

The Lenovo logo, consisting of the word "Lenovo" in white sans-serif font on a black rectangular background.

# Power Management with Lenovo Efficiency Mode

---

**Introduces Lenovo Efficiency Mode and how it operates in conjunction with the operating system**

---

**Explains the tradeoffs of standard operating system power plans**

---

**Describes the tests performed to quantify the benefit of Lenovo Efficiency Mode**

---

**Provides results showing the merits of using Lenovo Efficiency Mode**

**Robert Wolford**  
**Jenseng Chen**



# Abstract

As focus is increased on datacenter efficiency, understanding and investing in energy management has never been more important. Lenovo® has developed a unique feature called *Lenovo Efficiency Mode* in the System x® M5 series of servers. This new UEFI setting is designed to operate autonomously. It optimizes the processor frequency based on the dynamic workload changes and obtains the best performance/power ratio.

This paper is intended for use by datacenter administrators, end users, and technical sales personal who would like to learn about the benefits of Lenovo Efficiency Mode and how it operates.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

<http://lenovopress.com>

**Do you have the latest version?** We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

# Contents

Introduction .....	3
Methods of controlling CPU P-states .....	3
Comparing P-state control under Windows .....	5
Comparing P-state control under Linux .....	8
Conclusions .....	10
Change history .....	11
Authors .....	11
Notices .....	12
Trademarks .....	13

# Introduction

Energy and power utility costs continue to rise and have become a significant portion of the total data center expense. Energy efficiency enables data centers to use less energy while providing the same level of performance, and therefore is a critical priority.

In a typical data center, the cost for power and cooling a server for three years is 1.5 to 2 times the cost of purchasing the server hardware itself<sup>1</sup>. At the same time, IT managers are challenged to increase the workload capacity of data centers while maintaining or reducing the power consumption. This can only be achieved by improving the overall efficiency of the data center.

With data centers focusing on energy efficiency, understanding and investing in energy management has never been more important. Lenovo has developed a unique feature called *Lenovo Efficiency Mode* in the System x M5 series of servers. Lenovo Efficiency Mode is designed to optimize the processor frequency based on the dynamic workload and obtain the best performance/power ratio.

Lenovo Efficiency Mode is a setting in UEFI and is available on the following systems, with either Xeon E5-2600 v3 or E5-2600 v4 processors:

- ▶ System x3650 M5
- ▶ System x3550 M5
- ▶ System x3500 M5
- ▶ NeXtScale nx360 M5
- ▶ Flex System x240 M5

Standard Performance Evaluation Corporation (SPEC) has developed the SPECpower\_ssj2008 benchmark suite (referred to hereafter as SPECpower) to measure the power and performance characteristics of a server and report the overall performance/watt efficiency metric of the server. The benchmark can be used to compare power and performance among different servers and serves as an industry standard toolset for measuring server efficiency.

The SPECpower benchmark is designed for use by hardware vendors, IT industry, computer manufacturers, and governments. Also, because it is an industry-standard benchmark, the results from SPECpower are more credible compared to custom efficiency benchmarks. In the following sections, SPECpower will be used to compare the performance/watt efficiency in different scenarios.

With respect to performance/watt efficiency, one of the key factors that affects efficiency is CPU frequency. The CPUs in a server typically consume the largest portion of power used by the server. As the CPU frequency is increased, the CPU power generally increases exponentially which leads to higher system power. Therefore, monitoring the CPU frequency is critical in analyzing the overall efficiency of the server. However, the SPECpower benchmark suite does not monitor or report CPU frequency. To collect the CPU frequency while running the scenarios in this paper, Intel's Power Thermal Utility was used.

## Methods of controlling CPU P-states

Modern CPUs are capable of changing their operating frequency on-the-fly -- operating modes known as *P-states*. Running a CPU at a lower frequency (a higher-numbered P-state)

<sup>1</sup> United States Data Center Energy Usage Report, Ernest Orlando Lawrence, Berkeley National Laboratory, Table 4, <http://eta.lbl.gov/publications/united-states-data-center-energy-usag>

when the workload is light saves power. However, running a CPU at too low of a frequency can reduce the performance below an acceptable level.

Lenovo UEFI supports multiple P-state control mechanisms. The P-state setting is on the Processors page in UEFI as shown in Figure 1.

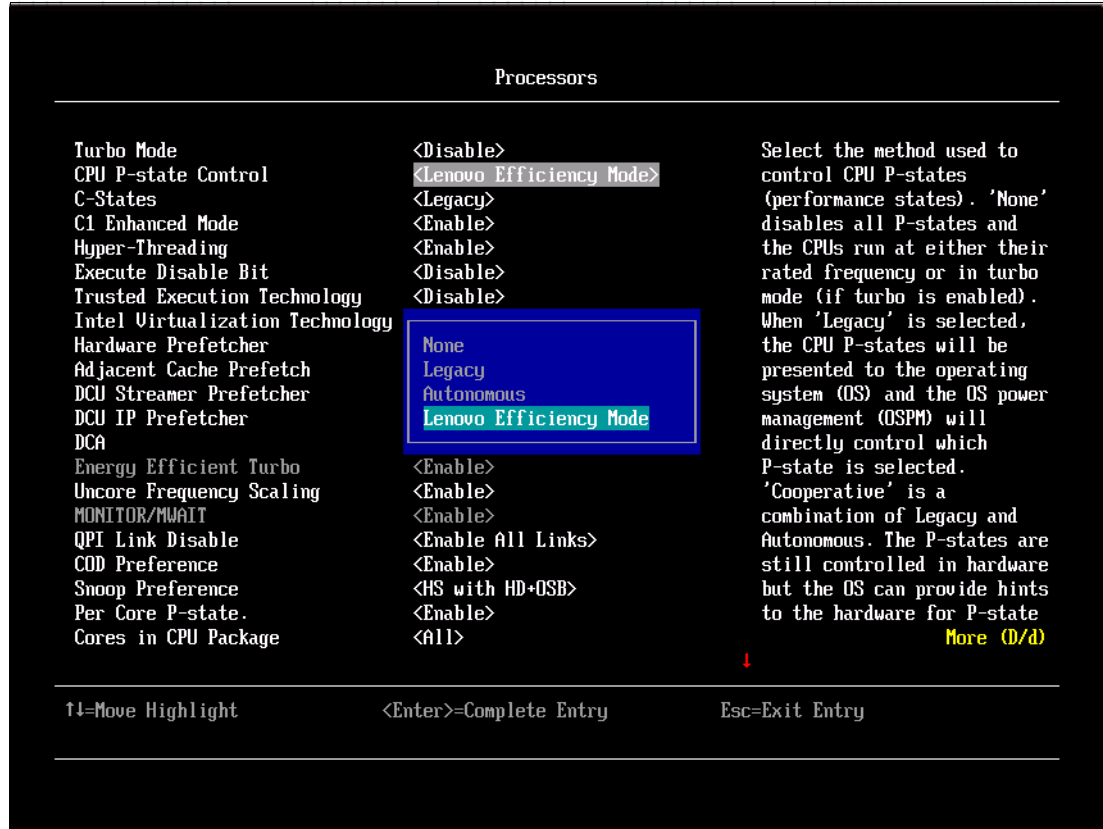


Figure 1 UEFI setting for P-state control

The choices are as follows:

► **None:**

The None selection disables all P-states and the CPUs run at either their rated frequency or in Turbo mode (if Turbo is supported by the processor and is enabled in UEFI). The trade-off in this mode is that CPUs will run at their top frequency regardless of the level of workload. This results in more CPU power being consumed, but there is a benefit to performance.

► **Legacy Mode:**

When Legacy is selected, the CPU P-states will be presented to the operating system (OS) via the ACPI table. The OS power management (OSPM) function will directly control which P-state is selected.

► **Autonomous Mode:**

For Autonomous mode, the server ignores all OS requests for P-states and decides on its own what the best P-state is to use.

► **Lenovo Efficiency Mode:**

Lenovo Efficiency Mode is similar to Legacy mode in that the OSPM chooses the desired P-state. But, unlike the Legacy option, the OSPM's P-state choice only serves as a suggestion when Lenovo Efficiency Mode is enabled. Lenovo Efficiency Mode then uses

the OSPM's P-state suggestion along with the CPU utilization level to make a final determination of the appropriate P-state that will achieve maximum efficiency. Depending on the CPU utilization, the actual P-state could be different from the P-state requested by the OSPM.

This paper compares Legacy Mode and Lenovo Efficiency Mode for energy efficiency. Based on the above description, it is clear that disabling P-states entirely (the None option) will have a negative effect on energy efficiency. Autonomous has proved to have lower energy efficiency than Legacy Mode, and therefore is not covered in this paper. Instead, this paper will compare the P-state Legacy option vs. the Lenovo Efficiency Mode option to demonstrate which is more efficient under either a Windows or Linux environment.

## Comparing P-state control under Windows

For our tests under Windows we used the following server configuration:

- ▶ Operating system: Windows Server 2012 R2
- ▶ Java virtual machine: Oracle Java7 update 80
- ▶ Server: Lenovo System x3650 M5
- ▶ Processors: 2x Intel Xeon E5-2699 v4 (2.2 GHz, Turbo frequency up to 3.6 GHz)
- ▶ Memory: 8x 8 GB 2Rx8 DDR4 DIMMs
- ▶ Local storage: 1x 480GB SSD

In order to compare Legacy vs. Lenovo Efficiency Mode P-state control, four SPECpower benchmark runs were completed and the results recorded under Windows, with and without Turbo Mode, as listed in Table 1.

*Table 1 Windows tests performed*

	<b>Turbo Mode disabled</b>	<b>Turbo Mode enabled</b>
Legacy Mode - P-state control	Run 1	Run 3
Lenovo Efficiency Mode - P-state control	Run 2	Run 4

When the CPU frequency increases, the power increases exponentially, and this is particularly true when Turbo Mode is engaged. Turbo Mode allows for extreme performance at the expense of increased power and reduced efficiency. To fully understand the impact of Turbo Mode, the legacy vs. Lenovo Efficiency Mode comparison was done both with Turbo Mode disabled and enabled.

After all four benchmark runs were completed, they were normalized to the lowest score and compared as shown in Figure 2 on page 6.




Figure 2 Impact on energy efficiency of the use of Lenovo Efficiency Mode

As shown in Figure 2, Lenovo Efficiency Mode improves average efficiency under Windows by approximately 2.5%. To further narrow down exactly where the efficiency boost is occurring, the SPECpower results need to be analyzed in greater detail. This is easy to do because the SPECpower benchmark stresses the server at 11 different utilization levels ranging from 100 to 0% and records the performance and power at each utilization level.

When the SPECpower benchmark runs, it drives the system to a predetermined performance level for each of the 11 utilization levels measured. As shown in Figure 3, the relative performance difference for both runs with Turbo Mode disabled is less than 1%.

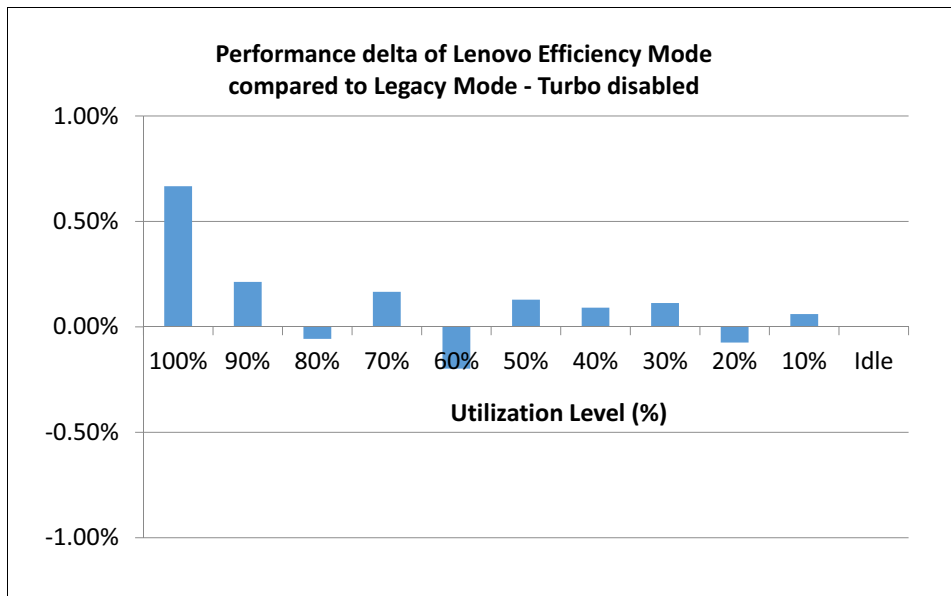


Figure 3 Performance delta - Turbo Mode disabled

The same is true when Turbo Mode is enabled as shown in Figure 4 on page 7 - the relative performance delta is less than 1%.




Figure 4 Performance delta - Turbo Mode enabled

Because the performance is virtually identical between the two non-Turbo runs shown in Figure 3 on page 6 and between the two Turbo runs shown in Figure 4, it can be concluded that the difference in the performance/watt efficiency between Legacy Mode and Lenovo Efficiency Mode P-state control must be due to differences in power consumption.

Figure 5 shows the power delta between Legacy Mode and Lenovo Efficiency Mode P-state control for each of the 11 utilization levels run by SPECpower.

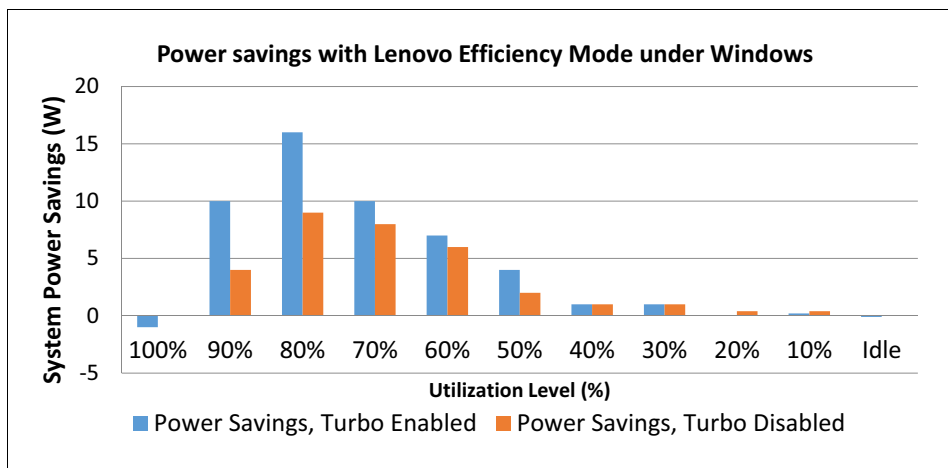


Figure 5 Power savings with Lenovo Efficiency Mode under Windows

As shown in Figure 5, Lenovo Efficiency Mode boosts power savings either with Turbo enabled or disabled. With Turbo disabled, there is a maximum 9 watts power savings. And with Turbo enabled, the maximum power savings increases to 16W.

Figure 6 on page 8 shows the frequency that the processor is driven at for each SPECpower utilization level. For a given load, the CPU frequency is lower when the system is in Lenovo Efficiency Mode than when the system is in Legacy Mode.

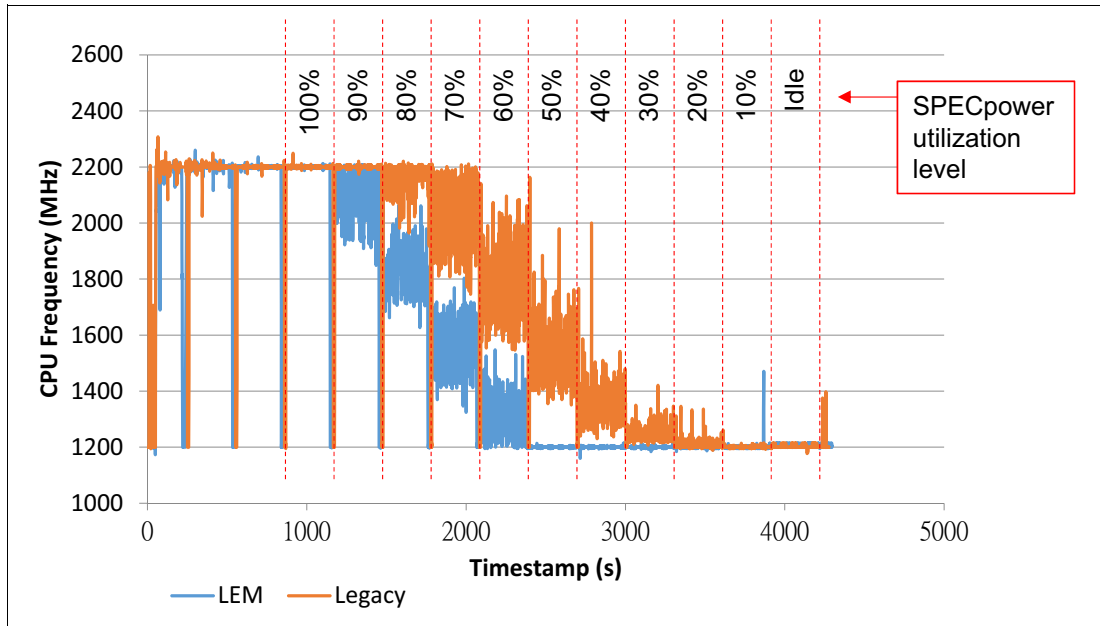


Figure 6 CPU frequency at each SPECpower utilization level (Turbo disabled)

With the analysis shown in Figure 6, Turbo Mode was disabled. Doing the same analysis with Turbo enabled results in a similar behavior.

## Comparing P-state control under Linux

For our tests under Linux we used the following server configuration:

- ▶ Operating system: Red Hat Enterprise Linux 7 Update 2
- ▶ Java virtual machine: Oracle Java7 update 80
- ▶ Server: Lenovo System x3650 M5
- ▶ Processors: 2x Intel Xeon E5-2699 v4 (2.2 GHz, Turbo frequency up to 3.6 GHz)
- ▶ Memory: 8x 8 GB 2Rx8 DDR4 DIMMs
- ▶ Local storage: 1x 480GB SSD

Under Linux, one of the most effective ways to reduce power consumption and heat output on a server is to use the cpufreq driver. Cpufreq can be enabled by adding the following line to `/boot/grub2/grub.cfg`:

```
intel_pstate=disable
```

There are four preconfigured governors in the operating system's acpi\_cpufreq driver:

- ▶ Performance
- ▶ Powersave
- ▶ Conservative
- ▶ Ondemand

In addition, end users can also choose the Userspace governor which allows more customization. The end user can select any of the governors for controlling the CPU frequency. A comparison of the governors are shown in Table 2 on page 9.



Table 2 A comparison of the different governors in the Linux acpi\_cpufreq driver

Governor selection	Efficiency impact
Performance	Not good for efficiency because it forces the CPUs to stay at the highest frequency regardless of the loadings. Overall, power consumption will be higher.
Powersave	Not good for efficiency either , because it forces the CPUs to stay at the lowest frequency no matter how much workload is present. Overall, performance will be lower.
Ondemand	Scales the CPU frequency between minimum and maximum. The final operating frequency is based on the workload. Efficiency will benefit.
Conservative	Operates the same as the ondemand governor but switches between P-states more gradually. This can have a detrimental effect on system latency.
Userspace	Requires a lot of intensive workload tuning to determine the optimal CPU frequency that achieves the highest efficiency. It's an iterative process and is not automatic. Most data center administrators don't have the time to manually optimize the userspace governor.

Figure 7 illustrates how each governor controls the CPU frequency during a SPECpower run while using legacy P-state control. It can be seen that the operation of each matches the descriptions above.

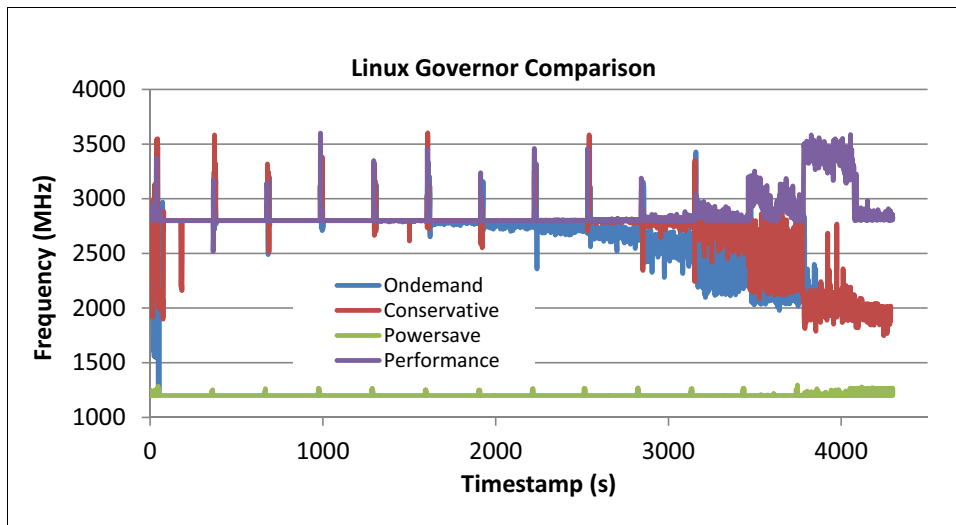


Figure 7 Linux governor comparison

Figure 7 shows a comparison of the CPU frequency control for each of the Linux governors. Based on our experience and several SPECpower runs, we know that the Ondemand governor produces the highest efficiency when Lenovo Efficiency Mode is disabled. Therefore, Ondemand was used as the comparison point when Lenovo Efficiency Mode was enabled. Lenovo Efficiency Mode complements the Linux governor selected and it will be shown that enabling Lenovo Efficiency Mode boosts the overall system efficiency when used along with the Ondemand governor.

As was done under Windows (see Figure 2 on page 6), overall efficiency comparisons were done with Turbo Mode both disabled and enabled. However, as shown in Figure 8 on page 10, the advantage of Lenovo Efficiency Mode is significantly more pronounced under Linux. When Lenovo Efficiency Mode is enabled there is a 16-19% efficiency advantage, compared to Legacy Mode P-state control.

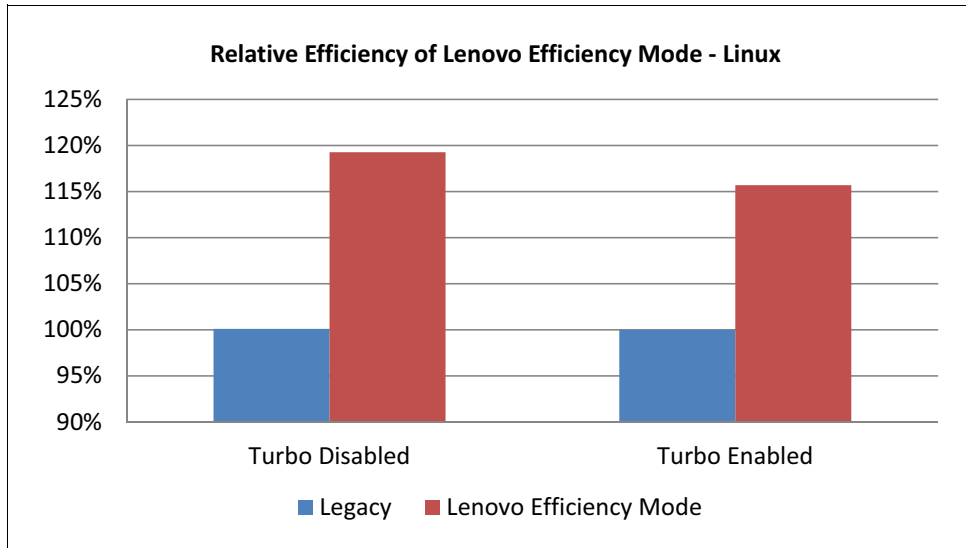


Figure 8 Impact on energy efficiency of the use of Lenovo Efficiency Mode

Drilling down into the details, the power savings across the full utilization curve can be studied. From Figure 9, it can be seen that with Lenovo Efficiency Mode, there is a maximum of 74W savings when Turbo is enabled and 60W when Turbo is disabled.




Figure 9 Power savings with Lenovo Efficiency Mode under Linux

## Conclusions

It has been shown how Lenovo Efficiency Mode optimizes the CPU frequency in real time to improve the performance/watt efficiency on Lenovo servers. In essence, Lenovo Efficiency Mode prevents the CPU from running faster than necessary to complete a job. It is analogous to selecting the optimal transmission gear in a vehicle based on the speed of the vehicle. Selecting the optimal gear and engine speed boosts efficiency and saves fuel. Similarly, selecting the optimized CPU frequency saves power and improves efficiency on Lenovo servers.

Whether running Windows or Linux, Lenovo Efficiency Mode can improve a server's performance/watt efficiency with one change in a UEFI setting. Unlike other efficiency tuning methods, Lenovo Efficiency Mode does not require extensive tuning or intricate knowledge of the server environment.

Data centers can realize significant savings by enabling Lenovo Efficiency Mode. Under Windows, efficiency is boosted by an average of 3% and a maximum power savings of 16W is realized. Under Linux, efficiency is boosted by an average of 19% and a maximum power savings of 74W is realized. The power savings leads to real dollar savings at the data center level.

## Change history

December 7, 2016:

- ▶ Grammar and style corrections

## Authors

This paper was produced by the following specialists:

**Robert R. Wolford** is a Senior Engineering Staff Member for power management and efficiency at Lenovo. He covers all technical aspects that are associated with power metering, management, and efficiency. Previously, he was a lead systems engineer for workstation products, video subsystems building block owner for System x, System p, and desktops, signal quality and timing analysis engineer, gate array designer, and product engineering. In addition, he was a Technical Sales Specialist for System x. Bob has 11 issued patents and three pending patent applications. He holds a Bachelor of Science degree in Electrical Engineering with Distinction from the Pennsylvania State University.

**Jenseng Chen** is an Advisory Software Engineer in the Lenovo Server Performance lab. He joined Lenovo in October 2014, after having worked at IBM for 8 years. He started his career in IBM as a operating system engineer, having worked on Linux development and support. Since then, he has been working on Server Performance Validation in energy efficiency. He holds three patents in server performance and system energy efficiency areas. Jenseng holds a Master degree in Computer Information Science from National Chiao Tung University in Taiwan.

Thanks to the following people for their contributions to this project:

- ▶ David Watts, Lenovo Press
- ▶ Mark T. Chapman, Lenovo Marketing Enablement

# Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.  
1009 Think Place - Building One  
Morrisville, NC 27560  
U.S.A.  
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on December 7, 2016.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/1p0548>

## Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

Lenovo(logo)®

System x®

The following terms are trademarks of other companies:

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, Windows Server, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.