



## Server-Related Policies

---

- [BIOS Settings, on page 1](#)
- [Trusted Platform Module, on page 96](#)
- [SPDM Security, on page 98](#)
- [Creating a SPDM Security Policy, on page 99](#)
- [Associating the Security Policy with a Server, on page 100](#)
- [Viewing the Fault Alert Settings, on page 100](#)
- [Consistent Device Naming, on page 101](#)
- [CIMC Security Policies, on page 104](#)
- [Graphics Card Policies, on page 107](#)
- [Local Disk Policies, on page 108](#)
- [Persistent Memory Modules, on page 120](#)
- [Scrub Policy, on page 121](#)
- [DIMM Error Management, on page 125](#)
- [Serial over LAN Policy Settings, on page 127](#)
- [Server Autoconfiguration Policies, on page 128](#)
- [Server Discovery Policy Settings, on page 130](#)
- [Server Inheritance Policy Settings, on page 133](#)
- [Server Pool Policy Settings, on page 134](#)
- [Server Pool Policy Qualifications Settings, on page 136](#)
- [vNIC/vHBA Placement Policy Settings, on page 141](#)
- [CIMC Mounted vMedia, on page 154](#)

## BIOS Settings

### Server BIOS Settings

#### Server BIOS Settings

Cisco UCS provides two methods for making global modifications to the BIOS settings on servers in an Cisco UCS domain. You can create one or more BIOS policies that include a specific grouping of BIOS settings that match the needs of a server or set of servers, or you can use the default BIOS settings for a specific server platform.

Both the BIOS policy and the default BIOS settings for a server platform enable you to fine tune the BIOS settings for a server managed by Cisco UCS Manager.

Depending upon the needs of the data center, you can configure BIOS policies for some service profiles and use the BIOS defaults in other service profiles in the same Cisco UCS domain, or you can use only one of them. You can also use Cisco UCS Manager to view the actual BIOS settings on a server and determine whether they are meeting current needs.



**Note** Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

## Main BIOS Settings

The following table lists the main server BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Properties</b>	
<b>Reboot on BIOS Settings Change</b>	<p>When the server is rebooted after you change one or more BIOS settings.</p> <p>If you enable this setting, the server is rebooted according to the maintenance policy in the server's service profile. For example, if the maintenance policy requires user acknowledgment, the server is not rebooted and the BIOS changes are not applied until a user acknowledges the pending activity.</p> <p>If you do not enable this setting, the BIOS changes are not applied until the next time the server is rebooted, whether as a result of another server configuration change or a manual reboot.</p>
<b>BIOS Setting</b>	
<b>Quiet Boot</b>	<p>What the BIOS displays during Power On Self-Test (POST). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS displays all messages and Option ROM information during boot.</li> <li>• <b>Enabled</b>—The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>POST error pause</b>	<p>What happens when the server encounters a critical error during POST. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS continues to attempt to boot the server.</li> <li>• <b>Enabled</b>—The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Resume on AC power loss</b>	<p>How the server behaves when power is restored after an unexpected power loss. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Stay Off</b>—The server remains off until manually powered on.</li> <li>• <b>Last State</b>—The server is powered on and the system attempts to restore its last state.</li> <li>• <b>Reset</b>—The server is powered on and automatically reset.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Front panel lockout</b>	<p>Whether the power and reset buttons on the front panel are ignored by the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The power and reset buttons on the front panel are active and can be used to affect the server.</li> <li>• <b>Enabled</b>—The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>CDN Control</b>	<p>Consistent Device Naming allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more uniform, easy to identify, and persistent when adapter or other configuration changes are made.</p> <p>Whether consistent device naming is enabled or not. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Consistent device naming is disabled for the BIOS policy.</li> <li>• <b>Enabled</b>—Consistent device naming is enabled for the BIOS policy. This enables Ethernet interfaces to be named consistently. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slots CDN Control</b>	<p>PCIe Slots Consistent Device Naming (CDN) control allows PCIe slots to be named in a consistent manner. This makes PCIe slot names more uniform, easy to identify, and persistent when the configuration changes are made. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Consistent device naming is disabled. This is the default option.</li> <li>• <b>Enabled</b>—Consistent device naming is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### Processor BIOS Settings

The following table lists the processor BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Intel Turbo Boost Tech</b>	<p>Whether the processor uses Intel Turbo Boost Technology, which allows the processor to automatically increase its frequency if it is running below power, temperature, or voltage specifications. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not increase its frequency automatically.</li> <li>• <b>Enabled</b>—The processor uses Turbo Boost Technology if required.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Enhanced Intel SpeedStep Tech</b>	<p>Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor never dynamically adjusts its voltage or frequency.</li> <li>• <b>Enabled</b>—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
<b>Intel HyperThreading Tech</b>	<p>Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not permit hyperthreading.</li> <li>• <b>Enabled</b>—The processor allows for the parallel execution of multiple threads.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>

Name	Description
<b>Intel Speed Select</b>	<p>Allows improved CPU performance by using Intel Speed Select technology to tune the CPU to run at one of three operating profiles, based on number of logical processor cores, frequency, and TDP thread setting, to improve performance over the basic Platform Default setting. These profiles correspond to High, Medium, and Low Core settings and can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Base</b>—The processor uses Base.</li> <li>• <b>Config 1</b>—The processor uses Config 1.</li> <li>• <b>Config 2</b>—The processor uses Config 2.</li> <li>• <b>Config 3</b>—The processor uses Config 3.</li> <li>• <b>Config 4</b>—The processor uses Config 4.</li> </ul> <p><b>Note</b>           The values <b>Config 1</b> and <b>Config 2</b> are not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Core Multi Processing</b>	<p>Sets the state of logical processor cores per CPU in a package. If you disable this setting, Intel Hyper Threading technology is also disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>All</b>—Enables multiprocessing on all logical processor cores.</li> <li>• <b>1 through n</b>—Specifies the number of logical processor cores per CPU that can run on the server. To disable multiprocessing and have only one logical processor core per CPU running on the server, choose 1.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>

Name	Description
<b>Execute Disable Bit</b>	<p>Classifies memory areas on the server to specify where the application code can execute. As a result of this classification, the processor disables code execution if a malicious worm attempts to insert code in the buffer. This setting helps to prevent damage, worm propagation, and certain classes of malicious buffer overflow attacks. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not classify memory areas.</li> <li>• <b>Enabled</b>—The processor classifies memory areas.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
<b>Intel Virtualization Technology</b>	<p>Whether the processor uses Intel Virtualization Technology, which allows a platform to run multiple operating systems and applications in independent partitions. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not permit virtualization.</li> <li>• <b>Enabled</b>—The processor allows multiple operating systems in independent partitions.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> If you change this option, you must power cycle the server before the setting takes effect.</p>
<b>Hardware Prefetcher</b>	<p>Whether the processor allows the Intel hardware prefetcher to fetch streams of data and instruction from memory into the unified second-level cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The hardware prefetcher is not used.</li> <li>• <b>Enabled</b>—The processor uses the hardware prefetcher when cache issues are detected.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> <b>CPU Performance</b> must be set to <b>Custom</b> in order to specify this value. For any value other than <b>Custom</b>, this option is overridden by the setting in the selected CPU performance profile.</p>

Name	Description
<b>Adjacent Cache Line Prefetcher</b>	<p>Whether the processor fetches cache lines in even/odd pairs instead of fetching just the required line. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor only fetches the required line.</li> <li>• <b>Enabled</b>—The processor fetches both the required line and its paired line.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b>      <b>CPU Performance</b> must be set to <b>Custom</b> in order to specify this value. For any value other than <b>Custom</b>, this option is overridden by the setting in the selected CPU performance profile.</p>
<b>DCU Streamer Prefetch</b>	<p>Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not try to anticipate cache read requirements and only fetches explicitly requested lines.</li> <li>• <b>Enabled</b>—The DCU prefetcher analyzes the cache read pattern and prefetches the next line in the cache if it determines that it may be needed.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>DCU IP Prefetcher</b>	<p>Whether the processor uses the DCU IP Prefetch mechanism to analyze historical cache access patterns and preload the most relevant lines in the L1 cache. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not preload any cache data.</li> <li>• <b>Enabled</b>—The DCU IP prefetcher preloads the L1 cache with the data it determines to be the most relevant.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



Name	Description
<b>KTI Prefetch</b>	<p>KTI prefetch is a mechanism to get the memory read started early on a DDR bus. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not preload any cache data.</li> <li>• <b>Enabled</b>—The KTI prefetcher preloads the L1 cache with the data it determines to be the most relevant.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>LLC Prefetch</b>	<p>Whether the processor uses the LLC Prefetch mechanism to fetch the date into the LLC. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not preload any cache data.</li> <li>• <b>Enabled</b>—The LLC prefetcher preloads the L1 cache with the data it determines to be the most relevant.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>XPT Prefetch</b>	<p>Whether XPT prefetch is used to enable a read request sent to the last level cache to issue a copy of that request to the memory controller prefetcher. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The CPU does not use the XPT Prefetch option.</li> <li>• <b>Enabled</b>—The CPU enables the XPT prefetcher option.</li> <li>• <b>Auto</b>—The CPU auto enables the XPT prefetcher option.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Direct Cache Access</b>	<p>Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how to place data from I/O devices into the processor cache.</li> <li>• <b>Disabled</b>—Data from I/O devices is not placed directly into the processor cache.</li> <li>• <b>Enabled</b>—Data from I/O devices is placed directly into the processor cache.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Processor C State</b>	<p>Whether the system can enter a power savings mode during idle periods. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The system remains in a high-performance state even when idle.</li> <li>• <b>Enabled</b>—The system can reduce power to system components such as the DIMMs and CPUs.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
<b>Processor C1E</b>	<p>Allows the processor to transition to its minimum frequency upon entering C1. This setting does not take effect until after you have rebooted the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The CPU continues to run at its maximum frequency in the C1 state.</li> <li>• <b>Enabled</b>—The CPU transitions to its minimum frequency. This option saves the maximum amount of power in the C1 state.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Processor C3 Report</b>	<p>Whether the processor sends the C3 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The processor sends the C3 report to the OS.</li> <li>• <b>Disabled</b>—The processor does not send the C3 report.</li> <li>• <b>ACPI C2</b>—The processor sends the C3 report using the advanced configuration and power interface (ACPI) C2 format.</li> <li>• <b>ACPI C3</b>—The processor sends the C3 report using the ACPI C3 format.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>On the Cisco UCS B440 Server, the BIOS Setup menu uses enabled and disabled for these options. If you specify acpi-c2 or acpi-c2, the server sets the BIOS value for that option to enabled.</p>

Name	Description
<b>Processor C6 Report</b>	<p>Whether the processor sends the C6 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not send the C6 report.</li> <li>• <b>Enabled</b>—The processor sends the C6 report.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Processor C7 Report</b>	<p>Whether the processor sends the C7 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>C7</b>—The processor sends the report using the C7 format.</li> <li>• <b>C7s</b>—The processor sends the report using the C7s format.</li> <li>• <b>Disabled</b>—The processor does not send the C7 report.</li> <li>• <b>Enabled</b>—The processor sends the C7 report.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Processor CMCI</b>	<p>Enables CMCI generation. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor disables CMCI.</li> <li>• <b>Enabled</b>—The processor enables CMCI.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>CPU Performance</b>	<p>Sets the CPU performance profile for the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Custom</b>—All performance profile options can be configured from the BIOS setup on the server. In addition, the Hardware Prefetcher and Adjacent Cache-Line Prefetch options can be configured as well.</li> <li>• <b>Enterprise</b>—For M4 and higher servers, all prefetchers and data reuse are enabled.</li> <li>• <b>High Throughput</b>—Data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled.</li> <li>• <b>HPC</b>—All prefetchers are enabled and data reuse is disabled. This setting is also known as high-performance computing.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Max Variable MTRR Setting</b>	<p>Allows you to select the number of mean time to repair (MTRR) variables. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto Max</b>—BIOS uses the default value for the processor.</li> <li>• <b>8</b>—BIOS uses the number specified for the variable MTRR.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Local X2 APIC</b>	<p>Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Processor disables Local X2 APIC.</li> <li>• <b>Enabled</b>—Processor enables Local X2 APIC.</li> <li>• <b>XAPIC</b>—Uses the standard xAPIC architecture.</li> <li>• <b>X2APIC</b>—Uses the enhanced x2APIC architecture to support 32 bit addressability of processors.</li> <li>• <b>AUTO</b>—Automatically uses the xAPIC architecture that is detected.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Power Technology</b>	<p>Enables you to configure the CPU power management settings for the following options:</p> <ul style="list-style-type: none"><li>• Enhanced Intel Speedstep Technology</li><li>• Intel Turbo Boost Technology</li><li>• Processor Power State C6</li></ul> <p>Power Technology can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>Disabled</b>—The server does not perform any CPU power management and any settings for the BIOS parameters mentioned above are ignored.</li><li>• <b>Energy Efficient</b>—The server determines the best settings for the BIOS parameters mentioned above and ignores the individual settings for these parameters.</li><li>• <b>Performance</b>—The server automatically optimizes the performance for the BIOS parameters mentioned above.</li><li>• <b>Custom</b>—The server uses the individual settings for the BIOS parameters mentioned above. You must select this option if you want to change any of these BIOS parameters.</li><li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li></ul>

Name	Description
<b>Energy Performance</b>	<p>Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Performance</b> — The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power.</li> <li>• <b>Balanced Performance</b> — The server provides all server components with enough power to keep a balance between performance and power.</li> <li>• <b>Balanced Energy</b> — The server provides all server components with enough power to keep a balance between performance and power.</li> <li>• <b>Energy Efficient</b> — The server provides all server components with less power to keep reduce power consumption.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b>      <b>Power Technology</b> must be set to <b>Custom</b> or the server ignores the setting for this parameter.</p>
<b>Frequency Floor Override</b>	<p>Whether the CPU is allowed to drop below the maximum non-turbo frequency when idle. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>— The CPU can drop below the maximum non-turbo frequency when idle. This option decreases power consumption but may reduce system performance.</li> <li>• <b>Enabled</b>— The CPU cannot drop below the maximum non-turbo frequency when idle. This option improves system performance but may increase power consumption.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>P STATE Coordination</b>	<p>Allows you to define how BIOS communicates the P-state support model to the operating system. There are 3 models as defined by the Advanced Configuration and Power Interface (ACPI) specification.</p> <ul style="list-style-type: none"> <li>• <b>HW ALL</b>—The processor hardware is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package).</li> <li>• <b>SW ALL</b>—The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a physical package), and must initiate the transition on all of the logical processors.</li> <li>• <b>SW ANY</b>—The OS Power Manager (OSPM) is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package), and may initiate the transition on any of the logical processors in the domain.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b>      <b>Power Technology</b> must be set to <b>Custom</b> or the server ignores the setting for this parameter.</p>
<b>DRAM Clock Throttling</b>	<p>Allows you to tune the system settings between the memory bandwidth and power consumption. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b> — CPU determines the DRAM Clock Throttling settings.</li> <li>• <b>Balanced</b>— DRAM clock throttling is reduced, providing a balance between performance and power.</li> <li>• <b>Performance</b>—DRAM clock throttling is disabled, providing increased memory bandwidth at the cost of additional power.</li> <li>• <b>Energy Efficient</b>—DRAM clock throttling is increased to improve energy efficiency.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>External SSC enable</b>	<p>This option allows you to Enable/Disable the Clock Spread Spectrum of the external clock generators.</p> <p>For Cisco B-Series M5 and M6 servers and S-Series M5 servers, this option is Disabled by default. For Cisco C-Series rack servers, it is enabled by default.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>— Clock Spread Spectrum support is not available.</li> <li>• <b>Enabled</b>— Clock Spread Spectrum support is always available.</li> <li>• <b>Platform Default</b> — The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Channel Interleaving</b>	<p>Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines what interleaving is done.</li> <li>• <b>1 Way</b>—</li> <li>• <b>2 Way</b></li> <li>• <b>3 Way</b></li> <li>• <b>4-way</b>—The maximum amount of channel interleaving is used.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Rank Interleaving</b>	<p>Whether the CPU interleaves physical ranks of memory so that one rank can be accessed while another is being refreshed. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines what interleaving is done.</li> <li>• <b>1 Way</b>—</li> <li>• <b>2 Way</b></li> <li>• <b>4-way</b></li> <li>• <b>8 Way</b>—The maximum amount of rank interleaving is used.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



Name	Description
<b>Sub NUMA Clustering</b>	<p>Whether the CPU supports sub NUMA clustering, in which the tag directory and the memory channel are always in the same region. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>— The BIOS determines what Sub NUMA clustering is done.</li> <li>• <b>Disabled</b>— Sub NUMA clustering does not occur. This is the default option.</li> <li>• <b>Enabled</b>— Sub NUMA clustering occurs.</li> <li>• <b>Platform Default</b> — The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IMC Interleaving</b>	<p>This BIOS option controls the interleaving between the Integrated Memory Controllers (IMCs).</p> <ul style="list-style-type: none"> <li>• <b>1-way Interleave</b>—There is no interleaving.</li> <li>• <b>2-way Interleave</b>—Addresses are interleaved between the two IMCs.</li> <li>• <b>Auto</b> —CPU determines the IMC Interleaving mode.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Memory Interleaving</b>	<p>Whether the CPU interleaves the physical memory so that the memory can be accessed while another is being refreshed. This controls fabric level memory interleaving. Channel, die and socket have requirements based on memory populations and will be ignored if the memory does not support the selected option. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>None</b></li> <li>• <b>Channel</b></li> <li>• <b>Die</b></li> <li>• <b>Socket</b></li> <li>• <b>Auto</b>—This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Demand Scrub</b>	<p>Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>— Single bit memory errors are not corrected.</li> <li>• <b>Enabled</b>— Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Patrol Scrub</b>	<p>Whether the system actively searches for, and corrects, single bit memory errors even in unused portions of the memory on the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The system checks for memory ECC errors only when the CPU reads or writes a memory address.</li> <li>• <b>Enabled</b>—The system periodically reads and writes memory searching for ECC errors. If any errors are found, the system attempts to fix them. This option may correct single bit errors before they become multi-bit errors, but it may adversely affect performance when the patrol scrub is running.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>DCPMM Firmware Downgrade</b>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Support is disabled.</li> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Configurable TDP Control</b>	<p>Allows you to set customized value for Thermal Design Power (TDP). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>— Uses the rated TDP value of the processor.</li> <li>• <b>Manual</b>—Allows you to customize the TDP value.</li> </ul>

Name	Description
<b>Altitude</b>	<p>The approximate number of meters above sea level at which the physical server is installed. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines the physical elevation.</li> <li>• <b>300 M</b>—The server is approximately 300 meters above sea level.</li> <li>• <b>900 M</b>—The server is approximately 900 meters above sea level.</li> <li>• <b>1500 M</b>—The server is approximately 1500 meters above sea level.</li> <li>• <b>3000 M</b>—The server is approximately 3000 meters above sea level.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Package C State</b>	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <p><b>Note</b> Cisco recommends setting <b>Package C State Limit</b> as <b>No Limit</b> or <b>Platform Default</b> for C220 M4, C240 M4, C460 M4, and S3260 M4 servers.</p> <p>If you are changing the <b>Package C State Limit</b> token from any other value to <b>No Limit</b>, then ensure that the <b>Power Technology</b> is set to <b>Custom</b>.</p>
<b>CPU Hardware Power Management</b>	<p>Enables processor Hardware Power Management (HWPM). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> <li>• <b>Disabled</b>—HWPM is disabled.</li> <li>• <b>HWPM Native Mode</b>—HWPM native mode is enabled.</li> <li>• <b>HWPM OOB Mode</b>—HWPM Out-Of-Box mode is enabled.</li> <li>• <b>Native Mode with no Legacy</b> (only GUI)</li> </ul>

Name	Description
<b>Energy Performance Tuning</b>	<p>Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS.</p> <ul style="list-style-type: none"> <li>• <b>BIOS</b>—</li> <li>• <b>OS</b>—</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Workload Configuration</b>	<p>This feature allows for workload optimization. The options are Balanced and I/O Sensitive:</p> <ul style="list-style-type: none"> <li>• <b>Balanced</b></li> <li>• <b>IO Sensitive</b>—This is the default option.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>Cisco recommends using Balanced.</p>
<b>Core Performance Boost</b>	<p>Whether the AMD processor increases its frequency on some cores when it is idle or not being used much. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically determines how to boost performance.</li> <li>• <b>Disabled</b>—Core performance boost is disabled.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Uncore Frequency Scaling</b>	<p>Allows you configure the scaling of the uncore frequency of the processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Uncore frequency of the processor scales up or down based on the load. (Default.)</li> <li>• <b>Disabled</b>—Uncore frequency of the processor remains fixed.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>Refer to the Intel Dear Customer Letter (DCL) to know the fixed higher and lower values for Uncore Frequency Scaling.</p>

Name	Description
<b>Configurable TDP Level</b>	<p>Allows adjustments in processor thermal design power (TDP) values. By modifying the processor behavior and the performance levels, power consumption of a processor can be configured and TDP can be adjusted at the same time. Hence, a processor operates at higher or lower performance levels, depending on the available cooling capacities and desired power consumption.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Normal</b>—The CPU operates at its normal performance level. (Default.)</li> <li>• <b>Level 1</b></li> <li>• <b>Level 2</b></li> </ul> <p><b>Note</b> Refer to the Intel Dear Customer Letter (DCL) for the values for TDP level.</p>
<b>UPI Link Speed</b>	<p>Allows you to configure the Intel Ultra Path Interconnect (UPI) link speed between multiple sockets. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Automatically configures the optimal link speed. (Default)</li> <li>• <b>9.6GT/s (gigatransfers per second)</b>—Configures the optimal link speed at 9.6GT/s</li> <li>• <b>10.4GT/s</b>—Configures the optimal link speed at 10.4GT/s</li> <li>• <b>11.2GT/s</b>—Configures the optimal link speed at 11.2GT/s</li> <li>• <b>Use Per Link Setting</b></li> </ul> <p><b>Note</b> The value <b>Use Per Link Setting</b> is not supported on UCS M6 servers.</p>
<b>Global C-state Control</b>	<p>Whether the AMD processors control IO-based C-state generation and DF C-states. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically determines how to control IO-based C-state generation.</li> <li>• <b>Disabled</b>—Global C-state control is disabled.</li> <li>• <b>Enabled</b>—Global C-state control is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>L1 Stream HW Prefetcher</b>	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L1 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how to place data from I/O devices into the processor cache.</li> <li>• <b>Disabled</b>—The hardware prefetcher is not used.</li> <li>• <b>Enabled</b>—The processor uses the hardware prefetcher when cache issues are detected.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>L2 Stream HW Prefetcher</b>	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L2 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how to place data from I/O devices into the processor cache.</li> <li>• <b>Disabled</b>—The hardware prefetcher is not used.</li> <li>• <b>Enabled</b>—The processor uses the hardware prefetcher when cache issues are detected.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>AMD Memory Interleaving Size</b>	<p>Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8,9,10 or 11). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• 1 KB</li> <li>• 2 KB</li> <li>• 256 Bytes</li> <li>• 512 Bytes</li> <li>• <b>Auto</b>—The CPU determines the size of the memory block.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Chipselect Interleaving</b>	<p>Whether memory blocks across the DRAM chip selects for node 0 are interleaved. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically determines how to interleave chip selects.</li> <li>• <b>Disabled</b>—Chip selects are not interleaved within the memory controller.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Bank Group Swap</b>	<p>Determines how physical addresses are assigned to applications. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically determines how to assign physical addresses to applications.</li> <li>• <b>Disabled</b>—Bank group swap is not used.</li> <li>• <b>Enabled</b>—Bank group swap is used to improve the performance of applications.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Determinism Slider</b>	<p>Allows AMD processors to determine how to operate. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically uses default power determinism settings.</li> <li>• <b>Performance</b>—Processor operates at the best performance in a consistent manner.</li> <li>• <b>Power</b>—Processor operates at the maximum allowable performance on a per die basis.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>IOMMU</b>	<p>Input Output Memory Management Unit (IOMMU) allows AMD processors to map virtual addresses to physical addresses. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how map these addresses.</li> <li>• <b>Disabled</b>—IOMMU is not used.</li> <li>• <b>Enabled</b>—Address mapping takes place through the IOMMU.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SVM Mode</b>	<p>Whether the processor uses AMD Secure Virtual Machine Technology. This can be one of the following: This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not use SVM Technology.</li> <li>• <b>Enabled</b>—The processor uses SVM Technology. This is the default option.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SMEE</b>	<p>Whether the processor uses the Secure Memory Encryption Enable (SMEE) function, which provides memory encryption support. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—This is the default option.</li> <li>• <b>Disabled</b>—The processor does not use the SMEE function.</li> <li>• <b>Enabled</b>—The processor uses the SMEE function.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>UPI Prefetch</b>	<p>UPI prefetch is a mechanism to get the memory read started early on a DDR bus. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The UPI prefetcher preloads the L1 cache with the data it determines to be the most relevant.</li> <li>• <b>Disabled</b>—The processor does not preload any cache data.</li> <li>• <b>Auto</b>—The processor enables the UPI prefetcher option.</li> </ul>



Name	Description
<b>SGX Auto MP Registration Agent</b>	<p>Allows you to enable the registration authority service to store the platform keys. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>SProcessor Epoch <i>n</i></b>	<p>Allows you to define the SGX EPOCH owner value for the EPOCH number designated by <i>n</i>.</p>
<b>SGX Factory Reset</b>	<p>Allows the system to perform SGX factory reset on subsequent boot. This deletes all registration data. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>SGX PBUKEY HASH<sub><i>n</i></sub></b>	<p>Allows you to set the Software Guard Extensions (SGX) value. This value can be set between:</p> <ul style="list-style-type: none"> <li>• SGX PUBKEY HASH0—Between 7-0</li> <li>• SGX PUBKEY HASH1—Between 15-8</li> <li>• SGX PUBKEY HASH2—Between 23-16</li> <li>• SGX PUBKEY HASH3—Between 31-24</li> </ul>
<b>SGX Write Enable</b>	<p>Allows you to enable SGX Write feature. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>SGX Pkg info In-Band Access</b>	<p>Allows you to enable SGX Package Info In-Band Access. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>SGX QoS</b>	<p>Allows you to enable SGX QoS. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>— Support is enabled.</li> <li>• <b>Disabled</b>— Support is disabled.</li> </ul>

Name	Description
<b>Intel Dynamic Speed Select</b>	<p>Intel Dynamic Speed Select modes allow you to run the CPU with different speed and cores in auto mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Intel Dynamic Speed Select is enabled.</li> <li>• <b>Disabled</b>—Intel Dynamic Speed Select is disabled.</li> </ul>
<b>IIO eDPC Support</b>	<p>eDPC allows a downstream link to be disabled after an uncorrectable error, making recovery possible in a controlled and robust manner. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—eDPC support is disabled.</li> <li>• <b>On Fatal Errors</b>—eDPC is enabled only for fatal errors.</li> <li>• <b>On Fatal and Non-Fatal Errors</b>—eDPC is enabled for both fatal and non-fatal errors.</li> </ul>
<b>Multikey Total Memory Encryption (MK-TME)</b>	<p>MK-TME allows you to have multiple encryption domains with one with own key. Different memory pages can be encrypted with different keys. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled. This is the default option.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>SW Guard Extensions (SGX)</b>	<p>Allows you to enable Software Guard Extensions (SGX) feature. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>Total Memory Encryption (TME)</b>	<p>Allows you to provide the capability to encrypt the entirety of the physical memory of a system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled. This is the default option.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>Select Owner EPOCH input type</b>	<p>Allows you to change the seed for the security key used for the locked memory region that is created. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>SGX Owner EPOCH activated</b>— Does not change the current input type.</li> <li>• <b>Change to New Random Owner EPOCHs</b>—Changes EPOCH to a system generated random number</li> <li>• <b>Manual User Defined Owner EPOCHs</b>—Changes the EPOCH seed to a hexadecimal value that you enter.</li> </ul>

Name	Description
<b>Enhanced CPU Performance</b>	<p>Enhances CPU performance by adjusting server settings automatically. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not run with this functionality. This is the default option.</li> <li>• <b>Auto</b>—Allows to adjust server settings to increase the processor performance.</li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• Enabling this functionality may increase power consumption.</li> <li>• The server should meet the following requirements in order to use this functionality: <ul style="list-style-type: none"> <li>• The server should not contain Barlow Pass DIMMs.</li> <li>• DIMM module size present in the Cisco UCS C220 M6 server should be less than 64GB and in Cisco UCS C240 M6 server should be less than 256GB.</li> <li>• No GPU cards are present in the server.</li> </ul> </li> </ul>
<b>UPI Link Enablement</b>	<p>Enables the number of Ultra Path Interconnect (UPI) links required by the processor. This can be one of the following</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—This is the default option.</li> <li>• <b>1</b></li> <li>• <b>2</b></li> </ul>
<b>UPI Power Management</b>	<p>The UPI power management can be used for conserving power on the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Enables the processor to support this functionality.</li> <li>• <b>Disabled</b>—Disables the processor to support this functionality. This is the default option.</li> </ul>
<b>C1 Auto UnDemotion</b>	<p>Select whether to enable processors to automatically undemote from C1. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>. This is the default option.</li> <li>• <b>Enabled</b>—Enables the processor to support this functionality.</li> <li>• <b>Disabled</b>—Disables the processor to support this functionality.</li> </ul>

Name	Description
<b>C1 Auto Demotion</b>	<p>If enabled, CPU automatically demotes to C1 based on un-core auto-demote information. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b> . This is the default option.</li> <li>• <b>Enabled</b>—Enables the processor to support this functionality.</li> <li>• <b>Disabled</b>—Disables the processor to support this functionality.</li> </ul>
<b>CPU Downcore control 7xx3</b>	<p>Provides the ability to remove one or more cores from operation is supported in the silicon. It may be desirable to reduce the number of cores due to OS restrictions, or power reduction requirements of the system. This item allows the control on the number of cores that are running. This setting can only reduce the number of cores from only those available in the processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how many cores need to be enabled. This is the default option</li> <li>• <b>ONE (1+0)</b>—One core enabled on one CPU complex</li> <li>• <b>Two (2+0)</b>—Two core enabled on one CPU complex</li> <li>• <b>Three (3+0)</b>—Three core enabled on one CPU complex.</li> <li>• <b>Four (4+0)</b>—Four core enabled on one CPU complex.</li> <li>• <b>Five (5+0)</b>—Five core enabled on one CPU complex</li> <li>• <b>Six (6+0)</b>—Six core enabled on one CPU complex</li> <li>• <b>Seven (7+0)</b>—Seven core enabled on one CPU complex</li> </ul> <p><b>Note</b> This token is applicable only for the servers with 7xx3 Model processors.</p>
<b>Fixed SOC P-State</b>	<p>This option defines the target P-state when APBDIS (to disable Algorithm Performance Boost (APB)) is set. The <b>P-x</b> specify a valid P-state for the processor installed. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Sets a valid P-state suitable for the processor. This is the default option.</li> <li>• <b>P0</b>—Highest-performing SOC P-state</li> <li>• <b>P1</b>—Next-highest-performing SOC P-state</li> <li>• <b>P2</b>—Next-highest-performing SOC P-state</li> <li>• <b>P3</b>—Minimum SOC power P-state</li> </ul>

Name	Description
<b>APBDIS</b>	<p>Allows you to select the Algorithm Performance Boost (APB) Disable value for the SMU. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Sets an auto ApbDis for the SMU. This is the default option.</li> <li>• <b>0</b>—Clear ApbDis to SMU</li> <li>• <b>1</b>—Set ApbDis to SMU</li> </ul>
<b>CCD Control</b>	<p>Allows you to specify the number of charge-coupled device CCDs that are desired to be enable in the system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The maximum CCDs provided by the processor is enabled. This is the default option.</li> <li>• <b>2 CCDS</b></li> <li>• <b>3 CCDS</b></li> <li>• <b>4 CCDS</b></li> <li>• <b>6 CCDS</b></li> </ul>
<b>Cisco xGMI Max Speed</b>	<p>This option enables 18 Gbps XGMI link speed. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The feature is disabled. This is the default option.</li> <li>• <b>Enabled</b>—The feature is enabled.</li> </ul>
<b>ACPI SRAT L3 Cache As NUMA Domain</b>	<p>Creates a layer of virtual domains on top of the physical domains in which each CCX is declared to be in its on domain. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Set to auto mode. This is the default option.</li> <li>• <b>Disabled</b>—Use NPS settings for domain configuration.</li> <li>• <b>Enabled</b>—Each CCX is declared to be in its own domain.</li> </ul>
<b>Streaming Stores Control</b>	<p>Enables the streaming stores functionality. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Set to auto mode. This is the default option.</li> <li>• <b>Disabled</b>—Feature is disabled.</li> <li>• <b>Enabled</b>—Feature is enabled.</li> </ul>

Name	Description
<b>DF C-States</b>	<p>When long duration idleness is expected in a system, this control allows the system to transition into a DF Cstate which can set the system into an even lower power state. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Set to auto mode. This is the default option.</li> <li>• <b>Disabled</b>—This option is turned off, long period of idleness are not expected so no power savings would be achieved.</li> <li>• <b>Enabled</b>—This option is active, saving power when the system is idle.</li> </ul>
<b>SEV-SNP Support</b>	<p>Allows you to enable Secure Nested Paging feature. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not use the SEV-SNP function. This is the default option.</li> <li>• <b>Enabled</b>—The processor uses the SEV-SNP function.</li> </ul>
<b>Efficiency Mode Enable</b>	<p>Allows you to configure power consumption based on efficiency. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically uses default settings. This is the default option.</li> <li>• <b>Enabled</b>—Efficiency mode is enabled.</li> </ul>
<b>SNP Memory Coverage</b>	<p>Allows you to configure SNP memory coverage. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—System decides the memory coverage. This is the default option.</li> <li>• <b>Disabled</b>—The processor does not use this function.</li> <li>• <b>Enabled</b>—This feature is enabled.</li> <li>• <b>Custom</b>—Custom size can be defined in SNP Memory Size to Cover.</li> </ul>
<b>SNP Memory Size to Cover in MB</b>	<p>Allows you to configure SNP memory size. The value can range from 0-1048576. 0 is the default option.</p>

Name	Description
<b>SMT Mode</b>	<p>Whether the processor uses AMD Simultaneous MultiThreading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The processor allows for the parallel execution of multiple threads.</li> <li>• <b>Enabled</b>—The processor allows permit multithreading. This is the default option.</li> <li>• <b>Disabled</b>—The processor allows permit multithreading.</li> </ul>
<b>CPCC</b>	<p>Allows you to configure Collaborative Processor Performance Control. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU automatically uses default CPCC settings. This is the default option.</li> <li>• <b>Disabled</b>—Feature is disabled.</li> <li>• <b>Enabled</b>—Collaborative Processor Performance is enabled.</li> </ul>
<b>Downcore control 7xx2</b>	<p>The ability to remove one or more cores from operation is supported in the silicon. It may be desirable to reduce the number of cores due to OS restrictions, or power reduction requirements of the system. This item allows the control of how many cores are running. This setting can only reduce the number of cores from those available in the processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines how many cores need to be enabled. This is the default option.</li> <li>• <b>Two (1+1)</b>—Two cores enabled on one CPU complex.</li> <li>• <b>Four (2+2)</b>—Four cores enabled on one CPU complex.</li> <li>• <b>Six (3+3)</b>—Six cores enabled on one CPU complex.</li> </ul>
<b>Processor EPP Profile</b>	<p>Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Performance</b></li> <li>• <b>Balanced Performance</b>—This is the default option.</li> <li>• <b>Balanced Power</b></li> <li>• <b>Power</b></li> </ul>

Name	Description
<b>Autonomous Core C-state</b>	<p>Enables CPU Autonomous C-State, which converts the HALT instructions to the MWAIT instructions. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—This is the default option.</li> <li>• <b>Enabled</b></li> </ul>
<b>Energy Efficient Turbo</b>	<p>When energy efficient turbo is enabled, the optimal turbo frequency of the CPU turns dynamic based on CPU utilization. The power/performance bias setting also influences energy efficient turbo. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—This is the default option.</li> <li>• <b>Enabled</b></li> </ul>
<b>Hardware P-States</b>	<p>Enables processor Hardware P-State. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—HWPM is disabled.</li> <li>• <b>HWPM Native Mode</b>—HWPM native mode is enabled. This is the default option.</li> <li>• <b>HWPM OOB Mode</b>—HWPM Out-of-Box mode is enabled.</li> <li>• <b>Native Mode with no Legacy</b></li> </ul>
<b>Energy/Performance Bias Config</b>	<p>Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Performance</b>—The server provides all server components with full power at all times. This option maintains the highest level of performance and requires the greatest amount of power.</li> <li>• <b>Balanced Performance</b>—The server provides all server components with enough power to keep a balance between performance and power. This is the default option.</li> <li>• <b>Balanced Power</b>—The server provides all server components with enough power to keep a balance between performance and power.</li> <li>• <b>Power</b>—The server provides all server components with maximum power to keep reduce power consumption.</li> </ul>



Name	Description
<b>Power Performance Tuning</b>	<p>Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>BIOS</b>—Chooses BIOS for energy performance tuning.</li> <li>• <b>OS</b>—Chooses OS for energy performance tuning. This is the default option.</li> <li>• <b>PECI</b>—Chooses Peci for energy performance tuning.</li> </ul>
<b>Cores Enabled</b>	<p>Allows you to disable one or more of the physical cores on the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>All</b>—Enables all physical cores. This also enables Hyper Threading on the associated logical processor cores.</li> <li>• <b>1 through 48</b>—Specifies the number of physical processor cores that can run on the server. Each physical core has an associated logical core.</li> </ul>
<b>Hyper-Threading [All]</b>	<p>Whether the processor uses Intel Hyper-Threading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not permit hyperthreading.</li> <li>• <b>Enabled</b>—The processor allows for the parallel execution of multiple threads.</li> </ul>
<b>SpeedStep (Pstates)</b>	<p>Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor never dynamically adjusts its voltage or frequency.</li> <li>• <b>Enabled</b>—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power.</li> </ul>
<b>Boot Performance Mode</b>	<p>Allows the user to select the BIOS performance state that is set before the operating system handoff. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Max Performance</b>—Processor P-state ratio is maximum.</li> <li>• <b>Max Efficient</b>—Processor P-state ratio is minimum.</li> <li>• <b>Set by Intel NM</b>—Processor P-state ratio is set by Intel.</li> </ul>

Name	Description
<b>EIST PSD Function</b>	<p>EIST reduces the latency inherent with changing the voltage-frequency pair (P-state), thus allowing those transitions to occur more frequently. This allows for more granular, demand-based switching and can optimize the power-to-performance balance, based on the demands of the applications. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>HW All</b>—The processor is coordinates the P-state among logical processors dependencies. The OS keeps the P-state request up to date on all logical processors. This is the default option.</li> <li>• <b>SW All</b>—The OS Power Manager coordinates the P-state among logical processors with dependencies and initiates the transition on all of those Logical Processors.</li> </ul>
<b>Turbo Mode</b>	<p>Whether the processor uses Intel Turbo Boost Technology, which allows the processor to automatically increase its frequency if it is running below power, temperature, or voltage specifications. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not increase its frequency automatically.</li> <li>• <b>Enabled</b>—The processor utilizes Turbo Boost Technology if required. This is the default option.</li> </ul>
<b>Extended APIC</b>	<p>Allows you to enable or disable extended APIC support. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—This is the default option.</li> <li>• <b>Enabled</b>.</li> </ul>
<b>Memory Interleaving Size</b>	<p>Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8, 9, 10 or 11). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• 1 KB</li> <li>• 2 KB</li> <li>• 4 KB</li> <li>• 256 Bytes</li> <li>• 512 Bytes</li> <li>• <b>Auto</b>—The CPU determines the size of the memory block.</li> <li>• <b>Platform Default</b> —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>UPI Link Frequency Select</b>	<p>Allows you to enable or disable extended APIC support. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—This option configures the optimal link speed automatically. This is the default option.</li> <li>• <b>9.6GT/S</b>—This option configures the optimal link speed at 9.6GT/s.</li> <li>• <b>10.4GT/S</b>—This option configures the optimal link speed at 10.4GT/s.</li> <li>• <b>11.2GT/S</b>—This option configures the optimal link speed at 10.4GT/s.</li> <li>• <b>12.8GT/S</b>—This option configures the optimal link speed at 12.8GT/s.</li> <li>• <b>14.4GT/S</b>—This option configures the optimal link speed at 14.4GT/s.</li> <li>• <b>16.0GT/S</b>—This option configures the optimal link speed at 16.0GT/s.</li> <li>• <b>20.0GT/S</b>—This option configures the optimal link speed at 20.0GT/s.</li> </ul>
<b>X2APIC Opt Out</b>	<p>Prevents the OS from enabling extended xAPIC (x2APIC) mode when the OS is not working with x2APIC. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Use the Extended xAPIC (x2APIC) mode. This is the default option.</li> <li>• <b>Enabled</b>—Opt out from Extended xAPIC (x2APIC) mode.</li> </ul>

### I/O BIOS Settings for Intel

The following table lists the Intel Directed I/O BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Intel VT for directed IO</b>	<p>Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not use virtualization technology.</li> <li>• <b>Enabled</b>—The processor uses virtualization technology.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.</p>

Name	Description
<b>Intel VTD interrupt Remapping</b>	<p>Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not support remapping.</li> <li>• <b>Enabled</b>—The processor uses VT-d Interrupt Remapping as required.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Intel VTD coherency support</b>	<p>Whether the processor supports Intel VT-d Coherency. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not support coherency.</li> <li>• <b>Enabled</b>—The processor uses VT-d Coherency as required.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Intel VTD ATS support</b>	<p>Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not support ATS.</li> <li>• <b>Enabled</b>—The processor uses VT-d ATS as required.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Intel VTD pass through DMA support</b>	<p>Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The processor does not support pass-through DMA.</li> <li>• <b>Enabled</b>—The processor uses VT-d Pass-through DMA as required.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

## I/O BIOS Settings for AMD

The following table lists the Input/Output BIOS settings that you can configure through a BIOS policy for AMD:

Name	Description
<b>PCIe ARI Support</b>	<p>The PCIe Alternative Routing ID (ARI) Interpretation feature specification supports greater numbers of virtual functions through the implementation of ARI, which reinterprets the device number field in the PCIe header allowing for more than eight functions. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—PCIe ARI Support is not available.</li> <li>• <b>Enabled</b>—PCIe ARI Support is available.</li> <li>• <b>Auto</b>—PCIe ARI Support is in auto mode. This is the default option.</li> </ul>
<b>IPv4 PXE Support</b>	<p>Enables or disables IPv4 support for PXE. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—IPv6 PXE support is not available.</li> <li>• <b>Enabled</b>—IPv6 PXE support is available. This is the default option.</li> </ul>
<b>IPv4 HTTP Support</b>	<p>Enables or disables IPv4 support for HTTP. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—IPv4 HTTP support is not available.</li> <li>• <b>Enabled</b>—IPv4 HTTP support is available. This is the default option.</li> </ul>
<b>IPv6 HTTP Support</b>	<p>Enables or disables IPv6 support for HTTP. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—IPv6 HTTP support is not available.</li> <li>• <b>Enabled</b>—IPv6 HTTP support is available. This is the default option.</li> </ul>

Name	Description
<b>Network Stack</b>	<p>This option allows you to monitor IPv6 and IPv4. This can be one of the following</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Network Stack support is not available.</li> </ul> <p><b>Note</b> When disabled, the value set for <b>IPV4 PXE Support</b> does not impact the system.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Network Stack support is available. This is the default option.</li> </ul> <p><b>Note</b> When Network Stack token value is Disabled, the below tokens and their values are also set</p> <ul style="list-style-type: none"> <li>• IPV4PXE - Disabled</li> <li>• IPV4HTTP - Disabled</li> <li>• IPV6HTTP - Disabled</li> </ul>
<b>SR-IOV Support</b>	<p>Whether SR-IOV (Single Root I/O Virtualization) is enabled or disabled on the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—SR-IOV is enabled. This is the default option.</li> <li>• <b>Disabled</b>—SR-IOV is disabled.</li> </ul>

## RAS Memory BIOS Settings

The following table lists the RAS memory BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Partial Cache Line Sparing</b>	<p>Partial cache line sparing (PCLS) is an error-prevention mechanism in memory controllers. PCLS statically encodes the locations of the faulty nibbles of bits into a sparing directory along with the corresponding data content for replacement during memory accesses. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Support is disabled.</li> <li>• <b>Enabled</b>—Support is enabled.</li> </ul>
<b>UMA</b>	<p>Allows you to set UMA settings. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disable-All-2All</b></li> <li>• <b>Hemisphere-2-clusters</b></li> </ul>

Name	Description
<b>Memory Thermal Throttling Mode</b>	<p>Provides a protective mechanism to ensure the memory temperature is within the limits. When the temperature exceeds the maximum threshold value, the memory access rate is reduced and Baseboard Management Controller (BMC) adjusts the fan to cool down the memory to avoid DIMM damage due to overheat. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>CLTT with PECEI</b>—Closed Loop Thermal Throttling (CLTT) with Platform Environment Control Interface (PECEI). This is the default option.</li> <li>• <b>Disabled</b>.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>CLTT with PECEI</b></p>
<b>Enhanced Memory Test</b>	<p>Enables enhanced memory tests during the system boot and increases the boot time based on the memory. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—This is the default option.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>Auto</b>.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b></li> </ul> <p><b>Note</b> • This BIOS token name modified from <b>Advanced Memory Test</b> to <b>Enhanced Memory Test</b> for M6 servers.</p>
<b>Transparent Secure Memory Encryption (TSME)</b>	<p>Provides transparent hardware memory encryption of all data stored on system memory. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b></li> <li>• <b>Auto</b>—This is the default option.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>Auto</b> to mitigate Rowhammer-style attacks.</p>

Name	Description
<b>Secure Encrypted Virtualization (SEV)</b>	<p>Enables running encrypted virtual machines (VMs) in which the code and data of the VM are isolated. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>253 ASIDs</b></li> <li>• <b>509 ASIDs</b></li> <li>• <b>Auto</b>—This is the default option.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>Auto</b> to mitigate Rowhammer-style attacks.</p>
<b>DRAM SW Thermal Throttling</b>	<p>Provides a protective mechanism to ensure that the software functions within the temperature limits. When the temperature exceeds the maximum threshold value, the performance is permitted to drop allowing to cool down to the minimum threshold value. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b>—This is the default option.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>Disabled</b> to mitigate Rowhammer-style attacks.</p>
<b>Memory Refresh Rate</b>	<p>Controls the refresh rate of the memory controller and might affect the memory performance and power depending on memory configuration and workload. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>1x Refresh</b></li> <li>• <b>2x Refresh</b>—This is the default option.</li> </ul>
<b>Panic and High Watermark</b>	<p>Controls the delayed refresh capability of the memory controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>High</b>—The memory controller is allowed to postpone up to a maximum of eight refresh commands. The memory controller executes all the postponed refreshes within the refresh interval. For the ninth refresh command, the refresh priority becomes <b>Panic</b> and the memory controller pauses the normal memory transactions until all the postponed refresh commands are executed.</li> <li>• <b>Low</b>—This is the default option. The memory controller is not allowed to postpone refresh commands.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state (<b>Low</b>) which will help to reduce susceptibility to Rowhammer-style attacks.</p>



Name	Description
<b>Memory RAS configuration</b>	<p>How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Maximum Performance</b>—Optimizes the system performance and disables all the advanced RAS features.</li> <li>• <b>Mirroring</b>—System reliability is optimized by using half the system memory as backup. This mode is used for UCS M4 and lower blade servers.</li> <li>• <b>Lockstep</b>—If the DIMM pairs in the server have an identical type, size, and organization and are populated across the SMI channels, you can enable lockstep mode to minimize memory access latency and provide better performance. Lockstep is enabled by default for B440 servers.</li> <li>• <b>Mirror Mode 1LM</b>—Mirror Mode 1LM will set the entire 1LM memory in the system to be mirrored, consequently reducing the memory capacity by half. This mode is used for UCS M5 and M6 blade servers.</li> <li>• <b>Partial Mirror Mode 1LM</b>—Partial Mirror Mode 1LM will set a part of the 1LM memory in the system to be mirrored, consequently reducing the memory capacity by half. This mode is used for UCS M5 and M6 blade servers.</li> <li>• <b>Sparing</b>—System reliability is optimized by holding memory in reserve so that it can be used in case other DIMMs fail. This mode provides some memory redundancy, but does not provide as much redundancy as mirroring.</li> <li>• <b>ADDDC Sparing</b>—System reliability is optimized by holding memory in reserve so that it can be used in case other DIMMs fail. This mode provides some memory redundancy, but does not provide as much redundancy as mirroring.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>NUMA optimized</b>	<p>Whether the BIOS supports NUMA. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS does not support NUMA.</li> <li>• <b>Enabled</b>—The BIOS includes the ACPI tables that are required for NUMA-aware operating systems. If you enable this option, the system must disable Inter-Socket Memory interleaving on some platforms.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Post Package Repair</b>	<p>Post Package Repair (PPR) provides the ability to repair faulty memory cells by replacing them with spare cells. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS does not support selecting PPR Type.</li> <li>• <b>Hard PPR</b>—This results in a permanent remapping of damaged storage cells. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Memory Size Limit in GB</b>	<p>Limits the capacity in Partial Memory Mirror Mode up to 50 percent of the total memory capacity. The memory size can range from 0 GB to 65535 GB in increments of 1 GB.</p>
<b>Mirroring Mode</b>	<p>Memory mirroring enhances system reliability by keeping two identical data images in memory.</p> <p>This option is only available if you choose the <b>mirroring</b> option for <b>Memory RAS Config</b>. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Inter-Socket</b>—Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets.</li> <li>• <b>Intra-Socket</b>—One IMC is mirrored with another IMC in the same socket.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Sparing Mode</b>	<p>Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population.</p> <p>This option is only available if you choose <b>sparing</b> option for <b>Memory RAS Config</b>. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>DIMM Sparing</b>—One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM.</li> <li>• <b>Rank Sparing</b>—A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>LV DDR Mode</b>	<p>Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The CPU determines whether to prioritize low voltage or high frequency memory operations.</li> <li>• <b>Power Saving Mode</b>—The system prioritizes low voltage memory operations over high frequency memory operations. This mode may lower memory frequency in order to keep the voltage low.</li> <li>• <b>Performance Mode</b>—The system prioritizes high frequency operations over low voltage operations.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>DRAM Refresh Rate</b>	<p>The refresh interval rate for internal memory. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>1x</b></li> <li>• <b>2x</b></li> <li>• <b>3x</b></li> <li>• <b>4x</b></li> <li>• <b>Auto</b></li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>DDR3 Voltage Selection</b>	<p>The voltage to be used by the dual-voltage RAM. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>DDR3-1500mv</b></li> <li>• <b>DDR3-1350mv</b></li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Partial Memory Mirror Mode</b>	<p>Partial Memory Mirroring enables you to partially mirror by GB or by a percentage of the memory capacity. Depending on the option selected here, you can define either a partial mirror percentage or a partial mirror capacity in GB in available fields. You can partially mirror up to 50 percent of the memory capacity. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Partial Memory Mode is disabled. This is the default option.</li> <li>• <b>Percentage</b>—The amount of memory to be mirrored in the Partial Memory Mode is defined as a percentage of the total memory.</li> <li>• <b>Value in GB</b>—The amount of memory to be mirrored in the Partial Memory Mode is defined in GB.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b>      <b>Partial Memory Mirror Mode</b> is mutually exclusive to standard <b>Mirroring Mode</b>.</p> <p>Partial Mirrors 1-4 can be used in any number or configuration, provided they do not exceed the capacity limit set in GB or Percentage in the related options.</p>
<b>Partial Mirror Percentage</b>	Limits the amount of available memory to be mirrored as a percentage of the total memory. This can range from 0.000.01 % to 50.00 % in increments of 0.01 %.
<b>Partial Mirror1 Size in GB</b>	Limits the amount of memory in Partial Mirror1 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
<b>Partial Mirror2 Size in GB</b>	Limits the amount of memory in Partial Mirror2 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
<b>Partial Mirror3 Size in GB</b>	Limits the amount of memory in Partial Mirror3 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.

Name	Description
<b>Partial Mirror4 Size in GB</b>	Limits the amount of memory in Partial Mirror4 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
<b>Volatile Memory Mode</b>	<p>Allows the memory mode configuration. This can be any of the following:</p> <ul style="list-style-type: none"> <li>• <b>1LM</b>—Configures 1 Layer Memory(1LM)</li> <li>• <b>2LM</b>—Configures 2 Layer Memory(1LM)</li> </ul>
<b>Memory Bandwidth Boost</b>	<p>Allows to boost the memory bandwidth. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b></li> </ul>
<b>Burst and Postponed Refresh</b>	<p>Allows the memory controller to defer the refresh cycles when the memory is active and accomplishes the refresh within a specified window. The deferred refresh cycles may run in a burst of several refresh cycles. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b>—This is the default option.</li> </ul> <p><b>Note</b> It is recommended to leave this setting in the default state of <b>Disabled</b> to mitigate Rowhammer-style attacks.</p>
<b>LLC Dead Line</b>	<p>In CPU non-inclusive cache scheme, Mid-Level Cache (MLC) evictions are filled into the Last-Level Cache (LLC). When lines are evicted from the MLC, the core can flag them as dead (not likely to be read again). The LLC has the option to drop dead lines and not fill them in the LLC. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Allows the LLC to fill dead lines into the LLC if there is free space available. This is the default option.</li> <li>• <b>Disabled</b>—The dead lines are always dropped and are never filled into the LLC.</li> <li>• <b>Auto</b>—The CPU determines the LLC dead line allocation</li> </ul>

Name	Description
<b>XPT Remote Prefetch</b>	<p>This feature allows an LLC request to be duplicated and sent to an appropriate memory controller in a remote machine based on the recent LLC history to reduce latency. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b></li> <li>• <b>Auto</b>—The CPU determines the functionality. This is the default option.</li> </ul>
<b>Virtual NUMA</b>	<p>The Virtual NUMA (virtual non-uniform memory access) is a memory-access optimization method for VMware virtual machines (VMs), which helps prevent memory-bandwidth bottlenecks. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The functionality is enabled.</li> <li>• <b>Disabled</b>—The functionality is disabled. This is the default option.</li> </ul>
<b>Above 4G Decoding</b>	<p>Enables or disables MMIO above 4GB or not. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The server maps I/O of 64-bit PCI devices to 4GB or greater address space. This is the default option.</li> <li>• <b>Disabled</b>—The server does not map I/O of 64-bit PCI devices to 4GB or greater address space.</li> </ul>
<b>NUMA Nodes per Socket</b>	<p>Allows you to configure the memory NUMA domains per socket. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Number of channels is set to auto. This is the default option.</li> <li>• <b>NPS0</b>—Zero NUMA node per socket.</li> <li>• <b>NPS1</b>—One NUMA node per socket.</li> <li>• <b>NPS2</b>—Two NUMA nodes per socket, one per Left/Right Half of the SoC.</li> <li>• <b>NPS4</b>—Four NUMA nodes per socket, one per Quadrant.</li> </ul>

Name	Description
<b>Select PPR Type</b>	<p>Supports <b>Hard-PPR</b>, which permanently remaps accesses from a designated faulty row to a designated spare row.</p> <ul style="list-style-type: none"> <li>• <b>Hard PPR</b>—Support is enabled. This is the default option.</li> </ul> <p><b>Note</b> Hard PPR can be used only when <b>Memory RAS Configuration</b> is set to <b>ADDDC Sparing</b>. For other RAS selections, this setting should be set to <b>Disabled</b>.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>Select Memory RAS Configuration</b>	<p>Determines how the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Mirror Mode 1LM</b>—System reliability is optimized by using half the system memory as backup.</li> <li>• <b>ADDDC Sparing</b>—Adaptive virtual lockstep is an algorithm implemented in the hardware and firmware to support the ADDDC mode. When selected, the system performance is optimized till the algorithm is activated. The algorithm is activated in case of DRAM device failure. Once the algorithm is activated, the virtual lockstep regions are activated to map out the failed region during run-time dynamically, and the performance impact is restricted at a region level. This is the default option.</li> <li>• <b>Partial Mirror Mode 1LM</b>—Partial DIMM Mirroring creates a mirrored copy of a specific region of memory cells, rather than keeping the complete mirror copy. Partial Mirroring creates a mirrored region in memory map with the attributes of a partial mirror copy. Up to 50% of the total memory capacity can be mirrored, using up to 4 partial mirrors.</li> <li>• <b>Maximum Performance</b>—System performance is optimized.</li> </ul>
<b>NUMA</b>	<p>Whether the BIOS supports Non-Uniform Memory Access (NUMA). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Support is enabled.</li> <li>• <b>Disabled</b>—Support is disabled.</li> </ul>
<b>Operation Mode</b>	<p>Allows you to set the Operation Mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Test Only</b>—Support is enabled.</li> <li>• <b>Test and Repair</b>—Support is disabled.</li> </ul>

## Intel® Optane™ DC Persistent Memory (DCPMM) BIOS Tokens

The following table lists the Intel® Optane™ DC Persistent Memory (DCPMM) BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>NVM Performance Setting</b>	<p><b>NVM Performance Setting</b> enables efficient major mode arbitration between DDR and DDRT transactions on the DDR channel to optimize channel BW and DRAM latency.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>BW Optimized</b>—Optimized for DDR and DDRT BW. This is the default option.</li> <li>• <b>Latency Optimized</b>—Better DDR latency in the presence of DDRT BW.</li> <li>• <b>Balanced Profile</b>—Optimized for Memory mode.</li> </ul>
<b>CR QoS</b>	<p>Prevents DRAM and overall system BW drop in the presence of concurrent DCPMM BW saturating threads, with minimal impact to homogenous DDRT-only usages, Good for multi-tenant use cases, VMs, etc. Targeted for App Direct, but also improves memory mode. Targets the “worst-case” degradations.</p> <p>Applies to all M5 and M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Feature disabled. This is the default option.</li> <li>• <b>Recipe 1</b>—6 modules, 4 modules per socket optimized</li> <li>• <b>Recipe 2</b>—2 modules per socket optimized</li> <li>• <b>Recipe 3</b>—1 module per socket optimized</li> <li>• <b>Mode 0 - Disable the PMem QoS Feature</b></li> <li>• <b>Mode 1 - M2M QoS Enable;CHA QoS Disable</b></li> <li>• <b>Mode 2 - M2M QoS Enable;CHA QoS Enable</b></li> </ul> <p><b>Note</b> The values <b>Disabled</b>, <b>Recipe 1</b>, <b>Recipe 2</b>, and <b>Recipe 3</b> are not supported on Cisco UCS M6 servers</p>



Name	Description
<b>CR FastGo Config</b>	<p><b>CR FastGo Config</b> improves DDRT non-temporal write bandwidth when FastGO is disabled. When FastGO is enabled, it gives faster flow of NT writes into the uncore, When FastGO is disabled, it lessens NT writes queueing up in the CPU uncore, thereby improving sequentially at DCPMM, resulting in improved bandwidth.</p> <p>Applies to all Cisco UCS M5 and Cisco UCS M6 servers.</p> <p>The values can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Auto</b>—Same as Option 1. Disables FastGO. Recommended for DDRT. This is the default option (not Default).</li> <li>• <b>Default</b>—Enables FastGO.</li> <li>• <b>Option 1</b>—Disables FastGO.</li> <li>• <b>Option 2, Option 3, Option 4, Option 5</b>—Not applicable.</li> <li>• <b>Enable Optimization</b></li> <li>• <b>Disable Optimization</b></li> </ul> <p><b>Note</b> The values <b>Enable Optimization</b>, <b>Disable Optimization</b>, and <b>Auto</b> are supported on Cisco UCS M6 servers</p>
<b>Snoopy mode for AD</b>	<p>Enables snoop-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for accesses to AD and instead snoops remote sockets to check for ownership. Directory is used only for DRAM accesses.</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b> This is the default option.</li> </ul>

Name	Description
<b>Snoopy mode for 2LM</b>	<p>Enables snoopy-mode for DCPMM accesses while maintaining directory on all DRAM accesses. Snoops maintain cache coherence between sockets. Directory reduces snoops by keeping the remote node information locally (in memory). Directory lookups and updates add memory traffic.</p> <p>Directory is a good tradeoff for DRAM, but not necessarily for DCPMM. For non-NUMA workload, when the feature is enabled, directory updates to DCPMM are eliminated, thereby helping DDRT bandwidth bound workloads. Directory is disabled for far memory accesses and instead snoops remote sockets to check for ownership. Directory is used only for DRAM (near memory).</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b> This is the default option.</li> </ul>
<b>eADR Support</b>	<p>Extended asynchronous DRAM refresh (eADR) ensures that CPU caches lines with data are flushed at the right time and in the desired order and are also included in the <b>power fail protected domain</b>. This can be any of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b></li> <li>• <b>Disabled</b></li> <li>• <b>Auto</b>—This is the default option.</li> </ul>

### Serial Port BIOS Settings

The following table lists the serial port BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Serial port A enable</b>	<p>Whether serial port A is enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The serial port is disabled.</li> <li>• <b>Enabled</b>—The serial port is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### USB BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Make Device Non Bootable</b>	<p>Whether the server can boot from a USB device. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The server can boot from a USB device.</li> <li>• <b>Enabled</b>—The server cannot boot from a USB device.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Legacy USB Support</b>	<p>Whether the system supports legacy USB devices. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—USB devices are only available to EFI applications.</li> <li>• <b>Enabled</b>—Legacy USB support is always available.</li> <li>• <b>Auto</b>—Disables legacy USB support if no USB devices are connected.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>USB Idle Power Optimizing Setting</b>	<p>Whether the USB Idle Power Optimizing setting is used to reduce USB EHCI idle power consumption. Depending upon the value you choose, this setting can have an impact on performance. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>high-performanceHigh Performance</b>—The USB System Idle Power Optimizing setting is disabled, because optimal performance is preferred over power savings.  Selecting this option can significantly improve performance. We recommend you select this option unless your site has server power restrictions.</li> <li>• <b>Lower Idle Power</b>—The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>USB Front Panel Access Lock</b>	<p>USB front panel access lock is configured to enable or disable the front panel access to USB ports. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b></li> <li>• <b>Enabled</b></li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Port 60/64 Emulation</b>	<p>Whether the system supports 60h/64h emulation for complete USB keyboard legacy support. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—60h/64 emulation is not supported.</li> <li>• <b>Enabled</b>—60h/64 emulation is supported.</li> </ul> <p>You should select this option if you are using a non-USB aware operating system on the server.</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>USB Port Front</b>	<p>Whether the front panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the front panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system.</li> <li>• <b>Enabled</b>—Enables the front panel USB ports. Devices connected to these ports are detected by the BIOS and operating system.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>USB Port Internal</b>	<p>Whether the internal USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the internal USB ports. Devices connected to these ports are not detected by the BIOS and operating system.</li> <li>• <b>Enabled</b>—Enables the internal USB ports. Devices connected to these ports are detected by the BIOS and operating system.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>USB Port KVM</b>	<p>Whether the vKVM ports are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the KVM keyboard and/or mouse devices. Keyboard and/or mouse will not work in the KVM window.</li> <li>• <b>Enabled</b>—Enables the KVM keyboard and/or mouse devices.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>USB Port Rear</b>	<p>Whether the rear panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the rear panel USB ports. Devices connected to these ports are not detected by the BIOS and operating system.</li> <li>• <b>Enabled</b>—Enables the rear panel USB ports. Devices connected to these ports are detected by the BIOS and operating system.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>USB Port SD Card</b>	<p>Whether the SD card drives are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the SD card drives. The SD card drives are not detected by the BIOS and operating system.</li> <li>• <b>Enabled</b>—Enables the SD card drives.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>USB Port VMedia</b>	<p>Whether the virtual media devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables the vMedia devices.</li> <li>• <b>Enabled</b>—Enables the vMedia devices.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>All USB Devices</b>	<p>Whether all physical and virtual USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—All USB devices are disabled.</li> <li>• <b>Enabled</b>—All USB devices are enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>xHCI Mode</b>	<p>Whether xHCI mode is enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—xHCI mode is disabled.</li> <li>• <b>Enabled</b>—xHCI mode is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>USB Port:M.2 Storage</b>	<p>Whether the USB Port:M.2 Storage are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables USB Port:M.2 Storage.</li> <li>• <b>Enabled</b>—Enables USB Port:M.2 Storage. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### PCI Configuration BIOS Settings

The following table lists the PCI configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Maximum memory below 4GB</b>	<p>Whether the BIOS maximizes memory usage below 4GB for an operating system without PAE support, depending on the system configuration. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Does not maximize memory usage. Choose this option for all operating systems with PAE support.</li> <li>• <b>Enabled</b>—Maximizes memory usage below 4GB for an operating system without PAE support.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Memory mapped IO above 4GB</b>	<p>Whether to enable or disable memory mapped I/O of 64-bit PCI devices to 4GB or greater address space. Legacy option ROMs are not able to access addresses above 4GB. PCI devices that are 64-bit compliant but use a legacy option ROM may not function correctly with this setting enabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Does not map I/O of 64-bit PCI devices to 4GB or greater address space.</li> <li>• <b>Enabled</b>—Maps I/O of 64-bit PCI devices to 4GB or greater address space.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>VGA Priority</b>	<p>Allows you to set the priority for VGA graphics devices if multiple VGA devices are found in the system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Onboard</b>—Priority is given to the onboard VGA device. BIOS post screen and OS boot are driven through the onboard VGA port.</li> <li>• <b>Offboard</b>—Priority is given to the PCIE Graphics adapter. BIOS post screen and OS boot are driven through the external graphics adapter port.</li> <li>• <b>Onboard VGA Disabled</b>—Priority is given to the PCIE Graphics adapter, and the onboard VGA device is disabled.</li> </ul> <p><b>Note</b>           The vKVM does not function when the onboard VGA is disabled.</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b>           Only onboard VGA devices are supported with Cisco UCS B-Series servers.</p>
<b>ASPM Support</b>	<p>Allows you to set the level of ASPM (Active Power State Management) support in the BIOS. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—ASPM support is disabled in the BIOS.</li> <li>• <b>Auto</b>—The CPU determines the power state.</li> <li>• <b>ForceL0</b>—Force all links to L0 standby (L0s) state.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>BME DMA Mitigation Support</b>	<p>Allows you to disable the PCI BME bit to mitigate the threat from an unauthorized external DMA. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—PCI BME bit is disabled in the BIOS.</li> <li>• <b>Enabled</b>—PCI BME bit is enabled in the BIOS.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



## QPI BIOS Settings

The following table lists the QPI BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>QPI Link Frequency Select</b>	<p>The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>20.0GT/s</b></li> <li>• <b>16.0GT/s</b></li> <li>• <b>14.4GT/s</b></li> <li>• <b>12.8GT/s</b></li> <li>• <b>6.4 GT/s</b></li> <li>• <b>7.2 GT/s</b></li> <li>• <b>8.0 GT/s</b></li> <li>• <b>9.6 GT/s</b></li> <li>• <b>Auto</b>—The CPU determines the QPI link frequency.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>QPI Snoop Mode</b>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Home Snoop</b>—The snoop is always spawned by the home agent (centralized ring stop) for the memory controller. This mode has a higher local latency than early snoop, but it provides extra resources for a larger number of outstanding transactions.</li> <li>• <b>Cluster On Die</b>—This mode is available only for processors that have 10 or more cores. It is the best mode for highly NUMA optimized workloads.</li> <li>• <b>Home Directory Snoop with OSB</b></li> <li>• <b>Early Snoop</b>—The distributed cache ring stops can send a snoop probe or a request to another caching agent directly. This mode has lower latency and it is best for workloads that have shared data sets across threads and can benefit from a cache-to-cache transfer, or for workloads that are not NUMA optimized.</li> <li>• <b>Auto</b> —The CPU determines the QPI Snoop mode.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

## Trusted Platform BIOS Settings

The following table lists the trusted platform BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Trusted Platform Module (TPM) Support</b>	<p>Whether to enable or disable the Trusted Platform Module (TPM), which is a component that securely stores artifacts that are used to authenticate the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables TPM.</li> <li>• <b>Enabled</b>—Enables TPM.</li> <li>• <b>Platform Default</b>—Enables TPM.</li> </ul>
<b>Intel Trusted Execution Technology (TXT) Support</b>	<p>Whether to enable or disable Intel Trusted Execution Technology (TXT), which provides greater protection for information that is used and stored on the business server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables TXT. This is default option.</li> <li>• <b>Enabled</b>—Enables TXT.</li> <li>• <b>Platform Default</b>—Disables TXT.</li> </ul> <p>When you only enable TXT, it implicitly enables TPM, VT, and VTDio.</p>
<b>SHA-1 PCR Bank</b>	<p>The Platform Configuration Register (PCR) is a memory location in the TPM. Multiple PCRs are collectively referred to as a PCR bank. A Secure Hash Algorithm 1 or SHA-1 PCR Bank allows to enable or disable TPM security. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables SHA-1 PCR Bank.</li> <li>• <b>Enabled</b>—Enables SHA-1 PCR Bank. This is the default option.</li> </ul>
<b>SHA-256 PCR Bank</b>	<p>The Platform Configuration Register (PCR) is a memory location in the TPM. Multiple PCRs are collectively referred to as a PCR bank. A Secure Hash Algorithm 256-bit or SHA-256 PCR Bank allows to enable or disable TPM security. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables SHA-256 PCR Bank.</li> <li>• <b>Enabled</b>—Enables SHA-256 PCR Bank. This is the default option.</li> </ul>
<b>Trusted Platform Module State</b>	<p>Trusted Platform Module (TPM ) is a microchip designed to provide basic security-related functions primarily involving encryption keys. This option allows you to control the TPM Security Device support for the system. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The server does not use the TPM.</li> <li>• <b>Enabled</b>—The server uses the TPM. This is the default option.</li> </ul>

Name	Description
<b>TPM Pending Operation</b>	Trusted Platform Module (TPM) Pending Operation option allows you to control the status of the pending operation. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>None</b>—No action. This is the default option.</li> <li>• <b>TPM Clear</b>—Clear the pending operations.</li> </ul>
<b>TPM Minimal Physical Presence</b>	Whether to enable or disable TPM Minimal Physical Presence, which enables or disables the communication between the OS and BIOS for administering the TPM without compromising the security. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables TPM Minimal Physical Presence. This is default option.</li> <li>• <b>Enabled</b>—Enables TPM Minimal Physical Presence.</li> <li>• <b>Platform Default</b>—Disables TPM Minimal Physical Presence.</li> </ul>
<b>DMA Control Opt-In Flag</b>	Enabling this token enables Windows 2022 Kernel DMA Protection feature. The OS treats this as a hint that the IOMMU should be enabled to prevent DMA attacks from possible malicious devices. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Disables DMA Control Opt-In Flag. This is default option.</li> <li>• <b>Enabled</b>—Enables DMA Control Opt-In Flag.</li> <li>• <b>Platform Default</b>—Disables DMA Control Opt-In Flag.</li> </ul>
<b>Security Dev. Support</b>	Enables or disables BIOS support for the security device. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Disabled</b>—OS will not show the security device.</li> <li>• <b>Enabled</b>—OS will show the security device. This is default option.</li> </ul>

### LOM and PCIe Slots BIOS Settings

The following table lists the USB BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>PCIe Slot SAS OptionROM</b>	<p>Whether Option ROM is available on the SAS port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot <i>n</i> Link Speed</b>	<p>This option allows you to restrict the maximum speed of an adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot <i>n</i> OptionROM</b>	<p>Whether Option ROM is available on the port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot HBA OptionROM</b>	<p>Whether Option ROM is available on the HBA port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>PCIe Slot MLOM OptionROM</b>	<p>Whether Option ROM is available on the MLOM port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot Nx OptionROM</b>	<p>Whether Option ROM is available on the port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe 10G LOM 2 Link</b>	<p>Whether Option ROM is available on the 10G LOM port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCI ROM CLP</b>	<p>PCI ROM Command Line Protocol (CLP) controls the execution of different Option ROMs such as PxE and iSCSI that are present in the card. By default, it is disabled.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>SIOC1 Option ROM</b>	<p>Whether the server can use Option ROM present in System IO Controller 1 (SIOC1). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SIOC2 Option ROM</b>	<p>Whether the server can use Option ROM present in System IO Controller 2 (SIOC2). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBMEZZ1 Option ROM</b>	<p>Whether the server can use Option ROM present in SBMezz1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBMEZZ2 Option ROM</b>	<p>Whether the server can use Option ROM present in SBMezz2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>IOESlot1 OptionROM</b>	<p>Whether option ROM is enabled on the IOE slot 1. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IOEMezz 1 OptionROM</b>	<p>Whether option ROM is enabled on the IOE Mezz1. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IOE Slot2 Option ROM</b>	<p>Whether option ROM is enabled on the IOE slot 2. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IO ENVME1 Option ROM</b>	<p>Whether option ROM is enabled on the IOE NVMe1. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>IO ENVME2 Option ROM</b>	<p>Whether option ROM is enabled on the IOE NVMe2. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBNVME1 Option ROM</b>	<p>Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot MRAID-<i>n</i> OptionROM</b>	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot RAID OptionROM</b>	<p>Whether Option ROM is available on the RAID port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



Name	Description
<b>Rear NVME <i>n</i> Link Speed</b>	<p>This option allows you to restrict the maximum speed of an NVME card installed in the rear PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• For <i>Rear NVME 1 Link Speed</i> and <i>Rear NVME 2Link Speed</i>, the value <b>Enabled</b> is not supported on Cisco UCS M6 servers.</li> <li>• For <i>Rear NVME 3 Link Speed</i> and <i>Rear NVME 4Link Speed</i>, the value <b>Enabled</b> is available but has no effect at the BIOS level if selected.</li> </ul> <ul style="list-style-type: none"> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Front NVME <i>n</i> Link Speed</b>	<p>This option allows you to restrict the maximum speed of an NVME card installed in the front PCIe slot. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically. This is the default option.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> </ul> <p><b>Note</b> For <i>Front NVME 1 Link Speed</i> and <i>Front NVME 2 Link Speed</i>, the value <b>Enabled</b> is available but not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> For <i>Front Nvme 13 Link Speed</i> to <i>Front Nvme 24 Link Speed</i>, the BIOS tokens and values are available but have no effect at the BIOS level if selected.</p>

Name	Description
<b>HBA Link Speed</b>	<p>This option allows you to restrict the maximum speed of an HBA card. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>MLOM Link Speed</b>	<p>This option allows you to restrict the maximum speed of an MLOM adapter. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> </ul> <p><b>Note</b>        The value <b>Enabled</b> is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>MRAID Link Speed</b>	<p>This option allows you to restrict the maximum speed of MRAID. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Enabled</b>—The maximum speed is not restricted.</li> </ul> <p><b>Note</b> The value <b>Enabled</b> is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>RAID-<i>n</i> Link Speed</b>	<p>This option allows you to restrict the maximum speed of RAID. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>All Onboard LOM</b>	<p>Whether all onboard LOM ports are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—All onboard LOM are enabled.</li> <li>• <b>Disabled</b>—All onboard LOM are disabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>LOM Port 1 OptionRom</b>	<p>Whether Option ROM is available on the LOM port 1. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>LOM Port 2 OptionRom</b>	<p>Whether Option ROM is available on the LOM port 2. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Slot <i>n</i> State</b>	<p>The state of the adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>UEFI Only</b>—The expansion slot is available for UEFI only.</li> <li>• <b>Legacy Only</b>—The expansion slot is available for legacy only.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBNVMe1 OptionROM</b>	<p>Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>SBNVMe2 OptionROM</b>	<p>Whether the server can use Option ROM present in SBNVMe2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SIOCNVMe1 OptionROM</b>	<p>Whether the server can use Option ROM present in SIOCNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SIOCNVMe2 OptionROM</b>	<p>Whether the server can use Option ROM present in SIOCNVMe2 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBLom1 OptionROM</b>	<p>Whether the server can use Option ROM present in the SBLom1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBNVMe<math>n</math> Link Speed</b>	<p>Link speed for SBNVMe slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>SIOCNVMe<math>n</math> Link Speed</b>	<p>Link speed for SIOCNVMe slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SIOC<math>n</math> Link Speed</b>	<p>Link speed for SIOC slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SBMezz<math>n</math> Link Speed</b>	<p>Link speed for SBMezz slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>IOESlot<math>n</math> Link Speed</b>	<p>Link speed for IOE slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IOEMezz<math>n</math> Link Speed</b>	<p>Link speed for IOEMezz slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IOENVMe<math>n</math> Link Speed</b>	<p>Link speed for IOENVMe slot <math>n</math>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Enabled</b>—The maximum speed is restricted.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>CDN Support for LOMs</b>	<p>Whether the Ethernet Networking Identifier naming convention is according to Consistent Device Naming (CDN) or the traditional way of naming conventions. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—OS Ethernet Network Identifier is named in a consistent device naming (CDN) convention according to the physical LAN on Motherboard (LOM) port numbering; LOM Port 0, LOM Port 1 and so on.</li> <li>• <b>Disabled</b>—OS Ethernet Networking Identifier is named in a default convention as ETH0, ETH1 and so on. By default, CDN option is disabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>VMD Enable</b>	<p>Whether NVMe SSDs that are connected to the PCIe bus can be hot swapped. It also standardizes the LED status light on these drives. LED status lights can be optionally programmed to display specific Failure indicator patterns.</p> <p>This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—Hot swap of NVMe SSDs that are connected to the PCIe bus is allowed.</li> <li>• <b>Disabled</b>—Hot swap of NVMe SSDs that are connected to the PCIe bus is not allowed.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>ACS Control SLOT-<i>n</i></b> <i>n</i> = 11 to 14	<p>Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for Control Slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>— Enables peer-to-peer communication between multiple devices for Control Slot <i>n</i>.</li> <li>• <b>Disabled</b>— Disables peer-to-peer communication between multiple devices for Control Slot <i>n</i>.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe Slot GPU<i>n</i> OptionROM</b>  Only for Cisco UCS C480 M5 ML Server	<p>Whether the Option ROM is enabled on GPU slot <i>n</i>. <i>n</i> is the slot number, which can be numbered 1 through 8. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



Name	Description
<b>ACS Control GPU-<i>n</i></b> <i>n</i> = 1 to 8	<p>Access Control Services (ACS) allow the processor to enable or disable peer-to-peer communication between multiple devices for GPUs. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>— Enables peer-to-peer communication between multiple devices for GPUs.</li> <li>• <b>Enabled</b>— Disables peer-to-peer communication between multiple devices for GPUs.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>PCIe PLL SSC</b>	<p>Reduces EMI interference by down-spreading the clock by 0.5%.            Disable this feature to centralize the clock without spreading.            For all Cisco UCS M5 and M6 servers, this option is Disabled by default.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>— Clock is centralized without spreading.</li> <li>• <b>Auto</b>— EMI interference is auto adjusted.</li> <li>• <b>ZeroPointFive</b>— EMI interference us reduced by down-spreading the clock by 0.5%.</li> <li>• <b>Platform Default</b>— The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Front Nvme <i>n</i> OptionROM</b>	<p>This options allows you to control the Option ROM execution of the PCIe adapter connected to the SSD:NVMe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—This is the default option.</li> <li>• <b>Disabled</b></li> </ul>
<b>PCIe Slot <i>n</i> Link Speed</b>	<p>Link speed for PCIe Slot designated by slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> </ul>

Name	Description
<b>MSTOR-RAID Link Speed</b>	<p>This option allows you to restrict the maximum speed of an MSTOR adapter. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> </ul> <p><b>Note</b> In this BIOS setting <i>MSTOR-RAID Link Speed</i>, the token and values are available but have no effect at the BIOS level if selected.</p>
<b>MSTOR-RAID OptionROM</b>	<p>Whether the server can use the Option ROMs present in the PCIe MSTOR RAID. This can be any of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Option ROM is available.</li> <li>• <b>Enabled</b>—Option ROM is not available. This is the default option.</li> </ul>
<b>MLOM OptionROM</b>	<p>Whether Option ROM is available on the MLOM port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>
<b>MRAID OptionROM</b>	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>
<b>Rear Nvme <i>n</i> OptionRom</b>	<p>Whether Option ROM is available on the Rear NVME<i>n</i> port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>PCIe slot MSTOR Link Speed</b>	<p>This option allows you to restrict the maximum speed of an MSTOR adapter. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically. This is the default option.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> </ul>
<b>PCIe Slot MSTOR RAID OptionROM</b>	<p>Whether the server can use the Option ROMs present in the PCIe MSTOR RAID. This can be any of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Option ROM is available.</li> <li>• <b>Enabled</b>—Option ROM is not available. This is the default option.</li> </ul>
<b>PCIe RAS Support</b>	<p>Whether PCIe RAS Support is available on the PCIe slot. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—PCIe RAS is available on the slot.</li> <li>• <b>Enabled</b>—PCIe RAS is not available on the slot. This is the default option.</li> </ul>
<b>MRAID<sub>n</sub> Link Speed</b>	<p>This option allows you to restrict the maximum speed of MRAID. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Gen 4</b>—16GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>MRAID<sub>n</sub> OptionROM</b>	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>

Name	Description
<b>NVME-<i>n</i> OptionROM</b>	<p>Whether Option ROM is available on the NVME port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>
<b>PCIe Slot OCP Link Speed</b>	<p>This option allows you to restrict the maximum speed of OCP. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically. This is the default option.</li> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>RAID<i>n</i> OptionROM</b>	<p>Whether Option ROM is available on the RAID port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>
<b>IOENVMe<i>n</i> OptionROM</b>	<p>Whether Option ROM is available on the IOENVMe port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>
<b>GPU<i>n</i> OptionRom</b>	<p>Whether Option ROM is available on the GPU port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The expansion slot is not available.</li> <li>• <b>Enabled</b>—The expansion slot is available. This is the default option.</li> </ul>

Name	Description
<b>RAID Link Speed</b>	<p>This option allows you to restrict the maximum speed of RAID. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Gen 1</b>—2.5GT/s (gigatransfers per second) is the maximum speed allowed.</li> <li>• <b>Gen 2</b>—5GT/s is the maximum speed allowed.</li> <li>• <b>Gen 3</b>—8GT/s is the maximum speed allowed.</li> <li>• <b>Auto</b>—The maximum speed is set automatically. This is the default option.</li> <li>• <b>Enabled</b>—The maximum speed is not restricted.</li> </ul> <p><b>Note</b> The value <b>Enabled</b> is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The maximum speed is not restricted.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### Graphics Configuration BIOS Settings

The following tables list the graphics configuration BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Integrated Graphics</b>	<p>Enables integrated graphics. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Integrated Graphics Aperture Size</b>	<p>Allows you to set the size of mapped memory for the integrated graphics controller. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Onboard Graphics</b>	<p>Enables onboard graphics (KVM). This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### Boot Options BIOS Settings

The following table lists the boot options BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

Name	Description
<b>Boot option retry</b>	<p>Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Waits for user input before retrying NON-EFI based boot options. This is the default option.</li> <li>• <b>Enabled</b>—Continually retries NON-EFI based boot options without waiting for user input.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SAS RAID</b>	<p>Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The Intel SAS Entry RAID Module is disabled.</li> <li>• <b>Enabled</b>—The Intel SAS Entry RAID Module is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>SAS RAID module</b>	<p>How the Intel SAS Entry RAID Module is configured. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>it-ir-raid</b>—Configures the RAID module to use Intel IT/IR RAID.</li> <li>• <b>intel-esrtii</b>—Configures the RAID module to use Intel Embedded Server RAID Technology II.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Onboard SCU Storage Support</b>	<p>Whether the onboard software RAID controller is available to the server. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The software RAID controller is not available.</li> <li>• <b>Enabled</b>—The software RAID controller is available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Cool Down Time (sec)</b>	<p>The time to wait (in seconds) before the next boot attempt. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>15</b>—System waits for 15 seconds before the next boot attempt.</li> <li>• <b>45</b>—System waits for 45 seconds before the next boot attempt.</li> <li>• <b>90</b>—System waits for 90 seconds before the next boot attempt. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>This token is valid only when the Boot Option Retry token has been enabled.</p>
<b>Number of Retries</b>	<p>Number of attempts to boot. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Infinite</b>—System tries all options to boot up.</li> <li>• <b>13</b>—System tries 13 times to boot up. This is the default option.</li> <li>• <b>5</b>—System tries 5 times to boot up</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>P-SATA mode</b>	<p>This options allows you to select the P-SATA mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—P-SATA mode is disabled.</li> <li>• <b>LSI SW RAID</b>—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Power On Password</b>	<p>This token requires that you set a BIOS password before using the F2 BIOS configuration. If enabled, password needs to be validated before you access BIOS functions such as IO configuration, BIOS set up, and booting to an operating system using BIOS. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Power On Password is disabled.</li> <li>• <b>Enabled</b>—Power On Password is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>IPV6 PXE Support</b>	<p>Enables or disables IPV6 support for PXE. This can be one of the following</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—IPV6 PXE support is not available.</li> <li>• <b>Enabled</b>—IPV6 PXE support is always available.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Adaptive Memory Training</b>	<p>When this token is enabled, the BIOS saves the memory training results (optimized timing/voltage values) along with CPU/memory configuration information and reuses them on subsequent reboots to save boot time. The saved memory training results are used only if the reboot happens within 24 hours of the last save operation. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—<b>Adaptive Memory Training</b> is disabled.</li> <li>• <b>Enabled</b>—<b>Adaptive Memory Training</b> is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>BIOS Tech Message Level Control</b> (for C125 M5)	<p>Enabling this token allows the BIOS Tech log output to be controlled at more a granular level. This reduces the number of BIOS Tech log messages that are redundant, or of little use. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—<b>BIOS Techlog Level</b> is disabled.</li> <li>• <b>Enabled</b>—<b>BIOS Techlog Level</b> is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>



Name	Description
<b>OptionROM Launch Optimization</b>	<p>The Option ROM launch is controlled at the PCI Slot level, and is enabled by default. In configurations that consist of a large number of network controllers and storage HBAs having Option ROMs, all the Option ROMs may get launched if the PCI Slot Option ROM Control is enabled for all. However, only a subset of controllers may be used in the boot process. When this token is enabled, Option ROMs are launched only for those controllers that are present in boot policy. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—<b>OptionROM Launch Optimization</b> is disabled.</li> <li>• <b>Enabled</b>—<b>OptionROM Launch Optimization</b> is enabled.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>BIOS Techlog Level</b>	<p>This option denotes the type of messages in <b>BIOS tech log</b> file. The log file can be any of the following types:</p> <ul style="list-style-type: none"> <li>• <b>Minimum</b>—Critical messages will be displayed in the log file. This is the default option.</li> <li>• <b>Normal</b>—Warning and loading messages will be displayed in the log file.</li> <li>• <b>Maximum</b>—Normal and information related messages will be displayed in the log file.</li> </ul>
<b>P-SATA OptionROM</b>	<p>This options allows you to select the P-SATA mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>LSI SW RAID</b>—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID. This is the default option.</li> <li>• <b>Disabled</b>—P-SATA mode is disabled.</li> <li>• <b>AHCI</b>—Sets the controllers to AHCI mode.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>M.2 SATA OptionROM</b>	<p>This options allows you to select the P-SATA mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>LSI SW RAID</b>—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID. This is the default option.</li> <li>• <b>Disabled</b>—P-SATA mode is disabled.</li> <li>• <b>AHCI</b>—Sets the controllers to AHCI mode.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>UEFI Boot Mode</b>	<p>This options allows you to select the UEFI Boot mode. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—UEFI Boot mode is disabled.</li> <li>• <b>Enabled</b>—UEFI Boot mode is enabled.</li> </ul>



**Note** BIOS parameter virtualization capability in Cisco UCS Manager maps a unified set of BIOS settings in a service profile to the actual BIOS supporting parameters. However, not all BIOS setting items are applicable to every server model/platform. When you create a custom BIOS policy and have the **Boot Option Retry** selected, and when there is no bootable option available, the reboot fails on the Cisco UCS B420 M4 servers and Cisco UCS Manager displays this message : *Reboot and Select proper Boot device or Insert Boot Media in selected Boot device and press a key*. You must manually set a boot option after the boot path is corrected, in order to enable the servers to reboot after a power outage. For more information about BIOS default server policies and the BIOS options and their default settings, see [BIOS Policy, on page 89](#) and [Server BIOS Settings, on page 1](#).

## Server Management BIOS Settings

The following tables list the server management BIOS settings that you can configure through a BIOS policy or the default BIOS settings:

## General Settings

Name	Description
<b>Assert NMI on SERR</b>	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS does not generate an NMI or log an error when a SERR occurs.</li> <li>• <b>Enabled</b>—The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable <b>Assert NMI on PERR</b>.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Assert NMI on PERR</b>	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The BIOS does not generate an NMI or log an error when a PERR occurs.</li> <li>• <b>Enabled</b>—The BIOS generates an NMI and logs an error when a PERR occurs. You must enable <b>Assert NMI on SERR</b> to use this setting.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>OS Boot Watchdog Timer Policy</b>	<p>What action the system takes if the watchdog timer expires. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Power Off</b>—The server is powered off if the watchdog timer expires during OS boot.</li> <li>• <b>Reset</b>—The server is reset if the watchdog timer expires during OS boot.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>

Name	Description
<b>OS Boot Watchdog Timer Timeout</b>	<p>What timeout value the BIOS uses to configure the watchdog timer. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>5-minutes</b>—The watchdog timer expires 5 minutes after the OS begins to boot.</li> <li>• <b>10-minutes</b>—The watchdog timer expires 10 minutes after the OS begins to boot.</li> <li>• <b>15-minutes</b>—The watchdog timer expires 15 minutes after the OS begins to boot.</li> <li>• <b>20-minutes</b>—The watchdog timer expires 20 minutes after the OS begins to boot.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>
<b>FRB-2 Timer</b>	<p>Whether the FRB-2 timer is used to recover the system if it hangs during POST. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The FRB-2 timer is not used.</li> <li>• <b>Enabled</b>—The FRB-2 timer is started during POST and used to recover the system if necessary.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

### Console Redirection Settings

Name	Description
<b>Console redirection</b>	<p>Allows a serial port to be used for console redirection during POST and BIOS booting. After the BIOS has booted and the operating system is responsible for the server, console redirection is irrelevant and has no effect. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—No console redirection occurs during POST.</li> <li>• <b>COM 0</b>—Enables serial port for console redirection during POST. This option is valid only for M6 blade servers and rack-mount servers.</li> </ul> <p><b>Note</b> The value <b>serial-port-a</b> is not supported on M6 servers.</p> <ul style="list-style-type: none"> <li>• <b>serial-port-b</b> or <b>COM 1</b>—Enables serial port B for console redirection and allows it to perform server management tasks. This option is only valid for rack-mount servers.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> If you enable this option, you also disable the display of the Quiet Boot logo screen during POST.</p>
<b>Flow Control</b>	<p>Whether a handshake protocol is used for flow control. Request to Send / Clear to Send (RTS/CTS) helps to reduce frame collisions that can be introduced by a hidden terminal problem. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>None</b>—No flow control is used.</li> <li>• <b>RTS-CTS</b>—RTS/CTS is used for flow control.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> This setting must match the setting on the remote terminal application.</p>

Name	Description
<b>Baud rate</b>	<p>What Baud rate is used for the serial port transmission speed. If you disable Console Redirection, this option is not available. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>9.6k</b>—A 9600 Baud rate is used.</li> <li>• <b>19.2k</b>—A 19200 Baud rate is used.</li> <li>• <b>38.4k</b>—A 38400 Baud rate is used.</li> <li>• <b>57.6k</b>—A 57600 Baud rate is used.</li> <li>• <b>115.2k</b>—A 115200 Baud rate is used. This is the default option.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> This setting must match the setting on the remote terminal application.</p>
<b>Terminal type</b>	<p>What type of character formatting is used for console redirection. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>PC-ANSI</b>—The PC-ANSI terminal font is used.</li> <li>• <b>VT100</b>—A supported vt100 video terminal and its character set are used.</li> <li>• <b>VT100-PLUS</b>—A supported vt100-plus video terminal and its character set are used.</li> <li>• <b>VT-UTF8</b>—A video terminal with the UTF-8 character set is used.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul> <p><b>Note</b> This setting must match the setting on the remote terminal application.</p>

Name	Description
<b>Legacy OS redirection</b>	<p>Whether redirection from a legacy operating system, such as DOS, is enabled on the serial port. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The serial port enabled for console redirection is hidden from the legacy operating system.</li> <li>• <b>Enabled</b>— The serial port enabled for console redirection is visible to the legacy operating system.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Putty KeyPad</b> <b>set console-redir-config</b> <b>putty-function-keypad</b>	<p>Allows you to change the action of the PuTTY function keys and the top row of the numeric keypad. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>VT100</b>—The function keys generate <b>ESC OP</b> through <b>ESC O[</b>.</li> <li>• <b>LINUX</b>—Mimics the Linux virtual console. Function keys F6 to F12 behave like the default mode, but F1 to F5 generate <b>ESC [ [A</b> through <b>ESC [ [E</b>.</li> <li>• <b>XTERMR6</b>—Function keys F5 to F12 behave like the default mode. Function keys F1 to F4 generate <b>ESC OP</b> through <b>ESC OS</b>, which are the sequences produced by the top row of the keypad on Digital terminals.</li> <li>• <b>SCO</b>—The function keys F1 to F12 generate <b>ESC [M</b> through <b>ESC [X</b>. The function and shift keys generate <b>ESC [Y</b> through <b>ESC [j</b>. The control and function keys generate <b>ESC [k</b> through <b>ESC [v</b>. The shift, control and function keys generate <b>ESC [w</b> through <b>ESC [ {</b>.</li> <li>• <b>ESCN</b>—The default mode. The function keys match the general behavior of Digital terminals. The function keys generate sequences such as <b>ESC [11~</b> and <b>ESC [12~</b>.</li> <li>• <b>VT400</b>—The function keys behave like the default mode. The top row of the numeric keypad generates <b>ESC OP</b> through <b>ESC OS</b>.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>

Name	Description
<b>Out of Band Management</b>	<p>Used for Windows Special Administration Control (SAC). This option allows you to configure the COM port 0 that can be used for Windows Emergency Management services. ACPI SPCR table is reported based on this setup option. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—Configures the COM port 0 as a general purpose port for use with the Windows Operating System.</li> <li>• <b>Enabled</b>—Configures the COM port 0 as a remote management port for Windows Emergency Management services.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>Redirection After BIOS POST</b>	<p>Whether BIOS console redirection should be active after BIOS POST is complete and control given to the OS bootloader. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Always Enable</b>—BIOS Legacy console redirection is active during the OS boot and run time.</li> <li>• <b>Bootloader</b>—BIOS Legacy console redirection is disabled before giving control to the OS boot loader.</li> <li>• <b>Platform Default</b>—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.</li> </ul>
<b>OS Watchdog Timer Policy</b>	<p>What action the system takes if the watchdog timer expires. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Power_Off</b>—The server is powered off if the watchdog timer expires during OS boot. This is the default option.</li> <li>• <b>Reset</b>—The server is reset if the watchdog timer expires during OS boot.</li> </ul>
<b>FRB 2 Timer</b>	<p>Whether the FRB2 timer is used for recovering the system if it hangs during POST. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The FRB2 timer is not used.</li> <li>• <b>Enabled</b>—The FRB2 timer is started during POST and used to recover the system if necessary. This is the default option.</li> </ul>



Name	Description
<b>OS Watchdog Timer</b>	<p>Whether the BIOS programs the watchdog timer with a specified timeout value. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Disabled</b>—The watchdog timer is not used to track how long the server takes to boot. This is the default option.</li> <li>• <b>Enabled</b>—The watchdog timer tracks how long the server takes to boot. This is the default option.</li> </ul>
<b>OS Watchdog Timer Timeout</b>	<p>If OS does not boot within the specified time, OS watchdog timer expires and system takes action according to timer policy. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>5 Minutes</b>—The OS watchdog timer expires 5 minutes after it begins to boot.</li> <li>• <b>10 Minutes</b>—The OS watchdog timer expires 10 minutes after it begins to boot. This is the default option.</li> <li>• <b>15 Minutes</b>—The OS watchdog timer expires 15 minutes after it begins to boot.</li> <li>• <b>20 Minutes</b>—The OS watchdog timer expires 20 minutes after it begins to boot.</li> </ul> <p><b>Note</b> This option is applicable only when you enable the OS Boot Watchdog Timer.</p>

## BIOS Policy

The BIOS policy is a policy that automates the configuration of BIOS settings for a server or group of servers. You can create global BIOS policies available to all servers in the root organization, or you can create BIOS policies in sub-organizations that are only available to that hierarchy.

To use a BIOS policy, do the following:

1. Create the BIOS policy in Cisco UCS Manager.
2. Assign the BIOS policy to one or more service profiles.
3. Associate the service profile with a server.

During service profile association, Cisco UCS Manager modifies the BIOS settings on the server to match the configuration in the BIOS policy. If you do not create and assign a BIOS policy to a service profile, the server uses the default BIOS settings for that server platform.

## Default BIOS Settings

Cisco UCS Manager includes a set of default BIOS settings for each type of server supported by Cisco UCS. The default BIOS settings are available only in the root organization and are global. Only one set of default BIOS settings can exist for each server platform supported by Cisco UCS. You can modify the default BIOS settings, but you cannot create an additional set of default BIOS settings.

Each set of default BIOS settings are designed for a particular type of supported server and are applied to all servers of that specific type which do not have a BIOS policy included in their service profiles.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

Cisco UCS Manager applies these server platform-specific BIOS settings as follows:

- The service profile associated with a server does not include a BIOS policy.
- The BIOS policy is configured with the platform-default option for a specific setting.

You can modify the default BIOS settings provided by Cisco UCS Manager. However, any changes to the default BIOS settings apply to all servers of that particular type or platform. If you want to modify the BIOS settings for only certain servers, we recommend that you use a BIOS policy.

The BIOS tokens for M5 servers and later are read-only and cannot be modified. For a complete and up to date list of BIOS tokens, defaults, and values, refer [Cisco UCS M5 Server BIOS Tokens](#).

The BIOS tokens for M6 servers and later are read-only and cannot be modified. For a complete and up to date list of BIOS tokens, defaults, and values, refer [Cisco UCS M6 Server BIOS Tokens](#).

## Creating a BIOS Policy



**Note** Cisco UCS Manager pushes BIOS configuration changes through a BIOS policy or default BIOS settings to the Cisco Integrated Management Controller (CIMC) buffer. These changes remain in the buffer and do not take effect until the server is rebooted.

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **BIOS Policies** and select **Create BIOS Policy**.
- Step 5** On the **Main** page of the **Create BIOS Policy** wizard, enter a name for the BIOS policy in the **Name** field.  
This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), \_ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
- Step 6** In the **Create BIOS Policy** wizard, do the following to configure the BIOS settings:
  - a) If you want to change a BIOS setting, click the desired radio button or make the appropriate choice from the drop-down list.

For descriptions and information about the options for each BIOS setting, see the following topics:

- **Main** page: [Main BIOS Settings, on page 2](#)
- **Advanced** page: [Main BIOS Settings, on page 2](#)
- **Processor** page: [Processor BIOS Settings, on page 4](#)
- **IO BIOS for Intel** page: [I/O BIOS Settings for Intel, on page 35](#)
- **IO BIOS for AMD** page: [I/O BIOS Settings for AMD, on page 36](#)
- **RAS Memory** page: [RAS Memory BIOS Settings, on page 38](#)
- **Serial Port** page: [Serial Port BIOS Settings, on page 50](#)
- **USB** page: [USB BIOS Settings, on page 50](#)
- **PCI Configuration** page: [PCI Configuration BIOS Settings, on page 55](#)
- **QPI** page: [QPI BIOS Settings, on page 57](#)
- **LOM and PCIe Slots** subtab: [LOM and PCIe Slots BIOS Settings, on page 59](#)
- **Trusted Platform** subtab: [Trusted Platform BIOS Settings, on page 58](#)
- **Graphics Configuration** subtab: [Graphics Configuration BIOS Settings, on page 77](#)
- **Boot Options** page: [Boot Options BIOS Settings, on page 77](#)
- **Server Management** page: [Server Management BIOS Settings, on page 82](#)

b) Click **Next** after each page.

**Step 7** After you configure all of the BIOS settings for the policy, click **Finish**.

---

## Modifying the BIOS Defaults

We recommend that you verify the support for BIOS settings in the server that you want to configure. Some settings, such as Mirroring Mode for RAS Memory, are not supported by all Cisco UCS servers.

Unless a Cisco UCS implementation has specific needs that are not met by the server-specific settings, we recommend that you use the default BIOS settings that are designed for each type of server in the Cisco UCS domain.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand **BIOS Defaults** and select the server model number or desired policy for which you want to modify the default BIOS settings.

**Step 5** In the **Work** pane, click the appropriate tab and then click the desired radio button or make a choice from the drop-down list to modify the default BIOS settings:

For descriptions and information about the options for each BIOS setting, see the following topics. Not all BIOS settings are available for each type of server.

- **Main** tab: [Main BIOS Settings, on page 2](#)
- **Advanced** tab:
  - **Processor** subtab: [Processor BIOS Settings, on page 4](#)
  - **IO BIOS for Intel** subtab: [I/O BIOS Settings for Intel, on page 35](#)
  - **IO BIOS for AMD** page: [I/O BIOS Settings for AMD, on page 36](#)
  - **RAS Memory** subtab: [RAS Memory BIOS Settings, on page 38](#)
  - **Serial Port** subtab: [Serial Port BIOS Settings, on page 50](#)
  - **USB** subtab: [USB BIOS Settings, on page 50](#)
  - **PCI Configuration** subtab: [PCI Configuration BIOS Settings, on page 55](#)
  - **QPI** subtab: [QPI BIOS Settings, on page 57](#)
  - **LOM and PCIe Slots** subtab: [LOM and PCIe Slots BIOS Settings, on page 59](#)
  - **Trusted Platform** subtab: [Trusted Platform BIOS Settings, on page 58](#)
  - **Graphics Configuration** subtab: [Graphics Configuration BIOS Settings, on page 77](#)
- **Boot Options** tab: [Boot Options BIOS Settings, on page 77](#)
- **Server Management** tab: [Server Management BIOS Settings, on page 82](#)

**Step 6** Click **Save Changes**.

---

## Viewing the Actual BIOS Settings for a Server

Follow this procedure to see the actual BIOS settings on a server.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 3** Choose the server for which you want to view the actual BIOS settings.
- Step 4** On the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Motherboard** subtab.
- Step 6** In the **BIOS Settings** area, click the **Expand** icon to the right of the heading to open that area.

Each tab in the **BIOS Settings** area displays the settings for that server platform. Some of the tabs contain subtabs with additional information.

---

## Memory RAS Features

The Intel® Xeon® processor supports additional RAS memory features via the BIOS. These features expand on the capabilities of the processor to increase the performance and reliability of memory DIMMs.

## Post-Package Repair (PPR)

Post Package Repair (PPR) allows you to use spare rows in the DRAM bank within the DDR4 DRAM to replace faulty rows detected during system boot time. Cisco UCS M5 and M6 platforms apply hard PPR. In hard PPR, the repair is permanent. The remapping of a faulty row to a spare row cannot be reverted. The remapping persists even after removal of power. If a PPR event occurs, the platform firmware generates a customer visible fault to schedule for system reboot for the repair to take effect.

The number of spare rows in the DRAM bank varies based on DIMM manufactures and models. The spare rows that are available after executing the PPR event are not visible to the platform firmware. Thereby, when all the spare rows that are available in the platform firmware visibility are utilized, the repair will not take effect and the memory errors may reoccur on the same DIMM.

## Enabling Post Package Repair

When enabled, the repair process is irrevocable.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Expand **Servers > Policies**.
  - Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
  - Step 4** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
  - Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
  - Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
  - Step 7** To enable automatic repair of faulty cell areas detected during system boot, in **Select PPR Type Configuration** select **Hard PPR**.
  - Step 8** Click **Save Changes**.
-

## Limiting Presented Memory

The amount of memory presented to the user can be limited in the BIOS. When the system is fully populated with high capacity DIMM modules, it may be desirable to reduce the amount of memory actually presented for use.

The memory limit will be applied evenly across all installed and available DIMMs to the extent possible. The minimum amount of presented memory you can specify is 1 GB. The following parameters apply:

0 = No limit. Full amount of installed memory is presented.

1 to  $2^{31} - 1$  = Size of presented memory in gigabytes (GB)

Actual presented memory size will always be equal or less than specified memory size.

## Limiting Memory Size

Actual presented memory size will always be equal or less than specified memory size.

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Expand **Servers > Policies**.
  - Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
  - Step 4** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
  - Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
  - Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
  - Step 7** To limit the amount of presented memory to be mirrored, go to **Memory Size Limit in GB** and enter a value (in GB) for the desired amount of memory to be presented to the user.
  - Step 8** Click **Save Changes**.
- 

## Partial Memory Mirroring

Partial Memory Mirroring if DIMMs is an advanced RAS feature. Only Gold and Platinum SKU CPUs support this feature

Partial DIMM Mirroring creates a mirrored copy of a specific region of memory cells, rather than keeping the complete mirror copy. Partial Memory Mirroring can be performed from either BIOS policy setup menu or from the Linux Operating System. Partial Mirroring creates a mirrored region in memory map with the attributes of a partial mirror copy. Up to 50% of the total memory capacity can be mirrored, using up to 4 partial mirrors

For mirroring, at least two DDR channels must be populated in each IMC. Partial mirroring supports one DDR4 mirror region per IMC, with a maximum of four mirror regions.

In a two-way channel interleave, two channels are populated in each IMC. In a three-way channel interleave, three channels are populated in each IMC.

Partial mirroring is incompatible with rank sparing and ADDDC.

The following rules apply to partial mirroring:

- The DIMM population must be identical for the mirrored channels.
- The mirror pair must be in the same M2M, within an IMC DDR channel.
- DDR4 partial mirror regions within one iMC must be either two-way channel interleaves or three-way channel interleave. Two and three-way channel interleaves cannot be mixed. When the mirror region spans across iMCs, the channel interleaves must be the same.

## Enabling Partial Memory Mirroring

The amount of partial DIMM memory mirroring can be configured either in percentage of available memory resources or in gigabytes..

### Before you begin




---

**Note** Partial Memory Mirror Mode is mutually exclusive to standard Mirror Mode.

---

Partial Mirroring is incompatible with rank sparing and ADDDC. Verify that these are not selected.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** In the Policies section, right-click the BIOS Policy section and select **Create BIOS Policy** from the popup. In the BIOS Policy form, enter a name and optional description. Click **OK** to create the policy.
- Step 5** Go to **Policies > Root > BIOS Policies** and select the new policy.
- Step 6** In the main work pane, select the **Advanced** tab, then select the **RAS Memory** tab.
- Step 7** Go to **Memory RAS Configuration** and select **Partial Mirror Mode 1LM** from the dropdown list.
- Step 8** To configure the partial mirror in percentage, go to **Partial Memory Mirror Mode** and select **Percentage** from the dropdown.
- Step 9** Go to **Partial Mirror percentage** and type a value between 0.01 and 50.00, representing the desired percentage of memory to be mirrored.
- Step 10** To configure the partial mirror in gigabytes, go to **Partial Memory Mirror Mode** and select **Value in GB** from the dropdown.
- Step 11** Enter a value between 1 and GB of memory displayed in the limits field for **Partial Mirror 1**.
- Step 12** If desired, enter additional values into **Partial Mirror 2**, **Partial Mirror 3**, and **Partial Mirror 4**. The total of values entered into these mirrors cannot exceed the total memory available.
- Step 13** Click **Save Changes**.
-

**What to do next**

Reboot the system.

# Trusted Platform Module

## Trusted Platform Module

The Trusted Platform Module (TPM) is a component that can securely store artifacts that are used to authenticate the server. These artifacts can include passwords, certificates, or encryption keys. A TPM can also be used to store platform measurements that help ensure that the platform remains trustworthy. Authentication (ensuring that the platform can prove that it is what it claims to be) and attestation (a process helping to prove that a platform is trustworthy and has not been breached) are necessary steps to ensure safer computing in all environments. It is a requirement for the Intel Trusted Execution Technology (TXT) security feature, which must be enabled in the BIOS settings for a server equipped with a TPM. Cisco UCS M4 and higher blade and rack-mount servers include support for TPM. TPM is enabled by default on these servers.

**Important**

- If you upgrade Cisco UCS Manager to Release 2.2(4) and higher, TPM is enabled.
- When TPM is enabled and you downgrade Cisco UCS Manager from Release 2.2(4), TPM is disabled.

## Intel Trusted Execution Technology

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, invisible 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. Cisco UCS M4 and higher blade and rack-mount servers include support for TXT. TXT is disabled by default on these servers.

TXT can be enabled only after TPM, Intel Virtualization technology (VT) and Intel Virtualization Technology for Directed I/O (VT-d) are enabled. When you only enable TXT, it also implicitly enables TPM, VT, and VT-d.

## Configuring Trusted Platform

**Procedure**

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to configure TPM.
- Step 4** Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.



**Step 5** In the **Work** pane, click the **Advanced** tab.

**Step 6** Click the **Trusted Platform** subtab.

**Step 7** To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enabled	Enables TPM
<b>Platform Default</b>	Enables TPM

**Step 8** To configure TXT, click one of the following:

Option	Description
disabled	Disables TXT
enabled	Enables TXT
<b>Platform Default</b>	Disables TXT

**Step 9** Click **Save Changes**.

## Configuring Trusted Platform

### Procedure

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

**Step 2** Expand **Servers > Policies**.

**Step 3** Expand the node for the organization where you want to configure TPM.

**Step 4** Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.

**Step 5** In the **Work** pane, click the **Advanced** tab.

**Step 6** Click the **Trusted Platform** subtab.

**Step 7** To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enable	Enables TPM
<b>Platform Default</b>	Enables TPM

**Step 8** To configure TXT, click one of the following:

Option	Description
disabled	Disables TXT

Option	Description
enable	Enables TXT
Platform Default	Disables TXT

**Step 9** Click **Save Changes**.

## Viewing TPM Properties

### Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > **Chassis Number** > **Cartridges** > **Cartridge Number** > **Servers**
- Step 3** Choose the server for which you want to view the TPM settings.
- Step 4** On the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Motherboard** subtab.

## SPDM Security

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



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

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

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

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

- Full Security:

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

- Partial Security (default):

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

- No Security

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

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

Certificates are listed in all user interfaces on a system.

## Creating a SPDM Security Policy

This step creates a SPDM policy.




---

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

---

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Go to **Policies**. Expand the root node.
  - Step 3** Right-click **SPDM Certificate Policies** and select **Create SPDM Policy**.
  - Step 4** Enter a name for this policy and select a **Fault Alert Setting** for the security level: **Disabled**, **Partial**, or **Full**.
    - Full—If you select this option, then a fault is generated when there is any endpoint authentication failure for both supported and unsupported endpoints.
    - Partial—If you select this option then a fault is generated when there is any endpoint authentication failure to only supported endpoints. No fault is generated when the endpoint does not support authentication.
    - Disabled—If you select this option then no fault is generated for endpoint authentication failure for both supported and unsupported endpoints.
- The default is **Partial**.

**Note** To perform SPDM re-authentication and update the faults, Cisco IMC reboot or host reboot is required when the fault alert value is changed for an associated profile.

- Step 5** Click on **Add** in the **Create Policy** window. The **Add SPDM Certificate** window will open.
- Step 6** Name the certificate.  
UCS Manager supports only **Pem**certificates.
- Step 7** Paste the contents of the certificate into the Certificate field.
- Step 8** Click **OK** to add the certificate and return to the **Create SPDM Policy** window.  
You can add up to 40 certificates.
- Step 9** In the **Create SPDM Policy** menu, click **Okay**.  
After the SPDM policy is created, it will be listed immediately, along with its Alert setting, when you select **SPDM Certificate Policy** under the Server root Policies.
- 

#### What to do next

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

## Associating the Security Policy with a Server

#### Before you begin

Create the SPDM security policy.

#### Procedure

---

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

#### What to do next

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

## Viewing the Fault Alert Settings

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

### Before you begin

Create a policy and associate it with a Service Profile.

### Procedure

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

## Consistent Device Naming

When there is no mechanism for the Operating System to label Ethernet interfaces in a consistent manner, it becomes difficult to manage network connections with server configuration changes. Consistent Device Naming (CDN), introduced in Cisco UCS Manager Release 2.2(4), allows Ethernet interfaces to be named in a consistent manner. This makes Ethernet interface names more persistent when adapter or other configuration changes are made.

To configure CDN for a vNIC, do the following:

- Enable consistent device naming in the BIOS policy.
- Associate the BIOS policy with a service profile.
- Configure consistent naming for a vNIC.

## Guidelines and Limitations for Consistent Device Naming (CDN)

- CDN is supported on the following Operating Systems:
  - Windows 2016 and later Windows releases
  - Windows Server 2019
  - Red Hat Enterprise Linux (RHEL) 7.x and later RHEL releases
  - SLES 12 SP3, SLES 12 SP4, and SLES 15 (for 4.0(4a) and later)
  - ESXi 6.7
- Consistent device naming (CDN) is supported on all M4 and higher blade and rack-mount servers.
- BIOS and adapter firmware must be part of the Release 2.2(4) or higher bundle to support CDN.
- If the RHEL Operating System is installed on the server, CDN will appear when running the command "**biosdevname -d**" as "**sysfs label**". CDN will not change the kernel name.
- CDN is supported for vNIC template.

- Multiple vNICs within the same service profile cannot have the same CDN name.
- When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.
- The CDN name that you configure for a vNIC appears as **Admin CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **Admin CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **Admin CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.
- In Cisco UCS Manager Release 3.1 and older releases, downgrade of the adapter firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- In Cisco UCS Manager Release 2.2(4), downgrade of Cisco UCS Manager or BIOS is prevented, if CDN enabled BIOS policy is assigned on the associated server profile.
- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled, the host reboots with a warning, irrespective of whether reboot on BIOS update is enabled or not.
- It is recommended that you enable CDN in the BIOS policy and add CDN names to the vNICs before the Windows Operating System is installed.
- If the Windows Operating System is already installed on the server and CDN is then enabled in the BIOS policy, do the following:
  1. Uninstall the network drivers.
  2. Scan the system for hidden devices and uninstall them.
  3. Rescan the system for new hardware and install the network drivers again.




---

**Note** If this is not done, the vNICs will not come up with the configured CDN names.

---

- When the applied BIOS policy is changed from CDN-disabled to CDN-enabled or from CDN-enabled to CDN-disabled on a service profile, do the following:
  1. Uninstall the network drivers.
  2. Scan the system for hidden devices and delete them.
  3. Re-scan the system for new hardware and install the network drivers again.




---

**Note** When the BIOS policy is changed from CDN-enabled to CDN-disabled, ensure that the CDN names are removed from all the vNICs on the system.

---

- If any change is made to the vNICs, the BDF of all the devices on the system also changes. Following are some of the scenarios that trigger a change in the BDF of all the vNICs present on the system:
  - When a vNIC is added or deleted
  - When a vNIC is moved from one adapter on the system to another adapter on the system

When these changes are made to the system, do the following:

1. Uninstall the network driver from all the present network interfaces.
2. Scan the system for hidden devices and uninstall them.
3. Re-scan the system for new hardware and install the network driver on the network controllers again.

If the hidden devices are not deleted, the CDN names of the network adapters will not appear as configured on Cisco UCS Manager.

### CDN with a Mixed Set of Adapters

When a CDN name is configured for a vNIC in a system with a mixed set of CDN-supported adapters and CDN-unsupported adapters, then system placement may not place CDN-configured vNICs on adapters that support CDN.

If CDN is enabled in the BIOS policy, and system placement places a CDN-configured vNIC (Admin CDN configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored.

If CDN is enabled in the BIOS policy, and system placement places a vNIC (Admin CDN not configured) on an adapter that does not support CDN, an info fault will be raised, but the configuration issue for the service profile will be ignored. The **Oper CDN Name** in this case will be empty and will not be derived from the vNIC name.

If you want to deploy the CDN name as the host network interface name for a server, you must manually place a vNIC on a supported adapter.

## Configuring Consistent Device Naming in a BIOS Policy

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand **root**.
- Step 4** Expand **BIOS Policies**.
- Step 5** Select the BIOS policy for which you want to configure CDN.

**Note** Because the default BIOS policy does not store the CDN values on B200 M6 servers, the UCS Manager does not push the custom CDN values to vNIC. To configure the CDN values for the BIOS policy, create a BIOS policy with the required values that includes a CDN value.

- Step 6** Under the **Main** tab, click one of the following in the **Consistent Device Naming** field to configure CDN:

Option	Description
<b>disabled</b>	Disables CDN in the BIOS policy
<b>enabled</b>	Enables CDN in the BIOS policy
<b>Platform Default</b>	The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

**Step 7** Click **Save Changes**.

---

## Configuring a CDN Name for a vNIC

When a CDN name is not specified for a vNIC, the vNIC name is used as the CDN name.

### Procedure

---

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

**Step 2** Expand **Servers > Service Profiles**.

**Step 3** Expand the node for the organization that contains the vNIC for which you want to configure a CDN name.

**Step 4** Expand the service profile and **vNICs** node that contain the vNIC for which you want to configure a CDN name.

**Step 5** Select the vNIC.

**Step 6** Click on the **General** tab.

**Step 7** In the **Properties** area, choose **User Defined** as the **CDN Source**.

**Step 8** Enter the CDN name for the vNIC in the **CDN Name** field.

**Note** The CDN name that you configure for a vNIC appears as **CDN Name**. The CDN name that is finally applied to the vNIC appears as **Oper CDN Name**. For example, if the **CDN Name** for a vNIC called "vnic0" is cdn0, then the **Oper CDN Name** for this vNIC will be cdn0, but if the **CDN Name** for the same vNIC is not specified, the **Oper CDN Name** will be vnic0.

**Step 9** Click **Save Changes**.

---

## CIMC Security Policies

Cisco UCS Manager provides the following policies to increase security:

- KVM Management Policy
- IPMI Access Profile

### IPMI Access Profile

This policy allows you to determine whether IPMI commands can be sent directly to the server, using the IP address. For example, you can send commands to retrieve sensor data from the CIMC. This policy defines the IPMI access, including a username and password that can be authenticated locally on the server, and whether the access is read-only or read-write.

You can also restrict remote connectivity by disabling or enabling IPMI over LAN in the IPMI access profile. IPMI over LAN is disabled by default on all unassociated servers, and on all servers without an IPMI access policy. When an IPMI access policy is created, the IPMI over LAN is set to enabled by default. If you do not change the value to disabled, IPMI over LAN will be enabled on all associated servers.



You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

## Creating an IPMI Access Profile

### Before you begin

An IPMI profile requires that one or more of the following resources already exist in the system:

- Username with appropriate permissions that can be authenticated by the operating system of the server
- Password for the username
- Permissions associated with the username

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **IPMI Access Profiles** and select **Create IPMI Access Profile**.
- Step 5** In the **Create IPMI Access Profile** dialog box:
- Enter a unique name and description for the profile.
  - In the **IPMI Over LAN** field, choose whether to allow or restrict remote connectivity.
  - Click **OK**.
- Step 6** In the **IPMI Users** area of the navigator, click +.
- Step 7** In the **Create IPMI User** dialog box:
- Complete the following fields:

Name	Description
<b>Name</b> field	The username to associate with this IPMI or Redfish profile. Enter 1 to 16 alphanumeric characters. You can also use @ (at sign), _ (underscore), and - (hyphen). You cannot change this name once the profile has been saved.
<b>Password</b> field	The password associated with this username. Enter 1 to 20 standard ASCII characters, except for = (equal sign), \$ (dollar sign), and   (vertical bar).
<b>Confirm Password</b> field	The password a second time for confirmation purposes.

Name	Description
<b>Role</b> field	The user role. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Admin</b></li> <li>• <b>Read Only</b></li> </ul>
<b>Description</b> field	User-defined description of the IPMI or Redfish user.

b) Click **OK**.

**Step 8** Repeat Steps 6 and 7 to add another user.

**Step 9** Click **OK** to return to the IPMI profiles in the **Work** pane.

#### What to do next

Include the IPMI profile in a service profile and/or template.

## Deleting an IPMI Access Profile

#### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **IPMI Profiles** node.
- Step 4** Right-click the profile you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

## KVM Management Policy

The KVM Management policy allows you to determine whether vMedia encryption is enabled when you access a server via KVM.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.



**Note** After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.

Before Cisco UCS Manager Release 4.0(4), port 2068 was the only KVM port. Beginning with Release 4.0(4), you can configure a port number between 1024 and 49151 as the KVM port. Port 2068 continues to be the default KVM port number.

## Creating a KVM Management Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **KVM Management Policies** and select **Create KVM Management Policy**.
- Step 5** In the **Create KVM Management Policy** dialog box:
- Enter a unique name and description for the policy.
  - In the **vMedia Encryption** field, choose whether to enable vMedia encryption.  
**Note** Starting with UCS Manager 4.2, vMedia Encryption is always enabled for security purposes. It cannot be modified by the user.
  - In the **KVM Port** field, enter a port number between 1024 and 49151 for KVM.  
The default KVM port number is 2068.
  - Click **OK**.
- Note** After a KVM vMedia session is mapped, if you change the KVM management policy, it will result in a loss of the vMedia session. You must re-map the KVM vMedia session again.
- 

## Graphics Card Policies

Cisco UCS Manager Release 3.1(3) extends graphics card support to include the ability to change the graphics card mode. You can now configure graphics card modes by using a graphics card policy. The graphics card modes are:

- Compute
- Graphics
- Any Configuration

## Creating a Graphics Card Policy



- Note** Cisco UCS Manager pushes the configuration changes to the GPU through the Graphics Card policy to the Processor Node Utility Operating System (PNuOS). These changes do not take effect until the server is rebooted.
-

## Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Graphics Card Policies** and select **Create Graphics Card Policy**.
- Step 5** On the **Main** page of the **Create Graphics Card Policy** dialog box:
- Enter a unique name for the policy.
  - (Optional) Enter a description for the policy.
  - In the **Graphics Card Mode** field, choose one of the options:
    - **Compute**
    - **Graphics**
    - **Any Configuration**
  - Click **OK**.
- 

# Local Disk Policies

## Local Disk Configuration Policy

This policy configures any optional SAS local drives that have been installed on a server through the onboard RAID controller of the local drive. This policy enables you to set a local disk mode for all servers that are associated with a service profile that includes the local disk configuration policy.

The local disk modes include the following:

- **No Local Storage**—For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.
- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.

If you choose **No RAID** and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the **No RAID** mode. This can lead to a mismatch between

the RAID configuration in the policy and the actual disk configuration shown in the **Inventory > Storage** tab for the server.

To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the **No RAID** configuration mode.

- **RAID 5 Striped Parity**—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.
- **RAID 6 Striped Dual Parity**—Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.
- **RAID 10 Mirrored and Striped**—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates.
- **RAID 50 Striped Parity and Striped**—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- **RAID 60 Striped Dual Parity and Striped**—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

You must include this policy in a service profile and that service profile must be associated with a server for the policy to take effect.



---

**Note** For a Cisco UCS C-Series server integrated with Cisco UCS Manager, with an embedded on-board RAID controller, the local disk mode should always be **Any Configuration**, and the RAID must be configured directly on the controller.

---

## Guidelines for all Local Disk Configuration Policies

Before you create a local disk configuration policy, consider the following guidelines:

### No Mixed HDDs and SSDs

Do not include HDDs and SSDs in a single server or RAID configuration.

## Guidelines for Local Disk Configuration Policies Configured for RAID

### Configure RAID Settings in Local Disk Configuration Policy for Servers with MegaRAID Storage Controllers

If a blade server or integrated rack-mount server has a MegaRAID controller, you must configure RAID settings for the drives in the Local Disk Configuration policy included in the service profile for that server. You can do this either by configuring the local disk configuration policy in the service profile using one of the defined RAID modes for that server, or you can use the **Any Configuration** mode with the LSI Utilities toolset to create the RAID volumes.

If you do not configure your RAID LUNs before installing the OS, disk discovery failures might occur during the installation and you might see error messages such as “No Device Found.”

### Server May Not Boot After RAID1 Cluster Migration if Any Configuration Mode Specified in Service Profile

After RAID1 clusters are migrated, you need to associate a service profile with the server. If the local disk configuration policy in the service profile is configured with **Any Configuration** mode rather than **RAID1**, the RAID LUN remains in "inactive" state during and after association. As a result, the server cannot boot.

To avoid this issue, ensure that the service profile you associate with the server contains the identical local disk configuration policy as the original service profile before the migration and does not include the **Any Configuration** mode.

### Do Not Use JBOD Mode on Servers with MegaRAID Storage Controllers

Do not configure or use JBOD mode or JBOD operations on any blade server or integrated rack-mount server with a MegaRAID storage controllers. JBOD mode and operations are not intended for nor are they fully functional on these servers.

### Maximum of One RAID Volume and One RAID Controller in Integrated Rack-Mount Servers

A rack-mount server that has been integrated with Cisco UCS Manager can have a maximum of one RAID volume irrespective of how many hard drives are present on the server.

All the local hard drives in an integrated rack-mount server must be connected to only one RAID Controller. Integration with Cisco UCS Manager does not support the connection of local hard drives to multiple RAID Controllers in a single rack-mount server. We therefore recommend that you request a single RAID Controller configuration when you order rack-mount servers to be integrated with Cisco UCS Manager.

In addition, do not use third party tools to create multiple RAID LUNs on rack-mount servers. Cisco UCS Manager does not support that configuration.

### Maximum of One RAID Volume and One RAID Controller in Blade Servers

A blade server can have a maximum of one RAID volume irrespective of how many drives are present in the server. All the local hard drives must be connected to only one RAID controller.

In addition, do not use third party tools to create multiple RAID LUNs on blade servers. Cisco UCS Manager does not support that configuration.

### License Required for Certain RAID Configuration Options on Some Servers

Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.

For RAID license information for a specific Cisco UCS server, see the *Hardware Installation Guide* for that server.

## Creating a Local Disk Configuration Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.

**Step 3** Expand the node for the organization where you want to create the policy.

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

**Step 4** Right-click **Local Disk Config Policies** and choose **Create Local Disk Configuration Policy**.

**Step 5** In the **Create Local Disk Configuration Policy** dialog box, complete the following fields:

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

Name	Description
Mode drop-down list	<p>This can be one of the following local disk policy modes:</p> <ul style="list-style-type: none"> <li>• <b>No Local Storage</b></li> <li>• <b>RAID 0 Striped</b></li> <li>• <b>RAID 1 Mirrored</b></li> <li>• <b>Any Configuration</b></li> <li>• <b>No RAID</b></li> </ul> <p>If you choose <b>No RAID</b> and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the <b>No RAID</b> mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the <b>Inventory &gt; Storage</b> tab for the server.</p> <p>To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the <b>No RAID</b> configuration mode.</p> <ul style="list-style-type: none"> <li>• <b>RAID 5 Striped Parity</b></li> <li>• <b>RAID 6 Striped Dual Parity</b></li> <li>• <b>RAID 10 Mirrored and Striped</b></li> <li>• <b>RAID 50 Striped Parity and Striped</b></li> <li>• <b>RAID 60 Striped Dual Parity and Striped</b></li> </ul> <p><b>Note</b> Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the <i>Hardware Installation Guide</i> for that server.</p>



Name	Description
<b>Protect Configuration</b> check box	<p>If checked, the server retains the configuration in the local disk configuration policy even if the server is disassociated from the service profile.</p> <p><b>Caution</b> Protect Configuration becomes non-functional if one or more disks in the server are defective or faulty.</p> <p>This property is checked by default.</p> <p>When a service profile is disassociated from a server and a new service profile associated, the setting for the Protect Configuration property in the new service profile takes precedence and overwrites the setting in the previous service profile.</p> <p>With this option enabled, the data on the disk is protected even after the server is decommissioned and then recommissioned. Hence, reassociation of the server with a service profile fails.</p> <p><b>Note</b> If you disassociate the server from a service profile with this option enabled and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>
<b>FlexFlash State</b> radio button	<p>To enable or disable the FlexFlash controller on the SD card, click the appropriate button.</p> <p><b>Note</b> This parameter only applies to a server with an SD card module.</p>
<b>FlexFlash RAID Reporting State</b> radio button	<p>To enable or disable RAID reporting, click the appropriate button. When RAID reporting is enabled, the RAID status is monitored and faults are enabled.</p> <p><b>Note</b> If only one SD card is installed, the RAID state will be displayed as Disabled and the RAID health as NA even if RAID reporting is enabled.</p>
<b>FlexFlash Removable State</b> radio button	<p>To select the removable state of the FlexFlash SD card, click the appropriate button.</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—Use this option to define the SD card as removable.</li> <li>• <b>No</b>—Use this option to define the SD card as fixed or non-removable.</li> <li>• <b>No Change</b>—Use this option if the hypervisor does not require a preset state for the SD card.</li> </ul>

**Step 6** Click **OK**.

## Changing a Local Disk Configuration Policy

This procedure describes how to change a local disk configuration policy from an associated service profile. You can also change a local disk configuration policy from the **Policies** node from **Servers**.

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the organization that includes the service profile with the local disk configuration policy you want to change.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Click the service profile that contains the local disk configuration policy you want to change.
- Step 5** In the **Work** pane, click the **Storage** tab.
- Step 6** In the **Actions** area, click **Change Local Disk Configuration Policy**.
- Step 7** In the **Change Local Disk Configuration Policy** dialog box, choose one of the following options from the **Select the Local Disk Configuration Policy** drop-down list.

Option	Description
<b>Use a Disk Policy</b>	Select an existing local disk configuration policy from the list below this option. Cisco UCS Manager assigns this policy to the service profile.
<b>Create a Local Disk Policy</b>	Enables you to create a local disk configuration policy that can only be accessed by the selected service profile.
<b>No Disk Policy</b>	Selects the default local disk policy.  <b>Note</b> If a UCS server is attached to the Cisco UCS Manager, selecting the No Disk Policy can erase and replace the RAID with individual RAID 0 disks if the default RAID configuration is not supported on the attached server.

- Step 8** Click **OK**.
- Step 9** (Optional) Expand the **Local Disk Configuration Policy** area to confirm that the change has been made.
- 

## Deleting a Local Disk Configuration Policy

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **Local Disk Config Policies** node.

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

## FlexFlash Support

### Overview

Cisco UCS B-Series, C-Series M4 and higher, and S-Series M4 servers support internal Secure Digital (SD) memory cards. The SD cards are hosted by the Cisco Flexible Flash storage controller, a PCI-based controller which has two slots for SD cards. The cards contain a single partition called HV. When FlexFlash is enabled, Cisco UCS Manager displays the HV partition as a USB drive to both the BIOS and the host operating system.


You can populate one or both the SD card slots that are provided. If two SD cards are populated, you can use them in a mirrored mode.




**Note** Do not mix different capacity cards in the same server.

The SD cards can be used to store operating system boot images or other information. The following figure illustrates the SD card slots.

**Figure 1: SD Card Slots**





FlexFlash is disabled by default. You can enable FlexFlash in a local disk policy used in a service profile. When FlexFlash is enabled in a local disk policy, and the server is capable of supporting SD cards, the FlexFlash controller is enabled during service profile association. If a server is not capable of supporting SD cards or has an older CIMC version, a config failure message is displayed.

If you disable FlexFlash in a supported server, the Hypervisor or HV partition is immediately disconnected from the host. The FlexFlash controller will also be disabled as part of a related service profile disassociation.

The FlexFlash controller supports RAID-1 for dual SD cards. The FlexFlash scrub policy erases the HV partition in both cards, and brings the cards to a healthy RAID state.

You can configure new SD cards in a RAID pair and format them using one of the following methods:

- Format the SD cards. [Formatting the SD Cards, on page 120](#) provides detailed information.
- Disassociate the service profile from the server. Re-acknowledge the server after modifying the default scrub policy and then associate the server profile back to the server.

The *Scrub Policy Settings* section in the *Cisco UCS Manager Server Management Guide* provides more details about the usage of the scrub policy.




---

**Note** Disable the scrub policy as soon as the pairing is complete.

---

To boot from the HV partition, the SD card must be present in the boot policy used in the service profile.

### FlexFlash Firmware Management

The FlexFlash controller firmware is bundled as part of the CIMC image. When you upgrade the CIMC, if a newer firmware version is available for the FlexFlash controller, the controller can no longer be managed, and the FlexFlash inventory displays the **Controller State** as **Waiting For User Action** and the **Controller Health** as **Old Firmware Running**. To upgrade the FlexFlash controller firmware, you need to perform a board controller update. For more information, see the appropriate *Cisco UCS B-Series Firmware Management Guide*, available at the following URL:

[http://www.cisco.com/en/US/products/ps10281/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html).

**Limitations for the Cisco Flexible Flash Storage Controller:**

- The Cisco Flexible Flash storage controller only supports 16 GB, 32 GB, and 64 GB SD cards.



---

**Note** The 64 GB SD cards are supported only on the M4M5 blade servers.

---

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- Some Cisco UCS C-Series rack-mount servers have SD cards with four partitions: HV, HUU, SCU, and Drivers. Only the HV partition is visible in Cisco UCS Manager. You can migrate a four-partition SD card to a single HV partition card with a FlexFlash scrub policy but there may be data loss.
- The FlexFlash controller does not support RAID-1 sync (mirror rebuild). If the SD cards are in a degraded RAID state, or if any metadata errors are reported by the controller, you must run the FlexFlash scrub policy to pair the cards for RAID. For more information about the FlexFlash scrub policy, see [Server-Related Policies](#). The following conditions might result in degraded RAID or metadata errors:
  - Inserting a new or used SD card in one slot, when the server already has an SD card populated in the second slot.
  - Inserting two SD cards from different servers.
- The server firmware version must be at 2.2(1a) or higher.

**FlexFlash FX3S Support**

Beginning with Release 2.2(3), Cisco UCS Manager allows additional FlexFlash support with the FX3S controller. The FX3S controller is present on the following servers:

- Cisco UCS B200 M4 and M5 blade server
- Cisco UCS C220 M4 and M5 rack server
- Cisco UCS C240 M4 and M5 rack server
- C480 M5 rack server
- C480 M5 ML blade server
- B480 M5 blade server
- Cisco UCS C125 M5 Server

FlexFlash operations with the FX3S control are similar to those with the Cisco Flexible Flash storage controller. FlexFlash is disabled by default, and is enabled using a local disk policy. You can also reset the controller, format the SD cards, and enable automatic synchronization of your paired SD cards.

The SD cards for the FX3S controller contain a single partition called Hypervisor.

**Limitations for the Cisco FX3S Controller:**

- The FX3S controller supports only 32 GB and 64 GB SD cards. 16 GB cards are not supported.

- The FX3S controller supports 128 GB cards on M5 blades and above.
- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- The server firmware version must be at 2.2(3a) or higher.

## Starting Up Blade Servers with FlexFlash SD Cards

Use this procedure to start up blade servers using FlexFlash cards 16 GB and larger. This procedure requires that you know how to setup the blade server, software, and the associated infrastructure, and ensure that they are working. This Cisco UCS Manager controlled procedure is applicable to all blade servers, running any version of firmware. This procedure does not apply to rack servers. Follow this procedure before you enable FlexFlash cards in a working environment.




---

**Caution** If you use the following procedure with FlexFlash cards already in use, you will lose all data from the cards.

---




---

**Note** This procedure does not cover FlexFlash card usage or other functions of the FlexFlash system.

---

### Procedure

- 
- Step 1** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 2** In the **Work** pane, check the details of the FlexFlash cards in the **FlexFlash Controller** window.
- Step 3** Expand **Servers** > **Service Profiles**.
- Step 4** Expand the node for the organization containing the pool.  
If the system does not include multi tenancy, expand the **root** node.
- Step 5** Expand the node for the organization containing the service profile and click **Storage**.
- Step 6** In the **Work** pane, click **Change Local Disk Configuration Policy** in the **Actions** area and expand **Create Local Disk Configuration Policy** link. Follow the procedure in [Creating a Local Disk Configuration Policy, on page 110](#) to create a Local Disk Configuration Policy.  
The FlexFlash policy name must not contain empty spaces or special characters.
- Step 7** Expand **Change Disk Local Configuration Policy**, and select the policy you just created and click **OK**.
- Step 8** Expand **Servers** > **Policies**.
- Step 9** Follow the procedure in [Creating a Scrub Policy, on page 123](#) and create a policy with a name such as *Scrub-FF-name* and click **OK**.  
The Scrub policy name must not contain empty spaces or special characters.
- Step 10** Select the policy you created from the drop-down box.
- Step 11** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
- Step 12** In the **Work** pane, click the **General** tab and select **Server Maintenance** from the **Actions** area.

- Step 13** In the **Maintenance Server** dialogue box, click on the **Re-acknowledge** radio button, and then click **OK**.
- Step 14** Click **Server Maintenance** in the **Action** area and click on the **Re-acknowledge** radio button again.
- Step 15** From the **Inventory** tab, select the **Storage** sub-tab.  
You can verify details of the enabled FlexFlash cards from the **FlexFlash Controller** window in the **Work** area.
- Step 16** Launch KVM Manager and log on to the operating system. Verify details of the Hypervisor partition from the Devices and drives folder. Depending on the card size, the HV partition displays details of 32GB, 64GB, or 128 GB.  
The FlexFlash cards are now synced and ready to use.
- 

## Enabling FlexFlash SD Card Support

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand **Local Disk Config Policies** and choose the local disk config policy for which you want to enable FlexFlash support.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** In the **FlexFlash State** field, click the **Enable** radio button.
- Step 7** In the **FlexFlash RAID Reporting State** field, click the **Enable** radio button.
- Step 8** Click **Save Changes**.
- 

## Enabling Auto-Sync

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Click the server for which you want to enable auto-sync.
- Step 4** In the **Work** pane, click the **Inventory** tab.
- Step 5** Click the **Storage** subtab.
- Step 6** In the **Actions** area, click **Enable Auto-sync**.
- Step 7** In the **Enable Auto-sync** dialog box, choose the **Admin Slot Number** for the SD card that you want to use as the primary.
- Step 8** Click **OK**.
-

## Formatting the SD Cards

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
  - Step 3** Click the server for which you want to format the SD cards.
  - Step 4** In the **Work** pane, click the **Inventory** tab.
  - Step 5** Click the **Storage** subtab.
  - Step 6** In the **Actions** area, click **Format SD Cards**.
  - Step 7** Click **Yes** to format the SD cards.
- 

## Resetting the FlexFlash Controller

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
  - Step 3** Click the server for which you want to reset the FlexFlash controller.
  - Step 4** In the **Work** pane, click the **Inventory** tab.
  - Step 5** Click the **Storage** subtab.
  - Step 6** In the **Actions** area, click **Reset FlexFlash Controller**.
  - Step 7** Click **Yes** to reset the FlexFlash controller.
- 

## Persistent Memory Modules

Cisco UCS Manager Release 4.0(4) introduces support for the Intel<sup>®</sup> Optane<sup>™</sup> Data Center persistent memory modules on the UCS M5 servers that are based on the Second Generation Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors. Starting with Cisco UCS Manager Release 4.2, the support for the Intel<sup>®</sup> Optane<sup>™</sup> Data Center persistent memory modules on the UCS M6 servers that are based on the Second Generation Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors are also provided. These persistent memory modules can be used only with the Second Generation Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors.

Persistent memory modules are non-volatile memory modules that bring together the low latency of memory and the persistence of storage. Data stored in persistent memory modules can be accessed quickly compared to other storage devices, and is retained across power cycles.

For detailed information about configuring persistent memory modules, see the *Cisco UCS: Configuring and Managing Intel<sup>®</sup> Optane<sup>™</sup> Data Center Persistent Memory Modules Guide*.



# Scrub Policy

## Scrub Policy Settings

This policy determines what happens to local data and to the BIOS settings on a server during the discovery process, when the server is re-acknowledged, or when the server is disassociated from a service profile.



---

**Note** Local disk scrub policies only apply to hard drives that are managed by Cisco UCS Manager and do not apply to other devices such as USB drives.

---

Depending upon how you configure a scrub policy, the following can occur at those times:

### Disk scrub

One of the following occurs to the data on any local drives on disassociation:

- If enabled, deletes initial 200MB of data from master boot record or the boot sectors. Thus, preventing the system to boot from an already installed OS if any. For secure deletion of data on drives, refer [UCS Secure Data Deletion For Commission Regulation \(EU\) 2019 /424 Users Guide](#).



---

**Note** Though the disk scrub policy is not intended to delete the user data that exceeds 200MB, Cisco UCS Manager cannot guarantee against data loss.

---

- If disabled (default), preserves all data on any local drives, including local storage configuration.

For a server associated with a service profile, disk scrub occurs during disassociation, based on the scrub policy used in the service profile. For an un-associated server, disk scrub occurs during the server discovery process, based on the default scrub policy.

Scrub policies are supported on all B-Series platforms and only on the following C-Series platforms:

- Cisco UCS C220 M4 Server
- Cisco UCS C240 M4 Server
- Cisco UCS C460 M4 Storage Server
- Cisco UCS C220 M5 Server
- Cisco UCS C240 M5 Server
- Cisco UCS C480 M5 Server
- Cisco UCS C220 M6 Server
- Cisco UCS C240 M6 Server
- Cisco UCS C225 M6 Server
- Cisco UCS C245 M6 Server

- Cisco UCS C480 M5 ML Server
- Cisco UCS S3260 M4 Storage Server—You can scrub only the boot drives and VDs created using the same drives.
- Cisco UCS S3260 M5 Storage Server—You can scrub only the boot drives and VDs created using the same drives.



- 
- Note** You must re-acknowledge the server to see the changes related to LUN deletion if:
- you are scrubbing boot drives which have LUNs under the SAS controller in a set up with Cisco UCS S3260 M4 or Cisco UCS S3260 M5 Storage Server.
  - you are scrubbing the LUNs on Cisco boot optimized M.2 RAID controller.
- 

### BIOS Settings Scrub

One of the following occurs to the BIOS settings when a service profile containing the scrub policy is disassociated from a server:

- If enabled, erases all BIOS settings for the server and resets them to the BIOS defaults for that server type and vendor.
- If disabled (default), preserves the existing BIOS settings on the server.

### FlexFlash Scrub

FlexFlash Scrub enables you to pair new or degraded SD cards, resolve FlexFlash metadata configuration failures, and migrate older SD cards with 4 partitions to single partition SD cards. One of the following occurs to the SD card when a service profile containing the scrub policy is disassociated from a server, or when the server is reacknowledged:

- If enabled, the HV partition on the SD card is formatted using the PNUOS formatting utility. If two SD cards are present, the cards are RAID-1 paired, and the HV partitions in both cards are marked as valid. The card in slot 1 is marked as primary, and the card in slot 2 is marked as secondary.
- If disabled (default), preserves the existing SD card settings.

**Note**

- For a server associated with a service profile, FlexFlash scrub occurs during disassociation, based on the scrub policy used in the service profile. For an un-associated server, FlexFlash scrub occurs during the server discovery process, based on the default scrub policy.
- Because the FlexFlash scrub erases the HV partition on the SD cards, we recommend that you take a full backup of the SD card(s) using your preferred host operating system utilities before performing the FlexFlash scrub.
- To resolve metadata config failures in a service profile, you need to disable FlexFlash in the local disk config policy before you run the FlexFlash scrub, then enable FlexFlash after the server is reacknowledged.
- Disable the scrub policy as soon as the pairing is complete or the metadata failures are resolved.
- FlexFlash scrub is not supported for Cisco UCS S3260 Storage Server.

**Persistent Memory Scrub**

Persistent memory scrub enables you to preserve or remove the persistent memory configuration and data on a server.

- If enabled:
  - erases all the persistent memory data
  - resets the configuration to factory default
  - disables DIMM security
- If disabled (default), preserves the existing persistent memory configuration and data on the server. It does not change the DIMM lock state.

## Creating a Scrub Policy

**Procedure**

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Scrub Policies** and select **Create Scrub Policy**.  
**Note** Cisco UCS Manager does not support NVME local disk scrub.
- Step 5** In the **Create Scrub Policy** wizard, complete the following fields:

Name	Description
<b>Name</b> field	<p>The name of the policy.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>
<b>Description</b> field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), &gt; (greater than), &lt; (less than), or ' (single quote).</p>
<b>Disk Scrub</b> field	<p>If this field is set to <b>Yes</b>, when a service profile containing this scrub policy is disassociated from a server, the initial 200MB of data is deleted from master boot record or the boot sectors. Thus, preventing the system to boot from an already installed OS if any. For secure deletion of data on drives, refer <a href="#">UCS Secure Data Deletion For Commission Regulation (EU) 2019 /424 Users Guide</a>. If this field is set to <b>No</b>, the data on the local drives is preserved, including all local storage configuration.</p> <p><b>Note</b>      Though the disk scrub policy is not intended to delete the user data that exceeds 200MB, Cisco UCS Manager cannot guarantee against data loss.</p>
<b>BIOS Settings Scrub</b> field	<p>If the field is set to <b>Yes</b>, when a service profile containing this scrub policy is disassociated from a server, the BIOS settings for that server are erased and reset to the defaults for that server type and vendor. If this field is set to <b>No</b>, the BIOS settings are preserved.</p>
<b>FlexFlash Scrub</b> field	<p>If the field is set to <b>Yes</b>, the HV partition on the SD card is formatted using the PNUOS formatting utility when the server is reacknowledged. If this field is set to <b>No</b>, the SD card is preserved.</p>
<b>Persistent Memory Scrub</b> field	<p>If the field is set to <b>Yes</b>, when a service profile containing this scrub policy is disassociated from a server, all persistent memory modules for that server are erased and reset to the defaults for that server type and vendor. If this field is set to <b>No</b>, the persistent memory modules are preserved.</p>

**Step 6** Click **OK**.

**Note** Disk Scrub option will scrub only boot-lun/boot-disk for Cisco UCS S3260 Storage Server and it will not scrub data-lun's/data-disks. The FlexFlash Scrub option is not supported for Cisco UCS S3260 Storage Server.

## Deleting a Scrub Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Expand **Servers > Policies > Organization\_Name**.
  - Step 3** Expand the **Scrub Policies** node.
  - Step 4** Right-click the policy you want to delete and select **Delete**.
  - Step 5** If a confirmation dialog box displays, click **Yes**.
- 

## DIMM Error Management

### DIMM Correctable Error Handling

In Cisco UCS Manager, when a DIMM encounters a significant correctable error in a given predefined window, it is stated as degraded and considered as a non-functional device.

The DIMM correctable error handling feature enables you to reset all the correctable and uncorrectable memory errors on all the DIMMs in a server. When you reset the error configuration, the error count of a given DIMM is cleared, the status changes to operable, and it resets the sensor state of the given DIMM.

### Resetting Memory Errors

Use this procedure to reset all correctable and uncorrectable memory errors encountered by Cisco UCS Manager and the baseboard management controller (BMC).

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
  - Step 3** Right-click on the server for which you want to reset the error configuration, and select **Reset All Memory Errors**. You can also select **Reset All Memory Errors** from the **Actions** area.
  - Step 4** If the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.
- 

### DIMM Blacklisting

In Cisco UCS Manager, the state of the Dual In-line Memory Module (DIMM) is based on SEL event records. When the BIOS encounters a noncorrectable memory error during memory test execution, the DIMM is marked as faulty. A faulty DIMM is considered a nonfunctional device.

If you enable DIMM blacklisting, Cisco UCS Manager monitors the memory test execution messages and blacklists any DIMMs that encounter memory errors in the DIMM SPD data. To allow the host to map out any DIMMs that encounter uncorrectable ECC errors.

## Enabling DIMM Blacklisting

The memory policy is a global policy that you can apply to existing servers on a Cisco UCS domain and also to the servers that are added after you set the memory policy.



- 
- Note**
- This feature is supported both on the Cisco UCS B-Series blade servers and UCS C-Series rack servers.
  - This global policy cannot be added to a service profile.
- 

### Before you begin

- For Cisco B-Series blade server, the server firmware must be at Release 2.2(1) or a later release.
- For Cisco C-Series and S-Series rack server, the server firmware must be at Release 2.2(3).
- You must be logged in with one of the following privileges:
  - Admin
  - Server policy
  - Server profile server policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to enable the blacklisting.  
If the system does not include multitenancy, expand the **root** node.
- Step 4** Expand **Memory Policy** and choose **default**.
- Step 5** In the **Blacklisting** area, click the **Enabled** radio button.
- 

The DIMM blacklisting is enabled for the domain level policy and these changes apply to all the servers on that particular domain.



- 
- Note** If the Cisco IMC of a server does not support DIMM blacklisting, an information level fault is generated.
-

# Serial over LAN Policy Settings

## Serial over LAN Policy Overview

This policy sets the configuration for the serial over LAN connection for all servers associated with service profiles that use the policy. By default, the serial over LAN connection is disabled.

If you implement a serial over LAN policy, we recommend that you also create an IPMI profile.

You must include this policy in a service profile and that service profile must be associated with a server for it to take effect.

## Creating a Serial over LAN Policy

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Serial over LAN Policies** and select **Create Serial over LAN Policy**.
- Step 5** In the **Create Serial over LAN Policy** wizard, complete the following fields:

Name	Description
Name field	The name of the policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy.  Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Serial over LAN State field	This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Disable</b>—Serial over LAN access is blocked.</li> <li>• <b>Enable</b>—Serial over LAN access is permitted.</li> </ul>

Name	Description
Speed drop-down list	This can be one of the following: <ul style="list-style-type: none"> <li>• 9600</li> <li>• 19200</li> <li>• 38400</li> <li>• 57600</li> <li>• 115200</li> </ul>

**Step 6** Click **OK**.

---

## Deleting a Serial over LAN Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Expand **Servers > Policies > Organization\_Name**.
  - Step 3** Expand the **Serial over LAN Policies** node.
  - Step 4** Right-click the policy you want to delete and select **Delete**.
  - Step 5** If a confirmation dialog box displays, click **Yes**.
- 

## Server Autoconfiguration Policies

### Server Autoconfiguration Policy Overview

Cisco UCS Manager uses this policy to determine how to configure a new server. If you create a server autoconfiguration policy, the following occurs when a new server starts:

1. The qualification in the server autoconfiguration policy is executed against the server.
2. If the server meets the required qualifications, the server is associated with a service profile created from the service profile template configured in the server autoconfiguration policy. The name of that service profile is based on the name given to the server by Cisco UCS Manager.
3. The service profile is assigned to the organization configured in the server autoconfiguration policy.



## Creating an Autoconfiguration Policy

### Before you begin

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

- Server pool policy qualifications
- Service profile template
- Organizations, if a system implements multitenancy

### Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Click the **Equipment** node.
- Step 3** In the **Work** pane, click the **Policies** tab.
- Step 4** Click the **Autoconfig Policies** subtab.
- Step 5** On the icon bar to the right of the table, click +.
- If the + icon is disabled, click an entry in the table to enable it.
- Step 6** In the **Create Autoconfiguration Policy** dialog box, complete the following fields:

Name	Description
<b>Name</b> field	The name of the policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
<b>Description</b> field	A description of the policy. Cisco recommends including information about where and when to use the policy.  Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
<b>Qualification</b> drop-down list	The server pool policy qualification associated with this auto-configuration policy.  If a new server is discovered that matches the criteria specified in the server pool policy qualification, Cisco UCS automatically creates a service profile based on the service profile template selected in the <b>Service Profile Template Name</b> drop-down list and associates the newly created service profile with the server.

Name	Description
Org drop-down list	The organization associated with this autoconfiguration policy.  If Cisco UCS automatically creates a service profile to associate with a server, it places the service profile under the organization selected in this field.
Service Profile Template Name drop-down list	The service profile template associated with this policy.

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

---

## Deleting an Autoconfiguration Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** Click the **Equipment** node.
  - Step 3** In the **Work** pane, click the **Policies** tab.
  - Step 4** Click the **Autoconfig Policies** subtab.
  - Step 5** Right-click the autoconfiguration policy that you want to delete and choose **Delete**.
  - Step 6** If a confirmation dialog box displays, click **Yes**.
- 

## Server Discovery Policy Settings

### Server Discovery Policy Overview

The server discovery policy determines how the UCS Manager reacts when you add a new UCS Blade Server and UCS Mini. If you create a server discovery policy, you can control whether the system conducts a deep discovery when a server is added to a chassis, or whether a user must first acknowledge the new server. By default, the system conducts a full discovery.

If you create a server discovery policy, the following occurs when a new server starts:

1. The server discovery policy qualification is executed against the server.
2. If the server meets the required qualifications, Cisco UCS Manager applies the following to the server:
  - Depending on the option that you select for the action, UCS Manager discovers the new server immediately, or waits for a user acknowledgment of the new server
  - Applies the scrub policy to the server

If automatic deep discovery is triggered by any hardware insertion, removal, or replacement, the following occurs:

1. The server is moved to a “pending activities” list.
2. A critical hardware mismatch fault is raised on the server, indicating that UCSM has detected a hardware mismatch.
3. User must explicitly acknowledge the server to trigger the deep discovery.



---

**Important** In Cisco UCS Manager Release 2.2 (4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears:

```
Unable to get Scsi Device Information from the system
```

If this error occurs, do the following:

1. Remove the 4K drive.
2. Reacknowledge the server.

Reacknowledging the server causes the server to reboot and results in loss of service.

---

## Creating a Server Discovery Policy

### Before you begin

If you plan to associate this policy with a server pool, create server pool policy qualifications.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** In the **Work** pane, click the **Policies** tab.
  - Step 3** Click the **Server Discovery Policies** subtab.
  - Step 4** Click the + icon on the table icon bar to open the **Create Server Discovery Policy** dialog box.
  - Step 5** In the **Description** field, enter a description for the discovery policy.
  - Step 6** In the **Action** field, select one of the following options:
    - **Immediate**—Cisco UCS Manager attempts to discover new servers automatically
    - **User Acknowledged**—Cisco UCS Manager waits until the user tells it to search for new servers
  - Step 7** (Optional) To associate this policy with a server pool, select server pool policy qualifications from the **Qualification** drop-down list.
  - Step 8** (Optional) To include a scrub policy, select a policy from the **Scrub Policy** drop-down list.
  - Step 9** Click **OK**.
-

**What to do next**

Include the server discovery policy in a service profile and/or template.

## Deleting a Server Discovery Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** In the **Work** pane, click the **Policies** tab.
  - Step 3** Click the **Server Discovery Policies** subtab.
  - Step 4** Right-click the server discover policy that you want to delete and choose **Delete**.
  - Step 5** If a confirmation dialog box displays, click **Yes**.
- 

## Hardware Change Discovery Policy

The Hardware Change Discovery is a global policy used to set the how Cisco UCS Manager behaves when there is a hardware component change. The policy has two values:

- **User Acknowledged:** You must acknowledge the server to clear all the hardware inventory mismatch faults.
- **Auto Acknowledged:** Triggers automatic deep discovery when a hardware component change is detected.

When UCSM detects any change in the server hardware component, a critical hardware inventory mismatch fault is raised on the server. You must manually acknowledge the server to clear the fault and complete the hardware inventory. Once you have acknowledged the server, deep discovery and deep association is triggered.

For rack servers, you must decommission and recommission the server to clear the fault and complete the hardware inventory.

You cannot make changes to the policy if there is a hardware inventory mismatch fault.

## Configuring Hardware Change Discovery Policy

### Procedure

---

- Step 1** Navigate to **Equipment > Policies > Global Policies**
- Step 2** Under **Hardware Change Discovery Policy**, choose one of the following:
  - **User Acknowledged:** You must acknowledge the server to clear all the hardware inventory mismatch faults.
  - **Auto Acknowledged:** Triggers automatic deep discovery when a hardware component change is detected.

**Step 3** Click **Save Changes**.

---

## Server Inheritance Policy Settings

### Server Inheritance Policy Overview

This policy is invoked during the server discovery process to create a service profile for the server. All service profiles created from this policy use the values burned into the blade at manufacture. The policy performs the following:

- Analyzes the inventory of the server
- If configured, assigns the server to the selected organization
- Creates a service profile for the server with the identity burned into the server at manufacture

You cannot migrate a service profile created with this policy to another server.

### Creating a Server Inheritance Policy

A blade server or rack-mount server with a VIC adapter, such as the Cisco UCS M81KR Virtual Interface Card, does not have server identity values burned into the server hardware at manufacture. As a result, the identity of the adapter must be derived from default pools. If the default pools do not include sufficient entries for one to be assigned to the server, service profile association fails with a configuration error.

#### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** In the **Work** pane, click the **Policies** tab.
- Step 3** Click the **Server Inheritance Policies** subtab.
- Step 4** On the icon bar at the bottom of the table, click + **Add**.  
If + **Add** is disabled, click an entry in the table to enable it.
- Step 5** In the **Create Server Inheritance Policy** dialog box, complete the following fields:

Name	Description
Name field	The name of the policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
<b>Description</b> field	A description of the policy. Cisco recommends including information about where and when to use the policy.  Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
<b>Qualification</b> drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.
<b>Org</b> drop-down list	If you want to associate an organization with this policy, or if you want to change the current association, choose the organization from the drop-down list.

**Step 6** Click **OK**.

---

## Deleting a Server Inheritance Policy

### Procedure

---

- Step 1** In the **Navigation** pane, click **Equipment**.
  - Step 2** In the **Work** pane, click the **Policies** tab.
  - Step 3** Click the **Server Inheritance Policies** subtab.
  - Step 4** Right-click the server inheritance policy that you want to delete and choose **Delete**.
  - Step 5** If a confirmation dialog box displays, click **Yes**.
- 

## Server Pool Policy Settings

### Server Pool Policy Overview

This policy is invoked during the server discovery process. It determines what happens if server pool policy qualifications match a server to the target pool specified in the policy.

If a server qualifies for more than one pool and those pools have server pool policies, the server is added to all those pools.

### Creating a Server Pool Policy

#### Before you begin

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

- A minimum of one server pool
- Server pool policy qualifications, if you choose to have servers automatically added to pools

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **Server Pool Policies** and select **Create Server Pool Policy**.
- Step 5** In the **Create Server Pool Policy** dialog box, complete the following fields:

Name	Description
<b>Name</b> field	The name of the policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
<b>Description</b> field	A description of the policy. Cisco recommends including information about where and when to use the policy.  Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
<b>Target Pool</b> drop-down list	If you want to associate this policy with a server pool, select that pool from the drop-down list.
<b>Qualification</b> drop-down list	To associate this policy with one or more specific server pools, choose the server pool qualification policy that identifies these pools.

- Step 6** Click **OK**.

## Deleting a Server Pool Policy

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **Server Pool Policies** node.
- Step 4** Right-click the policy you want to delete and select **Delete**.

**Step 5** If a confirmation dialog box displays, click **Yes**.

---

## Server Pool Policy Qualifications Settings

### Server Pool Policy Qualification Overview

This policy qualifies servers based on the inventory of a server conducted during the discovery process. The qualifications are individual rules that you configure in the policy to determine whether a server meets the selection criteria. For example, you can create a rule that specifies the minimum memory capacity for servers in a data center pool.

Qualifications are used in other policies to place servers, not just by the server pool policies. For example, if a server meets the criteria in a qualification policy, it can be added to one or more server pools or have a service profile automatically associated with it.

You can use the server pool policy qualifications to qualify servers according to the following criteria:

- Adapter type
- Chassis location
- Memory type and configuration
- Power group
- CPU cores, type, and configuration
- Storage configuration and capacity
- Server model

Depending upon the implementation, you might need to configure several policies with server pool policy qualifications including the following:

- Autoconfiguration policy
- Chassis discovery policy
- Server discovery policy
- Server inheritance policy
- Server pool policy

## Creating Server Pool Policy Qualifications

### Procedure

---

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

**Step 2** Expand **Servers > Policies**.



**Step 3** Expand the node for the organization where you want to create the policy.

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

**Step 4** Right-click the **Server Pool Policy Qualifications** node and select **Create Server Pool Policy Qualification**.

**Step 5** In the **Create Server Pool Policy Qualification** dialog box, enter a unique name and description for the policy.

**Step 6** (Optional) To use this policy to qualify servers according to their adapter configuration, do the following:

- a) Click **Create Adapter Qualifications**.
- b) In the **Create Adapter Qualifications** dialog box, complete the following fields:

Name	Description
<b>Type</b> drop-down list	The adapter type. Once you save the adapter qualification, this type cannot be changed.
<b>PID</b> field	A regular expression that the adapter PID must match.
<b>Maximum Capacity</b> field	The maximum capacity for the selected type. To specify a capacity, choose <b>select</b> and enter the desired maximum capacity. You can enter an integer between 1 and 65535.

- c) Click **OK**.

**Step 7** (Optional) To use this policy to qualify servers according to the chassis in which they physically reside, do the following:

- a) Click **Create Chassis/Server Qualifications**.
- b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:
  - **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
  - **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.

**Example:**

For example, if you want to use chassis 5, 6, 7, and 8, enter **5** in the **First Chassis ID** field and **4** in the **Number of Chassis** field. If you want to use only chassis 3, enter **3** in the **First Chassis ID** field and **1** in the **Number of Chassis** field.

**Tip** If you want to use chassis 5, 6, and 9, create a chassis/server qualification for the range 5-6 and another qualification for chassis 9. You can add as many chassis/server qualifications as needed.

- c) Click **Finish**.

**Step 8** (Optional) To use this policy to qualify servers according to both the chassis and slot in which they physically reside, do the following:

- a) Click **Create Chassis/Server Qualifications**.
- b) In the **Chassis Qualifications** area of the **Create Chassis and Server Qualifications** dialog box, complete the following fields to specify the range of chassis you want to use:

- **First Chassis ID** field—The first chassis ID from which server pools associated with this policy can draw.
  - **Number of Chassis** field—The total number of chassis to include in the pool, starting with the chassis identified in the **First Chassis ID** field.
- c) In the **Server Qualifications** table, click **Add**.
- d) In the **Create Server Qualifications** dialog box, complete the following fields to specify the range of server locations you want to use:
- **First Slot ID** field—The first slot ID from which server pools associated with this policy can draw.
  - **Number of Slots** field—The total number of slots from which server pools associated with this policy can draw.
- e) Click **Finish Stage**.
- f) To add another range of slots, click **Add** and repeat steps d and e.
- g) When you have finished specifying the slot ranges, click **Finish**.

**Step 9**

(Optional) To use this policy to qualify servers according to their memory configuration, do the following:

- a) Click **Create Memory Qualifications**.
- b) In the **Create Memory Qualifications** dialog box, complete the following fields:

Name	Description
<b>Clock</b> field	The minimum clock speed required, in megahertz.
<b>Latency</b> field	The maximum latency allowed, in nanoseconds.
<b>Min Cap</b> field	The minimum memory capacity required, in megabytes.
<b>Max Cap</b> field	The maximum memory capacity allowed, in megabytes.
<b>Width</b> field	The minimum width of the data bus.
<b>Units</b> field	The unit of measure to associate with the value in the <b>Width</b> field.

- c) Click **OK**.

**Step 10**

(Optional) To use this policy to qualify servers according to their CPU/Cores configuration, do the following:

- a) Click **Create CPU/Cores Qualifications**.
- b) In the **Create CPU/Cores Qualifications** dialog box, complete the following fields:

Name	Description
<b>Processor Architecture</b> drop-down list	The CPU architecture to which this policy applies.
<b>PID</b> field	A regular expression that the processor PID must match.
<b>Min Number of Cores</b> field	The minimum number of CPU cores required.  To specify a capacity, choose <b>select</b> and enter an integer between 1 and 65535 in the associated text field.

Name	Description
<b>Max Number of Cores</b> field	The maximum number of CPU cores allowed. To specify a capacity, choose <b>select</b> and enter an integer between 1 and 65535 in the associated text field.
<b>Min Number of Threads</b> field	The minimum number of CPU threads required. To specify a capacity, choose <b>select</b> and enter an integer between 1 and 65535 in the associated text field.
<b>Max Number of Threads</b> field	The maximum number of CPU threads allowed. To specify a capacity, choose <b>select</b> and enter an integer between 1 and 65535 in the associated text field.
<b>CPU Speed</b> field	The minimum CPU speed required. To specify a capacity, choose <b>select</b> and enter the minimum CPU speed.
<b>CPU Stepping</b> field	The minimum CPU version required. To specify a capacity, choose <b>select</b> and enter the maximum CPU speed.

c) Click **OK**.

### Step 11

(Optional) To use this policy to qualify servers according to their storage configuration and capacity, do the following:

- a) Click **Create Storage Qualifications**.
- b) In the **Create Storage Qualifications** dialog box, complete the following fields:

Name	Description
<b>Diskless</b> field	Whether the available storage must be diskless. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Unspecified</b>—Either storage type is acceptable.</li> <li>• <b>Yes</b>—The storage must be diskless.</li> <li>• <b>No</b>—The storage cannot be diskless.</li> </ul>
<b>Number of Blocks</b> field	The minimum number of blocks required. To specify a capacity, choose <b>select</b> and enter the number of blocks.
<b>Block Size</b> field	The minimum block size required, in bytes. To specify a capacity, choose <b>select</b> and enter the block size.
<b>Min Cap</b> field	The minimum storage capacity across all disks in the server, in megabytes. To specify a capacity, choose <b>select</b> and enter the minimum storage capacity.

Name	Description
<b>Max Cap</b> field	The maximum storage capacity allowed, in megabytes. To specify a capacity, choose <b>select</b> and enter the maximum storage capacity.
<b>Per Disk Cap</b> field	The minimum storage capacity per disk required, in gigabytes. To specify a capacity, choose <b>select</b> and enter the minimum capacity on each disk.
<b>Units</b> field	The number of units. To specify a capacity, choose <b>select</b> and enter the desired units.
<b>Number of Flex Flash Cards</b> field	The number of FlexFlash Cards. To specify a capacity, choose <b>select</b> and enter the desired units.
<b>Disk Type</b> field	The disk type. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Unspecified</b>—Either disk type is acceptable.</li> <li>• <b>HDD</b>—The disk must be HDD.</li> <li>• <b>SSD</b>—The disk must be SSD (SATA or SAS).</li> </ul>

c) Click **OK**.

**Step 12** (Optional) To use this policy to qualify servers according to the model of the server, do the following:

- Click **Create Server Model Qualifications**.
- In the **Create Server Model Qualifications** dialog box, enter a regular expression that the server model must match in the **Model** field.
- Click **OK**.

**Step 13** (Optional) To use this policy to qualify servers according to power group, do the following:

- Click **Create Power Group Qualifications**.
- In the **Create Power Group Qualifications** dialog box, choose a power group from the **Power Group** drop-down list.
- Click **OK**.

**Step 14** (Optional) To use this policy to qualify the rack-mount servers that can be added to the associated server pool, do the following:

- Click **Create Rack Qualifications**.
- In the **Create Rack Qualifications** dialog box, complete the following fields:

Name	Description
<b>First Slot ID</b> field	The first rack-mount server slot ID from which server pools associated with this policy can draw.
<b>Number of Slots</b> field	The total number of rack-mount server slots from which server pools associated with this policy can draw.

- Step 15** Verify the qualifications in the table and correct if necessary.
- Step 16** Click **OK**.
- 

## Deleting Server Pool Policy Qualifications

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **Server Pool Policy Qualifications** node.
- Step 4** Right-click the policy qualifications you want to delete and select **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.
- 

## Deleting Qualifications from Server Pool Policy Qualifications

Use this procedure to modify Server Pool Policy Qualifications by deleting one or more sets of qualifications.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **Server Pool Policy Qualifications** node.
- Step 4** Choose the policy you want to modify.
- Step 5** In the **Work** pane, choose the **Qualifications** tab.
- Step 6** To delete a set of qualifications:
- In the table, choose the row that represents the set of qualifications.
  - Right-click the row and select **Delete**.
- Step 7** Click **Save Changes**.
- 

## vNIC/vHBA Placement Policy Settings

### vNIC/vHBA Placement Policies

vNIC/vHBA placement policies are used to determine the following:

- How the virtual network interface connections (vCons) are mapped to the physical adapters on a server.

- What types of vNICs or vHBAs can be assigned to each vCon.

Each vNIC/vHBA placement policy contains four vCons that are virtual representations of the physical adapters. When a vNIC/vHBA placement policy is assigned to a service profile, and the service profile is associated with a server, the vCons in the vNIC/vHBA placement policy are assigned to the physical adapters and the vNICs and vHBAs are assigned to those vCons.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the type of server and the selected virtual slot mapping scheme, which can be **Round Robin** or **Linear Ordered**. For details about the available mapping schemes, see [vCon to Adapter Placement, on page 142](#).

After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the **Selection Preference** for each vCon. This can be one of the following:




---

**Note** You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

---

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.
- **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
- **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
- **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
- **Exclude usNIC**—Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.




---

**Note** An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

---

If you do not include a vNIC/vHBA placement policy in the service profile, Cisco UCS Manager defaults to the **Round Robin** vCon mapping scheme and the **All** vNIC/vHBA selection preference, distributing the vNICs and vHBAs between the adapters based on the capabilities and relative capacities of each adapter.

## vCon to Adapter Placement

Cisco UCS maps every vCon in a service profile to a physical adapter on the server. How that mapping occurs and how the vCons are assigned to a specific adapter in a server depends on the following:

- The type of server. N20-B6620-2 and N20-B6625-2 blade servers with two adapter cards use a different mapping scheme than other supported rack or blade servers.
- The number of adapters in the server.
- The setting of the virtual slot mapping scheme in the vNIC/vHBA placement policy, if applicable.

You must consider this placement when you configure the vNIC/vHBA selection preference to assign vNICs and vHBAs to vCons.



**Note** vCon to adapter placement is not dependent upon the PCIE slot number of the adapter. The adapter numbers used for the purpose of vCon placement are not the PCIE slot numbers of the adapters, but the ID assigned to them during server discovery.

## vCon to Adapter Placement for N20-B6620-2 and N20-B6625-2 Blade Servers

In N20-B6620-2 and N20-B6625-2 blade servers, the two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:

- **Round Robin**—Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default.
- **Linear Ordered**—Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2.

## vCon to Adapter Placement for All Other Supported Servers

For all other servers supported by Cisco UCS in addition to the N20-B6620-2 and N20-B6625-2 blade servers, the vCon assignment depends on the number of adapters in the server and the virtual slot mapping scheme.

For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme: Round Robin or Linear Ordered.

**Table 1: vCon to Adapter Placement Using the Round - Robin Mapping Scheme**

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter2	Adapter1	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter2
4	Adapter1	Adapter2	Adapter3	Adapter4

Round Robin is the default mapping scheme.

Table 2: vCon to Adapter Placement Using the Linear Ordered Mapping Scheme

Number of Adapters	vCon1 Assignment	vCon2 Assignment	vCon3 Assignment	vCon4 Assignment
1	Adapter1	Adapter1	Adapter1	Adapter1
2	Adapter1	Adapter1	Adapter2	Adapter2
3	Adapter1	Adapter2	Adapter3	Adapter3
4	Adapter1	Adapter2	Adapter3	Adapter4

## vNIC/vHBA to vCon Assignment

Cisco UCS Manager provides two options for assigning vNICs and vHBAs to vCons through the vNIC/vHBA placement policy: explicit assignment and implicit assignment.

### Explicit Assignment of vNICs and vHBAs

With explicit assignment, you specify the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned. Use this assignment option when you need to determine how the vNICs and vHBAs are distributed between the adapters on a server.

To configure a vCon and the associated vNICs and vHBAs for explicit assignment, do the following:

- Set the vCon configuration to any of the available options. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server. If a vCon is configured for **All**, you can still explicitly assign a vNIC or vHBA to that vCon.
- Assign the vNICs and vHBAs to a vCon. You can make this assignment through the virtual host interface placement properties of the vNIC or vHBA or in the service profile associated with the server.

If you attempt to assign a vNIC or vHBA to a vCon that is not configured for that type of vNIC or vHBA, Cisco UCS Manager displays a message advising you of the configuration error.

During service profile association, Cisco UCS Manager validates the configured placement of the vNICs and vHBAs against the number and capabilities of the physical adapters in the server before assigning the vNICs and vHBAs according to the configuration in the policy. Load distribution is based upon the explicit assignments to the vCons and adapters configured in this policy.

If the adapters do not support the assignment of one or more vNICs or vHBAs, Cisco UCS Manager raises a fault against the service profile.




---

**Note** You can specify the PCI order for the vHBA; however, the desired order works within a class of devices, such as vNICs or vHBAs and not across them. Within an adapter, vNICs are always placed ahead of the vHBAs.

---

### Implicit Assignment of vNICs and vHBAs

With implicit assignment, Cisco UCS Manager determines the vCon and, therefore, the adapter to which a vNIC or vHBA is assigned according to the capability of the adapters and their relative capacity. Use this assignment option if the adapter to which a vNIC or vHBA is assigned is not important to your system configuration.



To configure a vCon for implicit assignment, do the following:

- Set the vCon configuration to **All**, **Exclude Dynamic**, or **Exclude Unassigned**. You can configure the vCons through a vNIC/vHBA placement policy or in the service profile associated with the server.
- Do not set the vCon configuration to **Assigned Only**. Implicit assignment cannot be performed with this setting.
- Do not assign any vNICs or vHBAs to a vCon.

During service profile association, Cisco UCS Manager verifies the number and capabilities of the physical adapters in the server and assigns the vNICs and vHBAs accordingly. Load distribution is based upon the capabilities of the adapters, and placement of the vNICs and vHBAs is performed according to the actual order determined by the system. For example, if one adapter can accommodate more vNICs than another, that adapter is assigned more vNICs.

If the adapters cannot support the number of vNICs and vHBAs configured for that server, Cisco UCS Manager raises a fault against the service profile.

### Implicit Assignment of vNICs in a Dual Adapter Environment

When you use implicit vNIC assignment for a dual slot server with an adapter card in each slot, Cisco UCS Manager typically assigns the vNICs/vHBAs as follows:

- If the server has the same adapter in both slots, Cisco UCS Manager assigns half the vNICs and half the vHBAs to each adapter.
- If the server has one non-VIC adapter and one VIC adapter, Cisco UCS Manager assigns two vNICs and two vHBAs to the non-VIC adapter and the remaining vNICs and vHBAs to the VIC adapter.
- If the server has two different VIC adapters, Cisco UCS Manager assigns the vNICs and vHBAs proportionally, based on the relative capabilities of the two adapters.

The following examples show how Cisco UCS Manager would typically assign the vNICs and vHBAs with different combinations of supported adapter cards:

- If you want to configure four vNICs and the server contains two Cisco UCS M51KR-B Broadcom BCM57711 adapters (with two vNICs each), Cisco UCS Manager assigns two vNICs to each adapter.
- If you want to configure 50 vNICs and the server contains a Cisco UCS CNA M72KR-E adapter (2 vNICs) and a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs), Cisco UCS Manager assigns two vNICs to the Cisco UCS CNA M72KR-E adapter and 48 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter.
- If you want to configure 150 vNICs and the server contains a Cisco UCS M81KR Virtual Interface Card adapter (128 vNICs) and a Cisco UCS VIC-1240 Virtual Interface Card adapter (256 vNICs), Cisco UCS Manager assigns 50 vNICs to the Cisco UCS M81KR Virtual Interface Card adapter and 100 vNICs to the Cisco UCS VIC-1240 Virtual Interface Card adapter.



---

**Note** Exceptions to this implicit assignment occur if you configure the vNICs for fabric failover and if you configure dynamic vNICs for the server.

---

For a configuration that includes vNIC fabric failover where one adapter does not support vNIC failover, Cisco UCS Manager implicitly assigns all vNICs that have fabric failover enabled to the adapter that supports them. If the configuration includes only vNICs that are configured for fabric failover, no vNICs are implicitly assigned to the adapter that does not support them. If some vNICs are configured for fabric failover and some are not, Cisco UCS Manager assigns all failover vNICs to the adapter that supports them and a minimum of one nonfailover vNIC to the adapter that does not support them, according to the ratio above.

For a configuration that includes dynamic vNICs, the same implicit assignment would occur. Cisco UCS Manager assigns all dynamic vNICs to the adapter that supports them. However, with a combination of dynamic vNICs and static vNICs, at least one static vNIC is assigned to the adapter that does not support dynamic vNICs.

## Creating a vNIC/vHBA Placement Policy

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **vNIC/vHBA Placement Policies** and choose **Create Placement Policy**.
- Step 5** In the **Create Placement Policy** dialog box, do the following:
- Complete the following fields:

Name	Description
Name field	<p>The name for this placement policy.</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.</p>

Name	Description
<b>Virtual Slot Mapping Scheme</b> field	<p>Cisco UCS assigns virtual network interface connections (vCons) to the PCIe adapter cards in the server. Each vCon is a virtual representation of a physical adapter that can be assigned vNICs and vHBAs.</p> <p>For blade or rack servers that contain one adapter, Cisco UCS assigns all vCons to that adapter. For servers that contain four adapters, Cisco UCS assigns vCon1 to Adapter1, vCon2 to Adapter2, vCon3 to Adapter3, and vCon4 to Adapter4.</p> <p>For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the selected virtual slot mapping scheme. This can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>Round Robin</b>— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon3 to Adapter1, then assigns vCon2 and vCon4 to Adapter2.</li> </ul> <p>In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1, vCon2 and vCon4 to Adapter2, and vCon3 to Adapter3.</p> <p>This is the default scheme.</p> <ul style="list-style-type: none"> <li>• <b>Linear Ordered</b>— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon2 to Adapter1, then assigns vCon3 and vCon4 to Adapter2.</li> </ul> <p>In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1 and vCon2 to Adapter2, then assigns vCon3 and vCon4 to Adapter3.</p> <p><b>Note</b> In N20-B6620-2 and N20-B6625-2 blade servers, the two adapters are numbered left to right while vCons are numbered right to left. If one of these blade servers has a single adapter, Cisco UCS assigns all vCons to that adapter. If the server has two adapters, the vCon assignment depends upon the virtual slot mapping scheme:</p> <ul style="list-style-type: none"> <li>• <b>Round Robin</b>—Cisco UCS assigns vCon2 and vCon4 to Adapter1 and vCon1 and vCon3 to Adapter2. This is the default.</li> <li>• <b>Linear Ordered</b>—Cisco UCS assigns vCon3 and vCon4 to Adapter1 and vCon1 and vCon2 to Adapter2.</li> </ul> <p>After Cisco UCS assigns the vCons, it assigns the vNICs and vHBAs based on the <b>Selection Preference</b> for each vCon.</p>

- b) In the **Selection Preference** column for each **Virtual Slot**, choose one of the following from the drop-down list:

- **All**—All configured vNICs and vHBAs can be assigned to the vCon, whether they are explicitly assigned to it, unassigned, or dynamic. This is the default.
  - **Assigned Only**—vNICs and vHBAs must be explicitly assigned to the vCon. You can assign them explicitly through the service profile or the properties of the vNIC or vHBA.
  - **Exclude Dynamic**—Dynamic vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for all static vNICs and vHBAs, whether they are unassigned or explicitly assigned to it.
  - **Exclude Unassigned**—Unassigned vNICs and vHBAs cannot be assigned to the vCon. The vCon can be used for dynamic vNICs and vHBAs and for static vNICs and vHBAs that are explicitly assigned to it.
  - **Exclude usNIC**—Cisco usNICs cannot be assigned to the vCon. The vCon can be used for all other configured vNICs and vHBAs, whether they are explicitly assigned to it, unassigned, or dynamic.
- Note** An SRIOV usNIC that is explicitly assigned to a vCon set to **Exclude usNIC** will remain assigned to that vCon.

c) Click **OK**.

## Deleting a vNIC/vHBA Placement Policy

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization\_Name**.
- Step 3** Expand the **vNIC/vHBA Placement Policies** node.
- Step 4** Right-click the policy you want to delete and choose **Delete**.
- Step 5** If a confirmation dialog box displays, click **Yes**.

## Explicitly Assigning a vNIC to a vCon

### Before you begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

## Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization which contains the service profile whose vNICs you want to explicitly assign to a vCon.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand *Service\_Profile\_Name* > **vNICs**.
- Step 5** Click on the vNIC that you want to explicitly assign to a vCon.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
<b>Desired Placement</b> drop-down list	The user-specified virtual network interface connection (vCon) placement for the vNIC. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Any</b>—Allows Cisco UCS Manager to determine the vCon to which the vNIC is assigned.</li> <li>• <b>1</b>—Explicitly assigns the vNIC to vCon1.</li> <li>• <b>2</b>—Explicitly assigns the vNIC to vCon2.</li> <li>• <b>3</b>—Explicitly assigns the vNIC to vCon3.</li> <li>• <b>4</b>—Explicitly assigns the vNIC to vCon4.</li> </ul>
<b>Actual Assignment</b> field	The actual vCon assignment of the vNIC on the server.

If you attempt to assign a vNIC to a vCon that is not configured for that type of vNIC, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vNIC to another vCon or change the vCon configuration in the service profile.

- Step 8** In the **Order** section, complete the following fields:

Name	Description
<b>Desired Order</b> field	The user-specified PCI order for the vNIC. Enter an integer between 0 and 128. You cannot create more than 128 vNICs for a server.
<b>Actual Order</b> field	The actual PCI order of the vNIC on the server.

- Step 9** Click **Save Changes**.

## Explicitly Assigning a vHBA to a vCon

### Before you begin

Configure the vCons through a vNIC/vHBA placement policy or in the service profile with one of the following values:

- **Assigned Only**
- **Exclude Dynamic**
- **Exclude Unassigned**

If a vCon is configured for **All**, you can explicitly assign a vNIC or vHBA to that vCon. However, there is less control with this configuration.

### Procedure

- 
- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization which contains the service profile whose vHBAs you want to explicitly assign to a vCon.
- If the system does not include multi tenancy, expand the **root** node.
- Step 4** Expand *Service\_Profile\_Name* > **vHBAs**.
- Step 5** Click on the vHBA that you want to explicitly assign to a vCon.
- Step 6** In the **Work** pane, click the **General** tab.
- Step 7** In the **Virtual Host Interface Placement** section, complete the following fields:

Name	Description
<b>Desired Placement</b> field	The user-specified virtual network interface connection (vCon) placement for the vHBA. This can be one of the following: <ul style="list-style-type: none"> <li>• <b>Any</b>—Allows Cisco UCS Manager to determine the vCon to which the vHBA is assigned.</li> <li>• <b>1</b>—Explicitly assigns the vHBA to vCon1.</li> <li>• <b>2</b>—Explicitly assigns the vHBA to vCon2.</li> <li>• <b>3</b>—Explicitly assigns the vHBA to vCon3.</li> <li>• <b>4</b>—Explicitly assigns the vHBA to vCon4.</li> </ul>
<b>Actual Assignment</b> field	The actual vCon assignment of the vHBA on the server.

If you attempt to assign a vHBA to a vCon that is not configured for that type of vHBA, Cisco UCS Manager displays a message box to advise you of the configuration error. You must either assign the vHBA to another vCon or change the vCon configuration in the service profile.

**Step 8** In the **Order** section, complete the following fields:

Name	Description
<b>Desired Order</b> field	The user-specified PCI order for the vHBA. Enter an integer between 0 and 128. You cannot create more than 128 vHBAs for a server.
<b>Actual Order</b> field	The actual PCI order of the vHBA on the server.

**Step 9** Click **Save Changes**.

## Placing Static vNICs Before Dynamic vNICs

For optimal performance, static vNICs and vHBAs should be placed before dynamic vNICs on the PCIe bus. Static vNICs refer to both static vNICs and vHBAs. Cisco UCS Manager Release 2.1 provides the following functionality regarding the order of static and dynamic vNICs:

- After upgrading to Cisco UCS Manager Release 2.1, if no change is made to existing service profiles (profiles that are defined in releases prior to Cisco UCS Manager Release 2.1), the vNIC order does not change.
- After an upgrade to Cisco UCS Manager Release 2.1, any vNIC-related change would reorder the vNIC map. As a result, all dynamic vNICs would be placed after the static vNICs.
- For newly created service profiles in Cisco UCS Manager Release 2.1, static vNICs are always ordered before dynamic vNICs.
- The above behavior is independent of the sequence of creating or deleting static or dynamic vNICs.
- For SRIOV-enabled service profiles, UCSM places the vNIC Physical Function(PF) before the corresponding Virtual Functions (VFs). This scheme guarantees that the VFs are placed close to the parent PF vNIC on the PCIe bus and BDFs are in successive incremental order for the VFs.

### Example

Beginning Device Order in Cisco UCS Manager Release 2.0:

```
dyn-vNIC-1 1
dyn-vNIC-2 2
```

New Device Order in Cisco UCS Manager Release 2.0 (Add 2 static vNICs):

```
dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4
```

After upgrading to Cisco UCS Manager Release 2.1, (Before any vNIC-related change is made to the service profile.)

```
dyn-vNIC-1 1
dyn-vNIC-2 2
eth-vNIC-1 3
eth-vNIC-2 4
```

New Device Order in Cisco UCS Manager Release 2.1 (Add 2 dynamic vNICs by changing the policy count from 2 to 4.)

```

dyn-vNIC-1 3
dyn-vNIC-2 4
eth-vNIC-1 1
eth-vNIC-2 2
dyn-vNIC-3 5
dyn-vNIC-4 6

```

### Dynamic vNICs as Multifunction PCIe Devices

Cisco UCS Manager Version 2.1 provisions static vNICs as 0-function devices (new BUS for every static vNIC). Multifunction dynamic vNICs are placed from the new Bus-slot after the last static vNIC/vHBA.



**Note** Cisco UCS Manager Version 2.1 supports the new StaticZero mode.

**Table 3: Version Compatibility**

Cisco UCS Manager		
Version 1.4	Version 2.0	Version 2.1
<b>Scheme: ZeroFunction</b>	<b>Scheme: ZeroFunction / MultiFunction</b>	<b>Scheme: ZeroFunction / MultiFunction / StaticZero</b>
Static and Dynamic vNICs are all on Bus [0-57], Function [0] < ZeroFunction Mode >	Static vNICs and Dynamic vNICs are on Bus [0-57], Function [0-7]. Bus 0, Function 0 Bus 0, Function 7 Bus 1, Function 0 < MultiFunction Mode >	Static vNICs or PFs will be on Bus [0-57], Function [0]. SRIOV: Corresponding VFs will be on the same Bus and Functions [1-255] No-SRIOV: Dynamic vNICs are on Bus [0-57], Function [0-7] < StaticZero Mode >
	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <= 57. Once devices exceed 58, switch to MultiFunction mode.	Upgrade from Balboa will not renumber BDFs (remain in ZeroFunction mode) until Bus <=57. Once devices exceed 58 or Platform specific maximum PCIe Bus number or change to SRIOV configuration, switch to StaticZero mode.
		Upgrade from Cisco UCS Manager Version 2.0 will not renumber BDFs (remain in ZeroFunction / MultiFunction mode). Once devices exceed 58 or Platform specific maximum PCIe Bus number OR Change to SRIOV configuration, switch to StaticZero mode.



## vNIC/vHBA Host Port Placement

After a vNIC/vHBA is assigned to a vCON, it can be placed on one of the host ports of a specific adapter. You can either explicitly specify the host port for placement, or allow Cisco UCS Manager to automatically assign vNICs/vHBAs to host ports.



---

**Note** You can perform vNIC/vHBA host port placement on servers that support Cisco UCS VIC 1340/1380 adapters.

---

Cisco UCS 1300 Series adapters have 2x8 PCIe third generation host ports. Each PCIe host port is capable of a maximum of 64 Gbps bandwidth.

The host port placement of the vNIC/vHBA determines the order of the vNIC/vHBA on the adapter. The vNICs/vHBAs placed on the first host port will be enumerated first, followed by the vNICs/vHBAs on the second host port.



---

**Note** The 64 Gbps maximum is a theoretical maximum, actual data transfer is limited to around 40 Gbps.

---

All the vNICs sharing the same PCIe Host Port will share this bandwidth. To make the optimal use of PCIe host port bandwidth, vNICs should be distributed across the two host ports.

## Configuring Host Port Placement

You can configure host port placement for vNICs on servers that support Cisco UCS VIC 1340 and VIC 1380 adapters.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
  - Step 2** Expand **Servers > Service Profiles**.
  - Step 3** Select the service profile which is associated with the vNIC that you want to place on a host port.
  - Step 4** Expand **Service\_Profile\_Name > vNICs**
  - Step 5** Under the **Network** tab, in the **vNICs** summary table, double-click the **Admin Host Port** value of the vNIC which you want to configure and select one of the following:
    - **Any**—Allows Cisco UCS Manager to determine the host port to which the vNIC is assigned.
    - **1**—Explicitly assigns the vNIC to host port 1.
    - **2**—Explicitly assigns the vNIC to host port 2.
- Actual Host Port** displays the actual assignment of the vNIC on a host port. When this feature is not supported, this will appear as **None**.
- Step 6** Click **Save Changes**.
-

# CIMC Mounted vMedia

## Using Scriptable vMedia

Cisco UCS Manager allows provisioning of vMedia devices iso images for remote UCS servers. Using Scriptable vMedia, you can programmatically mount an IMG or an ISO image on a remote server. CIMC mounted vMedia provide communications between other mounted media inside your datacenter with no additional requirements media connection. Scriptable vMedia allows you to control virtual media devices without using a browser to manually map each UCS server individually.

**Scriptable vMedia** supports multiple share types including NFS, CIFS, HTTP, and HTTPS shares. **Scriptable vMedia** is enabled through BIOS configuration and configured through a Web GUI and CLI interface.

Cisco UCS Manager Scriptable vMedia supports the following functionality:

- Booting from a specific vMedia device
- Copying files from a mounted share to a local disk
- Installation and updating OS drivers



---

**Note** Cisco UCS Manager support for Scriptable vMedia is applicable for CIMC mapped devices only. Existing KVM based vMedia devices are not supported.

---

vMedia mount fails when the following conditions are met:

1. The remote vMedia image filename in the vMedia policy is set to **Service-Profile-Name**.
2. The service profile is renamed.

This is because the change in the name of the service profile does not change the remote vMedia image filename in the vMedia policy. The image filename still points to the older image on the remote device, which cannot be found.

## Creating a vMedia Policy

A vMedia policy is used to configure the mapping information for remote vMedia devices. Two vMedia devices and mappings for CD and HDD are allowed in a vMedia policy. You can configure one ISO and one IMG at a time. ISO configurations maps to a CD drive and IMG configurations maps to a HDD device.



---

**Note** If you want to map a device to a remote folder, you must create an IMG and map it as a HDD device.

---

### Before you begin

Make sure that you have access to the following:

- Remote vMedia Server

- vMedia Devices

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click **vMedia Policies** and select **Create vMedia Policy**.
- Step 5** In the **Create vMedia Policy** dialog box, complete the following fields:

Name	Description
<b>Name</b>	The name of the vMedia policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
<b>Description</b>	A description of the policy. We recommend including information about where and when the policy should be used. Maximum 115 characters.
<b>Retry on Mount Failure</b>	Designates if the vMedia will continue mounting when a mount failure occurs. This can be: <ul style="list-style-type: none"> <li>• <b>Yes</b></li> <li>• <b>No</b></li> </ul> <p><b>Note</b> The default setting is <b>Yes</b>. When <b>Yes</b> is selected the remote server will continue to try to mount the vMedia mount process until it is successful or you disable this option. If you select <b>No</b>, a warning message will appear indicating retry on mount failure will not work in case of mount failure.</p>

- Step 6** On the icon bar to the right of the table, click +.
- Step 7** In the **Create vMedia Mount** dialog box, complete the following fields:

Name	Description
<b>Name</b>	Name of the vMedia Mount policy.  This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.

Name	Description
<b>Device Type</b>	<p>The type of remote vMedia you plan to mount. This can be:</p> <ul style="list-style-type: none"> <li>• <b>CDD</b>—Scriptable vMedia CD.</li> <li>• <b>HDD</b>—Scriptable vMedia HDD.</li> </ul>
<b>Protocol</b>	<p>The protocol to use when communicating with the remote server. Click one of the following radio buttons to indicate the protocol you want to use to communicate with the mounted remote server. This can be:</p> <ul style="list-style-type: none"> <li>• <b>NFS</b> - Network Files System.</li> <li>• <b>CIFS</b> - Common Internet File System.</li> <li>• <b>HTTP</b> - Hypertext Transfer Protocol.</li> <li>• <b>HTTPS</b> - Hypertext Transfer Protocol over Secure.</li> </ul>
<b>Authentication Protocol</b>	<p>The protocol to use for authentication when you use CIFS as the protocol for communicating with the remote server. When you use any protocol other than CIFS, this field is not available. Select one of the following from the drop-down list to specify the authentication protocol.</p> <ul style="list-style-type: none"> <li>• <b>Default</b>—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2.</li> <li>• <b>None</b>—No authentication is used</li> <li>• <b>Ntlm</b>—NT LAN Manager (NTLM) security protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2.</li> <li>• <b>Ntlmi</b>—NTLMi security protocol. Use this option only when you enable Digital Signing in the CIFS Windows server.</li> <li>• <b>Ntlmssp</b>—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2.</li> <li>• <b>Ntlmsspi</b>—NTLMSSPi protocol. Use this option only when you enable Digital Signing in the CIFS Windows server.</li> <li>• <b>Ntlmv2</b>—NTLMv2 security protocol. Use this option only with Samba Linux.</li> <li>• <b>Ntlmv2i</b>—NTLMv2i security protocol. Use this option only with Samba Linux.</li> </ul> <p><b>Note</b> The authentication protocol options are available only when you select <b>CIFS</b> as the protocol. For all other protocols, the <b>Authentication Protocol</b> field is disabled.</p>

Name	Description
<b>Hostname/IPAddress</b>	<p>Enter the IP address or hostname of the location where the backup file is to be stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.</p> <p>If you use a hostname, you must configure Cisco UCS Manager to use a DNS server. The hostname (DNS) can be used when <b>Inband</b> network is configured for that server.</p>
<b>Image Name Variable</b>	<p>The name to be used for the image. This can be:</p> <ul style="list-style-type: none"> <li>• <b>None</b>—Filename must be entered in the <b>Remote File</b> field.</li> <li>• <b>Service Profile Name</b>—Filename automatically becomes the name of the service profile that the vMedia Policy is associated with.</li> </ul> <p><b>Note</b></p> <ul style="list-style-type: none"> <li>• If you select <b>Service Profile Name</b> as the <b>Image Name</b> variable, the <b>Remote File</b> field is disabled.</li> <li>• If you select <b>Service Profile Name</b> as the <b>Image Name</b> variable, do not rename the service profile. Renaming the service profile can result in vMedia mount failure.</li> </ul>
<b>Remote File</b>	<p>Enter the full path to the ISO or other image file.</p> <p><b>Note</b> Ensure that the full path to the file begins with “/” after the share name.</p> <p>This field can contain the filename [with the file extension] only.</p>
<b>Remote Path</b>	<p>Enter the share name on the remote server, for example “Share”.</p>
<b>Username</b>	<p>Enter the username that Cisco UCS Manager should use to log in to the remote server.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>
<b>Password</b>	<p>Enter the password associated with the username.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>
<b>Remap on Eject</b>	<p>Click this checkbox to remap mounted vMedia after it is ejected.</p>

Name	Description
<b>Writable</b>	<p>Click this checkbox to configure the vMedia mount as writable. If this checkbox is cleared, the vMedia mount remains read-only.</p> <p>vMedia mounts are read-only by default.</p> <p>You can configure a vMedia mount as writable only when both the following conditions are met:</p> <ul style="list-style-type: none"> <li>• <b>Device Type</b> is <b>HDD</b></li> <li>• <b>Protocol</b> is <b>NFS</b> or <b>CIFS</b></li> </ul>

- Step 8** Click **OK**.  
The remote server details are listed in the **vMedia Mounts** area of the **Create vMedia Mount** dialog box.

---

### What to do next

Create a vMedia boot policy.

## Adding a vMedia Policy to a Service Profile

Before you can use Scriptable vMedia, you must add the vMedia and Boot Policies to a Service Profile. After the vMedia and Boot Policies are added to a service profile you can associate the service profile with a Cisco UCS server. The following procedure describes how to add a vMedia policy to a Service Profile.

### Before you begin

Configure the vMedia Policy you want to add to a service profile.

### Procedure

---

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization where you want to create the service profile.  
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Service Profile (expert)**.  
**The Unified Computing System Manager** pane displays.
- Step 5** In the **Name** field, enter a unique name that you can use to identify the service profile.  
  
This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), \_ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.  
  
This name must be unique within the organization or sub-organization in which you are creating the service profile.
- Step 6** From the **UUID Assignment** drop-down list, do one of the following:

Option	Description
<b>Select (pool default used by default)</b>	Assigns a UUID from the default UUID Suffix pool.  Continue with Step 8.
<b>Hardware Default</b>	Uses the UUID assigned to the server by the manufacturer.  If you choose this option, the UUID remains unassigned until the service profile is associated with a server. At that point, the UUID is set to the UUID value assigned to the server by the manufacturer. If the service profile is later moved to a different server, the UUID is changed to match the new server.  Continue with Step 8.
<b>XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX</b>	Uses the UUID that you manually assign.  Continue with Step 7.
<b>Pools <i>Pool_Name</i></b>	Assigns a UUID from the UUID Suffix pool that you select from the list at the bottom of the drop-down list.  Each pool name is followed by two numbers in parentheses that show the number of UUIDs still available in the pool and the total number of UUIDs in the pool.  If you do not want use any of the existing pools, but instead want to create a pool that all service profiles can access, continue with Step 4. Otherwise, continue with Step 8.

- Step 7** (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:
- a) In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.
  - b) To verify that the selected UUID is available, click the **here** link.
- Step 8** (Optional) If you want to create a new UUID Suffix pool to use to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.
- Step 9** (Optional) In the text box, enter a description of this service profile.
- The user-defined description for this service profile.
- Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
- Step 10** Click **Next**.
- Step 11** From the **vMedia** drop down list, choose one of the following:

Option	Description
Select vMedia Policy to use	Enables you to assign a vMedia policy to this service profile. Continue with Step 12.
Create a Specific vMedia Policy	Enables you to create a local vMedia policy that can only be accessed by this service profile.
vMedia Policies <i>Policy_Name</i>	Assigns an existing vMedia policy to the service profile. If you choose this option, Cisco UCS Manager displays the details of the policy.  If you do not want use any of the existing policies but instead want to create a policy that all service profiles can access, click <b>Create vMedia Policy</b> . Otherwise, choose a policy from the list and continue with Step 13.

**Step 12** If you created a new vmedia policy accessible to all service profiles and template, choose that policy from the **vMedia** drop down list .

**Step 13** Click **Next**.

## Viewing CIMC vMedia Policy

### Before you begin

vMedia Policies are configured.

### Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Policies > vMedia Policies**.
- Step 3** Expand the **vMedia Policies** node to view the list of **vMedia Policies**.
- Step 4** Double-click the name of a vMedia policy to view the properties for the selected **vMedia Mount**.  
On the **Properties** page, you can modify the properties used for the **vMedia Mounts**.