



Storage Inventory

- [NVMe-optimized M5 Servers, on page 1](#)
- [NVMe Replacement Considerations for B200 M6, Servers, on page 3](#)
- [Volume Management Device \(VMD\) Setup, on page 3](#)

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

	Command or Action	Purpose
Step 1	UCS-A# scope server [<i>chassis-num/server-num</i> <i>dynamic-uuid</i>]	Enters server mode for the specified server.
Step 2	UCS-A /server # scope nvme-switch <i>nvme_switch</i>	Enters the specified NVMe switch.
Step 3	UCS-A /server/nvme-switch # set recover-nvme-switch	Deletes the LUN Set with the specified name.
Step 4	UCS-A /server/nvme-switch* # commit-buffer	Commits the transaction to the system configuration.
Step 5	UCS-A /server/nvme-switch # exit	Exits the MSwitch mode.

	Command or Action	Purpose
Step 6	UCS-A /server # ack-nvme-switch-recovery acknowledge	Acknowledges the MSwitch recovery.
Step 7	UCS-A /server* # commit-buffer	Commits the transaction to the system configuration. Note Do not reset the server during the disaster recovery process.

Example

The following example recovers the MSwitch on server1:

```
UCS-A# scope server 1
UCS-A/server # scope nvme-switch 1
UCS-A/server/nvme-switch # set recover-nvme-switch
UCS-A/server/nvme-switch* # commit-buffer
UCS-A/server/nvme-switch # exit
UCS-A/server # ack-nvme-switch-recovery acknowledge
UCS-A/server* # commit-buffer
```

NVMe Replacement Considerations for B200 M6, Servers

Swapping or replacing NVMe storage devices on the Cisco B200 M6, servers while the system is powered off can result in an error condition. To avoid encountering this error, take the following precautions:

- Replace or hot-swap NVMe SSD storage devices without powering off the server.
- If it is necessary to replace NVMe storage with the server powered off, decommission the server and remove or replace the hardware, then reboot the server. This will recommission the server and NVMe storage will be correctly discovered.

If NVMe storage is replaced when the system is powered off, the controller will be marked as unresponsive. To recover from this condition, re-acknowledge the server.

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.