



Server-Related Policies

- [BIOS Settings, on page 1](#)
- [Trusted Platform Module, on page 57](#)
- [Consistent Device Naming, on page 60](#)
- [CIMC Security Policies, on page 64](#)
- [Graphics Card Policies, on page 66](#)
- [Local Disk Policies, on page 67](#)
- [Scrub Policy, on page 80](#)
- [DIMM Error Management, on page 83](#)
- [Serial over LAN Policy Settings, on page 85](#)
- [Server Autoconfiguration Policies, on page 86](#)
- [Server Discovery Policy Settings, on page 88](#)
- [Server Inheritance Policy Settings, on page 91](#)
- [Server Pool Policy Settings, on page 92](#)
- [Server Pool Policy Qualifications Settings, on page 94](#)
- [vNIC/vHBA Placement Policy Settings, on page 99](#)
- [CIMC Mounted vMedia, on page 111](#)

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

Name	Description
POST error pause	<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	<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	<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	<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	<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
Enhanced Intel SpeedStep Tech	<p>Whether the processor uses Enhanced Intel SpeedStep Technology, which allows the system to dynamically adjust processor voltage and core frequency. This technology can result in decreased average power consumption and decreased average heat production. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor never dynamically adjusts its voltage or frequency. • Enabled—The processor utilizes Enhanced Intel SpeedStep Technology and enables all supported processor sleep states to further conserve power. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>We recommend that you contact your operating system vendor to make sure your operating system supports this feature.</p>
Intel HyperThreading Tech	<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>
Core Multi Processing	<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 <i>n</i>—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
Execute Disable Bit	<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>
Intel Virtualization Technology	<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>
Hardware Prefetcher	<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 CPU Performance 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
Adjacent Cache Line Prefetcher	<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 CPU Performance 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>
DCU Streamer Prefetch	<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.
DCU IP Prefetcher	<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	<p>KTI prefetch is a mechanism to get the memory read started early on a DDR bus. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not preload any cache data. • Enabled—The KTI prefetcher preloads the L1 cache with the data it determines to be the most relevant. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
LLC Prefetch drop-down list	<p>Whether the processor uses the LLC Prefetch mechanism to fetch the date into the LLC. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether XPT prefetch is used to enable a read request sent to the last level cache to issue a copy of that request to the memory controller prefetcher. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The CPU does not use the XPT Prefetch option. • Enabled—The CPU enables the XPT prefetcher option. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Direct Cache Access	<p>Allows processors to increase I/O performance by placing data from I/O devices directly into the processor cache. This setting helps to reduce cache misses. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. • Disabled—Data from I/O devices is not placed directly into the processor cache. • Enabled—Data from I/O devices is placed directly into the processor cache. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Processor C State	<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>
Processor C1E	<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.
Processor C3 Report	<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	<p>Whether the processor sends the C6 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the processor sends the C7 report to the operating system. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Enables CMCI generation. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor disables CMCI. • Enabled—The processor enables CMCI. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
CPU Performance	<p>Sets the CPU performance profile for the server. This can be one of the following:</p> <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.
Max Variable MTRR Setting	<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.
Local X2 APIC	<p>Allows you to set the type of Application Policy Infrastructure Controller (APIC) architecture. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—Processor disables Local X2 APIC. • Enabled—Processor enables Local X2 APIC. • XAPIC—Uses the standard xAPIC architecture. • X2APIC—Uses the enhanced x2APIC architecture to support 32 bit addressability of processors. • AUTO—Automatically uses the xAPIC architecture that is detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Power Technology	<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
Energy Performance	<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 Power Technology must be set to Custom or the server ignores the setting for this parameter.</p>
Frequency Floor Override	<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 STATE Coordination	<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 Power Technology must be set to Custom or the server ignores the setting for this parameter.</p>
DRAM Clock Throttling	<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
Channel Interleaving	<p>Whether the CPU divides memory blocks and spreads contiguous portions of data across interleaved channels to enable simultaneous read operations. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines what interleaving is done. • 1 Way— • 2 Way • 3 Way • 4-way—The maximum amount of channel interleaving is used. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Rank Interleaving	<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.
Sub NUMA Clustering drop-down list	<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
IMC Interleave drop-down list	<p>This BIOS option controls the interleaving between the Integrated Memory Controllers (IMCs).</p> <ul style="list-style-type: none"> • 1-way Interleave—There is no interleaving. • 2-way Interleave—Addresses are interleaved between the two IMCs. • Auto —CPU determines the IMC Interleaving mode. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Memory Interleaving	<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 interleaving Inter Socket • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Demand Scrub	<p>Whether the system corrects single bit memory errors encountered when the CPU or I/O makes a demand read. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled— Single bit memory errors are not corrected. • Enabled— Single bit memory errors are corrected in memory and the corrected data is set in response to the demand read. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
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.
Altitude	<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
Package C State Limit	<p>The amount of power available to the server components when they are idle. This can be one of the following:</p> <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
CPU Hardware Power Management drop-down list	<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)
Energy Performance Tuning drop-down list	<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.
Workload Configuration drop-down list	<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>
Core Performance Boost drop-down list	<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	<p>Whether the AMD processors control IO-based C-state generation and DF C-states This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to control IO-based C-state generation. • Disabled—Global C-state control is disabled. • Enabled—Global C-state control is enabled. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
L1 Stream HW Prefetcher drop-down list	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L1 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. • Disabled—The hardware prefetcher is not used. • Enabled—The processor uses the hardware prefetcher when cache issues are detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
L2 Stream HW Prefetcher drop-down list	<p>Whether the processor allows the AMD hardware prefetcher to speculatively fetch streams of data and instruction from memory into the L2 cache when necessary. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how to place data from I/O devices into the processor cache. • Disabled—The hardware prefetcher is not used. • Enabled—The processor uses the hardware prefetcher when cache issues are detected. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
AMD Memory Interleaving drop-down list	<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.
AMD Memory Interleaving Size drop-down list	<p>Determines the size of the memory blocks to be interleaved. It also determines the starting address of the interleave (bit 8,9,10 or 11). This can be one of the following:</p> <ul style="list-style-type: none"> • 1 KB • 2 KB • 256 Bytes • 512 Bytes • Auto—The CPU determines the size of the memory block. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Chipselect Interleaving drop-down list	<p>Whether memory blocks across the DRAM chip selects for node 0 are interleaved. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to interleave chip selects. • Disabled—Chip selects are not interleaved within the memory controller. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Bank Group Swap drop-down list	<p>Determines how physical addresses are assigned to applications. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU automatically determines how to assign physical addresses to applications. • Disabled—Bank group swap is not used. • Enabled—Bank group swap is used to improve the performance of applications. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Determinism Slider drop-down list	<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	<p>Input Output Memory Management Unit (IOMMU) allows AMD processors to map virtual addresses to physical addresses. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto—The CPU determines how map these addresses. • Disabled—IOMMU is not used. • Enabled—Address mapping takes place through the IOMMU. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Downcore control drop-down list	<p>Allows AMD processors to disable cores and, thus, select how many cores to enable. This can be one of the following:</p> <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	<p>Whether the processor uses AMD Secure Virtual Machine Technology. This can be one of the following: This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use SVM Technology. • Enabled—The processor uses SVM Technology. • Platform Default —The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
SMT Mode drop-down list	<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.
SMEE drop-down list	<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.
DCPMM Firmware Downgrade drop-down list	<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.
Intel Speed Select drop-down list	<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. • Config 1—The processor uses Config 1. • Config 2—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	<p>Whether the processor uses Intel Virtualization Technology for Directed I/O (VT-d). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not use virtualization technology. • Enabled—The processor uses virtualization technology. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor. <p>Note This option must be enabled if you want to change any of the other Intel Directed I/O BIOS settings.</p>
Intel VTD interrupt Remapping	<p>Whether the processor supports Intel VT-d Interrupt Remapping. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support remapping. • Enabled—The processor uses VT-d Interrupt Remapping as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD coherency support	<p>Whether the processor supports Intel VT-d Coherency. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support coherency. • Enabled—The processor uses VT-d Coherency as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Intel VTD ATS support	<p>Whether the processor supports Intel VT-d Address Translation Services (ATS). This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The processor does not support ATS. • Enabled—The processor uses VT-d ATS as required. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
Intel VTD pass through DMA support	<p>Whether the processor supports Intel VT-d Pass-through DMA. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>How the memory reliability, availability, and serviceability (RAS) is configured for the server. This can be one of the following:</p> <ul style="list-style-type: none"> • Maximum Performance—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
NUMA optimized	<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.
Mirroring Mode	<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.
Sparing Mode	<p>Sparing optimizes reliability by holding memory in reserve so that it can be used in case other DIMMs fail. This option provides some memory redundancy, but does not provide as much redundancy as mirroring. The available sparing modes depend on the current memory population.</p> <p>This option is only available if you choose sparing option for Memory RAS Config. It can be one of the following:</p> <ul style="list-style-type: none"> • DIMM Sparing—One DIMM is held in reserve. If a DIMM fails, the contents of a failing DIMM are transferred to the spare DIMM. • Rank Sparing—A spare rank of DIMMs is held in reserve. If a rank of DIMMs fails, the contents of the failing rank are transferred to the spare rank. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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

USB BIOS Settings

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

Name	Description
Make Device Non Bootable	<p>Whether the server can boot from a USB device. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The server can boot from a USB device. • Enabled—The server cannot boot from a USB device. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
Legacy USB Support	<p>Whether the system supports legacy USB devices. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—USB devices are only available to EFI applications. • Enabled—Legacy USB support is always available. • Auto—Disables legacy USB support if no USB devices are connected. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

Name	Description
USB Idle Power Optimizing Setting	<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-performanceHigh 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	<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	<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	<p>Whether the front panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the internal USB devices are enabled or disabled. This can be one of the following:</p> <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	<p>Whether the KVM ports are enabled or disabled. This can be one of the following:</p> <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	<p>Whether the rear panel USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the SD card drives are enabled or disabled. This can be one of the following:</p> <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	<p>Whether the virtual media devices are enabled or disabled. This can be one of the following:</p> <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	<p>Whether all physical and virtual USB devices are enabled or disabled. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether xHCI mode is enabled or disabled. This can be one of the following:</p> <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	<p>Whether the BIOS maximizes memory usage below 4GB for an operating system without PAE support, depending on the system configuration. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether to enable or disable memory mapped I/O of 64-bit PCI devices to 4GB or greater address space. Legacy option ROMs are not able to access addresses above 4GB. PCI devices that are 64-bit compliant but use a legacy option ROM may not function correctly with this setting enabled. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<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	<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	<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	<p>The Intel QuickPath Interconnect (QPI) link frequency, in megatransfers per second (MT/s). This can be one of the following:</p> <ul style="list-style-type: none"> • 6.4 GT/s • 7.2 GT/s • 8.0 GT/s • 9.6 GT/s • 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	<p>This can be one of the following:</p> <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	<p>Whether Option ROM is available on the SAS 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.
PCIe Slot <i>n</i> Link Speed	<p>This option allows you to restrict the maximum speed of an adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<p>Whether Option ROM is available on the 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.
PCIe Slot HBA OptionROM	<p>Whether Option ROM is available on the HBA 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
PCIe Slot MLOM OptionROM	<p>Whether Option ROM is available on the MLOM port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • 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	<p>Whether Option ROM is available on the 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.
PCIe 10G LOM 2 Link	<p>Whether Option ROM is available on the 10G LOM port. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>PCI ROM Command Line Protocol (CLP) controls the execution of different Option ROMs such as PxE and iSCSI that are present in the card. By default, it is disabled.</p> <ul style="list-style-type: none"> • 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	<p>Whether the server can use Option ROM present in System IO Controller 1 (SIOC1). This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the server can use Option ROM present in System IO Controller 2 (SIOC2). This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the server can use Option ROM present in SBMezz1 controller. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether option ROM is enabled on the IOE slot 1. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether option ROM is enabled on the IOE Mezz1. 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.
IOE Slot2 Option ROM	<p>Whether option ROM is enabled on the IOE slot 2. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether option ROM is enabled on the IOE NVMe1. 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.
IO ENVME2 Option ROM	<p>Whether option ROM is enabled on the IOE NVMe2. 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
SBNVME1 Option ROM	<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.
PCIe Slot MRAID OptionROM drop-down list	<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.
PCIe Slot RAID OptionROM	<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.
PCIe Slot Rear Nvmen OptionRom	<p>Whether Option ROM is available on the Rear NVMEn port. This can be one of the following:</p> <ul style="list-style-type: none"> • Disabled—The expansion slot is not available. • Enabled—The expansion slot is available. • 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	<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"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<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"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<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"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<p>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:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<p>This option allows you to restrict the maximum speed of an HBA card. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<p>This option allows you to restrict the maximum speed of an MLOM adapter. This can be one of the following:</p> <ul style="list-style-type: none"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<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"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<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"> • Gen 1—2.5GT/s (gigatransfers per second) is the maximum speed allowed. • Gen 2—5GT/s is the maximum speed allowed. • Gen 3—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	<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	<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	<p>Whether Option ROM is available on the LOM port 2. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>The state of the adapter card installed in PCIe slot <i>n</i>. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the Ethernet Networking Identifier naming convention is according to Consistent Device Naming (CDN) or the traditional way of naming conventions. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—OS Ethernet Network Identifier is named in a consistent device naming (CDN) convention according to the physical LAN on Motherboard(LOM) port numbering; LOM Port 0, LOM Port 1 and so on. • Disabled—OS Ethernet Networking Identifier is named in a default convention as ETH0, ETH1 and so on. By default, CDN option is disabled. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.
VMD Enable	<p>Whether NVMe SSDs that are connected to the PCIe bus can be hot swapped. It also standardizes the LED status light on these drives. This can be one of the following:</p> <ul style="list-style-type: none"> • Enabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is allowed. • Disabled—Hot swap of NVMe SSDs that are connected to the PCIe bus is not allowed. • Platform Default—The BIOS uses the value for this attribute contained in the BIOS defaults for the server type and vendor.

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	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. • —Integrated graphic is enabled. • —Integrated graphics is disabled.
Integrated Graphics Aperture Size	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. • • • • • •
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. • —Onboard graphics is enabled. • —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	<p>Whether the BIOS retries NON-EFI based boot options without waiting for user input. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the Intel SAS Entry RAID Module is enabled. This can be one of the following:</p> <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	<p>How the Intel SAS Entry RAID Module is configured. This can be one of the following:</p> <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	<p>Whether the onboard software RAID controller is available to the server. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a system error (SERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • 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	<p>Whether the BIOS generates a non-maskable interrupt (NMI) and logs an error when a processor bus parity error (PERR) occurs. This can be one of the following:</p> <ul style="list-style-type: none"> • 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
OS Boot Watchdog Timer	<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>
OS Boot Watchdog Timer Policy	<p>What action the system takes if the watchdog timer expires. This can be one of the following:</p> <ul style="list-style-type: none"> • Power Off—The server is powered off if the watchdog timer expires during OS boot. • 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
OS Boot Watchdog Timer Timeout	<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>
FRB-2 Timer	<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
Console redirection	<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>
Flow Control	<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
Baud rate	<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"> • 9.6k—A 9600 Baud rate is used. • 19.2k—A 19200 Baud rate is used. • 38.4k—A 38400 Baud rate is used. • 57.6k—A 57600 Baud rate is used. • 115.2k—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>
Terminal type	<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
Legacy OS redirection	<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.
Putty Keypad set console-redir-config putty-function-keypad	<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	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: <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	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: <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

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

- a) If you want to change a BIOS setting, click the desired radio button or make the appropriate choice from the drop-down list.

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

- **Main** page: [Main BIOS Settings, on page 2](#)
- **Processor** page: [Processor BIOS Settings, on page 4](#)
- **Intel Directed IO** page: [Intel Directed I/O BIOS Settings, on page 25](#)
- **RAS Memory** page: [RAS Memory BIOS Settings, on page 26](#)
- **Serial Port** page: [Serial Port BIOS Settings, on page 28](#)
- **USB** page: [USB BIOS Settings, on page 29](#)
- **PCI Configuration** page: [PCI Configuration BIOS Settings, on page 33](#)
- **Boot Options** page: [Boot Options BIOS Settings, on page 45](#)
- **Server Management** page: [Server Management BIOS Settings, on page 48](#)

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

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

Modifying the BIOS Defaults

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

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

Procedure

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

Step 2 Expand **Servers > Policies**.

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

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

Step 4 Expand **BIOS Defaults** and select the server model number for which you want to modify the default BIOS settings.

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

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

- **Main** tab: [Main BIOS Settings, on page 2](#)

- **Advanced** tab:
 - **Processor** subtab: [Processor BIOS Settings, on page 4](#)
 - **Intel Directed IO** subtab: [Intel Directed I/O BIOS Settings, on page 25](#)
 - **RAS Memory** subtab: [RAS Memory BIOS Settings, on page 26](#)
 - **Serial Port** subtab: [Serial Port BIOS Settings, on page 28](#)
 - **USB** subtab: [USB BIOS Settings, on page 29](#)
 - **PCI Configuration** subtab: [PCI Configuration BIOS Settings, on page 33](#)
- **Boot Options** tab: [Boot Options BIOS Settings, on page 45](#)
- **Server Management** tab: [Server Management BIOS Settings, on page 48](#)

Step 6 Click **Save Changes**.

Viewing the Actual BIOS Settings for a Server

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

Procedure

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

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

Trusted Platform Module

Trusted Platform Module

The Trusted Platform Module (TPM) is a component that can securely store artifacts that are used to authenticate the server. These artifacts can include passwords, certificates, or encryption keys. A TPM can also be used to store platform measurements that help ensure that the platform remains trustworthy. Authentication (ensuring that the platform can prove that it is what it claims to be) and attestation (a process helping to prove that a platform is trustworthy and has not been breached) are necessary steps to ensure safer computing in all

environments. It is a requirement for the Intel Trusted Execution Technology (TXT) security feature, which must be enabled in the BIOS settings for a server equipped with a TPM. Cisco UCS M4 and higher blade and rack-mount servers include support for TPM. TPM is enabled by default on these servers.


Important

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

Intel Trusted Execution Technology

Intel Trusted Execution Technology (TXT) provides greater protection for information that is used and stored on the business server. A key aspect of that protection is the provision of an isolated execution environment and associated sections of memory where operations can be conducted on sensitive data, invisible to the rest of the system. Intel TXT provides for a sealed portion of storage where sensitive data such as encryption keys can be kept, helping to shield them from being compromised during an attack by malicious code. Cisco UCS M4 and higher blade and rack-mount servers include support for TXT. TXT is disabled by default on these servers.

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

Configuring Trusted Platform

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to configure TPM.
- Step 4** Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.
- Step 5** In the **Work** pane, click the **Advanced** tab.
- Step 6** Click the **Trusted Platform** subtab.
- Step 7** To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enabled	Enables TPM
Platform Default	Enables TPM

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

Option	Description
disabled	Disables TXT

Option	Description
enabled	Enables TXT
Platform Default	Disables TXT

Step 9 Click **Save Changes**.

Configuring Trusted Platform

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to configure TPM.
- Step 4** Expand **BIOS Policies** and select the BIOS policy for which you want to configure TPM.
- Step 5** In the **Work** pane, click the **Advanced** tab.
- Step 6** Click the **Trusted Platform** subtab.
- Step 7** To configure TPM, click one of the following:

Option	Description
disabled	Disables TPM
enable	Enables TPM
Platform Default	Enables TPM

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

Option	Description
disabled	Disables TXT
enable	Enables TXT
Platform Default	Disables TXT

Step 9 Click **Save Changes**.

Viewing TPM Properties

Procedure

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

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.

Configuring Consistent Device Naming in a BIOS Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand **root**.
- Step 4** Expand **BIOS Policies**.
- Step 5** Select the BIOS policy for which you want to configure CDN.
- Step 6** Under the **Main** tab, click one of the following in the **Consistent Device Naming** field to configure CDN:

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

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

Configuring a CDN Name for a vNIC

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

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the vNIC for which you want to configure a CDN name.
- Step 4** Expand the service profile and **vNICs** node that contain the vNIC for which you want to configure a CDN name.
- Step 5** Select the vNIC.
- Step 6** Click on the **General** tab.
- Step 7** In the **Properties** area, choose **User Defined** as the **CDN Source**.
- Step 8** Enter the CDN name for the vNIC in the **CDN Name** field.

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

Step 9 Click **Save Changes**.

CIMC Security Policies

Cisco UCS Manager provides the following policies to increase security:

- KVM Management Policy
- IPMI Access Profile

IPMI Access Profile

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

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

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

Creating an IPMI Access Profile

Before you begin

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

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

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.
- Step 4** Right-click **IPMI Access Profiles** and select **Create IPMI Access Profile**.
- Step 5** In the **Create IPMI Access Profile** dialog box:

- a) Enter a unique name and description for the profile.
- b) In the **IPMI Over LAN** field, choose whether to allow or restrict remote connectivity.
- c) Click **OK**.

Step 6 In the **IPMI Users** area of the navigator, click +.

Step 7 In the **Create IPMI User** dialog box:

- a) Complete the following fields:

Name	Description
Name field	The username to associate with this IPMI profile. Enter 1 to 16 alphanumeric characters. You can also use @ (at sign), _ (underscore), and - (hyphen). You cannot change this name once the profile has been saved.
Password field	The password associated with this username. Enter 1 to 20 standard ASCII characters, except for = (equal sign), \$ (dollar sign), and (vertical bar).
Confirm Password field	The password a second time for confirmation purposes.
Role field	The user role. This can be one of the following: <ul style="list-style-type: none"> • Admin • Read Only
Description field	User-defined description of the IPMI user.

- b) Click **OK**.

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

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

What to do next

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

Deleting an IPMI Access Profile

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies > Organization_Name**.
- Step 3** Expand the **IPMI Profiles** node.
- Step 4** Right-click the profile you want to delete and select **Delete**.

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

KVM Management Policy

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

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



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

Creating a KVM Management Policy

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the **root** node.
- Step 4** Right-click **KVM Management Policies** and select **Create KVM Management Policy**.
- Step 5** In the **Create KVM Management Policy** dialog box:
- Enter a unique name and description for the policy.
 - In the **vMedia Encryption** field, choose whether to enable vMedia encryption.
 - Click **OK**.

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

Graphics Card Policies

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

- Compute
- Graphics

- Any Configuration

Creating a Graphics Card Policy



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

Procedure

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

Local Disk Policies

Local Disk Configuration Policy

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

The local disk modes include the following:

- **No Local Storage**—For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk.

- **RAID 0 Striped**—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- **RAID 1 Mirrored**—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- **Any Configuration**—For a server configuration that carries forward the local disk configuration without any changes.
- **No RAID**—For a server configuration that removes the RAID and leaves the disk MBR and payload unaltered.

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

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

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

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



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

Guidelines for all Local Disk Configuration Policies

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

No Mixed HDDs and SSDs

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

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

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

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

Name	Description
Mode drop-down list	<p>This can be one of the following local disk policy modes:</p> <ul style="list-style-type: none"> • No Local Storage • RAID 0 Striped • RAID 1 Mirrored • Any Configuration • No RAID <p>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.</p> <p>To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode.</p> <ul style="list-style-type: none"> • RAID 5 Striped Parity • RAID 6 Striped Dual Parity • RAID 10 Mirrored and Striped • RAID 50 Striped Parity and Striped • RAID 60 Striped Dual Parity and Striped <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the <i>Hardware Installation Guide</i> for that server.</p>

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

Step 6 Click **OK**.

Changing a Local Disk Configuration Policy

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

Procedure

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

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

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

Deleting a Local Disk Configuration Policy

Procedure

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

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

FlexFlash Support

Overview

Cisco UCS B-Series, C-Series 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. [Formatting the SD Cards, on page 80](#) provides detailed information.
- Disassociate the service profile from the server. Re-acknowledge the server after modifying the default scrub policy and then associate the server profile back to the server.

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



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

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

FlexFlash Firmware Management

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

http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html.

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 M3 blade servers, and the 64 GB SD cards are supported only on the 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 but there may be data loss.
- The FlexFlash controller does not support RAID-1 sync (mirror rebuild). If the SD cards are in a degraded RAID state, or if any metadata errors are reported by the controller, you must run the FlexFlash scrub policy to pair the cards for RAID. For more information about the FlexFlash scrub policy, see [Server-Related Policies](#). The following conditions might result in degraded RAID or metadata errors:
 - Inserting a new or used SD card in one slot, when the server already has an SD card populated in the second slot.
 - Inserting two SD cards from different servers.
- The server firmware version must be at 2.2(1a) or higher.

FlexFlash FX3S Support

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

- Cisco UCS B200 M4 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

-
- Step 1** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
 - Step 2** In the **Work** pane, check the details of the FlexFlash cards in the **FlexFlash Controller** window.
 - Step 3** Expand **Servers** > **Service Profiles**.
 - Step 4** Expand the node for the organization containing the pool.
If the system does not include multitenancy, expand the **root** node.
 - Step 5** Expand the node for the organization containing the service profile and click **Storage**.
 - Step 6** In the **Work** pane, click **Change Local Disk Configuration Policy** in the **Actions** area and expand **Create Local Disk Configuration Policy** link. Follow the procedure in [Creating a Local Disk Configuration Policy, on page 71](#) to create a Local Disk Configuration Policy.
The FlexFlash policy name must not contain empty spaces or special characters.
 - Step 7** Expand **Change Disk Local Configuration Policy**, and select the policy you just created and click **OK**.
 - Step 8** Expand **Servers** > **Policies**.
 - Step 9** Follow the procedure in [Creating a Scrub Policy, on page 81](#) and create a policy with a name such as *Scrub-FF-name* and click **OK**.
The Scrub policy name must not contain empty spaces or special characters.
 - Step 10** Select the policy you created from the drop-down box.
 - Step 11** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
 - Step 12** In the **Work** pane, click the **General** tab and select **Server Maintenance** from the **Actions** area.
 - Step 13** In the **Maintenance Server** dialogue box, click on the **Re-acknowledge** radio button, and then click **OK**.
 - Step 14** Click **Server Maintenance** in the **Action** area and click on the **Re-acknowledge** radio button again.
 - Step 15** From the **Inventory** tab, select the **Storage** sub-tab.
You can verify details of the enabled FlexFlash cards from the **FlexFlash Controller** window in the **Work** area.

- Step 16** Launch KVM Manager and log on to the operating system. Verify details of the Hypervisor partition from the Devices and drives folder. Depending on the card size, the HV partition displays details of 32GB, 64GB, or 128 GB.
The FlexFlash cards are now synced and ready to use.
-

Enabling FlexFlash SD Card Support

Procedure

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

Enabling Auto-Sync

Procedure

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

Formatting the SD Cards

Procedure

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

Resetting the FlexFlash Controller

Procedure

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

Scrub Policy

Scrub Policy Settings

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



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

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

Disk scrub

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

- If enabled, 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

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

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

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

Step 4 Right-click **Scrub Policies** and select **Create Scrub Policy**.

Step 5 In the **Create Scrub Policy** wizard, complete the following fields:

Name	Description
Name field	The name of the policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).
Disk Scrub field	If this field is set to Yes , when a service profile containing this scrub policy is disassociated from a server, all data on the server local drives is completely erased. If this field is set to No , the data on the local drives is preserved, including all local storage configuration.
BIOS Settings Scrub field	If the field is set to Yes , when a service profile containing this scrub policy is disassociated from a server, the BIOS settings for that server are erased and reset to the defaults for that server type and vendor. If this field is set to No , the BIOS settings are preserved.
FlexFlash Scrub field	If the field is set to Yes , the HV partition on the SD card is formatted using the PNUOS formatting utility when the server is reacknowledged. If this field is set to No , the SD card is preserved.

Step 6 Click **OK**.

Note Disk scrub and FlexFlash Scrub options are not supported for Cisco UCS S3260 Storage Server.

Deleting a Scrub Policy

Procedure

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

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

Step 3 Expand the **Scrub Policies** node.

Step 4 Right-click the policy you want to delete and select **Delete**.

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

DIMM Error Management

DIMM Correctable Error Handling

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

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

Resetting Memory Errors

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

Procedure

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

DIMM Blacklisting

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

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

Enabling DIMM Blacklisting

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

**Note**

- This feature is supported both on the Cisco UCS B-Series blade servers and UCS C-Series rack servers.

**Note**

Cisco UCS C-Series 420 M3 rack servers 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 and S-Series rack server, the server firmware must be at Release 2.2(3).
- You must be logged in with one of the following privileges:
 - Admin
 - Server policy
 - Server profile server policy

Procedure

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

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

**Note**

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

Serial over LAN Policy Settings

Serial over LAN Policy Overview

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

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

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

Creating a Serial over LAN Policy

Procedure

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

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

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

Step 6 Click **OK**.

Deleting a Serial over LAN Policy

Procedure

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

Server Autoconfiguration Policies

Server Autoconfiguration Policy Overview

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

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

Creating an Autoconfiguration Policy

Before you begin

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

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

Procedure

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

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

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

Step 7 Click **OK**.

Deleting an Autoconfiguration Policy

Procedure

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

Server Discovery Policy Settings

Server Discovery Policy Overview

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

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

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

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

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

**Important**

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

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

If this error occurs, do the following:

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

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

Creating a Server Discovery Policy

Before you begin

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

Procedure

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

What to do next

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

Deleting a Server Discovery Policy

Procedure

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

Hardware Change Discovery Policy

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

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

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

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

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

Configuring Hardware Change Discovery Policy

Procedure

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

Step 3 Click **Save Changes**.

Server Inheritance Policy Settings

Server Inheritance Policy Overview

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

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

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

Creating a Server Inheritance Policy

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

Procedure

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

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

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

Step 6 Click **OK**.

Deleting a Server Inheritance Policy

Procedure

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

Server Pool Policy Settings

Server Pool Policy Overview

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

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

Creating a Server Pool Policy

Before you begin

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

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

Procedure

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

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

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

Deleting a Server Pool Policy

Procedure

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

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

Server Pool Policy Qualifications Settings

Server Pool Policy Qualification Overview

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

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

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

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

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

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

Creating Server Pool Policy Qualifications

Procedure

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

Step 2 Expand **Servers > Policies**.

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

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

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

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

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

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

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

- c) Click **OK**.

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

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

Example:

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

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

- c) Click **Finish**.

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

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

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

Step 9

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

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

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

- c) Click **OK**.

Step 10

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

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

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

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

c) Click **OK**.

Step 11

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

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

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

Name	Description
Max Cap field	The maximum storage capacity allowed, in megabytes. To specify a capacity, choose select and enter the maximum storage capacity.
Per Disk Cap field	The minimum storage capacity per disk required, in gigabytes. To specify a capacity, choose select and enter the minimum capacity on each disk.
Units field	The number of units. To specify a capacity, choose select and enter the desired units.
Number of Flex Flash Cards field	The number of FlexFlash Cards. To specify a capacity, choose select and enter the desired units.
Disk Type field	The disk type. This can be one of the following: <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).

c) Click **OK**.

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

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

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

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

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

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

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

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

Deleting Server Pool Policy Qualifications

Procedure

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

Deleting Qualifications from Server Pool Policy Qualifications

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

Procedure

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

vNIC/vHBA Placement Policy Settings

vNIC/vHBA Placement Policies

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

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

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

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

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

For blade or rack servers that contain two or three adapters, Cisco UCS assigns the vCons based on the type of server and the selected virtual slot mapping scheme, which can be **Round Robin** or **Linear Ordered**. For details about the available mapping schemes, see [vCon to Adapter Placement, on page 100](#).

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.

Creating a vNIC/vHBA Placement Policy

Procedure

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

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

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

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

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

c) Click **OK**.

Deleting a vNIC/vHBA Placement Policy

Procedure

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

Explicitly Assigning a vNIC to a vCon

Before you begin

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

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

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

Procedure

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

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

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

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

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

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

Explicitly Assigning a vHBA to a vCon

Before you begin

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

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

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

Procedure

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

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

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

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

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

Step 9 Click **Save Changes**.

Placing Static vNICs Before Dynamic vNICs

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

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

Example

Beginning Device Order in Cisco UCS Manager Release 2.0:

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

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

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

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

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

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

```

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

```

Dynamic vNICs as Multifunction PCIe Devices

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



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

Table 3: Version Compatibility

Cisco UCS Manager		
Version 1.4 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, and the Cisco UCS VIC 1385 and Cisco UCS VIC 1387 adapters.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Select the service profile which is associated with the vNIC that you want to place on a host port.
- Step 4** Expand **Service_Profile_Name > vNICs**
- Step 5** Under the vNICs summary table, double-click the **Admin Host Port** value of the vNIC which you want to configure and select one of the following:
- **Any**—Allows Cisco UCS Manager to determine the host port to which the vNIC is assigned.
 - **1**—Explicitly assigns the vNIC to host port 1.
 - **2**—Explicitly assigns the vNIC to host port 2.

Actual Host Port displays the actual assignment of the vNIC on a host port. When this feature is not supported, this will appear as **None**.

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

CIMC Mounted vMedia

Using Scriptable vMedia

Cisco UCS Manager allows provisioning of vMedia devices iso images for remote UCS servers. Using Scriptable vMedia, you can programmatically mount an IMG or an ISO image on a remote server. CIMC mounted vMedia provide communications between other mounted media inside your datacenter with no

additional requirements media connection. Scriptable vMedia allows you to control virtual media devices without using a browser to manually map each UCS server individually.

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

Cisco UCS Manager Scriptable vMedia supports the following functionality:

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



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

vMedia mount fails when the following conditions are met:

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

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



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 vMedia Policy

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



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

Before you begin

Make sure that you have access to the following:

- Remote vMedia Server
- vMedia Devices

Procedure

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

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

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

Name	Description
Name	Name of the vMedia Mount policy. This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object is saved.
Device Type	The type of remote vMedia you plan to mount. This can be: <ul style="list-style-type: none"> • CDD—Scriptable vMedia CD. • HDD—Scriptable vMedia HDD.

Name	Description
Protocol	<p>The protocol to use when communicating with the remote server. Click one of the following radio buttons to indicate the protocol you want to use to communicate with the mounted remote server. This can be:</p> <ul style="list-style-type: none"> • NFS - Network Files System. • CIFS - Common Internet File System. • HTTP - Hypertext Transfer Protocol. • HTTPS - Hypertext Transfer Protocol over Secure.
Authentication Protocol	<p>The protocol to use for authentication when you use CIFS as the protocol for communicating with the remote server. When you use any protocol other than CIFS, this field is not available. Select one of the following from the drop-down list to specify the authentication protocol.</p> <ul style="list-style-type: none"> • Default—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • None—No authentication is used • Ntlm—NT LAN Manager (NTLM) security protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • Ntlmi—NTLMI security protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • Ntlmssp—NT LAN Manager Security Support Provider (NTLMSSP) protocol. Use this option only with Windows 2008 R2 and Windows 2012 R2. • Ntlmsspi—NTLMSSPi protocol. Use this option only when you enable Digital Signing in the CIFS Windows server. • Ntlmv2—NTLMv2 security protocol. Use this option only with Samba Linux. • Ntlmv2i—NTLMv2i security protocol. Use this option only with Samba Linux. <p>Note The authentication protocol options are available only when you select CIFS as the protocol. For all other protocols, the Authentication Protocol field is disabled.</p>
Hostname/IPAddress	<p>Enter the IP address or hostname of the location where the backup file is to be stored. This can be a server, storage array, local drive, or any read/write media that the fabric interconnect can access through the network.</p> <p>If you use a hostname, you must configure Cisco UCS Manager to use a DNS server. The hostname (DNS) can be used when Inband network is configured for that server.</p>

Name	Description
Image Name Variable	<p>The name to be used for the image. This can be:</p> <ul style="list-style-type: none"> • None—Filename must be entered in the Remote File field. • Service Profile Name—Filename automatically becomes the name of the service profile that the vMedia Policy is associated with. <p>Note</p> <ul style="list-style-type: none"> • If you select Service Profile Name as the Image Name variable, the Remote File field is disabled. • If you select Service Profile Name as the Image Name variable, do not rename the service profile. Renaming the service profile can result in vMedia mount failure.
Remote File	<p>Enter the full path to the ISO or other image file.</p> <p>Note Ensure that the full path to the file begins with “/” after the share name.</p> <p>This field can contain the filename [with the file extension] only.</p>
Remote Path	<p>Enter the share name on the remote server, for example “Share”.</p>
Username	<p>Enter the username that Cisco UCS Manager should use to log in to the remote server.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>
Password	<p>Enter the password associated with the username.</p> <p>This field does not apply if the protocol is NFS. This field is optional if the protocol is HTTP.</p>

Step 8

Click **OK**.

The remote server details are listed in the **vMedia Mounts** area of the **Create vMedia Mount** dialog box.

What to do next

Create a vMedia boot policy.

Adding a vMedia Policy to a Service Profile

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

Before you begin

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

Procedure

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

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

Step 3 Expand the node for the organization where you want to create the service profile.

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

Step 4 Right-click the organization and select **Create Service Profile (expert)**.
The **The Unified Computing System Manager** pane displays.

Step 5 In the **Name** field, enter a unique name that you can use to identify the service profile.

This name can be between 2 and 32 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and this name must be unique across all service profiles and service profile templates within the same organization.

This name must be unique within the organization or sub-organization in which you are creating the service profile.

Step 6 From the **UUID Assignment** drop-down list, do one of the following:

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

Option	Description
	service profiles can access, continue with Step 4. Otherwise, continue with Step 8.

Step 7 (Optional) If you selected the **XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX** option, do the following:

- In the **UUID** field, enter the valid UUID that you want to assign to the server which uses this service profile.
- To verify that the selected UUID is available, click the **here** link.

Step 8 (Optional) If you want to create a new UUID Suffix pool to use to use in this service profile, click **Create UUID Suffix Pool** and complete the fields in the **Create UUID Suffix Pool** wizard.

Step 9 (Optional) In the text box, enter a description of this service profile.

The user-defined description for this service profile.

Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).

Step 10 Click Next.

Step 11 From the **vMedia** drop down list, choose one of the following:

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

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

Step 13 Click Next.

Viewing CIMC vMedia Policy

Before you begin

vMedia Policies are configured.

Procedure

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