

# Cisco Unified Computing System BIOS Settings

Tuning for Performance: Cisco UCS B200 M3 & B420 M3 Blade Servers and C220 M3, C240 M3 & C420 M3 Rack Servers



---

# Contents

<b>What You Will Learn .....</b>	<b>3</b>
<b>Processor Configuration: Intel SpeedStep and Turbo Boost .....</b>	<b>3</b>
<b>Intel Data Direct I/O and Direct Cache Access Support .....</b>	<b>3</b>
<b>Processor C3 and C6 States .....</b>	<b>3</b>
<b>CPU Hyperthreading .....</b>	<b>3</b>
<b>Core Multiprocessing and Latency-Sensitive Single-Threaded Applications .....</b>	<b>3</b>
<b>Energy or Performance Bias .....</b>	<b>4</b>
<b>Power Technology Setting .....</b>	<b>4</b>
<b>CPU Performance Settings and Prefetchers .....</b>	<b>4</b>
Hardware Prefetcher .....	4
Adjacent Cache Line Prefetcher (Buddy Fetch) .....	4
DCU Prefetcher .....	4
DCU-IP Prefetcher .....	4
<b>Memory Performance Settings .....</b>	<b>5</b>
Memory Reliability, Availability, and Serviceability Configuration .....	5
Nonuniform Memory Access .....	5
Low-Voltage Double Data Rate Mode .....	5
ISOC (Isochronous Mode) .....	5
<b>Conclusion .....</b>	<b>5</b>
<b>For More Information .....</b>	<b>5</b>

---

## What You Will Learn

Setting performance options in your system BIOS can be a daunting and confusing task given some of the obscure options you can choose. For most options, you must choose between optimizing a server for power savings or for performance. This document provides some general guidelines and suggestions to help you achieve optimal performance from your Cisco UCS® B200 M3 & B420 M3 Blade Servers and Cisco UCS C220 M3, C240 M3 & C420 M3 Rack Servers.

## Processor Configuration: Intel SpeedStep and Turbo Boost

Intel SpeedStep Technology is designed to save energy by adjusting the CPU clock frequency up or down depending on how busy the system is. Intel Turbo Boost Technology provides the capability for the CPU to overclock itself higher than its stated clock speed if there is enough power to do so. Intel Turbo Boost depends on Intel SpeedStep: if you want to enable Intel Turbo Boost, you must enable Intel SpeedStep first. If you disable Intel SpeedStep, you lose the ability to use Intel Turbo Boost.

Intel Turbo Boost is especially useful for latency-sensitive applications or for scenarios in which the system is nearing saturation and the need arises to accelerate the system by overclocking the CPU. If your system is not running at this saturation level and you want the best performance at a utilization rate of less than 90 percent, you should disable Intel SpeedStep (on systems that allow it) to help ensure that the system is running at its stated clock speed at all times.

## Intel Data Direct I/O and Direct Cache Access Support

The Intel Data Direct I/O and Direct Cache Access option allows network packets to be dropped directly into the Layer 3 processor cache instead of main memory, thereby reducing the number of cycles needed for I/O processing when certain Ethernet adapters are used. This option typically is enabled.

## Processor C3 and C6 States

C3 and C6 are power saving halt and sleep states that a CPU can enter when it is not busy. Unfortunately, it can take some time for the CPU to leave these states and return to a running condition. If you are concerned about performance (for all but latency-sensitive single-threaded applications), and if you have the option, disable anything related to C states.

## CPU Hyperthreading

You should test the CPU hyperthreading option both enabled and disabled in your specific environment. If you are running a single-threaded application, you should disable hyperthreading.

## Core Multiprocessing and Latency-Sensitive Single-Threaded Applications

The core multiprocessing option is designed to give the user the capability to disable cores. This option may affect the pricing of certain software packages that license by the core. You should consult your software license and software vendor about whether disabling cores qualifies you for any particular pricing policies. Set core multiprocessing to "All" if pricing policy is not an issue for you. For latency-sensitive single-threaded applications, you can optimize performance by disabling unnecessary cores, disabling hyperthreading, enabling all C states, enabling Intel SpeedStep, and enabling Intel Turbo Boost. With this configuration, the remaining cores often will benefit from higher turbo speeds and better use of the shared Layer 3 cache.

## Energy or Performance Bias

You can use the power-saving mode to reduce system power consumption when the turbo mode is enabled. The mode can be set to Maximum Performance, Balanced Performance, Balanced Power, or Power Saver. Testing has shown that most applications run best with the Balanced Performance setting.

## Power Technology Setting

For best performance, always set the power technology option to Custom. If it is not set to Custom, the individual settings for Intel SpeedStep and Turbo Boost and the C6 power state are ignored.

## CPU Performance Settings and Prefetchers

Intel Xeon processors have several layers of cache. Each core has a tiny Layer 1 cache sometimes referred to as a data cache unit (DCU) that has 32 KB for instructions and 32 KB for data. Slightly bigger is the Layer 2 cache, with 256 KB shared between data and instructions per core. In addition, all cores on a chip share a much larger Layer 3 cache, which is about 10 to 20 MB in size (depending on the processor model). The prefetcher settings provided by Intel primarily affect the Layer 1 and Layer 2 cache on a processor core (Table 1). You will likely need to perform some testing with your individual workload to find the combination that works best for you. See Table 2 for guidance.

**Table 1.** CPU Performance and Prefetch Options from Intel

Performance Option	Cache Affected
Hardware prefetcher	Layer 2
Adjacent cache line prefetcher	Layer 2
DCU prefetcher	Layer 1
DCU Instruction Pointer (DCU-IP) prefetcher	Layer 1

### Hardware Prefetcher

The hardware prefetcher prefetches additional streams of instructions or data into the Layer 2 cache upon detection of an access stride. This behavior is more likely to happen when sorting through sequential data, such as database table scans or clustered index scans, or when running a tight loop in code.

### Adjacent Cache Line Prefetcher (Buddy Fetch)

The adjacent cache line prefetcher always prefetches the next cache line. Although this approach works well when accessing data sequentially in memory, it can quickly pollute the small Layer 2 cache with unneeded instructions or data if the system is not accessing data sequentially, causing frequently accessed instructions and code to leave the cache to make room for the “buddy” data or instructions.

### DCU Prefetcher

Like the hardware prefetcher, the DCU prefetcher prefetches additional streams of instructions or data upon detection of an access stride, except it stores the streams in the tiny Layer 1 cache instead of the Layer 2 cache.

### DCU-IP Prefetcher

The DCU-IP prefetcher predictably prefetches data into the Layer 1 cache on the basis of the recent instruction pointer load instruction history.

**Table 2.** Cisco UCS CPU Prefetcher Options and Target Benchmarks and Workloads

Prefetchers	Targeted Benchmarks and Workloads
Hardware prefetcher disabled; all others enabled	TPC-C online transaction processing (OLTP) database
All enabled	High-performance computing (HPC) benchmarks, webserver, SAP Application Server, and TPC-E
DCU-IP enabled; all others disabled	SPECjbb2005 and server-side Java application server
All disabled	VMmark and cloud virtualization benchmarks

## Memory Performance Settings

### Memory Reliability, Availability, and Serviceability Configuration

Always set the memory reliability, availability, and serviceability (RAS) configuration to Maximum Performance for systems that require the highest performance and do not require memory fault-tolerance options.

### Nonuniform Memory Access

Most modern operating systems, particularly virtualization hypervisors, support nonuniform memory access (NUMA) because in the latest server designs, a processor is attached to a memory controller: meaning that half the memory belongs to one processor, and half belongs to the other processor. If a core needs to access memory that resides in another processor, a longer latency period is needed to access that part of memory. Operating systems and hypervisors recognize this architecture and are designed to reduce such trips. For hypervisors such as those from VMware and for modern applications designed for NUMA, keep this option enabled.

### Low-Voltage Double Data Rate Mode

Low-voltage double data rate (LV DDR) memory can typically run at two voltages: 1.35V and 1.5V. At 1.5V, the memory will run up to 1600 MHz (assuming that you have dual-rank 1600-MHz DIMMs in the system and that the installed CPU supports memory running at 1600 MHz). For best performance, set this option to Performance Mode. The goal is to ensure that the memory is running at the maximum speed permitted by the DIMMs and the system. This performance can be verified at boot time, or by viewing the F2 BIOS and selecting the Advanced screen and then the Memory Configuration screen.

### ISOC (Isochronous Mode)

Enabling the Isochronous Mode ("ISOC") option reduces the credits available for memory traffic. For memory requests, this option reduces latency at the expense of throughput under heavy loads.

## Conclusion

When tuning system BIOS settings for performance, you need to consider a number of processor and memory options. If performance is your goal, make sure to choose options that optimize for performance in preference to power savings, and experiment with other options such as CPU prefetchers and CPU hyperthreading.

## For More Information

Cisco UCS B200 M3 Blade Server: <http://www.cisco.com/en/US/products/ps12288/index.html>

Cisco UCS B420 M3 Rack Server: <http://www.cisco.com/en/US/products/ps12527/index.html>

Cisco UCS C220 M3 Rack Server: <http://www.cisco.com/en/US/products/ps12369/index.html>

Cisco UCS C240 M3 Rack Server: <http://www.cisco.com/en/US/products/ps12370/index.html>

---

Cisco UCS C420 M3 Rack Server: <http://www.cisco.com/en/US/products/ps12770/index.html>



---

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C07-614438-02 05/13