



Cisco UCS S3260 M5 Server Node For Cisco UCS S3260 Storage Server Service Note

First Published: 2018-01-24

Last Modified: 2020-09-08

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2018–2020 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

Overview

This document describes server node installation and replacement of internal server node components.

This section gives an overview of components and features.

- [Cisco UCS S3260 M5 Software/Firmware Requirements, on page 1](#)
- [Cisco UCS S3260 M5 Server Node External Features, on page 2](#)
- [Externally Viewable LEDs, on page 2](#)
- [Rear Panel Buttons, on page 3](#)
- [Local Console Connection, on page 3](#)
- [S3260 M5 Server Node Internal Component Locations, on page 4](#)
- [I/O Expander Internal Component Locations, on page 4](#)

Cisco UCS S3260 M5 Software/Firmware Requirements


The Cisco UCS S3260 system firmware and software requirements for using the Cisco UCS S3260 M5 server nodes are listed in the following table.

Table 1: S3260 System Minimum Levels for S3260 M5 Server Nodes

Software or Firmware	Minimum Version
Cisco IMC	3.1(3)
BIOS	3.1(3)
Cisco UCS Manager (UCS Manager-controlled system only)	3.2(3)

Cisco UCS S3260 M5 Server Node External Features

Figure 1: M5 Server Node Rear Panel



1	Server node reset button	5	1 Gb Ethernet LOM port
2	Server node power button/LED	6	KVM cable connector Used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB connectors.
3	Unit identification button/LED	7	S3260 M5 server node label (M5 SVRN)
4	Captive thumbscrew for internal NVMe SSD sled	8	Ejector levers (two)

Externally Viewable LEDs

Table 2: S3260 M5 Server Node Externally Viewable LEDs

LED	Definition of States
Server node power	<ul style="list-style-type: none"> Off—There is no power to the server node. Amber—The server node is in standby power mode. Power is supplied only to the Cisco IMC. Green—The server node is in main power mode. Power is supplied to all server node components.
Unit identification	<ul style="list-style-type: none"> Off—The identification LED is not activated. Blue, blinking—The identification LED is activated.
1 Gb Ethernet management port link speed (the LED on the left side of the port)	<ul style="list-style-type: none"> Off—Link speed is 10 Mbps. Amber—Link speed is 100 Mbps. Green—Link speed is 1 Gbps.

LED	Definition of States
1 Gb Ethernet management port link status (the LED on the right side of the port)	<ul style="list-style-type: none"> • Off—No link is present. • Green—Link is active. • Green, Blinking—Traffic is present on the active link.


Rear Panel Buttons

- Reset button—You can hold this button down for 5 seconds and then release it to restart the server node controller chipset if other methods of restarting do not work.
- Server node power button/LED—You can press this button to put the server node in a standby power state or return it to full power instead of shutting down the entire system.
- Unit identification button/LED—This LED can be activated by pressing the button or by activating it from the software interface. This helps to locate a specific server node.

Local Console Connection

The local console connector allows a direct connection to a server node to allow management tasks to be done directly rather than remotely. The port uses the KVM cable N20-BKVM; it has a DB9 serial connector, a VGA connector for a monitor, and dual USB ports for a keyboard and mouse.


Figure 2: KVM Cable for Blade Servers



1	Connector to blade server local console connection	2	DB9 serial connector
3	VGA connector for a monitor	4	2-port USB connector for a mouse and keyboard

S3260 M5 Server Node Internal Component Locations

Figure 3: M5 Server Node Internal Components




1	DIMM sockets (7 sockets for each CPU, 14 total)	5	Trusted platform module (TPM) socket location on server board
2	CPUs and heatsinks (two)	6	Mezzanine-style connector for optional I/O expander module
3	RAID supercap unit on RAID card bracket (two, one on each side)	7	RTC battery CR2032 socket location on server board
4	Cisco storage controller card, dual RAID or dual HBA There is one mezzanine-style connector for the card, but the software designations of the dual controllers are SBMezz1 and SBMezz2.	8	NVMe SSD socket locations, not visible in this view on server board under storage controller card Software designation of connectors SBNVMe1, SBNVMe2) The NVMe SSD sled, which holds up to two 2.5-inch NVMe PCIe SSDs, plugs into both sockets.

I/O Expander Internal Component Locations



Note When using S3260 M5 server nodes in the chassis, the supported storage controllers are supported in the server nodes only. The storage controllers are not supported in the I/O expander.

Figure 4: I/O Expander Internal Components



<p>1</p>	<p>PCIe slot 1 (software designation IOESlot1) x8 electrical lane, x16 physical</p>	<p>3</p>	<p>NVMe SSD sockets (software designations IOENVMe1 and IOENVMe2)</p>
<p>2</p>	<p>PCIe slot 2 (software designation IOESlot2) x8 electrical lane, x16 physical</p>	<p>-</p>	



CHAPTER 2

Replacing an S3260 M5 Server Node

This chapter describes how to install or shut down and replace an M5 server node in the S3260 chassis.

- [Shutting Down an S3260 M5 Server Node, on page 7](#)
- [Replacing an S3260 M5 Server Node, on page 8](#)

Shutting Down an S3260 M5 Server Node

You can invoke a graceful shutdown or a hard shutdown of a server node by using either the Cisco Integrated Management Controller (Cisco IMC) interface, or the power button that is on the face of the server node. Both procedures are included in this section.

Shutting Down a Server Node By Using the Cisco IMC GUI

Step 1 Use a browser and the management IP address of the system to log in to the Cisco IMC GUI.

Step 2 In the **Navigation** pane, click the **Chassis** menu.

Step 3 In the **Chassis** menu, click **Summary**.

Step 4 In the toolbar above the work pane, click the **Host Power** link.

The Server Power Management dialog opens. This dialog lists all server nodes that are present in the system.

Step 5 In the Server Power Management dialog, select one of the following buttons for the server node that you want to shut down:

Caution To avoid data loss or damage to your operating system, you should always invoke a graceful shutdown of the operating system. Do not power off a server if any firmware or BIOS updates are in progress.

- **Shut Down**—Performs a graceful shutdown of the operating system and then powers off the selected server node.
- **Power Off**—Powers off the selected server node, even if tasks are running on that server node.

Step 6 Check that the Chassis Status pane shows the **Power State** as **Off** for the selected server node.

Note It is safe to remove the server node from the chassis when the physical power button on the rear panel turns amber.

Shutting Down a Server Node By Using the Power Button on the Server Node

Step 1 Check the color of the server node power button/status LED:

- Amber—the server node is powered off. It is safe to remove the server node from the chassis.
- Green—The server node is powered on. Continue with the next step.

Step 2 Invoke either a graceful shutdown or an emergency shutdown:

Caution To avoid data loss or damage to your operating system, you should always invoke a graceful shutdown of the operating system. Do not power off a server if any firmware or BIOS updates are in progress.

- Graceful shutdown—Press and release the Power button on the server node. The software performs a graceful shutdown of the operating system and then powers off the server node.
- Emergency shutdown—Press and hold the Power button for four seconds to force the power off the server node.

Note It is safe to remove the server node from the chassis when the physical power button on the rear panel turns amber.

Replacing an S3260 M5 Server Node

Use the procedure in this section to remove or install a server node in a chassis.

To replace components inside a server node or its attached I/O expander, see [Replacing Server Node and I/O Expander Components](#), on page 13.

S3260 M5 Server Node Population Rules

Observe the following rules:

- Do not mix an S3260 M5 server node with a different generation server node (M3 or M4) in the same Cisco UCS S3260 system. An M5 server node can be identified by the “M5 SVRN” label on the rear panel.
- Single-node systems:
 - Cisco IMC releases earlier than 2.0(13): If your S3260 system has only one server node, it must be installed in bay 1.
 - Cisco IMC releases 2.0(13) and later: If your S3260 system has only one server node, it can be installed in either server bay.



Note Whichever bay a server node is installed to, it must have a corresponding SIOC. That is, a server node in bay 1 must be paired with a SIOC in SIOC slot 1; a server node in bay 2 must be paired with a SIOC in SIOC bay 2.

Replacing a Node in a Chassis

The server node is accessed from the rear of the system, so you do not have to pull the system out from the rack.



Note You do not have to power off the system chassis in this procedure. Replacement with chassis powered on is supported if you shut down the server node before removal.



Caution Before you replace a server node, export and save the Cisco IMC configuration from the node if you want that same configuration on the new node. You can import the saved configuration to the new replacement node after you install it. This chapter contains procedures for exporting and importing a configuration.

-
- Step 1** Optional—Export the Cisco IMC configuration from the server node that you are replacing so that you can import it to the replacement server node. If you choose to do this, use the procedure in [Exporting Cisco IMC Configuration From a Server Node, on page 10](#), then return to the next step.
- Step 2** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 3** Remove the server node from the system:
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 4** If the server node has an I/O expander attached, use the procedures in [Disassembling the I/O Expander Assembly, on page 15](#) to remove it from the server node. Then use [Reassembling the I/O Expander Assembly, on page 17](#) to reinstall it on the new server node before you continue with the next step.
- If there is no I/O expander attached, continue with the next step.
- Step 5** Install a new server node:
- With the two ejector levers open, align the new server node with the empty bay. Note these configuration rules:
 - Cisco IMC releases earlier than 2.0(13): If your S3260 system has only one server node, it must be installed in bay 1.
 - Cisco IMC releases 2.0(13) and later: If your S3260 system has only one server node, it can be installed in either server bay.

Note Whichever bay a server node is installed to, it must have a corresponding SIOC. That is, a server node in bay 1 must be paired with a SIOC in SIOC slot 1; a server node in bay 2 must be paired with a SIOC in SIOC bay 2.

- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

- Step 6** Power on the server node.
- Step 7** Perform initial setup on the new server node to assign an IP address and your other preferred network settings. See Initial System Setup in the [Cisco UCS S3260 Storage Server Installation and Service Guide](#).
- Step 8** Optional—Import the Cisco IMC configuration that you saved in step 1. If you choose to do this, use the procedure in [Importing Cisco IMC Configuration to a Server Node, on page 11](#).
- Step 9** Update the node firmware using the Host Upgrade Utility. This will ensure that the firmware version on the node is compatible with the latest system firmware.

See the [Cisco Host Upgrade Utility User Guide For S3260 Storage Servers](#) for instructions on updating the firmware.

Exporting Cisco IMC Configuration From a Server Node

This operation can be performed using either the GUI or CLI interface of the Cisco IMC. The example in this procedure uses the CLI commands. For more information, see *Exporting a Cisco IMC Configuration* in the CLI and GUI guides for S3260 Storage Servers here: [Configuration Guides](#).

Step 1 Log in to the IP address and CLI interface of the server node that you are replacing.

Step 2 Enter the following commands as you are prompted:

```
Server# scope cimc
Server /cimc# scope import-export
Server /cimc/import-export# export-config <protocol> <ip-address> <path-and-filename>
```

Step 3 Enter the user name, password, and pass phrase.

This sets the user name, password, and pass phrase for the file that you are exporting. The export operation begins after you enter a pass phrase, which can be anything that you choose.

To determine whether the export operation has completed successfully, use the **show detail** command. To abort the operation, type CTRL+C.

The following is an example of an export operation. In this example, the TFTP protocol is used to export the configuration to IP address 192.0.2.34, in file /ucs/backups/cimc5.xml.

```
Server# scope cimc
Server /cimc # scope import-export
Server /cimc/import-export # export-config tftp 192.0.2.34 /ucs/backups/cimc5.xml
Username:xxxx
Password:****
Passphrase:***
Export config started. Please check the status using "show detail".
Server /cimc/import-export # show detail
Import Export:
Operation: EXPORT
Status: COMPLETED
```

```
Error Code: 100 (No Error)
Diagnostic Message: NONE
```

Importing Cisco IMC Configuration to a Server Node

This operation can be performed using either the GUI or CLI interface of the Cisco IMC. The example in this procedure uses the CLI commands. For more information, see *Importing a Cisco IMC Configuration* in the CLI and GUI guides for S3260 Storage Servers here: [Configuration Guides](#).

Step 1 SSH into the CLI interface of the new server node.

Step 2 Enter the following commands as you are prompted:

```
Server# scope cimc
Server /cimc# scope import-export
Server /cimc/import-export# import-config <protocol> <ip-address> <path-and-filename>
```

Step 3 Enter the user name, password, and pass phrase.

This should be the user name, password, and pass phrase that you used during the export operation. The import operation begins after you enter the pass phrase.

The following is an example of an import operation. In this example, the TFTP protocol is used to import the configuration to the server node from IP address 192.0.2.34, from file /ucs/backups/cimc5.xml.

```
Server# scope cimc
Server /cimc # scope import-export
Server /cimc/import-export # import-config tftp 192.0.2.34 /ucs/backups/cimc5.xml
Username:xxxx
Password:****
Passphrase:***
Export config started. Please check the status using "show detail".
Server /cimc/import-export # show detail
Import Export:
Operation: IMPORT
Status: COMPLETED
Error Code: 100 (No Error)
Diagnostic Message: NONE
```



CHAPTER 3

Replacing Server Node and I/O Expander Components

This chapter contains procedures for replacing components internal to the server node and the optional I/O expander.

- [Preparing For Component Replacement, on page 13](#)
- [Component Replacement Procedures, on page 20](#)
- [Service Headers on the Server Node Board, on page 69](#)


Preparing For Component Replacement

This section contains procedures that are referred to from the replacement procedures.

Removing an S3260 M5 Server Node or I/O Expander Top Cover

The optional I/O expander and the server node use the same top cover. If an I/O expander is attached to the top of the server node, the top cover is on the I/O expander, as shown in the side view in the following figure. In this case, there is also an intermediate cover between the server node and the I/O expander.

Figure 5: Side View, Server Node With Attached I/O Expander




1	Top Cover In this view, the top cover is on the attached I/O expander.	3	Intermediate cover securing screws (two on each side)
2	Intermediate cover In this view, the intermediate cover is attached to the server node.	-	



Note You do not have to power off the chassis in the next step. Replacement with the chassis powered on is supported if you shut down the server node before removal.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node from the system (or server node with attached I/O expander, if present):
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 3** Remove the top cover from the server node or the I/O expander (if present):
- Lift the cover latch handle to an upright position.
 - Turn the cover latch handle 90 degrees to release the lock.
 - Slide the cover toward the rear (toward the rear-panel buttons) and then lift it from the server node or I/O expander (if present).

Figure 6: Top View, Server Node or I/O Expander Top Cover



1	Top cover on server node or I/O expander	2	Cover latch handle (shown in closed, flat position)
---	--	---	---

Step 4 Reinstall the top cover:

- a) Set the cover in place on the server node or I/O expander (if present), offset about one inch toward the rear. Pegs on the inside of the cover must engage the tracks on the server node or I/O expander base.
- b) Push the cover forward until it stops.
- c) Turn the latch handle 90 degrees to close the lock.
- d) Fold the latch handle flat.

Step 5 Reinstall the server node to the chassis:

- a) With the two ejector levers open, align the server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 6 Power on the server node.

Removing an I/O Expander From a Server Node

This topics in this section describe how to remove and then reinstall an I/O expander and intermediate cover from a server node so that you can access the components inside the server node.



Note You do not have to power off the chassis in the next step. Replacement with the chassis powered on is supported if you shut down the server node before removal.

Disassembling the I/O Expander Assembly

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).


Step 2 Remove the server node with attached I/O expander from the system:

- a) Grasp the two ejector levers and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.

Step 3 Remove the top cover from the I/O expander:

- a) Lift the cover latch handle to an upright position.
- b) Turn the cover latch handle 90 degrees to release the lock.
- c) Slide the cover toward the rear (toward the rear-panel buttons) and then lift it from the I/O expander.

Figure 7: Top View, Server Node or I/O Expander Top Cover




1	Top cover on server node or I/O expander	2	Cover latch handle (shown in closed, flat position)
---	--	---	---

Step 4 Remove the I/O expander from the server node:

- a) Remove the five screws that secure the I/O expander to the top of the server node.

Figure 8: I/O Expander Securing Screws (Five)




- b) Use two small flat-head screwdrivers (1/4-inch or equivalent) to help separate the connector on the underside of the I/O expander from the socket on the server node board.

Insert a screwdriver about 1/2 inch into the “REMOVAL SLOT” that is marked with an arrow on each side of the I/O expander. Then lift up evenly on both screwdrivers at the same time to separate the connectors and lift the I/O expander about 1/2 inch.

- c) Grasp the two handles on the I/O expander board and lift it straight up.

Figure 9: Separating the I/O Expander From the Server Node



1	Side view showing REMOVAL SLOT for screwdriver insertion (one on each side of the I/O expander)	2	Rear view of server node with I/O expander
----------	---	----------	--

- Step 5** Remove the intermediate cover from the server node:
- a) Remove the four screws that secure the intermediate cover and set them aside. There are two screws on each side of the intermediate cover.
 - b) Slide the intermediate cover toward the rear (toward the rear-panel buttons) and then lift it from the server node.
- Step 6** Install the server node (with attached I/O expander, if present) to the chassis:
- a) With the two ejector levers open, align the new server node with the empty bay.
 - b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.
- Step 7** Power on the server node.

Reassembling the I/O Expander Assembly

Before you begin

This procedure assumes that you have already taken apart the I/O expander/server node assembly.

- Step 1** Reinstall the intermediate cover to the server node:
- a) Set the intermediate cover in place on the server node, offset about one inch toward the rear. Pegs on the inside of the cover must engage the tracks on the server node base.


- b) Push the cover forward until it stops.
- c) Reinstall the four screws that secure the intermediate cover.

Step 2 Reinstall the I/O expander to the server node:

Caution Use caution to align all features of the I/O expander with the intermediate cover and server node before mating the connector on the underside of the expander with the socket on the server board. The connector can be damaged if alignment is not correct.

- a) Carefully align the I/O expander with the alignment pegs on the top of the intermediate cover.
- b) Set the I/O expander down on the intermediate cover and lower it gently to mate the mezzanine connector and the socket on the server board.


Figure 10: Reinstalling the I/O Expander to the Server Node



1	Intermediate cover attached to server node	3	I/O expander mezzanine connector
2	Intermediate cover screws (two on each side of cover)	4	Alignment pegs on intermediate cover

- c) If an NVMe SSD is present in the right-hand socket (IOENVMe2) of the I/O expander, you must remove it to access the PRESS HERE plate in the next step.
See [Replacing an NVMe SSD in the I/O Expander, on page 68](#) to remove it, then return to the next step.
- d) Press down firmly on the plastic plate marked “PRESS HERE” to fully seat the connector to the server node board.

Figure 11: I/O Expander, Showing PRESS HERE Plate



e) If you removed an NVMe SSD to access the PRESS HERE plate, reinstall it now.

See [Replacing an NVMe SSD in the I/O Expander](#), on page 68.


Caution Before you reinstall the I/O expander securing screws, you must use the alignment tool (UCSC-C3K-M4IOTOOL) in the next step to ensure alignment of the connectors that plug into the internal chassis backplane. Failure to ensure alignment might damage the sockets on the backplane.

Step 3

Insert the four pegs of the alignment tool into the holes that are built into the connector side of the server node and I/O expander. Ensure that the alignment tool fits into all four holes and lies flat.

The alignment tool is shipped with systems that are ordered with an I/O expander. It is also shipped with I/O expander replacement spares. You can order the tool using Cisco PID UCSC-C3K-M4IOTOOL.

Figure 12: Using the I/O Expander Alignment Tool



1	Alignment Tool	3	Alignment tool installed
2	Connector side of server node and I/O expander	-	

- Step 4** Reinstall and tighten the five screws that secure the I/O expander to the top of the server node.
- Step 5** Remove the alignment tool.
- Step 6** Reinstall the top cover to the I/O expander:
- Set the cover in place on the server node or I/O expander (if present), offset about one inch toward the rear. Pegs on the inside of the cover must engage the tracks on the server node or I/O expander base.
 - Push the cover forward until it stops.
 - Turn the latch handle 90 degrees to close the lock.
 - Fold the latch handle flat.
- Step 7** Reinstall the server node to the chassis:
- With the two ejector levers open, align the server node with the empty bay.
 - Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.
- Step 8** Power on the server node.

Component Replacement Procedures

This section contains procedures for replacing components in the S3260 M5 server node and the optional I/O expander.

Replacing CPUs and Heatsinks

Each server node contains two CPUs.

Special Information For *Upgrades* to Second Generation Intel Xeon Scalable Processors



Caution

You must upgrade your server firmware to the required minimum level before you upgrade to the Second Generation Intel Xeon Scalable Processors that are supported in this server. Older firmware versions cannot recognize the new CPUs, so this would result in a non-bootable server.

The minimum software and firmware versions required for this server to support Second Generation Intel Xeon Scalable processors are as follows:

Table 3: Second Generation Intel Xeon Scalable Processor Minimum Requirements

Software or Firmware	Minimum Version
Server Cisco IMC	4.0(4)
Server BIOS	4.0(4)
Cisco UCS Manager (UCS-integrated servers only)	4.0(4)

Do one of the following actions:

- If your server's firmware and Cisco UCS Manager software are already at the required minimums shown above (or later), you can replace the CPU hardware by using the procedure in this section.
- If your server's firmware and Cisco UCS Manager software are earlier than the required levels, use the instructions in the [Cisco UCS C- and S-Series M5 Servers Upgrade Guide For Next Gen Intel Xeon Processors](#) to upgrade your software. After you upgrade the software, return to this section as directed to replace the CPU hardware.

CPU Configuration Rules

- A server node must have two CPUs to operate.
- **For Intel Xeon Scalable processors (first generation):** The maximum combined memory allowed in the DIMM slots controlled by any one CPU is 768 GB. To populate the DIMM slots with more than 768 GB of combined memory, you must use a high-memory CPU that has a PID that ends with an "M", for example, UCS-CPU-6134M.
- **For Second Generation Intel Xeon Scalable processors:** These Second Generation CPUs have three memory tiers. These rules apply on a *per-socket* basis:
 - If the CPU socket has up to 1 TB of memory installed, a CPU with no suffix can be used (for example, Gold 6240).
 - If the CPU socket has 1 TB or more (up to 2 TB) of memory installed, you must use a CPU with an M suffix (for example, Platinum 8276M).
 - If the CPU socket has 2 TB or more (up to 4.5 TB) of memory installed, you must use a CPU with an L suffix (for example, Platinum 8270L).

Tools Required For CPU Replacement

You need the following tools and equipment for this procedure:

- T-30 Torx driver—Supplied with replacement CPU.
- #1 flat-head screwdriver—Supplied with replacement CPU.
- CPU assembly tool—Supplied with replacement CPU. Orderable separately as Cisco PID UCS-CPUAT=.
- Heatsink cleaning kit—Supplied with replacement CPU. Orderable separately as Cisco PID UCSX-HSCK=.
- Thermal interface material (TIM)—Syringe supplied with replacement CPU. Use only if you are reusing your existing heatsink (new heatsinks have a pre-applied pad of TIM). Orderable separately as Cisco PID UCS-CPU-TIM=.

See also [Additional CPU-Related Parts to Order with RMA Replacement CPUs, on page 28](#).

Replacing a CPU and Heatsink




Caution CPUs and their sockets are fragile and must be handled with extreme care to avoid damaging pins. The CPUs must be installed with heatsinks and thermal interface material to ensure cooling. Failure to install a CPU correctly might result in damage to the server.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- a) Grasp the two ejector levers and pinch their latches to release the levers.
 - b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - c) Pull the server node straight out from the system.
- Step 3** Do one of the following to access the component inside the server node:
- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.
- Step 4** Remove the existing CPU/heatsink assembly from the server:
- a) Use the T-30 Torx driver that is supplied with the replacement CPU to loosen the four captive nuts that secure the assembly to the motherboard standoffs.

Note Alternate loosening the heatsink nuts evenly so that the heatsink remains level as it is raised. Loosen the heatsink nuts in the order shown on the heatsink label: 4, 3, 2, 1.
 - b) Lift straight up on the CPU/heatsink assembly and set it heatsink-down on an antistatic surface.

Figure 13: Removing the CPU/Heatsink Assembly



1	Heatsink	4	CPU socket on motherboard
2	Heatsink captive nuts (two on each side)	5	T-30 Torx driver
3	CPU carrier (below heatsink in this view)	-	


Step 5

Separate the heatsink from the CPU assembly (the CPU assembly includes the CPU and the CPU carrier):

- a) Place the heatsink with CPU assembly so that it is oriented upside-down as shown below.

Note the thermal-interface material (TIM) breaker location. TIM BREAKER is stamped on the CPU carrier next to a small slot.

Figure 14: Separating the CPU Assembly From the Heatsink



1	CPU carrier	4	CPU-carrier inner-latch nearest to the TIM breaker slot
2	CPU	5	#1 flat-head screwdriver inserted into TIM breaker slot
3	TIM BREAKER slot in CPU carrier	-	

- b) Pinch inward on the CPU-carrier inner-latch that is nearest the TIM breaker slot and then push up to disengage the clip from its slot in the heatsink corner.
- c) Insert the blade of a #1 flat-head screwdriver into the slot marked TIM BREAKER.

Caution In the following step, do not pry on the CPU surface. Use gentle rotation to lift on the plastic surface of the CPU carrier at the TIM breaker slot. Use caution to avoid damaging the heatsink surface.

- d) Gently rotate the screwdriver to lift up on the CPU until the TIM on the heatsink separates from the CPU.

Note Do not allow the screwdriver tip to touch or damage the green CPU substrate.

- e) Pinch the CPU-carrier inner-latch at the corner opposite the TIM breaker and push up to disengage the clip from its slot in the heatsink corner.
- f) On the remaining two corners of the CPU carrier, gently pry outward on the outer-latches and then lift the CPU-assembly from the heatsink.

Note Handle the CPU-assembly by the plastic carrier only. Do not touch the CPU surface. Do not separate the CPU from the carrier.


Step 6

The new CPU assembly is shipped on a CPU assembly tool. Take the new CPU assembly and CPU assembly tool out of the carton.

If the CPU assembly and CPU assembly tool become separated, note the alignment features shown below for correct orientation. The pin 1 triangle on the CPU carrier must be aligned with the angled corner on the CPU assembly tool.

Caution CPUs and their sockets are fragile and must be handled with extreme care to avoid damaging pins.

Figure 15: CPU Assembly Tool, CPU Assembly, and Heatsink Alignment Features



1	CPU assembly tool	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU assembly (CPU in plastic carrier)	5	Triangle cut into carrier (pin 1 alignment feature)
3	Heatsink	6	Angled corner on CPU assembly tool (pin 1 alignment feature)


Step 7

Apply new TIM to the heatsink:

Note The heatsink must have new TIM on the heatsink-to-CPU surface to ensure proper cooling and performance.

- If you are installing a new heatsink, it is shipped with a pre-applied pad of TIM. Go to step 5.
 - If you are reusing a heatsink, you must remove the old TIM from the heatsink and then apply new TIM to the CPU surface from the supplied syringe. Continue with step a below.
- a) Apply the cleaning solution that is included with the heatsink cleaning kit (UCSX-HSCK=) to the old TIM on the heatsink and let it soak for a least 15 seconds.
 - b) Wipe all of the TIM off the heatsink using the soft cloth that is included with the heatsink cleaning kit. Be careful to avoid scratching the heatsink surface.
 - c) Using the syringe of TIM provided with the new CPU (UCS-CPU-TIM=), apply 1.5 cubic centimeters (1.5ml) of thermal interface material to the top of the CPU. Use the pattern shown below to ensure even coverage.

Figure 16: Thermal Interface Material Application Pattern



Step 8 With the CPU assembly on the CPU assembly tool, set the heatsink onto the CPU assembly. Note the pin 1 alignment features for correct orientation. Push down gently until you hear the corner clips of the CPU carrier click onto the heatsink corners.


Caution In the following step, use extreme care to avoid touching or damaging the CPU contacts or the CPU socket pins.

Step 9 Install the CPU/heatsink assembly to the server:

- a) Lift the heatsink with attached CPU assembly from the CPU assembly tool.
- b) Align the CPU with heatsink over the CPU socket on the motherboard, as shown below.

Note the alignment features. The pin 1 angled corner on the heatsink must align with the pin 1 angled corner on the CPU socket. The CPU-socket posts must align with the guide-holes in the assembly.

Figure 17: Installing the Heatsink/CPU Assembly to the CPU Socket



1	Guide hole in assembly (two)	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU socket alignment post (two)	5	Angled corner on socket (pin 1 alignment feature)
3	CPU socket leaf spring	-	

- c) Set the heatsink with CPU assembly down onto the CPU socket.
- d) Use the T-30 Torx driver that is supplied with the replacement CPU to tighten the four captive nuts that secure the heatsink to the motherboard standoffs.

Caution Alternate tightening the heatsink nuts evenly so that the heatsink remains level while it is lowered. Tighten the heatsink nuts in the order shown on the heatsink label: 1, 2, 3, 4. The captive nuts must be fully tightened so that the leaf springs on the CPU socket lie flat.

Step 10

Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 11 Power on the server node.

Additional CPU-Related Parts to Order with RMA Replacement CPUs

When a return material authorization (RMA) of the CPU is done on a server node, additional parts might not be included with the CPU spare. The TAC engineer might need to add the additional parts to the RMA to help ensure a successful replacement.



Note The following items apply to CPU *replacement* scenarios. If you are replacing a server node and *moving* existing CPUs to the new board, you do not have to separate the heatsink from the CPU. See [Additional CPU-Related Parts to Order with RMA Replacement Server Node, on page 28](#).

- Scenario 1—You are reusing the existing heatsinks:
 - Heatsink cleaning kit (UCSX-HSCK=)
 - Thermal interface material (TIM) kit for M5 servers (UCS-CPU-TIM=)
- Scenario 2—You are replacing the existing heatsinks:



Caution Use only the correct heatsink for your CPUs to ensure proper cooling.

- Heatsink: UCS-S3260-M5HS=
- Heatsink cleaning kit (UCSX-HSCK=)
- Scenario 3—You have a damaged CPU carrier (the plastic frame around the CPU):
 - CPU Carrier: UCS-M5-CPU-CAR=
 - #1 flat-head screwdriver (for separating the CPU from the heatsink)
 - Heatsink cleaning kit (UCSX-HSCK=)
 - Thermal interface material (TIM) kit for M5 servers (UCS-CPU-TIM=)

A CPU heatsink cleaning kit is good for up to four CPU and heatsink cleanings. The cleaning kit contains two bottles of solution, one to clean the CPU and heatsink of old TIM and the other to prepare the surface of the heatsink.

New heatsink spares come with a pre-applied pad of TIM. It is important to clean any old TIM off of the CPU surface prior to installing the heatsinks. Therefore, even when you are ordering new heatsinks, you must order the heatsink cleaning kit.

Additional CPU-Related Parts to Order with RMA Replacement Server Node

When a return material authorization (RMA) of a server node is done, you move existing CPUs to the new chassis.



Note Unlike previous generation CPUs, the M5 server CPUs do not require you to separate the heatsink from the CPU when you *move* the CPU-heatsink assembly. Therefore, no additional heatsink cleaning kit or thermal-interface material items are required.

- The only tool required for moving a CPU/heatsink assembly is a T-30 Torx driver.

To move a CPU to a new chassis, use the procedure in [Moving an M5 Generation CPU, on page 29](#).

Moving an M5 Generation CPU

Tool required for this procedure: T-30 Torx driver



Caution When you receive a replacement server for an RMA, it includes dust covers on all CPU sockets. These covers protect the socket pins from damage during shipping. You must transfer these covers to the system that you are returning, as described in this procedure.

Step 1


When moving an M5 CPU to a new server, you do not have to separate the heatsink from the CPU. Perform the following steps:

- a) Use a T-30 Torx driver to loosen the four captive nuts that secure the assembly to the board standoff.

Note Alternate loosening the heatsink nuts evenly so that the heatsink remains level as it is raised. Loosen the heatsink nuts in the order shown on the heatsink label: 4, 3, 2, 1.

- b) Lift straight up on the CPU/heatsink assembly to remove it from the board.
- c) Set the CPUs with heatsinks aside on an anti-static surface.

Figure 18: Removing the CPU/Heatsink Assembly



1	Heatsink	4	CPU socket on motherboard
2	Heatsink captive nuts (two on each side)	5	T-30 Torx driver
3	CPU carrier (below heatsink in this view)	-	


Step 2

Transfer the CPU socket covers from the new system to the system that you are returning:

- a) Remove the socket covers from the replacement system. Grasp the two recessed finger-grip areas marked "REMOVE" and lift straight up.

Note Keep a firm grasp on the finger-grip areas at both ends of the cover. Do not make contact with the CPU socket pins.

Figure 19: Removing a CPU Socket Dust Cover



1	Finger-grip areas marked "REMOVE"	-	
---	-----------------------------------	---	--


- b) With the wording on the dust cover facing up, set it in place over the CPU socket. Make sure that all alignment posts on the socket plate align with the cutouts on the cover.

Caution In the next step, do not press down anywhere on the cover except the two points described. Pressing elsewhere might damage the socket pins.

- c) Press down on the two circular markings next to the word "INSTALL" that are closest to the two threaded posts (see the following figure). Press until you feel and hear a click.

Note You must press until you feel and hear a click to ensure that the dust covers do not come loose during shipping.

Figure 20: Installing a CPU Socket Dust Cover



-	Press down on the two circular marks next to the word INSTALL.	-	
---	--	---	--


Step 3

Install the CPUs to the new system:

- a) On the new board, align the assembly over the CPU socket, as shown below.

Note the alignment features. The pin 1 angled corner on the heatsink must align with the pin 1 angled corner on the CPU socket. The CPU-socket posts must align with the guide-holes in the assembly.

Figure 21: Installing the Heatsink/CPU Assembly to the CPU Socket



1	Guide hole in assembly (two)	4	Angled corner on heatsink (pin 1 alignment feature)
2	CPU socket alignment post (two)	5	Angled corner on socket (pin 1 alignment feature)
3	CPU socket leaf spring	-	

- b) On the new board, set the heatsink with CPU assembly down onto the CPU socket.
- c) Use a T-30 Torx driver to tighten the four captive nuts that secure the heatsink to the board standoffs.

Note Alternate tightening the heatsink nuts evenly so that the heatsink remains level while it is lowered. Tighten the heatsink nuts in the order shown on the heatsink label: 1, 2, 3, 4. The captive nuts must be fully tightened so that the leaf springs on the CPU socket lie flat.

Replacing Memory DIMMs

There are 14 DIMM sockets on the server node board, 7 DIMMs controlled by each CPU.



Caution DIMMs and their sockets are fragile and must be handled with care to avoid damage during installation.



Caution Cisco does not support third-party DIMMs. Using non-Cisco DIMMs in the system might result in system problems or damage to the motherboard.



Note To ensure the best system performance, it is important that you are familiar with memory performance guidelines and population rules before you install or replace the memory.


For additional information about troubleshooting DIMM memory issues, see the document [Troubleshoot DIMM Memory Issues in UCS](#).

DIMM Sockets

The following figure shows the DIMM sockets and how they are numbered on an S3260 M5 server node board.

- A server node has 14 DIMM sockets (7 for each CPU).
- Channels are labeled with letters as shown in the following figure. For example, channel A = DIMM sockets A1, A2.
- Channels A and G use two DIMMs per channel (DPC); all other channels use one DPC.

Figure 22: S3260 M5 DIMM and CPU Numbering



DIMM Population Rules

Observe the following guidelines when installing or replacing DIMMs:

- For optimal performance, spread DIMMs evenly across both CPUs and all channels. Populate the DIMM slots of each CPU identically.
- For optimal performance, populate DIMMs in the order shown in the following table, depending on the number of DIMMs per CPU.



Note The table below lists recommended configurations. Using 5 DIMMs per CPU is not recommended.

Table 4: DIMM Population Order

Number of DIMMs per CPU (Recommended Configurations)	Populate CPU 1 Slots		Populate CPU 2 Slots	
	Blue #1 Slots	Black #2 Slot	Blue #1 Slots	Black #2 Slot
1	(A1)	-	(G1)	-
2	(A1, B1)	-	(G1, H1)	-
3	(A1, B1, C1)	-	(G1, H1, J1)	-
4	(A1, B1); (D1, E1)	-	(G1, H1); (K1, L1)	-
6	(A1, B1); (C1, D1); (E1, F1)	-	(G1, H1); (J1, K1); (L1, M1)	-
7	(A1, B1); (C1, D1); (E1, F1)	(A2)	(G1, H1); (J1, K1); (L1, M1)	(G2)

- Observe the DIMM mixing rules in the following table.

Table 5: DIMM Mixing Rules

DIMM Parameter	DIMMs in the same Channel	DIMMs in the Same Bank
DIMM capacity	You can mix different capacity DIMMs in the same channel (for example, A1, A2).	You can mix different capacity DIMMs in the same bank (for example, A1, B1, C1). However, for optimal performance DIMMs in the same bank should have the same capacity.
DIMM speed	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPU installed in the channel.	You can mix speeds, but DIMMs will run at the speed of the slowest DIMMs/CPU installed in the bank.

DIMM Parameter	DIMMs in the same Channel	DIMMs in the Same Bank
DIMM type	You cannot mix DIMM types in a channel.	You cannot mix DIMM types in a bank.

Memory Mirroring Mode

When you enable memory mirroring mode, the memory subsystem simultaneously writes identical data to two channels. If a memory read from one of the channels returns incorrect data due to an uncorrectable memory error, the system automatically retrieves the data from the other channel. A transient or soft error in one channel does not affect the mirrored data and operation continues.

Memory mirroring reduces the amount of memory available to the operating system by 50 percent because only one of the two populated channels provides data.

Replacing DIMMs

-
- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 3** Do one of the following to access the component inside the server node:
- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.
- Step 4** Locate the faulty DIMM and remove it from the socket on the riser by opening the ejector levers at both ends of the DIMM socket.
- Step 5** Install a new DIMM:
- Note** Before installing DIMMs, refer to the population guidelines. See [DIMM Population Rules, on page 35](#)
- Align the new DIMM with the socket on the riser. Use the alignment key in the DIMM socket to correctly orient the DIMM.
 - Push the DIMM into the socket until it is fully seated and the ejector levers on either side of the socket lock into place.
- Step 6** Do one of the following:
- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.
- Step 7** Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 8 Power on the server node.

Replacing Intel Optane DC Persistent Memory Modules

This topic contains information for replacing Intel Optane Data Center Persistent Memory modules (DCPMMs), including population rules. DCPMMs have the same form-factor as DDR4 DIMMs and they install to DIMM slots.



Note DCPMMs require Second Generation Intel Xeon Scalable processors. You must upgrade the server firmware and BIOS to version 4.0(4) or later and install the supported Second Generation Intel Xeon Scalable processors before installing DCPMMs.



Caution DCPMMs and their sockets are fragile and must be handled with care to avoid damage during installation.



Note To ensure the best server performance, it is important that you are familiar with memory performance guidelines and population rules before you install or replace DCPMMs.

In this server, DCPMMs can be configured to operate in one mode at this time:

- App Direct Mode: The module operates as a solid-state disk storage device. Data is saved and is non-volatile.



Note In this server, App Direct Mode can be used non-interleaved only.


Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines

This topic describes the rules and guidelines for maximum memory performance when using Intel Optane DC persistent memory modules (DCPMMs) with DDR4 DRAM DIMMs.

DIMM Slot Numbering

The following figure shows the numbering of the DIMM slots on the server motherboard.

Figure 23: DIMM Slot Numbering



Configuration Rules

Observe the following rules and guidelines:

- To use DCPMMs in this server node, two CPUs must be installed.
- Intel Optane DC persistent memory modules require Second Generation Intel Xeon Scalable processors. You must upgrade the server firmware and BIOS to version 4.0(4) or later and then install the supported Second Generation Intel Xeon Scalable processors before installing DCPMMs.
- When using DCPMMs in a server:
 - The DDR4 DIMMs installed in the server must all be the same size.
 - The DCPMMs installed in the server must all be the same size and must have the same SKU.
- The DCPMMs run at 2666 MHz. If you have 2933 MHz RDIMMs or LRDIMMs in the server and you add DCPMMs, the main memory speed clocks down to 2666 MHz to match the speed of the DCPMMs.
- Each DCPMM draws 18 W sustained, with a 20 W peak.
- When App Direct mode is used in this server, it must be non-interleaved.
- The following table shows the supported DCPMM configuration for this server. Fill the DIMM slots for CPU 1 and CPU2 as shown.

Figure 24: Supported DCPMM Configuration For App Direct Mode (Non-Interleaved), Dual-CPU

DIMM to DCPMM Count	CPU 1									
	IMC1					IMC0				
	Channel 2		Channel 1		Channel 0	Channel 2		Channel 1		Channel 0
	F1	E1	D1	C1	B1	A2	A1			
6 to 1	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DCPMM	DIMM		

DIMM to DCPMM Count	CPU 2									
	IMC1					IMC0				
	Channel 2		Channel 1		Channel 0	Channel 2		Channel 1		Channel 0
	M1	L1	K1	J1	H1	G2	G1			
6 to 1	DIMM	DIMM	DIMM	DIMM	DIMM	DIMM	DCPMM	DIMM		

Installing Intel Optane DC Persistent Memory Modules



Note DCPMM configuration is always applied to all DCPMMs in a region, including a replacement DCPMM. You cannot provision a specific replacement DCPMM on a preconfigured server.

Understand which mode your DCPMM is operating in. App Direct mode has some additional considerations in this procedure.



Caution Replacing a DCPMM in App-Direct mode requires all data to be wiped from the DCPMM. Make sure to backup or offload data before attempting this procedure.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).

Step 2 Remove the server node (with attached I/O expander, if present) from the system:

- a) Grasp the two ejector levers and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.

Step 3 Do one of the following to access the component inside the server node:

- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.

Caution If you are moving DCPMMs with active data (persistent memory) from one server to another as in an RMA situation, each DCPMM must be installed to the identical position in the new server. Note the positions of each DCPMM or temporarily label them when removing them from the old server.

Step 4 For App Direct mode, backup the existing data stored in all Optane DIMMs to some other storage.

Step 5 For App Direct mode, remove the Persistent Memory policy which will remove goals and namespaces automatically from all Optane DIMMs.

Step 6 Locate the faulty DCPMM and remove it from the socket on the riser by opening the ejector levers at both ends of the DIMM socket.

Step 7 Install a new DCPMM:

Note Before installing DCPMM, refer to the population guidelines. See [Intel Optane DC Persistent Memory Module Population Rules and Performance Guidelines, on page 37](#).

- a) Align the new DCPMM with the socket on the riser. Use the alignment key in the DIMM socket to correctly orient the DCPMM.
- b) Push the DIMM into the socket until it is fully seated and the ejector levers on either side of the socket lock into place.

Step 8 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 9 Return the server node to the chassis:

- With the two ejector levers open, align the new server node with the empty bay.
- Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 10 Power on the server node.

Step 11 Perform post-installation actions:

Note If your Persistent Memory policy is Host Controlled, you must perform the following actions from the OS side.

- If the existing configuration is fully or partly in App-Direct mode and new DCPMM is also in App-Direct mode, then ensure that all DCPMMs are at the latest matching firmware level and also re-provision the DCPMMs by creating a new goal.
 - For App Direct mode, reapply the Persistent Memory policy.
 - For App Direct mode, restore all the offloaded data to the DCPMMs.

There a number of tools for configuring goals, regions, and namespaces. See [Server BIOS Setup Utility Menu for DCPMM, on page 40](#).

Server BIOS Setup Utility Menu for DCPMM

DCPMMs can be configured by using the server's BIOS Setup Utility, Cisco IMC, Cisco UCS Manager, or OS-related utilities.

- To use the BIOS Setup Utility, see the section below.
- To use Cisco IMC, see the configuration guides for Cisco IMC 4.0(4) or later: [Cisco IMC CLI and GUI Configuration Guides](#)
- To use Cisco UCS Manager, see the configuration guides for Cisco UCS Manager 4.0(4) or later: [Cisco UCS Manager CLI and GUI Configuration Guides](#)

The server BIOS Setup Utility includes menus for DCPMMs. They can be used to view or configure DCPMM regions, goals, and namespaces, and to update DCPMM firmware.

To open the BIOS Setup Utility, press **F2** when prompted onscreen during a system boot.

The DCPMM menu is on the Advanced tab of the utility:

Advanced > Intel Optane DC Persistent Memory Configuration

From this tab, you can access other menus:

- **DIMMs**: Displays the installed DCPMMs. From this page, you can update DCPMM firmware and configure other DCPMM parameters.

- Monitor health
- Update firmware
- Configure security

You can enable security mode and set a password so that the DCPMM configuration is locked. When you set a password, it applies to all installed DCPMMs. Security mode is disabled by default.

- Configure data policy
- Regions: Displays regions and their persistent memory types. When using App Direct mode with interleaving, the number of regions is equal to the number of CPU sockets in the server. When using App Direct mode without interleaving, the number of regions is equal to the number of DCPMMs in the server.

From the Regions page, you can configure memory goals that tell the DCPMM how to allocate resources.

- Create goal config
- Namespaces: Displays namespaces and allows you to create or delete them when persistent memory is used. Namespaces can also be created when creating goals. A namespace provisioning of persistent memory applies only to the selected region.
Existing namespace attributes such as the size cannot be modified. You can only add or delete namespaces.
- Total capacity: Displays the total DCPMM resource allocation across the server.

Updating the DCPMM Firmware Using the BIOS Setup Utility

You can update the DCPMM firmware from the BIOS Setup Utility if you know the path to the .bin files. The firmware update is applied to all installed DCPMMs.

1. Navigate to **Advanced > Intel Optane DC Persistent Memory Configuration > DIMMs > Update firmware**
2. Under **File:**, provide the file path to the .bin file.
3. Select **Update**.

Replacing a Storage Controller

The Cisco storage controller card (RAID or HBA) connects to a mezzanine-style socket on the server board. See [Supported Storage Controllers, on page 73](#) for information about the controllers supported in the S3260 M5 server node.




Note When using S3260 M5 server nodes in the chassis, the supported storage controllers are supported in the server nodes only. The controllers are not supported in the I/O expander.



Note Do not mix different storage controllers in the same system. If the system has two server nodes, they must both contain the same controller.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 3** Do one of the following to access the component inside the server node:
- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.
- Step 4** Remove a storage controller card:
- Loosen the captive thumbscrews that secure the card and its bracket to the board.
 - Grasp the card at both ends and lift it evenly to disengage the connector on the underside of the card from the mezzanine socket.
- Note** For the supported Cisco RAID controller card, the two supercap backup units come already attached to the bracket of a new RAID card, so you do not have to remove them separately.

Figure 25: Screw Locations on Storage Controller (Six)



- Step 5** Install a new storage card:

- Note** The storage controllers supported in this S3260 M5 server node require chassis part number 68-5286-06 or later. The chassis motherboard in earlier chassis versions does not support this controller. You can determine the chassis part number by looking on the part-number label on the top-front of the chassis, or by using the inventory-all command, as shown in the following example using the Cisco IMC CLI:

```

Server# scope chassis
Server/chassis# inventory-refresh
Server/chassis# inventory-all
|
CHS_FRU (ID1)
Board Mfg           : Cisco Systems Inc
Board Product       : N/A
Board Serial        : FCH20317RVA
Board Part Number   : 73-16125-03
Board Extra         : A12V02
Board Extra         : 0000000000
Product Manufacturer : Cisco Systems Inc
Product Name        : N/A
Product Part Number : 68-5286-07
Product Serial      : FCH20317RVA

```

- Align the card over the mezzanine socket on the server board while aligning the six thumbscrews with the standoffs.
- Press down on both ends of the card to engage the connector on the underside of the card with the mezzanine socket.
- Tighten the six captive screws that secure the card to the board.

Step 6 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 7 Return the server node to the chassis:

- With the two ejector levers open, align the new server node with the empty bay.
- Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 8 Power on the server node.

Step 9 Update the controller firmware using the Host Upgrade Utility. This will ensure that the firmware version on the card is compatible with the latest system firmware.

See the [Cisco Host Upgrade Utility User Guide For S3260 Storage Servers](#) for instructions on updating the firmware.

Replacing a Supercap Unit on a RAID Controller

The dual RAID controller UCS-S3260-DRAID uses two supercap units (RAID backup) that mount to the controller bracket.

Each supercap provides approximately 3 years of backup for the disk write-back cache DRAM in the case of sudden power loss by offloading the cache to the NAND flash.

The PID for the spare supercap is UCSC-SCAP-M5=.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).

Step 2 Remove the server node (with attached I/O expander, if present) from the system:

- Grasp the two ejector levers and pinch their latches to release the levers.

Replacing a Supercap Unit on a RAID Controller

- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.


Step 3 Do one of the following to access the component inside the server node:

- If the server node does not have an I/O expander attached—Remove the server node top cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.

Step 4 Remove a supercap unit:

- a) Press the securing clip on the bracket toward the center of the board and then lift the supercap from the bracket.
- b) Disconnect the supercap cable from the RAID controller cable.

Figure 26: Supercap Units on RAID Controller



1	Supercap cable connector	3	Supercap unit in clips on bracket (one each side)
2	Supercap release clip	-	

Step 5 Install a new supercap unit:

- a) Connect the new supercap cable to the RAID controller cable to which the old supercap was connected.
- b) Press the securing clip on the bracket toward the center of the board while you set the supercap in place. Release the securing clip.

Step 6 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.

- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 7 Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 8 Power on the server node.

Replacing an NVMe SSD in the Server Node

The optional NVMe SSD sled for the server node can hold up to two NVMe SSDs. The sled might be under a storage controller card, if one is installed in the server node.

To replace an NVMe SSD in the I/O expander, see [Replacing an NVMe SSD in the I/O Expander, on page 68](#).



Note At the current time, you can have NVMe SSDs in either the I/O expander or the server node, but not both.

All NVMe SSDs in the system (in the server nodes and/or the I/O expander) must be of the same partner brand. For example, two *Intel* NVMe SSDs in the server node and two *HGST* NVMe SSDs in the I/O expander is an invalid configuration because of driver incompatibility.



Note NVMe SSDs are bootable in UEFI mode; legacy booting is not supported.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).

Step 2 Remove the server node (with attached I/O expander, if present) from the system:

- a) Grasp the two ejector levers and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.

Step 3 Do one of the following to access the component inside the server node:

- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.

Step 4 Remove the NVMe SSD sled from the server node:

- a) Do one of the following:
 - If there is no storage controller card installed in the server node, continue with the next step.


- If there is a storage controller card installed in the server node, you must remove it to provide clearance to the sled. Remove the storage controller card as described in [Replacing a Storage Controller, on page 41](#) before you continue with the next step.

- Loosen the thumbscrew on the rear panel of the server node until the threads clear the sled.
- Loosen the thumbscrew on the top of the sled.
- Lift on the ribbons labeled "LIFT HERE" to disengage the two connectors on the underside of the sled from the sockets on the board.

Step 5 Remove the NVMe SSD (or a filler panel) from the sled:

- Remove the four screws that secure the SSD to the sled. There are two screws on each side of the SSD.
- Lift gently on the rear of the SSD and then pull it free from the connector in the sled.

Figure 27: NVMe SSD Sled in the Server Node



1	Thumbscrew on rear panel of server node	3	Sled removed from server, showing securing screws for upper NVMe SSD 2 (two on each side of sled)
2	Thumbscrew on sled	4	Sled removed from server, showing securing screws for lower NVMe SSD 1 (two on each side of sled)

- Step 6** Install a new NVMe SSD to the sled:
- Set the SSD in place, then gently push its connector into the connector on the sled. Ensure that the SSD is fully seated and that the screw holes on the sides of the sled align with the screw holes on the SSD.
 - Install the four securing screws.
- Step 7** Install the NVMe SSD sled to the server node:
- Gently set the sled in place so that the two connectors on the underside of the sled align with the two sockets on the server node board.
 - Press down on the ribbon labeled “PRESS HERE TO INSTALL” to fully seat the connectors in the sockets.
 - Tighten the thumbscrew on the top of the sled.
 - Tighten the thumbscrew on the rear panel of the server node.
- Step 8** If you removed a storage controller to provide clearance, reinstall it to the server node using the procedure in [Replacing a Storage Controller, on page 41](#).
- Step 9** Do one of the following:
- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.
- Step 10** Return the server node to the chassis:
- With the two ejector levers open, align the new server node with the empty bay.
 - Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.
- Step 11** Power on the server node.

Replacing the RTC Battery

The real-time clock (RTC) battery retains system settings when the server is disconnected from power. The battery type is CR2032. Cisco supports the industry-standard CR2032 battery, which can be purchased from Cisco or most electronic stores.



Note When the RTC battery is removed or it completely loses power, settings that were stored in the BMC of the server node are lost. You must reconfigure the BMC settings after installing a new battery.



Warning **Recyclers:** Do not shred the battery! Make sure you dispose of the battery according to appropriate regulations for your country or locale.

-
- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:

Replacing the RTC Battery

- a) Grasp the two ejector levers and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.


Step 3 Do one of the following to access the component inside the server node:

- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.

Step 4 Remove the server node RTC battery:

- a) Locate the RTC battery in its vertical socket.

Figure 28: RTC Battery Socket Location



1	RTC battery vertical socket	-	
---	-----------------------------	---	--

- b) Do one of the following:
 - If there is no storage controller card installed in the server node, continue with the next step.
 - If there is a storage controller card installed in the server node, you must remove it to provide clearance to the battery, which is under the storage controller card bracket. Remove the storage controller card as described in [Replacing a Storage Controller, on page 41](#) before you continue with the next step.
- c) Pull the battery retaining clip away from the battery and pull the battery from the socket.

Step 5 Install a new RTC battery:

- a) Pull the retaining clip away from the battery socket and insert the battery in the socket.

Note The flat, positive side of the battery marked “+” should face the retaining clip.

b) Push the battery into the socket until it is fully seated and the retaining clip clicks over the top of the battery.

Step 6 If you removed a storage controller to provide clearance, reinstall it to the server node using the procedure in [Replacing a Storage Controller, on page 41](#).

Step 7 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 8 Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 9 Power on the server node.

Step 10 Reconfigure the BMC settings for this node.

Installing a Trusted Platform Module (TPM)

The trusted platform module (TPM) is a small circuit board that plugs into a server board socket and is then permanently secured with a one-way screw.

TPM Considerations

- This server supports either TPM version 1.2 or TPM version 2.0. The TPM 2.0, UCSX-TPM2-002B(=), is compliant with Federal Information Processing (FIPS) Standard 140-2. FIPS support has existed, but FIPS 140-2 is now supported.
- Field *replacement* of a TPM is not supported; you can install a TPM after-factory only if the server does not already have a TPM installed.
- If there is an existing TPM 1.2 installed in the server, you cannot upgrade to TPM 2.0. If there is no existing TPM in the server, you can install TPM 2.0.
- If the TPM 2.0 becomes unresponsive, reboot the server.

Installing and Enabling a TPM



Note Field replacement of a TPM is not supported; you can install a TPM after-factory only if the server does not already have a TPM installed.

This topic contains the following procedures, which must be followed in this order when installing and enabling a TPM:

1. Installing the TPM Hardware

2. Enabling the TPM in the BIOS
3. Enabling the Intel TXT Feature in the BIOS

Installing TPM Hardware




Note For security purposes, the TPM is installed with a one-way screw. It cannot be removed with a standard screwdriver.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- a) Grasp the two ejector levers and pinch their latches to release the levers.
 - b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - c) Pull the server node straight out from the system.
- Step 3** Do one of the following to access the component inside the server node:
- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.
- Step 4** Install a TPM:
- a) Locate the TPM socket on the server board, as shown below.
 - b) Do one of the following:
 - If there is no storage controller card installed in the server node, continue with the next step.
 - If there is a storage controller card installed in the server node, you must remove it to provide clearance to the TPM socket, which is under the storage controller card bracket. Remove the storage controller card as described in [Replacing a Storage Controller, on page 41](#) before you continue with the next step.
 - c) Align the connector that is on the bottom of the TPM circuit board with the server board TPM socket. Align the screw hole on the TPM board with the screw hole that is adjacent to the TPM socket.
 - d) Push down evenly on the TPM to seat it in the socket.
 - e) Install the single one-way screw that secures the TPM to the motherboard.
- Step 5** If you removed a storage controller to provide clearance, reinstall it to the server node using the procedure in [Replacing a Storage Controller, on page 41](#).
- Step 6** Do one of the following:
- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

- Step 7** Return the server node to the chassis:
- With the two ejector levers open, align the new server node with the empty bay.
 - Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 8 Power on the server node.

Figure 29: Location of the TPM Socket



1	TPM socket location on server board	-	
---	-------------------------------------	---	--

Step 9 Continue with [Enabling the TPM in the BIOS, on page 51](#).

Enabling the TPM in the BIOS

After hardware installation, you must enable TPM support in the BIOS.



Note You must set a BIOS Administrator password before performing this procedure. To set this password, press the **F2** key when prompted during system boot to enter the BIOS Setup utility. Then navigate to **Advanced > Security > Set Administrator Password** and enter the new password twice as prompted.

Step 1 Enable TPM Support:

- Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
- Log in to the BIOS Setup Utility with your BIOS Administrator password.
- On the BIOS Setup Utility window, choose the **Advanced** tab.
- Choose **Trusted Computing** to open the TPM Security Device Configuration window.

- e) Change TPM SUPPORT to **Enabled**.
- f) Press **F10** to save your settings and reboot the server.

Step 2 Verify that TPM support is now enabled:

- a) Watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup.
- b) Log into the BIOS Setup utility with your BIOS Administrator password.
- c) Choose the **Advanced** tab.
- d) Choose **Trusted Computing** to open the TPM Security Device Configuration window.
- e) Verify that TPM SUPPORT and TPM State are Enabled.

Step 3 Continue with [Enabling the Intel TXT Feature in the BIOS, on page 52](#).

Enabling the Intel TXT Feature in the BIOS

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisibly to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code.

Step 1 Reboot the server and watch for the prompt to press F2.

Step 2 When prompted, press **F2** to enter the BIOS Setup utility.

Step 3 Verify that the prerequisite BIOS values are enabled:

- a) Choose the **Advanced** tab.
- b) Choose **Intel TXT(LT-SX) Configuration** to open the Intel TXT(LT-SX) Hardware Support window.
- c) Verify that the following items are listed as Enabled:
 - VT-d Support (default is Enabled)
 - VT Support (default is Enabled)
 - TPM Support
 - TPM State
- d) Do one of the following:
 - If VT-d Support and VT Support are already enabled, skip to step 4.
 - If VT-d Support and VT Support are not enabled, continue with the next steps to enable them.
- e) Press **Escape** to return to the BIOS Setup utility **Advanced** tab.
- f) On the Advanced tab, choose **Processor Configuration** to open the Processor Configuration window.
- g) Set Intel (R) VT and Intel (R) VT-d to **Enabled**.

Step 4 Enable the Intel Trusted Execution Technology (TXT) feature:

- a) Return to the Intel TXT(LT-SX) Hardware Support window if you are not already there.
- b) Set TXT Support to **Enabled**.

Step 5 Press **F10** to save your changes and exit the BIOS Setup utility.

Replacing an I/O Expander

The server node with optional I/O expander is accessed from the rear of the system, so you do not have to pull the system out from the rack.



Note You do not have to power off the chassis in this procedure. Replacement with the chassis powered on is supported if you shut down the server node before removal.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node](#), on page 7.

Step 2 Remove the server node with attached I/O expander from the system:


- Grasp the two ejector levers and pinch their latches to release the levers.
- Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- Pull the server node straight out from the system.

Step 3 Remove the I/O expander top cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover](#), on page 13.

Step 4 Remove the I/O expander from the server node:

- Remove the five screws that secure the I/O expander to the top of the server node.

Figure 30: I/O Expander Securing Screws (Five)




- Use two small flat-head screwdrivers (1/4-inch or equivalent) to help separate the connector on the underside of the I/O expander from the socket on the server node board.

Insert a screwdriver about ½ inch into the “REMOVAL SLOT” that is marked with an arrow on each side of the I/O expander. Then lift up evenly on both screwdrivers at the same time to separate the connectors and lift the I/O expander about ½ inch.

- c) Grasp the two handles on the I/O expander board and lift it straight up.

Figure 31: Separating the I/O Expander From the Server Node




1	Side view showing REMOVAL SLOT for screwdriver insertion (one on each side of the I/O expander)	2	Rear view of server node with I/O expander
---	---	---	--

Step 5 Install a new I/O expander to the server node:

Caution Use caution to align all features of the I/O expander with the intermediate cover and server node before mating the connector on the underside of the expander with the socket on the server board. The connector can be damaged if alignment is not correct.

- a) Carefully align the I/O expander with the alignment pegs on the top of the intermediate cover.
- b) Set the I/O expander down on the intermediate cover and lower it gently to mate the mezzanine connector and the socket on the server board.


Figure 32: Reinstalling the I/O Expander to the Server Node



1	Intermediate cover on server node	3	I/O expander mezzanine connector
2	Intermediate cover screws (two on each side of cover)	4	Alignment pegs on intermediate cover

- c) If an NVMe SSD is present in the right-hand socket (IOENVMe2) of the I/O expander, you must remove it to access the PRESS HERE plate in the next step.
See [Replacing an NVMe SSD in the I/O Expander, on page 68](#) to remove it, then return to the next step.
- d) Press down firmly on the plastic plate marked “PRESS HERE” to fully seat the connector to the server node board.

Figure 33: I/O Expander, Showing PRESS HERE Plate



e) If you removed an NVMe SSD to access the PRESS HERE plate, reinstall it now.

See [Replacing an NVMe SSD in the I/O Expander, on page 68](#).


Caution Before you reinstall the I/O expander securing screws, you must use the alignment tool (UCSC-C3K-M4IOTOOL) in the next step to ensure alignment of the connectors that plug into the internal chassis backplane. Failure to ensure alignment might damage the sockets on the backplane.

Step 6

Insert the four pegs of the alignment tool into the holes that are built into the connector side of the server node and I/O expander. Ensure that the alignment tool fits into all four holes and lies flat.

The alignment tool is shipped with systems that are ordered with an I/O expander. It is also shipped with I/O expander replacement spares. You can order the tool using Cisco PID UCSC-C3K-M4IOTOOL.

Figure 34: Using the I/O Expander Alignment Tool



1	Alignment Tool	3	Alignment tool installed
2	Connector side of server node and I/O expander	-	

- Step 7** Reinstall and tighten the five screws that secure the I/O expander to the top of the server node.
- Step 8** Remove the alignment tool.
- Step 9** Reinstall the top cover to the I/O expander:
- Set the cover in place on the server node or I/O expander (if present), offset about one inch toward the rear. Pegs on the inside of the cover must engage the tracks on the server node or I/O expander base.
 - Push the cover forward until it stops.
 - Turn the latch handle 90 degrees to close the lock.
 - Fold the latch handle flat.
- Step 10** Reinstall the server node to the chassis:
- With the two ejector levers open, align the server node with the empty bay.
 - Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.
- Step 11** Power on the server node.
- Step 12** Update the IOE firmware using the Host Upgrade Utility. This will ensure that the firmware version on the IOE is compatible with the latest system firmware.

See the [Cisco Host Upgrade Utility User Guide For S3260 Storage Servers](#) for instructions on updating the firmware.

Adding an I/O Expander After-Factory



Note This procedure is for adding an I/O expander to an S3260 M5 server node. If you are replacing an existing I/O expander, see [Replacing an I/O Expander, on page 53](#).

Tools Required

This procedure requires:


- I/O expander alignment tool UCSC-C3K-M4IOTOOL. This tool is included with the I/O expander spare.
- I/O expander kit UCS-S3260-IOLID. This kit includes:
 - One intermediate cover for the server node, including 4 cover screws
 - One threaded support post, with screw

Procedure

When an I/O expander is installed, the server node must occupy lower server bay 1. The I/O expander occupies upper server bay 2.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove any server node (or disk expansion tray) from upper server bay 2 and set it on an antistatic work surface.
- Step 3** Remove any server node (or disk expansion tray) from lower server bay 1 and set it on an antistatic work surface.
- Step 4** Remove the top cover from the server node to which you will install an I/O expander:
- a) Lift the latch handle to an upright position, and then turn the latch handle 90 degrees to release the lock.
 - b) Slide the cover toward the rear (toward the rear-panel buttons) and then lift it from the server node.

Figure 35: Top View, Server Node Top Cover




1	Top cover on server node	2	Cover latch handle (shown in closed, flat position)
----------	--------------------------	----------	---

Step 5

If there is a storage controller card installed in the server node, remove it to provide clearance in the next step. (If there is no storage controller card, skip to the next step.)

- a) Loosen the captive thumbscrews that secure the controller card to the board.
- b) Grasp the card at both ends and lift it evenly to disengage the connector on the underside of the card from the mezzanine socket.

Figure 36: Screw Locations on Storage Controller (Six)




Step 6

Move the cover latch assembly from the controller-card bracket to a bracket in the IO expander. You will install the top cover to the I/O expander at the end of this procedure and so it requires the cover latch.

Remove the two Phillips-head screws that secure the cover latch assembly to the controller card bracket in the server node. See the following figure.

- If your I/O expander has a storage controller card, use the two Phillips-head screws to install the cover latch assembly to the bracket of that controller card.
- If your I/O expander does not have a storage controller card, use the two Phillips-head screws to install the cover latch assembly to the bracket on the I/O expander board. See the following figure.

Figure 37: Cover Latch Assembly Screws in Server Node and I/O Expander




1	Two securing screws on cover latch assembly in server node (shown installed on controller card bracket)	2	Two securing screws on cover latch assembly in I/O expander (shown installed to bracket on I/O expander board)
---	---	---	--

Step 7 Remove one Phillips-head screw from the server node board. See the following figure for the screw location.

Step 8 Install the threaded metal support post from the kit, using one screw.

Set the post against the edge of the server board where you removed the screw in the prior step. The flange with the screw hole must sit flat on top of the board. See the following figure.

Figure 38: Support Post and Screw, S3260 M5 Server Node



1	Phillips-head screw location (remove this screw)	3	Support post screw
2	Side view of support post, showing installation position on edge of server board	-	

- Step 9** If you removed a storage controller card, install it back to the server node now:
- a) Align the card over the mezzanine socket and the standoffs.
 - b) Press down on both ends of the card to engage the connector on the underside of the card with the mezzanine socket.
 - c) Tighten the captive screws that secure the card to the board.


Step 10 Install the intermediate cover from the kit to the server node. Set the intermediate cover in place and then install its four securing screws (two on each side).

Step 11 Install the I/O expander to the server node:

Caution Use caution to align all features of the I/O expander with the server node before mating the connector on the underside of the expander with the socket on the server board. The connector can be damaged if correct alignment is not used.

- a) Carefully align the I/O expander with the alignment pegs on the top of the intermediate cover.
- b) Set the I/O expander down on the server node intermediate cover and push down gently to mate the connectors.


Figure 39: Reinstalling the I/O Expander to the Server Node



1	Intermediate cover on server node	3	I/O expander mezzanine connector
2	Intermediate cover screws (two on each side of cover)	4	Alignment pegs on intermediate cover

- c) If an NVMe SSD is present in the right-hand socket (IOENVMe2) of the I/O expander, you must remove it to access the PRESS HERE plate in the next step.
See [Replacing an NVMe SSD in the I/O Expander, on page 68](#) to remove it, then return to the next step.
- d) Press down firmly on the plastic plate marked “PRESS HERE” to fully seat the connectors.

Figure 40: I/O Expander, Showing PRESS HERE Plate



- e) If you removed an NVMe SSD to access the PRESS HERE plate, reinstall the SSD now.


Caution Before you install the I/O expander securing screws, you must use the supplied alignment tool (UCSC-C3K-M4IOTOOL) in the next step to ensure alignment of the connectors that connect to the internal chassis backplane. Failure to ensure alignment might damage the sockets on the backplane.

Step 12

Insert the four pegs of the alignment tool into the holes that are built into the connector side of the server node and I/O expander. Ensure that the alignment tool fits into all four holes and lies flat.

The alignment tool is shipped with I/O expander spares. You can also order the tool using Cisco PID UCSC-C3K-M4IOTOOL.


Figure 41: Using the I/O Expander Alignment Tool



1	Alignment Tool	3	Alignment tool installed
2	Connector side of server node and I/O expander	-	

Step 13 Install and tighten the five screws that secure the I/O expander to the top of the server node. The screw at the center of the board-edge screws into the support post that you installed earlier in this procedure.

Figure 42: I/O Expander Securing Screws (Five)



Step 14 Remove the alignment tool.

Step 15 Install the top cover that you removed from the server node to the I/O expander.

- a) Set the cover in place on the I/O expander, offset about one inch toward the rear. Pegs on the inside of the cover must set into the tracks on the I/O expander base.
- b) Push the cover forward until it stops.
- c) Turn the latch handle 90 degrees to close the lock and then fold the latch handle flat.

Step 16 Install the server node with attached I/O expander to the chassis:

- a) With the two ejector levers open, align the server node and I/O expander with the two empty bays.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 17 Power on the server node.

Step 18 Update the IOE firmware using the Host Upgrade Utility. This will ensure that the firmware version on the IOE is compatible with the latest system firmware.

See the [Cisco Host Upgrade Utility User Guide For S3260 Storage Servers](#) for instructions on updating the firmware.

Replacing a PCIe Card in the I/O Expander

The optional I/O expander has two horizontal PCIe sockets.

PCIe Slot Specifications

The following table describes the specifications for the PCIe slots.


Table 6: PCIe Expansion Slots in the I/O Expander

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support
1 (IOESlot1)	Gen-3 x8	x16 connector	$\frac{3}{4}$ length	$\frac{1}{2}$ height	No
2 (IOESlot2)	Gen-3 x8	x16 connector	$\frac{3}{4}$ length	$\frac{1}{2}$ height	No

Replacing a PCIe Card

-
- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 3** Remove the I/O expander top cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#).
- Step 4** Remove an existing PCIe card (or a filler panel if no card is present):
- Release the card-tab retainer. On the inside of the I/O expander, pull the spring-loaded plunger on the card-tab retainer inward and then rotate the card-tab retainer 90 degrees to the open position.
 - Slide the PCIe card horizontally to free its edge connector from the socket, and then lift the card out from the I/O expander.

Figure 43: PCIe Card Sockets in the I/O Expander



1	PCIe card sockets	3	Card-tab retainer (outside the I/O expander)
2	Spring-loaded plunger on card-tab retainer (inside the I/O expander)		

Step 5

Install a new PCIe card:

- a) With the card-tab retainer in the open position, set the card in the I/O expander and align its edge connector with the socket.
- b) Slide the card horizontally to fully engage the edge connector with the socket. The card’s tab should sit flat against the rear-panel opening.
- c) Close the card-tab retainer. Rotate the retainer 90 degrees until it clicks and locks.

Step 6

Replace the I/O expander top cover.

- a) Set the cover in place on the I/O expander, offset about one inch toward the rear. Pegs on the inside of the cover must engage the tracks on the I/O expander base.
- b) Push the cover forward until it stops.
- c) Turn the latch handle 90 degrees to close the lock.
- d) Fold the latch handle flat.
- e) Reinstall the four screws that secure the top cover to the I/O expander.

Step 7

Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 8 Power on the server node.

Replacing an NVMe SSD in the I/O Expander

The I/O expander has two sockets for NVMe SSDs. The software designations of the NVMe PCIe SSD sockets inside the I/O expander are IOENVMe1 and IOENVMe2.



Note At the current time, you can have NVMe SSDs in either the I/O expander or the server node, but not both. All NVMe SSDs in the system (in the server nodes and/or the I/O expander) must be of the same partner brand. For example, two *Intel* NVMe SSDs in the server node and two *HGST* NVMe SSDs in the I/O expander is an invalid configuration because of driver incompatibility.



Note NVMe SSDs are bootable in UEFI mode; legacy booting is not supported.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).

Step 2 Remove the server node (with attached I/O expander, if present) from the system:


- Grasp the two ejector levers and pinch their latches to release the levers.
- Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- Pull the server node straight out from the system.

Step 3 Remove the I/O expander top cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#).

Step 4 Remove an NVMe SSD:

- Remove the single screw that secures the drive to its bracket.
- Slide the drive horizontally to disengage it from its socket, then lift it from the I/O expander.

Figure 44: NVMe SSDs in the I/O Expander



1	Single securing screw (one each SSD)	2	NVMe PCIe SSD sockets IOENVM1 and IOENVM2
---	--------------------------------------	---	---

- Step 5** Install a new NVMe SSD:
- Set the drive in its bracket, then slide it forward to engage its connector with the socket.
 - Install the single screw that secures the drive to the bracket.
- Step 6** Reinstall the I/O expander top cover.
- Step 7** Reinstall the server node to the chassis:
- With the two ejector levers open, align the server node with the empty bay.
 - Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
 - Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.
- Step 8** Power on the server node.

Service Headers on the Server Node Board

The server node board includes headers that you can jumper for certain service functions.


Service Header Locations

There are two 2-pin service headers on the server node board that are supported for use.

- Header P10 = Password reset

- Header P11 = Clear CMOS

Figure 45: Service Headers on the Server Node Board <PLACEHOLDER, NEED NEW ILLU>



Using the Password Reset Header P10

You can use a jumper on header P10 to clear the administrator password.

- Step 1** Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).
- Step 2** Remove the server node (with attached I/O expander, if present) from the system:
- Grasp the two ejector levers and pinch their latches to release the levers.
 - Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
 - Pull the server node straight out from the system.
- Step 3** Do one of the following to access the component inside the server node:
- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
 - If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.
- Step 4** Locate header P10 and install a jumper across pins 1 and 2.
- Step 5** Return the server node to the chassis (do not reinstall an I/O expander to the server node for this reboot):
- With the two ejector levers open, align the new server node with the empty bay.

- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 6 Power on the server node.

Step 7 After the server node has booted fully, shut it down again.

Step 8 Remove the server node from the system.

Step 9 Remove the jumper from the header pins.

Note If you do not remove the jumper, the Cisco IMC resets the password each time that you boot the server node.

Step 10 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 11 Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 12 Power on the server node.

Using the Clear CMOS Header P11

You can use a jumper on header P11 to clear the CMOS settings.

Step 1 Shut down the server node by using the software interface or by pressing the node power button, as described in [Shutting Down an S3260 M5 Server Node, on page 7](#).

Step 2 Remove the server node (with attached I/O expander, if present) from the system:

- a) Grasp the two ejector levers and pinch their latches to release the levers.
- b) Rotate both levers to the outside at the same time to evenly disengage the server node from its midplane connectors.
- c) Pull the server node straight out from the system.

Step 3 Do one of the following to access the component inside the server node:

- If the server node does not have an I/O expander attached—Remove the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node has an I/O expander attached—Remove the I/O expander and intermediate cover as described in [Disassembling the I/O Expander Assembly, on page 15](#) before you continue with the next step.

Step 4 Locate header P11 and install a jumper across pins 1 and 2.

Step 5 Return the server node to the chassis (do not reinstall an I/O expander to the server node for this reboot):

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.

- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 6 Power on the server node.

Step 7 After the server node has booted fully, shut it down again.

Step 8 Remove the server node from the system.

Step 9 Remove the jumper from the header pins.

Note If you do not remove the jumper, the Cisco IMC clears the settings each time that you boot the server node.

Step 10 Do one of the following:

- If the server node did not have an I/O expander attached—Reinstall the server node cover as described in [Removing an S3260 M5 Server Node or I/O Expander Top Cover, on page 13](#) before continuing with the next step.
- If the server node had an I/O expander attached—Reinstall the I/O expander and intermediate cover as described in [Reassembling the I/O Expander Assembly, on page 17](#) before you continue with the next step.

Step 11 Return the server node to the chassis:

- a) With the two ejector levers open, align the new server node with the empty bay.
- b) Push the server node into the bay until it engages with the midplane connectors and is flush with the chassis.
- c) Rotate both ejector levers toward the center until they lay flat and their latches lock into the rear of the server node.

Step 12 Power on the server node.



APPENDIX **A**

Storage Controller Information

This appendix contains information about the supported RAID and HBA storage controllers, including the embedded software RAID controller.

- [Supported Storage Controllers, on page 73](#)
- [Cisco UCS S3260 Dual RAID Controller Information, on page 74](#)
- [Cisco UCS S3260 Dual Pass-Through Controller Information, on page 74](#)
- [Best Practices For Configuring RAID Controllers, on page 74](#)
- [Embedded Software RAID, on page 76](#)
- [For More RAID Utility Information, on page 84](#)

Supported Storage Controllers

Each S3260 M5 server node supports a single RAID or HBA controller card.

Each server node also includes an embedded MegaRAID controller that can control two rear-panel SATA solid state drives (SSDs) in a RAID 0 or 1 configuration. This embedded software RAID is available only when the server node has the UCS S3260 12G Dual Pass-Through Controller (UCS-S3260-DHBA) installed.

Table 7: S3260 M5 Storage Controller Options

Controller	Maximum Drives Controlled	RAID Levels	Optional Supercap Backup?
Embedded software RAID (PCH SATA) (Requires that UCS-S3260-DHBA is installed in the server node.)	2 rear-panel SATA SSDs	0, 1, 10	No
Cisco UCS S3260 Dual RAID Controller UCS-S3260-DRAID	66	0, 1, 5, 6, 10, 50, 60	Yes
Cisco UCS S3260 Dual Pass-Through Controller UCS-S3260-DHBA	66	Not applicable	No

Cisco UCS S3260 Dual RAID Controller Information

The Cisco UCS S3260 Dual RAID Controller (UCS-S3260-DRAID) is based on the Broadcom 3316 SAS/SATA, 16-port RAID-on-chip (RoC). This controller is divided into two boards: a base board and a power board. The power board supplies power to the base board.

This RAID controller has the following features:

- Provides a sixteen-port 12Gb/s SAS and 6Gb/s SATA interface.
- Provides a 40-bit or a 72-bit 1866-MHz DDR3/3L SDRAM/MRAM interface with a hardware RAID assist-engine for parity calculations.
- Provides a full-featured hardware-based RAID solution that supports RAID levels 0, 1, 5, 6, 10, 50, and 60.
- Increases system performance and provides fault-tolerant data storage.
- Supports data striping across multiple disks, which reduces disk access time because multiple disks simultaneously read or write data.
- Backs up data with either data mirroring or a parity block. You can use either backup method to recover lost data in the event of a disk failure.
- Maximum 50 W power consumption.

Cisco UCS S3260 Dual Pass-Through Controller Information

The Cisco UCS S3260 Dual Pass-Through Controller (UCS-S3260-DHBA) is based on the Broadcom 3316 SAS/SATA, 16-port RAID-on-chip (RoC). This controller is divided into two boards: a base board and a power board. The power board supplies power to the base board.

This pass-through controller has the following features:

- Dual Broadcom SAS3316-based subsystems.
- Dual 8x SAS-3 lanes (12G, 6G, 3G, and SATA), one for each subsystem.
- Dual 8x PCIe Gen-3 lanes via the server-board mezzanine connector, one for each subsystem.
- Maximum 50 W power consumption.

Best Practices For Configuring RAID Controllers

This section contains recommendations for best practices when configuring RAID controllers.

4K Sector Format Drives

Do not configure 4K sector format and 512-byte sector format drives as part of the same RAID volume.

Storage Controller Card Firmware Compatibility

Firmware on the storage controller (RAID or HBA) must be verified for compatibility with the current Cisco IMC and BIOS versions that are installed on the server. If not compatible, upgrade or downgrade the storage controller firmware using the Host Upgrade Utility (HUU) for your firmware release to bring it to a compatible level.

See the HUU guide for your Cisco IMC release for instructions on downloading and using the utility to bring server components to compatible levels: [HUU Guides](#).

Choosing Between RAID 0 and JBOD

The RAID controller supports JBOD mode (non-RAID) on physical drives that are in pass-through mode and directly exposed to the OS. We recommend that you use JBOD mode instead of individual RAID 0 volumes when possible.

RAID 5/RAID 6 Volume Creation

The RAID controller allows you to create a large RAID 5 or 6 volume by including all the drives in the system in a spanned array configuration (RAID 50/RAID 60). Where possible, we recommend you create multiple, smaller RAID 5 or 6 volumes with fewer drives per RAID array. This provides redundancy and reduces the operations time for initialization, RAID rebuilds, and other operations.

Choosing I/O Policy

The I/O policy applies to reads on a specific virtual drive. It does not affect the read-ahead cache. RAID volumes can be configured in two types of I/O policies. These are:

- **Cached I/O**—In this mode, all reads are buffered in cache memory. Cached I/O provides faster processing.
- **Direct I/O**—In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. Direct I/O makes sure that the cache and the host contain the same data.

Although Cached I/O provides faster processing, it is useful only when the RAID volume has a small number of slower drives. With the S3260 4-TB SAS drives, Cached I/O has not shown any significant advantage over Direct I/O. Instead, Direct I/O has shown better results over Cached I/O in a majority of I/O patterns. We recommend you use Direct I/O (the default) in all cases.

Background Operations (BGOPs)

The RAID controller conducts different background operations like Consistency Check (CC), Background Initialization (BGI), Rebuild (RBLD), Volume Expansion & Reconstruction (RLM), and Patrol Real (PR)

While these BGOPs are expected to limit their impact to I/O operations, there have been cases of higher impact during some of the I/O operations such as Format. In these cases, both the I/O operation and the BGOPs may take more time to complete. In such cases, we recommend you limit concurrent BGOPs and other intensive I/O operations where possible.

Embedded Software RAID

Each server node includes an embedded MegaRAID controller that can control two rear-panel SATA solid state drives (SSDs) in a RAID 0 or 1 configuration. This embedded software RAID is available only when the server node has the UCS S3260 12G Dual Pass-Through Controller (UCS-S3260-DHBA) installed. When this HBA is installed, the two rear-panel SSDs can be controlled through software RAID mode or AHCI mode, when selected in the server BIOS.



Note Embedded software RAID is not available when the HW RAID Cisco UCS S3260 Dual RAID Controller (UCS-S3260-DRAID) is installed. In that case, the rear-panel SSDs are controlled by hardware RAID.



Note The VMware ESX/ESXi operating system is not supported with the embedded SATA MegaRAID controller in SW RAID mode. You can use VMWare in AHCI mode.



Note The Microsoft Windows Server 2016 Hyper-V hypervisor is supported for use with the embedded MegaRAID controller in SW RAID mode, but all other hypervisors are not supported. All Hypervisors are supported in AHCI mode.



Note The embedded RAID controller in server node 1 can control the upper two rear-panel SSDs; the embedded RAID controller in server node 2 can control the lower two rear-panel SSDs.

Accessing the Software RAID Configuration Utility

To configure RAID settings for the embedded SATA RAID controllers, use the utility that is built into the BIOS.

-
- Step 1** Boot the server and press **F2** when prompted to enter the BIOS Setup utility.
 - Step 2** Choose the **Advanced** tab.
 - Step 3** Select the utility: **LSI Software RAID Configuration Utility (SATA)**
-

Enabling SATA Mode For the Embedded Controller

This procedure uses the server's BIOS Setup Utility

-
- Step 1** Set the SATA mode:

- a) Boot the server and press **F2** when prompted to enter the BIOS Setup utility.
- b) Choose the **Advanced** tab, and then choose **LOM and PCIe Slots Configuration**.
- c) Choose one of the options from the dialog:
 - LSI SW RAID—Enable the embedded SATA RAID controller for control of rear-panel SATA SSDs.
Note This menu option does not appear when the server is set to boot in Legacy mode (UEFI mode is required). To change the boot mode, use the BIOS setting for **Boot Options > Boot Mode**.
 - AHCI—Enable control of the rear-panel SSDs by AHCI through your OS rather than the embedded RAID controller.
 - Disabled—Disable the embedded RAID controller.

Step 2 Press **F10** to save your changes and exit the utility.

Installing LSI MegaSR Drivers For Windows and Linux



Note The required drivers for this controller are already installed and ready to use. However, if you will use this controller with Windows or Linux, you must download and install additional drivers for those operating systems.

This section explains how to install the LSI MegaSR drivers for the following supported operating systems:

- Microsoft Windows Server
- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

Downloading the MegaSR Drivers

The MegaSR drivers are included in the driver ISO for your server and OS.

- Step 1** Find the drivers ISO file download for your server online and download it to a temporary location on your workstation:
- a) See the following URL: <http://www.cisco.com/cisco/software/navigator.html>.
 - b) Type the name of your server in the **Select a Product** search field and then press **Enter**.
 - c) Click **Unified Computing System (UCS) Drivers**.
 - d) Click the release number that you are downloading.
 - e) Click the Download icon to download the drivers ISO file.

Step 2 Continue through the subsequent screens to accept the license agreement and then browse to a location where you want to save the driver ISO file.

Microsoft Windows Server Drivers

Installing Microsoft Windows Server Drivers

The Windows Server operating system automatically adds the driver to the registry and copies the driver to the appropriate directory.

Before you begin

Before you install this driver on an embedded controller, you must configure a RAID drive group on the embedded controller for the drives where you will install the OS.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility for the embedded controller: **LSI Software RAID Configuration Utility (SATA)**

-
- Step 1** Download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 77](#).
- Step 2** Prepare the drivers on a USB thumb drive:
- Burn the ISO image to a disk.
 - Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
/OS>/Storage/Intel/C600/
 - Expand the Zip file, which contains the folder with the MegaSR driver files.
 - Copy the expanded folder to a USB thumb drive.
- Step 3** Start the Windows driver installation using one of the following methods:
- To install from local media, connect an external USB DVD drive to the server (if the server does not have a DVD drive installed) and then insert the first Windows installation disk into the DVD drive. Skip to Step 6.
 - To install from remote ISO, log in to the server's Cisco IMC interface and continue with the next step.
- Step 4** Launch a Virtual KVM console window and click the **Virtual Media** tab.
- Click **Add Image** and browse to select your remote Windows installation ISO file.
 - Check the check box in the **Mapped** column for the media that you just added, and then wait for mapping to complete.
- Step 5** Power cycle the server.
- Step 6** Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.
- Step 7** On the Boot Manager window, choose the physical disk or virtual DVD and press **Enter**. The Windows installation begins when the image is booted.
- Step 8** Press **Enter** when you see the prompt, "Press any key to boot from CD."
- Step 9** Observe the Windows installation process and respond to prompts in the wizard as required for your preferences and company standards.
- Step 10** When Windows prompts you with "Where do you want to install Windows," install the drivers for embedded MegaRAID:
- Click **Load Driver**. You are prompted by a Load Driver dialog box to select the driver to be installed.
 - Connect the USB thumb drive that you prepared in Step 3 to the target server.
 - On the Windows Load Driver dialog, click **Browse**.
 - Use the dialog box to browse to the location of the drivers folder on the USB thumb drive, and then click **OK**.
- Windows loads the drivers from the folder and when finished, the driver is listed under the prompt, "Select the driver to be installed."

- e) Click **Next** to install the drivers.

Updating Microsoft Windows Server Drivers

- Step 1** Click **Start**, point to **Settings**, and then click **Control Panel**.
- Step 2** Double-click **System**, click the **Hardware** tab, and then click **Device Manager**. Device Manager starts.
- Step 3** In Device Manager, double-click **SCSI and RAID Controllers**, right-click the device for which you are installing the driver, and then click **Properties**.
- Step 4** On the Driver tab, click **Update Driver** to open the Update Device Driver wizard, and then follow the wizard instructions to update the driver.
-

Linux Drivers

Downloading the Driver Image File

See [Downloading the MegaSR Drivers, on page 77](#) for instructions on downloading the drivers. The Linux driver is included in the form of `dud-[driver version].img`, which is the boot image for the embedded MegaRAID stack.



Note The LSI MegaSR drivers that Cisco provides for RHEL and SLES are for the original GA versions of those distributions. The drivers do not support updates to those OS kernels.

Preparing Physical Thumb Drive for Linux

This topic describes how to prepare physical Linux thumb drive from the driver image files.

This procedure requires a CD or DVD drive that you can use to burn the ISO image to disk; and a USB thumb drive.

Alternatively, you can mount the `dud.img` file as a virtual floppy disk, as described in the installation procedures.

For RHEL and SLES, you can use a driver disk utility to create disk images from image files.

- Step 1** Download the drivers ISO, as described in [Downloading the MegaSR Drivers, on page 77](#) and save it to your Linux system.
- Step 2** Extract the `dud.img` or `dd.iso` driver file:
- Note** For RHEL 7.1 and later, there is no `dud.img` file--the driver is contained in a `dd.iso` file.
- a) Burn the Cisco UCS C-Series Drivers ISO image to a disc.
 - b) Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
`<OS>/Storage/Intel/C600-M5/`
 - c) Expand the Zip file, which contains the folder with the driver files.
- Step 3** Copy the driver update disk image `dud-[driver version].img` (or `dd.iso`) to your Linux system.
- Step 4** Insert a blank USB thumb drive into a port on your Linux system.

Step 5 Create a directory and mount the dud.img or dd.iso image to that directory:

Example:

```
mkdir <destination_folder>
mount -o loop <driver_image> <destination_folder>
```

Step 6 Copy the contents in the directory to your USB thumb drive.

Installing the Red Hat Enterprise Linux Driver

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

This topic describes the fresh installation of the RHEL device driver on systems that have the embedded MegaRAID stack.



Note If you use an embedded RAID controller with Linux, both the pSATA and the sSATA controller must be set to `LSI SW RAID` mode.

Before you begin

Before you install this driver on an embedded controller, you must configure a RAID drive group on the embedded controller.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility for the embedded controller: **LSI Software RAID Configuration Utility (SATA)**

Step 1 Prepare the dud.img (or .iso) file using one of the following methods:

Note For RHEL 7.1 and later, there is no dud.img file--the driver is contained in a dd.iso file.

- To install from physical drive, use the procedure in [Preparing Physical Thumb Drive for Linux, on page 79](#), then continue with step 4.
- To install from *virtual* disk, download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 77](#), then continue with the next step.

Step 2 Extract the dud.img (or dd.iso) file:

- a) Burn the Cisco UCS C-Series Drivers ISO image to a disk.
- b) Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:
/`<OS>/Storage/Intel/C600-M5/`
- c) Copy the dud-`<driver version>`.img (or dd.iso) file to a temporary location on your workstation.
- d) If you are using RHEL 7.x, rename the saved dd.iso to dd.img.

Note If you are using RHEL 7.x, renaming the dd.iso file to dd.img simplifies this procedure and saves time. The Cisco UCS virtual drive mapper can map only one .iso at a time, and only as a virtual CD/DVD. Renaming the file to dd.img allows you to mount the RHEL installation ISO as a virtual CD/DVD and the renamed dd.img as a virtual floppy disk or removable disk at the same time. This avoids the steps of unmounting and remounting the RHEL ISO when the dd.iso driver file is prompted for.

- Step 3** Start the Linux driver installation using one of the following methods:
- To install from local media, connect an external USB CD/DVD drive to the server and then insert the first RHEL installation disk into the drive. Then continue with Step 5.
 - To install from virtual disk, log in to the server's Cisco IMC interface. Then continue with the next step.
- Step 4** Launch a Virtual KVM console window and click the **Virtual Media** tab.
- a) Click **Add Image** and browse to select your remote RHEL installation ISO image.

Note An .iso file can be mapped only as a virtual CD/DVD.
 - b) Click **Add Image** again and browse to select your RHEL 6.x `dud.img` or the RHEL 7.x `dd.img` file that you renamed in step 2.

Note Map the .img file as a virtual floppy disk or virtual removable disk.
 - c) Check the check boxes in the **Mapped** column for the media that you just added, then wait for mapping to complete.
- Step 5** Power-cycle the target server.
- Step 6** Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.
- Note** Do not press Enter in the next step to start the installation. Instead, press **e** to edit installation parameters.
- Step 7** On the Boot Menu window, use the arrow keys to select **Install Red Hat Enterprise Linux** and then press **e** to edit installation parameters.
- Step 8** Append one of the following blacklist commands to the end of the line that begins with **linuxefi**:
- For RHEL 6.x (32- and 64-bit), type:


```
linux dd blacklist=iscsi blacklist=ahci nodmraid noprobe=<atadriver number>
```

Note The noprobe values depend on the number of drives. For example, to install RHEL 6.x on a RAID 5 configuration with three drives, type:

```
Linux dd blacklist=iscsi blacklist=ahci nodmraid noprobe=ata1 noprobe=ata2
```
 - For RHEL 7.x (32- and 64-bit), type:


```
linux dd modprobe.blacklist=ahci nodmraid
```
- Step 9** **Optional:** To see full, verbose installation status steps during installation, delete the **Quiet** parameter from the line.
- Step 10** On the Boot Menu window, press **Ctrl+x** to start the interactive installation.
- Step 11** Below **Driver disk device selection**, select the option to install your driver .img file. (Type **r** to refresh the list if it is not populated.)
- Note** The installer recognizes the driver file as an .iso file, even though you renamed it to dd.img for mapping.
- Type the number of the driver device ISO in the list. Do *not* select the RHEL ISO image. In the following example, type **6** to select device sdb:
- ```
5) sr0 iso9660 RHEL-7.6\x20Server.x
6) sdb iso9660 CDROM
to select, 'r' - refresh, or 'c' -continue: 6
```

The installer reads the driver file and lists the drivers.

**Step 12** Under **Select drivers to install**, type the number of the line that lists the megasr driver. In the following example, type **1**:

```
1) [] /media/DD-1/rpms/x86_61/kmod-megasr-18.01.2010.1107_e17.6-1.x86_61.rpm
```

```
to toggle selection, or 'c' -continue: 1
```

Your selection is displayed with an X in brackets.

```
1) [X] /media/DD-1/rpms/x86_61/kmod-megasr-18.01.2010.1107_e17.6-1.x86_61.rpm
```

**Step 13** Type **c** to continue.

**Step 14** Follow the RHEL installation wizard to complete the installation.

**Step 15** When the wizard's Installation Destination screen is displayed, ensure that **LSI MegaSR** is listed as the selection. If it is not listed, the driver did not load successfully. In that case, select **Rescan Disc**.

**Step 16** After the installation completes, reboot the target server.

## Installing the SUSE Linux Enterprise Server Driver

For the specific supported OS versions, see the [Hardware and Software Compatibility Matrix](#) for your server release.

This topic describes the fresh installation of the SLES driver on systems that have the embedded MegaRAID stack.



**Note** If you use an embedded RAID controller with Linux, both the pSATA and the sSATA controller must be set to `LSI SW RAID` mode.

### Before you begin

Before you install this driver on an embedded controller, you must configure a RAID drive group on the embedded controller.

To access the configuration utility, open the BIOS Setup Utility, go to the **Advanced** tab, and then choose the utility for the embedded controller: **LSI Software RAID Configuration Utility (SATA)**

**Step 1** Prepare the `dud.img` file using one of the following methods:

- To install from physical disk, use the procedure in [Preparing Physical Thumb Drive for Linux, on page 79](#), then continue with step 4.
- To install from *virtual* disk, download the Cisco UCS C-Series drivers' ISO, as described in [Downloading the MegaSR Drivers, on page 77](#), then continue with the next step.

**Step 2** Extract the `dud.img` file that contains the driver:

- a) Burn the ISO image to a disk.
- b) Browse the contents of the drivers folders to the location of the embedded MegaRAID drivers:  
/OS/Storage/Intel/C600-M5/...

- c) Within the SLES folder for your version, the `dud-<driver version>.img` file is packaged in a compressed `.gz` file. Extract the `.img` file from the `.gz` file.
- d) Copy the `dud-<driver version>.img` file to a temporary location on your workstation.

**Step 3** Start the Linux driver installation using one of the following methods:

- To install from local media, connect an external USB DVD drive to the server and then insert the first SLES installation disk into the drive. Then continue with Step 5.
- To install from remote ISO, log in to the server's Cisco IMC interface. Then continue with the next step.

**Step 4** Launch a Virtual KVM console window and click the **Virtual Media** tab.

- a) Click **Add Image** and browse to select your remote SLES installation ISO file.
- b) Click **Add Image** again and browse to select your `dud-<driver version>.img` file.
- c) Check the check boxes in the **Mapped** column for the media that you just added, then wait for mapping to complete.

**Step 5** Power-cycle the target server.

**Step 6** Press **F6** when you see the F6 prompt during bootup. The Boot Menu window opens.

**Step 7** On the Boot Manager window, select the physical or virtual SLES installation ISO and press **Enter**.

The SLES installation begins when the image is booted.

**Step 8** When the first SLES screen appears, select **Installation**.

**Step 9** Press **e** to edit installation parameters.

**Step 10** Append the following parameter to the end of the line that begins with **linuxefi**:

```
brokenmodules=ahci
```

**Step 11** **Optional:** To see detailed status information during the installation, add the following parameter to the line that begins with **linuxefi**:

```
splash=verbose
```

**Step 12** Press **Ctrl+x** to start the installation.

The installation proceeds. The installer finds the LSI driver automatically in the `dud-<driver version>.img` file that you provided. With verbose status messages, you see the driver being installed when `LSI MegaRAID SW RAID Module` is listed.

**Step 13** Follow the SLES installation wizard to complete the installation. Verify installation of the driver when you reach the **Suggested Partitioning** screen:

- a) On the **Suggested Partitioning** screen, select **Expert Partitioner**.
- b) Navigate to **Linux > Hard disks** and verify that there is a device listed for the `LSI - LSI MegaSR` driver. The device might be listed as a type other than `sda`. For example:

```
dev/sdd: LSI - LSI MegaSR
```

If no device is listed, the driver did not install properly. In that case, repeat the steps above.

**Step 14** When installation is complete, reboot the target server.

## For More RAID Utility Information

The Broadcom utilities have help documentation for more information about using the utilities.

- For basic information about RAID and for using the utilities for the RAID controller cards that are supported in Cisco servers, see the [Cisco UCS Servers RAID Guide](#).
- For hardware SAS MegaRAID configuration—[Broadcom 12Gb/s MegaRAID SAS Software User Guide, Version 2.8](#).
- For embedded software MegaRAID and the utility that is accessed via the server BIOS—[Broadcom Embedded MegaRAID Software User Guide, March 2018](#).



## APPENDIX **B**

### Related Information

---

The installation guide for each server node generation is separate from the installation guide for the Cisco UCS S3260 chassis.

- [Cisco UCS S3260 Storage Server Installation and Service Guide](#)
- [Regulatory Compliance and Safety Information For Cisco UCS S-Series Hardware](#)

