



Server-Related Policy Configuration

- [BIOS Settings, on page 1](#)
- [CIMC Security Policies, on page 107](#)
- [SPDM Security, on page 114](#)
- [Creating and Configuring a SPDM Security Certificate Policy using CLI, on page 115](#)
- [Loading an Outside SPDM Security Certificate Policy, on page 117](#)
- [Viewing the Certificate Inventory, on page 117](#)
- [Deleting a SPDM Policy, on page 119](#)
- [Graphics Card Policies, on page 119](#)
- [Configuring Local Disk Configuration Policies, on page 122](#)
- [Persistent Memory Modules, on page 137](#)
- [Scrub Policies, on page 137](#)
- [Configuring DIMM Error Management, on page 142](#)
- [Serial over LAN Policy, on page 144](#)
- [Server Autoconfiguration Policy, on page 146](#)
- [Server Discovery Policy, on page 148](#)
- [Server Inheritance Policies, on page 152](#)
- [Server Pool Policy, on page 154](#)
- [Server Pool Policy Qualification, on page 156](#)
- [Configuring vNIC/vHBA Placement Policies, on page 170](#)
- [CIMC Mounted vMedia, on page 183](#)

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
Properties	
Reboot on BIOS Settings Change set reboot-on-update	<p>When the server is rebooted after you change one or more BIOS settings.</p> <p>yes—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>no—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>
BIOS Setting	
Quiet Boot set quiet-boot-config quiet-boot	<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"> • disabled—The BIOS displays all messages and Option ROM information during boot. • enabled—The BIOS displays the logo screen, but does not display any messages or Option ROM information during boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
POST error pause set post-error-pause-config post-error-pause	<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"> • disabled—The BIOS continues to attempt to boot the server. • enabled—The BIOS pauses the attempt to boot the server and opens the Error Manager when a critical error occurs during POST. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Resume on AC power loss set resume-ac-on-power-loss-config resume-action	<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"> • stay-off—The server remains off until manually powered on. • last-state—The server is powered on and the system attempts to restore its last state. • reset—The server is powered on and automatically reset. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front panel lockout set front-panel-lockout-config front-panel-lockout	<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"> • disabled—The power and reset buttons on the front panel are active and can be used to affect the server. • enabled—The power and reset buttons are locked out. The server can only be reset or powered on or off from the CIMC GUI. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>CDN Control</p> <p>set consistent-device-name-control cdn-name</p>	<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"> • disabled—Consistent device naming is disabled for the BIOS policy. • enabled—Consistent device naming is enabled for the BIOS policy. This enables Ethernet interfaces to be named consistently. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>PCIe Slots CDN Control</p> <p>set consistent-device-name-control pcie-slot-cdn-name</p>	<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"> • disabled—Consistent device naming is disabled. This is the default option. • enabled—Consistent device naming is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
Intel Turbo Boost Tech set intel-turbo-boost-config turbo-boost	<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"> • disabled—The processor does not increase its frequency automatically. • enabled—The processor uses Turbo Boost Technology if required. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Enhanced Intel SpeedStep Tech set enhanced-intel-speedstep-config speed-step	<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"> • disabled—The processor never dynamically adjusts its voltage or frequency. • enabled—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
Intel HyperThreading Tech set hyper-threading-config hyper-threading	<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"> • disabled—The processor does not permit hyperthreading. • enabled—The processor allows for the parallel execution of multiple threads. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure the operating system supports this feature.</p>

Name	Description
<p>Intel Speed Select set-IntelSpeedSelect</p>	<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"> • base—The processor uses Base. • config1—The processor uses Config 1. • config2—The processor uses Config 2. • config3—The processor uses Config 3. • config4—The processor uses Config 4. <p>Note The values config1 and config2 are not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Core Multi Processing set core-multi-processing-config multi-processing</p>	<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"> • All—Enables multiprocessing on all logical processor cores. • 1 through n—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. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>

Name	Description
<p>Execute Disable Bit set execute-disable bit</p>	<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"> • disabled—The processor does not classify memory areas. • enabled—The processor classifies memory areas. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
<p>Intel Virtualization Technology set intel-vt-config vt</p>	<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"> • disabled—The processor does not permit virtualization. • enabled—The processor allows multiple operating systems in independent partitions. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note If you change this option, you must power cycle the server before the setting takes effect.</p>
<p>Hardware Prefetcher set processor-prefetch-config hardware-prefetch</p>	<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"> • disabled—The hardware prefetcher is not used. • enabled—The processor uses the hardware prefetcher when cache issues are detected. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note CPUPerformance must be set to Custom in order to specify this value. For any value other than Custom, this option is overridden by the setting in the selected CPU performance profile.</p>

Name	Description
<p>Adjacent Cache Line Prefetcher set processor-prefetch-config adjacent-cache-line-prefetch</p>	<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"> • disabled—The processor only fetches the required line. • enabled—The processor fetches both the required line and its paired line. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note must be set to Custom in order to specify this value. For any value other than Custom, this option is overridden by the setting in the selected CPU performance profile.</p>
<p>DCU Streamer Prefetch set processor-prefetch-config dcu-streamer-prefetch</p>	<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"> • disabled—The processor does not try to anticipate cache read requirements and only fetches explicitly requested lines. • enabled—The DCU prefetcher analyzes the cache read pattern and prefetches the next line in the cache if it determines that it may be needed. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>DCU IP Prefetcher set processor-prefetch-config dcu-ip-prefetch</p>	<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"> • disabled—The processor does not preload any cache data. • enabled—The DCU IP prefetcher preloads the L1 cache with the data it determines to be the most relevant. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
KTI Prefetch	<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"> • disabled—The processor does not preload any cache data. • enabled—The KTI prefetcher preloads the L1 cache with the data it determines to be the most relevant. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LLC Prefetch	<p>Whether the processor uses the LLC Prefetch mechanism to fetch the data into the LLC. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not preload any cache data. • enabled—The LLC prefetcher preloads the L1 cache with the data it determines to be the most relevant. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
XPT Prefetch	<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"> • disabled—The CPU does not use the XPT Prefetch option. • enabled—The CPU enables the XPT prefetcher option. • auto—The CPU auto enables the XPT prefetcher option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Direct Cache Access set direct-cache-access-config access	<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"> • auto—The CPU determines how to place data from I/O devices into the processor cache. • disabled—Data from I/O devices is not placed directly into the processor cache. • enabled—Data from I/O devices is placed directly into the processor cache. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>Processor C State set processor-c-state-config c-state</p>	<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"> • disabled—The system remains in a high-performance state even when idle. • enabled—The system can reduce power to system components such as the DIMMs and CPUs. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
<p>Processor C1E set processor-c1e-config c1e</p>	<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"> • disabled—The CPU continues to run at its maximum frequency in the C1 state. • enabled—The CPU transitions to its minimum frequency. This option saves the maximum amount of power in the C1 state. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Processor C3 Report set processor-c3-report-config processor-c3-report</p>	<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"> • enabled—The processor sends the C3 report to the OS. • disabled—The processor does not send the C3 report. • acpi-c2—The processor sends the C3 report using the advanced configuration and power interface (ACPI) C2 format. • acpi-c3—The processor sends the C3 report using the ACPI C3 format. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <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-c3, the server sets the BIOS value for that option to enabled.</p>

Name	Description
Processor C6 Report set processor-c6-report-config processor-c6-report	Whether the processor sends the C6 report to the operating system. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not send the C6 report. • enabled—The processor sends the C6 report. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor C7 Report set processor-c7-report-config processor-c7-report	Whether the processor sends the C7 report to the operating system. This can be one of the following: <ul style="list-style-type: none"> • c7—The processor sends the report using the C7 format. • c7s—The processor sends the report using the C7s format. • disabled—The processor does not send the C7 report. • enabled—The processor sends the C7 report. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Processor CMCI	Enables CMCI generation. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor disables CMCI. • enabled—The processor enables CMCI. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>CPU Performance set cpu-performance-config cpu-performance</p>	<p>Sets the CPU performance profile for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Custom—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. • enterprise—For M4 and higher servers, all prefetchers and data reuse are enabled. • high-throughput—Data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • hpc—All prefetchers are enabled and data reuse is disabled. This setting is also known as high-performance computing. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Max Variable MTRR Setting set max-variable-mtrr-setting-config processor-mtrr</p>	<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"> • auto-max—BIOS uses the default value for the processor. • 8—BIOS uses the number specified for the variable MTRR. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Local X2 APIC set local-x2-apic-config localx2-apic</p>	<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"> • disabled—Processor disables Local X2 APIC. • enabled—Processor enables Local X2 APIC. • xapic—Uses the standard xAPIC architecture. • x2apic—Uses the enhanced x2APIC architecture to support 32 bit addressability of processors. • auto—Automatically uses the xAPIC architecture that is detected. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Power Technology set processor-energy-config cpu-power-management	<p>Enables you to configure the CPU power management settings for the following options:</p> <ul style="list-style-type: none">• Enhanced Intel Speedstep Technology• Intel Turbo Boost Technology• Processor Power State C6 <p>Power Technology can be one of the following:</p> <ul style="list-style-type: none">• disabled—The server does not perform any CPU power management and any settings for the BIOS parameters mentioned above are ignored.• Energy_Efficient—The server determines the best settings for the BIOS parameters mentioned above and ignores the individual settings for these parameters.• performance—The server automatically optimizes the performance for the BIOS parameters mentioned above.• custom—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.• platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>Energy Performance set processor-energy-config energy-performance</p>	<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"> • performance — 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. • balanced-performance — The server provides all server components with enough power to keep a balance between performance and power. • balanced-energy — The server provides all server components with enough power to keep a balance between performance and power. • energy-efficient — The server provides all server components with less power to keep reduce power consumption. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note CPUPowerManagement must be set to Custom or the server ignores the setting for this parameter.</p>
<p>Frequency Floor Override set frequency-floor-override-config cpu-frequency</p>	<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"> • disabled— The CPU can drop below the maximum non-turbo frequency when idle. This option decreases power consumption but may reduce system performance. • enabled— The CPU cannot drop below the maximum non-turbo frequency when idle. This option improves system performance but may increase power consumption. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>P STATE Coordination set p-state-coordination-config p-state</p>	<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"> • hw-all—The processor hardware is responsible for coordinating the P-state among logical processors with dependencies (all logical processors in a package). • sw-all—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. • sw-any—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. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note CPUPowerManagement must be set to Custom or the server ignores the setting for this parameter.</p>
<p>DRAM Clock Throttling set dram-clock-throttling-config dram-clock-throttling</p>	<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"> • auto — CPU determines the DRAM Clock Throttling settings. • balanced— DRAM clock throttling is reduced, providing a balance between performance and power. • performance—DRAM clock throttling is disabled, providing increased memory bandwidth at the cost of additional power. • Energy_Efficient—DRAM clock throttling is increased to improve energy efficiency. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
External SSC enable	<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"> • disabled— Clock Spread Spectrum support is not available. • enabled— Clock Spread Spectrum support is always available. • platform-default — The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Channel Interleaving set interleave-config channel-interleave	<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"> • auto—The CPU determines what interleaving is done. • 1-way— • 2-way • 3-way • 4-way—The maximum amount of channel interleaving is used. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Rank Interleaving set interleave-config rank-interleave	<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"> • auto—The CPU determines what interleaving is done. • 1-way— • 2-way • 4-way • 8-way—The maximum amount of rank interleaving is used. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Sub NUMA Clustering	<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"> • auto— The BIOS determines what Sub NUMA clustering is done. • disabled— Sub NUMA clustering does not occur. This is the default option. • enabled— Sub NUMA clustering occurs. • platform-default — The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Interleaving set interleave-config memory-interleave	<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"> • none • channel • die • socket • auto—This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Demand Scrub set scrub-policies-config demand-scrub	<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"> • disabled— Single bit memory errors are not corrected. • enabled— Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>Patrol Scrub set scrub-policies-config patrol-scrub</p>	<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"> • disabled—The system checks for memory ECC errors only when the CPU reads or writes a memory address. • enabled—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. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>DCPMM Firmware Downgrade</p>	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—Support is disabled. • enabled—Support is enabled. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Configurable TDP Control</p>	<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"> • auto— Uses the rated TDP value of the processor. • manual—Allows you to customize the TDP value.

Name	Description
Altitude set altitude altitude-config	<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"> • auto—The CPU determines the physical elevation. • 300-m—The server is approximately 300 meters above sea level. • 900-m—The server is approximately 900 meters above sea level. • 1500-m—The server is approximately 1500 meters above sea level. • 3000-m—The server is approximately 3000 meters above sea level. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Package C State set package-c-state-limit-config package-c-state-limit	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <p>Note Cisco recommends setting Package C State Limit as no-limit or platform-default for C220 M4, C240 M4, C460 M4, and S3260 M4 servers.</p> <p>If you are changing the Package C State Limit token from any other value to no-limit, then ensure that the Power Technology is set to custom.</p>
CPU Hardware Power Management set cpu-hardware-power-management-config cpu-hardware-power-management	<p>Enables processor Hardware Power Management (HWPM). This can be one of the following:</p> <ul style="list-style-type: none"> • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • disabled—HWPM is disabled. • hwpm-native-mode—HWPM native mode is enabled. • hwpm-oob-mode—HWPM Out-Of-Box mode is enabled. • Native Mode with no Legacy (only GUI)

Name	Description
Energy Performance Tuning set power-performance-tuning-support power-performance-tuning-config	Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS. <ul style="list-style-type: none"> • bios— • os— • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Workload Configuration	This feature allows for workload optimization. The options are Balanced and I/O Sensitive: <ul style="list-style-type: none"> • balanced • io-sensitive—This is the default option. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Cisco recommends using Balanced.</p>
Core Performance Boost	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: <ul style="list-style-type: none"> • auto—The CPU automatically determines how to boost performance. • disabled—Core performance boost is disabled. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Uncore Frequency Scaling	Allows you configure the scaling of the uncore frequency of the processor. This can be one of the following: <ul style="list-style-type: none"> • enabled—Uncore frequency of the processor scales up or down based on the load. (Default.) • disabled—Uncore frequency of the processor remains fixed. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Refer to the Intel Dear Customer Letter (DCL) to know the fixed higher and lower values for Uncore Frequency Scaling.</p>

Name	Description
Configurable TDP Level	<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"> • normal—The CPU operates at its normal performance level. (Default.) • level1 • level1 <p>Note Refer to the Intel Dear Customer Letter (DCL) for the values for TDP level.</p>
UPI Link Speed set-qpilinkspeed	<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"> • auto—Automatically configures the optimal link speed. (Default) • 9.6gt/s (gigatransfers per second)—Configures the optimal link speed at 9.6GT/s • 10.4gt/s—Configures the optimal link speed at 10.4GT/s • 11.2gt/s—Configures the optimal link speed at 11.2GT/s • use per link setting <p>Note The value use per link setting is not supported on UCS M6 servers.</p>
Global C-state Control	<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"> • auto—The CPU automatically determines how to control IO-based C-state generation. • disabled—Global C-state control is disabled. • enabled—Global C-state control is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
L1 Stream HW Prefetcher	<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"> • auto—The CPU determines how to place data from I/O devices into the processor cache. • disabled—The hardware prefetcher is not used. • enabled—The processor uses the hardware prefetcher when cache issues are detected. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
L2 Stream HW Prefetcher	<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"> • auto—The CPU determines how to place data from I/O devices into the processor cache. • disabled—The hardware prefetcher is not used. • enabled—The processor uses the hardware prefetcher when cache issues are detected. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
AMD Memory Interleaving Size	<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"> • 1 KB • 2 KB • 256 Bytes • 512 Bytes • auto—The CPU determines the size of the memory block. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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

Name	Description
IOMMU	<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"> • auto—The CPU determines how map these addresses. • disabled—IOMMU is not used. • enabled—Address mapping takes place through the IOMMU. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SVM Mode	<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"> • disabled—The processor does not use SVM Technology. • enabled—The processor uses SVM Technology. This is the default option. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SMEE	<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"> • auto—This is the default option. • disabled—The processor does not use the SMEE function. • enabled—The processor uses the SMEE function. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
UPI Prefetch set-upi-prefetch	<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"> • enabled—The UPI prefetcher preloads the L1 cache with the data it determines to be the most relevant. • disabled—The processor does not preload any cache data. • auto—The processor enables the UPI prefetcher option.

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

Name	Description
Intel Dynamic Speed Select scope token-feature "IntelSpeedSelect Configuration" scope token-param IntelDynamicSpeedSelect	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: <ul style="list-style-type: none"> • enabled—Intel Dynamic Speed Select is enabled. • disabled—Intel Dynamic Speed Select is disabled.
IIO eDPC Support scope token-feature "Processor" scope token-param EdpcEn	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: <ul style="list-style-type: none"> • disabled—eDPC support is disabled. • on fatal errors—eDPC is enabled only for fatal errors. • on fatal and non-fatal errors—eDPC is enabled for both fatal and non-fatal errors.
Multikey Total Memory Encryption (MK-TME) scope token-feature "Processor" scope token-param EnableMktme	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: <ul style="list-style-type: none"> • enabled—Support is enabled. • disabled—Support is disabled.
SW Guard Extensions (SGX) scope token-feature "Processor" scope token-param EnableSgx	Allows you to enable Software Guard Extensions (SGX) feature. This can be one of the following: <ul style="list-style-type: none"> • enabled—Support is enabled. • disabled—Support is disabled.
Total Memory Encryption (TME) scope token-feature "Processor" scope token-param EnableTme	Allows you to provide the capability to encrypt the entirety of the physical memory of a system. This can be one of the following: <ul style="list-style-type: none"> • enabled—Support is enabled. • disabled—Support is disabled.
Select Owner EPOCH input type scope token-feature "Processor" scope token-param EpochUpdate	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: <ul style="list-style-type: none"> • sgx owner epoch activated— Does not change the current input type. • change to new random owner epochs—Changes EPOCH to a system generated random number • manual user defined owner epochs—Changes the EPOCH seed to a hexadecimal value that you enter.

Name	Description
<p>Enhanced CPU Performance scope token-feature "CpuPerfEnhancement" scope token-param CpuPerfEnhancement</p>	<p>Enhances CPU performance by adjusting server settings automatically. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not run with this functionality. This is the default option. • auto—Allows to adjust server settings to increase the processor performance. <p>Note</p> <ul style="list-style-type: none"> • Enabling this functionality may increase power consumption. • The server should meet the following requirements in order to use this functionality: <ul style="list-style-type: none"> • The server should not contain Barlow Pass DIMMs. • 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. • No GPU cards are present in the server.
<p>UPI Link Enablement scope token-feature "UPI Link Enablement" scope token-param UPILinkEnablement</p>	<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"> • auto—This is the default option. • 1 • 2
<p>UPI Power Manangement scope token-feature "UPI Power Manangement" scope token-param UPIPowerManagement</p>	<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"> • enabled—Enables the processor to support this functionality. • disabled—Disables the processor to support this functionality. This is the default option.
<p>C1 Auto UnDemotion scope token-feature "C1 Auto UnDemotion" scope token-param C1AutoDemotion</p>	<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"> • enabled—Enables the processor to support this functionality. This is the default option. • disabled—Disables the processor to support this functionality.

Name	Description
C1 Auto Demotion scope token-feature "C1 Auto Demotion" scope token-param C1AutoDemotion	If enabled, CPU automatically demotes to C1 based on un-core auto-demote information. This can be one of the following: <ul style="list-style-type: none"> • enabled—Enables the processor to support this functionality. This is the default option. • disabled—Disables the processor to support this functionality.
CPU Downcore control 7xx3 scope token-feature "Processor" scope token-param CbsCpuCoreCtrl	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: <ul style="list-style-type: none"> • auto—The CPU determines how many cores need to be enabled. This is the default option • one (1+0)—One core enabled on one CPU complex • two (2+0)—Two core enabled on one CPU complex • three (3+0)—Three core enabled on one CPU complex. • four (4+0)—Four core enabled on one CPU complex. • five (5+0)—Five core enabled on one CPU complex • six (6+0)—Six core enabled on one CPU complex • seven (7+0)—Seven core enabled on one CPU complex <p>Note This token is applicable only for the servers with 7xx3 Model processors.</p>
Fixed SOC P-State scope token-feature "Processor" scope token-param CbsCmnFixedSocPstate	This option defines the target P-state when APBDIS (to disable Algorithm Performance Boost (APB)) is set. The P-x specify a valid P-state for the processor installed. This can be one of the following: <ul style="list-style-type: none"> • auto—Sets a valid P-state suitable for the processor. This is the default option. • p0—Highest-performing SOC P-state • p1—Next-highest-performing SOC P-state • p2—Next-highest-performing SOC P-state • p3—Minimum SOC power P-state

Name	Description
APBDIS scope token-feature "Processor" scope token-param CbsCmnApbdis	Allows you to select the Algorithm Performance Boost (APB) Disable value for the SMU. This can be one of the following: <ul style="list-style-type: none"> • auto—Sets an auto ApbDis for the SMU. This is the default option. • 0—Clear ApbDis to SMU • 1—Set ApbDis to SMU
CCD Control scope token-feature "Processor" scope token-param CbsCpuCcdCtrlSsp	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: <ul style="list-style-type: none"> • auto—The maximum CCDs provided by the processor is enabled. This is the default option. • 2 ccds • 3 ccds • 4 ccds • 6 ccds
Cisco xGMI Max Speed scope token-feature "Processor" scope token-param CiscoXgmiMaxSpeed	This option enables 18 Gbps XGMI link speed. This can be one of the following: <ul style="list-style-type: none"> • disabled—The feature is disabled. This is the default option. • enabled—The feature is enabled.
ACPI SRAT L3 Cache As NUMA Domain scope token-feature "Processor" scope token-param CbsDfCmnAcpiSratL3Numa	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: <ul style="list-style-type: none"> • auto—Set to auto mode. This is the default option. • disabled—Use NPS settings for domain configuration. • enabled—Each CCX is declared to be in its own domain.
Streaming Stores Control scope token-feature "Processor" scope token-param CbsCmnCpuStreamingStoresCtrl	Enables the streaming stores functionality. This can be one of the following: <ul style="list-style-type: none"> • auto—Set to auto mode. This is the default option. • disabled—Feature is disabled. • enabled—Feature is enabled.

Name	Description
DF C-States scope token-feature "Processor" scope token-param CbsCmnGnbSMUDfCstates	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: <ul style="list-style-type: none"> • auto—Set to auto mode. This is the default option. • disabled—This option is turned off, long period of idleness are not expected so no power savings would be achieved. • enabled—This option is active, saving power when the system is idle.
SEV-SNP Support scope token-feature "Processor" scope token-param CbsSevSnpSupport	Allows you to enable Secure Nested Paging feature. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not use the SEV-SNP function. This is the default option. • enabled—The processor uses the SEV-SNP function.
Efficiency Mode Enable scope token-feature "Processor" scope token-param CbsCmnEfficiencyModeEn	Allows you to configure power consumption based on efficiency. This can be one of the following: <ul style="list-style-type: none"> • auto—The CPU automatically uses default settings. This is the default option. • enabled—Efficiency mode is enabled.
SNP Memory Coverage scope token-feature "Processor" scope token-param CbsDbgCpuSnpMemCover	Allows you to configure SNP memory coverage. This can be one of the following: <ul style="list-style-type: none"> • auto—System decides the memory coverage. This is the default option. • disabled—The processor does not use this function. • enabled—This feature is enabled. • custom—Custom size can be defined in SNP Memory Size to Cover.
SNP Memory Size to Cover in MB scope token-feature "Processor" scope token-param CbsDbgCpuSnpMemSizeCover	Allows you to configure SNP memory size. The value can range from 0-1048576. 0 is the default option.

Name	Description
SMT Mode scope token-feature "Processor" scope token-param SmtMode	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: <ul style="list-style-type: none"> • auto—The processor allows for the parallel execution of multiple threads. • enabled—The processor allows permit multithreading. This is the default option. • disabled—The processor allows permit multithreading.
CPCC scope token-feature "Processor" scope token-param CbsCmnGnbSMUCPPC	Allows you to configure Collaborative Processor Performance Control. This can be one of the following: <ul style="list-style-type: none"> • auto—The CPU automatically uses default CPPC settings. This is the default option. • disabled—Feature is disabled. • enabled—Collaborative Processor Performance is enabled.
Downcore control 7xx2 scope token-feature "Processor" scope token-param CbsCmnCpuGenDowncoreCtrl	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: <ul style="list-style-type: none"> • auto—The CPU determines how many cores need to be enabled. This is the default option. • two (1+1)—Two cores enabled on one CPU complex. • four (2+2)—Four cores enabled on one CPU complex. • six (3+3)—Six cores enabled on one CPU complex.
Processor EPP Profile set processor epp profile	Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following: <ul style="list-style-type: none"> • performance • balanced performance—This is the default option. • balanced power • power

Name	Description
Autonomous Core C-state set processor autonomous core c-state	Enables CPU Autonomous C-State, which converts the HALT instructions to the MWAIT instructions. This can be one of the following: <ul style="list-style-type: none"> • disabled—This is the default option. • enabled
Energy Efficient Turbo set energy efficient turbo	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: <ul style="list-style-type: none"> • disabled—This is the default option. • enabled
Hardware P-States set hardware p-states	Enables processor Hardware P-State. This can be one of the following: <ul style="list-style-type: none"> • disabled—HWPM is disabled. • hwpm native modeHWPM Native Mode—HWPM native mode is enabled. This is the default option. • hwpm oob modeHWPM OOB Mode—HWPM Out-of-Box mode is enabled. • native mode with no legacyNative Mode with no Legacy
Energy/Performance Bias Config set energy/performance	Allows you to determine whether system performance or energy efficiency is more important on this server. This can be one of the following: <ul style="list-style-type: none"> • performance—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. • balaced performanceBalanced Performance—The server provides all server components with enough power to keep a balance between performance and power. This is the default option. • balanced powerBalanced Power—The server provides all server components with enough power to keep a balance between performance and power. • powerPower—The server provides all server components with maximum power to keep reduce power consumption.

Name	Description
Power Performance Tuning set power performance	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: <ul style="list-style-type: none"> • bios—Chooses BIOS for energy performance tuning. • osOS—Chooses OS for energy performance tuning. This is the default option. • peciPECI—Chooses Peci for energy performance tuning.
Cores Enabled set cores enabled	Allows you to disable one or more of the physical cores on the server. This can be one of the following: <ul style="list-style-type: none"> • all—Enables all physical cores. This also enables Hyper Threading on the associated logical processor cores. • 1 through 481 through 48—Specifies the number of physical processor cores that can run on the server. Each physical core has an associated logical core.
Hyper-Threading [All] set hyper-threading-all	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: <ul style="list-style-type: none"> • disabled—The processor does not permit hyperthreading. • enabledEnabled—The processor allows for the parallel execution of multiple threads.
SpeedStep (Pstates) set speedstep (pstates)	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: <ul style="list-style-type: none"> • disabled—The processor never dynamically adjusts its voltage or frequency. • enabledEnabled—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power.

Name	Description
Boot Performance Mode set boot performance mode	Allows the user to select the BIOS performance state that is set before the operating system handoff. This can be one of the following: <ul style="list-style-type: none"> • max performance—Processor P-state ratio is maximum. • max efficientMax Efficient—Processor P-state ratio is minimum. • set by intel nmSet by Intel NM—Processor P-state ratio is set by Intel.
EIST PSD Function set eist psd function	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: <ul style="list-style-type: none"> • hw all—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. • sw all—The OS Power Manager coordinates the P-state among logical processors with dependencies and initiates the transition on all of those Logical Processors.
Turbo Mode set eist psd function	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: <ul style="list-style-type: none"> • disabled—The processor does not increase its frequency automatically. • enabled—The processor utilizes Turbo Boost Technology if required. This is the default option.
Extended APIC set extended apic	Allows you to enable or disable extended APIC support. This can be one of the following: <ul style="list-style-type: none"> • disabled—This is the default option. • enabled.

Name	Description
Memory Interleaving Size set memory interleaving	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: <ul style="list-style-type: none"> • 1 KB • 2 KB • 4 KB • 256 Bytes • 512 Bytes • auto—The CPU determines the size of the memory block. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
UPI Link Frequency Select set upi link frequency select	Allows you to enable or disable extended APIC support. This can be one of the following: <ul style="list-style-type: none"> • auto—This option configures the optimal link speed automatically. This is the default option. • 9.6gt/s—This option configures the optimal link speed at 9.6GT/s. • 10.4gt/s—This option configures the optimal link speed at 10.4GT/s. • 11.2gt/s—This option configures the optimal link speed at 10.4GT/s.
X2APIC Opt Out set X2ApicOptOut	Prevents the OS from enabling extended xAPIC (x2APIC) mode when the OS is not working with x2APIC. This can be one of the following: <ul style="list-style-type: none"> • disabled—Use the Extended xAPIC (x2APIC) mode. This is the default option. • enabled—Opt out from Extended xAPIC (x2APIC) mode.

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
Intel VT for directed IO set intel-vt-directed-io-config vtd	<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"> • disabled—The processor does not use virtualization technology. • enabled—The processor uses virtualization technology. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.</p>
Intel VTD interrupt Remapping set intel-vt-directed-io-config interrupt-remapping	<p>Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not support remapping. • enabled—The processor uses VT-d Interrupt Remapping as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD coherency support set intel-vt-directed-io-config coherency-support	<p>Whether the processor supports Intel VT-d Coherency. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not support coherency. • enabled—The processor uses VT-d Coherency as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD ATS support set intel-vt-directed-io-config ats-support	<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"> • disabled—The processor does not support ATS. • enabled—The processor uses VT-d ATS as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Intel VTD pass through DMA support set intel-vt-directed-io-config passthrough-dma	Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not support pass-through DMA. • enabled—The processor uses VT-d Pass-through DMA as required. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
PCIe ARI Support scope token-feature "PCIe ARI Support" scope token-param "PCIeARISupport"	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: <ul style="list-style-type: none"> • disabled—PCIe ARI Support is not available. • enabled—PCIe ARI Support is available. • auto—PCIe ARI Support is in auto mode. This is the default option.
IPv4 PXE Support scope token-feature "IPv4 PXE Support" scope token-param "IPv4PXEsupport"	Enables or disables IPv4 support for PXE. This can be one of the following: <ul style="list-style-type: none"> • disabled—IPv6 PXE support is not available. • enabled—IPv6 PXE support is available. This is the default option.
IPv4 HTTP Support scope token-feature "HTTP BOOT" scope token-param "IPV4HTTP"	Enables or disables IPv4 support for HTTP. This can be one of the following: <ul style="list-style-type: none"> • disabled—IPv4 HTTP support is not available. • enabled—IPv4 HTTP support is available. This is the default option.

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

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
Partial Cache Line Sparing scope token-feature "Partial Cache Line Sparing" scope token-param PartialCacheLineSparing	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: <ul style="list-style-type: none"> • disabled—Support is disabled. • enabled—Support is enabled.

Name	Description
UMA scope token feature "UMA" scope token-param UmaBasedClustering	Allows you to set UMA settings. This can be one of the following: <ul style="list-style-type: none"> • disable-all2-all • hemisphere-2-clusters
Memory Thermal Throttling Mode scope token-feature "Memory Thermal Throttling Mode" scope token-param MemoryThermalThrottling	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: <ul style="list-style-type: none"> • CLTT with PECCI—Closed Loop Thermal Throttling (CLTT) with Platform Environment Control Interface (PECCI). This is the default option. • disabled. <p>Note It is recommended to leave this setting in the default state of CLTT with PECCI</p>
Enhanced Memory Test scope token-feature "Advanced Memory Test" scope token-param AdvancedMemTest	Enables enhanced memory tests during the system boot and increases the boot time based on the memory. This can be one of the following: <ul style="list-style-type: none"> • auto—This is the default option. <p>Note It is recommended to leave this setting in the default state of auto.</p> <ul style="list-style-type: none"> • enabled • disabled <p>Note <ul style="list-style-type: none"> • This BIOS token name modified from Advanced Memory Test to Enhanced Memory Test for M6 servers. </p>
Transparent Secure Memory Encryption (TSME) scope token-feature "Processor" scope token-param TSME	Provides transparent hardware memory encryption of all data stored on system memory. This can be one of the following: <ul style="list-style-type: none"> • enabled • disabled • auto—This is the default option. <p>Note It is recommended to leave this setting in the default state of auto to mitigate Rowhammer-style attacks.</p>

Name	Description
Secure Encrypted Virtualization (SEV) scope token-feature "Processor" scope token-param SEV	Enables running encrypted virtual machines (VMs) in which the code and data of the VM are isolated. This can be one of the following: <ul style="list-style-type: none"> • 253 ASIDs • 509 ASIDs • auto—This is the default option. <p>Note It is recommended to leave this setting in the default state of auto to mitigate Rowhammer-style attacks.</p>
DRAM SW Thermal Throttling scope token-feature "Processor" scope token-param DramSwThermalThrottling	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: <ul style="list-style-type: none"> • enabled • disabled—This is the default option. <p>Note It is recommended to leave this setting in the default state of disabled to mitigate Rowhammer-style attacks.</p>
Memory Refresh Rate scope token-feature "Memory Refresh Rate" scope token-param MemoryRefreshRate	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: <ul style="list-style-type: none"> • 1x-Refresh • 2x-Refresh—This is the default option.
Panic and High Watermark scope token-feature "Panic and High Watermark" scope token-param PanicHighWatermark	Controls the delayed refresh capability of the memory controller. This can be one of the following: <ul style="list-style-type: none"> • High—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 Panic and the memory controller pauses the normal memory transactions until all the postponed refresh commands are executed. • Low—This is the default option. The memory controller is not allowed to postpone refresh commands. <p>Note It is recommended to leave this setting in the default state (Low) which will help to reduce susceptibility to Rowhammer-style attacks.</p>

Name	Description
<p>Memory RAS configuration set memory-ras-config ras-config</p>	<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"> • maximum-performance—Optimizes the system performance and disables all the advanced RAS features. • mirroring—System reliability is optimized by using half the system memory as backup. This mode is used for UCS M4 and lower blade servers. • lockstep—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. • Mirror Mode 1LM—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. • Partial Mirror Mode 1LM—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. • sparing—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. • adddc-sparing—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. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
NUMA optimized set numa-config numa-optimization	Whether the BIOS supports NUMA. This can be one of the following: <ul style="list-style-type: none"> • disabled—The BIOS does not support NUMA. • enabled—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. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Post Package Repair scope token-feature "PostPackageRepair" scope token-param PostPackageRepair	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: <ul style="list-style-type: none"> • disabled—The BIOS does not support selecting PPR Type. • hard-ppr—This results in a permanent remapping of damaged storage cells. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Size Limit in GB set memory-size-limit	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.
Mirroring Mode set memory-mirroring-mode mirroring-mode	Memory mirroring enhances system reliability by keeping two identical data images in memory. This option is only available if you choose the mirroring option for Memory RAS Config . It can be one of the following: <ul style="list-style-type: none"> • inter-socket—Memory is mirrored between two Integrated Memory Controllers (IMCs) across CPU sockets. • intra-socket—One IMC is mirrored with another IMC in the same socket. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>Sparing Mode set memory-sparing-mode sparing-mode</p>	<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 sparing option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • dimm-sparing—One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • rank-sparing—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. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>LV DDR Mode set lv-dimm-support-config lv-ddr-mode</p>	<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"> • auto—The CPU determines whether to prioritize low voltage or high frequency memory operations. • power-saving-mode—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. • performance-mode—The system prioritizes high frequency operations over low voltage operations. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>DRAM Refresh Rate set dram-refresh-rate-config dram-refresh</p>	<p>The refresh interval rate for internal memory. This can be one of the following:</p> <ul style="list-style-type: none"> • 1x • 2x • 3x • 4x • auto • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
DDR3 Voltage Selection set ddr3-voltage-config ddr3-voltage	The voltage to be used by the dual-voltage RAM. This can be one of the following: <ul style="list-style-type: none"> • ddr3-1500mv • ddr3-1350mv • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Partial Memory Mirror Mode set memory-mirroring-mode mirroring-mode	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: <ul style="list-style-type: none"> • Disabled—Partial Memory Mode is disabled. This is the default option. • Percentage—The amount of memory to be mirrored in the Partial Memory Mode is defined as a percentage of the total memory. • Value in GB—The amount of memory to be mirrored in the Partial Memory Mode is defined in GB. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Partial Memory Mirror Mode is mutually exclusive to standard Mirroring Mode.</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>
Partial Mirror Percentage	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 %.
Partial Mirror1 Size in GB	Limits the amount of memory in Partial Mirror1 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
Partial Mirror2 Size in GB	Limits the amount of memory in Partial Mirror2 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
Partial Mirror3 Size in GB	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
Partial Mirror4 Size in GB	Limits the amount of memory in Partial Mirror4 in GB. This can range from 0 GB to 65535 GB in increments of 1 GB.
Volatile Memory Mode scope token-feature "VolMemoryMode" scope token-param VolMemoryMode	Allows the memory mode configuration. This can be any of the following: <ul style="list-style-type: none"> • 1lm—Configures 1 Layer Memory(1LM) • 2lm—Configures 2 Layer Memory(1LM)
Memory Bandwidth Boost scope token-feature "MemoryBandwidthBoost" scope token-param MemoryBandwidthBoost	Allows to boost the memory bandwidth. This can be one of the following: <ul style="list-style-type: none"> • enabled • disabled
Burst and Postponed Refresh scope token-feature "Processor" scope token-param BurstAndPostponedRefresh	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: <ul style="list-style-type: none"> • enabled • disabled—This is the default option. <p>Note It is recommended to leave this setting in the default state of disabled to mitigate Rowhammer-style attacks.</p>
LLC Dead Line scope token-feature "LLC Dead Line" scope token-param LLCAlloc	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: <ul style="list-style-type: none"> • enabled—Allows the LLC to fill dead lines into the LLC if there is free space available. This is the default option. • disabled—The dead lines are always dropped and are never filled into the LLC. • auto—The CPU determines the LLC dead line allocation

Name	Description
<p>XPT Remote Prefetch scope token-feature "XPT Remote Prefetch" scope token-param XPTRemotePrefetch</p>	<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"> • enabled • disabled • auto—The CPU determines the functionality. This is the default option.
<p>Virtual NUMA scope token-feature "Virtual Numa" scope token-param VirtualNuma</p>	<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"> • enabled—The functionality is enabled. • disabled—The functionality is disabled. This is the default option.
<p>Above 4G Decoding scope token-feature "Above 4G Decoding" scope token-param Above 4G Decoding</p>	<p>Enables or disables MMIO above 4GB or not. This can be one of the following:</p> <ul style="list-style-type: none"> • enabled—The server maps I/O of 64-bit PCI devices to 4GB or greater address space. This is the default option. • disabled—The server does not map I/O of 64-bit PCI devices to 4GB or greater address space.
<p>NUMA Nodes per Socket scope token-feature "nodes-per-socket" scope token-param nodes-per-socket</p>	<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"> • auto—Number of channels is set to auto. This is the default option. • nps0—Zero NUMA node per socket. • nps1—One NUMA node per socket. • nps2—Two NUMA nodes per socket, one per Left/Right Half of the SoC. • nps4—Four NUMA nodes per socket, one per Quadrant.

Name	Description
Select PPR Type scope token-feature "select ppr type"	Supports Hard-PPR , which permanently remaps accesses from a designated faulty row to a designated spare row. <ul style="list-style-type: none"> • hard ppr—Support is enabled. This is the default option. <p>Note Hard PPR can be used only when Memory RAS Configuration is set to ADDDC Sparing. For other RAS selections, this setting should be set to Disabled.</p> <ul style="list-style-type: none"> • disabled—Support is disabled.
Select Memory RAS Configuration scope token-feature "select memory ras configuration"	Determines how the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following: <ul style="list-style-type: none"> • Mirror Mode 1LM—System reliability is optimized by using half the system memory as backup. • ADDDC sparing—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. • Partial Mirror Mode 1LM—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. • maximum performance—System performance is optimized.
NUMA scope token-feature "numa"	Whether the BIOS supports Non-Uniform Memory Access (NUMA). This can be one of the following: <ul style="list-style-type: none"> • enabled—Support is enabled. • disabled—Support is disabled.
Operation Mode scope token-feature "operation mode"	Allows you to set the Operation Mode. This can be one of the following: <ul style="list-style-type: none"> • test only—Support is enabled. • test and repair—Support is disabled.

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
NVM Performance Setting set NvmdimmPerformConfig	<p>NVM Performance Setting 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"> • BW Optimized—Optimized for DDR and DDRT BW. This is the default option. • Latency Optimized—Better DDR latency in the presence of DDRT BW. • Balanced Profile—Optimized for Memory mode.
CR QoS set crqos	<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"> • disabled—Feature disabled. This is the default option. • recipe 1—6 modules, 4 modules per socket optimized • recipe 2—2 modules per socket optimized • recipe 3—1 module per socket optimized • mode 0 - disable the pmem qos feature • mode 1 - m2m qos enable;cha qos disable • mode 2 - m2m qos enable;cha qos enable <p>Note The values disabled, recipe 1, recipe 2, and recipe 3 are not supported on Cisco UCS M6 servers</p>

Name	Description
<p>CR FastGo Config set CrfastgoConfig</p>	<p>CR FastGo Config 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"> • auto—Same as Option 1. Disables FastGO. Recommended for DDRT. This is the default option (not Default). • default—Enables FastGO. • option 1—Disables FastGO. • option 2, option 3, option 4, option 5—Not applicable. • enable optimization • disable optimization <p>Note The values enable optimization, disable optimization, and auto are supported on Cisco UCS M6 servers</p>
<p>Snoopy mode for AD set SnoopyModeForAD</p>	<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"> • enabled • disabled This is the default option.

Name	Description
Snoopy mode for 2LM set SnoopyModeFor2LM	<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 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"> • enabled • disabled This is the default option.
eADR Support scope token-feature "EadrSupport" scope token-param EadrSupport	<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 power fail protected domain. This can be any of the following:</p> <ul style="list-style-type: none"> • enabled • disabled • auto—This is the default option.

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
Serial port A enable set serial-port-a-config serial-port-a	<p>Whether serial port A is enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The serial port is disabled. • enabled—The serial port is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
<p>Make Device Non Bootable set usb-boot-config make-device-non-bootable</p>	<p>Whether the server can boot from a USB device. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The server can boot from a USB device. • enabled—The server cannot boot from a USB device. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Legacy USB Support set usb-boot-config legacy-support</p>	<p>Whether the system supports legacy USB devices. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—USB devices are only available to EFI applications. • enabled—Legacy USB support is always available. • auto—Disables legacy USB support if no USB devices are connected. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>USB Idle Power Optimizing Setting set usb-system-idle-power-optimizing-setting-config usb-idle-power-optimizing</p>	<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"> • high-performance—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. • lower-idle-power—The USB System Idle Power Optimizing setting is enabled, because power savings are preferred over optimal performance. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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

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

Name	Description
USB Port SD Card set usb-port-config usb-sdcard	Whether the SD card drives are enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—Disables the SD card drives. The SD card drives are not detected by the BIOS and operating system. • enabled—Enables the SD card drives. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
USB Port VMedia set usb-port-config usb-vmedia	Whether the virtual media devices are enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—Disables the vMedia devices. • enabled—Enables the vMedia devices. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
All USB Devices set all-usb-devices-config all-usb	Whether all physical and virtual USB devices are enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—All USB devices are disabled. • enabled—All USB devices are enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
xHCI Mode set usb-configuration-select-config xhci-enable-disable	Whether xHCI mode is enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—xHCI mode is disabled. • enabled—xHCI mode is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Port:M.2 Storage set usb port:m.2	Whether the USB Port:M.2 Storage are enabled or disabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—Disables USB Port:M.2 Storage. • enabled—Enables USB Port:M.2 Storage. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
Maximum memory below 4GB set max-memory-below-4gb-config max-memory	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: <ul style="list-style-type: none"> • disabled—Does not maximize memory usage. Choose this option for all operating systems with PAE support. • enabled—Maximizes memory usage below 4GB for an operating system without PAE support. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory mapped IO above 4GB set memory-mapped-io-above-4gb-config memory-mapped-io	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: <ul style="list-style-type: none"> • disabled—Does not map I/O of 64-bit PCI devices to 4GB or greater address space. • enabled—Maps I/O of 64-bit PCI devices to 4GB or greater address space. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>VGA Priority set vga-priority-config vga-priority</p>	<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"> • onboard—Priority is given to the onboard VGA device. BIOS post screen and OS boot are driven through the onboard VGA port. • offboard—Priority is given to the PCIE Graphics adapter. BIOS post screen and OS boot are driven through the external graphics adapter port. • onboard-vga-disabled—Priority is given to the PCIE Graphics adapter, and the onboard VGA device is disabled. <p>Note The vKVM does not function when the onboard VGA is disabled.</p> <ul style="list-style-type: none"> • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note Only onboard VGA devices are supported with Cisco UCS B-Series servers.</p>
<p>ASPM Support set aspm-support-config aspm-support</p>	<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"> • disabled—ASPM support is disabled in the BIOS. • auto—The CPU determines the power state. • forcel0—Force all links to L0 standby (L0s) state. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>BME DMA Mitigation Support set bme-dma-config</p>	<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"> • disabled—PCI BME bit is disabled in the BIOS. • enabled—PCI BME bit is enabled in the BIOS. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
QPI Link Frequency Select set qpi-link-frequency-select-config qpi-link-frequency-mt-per-sec	The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s). This can be one of the following: <ul style="list-style-type: none"> • 6400 • 7200 • 8000 • 9600 • auto—The CPU determines the QPI link frequency. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
QPI Snoop Mode set qpi-snoop-mode vpqpisnoopmode	This can be one of the following: <ul style="list-style-type: none"> • home-snoop—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. • cluster-on-die—This mode is available only for processors that have 10 or more cores. It is the best mode for highly NUMA optimized workloads. • home-directory-snoop-with-osb • early-snoop—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. • auto —The CPU determines the QPI Snoop mode. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
Trusted Platform Module (TPM) Support set trusted-platform-module-config tpm-support	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: <ul style="list-style-type: none"> • disabled—Disables TPM. • enabled—Enables TPM. • platform-default—Enables TPM.
Intel Trusted Execution Technology (TXT) Support set intel-trusted-execution-technology-config txt-support	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: <ul style="list-style-type: none"> • disabled—Disables TXT. This is default option. • enabled—Enables TXT. • platform-default—Disables TXT. When you only enable TXT, it implicitly enables TPM, VT, and VTDio.
SHA-1 PCR Bank scope token-feature "Trusted Platform Module" scope token-param SHA1PCRBank	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: <ul style="list-style-type: none"> • disabled—Disables SHA-1 PCR Bank. • enabled—Enables SHA-1 PCR Bank. This is the default option.
SHA-256 PCR Bank scope token-feature "Trusted Platform Module" scope token-param SHA256PCRBank	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: <ul style="list-style-type: none"> • disabled—Disables SHA-256 PCR Bank. • enabled—Enables SHA-256 PCR Bank. This is the default option.
Trusted Platform Module State scope token-feature "Trusted Platform Module" scope token-param "Trusted Platform Module state"	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: <ul style="list-style-type: none"> • disabled—The server does not use the TPM. • enabled—The server uses the TPM. This is the default option.

Name	Description
TPM Pending Operation scope token-feature "TPM Pending Operation" scope token-param "TPM Pending Operation"	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"> • none—No action. This is the default option. • tpmclear—Clear the pending operations.
TPM Minimal Physical Presence scope token-feature "Trusted Platform Module" # scope token-param TpmPpiRequired # show token-settings expand	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"> • disabled—Disables TPM Minimal Physical Presence. This is default option. • enabled—Enables TPM Minimal Physical Presence. • platform-default—Disables TPM Minimal Physical Presence.
DMA Control Opt-In Flag scope token-feature "Trusted Platform Module" # scope token-param "DmaCtrlOptIn" token-param # show token-settings	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"> • disabled—Disables DMA Control Opt-In Flag. This is default option. • enabled—Enables DMA Control Opt-In Flag. • platform-default—Disables DMA Control Opt-In Flag.
Security Dev. Support set TpmSupport	Enables or disables BIOS support for the security device. This can be one of the following: <ul style="list-style-type: none"> • disabled—OS will not show the security device. • enabled—OS will show the security device. This is default option.

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
PCIe Slot SAS OptionROM set slot-option-rom-enable-config pcie-sas	Whether Option ROM is available on the SAS port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot <i>n</i> Link Speed set slot-link-speed-config pcie-slot <i>n</i> -link-speed	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: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot <i>n</i> OptionROM set slot-option-rom-enable-config slot <i>n</i> -option-rom-enable	Whether Option ROM is available on the port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot HBA OptionROM set slot-option-rom-enable-config pcie-hba	Whether Option ROM is available on the HBA port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
PCIe Slot MLOM OptionROM set slot-option-rom-enable-config pcie-mlom	Whether Option ROM is available on the MLOM port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot Nx OptionROM set slot-option-rom-enable-config pcie-nx	Whether Option ROM is available on the port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe 10G LOM 2 Link set lom-ports-config pcie-lom2-link	Whether Option ROM is available on the 10G LOM port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCI ROM CLP set pci-rom-clp-support pci-rom-clp-config	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. <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SIOC1 Option ROM set sioc1-optionrom-config sioc1-optionrom	Whether the server can use Option ROM present in System IO Controller 1 (SIOC1). This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOC2 Option ROM set sioc2-optionrom-config sioc2-optionrom	Whether the server can use Option ROM present in System IO Controller 2 (SIOC2). This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBMEZZ1 Option ROM set sbmezz1-optionrom-config sbmezz1-optionrom	Whether the server can use Option ROM present in SBMezz1 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBMEZZ2 Option ROM set sbmezz2-optionrom-config sbmezz2-optionrom	Whether the server can use Option ROM present in SBMezz2 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IOESlot1 OptionROM set ioeslot1-optionrom-config ioeslot1-optionrom	Whether option ROM is enabled on the IOE slot 1. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOEMEZ1 OptionROM set ioemezz1-optionrom-config ioemezz1-optionrom	Whether option ROM is enabled on the IOE Mezz1. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOE Slot2 Option ROM set ioeslot2-optionrom-config ioeslot2-optionrom	Whether option ROM is enabled on the IOE slot 2. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IO ENVME1 Option ROM set ioenvme1-optionrom-config ioenvme1-optionrom	Whether option ROM is enabled on the IOE NVMe1. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IO ENVME2 Option ROM set ioenvme2-optionrom-config ioenvme2-optionrom	Whether option ROM is enabled on the IOE NVMe2. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBNVME1 Option ROM set sbnvme1-optionrom-config sbnvme1-optionrom	Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot MRAID-<i>n</i> OptionROM set Pcie SlotMRAID<i>n</i>OptionROM	Whether Option ROM is available on the MRAID port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot RAID OptionROM set Pcie SlotRAIDOptionROM	Whether Option ROM is available on the RAID port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Rear NVME <i>n</i> Link Speed set Pcie SlotRearNvmeILinkSpeed	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. <p>Note</p> <ul style="list-style-type: none"> • For <i>Rear NVME 1 Link Speed</i> and <i>Rear NVME 2Link Speed</i>, the value enabled is not supported on Cisco UCS M6 servers. • For <i>Rear NVME 3 Link Speed</i> and <i>Rear NVME 4Link Speed</i>, the value enabled is available but has no effect at the BIOS level if selected. <ul style="list-style-type: none"> • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front NVME <i>n</i> Link Speed set Pcie SlotFrontNvmenLinkSpeed	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. This is the default option. • enabled—The maximum speed is restricted. <p>Note For <i>Front NVME 1 Link Speed</i> and <i>Front NVME 2 Link Speed</i>, the value enabled is available but not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note 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
HBA Link Speed set HBALinkSpeed	This option allows you to restrict the maximum speed of an HBA card. This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
MLOM Link Speed set Pcie SlotMLOMLinkSpeed	This option allows you to restrict the maximum speed of an MLOM adapter. This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • enabled—The maximum speed is restricted. <p>Note The value enabled is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
MRAID Link Speed scope token-feature "Pcie Slot Link Speed" scope token-param PcieSlotMRAIDLinkSpeed	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • enabled—The maximum speed is not restricted. <p>Note The value enabled is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
RAID-<i>n</i> Link Speed set Pcie SlotRAIDLinkSpeed	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
All Onboard LOM set AllLomPortControl	<p>Whether all onboard LOM ports are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • enabled—All onboard LOM are enabled. • disabled—All onboard LOM are disabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
LOM Port 1 OptionRom set LomOpromControlPort0	Whether Option ROM is available on the LOM port 1. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LOM Port 2 OptionRom set LomOpromControlPort1	Whether Option ROM is available on the LOM port 2. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Slot <i>n</i> State set SlotnState	The state of the adapter card installed in PCIe slot <i>n</i> . This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • uefi-only—The expansion slot is available for UEFI only. • legacy-only—The expansion slot is available for legacy only. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBNVMe1 OptionROM set SBNVMe1OptionROM	Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SBNVMe2 OptionROM set SBNVMe2OptionROM	Whether the server can use Option ROM present in SBNVMe2 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCNVMe1 OptionROM set SIOCNVMe1OptionROM	Whether the server can use Option ROM present in SIOCNVMe1 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCNVMe2 OptionROM set SIOCNVMe2OptionROM	Whether the server can use Option ROM present in SIOCNVMe2 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBLom1 OptionROM set SBLom1OptionROM	Whether the server can use Option ROM present in the SBLom1 controller. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBNVMen Link Speed set SBNVMenLinkSpeed	Link speed for SBNVMe slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SIOCNVMen Link Speed set SIOCNVMenLinkSpeed	Link speed for SIOCNVMe slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SIOCn Link Speed set SIOCnLinkSpeed	Link speed for SIOC slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SBMezzn Link Speed set SBMezznLinkSpeed	Link speed for SBMezz slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
IOESlotn Link Speed set IOESlotnLinkSpeed	Link speed for IOE slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOEMezzn Link Speed set IOEMezznLinkSpeed	Link speed for IOEMezz slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IOENVMen Link Speed set IOENVMenLinkSpeed	Link speed for IOENVMe slot n . This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • enabled—The maximum speed is restricted. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CDN Support for LOMs set CdnSupport	<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"> • enabled—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. • disabled—OS Ethernet Networking Identifier is named in a default convention as ETH0, ETH1 and so on. By default, CDN option is disabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
VMD Enable set VMDEnable	<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"> • enabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is allowed. • disabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is not allowed. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
ACS Control SLOT-<i>n</i> set ACSCtlSlot<i>n</i> <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"> • enabled— Enables peer-to-peer communication between multiple devices for Control Slot <i>n</i>. • disabled— Disables peer-to-peer communication between multiple devices for Control Slot <i>n</i>. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe Slot GPU<i>n</i> OptionROM 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"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
ACS Control GPU-<i>n</i> set ACSctlGpu<i>n</i> <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"> • disabled— Enables peer-to-peer communication between multiple devices for GPUs. • enabled— Disables peer-to-peer communication between multiple devices for GPUs. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
PCIe PLL SSC	<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"> • disabled— Clock is centralized without spreading. • auto— EMI interference is auto adjusted. • zeropointfive— EMI interference us reduced by down-spreading the clock by 0.5%. • platform-default— The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Front Nvme <i>n</i> OptionROM scope token-feature "PCI Slot OptionROM Enable" scope token-param PcieSlotFrontNvmeOptionROM	<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"> • enabled—This is the default option. • disabled
PCIe Slot <i>n</i> Link Speed scope token-feature "PCI Slot LINK Speed" scope token-param PcieSlotLinkSpeed	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted.

Name	Description
<p>MSTOR-RAID Link Speed</p> <p>sc token-feature "PCI Slot LINK Speed"</p> <p>sc token-param PciSlotMSTORRAIDLInkSpeed</p>	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. <p>Note 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>
<p>MSTOR-RAID OptionROM</p> <p>sc token-feature "MSTOR-RAID OptionROM"</p> <p>sc token-param PciSlotMSTORRAIDOptnROM</p>	<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"> • disabled—Option ROM is available. • enabled—Option ROM is not available. This is the default option.
<p>MLOM OptionROM</p> <p>set slot-option-rom-enable-config pcie-mlom</p>	<p>Whether Option ROM is available on the MLOM port. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.
<p>MRAID OptionROM</p> <p>set Pcie SlotMRAID OptionROM</p>	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.
<p>Rear Nvme n OptionRom</p> <p>set RearNvmenOptionROM</p>	<p>Whether Option ROM is available on the Rear NVMEn port. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>PCIe slot MSTOR Link Speed</p> <p>sc token-feature "PCI Slot LINK Speed"</p> <p>sc token-param PciSlotMSTORRAIDLinkSpeed</p>	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. This is the default option. • disabled—The maximum speed is not restricted.
<p>PCIe Slot MSTOR RAID OptionROM</p> <p>scope token-feature "pcie MSTOR-RAID OptionROM"</p> <p>sc token-param PciSlotMSTORRAIDOptnROM</p>	<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"> • disabled—Option ROM is available. • enabled—Option ROM is not available. This is the default option.
<p>PCIe RAS Support</p> <p>sc token-feature "pcie ras-support"</p>	<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"> • disabled—PCIe RAS is available on the slot. • enabled—PCIe RAS is not available on the slot. This is the default option.
<p>MRAID_n Link Speed</p> <p>scope token-feature "Pcie Slot Link Speed"</p> <p>scope token-param PcieSlotMRAIDLinkSpeed</p>	<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"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • gen4—16GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>MRAID_n OptionROM</p> <p>scope token-feature "Pcie Slot OptionROM"</p> <p>scope token-param PcieSlotOptionROM</p>	<p>Whether Option ROM is available on the MRAID port. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.

Name	Description
NVME-<i>n</i> OptionROM scope token-feature "Pcie Slot OptionROM" scope token-param PcieSlotOptionROM	Whether Option ROM is available on the NVME port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.
PCIe Slot OCP Link Speed scope token-feature "Pcie Slot ocp Link Speed" scope token-param PcieSlotocpLinkSpeed	This option allows you to restrict the maximum speed of OCP. This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. This is the default option. • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
RAID<i>n</i> OptionROM scope token-feature "raid optionrom" scope token-param raidoptionrom	Whether Option ROM is available on the RAID port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.
IOENVMe<i>n</i> OptionROM scope token-feature "ioenvme optionrom" scope token-param ioenvmeoptionrom	Whether Option ROM is available on the IOENVMe port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.
GPU<i>n</i> OptionRom scope token-feature "ioemezz1 optionrom" scope token-param ioemezz1optionrom	Whether Option ROM is available on the GPU port. This can be one of the following: <ul style="list-style-type: none"> • disabled—The expansion slot is not available. • enabled—The expansion slot is available. This is the default option.

Name	Description
RAID Link Speed scope token-feature "raid link speed" scope token-param RAIDLinkSpeed	This option allows you to restrict the maximum speed of RAID. This can be one of the following: <ul style="list-style-type: none"> • gen1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • gen2—5GT/s is the maximum speed allowed. • gen3—8GT/s is the maximum speed allowed. • auto—The maximum speed is set automatically. This is the default option. • enabled—The maximum speed is not restricted. <p>Note The value enabled is not supported on Cisco UCS M6 servers.</p> <ul style="list-style-type: none"> • disabled—The maximum speed is not restricted. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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
Integrated Graphics set integrated-graphics-config integrated-graphics	Enables integrated graphics. This can be one of the following: <ul style="list-style-type: none"> • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • enabled—Integrated graphic is enabled. • disabled—Integrated graphics is disabled.
Integrated Graphics Aperture Size set integrated-graphics-aperture-config integrated-graphics-aperture	Allows you to set the size of mapped memory for the integrated graphics controller. This can be one of the following: <ul style="list-style-type: none"> • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • 128mb • 256mb • 512mb • 1024mb • 2048mb • 4096mb

Name	Description
Onboard Graphics set onboard-graphics-config onboard-graphics	Enables onboard graphics (KVM). This can be one of the following: <ul style="list-style-type: none"> • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. • enabled—Onboard graphics is enabled. • disabled—Onboard graphics is disabled.

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
Boot option retry set boot-option-retry-config retry	Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following: <ul style="list-style-type: none"> • disabled—Waits for user input before retrying NON-EFI based boot options. This is the default option. • enabled—Continually retries NON-EFI based boot options without waiting for user input. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SAS RAID set intel-entry-sas-raid-config sas-raid	Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following: <ul style="list-style-type: none"> • disabled—The Intel SAS Entry RAID Module is disabled. • enabled—The Intel SAS Entry RAID Module is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SAS RAID module set intel-entry-sas-raid-config sas-raid-module	How the Intel SAS Entry RAID Module is configured. This can be one of the following: <ul style="list-style-type: none"> • it-ir-raid—Configures the RAID module to use Intel IT/IR RAID. • intel-esrtii—Configures the RAID module to use Intel Embedded Server RAID Technology II. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Onboard SCU Storage Support set onboard-sas-storage-config onboard-sas-ctrl	Whether the onboard software RAID controller is available to the server. This can be one of the following: <ul style="list-style-type: none"> • disabled—The software RAID controller is not available. • enabled—The software RAID controller is available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Cool Down Time (sec)	The time to wait (in seconds) before the next boot attempt. This can be one of the following: <ul style="list-style-type: none"> • 15—System waits for 15 seconds before the next boot attempt. • 45—System waits for 45 seconds before the next boot attempt. • 90—System waits for 90 seconds before the next boot attempt. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. This token is valid only when the Boot Option Retry token has been enabled.
Number of Retries	Number of attempts to boot. This can be one of the following: <ul style="list-style-type: none"> • infinite—System tries all options to boot up. • 13—System tries 13 times to boot up. This is the default option. • 5—System tries 5 times to boot up • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
P-SATA mode	This options allows you to select the P-SATA mode. This can be one of the following: <ul style="list-style-type: none"> • disabled—P-SATA mode is disabled. • lsi-sw-raid—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Power On Password	<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"> • disabled—Power On Password is disabled. • enabled—Power On Password is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
IPV6 PXE Support	<p>Enables or disables IPV6 support for PXE. This can be one of the following</p> <ul style="list-style-type: none"> • disabled—IPV6 PXE support is not available. • enabled—IPV6 PXE support is always available. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Adaptive Memory Training	<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"> • disabled—Adaptive Memory Training is disabled. • enabled—Adaptive Memory Training is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
BIOS Tech Message Level Control (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"> • disabled—BIOS Techlog Level is disabled. • enabled—BIOS Techlog Level is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
OptionROM Launch Optimization	<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"> • disabled—OptionROM Launch Optimization is disabled. • enabled—OptionROM Launch Optimization is enabled. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
BIOS Techlog Level BIOSTechlogLevel	<p>This option denotes the type of messages in BIOS tech log file. The log file can be any of the following types:</p> <ul style="list-style-type: none"> • minimum—Critical messages will be displayed in the log file. This is the default option. • normal—Warning and loading messages will be displayed in the log file. • maximum—Normal and information related messages will be displayed in the log file.
P-SATA OptionROM	<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"> • lsi-sw-raid—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID. This is the default option. • disabled—P-SATA mode is disabled. • ahci—Sets the controllers to AHCI mode. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
M.2 SATA OptionROM	<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"> • lsi-sw-raid—Sets both SATA and sSATA controllers to RAID mode for LSI SW RAID. This is the default option. • disabled—P-SATA mode is disabled. • ahci—Sets the controllers to AHCI mode. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
UEFI Boot Mode	<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"> • disabled—UEFI Boot mode is disabled. • enabled—UEFI Boot mode is enabled.



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
Assert NMI on SERR set assert-nmi-on-serr-config assertion	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: <ul style="list-style-type: none"> • disabled—The BIOS does not generate an NMI or log an error when a SERR occurs. • enabled—The BIOS generates an NMI and logs an error when a SERR occurs. You must enable this setting if you want to enable Assert NMI on PERR. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Assert NMI on PERR set assert-nmi-on-perr-config assertion	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: <ul style="list-style-type: none"> • disabled—The BIOS does not generate an NMI or log an error when a PERR occurs. • enabled—The BIOS generates an NMI and logs an error when a PERR occurs. You must enable Assert NMI on SERR to use this setting. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
OS Boot Watchdog Timer Policy set os-boot-watchdog-timer-policy-config os-boot-watchdog-timer-policy	What action the system takes if the watchdog timer expires. This can be one of the following: <ul style="list-style-type: none"> • power-off—The server is powered off if the watchdog timer expires during OS boot. • reset—The server is reset if the watchdog timer expires during OS boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>

Name	Description
<p>OS Boot Watchdog Timer Timeout set os-boot-watchdog-timer-timeout-config os-boot-watchdog-timer-timeout</p>	<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"> • 5-minutes—The watchdog timer expires 5 minutes after the OS begins to boot. • 10-minutes—The watchdog timer expires 10 minutes after the OS begins to boot. • 15-minutes—The watchdog timer expires 15 minutes after the OS begins to boot. • 20-minutes—The watchdog timer expires 20 minutes after the OS begins to boot. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This option is only available if you enable the OS Boot Watchdog Timer.</p>
<p>FRB-2 Timer set frb-2-timer-config frb-2-timer</p>	<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"> • disabled—The FRB-2 timer is not used. • enabled—The FRB-2 timer is started during POST and used to recover the system if necessary. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Console Redirection Settings

Name	Description
<p>Console redirection set console-redir-config console-redir</p>	<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"> • disabled—No console redirection occurs during POST. • com 0—Enables serial port for console redirection during POST. This option is valid only for M6 blade servers and rack-mount servers. <p>Note The value serial-port-a is not supported on M6 servers.</p> <ul style="list-style-type: none"> • serial-port-b or COM 1—Enables serial port B for console redirection and allows it to perform server management tasks. This option is only valid for rack-mount servers. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note If you enable this option, you also disable the display of the Quiet Boot logo screen during POST.</p>
<p>Flow Control set console-redir-config flow-control</p>	<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"> • none—No flow control is used. • rts-cts—RTS/CTS is used for flow control. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>

Name	Description
<p>Baud rate set console-redir-config baud-rate</p>	<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"> • 9600—A 9600 Baud rate is used. • 19200—A 19200 Baud rate is used. • 38400—A 38400 Baud rate is used. • 57600—A 57600 Baud rate is used. • 115200—A 115200 Baud rate is used. This is the default option. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>
<p>Terminal type set console-redir-config terminal-type</p>	<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"> • pc-ansi—The PC-ANSI terminal font is used. • vt100—A supported vt100 video terminal and its character set are used. • vt100-plus—A supported vt100-plus video terminal and its character set are used. • vt-utf8—A video terminal with the UTF-8 character set is used. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This setting must match the setting on the remote terminal application.</p>

Name	Description
<p>Legacy OS redirection set console-redir-config legacy-os-redir</p>	<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"> • disabled—The serial port enabled for console redirection is hidden from the legacy operating system. • enabled— The serial port enabled for console redirection is visible to the legacy operating system. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Putty Keypad set console-redir-config putty-function-keypad</p>	<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"> • vt100—The function keys generate ESC OP through ESC O[. • linux—Mimics the Linux virtual console. Function keys F6 to F12 behave like the default mode, but F1 to F5 generate ESC [[A through ESC [[E. • xtermr6—Function keys F5 to F12 behave like the default mode. Function keys F1 to F4 generate ESC OP through ESC OS, which are the sequences produced by the top row of the keypad on Digital terminals. • sco—The function keys F1 to F12 generate ESC [M through ESC [X. The function and shift keys generate ESC [Y through ESC [j. The control and function keys generate ESC [k through ESC [v. The shift, control and function keys generate ESC [w through ESC [{. • escn—The default mode. The function keys match the general behavior of Digital terminals. The function keys generate sequences such as ESC [11~ and ESC [12~. • vt400—The function keys behave like the default mode. The top row of the numeric keypad generates ESC OP through ESC OS. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Out of Band Management	<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"> • disabled—Configures the COM port 0 as a general purpose port for use with the Windows Operating System. • enabled—Configures the COM port 0 as a remote management port for Windows Emergency Management services. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Redirection After BIOS POST set console-redir-config putty-function-keypad	<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"> • always_enable—BIOS Legacy console redirection is active during the OS boot and run time. • bootloader—BIOS Legacy console redirection is disabled before giving control to the OS boot loader. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
OS Watchdog Timer Policy scope token-feature "OS Boot Watchdog Timer Policy" scope token-param "OS Boot Watchdog Timer Policy"	<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"> • power_off—The server is powered off if the watchdog timer expires during OS boot. This is the default option. • reset—The server is reset if the watchdog timer expires during OS boot.
FRB 2 Timer scope token-feature "FRB 2 Timer" scope token-param "FRB 2 Timer"	<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"> • disabled—The FRB2 timer is not used. • enabled—The FRB2 timer is started during POST and used to recover the system if necessary. This is the default option.

Name	Description
OS Watchdog Timer scope token-feature "OS Boot Watchdog Timer" scope token-param "OS Boot Watchdog Timer"	Whether the BIOS programs the watchdog timer with a specified timeout value. This can be one of the following: <ul style="list-style-type: none"> • disabled—The watchdog timer is not used to track how long the server takes to boot. This is the default option. • enabled—The watchdog timer tracks how long the server takes to boot. This is the default option.
OS Watchdog Timer Timeout scope token-feature "OS Boot Watchdog Timer Timeout" scope token-param "OS Boot Watchdog Timer Timeout"	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: <ul style="list-style-type: none"> • 5 minutes—The OS watchdog timer expires 5 minutes after it begins to boot. • 10 minutes—The OS watchdog timer expires 10 minutes after it begins to boot. This is the default option. • 15 minutes—The OS watchdog timer expires 15 minutes after it begins to boot. • 20 minutes—The OS watchdog timer expires 20 minutes after it begins to boot. <p>Note 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters org mode for the specified organization. To enter the default org mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create bios-policy <i>policy-name</i>	Creates a BIOS policy with the specified policy name, and enters org BIOS policy mode.
Step 3	Configure the BIOS settings.	
Step 4	UCS-A /org/bios-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a BIOS policy under the root organization and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create bios-policy biosPolicy3
UCS-A /org/bios-policy* # set numa-config numa-optimization enabled
UCS-A /org/bios-policy* # commit-buffer
UCS-A /org/bios-policy #
```

Modifying BIOS Defaults

You can use the following procedure to modify and configure the BIOS defaults for UCS M4 and earlier servers. The new BIOS settings that are introduced with the UCS M5 servers cannot be configured using this procedure.

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

	Command or Action	Purpose
Step 1	UCS-A# scope system	Enters system mode.
Step 2	UCS-A /system # scope server-defaults	Enters server defaults mode.
Step 3	UCS-A /system/server-defaults # show platform	(Optional) Displays platform descriptions for all servers.
Step 4	UCS-A /system/server-defaults # scope platform platform-description	Enters server defaults mode for the server specified. For the <i>platform-description</i> argument, enter the server description displayed by the show platform command using the following format: " <i>vendor</i> " <i>model revision</i> . Tip You must enter the vendor exactly as shown in the show platform command, including all punctuation marks.
Step 5	UCS-A /system/server-defaults/platform # scope bios-settings	Enters server defaults BIOS settings mode for the server.
Step 6	Reconfigure the BIOS settings.	
Step 7	UCS-A /system/server-defaults/platform/bios-settings # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to change the NUMA default BIOS setting for a platform and commit the transaction:

```
UCS-A# scope system
UCS-A /system # scope server-defaults
UCS-A /system/server-defaults # show platform

Platform:
  Product Name Vendor      Model      Revision
  -----
Cisco B200-M1
          Cisco Systems, Inc.
          N20-B6620-1
          0

UCS-A /system/server-defaults # scope platform "Cisco Systems, Inc." N20-B6620-1 0
UCS-A /system/server-defaults/platform # scope bios-settings
UCS-A /system/server-defaults/platform/bios-settings # set numa-config numa-optimization
disabled
UCS-A /system/server-defaults/platform/bios-settings* # commit-buffer
UCS-A /system/server-defaults/platform/bios-settings #

UCS-A# scope system
UCS-A /system # scope server-defaults
UCS-A /system/server-defaults # show platform

Platform:
  Product Name Vendor      Model      Revision
  -----
Cisco UCS B230-M2
          Cisco Systems, Inc.
          B230-BASE-M2
          0
Cisco UCS B440 M2
          Cisco Systems, Inc.
          B440-BASE-M2
          0
Cisco C260-M2
          Cisco Systems, Inc.
          C260-BASE-2646
          0
Cisco B200-M1
          Cisco Systems, Inc.
          N20-B6620-1
          0
Cisco B250-M1
          Cisco Systems, Inc.
          N20-B6620-2

UCS-A /system/server-defaults # scope platform "Cisco Systems, Inc." B230-BASE-M2 0
UCS-A /system/server-defaults/platform # scope bios-settings
UCS-A /system/server-defaults/platform/bios-settings # set numa-config numa-optimization
disabled
UCS-A /system/server-defaults/platform/bios-settings* # committ-buffer
UCS-A /system/server-defaults/platform/bios-settings* #
```

Configuring BIOS Settings for M5 Servers

You can configure BIOS settings for UCS M5 and earlier servers through Cisco UCS Manager CLI. The new BIOS settings that are introduced with the UCS M5 servers can be configured only by using this procedure.

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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope bios-policy <i>bios-policy-name</i>	Enters the bios-policy mode for the specified BIOS policy.
Step 3	(Optional) UCS-A /org/bios-policy # show token-feature	Displays the complete list of BIOS token features in the specified BIOS policy.
Step 4	UCS-A /org/bios-policy # scope token-feature " <i>bios-token-feature-name</i> "	Enters the token feature mode for the specified BIOS token feature.
Step 5	(Optional) UCS-A /org/bios-policy/token-feature # show token-param	Displays the complete list of BIOS token parameters for the specified BIOS token feature.
Step 6	UCS-A /org/bios-policy/token-feature # scope token-param <i>bios-token-parameter-name</i>	Enters the token parameter mode for the specified BIOS token parameter name.
Step 7	(Optional) UCS-A /org/bios-policy/token-feature/token-param # show token-settings	Displays the complete list of token settings for the specified BIOS token parameter.
Step 8	UCS-A /org/bios-policy/token-feature/token-param # scope token-settings <i>token-setting</i>	Enters the token settings mode for the specified BIOS token parameter name.
Step 9	UCS-A /org/bios-policy/token-feature/token-param/token-settings # set is-selected <i>yes no</i>	Set the specified token setting as selected or not by using the yes or no keyword.
Step 10	UCS-A /org/bios-policy/token-feature/token-param/token-settings # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to configure BIOS token settings:

```
UCS-A# scope org
UCS-A /org # scope bios-policy bp
UCS-A /org/bios-policy # scope token-feature "Consistent Device Name Control"
```

```
UCS-A /org/bios-policy/token-feature # scope token-param cdnEnable
UCS-A /org/bios-policy/token-feature/token-param # scope token-settings Enabled
UCS-A /org/bios-policy/token-feature/token-param/token-settings # set is-selected yes
UCS-A /org/bios-policy/token-feature/token-param/token-settings* # commit-buffer
UCS-A /org/bios-policy/token-feature/token-param/token-settings #
```

Viewing the Actual BIOS Settings for M4 Servers

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope bios	Enters BIOS mode for the specified server.
Step 3	UCS-A /chassis/server/bios # scope bios-settings	Enters BIOS settings mode for the specified server.
Step 4	UCS-A /chassis/server/bios/bios-settings # show setting	Displays the BIOS setting. Enter show ? to display a list of allowed values for <i>setting</i> . Note The show setting command is not supported on M5 and higher servers. For M5 and higher servers, see Viewing the Actual BIOS Settings for M5 and Higher Servers , on page 94.

Example

The following example displays a BIOS setting for blade 3 in chassis 1:

```
UCS-A# scope server 1/3
UCS-A /chassis/server # scope bios
UCS-A /chassis/server/bios # scope bios-settings
UCS-A /chassis/server/bios/bios-settings # show intel-vt-config

Intel Vt Config:
  Vt
  --
  Enabled

UCS-A /chassis/server/bios/bios-settings #
```

Viewing the Actual BIOS Settings for M5 and Higher Servers

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope bios	Enters BIOS mode for the specified server.
Step 3	UCS-A /chassis/server/bios # scope bios-settings	Enters BIOS settings mode for the specified server.
Step 4	UCS-A /chassis/server/bios/bios-settings # show token-feature <i>"BIOS-Token-Feature-Name"</i> detail expand	Displays the BIOS setting for a specific BIOS token feature name. Enter show ? to display a list of allowed values for <i>BIOS-Token-Feature-Name</i> .
Step 5	(Optional) UCS-A /chassis/server/bios/bios-settings # show detail	Displays the BIOS setting for all the BIOS tokens.

Example

The following example displays BIOS setting for Consistent Device Name Control on blade 4 in chassis 1:

```
UCS-A# scope server 1/4
UCS-A /chassis/server # scope bios
UCS-A /chassis/server/bios # scope bios-settings
UCS-A /chassis/server/bios/bios-settings # show token-feature "Consistent Device Name
Control" detail expand

Token Feature:
  Bios Token Feature Name: Consistent Device Name Control

  Token Parameter:
    Bios Token Parameter Name: cdnEnable
    UI Display Name: CDN Control

  Token Settings:
    Bios Token Settings Name: Disabled
    BIOS Returned Setting Name: Disabled
    Selected: Yes
UCS-A /chassis/server/bios/bios-settings #
```

Displaying Details of BIOS Tokens in a BIOS Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # scope bios-policy <i>bios-policy-name</i>	Enters the bios-policy mode for the specified BIOS policy.
Step 3	UCS-A /org/bios-policy # show detail expand	Displays detailed information about all the BIOS token features, parameters, and settings that are configured for the specified BIOS policy.
Step 4	(Optional) UCS-A /org/bios-policy # scope token-feature " <i>bios-token-feature-name</i> "	Enters the token feature mode for the specified BIOS token feature.
Step 5	(Optional) UCS-A /org/bios-policy/token-feature # show detail [expand]	Displays the complete list of BIOS token parameters for the specified BIOS token feature.
Step 6	(Optional) UCS-A /org/bios-policy/token-feature # scope token-param <i>bios-token-parameter-name</i>	Enters the token parameter mode for the specified BIOS token parameter name.
Step 7	(Optional) UCS-A /org/bios-policy/token-feature/token-param # show detail [expand]	Displays the complete list of token settings for the specified BIOS token parameter.

Example

This example shows how to display detailed information about a BIOS policy, including all the BIOS token features, parameters, and settings:

```
UCS-A# scope org
UCS-A /org # scope bios-policy bp
UCS-A /org/bios-policy # show detail expand

BIOS Policy:
  Name: bp
  Description: Recommended bios settings for bp
  Reboot on BIOS Policy Change: No
  Policy Owner: Local

  Token Feature:
    Bios Token Feature Name: All USB Devices

    Token Parameter:
      Bios Token Parameter Name: AllUsbDevices
      UI Display Name: All USB Devices

    Token Settings:
      Bios Token Settings Name: Disabled
      BIOS Returned Setting Name: Disabled
      Selected: No

      Bios Token Settings Name: Enabled
      BIOS Returned Setting Name: Enabled
      Selected: No

  Bios Token Feature Name: Altitude
```



```

Token Parameter:
  Bios Token Parameter Name: Altitude
  UI Display Name: Altitude

Token Settings:
  Bios Token Settings Name: 1500-M
  BIOS Returned Setting Name: 1500 M
  Selected: No

  Bios Token Settings Name: 300-M
  BIOS Returned Setting Name: 300 M
  Selected: No

  Bios Token Settings Name: 3000-M
  BIOS Returned Setting Name: 3000 M
  Selected: No

  Bios Token Settings Name: 900-M
  BIOS Returned Setting Name: 900 M
  Selected: No

  Bios Token Settings Name: Auto
  BIOS Returned Setting Name: Auto
  Selected: No
...

```

This example shows how to display detailed information about the BIOS token parameters for a specific BIOS token feature:

```

UCS-A# scope org
UCS-A /org # scope bios-policy bp
UCS-A /org/bios-policy # scope token-feature "Console redirection"
UCS-A /org/bios-policy/token-feature # show detail expand

```

```

Token Feature:
  Bios Token Feature Name: Console redirection

Token Parameter:
  Bios Token Parameter Name: BaudRate
  UI Display Name: Baud rate

Token Settings:
  Bios Token Settings Name: 115.2k
  BIOS Returned Setting Name: 115.2k
  Selected: No

  Bios Token Settings Name: 19.2k
  BIOS Returned Setting Name: 19.2k
  Selected: No

  Bios Token Settings Name: 38.4k
  BIOS Returned Setting Name: 38.4k
  Selected: No

  Bios Token Settings Name: 57.6k
  BIOS Returned Setting Name: 57.6k
  Selected: No

  Bios Token Settings Name: 9.6k
  BIOS Returned Setting Name: 9.6k
  Selected: No

```

```

Bios Token Parameter Name: FlowCtrl
UI Display Name: Flow Control

Token Settings:
  Bios Token Settings Name: None
  BIOS Returned Setting Name: None
  Selected: No

  Bios Token Settings Name: RTS-CTS
  BIOS Returned Setting Name: RTS-CTS
  Selected: No

```

This example shows how to display detailed information about the BIOS token settings for a specific BIOS token parameter:

```

UCS-A# scope org
UCS-A /org # scope bios-policy bp
UCS-A /org/bios-policy # scope token-feature "Console redirection"
UCS-A /org/bios-policy/token-feature # scope token-param BaudRate
UCS-A /org/bios-policy/token-feature/token-param # show detail expand

```

```

Token Parameter:
  Bios Token Parameter Name: BaudRate
  UI Display Name: Baud rate

Token Settings:
  Bios Token Settings Name: 115.2k
  BIOS Returned Setting Name: 115.2k
  Selected: No

  Bios Token Settings Name: 19.2k
  BIOS Returned Setting Name: 19.2k
  Selected: No

  Bios Token Settings Name: 38.4k
  BIOS Returned Setting Name: 38.4k
  Selected: No

  Bios Token Settings Name: 57.6k
  BIOS Returned Setting Name: 57.6k
  Selected: No

  Bios Token Settings Name: 9.6k
  BIOS Returned Setting Name: 9.6k
  Selected: No

```

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 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) and higher, 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 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.

Enabling or Disabling TPM

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create bios-policy <i>policy-name</i>	Creates a BIOS policy with the specified policy name, and enters org BIOS policy mode.
Step 3	UCS-A /org/bios-policy* # set trusted-platform-module-config tpm-state {disabled enabled platform-default}	Specifies whether TPM is enabled or disabled . platform-default is TPM enabled.
Step 4	UCS-A /org/bios-policy* # commit-buffer	Commits the transaction to the system configuration.
Step 5	UCS-A /org # create service-profile <i>sp-name</i> }	Creates the service profile specified and enters service profile configuration mode.
Step 6	UCS-A /org/service-profile* # set bios-policy <i>policy-name</i>	Associates the specified BIOS policy with the service profile.
Step 7	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.

	Command or Action	Purpose
Step 8	UCS-A /org/service-profile # associate server <i>chassis-id / slot-id</i>	Associates the service profile with a single server.

Example

The following example shows how to enable TPM:

```
UCS-A # scope org
UCS-A /org # create bios-policy bp1
UCS-A /org/bios-policy* # set trusted-platform-module-config tpm-state enabled
UCS-A /org/bios-policy* # commit-buffer
UCS-A /org # create service-profile sp1
UCS-A /org/service-profile* # set bios-policy bp1
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile # associate server 1/2
```

Viewing TPM Properties

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id/cartridge-id/server-id</i>	Enters server mode for the specified server.
Step 2	UCS-A /chassis/cartridge/server # scope tpm <i>tpm-id</i>	Enters TPM mode for the specified TPM ID.
Step 3	UCS-A /chassis/cartridge/server/tpm # show	Displays the TPM properties.
Step 4	UCS-A /chassis/cartridge/server/tpm # show detail	Displays detailed TPM properties.

Example

The following example shows how to display the TPM properties a modular server:

```
UCS-A# scope server 1/3/1
UCS-A /chassis/cartridge/server # scope tpm 1
UCS-A /chassis/cartridge/server/tpm # show

Trusted Platform Module:
  Presence: Equipped
  Enabled Status: Enabled
  Active Status: Activated
  Ownership: Unowned
UCS-A /chassis/cartridge/server/tpm # show detail

Trusted Platform Module:
  Enabled Status: Enabled
  Active Status: Activated
  Ownership: Unowned
  Tpm Revision: 2
```

```

Model: UCSX-TPM2-001
Vendor: Cisco Systems Inc
Serial: FCH19257E58
Admin Action: Unspecified
Config State: Not Applied
UCS-A /chassis/cartridge/server/tpm #

```

Enabling or Disabling TXT

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create bios-policy <i>policy-name</i>	Creates a BIOS policy with the specified policy name, and enters org BIOS policy mode.
Step 3	UCS-A /org/bios-policy* # set intel-trusted-execution-technology-config txt-support {enabled disabled platform-default}	Specifies whether TXT is enabled or disabled . platform-default is TXT disabled.
Step 4	UCS-A /org/bios-policy* # commit-buffer	Commits the transaction to the system configuration.
Step 5	UCS-A /org # create service-profile <i>sp-name</i> }	Creates the service profile specified and enters service profile configuration mode.
Step 6	UCS-A /org/service-profile* # set bios-policy <i>policy-name</i>	Associates the specified BIOS policy with the service profile.
Step 7	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.
Step 8	UCS-A /org/service-profile # associate server <i>chassis-id / slot-id</i>	Associates the service profile with a single server.

Example

The following example shows how to enable TXT:

```

UCS-A # scope org
UCS-A /org # create bios-policy bp1
UCS-A /org/bios-policy* # set intel-trusted-execution-technology-config txt-support enabled
UCS-A /org/bios-policy* # commit-buffer
UCS-A /org # create service-profile sp1
UCS-A /org/service-profile* # set bios-policy bp1
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile # associate server 1/2

```

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 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.

Enabling Consistent Device Naming in a BIOS Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # create bios-policy <i>policy-name</i>	Creates a BIOS policy with the specified policy name, and enters org BIOS policy mode.
Step 3	UCS-A /org/bios-policy* # set consistent-device-name-control cdn-name { enabled disabled platform-default }	Specifies whether consistent device naming (CDN) is enabled or disabled .
Step 4	UCS-A /org/bios-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to enable CDN in a BIOS policy:

```
UCS-A # scope org
UCS-A /org # create bios-policy cdn-bios-policy
UCS-A /org/bios-policy* # set consistent-device-name-control cdn-name enabled
UCS-A /org/bios-policy* # commit-buffer
```

Associating a BIOS Policy with a Service Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # scope service-profile <i>sp-name</i> }	Enters service profile configuration mode for the specified service profile.
Step 3	UCS-A /org/service-profile # set bios-policy <i>policy-name</i>	Associates the specified BIOS policy with the service profile.
Step 4	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to associate a CDN-enabled BIOS policy with a service profile:

```
UCS-A # scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile # set bios-policy cdn-bios-policy
UCS-A /org/service-profile* # commit-buffer
```

Configuring Consistent Device Naming for a vNIC

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>sp-name</i>	Enters service profile configuration mode for the specified service profile.
Step 3	UCS-A /org/service-profile # scope vnic <i>vnic-name</i>	Enters vNIC configuration mode for the specified vNIC.
Step 4	UCS-A /org/service-profile/vnic # set cdn-name <i>cdn-name</i>	Specifies the CDN name for the vNIC.
Step 5	UCS-A /org/service-profile/vnic* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to configure CDN for a vNIC:

```
UCS-A # scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile # scope vnic vn1
```

```
UCS-A /org/service-profile/vnic # set cdn-name eth0
UCS-A /org/service-profile/vnic* # commit-buffer
```

Displaying the CDN Name of a vNIC

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server-num</i>	Enters server mode for the specified server.
Step 2	UCS-A /server # scope adapter <i>adapter-id</i>	Enters adapter mode for the specified adapter.
Step 3	UCS-A /server/adapter # show host-eth-if [detail] [expand]	Displays the details of the host Ethernet interface for the specified adapter.

Example

The following example shows how to display the CDN name of a vNIC:

```
UCS-A # scope server 3
UCS-A /server # scope adapter 1
UCS-A /server/adapter # show host-eth-if detail expand
```

```
Eth Interface:
  ID: 1
  Dynamic MAC Address: 00:25:B5:00:00:99
  Burned-In MAC Address: 00:00:00:00:00:00
  Model: UCSC-PCIE-CSC-02
  Name: vnic1
  Cdn Name: cdn0
  Admin State: Enabled
  Operability: Operable
  Order: 1
```

Displaying the Status of a vNIC

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>sp-name</i>	Enters service profile configuration mode for the specified service profile.
Step 3	UCS-A /org/service-profile # show vnic [detail] [expand]	Displays the details of the vNIC in the specified service profile.

Example

This example shows how to display the status of a vNIC.



Note The CDN name that you configured for the vNIC appears as the **Admin CDN Name**. The CDN name that is finally applied to the BIOS policy appears as the **Oper CDN Name**.

```
UCS-A# scope org
UCS-A /org # scope service-profile spl
UCS-A /org/service-profile # show vnic detail expand
```

```
vNIC:
  Name: vnic1
  Fabric ID: B
  Dynamic MAC Addr: 00:25:B5:17:47:01
  Desired Order: Unspecified
  Actual Order: 1
  Desired VCon Placement: 2
  Actual VCon Placement: 2
  Desired Host Port: ANY
  Actual Host Port: NONE
  Equipment: sys/chassis-2/blade-5/adaptor-3/host-eth-2
  Host Interface Ethernet MTU: 1500
  Ethernet Interface Admin CDN Name:cdn0
  Ethernet Interface Oper CDN Name:cdn0
  Template Name:
```

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

Obtain the following:

- 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create ipmi-access-profile <i>profile-name</i>	Creates the specified IPMI access profile and enters organization IPMI access profile mode.
Step 3	UCS-A /org/ipmi-access-profile # set ipmi-over-lan { disable enable }	Determines whether remote connectivity can be established. Note 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.
Step 4	UCS-A /org/ipmi-access-profile # create ipmi-user <i>ipmi-user-name</i>	Creates the specified endpoint user and enters organization IPMI access profile endpoint user mode. Note More than one endpoint user can be created within an IPMI access profile, with each endpoint user having its own password and privileges.
Step 5	UCS-A /org/ipmi-access-profile/ipmi-user # set password	Sets the password for the endpoint user. After entering the set password command, you are prompted to enter and confirm the password. For security purposes, the password that you type does not appear in the CLI.

	Command or Action	Purpose
Step 6	UCS-A /org/ipmi-access-profile/ipmi-user # set privilege {admin readonly}	Specifies whether the endpoint user has administrative or read-only privileges.
Step 7	UCS-A /org/ipmi-access-profile/ipmi-user # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates an IPMI access profile named ReadOnly, creates an endpoint user named bob, sets the password and the privileges for bob, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile* # create ipmi-user bob
UCS-A /org/ipmi-access-profile/ipmi-user* # set password
Enter a password:
Confirm the password:
UCS-A /org/ipmi-access-profile/ipmi-user* # set privilege readonly
UCS-A /org/ipmi-access-profile/ipmi-user* # commit-buffer
UCS-A /org/ipmi-access-profile/ipmi-user #
```

What to do next

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

Deleting an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete ipmi-access-profile <i>profile-name</i>	Deletes the specified IPMI access profile.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the IPMI access profile named ReadOnly and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete ipmi-access-profile ReadOnly
UCS-A /org* # commit-buffer
UCS-A /org #
```

Adding an Endpoint User to an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope ipmi-access-profile <i>profile-name</i>	Enters organization IPMI access profile mode for the specified IPMI access profile.
Step 3	UCS-A /org/ipmi-access-profile # create ipmi-user <i>ipmi-user-name</i>	Creates the specified endpoint user and enters organization IPMI access profile endpoint user mode. Note More than one endpoint user can be created within an IPMI access profile, with each endpoint user having its own password and privileges.
Step 4	UCS-A /org/ipmi-access-profile/ipmi-user # set password	Sets the password for the endpoint user. After entering the set password command, you are prompted to enter and confirm the password. For security purposes, the password that you type does not appear in the CLI.
Step 5	UCS-A /org/ipmi-access-profile/ipmi-user # set privilege { admin readonly }	Specifies whether the endpoint user has administrative or read-only privileges.
Step 6	UCS-A /org/ipmi-access-profile/ipmi-user # commit-buffer	Commits the transaction to the system configuration.

Example

The following example adds an endpoint user named `alice` to the IPMI access profile named `ReadOnly` and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile* # create ipmi-user alice
UCS-A /org/ipmi-access-profile/ipmi-user* # set password
Enter a password:
Confirm the password:
UCS-A /org/ipmi-access-profile/ipmi-user* # set privilege readonly
UCS-A /org/ipmi-access-profile/ipmi-user* # commit-buffer
UCS-A /org/ipmi-access-profile/ipmi-user #
```

Deleting an Endpoint User from an IPMI Access Profile

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope ipmi-access-profile <i>profile-name</i>	Enters organization IPMI access profile mode for the specified IPMI access profile.
Step 3	UCS-A /org/ipmi-access-profile # delete ipmi-user <i>epuser-name</i>	Deletes the specified endpoint user from the IPMI access profile.
Step 4	UCS-A /org/ipmi-access-profile # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the endpoint user named alice from the IPMI access profile named ReadOnly and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope ipmi-access-profile ReadOnly
UCS-A /org/ipmi-access-profile # delete ipmi-user alice
UCS-A /org/ipmi-access-profile* # commit-buffer
UCS-A /org/ipmi-access-profile #
```

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.

Configuring a KVM Management Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create kvm-mgmt-policy <i>policy-name</i>	Creates the specified KVM management policy and enters organization KVM management policy mode.
Step 3	(Optional) UCS-A /org/kvm-mgmt-policy* # set descr <i>description</i>	Provides a description for the policy.
Step 4	UCS-A /org/kvm-mgmt-policy* # set vmedia-encryption { disable enable }	Specifies vMedia encryption is enabled or disabled. Note Starting with UCS Manager 4.2, vMedia Encryption is always enabled for security purposes. It cannot be modified by the user.
Step 5	UCS-A /org/kvm-mgmt-policy* # set kvm-port <i>port-num</i>	Specifies the KVM port. This can be a port number between 1024 and 49151. The default port number is 2068.
Step 6	UCS-A /org/kvm-mgmt-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create a KVM management policy named KVM_Policy1, enable vMedia encryption, set the KVM port number, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # create kvm-mgmt-policy KVM_Policy1
UCS-A /org/kvm-mgmt-policy* # set vmedia-encryption enable
UCS-A /org/kvm-mgmt-policy* # set kvm-port 2078
UCS-A /org/kvm-mgmt-policy* # commit-buffer
UCS-A /org/kvm-mgmt-policy #
```


Modifying a KVM Management Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope kvm-mgmt-policy <i>policy-name</i>	Enters organization KVM management policy mode for the specified KVM management policy.
Step 3	(Optional) UCS-A /org/kvm-mgmt-policy # set descr <i>description</i>	Provides a description for the policy.
Step 4	UCS-A /org/kvm-mgmt-policy* # set vmedia-encryption { disable enable }	Specifies whether vMedia encryption is enabled or disabled.
Step 5	UCS-A /org/kvm-mgmt-policy* # set kvm-port <i>port-num</i>	Specifies the KVM port. This can be a port number between 1024 and 49151. The default port number is 2068.
Step 6	UCS-A /org/kvm-mgmt-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to modify a KVM management policy named KVM_Policy1, and commit the transaction:

```
UCS-A# scope org /
UCS-A /org # scope kvm-mgmt-policy KVM_Policy1
UCS-A /org/kvm-mgmt-policy # show detail
Kvm Mgmt Policy:
  Name: KVM_Policy1
  Description:
  Vmedia Encryption: Disable
  Kvm Port: 2078
UCS-A /org/kvm-mgmt-policy* # set vmedia-encryption enable
UCS-A /org/kvm-mgmt-policy* # set kvm-port 2088
UCS-A /org/kvm-mgmt-policy* # commit-buffer
UCS-A /org/kvm-mgmt-policy # show detail
Kvm Mgmt Policy:
  Name: KVM_Policy1
  Description:
  Vmedia Encryption: Enable
  Kvm Port: 2088
```

Displaying KVM Management Policy Properties

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope kvm-mgmt-policy <i>policy-name</i>	Enters organization KVM management policy mode for the specified KVM management policy.
Step 3	UCS-A /org/kvm-mgmt-policy # show detail	Displays details of the specified policy.

Example

The following example shows how to display details of a KVM management policy named KVM_Policy1:

```
UCS-A# scope org /
UCS-A /org # scope kvm-mgmt-policy KVM_Policy1
UCS-A /org/kvm-mgmt-policy # show detail
Kvm Mgmt Policy:
  Name: KVM_Policy1
  Description:
  Vmedia Encryption: Enable
  Kvm Port: 2088
UCS-A /org/kvm-mgmt-policy #
```

SPDM Security

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



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

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

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

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

- Full Security:

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

- Partial Security (default):

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

- No Security

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

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

Certificates are listed in all user interfaces on a system.

Creating and Configuring a SPDM Security Certificate Policy using CLI

A Security Protocol and Data Model (SPDM) policy can be created to present security alert-level and certificate contents to BMC for authentication.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create spdm-certificate-policy <i>policy-name</i>	Creates a SPDM security certificate policy with the specified policy name, and enters organization SPDM certificate policy mode. Note The only supported certificate type is pem .
Step 3	UCS-A /org/spdm-certificate-policy* # set fault-alert { full partial no }	Configures the fault alert level for this policy.

	Command or Action	Purpose
Step 4	(Optional) UCS-A /org/spdm-certificate-policy* # set descr <i>description</i>	Provides a description for the SPDM security certificate policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/spdm-certificate-policy* # create certificate <i>certificate-name</i>	
Step 6	UCS-A /org/spdm-certificate-policy* # set content	This prompts for the content of the outside certificate. Enter certificate content one line at a time. After End of Certificate, enter ENDOFBUF at the prompt to return to the command line. Note To exit without committing the certificate content, enter C .
Step 7	UCS-A /org/spdm-certificate-policy # commit-buffer	Commits the transaction to the system configuration.

What to do next

Assign outside security certificates, if desired.

Displaying the Security Policy Fault Alert Level

After the policy is created, you can check the alert level for the SPDM policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /org/spdm-certificate-policy # show fault-alert Example: UCS-A /server/cimc/spdm-certificate #show fault-alert	The returned result shows that the setting for this SPDM policy is Partial, the default. SPDM Fault Alert Setting: Partial

Loading an Outside SPDM Security Certificate Policy

The SPDM allows you to download an outside security certificate.

Before you begin

Create a SPDM security certificate policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A /org # scope spdm-certificate-policy	Enters SPDM security certificate policy mode.
Step 2	UCS-A org/spdm-certificate-policy# create spdm-cert <i>Certificate name</i>	Creates a SPDM security certificate policy for the specified external certificate,.
Step 3	UCS-A /org/spdm-certificate-policy* # set <i>{certificate }</i>	Specifying certificate prompts for the content of the outside certificate. The only supported certificate type is pem .
Step 4	UCS-A /org/spdm-certificate-policy # commit-buffer	Commits the transaction to the system configuration.

The following example shows loading a certificate for Broadcom of type PEM.

Example

```
UCS-A-FI-A /org/spdm-certificate-policy# create spdm-cert?
Name - Certificate name

UCS-A-FI-A /org/spdm-certificate-policy# create spdm-cert Broadcom
UCS-A-FI-A /org/spdm-certificate-policy/spdm-cert* # set?
certificate - Certificate content

UCS-A-FI-A /org/spdm-certificate-policy/spdm-cert* # set certificate
{enter certificate content}
UCS-A-FI-A /org/spdm-certificate-policy/spdm-cert* # commit-buffer
UCS-A-FI-A /org/spdm-certificate-policy/spdm-cert# show detail
SPDM Certificate:
Name: Broadcom
Certificate Type: pem
Certificate Content:
```

Viewing the Certificate Inventory

You can view what SPDM certificates have been uploaded and also request further details for a specified certificate.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope server <i>server</i>	
Step 2	UCS-A/server # scope cimc <i>server</i>	
Step 3	UCS-A/server/cimc # scope spdm <i>server</i>	
Step 4	UCS-A/server/cimc/spdm # show certificate	The returned result shows the certificate inventory.
Step 5	UCS-A/server/cimc/spdm # show certificate <i>certificate-iddetail</i> Example: UCS-A /server/cimc/spdm-certificate #show certificate 3 detail Certificate Information Certificate Id : 3 Subject Country Code (C) : US Subject State (ST) : Colorado Subject Organization (O) : Broadcom Inc. Subject Organization Unit (OU) : NA Subject Common Name (CN) : NA Issuer Country Code (C) : US Issuer State (ST) : Colorado Issuer City (L) : Colorado Springs Issuer Organization (O) : Broadcom Inc. Issuer Organization Unit (OU) : NA Issuer Common Name (CN) : NA Valid From : Oct 23 00:25:13 2019 GMT Valid To : Apr 8 10:36:14 2021 GMT UserUploaded : Yes Certificate Content : <Certificate String> Certificate Type : PEM	The returned result shows the certificate ID, identifiers, and expiration date.
Step 6	UCS-A /org/spdm-certificate-policy/certificate # show Example: SPDM Certificate: Name SPDM Certificate Type ----- ----- cert1 Pem Example: UCS-A /server/cimc/spdm-certificate/certificate #up	The returned result shows the type of certificate details. The returned result shows the fault alert setting.

	Command or Action	Purpose
	<pre>UCS-A /server/cimc/spdm-certificate #show SPDM Certificate Policy: Name Fault Alert Setting ----- Broadcom Full</pre>	

Deleting a SPDM Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # delete spdm-certificate-policy <i>policy-name</i>	Deletes the specified SPDM control policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes a power control policy called VendorPolicy2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete spdm-certificate-policy VendorPolicy2
UCS-A /org* # commit-buffer
UCS-A /org #
```

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org	Enters the root organization mode
Step 2	UCS-A /org # create graphicscard-policy <i>policy name</i>	Creates a graphics card policy with the specified policy name, and enters the graphics card policy mode.
Step 3	UCS-A /org/graphicscard-policy # commit buffer	Commits the transaction to the system configuration.

Example

This example shows how to create a graphics card policy:

```
UCS-A# scope org
UCS-A /org # create graphicscard-policy sample
UCS-A /org/graphicscard-policy* # commit-buffer
UCS-A /org/graphicscard-policy #
```

Setting Mode of the Graphics Card Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org	Enters the root organization mode
Step 2	UCS-A /org # scope graphicscard-policy <i>policy name</i>	Enters organization graphics card policy mode.
Step 3	UCS-A /org/graphicscard-policy # set graphicscard-policy-mode [compute] [graphic] [any configuration]	Specifies the mode for the graphics card policy.
Step 4	UCS-A /org/graphicscard-policy # commit buffer	Commits the transaction to the system configuration.

Example

This example shows how to set the mode of a graphics card policy:

```
UCS-A# scope org
UCS-A /org # scope graphicscard-policy sample
UCS-A /org/graphicscard-policy # set graphicscard-policy-mode graphics
UCS-A /org/graphicscard-policy* # commit-buffer
```



```
UCS-A /org/graphicscard-policy #
```

Displaying Details of the Graphics Card

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>server number</i>	Enters the chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # scope graphics-card <i>identifier</i>	Enters the graphics card configuration mode for the specified server.
Step 3	UCS-A /chassis/server/graphics-card # show graphics-card [detail] [expand]	Displays the details of the graphics card for the specified server.

Example

This example shows how to display the details of a graphics card:

```
UCS-A# scope server 1/3
UCS-A /chassis/server # scope graphics-card 2
UCS-A /chassis/server/graphics-card* # show detail
```

```
Graphics Card:
  ID: 2
  Slot Id: 2
  Magma Expander Slot Id:
  Is Supported: Yes
  Vendor: Cisco Systems Inc
  Model: UCSB-GPU-M6
  Serial: FHH1924002B
  Mode: Graphics
  PID: UCSB-GPU-M6
  Firmware Version: 84.04.89.00.01|2754.0200.01.02
  Vendor Id: 0x10de
  Subvendor Id: 0x10de
  Device Id: 0x13f3
  Subdevice Id: 0x1143
UCS-A /chassis/server/graphics-card #
```

Displaying Details of the Graphics Card Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org	Enters the root organization mode
Step 2	UCS-A /org # show graphicscard-policy detail	Displays the details of the graphics card policy.

Example

This example shows how to display the details of a graphics card policy:

```
UCS-A# scope org
UCS-A /org # show graphicscard-policy detail

Graphics Card Policy:
  Name: sample
  Description:
  Graphics Card Policy Mode: Compute

  Name: default
  Description:
  Graphics Card Policy Mode: Any Configuration

  Name: graphics
  Description:
  Graphics Card Policy Mode: Graphics
UCS-A /org #
```

Configuring Local Disk Configuration 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create local-disk-config-policy <i>policy-name</i>	Creates a local disk configuration policy and enters local disk configuration policy mode.

	Command or Action	Purpose
Step 3	(Optional) UCS-A /org/local-disk-config-policy # set descr <i>description</i>	Provides a description for the local disk configuration policy.
Step 4	UCS-A /org/local-disk-config-policy # set mode { any-configuration no-local-storage no-raid raid-0-striped raid-1-mirrored raid-5-striped-parity raid-6-striped-dual-parity raid-10-mirrored-and-striped }	Specifies the mode for the local disk configuration policy.
Step 5	UCS-A /org/local-disk-config-policy # set protect { yes no }	<p>Specifies whether the server retains the configuration in the local disk configuration policy even if the server is disassociated from the service profile.</p> <p>Caution Protect Configuration becomes non-functional if one or more disks in the server are defective or faulty.</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>Note 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>
Step 6	UCS-A /org/local-disk-config-policy # set flexflash-state { enable disable }	Specifies whether FlexFlash SD card support is enabled.
Step 7	UCS-A /org/local-disk-config-policy # set flexflash-raid-reporting-state { enable disable }	Specifies whether FlexFlash RAID reporting support is enabled.

	Command or Action	Purpose
		Note If only one SD card is installed, the FlexFlash inventory displays the RAID State as Disabled and the RAID Health as NA.
Step 8	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a local disk configuration policy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create local-disk-config-policy DiskPolicy7
UCS-A /org/local-disk-config-policy* # set mode raid-1-mirrored
UCS-A /org/local-disk-config-policy* # set protect yes
UCS-A /org/local-disk-config-policy* # commit-buffer
UCS-A /org/local-disk-config-policy #
```

Viewing a Local Disk Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # show local-disk-config-policy policy-name	Displays the local disk policy. If you have not configured a local disk policy, the local disk configuration (created by the create local-disk-config command) displays. Displays the local disk definition (set by the create local-disk-config command). If the serial over LAN definition is not set, and if a policy is set (using the set local-disk-config-policy command), then the policy will be displayed.

Example

The following example shows how to display local disk policy information for a local disk configuration policy called DiskPolicy7:

```
UCS-A# scope org /
UCS-A /org # show local-disk-config-policy DiskPolicy7
```

```
Local Disk Config Policy:
Name: DiskPolicy7
```

```
Mode: Raid 1 Mirrored
Description:
Protect Configuration: Yes
```

Deleting a Local Disk Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete local-disk-config-policy <i>policy-name</i>	Deletes the specified local disk configuration policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the local disk configuration policy named DiskPolicy7 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete local-disk-config-policy DiskPolicy7
UCS-A /org* # commit-buffer
UCS-A /org #
```

FlexFlash Secure Digital Card 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.
- For an associated server, create a FlexFlash scrub policy and disassociate the service profile from the server. For an unassociated server, create a FlexFlash scrub policy and reacknowledge the server after modifying the default scrub policy.

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.

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.
- 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.
- 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope local-disk-config-policy <i>FlexFlash-name</i>	Enters the specified local disk configuration policy mode.
Step 3	UCS-A /org/local-disk-config-policy # set flexflash-state {enable disable}	Specifies whether FlexFlash SD card support is enabled.
Step 4	UCS-A /org/local-disk-config-policy # set flexflash-raid-reporting-state {enable disable}	Specifies whether FlexFlash RAID reporting support is enabled. Note If only one SD card is installed, the FlexFlash inventory displays the RAID State as Disabled and the RAID Health as NA.
Step 5	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system.
Step 6	UCS-A/org/local-disk-config-policy # show detail	Displays the detailed FlexFlash controller properties. <i>/</i> as the <i>org-name</i> .
Step 7	UCS-A# top	

	Command or Action	Purpose
Step 8	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 9	UCS-A /org # scope service-profile <i>slot-3-name</i>	Enters organization service profile mode for the specified service. Slot 3 represents the particular blade's service profile.
Step 10	UCS-A /org/scope service-profile# set local-disk-policy-state <i>FlexFlash-name</i>	Associates the specified local disk policy with the service profile. FlexFlash represents the particular local disk policy.
Step 11	UCS-A /org/scope service-profile# associate server <i>1/1</i>	Associates the service profile with the specified blade server. 1 represents the blade number and the other represents the chassis number.
Step 12	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system.
Step 13	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 14	UCS-A /org # create scrub-policy <i>Scrub-FF-name</i>	Creates a scrub policy with the specified policy name, and enters organization scrub policy mode.
Step 15	(Optional) UCS-A /org/scrub-policy # set descr <i>Scrub FlexFlash ONLY-name</i>	Provides a description for the scrub policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 16	UCS-A /org/scrub-policy # set disk-scrub no	Disables disk scrubbing on servers using this scrub policy.
Step 17	UCS-A /org/scrub-policy # set bios-settings-scrub no	Disables BIOS settings scrubbing on servers using this scrub policy.
Step 18	UCS-A /org/scrub-policy # set flexflash-scrub yes	Enables FlexFlash settings scrubbing on servers using this scrub policy.
Step 19	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system.
Step 20	UCS-A# top	

	Command or Action	Purpose
Step 21	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 22	UCS-A /org # scope service-profile <i>slot-3-name</i>	Enters organization service profile mode for the specified service. Slot 3 represents the particular blade's service profile.
Step 23	UCS-A # acknowledge server <i>1/3-name</i>	Acknowledges the specified blade server. 1 represents the chassis-num and 3 represents the server number.
Step 24	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system. Wait for the blade server to complete committing the transaction.
Step 25	UCS-A # acknowledge server <i>1/3-name</i>	Acknowledges the specified blade server. 1 represents the chassis-num and 3 represents the server number.
Step 26	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system. The FlexFlash cards are now synced and ready to use.

Example

The following example shows the output from the console for starting up the FlexFlash card and creating the policies:

```
#Creating the FlexFlash off policy

UCS-A# scope org
UCS-A /org # create local-disk-config-policy FF-off
UCS-A /org/local-disk-config-policy* # set flexflash-state disable
UCS-A/org/local-disk-config-policy* # commit-buffer
UCS-A/org/local-disk-config-policy # show detail

#Creating a Local Disk Configuration Policy

UCS-A# scope org
UCS-A /org # scope service-profile slot_4
UCS-A /org/service-profile # set local-disk-policy FF-off
UCS-A /org/service-profile* #

UCS-A/org/service-profile* # associate server 1/4
UCS-A/org/service-profile* # commit-buffer
UCS-A /org/service-profile # show detail

#Creating a FlexFlash On policy

UCS-A /org # top
UCS-A# scope org
UCS-A /org # create local-disk-config-policy FF-ON
UCS-A /org/local-disk-config-policy* # set flexflash-state enable
UCS-A /org/local-disk-config-policy* # set flexflash-raid-reporting-state enable
```

```

UCS-A /org/local-disk-config-policy* # commit-buffer
UCS-A /org/local-disk-config-policy #
UCS-A /org/local-disk-config-policy #
UCS-A /org/local-disk-config-policy # show detail

UCS-A /org # top
UCS-A# scope org
UCS-A /org # scope service-profile slot_4
UCS-A /org/service-profile # set local-disk-policy FF-ON
UCS-A /org/service-profile* #

UCS-A /org/service-profile* # associate server 1/4
UCS-A /org/service-profile* # commit-buffer
UCS-A /org/service-profile # show detail

```

Enabling Auto-Sync

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope server <i>server-num</i>	Enters server chassis mode.
Step 3	UCS-A /chassis/server # scope flexflash-controller <i>controller-id</i>	Enters flexflash controller server chassis mode.
Step 4	UCS-A /chassis/server/flexflash-controller # pair <i>primary_slot_number</i>	Resyncs the SD cards if they are out of sync, using the card in the selected slot number as the primary. This can be one of the following: <ul style="list-style-type: none"> • 1—The SD card in slot 1 will be used as the primary. • 2—The SD card in slot 2 will be used as the primary.
Step 5	UCS-A /chassis/server/flexflash-controller # commit-buffer	Commits the transaction to the system configuration.

Example

The following example resyncs the SD cards using the SD card in slot 2 as the primary:

```

UCS-A# scope chassis 1
UCS-A /chassis # scope server 1
UCS-A /chassis/server # scope flexflash-controller 1
UCS-A /chassis/server/flexflash-controller # pair 2
UCS-A /chassis/server/flexflash-controller* # commit-buffer
UCS-A /chassis/server/flexflash-controller #

```

Formatting the FlexFlash Cards

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope server <i>server-num</i>	Enters server chassis mode.
Step 3	UCS-A /chassis/server # scope flexflash-controller <i>controller-id</i>	Enters flexflash controller server chassis mode.
Step 4	UCS-A /chassis/server/flexflash-controller # format	Formats the SD cards.
Step 5	UCS-A /chassis/server/flexflash-controller # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to format the FlexFlash controller:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope server 1
UCS-A /chassis/server # scope flexflash-controller 1
UCS-A /chassis/server/flexflash-controller # format
Warning: When committed, UCSM will format the SD Cards.
This will completely erase the data on the SD Cards!!

UCS-A /chassis/server/flexflash-controller* # commit-buffer
UCS-A /chassis/server/flexflash-controller #
```

Resetting the FlexFlash Controller

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope server <i>server-num</i>	Enters server chassis mode.
Step 3	UCS-A /chassis/server # scope flexflash-controller <i>controller-id</i>	Enters flexflash controller server chassis mode.
Step 4	UCS-A /chassis/server/flexflash-controller # reset	Resets the specified FlexFlash controller.
Step 5	UCS-A /chassis/server/flexflash-controller # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to reset the FlexFlash controller:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope server 1
UCS-A /chassis/server # scope flexflash-controller 1
UCS-A /chassis/server/flexflash-controller # reset
Warning: When committed, UCSM will reset the FlexFlash Controller.
This will cause the host OS to lose connectivity to the SD Cards.

UCS-A /chassis/server/flexflash-controller* # commit-buffer
UCS-A /chassis/server/flexflash-controller #
```

Viewing the FlexFlash Controller Status

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A /chassis # scope server <i>server-num</i>	Enters server chassis mode.
Step 3	UCS-A /chassis/server # scope flexflash-controller <i>controller-id</i>	Enters flexflash controller server chassis mode.
Step 4	UCS-A /chassis/server/flexflash-controller # show detail expand	Displays the detailed FlexFlash controller properties.

Example

The following example shows the status of the FlexFlash controller and SD cards:

```
UCS-A# scope chassis 1
UCS-A /chassis # scope server 1
UCS-A /chassis/server # scope flexflash-controller 1
UCS-A /chassis/server/flexflash-controller # show detail expand
```

```
FlexFlash Controller:
  ID: 1
  Type: SD
  FlexFlash Type: FX3S
  Vendor: Cypress
  Model: FX3S
  Serial: NA
  Firmware Version: 1.3.2 build 158
  Controller State: Connected Partition Over USB To Host
  Controller Health: Old Firmware Running
  RAID State: Enabled Paired
  RAID Health: OK
  Physical Drive Count: 2
  Virtual Drive Count: 1
  RAID Sync Support: Supported
  Operability: Operable
  Oper Qualifier Reason:
  Presence: Equipped
```

Current Task:

FlexFlash Card:

Controller Index: 1
Slot Number: 1
Vendor: SE32G
Model: SE32G
HW Rev: 8.0
Serial: Oxa2140794
Manufacturer ID: 3
OEM ID: SD
Manufacturer Date: 2/14
Size (MB): 30436
Block Size: 512
Card Type: FX3S configured
Write Enabled: Not Write Protected
Card Health: OK
Card Mode: Secondary Active
Operation State: Raid Partition
Card State: Active
Write IO Error Count: 0
Read IO Error Count: 0
Operability: Operable
Oper Qualifier Reason:
Presence: Equipped

FlexFlash Card Drive:

Name: Hypervisor
Size (MB): 30432
Removable: Yes
Operability: Operable
Operation State: Raid Partition

Controller Index: 1
Slot Number: 2
Vendor: SE32G
Model: SE32G
HW Rev: 8.0
Serial: Oxa2140742
Manufacturer ID: 3
OEM ID: SD
Manufacturer Date: 2/14
Size (MB): 30436
Block Size: 512
Card Type: FX3S configured
Write Enabled: Not Write Protected
Card Health: OK
Card Mode: Primary
Operation State: Raid Partition
Card State: Active
Write IO Error Count: 0
Read IO Error Count: 0
Operability: Operable
Oper Qualifier Reason:
Presence: Equipped

FlexFlash Card Drive:

Name: Hypervisor
Size (MB): 30432
Removable: Yes
Operability: Operable
Operation State: Raid Partition

Local Disk Config Definition:


```
Mode: Any Configuration
Description:
Protect Configuration: Yes
```

```
UCS-A /chassis/server/flexflash-controller #
```

Persistent Memory Modules

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

Scrub Policies

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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # create scrub-policy <i>policy-name</i>	Creates a scrub policy with the specified policy name, and enters organization scrub policy mode.
Step 3	(Optional) UCS-A /org/scrub-policy # set descr <i>description</i>	Provides a description for the scrub policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/scrub-policy # set disk-scrub {no yes}	Disables or enables disk scrubbing on servers using this scrub policy as follows: <ul style="list-style-type: none"> • 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. <ul style="list-style-type: none"> • If disabled (default), preserves all data on any local drives, including local storage configuration.
Step 5	UCS-A /org/scrub-policy # set bios-settings-scrub {no yes}	Disables or enables BIOS settings scrubbing on servers using this scrub policy as follows: <ul style="list-style-type: none"> • 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.
Step 6	UCS-A /org/scrub-policy # set flexflash-scrub {no yes}	Disables or enables flexflash scrubbing on servers using this scrub policy as follows:

	Command or Action	Purpose
		<ul style="list-style-type: none"> 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.
Step 7	UCS-A /org/scrub-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a scrub policy named ScrubPolicy2, enables disk scrubbing on servers using the scrub policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create scrub-policy ScrubPolicy2
UCS-A /org/scrub-policy* # set descr "Scrub disk but not BIOS."
UCS-A /org/scrub-policy* # set disk-scrub yes
UCS-A /org/scrub-policy* # set bios-settings-scrub no
UCS-A /org/scrub-policy* # set flexflash-scrub no
UCS-A /org/scrub-policy* # commit-buffer
UCS-A /org/scrub-policy #
```

Deleting a Scrub Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete scrub-policy <i>policy-name</i>	Deletes the specified scrub policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the scrub policy named ScrubPolicy2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete scrub-policy ScrubPolicy2
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope chassis <i>chassis-num</i>	Enters chassis mode for the specified chassis.
Step 2	UCS-A/chassis # scope server <i>server-num</i>	Enters server mode for the specified server.
Step 3	UCS-A/chassis/server # reset-all-memory-errors	Resets the correctable and uncorrectable errors on all the DIMMs in a server.
Step 4	UCS-A /chassis/server* # commit-buffer	Commits any pending transactions.

Example

This example shows how to reset the memory errors for the selected memory unit(s):

```
UCS-A# scope chassis 1
UCS-A/chassis # scope server 1
UCS-A/chassis/server # reset-all-memory-errors
UCS-A/chassis/server* # commit-buffer
UCS-A/chassis/server #
```

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 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

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters root organization mode.
Step 2	UCS-A /org # scope memory-config-policy default	Enters memory policy mode for the global memory policy.
Step 3	UCS-A /org/memory-config-policy # set blacklisting enabled	Enables DIMM blacklisting for the domain level policy and these changes applies 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.
Step 4	UCS-A /org/memory-config-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to enable DIMM blacklisting:

```
UCS-A# scope org /
UCS-A /chassis/org # scope memory-config-policy default
UCS-A /chassis/org/memory-config-policy # set blacklisting enabled
UCS-A /chassis/org/memory-config-policy* # commit-buffer
UCS-A /chassis/org/memory-config-policy #
UCS-A /chassis/org/memory-config-policy # show detail
```

```
Memory Config Policy:
  Blacklisting: enabled
```

Serial over LAN Policy

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.

Configuring a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org# create sol-policy <i>policy-name</i>	Creates a serial over LAN policy and enters organization serial over LAN policy mode.
Step 3	(Optional) UCS-A /org/sol-policy # set descr <i>description</i>	Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.

	Command or Action	Purpose
Step 4	UCS-A /org/sol-policy # set speed {9600 19200 38400 57600 115200 }	Specifies the serial baud rate.
Step 5	UCS-A /org/sol-policy # {disable enable}	Disables or enables the serial over LAN policy. By default, the serial over LAN policy is disabled; you must enable it before it can be applied.
Step 6	UCS-A /org/sol-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a serial over LAN policy named Sol115200, provides a description for the policy, sets the speed to 115200 baud, enables the policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create sol-policy Sol115200
UCS-A /org/sol-policy* # set descr "Sets serial over LAN policy to 115200 baud."
UCS-A /org/sol-policy* # set speed 115200
UCS-A /org/sol-policy* # enable
UCS-A /org/sol-policy* # commit-buffer
UCS-A /org/sol-policy #
```

Viewing a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # show sol-policy <i>policy-name</i>	Displays the serial over LAN definition (set by the create sol-config command). If the serial over LAN definition is not set, and if a policy is set (using the set sol-policy command), then the policy will be displayed.

Example

The following example shows how to display serial over LAN information for a serial over LAN policy called Sol115200:

```
UCS-A# scope org /
UCS-A /org # show sol-policy Sol115200 detail

SOL Policy:
  Name: Sol115200
```

```

SOL State: Enable
Speed: 115200
Description:
Policy Owner: Local

UCS-A /org # show sol-policy Sol115200
SOL Policy:
  Name                               SOL State Speed
  -----
Sol115200                             Enable   115200
UCS-A /org #

```

Deleting a Serial over LAN Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete sol-policy <i>policy-name</i>	Deletes the specified serial over LAN policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the serial over LAN policy named Sol115200 and commits the transaction:

```

UCS-A# scope org /
UCS-A /org* # delete sol-policy Sol115200
UCS-A /org* # commit-buffer
UCS-A /org #

```

Server Autoconfiguration Policy

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.

Configuring a Server Autoconfiguration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create server-autoconfig-policy <i>policy-name</i>	Creates a server autoconfiguration policy with the specified policy name, and enters organization server autoconfiguration policy mode.
Step 3	(Optional) UCS-A /org/server-autoconfig-policy # set descr <i>description</i>	Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	(Optional) UCS-A /org/server-autoconfig-policy # set destination org <i>org-name</i>	Specifies the organization for which the server is to be used.
Step 5	(Optional) UCS-A /org/server-autoconfig-policy # set qualifier <i>server-qual-name</i>	Specifies server pool policy qualification to use for qualifying the server.
Step 6	(Optional) UCS-A /org/server-autoconfig-policy # set template <i>profile-name</i>	Specifies a service profile template to use for creating a service profile instance for the server.
Step 7	UCS-A /org/server-autoconfig-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a server autoconfiguration policy named AutoConfigFinance, provides a description for the policy, specifies finance as the destination organization, ServPoolQual22 as the server pool policy qualification, and ServTemp2 as the service profile template, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create server-autoconfig-policy AutoConfigFinance
UCS-A /org/server-autoconfig-policy* # set descr "Server Autoconfiguration Policy for Finance"
UCS-A /org/server-autoconfig-policy* # set destination org finance
UCS-A /org/server-autoconfig-policy* # set qualifier ServPoolQual22
UCS-A /org/server-autoconfig-policy* # set template ServTemp2
```

```
UCS-A /org/server-autoconfig-policy* # commit-buffer
UCS-A /org/server-autoconfig-policy #
```

Deleting a Server Autoconfiguration Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A/org# delete server-autoconfig-policy <i>policy-name</i>	Deletes the specified server autoconfiguration policy.
Step 3	UCS-A/org# commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the server autoconfiguration policy named AutoConfigFinance and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete server-autoconfig-policy AutoConfigFinance
UCS-A /org* # commit-buffer
UCS-A /org #
```

Server Discovery Policy

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.

Configuring a Server Discovery Policy

Before you begin

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

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode. Note Chassis discovery policies can only be accessed from the root organization.
Step 2	UCS-A /org # create server-disc-policy <i>policy-name</i>	Creates a server discovery policy with the specified policy name, and enters org server discovery policy mode.
Step 3	UCS-A /org/server-disc-policy # set action { diag immediate user-acknowledged }	Specifies when the system will attempt to discover new servers.
Step 4	(Optional) UCS-A /org/chassis-disc-policy # set descr <i>description</i>	Provides a description for the server discovery policy.

	Command or Action	Purpose
		Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	(Optional) UCS-A /org/server-disc-policy # set qualifier <i>qualifier</i>	Uses the specified server pool policy qualifications to associates this policy with a server pool.
Step 6	UCS-A /org/server-disc-policy # set scrub-policy	Specifies the scrub policy to be used by this policy. The scrub policy defines whether the disk drive on a server should be scrubbed clean upon discovery.
Step 7	UCS-A /org/server-disc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a server discovery policy named ServDiscPolExample, sets it to immediately discover new servers, provides a description for the policy, specifies the server pool policy qualifications and scrub policy, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create server-disc-policy ServDiscPolExample
UCS-A /org/server-disc-policy* # set action immediate
UCS-A /org/server-disc-policy* # set descr "This is an example server discovery policy."
UCS-A /org/server-disc-policy* # set qualifier ExampleQual
UCS-A /org/server-disc-policy* # set scrub-policy NoScrub
UCS-A /org/server-disc-policy # commit-buffer
```

What to do next

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

Deleting a Server Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .

	Command or Action	Purpose
Step 2	UCS-A /org # Delete server-disc-policy <i>policy-name</i>	Deletes the specified server discovery policy.
Step 3	UCS-A /org/server-disc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the server discovery policy named ServDiscPolExample and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete server-disc-policy ServDiscPolExample
UCS-A /org* # commit-buffer
UCS-A /org #
```

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 a Hardware Change Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org /	Enters the root organization mode.
Step 2	UCS-A /org # scope server-hwchange-disc-policy <i>policy-name</i>	Enters org hardware change discovery policy mode.
Step 3	UCS-A /org/server-hwchange-disc-policy # set action { auto-acknowledged user-acknowledged }	Specifies when the system will attempt to discover new servers.

	Command or Action	Purpose
Step 4	UCS-A /org/server-hwchange-disc-policy # set action auto-acknowledged	Specifies the hardware change discovery policy to be used..
Step 5	UCS-A /org/server-hwchange-disc-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a hardware change discovery policy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-hwchange-disc-policy
UCS-A /org/server-hwchange-disc-policy # set action
UCS-A /org/server-hwchange-disc-policy # set action auto-acknowledged
UCS-A /org/server-hwchange-disc-policy # commit-buffer
```

Viewing a Hardware Change Discovery Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org .	Enters the root organization mode.
Step 2	UCS-A /org # scope server-hwchange-disc-policy <i>policy-name</i>	Enters org hardware change discovery policy mode.
Step 3	UCS-A /org/server-hwchange-disc-policy # show detail	Displays the Hardware Change Discovery Policy setting.

Example

The following example shows to to view the policy setting:

```
UCS-A# scope org /
UCS-A /org # scope server-hwchange-disc-policy
UCS-A /org/server-hwchange-disc-policy # show detail
Server Hardware Change Discovery Policy:
    Action: User Acknowledged
```

Server Inheritance Policies

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.

Configuring 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create server-inherit-policy <i>policy-name</i>	Creates a server inheritance policy with the specified policy name, and enters organization server inheritance policy mode.
Step 3	(Optional) UCS-A /org/server-inherit-policy # set descr <i>description</i>	Provides a description for the policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	(Optional) UCS-A /org/server-inherit-policy # set destination org <i>org-name</i>	Specifies the organization for which the server is to be used.
Step 5	(Optional) UCS-A /org/server-inherit-policy # set qualifier <i>server-qual-name</i>	Specifies server pool policy qualification to use for qualifying the server.
Step 6	UCS-A /org/server-inherit-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a server inheritance policy named InheritEngineering, provides a description for the policy, specifies engineering as the destination organization and ServPoolQual22 as the server pool policy qualification, and commits the transaction:

```

UCS-A# scope org /
UCS-A /org* # create server-inherit-policy InheritEngineering
UCS-A /org/server-inherit-policy* # set descr "Server Inheritance Policy for Engineering"
UCS-A /org/server-inherit-policy* # set destination org engineering
UCS-A /org/server-inherit-policy* # set qualifier ServPoolQual22
UCS-A /org/server-inherit-policy* # commit-buffer
UCS-A /org/server-inherit-policy #

```

Deleting a Server Inheritance Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-inherit-policy <i>policy-name</i>	Deletes the specified server inheritance policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the server inheritance policy named InheritEngineering and commits the transaction:

```

UCS-A# scope org /
UCS-A /org* # delete server-inherit-policy InheritEngineering
UCS-A /org* # commit-buffer
UCS-A /org #

```

Server Pool Policy

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.

Configuring a Server Pool Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create pooling-policy <i>policy-name</i>	Creates a server pool policy with the specified name, and enters organization pooling policy mode.
Step 3	(Optional) UCS-A /org/pooling-policy # set descr <i>description</i>	Provides a description for the server pool policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 4	UCS-A /org/pooling-policy # set pool <i>pool-distinguished-name</i>	Specifies the server pool to use with the server pool policy. You must specify the full distinguished name for the pool.
Step 5	UCS-A /org/pooling-policy # set qualifier <i>qualifier-name</i>	Specifies the server pool qualifier to use with the server pool policy.
Step 6	UCS-A /org/pooling-policy # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a server pool policy named ServerPoolPolicy4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create pooling-policy ServerPoolPolicy4
UCS-A /org/pooling-policy* # set pool org-root/compute-pool-pool3
UCS-A /org/pooling-policy* # set qualifier ServPoolQual8
UCS-A /org/pooling-policy* # commit-buffer
UCS-A /org/pooling-policy #
```

Deleting a Server Pool Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # delete pooling-policy <i>policy-name</i>	Deletes the specified server pool policy.
Step 3	UCS-A /org # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the server pool policy named ServerPoolPolicy4 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete pooling-policy ServerPoolPolicy4
UCS-A /org/pooling-policy* # commit-buffer
UCS-A /org/pooling-policy #
```

Server Pool Policy Qualification

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 a Server Pool Policy Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create server-qual <i>server-qual-name</i>	Creates a server pool qualification with the specified name, and enters organization server qualification mode.
Step 3	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a server pool qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # create server-qual ServPoolQual22
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

What to do next

Configure one or more of the following server component qualifications:

- Adapter qualification
- Chassis qualification
- Memory qualification

- Power group qualification
- Processor qualification
- Storage qualification

Deleting a Server Pool Policy Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete server-qual <i>server-qual-name</i>	Deletes the specified server pool qualification.
Step 3	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the server pool qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # delete server-qual ServPoolQual22
UCS-A /org* # commit-buffer
UCS-A /org #
```

Creating an Adapter Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create adapter	Creates an adapter qualification and enters organization server qualification adapter mode.

	Command or Action	Purpose
Step 4	UCS-A /org/server-qual/adapter # create cap-qual <i>adapter-type</i>	<p>Creates an adapter capacity qualification for the specified adapter type and enters organization server qualification adapter capacity qualification mode. The <i>adapter-type</i> argument can be any of the following values:</p> <ul style="list-style-type: none"> • fcoe —Fibre Channel over Ethernet • non-virtualized-eth-if —Non-virtualized Ethernet interface • non-virtualized-fc-if —Non-virtualized Fibre Channel interface • path-encap-consolidated —Path encapsulation consolidated • path-encap-virtual —Path encapsulation virtual • protected-eth-if —Protected Ethernet interface • protected-fc-if —Protected Fibre Channel interface • protected-fcoe —Protected Fibre Channel over Ethernet • virtualized-eth-if —Virtualized Ethernet interface • virtualized-fc-if —Virtualized Fibre Channel interface • virtualized-scsi-if —Virtualized SCSI interface
Step 5	UCS-A /org/server-qual/adapter/cap-qual # set maximum { <i>max-cap</i> unspecified }	Specifies the maximum capacity for the selected adapter type.
Step 6	UCS-A /org/server-qual/adapter/cap-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates and configures an adapter qualification for a non-virtualized Ethernet interface and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create adapter
UCS-A /org/server-qual/adapter* # create cap-qual non-virtualized-eth-if
UCS-A /org/server-qual/adapter/cap-qual* # set maximum 2500000000
```

```
UCS-A /org/server-qual/adapter/cap-qual* # commit-buffer
UCS-A /org/server-qual/adapter/cap-qual #
```

Deleting an Adapter Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete adapter	Deletes the adapter qualification from the server pool policy qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the adapter qualification from the server pool policy qualification named ServPoolQual22 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete adapter
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Configuring a Chassis Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.

	Command or Action	Purpose
Step 3	UCS-A /org/server-qual # create chassis <i>min-chassis-num max-chassis-num</i>	Creates a chassis qualification for the specified chassis range and enters organization server qualification chassis mode.
Step 4	UCS-A /org/server-qual/chassis # create slot <i>min-slot-num max-slot-num</i>	Creates a chassis slot qualification for the specified slot range and enters organization server qualification chassis slot mode.
Step 5	UCS-A /org/server-qual/chassis/slot # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a chassis qualification for slots 1 to 4 on chassis 1 and 2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org* # scope server-qual ServPoolQual22
UCS-A /org/server-qual* # create chassis 1 2
UCS-A /org/server-qual/chassis* # create slot 1 4
UCS-A /org/server-qual/chassis/slot* # commit-buffer
UCS-A /org/server-qual/chassis/slot #
```

Deleting a Chassis Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete chassis <i>min-chassis-num max-chassis-num</i>	Deletes the chassis qualification for the specified chassis range.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the chassis qualification for chassis 1 and 2 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
```

```
UCS-A /org/server-qual # delete chassis 1 2
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a CPU Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create cpu	Creates a CPU qualification and enters organization server qualification processor mode.
Step 4	UCS-A /org/server-qual/cpu # set arch { any dual-core-opteron intel-p4-c opteron pentium-4 turion-64 xeon xeon-mp }	Specifies the processor architecture type.
Step 5	UCS-A /org/server-qual/cpu # set maxcores { <i>max-core-num</i> unspecified }	Specifies the maximum number of processor cores.
Step 6	UCS-A /org/server-qual/cpu # set mincores { <i>min-core-num</i> unspecified }	Specifies the minimum number of processor cores.
Step 7	UCS-A /org/server-qual/cpu # set maxprocs { <i>max-proc-num</i> unspecified }	Specifies the maximum number of processors.
Step 8	UCS-A /org/server-qual/cpu # set minprocs { <i>min-proc-num</i> unspecified }	Specifies the minimum number of processors.
Step 9	UCS-A /org/server-qual/cpu # set maxthreads { <i>max-thread-num</i> unspecified }	Specifies the maximum number of threads.
Step 10	UCS-A /org/server-qual/cpu # set minthreads { <i>min-thread-num</i> unspecified }	Specifies the minimum number of threads.
Step 11	UCS-A /org/server-qual/cpu # set stepping { <i>step-num</i> unspecified }	Specifies the processor stepping number.
Step 12	UCS-A /org/server-qual/cpu # set model-regex <i>regex</i>	Specifies a regular expression that the processor name must match.

	Command or Action	Purpose
Step 13	UCS-A /org/server-qual/cpu # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates and configures a CPU qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual122
UCS-A /org/server-qual # create processor
UCS-A /org/server-qual/cpu* # set arch xeon
UCS-A /org/server-qual/cpu* # set maxcores 8
UCS-A /org/server-qual/cpu* # set mincores 4
UCS-A /org/server-qual/cpu* # set maxprocs 2
UCS-A /org/server-qual/cpu* # set minprocs 1
UCS-A /org/server-qual/cpu* # set maxthreads 16
UCS-A /org/server-qual/cpu* # set minthreads 8
UCS-A /org/server-qual/cpu* # set stepping 5
UCS-A /org/server-qual/cpu* # commit-buffer
UCS-A /org/server-qual/cpu #
```

Deleting a CPU Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete cpu	Deletes the processor qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the processor qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual122
UCS-A /org/server-qual # delete cpu
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Power Group Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create power-group <i>power-group-name</i>	Creates a power group qualification for the specified power group name.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example configures a power group qualification for a power group called powergroup1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create power-group powergroup1
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Deleting a Power Group Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete power-group <i>power-group-name</i>	Deletes the specified power group qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes a power group qualification for a power group called powergroup1 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete power-group powergroup1
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Memory Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create memory	Creates a memory qualification and enters organization server qualification memory mode.
Step 4	UCS-A /org/server-qual/memory # set clock { <i>clock-num</i> unspec }	Specifies the memory clock speed.
Step 5	UCS-A /org/server-qual/memory # set maxcap { <i>max-cap-num</i> unspec }	Specifies the maximum capacity of the memory array.
Step 6	UCS-A /org/server-qual/memory # set mincap { <i>min-cap-num</i> unspec }	Specifies the minimum capacity of the memory array.
Step 7	UCS-A /org/server-qual/memory # set speed { <i>speed-num</i> unspec }	Specifies the memory data rate.
Step 8	UCS-A /org/server-qual/memory # set units { <i>unit-num</i> unspec }	Specifies the number of memory units (DRAM chips mounted to the memory board).
Step 9	UCS-A /org/server-qual/memory # set width { <i>width-num</i> unspec }	Specifies the bit width of the data bus.
Step 10	UCS-A /org/server-qual/memory # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates and configures a memory qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create memory
UCS-A /org/server-qual/memory* # set clock 1067
UCS-A /org/server-qual/memory* # set maxcap 4096
UCS-A /org/server-qual/memory* # set mincap 2048
UCS-A /org/server-qual/memory* # set speed unspec
UCS-A /org/server-qual/memory* # set units 16
UCS-A /org/server-qual/memory* # set width 64
UCS-A /org/server-qual/memory* # commit-buffer
UCS-A /org/server-qual/memory #
```

Deleting a Memory Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete memory	Deletes the memory qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the memory qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete memory
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Physical Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create physical-qual	Creates a physical qualification and enters organization server qualification physical mode.
Step 4	UCS-A /org/server-qual/physical-qual # set model-regex <i>regex</i>	Specifies a regular expression that the model name must match.
Step 5	UCS-A /org/server-qual/physical-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates and configures a physical qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # create physical-qual
UCS-A /org/server-qual/physical-qual* # set model-regex
UCS-A /org/server-qual/physical-qual* # commit-buffer
UCS-A /org/server-qual/physical-qual #
```

Deleting a Physical Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete physical-qual	Deletes the physical qualification.
Step 4	UCS-A /org/server-qual # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes a physical qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual22
UCS-A /org/server-qual # delete physical-qual
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Creating a Storage Qualification

Before you begin

Create a server pool policy qualification.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # create storage	Creates a storage qualification and enters organization server qualification storage mode.
Step 4	UCS-A /org/server-qual/storage # set blocksize { <i>block-size-num</i> unknown }	Specifies the storage block size.
Step 5	UCS-A /org/server-qual/storage # set diskless { no unspecified yes }	Specifies whether the available storage must be diskless.
Step 6	UCS-A /org/server-qual/storage # set disktype { hdd ssd unspecified }	Specifies the type of disk that can be used. The options are: <ul style="list-style-type: none"> • Unspecified—Either disk type is acceptable. • HDD—The disk must be HDD. • SSD—The disk must be SSD (SATA or SAS).
Step 7	UCS-A /org/server-qual/storage # set flexflash-num-cards { <i>ff_card-num</i> unknown }	Specifies the number of FlexFlash cards.
Step 8	UCS-A /org/server-qual/storage # set maxcap { <i>max-cap-num</i> unknown }	Specifies the maximum capacity of the storage array.

	Command or Action	Purpose
Step 9	UCS-A /org/server-qual/storage # set mincap { <i>min-cap-num</i> unknown }	Specifies the minimum capacity of the storage array.
Step 10	UCS-A /org/server-qual/storage # set numberofblocks { <i>block-num</i> unknown }	Specifies the number of blocks.
Step 11	UCS-A /org/server-qual/storage # set perdiskcap { <i>disk-cap-num</i> unknown }	Specifies the per-disk capacity.
Step 12	UCS-A /org/server-qual/storage # set units { <i>unit-num</i> unspecified }	Specifies the number of storage units.
Step 13	UCS-A /org/server-qual/storage # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to create and configure a storage qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual122
UCS-A /org/server-qual # create storage
UCS-A /org/server-qual/storage* # set blocksize 512
UCS-A /org/server-qual/storage* # set disktype hdd
UCS-A /org/server-qual/storage* # set maxcap 420000
UCS-A /org/server-qual/storage* # set mincap 140000
UCS-A /org/server-qual/storage* # set numberofblocks 287277984
UCS-A /org/server-qual/storage* # set perdiskcap 140000
UCS-A /org/server-qual/storage* # set units 1
UCS-A /org/server-qual/storage* # set flexflash-num-cards 2
UCS-A /org/server-qual/storage* # commit-buffer
UCS-A /org/server-qual/storage #
```

Deleting a Storage Qualification

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # scope server-qual <i>server-qual-name</i>	Enters organization server qualification mode for the specified server pool policy qualification.
Step 3	UCS-A /org/server-qual # delete storage	Deletes the storage qualification.
Step 4	UCS-A /org/server-qual/ # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the storage qualification and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope server-qual ServPoolQual122
UCS-A /org/server-qual # delete storage
UCS-A /org/server-qual* # commit-buffer
UCS-A /org/server-qual #
```

Configuring vNIC/vHBA Placement Policies

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 171](#).

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.

Configuring a vNIC/vHBA Placement Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 2	UCS-A /org # create vcon-policy <i>policy-name</i>	Creates the specified vNIC/vHBA placement profile and enters organization vcon policy mode.
Step 3	(Optional) UCS-A /org/vcon-policy # set descr <i>description</i>	Provides a description for the vNIC/vHBA Placement Profile.

	Command or Action	Purpose
		<p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p> <p>Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.</p>
Step 4	(Optional) UCS-A /org/vcon-policy # set mapping-scheme {round-robin linear-ordered}	<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"> • round-robin— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon3 to Adapter1, then assigns vCon2 and vCon4 to Adapter2. In a server with three adapter cards, Cisco UCS assigns vCon1 to Adapter1, vCon2 and vCon4 to Adapter2, and vCon3 to Adapter3. This is the default scheme. • linear-ordered— In a server with two adapter cards, Cisco UCS assigns vCon1 and vCon2 to Adapter1, then assigns vCon3 and vCon4 to Adapter2. 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>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,</p>

	Command or Action	Purpose
		<p>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"> • 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.
Step 5	<pre>UCS-A /org/vcon-policy # set vcon {1 2 3 4} selection {all assigned-only exclude-dynamic exclude-unassigned}</pre>	<p>Specifies the selection preference for the specified vCon. The options are:</p> <ul style="list-style-type: none"> • 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. <p>Note An SRIOV usNIC that is explicitly assigned to a vCon set to exclude-usnic will remain assigned to that vCon.</p>

	Command or Action	Purpose
Step 6	UCS-A /org/vcon-policy # commit-buffer	Commits the transaction.

Example

The following example creates a vNIC/vHBA placement policy named Adapter1All, sets the vCon mapping scheme to Linear Ordered, specifies that only assigned vNICs and vHBAs can be placed on adapter 1, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create vcon-policy Adapter1
UCS-A /org/vcon-policy* # set descr "This profile places all vNICs and vHBAs on adapter 1."
UCS-A /org/vcon-policy* # set mapping-scheme linear-ordered
UCS-A /org/vcon-policy* # set vcon 1 selection assigned-only
UCS-A /org/vcon-policy* # commit-buffer
UCS-A /org/vcon-policy* #
UCS-A /org #
```

Deleting a vNIC/vHBA Placement Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # delete vcon-policy <i>policy-name</i>	Deletes the specified vNIC/vHBA placement profile.
Step 3	UCS-A /org # commit-buffer	Commits the transaction.

Example

The following example deletes the vNIC/vHBA placement profile named Adapter1All and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete vcon-policy Adapter1All
UCS-A /org* # commit-buffer
UCS-A /org #
```

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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the organization which contains the service profile whose vNICs you want to explicitly assign to a vCon. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope vnic <i>vnic-name</i>	Enters organization service profile mode for the specified vnic.
Step 4	UCS-A /org/service-profile/vnic # set vcon { 1 2 3 4 any }	Sets the virtual network interface connection (vCon) placement for the specified vNIC. Entering a value of any allows Cisco UCS Manager to determine the vCon to which the vNIC is assigned.
Step 5	UCS-A /org/service-profile/vnic # set order { <i>order-num</i> unspecified }	Specifies the desired PCI order for the vNIC. Valid values include 0-128 and unspecified.
Step 6	UCS-A /org/service-profile/vnic # commit-buffer	Commits the transaction to the system configuration.

Example

The following example sets the vCon placement for a vNIC called vnic3 to 2, sets the desired order to 10, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
UCS-A /org/service-profile # scope vnic vnic3
UCS-A /org/service-profile/vnic # set vcon 2
UCS-A /org/service-profile/vnic* # set order 10
UCS-A /org/service-profile/vnic* # commit-buffer
UCS-A /org/service-profile/vnic #
```

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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the organization which contains the service profile whose vHBAs you want to explicitly assign to a vCon. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters organization service profile mode for the specified service.
Step 3	UCS-A /org/service-profile # scope vhma <i>vhba-name</i>	Enters organization service profile mode for the specified vHBA.
Step 4	UCS-A /org/service-profile/vhma # set vcon { 1 2 3 4 any }	Sets the virtual network interface connection (vCon) placement for the specified vHBA. Entering a value of any allows Cisco UCS Manager to determine the vCon to which the vHBA is assigned.
Step 5	UCS-A /org/service-profile/vhma # set order { <i>order-num</i> unspecified }	Specifies the desired PCI order for the vHBA. Valid desired order number values include 0-128 and unspecified.
Step 6	UCS-A /org/service-profile/vhma # commit-buffer	Commits the transaction to the system configuration.

Example

The following example sets the vCon placement for a vHBA called vhma3 to 2, sets the desired order to 10, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope service-profile accounting
```

```
UCS-A /org/service-profile # scope vhba vhba3
UCS-A /org/service-profile/vhba # set vcon 2
UCS-A /org/service-profile/vhba* # set order 10
UCS-A /org/service-profile/vhba* # commit-buffer
UCS-A /org/service-profile/vhba #
```

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 Scheme: ZeroFunction	Version 2.0 Scheme: ZeroFunction / MultiFunction	Version 2.1 Scheme: ZeroFunction / MultiFunction / StaticZero
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 specific adapters. 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 and VIC 1380 adapters.

Cisco UCS 13xx 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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters the organization mode for the specified organization. To enter the root organization mode, enter / as the <i>org-name</i> .
Step 2	UCS-A /org # scope service-profile <i>profile-name</i>	Enters service profile organization mode for the service profile.
Step 3	UCS-A /org/service-profile # scope vnic <i>vnic-name</i>	Enters organization service profile mode for the specified vNIC.
Step 4	UCS-A /org/service-profile/vnic # set host-port {1 2 any}	Sets the host port for the specified vNIC. Entering a value of any allows Cisco UCS Manager to determine the host port to which the vNIC is assigned. If you set the host port for a vNIC on an adapter that does not support host port placement, the Actual Host Port parameter displays None .
Step 5	UCS-A /org/service-profile/vnic* # commit-buffer	Commits the transaction to the system configuration.
Step 6	UCS-A /org/service-profile/vnic # show detail	Displays details about the specified vNIC.

Example

The following example places a vNIC called vnic3 to host port 2, commits the transaction, and displays the host port information:

```
UCS-A# scope org
UCS-A /org # scope service-profile SP-2
UCS-A /org/service-profile # scope vnic vnic3
UCS-A /org/service-profile/vnic # set host-port 2
UCS-A /org/service-profile/vnic* # commit-buffer
UCS-A /org/service-profile/vnic # show detail
vNIC:
  Name: vnic3
  Fabric ID: A
  Dynamic MAC Addr: 00:25:B5:13:13:11
  Desired Order: 2
  Actual Order: 3
  Desired VCon Placement: 1
  Actual VCon Placement: 1
  Desired Host Port: 2
  Actual Host Port: 2
...
UCS-A /org/service-profile/vnic #
```

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**.

- 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 CIMC vMedia Policy

Before you begin

Make sure that you have access to the following:

- Remote vMedia Server
- vMedia Devices

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
Step 2	UCS-A /org # create vmedia-policy <i>policy-name</i>	Creates a vMedia policy with the specified policy name. 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 3	UCS-A /org/vmedia-policy* # create vmedia-mapping <i>mapping -name</i>	Creates a vMedia policy sub-directory with the specified mapping name.
Step 4	(Optional) UCS-A /org/vmedia-policy/vmedia-mapping # set descr <i>description</i>	Provides a description for the vMedia policy. Note If your description includes spaces, special characters, or punctuation, you must begin and end your description with quotation marks. The quotation marks will not appear in the description field of any show command output.
Step 5	UCS-A /org/vmedia-policy/vmedia-mapping* # set device type <i>device-type</i>	Specifies the remote vMedia image type you wish to mount. Options are: <ul style="list-style-type: none"> • CDD - Scriptable vMedia CD. • HDD - Scriptable vMedia HDD.

	Command or Action	Purpose
Step 6	UCS-A /org/vmedia-policy/vmedia-mapping* # set image-file <i>image-file-name</i>	Specifies the type of remote vMedia image file name. Enter the full path to the backup configuration file. This field can contain the filename [with the file extension] only. Note Ensure that the full path to the file begins with “/” after the share name.
Step 7	UCS-A /org/vmedia-policy/vmedia-mapping* # set image-path <i>image-path</i>	Specifies the remote vMedia image path. Enter the full path to the remote vMedia configuration file.
Step 8	UCS-A /org/vmedia-policy/vmedia-mapping* # set mount-protocol <i>mount-protocol</i>	Specifies the remote vMedia mount protocol. Options are: <ul style="list-style-type: none"> • CIFS • NFS • HTTP • HTTPS
Step 9	UCS-A /org/vmedia-policy/vmedia-mapping* # set password	Specifies the remote vMedia image password.
Step 10	UCS-A /org/vmedia-policy/vmedia-mapping* # set remote-ip <i>remote-ip</i>	Specifies the remote vMedia image IP address.
Step 11	UCS-A /org/vmedia-policy/vmedia-mapping* # set user-id <i>user-id</i>	Specifies the user id for mounting the vMedia device. Enter the username that Cisco UCS Manager should use to log in to the remote server. This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.
Step 12	UCS-A /org/vmedia-policy/vmedia-mapping* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a vMedia policy named vMediaPolicy2, selects remote vMedia device type, mount protocol, image location, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create vmedia-policy vmediapolicy2
UCS-A /org/vmedia-policy* # create vmedia-mapping map1
UCS-A /org/vmedia-policy/vmedia-mapping* # set descr vmedia-map
UCS-A /org/vmedia-policy/vmedia-mapping* # set device-type cdd
UCS-A /org/vmedia-policy/vmedia-mapping* # set image-file-name win2011.iso
```

```

UCS-A /org/vmedia-policy/vmedia-mapping* # set image-path cifs
UCS-A /org/vmedia-policy/vmedia-mapping* # set image-variable-name service-profile-name
UCS-A /org/vmedia-policy/vmedia-mapping* # set mount-protocol cifs
UCS-A /org/vmedia-policy/vmedia-mapping* # set auth-option default
UCS-A /org/vmedia-policy/vmedia-mapping* # set password Password:
UCS-A /org/vmedia-policy/vmedia-mapping* # set remote-ip 172.41.1.158
UCS-A /org/vmedia-policy/vmedia-mapping* # set user-id Administrator
UCS-A /org/vmedia-policy/vmedia-mapping* # commit-buffer

```



Note When vMedia policy is created the **Retry on Mount Fail** option is set to **Yes**. The following example changes the **Retry on Mount Fail** option to **No**.

```

UCS-A# scope org /
UCS-A /org # create vmedia-policy vmediapolicy2
UCS-A /org/vmedia-policy* # set retry-on-mount-fail No
UCS-A /org/vmedia-policy* # commit-buffer

```



Warning When you set the **Retry on Mount Fail** option to **No**, a warning message appears stating: **This will disable automatic retry of mount in case of any vMedia mount failure.**
