



Hewlett Packard
Enterprise

UEFI Workload-based Performance and Tuning Guide for HPE ProLiant Gen10 Servers and HPE Synergy

Abstract

This guide details the use of Workload Profiles in performance tuning. System provided profiles are listed and described, profile dependencies listed for each profile, and tasks for applying and changing options are provided.

Part Number: 881335-001
Published: July 2017
Edition: 1

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

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

UEFI® is a registered trademark of the UEFI Forum, Inc.

Contents

- Workload Profiles..... 4**
 - Workload Profiles and performance options.....4
 - Workload Profiles dependencies 6
 - Applying a Workload Profile 9
 - Changing dependent options after applying a profile..... 10

- Websites..... 11**

- Support and other resources..... 12**
 - Accessing Hewlett Packard Enterprise Support..... 12
 - Accessing updates..... 12
 - Customer self repair..... 12
 - Remote support..... 13
 - Warranty information..... 13
 - Regulatory information..... 14
 - Documentation feedback..... 14

Workload Profiles

Workload Profiles and performance options

Workload Profiles are a configuration option to deploy BIOS settings based on the workload customer intends to run on the server. Workload Profiles are a configuration option to deploy BIOS settings to accommodate the intended application of the server.

This document lists and describes the provided Workload Profiles and lists the corresponding settings for each profile. Use this information to determine which profile is best suited for your deployment.

System provided Workload Profiles

The system provides these Workload Profiles:

General Power Efficient Compute

This profile is the default profile for most ProLiant servers and HPE Synergy compute modules.

This profile applies the most common performance settings that benefit most application workloads while also enabling power management settings that have minimal impact to overall performance. The settings that are applied heavily favor a balanced approach between general application performances versus power efficiency.

This profile is recommended for customers that do not typically tune their BIOS for their workload.

General Peak Frequency Compute

This profile is intended for workloads that generally benefit from processors or memory that must achieve the maximum frequency possible, for any individual core, at any time. Power management settings are applied when they ensure that any component frequency upside can be readily achieved. Processing speed is favored over any latencies that might occur. This profile is a general-purpose profile, so optimizations are done generically to increase processor core and memory speed.

This profile benefits workloads that typically benefit from faster compute time.

General Throughput Compute

This profile is intended to be used for workloads where the total maximum sustained workload throughput is needed. Increased throughput does not always occur when the processor runs at the highest individual core speed. Increased throughput can occur when the processor is able to perform sustained work across all available cores during maximum utilization. Power management settings are disabled when they are known to have impact on maximum achievable bandwidth.

Best throughput is achieved when the workload is also (Non-uniformed Memory Access) NUMA aware and optimized so settings that benefit NUMA awareness are applied.

Virtualization - Power Efficient

This profile is intended to be used for virtualization environments. The profile ensures that all available virtualization options are enabled. Certain virtualization technologies can have possible performance impacts to nonvirtualized environments and can be disabled in other profiles. Power management settings can have an impact on performance when running virtualization operating systems and this profile applies power management settings that are virtualization friendly.

Virtualization - Max Performance

This profile is intended to be used for virtualization environments. The profile ensures that all available virtualization options are enabled. Power management settings are disabled in favor of delivering maximum performance.

Low Latency

This profile is intended to be used by customers who desire the least amount of computational latency for their workloads. This profile follows the most common best practices that are documented in the HPE Low Latency Whitepaper. Maximum speed and throughput are often sacrificed to lower overall computational latency. Power management and other management features that might introduce computational latency are also disabled.

The profile benefits customers running Real-Time Operating Systems (RTOS) or other transactional latency sensitive workloads.

Mission Critical

This profile is intended to be used by customers who trade off performance for server reliability above the basic server defaults. The profile enables advanced memory reliability, availability, and serviceability (RAS) features that are known to have more than a measurable impact to computational performance. Enabling this profile will have an impact to maximum memory bandwidth and will increase memory latency.

Transactional Application Processing

This profile is intended to be used for business processing environments, such as online transaction processing (OLTP) applications that require a database back-end. For example, workloads typically comprised of a high number of user-based, transactional applications running on a single server with cohosted database component. The profile balances the requirement of managing both peak frequency and throughput.

High Performance Compute (HPC)

This profile is intended for customers running in a traditional HPC environment. Typically, these environments are clustered environments where each node performs at maximum utilization for extended periods of time to solve large-scale scientific and engineering workloads. The default for our Apollo series servers, power management is typically disabled in favor of sustained available bandwidth and processor compute capacity. This profile is similar to the Low Latency profile except that some latency is accepted to achieve maximum throughput.

Decision Support

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

Graphic Processing

This profile is intended for workloads that are run on server configurations which utilize Graphics Processing Units (GPUs.) GPUs typically depend on maximum bandwidth between I/O and Memory. Power management features that have impact on the links between I/O and memory are disabled. Peer to Peer traffic is also critical and therefore virtualization is also disabled.

I/O Throughput

This profile is intended to be used for configurations that depend on maximum throughput between I/O and memory. Processor utilization driven power management features that have performance impact to the links between I/O and memory are disabled.

Custom

This option on the Workload Profiles menu disables Workload Profiles. Use this option if you want to set specific BIOS options for your deployment manually. When you select Custom, all the settings for the previously selected profile are carried forward. You can edit all or some of the options.

Custom is not a profile and settings that you specify are not saved as a template.

Default profiles for servers

Workload Profile options support a variety of power and performance requirements. For most HPE ProLiant Gen10 servers and HPE Synergy compute modules, Workload Profile is set to **General Power Efficient Compute** by default. This Workload Profile provides common performance and power settings suitable for most application workloads. For ProLiant XL servers in an HPE Apollo system, the Workload Profile is set to **HighPerformance Compute** by default.

Selecting a Workload Profile other than the Custom profile affects other setting options. For example, selecting the **General Peak Frequency Compute** profile automatically sets **Power Regulator** mode to **Static High Performance**. This setting cannot be changed and is grayed out.

Workload Profiles dependencies

Dependencies

There are multiple options that are available for BIOS configuration. Not all profiles set the same options to specific settings. Each profile is designed to obtain specific performance results and sets different options to meet those results. The options that a profile sets are called dependencies. All other options are unaffected by the Workload Profile and are referred to as nondependent settings.

Dependencies and switching profiles

When you change a profile, only the dependent settings for that profile are changed. Nondependent settings remain what they were before you changed your profile.

For example:

1. Select the General Power Efficient Compute profile, which has the Energy Performance Bias set to Balanced Performance.
2. Select the General Peak Frequency Compute profile, which has no dependency on Energy Performance Bias. The Energy Performance option is set to Balanced Performance because that setting is carried forward from the General Power Efficient Compute profile.
3. Select the General Throughput Compute profile, which has the Energy Performance Bias set to Maximum Performance.
4. Select the General Peak Frequency Compute profile which has no dependency on Energy Performance Bias. Energy Performance Bias is set to Maximum Performance because that setting is carried forward from the General Throughput Compute profile.

There is no way to revert to a previous profile and dependencies. Once you change to a new profile, the new dependencies are applied. The only way to revert to older profiles, is to exit without saving your changes. Exiting without saving reverts to where you were when you entered RBSU. Once you save a profile, you cannot revert from that profile to any intermediate dependencies.

Dependencies and options matrix

The following tables show the Workload Profiles and their dependencies. The Workload Profiles are listed in the order that they are listed on the user interface. In the table, "X" means that the option setting has no requirement for the profile and can be edited. Dependencies cannot be edited.

NOTE:

Not all the options listed are adjustable on all servers. However, even if you do not have the option of adjusting some of these settings, they default to the values shown here.

Table 1: Workload Profiles General Power Efficient Compute — Low Latency

	General Power Efficient Compute	General Peak Frequency Compute	General Throughput Compute	Virtualization - Power Efficient	Virtualization - Max Performance	Low Latency
SR-IOV	X	X	X	Enabled	Enabled	Disabled
VT-D	X	X	X	Enabled	Enabled	Disabled
VT-x	X	X	X	Enabled	Enabled	Disabled
Power Regulator	Dynamic Power Savings	Static High Performance	Static High Performance	OS Control	Static High Performance	Static High Performance
Minimum Processor Idle Power Core C-state	C6	X	X	C6	No C-states	No C-states
Minimum Processor Idle Power Package C-state	Package C6 Retention	Package C6 Retention	Package C6 Retention	Package C6 Retention	No C-states	No C-states
Energy Performance Bias	Balanced Performance	X	Max Performance	Balanced Performance	Max Performance	Max Performance
Collaborative Power Control	Enabled	Disabled	Disabled	Enabled	Disabled	Disabled
Intel DMI Link Frequency	Auto	Auto	Auto	Auto	Auto	Auto
Intel Turbo Boost Technology	Enabled	Enabled	Enabled	X	Enabled	Disabled
Intel NIC DMA Channels (IOAT)	Enabled	X	X	X	X	X
HW Prefetcher	Enabled	Enabled	Enabled	X	X	Enabled
Adjacent Sector Prefetch	Enabled	Enabled	Enabled	X	X	Enabled
DCU Stream Prefetcher	Enabled	Enabled	Enabled	X	X	Enabled
DCU IP Prefetcher	Enabled	Enabled	Enabled	X	X	Enabled
NUMA Group Size Optimization	Flat	Clustered	Clustered	Clustered	Clustered	Clustered

Table Continued

Memory Patrol Scrubbing	X	X	X	X	X	Disabled
Memory Refresh Rate	X	1X	1X	X	X	1X
UPI Link Power Management	Enabled	Disabled	Disabled	Enabled	Disabled	Disabled
Sub-NUMA Clustering	Disabled	X	Enabled	Disable	Enabled	X
Energy-Efficient Turbo	Enabled	Disabled	Disabled	Enabled	Disabled	Disabled
Uncore Frequency Shifting	Auto	Max	X	Auto	Max	Max
x2APIC	X	X	X	X	X	Disabled
Channel Interleaving	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Memory Bus Frequency	X	X	X	X	X	X
A3DC	X	X	X	X	X	Disabled

Table 2: Workload Profiles Mission Critical — I/O Throughput

	Mission Critical	Transactional Application Processing	High Performance Compute (HPC)	Decision Support	Graphic Processing	I/O Throughput
SR-IOV	X	X	Disabled	X	Disabled	X
VT-D	X	X	Disabled	X	Disabled	X
VT-x	X	X	Disabled	X	Disabled	X
Power Regulator	X	Static High Performance	Static High Performance	X	X	X
Minimum Processor Idle Power Core C-state	X	No C-states	No C-states	X	X	X
Minimum Processor Idle Power Package C-state	X	No C-states	No C-states	X	X	X
Energy Performance Bias	X	Max Performance	Max Performance	X	Max Performance	Max Performance

Table Continued

Collaborative Power Control	X	X	Disabled	X	X	X
Intel DMI Link Frequency	Auto	Auto	Auto	Auto	Auto	Auto
Intel Turbo Boost Technology	X	Enabled	Enabled	X	X	X
Intel NIC DMA Channels (IOAT)	X	Enabled	Enabled	X	X	Enabled
HW Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Adjacent Sector Prefetch	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
DCU Stream Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
DCU IP Prefetcher	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
NUMA Group Size Optimization	X	Clustered	Clustered	Clustered	Clustered	Clustered
Memory Refresh Rate	2X	X	1X	X	X	X
UPI Link Power Management	X	Disabled	Disabled	X	X	X
Sub-NUMA Clustering	X	X	X		X	X
Energy-Efficient Turbo	X	X	Disabled	X	X	X
Uncore Frequency Shifting	X	X	Max	X	Max	Max
x2APIC	X	X	Disabled	X	Disabled	X
Channel Interleaving	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Memory Bus Frequency	X	X	X	X	X	X
A3DC	Enabled	X	Disabled	X	X	X

Applying a Workload Profile

You apply a Workload Profile to have the system manage your workload according to predefined settings provided with the system. Dependent options cannot be changed and are grayed out. You can change any nondependent options in a profile.

Procedure

1. From the **System Utilities** screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Workload Profile**.
2. Select a **Workload Profile**.
3. Optional: Change any nondependent options that you want to change.
4. Save and reboot to apply your Workload Profile.

Changing dependent options after applying a profile

There may be one or more dependent options that you want to change in your Workload Profile. Dependent options cannot be changed for a predefined profile. You can change the dependent options in Custom mode. When you are in Custom mode, your deployment is no longer in profile mode and you can manually adjust option settings. When you enter Custom mode, all the settings from the previously applied profile are shown.

The easiest way to change dependent settings is to modify an applied profile. First apply a Workload Profile that has most of the settings that you want to use then change to Custom mode. Then change only the settings you want to have new values.

Prerequisites

Apply a Workload Profile before you do this task.

Procedure

1. From the **System Utilities** screen, select **System Configuration > BIOS/Platform Configuration (RBSU) > Workload Profile**.
2. Select the Custom profile option.
All of the settings from the previously applied Workload Profile are shown. All options are editable.
3. Change the options that you want to have new values.
4. Save and reboot to apply the changes.

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

www.hpe.com/storage/spock

Storage white papers and analyst reports

www.hpe.com/storage/whitepapers

UEFI Specification

www.uefi.org/specifications

UEFI Learning Resources

www.uefi.org/learning_center

Contact Hewlett Packard Enterprise Worldwide

<http://www.hpe.com/assistancecenter>

Subscription Service/Support Alerts

<http://www.hpe.com/support/e-updates>

Software Depot

<http://www.hpe.com/support/softwaredepot>

Customer Self Repair

<http://www.hpe.com/support/selfrepair>

Insight Remote Support

<http://www.hpe.com/info/insightremotesupport/docs>

For additional websites, see [Support and other resources](#).

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<http://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<http://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

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

Software Depot

www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:
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:
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.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience.

Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

<http://www.hpe.com/support/selfrepair>

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

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

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

Additional warranty information

HPE ProLiant and x86 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

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:

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:

www.hpe.com/info/reach

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

www.hpe.com/info/ecodata

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

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.