



IBM EnergyScale for Power10 Processor-Based Systems

October 2021

Martha Broyles

Christopher J. Cain

Chris Francois

Stuart Jacobs

Todd Rosedahl

Vaidyanathan Srinivasan

Brian Veale

Executive Overview	3
EnergyScale Features	4
Performance Boost.....	4
Energy-Optimized Fan Control.....	4
Processor Stop States	4
Processor Folding.....	4
Processor Folding in AIX.....	4
Processor Folding in IBM i.....	5
Processor Folding in Linux	6
Safe Mode	7
Reasons for Temporary Safe Mode.....	7
Permanent Safe Mode	7
System Power and Performance Mode	8
Query System Power and Performance Mode in AIX	8
Static Mode	8
Power Saving Mode.....	8
Maximum Performance Mode (MPM)	10
EnergyScale for I/O.....	10
Frequency and Performance Measurement.....	10
AIX.....	10
IBM i.....	11
ASMI.....	12
Setting System Power and Performance Mode.....	12
HMC.....	13
Setting System Power and Performance Mode.....	13
Appendix I: System Requirements	14
Release Level FW1010	15
Feature Support	15
Frequency.....	16
Appendix II: Processor Usage and Accounting	17
Appendix III: Resources	19

Executive Overview

The energy required to power and cool computers can be a significant cost to a business – reducing profit margins and consuming resources. In addition, the cost of creating power and cooling infrastructure can be prohibitive to business growth. In response to these challenges, IBM developed EnergyScale™ Technology for IBM® Power®. EnergyScale provides functions that help the user to understand and control IBM server power and cooling usage. This enables better facility planning, provides energy and cost savings, enables peak energy usage control, and increases system availability. Administrators may leverage EnergyScale capabilities to control the power consumption and performance of Power processor-based systems to meet their particular data center needs.

In this paper, the functions provided by EnergyScale are described along with usage examples, and hardware and software requirements. Support and awareness of EnergyScale extends throughout the system software stack, and is included in the AIX, IBM i, and Linux operating systems. This paper focuses on features and functions found in systems based on Power10 processors. For previous-generation IBM Power, please refer to companion paper, “IBM EnergyScale for Power9 Processor-based Systems”.

EnergyScale Features

EnergyScale provides features to control the power consumption, energy efficiency, and performance of Power10 Systems.

Performance Boost

EnergyScale provides additional performance. The system is designed to provide enough power and cooling for the processors to run most workloads at a specified base frequency. However, most workloads do not consume all available power at the base frequency. In these cases, the frequency can be increased above that base frequency to take advantage of the additional power and thermal headroom in the system.

Energy-Optimized Fan Control

Cooling fans contribute significantly to the overall power consumption of a given computer. In order to minimize energy expended on cooling and to minimize the energy wasted “over-cooling” a system, firmware on all Power10 systems will adjust fan speed in response to real-time temperatures of the system components.

Processor Stop States

The IBM Power processors uses a combination of clock and power gating to achieve different levels of power saving and associated latency to enter and exit these levels known as stop states. For more information see

<https://community.ibm.com/community/user/power/blogs/amit-tendolkar1/2020/09/20/an-overview-of-idle-states-in-the-power9-processor>

Processor Folding

While Processor Core Nap and sleep provide energy savings and performance boost when processors become idle, additional savings/boost can be realized if processors remain idle by intent rather than by happenstance. Processor Folding is a consolidation technique that dynamically adjusts, over the short-term, the number of processors available for dispatch to match the number of processors demanded by the workload. As the workload increases, the number of processors made available increases; as the workload decreases, the number of processors made available decreases. Processor Folding increases energy savings/boost during periods of low to moderate workload because unavailable processors remain in low-power idle states longer than they otherwise would. Processor Folding achieves power savings similar to those that could be achieved by intelligent, utilization-based logical partition (LPAR) configuration changes, but it does so with much greater efficiency and fidelity, and without impacting the configuration or processor utilization of the LPAR.

Processor Folding in AIX

In AIX the processor folding policy can be configured from the command line via the `schedo` command. The `vpm_fold_policy` tunable is a 4-bit value where each bit indicates the configuration of a different setting. The following table shows the various settings that are controlled.

<i>Bit</i>	<i>Setting</i>
0	=1 processor folding is enabled when the partition is using shared processors
1	=1 processor folding is enabled when the partition is using dedicated processors
2	=1 disables the automatic setting of processor folding when the partition is in Power Saving mode
3	=1 processor affinity will be ignored when making folding decisions

Table 1: `vpm_fold_policy` is a 4-bit value, in which each bit controls an aspect of folding.

The following command displays the current setting of `vpm_fold_policy`:

```
# schedo -L vpm_fold_policy
NAME                CUR    DEF    BOOT    MIN    MAX    UNIT  TYPE
DEPENDENCIES
-----
vpm_fold_policy      1      1      1       0     15     D
-----
```

To enable processor folding on a partition using dedicated partitions when the current value of `vpm_fold_policy` is set to 1, the following command would be issued to set the value to 3:

```
# schedo -o vpm_fold_policy=3
```

To disable processor folding, the value of `vpm_fold_policy` can be set to the value 4 using the following command:

```
# schedo -o vpm_fold_policy=4
```

By default, AIX will attempt to consider processor affinity (or topology) information when making processor folding decisions. This allows for the workload to remain spread across the processor nodes (e.g., chips depending on the system) and benefit from improved performance. Bit 3 of the `vpm_fold_policy` tunable allows this default behavior to be disabled. For example, if `vpm_fold_policy` is currently set to 6, indicating that processor folding is enabled when the partition is using dedicated partitions and that the partition will automatically enable processor folding when in Power Saving mode, the following command would change the setting to indicate that the operating system should no longer consider processor affinity when making folding decisions:

```
# schedo -o vpm_fold_policy=14
```

For more information see the help information via the `-h` option of the `schedo` command:

```
# schedo -h vpm_fold_policy
```

Processor Folding in IBM i

In an IBM i partition, processor folding is configured and controlled by the operating system by default. On Power10 servers, the operating system enables processor folding by default in shared processor LPARs or when Power Saving mode is enabled. Operating system control of processor folding may be overridden via the **QWCCTLSW** limited availability API, which provides a key-based control language programming interface to a limited set of IBM i tunable parameters. Processor folding control is accessed via **QWCCTLSW** key 1060. The following sequence of calls cycles

through the various options. Changes to key 1060 take effect immediately but do not persist across IPLs.

Get current status:

```
> CALL QWCCTLSW PARM('1060' '1')
KEY 1060 IS *SYSCTL.
KEY 1060 IS SUPPORTED ON THE CURRENT IPL.
KEY 1060 IS CURRENTLY ENABLED.
```

Explicitly disable processor folding:

```
> CALL QWCCTLSW PARM('1060' '3' )
KEY 1060 SET TO *OFF.
```

Explicitly enable processor folding:

```
> CALL QWCCTLSW PARM('1060' '2' 1)
KEY 1060 SET TO *ON.
```

Re-establish system control of processor folding:

```
> CALL QWCCTLSW PARM('1060' '2' 2)
KEY 1060 SET TO *SYSCTL.
```

Processor Folding in Linux

It is essential to install a daemon package based on the host OS to enable utilization-based processor folding for Power Saving:

```
pseries-energy-1.4.0-1.el7.ppc64.rpm
pseries-energy-1.4.0-1.el6.ppc64.rpm
pseries-energy-1.4.0-1.sles11.ppc64.rpm
```

Version 5.4 has the necessary user space tools required to enable CPU Folding.¹

Once this package is installed, the `energyd` daemon will monitor the system power mode and activate processor folding when system power mode is set to "Power Saving" and deactivate processor folding in all other modes. The utilization-based CPU folding daemon will deactivate unused cores and transition them to low power idle states until the CPU utilization increases and those cores are activated to run a workload.

Utilization-based processor folding can be manually disabled using the following commands:

```
/etc/init.d/energyd stop #Stop daemon now, activate all cores
chkconfig energyd off    #Do not restart daemon on startup
```

-or-

```
rpm -e pseries-energy    #un-install the package completely
```

¹Refer to <http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/installtools/home.html> for more details.

Alternatively, CPU cores can be folded or set to low power idle state in any power mode manually using the following command line:

```
echo 0 > /sys/devices/system/cpu/cpuN/online #Where N is the logical CPU number
```

Please note that all active hardware threads of a core need to be taken off-line using the above command in order to move the core to a low power idle state.

The cores can be activated again with the following command:

```
echo 1 > /sys/devices/system/cpu/cpuN/online #Where N is the logical CPU number
```

Safe Mode

All systems and firmware releases support safe mode. “Safe Mode” is a system mode where the firmware will automatically drop to a fixed lower frequency to keep the system thermal and power safe

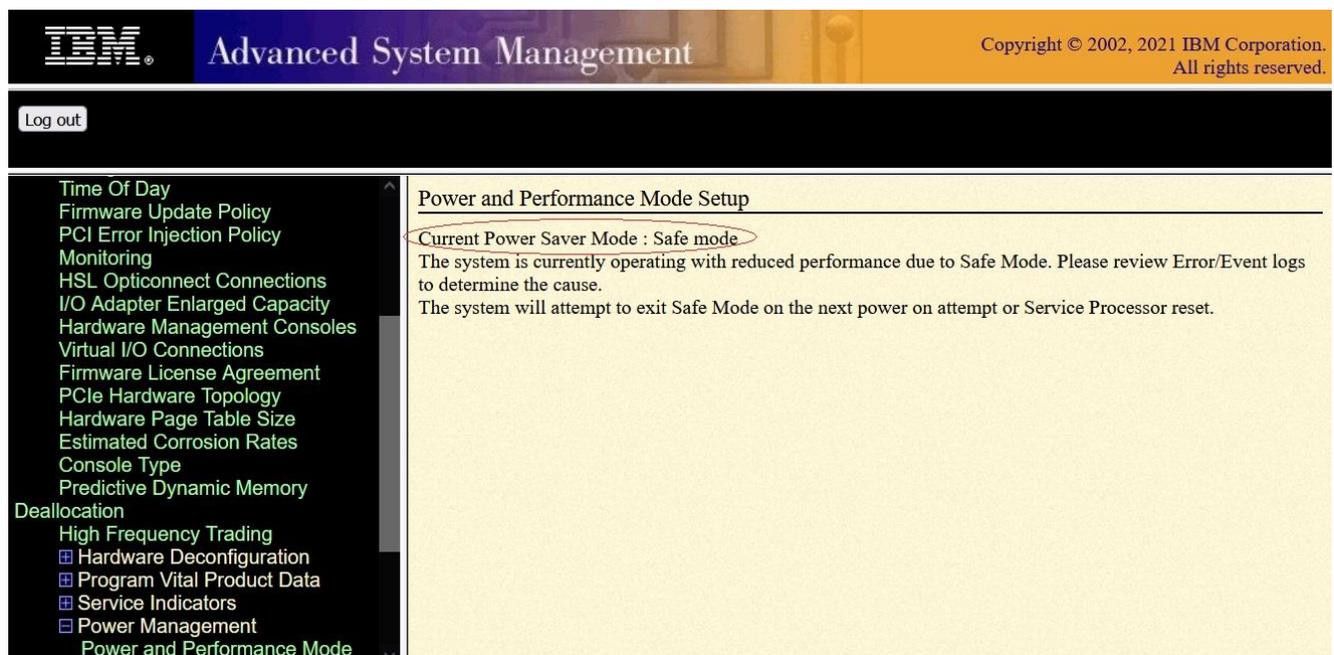
Reasons for Temporary Safe Mode

The following *may* cause the system to temporarily enter safe mode, no user action is required to exit safe mode from these conditions and no user notification is provided when it occurs.

- Concurrent code update
- Error recovery

Permanent Safe Mode

Certain hardware and firmware failures that cannot be recovered from can cause the system firmware to stay in safe mode. There will always be an error log generated when safe mode is entered due to a hard failure. In addition to a posted error log safe mode will also be indicated in the ASMI Power and Performance Mode Setup menu:



To re-enable full EnergyScale functionality, a reboot, firmware update, FRU replacement, or complete A/C power cycle of the system must be completed. If the problem which originally caused the system to enter safe mode is still present, the system will re-enter safe mode and generate an additional error log entry.

System Power and Performance Mode

The mode that a system is in can be queried from ASMI, HMC or the OS but it can only be changed from ASMI or the HMC.

<i>Power9</i>	<i>Power10</i>	<i>Notes</i>
Modes Disabled also referred to as "nominal"	Static	Re-name only. Functionally the same between Power9 and Power10
Static Power Saver	Power Saving	Re-name only. Functionally the same between Power9 and Power10
Dynamic Performance	Not supported	
Maximum Performance	Maximum Performance	No change. All Power10 systems will be shipped in Maximum Performance mode by default

The System Power and Performance mode setting persists across system boots, service processor resets and loss of ac power (unless the power outage is long enough to drain the Service Processor NVRAM battery).

Query System Power and Performance Mode in AIX

Starting with AIX 7.2 TL 3 SP 5 and AIX 7.1 TL 5 SP 9, lparstat -N displays information about Power and Performance Modes, Processor Frequency, and Processor Folding.

```
# lparstat -N
```

```
Power and Performance Mode      : Maximum Performance
Idle Power Saver Status        : Not Supported
Minimum Frequency (Mhz)       : 3251
Static Frequency (Mhz)        : 3550
Maximum Frequency (Mhz)       : 3900
Processor Folding Status       : Disabled
```

Static Mode

This is also known as "Nominal" mode. In this mode, the frequency is set to fixed value that can run all normal workloads under all normal environmental conditions. There is no frequency boost based on idle cores, lower workloads, nor favorable ambient conditions.

Power Saving Mode

Power Saving mode lowers the processor frequency and voltage to the minimum, reducing the power consumption of the system while still delivering predictable performance. In addition, this mode automatically enables processor folding in dedicated processor partitions. See [Appendix I](#) for details on the actual Power Saving frequency by system and firmware release.

Power Saving could be enabled based on regular variations in workloads, such as predictable dips in utilization overnight, or over weekends. Power Saving can be used to reduce peak energy consumption, which can lower the cost of all power used. Please note that when Power Saving is enabled for certain workloads with low CPU utilization, workload performance will not be impacted, though CPU utilization may increase due to the reduced processor frequency.

The only time that a system does not support operating at the Power Saving voltage and frequency is during a system boot or re-boot. Power Saving may be enabled at any time; however, if Power Saving was enabled prior to a system boot the voltage and frequency will remain at the default boot values until the platform firmware reaches a standby or running state. Immediately before the platform firmware starts executing on the system's processors, the voltage and frequency will drop to the Power Saving mode values. If a system re-boot occurs while in Power Saving mode, the voltage and frequency will return back to boot values and following a successful system re-boot the voltage and frequency will be dropped back to Power Saving mode.

Maximum Performance Mode (MPM)

Maximum Performance mode will allow the system to reach the maximum frequency by taking advantage of the thermal and power headroom provided by idle cores, lower workloads, and favorable environmental conditions.

EnergyScale for I/O

IBM Power automatically power off pluggable PCI adapter slots that are not being used to save approximately 14 watts per slot. A PCI adapter slot is considered not being used when the slot is empty, when the slot is not assigned to a partition, or when the partition to which the slot is assigned is not powered on. A PCI slot is powered off immediately by system firmware when it is dynamically removed from the partition to which it was assigned, and when the partition to which it is assigned is powered off. Furthermore, system firmware automatically scans all pluggable PCI slots at regular intervals looking for those that meet the criteria for being not in use and powers them off. This ensures among other things that slots left on after platform power-on are subsequently powered off if they are not in use. This is supported on all Power10 processor-based systems, and the expansion units that they support. Note that it applies to hot-pluggable PCI slots only. Power controls for other types of I/O features and built-in, or embedded, PCI adapters are not available and so they cannot be powered off independently from their enclosure power.

Frequency and Performance Measurement

Methods to measure the frequency and performance vary by operating system.

AIX

`lparstat -E` and `mpstat -E` are used to view the current processor frequency

- `lparstat -E` reports the frequency averaged across all of the Virtual Processors assigned to the LPAR
- `mpstat -E` reports the frequency per Virtual Processor

Usage: `lparstat/mpstat -E [Interval [Count]]`

```
# lparstat -E 1 1
System configuration: type=Dedicated mode=Capped smt=4 lcpu=4 mem=4096MB Power=Disabled
Physical Processor Utilization:
-----Actual-----          -----Normalised-----
user  sys  wait  idle  freq          user  sys  wait  idle
----  ---  ----  ----  -
0.001 0.002 0.000 0.997  3.0GHz [100%] 0.001 0.002 0.000 0.997
# mpstat -E 1 1
System configuration: lcpu=4 mode=Capped
vcpu    pbusy          physc          freq          scaled physc
----    -
0       0.0022 [0%]    1.0005 [100%] 3.0GHz [100%] 1.0048 [100%]
```

The *Interval* parameter must be supplied to `lparstat/mpstat -E` to read the current frequency.

Without the interval parameter, the commands generate a single report containing statistics since the last boot of the LPAR by default.

NOTE: For AIX version 7.2 and earlier, `pmcycles` command in AIX should NOT be used for reading the current processor frequency. The recommended approach is to use `lparstat -E` and `mpstat -E` as discussed above.

References:

https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.cmds3/lparstat.htm

https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.cmds3/mpstat.htm

IBM i

IBM iDoctor for IBM i

IBM iDoctor for IBM i displays the CPU rate for the IBM i partition over time on the Collection Overview graph. The CPU rate for the partition is the ratio of scaled to unscaled processor utilized time, expressed as a percentage. The processor utilized time is the accumulation of non-idle virtual processor SPURR and PURR² over each time interval. The ratio of SPURR and PURR accumulated over an interval represents the processor frequency versus nominal over the interval.

WRKSYSACT

The Work with System Activity (WRKSYSACT) command displays the Average CPU rate since last refresh for the partition in output shown on the display station. The Average CPU rate for the partition is the ratio of scaled to unscaled processor utilized time, expressed as a percentage. The processor utilized time is accumulation of non-idle virtual processor SPURR and PURR for the interval since the last refresh.

IBM i Collection Services

Database file QAPMJOBMI contains time series data by task, primary thread, and secondary thread. Scaled and unscaled CPU times, both charged and used, are available to calculate average CPU rate for processing activity of tasks and threads.

Database file QAPMSYSTEM contains time series system-wide (i.e. partition) accumulations of performance data. Scaled and unscaled CPU times are accumulated for various categories of processor usage. The ratio of scaled to unscaled time is the average CPU rate for the category of time accumulation. The processor utilized time is accumulation of non-idle virtual processor SPURR and PURR for the time interval.

Note: As of IBM i 7.3, the QAPMCONF database file key "NF" contains the processor nominal frequency in MHz. The processor nominal frequency can be used to convert average CPU rate to average processor frequency.

² See [Appendix II](#) for more info.

ASMI

Supported EnergyScale features can be found on the Advanced System Management Interface (ASMI) under the “System Configuration” → “Power Management” menu.

Setting System Power and Performance Mode

The screenshot displays the ASMI interface. At the top, the IBM logo and 'Advanced System Management' title are visible, along with a copyright notice: 'Copyright © 2002, 2021 IBM Corporation. All rights reserved.' A 'Log out' button is located in the top left corner of the main content area. On the left side, there is a navigation menu with several items. The 'Power Management' item is highlighted with a red oval. The main content area is titled 'Power and Performance Mode Setup'. It shows the current mode as 'Maximum Performance mode'. Below this, there are three radio button options: 'Static mode', 'Power Saving mode', and 'Maximum Performance mode'. The 'Maximum Performance mode' option is selected. A note below the options states: 'Note: Enabling any of the Power Saver modes will cause changes in the processor frequencies, changes in processor utilization, changes in power consumption, and performance to vary. Other effects are possible as well. Please see the EnergyScale™ white paper for more information on power saving modes.' At the bottom of the main content area, there is a 'Continue' button.

HMC

The HMC (Hardware Management Console) is a management console that controls managed systems, logical partitions, managed frames, other features provided through the managed objects, and the HMC itself. The HMC provides both graphical user interface (GUI) and command line interfaces. Users can use the user interfaces to configure or manage various features offered by the managed objects.

Setting System Power and Performance Mode

The user can list which power management modes are supported using the `lspwrmgmt` command on the command line:

```
lspwrmgmt -m <managed system name> -r sys -F supported_power_saver_mode_types
```

Example output:

```
"static,power_saving,max_perf"
```

A user can then enable one of the supported power management modes using the `chpwrmgmt` command on the command line as illustrated below enabling Maximum Performance mode:

```
chpwrmgmt -m <managed system name> -r sys -o enable -t max_perf
```

To query the current power management mode use `lspwrmgmt` command on the command line as illustrated below:

```
lspwrmgmt -m <managed system name> -r sys
```

Example output:

```
curr_power_saver_mode=Enabled,curr_power_saver_mode_type=max_perf,desired_power_saver_mode=Enabled,desired_power_saver_mode_type=max_perf,"supported_power_saver_mode_types=static,power_saving,max_perf",idle_power_saver_mode=unavailable
```

NOTE: The new mode may not take effect immediately. Normally, if the operation is performed before the system is powered on, the desired mode won't take effect until the system is up and running. If the mode is in transition, any changes will be blocked.

With the HMC GUI, users can reach this task by selecting the managed system -> Operations -> Power Management.

Appendix I: System Requirements

Due to differences in each release, this appendix details the systems and EnergyScale features supported by release including the actual frequency limits for various power management modes.

Release Level FW1010

Feature Support

Refer to the EnergyScale Features chapter earlier in this document for a definition of each feature.

	Power E1080
Power Saving Mode	ASM, HMC
Maximum Performance Mode	Enabled by Default ASM, HMC
Static Mode	ASM, HMC

Frequency

Depending upon the Power savings setting selected, the maximum and minimum frequency may change. For a definition of each setting, please refer to the EnergyScale Features chapter earlier in this document.

Support Notes

¹ Note that CPU frequencies greater than the system base are *not* guaranteed. The actual maximum frequency may vary based on environmental conditions, system configuration, firmware version, component tolerances, and workload.

² For frequency to be guaranteed to be in the normal operating frequency range the ambient must be below this temperature when in Maximum Performance mode

	<i>Guaranteed Ambient Temperature²</i>	<i>Maximum Performance Mode Typical Operating Frequency Range¹</i>	<i>Power Saving Mode / Minimum Frequency</i>
E1080 @3.55GHz	27C	3.55 - 4.00 GHz	3.25 GHz
E1080 @3.60GHz	27C	3.60 – 4.15 GHz	3.4 GHz
E1080 @3.65GHz	27C	3.65 - 3.90 GHz	3.25 GHz

Appendix II: Processor Usage and Accounting

For historical reasons, process accounting charges and processor utilization are usually formulated in terms of time. The processors used in early time-sharing computer systems were fixed-frequency and single-threaded; processing capacity per unit of time was relatively constant. Consequently, it was convenient to express process accounting charges in terms of processor time and processor utilization as the percentage of time that the processor was not idle over an interval of interest.

With the introduction of multi-threaded processors (i.e. processors capable of executing multiple programs simultaneously), it was no longer desirable to report processor utilization based on the percentage of time that the processor was not idle. Consider a processor that supports 2-way Simultaneous Multi-Threading (SMT). When both threads are idle, utilization should be 0%. When both threads are not idle, utilization should be 100%. When one thread is idle and the other is not idle, there are several options:

1. Treat the processor as idle (0% utilization). This is obviously wrong, as the process is consuming more than 50% of the processing capacity.
2. Continue to treat the processor as not idle and charge the process with 100% of the time (100% utilization). This is a little better than the first option, but it causes utilization to be over-reported since it does not recognize the capacity available in the idle thread. The process is also significantly over-charged compared to when it shares the processor with another process.
3. Treat processor threads as individual processors, charging the process with 100% of the time and reporting processor-thread based utilization (50% utilization). This causes utilization to be under-reported because SMT efficiencies are much less than 100%, that is, the not idle thread typically represents only 15%-30% additional capacity. The process is also significantly over-charged compared to when it shares the processor with another process.

To address this problem, the Processor Utilization Resource Register (PURR) was introduced on Power5™ for the purpose of apportioning processor time among the processor's threads. The PURR is defined such that $\sum (\Delta \text{PURR}) = \Delta \text{TIME}$ over any interval. For each SMT context, PURR ticks are apportioned among the idle and non-idle processor threads, so that non-idle PURR as a portion of available time for the processor reflects the relationship between throughput and processor utilization observed for a representative commercial processing workload. By expressing processor utilization as the ratio of not idle PURR ticks to available time and by expressing process accounting in terms of PURR ticks, the historical definition and relationship between processor utilization and process accounting was maintained.

IBM Power10 processor-based systems with EnergyScale employ variable processor speed technology to dynamically optimize the speed and energy usage of the processor to the demand of the workload. Because the PURR ticks at a constant rate independent of processor speed, PURR-based processor utilization remains a useful and accurate metric, but PURR-based process accounting charges can vary depending on processor speed. To address this problem, the IBM Power6 processor included a new per-thread processor timekeeping facility to normalize the relationship between processor time and processor speed. The new facility was named the Scaled Processor Utilization Resource Register (SPURR), and it represented processor time at nominal (i.e. 100%) speed.³ The SPURR was primarily intended to improve process accounting consistency, but it can also be used in conjunction with the

³When the Power10 processor is operating at full speed, the PURR and SPURR tick in lockstep; when the Power10 processor is operating at reduced speed, the SPURR ticks slower than the PURR; when the Power10 processor is operating in excess of full speed, the SPURR ticks faster than the PURR.

PURR in processor speed and capacity calculations. For example, a SPURR to PURR percentage of 85% indicates that the processor operated at 85% nominal speed over a sample interval.

While the PURR and SPURR provide the information necessary to measure accurate processor utilization and to perform consistent process accounting in the variable processor speed environment, they do not address the ambiguity of CPU time in the historical context. Simply stated, whether a CPU time value in a legacy software interface should represent PURR-based CPU time or SPURR-based CPU time is open to some interpretation. There is no single best choice to handle all cases, and in fact, the issue has not been uniformly addressed by all operating systems or even among versions of the same system. There is general agreement that SPURR-based process accounting is preferable to PURR-based accounting since the results are more consistent across EnergyScale modes and within modes that vary the processor speed dynamically.

Appendix III: Resources

AIX:

<http://publib.boulder.ibm.com/infocenter/pseries/v6r1/index.jsp>

(Search for "spurr" to discover references regarding enablement in APIs and/or tools)

IBM i:

[IBM Docs – IBM i Documentation](#)

(Search for “Energy management”, “Processor folding”, “Scaled processor time attribute”, “Processor time”, “Scaled processor time”, “Processor utilized time”, “Processor scaled utilized time”, “Processor interrupt time”, “Processor scaled interrupt time”, “Processor stolen time”, “Processor scaled stolen time”, “Processor donated time”, “Processor scaled donated time”, “Processor idle time”, “Processor scaled idle time”)

IBM EnergyScale for Power9 Processor-Based Systems:

https://www.ibm.com/downloads/cas/6GZMODN3?mhsrc=ibmsearch_a&mhq=EnergyScale

Power Systems:

<https://www.ibm.com/it-infrastructure/power>

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

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

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. In the United States and/or other countries.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.



© IBM Corporation 2021
IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
October 2021
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, ibm.com, AlX, EnergyScale, i5/OS, Power, Power6, Power7, Power8, Power9, Power10, System i and System p are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

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

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

When referring to storage capacity, 1 TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at: <http://www.ibm.com>.

The IBM Power home page on the Internet can be found at: <http://www.ibm.com/systems/power/>