



eServer i5 Uncapped Partition CPU Utilization & Queuing

i5 Uncapped Partition Utilization & Queuing

Table of Contents

Trademarks and Disclaimer	3
Audience.....	4
Objectives.....	4
Introduction to Uncapped Mode	5
CPU Utilization Data	6
Partition 5250 Capacity.....	6
Partition Total Capacity	6
Shared Pool Capacity	6
Collecting Shared Pool Information.....	7
CPU Utilization Calculation (V5R3).....	8
Total CPU Utilization Calculation Options	8
Queuing and CPU Utilization	8
Available Capacity	10
Conclusions.....	11

i5 Uncapped Partition Utilization & Queuing

Trademarks and Disclaimer

IBM Corporation 1994-2005. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

- IBM
- IBM (logo)
- AS/400
- AS/400e
- eServer
- iSeries
- OS/400
- i5/OS

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

The examples given in this paper are hypothetical examples of how a customer can use the products described herein and examples of potential cost or efficiency savings are not based on any actual case study. There is no guarantee of comparable results. Many factors determine the sizing requirements and performance of a system's architecture. IBM assumes no liability for the methodology used for determining the configurations recommended in this document nor for the results it provides. Any performance data contained in this presentation was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this presentation should verify the applicable data for their specific environment.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

This paper is intended to provide information regarding queuing for processors in an eServer i5 uncapped partition. These findings may not be realized in all customer environments, and implementation in such environments may require additional steps, configurations, and performance analysis. The information herein is provided "AS IS" with no warranties, express or implied. This information does not constitute a specification or form part of the warranty for any IBM products.

i5 Uncapped Partition Utilization & Queuing

Audience

This review is intended for use by technical professionals interested in understanding the concepts of capped and uncapped modes of operation of i5 partitions using the shared processor pool of processors, and the impact of queuing on performance in an uncapped environment.

It is assumed that the reader has a good understanding of the concepts of logical partitioning, including virtual processor considerations in an uncapped environment.

Objectives

The purpose of this document is to provide a technical discussion on the effects of queuing in i5 partitions configured to process in an uncapped mode. The concept of two alternatives to calculate CPU utilization (based on the configured capacity and the available capacity) is presented. The following documents discuss other POWER5 and partition related topics:

- i5/OS LPAR Performance on POWER4 & POWER5
<http://www-1.ibm.com/servers/eserver/series/perfmgmt/pdf/lparperf.pdf>
- Memory Affinity
<http://www-1.ibm.com/servers/eserver/series/perfmgmt/pdf/memaffin.pdf>
- Simultaneous Multithreading on POWER5
<http://www-1.ibm.com/servers/eserver/series/perfmgmt/pdf/SMT.pdf>

NOTE: The algorithms described here are at best, approximations of a real, live iSeries environment. Any predictions made using these algorithms are rough approximations and may or may not be close to what will be observed in a real-world environment. It is up to the tool developer - and most importantly - the user of the tool and the person who analyzes the data to judge the accuracy of the predictions made by the model. IBM makes no guarantees regarding the accuracy of predictions made using these algorithms and formulae. Please read the disclaimer at the beginning of this document.

i5 Uncapped Partition Utilization & Queuing

Introduction to Uncapped Mode

Partitions using processors capacity in the shared processor pool are defined either as **capped** or **uncapped**. Capped partitions cannot exceed the configured processor capacity. Uncapped support for logical partitions enables partitions to exceed their configured capacity in situations where there is unused capacity in the shared processor pool because either (a) other partitions assigned capacity from the shared processor pool are not using all of their configured capacity or (b) the capacity of the shared pool is not completely allocated to partitions or both (a) and (b).

When the processor utilization of an uncapped partition exceeds its configured capacity, the processor utilization metric for this partition (as indicated by many IBM tools) will be greater than 100% because the computation is based on the **configured capacity**.

A logical partition that uses processor capacity from the shared processor pool can be assigned one of two modes:

Capped: A shared processor logical partition can only use the amount of configured processing units. Thus, a partition in capped mode cannot use additional processor capacity even if there is unused processor capacity in the shared processor pool. Only capped mode logical partitions were supported prior to V5R3.

Uncapped: Allows a shared processor logical partition to utilize processor capacity in excess of its configured processor capacity, provided there is unused capacity in the shared processor pool. The number of virtual processors is one of the factors that determine the maximum processor capacity that can be used by the partition. This capability is available with V5R3 and is supported on eServer i5 hardware.

The upper limit of processor capacity that may be used by an uncapped partition at any time is determined by the lower of

- The number of processors in the shared processor pool
- The number of virtual processors
- The actual usage plus the unutilized capacity in the shared processor pool

Note: In this document, the number of processing units assigned to a partition at a particular time is referred to as “configured capacity”. This is also known as:

- entitled capacity
- current capacity (on the HMC)
- guaranteed capacity (now less frequently used)
- assured capacity

Keep in mind that this value can change while a partition is active though dynamic logical partition (DLPAR) facilities.

i5 Uncapped Partition Utilization & Queuing

CPU Utilization Data

Since i5/OS (OS/400 V5R3) CPU configuration and usage information is recorded in the QAPMSYSTEM file (in milliseconds) for each performance interval. For completeness, many of the new metrics are listed here, but this discussion will focus on the Total Capacity of a partition:

Partition 5250 Capacity

- **SYIFUS:** Interactive CPU time used. The total CPU time charged to 5250-OLTP capacity.
- **SYIFTA:** Interactive CPU time available. This is the interactive (5250-OLTP) capacity configured for use by the partition.

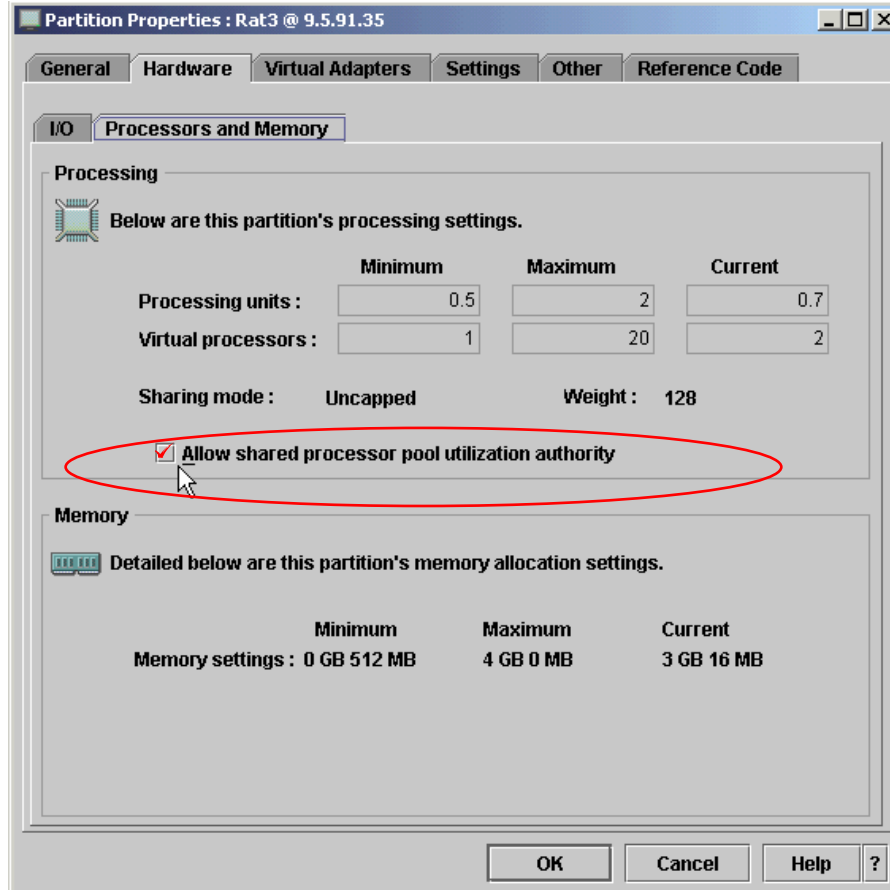
Partition Total Capacity

- **SYSPTU:** Total CPU time used by the partition.
- **SYSCTA:** Total CPU time configured for the partition. This is the system processing capacity as determined by processor unit allocations during the interval. Note: For capped partitions, the actual utilization cannot exceed this value, but for uncapped partitions, the actual CPU used can exceed this value.
- **SYSUTC:** Uncapped CPU time configured. The maximum amount of CPU time that the partition is allowed to use within the shared pool (adjusted for configuration changes over time). This field represents the minimum of the number of virtual processors configured or the number of processors in the shared pool. In spite of its name, this field contains values for capped and dedicated partitions too. For capped and dedicated partitions, this value is the same as the CPU time configured (SYSCTA).
- **SYSUTA:** Total CPU time that could have been used by the partition (adjusted for configuration changes over time). For uncapped partitions, it considers the configured processor capacity and configured virtual processors, as well as the shared processor pool capacity that was not used. For capped and dedicated partitions (or if shared pool data is not available), this is the same as Shared Pool Capacity
- **SYSPLU:** Shared pool CPU time used. Total amount of CPU used within the shared pool by all partitions that share the pool. Set to zero if a shared pool is not used or the data is not available.
- **SYSPLA:** Shared pool CPU time available. This is the total amount of CPU available within the shared pool. This value is determined based on the number of physical processors that are allocated to the pool. This is set to zero if a shared pool is not used or the data is not available.

i5 Uncapped Partition Utilization & Queuing

Collecting Shared Pool Information

The shared pool utilization information will only be in a partition's performance data if the check-box titled “**allow shared processor pool utilization authority**” is selected in the partition properties parameter dialog accessed through the HMC:



If this data is not collected, evaluation of uncapped partition profile data may lead to erroneous conclusions.

i5 Uncapped Partition Utilization & Queuing

CPU Utilization Calculation (V5R3)

Beginning with V5R3 of OS/400 (i5/OS), IBM defined the percentage processor utilization (5250 OLTP or total) as the CPU time used (in milliseconds) as a **percentage of the configured capacity** (in milliseconds). This definition allows the CPU Utilization to exceed 100%, during periods when the CPU time used exceeds the configured CPU time (Uncapped partitions).

Total CPU Utilization Calculation Options

However, the performance data presents multiple options to calculate Total CPU utilization, which include:

1. **Total CPU Utilization (configured) = (SYSPTU/SYSCTA) * 100**
2. **Total CPU Utilization (available) = (SYSPTU/SYSUTA) * 100**

In most IBM performance reports and displays, the CPU Utilization is based on the configured capacity. There are definite benefits in reviewing utilization based on a specific value defined within the partition, particularly when managing the partition configuration parameters.

Queuing and CPU Utilization

A basic concept of Queuing Theory is that as the utilization of a server increases, the time a request spends waiting to be serviced increases according to the formula

- Queuing Multiplier (QM) = $1/(1-U)$ or
- Queuing Multiplier (QM) = $1/(1-U^p)$ and
- Server Response Time (T) = $S*QM$

Where

- U = Server Utilization %
- p = the number of servers
- S = Server Service Time

Note: This discussion is confined to the processors only (servers responding to requests for processing time). The impact of CPU queuing will affect CPU response time only. The user perceived response time would include service times of other servers like disk and communications, wait time etc.

The impact of queuing (for a server) on response time of a request is a function of usage compared to the available capacity. In an uncapped partition, the available capacity at the time will determine the percentage of utilization at the time, and the computation of CPU utilization and the Queuing Multiplier must be based on the available capacity and not on the configured value. This is particularly true when it exceeds 100%.

i5 Uncapped Partition Utilization & Queuing

Thus, computing the impact of queuing on processor response time may be based on:

- **CPU Utilization (available)** at the time (SYSUTA). The utilization will never exceed 100%.
- Number of processor servers (p) is the **number of processors in the shared processor pool**.

Capped Partition: The maximum available CPU capacity in a given period (measured time interval) for a capped partition is the number of milliseconds configured for the partition (adjusted for configuration changes over time). In a capped partition the available CPU capacity is the same as the configured CPU capacity.

Uncapped Partition: The maximum CPU capacity available in a given period (measured time interval) for an uncapped partition is the number of milliseconds that could have been used by the partition (adjusted for configuration changes over time). This value is the lower of

- The number of processors in the shared processor pool
- The number of virtual processors
- The actual usage plus the unutilized capacity in the shared processor pool

However, a partition is only assured of its configured capacity. Thus, a partition may exceed its configured capacity only if there is unused capacity in the shared processor pool.

Examples:

1. Limited by the size of the shared pool - If a shared processors pool has 4 processors, then the maximum capacity a partition may use is limited to 4 processors units, even if 6 virtual processors are defined. (Ensure that there is good reason to specify the number of virtual processors greater than the number processors in the pool!).
2. Limited by the number of virtual processors - If a shared processors pool has 4 processors, and 3 virtual processors are specified, the maximum capacity of the partition cannot exceed 3 processing units.
3. Limited by the unused capacity in the shared pool - If a shared processors pool has 4 processors, and 6 virtual processors are defined, but there is only 0.5 processors of capacity unused in the shared processor pool at this time, the maximum capacity available to the partition is the sum of its configured capacity and the unused 0.5 processor units in the shared pool.

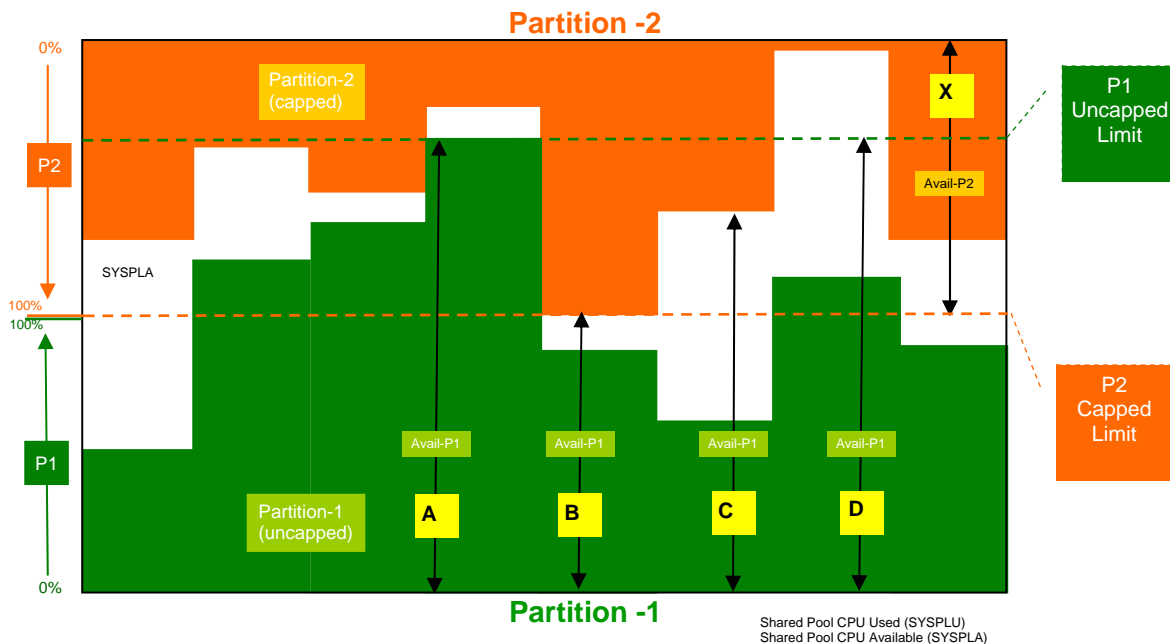
i5 Uncapped Partition Utilization & Queuing

Available Capacity

The diagram below attempts to clarify the amount of processing capacity that would be available with two partitions (one capped and the other uncapped) using a shared processor pool.

- The **capped partition** (Partition-2, orange in the diagram), has its usage shown from the top of the chart (0%) and increases downwards. The upper limit of utilization by the capped partition is determined by the configure capacity – and is represented by the dotted orange line.
- The **uncapped partition** (Partition-1, green in the diagram) has its usage represented from the bottom of the chart (0%) and increases upwards. The upper limit of utilization by the uncapped partition is determined (in this example) by the number of virtual processors – and is represented by the dotted green line.
- The unused capacity of the shared processor pool is in the mid-region of the chart (in white where available). The information is recorded in the field SYSPLA (provided the appropriate authority is granted. See “allow shared processor pool utilization authority on page 7 of this document).

For simplicity, the chart assumes that the number of virtual processors assigned does not exceed the shared processor pool capacity!



In a capped partition, the available CPU capacity (SYSUTA) is always equal to the configured capacity (SYSCTA) (represented by X above).

In an uncapped partition, the available CPU capacity (SYSUTA) depends on the amount of unused capacity in the shared pool. In the diagram, the following points represent different examples of available CPU capacity in an uncapped partition:

i5 Uncapped Partition Utilization & Queuing

A – The uncapped partition uses all of the capacity defined by the number of virtual processors. The CPU utilization (configured) is **greater** than 100%. Even though there is more unutilized capacity in the shared pool (the “white space” in the diagram), the uncapped partition is unable to take advantage of this. The available CPU capacity (SYSUTA) is equal to the number of virtual processors.

B – The uncapped partition uses less than its configured capacity. The CPU utilization (configured) is **less** than 100%. The available CPU capacity (SYSUTA) is the sum of the capacity used and the unused capacity in the shared pool. But in this case, it is also equal to the configured capacity of the uncapped partition (because the capped partition is using all of its capacity and the only unused capacity in the shared pool is configured in the uncapped partition).

C – The uncapped partition uses less than its configured capacity. The CPU utilization (configured) is **less** than 100%. The available CPU capacity (SYSUTA) is the sum of the capacity used and the unused capacity in the shared pool.

D – The uncapped partition uses more than its configured capacity. The CPU utilization (configured) is **greater** than 100%. The available CPU capacity (SYSUTA) is the sum of the capacity used and the unused capacity in the shared pool (but not exceeding the number of virtual processors). So, in this case, it is represented by the number of virtual processors specified.

Conclusions

- IBM iSeries performance reports and displays present CPU Utilization based on the configured value in the partition.
- In uncapped mode, partitions may exceed its allocated CPU Utilization, and result in reported CPU Utilization values in excess of 100%.
- Referencing processor utilization to a value defined in the partition provides useful information for managing the capacity of the partition.
- However, for queuing calculation, the available CPU utilization (SYSUTA) may be used to obtain an acceptable result.