

# **Sizing IBM eServer i5 Servers for AIX 5L Applications**

**James Cioffi**  
**iSeries Technical Support Marketing**  
**Version 1.3**  
**8/16/04**

## 1.0 Overview

With the introduction of the IBM @server i5 and @server p5 line of servers, there is now a common POWER 5 based hardware platform for AIX 5L, i5/OS, and Linux on POWER applications. To take advantage of this capability, there is a need to be able to estimate the amount of server capacity needed to support various applications on logically partitioned systems. The purpose of this paper is to provide a guide for how to approach sizing servers to run AIX 5L applications on eServer i5 servers.

The methodology is to:

1. Use existing pSeries and eServer p5 estimating procedures and tools.
2. Apply these estimates to the eServer i5 server environment while taking into consideration the configuration differences that might be present.
3. Incorporate the estimate into the eServer i5 server sizing process.

**Note:** Because every eServer i5 server that runs AIX 5L will be a partitioned server, this process is best performed by a capacity planning specialist with LPAR knowledge. There are too many variables involved with partitioned servers to cover all the details here. Capacity planning and LPAR knowledge is assumed.

## 2.0 IBM @server i5 and IBM @server p5 Performance & Sizing Metrics

IBM has historically used different metrics to describe the capacity of the iSeries and pSeries server lines. These metrics are used to compare the relative performance of various server models based on certain predefined workloads. While these workloads are not the same as “real world” workloads, the metrics that are based on them are frequently used to do estimates of server sizing.

The most frequently used metric in the iSeries world is CPW (Commercial Processing Workload). iSeries also publishes measurements for MCU (Domino Mail and Calendar Users). The CPW and MCU ratings for iSeries servers can be found in the *IBM @server i5 and iSeries System Handbook*. The latest version can be found at:

[Http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/ga195486.html?Open](http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/ga195486.html?Open)

CPW is described in the handbook as follows:

“The performance capacity of all iSeries servers is represented by a workload measurement called Commercial Processing Workload (CPW). CPW values are given to all iSeries processors and are derived by performing various monitored and measured workloads on iSeries servers. The results (reported values) can be used to compare relative performance characteristics of processor features offered for iSeries servers. The reported values for CPW do not represent a guaranteed level of capacity to perform a given workload. They can serve as a quick means to compare performance.”

A more detailed description of CPW can be found in Appendix A of the *iSeries Performance Capabilities Reference i5/OS V5R3* found here:

<http://www-1.ibm.com/servers/eserver/iseries/perfmgmt/resource.htm>

pSeries publishes a number of different performance benchmarks and metrics, the most notable being “rPerf” or “Relative Performance”. These are published in the *IBM @server pSeries and IBM RS/6000 Performance Report* which describes rPerf as follows:

“The rPerf (Relative Performance) is an estimate of commercial processing performance. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory available. However, the model does not simulate disk or network I/O operations. Although the model uses general database and operating system parameters, it does not reflect specific database or AIX version or releases.”

The latest version of this report can be found at:

[http://www-1.ibm.com/servers/eserver/pseries/hardware/system\\_perf.html](http://www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.html)

### 3.0 @server i5 rPerf values

IBM has published rPerf values for the eServer p5 520, 550, and 570 models. While IBM has not run the rPerf benchmark with eServer i5 configurations, the rPerf values for eServer i5 servers can be derived from the published eServer p5 rPerf values due to the fact that many eServer i5 server models have eServer p5 electronic equivalents. The table below lists the published rPerf values for several eServer p5 server models that have eServer i5 equivalents.

eServer p5 Server Model	No of CPUs	GHz	AIX 5L 5.3 rPerf
eServer p5 520	2	1.65	9.86
eServer p5 550	2	1.65	9.86
	4	1.65	19.66
eServer p5 570	2	1.65	9.86
	4	1.65	19.66
	8	1.65	37.22
	12	1.65	53.43
	16	1.65	68.40

Notes:

1. The rPerf results of IBM eServer p5 systems are based on AIX 5L V5.3 which supports simultaneous multithreading. Simultaneous multithreading results in a 30% boost in rPerf. For purposes of estimating rPerf for AIX 5L V5.2 on eServer p5 systems divide the published rPerf number by 1.3 and round down to the nearest 1/100th (i.e. 25.999 = 25.99).
2. These performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

Table 1. rPerf values for eServer p5 servers

The IBM pSeries and IBM RS/6000 Performance Report provides rPerf values for other pSeries servers [http://www-1.ibm.com/servers/eserver/pseries/hardware/system\\_perf.html](http://www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.html)

When considering how to use rPerf values for eServer i5 servers, it is important to remember that there are some differences between the environments in which AIX 5L applications may run on eServer i5

and eServer p5 servers. These differences and their performance and capacity impact must be accounted for when estimating the size of a server partition needed to run an AIX 5L application.

- Logical partitioning configuration
- Disk configuration differences
- Virtual Storage considerations
- Virtual LAN considerations

The rPerf benchmark results do not take into account these advanced configuration options. As a result, extra steps are required to develop sizing estimates for AIX 5L workloads on eServer i5 and eServer p5 servers using LPAR and Virtual I/O.

#### 4.0 pSeries sizing tools and procedures

Server sizing recommendations for specific applications typically come from the application vendors, either Independent Software Vendors (ISVs) or IBM Software Group. Guidelines and/or tools for sizing pSeries servers can be found at vendors' web sites and (for several ISV applications) at the IBM PartnerWorld for Developers web site at:

<http://www.developer.ibm.com/welcome/eserver/e3/CSFServlet?mvcid=main&packageid=3002> .

A list of resources that help with sizing some of the more popular AIX 5L applications is in Section 7.0, "Sizing Guides, Other Information Sources & References."

As with sizing applications for i5/OS, sizing for AIX 5L requires knowledge of the application, its operating environment, and general technical knowledge of AIX 5L and pSeries. If you are inexperienced with sizing pSeries applications, an excellent redbook on pSeries sizing can be found at: <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247071.html?Open>

If you are experienced with sizing pSeries applications or have help from someone who is, the first step in sizing eServer i5 server to run an AIX 5L application is to determine the eServer p5 server requirements for processor, memory, disk, and network I/O. It is reasonable to assume that these requirements will be the same for eServer i5. The next step is to determine what server model and how much CPU capacity will be needed for the AIX 5L partition on the eServer i5.

#### 5.0 LPAR Considerations

Once you have determined what processor, memory, disk, and network requirements are for an eServer p5 Server, the next step is to estimate the size of the partition needed to support the AIX 5L workload. The rPerf numbers can be used as a guide for determining the amount of CPU that should be allocated to the partition. However it is not simply a 1 to 1 relationship. There are many factors that affect capacity in an LPAR. This is true for AIX 5L partitions the same as it is for i5/OS and Linux partitions.

To deal effectively with these when sizing partitioned servers, LPAR education is strongly recommended. IBM offers an excellent class on LPAR on IBM eServer i5, the details of which can be found at:

<http://w3-3.ibm.com/education/CoursecodeSearchTypeServlet.wss?coursecode=as530>

Consider the impact of the following:

1. Partitioning - rPerf numbers are calculated using unpartitioned system. AIX 5L applications will never run on an eServer i5 server this way. They will always be in a partition or a portion of the eServer i5 server's total capacity. This must be accounted for manually through interpolation or automatically in the eServer Workload Estimator. The Workload Estimator is discussed in the next section.
2. LPAR Processing Capacity - There is overhead associated with partitioning that must be considered. The actual amount depends on a number of factors such as the number of processors in the system and the partition and the utilization of those processors. **The Workload Estimator will calculate this automatically for you.** Therefore to avoid double-counting this overhead, when sizing the AIX 5L workload for an eServer p5 server, you should size it as though it is running on a dedicated server. Details of LPAR performance are described in the white paper that can be found at: <http://www-1.ibm.com/servers/eserver/iserries/perfmgmt/resource.htm>
3. Uncapped processor partitioning - The use of uncapped processor partitioning can enable a system to handle brief peaks in demand for processor utilization by using the unused processor capacity in the shared pool. The use of this feature can provide a lot of flexibility in managing capacity but also increases the complexity when you are doing a sizing. When you size an AIX 5L workload for an eServer p5 server, the results assume a certain number of processors at some average or peak utilization. But your partition may or may not run at that utilization. This needs to be accounted for. This is another reason that the LPAR education is strongly recommended.
4. Memory configuration - Memory requirements for a partition should be the same as required for an eServer p5 server. However, the Workload Estimator requires you to know how much memory is needed to support the AIX 5L application workload separately from AIX 5L. The Workload Estimator automatically reserves 64MB for each AIX 5L partition.
5. Disk arms and configuration - Although disk I/O is not a factor in rPerf calculations, it is assumed that it is not a constraint. You must ensure that the disk configuration has:
  - I/O adapters that can handle the required throughput (ops per second)
  - enough disk arms for acceptable performance.
  - Although it is not necessarily a factor in sizing, high availability and backup/recovery must be considered too. If the AIX 5L partition is configured with dedicated, direct I/O, it requires a separate HA and backup strategy from the i5/OS partitions.
  - If the AIX 5L partition uses virtual I/O through an i5/OS hosting partition, some additional overhead is incurred in the hosting partition. More information on how to account for this will be published in a subsequent version of this paper.The Workload Estimator will automatically reserve 2.2GB of disk space per AIX 5L partition to hold the operating system files.
6. Network I/O - Network I/O requirements, including virtual LAN considerations should be similar to those for eServer p5. Network traffic peaks must be considered when adding up the throughput requirements. Virtual LAN introduces a very small amount of CPU overhead in a partition.
7. Adjustment for AIX 5L version - The published rPerf numbers for eServer p5 processors state that you must divide the rPerf rating for a server if the workload is AIX 5L V-5.2 and does not exploit SMT. To compensate for this when sizing an AIX 5L V-5.2 workload for eServer i5, you should **multiply the rPerf that the workload requires by 1.3** before entering it into the Workload Estimator. This provides the uplift needed to drive the additional processor needed to support the

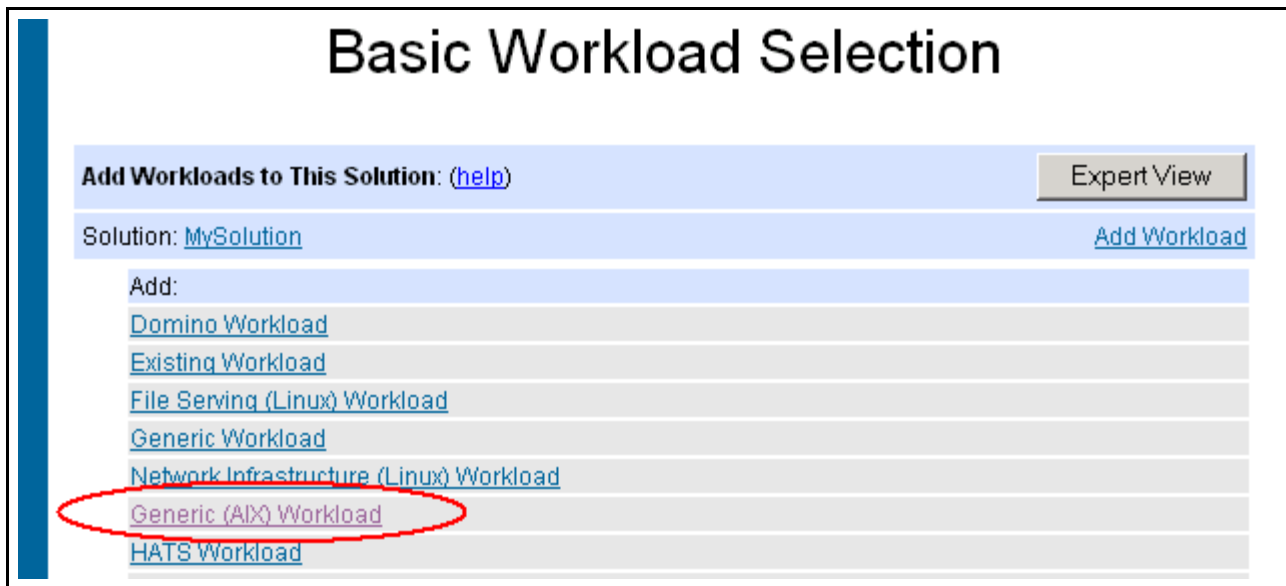
AIX 5L V-5.2 workload. Also, if the workload is AIX 5L V-5.3 but does not exploit SMT (for example a single-threaded job), then you should treat it as an AIX 5L V-5.2 workload. **Caution: Make sure that the rPerf value you are using has no *already* taken this into account. If the person who provided the rPerf value has sized the workload for an eServer p5, they may have already multiplied the rPerf by this factor.**

Once you have determined the rPerf requirement for the AIX 5L workload and factored in the factors mentioned above, you can derive an estimate of the number of eServer i5 processors needed to support the AIX 5L workload in a partition. You should also have the memory and disk requirements. Now you are ready to add this to your i5/OS workload(s) to produce a server sizing.

## 6.0 Sizing a Consolidated Server

The recommended tool for sizing eServer i5 servers for new workloads is the eServer Workload Estimator. The Workload Estimator is available on the web, here: <http://www-912.ibm.com/wle/EstimatorServlet>

In order to enter the AIX 5L workload into the Workload Estimator, you will create a “Generic Workload”. Start the Workload Estimator and you are presented with the following screen:



The screenshot displays the 'Basic Workload Selection' interface. At the top, there is a header 'Basic Workload Selection'. Below it, a navigation bar includes 'Add Workloads to This Solution: (help)' and an 'Expert View' button. The main area shows 'Solution: MySolution' with an 'Add Workload' link. A list of workload options is presented under the 'Add:' label. The options are: Domino Workload, Existing Workload, File Serving (Linux) Workload, Generic Workload, Network Infrastructure (Linux) Workload, Generic (AIX) Workload (circled in red), and HATS Workload.

Select the Generic (AIX) Workload and click on it in the configuration to edit it. Be sure to choose the correct option regarding the partition type, Shared Processor or Dedicated Processor. You will be presented with this form:

IBM @server Workload Estimator

GenAIX #1  
Generic (AIX) Workload Definition

Shared Processor LPAR  
Partition Name: [AIX Partition](#)  
OS: AIX - 5.3

1. [rPerf](#):

2. [Memory](#) : (MB)

3. [Disk Configuration](#):

Drive Grouping:  Group 1

Drive Attachment:

Drive Type:

Data Protection:

Drive Units:

Storage (GB):

Enter the information you have for the AIX 5L partition:

Field	Enter this:
rPerf	Enter the number rPerf required for this workload. Be sure to include adjustments for partitioning and relevant adjustments for the release of AIX 5L to be run. AIX 5L 5.2 workloads require more rPerf than AIX 5L 5.3 workloads, therefore if this is an AIX 5L 5.2 workload, be sure to multiply the required rPerf by 1.3 before entering it in this field. This workload assumes AIX 5L 5.3 applications.
Memory	Enter the amount of memory the workload requires. Enter <i>only</i> the amount needed for the workload you are sizing. Do not include AIX 5L overhead. The Workload Estimator automatically allocates 64MB per partition for AIX 5L. If you know that AIX 5L will need more than 64MB to support a workload you can add that additional amount to one of the AIX 5L workloads in the partition you are sizing.  <b>Note:</b> The Workload Estimator will always allocate AT LEAST 256MB per AIX 5L partition.

Disk Configuration	Enter the appropriate Disk Configuration for the AIX 5L partition. Specify the number of Disk Arms needed for adequate performance as well as the required amount of storage in GB. The Workload Estimator will automatically add 2.2GB per partition to account for the installation of the AIX 5L operating system.
Table 2. Generic Workload Calculation	

Finish creating the other workloads that will run on the eServer i5 server and WLE will produce a consolidated server sizing estimate.

## 7.0 Sizing Guides, Other Information Sources & References

### **General Information:**

iSeries Information Center:

<http://publib.boulder.ibm.com/pubs/html/as400/infocenter.html>

pSeries Information Center:

[http://publib16.boulder.ibm.com/pseries/en\\_US/infocenter/base/](http://publib16.boulder.ibm.com/pseries/en_US/infocenter/base/)

AIX 5L Web Site

<http://www.ibm.com/servers/aix>

AIX 5L on eServer i5

<http://www.ibm.com/series/aix>

### **Sizing Information**

Sizewise web site - Server sizing information for the IBMers in the Americas:

[http://w3-1.ibm.com/support/americas/techline/q\\_swcap.html](http://w3-1.ibm.com/support/americas/techline/q_swcap.html)

eServer Workload Estimator

<http://www-912.ibm.com/wle/EstimatorServlet>

iSeries Performance Capabilities Reference i5/OS V5R3 & other eServer i5 Performance-related documentation

<http://www-1.ibm.com/servers/eserver/series/perfmgmt/resource.htm>

IBM pSeries and IBM RS/6000 Performance Report

[http://www-1.ibm.com/servers/eserver/pseries/hardware/system\\_perf.html](http://www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.html)

eServer Solutions Sizing Tools on PartnerWorld

<http://www.developer.ibm.com/welcome/eserver/e3/CSFServlet?mvcid=Welcome&packageid=000>

Sizing Enterprise Resource Planning (ERP) workloads:

<http://www.ibm.com/erp/sizing>

Sizing SAP workloads:

<http://www.ibm.com/erp/sap/insight>

ISV Sizing Flash Techdoc (Americas):

<http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10116>

Search for Sizing Techdocs:

<http://www-1.ibm.com/support/techdocs/atsmastr.nsf/Web/Sizings>

**Pre-sales support for sizing**

Techline (IBMers)	Americas: <a href="http://w3-1.ibm.com/support/americas/techline/index.html">http://w3-1.ibm.com/support/americas/techline/index.html</a> EMEA: <a href="http://w3-5.ibm.com/support/emea/techline/index.shtml">http://w3-5.ibm.com/support/emea/techline/index.shtml</a> AP: <a href="http://w3-6.ibm.com/support/ap/ap_techline.html">http://w3-6.ibm.com/support/ap/ap_techline.html</a>
Worldwide Techline Contact Information for entitled Business Partners	<a href="http://www-1.ibm.com/partnerworld/pwhome.nsf/weblook/trs_wgd_contacts.html">http://www-1.ibm.com/partnerworld/pwhome.nsf/weblook/trs_wgd_contacts.html</a>

IBM, the IBM logo, the e-business logo, AIX 5L, DB2 Universal Database, eServer, i5/OS, iSeries, OS/400, POWER, POWER4, POWER Hypervisor, PowerPC, pSeries, Tivoli, WebSphere and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

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

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

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

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.