

Hewlett Packard
Enterprise

HPE Server Performance Management and Tuning Guide

Abstract

This guide is intended for server administrators, and describes how the HPE server performance management and tuning features improve performance in HPE Gen10 and later server environments.

Part Number: 882269-006
Published: December 2019
Edition: 1

© Copyright 2017-2019 Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

SPEC® and the benchmark names SPEC CPU®, SPECfp®, and SPECint® are registered trademarks of the Standard Performance Evaluation Corporation (SPEC). Results shown are from internal SPEC CPU2006 results as of 09-20-2017. For more information, see **http://www.spec.org**.

AMD is a trademark of Advanced Micro Devices, Inc.

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Contents

Performance management 5
- Performance management feature requirements 5

Workload matching 7
- Workload profiles 7
- Selecting a workload profile by using iLO 5 and Always On Intelligent Provisioning 8

Jitter smoothing 10
- Sources of jitter within the processor 10
  - P-states and power management 10
  - Turbo Boost 10
  - C-states 11
  - Power and thermal events 11
  - Advanced Vector Extensions 11
- Jitter and latency 11
- Processor Jitter Control 12
  - Auto Processor Jitter Control mode 12
  - Manual Processor Jitter Control mode 13
  - Configuring Processor Jitter Control through the UEFI System Utilities 13
  - Processor Jitter Control configuration options in the iLO RESTful API 14
  - Configuring jitter smoothing settings through iLO 5 15
- Increasing performance with Processor Jitter Control 17

Performance monitoring 20
- Performance monitoring in iLO 5 20
  - Performance data details 21
  - Viewing performance data 21
  - Configuring performance alerts 23

Workload advisor 25
- Workload advisor in iLO 5 25
  - Server workload details 25
  - Viewing server workload details 26
  - Configuring the performance tuning options 26

Core boosting 29
- Processor core boosting architecture 29
- The core boosting turbo profile 30
- Increasing performance with core boosting 31
- Configuring core boosting by using iLO 5 and Always On Intelligent Provisioning 32
  - Core boosting options 33
- Core boosting Integrated Management Log and error messages 33

Websites 34

Support and other resources 35

Accessing Hewlett Packard Enterprise Support 35
Accessing updates 35
Remote support 36
Warranty information 36
Regulatory information 36
Documentation feedback 37

Performance management

Selected HPE Gen10 and later servers support the following server performance management and tuning features:

- **Workload matching**—Use preconfigured server profiles to maximize application performance.
- **Jitter smoothing**—Use the **Processor Jitter Control Mode** setting to level and balance frequency fluctuation (jitter) resulting in lower latency.
- **Performance monitoring**—View performance data collected from supported sensors on servers with Innovation Engine support. You can configure alerts based on the collected data.
- **Workload advisor**—View selected server workload characteristics. You can view and configure recommended performance tuning settings based on the monitored data.
- **Core boosting**—Enable this feature to produce higher performance across more active processor cores.

  This feature is supported on Gen10 servers only. It is not supported on Gen10 Plus servers.

If you reset iLO to the factory default settings, all performance management settings and data are deleted.

When you use the iLO backup and restore feature, the performance management settings are retained. The collected performance data is not backed up or restored.

Performance management feature requirements

Requirement	Workload matching	Jitter smoothing	Core boosting	Performance monitoring	Workload advisor
HPE Gen10 Plus servers	✓	✓		✓	✓
HPE Gen10 servers	✓	✓	✓ [1]	✓	✓
Intel processor		✓	✓ [2]	✓	✓
iLO 5	✓	✓	✓	✓	✓
iLO Advanced license		✓	✓	✓	✓
Minimum System ROM	1.00	1.00 static 1.20 dynamic 1.40 optimizations	1.20	2.00	2.00
Minimum iLO firmware	N/A	1.15 iLO RESTful API 1.30 iLO web interface	1.15 iLO RESTful API 1.30 iLO web interface	1.40 iLO RESTful API 1.40 iLO web interface	1.40 iLO RESTful API 1.40 iLO web interface
Minimum HPE Innovation Engine Firmware[3]	N/A	1.2.4	1.2.4	2.0.11	2.0.11

[1] Select servers only; requires high performance heatsinks and fans.

[2] Select Intel processors only (end of life 2019)

[3] The iLO web interface **Performance** page is not available on servers without Innovation Engine support. To verify Innovation Engine support, look for the Innovation Engine firmware on the **Installed Firmware** page.

Workload matching

The default BIOS settings on Hewlett Packard Enterprise servers provide a balance between performance and power efficiency. These settings can be adjusted to match specific application workloads.

HPE Gen10 and later servers offer a UEFI configuration option to help customers tune their BIOS settings by using known workload-based tuning profiles. When matching your workload profile setting to your actual deployed workload, you can realize performance gains versus just using the out-of-box BIOS defaults.

For more information, see the *UEFI Workload-based Performance and Tuning Guide for HPE ProLiant Gen10, ProLiant Gen10 Plus Servers, and HPE Synergy* at **https://www.hpe.com/support/Workload-UG-en**.

Workload profiles

The following system-generated workload profiles can be used to increase server performance:

General Power Efficient Compute

Applies the most common performance and power management settings. Recommended for users who do not tune BIOS settings to match workloads.

General Peak Frequency Compute

Applies performance and power management settings to achieve the maximum frequency possible for any individual core. Recommended for workloads that benefit from faster compute time.

General Throughput Compute

Applies performance and power management settings to achieve total maximum sustained throughput. Optimized to support NUMA (non-uniform memory access) awareness.

Virtualization - Power Efficient

Applies performance settings to enable all virtualization options. Manages power settings so that they do not interfere with virtualization. Recommended for virtualization environments.

When this workload profile is selected, you cannot enable the **Processor Jitter Control Mode** feature.

Virtualization - Max Performance

Applies performance settings to enable all virtualization options. Disables power settings to deliver optimum performance. Recommended for virtualization environments.

Low Latency

Applies reductions to speed and throughput and disables power management to lower overall computational latency. Recommended for RTOS (Real-Time Operating Systems) workloads or other workloads that are latency sensitive.

Mission Critical

Manages advanced memory RAS (reliability, availability, and serviceability) features. This profile is intended to be used by customers who trade off performance for server reliability above the basic server defaults.

Transactional Application Processing

Manages peak frequency and throughput. Recommended for processing environments such as those using OLTP (online transaction processing) applications that require a back-end database.

High Performance Compute (HPC)

Disables power management to optimize sustained available bandwidth and processor compute capacity. Recommended for users running a traditional HPC environment.

Decision Support

This profile is intended for Enterprise Business Database (Business Intelligence) workloads that are focused on operating or accessing data warehouses, such as data mining or OLAP (online analytical processing).

Graphic Processing

Power management and virtualization are disabled to optimize the bandwidth between I/O and memory. Recommended for workloads running on servers that use GPUs (Graphics Processing Units).

I/O Throughput

Disables power management features that impact the links between I/O and memory. Recommended for configurations that depend on maximum throughput between I/O and memory.

Custom

Disables workload profiles. Recommended for users who want to set specific BIOS options.

Selecting a workload profile by using iLO 5 and Always On Intelligent Provisioning

Prerequisites

- Configure iLO Settings privilege
- Host BIOS privilege
- Remote Console privilege
- Version 1.20 or later of the System ROM (BIOS) is installed.
- The server has Innovation Engine support.

  The **Performance** page is not available on servers that do not have Innovation Engine support. To verify Innovation Engine support, look for the Innovation Engine firmware on the **Installed Firmware** page.
- MCTP discovery is enabled.
- The server is powered on and POST is complete.
- The latest version of Intelligent Provisioning is installed.

Procedure

1. Click **Performance** in the navigation tree.

   iLO displays the **Settings** tab.
2. Click **Settings**.
3. To start Intelligent Provisioning, click **Configure with Always On**.

   The Intelligent Provisioning web interface starts in a new browser window.
4. Click **Perform Maintenance**.
5. Click **BIOS/Platform Configuration**.

   The **BIOS/Platform Configuration** page opens.
6. Select a profile from the **Workload Profile** list.

NOTE: Selecting a workload profile automatically configures many options on the RBSU **Power and Performance Options** screen. To modify power and performance options yourself, select **Custom** from the **Workload Profile** list.

The change is pending. You can click **View Change** to view the old and new settings.

7. Click **Update**.

   Intelligent Provisioning applies the changes and notifies you that a server reboot is required for the changes to take effect.

8. Reboot the server.

Jitter smoothing

For the past several years, server-class customers have seen processor-based performance increase generation over generation. This increase is due in a large part to increases in core counts and more efficient instruction set architectures. Unlike the preceding decades, the base frequency of the CPU has stayed rather stable with performance improvements coming from increasing core counts and architectural enhancements. However, processor vendors began to realize that not all workloads benefit from increased core counts, so they introduced features that allow some cores to run opportunistically at higher frequencies if power headroom is available or other cores are underutilized.

Although these opportunistic frequency upsides can increase performance, they also introduce an unwanted side effect. Frequency shifting itself introduces computation jitter, or nondeterminism, and undesirable latency. Jitter and the latency associated with it create problems for several customer segments. For example, high-frequency traders rely on time-sensitive transactions. They cannot tolerate the microseconds of delay that can be added non-deterministically to a trade, caused by a frequency shift. These delays over time can cost a trader upwards of millions of dollars. In other environments, servers running RTOS (real-time operating systems) to control critical functionality cannot tolerate random latencies that happen when opportunistic-frequency features are left enabled.

The current trend for latency-sensitive customers is to disable the features that normally would result in increased application performance because of the associated jitter. A trade executes faster if the processor runs faster, but if it comes at the cost of random delay, the benefit of increased performance is lost.

Hewlett Packard Enterprise introduced the Processor Jitter Control feature in its Gen10 and later servers to enable customers to achieve both frequency upside and low jitter. This feature is available for Gen10 and later servers using Intel Xeon Scalable Processors. Servers using AMD processors do not support this feature. This feature allows the customer to remove or reduce jitter caused by opportunistic frequency management, which results in better latency response and higher throughput performance.

Enabling the Processor Jitter Control feature might require changes to power management settings under the operating system.

More information

Auto Processor Jitter Control mode

Sources of jitter within the processor

The Intel processor introduces jitter in cores at any time they change the operating frequency. Several possible reasons exist for a processor to dynamically change frequencies during run time. Some of the sources that request frequency changes are driven by software, while others are driven by the processor itself.

P-states and power management

P-states are predefined performance states that are made available by the processor for software to control how much performance the processor can deliver so that it can manage the power-performance efficiency of the platform. Performance states are mapped to specific frequencies at which the processor is capable of operating. Power management software instructs the processor to change P-states (frequency) to save power when processor utilization (demand) is low. A processor often offers several different P-states over a range of operating frequencies.

Turbo Boost

The Intel Turbo Boost feature allows the processor to run at higher frequencies than the base frequency stated in its specification, assuming it follows certain conditions. The conditions include the amount of heat being dissipated, the temperature of the part, and the number of cores active (enabled and not idle.) When a workload is run on these processors while Turbo Boost is enabled, the processor will opportunistically switch between frequencies in an attempt to achieve the highest possible performance. But as the demands of the workload change, so can the frequencies. When frequencies change, we get frequency jitter as well as a small amount of latency that occurs, which is required to electrically change frequencies.

The net effect of having Turbo Boost enabled is that while the processor attempts to provide the maximum amount of performance within its limits, it typically does so by changing frequency often.

C-states

C-states are predefined power-saving states that the processor offers to power management software to use when the operating system idles a processor core. The operating system puts the processor into one of a number of C-states that are made available. The deeper the C-state, the more power that is saved, but at the cost of longer exit latencies to return to the operating state. In an attempt to save power, C-states on Intel processors also lower the frequency of the processor. Upon exiting a C-state, the processor, running at the lowest frequency available to the C-state, must perform an additional frequency shift to return to the P-state previously requested by power management software.

C-states are useful in saving power when the processor is not being used. However, entering and exiting these states introduces a large amount of jitter. Some Linux distributions that leverage the `intel_idle` driver ignore ACPI C-state reporting, which might cause issues when Processor Jitter Control is configured to run in auto-tuned mode.

More information

Auto Processor Jitter Control mode

Power and thermal events

The processor, in an attempt to run within the constraints of its design, employs the use of frequency throttling to protect itself from thermal or overcurrent conditions.

Frequency throttling controls the amount of stress that the workload can introduce on the host. Higher stress levels result in higher heat and current draw. Several factors can lead to high operating temperature or overcurrent events. Server ambient temperature, air flow, and other factors all play an important role in processor temperature.

An overcurrent event can occur when the processor executes very high-demand, power-hungry workloads that can consume significant resources within the processor itself. Overcurrent can also occur if Turbo Boost is enabled and the processor attempts to maximize the amount of performance when a particularly aggressive workload executes and the power that is consumed is driven very high.

Advanced Vector Extensions

Server processors offer Advanced Vector Extensions (AVX) that perform complex math at the cost of utilizing logic, which can drive very high power usage inside the processor itself.

If left unchecked, overcurrent throttling is required when these instructions eventually drive the processor to consume higher power. Instead of reactively throttling, processors typically proactively force cores to run at a lower frequency to limit the chances of extreme power excursions whenever those instructions are executed.

On Intel processors, the use of AVX instructions causes the processor to limit the processor frequency automatically. Because these instructions cause the processor to automatically limit and potentially lower the frequency, their usage often introduces jitter.

Jitter and latency

Jitter and latency are directly related. Jitter induced by processor frequency changes introduces changes in latency as observed by a workload. When a processor executes a change in frequency, it goes through a process that causes thread execution to stop entirely before the processor is capable of running at the new chosen frequency. This process occurs regardless of whether the processor will shift to a faster or slower frequency. The amount of time that the processor is stopped can vary, but is typically between 10 microseconds and 15 microseconds. For a workload that depends on processor execution, a change in frequency will then always introduce an additional 10-15 microseconds of latency.

Because frequency shifts are often asynchronous to application tasks running on the server, these latencies are random and thus non-deterministic from an application standpoint. Also, it is important to note that a processor that varies its frequency also creates a non-deterministic level of performance for the running applications. Software will execute slower

as the frequency is lowered and vice versa. The difference in frequency itself also means that there is a variable amount of latency involved when an application depends on a certain amount of execution time.

Latency introduced by frequency changes can be illustrated by measuring latency when the processor is configured to allow for frequency shifts (for example, when Turbo Boost is enabled). **Figure 1: Latency from frequency shifts** illustrates how latency spikes in the 10-15 microsecond range occur when the processor is allowed to shift frequency.

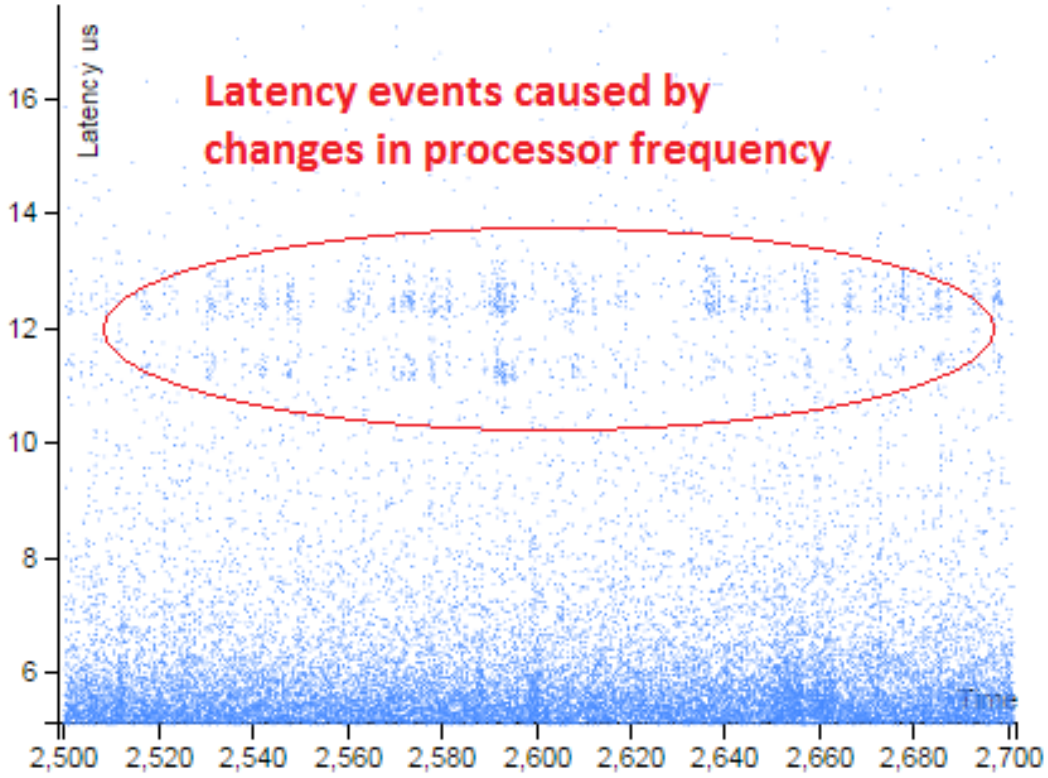

Figure 1: Latency from frequency shifts

Processor Jitter Control

The Processor Jitter Control feature uses the platform firmware in HPE ProLiant Gen10 and later servers with Intel Xeon Scalable Processors to reduce or remove processor jitter (jitter smoothing). Use the UEFI System Utilities, the iLO RESTful API, or the iLO 5 web interface to configure auto-tune or manual mode, or to disable the feature.

NOTE: Processor Jitter Control requires HPE Power Regulator to be set to any mode other than OS Control. The BIOS software does not automatically modify the Power Regulator setting when Processor Jitter Control is enabled because if it is disabled during runtime, the system switches back to the behavior of the Power Regulator setting.

Auto Processor Jitter Control mode

When Processor Jitter Control is configured to run in auto mode, HPE server firmware disables the impact of power management and dynamically adjusts the processor during runtime to eliminate frequency shift induced jitter. The result of running in auto-tuned mode is that the processor will eventually run at the highest frequency that can be achieved where the processor stops making frequency changes to stay within its thermal, power, and core usage constraints. Auto-tuned mode lowers the frequency upon detection of frequency changes caused by the following sources:

- C-state transitions
- AVX induced transitions
- Turbo transitions (due to power, thermal, and core usage)
- Thermal throttling

Jitter control equally affects all cores on all processors.

Using auto-tuned mode in Linux distributions that load the `intel_idle` driver

When you select auto-tuned mode through the UEFI System Utilities, C-state settings are set to disabled. Most operating systems rely on the BIOS to report C-state support. The reporting is done through the Advanced Configuration and Power Interface (ACPI).

WARNING: Certain Linux distributions that load the `intel_idle` driver ignore the ACPI reporting of C-state support. For auto-tuned mode to function properly with a Linux distribution that ignores ACPI reporting of C-state support, you must disable the `intel_idle` driver.

You can disable the `intel_idle` driver by adding **`intel_idle.max_cstate=0`** in the kernel boot command parameters.

Manual Processor Jitter Control mode

When Processor Jitter Control mode is set to manual, the processor is configured to run no faster than a user selectable frequency. In this mode, firmware does not lower the user-selected upper frequency limit dynamically even if processor frequency changes are detected. Manual mode is useful for users who want to manually tune for jitter reduction by setting a maximum operational frequency. In contrast to auto mode, if a frequency change occurs below the programmed frequency, the server will not reduce the operating frequency permanently. The processor returns to the user-selected maximum frequency when the limiting constraints no longer exist.

Configuring Processor Jitter Control through the UEFI System Utilities

Prerequisites

An iLO license that supports this feature is installed.

Procedure

1. Restart or power on the server.
2. When the POST screen appears, Press **F9**.

   The UEFI System Utilities screen appears.
3. Click **System Configuration** > **BIOS/Platform Configuration (RBSU)** > **Power and Performance Options** > **Advanced Performance Tuning Options**.

   The **Processor Jitter Control** and **Processor Jitter Control Frequency** options are available when the server supports jitter control.

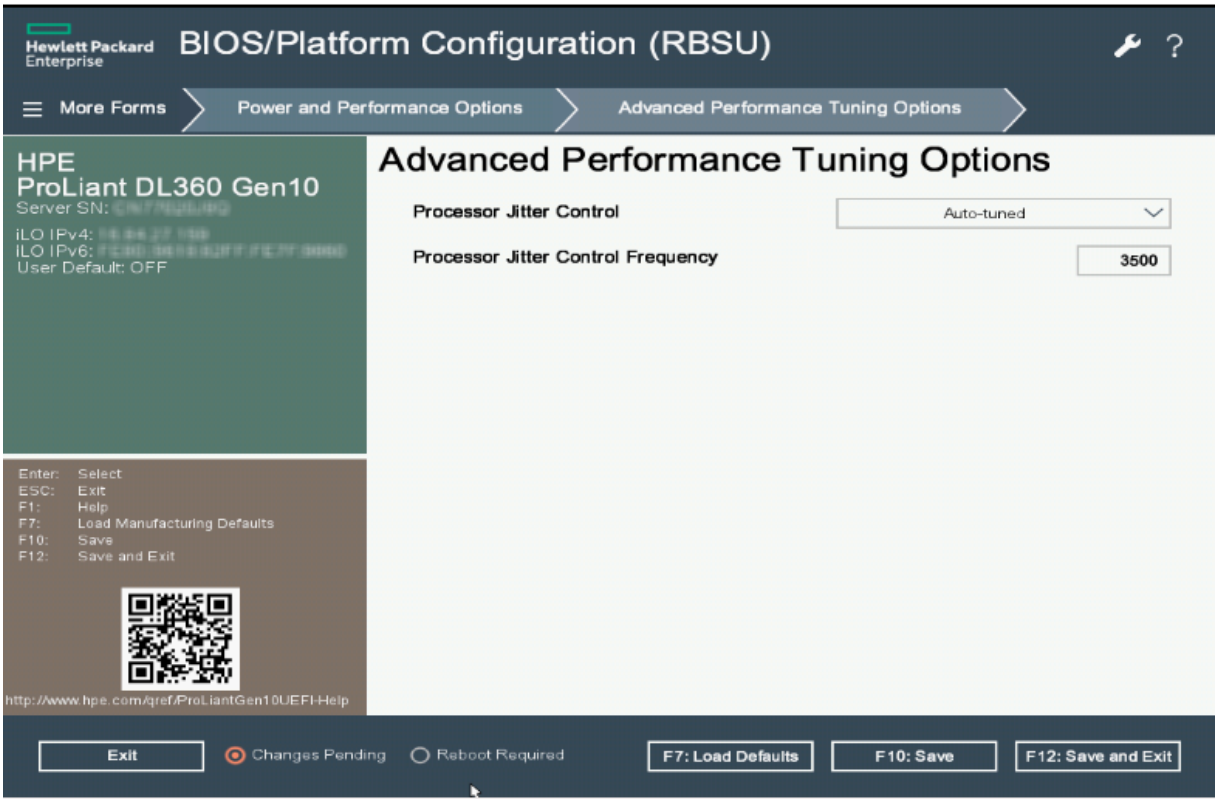

Figure 2: The Processor Jitter Control option in UEFI System Utilities (v1.10)

4. Configure the **Processor Jitter Control** option.

   The **Processor Jitter Control** option has three modes: **Disabled**, **Auto-tuned**, and **Manual-mode**.

   Selecting **Auto-tuned** or **Manual-mode** activates the **Processor Jitter Control Frequency** input option. This setting allows you to select the target frequency for manual-tuned mode or the starting maximum frequency for auto-tuned mode. Frequency is entered in units of MHz and the system firmware rounds up to the nearest frequency interval allowed by the processor. For example, Intel Xeon Scalable Processors support frequency programming in intervals of 100 MHz. If a user enters 2,050 MHz, the resulting frequency will be 2,100 MHz if supported by the installed processor.

5. Enter a **Processor Jitter Control Frequency**.

6. Follow the onscreen instructions to save the changes.

   For more information, see the UEFI System Utilities documentation.

More information

Manual Processor Jitter Control mode

Auto Processor Jitter Control mode

Processor Jitter Control configuration options in the iLO RESTful API

iLO 5 server management provides intelligent remote control automation using the iLO RESTful API, which integrates server management components and full compute power. You can use the iLO RESTful API to manage the complete life cycle of HPE Gen10 and later servers through Redfish API conformance.

You can use the iLO RESTful API to configure Processor Jitter Control during runtime operation of the server. Using a RESTful Interface Tool such as `ilorest`, Processor Jitter Control can be switched between Disabled, Auto-tuned, or Manual-tuned without requiring a server reboot. The iLO RESTful API allows this change in configuration to be applied

from a tool running on the server locally or remotely. Additionally, when in Auto-mode, the iLO RESTful API can be used to determine which frequency the server has settled at for jitter-free operation.

The properties that are provided by the iLO RESTful API that pertain to runtime Processor Jitter Control are located under the URI **`/redfish/v1/systems/1`**. The Processor Jitter Control property is at **`{"Oem":{"Hpe":{"ProcessorJitterControl"}}}`** and contains the following two configurable properties **`FrequencyLimitMHz`** and **`Mode`**. The following are examples of how to configure Processor Jitter Control using the iLO RESTful API:

- To configure Processor Jitter Control for Auto-mode to use the default frequency, the following properties are modified:

  `{"Oem":{"Hpe":{"ProcessorJitterControl":{"FrequencyLimitMHz": 0,"Mode": "Auto"}}}}`

- To configure Processor Jitter Control for Manual-mode using 2.6 GHz:

  `{"Oem":{"Hpe":{"ProcessorJitterControl":{"FrequencyLimitMHz": 2600,"Mode": "Manual"}}}}`

- To disable Processor Jitter Control:

  `{"Oem":{"Hpe":{"ProcessorJitterControl":{"Mode": "Disabled"}}}}`

For more information about using the iLO RESTful API, see the following website: **https://www.hpe.com/us/en/servers/restful-api.html**.

NOTE: CONREP is not a supported utility for performing dynamic updates during runtime.

More information

Manual Processor Jitter Control mode
Auto Processor Jitter Control mode

Configuring jitter smoothing settings through iLO 5

Prerequisites

- Configure iLO Settings privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.
- If used, HPE Power Regulator is set to any mode other than OS Control.
- The server has Innovation Engine support and version 1.2.4 or later of the Innovation Engine firmware is installed.
- MCTP discovery is enabled.
- Version 1.20 or later of the System ROM (BIOS) is installed. Version 1.40 or later is required for using the **Processor Jitter Control Optimization** setting.
- The server is powered on and POST is complete.
- For the **Processor Jitter Control Optimization** feature only: If iLO was reset to the factory default settings, the server OS was rebooted.

Procedure

1. Click **Performance** in the navigation tree.

iLO displays the **Settings** tab.

2. Click **Settings**.
3. Select a **Processor Jitter Control Mode**.
4. If applicable, enter the **Processor Jitter Control Frequency (in MHz)**.
5. Select a **Processor Jitter Control Optimization** value.
6. Click **Apply**.

   iLO prompts you to confirm the change.
7. Click **Yes**.

Jitter smoothing options

Processor Jitter Control Mode

This feature reduces or removes processor jitter (jitter smoothing). The following settings are available:

- **Auto**—Monitors frequency variance, and automatically adjusts the frequency to minimize variance over time.

  When you select **Auto**, consider the following:

  - Certain Linux distributions that load the `intel_idle` driver ignore the ACPI reporting of C-state support. For **Auto** mode to work with a Linux distribution that ignores ACPI reporting of C-state support, you must disable the `intel_idle` driver.

    To disable the `intel_idle` driver, add **`intel_idle.max_cstate=0`** in the kernel boot command parameters.
  - If you select **Auto** when the **Minimum Processor Idle Power Package C-state** has an enabled C-state value, the **Processor Jitter Control Frequency (in MHz)** automatically decreases to zero, and **Processor Jitter Control Mode** is set to **Disabled**. When a C-state value is enabled, using the **Auto** value is not supported.

- **Manual**—Operates the processor at a fixed frequency, and enables you to select lower or higher frequencies statically.
- **Disabled**—Disables **Processor Jitter Control Mode**.

You cannot set this option to **Auto** or **Manual** when the **Workload Profile** is set to **Virtualization - Power Efficient**.

Processor Jitter Control Frequency (in MHz)

Enter this value when **Processor Jitter Control Mode** is set to **Auto** or **Manual**.

- When **Auto** is configured, enter a starting frequency unit in MHz. To specify the maximum allowable speed, enter 0.
- When **Manual** is configured, enter a frequency unit in MHz.

You can enter a value from 0 to 10000. The supported frequency range depends on the processor model. Typically, it ranges from 1,000 MHz to 4,000 MHz.

Frequency is entered in units of MHz and the system firmware rounds down to the nearest frequency interval allowed by the processor. For example, Intel Xeon Scalable Processors support frequency programming in intervals of 100 MHz. If a user enters 2,050 MHz, the resulting frequency will be 2,000 MHz if supported by the installed processor.

Processor Jitter Control Optimization

- **Optimized For Throughput**—Threshold and polling rate are programmed for maximum throughput.
- **Optimized For Latency**—Threshold and polling rate are programmed for low latency.
- **Zero Latency**—Threshold and polling rate are programmed for zero latency.

This feature is disabled when **Processor Jitter Control Mode** is set to **Manual**.

When iLO is reset to the factory default settings, **Processor Jitter Control Optimization** is unavailable until the server OS is rebooted.

Increasing performance with Processor Jitter Control

Jitter control can be used to tune for best performance in workloads that are traditionally sensitive to changes in latency, as well as in workloads that are impacted by excessive amounts of frequency shifting.

Users that traditionally tune servers for low latency and follow the HPE technical white paper *Configuring and tuning HPE ProLiant Servers for low-latency applications* are familiar with the practice of disabling P-state power management, C-states, and Turbo mode to eliminate jitter caused by frequency shifting by the processor. However, a significant amount of performance can be gained by leaving Turbo Boost enabled and enabling Processor Jitter Control.

NOTE: The *Configuring and tuning HPE ProLiant Servers for low-latency applications* technical white paper is available on the following website: **http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c01804533**

Although the frequencies above the base frequency of the processor and up to the maximum Turbo Boost mode frequency are not guaranteed, a workload is not likely to encounter limiting constraints if the processor is limited to a specific frequency within the Turbo Boost frequency range. Processor Jitter Control, when configured for auto-mode, finds the frequency dynamically for any workload and environment in which the server is deployed. If the resulting auto-tuned frequency is higher than the base frequency, gains in performance are achieved over the conservative practice of disabling Turbo Boost. **Figure 3: Results of using Processor Jitter Control** demonstrates the gains that can be achieved using Processor Jitter Control with an HPE ProLiant DL360 Gen10 server with Intel Xeon Platinum 8180 Scalable Processor.

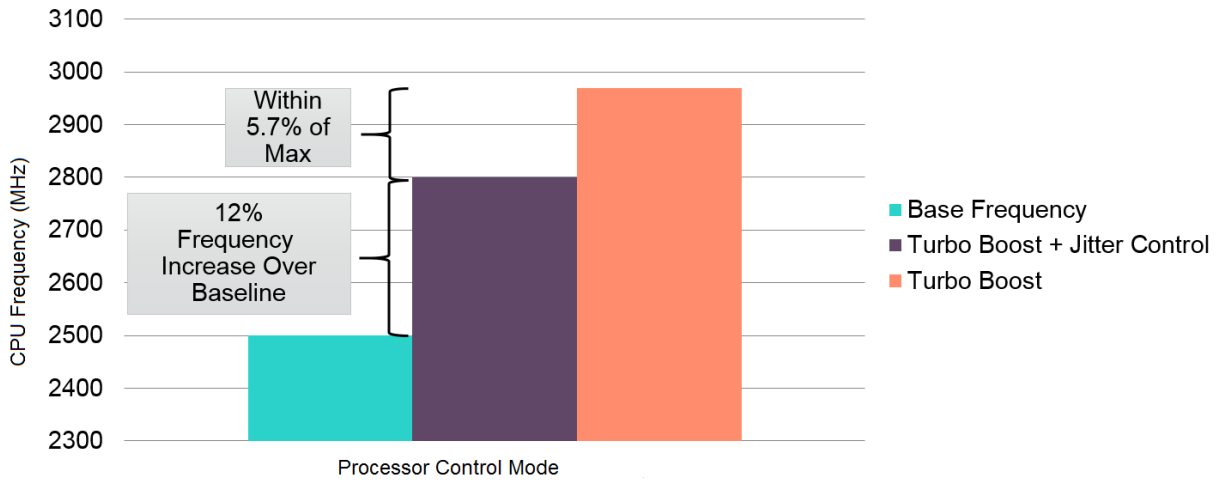

Figure 3: Results of using Processor Jitter Control

When using Processor Jitter Control, a workload might even surpass the performance seen by enabling Turbo Boost alone. If a workload drives the processor to a state where Turbo Boost is constantly making frequency shifts to maximize performance, but cannot stay at the higher frequencies for a substantial time period, then the gains by attempting to run

at higher frequencies are nullified, and, in some cases, performance might be reduced. When using the HPE ProLiant DL360 Gen10 with an Intel Xeon Platinum 8180 Scalable Processor running a server-side Java workload, it is possible to see gains beyond using Turbo Boost alone.

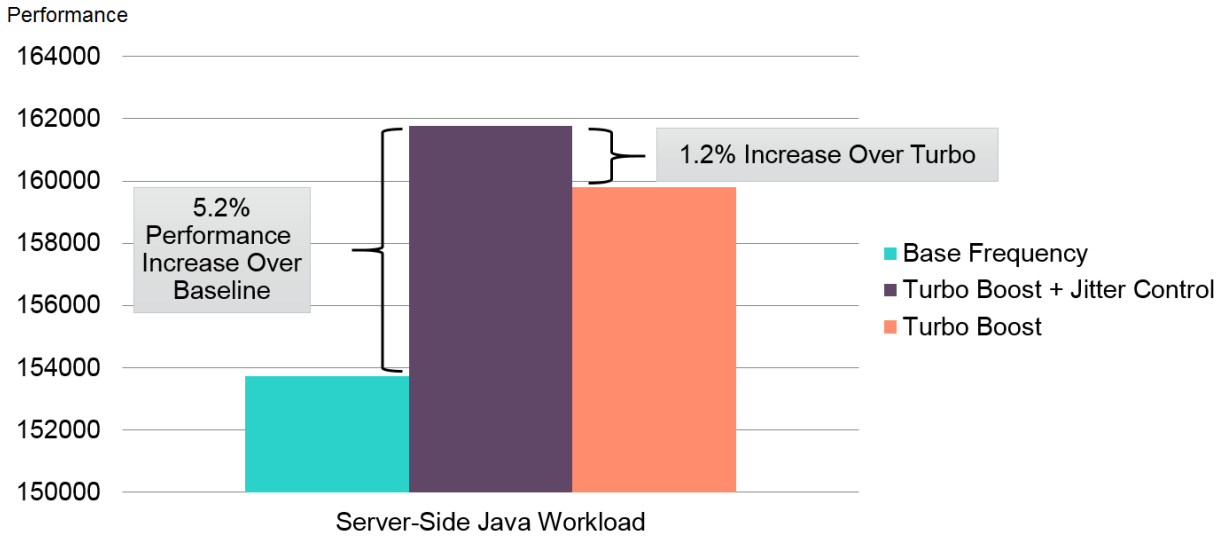

Figure 4: Results of Processor Jitter Control vs Turbo Boost

Jitter control is designed to reduce latency by limiting the causes for frequency changes. While throughput performance metrics are impacted by enabling Jitter Control, the removal of latency spikes is critical for workloads that either cannot tolerate varying latency or have costs associated with changing computational latencies. When using auto-tuned mode, jitter within the 10-15 µsec range caused by frequency changes can be eliminated.

The images in **Figure 5: Latency comparison with Processor Jitter Control** show the latencies measured by the HPE Timetest tool. Without jitter control and with turbo mode enabled, time test plots show a band of 10-15 µsec latencies that are introduced due to processor frequency changes. After running the same workload with auto-tuned mode, those latencies disappear.

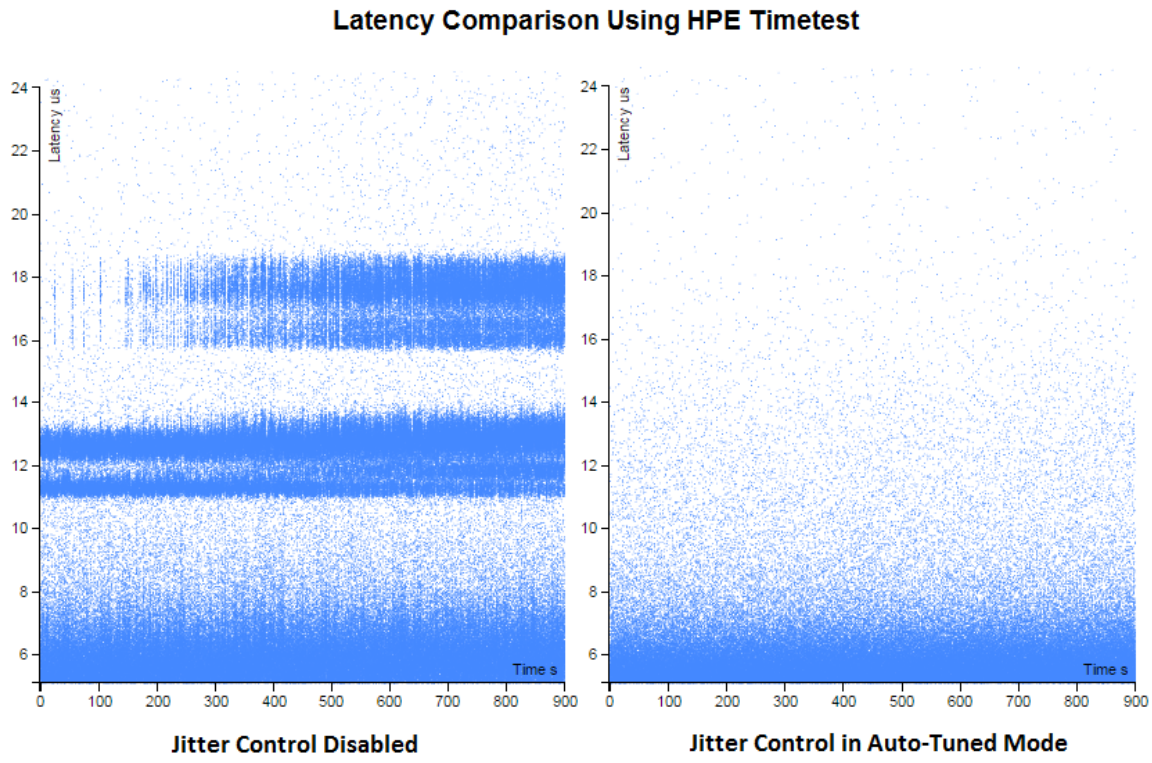

Figure 5: Latency comparison with Processor Jitter Control

Performance monitoring

Performance monitoring for HPE iLO 5 adds the ability to monitor key system metrics that measure the available system performance utilization by analyzing the workload running on the server. The metrics, provided by low level firmware, measure and analyze various components. Examples of the measured components include the processor and its interconnects to other resources, such as memory and I/O. iLO 5 offers a GUI to track these metrics over time as well as support for scripting through the iLO RESTful API.

There are important differences between the performance monitoring support provided by iLO 5 and the performance monitoring support that is typically provided by OS-based applications. CPU utilization is a good example of where obvious differences between reporting can be seen. Under the operating system, utilization is often measured by how often the processor is doing work on a time scale basis. The analysis is based on whether an application is running or not, and it often does not account for how fast the processor is executing. Additionally, cores and their threads are considered equal from a processing standpoint as well, based on whether an application is using them. In essence, while the CPU might be deemed busy, it might not be running at maximum compute capacity. CPU utilization as reported by iLO 5 is measured by the processor view of how much available execution capacity exists. So two different workloads that report the same utilization in the OS might have different utilization values as reported in iLO 5. This difference is not only important to understand when comparing the two sources, but also to understanding why iLO 5 tracks utilization differently.

Performance metrics reported by iLO 5 are useful when characterizing the workload usage of server resources and the amount of performance they can provide. iLO 5 performance monitoring, when paired with the workload advisor functionality, can help characterize the workload behavior and provide recommendations on tuning certain BIOS settings. The recommendations are based on how these resources are used by the workload that is running at the time.

More information

Workload advisor

Performance monitoring in iLO 5

The **Performance - Monitoring** page displays performance data collected from the following sensors on servers with Innovation Engine support.

CPU Utilization

This sensor reports the utilization of all processors installed in the system. The measurement is based on a percentage of the maximum compute capacity of the processor. It considers how slow or fast the processor runs when doing work. This measurement might differ from the values that some operating systems report for utilization, which is often calculated by how often the processor is not idle.

Memory Bus Utilization

This sensor reports the utilization of the total bandwidth of the memory bus. The measurement is based on a percentage of the maximum memory bandwidth of the configuration. This measurement might differ from the values that some operating systems report for memory utilization, which is often calculated by how much of the available system memory is being used or allocated.

I/O Bus Utilization

This sensor reports the utilization of all processors connected to I/O buses (total PCI-e bus bandwidth). The measurement is based on a percentage of the maximum total bandwidth of these buses. This measurement is not an indication of how busy an I/O device might be, but rather how much PCI-e bandwidth the device is using.

CPU Interconnect Utilization

This sensor reports the calculated bandwidth usage of the link connecting multiple processor sockets in the system. It is an aggregate of all the links within the system.

Jitter Count

This sensor reports the rate of processor frequency changes or "jitter" that occurs on a per second basis.

Average CPU Frequency

This sensor reports the average overall processor frequency. A value of zero means that the processor is idle. This value is different from the "running frequency" often seen under some operating systems that measures frequency only when the processor is not idle.

CPU Power

This sensor reports the power consumed by the processor. It is based on an energy accumulator within the processor and is the value that the processor uses to regulate power limits internally.

The information on this page might differ from the **Total CPU** power data on the **Power Meter** page, which is obtained without using the Innovation Engine.

Performance data details

The **Performance Data** section shows the following details:

Sensor

The name of the selected sensor.

Maximum

The maximum measured value.

Minimum

The minimum measured value.

Viewing performance data

Performance monitoring graphs display recent data collected from the Innovation Engine firmware.

When the server is powered off or in POST, a message is displayed, and measured performance values display the value **0**. Performance data is updated when the server is powered on and POST is complete. After a reset, there might be **0** values in the graph where data was not collected when the server was off or in POST. To confirm that these values are due to a server reset, check the IML.

When iLO is reset:

- Performance data for the **10 min** and **1 hr** intervals is cleared.
- Data from the **24 hr** and **1 week** graphs is saved, and can be viewed after a reset is complete.
- Hourly data might be missing when you view the **24 hr** and **1 week** graphs after a reset is complete.

Prerequisites

- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.

  If a license is not installed, a message is displayed, and only the 10 minute graph is available.
- The server has Innovation Engine support and version 2.0.11 or later of the Innovation Engine firmware is installed.
- MCTP Discovery is enabled.
- The **iLO Date/Time** is set correctly to ensure collection of valid performance telemetry samples.

Procedure

1. Click **Performance** in the navigation tree, and then click the **Monitoring** tab.
2. Select a sensor in the **Selected Sensor** menu.
3. Click one of the following options to select a graph interval:
   - **10 min**
   - **1 hr**
   - **24 hr**
   - **1 week**

   The graph is populated with data for the requested interval.
4. (Optional) To view data for a specific point on the graph, move the slider ○ beneath the graph to the point you want to view.

   When you move the slider, details for the selected point on the graph are displayed next to the graph.
5. (Optional) If you selected **CPU Power** or **Average CPU Frequency**, select or clear the check boxes in the CPU list next to the graph.

   Select a CPU check box to display it in the graph. Clear a CPU check box to remove it from the graph.
6. (Optional) Choose how to refresh data on this page.

   By default, the page data is not refreshed after you open the page.
   - To refresh the page immediately, click ↻.
   - To start refreshing the page automatically, click ▷. Depending on the selected graph type, the page refreshes at ten-second or five minute intervals. The page refreshes until you click □ or navigate to another page.

Performance monitoring graph display options

Selected Sensor menu

To view performance data for a sensor, select the sensor in the **Selected Sensor** menu.

Graph Type

To specify the graph time period, click a graph type name:

- **10 min**—Displays performance data for the last 10 minutes. The iLO firmware collects performance data for this graph every 20 seconds. The maximum number of samples displayed in the graph is 30.
- **1 hr**—Displays the performance data for the last hour. The iLO firmware collects performance data for this graph every 20 seconds. The maximum number of samples displayed in the graph is 180.
- **24 hr**—Displays the performance data for the last 24 hours. The iLO firmware collects performance data for this graph every 5 minutes. The maximum number of samples displayed in the graph is 288.
- **1 week**—Displays the performance data for the last week. The iLO firmware collects performance data for this graph every 30 minutes. The maximum number of samples displayed in the graph is 336.

Refreshing performance graphs

- To refresh the page immediately, click ↻.
- To start refreshing the page automatically, click ▷.

  The page refreshes automatically until you click ☐ or navigate to another page.

Viewing a specific data point on the graph

- To view data for a specific point on the graph, move the slider ○ beneath the graph to the point you want to view.

  You can also use the following methods to move the slider:

  - Click the slider track.
  - Click the slider icon, and then press the arrow keys on the keyboard.

  When you move the slider, details for the selected point on the graph are displayed next to the graph.
- When automatic refresh is running, move the slider ○ beneath the graph to focus on a data point that falls under a specific historical point along the x-axis.

Configuring performance alerts

You can configure performance alerts that will post an event in the IML when a configured threshold is reached.

Upper and lower thresholds are supported for the CPU Utilization, Memory Bus Utilization, and I/O Bus Utilization sensors.

Upper thresholds are supported for the CPU Interconnect Utilization, CPU Power, and Jitter Count sensors.

Prerequisites

- Configure iLO Settings privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.
- The server has Innovation Engine support and version 2.0.11 or later of the Innovation Engine firmware is installed.
- MCTP discovery is enabled.
- The **iLO Date/Time** is set correctly to ensure collection of valid performance telemetry samples.

Procedure

1. Click **Performance** in the navigation tree, and then click the **Monitoring** tab.
2. Select a sensor that supports performance alerts.
3. Enter the threshold settings and dwell time, and then click **Apply**.

   To disable an alert, set the dwell time to 0.

Performance alert settings options

Lower Threshold

The lowest value the sensor can report before an event is posted in the IML.

Enter a percentage of utilization.

Upper Threshold

The highest value the sensor can report before an event is posted in the IML.

- For utilization sensors, enter a percentage of utilization for the selected sensor.
- For CPU Power, enter a value in watts.
- For Jitter Count, enter the threshold count.

Dwell Time

The number of seconds the sensor reading is higher or lower than the configured value before the threshold is violated. When a threshold is violated, an event is posted in the IML.

For example, if you set an upper threshold to 70% with a dwell time of 40 seconds, an event is posted when the sensor reports readings over 70% for more than 40 seconds.

- To enable an alert, set the dwell time to a valid value in multiples of 20, between 20 and 64800 (20 seconds to 18 hours). If you enter a value that is not a multiple of 20, the value is rounded up to the next multiple of 20.
- To disable an alert, set the dwell time to 0.

Workload advisor

Understanding how to map server resource utilization to BIOS tuning settings for better workload performance is often not intuitive and can be cost prohibitive. The workload advisor feature monitors key server performance metrics and provides users with real-time tuning advice for select BIOS tuning settings. Tuning advice is based on actual server resources that are being used while running your workload. This functionality builds off of the workload matching feature and allows users to customize tuning settings ever further.

The main foundation for workload advisor is the ability in HPE iLO 5 to monitor key system performance metrics and analyze the results over an extended period. Whether it is minutes, hours, or a day, workload advisor uses observed system behavior to determine whether certain BIOS settings are optimally set.

Workload advisor takes a multistage approach to analyzing the system behavior while your workload is running:

1. Workload advisor leverages the performance monitoring support metrics to map your workload into a set of usage or behavior characteristics reports. These characteristics, which can be as simple as *high memory utilization*, are the beginning of understanding where resource usage bottlenecks might exist.

   For information about the performance monitoring support metrics, see **Performance monitoring**.

2. Workload advisor uses its analysis to recommend BIOS changes that might provide better results.

   Users are often pointed to documentation to learn how different workloads might react to a certain BIOS configuration. This approach assumes that the user is familiar with their workload characteristics.

   Workload advisor educates users about the characteristics of their workload, and then it recommends a configuration based on the workload characteristics.

Example: Benefits of using Workload advisor

Workload advisor uses real-time analysis of server feedback to determine how much the workload is NUMA (or non-uniform memory access) aware or optimized. Traditionally, a user would need to understand how an application was written to know if it can take advantage of the latest NUMA optimizations that processor vendors introduce in every generation.

Most documentation for NUMA-based settings provides simple recommendations based on presumed first-hand knowledge of how an application is written. This understanding, for the most part, is an unreasonable expectation. Users are not likely to know how much their workload is NUMA aware or optimized. The workload advisor and performance monitoring features perform this analysis by monitoring, analyzing, and providing tuning recommendations that might lead to performance improvements.

Workload advisor in iLO 5

iLO monitors selected server workload characteristics and provides recommended performance tuning settings based on the monitored data.

Server workload details

Workload characteristics are qualitative assessments of how the workload is using system resources. They are based on the quantitative measurements from the performance monitoring events and are useful as a reference when making tuning decisions. These observed characteristics are typically needed for making intelligent tuning decisions. For instance, a specific BIOS option might provide benefits only if the workload has a high degree of NUMA awareness.

The following workload characteristics are displayed:

- **CPU Utilization**—How busy the processors are in the server.
- **Memory Bus Utilization**—The amount of memory traffic observed by the server.
- **I/O Bus Utilization**— The amount of I/O traffic observed by the server.
- **NUMA Awareness**—How the workload is distributing memory and I/O accesses across multiple processors. A high degree of NUMA awareness means that I/O and memory traffic are directed more to local resources versus remote resources.

The possible values are **High**, **Medium**, and **Low**.

Server workload data for the **10 min** and **1 hr** intervals is cleared when iLO is reset.

Viewing server workload details

Prerequisites

- Host BIOS privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.
- The server is powered on and POST is complete.

  Make sure that the server was powered on for the time interval you want to monitor. For example, data for the 24-hour interval is not available until the server has been powered on for 24 hours.
- The server has Innovation Engine support and version 2.0.11 or later of the Innovation Engine firmware is installed.
- MCTP discovery is enabled.
- The **iLO Date/Time** is set correctly to ensure collection of valid performance telemetry samples.

Procedure

1. Click **Performance** in the navigation tree, and then click the **Workload Advisor** tab.
2. Review the details in the **Server Workload** section.

   If iLO was reset, information for the **10 min** and **1 hr** intervals will be available after the server has been powered on for 10 minutes or 1 hour.
3. (Optional) To update the table with the latest information, click ↻.

Configuring the performance tuning options

Prerequisites

- Host BIOS privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.
- The server is powered on and POST is complete.

  Make sure that the server was powered on for the time interval you want to monitor. For example, data and recommendations for the **24 hr** interval are not available until the server has been powered on for 24 hours.
- The server has Innovation Engine support and version 2.0.11 or later of the Innovation Engine firmware is installed.

- MCTP discovery is enabled.
- The **iLO Date/Time** is set correctly to ensure collection of valid performance telemetry samples.

Procedure

1. Click **Performance** in the navigation tree, and then click the **Workload Advisor** tab.
2. Select a value in the **Selected Duration** menu.

   You can review recommended settings based on data collected in **10 min**, **1 hr**, or **24 hr** intervals.
3. Review the recommendations in the **Recommended Setting** column.

   If iLO was reset, information for the **10 min** and **1 hr** intervals will be available after the server has been powered on for 10 minutes or 1 hour.
4. To change one or more settings, click **Settings**.
5. Change the tuning options as needed, and then click **Apply**.

   iLO notifies you that changing the tuning options will change the **Workload Profile** setting to **Custom**.
6. Click **Yes, apply**.

   iLO saves the settings and notifies you that a server reboot is required for the changes to take effect.
7. Reboot the server.

   You can click the link in the status banner to navigate to the **Server Power** page.

Performance tuning settings

Sub-NUMA Clustering

When this option is set to **Enabled**, this feature divides the processor cores, cache, and memory into multiple NUMA domains. Enabling this feature can increase performance for workloads that are NUMA-aware and optimized.

When this feature is enabled, up to 1 GB of system memory might become unavailable.

NUMA Group Size Optimization

This option sets how the system BIOS reports the size of a NUMA node (number of logical processors), which assists the OS in grouping processors for application use (Kgroups). The default value **Clustered** provides better performance because it optimizes the resulting groups along NUMA boundaries. Some applications might not be optimized to take advantage of processors that span multiple groups. In such cases, it might be necessary to select the **Flat** option to allow affected applications to use more logical processors.

Uncore Frequency Scaling

This option controls the frequency scaling of the internal processor buses (the uncore). Setting this option to **Auto** enables the processor to dynamically change frequencies based on workload. Setting the **Maximum** or **Minimum** frequency enables tuning for latency or power consumption.

Memory Refresh Rate

This option controls the refresh rate of the memory controller. It might affect the performance and resiliency of the server memory. Hewlett Packard Enterprise recommends using the default value (**1x Refresh**) unless changing this value is recommended in other documentation for the server.

Power Regulator

Use this option to configure Power Regulator support. The following values are available:

- **Dynamic Power Savings Mode**—Automatically varies processor speed and power usage based on processor utilization. This option allows the reduction of overall power consumption with little or no impact to performance. It does not require OS support.
- **Static Low Power Mode**—Reduces processor speed and power usage. This option guarantees a lower maximum power usage value for the system. Performance impacts are greater for environments with higher processor utilization.
- **Static High Performance Mode**—Processors will run at maximum power and performance at all times, regardless of the OS power management policy.
- **OS Control Mode**—Processors will run at maximum power and performance at all times, unless the OS enables a power management policy.

NOTE: The Power Regulator setting displayed on the **Workload Performance Advisor** page reflects the static boot time configuration. It does not reflect run-time changes to this setting that have been applied since system power-on. Applying recommended settings changes on the **Workload Performance Advisor** page changes only the boot time configuration of this setting. A system reboot is required for the change to take effect.

Minimum Processor Idle Power Package C-state

Use this option to select the lowest idle power state (C-state) of the processor that the operating system uses. The higher the C-state, the lower the power usage of that idle state. **C6 State** is the lowest power idle state supported by the processor.

Energy/Performance Bias

Use this option to configure several processor subsystems to optimize the performance and power usage of the processor. The following values are available:

- **Maximum Performance**—This setting is for environments that require the highest performance and lowest latency, but are not sensitive to power consumption.
- **Balanced Performance**—This setting provides optimum power efficiency. Hewlett Packard Enterprise recommends this setting for most environments.
- **Balanced Power**—Provides optimum power efficiency based on server utilization.
- **Power Savings Mode**—This setting is suitable for environments that are power sensitive and can accept reduced performance.

Core boosting

Core boosting technology uses a relaxed and optimized turbo profile that adapts the processor to specific use cases, configurations, and environments. Core boosting processors take advantage of extra server power and thermal headroom provided by an innovative HPE voltage regulator design and by cooling technologies. Consequently, systems that have core boosting processors can alleviate common setbacks and maximize processor computing power.

For example, a processor can have a number of cores and a base frequency at which the cores operate. Processors may also have a turbo mode that operates processor cores at a faster frequency than the base frequency. The turbo mode may use thermal and power capacity headroom opportunistically to operate processor cores at an increased frequency. Turbo mode can increase processor performance while maintaining the same TDP (Thermal Design Power) level.

Some processors are preconfigured with a TDP and a maximum power level. To maintain these parameters safely, the predefined settings are typically fused and locked into the processor. These settings ensure that the processor operates within its standard electrical, thermal, and power design specifications. A turbo profile for a processor is bounded by these constraints using fixed frequency registers and core-to-frequency ratio registers. The power limits can be fused to the TDP level. To maintain the specified TDP level, the CPU turbo frequencies are dictated by the number of active cores that were fused in core-to-frequency ratio register. The turbo frequency profile scales from all cores active to a single core active. Accordingly, the turbo frequency increases as the number of active cores being utilized by lower workload demands or by core parking or disabling technologies.

These fused frequency registers and core-to-frequency ratio registers cap the processor computing capacity at certain levels. On a general-purpose computing processor, however, you can set the turbo profile more conservatively to cover various workloads or a worst-case thermal condition. In other words, the turbo mode might have a one-size-fits-all profile that does not consider the specific configuration or environment in which the processor operates. Accordingly, the processor might not be tuned to use its full operating potential.

Processor core boosting architecture

The Processor core boosting feature uses the platform firmware to provide enhanced Turbo mode performance for select HPE ProLiant Gen10 servers. The platform firmware automatically detects the presence of supported processors and enables core boosting by default. You can also configure the core boosting functionality through the UEFI System Utilities, iLO RESTful API, or Always On Intelligent Provisioning through iLO 5.

Core boosting uses select Intel Xeon Scalable processors with an optimized turbo profile to address key Enterprise use cases, such as high performance computing, virtualization, and big data. Due to high computing power consumption from core boosting processors, high performance heatsinks and fans are required. The required hardware is included during the CTO process when you order servers with processors that support core boosting technology.

Table 1: Supported Intel Xeon Scalable Processors with core boosting

	Intel Xeon Scalable 6143 Gold	Intel Xeon Scalable 8165 Platinum
Launch Date	Q3 2017	Q4 2017
Cores	16	24
Threads	32	48
Base Frequency	2.80 GHz	2.30 GHz
Max Turbo Frequency	4.0 GHz	3.7 GHz

Table Continued

	Intel Xeon Scalable 6143 Gold	Intel Xeon Scalable 8165 Platinum
Core Boosting Max Turbo Profile	1-2 Cores (4.0 GHz) 3-16 Cores (3.9 GHz)	1-2 Cores (3.7 GHz) 3-4 Cores (3.5 GHz) 5-24 Cores (3.4 GHz)
Cache	22 MB L3	33 MB L3
UPI Links	3	3
TDP	205W	205W
Max Memory	768 GB	768 GB
Memory Type	DDR4-266	DDR4-266
Max Memory Speed	2666 MHz	2666 MHz
Intel Turbo Boost Technology	2.0	2.0
Intel Hyper-Threading	Yes	Yes

Table 2: Core-Boosting SKUs shows the core boosting SKUs available on the following Gen10 servers:

Table 2: Core-Boosting SKUs

Servers	Intel Xeon Scalable 6143 Gold (16C)	Intel Xeon Scalable 8165 Platinum (24C)
HPE ProLiant DL380 Gen10	✓	✓
HPE ProLiant DL360 Gen10	✓	✓
HPE ProLiant DL580 Gen10	✓	✓
HPE ProLiant DL560 Gen10	✓	✓
HPE Apollo XL230k Gen10	✓	✓
HPE Synergy SY660 Gen10	✓	✓
HPE Synergy SY480 Gen10	✓	✓

The core boosting turbo profile

The core boosting processor has relaxed the frequency and core-to-frequency ratio registers to optimize processor performance for integer and SSE/AVX-1 operations. Using the core boosting settings, the systems can opportunistically run workloads at higher clock frequencies.

Figure 6: Active Cores with core boosting illustrates the benefits of core boosting in MIPS (Millions of instructions per second) when compared to an Intel Xeon stock 16-core processor. The image shows that core boosting enables workloads running on more active cores to do so at a higher frequency compared to turbo, which improves system performance without overclocking, keeping the Intel warranty intact. The reliability specification is also maintained the same as without core boosting enabled. Core boosting also benefits single-threaded applications because its single-core active clock frequency is higher than without core boosting, as shown in the image.

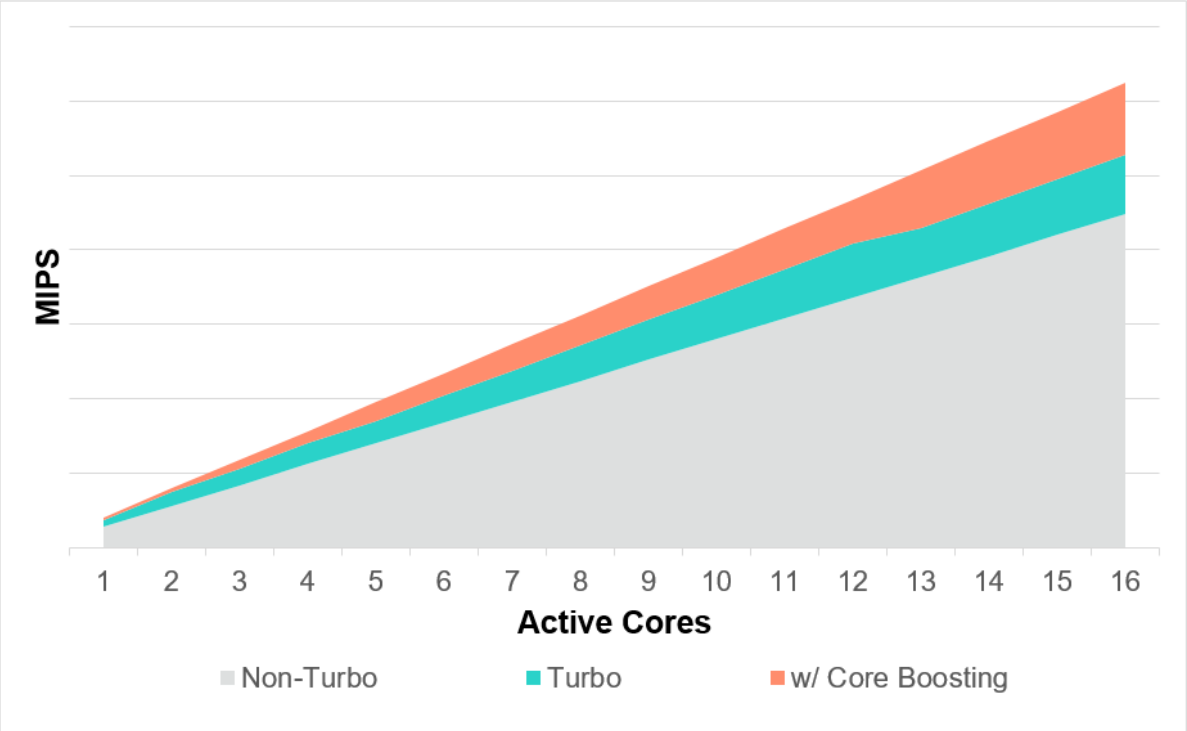

Figure 6: Active Cores with core boosting

For Intel Xeon v4 or later processors, AVX-2 and AVX-3 technologies were implemented with different turbo profiles than the profiles used for integer or AVX-1/SSE instructions. AVX-2 instructions are 256 bits and AVX-3 are 512 bits, extension of AVX-2. AVX-2 and AVX-3 instructions help accelerate performance but consume more power than non-AVX workloads. However, to maintain processor TDP, their base and turbo clock frequency are noticeably lower. See specifications on the Intel website for more information. For core boosting processors, AVX-2 and AVX-3 turbo profiles are maintained in a similar way to comparable Intel Xeon stock parts. However, with higher TDP, core boosting processor SKUs still yield higher performance than comparable Intel Xeon stock processors.

Increasing performance with core boosting

Core boosting enables higher performance across more processor cores. It can accelerate system performance by running workloads on more active cores at higher clock frequencies than comparable Intel Xeon stock processors. Furthermore, core boosting can effectively lower licensing costs for core and processor-based license models, yielding better total cost of ownership when compared to higher core-count Intel Xeon stock processors.

The following table shows SPEC CPU2006 benchmark results among three processors for various configurations: Intel stock 6142 (16 Core) and 6148 (20 Core), and core boosting 6143 (16 Core) SKUs. The results for SPECint_base2006 and SPECfp_base2006 clearly show core boosting advantages on single-threaded applications over Intel Xeon stock processors. On system throughput measurements, the SPECint_rate_base2006 and SPECfp_base2006 numbers show the core boosting performance to be superior over the comparable Intel Xeon processor. The competitiveness of core boosting performance compared with higher core-count processors is also apparent in the same benchmark results.

Table 3: SPEC CPU2006 benchmark results

DL380 Gen10	6143 vs. 6142 HT Disabled +sse4.2	6143 vs. 6148 HT Disabled +sse4.2
SPECint_base2006 (single threaded)	6.91%	5.69%
SPECfp_base2006 (single threaded)	5.97%	6.77%
SPECint_rate_base2006 (multi-threaded)	13.97%	-2.58%
SPECfp_rate_base2006 (multi-threaded)	9.48%	-1.57%

Configuring core boosting by using iLO 5 and Always On Intelligent Provisioning

Prerequisites

- Configure iLO Settings privilege
- Host BIOS privilege
- Remote Console privilege
- A license that supports this feature is installed. For information about the available license types and the features they support, see the licensing documentation at the following website: **https://www.hpe.com/support/ilo-docs**.
- The server is configured with core boosting processors.
- Version 1.20 or later of the System ROM (BIOS) is installed.
- The server has Innovation Engine support, and version 1.2.4 or later of the Innovation Engine firmware is installed.
- MCTP discovery is enabled.
- The server is powered on and POST is complete.
- The latest version of Intelligent Provisioning is installed.

Procedure

1. Click **Performance** in the navigation tree.

   iLO displays the **Settings** tab.
2. Click **Settings**.
3. To start Intelligent Provisioning, click **Configure with Always On**.

   The Intelligent Provisioning web interface starts in a new browser window.
4. Click **Perform Maintenance**.
5. Click **BIOS/Platform Configuration**.

   The **BIOS/Platform Configuration** page opens.
6. Click **Power and Performance Options**.

7. Click **Advanced Performance Tuning Options**.
8. Select a **Core Boosting** option.

   The change is pending.
9. Click **BIOS/Platform Configuration** to return to the overview page.
10. Click **Update**.

    Intelligent Provisioning applies the changes and notifies you that a server reboot is required for the changes to take effect.
11. Reboot the server.

Core boosting options

Enabled

Enabling this feature allows the server to use the enhanced performance capabilities of processors that support core boosting.

Enabled is the default value when the system detects that core boosting processors are installed.

Disabled

Disabling this feature causes the processor to use a restricted turbo frequency profile and lower maximum wattage capabilities.

Core boosting Integrated Management Log and error messages

HPE ProLiant Gen10 servers with processor SKUs that support core boosting generate Integrated Management Log (IML) events when an iLO Advanced license is not installed and when the core boosting feature is disabled.

The platform firmware issues an error message if an iLO license is not installed. You can resolve this issue by installing an iLO license.

The platform firmware issues an informational error message when the core boosting feature is disabled. This message is a warning that optimal processor performance is restricted. You can resolve this error by re-enabling the core boosting option in the UEFI System Utilities, the iLO RESTful API, or Always On Intelligent Provisioning through iLO 5.

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

For additional websites, see **Support and other resources**.

Product websites

iLO

https://www.hpe.com/info/ilo

iLO 5 Information Library

http://www.hpe.com/support/ilo-docs

iLO RESTful API and RESTful Interface Tool

https://www.hpe.com/support/restfulinterface/docs

HPE ProLiant Gen10 servers

https://www.hpe.com/info/proliantgen10-docs

HPE ProLiant Gen10 Plus servers

https://www.hpe.com/info/proliantgen10plus-docs

UEFI System Utilities

https://www.hpe.com/info/ProLiantUEFI/docs

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  https://www.hpe.com/info/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  https://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center

  https://www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  https://www.hpe.com/support/downloads

  Software Depot

  https://www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  https://www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

  https://www.hpe.com/support/AccessToSupportMaterials

IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

https://www.hpe.com/services/getconnected

HPE Proactive Care services

https://www.hpe.com/services/proactivecare

HPE Datacenter Care services

https://www.hpe.com/services/datacentercare

HPE Proactive Care service: Supported products list

https://www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

https://www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

https://www.hpe.com/services/proactivecarecentral

Proactive Care service activation

https://www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

https://www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

https://www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

https://www.hpe.com/support/Storage-Warranties

HPE Networking Products

https://www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

https://www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

https://www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

https://www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.