



Server-Related Policy Configuration

- [BIOS Settings, on page 1](#)
- [CIMC Security Policies, on page 67](#)
- [Graphics Card Policies, on page 72](#)
- [Configuring Local Disk Configuration Policies, on page 75](#)
- [Scrub Policies, on page 92](#)
- [Configuring DIMM Error Management, on page 95](#)
- [Serial over LAN Policy, on page 97](#)
- [Server Autoconfiguration Policy, on page 100](#)
- [Server Discovery Policy, on page 102](#)
- [Server Inheritance Policies, on page 106](#)
- [Server Pool Policy, on page 108](#)
- [Server Pool Policy Qualification, on page 110](#)
- [Configuring vNIC/vHBA Placement Policies, on page 124](#)
- [CIMC Mounted vMedia, on page 137](#)

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	When the server is rebooted after you change one or more BIOS settings. 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. 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.
BIOS Setting	
Quiet Boot set quiet-boot-config quiet-boot	What the BIOS displays during Power On Self-Test (POST). This can be one of the following: <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
CDN Control set consistent-device-name-control cdn-name	<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. • 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.

Name	Description
<p>Enhanced Intel SpeedStep Tech set enhanced-intel-speedstep-config speed-step</p>	<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>
<p>Intel HyperThreading Tech set hyper-threading-config hyper-threading</p>	<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>
<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 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>
<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 drop-down list set KTI Prefetch	KTI prefetch is a mechanism to get the memory read started early on a DDR bus. This can be one of the following: <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 drop-down list set LLC Prefetch	Whether the processor uses the LLC Prefetch mechanism to fetch the date into the LLC. This can be one of the following: <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 drop-down list set XPT Prefetch	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: <ul style="list-style-type: none"> • disabled—The CPU does not use the XPT Prefetch option. • enabled—The CPU 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	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: <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 drop-down list 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 drop-down list set ProcessorCMCI	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.
CPU Performance set cpu-performance-config cpu-performance	Sets the CPU performance profile for the server. This can be one of the following: <ul style="list-style-type: none"> • Custom — • enterprise—For M3 and higher servers, all prefetchers and data reuse are enabled. For M1 and M2 servers, data reuse and the DCU IP prefetcher are enabled, and all other prefetchers are disabled. • 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.

Name	Description
<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
<p>Power Technology set processor-energy-config cpu-power-management</p>	<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
<p>Channel Interleaving set interleave-config channel-interleave</p>	<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.
<p>Rank Interleaving set interleave-config rank-interleave</p>	<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.
<p>Sub NUMA Clustering drop-down list set SNC</p>	<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"> • disabled— Sub NUMA clustering does not occur. • 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.

Name	Description
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 can be one of the following:</p> <ul style="list-style-type: none"> • 2-way-node-interleave — • 4-way-node-interleave • numa-1-way-node-interleave • 8-way-node-interleave • 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.
Patrol Scrub set scrub-policies-config patrol-scrub	<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.

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.

Name	Description
<p>Package C State Limit set package-c-state-limit-config package-c-state-limit</p>	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <ul style="list-style-type: none"> • auto —The CPU determines the physical elevation. • no-limit—The server may enter any available C state. • c0—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. • c1—When the CPU is idle, the system slightly reduces the power consumption. This option requires less power than C0 and allows the server to return quickly to high performance mode. • c3—When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • c6—When the CPU is idle, the system reduces the power consumption further than with the C3 option. This option saves more power than C0, C1, or C3, but there may be performance issues until the server returns to full power. • c2—When the CPU is idle, the system reduces the power consumption further than with the C1 option. This requires less power than C1 or C0, but it takes the server slightly longer to return to high performance mode. • c7—When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves the maximum amount of power but it also requires the longest time for the server to return to high performance mode. • c7s—When the CPU is idle, the server makes a minimal amount of power available to the components. This option saves more power than C7, but it also requires the longest time for the server to return to high performance mode. • 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 Hardware Power Management drop-down list</p> <p>set cpu-hardware-power-management-config cpu-hardware-power-management</p>	<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)
<p>Energy Performance Tuning drop-down list</p> <p>set power-performance-tuning-support power-performance-tuning-config</p>	<p>Determines if the BIOS or Operating System can turn on the energy performance bias tuning. The options are BIOS and OS.</p> <ul style="list-style-type: none"> • bios— • os— • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Workload Configuration drop-down list</p> <p>set workload-configuration-config</p>	<p>This feature allows for workload optimization. The options are Balanced and I/O Sensitive:</p> <ul style="list-style-type: none"> • balanced— • io-sensitive— • NUMA — (Only GUI) • UMA — (Only GUI) • 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>
<p>Core Performance Boost drop-down list</p> <p>set CbsCmnCpuCpb</p>	<p>Whether the AMD processor increases its frequency on some cores when it is idle or not being used much. This can be one of the following:</p> <ul style="list-style-type: none"> • 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.

Name	Description
Global C-state Control drop-down list set CbsCmnCpuGlobalCstateCtrl	Whether the AMD processors control IO-based C-state generation and DF C-states This can be one of the following: <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.
L1 Stream HW Prefetcher drop-down list set CbsCmnCpuL1StreamHwPrefetcher	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: <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 drop-down list set CbsCmnCpuL2StreamHwPrefetcher	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: <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.

Name	Description
<p>AMD Memory Interleaving drop-down list</p> <p>set CbsDfCmnMemIntlv</p>	<p>Whether the AMD 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"> • auto—The CPU determines how to interleave memory. • channel—Interleaves the physical address space over multiple channels, as opposed to each channel owning single consecutive address spaces. • die—Interleaves the physical address space over multiple dies, as opposed to each die owning single consecutive address spaces. • none—Consecutive memory blocks are accessed from the same physical memory. • socket—Interleaves the physical address space over multiple sockets, as opposed to each socket owning single consecutive address spaces. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>AMD Memory Interleaving Size drop-down list</p> <p>set CbsDfCmnMemIntlvSize</p>	<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
<p>Chipselect Interleaving drop-down list set CbsCmnMemMapBankInterleaveDdr4</p>	<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.
<p>Bank Group Swap drop-down list set CbsCmnMemCtrlBankGroupSwapDdr4</p>	<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.
<p>Determinism Slider drop-down list set CbsCmnDeterminismSlider</p>	<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 drop-down list set CbsCmnGnbNbIOMMU	Input Output Memory Management Unit (IOMMU) allows AMD processors to map virtual addresses to physical addresses. This can be one of the following: <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.
Downcore control drop-down list set CbsCmnCpuGenDowncoreCtrl	Allows AMD processors to disable cores and, thus, select how many cores to enable. This can be one of the following: <ul style="list-style-type: none"> • FOUR (2+2)—Two cores enabled on each CPU complex. • FOUR (4+0)—Four cores enabled on one CPU complex. • SIX (3+3)—Three cores enabled on each CPU complex. • THREE (3+0)—Three cores enabled on one CPU complex. • TWO (1+1)—Two cores enabled on each CPU complex. • TWO (2+0)—Two cores enabled on one CPU complex. • auto—The CPU determines how many cores need to be enabled. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
SVM Mode drop-down list set SvmMode	Whether the processor uses AMD Secure Virtual Machine Technology. This can be one of the following: This can be one of the following: <ul style="list-style-type: none"> • disabled—The processor does not use SVM Technology. • enabled—The processor uses SVM Technology. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
<p>SMT Mode drop-down list set SmtMode</p>	<p>Whether the processor uses AMD Simultaneous MultiThreading Technology, which allows multithreaded software applications to execute threads in parallel within each processor. This can be one of the following:</p> <ul style="list-style-type: none"> • auto—The processor allows for the parallel execution of multiple threads. • off—The processor does not permit multithreading. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>SMEE drop-down list set SMEE</p>	<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"> • 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.
<p>DCPMM Firmware Downgrade drop-down list DCPMMFirmwareDowngrade</p>	<p>Whether the processor uses the DCPMM firmware downgrade function. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The processor does not use the DCPMM firmware downgrade function. • enabled—The processor uses the DCPMM firmware downgrade function. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
<p>Intel Speed Select drop-down list IntelSpeedSelect</p>	<p>How the Intel Speed Select function is configured. This 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. • platform-default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Intel Directed I/O BIOS Settings

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	Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following: <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	Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following: <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	Whether the processor supports Intel VT-d Coherency. This can be one of the following: <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	Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following: <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.

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
Memory RAS configuration set memory-ras-config ras-config	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"> • maximum-performance—System performance is optimized. • 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. • 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
<p>NUMA optimized set numa-config numa-optimization</p>	<p>Whether the BIOS supports NUMA. This can be one of the following:</p> <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.
<p>Mirroring Mode set memory-mirroring-mode mirroring-mode</p>	<p>Memory mirroring enhances system reliability by keeping two identical data images in memory.</p> <p>This option is only available if you choose the mirroring option for Memory RAS Config. It can be one of the following:</p> <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.
<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"> • dim-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.

Name	Description
LV DDR Mode set lv-dimm-support-config lv-ddr-mode	Whether the system prioritizes low voltage or high frequency memory operations. This can be one of the following: <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.
DRAM Refresh Rate set dram-refresh-rate-config dram-refresh	The refresh interval rate for internal memory. This can be one of the following: <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.
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.

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	Whether serial port A is enabled or disabled. This can be one of the following: <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
Make Device Non Bootable set usb-boot-config make-device-non-bootable	Whether the server can boot from a USB device. This can be one of the following: <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.
Legacy USB Support set usb-boot-config legacy-support	Whether the system supports legacy USB devices. This can be one of the following: <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.

Name	Description
USB Idle Power Optimizing Setting set usb-system-idle-power-optimizing-setting-config usb-idle-power-optimizing	<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.
USB Front Panel Access Lock set usb-front-panel-access-lock-config usb-front-panel-lock	<p>USB front panel access lock is configured to enable or disable the front panel access to USB ports. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the system supports 60h/64h emulation for complete USB keyboard legacy support. This can be one of the following:</p> <ul style="list-style-type: none"> • 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.

Name	Description
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.
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 KVM 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.

Name	Description
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.
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-vmmedia	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.

Name	Description
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.

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
VGA Priority set vga-priority-config vga-priority	<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>
ASPM Support set aspm-support-config aspm-support	<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.
BME DMA Mitigation Support set bme-dma-config	<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.

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. • 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. • 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 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.
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.

Name	Description
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.
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.

Name	Description
<p>SBNVME1 Option ROM</p> <p>set sbnvme1-optionrom-config sbnvme1-optionrom</p>	<p>Whether the server can use Option ROM present in SBNVMe1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • 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.
<p>PCIe Slot MRAID OptionROM drop-down list</p> <p>set Pcie SlotMRAIDOptionROM</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. • 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.
<p>PCIe Slot RAID OptionROM</p> <p>set Pcie SlotRAIDOptionROM</p>	<p>Whether Option ROM is available on the RAID 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. • 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.
<p>PCIe Slot Rear Nvme<i>n</i> OptionRom</p> <p>set Pcie SlotRearNvme<i>n</i>OptionROM</p>	<p>Whether Option ROM is available on the Rear NVMe<i>n</i> port. This can be one of the following:</p> <ul style="list-style-type: none"> • 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 NVME1 Link Speed set Pci SlotRearNvme1LinkSpeed	<p>This option allows you to restrict the maximum speed of an NVME card installed in the rear PCIe slot 1. 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. • 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.
Rear NVME2 Link Speed set Pci SlotRearNvme2LinkSpeed	<p>This option allows you to restrict the maximum speed of an NVME card installed in the rear PCIe slot 2. 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. • 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 NVME1 Link Speed set Pci SlotFrontNvme1LinkSpeed	<p>This option allows you to restrict the maximum speed of an NVME card installed in the front PCIe slot 1. 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. • 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.

Name	Description
Front NVME2 Link Speed set Pcie SlotFrontNvme2LinkSpeed	This option allows you to restrict the maximum speed of an NVME card installed in the front PCIe slot 2. 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.
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. • 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.

Name	Description
MRAID Link Speed set Pcie SlotMRAIDLinkSpeed	<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. • 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.
RAID 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. • 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.
LOM Port 1 OptionRom set LomOpromControlPort0	<p>Whether Option ROM is available on the LOM port 1. This can be one of the following:</p> <ul style="list-style-type: none"> • 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
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.
CDN Support for LOMs set CdnSupport	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: <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	Whether NVMe SSDs that are connected to the PCIe bus can be hot swapped. It also standardizes the LED status light on these drives. This can be one of the following: <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.

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
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. • 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.
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.

Name	Description
P-SATA mode drop-down list	<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"> • 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.
Power On Password drop-down list	<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 drop-down list	<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.



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 M3 or 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 54](#) 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.

Name	Description
<p>OS Boot Watchdog Timer set os-boot-watchdog-timer-config os-boot-watchdog-timer</p>	<p>Whether the BIOS programs the watchdog timer with a predefined timeout value. If the operating system does not complete booting before the timer expires, the CIMC resets the system and an error is logged. This can be one of the following:</p> <ul style="list-style-type: none"> • disabled—The watchdog timer is not used to track how long the server takes to boot. • enabled—The watchdog timer tracks how long the server takes to boot. If the server does not boot within the predefined length of time, the CIMC resets the system and logs an error. • platform-default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>This feature requires either operating system support or Intel Management software.</p>
<p>OS Boot Watchdog Timer Policy set os-boot-watchdog-timer-policy-config os-boot-watchdog-timer-policy</p>	<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. • 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. • serial-port-a or COM0 or Enabled—Enables serial port A for console redirection during POST. This option is valid for blade servers and rack-mount servers. • serial-port-b or COM1—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. • 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 drop-down list	<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 drop-down list 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.

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.

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.

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</i> <i>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* # commit-buffer
UCS-A /system/server-defaults/platform/bios-settings* #
```

Viewing the Actual BIOS Settings for M3 and 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> .

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 #
```

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.
Step 8	UCS-A /org/service-profile # associate server <i>chassis-id</i> / <i>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 2012 R2
 - Windows 2016
 - Red Hat Enterprise Linux (RHEL) 6.x and 7.x
- Consistent device naming (CDN) is supported on all M3 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.
- In Cisco UCS Manager Release 3.1 and later releases, CDN is supported on the following adapters:
 - Cisco UCS VIC 1385 (UCSC-PCIE-C40Q-03)
 - Cisco UCS VIC 1387 (UCSC-MLOM-C40Q-03)
- 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.
- In Cisco UCS Manager Release 2.2(4), CDN is supported only on the following adapters:
 - Cisco UCS VIC 1225 (UCSC-PCIE-CSC-02)
 - Cisco UCS MLOM 1227 (UCSC-MLOM-CSC-02)
 - Cisco UCS VIC 1225T (UCSC-PCIE-C10T-02)
 - Cisco UCS MLOM 1227T (UCSC-MLOM-C10T-02)
 - Cisco UCS VIC 1240 (UCSB-MLOM-40G-01)
 - Cisco UCS VIC 1280 (UCS-VIC-M82-8P)
 - Cisco UCS VIC 1340 (UCSB-MLOM-40G-03)
 - Cisco UCS VIC 1380 (UCSB-VIC-M83-8P)
- CDN is supported for vNIC template and dynamic vNIC.
- 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 later 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 is prevented if CDN is enabled in a BIOS policy that is assigned to an associated server.
- In Cisco UCS Manager Release 2.2(4), downgrade of the BIOS firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- In Cisco UCS Manager Release 2.2(4), downgrade of the adapter firmware is prevented if a CDN-enabled BIOS policy is assigned to a server.
- 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.

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> .
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 sp1
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.

	Command or Action	Purpose
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 sp1
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 / 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.

	Command or Action	Purpose
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.
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> .

	Command or Action	Purpose
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 / 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.

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 / 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.
Step 5	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, 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* # commit-buffer
UCS-A /org/kvm-mgmt-policy #
```

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

	Command or Action	Purpose
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.

Do Not Assign a Service Profile with the Default Local Disk Configuration Policy from a B200 M1 or M2 to a B200 M3

Due to the differences in the RAID/JBOD support provided by the storage controllers of B200 M1 and M2 servers and those of the B200 M3 server, you cannot assign or re-assign a service profile that includes the default local disk configuration policy from a B200M1 or M2 server to a B200 M3 server. The default local disk configuration policy includes those with Any Configuration or JBOD configuration.

JBOD Mode Support

The B200 M3 server supports JBOD mode for local disks.



Note Only B200 M1, B200 M2, B200 M3, B250 M1, B250 M2 and B22 M3 blade servers support the JBOD mode for local disks.

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. For example, a B200 M3 server has an LSI controller and an Intel Patsburg controller, but only the LSI controller can be used as a 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.

Number of Disks Selected in Mirrored RAID Should Not Exceed Two

If the number of disks selected in the Mirrored RAID exceed two, RAID 1 is created as a RAID 10 LUN. This issue can occur with the Cisco UCS B440 M1 and B440 M2 servers.

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.

B420 M3 Server Does Not Support All Configuration Modes

The B420 M3 server does not support the following configuration modes in a local disk configuration policy:

- No RAID
- RAID 6 Striped Dual Parity

In addition, the B420 M3 does not support JBOD modes or operations.

Single-Disk RAID 0 Configurations Not Supported on Some Blade Servers

A single-disk RAID 0 configuration is not supported in the following blade servers:

- Cisco UCS B200 M1
- Cisco UCS B200 M2
- Cisco UCS B250 M1
- Cisco UCS B250 M2

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.
Step 3	(Optional) UCS-A /org/local-disk-config-policy # set descr <i>description</i>	Provides a description for the local disk configuration policy.

	Command or Action	Purpose
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 }	<p>Specifies whether FlexFlash RAID reporting support is enabled.</p> <p>Note If only one SD card is installed, the FlexFlash inventory displays the RAID State as Disabled and the RAID Health as NA.</p>
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 <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 local-disk-config-policy <i>policy-name</i>	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 Support

Overview

Cisco UCS B-Series, C-Series M3 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.



Note 16 GB and 32 GB cards are supported only on the B200-M3 blade servers, and the 64 GB SD cards are supported only on the B200-M4 blade servers.

- We do not recommend using an SD card from a rack server in a blade server, or using an SD card from a blade server in a rack server. Switching SD cards between server types might result in data loss from the SD card.
- Some Cisco UCS C-Series rack-mount servers have SD cards with four partitions: HV, HUU, SCU, and Drivers. Only the HV partition is visible in Cisco UCS Manager. You can migrate a four-partition SD card to a single HV partition card with a FlexFlash scrub policy.
- 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 blade server
- Cisco UCS C220 M4 rack server
- Cisco UCS C240 M4 rack 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.

- 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 / 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. / 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 FlexFlash SD Card Support

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 local-disk-config-policy <i>policy-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.

Example

The following example shows how to enable FlexFlash SD card support and FlexFlash RAID reporting state on the local disk config policy default, and commits the transaction to the system:

```

UCS-A# scope org/
UCS-A /org # scope local-disk-config-policy default
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 #

```

Disabling FlexFlash SD Card Support

This procedure describes how to disable FlexFlash support in a local disk policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>{/ org-name}</i>)	Enters the organization mode for the organization you enter
Step 2	UCS-A /org # scope local-disk-config-policy <i>policy-name</i>	Enters the local disk policy configuration mode for the policy name you enter.
Step 3	UCS-A /org/local-disk-config-policy # set flexflash-state <i>disable</i>	Disables FlexFlash SD card support in the local disk policy.
Step 4	UCS-A /org/local-disk-config-policy # commit-buffer	Commits the transaction to the system.

Example

The following example scopes to the root organization, then to the default local disk configuration policy, sets the FlexFlash SD card support to disabled, and commits the transaction to the system:

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

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.

	Command or Action	Purpose
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 #
```

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

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

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

- 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.

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 / as the <i>org-name</i> .
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.

	Command or Action	Purpose
		<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	UCS-A /org/scrub-policy # set disk-scrub {no yes}	<p>Disables or enables disk scrubbing on servers using this scrub policy as follows:</p> <ul style="list-style-type: none"> • If enabled, destroys all data on any local drives. • 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}	<p>Disables or enables BIOS settings scrubbing on servers using this scrub policy as follows:</p> <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}	<p>Disables or enables flexflash scrubbing on servers using this scrub policy as follows:</p> <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.

**Note**

Cisco UCS C-Series 420 M3 rack server do not support this feature.

- 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 / 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.
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 <i>/</i> 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 <i>/</i> 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 / 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.

	Command or Action	Purpose
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. 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> .
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.
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 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.

	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 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 <code>/</code> 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.

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

	Command or Action	Purpose
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 / 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 / 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.
Step 4	UCS-A /org/server-qual/adapter # create cap-qual <i>adapter-type</i>	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: <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

	Command or Action	Purpose
		<ul style="list-style-type: none"> • 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.

	Command or Action	Purpose
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.
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 / 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</i> <i>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> .

	Command or Action	Purpose
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.
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 ServPoolQual22
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.

	Command or Action	Purpose
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 unspec}</i>	Specifies the memory clock speed.
Step 5	UCS-A /org/server-qual/memory # set maxcap <i>{max-cap-num unspec}</i>	Specifies the maximum capacity of the memory array.
Step 6	UCS-A /org/server-qual/memory # set mincap <i>{min-cap-num unspec}</i>	Specifies the minimum capacity of the memory array.
Step 7	UCS-A /org/server-qual/memory # set speed <i>{speed-num unspec}</i>	Specifies the memory data rate.
Step 8	UCS-A /org/server-qual/memory # set units <i>{unit-num unspec}</i>	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 unspec}</i>	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 / 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 / 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.

	Command or Action	Purpose
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 / 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 / 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.
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.

	Command or Action	Purpose
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 ServPoolQual22
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 ServPoolQual22
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 125](#).

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



Note If you are using a vCon policy with two adapters in the Cisco UCS B440 M2 Blade Server, be aware of the following mapping.

- vCon 2 to adapter 1 maps first
- vCon 1 to adapter 2 maps second ZXA Q

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, Cisco UCS assigns all vCons to that adapter. If</p>

	Command or Action	Purpose
		<p>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>
Step 6	<pre>UCS-A /org/vcon-policy # commit-buffer</pre>	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# <code>scope org org-name</code>	Enters organization mode for the specified organization. To enter the root organization mode, type <code>/</code> as the <code>org-name</code> .
Step 2	UCS-A /org # <code>delete vcon-policy policy-name</code>	Deletes the specified vNIC/vHBA placement profile.
Step 3	UCS-A /org # <code>commit-buffer</code>	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 / 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 / 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 vhma vhma3
UCS-A /org/service-profile/vhma # set vcon 2
UCS-A /org/service-profile/vhma* # set order 10
UCS-A /org/service-profile/vhma* # commit-buffer
UCS-A /org/service-profile/vhma #
```

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.

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.

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**.
2. The service profile is renamed.

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



Note Cisco UCS B200M2 Blade Server and Cisco UCS B230M2 Blade Server cannot use a vMedia policy as the policy is not supported on these blade servers.

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.
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.

	Command or Action	Purpose
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.**
