Enter an integer between 0 and 240000. The default is 30,000. We recommend you consult your storage array documentation for the optimal value for this parameter.</p> <p>Note This parameter only applies to a server with a VIC adapter running Windows.</p>
IO Throttle Count field	<p>The maximum number of data or control I/O operations that can be pending in the vHBA at one time. If this value is exceeded, the additional I/O operations wait in the queue until the number of pending I/O operations decreases and the additional operations can be processed.</p> <p>Note This parameter is not the same as the LUN queue depth, which is controlled by Cisco UCS Manager based on the operating system installed on the server.</p> <p>Enter an integer between 256 and 1024. The default is 256. We recommend you consult your storage array documentation for the optimal value for this parameter.</p>
Max LUNs Per Target field	<p>The maximum number of LUNs that the Fibre Channel driver will export or show. The maximum number of LUNs is usually controlled by the operating system running on the server.</p> <p>Enter an integer between 1 and 4096.</p> <p>We recommend you consult your operating system documentation for the optimal value for this parameter.</p> <p>Note</p> <ul style="list-style-type: none"> • This parameter only applies to a server with a VIC adapter or a network adapter. • This parameter is applicable only for FC-Initiator.
LUN Queue Depth field	<p>The number of commands that the HBA can send and receive in a single transmission per LUN.</p> <p>Enter an integer between 1 and 254. The default LUN queue depth is 20.</p> <p>Note This parameter is applicable only for FC-Initiator.</p>

Name	Description
Interrupt Mode radio button	<p>The method used to send interrupts to the operating system from the driver. This can be one of the following:</p> <ul style="list-style-type: none"> • MSI-X—Message Signaled Interrupts (MSI) with the optional extension. We recommend that you select this option if the operating system on the server supports it. • MSI—MSI only. • INTx—PCI INTx interrupts. <p>Note This parameter only applies to a server with a VIC adapter or a network adapter running an operating system other than Windows. The Windows operating system ignores this parameter.</p>
vHBA Type radio button	<p>The vHBA type used in this policy. vHBAs supporting FC and FC-NVMe can now be created on the same adapter. The vHBA type used in this policy can be one of the following:</p> <ul style="list-style-type: none"> • FC Initiator—Legacy SCSI FC vHBA initiator • FC Target—vHBA that supports SCSI FC target functionality <p>Note This option is available as a Tech Preview.</p> <ul style="list-style-type: none"> • FC NVME Initiator—vHBA that is an FC NVME initiator, which discovers FC NVME targets and connects to them • FC NVME Target—vHBA that acts as an FC NVME target and provides connectivity to the NVME storage <p>Note This option is available as a Tech Preview.</p> <p>vHBA type is supported only on UCS VIC 1400 adapters and UCS VIC 15000 adapters.</p>

Step 8 Click **OK**.

Step 9 If a confirmation dialog box displays, click **Yes**.

Deleting a Fibre Channel Adapter Policy

Procedure

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

Step 2 Expand **SAN > Policies > Organization_Name**.

Step 3 Expand the **Fibre Channel Policies** node.

- Step 4** Right-click the policy you want to delete and choose **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
-

About the Default vHBA Behavior Policy

Default vHBA Behavior Policy

Default vHBA behavior policy allow you to configure how vHBAs are created for a service profile. You can choose to create vHBAs manually, or you can allow them to be created automatically.

You can configure the default vHBA behavior policy to define how vHBAs are created. This can be one of the following:

- **None**—Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
- **HW Inherit**—If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.



Note If you do not specify a default behavior policy for vHBAs, **none** is used by default.

Configuring a Default vHBA Behavior Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** Expand the **root** node.

You can configure only the default vHBA behavior policy in the root organization. You cannot configure the default vHBA behavior policy in a sub-organization.

- Step 4** Click **Default vHBA Behavior**.
- Step 5** On the **General Tab**, in the **Properties** area, click one of the following radio buttons in the **Action** field:
- **None**—Cisco UCS Manager does not create default vHBAs for a service profile. All vHBAs must be explicitly created.
 - **HW Inherit**—If a service profile requires vHBAs and none have been explicitly defined, Cisco UCS Manager creates the required vHBAs based on the adapter installed in the server associated with the service profile.

Step 6 Click **Save Changes**.

SPDM Security Policy

SPDM Security

Cisco UCS M6 Servers can contain mutable components that could provide vectors for attack against a device itself or use of a device to attack another device within the system. To defend against these attacks, the Security Protocol and Data Model (SPDM) Specification enables a secure transport implementation that challenges a device to prove its identity and the correctness of its mutable component configuration. This feature is supported on Cisco UCS C220 and C240 M6 Servers starting with in Cisco UCS Manager, Release 4.2(1d).



Note SPDM is currently not supported on the Cisco UCS C225 M6 Server and Cisco UCS C245 M6 Server.

SPDM defines messages, data objects, and sequences for performing message exchanges between devices over a variety of transport and physical media. It orchestrates message exchanges between Baseboard Management Controllers (BMC) and end-point devices over a Management Component Transport Protocol (MCTP). Message exchanges include authentication of hardware identities accessing the BMC. The SPDM enables access to low-level security capabilities and operations by specifying a managed level for device authentication, firmware measurement, and certificate management. Endpoint devices are challenged to provide authentication, and BMC authenticates the endpoints and only allows access for trusted entities.

The UCS Manager optionally allows uploads of external security certificates to BMC. A maximum of 40 SPDM certificates is allowed, including native internal certificates. Once the limit is reached, no more certificates can be uploaded. User uploaded certificates can be deleted but internal/default certificates cannot.

A SPDM security policy allows you to specify one of three Security level settings. Security can be set at one of the three levels listed below:

- Full Security:

This is the highest MCTP security setting. When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. A fault will also be generated if any of the endpoints do not support either endpoint authentication or firmware measurements.

- Partial Security (default):

When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. There will NOT be a fault generated when the endpoint doesn't support endpoint authentication or firmware measurements.

- No Security

When you select this setting, there will NOT be a fault generated for any failure (either endpoint measurement or firmware measurement failures).

You can also upload the content of one or more external/device certificates into BMC. Using a SPDM policy allows you to change or delete security certificates or settings as desired. Certificates can be deleted or replaced when no longer needed.

Certificates are listed in all user interfaces on a system.

Creating a SPDM Security Policy

This step creates a SPDM policy.



Note You can upload up to 40 SPDM certificates (including native certificates).

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Go to **Policies**. Expand the root node.
- Step 3** Right-click **SPDM Certificate Policies** and select **Create SPDM Policy**.
- Step 4** Enter a name for this policy and select a **Fault Alert Setting** for the security level: **Disabled**, **Partial**, or **Full**.
- Full—If you select this option, then a fault is generated when there is any endpoint authentication failure for both supported and unsupported endpoints.
- Partial—If you select this option then a fault is generated when there is any endpoint authentication failure to only supported endpoints. No fault is generated when the endpoint does not support authentication.
- Disabled—If you select this option then no fault is generated for endpoint authentication failure for both supported and unsupported endpoints.
- The default is **Partial**.
- Note** To perform SPDM re-authentication and update the faults, Cisco IMC reboot or host reboot is required when the fault alert value is changed for an associated profile.
- Step 5** Click on **Add** in the **Create Policy** window. The **Add SPDM Certificate** window will open.
- Step 6** Name the certificate.
- UCS Manager supports only **Pem**certificates.
- Step 7** Paste the contents of the certificate into the Certificate field.
- Step 8** Click **OK** to add the certificate and return to the **Create SPDM Policy** window.
- You can add up to 40 certificates.
- Step 9** In the **Create SPDM Policy** menu, click **Okay**.
- After the SPDM policy is created, it will be listed immediately, along with its Alert setting, when you select **SPDM Certificate Policy** under the Server root Policies.
-

What to do next

Assign the Certificate to a Service Profile. The Service Profile must be associated with a server for it to take effect.

Associating the Security Policy with a Server

Before you begin

Create the SPDM security policy.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Go to **Service Profiles**. Expand the root node.
- Step 3** Select the Service Profile you want to associate with the Policy you created.
- a) On the **Policies** tab, scroll down and expand **SPDM Certificate Policy**. In the **SPDM Certificate Policy** dropdown, select the desired policy to associate with this Service Profile.
- Step 4** Click **OK**.
- The SPDM Policy will now be associated with the service profile.
-

What to do next

Check the fault alert level to make sure it is set to the desired setting.

Viewing the Fault Alert Settings

You can view the Fault Alert setting associated with a specific chassis.

Before you begin

Create a policy and associate it with a Service Profile.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Select a Rack-Mount Server.
- Step 3** On the **Inventory** tab, select **CIMC**.
- User uploaded certificates are listed and information for specific certificates can be selected and viewed.
-

SAN Connectivity Policies

About the LAN and SAN Connectivity Policies

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNS to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.



Note We do not recommend that you use static IDs in connectivity policies, because these policies are included in service profiles and service profile templates and can be used to configure multiple servers.

Privileges Required for LAN and SAN Connectivity Policies

Connectivity policies enable users without network or storage privileges to create and modify service profiles and service profile templates with network and storage connections. However, users must have the appropriate network and storage privileges to create connectivity policies.

Privileges Required to Create Connectivity Policies

Connectivity policies require the same privileges as other network and storage configurations. For example, you must have at least one of the following privileges to create connectivity policies:

- admin—Can create LAN and SAN connectivity policies
- ls-server—Can create LAN and SAN connectivity policies
- ls-network—Can create LAN connectivity policies
- ls-storage—Can create SAN connectivity policies

Privileges Required to Add Connectivity Policies to Service Profiles

After the connectivity policies have been created, a user with ls-compute privileges can include them in a service profile or service profile template. However, a user with only ls-compute privileges cannot create connectivity policies.

Interactions between Service Profiles and Connectivity Policies

You can configure the LAN and SAN connectivity for a service profile through either of the following methods:

- LAN and SAN connectivity policies that are referenced in the service profile
- Local vNICs and vHBAs that are created in the service profile
- Local vNICs and a SAN connectivity policy
- Local vHBAs and a LAN connectivity policy

Cisco UCS maintains mutual exclusivity between connectivity policies and local vNIC and vHBA configuration in the service profile. You cannot have a combination of connectivity policies and locally created vNICs or vHBAs. When you include a LAN connectivity policy in a service profile, all existing vNIC configuration is erased, and when you include a SAN connectivity policy, all existing vHBA configuration in that service profile is erased.

Creating a SAN Connectivity Policy

Procedure

-
- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies**.
- Step 3** 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 4** Right-click **SAN Connectivity Policies** and choose **Create SAN Connectivity Policy**.
- Step 5** In the **Create SAN Connectivity Policy** dialog box, enter a name and optional description.
- Step 6** From the **WWNN Assignment** drop-down list in the **World Wide Node Name** area, choose one of the following:
- Choose **Select (pool default used by default)** to use the default WWN pool.
 - Choose one of the options listed under **Manual Using OUI** and then enter the WWN in the **World Wide Node Name** field.

You can specify a WWNN in the range from 20:00:00:00:00:00:00:00 to 20:FF:FF:FF:FF:FF:FF:FF or from 50:00:00:00:00:00:00:00 to 5F:FF:FF:FF:FF:FF:FF:FF. You can click the **here** link to verify that the WWNN you specified is available.
 - Choose a WWN pool name from the list to have a WWN assigned from the specified pool. Each pool name is followed by two numbers in parentheses that show the number of WWNs still available in the pool and the total number of WWNs in the pool.
- Step 7** In the **vHBAs** table, click **Add**.
- Step 8** In the **Create vHBAs** dialog box, enter the name and optional description.
- Step 9** Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
- Step 10** In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
- Step 11** In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
- Step 12** After you have created all the vHBAs you need for the policy, click **OK**.
-

What to do next

Include the policy in a service profile or service profile template.

Creating a vHBA for a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** On the **SAN** tab, expand **SAN > Policies > Organization_Name > San Connectivity Policies**.
- Step 3** Choose the policy for which you want to create a vHBA.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the table icon bar, click the + button.
- Step 6** In the **Create vHBAs** dialog box, enter the name and optional description.
- Step 7** Choose the **Fabric ID**, **Select VSAN**, **Pin Group**, **Persistent Binding**, and **Max Data Field Size**.
You can also create a VSAN or SAN pin group from this area.
- Step 8** In the **Operational Parameters** area, choose the **Stats Threshold Policy**.
- Step 9** In the **Adapter Performance Profile** area, choose the **Adapter Policy** and **QoS Policy**.
You can also create a fibre channel adapter policy or QoS policy from this area.
- Step 10** Click **Save Changes**.
-

Deleting a vHBA from a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
- Step 2** Expand **SAN > Policies > Organization_Name**.
- Step 3** Choose the policy from which you want to delete the vHBA.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **vHBAs** table, do the following:
a) Click the vHBA that you want to delete.
b) On the icon bar, click **Delete**.
- Step 6** If a confirmation dialog box displays, click **Yes**.
-

Creating an Initiator Group for a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.

- Step 2** Expand **SAN > Policies > Organization_Name**.
- Step 3** Choose the policy for which you want to create an initiator group.
- Step 4** In the **Work** pane, click the **vHBA Initiator Groups** tab.
- Step 5** In the table icon bar, click the + button.
- Step 6** In the **Create vHBA Initiator Group** dialog box, complete the following fields:

Name	Description
Name field	<p>The name of the vHBA initiator group.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>
Description field	<p>A description of the group.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
Select vHBA Initiators table	<p>Check the check box in the Select column for each vHBA that you want to use.</p>
Storage Connection Policy drop-down list	<p>The storage connection policy associated with this vHBA initiator group. If you want to:</p> <ul style="list-style-type: none"> Use an existing storage connection policy, then choose that policy from the drop-down list. The Cisco UCS Manager GUI displays information about the policy and its FC target endpoints in the Global Storage Connection Policy area. Create a new storage connection policy that will be globally available, then click the Create Storage Connection Policy link. Create a local storage connection policy that is available only to this vHBA initiator group, then choose the Specific Storage Connection Policy option. The Cisco UCS Manager GUI displays the Specific Storage Connection Policy area that allows you to configure the local storage connection policy.
Create Storage Connection Policy link	<p>Click this link to create a new storage connection policy that will be available to all service profiles and service profile templates.</p>

- Step 7** Click **OK**.

Deleting an Initiator Group from a SAN Connectivity Policy

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Choose the policy from which you want to delete the initiator group
 - Step 4** In the **Work** pane, click the **vHBA Initiator Groups** tab.
 - Step 5** In the table, do the following:
 - a) Click the initiator group that you want to delete.
 - b) On the icon bar, click **Delete**.
 - Step 6** If a confirmation dialog box displays, click **Yes**.
-

Deleting a SAN Connectivity Policy

If you delete a SAN connectivity policy that is included in a service profile, it also deletes all vHBAs from that service profile and disrupts SAN data traffic for the server associated with the service profile.

Procedure

- Step 1** In the **Navigation** pane, click **SAN**.
 - Step 2** Expand **SAN > Policies > Organization_Name**.
 - Step 3** Expand the **SAN Connectivity Policies** node.
 - Step 4** Right-click the policy that you want to delete and choose **Delete**.
 - Step 5** If a confirmation dialog box displays, click **Yes**.
-

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 submenu.
 - 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 proceeding.

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 Assign 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 2: 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: <ul style="list-style-type: none"> 1. It is physically removed. or the RAID volume, that contains the failed drive, is either deleted or physically removed. <ul style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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 3: LED Blinking Patterns: Linux

Status LED	Behavior	Options
Skip/exclude controller BLACKLIST	<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.	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 <code>ledmon</code> scan interval INTERVAL	Defines the time interval between <code>ledmon sysfs scans</code> . 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 <code>ledmon</code> debug level LOG_LEVEL	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 RAID_MEMBERS_ONLY	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

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

Table 4: 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.