Large Systems Performance Reference

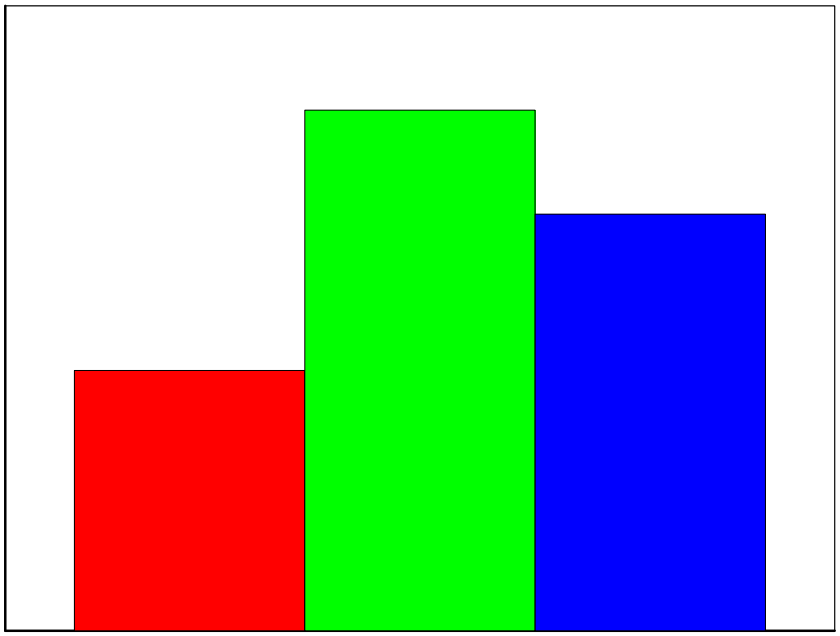

Large Systems Performance Reference

Document Number SC28-1187-14

Note

Before using this information and the product it supports, be sure to read the general information under "Notices".

Fourteenth (July 2010) This edition is a major update and the ITRR tables complement those included in SC28-1187-13. This release, SC28-1187-14, contains performance information for z/OS V1R11 for all z/Architecture models including the new System zEnterprise 196 (z196) models. Content other than ITRR tables in previous editions is superseded by the latest version.

Information contained in this publication is subject to change from time to time. Any such changes will be reported in subsequent revisions or Technical Newsletters.

This publication will only be available in PDF format at the following web address.

https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument

Table of Contents

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the user's responsibility.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

Trademarks and Service Marks

The following terms, denoted by the symbol (™) in this publication, are trademarks or service marks of the IBM Corporation in the United States or other countries or both:

z900	zSeries	IBM eServer
z/Architecture	z/VM	z/OS
ACF/VTAM	AIX	AS/400
CICS	CICS/ESA	CICS/MVS
CICS/VSE	DB/2	DB2
DFSMS	ES/3090	ES/4381
IBM S/390 Parallel Enterprise Server-Generation 6	ES/9000	ES/9370
IBM S/390 Parallel Enterprise Server-Generation 5	ESA/370	ESA/390
IBM S/390 Parallel Enterprise Server-Generation 3	IBM	3090
IBM S/390 Parallel Enterprise Server-Generation 4	IMS/ESA	GDDM
IBM S/390 Multiprise 2000	OS/390	VSE/ESA
Large Systems Performance Reference	LPAR	VTAM
LSPR	LSPR/PC	MVS/DFP
MVS/ESA	MVS/SP	MVS/XA
Processor Resource/Systems Manager	PR/SM	OS/400
System/370	System/390	RACF
S/390 Multiprise	VM/ESA	VM/XA
WebSphere	VisualAge	
IBM System z9 EC formerly IBM System z9 109		
IBM System z9 BC		
IBM System z10 EC		
IBM System z10 BC		
IBM System zEnterprise 196 or z196		

The following terms, also denoted the symbol (™) in this publication, are trademarks or service marks of other companies or individuals listed below:

MARK	**OWNER**
EXPLORE	Goal Systems International Inc
MSC/NASTRAN	MacNeal-Schwendler Corp.
UNIX	The Open Group in the United States and other countries
Linux	Linus Torvalds

Preface

The IBM™ Large System Performance Reference™ (LSPR™) ratios represent IBM's assessment of relative processor capacity in an unconstrained environment for the specific benchmark workloads and system control programs specified in the tables. Ratios are based on measurements and analysis. The amount of analysis as compared to measurement varies with each processor.

Many factors, including but not limited to the following, may result in the variances between the ratios provided herein and actual operating environments:

- Differences between the specified workload characteristics and your operating environment
- Differences between the specified system control program and your actual system control program
- I/O constraints in your environment
- Incorrect assumptions in the analysis
- Unknown hardware defects in processors used for measurement
- Inaccurate vendor claims.

IBM does not guarantee that your results will correspond to the ratios provided herein. This information is provided "as is", without warranty, express or implied.

Abstract

IBM's Large Systems Performance Reference (LSPR) method is designed to provide relative processor capacity data for IBM System/370™, System/390™, and z/Architecture™ processors. All LSPR data is based on a set of measured benchmarks and analysis, covering a variety of system control program (SCP) and workload environments. LSPR data is intended to be used to estimate the capacity expectation for a production workload when considering a move to a new processor.

IBM considers LSPR data to be a reliable set of relative processor capacity data. This is the only reference of its type that is based primarily on actual processor measurements. Because it is based on measurements, LSPR data takes into account individual SCP and workload sensitivities to the underlying design of each processor represented.

Chapter 1. Background

IBM's Large Systems Performance Reference™ (LSPR)™ method is intended to provide IBM System/370™ and System/390™ architecture, and z/Architecture™ processor capacity data across a wide variety of system control programs (SCP’s) and workload environments.

The LSPR's focus is solely on processor capacity, without regard to external resources such as storage size, or number of channels, control units, or I/O devices. To assure that the processor is the primary focus, the processor capacity data reported assume sufficient external resources so as to prevent any significant external resource constraints. With this approach, ***the LSPR is designed to represent each processor in its best light; that is, the processor itself is the only limiting factor to doing work***. Resulting LSPR capacity data is therefore meaningful for establishing a realistic view of relative capacity between specific processors for SCP’s and workload environments that have characteristics similar to those measured.

Rationale for Reliable Processor Capacity Data

When considering the acquisition of a large central processor, one needs to understand its capacity potential as precisely as possible. This capacity potential is generally expressed in terms relative to a currently installed processor. If expected capacity is understated or overstated, the cost of the error can be significant. That cost can be misspent dollars, or lost ability to accommodate work.

Any processor's ability to support work, either in terms of jobs, transactions, or end-users, is a function of the nature of the work to be performed. A processor's absolute capacity can be determined for any specific workload, but such absolute capacity information is not particularly useful to a capacity planner unless the information represents his exact production workload. It is difficult to produce tables that meaningfully represent processor capacity in absolute terms, except when they relate to a specific workload.

However, given that a table of absolute processor capacity values can be built for a given workload environment, we then have a basis for determining the relative capacity between those processors. These relative capacity values are meaningful, not only for the exact workload represented, but also for workloads of a similar nature.

Many processor acquisitions today are replacements for existing machines, made for the purpose of adding (or consolidating) capacity. Since the capacity of the current processor is normally well understood, the capacity of a potential new processor, relative to the current one, can be assessed by using known capacity relationships between those machines for the appropriate workload type.

Sources for Processor Capacity Data

There are several ways to establish a capacity expectation for a new processor, each with its advantages and disadvantages.

Customized Benchmark

A customized individual benchmark, which could be run on all processors of interest, will produce the most accurate capacity data on which to base a processor acquisition decision. However, ***for results to be meaningful, it is essential that a benchmark consist of representative work run in a representative way***.

Insight

If done properly, a customized representative benchmark will provide the most accurate view of relative processor capacity. However, customized benchmarks are expensive to create, maintain, and run.

A measured benchmark that is not representative of production work has little value in trying to understand processor capacity relationships.

There are many considerations to preparing a benchmark, including:

- SCP and related products (JES, RACF™, VTAM, RMF ...)
- Application subsystems (CICS™, DB2™, IMS, TSO, CMS, WAS ...)
- Other program offerings
- Application programs
- Performance monitors
- Data files (datasets), and databases
- Scripts (end-user commands), or jobs
- Working set sizes
- Terminal simulation
- Size of end-user population
- Average think time, and think time distribution
- Transaction rates
- Response time criteria
- Operational methodology
- Metrics to be used
- Repeatability and consistency
- Portability

The creation, maintenance, and measurement activity associated with a benchmark is likely to become an extremely resource intensive proposition given the complexities of the typical DP environment today. For that reason, individual customized benchmarks are not as common as they once were.

Benchmarks are all too often assembled without concern for representativeness. For example, "kernels" are used because they require minimal effort. Or a batch workload is created to represent an on-line environment, simply because batch is easier to construct and run. There is no short-cut approach to benchmarking that can assure reliable results. If the results of these types of benchmarks should happen to match those of a production workload, it would be purely coincidental, with no assurance that they would continue to match when a different set of processors is measured.

MIPS Tables

There are many published sources of processor capacity data available in the industry today. Most of these sources provide data in the form of MIPS tables. MIPS tables available from consultants and industry watchers are not based on independent measurements. Rather, they typically are developed using manufacturer's announced performance claims. Over time, some of these MIPS tables may include a subjective analysis of feedback from various clients of these systems.

Insight

While MIPS tables may be useful for rough processor positioning, they should not be used for capacity planning purposes. Single-number processor capacity tables are inherently prone to error because they are not sensitive to the type of work being processed or to the LPAR configuration of the processor.

Most published MIPS tables carry a single-number connotation; that is, the capacity of each processor can be represented with one number. Such tables are insensitive to the workload environment run on those processors.

A perceived advantage in the use of MIPS is that the implied scale is easily recognized; that is, a MIPS rating for a processor provides an easy visual positioning for that machine. This type of rating may be useful for "rough" processor positioning, but does not offer the accuracy necessary for doing capacity sizing. Studies of LSPR data for the various workload environments measured show that the potential for error, when using MIPS tables to assess relative capacity, can sometimes be significant. This is true even if the MIPS table is built from LSPR data. The problem relates to the fact that workload sensitivity and specific LPAR configurations are simply not considered in a single-number table.

Other Sources

Modeling Tools

Several modeling tools are available to aid in capacity planning for DP systems. Models are designed to provide a "big picture" view of performance, rather than to expose the capacity of any single component. In other words, models are designed for the purpose of analyzing overall system performance, given any number of resource tradeoff scenarios.

All performance models must carry some form of processor capacity data, since that element is a part of the total resource picture. Processor capacity data within a model is usually either carried as a single-number per processor table, or modeled in some way based on a set of processor and workload related characteristics.

Sophisticated modeling tools allow the use of alternative processor capacity data, instead of the built-in data or algorithm. Providers of such tools generally recommend that user-provided processor data be used unless better data is available. IBM considers LSPR data as a reliable alternative, since it is workload sensitive, and based primarily on measurements (see Figure 1).

Often a model will be used for other than its intended purpose, such as to extract processor capacity information. Modeling of this nature does nothing more than expose the underlying processor capacity data contained in the model.

System Resources Manager (SRM) Constants

One of the features of MVS, OS/390, and z/OS is that they attempt to offer somewhat consistent service units for the processor resource to do work, no matter what specific processor is being used. These service units are computed by applying the MVS, OS390, or z/OS built-in SRM constant to the CPU time consumed by each unit of work. There is a unique SRM constant defined for every processor model. IBM assigns the SRM constants for IBM processors; LSPR data is used as an input when developing the IBM assigned SRM constants. SRM constants for IBM compatible processors are supplied by each vendor.

SRM constants are sometimes used as a source of processor capacity data. By design, SRM constants are single-number metrics. Therefore, SRM constants have the same problems as MIPS when it comes to providing a precise view of processor capacity relationships, because there is no consideration for workload type.

The LSPR Alternative

Each of the above sources for processor capacity data is based on "single-number per processor" data. The interrelationships between workloads, LPAR configurations and processor design, today, are extremely complex. To be accurate, the performance of each processor must be assessed in an environment sensitive way. For this reason, ***no "single-number per processor" capacity table can necessarily provide an accurate view of relative processor capacity***.

As an alternative to the above sources, IBM offers the Large Systems Performance Reference (LSPR). LSPR data consists of a variety of workloads, each representing a type of production environment. LSPR results are based on measurements and analysis across the majority of contemporary System 370/390 and z/Architecture processors, both IBM and IBM compatible. Workloads include various batch and on-line environments that provide a reasonable representation of major types of data processing activity, such as WebSphere Application Serving, traditional on-line transactions processing, and commercial batch. These benchmarks are run using the same software that would be installed in a production environment.

The goal of the LSPR is to offer reliable relative capacity information, which takes into account processor design sensitivities to workload type.

zPCR – LSPR taken to the next level

The LSPR shows relative capacity ratios that are sensitive to workload type. However, LPAR configuration is also a very sensitive factor in capacity relationships. IBM offers a tool for customer use, zPCR, that takes the LSPR to the next level by estimating capacity relationships that are sensitive to workload type and LPAR configuration, processor configuration, as well as specialty engine configuration. All these factors may be customized to match a customer's configuration. The LSPR data is contained in the tool. For the most accurate capacity sizings, zPCR should be used.

Chapter 2. Metrics

Over time, various approaches to characterizing processor capacity have evolved. Early metrics tended to concentrate on the rate at which a processor executes instructions. One metric of this type that has survived is MIPS (millions of instructions per processor second).

As processors have grown larger and more complex in their design, the ability to characterize processor capacity relationships accurately with MIPS has diminished drastically. Processor sensitivity to different workload types must be considered as a factor in establishing capacity relationships. Therefore, metrics that relate directly to work done have been defined. These terms are **external throughput rate (ETR)** and **internal throughput rate (ITR).**

MIPS (IER or Instruction Execution Rate)

Processors are designed to execute instructions (OPCODES) that are in its inventory of functional activities. These instructions are processed at some average rate. That rate is quoted as MIPS (millions of instructions per processor second).

To express MIPS rates, IBM traditionally has used the term **instruction execution rate (IER)** instead. The capacity for IBM processors announced through the early 1980's was generally made in terms of relative IER. That is, the IER of the new processor was compared to the IER of a previously announced model. Then, the capacity of the new processor was expressed as an IER ratio relative to the older model.

In the early days of data processing, when processor design was simple, instruction execution rates correlated reasonably well with the ability of the processor to do work. Expressing capacity with instruction execution rates could then provide an adequate view of how one processor might perform relative to another. However, as processors got more complex in their internal design, and as user interactions got more varied and specialized, the usefulness of instruction execution rates as an expression of capacity began to diminish.

With today's high-performance processors, the actual MIPS rate achieved is extremely sensitive to the workload type being run, and its relationship to underlying processor design. LSPR workloads can be used to demonstrate this fact, since actual MIPS rates are frequently captured in the LSPR measurement process for IBM processors.

Figure 1 shows the ratio of MIP rates measured for several individual LSPR workloads on the same physical processor. Although the actual MIPS rates are not identified in this figure, you can see that, depending on workload, the ratio of MIPS for the same machine could exceed 2.

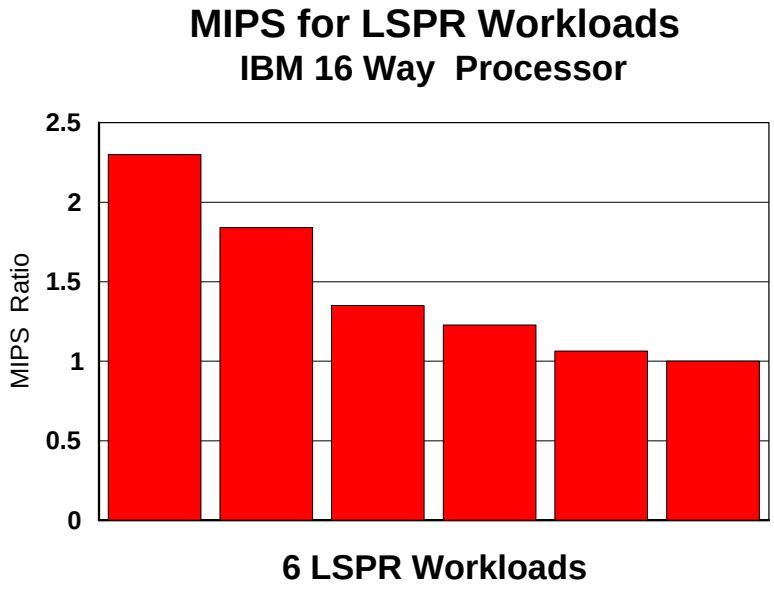

Figure 1. Relative MIP rates for LSPR workloads on a single processor

Generally, the workloads that generate the highest MIPS rates are batch, and the workloads that generate the lowest MIPS rates are on-line. This is true for all System/370 and System/390 architecture, and z/Architecture processors including IBM-compatible processors. The actual range from lowest to highest MIPS rate for LSPR workloads (or production workloads) on any given processor depends on that processor's design.

There are several high-level design factors on contemporary processors that prevent IER (or MIPS) from being meaningful as a capacity indicator.

- **Overlapped function**

  One of the techniques used to enhance the level of performance for large processors is to design high levels of overlap for the functions that a processor must perform to execute instructions. Various degrees of overlap are achieved in the instruction decode and execution process. The degree to which overlap is achieved is extremely dependent on the processor's specific design, and on the type of workload being processed. The greater the degree of overlap achieved the higher the effective instruction execution rate.

One of the design factors usually known about a processor is its basic cycle time. Two processors with the same cycle time are not necessarily equal in capacity because one may be able to accomplish more within a cycle than the other. Instruction efficiency relates the number of cycles an average instruction requires to execute. Instruction efficiency is a function of all the overlapped capability that can be realized by the processor.

Each different production workload environment tends to have its dominant instruction sets. Various processor designs can be sensitive to dominant instruction sets and instruction sequences in different ways. This relationship of design to dominant instruction sets has a significant influence on individual processor capacity, and therefore, on processor capacity relationships.

- **High-speed buffer (HSB) or Cache**

  All System/370, System/390, and z/Architecture high-end processors today have one or more high-speed buffers (HSBs) implemented in their design to enhance overall performance. By keeping data and instructions that are being referenced (or likely to be referenced) in the HSB, the effective time to access these items is reduced dramatically. The movement between the slower central storage and the HSB is managed automatically by the processor, being overlapped with normal instruction processing activity.

  The size and design of the HSB plays an important role in the ability of a machine to process a workload. Some workload types benefit more from certain HSB design implementations than do others. Workloads with the best buffer hit ratios will cause the processor to have higher MIPS rates. Workloads with lower buffer hit ratios will cause that processor to have lower MIPS rates. Processor capacity and processor capacity relationships, therefore, become very sensitive to the HSB implementations of the processors being considered.

- **N-way processing**

  Early data processing systems were primarily uniprocessors. Over time multiprocessors have evolved to the point where today there are as many as eighty tightly-coupled processors in a complex. System control program software is especially designed to manage multi-engine processors. Some portion of the instructions executed must go toward this N-way management function, as opposed to doing application work. Any use of instruction execution rates as a capacity indicator on these systems would include processing time that did not represent application work.

  There are also hardware performance considerations related to N-way processing. For most SCP operating environments, any work may be dispatched on any of the "N" engines, at any time. For most N-way designs, each engine has its own high speed buffer (HSB). The hardware must assure that any particular memory location that is changed is represented in only one engine's HSB at a time. Therefore, as work tends

to move around to different engines, so must any HSB-associated storage. When SCP dispatching decisions are frequent, as is typical with on-line workloads, this hardware N-way overhead will be the greatest.

- **Micro code**

  In the beginning, data processing systems were primarily hardware - based. Over time processor technology has evolved into extensive use of micro code to provide function. The advantage of micro code is that modifications to a design could be made without expensive hardware rework. As micro code flexibility and use grew, some functions normally performed by software were implemented more effectively in micro code. Usually, each micro coded software function becomes a new or extended instruction, more complex than typical OPCODEs, but doing the work that would normally be done with an entire routine in software. The use of these micro coded instructions has the tendency to lower the actual MIPS rate, while improving the processor's ability to do work.

Every workload (production or benchmark) has its own characteristics relative to how each of these hardware design features is exploited. For example, batch workloads tend to exploit N-way processors more efficiently than do on-line workloads, simply because of the difference in the rate of dispatching decisions that the operating system must make. On-line workloads tend to realize a greater performance benefit from larger and more sophisticated high speed buffer designs than do batch workloads. This is because the storage reference patterns of on-line workloads tend to be much more random.

Insight

Hardware design defines how a processor can perform. The software being exercised determines how a processor actually does perform.

The reason any particular processor performs as it does, lies in the interrelationship of the particular workload being run, to the underlying processor design.

Contemporary Use of MIPS

Various MIPS tables are used by organizations for the purpose of calculating relative capacity between processors. The individual MIPS values supplied for each processor are only loosely tied to the traditional meaning of the term. One perceived advantage in the use of MIPS tables is that the implied scale is easily recognized, that is, a MIPS rating for a processor provides an easy visual positioning of that machine on some grand scale of processor capability.

The use of MIPS tables produces a major problem when trying to understand relative processor capacity. The problem relates to the fact that different workload environments can have a significant effect on the way any particular processor design behaves. Therefore, the relative capacity of one processor to another will be very dependent on the type of work being run. As a result, ***it is***

often difficult to accurately position the processing capability of today's high-end processors with single-number tables.

The perception is that MIPS tables aid in understanding relative processor capacity across vendors, since these processor ratings are all tied to the same scale. However, just as IBM processor designs are extremely sensitive to the workload environment being run, so are those of the IBM compatible vendors. In fact, one can often see greater workload sensitivity when comparing processors from two different vendors, than when comparing two processors of the same vendor.

In today's world, it would be unwise to ignore the effects of different SCP and workload environments when making processor capacity comparisons. For this reason, IBM has chosen to provide capacity data in terms of work accomplished for a variety of workloads and SCP's, rather than in MIPS or instruction execution rates.

Workload Throughput Rates

Processors are purchased to do work rather than to do "MIPS". Given that the rate at which work is processed can be easily determined, it would seem natural that the best way to rate a processor is in terms of work units that it can do over time. To measure work done, two metrics have been defined. These metrics are external throughput rate (ETR) and internal throughput rate (ITR).

Assume that a benchmark workload is measured on each of two systems, with the results noted in Figure 2.

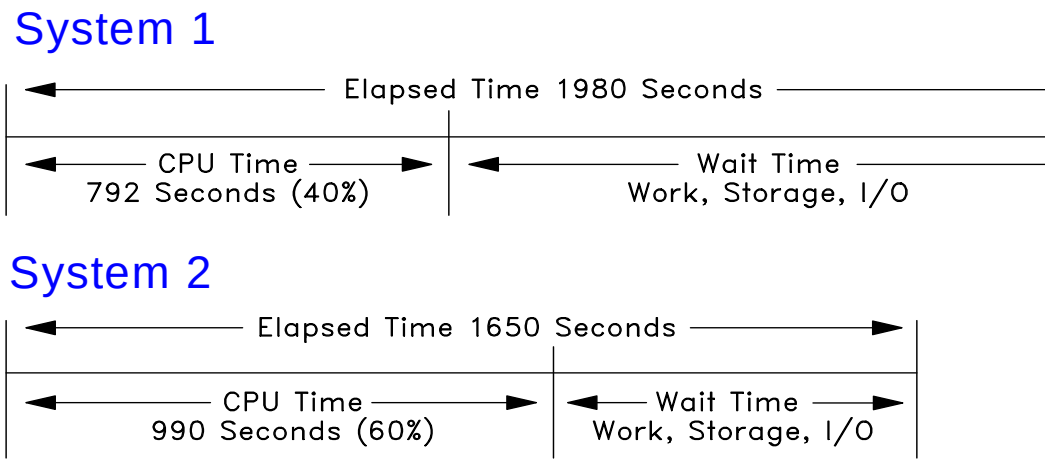

Figure 2. System capacity versus processor capacity

If the question being asked is:

"Which of the two systems is the better one for this workload?"

the correct answer is system number 2, because it processed the work in less elapsed time. ***External throughput rate (ETR), an elapsed time measure, focuses on system capacity.***

If the question being asked is:

"Which of the two systems has the better processor for this workload?"

the correct answer is system number 1, because it used less processor time to accomplish the same work. Because of the way that the question is posed here the focus must be changed to the processor itself. ***Internal throughput rate (ITR), a processor time measure, focuses on processor capacity.***

External Throughput Rate (ETR)

External throughput rate is computed as:

ETR = Units of Work / Elapsed Time

“Units of work" are normally expressed as jobs (or job-steps) for batch workloads, and as transactions or commands for on-line workloads (SCPs and most major software products have facilities to provide this information). To be useful, the "units of work" measured must represent a large and repeatable sample of the total workload, in order to best represent the average. Elapsed time is normally expressed in seconds.

ETR characterizes system capacity because it is an elapsed time measurement (system capacity encompasses the performance of the processor and all of its external resources, considered together). As such, ***ETR lends itself to the "system comparison methodology"***. This methodology requires the data processing system to be configured with all intended resources, including the processor, with appropriate amounts of central storage, expanded storage, channels, control units, I/O devices, TP network, and so on.

Once configured, the goal is to determine how much work the system, as a whole, can process over time. To do this, the system is loaded with the appropriate workload, until it cannot absorb work at any greater rate. The highest ETR achieved is the processing capability of the system.

When you make a system measurement of this type, all resources on the system are potential capacity inhibitors. If a resource other than the processor itself is, in fact, a capacity inhibitor, then it is likely that the processor will be running at something less than optimal utilization.

This **system comparison methodology** is a legitimate way to measure when the intent is to assess the capacity of the system as a whole. For on-line

systems, response time also becomes an important system related metric, as poor response times will inhibit a user's ability to do work. Therefore, system measurements for on-line work usually involve some type of response time criteria. If the response time criteria is not met, then it does not matter what ETR can be realized.

Internal Throughput Rate (ITR)

Internal throughput rate is computed as:

ITR = Units of Work / Processor Busy

As with ETR, "units of work" are normally expressed as jobs (or job-steps) for batch workloads, and as transactions or commands for on-line workloads (SCPs and most major software products have facilities to provide this information). To be useful, the "units of work" measured must represent a large and repeatable sample of the total workload, in order to best represent the average. Processor busy time is normally expressed in seconds.

For the purpose of computing an ITR, processor busy time should include all processing time to do work, including the operating system related overhead. On an N-way processor, processor busy time must represent the entire complex of engines as if it were a single resource. Therefore, processor busy time is the sum of the busy times for each individual engine, divided by the total number of engines. Since all processor time is included, "captured" and "UN-captured" time considerations are unnecessary.

ITR characterizes processor capacity, since it is a CPU busy time measurement. As such, ***ITR lends itself to the "processor comparison methodology"***. Because the LSPR's focus is on a single resource (the processor), you must modify the measurement approach from that used for a ***system*** comparison methodology.

To ensure that the processor is the primary point of focus, you must configure it with all necessary external resources (including central storage, expanded storage, channels, control units, I/O devices) in adequate quantities so that they do not become constraints. You need to avoid using processor cycles to manage external resource constraints in order to assure consistent and comparable measurement data across the spectrum of processors being tested.

There are many acceptance criteria for LSPR measurements that help assure that external resources are adequate. For example, internal response times should be sub second; if they are not, then there is some type of resource constraint that needs to be resolved. For various DASD device types, expected nominal service times are known. If the measured service times are high, then some type of queuing is occurring, indicating a constrained resource. When unexpected resource constraints are detected, they are fixed, and the measurement is redone.

Because the processor itself is also a resource which must be managed by the SCP, steps must be taken to ensure that excess queuing on it does not occur. The way to avoid this type of constraint is to make the measurements at preselected utilization levels that are less than 100%. Because the LSPR is designed to relate processor capacity, measurements must be made at reasonably high utilization, but without causing uncontrolled levels of processor queuing. Typically, LSPR measurements for on-line workloads are made at a utilization level of approximately 90%. Batch workloads are always measured with steady-state utilization's above 90%. Mixed workloads containing both an on-line and batch component are measured at utilizations near 99%.

One additional point needs to be made about processor utilization. Whenever two processors are to be compared for capacity purposes, they should both be viewed at the same loading point, or, in other words, at equal utilization. It is imprecise to assess relative capacity when one processor is running at low utilization and the other is running at high utilization. The LSPR methodology mandates that processor comparisons be made at equivalent utilization levels.

ITR/ETR Relationship

An ITR can be viewed as a special case ETR, that is, ***an ITR is the measured ETR normalized to full processor utilization***. Therefore, an alternate way to compute an ITR is:

ITR = ETR / Processor Utilization

To show that the arithmetic above works out to be the same as the previous ITR formula, consider the following. The formula for processor utilization is:

Processor Utilization = Processor Busy Time / Elapsed Time

Substituting in the above for *ETR* and for *Processor Utilization*, gives:

ITR = (Units of Work / Elapsed Time)/(Processor Busy Time / Elapsed Time)

You will see that the two **Elapsed Time** values factor out, giving the same formula as originally stated for ITR.

Insight

The sole purpose of computing an ITR is to normalize out the slightly unequal utilization that may be represented by the ETR, since measurement techniques cannot assure exactly equal utilization levels.

There is a reason to normalize ETRs when comparing processor capacity. If every benchmark measurement could be made at the **exact** same utilization, you could simply compare the two ETR values to determine relative processor capacity. However, measurement techniques seldom allow identical utilization levels to be achieved between runs.

Table 1 shows results from two online workload measurements made for the LSPR. The target utilization for these measurements was 90%. To achieve this target, the number of logged-on users is adjusted as necessary, so that we are within three percentage points of the target. As you can see, the measurements resulted in utilizations close to the target, but slightly different.

Table 1. **LSPR Measurement example for Online workload**.

	Processor A	**Processor B**	**Ratio**
Measured Data:			
Elapsed Seconds	720.54	720.32	
Processor Seconds	630.11	627.72	
Transaction count	727,736	1,273,150	
Calculated Data:			
Utilization (%)	87.45	87.14	
ETR	1001.0	1767.5	(1.77)
ITR	1154.9	2028.3	1.76

With the measured data, you can compute an ETR and an ITR for each processor. If you were to simply compare the two ETR values to determine relative capacity, the ratio would be flawed, because of this slightly unequal utilization. To make the results comparable to each other, we must normalize out the slightly unequal utilization values measured. It does not matter what we normalize to; for LSPR purposes, we chose to normalize the ETR to 100% utilization, and call it an ITR.

Throughput Rates Are Workload Unique

Both ITR values and ETR values are unique to the specific SCP and workload that was measured. ITRs are useful for determining the relative capacity between two processors running the exact same workload environment. ETRs are useful

for determining the relative capacity between two appropriately configured processing systems running the exact same workload environment. ***The absolute ITR and ETR values from one workload and SCP cannot be meaningfully compared to those of a different workload or SCP.*** Nor should absolute ITR and ETR values from a specific benchmark workload be compared to those of a production workload.

Table 2. **Example showing ITR values for 4 different processors**

Processor	**Workload 1**	**Workload 2**	**Workload 3**
Processor A	0.02506	0.901	41.69
Processor B	0.04835	2.091	102.16
Processor C	0.05305	2.334	118.23
Processor D	0.06296	2.713	137.36

Table 2 shows ITR values for four different processors for several LSPR workloads labeled workload 1, 2 and 3. It should be stated that the average unit-of-work is completely different between each of these workloads, and therefore the ITR scale is also unique for each workload.

Expressing Relative Capacity with ITRs

ITRs are useful for determining capacity relationships between processors for a given SCP and workload environment. This is done by dividing the ITR of one processor by the ITR of another to produce an ITR ratio (ITRR). For example, to determine the capacity of processor "B" relative to that of processor "A", use the formula:

ITR Ratio (or ITRR) = ITR for CPU-B / ITR for CPU-A

Note: ITR values used in this calculation must be for identical SCP and workload environments.

ITR values are intended to be used for calculating relative processor capacity. The benefit of using LSPR ITR data is that you are working with workload sensitive data. As such you will have a more reliable and accurate view of relative capacity than can be provided by any MIPS table, or any other single-number per processor source.

Table 3. **Example showing ITRR values relative to Processor A**.

Processor	Workload 1	Workload 2	Workload 3
Processor A	1.00	1.00	1.00
Processor B	1.93	2.32	2.45
Processor C	2.12	2.59	2.84
Processor D	2.47	3.01	3.30

Table 3 shows ITR ratios developed using the absolute ITR values in Table 2. By representing the capacity of a set of processors relative to Processor A, we can see how relative capacity varies with the different workload types.

Chapter 3. Workload Environments

Data processing systems are designed to provide a range of services, using a wide variety of software products and application programs. On-line systems provide services directly to the end-user, while batch systems offer deferred services. In most cases, production systems offer a combination of many different types of services.

Assuring Representativeness for a Benchmark

In order for any benchmark workload to be useful, it is essential that its instruction paths and the storage reference patterns be representative of actual production work. Because of the complex interrelationships between software and processor design, it is impossible to ascertain the instruction paths and the storage reference patterns for a production workload. Therefore, the only way to assure that a benchmark is truly representative of production work is to use actual production software and activities.

For this reason, non-representative workloads (including "kernels") are not considered useful as benchmarks. It would only be by sheer chance that capacity relationships derived from a non-representative benchmark would match up with the capacity realized when the actual production workload is moved to another processor.

Insight

There is no reasonable way to construct a benchmark that simulates instruction paths and storage reference patterns typical of a production workload without using actual production software and activities.

Software

Many types of software are used to take advantage of System/370, System/390 architecture, and z/Architecture processors.

System Control Program (SCP)

Basic processor support software is known as the system control program, or SCP. The SCP provides the routines to manage the processor and external resources such as storage and I/O devices. The SCP controls the dispatching of work and the allocation of resources on the processor complex.

Various software products are also associated with the SCP, including JES, RACF, VTAM™, TSO, and CMS. Performance monitors, which may also be associated with the SCP, are discussed below.

Each LSPR benchmark workload includes the use of the appropriate SCP, and associated components as applicable. SCPs used for the various LSPR benchmarks measured include z/OS, OS/390, MVS, z/VM, VM, VSE, and Linux on zSeries.

Subsystems

Most production systems include the use of one or more major application subsystems that are available for System/370, System/390 architecture, and z/Architecture processors. Examples of such subsystems include CICS, DB2, and IMS. Each of these subsystems is represented in one or more of the LSPR benchmark workloads.

Application Servers

To facilitate the rapid deployment of e-business applications, many production systems are increasing their exploitation of application server software. To reflect this trend, several LSPR benchmarks utilize the WebSphere Application Server (WAS).

Other Program Offerings

Products such as language compilers, linkage editors, and commercial or engineering/scientific programs are typically used by installations in either on-line or batch mode. LSPR measurements include the use of such program products where appropriate.

Application Software

Application software is the custom programming that must be done to make system software perform the specific functions necessary for a business enterprise. LSPR workloads include typical installation-written database application programs for use under CICS, DB2, and IMS (usually written by professional programmers), and typical end-user programs for use in batch, TSO and CMS (which may or may not be professionally written).

Performance Monitors

Software performance monitors are available, both at the SCP level, and at the application subsystem level. It is felt that the LSPR benchmark measurements should use the same performance monitor software as is commonly used in production environments. Doing so not only helps to assure that LSPR workload instruction paths are representative, but also provides a common basis for reporting detailed measurement results.

From the SCP standpoint, both RMF and SMF are used with MVS, OS/390, and z/OS, and the VM Monitor is used with z/VM. Subsystem-specific monitors are also used where applicable, such as the CICS Monitor, IMS/VS Performance Analysis Reporting System (IMS PARS), or DB2 Performance Monitor, for the MVS, OS/390 and z/OS environments.

Workload Content

Another aspect of representativeness is how the actual work is presented to the system. These are the activities (jobstreams or end-user commands) that cause the various forms of software discussed above to be exploited. The two basic ways that work enters a system is via batch submission of jobs, and online entries by an end-user at a terminal or client at a workstation.

Every individual unit of work in a production workload, whether it be a specific job or application, has its own characteristics, and therefore its own unique relationship to the hardware design of the processors being measured. Production workloads normally consist of a large and diverse cross-section of individual jobs and applications, all being managed by the operating system and related software products.

In order for any benchmark to serve a useful purpose, it must represent a rich cross-section of production workload activity. Benchmarks that focus on only one, or just a few individual types of work are very unlikely to provide the proper capacity perspective for an entire production workload.

Batch

Work associated with batch is presented to the system as jobs, read in through a job queue. Initiators select these jobs on a priority and class basis, and guide their progress through the system, obtaining all the resources required to complete the work. Typical batch work includes compile, link-edit, execution of batch oriented production applications, and utility programs (usually involving some form of data manipulation).

Batch benchmarks are generally measured as start-to-finish workloads. The job queue is loaded with a predetermined number of copies of the jobstream to assure a reasonable measurement window. Enough initiators are activated to allow steady-state processor utilization (the period when all initiators are active) to be as close to 100% as possible. The measurement starts when the job queue is released, and finishes when the last job is completed.

Job (or job step) count and processor busy time are combined to compute an ITR value (see formula 2 given earlier).

On-Line

Work associated with on-line systems is generated by end-users sitting at terminals or clients sitting at workstations, entering transactions or commands. Two different types of activity are represented by on-line workloads:

1. **Structured Work**

   End-user transactions are directed toward the manipulation of one or more databases, or some other form of organized data. This type of on-line workload generally consists of a limited set of fixed transaction types that can be requested by end-users, each relating to the purpose of manipulating the relevant data.

2. **Unstructured (Ad hoc) Work**

This type of on-line work is the effect of providing a wide array of data processing capabilities to end-users, such as program entry and/or testing, file input and editing, use of decision support products, office support and management, on-line queries, and so on.

End-user on-line interactions have attributes relating to their arrival, such as average think time and think time distribution. It is essential that, not only the content of the commands be representative, but that inter-arrival times be representative also.

To benchmark on-line systems, a terminal or client simulator must be used to generate end-user activity. User transactions that comprise the workload are organized into scripts, with each script representing a set of coordinated activities. For each active terminal or client, the simulator assigns a script. From that script, it selects each end-user input in sequence, applies a think time (using representative think time distribution tables), sends the input to the system, and waits for the response before starting on the next input. Terminal and client simulators normally continue to submit commands as a never-ending process.

On-line systems are generally measured as "steady-state" systems, with the processor running at some predetermined target utilization. To reach the desired state, an adequate number of users (terminals or clients) are connected, each immediately starting to execute a script. After the final terminal or client has logged on, and the system has stabilized, a measurement is taken over an elapsed period that is considered a repeatable window of work.

Transaction count and processor busy time are combined to compute an ITR value for on-line workloads (see formula 2 given earlier). Alternatively, external transaction rate (ETR) and processor utilization can be used (see formula 3 given earlier).

Data Considerations

Data processing systems, as the name implies, are designed to manage data. A benchmark workload cannot afford to ignore this aspect of production work. Data exists in many forms and formats. Data files (data sets) and databases are used by the LSPR workloads, as appropriate.

There are two special considerations about data, when performing benchmark measurements, if repeatability is to be assured:

- Data files (datasets) and databases used by a benchmark workload must be restored to their pristine state before each measurement.
- Data files (datasets) and databases must be used in such a way that changes made by the benchmark scripts will not cause the performance of the processor to change significantly over time.

Obviously, data files and databases on a production system do not remain constant. As they get updated and extended, processor (and system) performance can be affected. However, over time, a steady-state data condition is normally achieved.

A benchmark does not have the luxury of waiting for this steady-state data condition to occur. By assuring that the benchmark data is in the same state for each measurement, we know that the processor performance data obtained will be comparable to other measurements made the same way.

LSPR Workload Categories

Introduction

Historically, LSPR workload capacity curves (primitives and mixes) have had application names or been identified by a *software* characteristic. For example, past workload names have included CICS, IMS, OLTP-T, CB-L, LoIO-mix and TI-mix. However, capacity performance has always been more closely associated with how a workload uses and interacts with a particular processor *hardware* design. With the availability of CPU MF (SMF 113) data on z10, the ability to gain insight into the interaction of workload and hardware design in production workloads has arrived. The knowledge gained is still evolving, but the first step in the process is to produce LSPR workload capacity curves based on the underlying hardware sensitivities. Thus the LSPR introduces three new workload capacity categories which replace all prior primitives and mixes.

Fundamental Components of Workload Capacity Performance

Workload capacity performance is sensitive to three major factors: instruction path length, instruction complexity, and memory hierarchy. Let us examine each of these three.

Instruction Path Length

A transaction or job will need to execute a set of instructions to complete its task. These instructions are composed of various paths through the operating system, subsystems and application. The total count of instructions executed across these software components is referred to as the transaction or job path length. Clearly, the path length will be different for each transaction or job depending on the complexity of the task(s) that must be performed. For a particular transaction or job, the application path length tends to stay the same presuming the transaction or job is asked to perform the same task each time. However, the path length associated with the operating system or subsystem may vary based on a number of factors including: a) competition with other tasks in the system for shared resources – as the total number of tasks grows, more instructions are needed to manage the resources; b) the Nway (number of logical processors) of the image or LPAR – as the number of logical processors grows, more instructions are needed to manage resources serialized by latches and locks.

Instruction Complexity

The type of instructions and the sequence in which they are executed will interact with the design of a micro-processor to affect a performance component we can define as “instruction complexity.” There are many design alternatives that affect this component such as: cycle time (GHz), instruction architecture, pipeline, superscalar, out-of-order execution, branch prediction and others. As workloads are moved between micro-processors with different designs, performance will likely vary. However, once on a processor this component tends to be quite similar across all models of that processor.

Memory Hierarchy and “Nest”

The memory hierarchy of a processor generally refers to the caches (previously referred to as HSB or High Speed Buffer), data buses, and memory arrays that stage the instructions and data needed to be executed on the micro-processor to complete a transaction or job. There are many design alternatives that affect this component such as: cache size, latencies (sensitive to distance from the micro-processor), number of levels, MESI (management) protocol, controllers, switches, number and bandwidth of data buses and others. Some of the cache(s) are “private” to the micro-processor which means only that micro-processor may access them. Other cache(s) are shared by multiple micro-processors. We will define the term memory “nest” for a System z processor to refer to the shared caches and memory along with the data buses that interconnect them.

Workload capacity performance will be quite sensitive to how deep into the memory hierarchy the processor must go to retrieve the workload’s instructions and data for execution. Best performance occurs when the instructions and data are found in the cache(s) nearest the processor so that little time is spent waiting prior to execution; as instructions and data must be retrieved from farther out in the hierarchy, the processor spends more time waiting for their arrival.

As workloads are moved between processors with different memory hierarchy designs, performance will vary as the average time to retrieve instructions and data from within the memory hierarchy will vary. Additionally, once on a processor this component will continue to vary significantly as the location of a workload’s instructions and data within the memory hierarchy is affected by many factors including: locality of reference, IO rate, competition from other applications and/or LPARs, and more.

Relative Nest Intensity

The most performance sensitive area of the memory hierarchy is the activity to the memory nest, namely, the distribution of activity to the shared caches and

memory. We introduce a new term, “Relative Nest Intensity (RNI)” to indicate the level of activity to this part of the memory hierarchy. Using data from CPU MF, the RNI of the workload running in an LPAR may be calculated. The higher the RNI, the deeper into the memory hierarchy the processor must go to retrieve the instructions and data for that workload.

Many factors influence the performance of a workload. However, for the most part what these factors are influencing is the RNI of the workload. It is the interaction of all these factors that result in a net RNI for the workload which in turn directly relates to the performance of the workload.

The traditional factors that have been used to categorize workloads in the past are listed along with their RNI tendency in figure 3.

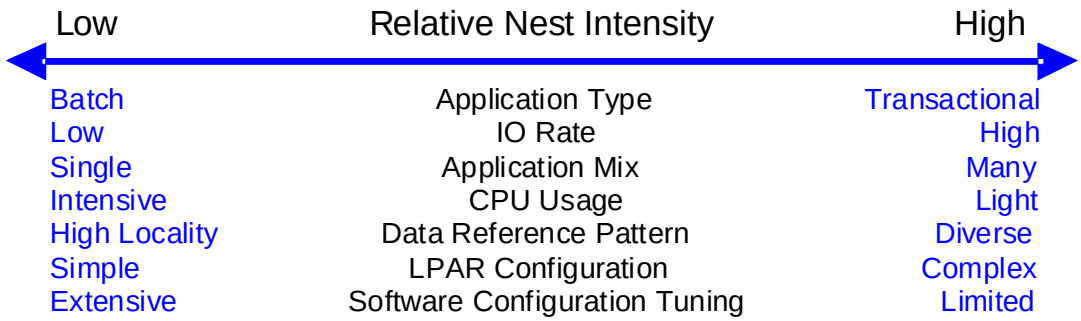

Figure 3. **Relative Nest Intensity Tendency**

It should be emphasized that these are simply tendencies and not absolutes. For example, a workload may have a low IO rate, intensive CPU use, and a high locality of reference – all factors that suggest a low RNI. But, what if it is competing with many other applications within the same LPAR and many other LPARs on the processor which tend to push it toward a higher RNI? It is the net effect of the interaction of all these factors that determines the RNI of the workload which in turn greatly influences its performance.

Note that there is little one can do to affect most of these factors. An application type is whatever is necessary to do the job. Data reference pattern and CPU usage tend to be inherent in the nature of the application. LPAR configuration and application mix are mostly a function of what needs to be supported on a system. IO rate can be influenced somewhat through buffer pool tuning.

However, one factor that can be affected, **software configuration tuning,** is often overlooked but can have a direct impact on RNI. Here we refer to the number of address spaces (such as CICS AORs or batch initiators) that are needed to support a workload. This factor has always existed but its sensitivity is higher with today’s high frequency microprocessors. Spreading the same workload over a larger number of address spaces than necessary can raise a workload’s RNI as the working set of instructions and data from each address

space increases the competition for the processor caches. Tuning to reduce the number of simultaneously active address spaces to the proper number needed to support a workload can reduce RNI and improve performance. In the LSPR, we tune the number of address spaces for each processor type and Nway configuration to be consistent with what is needed to support the workload. Thus, the LSPR workload capacity ratios reflect a presumed level of software configuration tuning. This suggests that re-tuning the software configuration of a production workload as it moves to a bigger or faster processor may be needed to achieve the published LSPR ratios.

Calculating Relative Nest Intensity

The RNI of a workload may be calculated using CPU MF data. For z10, three factors are used:

L2LP: percentage of L1 misses sourced from the local book L2 cache,
L2RP: percentage of L1 misses sourced from a remote book L2 cache, and
MEMP: percentage of L1 misses sourced from memory.

These percentages are multiplied by weighting factors and the result divided by 100. The formula for z10 is:

z10 RNI=(1.0xL2LP+2.4xL2RP+7.5xMEMP)/100.

Tools available from IBM (zPCR) and several vendors can extract these factors from CPU MF data. For z196 the CPU MF factors needed are:

L3P: percentage of L1 misses sourced from the shared chip-level L3 cache,
L4LP: percentage of L1 misses sourced from the local book L4 cache,
L4RP: percentage of L1 misses sourced from a remote book L4 cache,
MEMP (percentage of L1 misses sourced from memory).

The formula for z196 is:

z196 RNI=1.6x(0.4xL3P+1.0xL4LP+2.4xL4RP+7.5xMEMP)/100

Note these formulas may change in the future.

LSPR Workload Categories Based on Relative Nest Intensity

As discussed above, a workload's relative nest intensity is the most influential factor that determines workload performance. Other more traditional factors such as application type or IO rate have RNI tendencies, but it is the net RNI of the workload that is the underlying factor in determining the workload's capacity performance. With this in mind, the LSPR now runs various combinations of

former workload primitives such as CICS, DB2, IMS, OSAM, VSAM, WebSphere, COBOL and utilities to produce capacity curves that span the typical range of RNI. The three new workload categories represented in the LSPR tables are described below.

LOW (relative nest intensity): A workload category representing light use of the memory hierarchy. This would be similar to past high scaling primitives.

AVERAGE (relative nest intensity): A workload category representing average use of the memory hierarchy. This would be similar to the past LoIO-mix workload and is expected to represent the majority of production workloads.

HIGH (relative nest intensity): A workload category representing heavy use of the memory hierarchy. This would be similar to the past DI-mix workload.

LSPR Workload Primitives

Various combinations of LSPR workload "primitives" have been and continue to be run to create the capacity ratios given in the LSPR tables. Each individual LSPR workload is designed to focus on a major type of activity, such as interactive, on-line database, or batch. The LSPR does not focus on individual pieces of work such as a specific job or application. Instead, each LSPR workload includes a broad mix of activity related to that workload type. Focusing on a broad mix can help assure that resulting capacity comparisons are not skewed.

The LSPR workload suite is updated periodically to reflect changing production environments. High-level workload descriptions are provided below.

z/OS and OS/390

OLTP-T (formerly IMS) - Traditional On-line Workload

The OLTP-T workload consists of moderate to heavy IMS transactions from DLI applications covering diverse business functions, including order entry, stock control, inventory tracking, production specification, hotel reservations, banking, and teller system. These applications all make use of IMS functions such as logging and recovery. The workload contains sets of 12 (17 for OS/390 Version 1 Release 1 and earlier) unique transactions, each of which has different transaction names and IDs, and uses different databases. Conversational and wait-for-input transactions are included in the workload.

The number of copies of the workload and the number of Message Processing Regions (MPRs) configured is adjusted to ensure that the IMS subsystem is processing smoothly, with no unnecessary points of contention. No Batch Message Processing regions (BMPs) are run. IMS address spaces are non-swappable.

This IMS workload accesses both VSAM and OSAM databases, with VSAM indexes (primary and secondary). DLI HDAM and HIDAM access methods are used. The workload has a moderate I/O load, and data in memory is not implemented for the DLI databases.

Measurements are made with z/OS, OS/390, DFSMS, JES2, RMF, VTAM, and IMS/ESA. IMS coat-tailing (enabling reuse of a module already in storage) is not used; since this activity is so sensitive to processor utilization, it could cause distortion when comparing ITRs between faster and slower processors. Beginning with OS/390 Version 1 Release 1, measurements were done with one or more control regions. The number of data base copies, MPR's, and control regions (within the limits of granularity) are scaled with the processing power of a particular machine in-order to assure equal and normalized tuning. Performance data collected consists of IMS PARS, and the usual SMF data, including type 72 records (workload data), and RMF data.

OLTP-W -Web-enabled On-line Workload

The OLTP-W workload reflects a production environment that has web-enabled access to a traditional data base. For the LSPR, this has been accomplished by placing a WebSphere front-end to connect to the LSPR CICS/DB2 workload (described below).

The J2EE application for legacy CICS transactions was created using the CICS Transaction Gateway (CTG) external call interface (ECI) connector enabled in a J2EE server in WebSphere for z/OS Version 5.1. The application uses the J2EE architected Common Client Interface (CCI). Clients access WebSphere services using the HTTP Transport Handlers. Then, the appropriate servlet is run through the webcontainer, which calls EJB's in the EJB Container. Using the CTG External Call Interface (ECI) CICS is called to invoke DB2 to access the database and obtain the information for the client.

For a description of the CICS and DB2 components of this workload, please see the CICS/DB2 workload description further below.

WASDB - WebSphere Application Server and Data Base

The WASDB workload reflects a new e-business production environment that uses WebSphere applications and a DB2 data base all running in z/OS.

WASDB is a collection of Java classes, Java Servlets, Java Server Pages and Enterprise Java Beans integrated into a single application. It is designed to emulate an online brokerage firm. WASDB was developed using the IBM VisualAge™ for Java and WebSphere Studio tools. Each of the components is written to open Web and Java Enterprise APIs, making the WASDB application portable across J2EE-compliant application servers.
The WASDB application allows a user, typically using a web browser, to perform the
following actions:
· Register to create a user profile, user ID/password and initial account balance.
· Login to validate an already registered user.
· Browse current stock price for a ticker symbol.
· Purchase shares.
· Sell shares from holdings.
· Browse portfolio.
· Logout to terminate the user's active interval.
· Browse and update user account information.

CB-L (formerly CBW2)-Commercial Batch Long Job Steps

The CB-L workload is a commercial batch job stream reflective of large batch jobs with fairly heavy CPU processing. The job stream consists of 1 or more copies of a set of batch jobs. Each copy consists of 18 jobs, with 105 job steps. These jobs are more resource intensive than jobs in the CB-S workload (discussed below), use more current software, and exploit ESA features. The work done by these jobs includes various combinations of C, COBOL, FORTRAN, and PL/I compile, link-edit, and execute steps. Sorting, DFSMS utilities (e.g. dump/restore and IEBCOPY), VSAM and DB2 utilities, SQL processing, SLR processing, GDDM™ graphics, and FORTRAN engineering/scientific subroutine library processing are also included. Compared to CB-S, there is much greater use of JES processing, with more JCL statements processed and more lines of output spooled to the SYSOUT and HOLD queues. This workload is heavily DB2 oriented with about half of the processing time performing DB2 related functions.

Measurements are made with z/OS, OS/390, DFSMS, JES2, RMF, and RACF. C/370, COBOL II, DB2, DFSORT, FORTRAN II, GDDM, PL/I, and SLR software are also used by the job stream. Access methods include DB2, VSAM, and QSAM. SMS is used to manage all data. Performance data collected consists of the usual SMF data, including type 30 records (workload data), and RMF data.

The CB-L job stream contains sufficient copies of the job set to assure a reasonable measurement period, and the job queue is pre loaded. Enough initiators are activated to ensure a high steady-state utilization level of 90% or greater. The number of initiators is generally scaled with processing power to achieve comparable tuning across different machines. The measurement is started when the job queue is released, and ended as the last job completes.

Each copy of the job set uses its own datasets, but jobs within the job set share data.

ODE-B - On Demand Environment - Batch

The ODE-B workload reflects the billing process used in the telecommunications industry. This is a multi-step approach which includes the initial processing of Call Detail Records (CDR), the calculation of the telephone fees, and the insertion of the created telephone bills in a database. The CDRs contain the details of the telephone calls such as the source and target numbers along with the time and the duration of the call. The CDRs are stored in flat files within a zFS file system. A feeder application reads the CDRs from the files, converts them into XML format and sends them to a queue. An analyzer application reads the messages from the queue and performs analysis on the data. During the analysis further information is retrieved from the relational database, and the same database is subsequently updated with the newly created telephone bill and new records for each call. The feeder and the analyzer applications are implemented as enterprise java beans (EJB) in IBM WebSphere Application Server for z/OS. Using the concept of multi-servant regions, which is unique to the z/OS implementation of WebSphere Application Server, the threads of the feeder and the analyzer applications are distributed over several java virtual machines (JVM). The WebSphere internal queuing engine is used as the queue for the message transport between the feeder and analyzer.

CB-J - JavaBatch

The JavaBatch workload reflects the batch production environment of a clearing bank
that uses a collection of java classes working on a DB2 database and a set of flat files in z/OS. JavaBatch is a native, standard Java application that can be run standalone on a single JVM (Java Virtual Machine) or in parallel to itself on multiple JVMs.
Each of the parallel applications instances can be tuned separately.
All parallel applications are working on the same set of flat files and database tables.
The JavaBatch application is based on a Java-JDBC-framework
from an external banking software vendor and has been enhanced and adapted using the Websphere Application Developer tool.
Various properties such as number of banks, number of accounts,
and more can be adapted for the specific runtime environment. These are kept in a special properties file, keeping the java application unchanged.
The JavaBatch application allows a user to perform the following activities:
- initialize the working database
- create a set of flat files, each containing several hundreds to thousands of payments

- read the flat files, perform various syntax-checks and validation for each payment
  and store the payments to the working database
- read the payments from the database and route them to destination bank's flat files

CICS/DB2 - On-line Workload for pre-z/OS version 1 release 4

The CICS/DB2 workload is an LSPR workload that was designed to represent clients' daily business by simulating the placement of orders and delivery of products, as well as business function like supply and demand management, client demographics and item selling hit list information. The workload consists of ten unique transactions.

CICS is used as a transaction monitor system. It provides both an API for designing the dialogue panels and parameters to drive the interface to the DB2 database. The interface between the two subsystems is fully supported by S/390 and exploits N-Way designs. CICS functions like dynamic workload gathering and function shipping are not exploited in this workload. The CICS implementation uses an MRO model, which is managed by CP/SM. The number of AOR (Address Owning Region) and TOR (Terminal Owning Region) used, depends on the number of engines of the processor under test. The ratio between TOR and AOR is 1:3. The utilization of the TOR and the AOR regions is kept under 60%.

The application database is implemented in a DB2 subsystem. One of the major design efforts was to achieve a read-to-write ratio exhibited by OLTP clients. Several data center surveys indicate an average read-to-write ratio to be in the range of 4:1 - 6:1. The read-to-write ratio is an indication of how much of the accessed data are changed as well. For this CICS/DB2 workload implemented on a S/390 or z/Architecture system and using DB2 as database system, an approximation of the read-to-write ratio is the ratio of SQL statements performing 'read' operation, like select, fetch, open cursor to the 'write' SQL statements, like insert, update, delete.

To reduce the number of database locks and the inter system communication required for each database update and to preserve local buffer coherency in data sharing environments, DB2 type 2 indexes have been used. Additionally, row-level-locking has been introduced for some tables. Each table and index is buffered in separate buffer pools for easy sizing and control.

Linux™ on zSeries

WASDB/L - WebSphere Application Server and Data Base under Linux on zSeries

The WASDB/L workload reflects an e-business environment where a full function application is being run under Linux on zSeries in an LPAR partition. For LSPR this was accomplished by taking the WASDB workload (described above under z/OS), and converting it to run both application and data base server in a single

Linux on zSeries image. The WASDB/L workload is basically the same as the WASDB workload on z/OS with the exception of being enabled for Linux on zSeries. . See the 'WASDB - WebSphere Application Serving and Data Base' section for a detailed description.

z/VM

WASDB/LVm - many Linux on zSeries guests under z/VM running WebSphere Application Server and Data Base

The WASDB/LVm workload reflects a server consolidation environment where each server is running a full function application. For LSPR this was accomplished by taking the WASDB workload (described above under z/OS), and then replicating the Linux- guest a number of times based on the N-way of the processor. Guest pair activity was then adjusted to achieve a constant processor utilization for each N-way. Thus the ratios between processors of equal N-way are based on the throughput per guest rather than the number of guests.

LSPR Measurement Methodology

Each of the LSPR workloads is run according to a specific set of rules, to assure that results can be compared with other measurements of the same workload. Neither changes to setup or operation, nor unique tuning activities are done to favor any processor. Some of the measurement methodology concerns include:

- Assure adequate configuration (storage, channels, DASD)
- Distribution of system data
- Distribution of program libraries
- Distribution of data files (datasets) and databases
- Files and databases restored to pristine state
- Logon end-users (terminals or clients), or load job queue
- Determine measurement period to obtain a repeatable sample of work
- Adjust activity to realize target processor utilization level
- Assure steady-state has been achieved
- Capture appropriate performance monitor data
- Capture operator console logs
- Verify that no hardware errors occurred
- Verify measurement data against acceptance criteria
- Construct detailed measurement reports

When a suitable testing environment is not available, analysis is used to address these methodology concerns.

As stated earlier, on-line workloads require some type of terminal or client simulator to generate the workload of an end-user community. There are a variety of products available that can serve this purpose, including IBM Teleprocessing Network Simulator (TPNS). Products like TPNS generally run out-board on a stand-alone processor that is connected to the system being benchmarked with normal teleprocessing hardware. Alternatively, TPNS could be run on the host processor (in-board) along with the benchmark workload itself. LSPR will choose between an in-board and out-board network simulator based on the functionality required and the processing overhead associated with the simulation. Generally, out-board network simulators are used.

Chapter 4. Using LSPR Data

The purpose of the LSPR is to provide relative capacity data for System z Architecture processors on which a reliable capacity planning exercise may be based. As described in Chapter 3, Workload Environments, capacity ratios for a variety of workload environments will be presented. Since each SCP/workload combination may have characteristics that react differently with the design of a particular processor, the LSPR ratios offer insight into the variability that may be experienced by different production workloads. In this chapter, the background and suggested methodologies for using LSPR data are discussed. Note that the examples are drawn from the z/OS 1.4 data tables, but the underlying techniques are applicable using the other SCP tables as well.

LSPR Data Sources

To maximize its usefulness, the LSPR includes as many System z processors as possible. Understanding capacity for older processors is important because they are the ones being migrated from. Understanding capacity for newer processors is important because they are the ones being migrated to. Without accurate LSPR data on both, it would not be possible to develop reliable relative processor capacity expectations.

Insight

All LSPR data is based on measurement and analysis.

All LSPR data is based on measurement and analysis.

Measured Data

LSPR data reflects measurements and projections. LSPR data for IBM processors is generally made available at announce time. IBM's announcement claims are based on LSPR measurements.

- **Projections**

  Primary models of each processor series are always measured for the LSPR. It is not practical, however, for IBM to allocate the necessary resource to measure every individual processor model announced. For those that are not measured, ITR numbers are projected. A projected ITR is based on the known processor internal design, and on the known characteristics of the subject workload on similarly designed processor models. Projections have been shown to have accuracy comparable to that of measurements, where subsequent testing has occurred.

- **Estimates**

Many older processor models that have been measured in the past, are no longer accessible for testing. It is often desirable to maintain these processors in LSPR tables, for those SCPs that are supported. Therefore, as the LSPR moves ahead to more current software levels, ITRs for these processors must be estimated. Estimates are based on known deltas between the older software and the current software on similarly designed processor models.

As capacities of today's processors grow ever larger, it becomes more difficult to commit the time, and/or the external resources necessary for full, unconstrained LSPR measurements. Experience has shown that there are estimation techniques, based on making reduced resource measurements and analysis, that can yield reliable results.

Relating LSPR Data to MIPS

IBM views MIPS tables as an often imprecise way to express processor capacity. The major problem is that most commonly accepted MIPS tables are insensitive to the SCP and workload environment being run. We must recognize, however, that MIPS is still a commonly used metric in the industry.

When MIPS is the processor capacity metric of choice, LSPR ITR data can still be useful to relate capacity for a potential new machine. If you want to assume that your currently installed processor represents some MIPS value to you, you can directly apply an LSPR ITR ratio for the appropriate workload to assess what a new processor's MIPS value would be. The formula is:

Expected CPU-B MIPS = CPU-A MIPS * (ITR for CPU-B / ITR for CPU-A)

When using the LSPR to compute the MIPS expectation for a new processor, the result will not necessarily line up with those in currently accepted MIPS tables. This is because our MIPS calculation is based on workload sensitive LSPR data. Published MIPS tables do not consider any such SCP or workload sensitivity.

A processor MIPS number determined by the above calculation simply states that, if you believe that your current processor is capable of "X" MIPS, then a new processor should be capable of "Y" MIPS, for the workload environment assumed. You are simply assigning a relative MIPS value for the new processor based on LSPR measurement experience. It does not suggest that IBM rates any processor with MIPS numbers.

Resource Constrained Environments

As stated previously, LSPR measurements are made on processor configurations designed to prevent any significant external constraints. The reason is that we want to represent every processor in its best light. This is the

only reasonable way to assure that every processor contained in the LSPR is fairly represented. It is impossible to produce globally useful capacity tables that represent processors in various constrained situations, since there are so many different resource constraint scenarios that could exist.

Insight

LSPR data can be used to assess relative processor capacity for production workloads that are resource constrained.

Production workloads may, of course, incur some types of external constraint. Often these constraints are not terribly significant, while, in other cases they may be. Normally, an external constraint can be relieved by installing more of a resource, such as central storage, expanded storage, channels, controllers, or I/O devices. Generally, it is a price/performance tradeoff, whether to use processor cycles to manage the constraint, or to purchase more of the constrained resource.

LSPR data is useful for assessing capacity for a new processor, even if the current processor has resource constraints. If we assume that the current level of constraint will remain the same on a planned new processor, then the capacity relationship established from the LSPR will apply. If we expect to relieve the constraint when the new processor is installed, then the capacity relationship established from the LSPR is conservative, and should be elevated (this is often the case when moving to more advanced technology hardware). If we expect to incur additional constraint on the planned new processor, then the capacity relationship established from the LSPR is overstated, and should be lowered.

It is beyond the purpose of the LSPR to provide data to evaluate the many various resource constraint scenarios possible. From time to time, IBM does publish technical bulletins containing "case study" information for this purpose. To assess the cost or value of reducing any type of resource constraint, documentation outside the LSPR will need to be consulted. In the absence of such documentation, a study will be required.

New Function

Over time, new function will appear in hardware and/or software. Often a new function is directed at minimizing or eliminating resource constraints. One such example is ESA's ability to exploit data-in-memory (DIM) in various ways. Whatever the purpose, there is always an interest in any performance or capacity benefit related to the exploitation of new function.

The LSPR's stated purpose is to compare processor capacity when running the same software. It is beyond the purpose of the LSPR to provide data to evaluate new function. As function is introduced, IBM will normally publish technical

documents showing the benefits related to exploiting such function. As new function gets accepted in the data processing community, LSPR workloads may be updated to reflect such activity.

It is also beyond the purpose of the LSPR to assess the effects of a software migration. The LSPR ratios in a given table are all at the same software level. While you may be running other levels of SCP or subsystems than the LSPR, most of the software performance differences cancel out when you do a hardware only change. This will result in ratios like the LSPR. Generally, there are technical bulletins that describe software migration performance effects. A combined software and hardware migration should be approached as a two step process that multiplies the hardware ratio (i.e., LSPR) by the appropriate software ratio (from the technical bulletins).

Relating Production Workloads to LSPR Workloads

Historically, there have been a number of techniques used to match production workloads to LSPR workloads such as a) application name (a customer running CICS would use the CICS LSPR workload), b) application type (create a mix of the LSPR online and batch workloads), c) IO rate (low IO rates used a mix of the low IO rate LSPR workloads). However, as discussed in the “LSPR Workload Categories” section, the underlying performance sensitive factor is how a workload interacts with the processor hardware. These past techniques were simply trying to approximate the hardware characteristics that were not available through software performance reporting tools. Beginning with the z10 processor, the hardware characteristics can now be measured using CPU MF (SMF 113) COUNTERS data. Thus, the opportunity exists to be able to match a production workload to an LSPR workload category via these hardware characteristics (see the “LSPR Workload Categories” section for a discussion about RNI – Relative Nest Intensity).

The AVERAGE RNI LSPR workload is intended to match the majority of customer workloads. When no other data is available, it should be used for a capacity analysis.

DASD IO rate has been used for many years to separate workloads into two categories: those whose DASD IO per MSU (adjusted) is <30 (or DASD IO per PCI <5) and those higher than these values. The majority of production workloads fell into the “low IO” category and a LoIO-mix workload was used to represent them. Using the same IO test, these workloads would now use the AVERAGE RNI LSPR workload. Workloads with higher IO rates may use the HIGH RNI workload or the AVG-HIGH RNI workload that is included with zPCR.

For z10 and newer processors, the CPU MF data may be used to provide an additional “hint” as to workload selection. When available, this data allows the RNI for a production workload to be calculated. Using the RNI and another value from CPU MF, the L1 cache miss percentage, a workload may be classified as LOW, AVERAGE or HIGH RNI. This classification and resulting “hint” is automated in the zPCR tool. It is highly recommended to use zPCR for capacity sizing. For those wanting to create the “hint” by hand, the following table may be used for z10:

Table 4. **z10 RNI “hint”**

L1MP	RNI	Workload Hint
<3%	>= 0.75	AVERAGE
	< 0.75	LOW
3% to 6%	>1.0	HIGH
	0.6 to 1.0	AVERAGE
	< 0.6	LOW
>6%	>=0.75	HIGH
	< 0.75	AVERAGE

Note this table may change in the future.

Estimating Utilization for a New Processor

If both the relative capacity of a planned new processor and the utilization of the current processor are known, then that information can be used to estimate the utilization of the new processor after the workload has been moved. For example, if planning to move an existing workload from CPU-A, currently running at some utilization level, to CPU-B, we can approximate the processor utilization expected on CPU-B as follows:

CPU-B Utilization = CPU-A Utilization * (ITR for CPU-A / ITR for CPU-B)

As an example, let's assume that CPU-A has a mixed workload ITR rating of 10 and is running at 90% utilization today (a peak period average). We are considering moving this workload to CPU-B with a mixed workload ITR rating of 15. If the workload were moved over to CPU-B without change, you could expect the utilization on the new processor to be 60% (90% × 10 ÷ 15).

Estimating Utilization with Workload Growth

A new processor model is usually chosen so that it can absorb a certain amount of workload growth before another processor change is necessary. If workload growth can be estimated, rates can be applied to the current known production workload, to determine the anticipated utilization level, for the new processor, at points in the future. In this way, one can see when the capacity of the processor will be exceeded, thereby needing to be replaced.

- **Latent Demand**

  Frequently the need to consider a new processor is driven by the fact that the current machine has already become a constraint to getting work done. In other words, more work would be done if more processor resource were available (making changes to other resources could also produce such an effect). Here we have pent-up demand which will result in an instantaneous workload growth when the new processor is installed. Such growth is called "latent demand".

- **Annual Growth**

  Planned (and unplanned) growth over time is referred to as annual growth. As user population increases, or as new/enhanced applications are put into production, growth will occur. By historically tracking growth, one can make assumptions about future annual growth rates.

Growth rates can be applied to the average utilization of the current processor to determine what the utilization will be on a new processor as that growth occurs. The formula to adjust the current processor's utilization for either latent or annual growth is:

CPU-A Utilization = CPU-A Utilizaiton * (1 + Growth rate)

With workload growth

Figure 4 shows a graph, with three horizontal lines representing three potential new processors, each scaled with its capacity relative to a currently installed CPU-A, based on LSPR data for the appropriate workload mix. Workload growth is assumed to be 18%, compounded annually over a five year period.

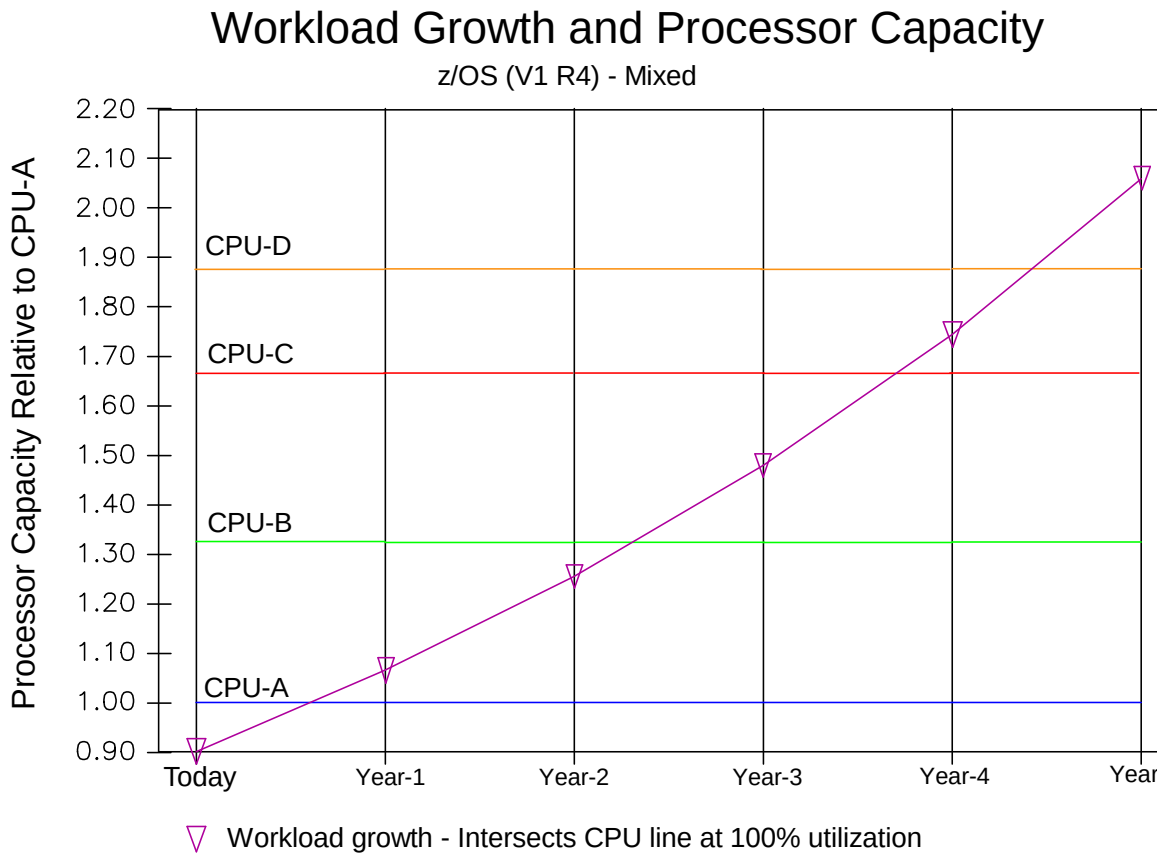

Figure 4**. Example plotting workload growth against relative processor capacity**

Knowing that the current utilization of CPU-A is 90%, we can start the growth curve at 0.90 of the capacity of the CPU-A. With an annual growth rate of 18%, the growth curve will be at 1.062 for the first year (0.90 × 1.18), 1.253 for the second year (1.062 × 1.18), and so on.

The expected utilization can be computed for any of the processors by dividing the growth value expected at any point in time by the relative capacity of the processor in question.

If latent demand is considered to be a factor in the move to a new processor, its growth should be applied to the growth curve's starting point (today). Annual growth should then be applied to that value.

When the growth curve crosses a processor line, the estimated utilization level on that machine will be 100%, meaning that its capacity will be exceeded. Many installations would consider themselves out of capacity when the average utilization level exceeds some threshold that is less than 100%. Here the growth line can be drawn at that threshold to determine when a processor's capacity will be exceeded.

Note: The utilization estimates above do not address any generic low utilization effects (LUEs), or effects of changing the logically partitioned (LPAR) mode setup during migrations.

Chapter 5. Validating a New Processor's Capacity Expectation

The decision about which processor to install as a replacement will be based on the relative capacity expectation for a potential new processor. No matter what source for relative processor capacity data was used, one would like to be able to show that the capacity ratios used, were, in fact, correct for the actual production workload.

As discussed in "Customized Benchmark", processor capacity data derived from a customized benchmark, which has been designed to be representative of the production workload, should be accurate. In this case, you have the desired data before the new processor is installed, and you can have a high degree of confidence in that data.

If a representative benchmark is not possible, processor capacity data must be accepted from one or more outside sources, such as consultant MIPS tables, vendor claims, capacity planning models, or preferably IBM's LSPR. Whatever source is used, there will be no easy means to validate the expected processor relationship until after the new processor is installed. Validating your expectation after the fact is useful, in that doing so can provide you with some level of confidence in the source of capacity data being used.

Note: Being correct once does not prove infallibility. A one-time validation of a new processor's capacity expectation is no guarantee that the source used will always be correct. It is always possible that the source happened to have provided the correct capacity relationship for the two processors only by chance (just as it is always possible that a "kernel" might happen to produce the exact capacity ratio that a production workload would realize). That same source might well be in error when a different set of processors are compared.

A Limited View

One way that relative capacity is typically measured is to look at specific pieces of production work, such as certain jobs or applications. By comparing the CPU time to do a single piece of work on the new processor, to the same data for the old processor, you can establish a capacity relationship for that specific piece of work. There are two major problems with this approach:

- **The sample of work tested is small**

  You are only looking at specific portions of the overall workload. Even though they may be running along with the normal workload mix, the numbers represent only that work. The relative capacity number that you

should be looking for represents the entire production workload, taken as a whole.

Each unit of work run has its own individual sensitivity to processor design. Processor capacity ratios are even more sensitive for specific individual units of work than they are to general workload types. Capacity ratios for individual work units have a significant potential to be higher or lower than the capacity relationship for the workload as a whole. To determine the overall capacity relationship, you would have to perform our test on all (or most) of the individual work units in the production workload.

Figure 5 shows an example of two processors, where the relative capacity of the second is exactly two times the first. This relationship is determined by computing a ratio between the two measured ITR values for a specific workload environment. When you examine the relative capacity for any given unit-of-work within the overall workload, there can be a significant variation from the average.

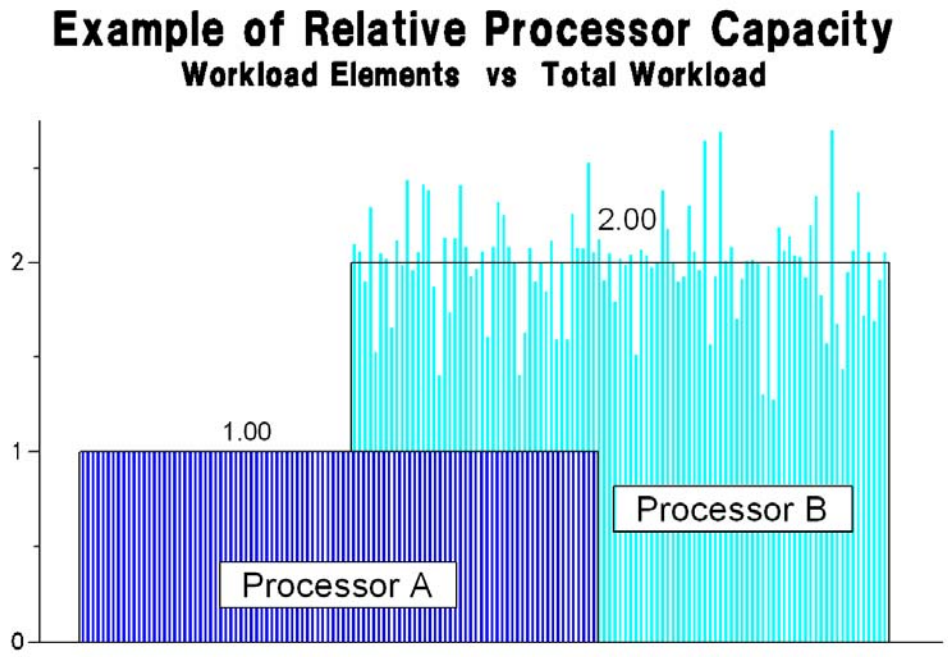

Figure 5. Workload element capacity vs total workload capacity

- **System task time is difficult to apply**

Many operating system functions are necessary to support a workload, including items such as JES, VTAM, RACF, RMF, and VM monitor. These functions also contribute to the overall capacity relationship that is realized, having their own relationship with the individual processor designs. By using only the time to do a specific job or transaction, you are focusing on only a portion of the overall processor time necessary to support the work.

Processor time used by system tasks is normally captured by the operating system. However, because these system tasks support various aspects of the entire workload, it is not easy to accurately apportion that time to the specific work being tested.

- **Un-captured time is ignored**

  The processing time for individual work is determined from accounting data, or from self-contained timing routines. Data obtained in this way does not include all of the processing time to do the work. In other words, you are dealing only with captured processor time, and are ignoring a portion of the processing time necessary to support the work. This Un-captured processor time is the portion that the SCP cannot assign to a specific job or application.

  The amount of Un-captured time that can occur on a processor varies greatly, being dependent on the overall nature of the workload, and on the level of resource management necessary. Uncaptured time is generally SCP time used for managing resources. Just as workloads have their individual behavior characteristics, so do these SCP routines. Therefore, this Un-captured time can have an influence on the overall capacity relationship between the two machines.

The Complete View

There is a better approach toward validating a relative capacity expectation, that will yield a more realistic view of the actual capacity relationship that was realized. With appropriate planning, a validation can be made with a modest effort and minimal impact.

Because both the old and the new processors are seldom available at the same time, you must capture data when they are available. This means that the old processor must be adequately measured before the upgrade, and the new processor must be adequately measured after the upgrade. Generally, there will be no opportunity to repeat measurements on the old processor after the new one is installed.

For a validation to work, there must be a commitment that the workload run on the new processor be the same as that on the old processor. In other words, there should be no shifting of workloads until after the validation is complete.

In order to provide a realistic view of the relative capacity that was realized, you will need to develop an ITR value for the production workload, on both the old processor and the new processor. As stated in "Internal Throughput Rate (ITR)", an ITR is computed as:

The factors used to compute the ITRs for our validation have some special considerations, since the production workload is not a controlled benchmark.

- **Units of work**

  Because a production workload generally consists of a mixture of different types of work (for example, on-line and batch), it becomes difficult to use traditional units, such as jobs or transactions, as the measure of work. Therefore, you need to come up with an alternate unit of work that can be used for this exercise.

  The best approach to solve this dilemma, is to take the average data I/O as the unit of work. These are logical data I/Os or EXCP counts as issued from subsystems and application software. (Physical data I/Os cannot be used, since a logical I/O does not always result in a physical I/O.) All other I/Os, such as paging, should be excluded from consideration. Using the logical data I/O is reasonable, in that the relationship of these I/Os to CPU time on any given processor should remain constant for a given production workload. It is important, however, that you capture logical data I/Os over a long enough period to get a representative view of the average workload.

  When using this approach you must insure that the workload remains relatively constant over the measurement period (there are statistical techniques to verify this). The use of I/Os as a constant work reference is only valid if the workload characteristics do not vary significantly over the measurement interval.

  Note: The use of data-in-memory may require some special considerations because of the way I/Os are reported.

- **Processor busy time**

  To be meaningful, an ITR must consider all processor busy time to do the work being measured; that is, Un-captured time must not be excluded. Accounting data does not include Un-captured system overhead time. If using accounting data to determine processor time, it must be adjusted to include any amount that is Un-captured.

  Accounting data usually provides processing time in terms of a single engine. If the machine happens to be an N-way processor, this data must be adjusted to represent the complex as a single entity, rather than simply one engine. Therefore, on an N-way processor, processor busy time is computed as the sum of the individual CP busy times divided by "N".

To compute an ITR for a production workload, formula 3 may be easier to work with, in that the units needed are more readily available. As stated in "ITR/ETR Relationship", an alternate way to compute an ITR is:

ITR = ETR / Processor Utilization

For the ETR, substitute the logical data I/O rate (I/Os per elapsed second). Processor utilization is provided directly by most SCP related performance monitors. This utilization value includes all processing time, both captured and uncaptured, and therefore meets our requirement of representing all processor time. Processor utilization should be expressed as a fraction relative to 1.00.

No matter what approach is taken to compare the capacity realized for the two processors, care must be taken that the workload measured on the old processor is the same workload that is measured on the new processor. Therefore, the installation plan for the new processor should exclude any intentional change in the day-to-day production workload during the period of test. There should be no new applications added. Nor should there be any workload balancing attempted (shifting applications or load across different processors). Once the validation is completed, you are free to add applications, and shift workload components around as desired.

For this validity check to be fair, you must be certain that you are looking at equivalent samples of work for both measurements. The best way to ensure repeatable work is to capture a large sample, such as a full weeks worth of data, probably during the normal prime shift hours. Take care that you are not encountering business cycle differences between the two measurement periods. (There are statistical approaches for analyzing the captured data to assure that the samples are, in fact, repeatable work over the measurement period).

A major advantage of using this validation approach is that you are determining relative capacity with the ultimate benchmark, the exact workload for which the processor was purchased. The ITR computed represents all of the processing time ("captured" and "uncaptured") to do our day-to-day production workload.

The primary disadvantage of this validation approach is that it can only be done after making a commitment to a new processor model. By doing this validation, however, you will be in a position to assess your confidence level in whatever processor capacity reference you used to make your processor decision. In fact, once the validation is completed, the results can also be used to help assess the accuracy of any other processor capacity reference data that could have been used.

Chapter 6. Summary

There is a need for reliable processor capacity planning data across high-end System z processors. Accurate capacity planning exercises and processor upgrade decisions depend on the availability of reliable data. Since workloads with different characteristics may have significantly different performance ratios when moved among various processors, the most reliable data must be workload sensitive. Additionally, LPAR configuration, specialty engines and processor configuration all should be factored into accurate capacity relationships. MIPS tables provide a single-number-metric reflecting average workloads and configurations. LSPR tables provide workload-sensitive ratios. zPCR allows customized estimates to be created that are sensitive to all the afore mentioned factors.

Appendixes

Appendix A. LSPR ITR Ratios for IBM Processors

This appendix is intended as a quick reference for Large Systems Performance Reference (LSPR) processor capacity data (based on the latest information available at date of publication of this manual). It provides capacity ratios for IBM processors, running the various LSPR workload environments under z/OS, , z/VM, and Linux on System z.

Data contained in this appendix is subject to change to reflect new or additional measurements, or to provide more accurate data based on the latest information available. It is your responsibility to insure that you are working with the latest version of the data which can be found at the following website:

https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument

The primary purpose of the LSPR is to provide relative capacity data for IBM processors, running a variety of SCP and workload environments. The LSPR provides an extensive set of relative capacity data for, z/Architecture processors, across the entire IBM processor line. LSPR ITR data is obtained using representative benchmark workloads, run as laboratory controlled tests, with objective analysis of the results. IBM believes that, with the LSPR data, it has the most exhaustive and accurate set of relative capacity data for z/Architecture processors in the industry.

In addition to all supported IBM processors, some of the previously issued sets of LSPR data include relative capacity information for System/370, System/390 archtectures for IBM and competitive processors such as Amdahl and HDS processors where it has been measured, or can be reasonably assessed. Only the System z machines are included in the newer sets of ITRR tables in this document; if information is needed for older processors, running older versions of the SCPs, older tables may be found at the previously referenced LSPR website.

A VSE ITRR table was not included in this release of the LSPR. However, in those instances that a VSE performance ratio between an existing machine and a z10 BC is required, it is expected that the VSE performance ratio will be similar to the performance ratios established for the two processors as determined by the z/OS V1R11 workloads in the z10 BC z/OS V1R11 ITRR table.

Using the ITR Ratio Values in the Tables

To determine the capacity of any specific processor relative to another for any given SCP and workload, divide the ITR ratio of the 2nd processor by the ITR ratio of the 1st. For the tables contained in this document this process is straightforward since all of the tables presented have the same “base” processor.

PR/SM LPAR Considerations / Multi-Image and Single-Image Tables

Depending on the level of the z/OS, the ratios in the LSPR tables are presented in several ways All measurements in support of the tables contained in this document were run in LPAR mode.

It has been observed that the vast majority of System z clients run fairly complex LPAR configurations on their processors, while some clients continue to grow their z/OS single-image size. To address these two environments, two tables of capacity ratios have been provided: 1) a table based on configuring multiple images of z/OS reflecting average configurations across the processor family and 2) a table based on configuring a single image of z/OS equal in size to the number of engines in the processor (subject to the z/OS 1.8 limit of 32).

Extensive profiling of client usage of System z processors has shown that over 95% of the processors have significantly exploited the virtualization capabilities of the System z platform, that is, they are configured with multiple images of z/OS. Thus, the multi-image tables are based on an average multi-image z/OS configuration. The main variables in the configuration are: 1) number of images, 2) size of each image (number of logical engines), 3) relative weight of each image, 4) overall ratio of logical engines to physical engines, 5) the number of books, and 6) the number of ICFs/IFLs. The configurations used for the multi-image table are based on the average values for these variables as observed across a processor family. For example, it was found that the average number of images ranged from 5 at low-end models to 8 at the high end. Most systems were configured with 2 major images (those defined with >20% relative weight). On low- to mid-range models, at least one of the major images tended to be configured with a number of logical engines close to the number of physical engines. On high-end boxes, the major images were generally configured with a number of logical engines well below the count of physical engines reflecting the more common use of these processors for consolidation. The overall ratio of logical to physical engines (often referred to as “the level of over commitment” in a virtualized environment) averaged as high as 5:1 on the smallest models, hovered around 2:1 across the majority of models, and dropped to 1.3:1 on the largest models. A majority of models were configured with an additional book beyond what would be required to hold the enabled processor engines, and the average model was configured with 2 ICFs/IFLs.

For high-level sizing, most users will find the multi-image table to reflect configurations closest to their own. This is simply due to the fact that most systems are run with multiple z/OS images. However, the most accurate sizings require the zPCR tool which can be customized to exactly match a specific multi-image configuration rather than the average configurations reflected in the multi-image table. The zPCR tool is publicly available.

z/OS 1.9 tables: The LSPR 1.9 tables add the z10 BC and continue the practice of providing both a Single Image table and a Multi-Image table. z/OS 1.9 support for more than 32way for a single OS image allows for the expansion of the Single Image table to include up to 64 processors. The highest N-ways of the single-image table include enhancements whose availability are expected to be consistent with customer requirements. Anyone contemplating running very large z/OS images should contact IBM for a review of potential constraints.

z/OS 1.11 table: The LSPR for z/OS 1.11 add the zEnterprise 196 and now provides only a Multi-Image table. Note the single-image measurements are still done and are input to the zPCR tool. However, only the Multi-Image table appears in the LSPR as it provides the most representative view of performance. Note that it is still recommended that anyone contemplating running very large z/OS images (>32way) should contact IBM for a review of potential constraints.

Your Mileage May Vary

Any processor running a multitude of logical partitions is at increased risk of performance variability from minute to minute or day to day as the workload demands of each partition can affect the performance achieved by all other partitions. The likelihood of significant performance variation increases in proportion to the size (capacity) of the processor and the number of logical partitions that are active. Performance variability may manifest itself in several forms, for example, the capacity realized by an individual logical partition may be impacted as may any charge back algorithm based on CPU time or service units. While the LSPR and zPCR tool can provide good "middle-of-the-road" capacity estimates, your mileage may vary (by minute, hour or day) particularly with larger LPAR configurations.

SMG/MSU Software Model Group and MSU Values

The column heading SMG contains the Software Model Group and the column headed MSU (Million of Service Units/hr.) contains the MSU value for each of the machines. The MSUs for z990 and later processor families include adjustments to provide increased IBM software price/performance for applicable IBM software that has MSU-based pricing. Therefore, MSU ratios among processor families do not necessarily reflect the capacity ratios among processor families. These values are provided for information only and the official source for MSU publication is at the following Web address:
http://www.ibm.com/zseries/library/swpriceinfo/

MSU's of non-IBM processors are based upon vendor claims of MSU's. They are not based upon Large Systems Performance Reference measurements.

Note: Throughout the tables, "NA" means ITR data is not available (there is no measured basis on which to make a relative capacity determination).

Table 5. ITR Ratio Table – z/OS V1 R11 Multi-Image

The table in this section presents LSPR capacity data in the form of ITR ratios for IBM processors where each model is configured with multiple z/OS V1R11 images based on an average LPAR profile of client systems. All capacity numbers are relative to the IBM 2094-701.

Part 1 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.

Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM e(logo)Server zSeries 800						
2066-0E1	1	35	7	0.06	0.06	0.06
2066-0A1	1	69	13	0.12	0.12	0.12
2066-0B1	1	99	20	0.18	0.18	0.18
2066-0C1	1	123	25	0.22	0.22	0.22
2066-0X2	2	141	28	0.26	0.25	0.24
2066-001	1	164	32	0.29	0.29	0.30
2066-0A2	2	229	44	0.42	0.41	0.39
2066-002	2	309	60	0.56	0.55	0.53
2066-003	3	445	84	0.82	0.80	0.75
2066-004	4	577	108	1.06	1.03	0.98

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 2 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM e(logo)Server zSeries 890						
2086-110	1	23	4	0.04	0.04	0.04
2086-210	2	45	8	0.08	0.08	0.08
2086-310	3	67	11	0.12	0.12	0.11
2086-410	4	90	15	0.16	0.16	0.15
2086-120	1	42	7	0.07	0.07	0.07
2086-220	2	80	13	0.15	0.14	0.14
2086-320	3	120	20	0.22	0.21	0.20
2086-420	4	160	26	0.29	0.29	0.27
2086-130	1	80	13	0.15	0.14	0.14
2086-230	2	155	26	0.29	0.28	0.26
2086-330	3	231	38	0.43	0.41	0.39
2086-430	4	309	49	0.57	0.55	0.52
2086-140	1	101	17	0.18	0.18	0.18
2086-240	2	194	32	0.36	0.35	0.33
2086-340	3	289	47	0.53	0.52	0.49
2086-440	4	387	62	0.71	0.69	0.65
2086-150	1	157	26	0.28	0.28	0.28
2086-250	2	303	50	0.56	0.54	0.52
2086-350	3	452	74	0.84	0.81	0.76
2086-450	4	605	97	1.11	1.08	1.02
2086-160	1	192	32	0.35	0.34	0.34
2086-260	2	370	62	0.68	0.66	0.63
2086-360	3	552	91	1.02	0.99	0.93
2086-460	4	738	119	1.36	1.32	1.25
2086-170	1	333	56	0.60	0.60	0.59
2086-270	2	644	107	1.18	1.15	1.11
2086-370	3	959	158	1.76	1.71	1.64
2086-470	4	1,282	208	2.35	2.29	2.19

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 3 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z9 BC R07						
2096-A01	1	25	4	0.05	0.05	0.04
2096-A02	2	49	7	0.09	0.09	0.08
2096-A03	3	73	10	0.13	0.13	0.12
2096-B01	1	36	5	0.06	0.06	0.06
2096-B02	2	70	10	0.13	0.13	0.12
2096-B03	3	104	15	0.19	0.19	0.18
2096-C01	1	44	6	0.08	0.08	0.08
2096-C02	2	86	12	0.16	0.15	0.15
2096-C03	3	128	18	0.23	0.23	0.22
2096-D01	1	57	8	0.10	0.10	0.10
2096-D02	2	110	16	0.20	0.20	0.19
2096-D03	3	163	23	0.30	0.29	0.28
2096-E01	1	67	10	0.12	0.12	0.12
2096-E02	2	131	19	0.24	0.23	0.23
2096-F01	1	84	12	0.15	0.15	0.15
2096-F02	2	162	24	0.30	0.29	0.28
2096-G01	1	106	15	0.19	0.19	0.19
2096-H01	1	125	18	0.22	0.22	0.22
2096-I01	1	144	21	0.26	0.26	0.26
2096-J01	1	165	24	0.30	0.30	0.30

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 4 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z9 BC S07						
2096-K04	4	215	30	0.39	0.38	0.36
2096-L03	3	193	28	0.35	0.35	0.33
2096-L04	4	255	36	0.47	0.46	0.43
2096-M03	3	240	34	0.44	0.43	0.41
2096-M04	4	316	45	0.58	0.56	0.54
2096-N02	2	205	30	0.37	0.37	0.36
2096-N03	3	303	43	0.55	0.54	0.52
2096-N04	4	400	56	0.73	0.71	0.68
2096-O02	2	242	36	0.44	0.43	0.42
2096-O03	3	358	52	0.65	0.64	0.61
2096-O04	4	473	67	0.86	0.84	0.80
2096-P02	2	279	41	0.51	0.50	0.49
2096-P03	3	413	59	0.75	0.74	0.71
2096-P04	4	545	77	0.99	0.97	0.93
2096-Q02	2	320	47	0.58	0.57	0.56
2096-Q03	3	474	68	0.87	0.85	0.81
2096-Q04	4	625	88	1.14	1.12	1.06
2096-R01	1	185	27	0.33	0.33	0.33
2096-R02	2	359	52	0.66	0.64	0.63
2096-R03	3	531	76	0.97	0.95	0.91
2096-R04	4	702	99	1.28	1.25	1.19
2096-S01	1	207	30	0.37	0.37	0.37
2096-S02	2	402	59	0.73	0.72	0.70
2096-S03	3	595	85	1.09	1.06	1.02
2096-S04	4	785	111	1.43	1.40	1.34
2096-T01	1	232	34	0.42	0.42	0.42
2096-T02	2	451	66	0.82	0.81	0.79
2096-T03	3	666	95	1.22	1.19	1.14
2096-T04	4	880	124	1.60	1.57	1.50
2096-U01	1	259	38	0.46	0.46	0.46
2096-U02	2	503	73	0.92	0.90	0.88
2096-U03	3	744	106	1.36	1.33	1.28
2096-U04	4	982	138	1.79	1.75	1.67
2096-V01	1	291	42	0.52	0.52	0.52
2096-V02	2	564	82	1.03	1.01	0.99
2096-V03	3	835	119	1.52	1.49	1.44
2096-V04	4	1,102	155	2.01	1.97	1.88
2096-W01	1	327	47	0.58	0.58	0.59
2096-W02	2	633	92	1.15	1.13	1.11
2096-W03	3	936	134	1.71	1.67	1.61
2096-W04	4	1,236	174	2.25	2.21	2.11
2096-X01	1	365	53	0.65	0.65	0.66
2096-X02	2	708	103	1.29	1.26	1.24
2096-X03	3	1,047	150	1.91	1.87	1.81
2096-X04	4	1,382	195	2.52	2.47	2.36
2096-Y01	1	405	59	0.73	0.72	0.73
2096-Y02	2	786	115	1.43	1.40	1.38
2096-Y03	3	1,163	166	2.12	2.08	2.01
2096-Y04	4	1,535	216	2.80	2.74	2.63
2096-Z01	1	462	67	0.83	0.82	0.83
2096-Z02	2	895	130	1.63	1.60	1.57
2096-Z03	3	1,323	189	2.42	2.36	2.29
2096-Z04	4	1,747	246	3.19	3.12	3.00

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 5 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z10 BC						
2098-A01	1	25	3	0.05	0.05	0.04
2098-A02	2	47	6	0.10	0.08	0.06
2098-A03	3	68	9	0.14	0.12	0.09
2098-A04	4	88	11	0.18	0.16	0.11
2098-A05	5	107	13	0.23	0.19	0.14
2098-B01	1	29	4	0.06	0.05	0.04
2098-B02	2	54	7	0.11	0.10	0.07
2098-B03	3	78	10	0.16	0.14	0.10
2098-B04	4	102	13	0.21	0.18	0.13
2098-B05	5	123	15	0.26	0.22	0.16
2098-C01	1	37	5	0.07	0.07	0.06
2098-C02	2	69	9	0.14	0.12	0.10
2098-C03	3	100	12	0.20	0.18	0.14
2098-C04	4	130	16	0.26	0.23	0.18
2098-C05	5	158	19	0.32	0.28	0.22
2098-D01	1	45	6	0.09	0.08	0.07
2098-D02	2	84	11	0.16	0.15	0.12
2098-D03	3	121	15	0.24	0.22	0.17
2098-D04	4	157	19	0.31	0.28	0.21
2098-D05	5	191	23	0.39	0.34	0.26
2098-E01	1	52	7	0.10	0.09	0.08
2098-E02	2	96	12	0.19	0.17	0.14
2098-E03	3	139	17	0.28	0.25	0.19
2098-E04	4	181	22	0.36	0.32	0.25
2098-E05	5	220	27	0.44	0.39	0.30
2098-F01	1	58	7	0.11	0.10	0.09
2098-F02	2	107	14	0.21	0.19	0.15
2098-F03	3	155	19	0.31	0.28	0.22
2098-F04	4	201	25	0.40	0.36	0.27
2098-F05	5	244	30	0.50	0.44	0.33
2098-G01	1	69	9	0.13	0.12	0.10
2098-G02	2	127	16	0.25	0.23	0.18
2098-G03	3	184	23	0.37	0.33	0.26
2098-G04	4	239	29	0.48	0.43	0.33
2098-G05	5	290	36	0.59	0.52	0.39
2098-H01	1	76	10	0.14	0.14	0.12
2098-H02	2	142	18	0.28	0.25	0.20
2098-H03	3	205	26	0.41	0.37	0.28
2098-H04	4	266	33	0.53	0.48	0.36
2098-H05	5	323	40	0.66	0.58	0.44
2098-I01	1	86	11	0.16	0.15	0.13
2098-I02	2	160	20	0.31	0.29	0.23
2098-I03	3	231	29	0.46	0.41	0.32
2098-I04	4	300	37	0.60	0.54	0.41
2098-I05	5	365	45	0.74	0.65	0.49
2098-J01	1	96	12	0.18	0.17	0.14
2098-J02	2	178	23	0.35	0.32	0.25
2098-J03	3	257	32	0.51	0.46	0.36
2098-J04	4	334	41	0.67	0.60	0.45
2098-J05	5	406	50	0.83	0.73	0.55
2098-K01	1	108	14	0.20	0.19	0.16
2098-K02	2	200	25	0.39	0.36	0.29
2098-K03	3	289	36	0.58	0.52	0.40
2098-K04	4	375	46	0.75	0.67	0.51
2098-K05	5	456	56	0.93	0.81	0.62
2098-L01	1	127	16	0.24	0.23	0.19
2098-L02	2	236	30	0.47	0.42	0.34
2098-L03	3	342	43	0.68	0.61	0.47
2098-L04	4	443	55	0.89	0.79	0.60
2098-L05	5	538	66	1.10	0.96	0.73
2098-M01	1	147	19	0.28	0.26	0.22
2098-M02	2	273	35	0.54	0.49	0.39
2098-M03	3	394	49	0.79	0.70	0.54
2098-M04	4	511	63	1.03	0.91	0.69
2098-M05	5	621	76	1.27	1.11	0.84

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 6 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z10 BC (continued)						
2098-N01	1	168	21	0.32	0.30	0.25
2098-N02	2	313	40	0.62	0.56	0.44
2098-N03	3	452	57	0.91	0.81	0.62
2098-N04	4	586	72	1.19	1.05	0.79
2098-N05	5	712	87	1.46	1.27	0.96
2098-O01	1	189	24	0.36	0.34	0.28
2098-O02	2	351	44	0.70	0.63	0.49
2098-O03	3	507	63	1.02	0.91	0.69
2098-O04	4	658	81	1.34	1.17	0.88
2098-O05	5	799	98	1.64	1.43	1.07
2098-P01	1	211	27	0.41	0.38	0.31
2098-P02	2	393	50	0.79	0.70	0.55
2098-P03	3	567	71	1.15	1.01	0.77
2098-P04	4	736	91	1.50	1.31	0.98
2098-P05	5	894	110	1.85	1.60	1.19
2098-Q01	1	237	30	0.46	0.42	0.35
2098-Q02	2	440	56	0.88	0.79	0.62
2098-Q03	3	636	80	1.29	1.14	0.87
2098-Q04	4	825	102	1.68	1.47	1.11
2098-Q05	5	1,002	123	2.07	1.79	1.34
2098-R01	1	264	33	0.51	0.47	0.40
2098-R02	2	491	62	0.98	0.88	0.69
2098-R03	3	710	89	1.43	1.27	0.97
2098-R04	4	921	113	1.87	1.64	1.24
2098-R05	5	1,119	137	2.30	2.00	1.50
2098-S01	1	297	38	0.57	0.53	0.45
2098-S02	2	551	70	1.10	0.98	0.78
2098-S03	3	797	100	1.60	1.42	1.10
2098-S04	4	1,034	127	2.09	1.85	1.40
2098-S05	5	1,256	154	2.58	2.24	1.69
2098-T01	1	333	42	0.64	0.59	0.50
2098-T02	2	619	78	1.23	1.11	0.88
2098-T03	3	895	112	1.80	1.60	1.24
2098-T04	4	1,161	143	2.35	2.07	1.57
2098-T05	5	1,410	173	2.89	2.52	1.91
2098-U01	1	373	47	0.71	0.67	0.57
2098-U02	2	694	88	1.37	1.24	0.99
2098-U03	3	1,003	125	2.01	1.79	1.39
2098-U04	4	1,301	160	2.62	2.32	1.77
2098-U05	5	1,581	194	3.23	2.83	2.15
2098-V01	1	421	53	0.80	0.75	0.64
2098-V02	2	781	99	1.54	1.40	1.13
2098-V03	3	1,130	141	2.26	2.02	1.58
2098-V04	4	1,466	180	2.94	2.62	2.01
2098-V05	5	1,782	218	3.62	3.18	2.44
2098-W01	1	471	60	0.89	0.84	0.72
2098-W02	2	875	111	1.72	1.56	1.27
2098-W03	3	1,265	158	2.52	2.26	1.78
2098-W04	4	1,642	202	3.28	2.93	2.27
2098-W05	5	1,995	245	4.04	3.56	2.75
2098-X01	1	528	67	1.00	0.94	0.82
2098-X02	2	981	124	1.92	1.75	1.44
2098-X03	3	1,419	177	2.81	2.54	2.01
2098-X04	4	1,842	226	3.67	3.29	2.56
2098-X05	5	2,239	274	4.52	4.00	3.10
2098-Y01	1	604	76	1.13	1.08	0.95
2098-Y02	2	1,123	142	2.19	2.01	1.66
2098-Y03	3	1,624	202	3.20	2.90	2.32
2098-Y04	4	2,108	258	4.18	3.76	2.96
2098-Y05	5	2,562	313	5.14	4.58	3.59
2098-Z01	1	661	83	1.24	1.18	1.04
2098-Z02	2	1,229	155	2.38	2.19	1.83
2098-Z03	3	1,777	221	3.49	3.17	2.56
2098-Z04	4	2,307	283	4.55	4.12	3.26
2098-Z05	5	2,805	342	5.60	5.01	3.96

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 7 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM e(logo)Server zSeries 900 (1XX Series)						
2064-101	1	210	41	0.37	0.37	0.38
2064-102	2	400	78	0.73	0.72	0.70
2064-103	3	589	112	1.07	1.05	1.01
2064-104	4	776	143	1.41	1.39	1.31
2064-105	5	949	173	1.75	1.69	1.60
2064-106	6	1,114	199	2.07	1.99	1.87
2064-107	7	1,273	225	2.39	2.27	2.13
2064-108	8	1,423	245	2.70	2.54	2.36
2064-109	9	1,568	265	3.00	2.80	2.59
2064-1C1	1	215	43	0.38	0.38	0.39
2064-1C2	2	417	83	0.76	0.74	0.73
2064-1C3	3	619	119	1.13	1.11	1.07
2064-1C4	4	822	153	1.50	1.47	1.41
2064-1C5	5	1,014	187	1.86	1.81	1.74
2064-1C6	6	1,201	217	2.22	2.15	2.06
2064-1C7	7	1,384	247	2.58	2.47	2.37
2064-1C8	8	1,562	276	2.93	2.79	2.67
2064-1C9	9	1,737	302	3.28	3.10	2.96
2064-110	10	1,909	327	3.62	3.41	3.25
2064-111	11	2,078	350	3.96	3.71	3.53
2064-112	12	2,243	372	4.30	4.01	3.81
2064-113	13	2,405	392	4.63	4.30	4.08
2064-114	14	2,564	410	4.96	4.58	4.34
2064-115	15	2,719	426	5.29	4.86	4.59
2064-116	16	2,870	441	5.61	5.13	4.84
IBM e(logo)Server zSeries 900 (2XX Series)						
2064-2C1	1	261	52	0.46	0.47	0.48
2064-2C2	2	506	100	0.92	0.90	0.88
2064-2C3	3	751	144	1.37	1.34	1.28
2064-2C4	4	996	184	1.82	1.78	1.69
2064-2C5	5	1,228	224	2.26	2.19	2.09
2064-2C6	6	1,453	260	2.69	2.60	2.47
2064-2C7	7	1,673	296	3.12	2.99	2.84
2064-2C8	8	1,886	330	3.54	3.37	3.20
2064-2C9	9	2,095	362	3.96	3.74	3.55
2064-210	10	2,300	392	4.37	4.11	3.89
2064-211	11	2,500	420	4.78	4.47	4.23
2064-212	12	2,696	445	5.18	4.82	4.55
2064-213	13	2,888	475	5.57	5.16	4.87
2064-214	14	3,074	497	5.96	5.49	5.18
2064-215	15	3,256	517	6.34	5.82	5.48
2064-216	16	3,434	535	6.72	6.13	5.77

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 8 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM e(logo)Server zSeries 990						
2084-301	1	413	70	0.74	0.74	0.75
2084-302	2	801	132	1.46	1.43	1.41
2084-303	3	1,186	191	2.17	2.12	2.05
2084-304	4	1,568	248	2.86	2.80	2.69
2084-305	5	1,923	302	3.54	3.44	3.31
2084-306	6	2,265	352	4.20	4.05	3.89
2084-307	7	2,595	402	4.85	4.64	4.44
2084-308	8	2,913	448	5.48	5.20	4.96
2084-309	9	3,224	492	6.11	5.76	5.47
2084-310	10	3,529	538	6.73	6.30	5.96
2084-311	11	3,827	580	7.34	6.84	6.45
2084-312	12	4,119	620	7.94	7.36	6.93
2084-313	13	4,405	661	8.54	7.87	7.39
2084-314	14	4,685	696	9.13	8.37	7.85
2084-315	15	4,959	730	9.72	8.86	8.29
2084-316	16	5,227	761	10.29	9.34	8.73
2084-317	17	5,491	799	10.86	9.81	9.16
2084-318	18	5,751	837	11.42	10.27	9.59
2084-319	19	6,009	878	11.98	10.73	10.02
2084-320	20	6,267	919	12.54	11.19	10.45
2084-321	21	6,523	959	13.09	11.65	10.87
2084-322	22	6,777	999	13.64	12.11	11.29
2084-323	23	7,031	1,037	14.19	12.56	11.71
2084-324	24	7,283	1,076	14.74	13.01	12.13
2084-325	25	7,534	1,114	15.28	13.46	12.54
2084-326	26	7,783	1,151	15.82	13.90	12.95
2084-327	27	8,032	1,188	16.36	14.35	13.35
2084-328	28	8,279	1,225	16.89	14.79	13.76
2084-329	29	8,524	1,261	17.42	15.23	14.16
2084-330	30	8,768	1,296	17.95	15.66	14.55
2084-331	31	9,010	1,332	18.48	16.09	14.95
2084-332	32	9,250	1,365	19.00	16.52	15.34

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 9 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z9 EC						
2094-401	1	193	28	0.34	0.34	0.34
2094-402	2	375	54	0.68	0.67	0.65
2094-403	3	554	78	1.01	0.99	0.95
2094-404	4	732	102	1.33	1.31	1.24
2094-405	5	897	124	1.65	1.60	1.52
2094-406	6	1,056	144	1.95	1.89	1.79
2094-407	7	1,209	164	2.25	2.16	2.04
2094-408	8	1,356	182	2.55	2.42	2.28
2094-501	1	374	53	0.67	0.67	0.67
2094-502	2	726	104	1.32	1.30	1.26
2094-503	3	1,074	152	1.96	1.92	1.84
2094-504	4	1,419	197	2.58	2.53	2.41
2094-505	5	1,739	240	3.19	3.11	2.96
2094-506	6	2,048	279	3.79	3.66	3.47
2094-507	7	2,344	317	4.37	4.19	3.96
2094-508	8	2,630	352	4.94	4.70	4.43
2094-601	1	454	65	0.81	0.81	0.81
2094-602	2	880	127	1.60	1.57	1.53
2094-603	3	1,303	184	2.38	2.33	2.23
2094-604	4	1,720	240	3.13	3.07	2.92
2094-605	5	2,109	292	3.87	3.77	3.58
2094-606	6	2,482	339	4.59	4.43	4.21
2094-607	7	2,842	385	5.30	5.08	4.81
2094-608	8	3,188	428	5.99	5.69	5.37
2094-701	1	560	81	1.00	1.00	1.00
2094-702	2	1,086	158	1.98	1.94	1.89
2094-703	3	1,607	229	2.93	2.87	2.75
2094-704	4	2,122	298	3.86	3.79	3.60
2094-705	5	2,601	363	4.78	4.65	4.42
2094-706	6	3,062	422	5.67	5.47	5.19
2094-707	7	3,505	479	6.54	6.26	5.93
2094-708	8	3,932	532	7.38	7.02	6.62
2094-709	9	4,349	584	8.22	7.77	7.30
2094-710	10	4,757	640	9.05	8.50	7.96
2094-711	11	5,156	690	9.87	9.21	8.61
2094-712	12	5,546	742	10.67	9.91	9.24
2094-713	13	5,928	795	11.47	10.59	9.86
2094-714	14	6,301	843	12.26	11.26	10.46
2094-715	15	6,667	893	13.03	11.91	11.06
2094-716	16	7,024	938	13.80	12.55	11.63
2094-717	17	7,373	985	14.55	13.17	12.20
2094-718	18	7,714	1,032	15.30	13.78	12.75
2094-719	19	8,054	1,077	16.04	14.39	13.30
2094-720	20	8,392	1,127	16.78	14.99	13.85
2094-721	21	8,729	1,177	17.51	15.59	14.39
2094-722	22	9,063	1,226	18.24	16.19	14.93
2094-723	23	9,396	1,274	18.97	16.79	15.47
2094-724	24	9,728	1,314	19.69	17.38	16.00
2094-725	25	10,058	1,353	20.41	17.97	16.53
2094-726	26	10,386	1,400	21.13	18.55	17.05
2094-727	27	10,712	1,436	21.84	19.14	17.57
2094-728	28	11,037	1,481	22.55	19.72	18.09
2094-729	29	11,360	1,524	23.26	20.29	18.60
2094-730	30	11,680	1,567	23.96	20.86	19.10
2094-731	31	11,998	1,609	24.66	21.43	19.61
2094-732	32	12,314	1,650	25.35	22.00	20.11
2094-733	33	12,627	1,691	26.04	22.56	20.60
2094-734	34	12,939	1,732	26.73	23.11	21.09
2094-735	35	13,248	1,772	27.41	23.67	21.57
2094-736	36	13,555	1,811	28.10	24.21	22.05
2094-737	37	13,860	1,850	28.77	24.76	22.53
2094-738	38	14,163	1,889	29.45	25.30	23.00
2094-739	39	14,462	1,927	30.12	25.83	23.47
2094-740	40	14,757	1,963	30.78	26.36	23.93
2094-741	41	15,048	1,998	31.43	26.88	24.39
2094-742	42	15,335	2,033	32.08	27.39	24.84
2094-743	43	15,619	2,067	32.72	27.90	25.29
2094-744	44	15,899	2,101	33.36	28.40	25.73
2094-745	45	16,175	2,135	33.98	28.89	26.17
2094-746	46	16,448	2,168	34.61	29.38	26.61
2094-747	47	16,717	2,201	35.22	29.86	27.04
2094-748	48	16,982	2,233	35.83	30.34	27.47
2094-749	49	17,245	2,265	36.44	30.81	27.89
2094-750	50	17,503	2,295	37.03	31.27	28.31
2094-751	51	17,759	2,324	37.63	31.72	28.72
2094-752	52	18,011	2,353	38.21	32.17	29.13
2094-753	53	18,260	2,381	38.79	32.62	29.54
2094-754	54	18,505	2,409	39.37	33.06	29.94

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 10 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z10 EC						
2097-401	1	214	27	0.39	0.38	0.35
2097-402	2	404	51	0.76	0.72	0.66
2097-403	3	589	75	1.13	1.05	0.95
2097-404	4	768	97	1.48	1.37	1.23
2097-405	5	943	118	1.84	1.68	1.50
2097-406	6	1,115	139	2.18	1.99	1.77
2097-407	7	1,284	160	2.52	2.29	2.03
2097-408	8	1,449	180	2.85	2.59	2.29
2097-409	9	1,610	199	3.18	2.88	2.55
2097-410	10	1,769	218	3.50	3.16	2.80
2097-411	11	1,924	237	3.81	3.44	3.05
2097-412	12	2,077	255	4.12	3.71	3.29
2097-501	1	462	58	0.85	0.83	0.76
2097-502	2	882	110	1.65	1.58	1.42
2097-503	3	1,281	160	2.43	2.29	2.05
2097-504	4	1,661	207	3.20	2.97	2.64
2097-505	5	2,031	252	3.94	3.63	3.22
2097-506	6	2,391	296	4.68	4.27	3.78
2097-507	7	2,743	340	5.40	4.90	4.33
2097-508	8	3,085	382	6.11	5.51	4.87
2097-509	9	3,420	422	6.80	6.11	5.39
2097-510	10	3,746	462	7.48	6.69	5.90
2097-511	11	4,063	500	8.15	7.26	6.39
2097-512	12	4,373	537	8.80	7.81	6.88
2097-601	1	626	79	1.15	1.12	1.03
2097-602	2	1,189	149	2.22	2.12	1.92
2097-603	3	1,729	215	3.27	3.09	2.76
2097-604	4	2,248	277	4.29	4.02	3.55
2097-605	5	2,751	339	5.30	4.91	4.32
2097-606	6	3,240	398	6.28	5.79	5.07
2097-607	7	3,713	455	7.25	6.63	5.80
2097-608	8	4,172	511	8.20	7.45	6.50
2097-609	9	4,618	565	9.13	8.25	7.19
2097-610	10	5,050	617	10.04	9.02	7.86
2097-611	11	5,469	668	10.93	9.77	8.51
2097-612	12	5,876	717	11.81	10.50	9.15

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 11 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System z10 EC (Continued)						
2097-701	1	902	115	1.65	1.61	1.49
2097-702	2	1,705	215	3.20	3.05	2.76
2097-703	3	2,468	312	4.70	4.41	3.95
2097-704	4	3,192	401	6.16	5.70	5.07
2097-705	5	3,895	488	7.60	6.96	6.16
2097-706	6	4,577	571	9.01	8.18	7.21
2097-707	7	5,238	651	10.38	9.36	8.23
2097-708	8	5,879	729	11.73	10.50	9.22
2097-709	9	6,502	804	13.06	11.61	10.18
2097-710	10	7,105	875	14.36	12.69	11.10
2097-711	11	7,691	944	15.63	13.74	12.00
2097-712	12	8,258	1,011	16.87	14.75	12.87
2097-713	13	8,809	1,076	18.09	15.74	13.72
2097-714	14	9,355	1,139	19.31	16.71	14.55
2097-715	15	9,894	1,202	20.51	17.68	15.38
2097-716	16	10,429	1,264	21.70	18.63	16.19
2097-717	17	10,958	1,329	22.88	19.58	16.99
2097-718	18	11,482	1,390	24.06	20.51	17.79
2097-719	19	12,001	1,451	25.22	21.44	18.57
2097-720	20	12,514	1,512	26.37	22.35	19.35
2097-721	21	13,022	1,571	27.52	23.26	20.12
2097-722	22	13,525	1,631	28.65	24.16	20.87
2097-723	23	14,024	1,690	29.78	25.05	21.62
2097-724	24	14,517	1,748	30.90	25.93	22.36
2097-725	25	15,005	1,805	32.01	26.80	23.09
2097-726	26	15,488	1,865	33.11	27.67	23.81
2097-727	27	15,969	1,922	34.20	28.53	24.53
2097-728	28	16,448	1,979	35.29	29.38	25.24
2097-729	29	16,924	2,037	36.38	30.23	25.95
2097-730	30	17,398	2,092	37.45	31.08	26.66
2097-731	31	17,869	2,146	38.53	31.92	27.37
2097-732	32	18,338	2,200	39.59	32.76	28.07
2097-733	33	18,805	2,257	40.66	33.59	28.76
2097-734	34	19,269	2,309	41.71	34.42	29.46
2097-735	35	19,731	2,366	42.76	35.25	30.14
2097-736	36	20,191	2,422	43.81	36.07	30.83
2097-737	37	20,648	2,478	44.85	36.89	31.51
2097-738	38	21,104	2,530	45.89	37.70	32.19
2097-739	39	21,557	2,585	46.92	38.51	32.87
2097-740	40	22,007	2,636	47.94	39.31	33.54
2097-741	41	22,455	2,687	48.96	40.11	34.21
2097-742	42	22,902	2,740	49.98	40.91	34.88
2097-743	43	23,345	2,789	50.99	41.70	35.54
2097-744	44	23,787	2,838	52.00	42.49	36.20
2097-745	45	24,227	2,886	53.00	43.28	36.85
2097-746	46	24,664	2,934	53.99	44.06	37.51
2097-747	47	25,099	2,981	54.98	44.84	38.16
2097-748	48	25,532	3,028	55.97	45.61	38.80
2097-749	49	25,963	3,075	56.95	46.38	39.45
2097-750	50	26,391	3,120	57.92	47.14	40.09
2097-751	51	26,818	3,166	58.89	47.91	40.72
2097-752	52	27,242	3,214	59.86	48.66	41.35
2097-753	53	27,664	3,262	60.82	49.42	41.98
2097-754	54	28,084	3,305	61.77	50.17	42.61
2097-755	55	28,502	3,352	62.73	50.92	43.24
2097-756	56	28,918	3,395	63.67	51.66	43.86
2097-757	57	29,321	3,438	64.60	52.38	44.46
2097-758	58	29,713	3,480	65.52	53.08	45.03
2097-759	59	30,092	3,525	66.43	53.76	45.59
2097-760	60	30,460	3,570	67.32	54.41	46.13
2097-761	61	30,818	3,611	68.20	55.05	46.65
2097-762	62	31,164	3,652	69.06	55.67	47.15
2097-763	63	31,500	3,695	69.91	56.27	47.63
2097-764	64	31,826	3,739	70.75	56.85	48.10

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 12 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System zEnterprise 196						
2817-401	1	240	30	0.43	0.43	0.41
2817-402	2	459	58	0.84	0.82	0.77
2817-403	3	672	85	1.24	1.20	1.12
2817-404	4	879	110	1.63	1.57	1.46
2817-405	5	1,082	135	2.02	1.93	1.80
2817-406	6	1,280	160	2.40	2.29	2.13
2817-407	7	1,476	183	2.78	2.64	2.45
2817-408	8	1,667	207	3.15	2.98	2.77
2817-409	9	1,855	229	3.51	3.31	3.08
2817-410	10	2,039	252	3.87	3.64	3.38
2817-411	11	2,220	274	4.23	3.97	3.68
2817-412	12	2,397	296	4.58	4.28	3.98
2817-413	13	2,571	318	4.92	4.59	4.26
2817-414	14	2,742	339	5.26	4.90	4.55
2817-415	15	2,909	359	5.59	5.20	4.83
2817-501	1	588	74	1.04	1.05	1.01
2817-502	2	1,123	140	2.03	2.01	1.88
2817-503	3	1,642	204	3.00	2.93	2.73
2817-504	4	2,147	265	3.95	3.84	3.55
2817-505	5	2,639	326	4.88	4.71	4.35
2817-506	6	3,118	385	5.80	5.57	5.14
2817-507	7	3,586	442	6.70	6.41	5.90
2817-508	8	4,042	498	7.60	7.22	6.65
2817-509	9	4,486	552	8.47	8.01	7.38
2817-510	10	4,919	604	9.34	8.79	8.09
2817-511	11	5,342	655	10.19	9.54	8.78
2817-512	12	5,754	705	11.03	10.28	9.46
2817-513	13	6,156	754	11.86	11.00	10.12
2817-514	14	6,547	806	12.67	11.70	10.77
2817-515	15	6,929	849	13.48	12.38	11.40
2817-601	1	768	97	1.36	1.37	1.31
2817-602	2	1,460	182	2.66	2.61	2.43
2817-603	3	2,131	263	3.93	3.81	3.52
2817-604	4	2,783	344	5.18	4.97	4.57
2817-605	5	3,417	421	6.40	6.10	5.60
2817-606	6	4,035	497	7.59	7.21	6.61
2817-607	7	4,635	569	8.77	8.28	7.59
2817-608	8	5,219	640	9.92	9.32	8.54
2817-609	9	5,788	709	11.05	10.34	9.47
2817-610	10	6,341	777	12.16	11.33	10.38
2817-611	11	6,879	843	13.25	12.29	11.27
2817-612	12	7,403	907	14.32	13.22	12.13
2817-613	13	7,913	969	15.37	14.14	12.97
2817-614	14	8,409	1,028	16.39	15.02	13.79
2817-615	15	8,891	1,084	17.40	15.88	14.59

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 13 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System zEnterprise 196 (continued)						
2817-701	1	1,202	150	2.14	2.15	2.06
2817-702	2	2,272	281	4.15	4.06	3.78
2817-703	3	3,311	408	6.13	5.92	5.46
2817-704	4	4,320	531	8.06	7.72	7.08
2817-705	5	5,300	650	9.96	9.47	8.66
2817-706	6	6,251	766	11.82	11.17	10.19
2817-707	7	7,175	879	13.65	12.82	11.68
2817-708	8	8,072	988	15.44	14.42	13.12
2817-709	9	8,943	1,091	17.19	15.97	14.52
2817-710	10	9,788	1,191	18.92	17.49	15.88
2817-711	11	10,609	1,286	20.61	18.95	17.21
2817-712	12	11,407	1,381	22.27	20.38	18.49
2817-713	13	12,181	1,473	23.89	21.76	19.74
2817-714	14	12,932	1,562	25.49	23.10	20.95
2817-715	15	13,662	1,648	27.06	24.41	22.12
2817-716	16	14,371	1,731	28.59	25.67	23.26
2817-717	17	15,076	1,816	30.12	26.93	24.40
2817-718	18	15,778	1,899	31.64	28.19	25.53
2817-719	19	16,476	1,983	33.15	29.43	26.66
2817-720	20	17,171	2,064	34.65	30.67	27.78
2817-721	21	17,862	2,144	36.14	31.91	28.90
2817-722	22	18,550	2,224	37.62	33.14	30.02
2817-723	23	19,234	2,306	39.09	34.36	31.12
2817-724	24	19,915	2,388	40.56	35.58	32.23
2817-725	25	20,592	2,469	42.01	36.79	33.33
2817-726	26	21,266	2,550	43.46	37.99	34.43
2817-727	27	21,937	2,627	44.89	39.19	35.52
2817-728	28	22,604	2,704	46.32	40.38	36.61
2817-729	29	23,268	2,780	47.74	41.57	37.69
2817-730	30	23,929	2,855	49.16	42.75	38.77
2817-731	31	24,586	2,927	50.56	43.92	39.84
2817-732	32	25,241	2,998	51.96	45.09	40.91

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Part 14 of 14

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE.
Information presented in this table is current as of May 2010.

Processor	#CP	PCI**	MSU***	z/OS V1R11 Low*	z/OS V1R11 Average*	z/OS V1R11 High*
IBM System zEnterprise 196 (continued)						
2817-733	33	25,893	3,068	53.35	46.25	41.98
2817-734	34	26,543	3,134	54.73	47.42	43.04
2817-735	35	27,191	3,207	56.12	48.57	44.09
2817-736	36	27,837	3,279	57.49	49.73	45.14
2817-737	37	28,482	3,351	58.87	50.88	46.18
2817-738	38	29,124	3,418	60.24	52.03	47.22
2817-739	39	29,765	3,489	61.61	53.17	48.26
2817-740	40	30,403	3,564	62.97	54.31	49.29
2817-741	41	31,040	3,639	64.33	55.45	50.31
2817-742	42	31,675	3,713	65.68	56.58	51.33
2817-743	43	32,308	3,788	67.03	57.71	52.34
2817-744	44	32,938	3,861	68.37	58.84	53.35
2817-745	45	33,568	3,935	69.72	59.96	54.35
2817-746	46	34,195	4,009	71.05	61.08	55.35
2817-747	47	34,820	4,082	72.39	62.20	56.34
2817-748	48	35,443	4,155	73.72	63.32	57.33
2817-749	49	36,065	4,228	75.04	64.43	58.31
2817-750	50	36,683	4,300	76.36	65.53	59.29
2817-751	51	37,297	4,372	77.68	66.63	60.25
2817-752	52	37,907	4,444	78.99	67.72	61.21
2817-753	53	38,514	4,515	80.29	68.80	62.16
2817-754	54	39,117	4,586	81.59	69.88	63.11
2817-755	55	39,717	4,656	82.89	70.95	64.04
2817-756	56	40,313	4,726	84.18	72.01	64.97
2817-757	57	40,905	4,795	85.46	73.07	65.90
2817-758	58	41,494	4,864	86.74	74.12	66.81
2817-759	59	42,079	4,933	88.01	75.17	67.72
2817-760	60	42,661	5,001	89.28	76.21	68.62
2817-761	61	43,239	5,069	90.54	77.24	69.51
2817-762	62	43,814	5,136	91.80	78.27	70.39
2817-763	63	44,385	5,203	93.06	79.29	71.27
2817-764	64	44,953	5,270	94.30	80.30	72.14
2817-765	65	45,517	5,336	95.55	81.31	73.01
2817-766	66	46,078	5,402	96.79	82.31	73.87
2817-767	67	46,622	5,466	98.00	83.28	74.70
2817-768	68	47,147	5,528	99.19	84.22	75.50
2817-769	69	47,656	5,588	100.35	85.13	76.29
2817-770	70	48,149	5,646	101.50	86.01	77.05
2817-771	71	48,626	5,702	102.62	86.86	77.78
2817-772	72	49,087	5,757	103.71	87.69	78.50
2817-773	73	49,534	5,810	104.79	88.49	79.19
2817-774	74	49,967	5,862	105.84	89.26	79.86
2817-775	75	50,385	5,912	106.87	90.01	80.51
2817-776	76	50,791	5,960	107.89	90.73	81.14
2817-777	77	51,183	6,007	108.88	91.43	81.76
2817-778	78	51,563	6,053	109.85	92.11	82.35
2817-779	79	51,931	6,097	110.80	92.77	82.93
2817-780	80	52,286	6,140	111.74	93.40	83.49

*** LOW, AVERAGE, HIGH workloads describe their Relative Nest Intensity (RNI) category - see workload description for more information**

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Average column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing only; they are not a capacity metric**

Table 6. ITR Ratio Table – z/OS V1 R9 Multi-Image

The tables in this section present LSPR capacity data in the form of ITR ratios for IBM processors where each model is configured with multiple z/OS V1R9 images based on an average LPAR profile of client systems. All capacity numbers are relative to the IBM 2094-701. All workloads were run in 64 bit real addressing mode.

Part 1 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM e(logo)Server zSeries 800											
2066-0E1	1	35	7	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
2066-0A1	1	71	13	0.12	0.12	0.12	0.12	0.12	0.13	0.12	0.12
2066-0B1	1	102	20	0.18	0.18	0.18	0.17	0.18	0.18	0.18	0.18
2066-0C1	1	126	25	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.23
2066-0X2	2	145	28	0.25	0.25	0.25	0.26	0.26	0.25	0.25	0.25
2066-001	1	170	32	0.30	0.29	0.30	0.29	0.30	0.29	0.30	0.31
2066-0A2	2	234	44	0.41	0.41	0.41	0.42	0.42	0.41	0.40	0.41
2066-002	2	319	60	0.56	0.56	0.56	0.57	0.57	0.56	0.55	0.57
2066-003	3	460	84	0.81	0.80	0.80	0.81	0.82	0.81	0.77	0.83
2066-004	4	590	108	1.04	1.03	1.02	1.03	1.05	1.07	0.99	1.05
IBM e(logo)Server zSeries 890											
2086-110	1	24	4	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.05
2086-210	2	47	8	0.08	0.08	0.08	0.09	0.09	0.08	0.08	0.09
2086-310	3	70	11	0.12	0.12	0.12	0.13	0.13	0.12	0.11	0.13
2086-410	4	92	15	0.16	0.16	0.16	0.17	0.17	0.15	0.15	0.17
2086-120	1	43	7	0.08	0.07	0.08	0.07	0.08	0.07	0.07	0.08
2086-220	2	84	13	0.15	0.15	0.15	0.15	0.15	0.15	0.14	0.16
2086-320	3	125	20	0.22	0.22	0.22	0.22	0.22	0.22	0.21	0.23
2086-420	4	165	26	0.29	0.29	0.29	0.29	0.29	0.29	0.27	0.30
2086-130	1	83	13	0.15	0.14	0.15	0.14	0.14	0.15	0.14	0.15
2086-230	2	162	26	0.28	0.28	0.28	0.28	0.28	0.29	0.27	0.29
2086-330	3	240	38	0.42	0.42	0.42	0.41	0.42	0.44	0.40	0.43
2086-430	4	317	49	0.56	0.55	0.55	0.55	0.56	0.58	0.53	0.56
2086-140	1	104	17	0.18	0.18	0.18	0.18	0.18	0.19	0.18	0.19
2086-240	2	202	32	0.36	0.35	0.35	0.35	0.36	0.37	0.34	0.37
2086-340	3	300	47	0.53	0.52	0.52	0.52	0.53	0.55	0.50	0.54
2086-440	4	396	62	0.69	0.69	0.68	0.69	0.70	0.73	0.66	0.70
2086-150	1	162	26	0.28	0.28	0.28	0.28	0.28	0.29	0.28	0.30
2086-250	2	316	50	0.55	0.55	0.55	0.55	0.55	0.57	0.53	0.57
2086-350	3	468	74	0.82	0.81	0.81	0.81	0.82	0.85	0.78	0.84
2086-450	4	618	97	1.08	1.08	1.07	1.08	1.09	1.13	1.03	1.10
2086-160	1	198	32	0.35	0.34	0.35	0.34	0.34	0.36	0.34	0.36
2086-260	2	386	62	0.68	0.67	0.67	0.66	0.67	0.70	0.65	0.70
2086-360	3	571	91	1.00	0.99	0.99	0.99	1.00	1.04	0.96	1.02
2086-460	4	754	119	1.32	1.32	1.30	1.31	1.33	1.39	1.26	1.34
2086-170	1	344	56	0.60	0.60	0.60	0.59	0.60	0.61	0.59	0.62
2086-270	2	670	107	1.18	1.16	1.17	1.15	1.17	1.21	1.14	1.21
2086-370	3	992	158	1.74	1.73	1.72	1.72	1.74	1.80	1.68	1.77
2086-470	4	1309	208	2.30	2.29	2.27	2.28	2.31	2.39	2.21	2.31

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 2 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 BC R07											
2096-A01	1	26	4	0.05	0.05	0.05	0.05	0.05	0.04	0.04	0.05
2096-A02	2	51	7	0.09	0.09	0.09	0.10	0.10	0.08	0.08	0.10
2096-A03	3	76	10	0.13	0.13	0.14	0.15	0.15	0.12	0.12	0.15
2096-B01	1	38	5	0.07	0.06	0.07	0.07	0.07	0.06	0.06	0.08
2096-B02	2	75	10	0.13	0.13	0.14	0.13	0.13	0.12	0.12	0.15
2096-B03	3	110	15	0.19	0.19	0.20	0.20	0.20	0.18	0.18	0.22
2096-C01	1	46	6	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08
2096-C02	2	89	12	0.16	0.16	0.16	0.16	0.16	0.15	0.15	0.16
2096-C03	3	131	18	0.23	0.23	0.23	0.24	0.24	0.23	0.22	0.24
2096-D01	1	58	8	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.11
2096-D02	2	115	16	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.21
2096-D03	3	168	23	0.30	0.29	0.30	0.30	0.30	0.29	0.28	0.30
2096-E01	1	69	10	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.13
2096-E02	2	136	19	0.24	0.24	0.24	0.24	0.24	0.24	0.23	0.25
2096-F01	1	86	12	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.16
2096-F02	2	168	24	0.30	0.29	0.30	0.29	0.30	0.30	0.29	0.30
2096-G01	1	109	15	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.20
2096-H01	1	129	18	0.23	0.22	0.23	0.22	0.22	0.23	0.22	0.23
2096-I01	1	148	21	0.26	0.26	0.26	0.26	0.26	0.26	0.26	0.27
2096-J01	1	170	24	0.30	0.30	0.30	0.29	0.29	0.30	0.30	0.31

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 3 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 BC S07											
2096-K04	4	220	30	0.39	0.38	0.38	0.40	0.40	0.38	0.37	0.40
2096-L03	3	199	28	0.35	0.35	0.35	0.35	0.35	0.35	0.34	0.36
2096-L04	4	260	36	0.46	0.46	0.45	0.47	0.47	0.46	0.44	0.47
2096-M03	3	247	34	0.43	0.43	0.43	0.43	0.43	0.44	0.42	0.45
2096-M04	4	323	45	0.57	0.56	0.56	0.57	0.57	0.57	0.54	0.58
2096-N02	2	213	30	0.37	0.37	0.37	0.37	0.37	0.37	0.37	0.38
2096-N03	3	313	43	0.55	0.55	0.55	0.55	0.55	0.55	0.53	0.56
2096-N04	4	409	56	0.72	0.71	0.71	0.72	0.73	0.73	0.69	0.73
2096-O02	2	252	36	0.44	0.44	0.44	0.44	0.44	0.44	0.43	0.45
2096-O03	3	370	52	0.65	0.64	0.65	0.65	0.65	0.66	0.63	0.66
2096-O04	4	483	67	0.85	0.84	0.84	0.85	0.85	0.86	0.81	0.86
2096-P02	2	290	41	0.51	0.50	0.51	0.50	0.50	0.51	0.50	0.52
2096-P03	3	427	59	0.75	0.74	0.75	0.75	0.75	0.76	0.72	0.77
2096-P04	4	557	77	0.98	0.97	0.97	0.98	0.99	0.99	0.94	1.00
2096-Q02	2	333	47	0.58	0.58	0.58	0.58	0.58	0.59	0.57	0.60
2096-Q03	3	489	68	0.86	0.85	0.85	0.86	0.86	0.87	0.83	0.88
2096-Q04	4	639	88	1.12	1.12	1.11	1.13	1.13	1.14	1.08	1.14
2096-R01	1	191	27	0.33	0.33	0.34	0.33	0.33	0.33	0.33	0.34
2096-R02	2	374	52	0.66	0.65	0.65	0.65	0.65	0.66	0.64	0.67
2096-R03	3	549	76	0.96	0.96	0.96	0.96	0.96	0.98	0.93	0.98
2096-R04	4	717	99	1.26	1.25	1.25	1.26	1.27	1.28	1.21	1.28
2096-S01	1	213	30	0.37	0.37	0.38	0.37	0.37	0.38	0.37	0.38
2096-S02	2	418	59	0.73	0.73	0.73	0.72	0.73	0.74	0.72	0.75
2096-S03	3	614	85	1.08	1.07	1.07	1.07	1.08	1.09	1.05	1.10
2096-S04	4	802	111	1.41	1.40	1.39	1.41	1.42	1.43	1.35	1.43
2096-T01	1	239	34	0.42	0.42	0.42	0.41	0.41	0.42	0.42	0.43
2096-T02	2	468	66	0.82	0.81	0.82	0.81	0.81	0.83	0.81	0.84
2096-T03	3	688	95	1.21	1.20	1.20	1.20	1.21	1.23	1.17	1.23
2096-T04	4	899	124	1.58	1.57	1.56	1.58	1.59	1.61	1.52	1.60
2096-U01	1	267	38	0.47	0.46	0.47	0.46	0.46	0.47	0.46	0.48
2096-U02	2	522	73	0.92	0.91	0.91	0.90	0.91	0.93	0.90	0.93
2096-U03	3	767	106	1.35	1.34	1.34	1.34	1.35	1.37	1.31	1.37
2096-U04	4	1002	138	1.76	1.75	1.74	1.76	1.78	1.80	1.69	1.78
2096-V01	1	299	42	0.52	0.52	0.53	0.51	0.52	0.53	0.52	0.54
2096-V02	2	586	82	1.03	1.02	1.03	1.01	1.02	1.04	1.01	1.05
2096-V03	3	861	119	1.51	1.50	1.50	1.50	1.51	1.54	1.47	1.53
2096-V04	4	1125	155	1.97	1.97	1.95	1.97	1.99	2.02	1.90	2.00
2096-W01	1	336	47	0.59	0.58	0.59	0.58	0.58	0.59	0.59	0.60
2096-W02	2	657	92	1.15	1.14	1.15	1.13	1.15	1.17	1.14	1.17
2096-W03	3	966	134	1.69	1.69	1.68	1.68	1.70	1.73	1.65	1.72
2096-W04	4	1262	174	2.21	2.21	2.19	2.21	2.24	2.27	2.14	2.24
2096-X01	1	375	53	0.66	0.65	0.66	0.64	0.64	0.66	0.66	0.67
2096-X02	2	735	103	1.29	1.28	1.29	1.27	1.27	1.31	1.27	1.31
2096-X03	3	1080	150	1.89	1.88	1.88	1.87	1.88	1.94	1.85	1.92
2096-X04	4	1410	195	2.47	2.47	2.45	2.47	2.48	2.54	2.39	2.49
2096-Y01	1	417	59	0.73	0.72	0.73	0.71	0.71	0.74	0.73	0.74
2096-Y02	2	816	115	1.43	1.42	1.43	1.40	1.41	1.46	1.42	1.45
2096-Y03	3	1199	166	2.10	2.09	2.09	2.08	2.09	2.15	2.06	2.12
2096-Y04	4	1566	216	2.75	2.74	2.72	2.73	2.76	2.83	2.66	2.76
2096-Z01	1	474	67	0.83	0.82	0.83	0.81	0.81	0.84	0.83	0.84
2096-Z02	2	928	130	1.63	1.62	1.62	1.59	1.61	1.66	1.61	1.64
2096-Z03	3	1364	189	2.39	2.38	2.37	2.36	2.38	2.46	2.35	2.41
2096-Z04	4	1782	246	3.12	3.12	3.09	3.10	3.14	3.23	3.03	3.14

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 4 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 BC											
2098-A01	1	26	3	0.05	0.05	0.05	0.06	0.05	0.04	0.04	0.06
2098-A02	2	48	6	0.08	0.08	0.08	0.12	0.10	0.07	0.06	0.10
2098-A03	3	69	9	0.12	0.12	0.12	0.17	0.15	0.11	0.09	0.14
2098-A04	4	88	11	0.15	0.16	0.15	0.22	0.19	0.14	0.11	0.18
2098-A05	5	107	13	0.19	0.19	0.18	0.28	0.24	0.17	0.14	0.21
2098-B01	1	30	4	0.05	0.05	0.05	0.07	0.06	0.05	0.04	0.06
2098-B02	2	56	7	0.10	0.10	0.10	0.13	0.12	0.09	0.08	0.11
2098-B03	3	80	10	0.14	0.14	0.14	0.19	0.17	0.13	0.11	0.16
2098-B04	4	102	13	0.18	0.18	0.17	0.25	0.22	0.17	0.13	0.20
2098-B05	5	123	15	0.22	0.22	0.21	0.30	0.27	0.21	0.16	0.24
2098-C01	1	38	5	0.07	0.07	0.07	0.08	0.07	0.06	0.06	0.07
2098-C02	2	71	9	0.12	0.12	0.12	0.15	0.14	0.12	0.10	0.13
2098-C03	3	101	12	0.18	0.18	0.17	0.22	0.20	0.18	0.14	0.19
2098-C04	4	129	16	0.23	0.23	0.22	0.28	0.26	0.23	0.18	0.23
2098-C05	5	156	19	0.27	0.28	0.26	0.35	0.31	0.28	0.21	0.28
2098-D01	1	46	6	0.08	0.08	0.08	0.09	0.09	0.08	0.07	0.09
2098-D02	2	85	11	0.15	0.15	0.15	0.18	0.17	0.15	0.12	0.16
2098-D03	3	122	15	0.21	0.22	0.21	0.26	0.24	0.22	0.17	0.22
2098-D04	4	156	19	0.27	0.28	0.26	0.34	0.31	0.28	0.22	0.28
2098-D05	5	189	23	0.33	0.34	0.31	0.42	0.38	0.34	0.26	0.34
2098-E01	1	53	7	0.09	0.09	0.09	0.11	0.10	0.09	0.08	0.10
2098-E02	2	99	12	0.17	0.17	0.17	0.21	0.19	0.17	0.14	0.18
2098-E03	3	141	17	0.25	0.25	0.24	0.30	0.28	0.25	0.20	0.26
2098-E04	4	180	22	0.32	0.32	0.30	0.39	0.36	0.32	0.25	0.33
2098-E05	5	218	27	0.38	0.39	0.36	0.48	0.44	0.40	0.30	0.39
2098-F01	1	59	7	0.10	0.10	0.10	0.12	0.11	0.10	0.09	0.11
2098-F02	2	110	14	0.19	0.19	0.19	0.23	0.21	0.19	0.16	0.21
2098-F03	3	156	19	0.27	0.28	0.26	0.34	0.31	0.28	0.22	0.29
2098-F04	4	200	25	0.35	0.36	0.33	0.44	0.40	0.36	0.28	0.36
2098-F05	5	242	30	0.42	0.44	0.40	0.54	0.49	0.44	0.33	0.43
2098-G01	1	70	9	0.12	0.12	0.12	0.14	0.13	0.12	0.10	0.13
2098-G02	2	130	16	0.23	0.23	0.22	0.27	0.25	0.23	0.19	0.24
2098-G03	3	186	23	0.33	0.33	0.31	0.40	0.37	0.33	0.26	0.34
2098-G04	4	237	29	0.42	0.43	0.39	0.52	0.48	0.43	0.33	0.43
2098-G05	5	287	36	0.50	0.52	0.47	0.64	0.58	0.53	0.39	0.51
2098-H01	1	78	10	0.14	0.14	0.14	0.16	0.15	0.13	0.12	0.15
2098-H02	2	145	18	0.25	0.26	0.25	0.30	0.28	0.25	0.21	0.27
2098-H03	3	207	26	0.36	0.37	0.35	0.44	0.41	0.37	0.29	0.38
2098-H04	4	264	33	0.46	0.47	0.44	0.58	0.53	0.48	0.37	0.48
2098-H05	5	320	40	0.56	0.58	0.53	0.71	0.65	0.59	0.44	0.57
2098-I01	1	88	11	0.15	0.15	0.15	0.18	0.17	0.15	0.13	0.17
2098-I02	2	163	20	0.29	0.29	0.28	0.34	0.32	0.29	0.23	0.30
2098-I03	3	233	29	0.41	0.42	0.39	0.50	0.46	0.42	0.33	0.43
2098-I04	4	298	37	0.52	0.54	0.49	0.65	0.60	0.55	0.41	0.54
2098-I05	5	361	45	0.63	0.65	0.59	0.80	0.73	0.67	0.49	0.64
2098-J01	1	98	12	0.17	0.17	0.17	0.20	0.18	0.17	0.14	0.19
2098-J02	2	182	23	0.32	0.32	0.31	0.38	0.35	0.32	0.26	0.34
2098-J03	3	260	32	0.46	0.46	0.43	0.55	0.52	0.47	0.37	0.47
2098-J04	4	332	41	0.58	0.60	0.55	0.72	0.67	0.61	0.46	0.60
2098-J05	5	402	50	0.70	0.73	0.66	0.88	0.82	0.75	0.55	0.72
2098-K01	1	110	14	0.19	0.19	0.19	0.22	0.21	0.19	0.16	0.21
2098-K02	2	204	25	0.36	0.36	0.35	0.42	0.40	0.36	0.29	0.38
2098-K03	3	292	36	0.51	0.52	0.49	0.62	0.58	0.53	0.41	0.53
2098-K04	4	372	46	0.65	0.67	0.62	0.81	0.75	0.69	0.51	0.67
2098-K05	5	451	56	0.79	0.81	0.74	0.99	0.92	0.85	0.62	0.80
2098-L01	1	130	16	0.23	0.23	0.22	0.26	0.25	0.23	0.19	0.25
2098-L02	2	241	30	0.42	0.43	0.41	0.50	0.47	0.44	0.34	0.45
2098-L03	3	345	43	0.60	0.61	0.57	0.73	0.69	0.64	0.48	0.63
2098-L04	4	440	55	0.77	0.79	0.73	0.95	0.89	0.83	0.61	0.79
2098-L05	5	533	66	0.93	0.96	0.87	1.16	1.09	1.02	0.73	0.94
2098-M01	1	150	19	0.26	0.26	0.26	0.30	0.28	0.26	0.22	0.29
2098-M02	2	279	35	0.49	0.49	0.47	0.57	0.55	0.51	0.40	0.51
2098-M03	3	398	49	0.70	0.71	0.66	0.83	0.80	0.74	0.56	0.72
2098-M04	4	508	63	0.89	0.91	0.83	1.09	1.03	0.97	0.70	0.91
2098-M05	5	615	76	1.08	1.11	1.00	1.33	1.26	1.19	0.84	1.09

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 5 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 BC (continued)											
2098-N01	1	172	21	0.30	0.30	0.29	0.34	0.33	0.31	0.25	0.33
2098-N02	2	319	40	0.56	0.57	0.53	0.65	0.63	0.59	0.45	0.59
2098-N03	3	456	57	0.80	0.81	0.75	0.95	0.92	0.86	0.63	0.82
2098-N04	4	582	72	1.02	1.05	0.95	1.24	1.19	1.12	0.80	1.04
2098-N05	5	705	87	1.24	1.27	1.15	1.52	1.45	1.38	0.96	1.24
2098-O01	1	193	24	0.34	0.34	0.33	0.38	0.37	0.35	0.28	0.36
2098-O02	2	358	44	0.63	0.63	0.60	0.73	0.71	0.67	0.51	0.66
2098-O03	3	511	63	0.90	0.91	0.84	1.06	1.03	0.98	0.71	0.92
2098-O04	4	653	81	1.15	1.17	1.07	1.38	1.34	1.27	0.89	1.16
2098-O05	5	791	98	1.39	1.43	1.28	1.70	1.63	1.56	1.07	1.39
2098-P01	1	216	27	0.38	0.38	0.37	0.42	0.41	0.40	0.31	0.41
2098-P02	2	401	50	0.70	0.71	0.67	0.81	0.80	0.76	0.57	0.73
2098-P03	3	572	71	1.00	1.02	0.94	1.18	1.16	1.11	0.79	1.03
2098-P04	4	731	91	1.28	1.31	1.19	1.54	1.50	1.45	1.00	1.29
2098-P05	5	885	110	1.55	1.60	1.43	1.89	1.84	1.78	1.19	1.55
2098-Q01	1	242	30	0.42	0.42	0.41	0.47	0.46	0.44	0.35	0.45
2098-Q02	2	449	56	0.79	0.80	0.75	0.90	0.89	0.85	0.64	0.82
2098-Q03	3	641	80	1.12	1.14	1.05	1.32	1.30	1.24	0.89	1.15
2098-Q04	4	819	102	1.44	1.47	1.33	1.72	1.68	1.62	1.12	1.44
2098-Q05	5	991	123	1.74	1.79	1.60	2.11	2.06	1.99	1.34	1.73
2098-R01	1	270	33	0.47	0.47	0.46	0.52	0.52	0.49	0.40	0.51
2098-R02	2	501	62	0.88	0.89	0.83	1.01	0.99	0.95	0.71	0.91
2098-R03	3	715	89	1.25	1.28	1.18	1.47	1.45	1.38	1.00	1.28
2098-R04	4	913	113	1.60	1.64	1.49	1.91	1.88	1.80	1.25	1.61
2098-R05	5	1106	137	1.94	2.00	1.79	2.35	2.29	2.21	1.50	1.93
2098-S01	1	303	38	0.53	0.53	0.51	0.59	0.58	0.55	0.45	0.57
2098-S02	2	563	70	0.99	1.00	0.94	1.13	1.11	1.06	0.80	1.02
2098-S03	3	803	100	1.41	1.43	1.32	1.65	1.62	1.55	1.12	1.43
2098-S04	4	1025	127	1.80	1.84	1.67	2.15	2.10	2.02	1.41	1.80
2098-S05	5	1241	154	2.18	2.24	2.01	2.64	2.57	2.48	1.69	2.16
2098-T01	1	340	42	0.60	0.59	0.58	0.66	0.65	0.62	0.50	0.63
2098-T02	2	631	78	1.11	1.12	1.05	1.26	1.25	1.19	0.90	1.14
2098-T03	3	901	112	1.58	1.61	1.48	1.84	1.82	1.74	1.27	1.60
2098-T04	4	1151	143	2.02	2.07	1.88	2.40	2.36	2.26	1.59	2.02
2098-T05	5	1393	173	2.44	2.52	2.26	2.95	2.88	2.78	1.91	2.42
2098-U01	1	381	47	0.67	0.67	0.65	0.73	0.73	0.69	0.57	0.71
2098-U02	2	707	88	1.24	1.25	1.18	1.41	1.40	1.33	1.02	1.28
2098-U03	3	1010	125	1.77	1.81	1.67	2.06	2.04	1.94	1.43	1.79
2098-U04	4	1289	160	2.26	2.32	2.11	2.69	2.64	2.53	1.79	2.26
2098-U05	5	1561	194	2.74	2.83	2.53	3.30	3.23	3.10	2.15	2.70
2098-V01	1	429	53	0.75	0.75	0.73	0.82	0.82	0.78	0.64	0.80
2098-V02	2	797	99	1.40	1.41	1.33	1.58	1.57	1.49	1.16	1.44
2098-V03	3	1137	141	1.99	2.03	1.88	2.32	2.29	2.18	1.62	2.01
2098-V04	4	1452	180	2.55	2.62	2.38	3.02	2.97	2.84	2.03	2.53
2098-V05	5	1758	218	3.08	3.18	2.86	3.71	3.63	3.48	2.44	3.03
2098-W01	1	480	60	0.84	0.84	0.82	0.92	0.92	0.87	0.72	0.89
2098-W02	2	891	111	1.56	1.58	1.49	1.77	1.76	1.66	1.30	1.60
2098-W03	3	1272	158	2.23	2.28	2.11	2.58	2.56	2.43	1.82	2.25
2098-W04	4	1625	202	2.85	2.93	2.66	3.37	3.32	3.16	2.29	2.83
2098-W05	5	1968	245	3.45	3.57	3.20	4.14	4.06	3.88	2.75	3.38
2098-X01	1	538	67	0.94	0.94	0.92	1.03	1.03	0.97	0.82	0.99
2098-X02	2	999	124	1.75	1.77	1.68	1.98	1.97	1.86	1.47	1.79
2098-X03	3	1426	177	2.50	2.55	2.37	2.89	2.87	2.71	2.06	2.51
2098-X04	4	1822	226	3.20	3.29	2.99	3.76	3.72	3.53	2.59	3.16
2098-X05	5	2206	274	3.87	4.00	3.60	4.62	4.55	4.34	3.10	3.78
2098-Y01	1	615	76	1.08	1.08	1.05	1.17	1.17	1.10	0.95	1.13
2098-Y02	2	1142	142	2.00	2.03	1.92	2.25	2.25	2.11	1.70	2.04
2098-Y03	3	1631	202	2.86	2.92	2.71	3.29	3.27	3.09	2.38	2.85
2098-Y04	4	2083	258	3.65	3.76	3.43	4.28	4.25	4.02	2.99	3.59
2098-Y05	5	2523	313	4.42	4.58	4.12	5.26	5.19	4.94	3.59	4.30
2098-Z01	1	673	83	1.18	1.18	1.15	1.27	1.28	1.20	1.04	1.23
2098-Z02	2	1250	155	2.19	2.22	2.11	2.45	2.46	2.30	1.88	2.22
2098-Z03	3	1784	221	3.13	3.20	2.97	3.58	3.58	3.36	2.63	3.11
2098-Z04	4	2279	283	4.00	4.12	3.76	4.67	4.64	4.38	3.30	3.91
2098-Z05	5	2760	342	4.84	5.01	4.52	5.74	5.67	5.38	3.95	4.69

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 6 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM e(logo)Server zSeries 900 (1XX Series)											
2064-101	1	215	41	0.38	0.37	0.38	0.37	0.37	0.37	0.38	0.38
2064-102	2	413	78	0.72	0.72	0.72	0.73	0.73	0.72	0.72	0.73
2064-103	3	603	112	1.06	1.06	1.05	1.08	1.08	1.06	1.03	1.06
2064-104	4	786	143	1.38	1.38	1.36	1.42	1.41	1.40	1.32	1.38
2064-105	5	959	173	1.68	1.70	1.66	1.74	1.74	1.73	1.60	1.67
2064-106	6	1125	199	1.97	1.99	1.93	2.06	2.06	2.05	1.86	1.95
2064-107	7	1282	225	2.25	2.28	2.18	2.36	2.36	2.37	2.09	2.20
2064-108	8	1430	245	2.51	2.55	2.42	2.64	2.65	2.68	2.31	2.44
2064-109	9	1571	265	2.76	2.81	2.64	2.92	2.93	2.99	2.51	2.66
2064-1C1	1	222	43	0.39	0.38	0.39	0.38	0.38	0.39	0.39	0.40
2064-1C2	2	433	83	0.76	0.75	0.76	0.75	0.75	0.77	0.75	0.78
2064-1C3	3	639	119	1.12	1.11	1.12	1.11	1.11	1.14	1.09	1.14
2064-1C4	4	840	153	1.47	1.47	1.46	1.47	1.48	1.51	1.42	1.49
2064-1C5	5	1036	187	1.82	1.81	1.79	1.83	1.84	1.87	1.74	1.84
2064-1C6	6	1226	217	2.15	2.15	2.12	2.18	2.20	2.22	2.04	2.17
2064-1C7	7	1412	247	2.48	2.48	2.43	2.52	2.56	2.57	2.33	2.49
2064-1C8	8	1592	276	2.79	2.80	2.73	2.86	2.91	2.91	2.61	2.80
2064-1C9	9	1768	302	3.10	3.11	3.03	3.20	3.26	3.25	2.87	3.11
2064-110	10	1940	327	3.40	3.42	3.31	3.53	3.61	3.58	3.13	3.40
2064-111	11	2108	350	3.70	3.73	3.59	3.86	3.95	3.90	3.38	3.69
2064-112	12	2271	372	3.98	4.02	3.86	4.18	4.30	4.22	3.62	3.96
2064-113	13	2431	392	4.26	4.31	4.12	4.51	4.64	4.54	3.86	4.23
2064-114	14	2586	410	4.53	4.59	4.37	4.82	4.97	4.84	4.08	4.48
2064-115	15	2736	426	4.80	4.87	4.61	5.13	5.30	5.15	4.29	4.73
2064-116	16	2882	441	5.05	5.14	4.85	5.44	5.63	5.44	4.50	4.96
IBM e(logo)Server zSeries 900 (2XX Series)											
2064-2C1	1	269	52	0.47	0.47	0.48	0.46	0.46	0.47	0.48	0.49
2064-2C2	2	524	100	0.92	0.91	0.92	0.91	0.92	0.93	0.90	0.93
2064-2C3	3	772	144	1.35	1.35	1.34	1.36	1.36	1.38	1.32	1.36
2064-2C4	4	1013	184	1.78	1.78	1.75	1.79	1.80	1.83	1.71	1.78
2064-2C5	5	1247	224	2.19	2.19	2.15	2.22	2.24	2.26	2.09	2.18
2064-2C6	6	1474	260	2.59	2.60	2.54	2.65	2.67	2.69	2.45	2.57
2064-2C7	7	1695	296	2.97	2.99	2.91	3.06	3.10	3.11	2.80	2.94
2064-2C8	8	1908	330	3.35	3.38	3.26	3.47	3.53	3.52	3.13	3.30
2064-2C9	9	2117	362	3.71	3.76	3.61	3.88	3.94	3.92	3.44	3.65
2064-210	10	2319	392	4.07	4.12	3.94	4.27	4.36	4.31	3.75	3.99
2064-211	11	2516	420	4.41	4.48	4.26	4.66	4.76	4.70	4.05	4.31
2064-212	12	2708	445	4.75	4.83	4.58	5.05	5.17	5.08	4.34	4.62
2064-213	13	2893	475	5.07	5.18	4.88	5.43	5.56	5.45	4.61	4.91
2064-214	14	3073	497	5.39	5.51	5.16	5.80	5.95	5.81	4.87	5.19
2064-215	15	3246	517	5.69	5.83	5.44	6.16	6.33	6.17	5.13	5.46
2064-216	16	3413	535	5.99	6.15	5.71	6.52	6.71	6.52	5.37	5.71

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 7 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

							z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9
Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	CB-L	ODE-B	WASDB	OLTP-W	OLTP-T
IBM e(logo)Server zSeries 990							0.679	0.684	0.715	0.704	0.710
2084-301	1	424	70	0.74	0.74	0.74	0.72	0.72	0.76	0.75	0.75
2084-302	2	831	132	1.46	1.45	1.45	1.42	1.43	1.49	1.45	1.47
2084-303	3	1222	191	2.14	2.14	2.12	2.11	2.13	2.21	2.11	2.15
2084-304	4	1597	248	2.80	2.80	2.77	2.78	2.81	2.91	2.72	2.80
2084-305	5	1957	302	3.43	3.44	3.38	3.44	3.48	3.58	3.30	3.42
2084-306	6	2303	352	4.04	4.05	3.96	4.08	4.14	4.24	3.85	4.01
2084-307	7	2635	402	4.62	4.65	4.51	4.70	4.78	4.88	4.36	4.58
2084-308	8	2953	448	5.18	5.22	5.04	5.31	5.41	5.50	4.84	5.11
2084-309	9	3264	492	5.72	5.78	5.55	5.91	6.02	6.11	5.30	5.63
2084-310	10	3566	538	6.25	6.33	6.04	6.51	6.63	6.71	5.75	6.14
2084-311	11	3861	580	6.77	6.86	6.52	7.09	7.23	7.29	6.18	6.63
2084-312	12	4148	620	7.27	7.38	6.99	7.67	7.83	7.87	6.59	7.10
2084-313	13	4428	661	7.77	7.90	7.44	8.23	8.41	8.44	6.99	7.56
2084-314	14	4700	696	8.24	8.40	7.88	8.79	8.99	9.00	7.38	8.01
2084-315	15	4966	730	8.71	8.88	8.31	9.34	9.56	9.54	7.75	8.45
2084-316	16	5225	761	9.16	9.36	8.72	9.88	10.13	10.08	8.11	8.87
2084-317	17	5480	799	9.61	9.83	9.13	10.42	10.69	10.60	8.46	9.28
2084-318	18	5731	837	10.05	10.29	9.53	10.94	11.23	11.12	8.81	9.70
2084-319	19	5981	878	10.49	10.75	9.93	11.47	11.78	11.63	9.16	10.11
2084-320	20	6230	919	10.93	11.21	10.33	11.99	12.32	12.14	9.50	10.52
2084-321	21	6477	959	11.36	11.66	10.73	12.51	12.86	12.64	9.85	10.93
2084-322	22	6724	999	11.79	12.11	11.12	13.02	13.40	13.15	10.19	11.33
2084-323	23	6970	1037	12.22	12.56	11.52	13.54	13.94	13.65	10.54	11.74
2084-324	24	7215	1076	12.65	13.01	11.91	14.05	14.47	14.15	10.88	12.15
2084-325	25	7459	1114	13.08	13.46	12.31	14.56	15.00	14.64	11.22	12.55
2084-326	26	7702	1151	13.51	13.91	12.70	15.07	15.52	15.14	11.56	12.95
2084-327	27	7944	1188	13.93	14.35	13.09	15.58	16.05	15.63	11.90	13.35
2084-328	28	8185	1225	14.36	14.79	13.48	16.08	16.57	16.12	12.24	13.75
2084-329	29	8425	1261	14.78	15.23	13.86	16.58	17.09	16.60	12.58	14.15
2084-330	30	8664	1296	15.20	15.67	14.25	17.08	17.60	17.09	12.92	14.55
2084-331	31	8902	1332	15.61	16.10	14.63	17.58	18.11	17.57	13.25	14.94
2084-332	32	9138	1365	16.03	16.54	15.01	18.07	18.61	18.05	13.58	15.33

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 8 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

	Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 EC												
	2094-401	1	197	28	0.34	0.34	0.35	0.34	0.34	0.35	0.34	0.35
	2094-402	2	385	54	0.68	0.67	0.67	0.67	0.67	0.68	0.66	0.69
	2094-403	3	566	78	0.99	0.99	0.99	0.98	0.99	1.01	0.97	1.01
	2094-404	4	739	102	1.30	1.29	1.28	1.30	1.31	1.33	1.25	1.32
	2094-405	5	905	124	1.59	1.58	1.57	1.60	1.62	1.63	1.51	1.61
	2094-406	6	1064	144	1.87	1.87	1.84	1.90	1.92	1.93	1.76	1.88
	2094-407	7	1217	164	2.13	2.14	2.09	2.18	2.22	2.22	2.00	2.15
	2094-408	8	1363	182	2.39	2.40	2.33	2.46	2.50	2.50	2.22	2.40
	2094-501	1	381	53	0.67	0.66	0.67	0.65	0.66	0.68	0.67	0.68
	2094-502	2	747	104	1.31	1.30	1.31	1.28	1.30	1.33	1.29	1.33
	2094-503	3	1097	152	1.92	1.92	1.91	1.90	1.92	1.97	1.88	1.94
	2094-504	4	1433	197	2.51	2.51	2.49	2.50	2.54	2.59	2.43	2.53
	2094-505	5	1755	240	3.08	3.08	3.03	3.09	3.14	3.19	2.95	3.09
	2094-506	6	2064	279	3.62	3.62	3.55	3.66	3.72	3.77	3.44	3.63
	2094-507	7	2359	317	4.14	4.15	4.05	4.22	4.30	4.34	3.89	4.13
	2094-508	8	2642	352	4.63	4.66	4.52	4.76	4.86	4.88	4.32	4.62
	2094-601	1	462	65	0.81	0.80	0.81	0.79	0.79	0.82	0.81	0.82
	2094-602	2	905	127	1.59	1.58	1.58	1.55	1.56	1.62	1.57	1.60
	2094-603	3	1330	184	2.33	2.32	2.31	2.30	2.32	2.40	2.29	2.35
	2094-604	4	1737	240	3.05	3.04	3.01	3.03	3.06	3.15	2.96	3.06
	2094-605	5	2128	292	3.73	3.73	3.67	3.74	3.78	3.88	3.58	3.74
	2094-606	6	2502	339	4.39	4.40	4.31	4.43	4.48	4.59	4.17	4.38
	2094-607	7	2860	385	5.02	5.03	4.90	5.10	5.17	5.27	4.73	5.00
	2094-608	8	3204	428	5.62	5.65	5.47	5.76	5.85	5.94	5.25	5.58

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 9 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 EC (continued)											
2094-701	1	570	81	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
2094-702	2	1117	158	1.96	1.96	1.95	1.98	1.98	1.97	1.94	1.95
2094-703	3	1641	229	2.88	2.89	2.86	2.93	2.93	2.91	2.82	2.86
2094-704	4	2144	298	3.76	3.79	3.72	3.86	3.86	3.82	3.64	3.72
2094-705	5	2625	363	4.60	4.65	4.53	4.77	4.78	4.71	4.42	4.54
2094-706	6	3087	422	5.41	5.48	5.31	5.65	5.67	5.57	5.14	5.32
2094-707	7	3530	479	6.19	6.28	6.05	6.52	6.54	6.40	5.82	6.07
2094-708	8	3954	532	6.93	7.04	6.75	7.36	7.38	7.21	6.46	6.78
2094-709	9	4366	584	7.66	7.80	7.43	8.19	8.22	8.00	7.08	7.47
2094-710	10	4769	640	8.36	8.53	8.10	9.01	9.05	8.78	7.67	8.13
2094-711	11	5160	690	9.05	9.24	8.74	9.81	9.87	9.54	8.25	8.78
2094-712	12	5541	742	9.72	9.94	9.36	10.60	10.67	10.28	8.80	9.41
2094-713	13	5913	795	10.37	10.63	9.96	11.38	11.47	11.01	9.33	10.01
2094-714	14	6274	843	11.00	11.29	10.55	12.15	12.26	11.73	9.84	10.60
2094-715	15	6626	893	11.62	11.94	11.11	12.90	13.03	12.43	10.34	11.18
2094-716	16	6969	938	12.22	12.58	11.66	13.64	13.80	13.12	10.81	11.73
2094-717	17	7303	985	12.81	13.20	12.20	14.37	14.55	13.79	11.27	12.27
2094-718	18	7628	1032	13.38	13.80	12.71	15.09	15.30	14.45	11.71	12.79
2094-719	19	7951	1077	13.95	14.41	13.23	15.80	16.04	15.11	12.16	13.32
2094-720	20	8274	1127	14.51	15.01	13.74	16.51	16.78	15.76	12.60	13.83
2094-721	21	8595	1177	15.07	15.60	14.25	17.22	17.52	16.41	13.04	14.35
2094-722	22	8914	1226	15.63	16.20	14.76	17.93	18.26	17.06	13.48	14.87
2094-723	23	9233	1274	16.19	16.79	15.27	18.63	18.99	17.70	13.92	15.38
2094-724	24	9550	1314	16.75	17.38	15.78	19.34	19.72	18.34	14.35	15.90
2094-725	25	9866	1353	17.30	17.97	16.28	20.03	20.45	18.98	14.79	16.41
2094-726	26	10181	1400	17.86	18.56	16.79	20.73	21.18	19.61	15.22	16.92
2094-727	27	10495	1436	18.41	19.14	17.29	21.42	21.91	20.25	15.66	17.43
2094-728	28	10808	1481	18.96	19.72	17.79	22.11	22.63	20.87	16.09	17.93
2094-729	29	11119	1524	19.50	20.30	18.29	22.80	23.35	21.50	16.53	18.44
2094-730	30	11428	1567	20.04	20.87	18.79	23.48	24.07	22.12	16.96	18.94
2094-731	31	11736	1609	20.58	21.44	19.28	24.16	24.78	22.74	17.38	19.44
2094-732	32	12042	1650	21.12	22.01	19.77	24.84	25.49	23.35	17.81	19.94
2094-733	33	12347	1691	21.65	22.58	20.26	25.51	26.19	23.97	18.23	20.43
2094-734	34	12650	1732	22.19	23.14	20.75	26.18	26.90	24.58	18.65	20.93
2094-735	35	12952	1772	22.72	23.70	21.23	26.84	27.59	25.18	19.07	21.41
2094-736	36	13252	1811	23.24	24.26	21.71	27.51	28.29	25.79	19.49	21.90
2094-737	37	13550	1850	23.76	24.81	22.19	28.17	28.98	26.39	19.90	22.39
2094-738	38	13847	1889	24.29	25.36	22.67	28.82	29.67	26.98	20.32	22.87
2094-739	39	14140	1927	24.80	25.91	23.14	29.47	30.35	27.57	20.72	23.34
2094-740	40	14429	1963	25.31	26.45	23.60	30.12	31.03	28.15	21.12	23.80
2094-741	41	14714	1998	25.81	26.98	24.05	30.77	31.71	28.72	21.52	24.26
2094-742	42	14995	2033	26.30	27.51	24.50	31.41	32.38	29.29	21.91	24.70
2094-743	43	15272	2067	26.79	28.04	24.94	32.05	33.05	29.85	22.30	25.14
2094-744	44	15546	2101	27.27	28.55	25.38	32.69	33.71	30.39	22.67	25.57
2094-745	45	15816	2135	27.74	29.07	25.81	33.32	34.37	30.94	23.05	25.99
2094-746	46	16082	2168	28.21	29.57	26.23	33.95	35.03	31.47	23.42	26.40
2094-747	47	16345	2201	28.67	30.08	26.65	34.58	35.68	32.00	23.78	26.80
2094-748	48	16604	2233	29.12	30.57	27.06	35.20	36.33	32.51	24.14	27.19
2094-749	49	16859	2265	29.57	31.06	27.46	35.83	36.97	33.03	24.49	27.58
2094-750	50	17111	2295	30.01	31.55	27.86	36.44	37.61	33.53	24.84	27.96
2094-751	51	17360	2324	30.45	32.03	28.25	37.06	38.25	34.03	25.19	28.33
2094-752	52	17605	2353	30.88	32.51	28.63	37.67	38.88	34.52	25.52	28.69
2094-753	53	17847	2381	31.30	32.98	29.02	38.28	39.51	35.00	25.86	29.05
2094-754	54	18085	2409	31.72	33.44	29.39	38.88	40.13	35.48	26.19	29.40

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 10 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 EC											
2097-401	1	219	27	0.38	0.38	0.38	0.39	0.40	0.38	0.38	0.39
2097-402	2	414	51	0.73	0.73	0.72	0.76	0.77	0.73	0.69	0.73
2097-403	3	602	75	1.06	1.06	1.04	1.12	1.13	1.08	0.98	1.06
2097-404	4	782	97	1.37	1.39	1.34	1.47	1.48	1.41	1.25	1.37
2097-405	5	957	118	1.68	1.70	1.63	1.82	1.82	1.74	1.52	1.67
2097-406	6	1129	139	1.98	2.01	1.92	2.16	2.16	2.06	1.77	1.97
2097-407	7	1295	160	2.27	2.31	2.20	2.49	2.49	2.38	2.02	2.26
2097-408	8	1458	180	2.56	2.60	2.46	2.82	2.82	2.69	2.26	2.54
2097-409	9	1617	199	2.84	2.89	2.73	3.14	3.14	2.99	2.49	2.81
2097-410	10	1772	218	3.11	3.17	2.98	3.45	3.46	3.29	2.71	3.08
2097-411	11	1923	237	3.37	3.44	3.23	3.76	3.76	3.59	2.93	3.33
2097-412	12	2070	255	3.63	3.71	3.47	4.06	4.07	3.88	3.14	3.58
2097-501	1	473	58	0.83	0.83	0.83	0.85	0.86	0.82	0.81	0.83
2097-502	2	894	110	1.57	1.58	1.55	1.64	1.65	1.58	1.48	1.58
2097-503	3	1296	160	2.27	2.29	2.23	2.42	2.43	2.32	2.10	2.28
2097-504	4	1681	207	2.95	2.98	2.88	3.17	3.19	3.04	2.68	2.95
2097-505	5	2055	252	3.60	3.65	3.50	3.91	3.94	3.74	3.24	3.60
2097-506	6	2418	296	4.24	4.30	4.10	4.63	4.67	4.42	3.78	4.23
2097-507	7	2771	340	4.86	4.94	4.69	5.34	5.39	5.10	4.29	4.83
2097-508	8	3114	382	5.46	5.56	5.25	6.04	6.09	5.76	4.79	5.42
2097-509	9	3447	422	6.05	6.16	5.80	6.72	6.78	6.40	5.27	5.99
2097-510	10	3771	462	6.61	6.75	6.33	7.39	7.46	7.04	5.73	6.54
2097-511	11	4086	500	7.17	7.32	6.84	8.04	8.13	7.66	6.17	7.07
2097-512	12	4391	537	7.70	7.88	7.33	8.69	8.79	8.27	6.60	7.59
2097-601	1	640	79	1.12	1.12	1.12	1.15	1.16	1.11	1.10	1.13
2097-602	2	1208	149	2.12	2.13	2.09	2.22	2.24	2.14	2.00	2.13
2097-603	3	1749	215	3.07	3.10	3.00	3.26	3.30	3.14	2.83	3.07
2097-604	4	2264	277	3.97	4.02	3.87	4.28	4.32	4.10	3.61	3.96
2097-605	5	2765	339	4.85	4.92	4.70	5.27	5.33	5.04	4.35	4.83
2097-606	6	3250	398	5.70	5.79	5.51	6.24	6.32	5.96	5.07	5.67
2097-607	7	3720	455	6.52	6.64	6.28	7.20	7.29	6.86	5.75	6.47
2097-608	8	4177	511	7.32	7.46	7.04	8.13	8.25	7.74	6.41	7.26
2097-609	9	4619	565	8.10	8.26	7.76	9.04	9.19	8.60	7.04	8.01
2097-610	10	5048	617	8.85	9.04	8.46	9.93	10.11	9.44	7.65	8.74
2097-611	11	5465	668	9.58	9.80	9.14	10.81	11.01	10.27	8.23	9.45
2097-612	12	5869	717	10.29	10.54	9.79	11.66	11.90	11.07	8.78	10.14

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 11 of 11

All capacity numbers are relative to the IBM 2094-701 running MULTI-IMAGE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	PCI**	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 EC (continued)											
2097-701	1	923	115	1.62	1.62	1.62	1.66	1.67	1.60	1.59	1.63
2097-702	2	1735	215	3.04	3.07	3.00	3.20	3.23	3.08	2.86	3.05
2097-703	3	2506	312	4.39	4.44	4.30	4.69	4.75	4.50	4.04	4.39
2097-704	4	3237	401	5.68	5.75	5.52	6.14	6.23	5.87	5.13	5.67
2097-705	5	3944	488	6.92	7.02	6.70	7.55	7.69	7.20	6.17	6.90
2097-706	6	4626	571	8.11	8.24	7.83	8.93	9.12	8.50	7.16	8.07
2097-707	7	5285	651	9.27	9.43	8.91	10.28	10.52	9.77	8.11	9.20
2097-708	8	5921	729	10.38	10.59	9.95	11.59	11.89	11.00	9.02	10.29
2097-709	9	6535	804	11.46	11.70	10.95	12.88	13.25	12.21	9.88	11.34
2097-710	10	7129	875	12.50	12.79	11.91	14.14	14.57	13.38	10.70	12.34
2097-711	11	7702	944	13.51	13.83	12.84	15.37	15.88	14.53	11.49	13.31
2097-712	12	8256	1011	14.48	14.85	13.73	16.57	17.16	15.64	12.24	14.24
2097-713	13	8790	1076	15.42	15.83	14.58	17.74	18.41	16.73	12.95	15.13
2097-714	14	9319	1139	16.34	16.81	15.42	18.91	19.66	17.80	13.66	16.01
2097-715	15	9842	1202	17.26	17.77	16.25	20.06	20.89	18.87	14.36	16.88
2097-716	16	10360	1264	18.17	18.73	17.08	21.22	22.12	19.92	15.05	17.75
2097-717	17	10873	1329	19.07	19.67	17.90	22.36	23.33	20.96	15.73	18.60
2097-718	18	11381	1390	19.96	20.61	18.71	23.50	24.53	21.99	16.41	19.44
2097-719	19	11883	1451	20.84	21.54	19.51	24.63	25.72	23.01	17.08	20.28
2097-720	20	12380	1512	21.71	22.46	20.30	25.75	26.90	24.01	17.74	21.10
2097-721	21	12873	1571	22.58	23.37	21.08	26.87	28.07	25.01	18.40	21.92
2097-722	22	13360	1631	23.43	24.27	21.86	27.98	29.22	25.99	19.05	22.73
2097-723	23	13842	1690	24.28	25.17	22.63	29.08	30.37	26.96	19.69	23.53
2097-724	24	14320	1748	25.11	26.05	23.39	30.18	31.51	27.92	20.32	24.32
2097-725	25	14793	1805	25.94	26.93	24.14	31.27	32.64	28.88	20.95	25.10
2097-726	26	15261	1865	26.76	27.80	24.89	32.36	33.75	29.82	21.57	25.87
2097-727	27	15725	1922	27.58	28.67	25.63	33.44	34.86	30.75	22.19	26.64
2097-728	28	16188	1979	28.39	29.53	26.36	34.52	35.97	31.68	22.80	27.40
2097-729	29	16647	2037	29.20	30.38	27.09	35.59	37.07	32.60	23.41	28.16
2097-730	30	17103	2092	30.00	31.23	27.82	36.66	38.17	33.51	24.01	28.91
2097-731	31	17557	2146	30.79	32.08	28.54	37.73	39.27	34.42	24.61	29.66
2097-732	32	18008	2200	31.58	32.92	29.26	38.79	40.35	35.32	25.21	30.40
2097-733	33	18456	2257	32.37	33.76	29.97	39.86	41.44	36.22	25.80	31.13
2097-734	34	18901	2309	33.15	34.59	30.68	40.91	42.52	37.11	26.39	31.86
2097-735	35	19343	2366	33.92	35.42	31.38	41.97	43.60	37.99	26.97	32.59
2097-736	36	19783	2422	34.70	36.24	32.08	43.02	44.67	38.87	27.55	33.31
2097-737	37	20220	2478	35.46	37.06	32.77	44.07	45.74	39.74	28.13	34.03
2097-738	38	20654	2530	36.22	37.87	33.46	45.12	46.80	40.60	28.70	34.74
2097-739	39	21085	2585	36.98	38.68	34.15	46.16	47.86	41.46	29.27	35.44
2097-740	40	21514	2636	37.73	39.49	34.83	47.20	48.91	42.31	29.83	36.14
2097-741	41	21939	2687	38.48	40.29	35.50	48.23	49.96	43.16	30.39	36.83
2097-742	42	22361	2740	39.22	41.08	36.17	49.27	51.01	44.00	30.95	37.51
2097-743	43	22780	2789	39.95	41.88	36.83	50.29	52.06	44.84	31.51	38.17
2097-744	44	23196	2838	40.68	42.67	37.49	51.32	53.10	45.68	32.06	38.83
2097-745	45	23608	2886	41.40	43.45	38.14	52.34	54.14	46.51	32.61	39.47
2097-746	46	24017	2934	42.12	44.23	38.78	53.35	55.17	47.33	33.16	40.10
2097-747	47	24422	2981	42.83	45.01	39.42	54.36	56.20	48.16	33.70	40.73
2097-748	48	24825	3028	43.54	45.79	40.05	55.37	57.23	48.97	34.25	41.34
2097-749	49	25224	3075	44.24	46.56	40.68	56.37	58.26	49.79	34.79	41.94
2097-750	50	25621	3120	44.93	47.33	41.30	57.37	59.28	50.60	35.32	42.53
2097-751	51	26014	3166	45.62	48.09	41.91	58.36	60.30	51.40	35.86	43.11
2097-752	52	26404	3214	46.31	48.85	42.52	59.36	61.32	52.20	36.39	43.68
2097-753	53	26791	3262	46.99	49.61	43.12	60.34	62.33	53.00	36.92	44.24
2097-754	54	27175	3305	47.66	50.36	43.72	61.32	63.34	53.79	37.45	44.79
2097-755	55	27557	3352	48.33	51.11	44.31	62.30	64.35	54.59	37.97	45.34
2097-756	56	27935	3395	48.99	51.86	44.90	63.28	65.35	55.37	38.49	45.87
2097-757	57	28309	3438	49.65	52.59	45.48	64.25	66.35	56.15	39.01	46.39
2097-758	58	28679	3480	50.30	53.33	46.05	65.21	67.34	56.92	39.52	46.91
2097-759	59	29044	3525	50.94	54.05	46.61	66.17	68.33	57.68	40.02	47.41
2097-760	60	29406	3570	51.57	54.77	47.17	67.12	69.31	58.43	40.52	47.91
2097-761	61	29763	3611	52.20	55.49	47.72	68.07	70.28	59.18	41.02	48.40
2097-762	62	30116	3652	52.82	56.20	48.27	69.01	71.25	59.92	41.51	48.88
2097-763	63	30466	3695	53.43	56.90	48.80	69.94	72.21	60.65	41.99	49.36
2097-764	64	30811	3739	54.04	57.59	49.33	70.87	73.17	61.38	42.47	49.82

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

**** PCI stands for Processor capacity Index. The PCI values were calculated by multiplying the LSPR Mixed column by a common scaling factor associated with a particular LSPR table. Note the values appearing here were generated using zPCR so the full precision of each ITRR ratio is represented.**

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Table 7. ITR Ratio Table -z/OS V1R9 Single-Image

The tables in this section present LSPR capacity data in the form of ITR ratios for IBM processors where each model is configured with a single z/OS V1R9 image equal in size to the physical Nway (#CP) of the model. All capacity numbers are relative to the IBM 2094-701 running a single z/OS image. All workloads were run in 64 bit real addressing mode.

Part 1 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM e(logo)Server zSeries 800										
2066-0E1	1	7	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
2066-0A1	1	13	0.12	0.12	0.12	0.12	0.12	0.13	0.12	0.12
2066-0B1	1	20	0.18	0.18	0.18	0.17	0.18	0.18	0.18	0.18
2066-0C1	1	25	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.23
2066-0X2	2	28	0.25	0.26	0.25	0.26	0.26	0.25	0.25	0.25
2066-001	1	32	0.30	0.29	0.30	0.29	0.30	0.29	0.30	0.31
2066-0A2	2	44	0.41	0.41	0.41	0.42	0.43	0.42	0.40	0.41
2066-002	2	60	0.56	0.56	0.56	0.57	0.57	0.56	0.55	0.57
2066-003	3	84	0.81	0.80	0.80	0.82	0.83	0.82	0.77	0.83
2066-004	4	108	1.04	1.04	1.03	1.04	1.06	1.08	0.99	1.06
IBM e(logo)Server zSeries 890										
2086-110	1	4	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.05
2086-210	2	8	0.08	0.08	0.08	0.09	0.09	0.08	0.08	0.09
2086-310	3	11	0.12	0.12	0.12	0.13	0.13	0.12	0.11	0.13
2086-410	4	15	0.16	0.16	0.16	0.17	0.17	0.16	0.15	0.17
2086-120	1	7	0.08	0.07	0.08	0.07	0.08	0.07	0.07	0.08
2086-220	2	13	0.15	0.15	0.15	0.15	0.15	0.15	0.14	0.16
2086-320	3	20	0.22	0.22	0.22	0.22	0.22	0.22	0.21	0.23
2086-420	4	26	0.29	0.29	0.29	0.29	0.30	0.29	0.28	0.30
2086-130	1	13	0.15	0.14	0.15	0.14	0.14	0.15	0.14	0.15
2086-230	2	26	0.28	0.28	0.28	0.28	0.28	0.29	0.27	0.29
2086-330	3	38	0.42	0.42	0.42	0.42	0.42	0.44	0.40	0.43
2086-430	4	49	0.56	0.55	0.55	0.55	0.57	0.59	0.53	0.57
2086-140	1	17	0.18	0.18	0.18	0.18	0.18	0.19	0.18	0.19
2086-240	2	32	0.36	0.35	0.35	0.35	0.36	0.37	0.34	0.37
2086-340	3	47	0.53	0.52	0.52	0.52	0.53	0.55	0.50	0.54
2086-440	4	62	0.70	0.69	0.69	0.69	0.71	0.73	0.66	0.71
2086-150	1	26	0.28	0.28	0.28	0.28	0.28	0.29	0.28	0.30
2086-250	2	50	0.56	0.55	0.55	0.55	0.56	0.57	0.53	0.57
2086-350	3	74	0.82	0.82	0.81	0.81	0.83	0.86	0.78	0.84
2086-450	4	97	1.09	1.08	1.07	1.08	1.10	1.14	1.04	1.10
2086-160	1	32	0.35	0.34	0.35	0.34	0.34	0.36	0.34	0.36
2086-260	2	62	0.68	0.67	0.67	0.67	0.68	0.70	0.65	0.70
2086-360	3	91	1.01	1.00	0.99	0.99	1.01	1.05	0.96	1.03
2086-460	4	119	1.33	1.32	1.31	1.32	1.34	1.40	1.27	1.34
2086-170	1	56	0.60	0.60	0.60	0.59	0.60	0.61	0.59	0.62
2086-270	2	107	1.18	1.17	1.17	1.16	1.17	1.21	1.14	1.21
2086-370	3	158	1.75	1.73	1.73	1.72	1.75	1.81	1.68	1.78
2086-470	4	208	2.31	2.30	2.27	2.29	2.33	2.41	2.22	2.32

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 2 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 BC R07										
2096-A01	1	4	0.05	0.05	0.05	0.05	0.05	0.04	0.04	0.05
2096-A02	2	7	0.09	0.09	0.09	0.10	0.10	0.08	0.08	0.10
2096-A03	3	10	0.13	0.13	0.14	0.15	0.15	0.12	0.12	0.15
2096-B01	1	5	0.07	0.06	0.07	0.07	0.07	0.06	0.06	0.08
2096-B02	2	10	0.13	0.13	0.14	0.13	0.13	0.12	0.12	0.15
2096-B03	3	15	0.19	0.19	0.20	0.20	0.20	0.18	0.18	0.22
2096-C01	1	6	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08
2096-C02	2	12	0.16	0.16	0.16	0.16	0.16	0.15	0.15	0.16
2096-C03	3	18	0.23	0.23	0.23	0.24	0.24	0.23	0.22	0.24
2096-D01	1	8	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.11
2096-D02	2	16	0.20	0.20	0.20	0.20	0.20	0.20	0.20	0.21
2096-D03	3	23	0.30	0.29	0.30	0.30	0.30	0.29	0.28	0.30
2096-E01	1	10	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.13
2096-E02	2	19	0.24	0.24	0.24	0.24	0.24	0.24	0.23	0.25
2096-F01	1	12	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.16
2096-F02	2	24	0.30	0.29	0.30	0.29	0.30	0.30	0.29	0.30
2096-G01	1	15	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.20
2096-H01	1	18	0.23	0.22	0.23	0.22	0.22	0.23	0.22	0.23
2096-I01	1	21	0.26	0.26	0.26	0.26	0.26	0.26	0.26	0.27
2096-J01	1	24	0.30	0.30	0.30	0.29	0.29	0.30	0.30	0.31

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 3 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 BC S07										
2096-K04	4	30	0.39	0.39	0.39	0.40	0.40	0.39	0.37	0.40
2096-L03	3	28	0.35	0.35	0.35	0.36	0.35	0.35	0.34	0.36
2096-L04	4	36	0.46	0.46	0.46	0.47	0.47	0.46	0.44	0.47
2096-M03	3	34	0.44	0.43	0.43	0.44	0.44	0.44	0.42	0.45
2096-M04	4	45	0.57	0.57	0.57	0.58	0.58	0.58	0.54	0.58
2096-N02	2	30	0.37	0.37	0.37	0.37	0.37	0.38	0.37	0.38
2096-N03	3	43	0.55	0.55	0.55	0.55	0.56	0.56	0.53	0.56
2096-N04	4	56	0.72	0.72	0.71	0.73	0.73	0.73	0.69	0.74
2096-O02	2	36	0.44	0.44	0.44	0.44	0.44	0.44	0.43	0.45
2096-O03	3	52	0.65	0.65	0.65	0.65	0.65	0.66	0.63	0.67
2096-O04	4	67	0.85	0.85	0.84	0.86	0.86	0.87	0.81	0.87
2096-P02	2	41	0.51	0.51	0.51	0.51	0.50	0.51	0.50	0.52
2096-P03	3	59	0.75	0.75	0.75	0.75	0.75	0.76	0.73	0.77
2096-P04	4	77	0.98	0.98	0.97	0.99	0.99	1.00	0.94	1.00
2096-Q02	2	47	0.58	0.58	0.58	0.58	0.58	0.59	0.57	0.60
2096-Q03	3	68	0.86	0.86	0.86	0.86	0.86	0.87	0.83	0.88
2096-Q04	4	88	1.13	1.12	1.12	1.13	1.14	1.15	1.08	1.15
2096-R01	1	27	0.33	0.33	0.34	0.33	0.33	0.33	0.33	0.34
2096-R02	2	52	0.66	0.65	0.66	0.65	0.65	0.66	0.64	0.67
2096-R03	3	76	0.97	0.96	0.96	0.96	0.97	0.98	0.94	0.99
2096-R04	4	99	1.26	1.26	1.25	1.27	1.28	1.29	1.21	1.29
2096-S01	1	30	0.37	0.37	0.38	0.37	0.37	0.38	0.37	0.38
2096-S02	2	59	0.73	0.73	0.73	0.73	0.73	0.74	0.72	0.75
2096-S03	3	85	1.08	1.07	1.07	1.08	1.08	1.10	1.05	1.10
2096-S04	4	111	1.41	1.41	1.40	1.42	1.43	1.45	1.36	1.44
2096-T01	1	34	0.42	0.42	0.42	0.41	0.41	0.42	0.42	0.43
2096-T02	2	66	0.82	0.82	0.82	0.81	0.82	0.83	0.81	0.84
2096-T03	3	95	1.21	1.20	1.20	1.21	1.21	1.23	1.18	1.23
2096-T04	4	124	1.58	1.58	1.57	1.59	1.60	1.62	1.52	1.61
2096-U01	1	38	0.47	0.46	0.47	0.46	0.46	0.47	0.46	0.48
2096-U02	2	73	0.92	0.91	0.92	0.91	0.91	0.93	0.90	0.94
2096-U03	3	106	1.35	1.34	1.34	1.34	1.35	1.38	1.31	1.37
2096-U04	4	138	1.77	1.76	1.75	1.77	1.79	1.81	1.70	1.79
2096-V01	1	42	0.52	0.52	0.53	0.51	0.52	0.53	0.52	0.54
2096-V02	2	82	1.03	1.02	1.03	1.01	1.02	1.04	1.01	1.05
2096-V03	3	119	1.52	1.51	1.51	1.51	1.52	1.55	1.47	1.54
2096-V04	4	155	1.98	1.98	1.96	1.99	2.01	2.04	1.91	2.01
2096-W01	1	47	0.59	0.58	0.59	0.58	0.58	0.59	0.59	0.60
2096-W02	2	92	1.16	1.15	1.15	1.14	1.15	1.17	1.14	1.17
2096-W03	3	134	1.70	1.69	1.69	1.69	1.71	1.74	1.66	1.72
2096-W04	4	174	2.22	2.22	2.20	2.22	2.26	2.29	2.14	2.25
2096-X01	1	53	0.66	0.65	0.66	0.64	0.64	0.66	0.66	0.67
2096-X02	2	103	1.29	1.28	1.29	1.27	1.27	1.31	1.27	1.31
2096-X03	3	150	1.90	1.89	1.89	1.88	1.89	1.95	1.85	1.92
2096-X04	4	195	2.49	2.48	2.46	2.48	2.50	2.56	2.40	2.51
2096-Y01	1	59	0.73	0.72	0.73	0.71	0.71	0.74	0.73	0.74
2096-Y02	2	115	1.43	1.42	1.43	1.41	1.41	1.46	1.42	1.45
2096-Y03	3	166	2.11	2.10	2.09	2.09	2.10	2.16	2.06	2.13
2096-Y04	4	216	2.76	2.75	2.73	2.75	2.78	2.85	2.67	2.78
2096-Z01	1	67	0.83	0.82	0.83	0.81	0.81	0.84	0.83	0.84
2096-Z02	2	130	1.63	1.62	1.62	1.60	1.61	1.67	1.62	1.65
2096-Z03	3	189	2.40	2.39	2.38	2.37	2.39	2.47	2.35	2.42
2096-Z04	4	246	3.14	3.14	3.10	3.12	3.16	3.25	3.04	3.15

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 4 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 BC										
2098-A01	1	3	0.05	0.05	0.05	0.06	0.05	0.04	0.04	0.06
2098-A02	2	6	0.09	0.09	0.08	0.12	0.10	0.08	0.07	0.10
2098-A03	3	9	0.12	0.12	0.12	0.18	0.15	0.11	0.09	0.14
2098-A04	4	11	0.16	0.16	0.15	0.23	0.20	0.15	0.12	0.18
2098-A05	5	13	0.19	0.19	0.18	0.28	0.24	0.18	0.14	0.21
2098-B01	1	4	0.05	0.05	0.05	0.07	0.06	0.05	0.04	0.06
2098-B02	2	7	0.10	0.10	0.10	0.13	0.12	0.09	0.08	0.11
2098-B03	3	10	0.14	0.14	0.14	0.19	0.17	0.13	0.11	0.16
2098-B04	4	13	0.18	0.18	0.17	0.25	0.22	0.17	0.14	0.20
2098-B05	5	15	0.22	0.23	0.21	0.31	0.27	0.21	0.16	0.24
2098-C01	1	5	0.07	0.07	0.07	0.08	0.07	0.06	0.06	0.07
2098-C02	2	9	0.12	0.13	0.12	0.15	0.14	0.12	0.10	0.13
2098-C03	3	12	0.18	0.18	0.17	0.22	0.20	0.18	0.14	0.19
2098-C04	4	16	0.23	0.24	0.22	0.29	0.26	0.24	0.18	0.24
2098-C05	5	19	0.28	0.29	0.26	0.36	0.32	0.29	0.22	0.28
2098-D01	1	6	0.08	0.08	0.08	0.09	0.09	0.08	0.07	0.09
2098-D02	2	11	0.15	0.15	0.15	0.18	0.17	0.15	0.12	0.16
2098-D03	3	15	0.22	0.22	0.21	0.27	0.24	0.22	0.17	0.23
2098-D04	4	19	0.28	0.29	0.26	0.35	0.32	0.29	0.22	0.29
2098-D05	5	23	0.34	0.35	0.32	0.43	0.39	0.35	0.26	0.34
2098-E01	1	7	0.09	0.09	0.09	0.11	0.10	0.09	0.08	0.10
2098-E02	2	12	0.17	0.18	0.17	0.21	0.19	0.17	0.14	0.19
2098-E03	3	17	0.25	0.25	0.24	0.31	0.28	0.25	0.20	0.26
2098-E04	4	22	0.32	0.33	0.30	0.40	0.37	0.33	0.25	0.33
2098-E05	5	27	0.39	0.40	0.37	0.50	0.45	0.41	0.30	0.39
2098-F01	1	7	0.10	0.10	0.10	0.12	0.11	0.10	0.09	0.11
2098-F02	2	14	0.19	0.20	0.19	0.23	0.21	0.19	0.16	0.21
2098-F03	3	19	0.28	0.28	0.27	0.34	0.31	0.28	0.22	0.29
2098-F04	4	25	0.36	0.37	0.34	0.45	0.41	0.37	0.28	0.37
2098-F05	5	30	0.43	0.45	0.41	0.55	0.50	0.46	0.34	0.44
2098-G01	1	9	0.12	0.12	0.12	0.14	0.13	0.12	0.10	0.14
2098-G02	2	16	0.23	0.23	0.22	0.27	0.26	0.23	0.19	0.24
2098-G03	3	23	0.33	0.34	0.31	0.40	0.37	0.34	0.26	0.34
2098-G04	4	29	0.42	0.43	0.40	0.53	0.49	0.44	0.33	0.43
2098-G05	5	36	0.51	0.53	0.48	0.65	0.60	0.54	0.40	0.52
2098-H01	1	10	0.14	0.14	0.14	0.16	0.15	0.13	0.12	0.15
2098-H02	2	18	0.26	0.26	0.25	0.31	0.28	0.26	0.21	0.27
2098-H03	3	26	0.37	0.37	0.35	0.45	0.42	0.38	0.29	0.38
2098-H04	4	33	0.47	0.48	0.45	0.59	0.54	0.49	0.37	0.48
2098-H05	5	40	0.57	0.59	0.54	0.72	0.66	0.61	0.45	0.58
2098-I01	1	11	0.16	0.15	0.15	0.18	0.17	0.15	0.13	0.17
2098-I02	2	20	0.29	0.29	0.28	0.34	0.32	0.29	0.24	0.31
2098-I03	3	29	0.41	0.42	0.39	0.51	0.47	0.43	0.33	0.43
2098-I04	4	37	0.53	0.55	0.50	0.66	0.61	0.56	0.42	0.54
2098-I05	5	45	0.65	0.67	0.60	0.81	0.75	0.69	0.50	0.65
2098-J01	1	12	0.17	0.17	0.17	0.20	0.19	0.17	0.15	0.19
2098-J02	2	23	0.32	0.32	0.31	0.38	0.36	0.33	0.26	0.34
2098-J03	3	32	0.46	0.47	0.44	0.56	0.52	0.48	0.37	0.48
2098-J04	4	41	0.59	0.61	0.56	0.73	0.68	0.63	0.47	0.60
2098-J05	5	50	0.72	0.74	0.67	0.90	0.84	0.77	0.56	0.72
2098-K01	1	14	0.19	0.19	0.19	0.22	0.21	0.19	0.16	0.21
2098-K02	2	25	0.36	0.36	0.35	0.43	0.40	0.37	0.29	0.38
2098-K03	3	36	0.52	0.53	0.49	0.63	0.59	0.54	0.41	0.54
2098-K04	4	46	0.66	0.68	0.62	0.82	0.77	0.71	0.52	0.68
2098-K05	5	56	0.81	0.83	0.75	1.01	0.94	0.87	0.63	0.81
2098-L01	1	16	0.23	0.23	0.22	0.26	0.25	0.23	0.19	0.25
2098-L02	2	30	0.43	0.43	0.41	0.50	0.48	0.44	0.35	0.45
2098-L03	3	43	0.61	0.62	0.58	0.74	0.70	0.65	0.49	0.63
2098-L04	4	55	0.79	0.81	0.74	0.97	0.91	0.85	0.62	0.80
2098-L05	5	66	0.95	0.98	0.89	1.19	1.11	1.04	0.74	0.96
2098-M01	1	19	0.26	0.26	0.26	0.30	0.29	0.27	0.22	0.29
2098-M02	2	35	0.49	0.50	0.47	0.58	0.55	0.51	0.40	0.52
2098-M03	3	49	0.71	0.72	0.67	0.85	0.81	0.75	0.56	0.73
2098-M04	4	63	0.91	0.93	0.85	1.11	1.05	0.99	0.71	0.92
2098-M05	5	76	1.10	1.13	1.02	1.37	1.29	1.22	0.85	1.10

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 5 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 BC (continued)										
2098-N01	1	21	0.30	0.30	0.29	0.34	0.33	0.31	0.25	0.33
2098-N02	2	40	0.56	0.57	0.54	0.66	0.64	0.60	0.46	0.59
2098-N03	3	57	0.81	0.82	0.76	0.97	0.93	0.88	0.64	0.83
2098-N04	4	72	1.04	1.07	0.97	1.26	1.21	1.15	0.81	1.05
2098-N05	5	87	1.26	1.30	1.16	1.56	1.48	1.41	0.97	1.26
2098-O01	1	24	0.34	0.34	0.33	0.38	0.37	0.35	0.28	0.37
2098-O02	2	44	0.63	0.64	0.60	0.73	0.72	0.68	0.51	0.66
2098-O03	3	63	0.91	0.93	0.85	1.08	1.05	0.99	0.72	0.93
2098-O04	4	81	1.16	1.20	1.08	1.41	1.37	1.30	0.91	1.17
2098-O05	5	98	1.41	1.46	1.30	1.74	1.67	1.61	1.09	1.40
2098-P01	1	27	0.38	0.38	0.37	0.42	0.42	0.40	0.32	0.41
2098-P02	2	50	0.71	0.72	0.67	0.82	0.80	0.77	0.57	0.73
2098-P03	3	71	1.02	1.04	0.95	1.20	1.18	1.13	0.80	1.03
2098-P04	4	91	1.30	1.34	1.20	1.57	1.53	1.48	1.01	1.30
2098-P05	5	110	1.58	1.63	1.45	1.93	1.88	1.82	1.21	1.56
2098-Q01	1	30	0.43	0.43	0.41	0.47	0.47	0.45	0.35	0.46
2098-Q02	2	56	0.79	0.80	0.75	0.91	0.90	0.86	0.64	0.82
2098-Q03	3	80	1.14	1.16	1.06	1.34	1.32	1.26	0.90	1.16
2098-Q04	4	102	1.46	1.50	1.35	1.75	1.72	1.66	1.14	1.46
2098-Q05	5	123	1.77	1.83	1.63	2.16	2.10	2.04	1.36	1.75
2098-R01	1	33	0.48	0.47	0.46	0.53	0.52	0.50	0.40	0.51
2098-R02	2	62	0.89	0.90	0.84	1.02	1.00	0.96	0.72	0.92
2098-R03	3	89	1.27	1.30	1.19	1.49	1.47	1.41	1.01	1.29
2098-R04	4	113	1.63	1.67	1.51	1.95	1.92	1.84	1.27	1.63
2098-R05	5	137	1.98	2.04	1.82	2.41	2.34	2.27	1.53	1.95
2098-S01	1	38	0.53	0.53	0.52	0.59	0.58	0.56	0.45	0.57
2098-S02	2	70	0.99	1.01	0.94	1.14	1.13	1.07	0.81	1.03
2098-S03	3	100	1.42	1.45	1.33	1.67	1.65	1.58	1.13	1.44
2098-S04	4	127	1.83	1.88	1.69	2.19	2.15	2.07	1.43	1.82
2098-S05	5	154	2.22	2.29	2.04	2.70	2.63	2.54	1.72	2.18
2098-T01	1	42	0.60	0.60	0.58	0.66	0.65	0.62	0.51	0.64
2098-T02	2	78	1.11	1.13	1.06	1.28	1.26	1.20	0.91	1.15
2098-T03	3	112	1.60	1.63	1.50	1.87	1.85	1.76	1.28	1.61
2098-T04	4	143	2.05	2.11	1.90	2.45	2.41	2.31	1.62	2.04
2098-T05	5	173	2.49	2.57	2.29	3.02	2.95	2.85	1.94	2.44
2098-U01	1	47	0.67	0.67	0.65	0.74	0.73	0.70	0.57	0.71
2098-U02	2	88	1.25	1.27	1.19	1.43	1.41	1.34	1.02	1.28
2098-U03	3	125	1.79	1.83	1.68	2.09	2.07	1.97	1.44	1.80
2098-U04	4	160	2.30	2.37	2.14	2.74	2.70	2.59	1.82	2.28
2098-U05	5	194	2.79	2.89	2.57	3.38	3.30	3.18	2.18	2.73
2098-V01	1	53	0.76	0.76	0.73	0.83	0.82	0.78	0.65	0.80
2098-V02	2	99	1.41	1.43	1.34	1.60	1.59	1.51	1.16	1.44
2098-V03	3	141	2.02	2.06	1.90	2.35	2.33	2.21	1.63	2.03
2098-V04	4	180	2.59	2.67	2.41	3.08	3.03	2.90	2.06	2.56
2098-V05	5	218	3.14	3.25	2.90	3.79	3.71	3.57	2.48	3.07
2098-W01	1	60	0.85	0.85	0.82	0.92	0.92	0.87	0.73	0.89
2098-W02	2	111	1.57	1.60	1.50	1.79	1.78	1.68	1.31	1.61
2098-W03	3	158	2.26	2.31	2.13	2.62	2.60	2.47	1.84	2.26
2098-W04	4	202	2.90	2.99	2.70	3.44	3.39	3.24	2.33	2.85
2098-W05	5	245	3.52	3.64	3.25	4.23	4.15	3.99	2.79	3.42
2098-X01	1	67	0.95	0.95	0.92	1.03	1.03	0.97	0.82	1.00
2098-X02	2	124	1.76	1.79	1.69	2.00	1.99	1.88	1.48	1.79
2098-X03	3	177	2.53	2.59	2.39	2.93	2.91	2.76	2.08	2.52
2098-X04	4	226	3.25	3.35	3.03	3.84	3.80	3.62	2.63	3.19
2098-X05	5	274	3.94	4.09	3.65	4.73	4.65	4.45	3.15	3.82
2098-Y01	1	76	1.08	1.08	1.06	1.17	1.18	1.11	0.95	1.13
2098-Y02	2	142	2.02	2.05	1.93	2.27	2.27	2.14	1.71	2.04
2098-Y03	3	202	2.89	2.96	2.74	3.34	3.33	3.14	2.40	2.87
2098-Y04	4	258	3.72	3.83	3.48	4.37	4.34	4.11	3.04	3.63
2098-Y05	5	313	4.51	4.68	4.19	5.39	5.30	5.07	3.65	4.35
2098-Z01	1	83	1.19	1.19	1.16	1.28	1.29	1.21	1.05	1.24
2098-Z02	2	155	2.21	2.24	2.12	2.48	2.48	2.33	1.89	2.23
2098-Z03	3	221	3.17	3.24	3.00	3.64	3.63	3.42	2.65	3.13
2098-Z04	4	283	4.06	4.20	3.81	4.77	4.74	4.48	3.35	3.95
2098-Z05	5	342	4.93	5.12	4.59	5.87	5.79	5.52	4.02	4.74

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 6 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM e(logo)Server zSeries 900 (1XX Series)										
2064-101	1	41	0.38	0.37	0.38	0.37	0.37	0.37	0.38	0.38
2064-102	2	78	0.73	0.73	0.73	0.73	0.73	0.72	0.72	0.73
2064-103	3	112	1.06	1.06	1.06	1.08	1.08	1.07	1.03	1.07
2064-104	4	143	1.38	1.39	1.37	1.43	1.42	1.41	1.33	1.38
2064-105	5	173	1.69	1.71	1.66	1.76	1.76	1.74	1.61	1.68
2064-106	6	199	1.99	2.01	1.94	2.08	2.08	2.07	1.86	1.96
2064-107	7	225	2.27	2.30	2.20	2.38	2.40	2.40	2.10	2.22
2064-108	8	245	2.54	2.58	2.44	2.68	2.70	2.73	2.32	2.46
2064-109	9	265	2.79	2.85	2.67	2.97	2.99	3.04	2.53	2.68
2064-1C1	1	43	0.39	0.38	0.39	0.38	0.38	0.39	0.39	0.40
2064-1C2	2	83	0.76	0.75	0.76	0.75	0.75	0.77	0.75	0.78
2064-1C3	3	119	1.13	1.12	1.12	1.12	1.12	1.15	1.10	1.14
2064-1C4	4	153	1.48	1.48	1.47	1.48	1.49	1.52	1.43	1.50
2064-1C5	5	187	1.83	1.83	1.80	1.84	1.86	1.89	1.75	1.85
2064-1C6	6	217	2.17	2.17	2.13	2.20	2.23	2.25	2.05	2.19
2064-1C7	7	247	2.50	2.51	2.45	2.55	2.60	2.61	2.34	2.51
2064-1C8	8	276	2.82	2.83	2.76	2.90	2.96	2.96	2.62	2.83
2064-1C9	9	302	3.14	3.16	3.06	3.25	3.32	3.31	2.89	3.14
2064-110	10	327	3.45	3.47	3.34	3.59	3.69	3.65	3.15	3.44
2064-111	11	350	3.74	3.78	3.62	3.93	4.04	3.98	3.39	3.72
2064-112	12	372	4.03	4.08	3.89	4.27	4.40	4.31	3.63	3.99
2064-113	13	392	4.31	4.37	4.14	4.60	4.75	4.63	3.85	4.25
2064-114	14	410	4.58	4.65	4.39	4.92	5.10	4.94	4.06	4.50
2064-115	15	426	4.84	4.93	4.62	5.25	5.45	5.24	4.26	4.74
2064-116	16	441	5.09	5.19	4.84	5.57	5.79	5.54	4.45	4.96
IBM e(logo)Server zSeries 900 (2XX Series)										
2064-2C1	1	52	0.47	0.47	0.48	0.46	0.46	0.47	0.48	0.49
2064-2C2	2	100	0.92	0.92	0.92	0.91	0.92	0.93	0.91	0.93
2064-2C3	3	144	1.36	1.36	1.35	1.36	1.37	1.39	1.32	1.37
2064-2C4	4	184	1.79	1.79	1.76	1.80	1.82	1.84	1.72	1.79
2064-2C5	5	224	2.20	2.21	2.16	2.24	2.26	2.28	2.10	2.19
2064-2C6	6	260	2.61	2.62	2.55	2.67	2.71	2.72	2.46	2.59
2064-2C7	7	296	3.00	3.03	2.93	3.10	3.15	3.15	2.81	2.97
2064-2C8	8	330	3.39	3.42	3.29	3.52	3.59	3.58	3.15	3.34
2064-2C9	9	362	3.76	3.81	3.64	3.94	4.02	3.99	3.47	3.69
2064-210	10	392	4.12	4.18	3.98	4.34	4.45	4.40	3.77	4.03
2064-211	11	420	4.47	4.55	4.30	4.75	4.87	4.79	4.06	4.35
2064-212	12	445	4.80	4.90	4.61	5.14	5.29	5.18	4.34	4.65
2064-213	13	475	5.13	5.24	4.90	5.54	5.70	5.56	4.60	4.94
2064-214	14	497	5.44	5.58	5.18	5.92	6.11	5.93	4.85	5.21
2064-215	15	517	5.74	5.90	5.45	6.30	6.51	6.28	5.09	5.47
2064-216	16	535	6.02	6.21	5.70	6.67	6.90	6.63	5.31	5.71

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 7 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

							z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9
	Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	CB-L	ODE-B	WASDB	OLTP-W	OLTP-T
IBM e(logo)Server zSeries 990											
	2084-301	1	70	0.74	0.74	0.74	0.72	0.72	0.76	0.75	0.75
	2084-302	2	132	1.46	1.45	1.45	1.43	1.44	1.50	1.45	1.47
	2084-303	3	191	2.15	2.14	2.13	2.12	2.14	2.22	2.11	2.16
	2084-304	4	248	2.82	2.81	2.78	2.80	2.83	2.93	2.73	2.81
	2084-305	5	302	3.46	3.46	3.39	3.46	3.52	3.62	3.32	3.44
	2084-306	6	352	4.07	4.09	3.98	4.12	4.19	4.29	3.87	4.04
	2084-307	7	402	4.67	4.70	4.55	4.76	4.85	4.95	4.39	4.62
	2084-308	8	448	5.24	5.29	5.09	5.39	5.50	5.59	4.87	5.17
	2084-309	9	492	5.80	5.86	5.60	6.01	6.14	6.22	5.34	5.69
	2084-310	10	538	6.33	6.42	6.10	6.62	6.77	6.84	5.78	6.20
	2084-311	11	580	6.85	6.96	6.58	7.22	7.40	7.44	6.20	6.69
	2084-312	12	620	7.36	7.49	7.04	7.82	8.02	8.03	6.60	7.16
	2084-313	13	661	7.85	8.00	7.48	8.40	8.63	8.61	6.98	7.61
	2084-314	14	696	8.32	8.50	7.91	8.98	9.23	9.17	7.34	8.05
	2084-315	15	730	8.78	8.98	8.32	9.55	9.83	9.72	7.69	8.47
	2084-316	16	761	9.22	9.46	8.71	10.11	10.42	10.26	8.02	8.87
	2084-317	17	799	9.65	9.92	9.10	10.67	11.00	10.78	8.34	9.26
	2084-318	18	837	10.08	10.37	9.48	11.21	11.57	11.29	8.65	9.65
	2084-319	19	878	10.50	10.82	9.85	11.75	12.14	11.80	8.96	10.04
	2084-320	20	919	10.91	11.26	10.22	12.28	12.70	12.29	9.27	10.41
	2084-321	21	959	11.31	11.69	10.58	12.80	13.24	12.77	9.57	10.79
	2084-322	22	999	11.71	12.11	10.94	13.31	13.78	13.24	9.87	11.15
	2084-323	23	1037	12.10	12.53	11.30	13.82	14.31	13.70	10.17	11.51
	2084-324	24	1076	12.49	12.94	11.64	14.32	14.83	14.16	10.47	11.87
	2084-325	25	1114	12.87	13.35	11.99	14.81	15.34	14.60	10.76	12.22
	2084-326	26	1151	13.25	13.75	12.33	15.30	15.85	15.04	11.05	12.57
	2084-327	27	1188	13.62	14.14	12.67	15.77	16.34	15.46	11.34	12.91
	2084-328	28	1225	13.98	14.53	13.00	16.24	16.83	15.88	11.62	13.24
	2084-329	29	1261	14.34	14.91	13.32	16.70	17.31	16.29	11.90	13.57
	2084-330	30	1296	14.69	15.29	13.64	17.16	17.78	16.70	12.18	13.89
	2084-331	31	1332	15.04	15.66	13.96	17.61	18.24	17.10	12.45	14.21
	2084-332	32	1365	15.38	16.03	14.27	18.05	18.70	17.49	12.72	14.52

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

Part 8 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008 ******see below**

							z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9	z/OS V1R9
	Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	CB-L	ODE-B	WASDB	OLTP-W	OLTP-T
IBM System z9 EC											
	2094-401	1	28	0.34	0.34	0.35	0.34	0.34	0.35	0.34	0.35
	2094-402	2	54	0.68	0.67	0.68	0.67	0.67	0.69	0.67	0.69
	2094-403	3	78	1.00	0.99	0.99	0.99	1.00	1.02	0.97	1.01
	2094-404	4	102	1.30	1.30	1.29	1.30	1.32	1.34	1.25	1.32
	2094-405	5	124	1.60	1.60	1.58	1.61	1.63	1.65	1.52	1.62
	2094-406	6	144	1.88	1.88	1.85	1.91	1.94	1.96	1.77	1.90
	2094-407	7	164	2.16	2.16	2.11	2.21	2.25	2.25	2.01	2.17
	2094-408	8	182	2.42	2.43	2.36	2.50	2.55	2.55	2.23	2.42
	2094-501	1	53	0.67	0.66	0.67	0.65	0.66	0.68	0.67	0.68
	2094-502	2	104	1.31	1.30	1.31	1.29	1.30	1.34	1.30	1.33
	2094-503	3	152	1.93	1.92	1.92	1.91	1.93	1.98	1.89	1.95
	2094-504	4	197	2.53	2.52	2.50	2.52	2.56	2.61	2.44	2.54
	2094-505	5	240	3.10	3.10	3.05	3.11	3.17	3.22	2.96	3.11
	2094-506	6	279	3.65	3.66	3.58	3.70	3.77	3.82	3.45	3.65
	2094-507	7	317	4.18	4.20	4.08	4.27	4.36	4.40	3.91	4.17
	2094-508	8	352	4.69	4.72	4.56	4.82	4.94	4.97	4.35	4.67
	2094-601	1	65	0.81	0.80	0.81	0.79	0.79	0.82	0.81	0.82
	2094-602	2	127	1.59	1.58	1.58	1.56	1.57	1.63	1.57	1.61
	2094-603	3	184	2.34	2.33	2.32	2.31	2.33	2.41	2.29	2.36
	2094-604	4	240	3.06	3.06	3.02	3.05	3.08	3.17	2.97	3.07
	2094-605	5	292	3.76	3.76	3.70	3.77	3.81	3.92	3.60	3.76
	2094-606	6	339	4.43	4.44	4.34	4.47	4.54	4.64	4.19	4.42
	2094-607	7	385	5.07	5.09	4.95	5.16	5.25	5.35	4.75	5.04
	2094-608	8	428	5.69	5.72	5.53	5.84	5.95	6.04	5.28	5.64

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 9 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z9 EC (continued)										
2094-701	1	81	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
2094-702	2	158	1.96	1.97	1.95	1.98	1.98	1.98	1.94	1.96
2094-703	3	229	2.89	2.90	2.86	2.94	2.95	2.93	2.82	2.87
2094-704	4	298	3.78	3.81	3.73	3.88	3.89	3.85	3.65	3.74
2094-705	5	363	4.64	4.68	4.56	4.80	4.82	4.76	4.43	4.57
2094-706	6	422	5.46	5.53	5.35	5.71	5.74	5.64	5.17	5.36
2094-707	7	479	6.25	6.34	6.10	6.59	6.63	6.50	5.86	6.12
2094-708	8	532	7.02	7.13	6.82	7.46	7.51	7.33	6.51	6.85
2094-709	9	584	7.75	7.90	7.51	8.32	8.38	8.15	7.12	7.54
2094-710	10	640	8.47	8.65	8.17	9.16	9.24	8.95	7.71	8.21
2094-711	11	690	9.16	9.37	8.81	9.99	10.09	9.73	8.27	8.86
2094-712	12	742	9.83	10.08	9.42	10.81	10.93	10.49	8.80	9.48
2094-713	13	795	10.48	10.76	10.01	11.61	11.76	11.23	9.31	10.08
2094-714	14	843	11.10	11.43	10.58	12.40	12.58	11.95	9.79	10.65
2094-715	15	893	11.71	12.08	11.13	13.18	13.39	12.66	10.25	11.20
2094-716	16	938	12.30	12.71	11.65	13.95	14.19	13.35	10.69	11.74
2094-717	17	985	12.86	13.32	12.16	14.71	14.98	14.02	11.11	12.25
2094-718	18	1032	13.41	13.91	12.64	15.46	15.76	14.68	11.50	12.74
2094-719	19	1077	13.95	14.50	13.12	16.19	16.53	15.32	11.90	13.22
2094-720	20	1127	14.48	15.07	13.59	16.92	17.29	15.95	12.29	13.70
2094-721	21	1177	15.01	15.64	14.06	17.63	18.04	16.57	12.67	14.17
2094-722	22	1226	15.52	16.20	14.52	18.33	18.77	17.18	13.06	14.63
2094-723	23	1274	16.03	16.74	14.97	19.03	19.50	17.77	13.43	15.09
2094-724	24	1314	16.53	17.28	15.42	19.71	20.22	18.35	13.81	15.54
2094-725	25	1353	17.02	17.82	15.86	20.38	20.92	18.93	14.18	15.98
2094-726	26	1400	17.51	18.34	16.30	21.04	21.62	19.48	14.55	16.42
2094-727	27	1436	17.98	18.86	16.73	21.69	22.31	20.03	14.92	16.85
2094-728	28	1481	18.45	19.36	17.15	22.33	22.99	20.57	15.28	17.27
2094-729	29	1524	18.92	19.87	17.57	22.96	23.65	21.10	15.64	17.68
2094-730	30	1567	19.37	20.36	17.98	23.59	24.31	21.62	15.99	18.09
2094-731	31	1609	19.82	20.84	18.39	24.20	24.96	22.13	16.33	18.49
2094-732	32	1650	20.26	21.32	18.78	24.80	25.60	22.64	16.67	18.88
2094-733	33	1691	20.69	21.79	19.17	25.40	26.23	23.13	17.01	19.26
2094-734	34	1732	21.12	22.26	19.56	25.99	26.85	23.62	17.34	19.64
2094-735	35	1772	21.54	22.71	19.94	26.56	27.46	24.11	17.67	20.01
2094-736	36	1811	21.95	23.16	20.31	27.13	28.06	24.58	17.99	20.37
2094-737	37	1850	22.36	23.61	20.67	27.69	28.65	25.05	18.30	20.73
2094-738	38	1889	22.76	24.04	21.03	28.24	29.23	25.51	18.62	21.08
2094-739	39	1927	23.15	24.47	21.38	28.79	29.80	25.97	18.92	21.42
2094-740	40	1963	23.54	24.89	21.73	29.32	30.37	26.42	19.22	21.75
2094-741	41	1998	23.91	25.31	22.07	29.85	30.93	26.86	19.52	22.07
2094-742	42	2033	24.29	25.72	22.40	30.37	31.48	27.29	19.81	22.39
2094-743	43	2067	24.65	26.12	22.73	30.88	32.02	27.72	20.10	22.70
2094-744	44	2101	25.01	26.52	23.05	31.38	32.55	28.14	20.39	23.00
2094-745	45	2135	25.36	26.91	23.36	31.88	33.07	28.56	20.67	23.29
2094-746	46	2168	25.71	27.30	23.67	32.37	33.59	28.96	20.94	23.57
2094-747	47	2201	26.05	27.68	23.97	32.85	34.10	29.37	21.22	23.85
2094-748	48	2233	26.38	28.06	24.27	33.32	34.60	29.76	21.48	24.12
2094-749	49	2265	26.71	28.42	24.56	33.79	35.09	30.15	21.75	24.39
2094-750	50	2295	27.03	28.79	24.85	34.25	35.58	30.54	22.01	24.64
2094-751	51	2324	27.35	29.15	25.13	34.71	36.06	30.92	22.26	24.90
2094-752	52	2353	27.66	29.50	25.40	35.15	36.53	31.29	22.52	25.14
2094-753	53	2381	27.97	29.84	25.67	35.59	36.99	31.66	22.76	25.38
2094-754	54	2409	28.27	30.19	25.94	36.02	37.45	32.02	23.01	25.61

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 10 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008 ****see below

	Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	V1R9 CB-L	V1R9 ODE-B	V1R9 WASDB	V1R9 OLTP-W	V1R9 OLTP-T
IBM System z10 EC											
	2097-401	1	27	0.39	0.39	0.39	0.39	0.40	0.38	0.38	0.39
	2097-402	2	51	0.73	0.74	0.72	0.77	0.77	0.74	0.69	0.73
	2097-403	3	75	1.06	1.07	1.04	1.13	1.14	1.09	0.98	1.06
	2097-404	4	97	1.38	1.40	1.35	1.49	1.50	1.43	1.26	1.38
	2097-405	5	118	1.70	1.72	1.65	1.84	1.86	1.77	1.53	1.69
	2097-406	6	139	2.01	2.04	1.94	2.19	2.21	2.10	1.79	1.99
	2097-407	7	160	2.31	2.35	2.22	2.53	2.56	2.42	2.04	2.28
	2097-408	8	180	2.60	2.65	2.50	2.87	2.90	2.75	2.28	2.57
	2097-409	9	199	2.89	2.94	2.77	3.21	3.23	3.06	2.52	2.85
	2097-410	10	218	3.17	3.23	3.03	3.54	3.57	3.38	2.75	3.13
	2097-411	11	237	3.44	3.52	3.28	3.87	3.89	3.68	2.96	3.39
	2097-412	12	255	3.71	3.80	3.53	4.19	4.21	3.99	3.18	3.66
	2097-501	1	58	0.83	0.83	0.83	0.85	0.86	0.82	0.82	0.84
	2097-502	2	110	1.58	1.59	1.56	1.65	1.67	1.59	1.49	1.58
	2097-503	3	160	2.29	2.31	2.24	2.44	2.47	2.34	2.11	2.29
	2097-504	4	207	2.98	3.01	2.90	3.21	3.25	3.08	2.70	2.97
	2097-505	5	252	3.65	3.70	3.53	3.96	4.02	3.80	3.27	3.63
	2097-506	6	296	4.30	4.37	4.15	4.71	4.78	4.51	3.82	4.26
	2097-507	7	340	4.93	5.02	4.75	5.44	5.53	5.20	4.34	4.89
	2097-508	8	382	5.55	5.66	5.32	6.16	6.26	5.88	4.85	5.49
	2097-509	9	422	6.15	6.28	5.89	6.88	6.99	6.55	5.33	6.07
	2097-510	10	462	6.74	6.89	6.43	7.58	7.70	7.21	5.80	6.64
	2097-511	11	500	7.31	7.48	6.96	8.27	8.41	7.86	6.25	7.20
	2097-512	12	537	7.87	8.06	7.47	8.95	9.10	8.49	6.69	7.74
	2097-601	1	79	1.13	1.13	1.13	1.15	1.16	1.12	1.11	1.14
	2097-602	2	149	2.13	2.15	2.10	2.24	2.26	2.15	2.01	2.14
	2097-603	3	215	3.10	3.13	3.03	3.30	3.34	3.17	2.85	3.10
	2097-604	4	277	4.02	4.07	3.91	4.34	4.40	4.16	3.64	4.01
	2097-605	5	339	4.92	4.99	4.76	5.36	5.44	5.13	4.40	4.89
	2097-606	6	398	5.79	5.88	5.58	6.36	6.47	6.08	5.12	5.74
	2097-607	7	455	6.63	6.76	6.37	7.34	7.48	7.01	5.82	6.57
	2097-608	8	511	7.46	7.61	7.14	8.31	8.48	7.92	6.49	7.37
	2097-609	9	565	8.26	8.44	7.88	9.26	9.46	8.82	7.13	8.14
	2097-610	10	617	9.03	9.24	8.60	10.20	10.43	9.69	7.74	8.89
	2097-611	11	668	9.79	10.03	9.30	11.12	11.38	10.55	8.33	9.62
	2097-612	12	717	10.52	10.79	9.97	12.02	12.32	11.39	8.90	10.33

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Part 11 of 11

All capacity numbers are relative to the IBM 2094-701 running A SINGLE z/OS IMAGE. Information presented in this table is current as of Oct. 2008

Processor	#CP	MSU***	Mixed*	LoIO-Mix*	TI-Mix*	z/OS V1R9 CB-L	z/OS V1R9 ODE-B	z/OS V1R9 WASDB	z/OS V1R9 OLTP-W	z/OS V1R9 OLTP-T
IBM System z10 EC (continued)										
2097-701	1	115	1.63	1.63	1.63	1.66	1.68	1.61	1.60	1.64
2097-702	2	215	3.07	3.09	3.02	3.23	3.26	3.11	2.88	3.08
2097-703	3	312	4.44	4.49	4.34	4.75	4.82	4.56	4.07	4.45
2097-704	4	401	5.76	5.83	5.59	6.24	6.34	5.98	5.18	5.74
2097-705	5	488	7.03	7.14	6.78	7.70	7.85	7.37	6.24	6.98
2097-706	6	571	8.26	8.41	7.94	9.13	9.33	8.72	7.25	8.18
2097-707	7	651	9.45	9.63	9.05	10.53	10.79	10.04	8.21	9.34
2097-708	8	729	10.60	10.83	10.12	11.90	12.23	11.33	9.13	10.46
2097-709	9	804	11.71	11.98	11.14	13.24	13.64	12.58	10.02	11.53
2097-710	10	875	12.79	13.11	12.13	14.56	15.04	13.81	10.86	12.57
2097-711	11	944	13.83	14.20	13.09	15.85	16.41	15.00	11.66	13.57
2097-712	12	1011	14.84	15.26	14.01	17.12	17.77	16.17	12.43	14.53
2097-713	13	1076	15.82	16.28	14.89	18.36	19.10	17.31	13.16	15.46
2097-714	14	1139	16.77	17.29	15.75	19.58	20.41	18.43	13.88	16.36
2097-715	15	1202	17.70	18.27	16.59	20.80	21.70	19.52	14.57	17.25
2097-716	16	1264	18.61	19.24	17.41	22.00	22.98	20.58	15.24	18.11
2097-717	17	1329	19.51	20.19	18.21	23.19	24.23	21.62	15.89	18.95
2097-718	18	1390	20.38	21.11	18.99	24.37	25.47	22.64	16.53	19.77
2097-719	19	1451	21.23	22.02	19.76	25.54	26.68	23.64	17.14	20.56
2097-720	20	1512	22.06	22.91	20.50	26.69	27.88	24.61	17.74	21.34
2097-721	21	1571	22.87	23.79	21.22	27.84	29.06	25.57	18.32	22.10
2097-722	22	1631	23.67	24.64	21.93	28.97	30.22	26.50	18.88	22.84
2097-723	23	1690	24.45	25.48	22.62	30.09	31.37	27.41	19.43	23.56
2097-724	24	1748	25.21	26.30	23.29	31.20	32.49	28.30	19.96	24.27
2097-725	25	1805	25.95	27.11	23.95	32.30	33.61	29.17	20.48	24.95
2097-726	26	1865	26.68	27.90	24.59	33.39	34.70	30.03	20.98	25.62
2097-727	27	1922	27.39	28.67	25.22	34.46	35.78	30.86	21.48	26.27
2097-728	28	1979	28.09	29.44	25.83	35.52	36.84	31.69	21.97	26.90
2097-729	29	2037	28.77	30.19	26.44	36.56	37.88	32.49	22.46	27.50
2097-730	30	2092	29.45	30.94	27.03	37.58	38.91	33.28	22.95	28.09
2097-731	31	2146	30.11	31.67	27.61	38.59	39.93	34.06	23.43	28.67
2097-732	32	2200	30.76	32.39	28.19	39.58	40.92	34.83	23.90	29.22
2097-733	33	2257	31.40	33.11	28.75	40.56	41.91	35.58	24.37	29.76
2097-734	34	2309	32.02	33.81	29.30	41.53	42.88	36.31	24.84	30.28
2097-735	35	2366	32.64	34.50	29.84	42.47	43.83	37.04	25.30	30.78
2097-736	36	2422	33.24	35.18	30.37	43.41	44.77	37.75	25.76	31.27
2097-737	37	2478	33.84	35.86	30.89	44.33	45.70	38.45	26.22	31.74
2097-738	38	2530	34.42	36.52	31.40	45.24	46.61	39.13	26.67	32.20
2097-739	39	2585	34.99	37.18	31.91	46.13	47.51	39.80	27.11	32.64
2097-740	40	2636	35.55	37.82	32.40	47.01	48.39	40.46	27.56	33.07
2097-741	41	2687	36.10	38.46	32.88	47.88	49.26	41.11	27.99	33.49
2097-742	42	2740	36.65	39.09	33.36	48.74	50.12	41.75	28.43	33.89
2097-743	43	2789	37.18	39.71	33.82	49.58	50.96	42.39	28.86	34.28
2097-744	44	2838	37.70	40.32	34.28	50.41	51.80	43.01	29.28	34.65
2097-745	45	2886	38.21	40.92	34.73	51.22	52.62	43.62	29.71	35.01
2097-746	46	2934	38.72	41.52	35.17	52.03	53.42	44.22	30.12	35.36
2097-747	47	2981	39.21	42.11	35.60	52.82	54.22	44.82	30.54	35.69
2097-748	48	3028	39.70	42.69	36.02	53.60	55.00	45.40	30.94	36.02
2097-749	49	3075	40.17	43.26	36.43	54.37	55.77	45.98	31.35	36.33
2097-750	50	3120	40.64	43.83	36.84	55.12	56.53	46.55	31.75	36.63
2097-751	51	3166	41.10	44.39	37.24	55.87	57.28	47.10	32.15	36.92
2097-752	52	3214	41.55	44.94	37.63	56.60	58.02	47.66	32.54	37.20
2097-753	53	3262	42.00	45.48	38.01	57.33	58.75	48.20	32.93	37.48
2097-754	54	3305	42.43	46.02	38.38	58.04	59.46	48.73	33.31	37.74
2097-755	55	3352	42.86	46.54	38.75	58.74	60.17	49.26	33.70	37.99
2097-756	56	3395	43.28	47.07	39.11	59.43	60.86	49.78	34.07	38.23
2097-757	57	3438	43.69	47.58	39.47	60.11	61.54	50.29	34.45	38.47
2097-758	58	3480	44.10	48.09	39.82	60.79	62.22	50.79	34.82	38.70
2097-759	59	3525	44.50	48.59	40.16	61.45	62.88	51.28	35.18	38.92
2097-760	60	3570	44.89	49.09	40.49	62.10	63.53	51.77	35.55	39.13
2097-761	61	3611	45.27	49.58	40.82	62.74	64.18	52.25	35.90	39.33
2097-762	62	3652	45.65	50.06	41.14	63.37	64.81	52.73	36.26	39.53
2097-763	63	3695	46.02	50.54	41.46	63.99	65.43	53.19	36.61	39.72
2097-764	64	3739	46.38	51.01	41.77	64.60	66.05	53.65	36.96	39.90

***LSPR default mixed workload.**

*** Mixed consists of an equal mix of OLTP-T, OLTP-W, WASDB, and CB-L.**
LoIO-Mix consists of 44% CB-L, 28% WASDB and 28% OLTP-W.
TI-Mix consists of 20% CB-L, 40% OLTP-W, and 40% OLTP-T

*****MSUs are used for software pricing and are not necessarily a direct indication of relative processor capacity.**
MSUs for the 2084, 2086, 2094, 2096, and 2098 model types are set to provide increased software price/performance for applicable programs which have MSU based pricing.

******The higher Nways of the Single-Image table include enahncements whose availabilities are expected to be consistent with customer requirements. anyone contemplating running a 48way or larger z/OS image should contact IBM for a review of potential constraints.**

Table 8. ITR Ratio Table - Linux on System z SuSE SLES 10

The tables in this section present LSPR capacity data in the form of ITR ratios relative to the 2094-701. The data represents Linux SuSE SLES – 10

Part 1 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM e(logo)Server zSeries 800		
2066-0E1	1	0.06
2066-0A1	1	0.12
2066-0B1	1	0.18
2066-0C1	1	0.22
2066-0X2	2	0.26
2066-001	1	0.29
2066-0A2	2	0.42
2066-002	2	0.57
2066-003	3	0.82
2066-004	4	1.07
IBM e(logo)Server zSeries 890		
2086-110	1	0.04
2086-210	2	0.08
2086-310	3	0.12
2086-410	4	0.17
2086-120	1	0.07
2086-220	2	0.15
2086-320	3	0.22
2086-420	4	0.29
2086-130	1	0.15
2086-230	2	0.29
2086-330	3	0.43
2086-430	4	0.57
2086-140	1	0.18
2086-240	2	0.36
2086-340	3	0.54
2086-440	4	0.72
2086-150	1	0.28
2086-250	2	0.56
2086-350	3	0.84
2086-450	4	1.12
2086-160	1	0.35
2086-260	2	0.69
2086-360	3	1.02
2086-460	4	1.37
2086-170	1	0.60
2086-270	2	1.18
2086-370	3	1.77
2086-470	4	2.36

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 2 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z9 BC R07		
2096-A01	1	0.05
2096-A02	2	0.09
2096-A03	3	0.13
2096-B01	1	0.06
2096-B02	2	0.13
2096-B03	3	0.19
2096-C01	1	0.08
2096-C02	2	0.16
2096-C03	3	0.23
2096-D01	1	0.10
2096-D02	2	0.20
2096-D03	3	0.30
2096-E01	1	0.12
2096-E02	2	0.24
2096-F01	1	0.15
2096-F02	2	0.30
2096-G01	1	0.19
2096-H01	1	0.22
2096-I01	1	0.26
2096-J01	1	0.30

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 3 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z9 BC S07		
2096-K04	4	0.40
2096-L03	3	0.35
2096-L04	4	0.47
2096-M03	3	0.44
2096-M04	4	0.58
2096-N02	2	0.37
2096-N03	3	0.56
2096-N04	4	0.73
2096-O02	2	0.44
2096-O03	3	0.66
2096-O04	4	0.87
2096-P02	2	0.51
2096-P03	3	0.76
2096-P04	4	1.00
2096-Q02	2	0.59
2096-Q03	3	0.87
2096-Q04	4	1.15
2096-R01	1	0.33
2096-R02	2	0.66
2096-R03	3	0.98
2096-R04	4	1.29
2096-S01	1	0.37
2096-S02	2	0.74
2096-S03	3	1.09
2096-S04	4	1.44
2096-T01	1	0.42
2096-T02	2	0.82
2096-T03	3	1.22
2096-T04	4	1.62
2096-U01	1	0.46
2096-U02	2	0.92
2096-U03	3	1.37
2096-U04	4	1.80
2096-V01	1	0.52
2096-V02	2	1.03
2096-V03	3	1.53
2096-V04	4	2.02
2096-W01	1	0.58
2096-W02	2	1.16
2096-W03	3	1.72
2096-W04	4	2.27
2096-X01	1	0.65
2096-X02	2	1.29
2096-X03	3	1.92
2096-X04	4	2.54
2096-Y01	1	0.73
2096-Y02	2	1.44
2096-Y03	3	2.13
2096-Y04	4	2.82
2096-Z01	1	0.83
2096-Z02	2	1.64
2096-Z03	3	2.43
2096-Z04	4	3.21

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 4 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z10 BC		
2098-A01	1	0.05
2098-A02	2	0.10
2098-A03	3	0.14
2098-A04	4	0.19
2098-A05	5	0.23
2098-B01	1	0.06
2098-B02	2	0.11
2098-B03	3	0.16
2098-B04	4	0.21
2098-B05	5	0.26
2098-C01	1	0.07
2098-C02	2	0.14
2098-C03	3	0.20
2098-C04	4	0.26
2098-C05	5	0.33
2098-D01	1	0.09
2098-D02	2	0.17
2098-D03	3	0.24
2098-D04	4	0.32
2098-D05	5	0.40
2098-E01	1	0.10
2098-E02	2	0.19
2098-E03	3	0.28
2098-E04	4	0.37
2098-E05	5	0.46
2098-F01	1	0.11
2098-F02	2	0.21
2098-F03	3	0.31
2098-F04	4	0.41
2098-F05	5	0.51
2098-G01	1	0.13
2098-G02	2	0.25
2098-G03	3	0.37
2098-G04	4	0.49
2098-G05	5	0.60
2098-H01	1	0.15
2098-H02	2	0.28
2098-H03	3	0.42
2098-H04	4	0.55
2098-H05	5	0.67
2098-I01	1	0.16
2098-I02	2	0.32
2098-I03	3	0.47
2098-I04	4	0.62
2098-I05	5	0.76
2098-J01	1	0.18
2098-J02	2	0.35
2098-J03	3	0.52
2098-J04	4	0.69
2098-J05	5	0.85
2098-K01	1	0.21
2098-K02	2	0.40
2098-K03	3	0.59
2098-K04	4	0.77
2098-K05	5	0.95
2098-L01	1	0.24
2098-L02	2	0.47
2098-L03	3	0.70
2098-L04	4	0.91
2098-L05	5	1.13
2098-M01	1	0.28
2098-M02	2	0.55
2098-M03	3	0.80
2098-M04	4	1.06
2098-M05	5	1.30

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 5 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z10 BC (continued)		
2098-N01	1	0.32
2098-N02	2	0.63
2098-N03	3	0.92
2098-N04	4	1.21
2098-N05	5	1.50
2098-O01	1	0.36
2098-O02	2	0.71
2098-O03	3	1.04
2098-O04	4	1.37
2098-O05	5	1.69
2098-P01	1	0.41
2098-P02	2	0.79
2098-P03	3	1.17
2098-P04	4	1.53
2098-P05	5	1.89
2098-Q01	1	0.46
2098-Q02	2	0.89
2098-Q03	3	1.31
2098-Q04	4	1.72
2098-Q05	5	2.12
2098-R01	1	0.51
2098-R02	2	0.99
2098-R03	3	1.46
2098-R04	4	1.91
2098-R05	5	2.36
2098-S01	1	0.57
2098-S02	2	1.11
2098-S03	3	1.63
2098-S04	4	2.14
2098-S05	5	2.64
2098-T01	1	0.64
2098-T02	2	1.24
2098-T03	3	1.83
2098-T04	4	2.40
2098-T05	5	2.96
2098-U01	1	0.72
2098-U02	2	1.39
2098-U03	3	2.04
2098-U04	4	2.68
2098-U05	5	3.31
2098-V01	1	0.80
2098-V02	2	1.56
2098-V03	3	2.29
2098-V04	4	3.01
2098-V05	5	3.72
2098-W01	1	0.90
2098-W02	2	1.74
2098-W03	3	2.56
2098-W04	4	3.36
2098-W05	5	4.15
2098-X01	1	1.00
2098-X02	2	1.94
2098-X03	3	2.86
2098-X04	4	3.75
2098-X05	5	4.63
2098-Y01	1	1.14
2098-Y02	2	2.21
2098-Y03	3	3.25
2098-Y04	4	4.27
2098-Y05	5	5.27
2098-Z01	1	1.24
2098-Z02	2	2.41
2098-Z03	3	3.55
2098-Z04	4	4.66
2098-Z05	5	5.75

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 6 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM e(logo)Server zSeries 900 (1XX Series)		
2064-101	1	0.37
2064-102	2	0.73
2064-103	3	1.08
2064-104	4	1.43
2064-105	5	1.77
2064-106	6	2.10
2064-107	7	2.43
2064-108	8	2.75
2064-109	9	3.06
2064-1C1	1	0.38
2064-1C2	2	0.76
2064-1C3	3	1.14
2064-1C4	4	1.51
2064-1C5	5	1.88
2064-1C6	6	2.25
2064-1C7	7	2.62
2064-1C8	8	2.98
2064-1C9	9	3.34
2064-110	10	3.70
2064-111	11	4.05
2064-112	12	4.40
2064-113	13	4.75
2064-114	14	5.09
2064-115	15	5.43
2064-116	16	5.77
IBM e(logo)Server zSeries 900 (2XX Series)		
2064-2C1	1	0.46
2064-2C2	2	0.92
2064-2C3	3	1.38
2064-2C4	4	1.83
2064-2C5	5	2.28
2064-2C6	6	2.73
2064-2C7	7	3.17
2064-2C8	8	3.60
2064-2C9	9	4.04
2064-210	10	4.46
2064-211	11	4.88
2064-212	12	5.30
2064-213	13	5.71
2064-214	14	6.12
2064-215	15	6.52
2064-216	16	6.91

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 7 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM e(logo)Server zSeries 990		
2084-301	1	0.74
2084-302	2	1.46
2084-303	3	2.18
2084-304	4	2.88
2084-305	5	3.57
2084-306	6	4.25
2084-307	7	4.92
2084-308	8	5.58
2084-309	9	6.23
2084-310	10	6.87
2084-311	11	7.51
2084-312	12	8.13
2084-313	13	8.76
2084-314	14	9.37
2084-315	15	9.98
2084-316	16	10.58

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 8 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z9 EC		
2094-401	1	0.34
2094-402	2	0.68
2094-403	3	1.02
2094-404	4	1.34
2094-405	5	1.66
2094-406	6	1.98
2094-407	7	2.29
2094-408	8	2.59
2094-501	1	0.67
2094-502	2	1.33
2094-503	3	1.97
2094-504	4	2.60
2094-505	5	3.23
2094-506	6	3.84
2094-507	7	4.44
2094-508	8	5.03
2094-601	1	0.81
2094-602	2	1.61
2094-603	3	2.39
2094-604	4	3.16
2094-605	5	3.91
2094-606	6	4.65
2094-607	7	5.38
2094-608	8	6.09
2094-701	1	1.00
2094-702	2	1.98
2094-703	3	2.95
2094-704	4	3.89
2094-705	5	4.82
2094-706	6	5.74
2094-707	7	6.63
2094-708	8	7.51
2094-709	9	8.38
2094-710	10	9.24
2094-711	11	10.09
2094-712	12	10.93
2094-713	13	11.76
2094-714	14	12.58
2094-715	15	13.39
2094-716	16	14.19

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 9 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z10 EC		
2097-401	1	0.39
2097-402	2	0.78
2097-403	3	1.15
2097-404	4	1.52
2097-405	5	1.88
2097-406	6	2.23
2097-407	7	2.58
2097-408	8	2.93
2097-409	9	3.27
2097-410	10	3.60
2097-411	11	3.93
2097-412	12	4.26
2097-501	1	0.85
2097-502	2	1.67
2097-503	3	2.47
2097-504	4	3.25
2097-505	5	4.02
2097-506	6	4.78
2097-507	7	5.53
2097-508	8	6.27
2097-509	9	6.99
2097-510	10	7.71
2097-511	11	8.41
2097-512	12	9.11
2097-601	1	1.15
2097-602	2	2.24
2097-603	3	3.32
2097-604	4	4.38
2097-605	5	5.43
2097-606	6	6.45
2097-607	7	7.45
2097-608	8	8.44
2097-609	9	9.41
2097-610	10	10.36
2097-611	11	11.29
2097-612	12	12.21
2097-701	1	1.66
2097-702	2	3.22
2097-703	3	4.76
2097-704	4	6.27
2097-705	5	7.76
2097-706	6	9.22
2097-707	7	10.65
2097-708	8	12.06
2097-709	9	13.45
2097-710	10	14.81
2097-711	11	16.15
2097-712	12	17.47
2097-713	13	18.77
2097-714	14	20.04
2097-715	15	21.30
2097-716	16	22.54

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Part 10 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a Linux workload running on System z. Information presented in this table is current as of May 2010

Processor	#CP	SLES 10 Low*/L
IBM System z196		
2817-401	1	0.43
2817-402	2	0.85
2817-403	3	1.27
2817-404	4	1.68
2817-405	5	2.08
2817-406	6	2.48
2817-407	7	2.87
2817-408	8	3.26
2817-409	9	3.64
2817-410	10	4.02
2817-411	11	4.39
2817-412	12	4.76
2817-413	13	5.12
2817-414	14	5.48
2817-415	15	5.83
2817-501	1	1.06
2817-502	2	2.07
2817-503	3	3.07
2817-504	4	4.05
2817-505	5	5.02
2817-506	6	5.98
2817-507	7	6.92
2817-508	8	7.85
2817-509	9	8.78
2817-510	10	9.68
2817-511	11	10.58
2817-512	12	11.47
2817-513	13	12.34
2817-514	14	13.20
2817-515	15	14.05
2817-601	1	1.38
2817-602	2	2.72
2817-603	3	4.03
2817-604	4	5.32
2817-605	5	6.59
2817-606	6	7.84
2817-607	7	9.07
2817-608	8	10.28
2817-609	9	11.46
2817-610	10	12.63
2817-611	11	13.78
2817-612	12	14.90
2817-613	13	16.01
2817-614	14	17.10
2817-615	15	18.17
2817-701	1	2.16
2817-702	2	4.23
2817-703	3	6.27
2817-704	4	8.27
2817-705	5	10.24
2817-706	6	12.17
2817-707	7	14.07
2817-708	8	15.94
2817-709	9	17.78
2817-710	10	19.58
2817-711	11	21.36
2817-712	12	23.10
2817-713	13	24.82
2817-714	14	26.51
2817-715	15	28.16
2817-716	16	29.79

*** Low RNI (Relative Nest Intensity) describes a workload category representing light use of the memory hierarchy - see workload description for more information**

Table 9. ITR Ratio Table - Linux on System z SuSE SLES 9

The tables in this section present LSPR capacity data in the form of ITR ratios relative to the 2094-701. The data represents Linux SuSE SLES - 9 workload running in 64-bit real addressing mode.

	Part 1 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			Linux						
	Model	# CP	WASDB/L						
zSeries 800									
	2066-0E1	1W	0.07						
	2066-0A1	1W	0.13						
	2066-0B1	1W	0.19						
	2066-0X2	2W	0.28						
	2066-0C1	1W	0.23						
	2066-0A2	2W	0.46						
	2066-001	1W	0.31						
	2066-002	2W	0.61						
	2066-003	3W	0.86						
	2066-004	4W	1.08						
zSeries 890									
	2086-110	1W	0.04						
	2086-210	2W	0.08						
	2086-310	3W	0.12						
	2086-410	4W	0.15						
	2086-120	1W	0.07						
	2086-220	2W	0.15						
	2086-320	3W	0.22						
	2086-420	4W	0.29						
	2086-130	1W	0.15						
	2086-230	2W	0.29						
	2086-330	3W	0.43						
	2086-430	4W	0.57						
	2086-140	1W	0.18						
	2086-240	2W	0.37						
	2086-340	3W	0.54						
	2086-440	4W	0.72						
	2086-150	1W	0.29						
	2086-250	2W	0.58						
	2086-350	3W	0.85						
	2086-450	4W	1.13						
	2086-160	1W	0.35						
	2086-260	2W	0.70						
	2086-360	3W	1.04						
	2086-460	4W	1.38						
	2086-170	1W	0.61						
	2086-270	2W	1.23						
	2086-370	3W	1.81						
	2086-470	4W	2.41						

	Part 2 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			Linux						
	Model	# CP	WASDB/L						
==========									
System z9 BC R07									
	2096-A01	1W	0.04						
	2096-A02	2W	0.09						
	2096-A03	3W	0.13						
	2096-B01	1W	0.07						
	2096-B02	2W	0.13						
	2096-B03	3W	0.19						
	2096-C01	1W	0.08						
	2096-C02	2W	0.16						
	2096-C03	3W	0.23						
	2096-D01	1W	0.10						
	2096-D02	2W	0.20						
	2096-D03	3W	0.30						
	2096-E01	1W	0.12						
	2096-E02	2W	0.24						
	2096-F01	1W	0.15						
	2096-F02	2W	0.30						
	2096-G01	1W	0.19						
	2096-H01	1W	0.23						
	2096-I01	1W	0.26						
	2096-J01	1W	0.30						

Part 3 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

Model	# CP	Linux WASDB/L
System z9 BC S07		
2096-K04	4W	0.39
2096-L03	3W	0.35
2096-L04	4W	0.46
2096-M03	3W	0.44
2096-M04	4W	0.57
2096-N02	2W	0.37
2096-N03	3W	0.55
2096-N04	4W	0.72
2096-O02	2W	0.44
2096-O03	3W	0.65
2096-O04	4W	0.86
2096-P02	2W	0.51
2096-P03	3W	0.75
2096-P04	4W	0.99
2096-Q02	2W	0.59
2096-Q03	3W	0.86
2096-Q04	4W	1.13
2096-R01	1W	0.34
2096-R02	2W	0.66
2096-R03	3W	0.97
2096-R04	4W	1.27
2096-S01	1W	0.37
2096-S02	2W	0.73
2096-S03	3W	1.08
2096-S04	4W	1.42
2096-T01	1W	0.42
2096-T02	2W	0.82
2096-T03	3W	1.21
2096-T04	4W	1.59
2096-U01	1W	0.47
2096-U02	2W	0.92
2096-U03	3W	1.35
2096-U04	4W	1.77
2096-V01	1W	0.53
2096-V02	2W	1.03
2096-V03	3W	1.52
2096-V04	4W	1.99
2096-W01	1W	0.59
2096-W02	2W	1.16
2096-W03	3W	1.70
2096-W04	4W	2.23
2096-X01	1W	0.66
2096-X02	2W	1.29
2096-X03	3W	1.90
2096-X04	4W	2.50
2096-Y01	1W	0.73
2096-Y02	2W	1.43
2096-Y03	3W	2.11
2096-Y04	4W	2.77
2096-Z01	1W	0.83
2096-Z02	2W	1.63
2096-Z03	3W	2.40
2096-Z04	4W	3.15

Part 4 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

Model	# CP	Linux WASDB/L
System z10 BC		
2098-A01	1	0.05
2098-A02	2	0.10
2098-A03	3	0.15
2098-A04	4	0.19
2098-A05	5	0.23
2098-B01	1	0.06
2098-B02	2	0.11
2098-B03	3	0.16
2098-B04	4	0.21
2098-B05	5	0.25
2098-C01	1	0.07
2098-C02	2	0.13
2098-C03	3	0.19
2098-C04	4	0.24
2098-C05	5	0.29
2098-D01	1	0.08
2098-D02	2	0.16
2098-D03	3	0.23
2098-D04	4	0.29
2098-D05	5	0.35
2098-E01	1	0.10
2098-E02	2	0.18
2098-E03	3	0.26
2098-E04	4	0.33
2098-E05	5	0.40
2098-F01	1	0.11
2098-F02	2	0.20
2098-F03	3	0.29
2098-F04	4	0.37
2098-F05	5	0.45
2098-G01	1	0.13
2098-G02	2	0.24
2098-G03	3	0.35
2098-G04	4	0.45
2098-G05	5	0.54
2098-H01	1	0.14
2098-H02	2	0.27
2098-H03	3	0.39
2098-H04	4	0.50
2098-H05	5	0.60
2098-I01	1	0.16
2098-I02	2	0.31
2098-I03	3	0.44
2098-I04	4	0.57
2098-I05	5	0.69
2098-J01	1	0.18
2098-J02	2	0.35
2098-J03	3	0.50
2098-J04	4	0.64
2098-J05	5	0.77
2098-K01	1	0.21
2098-K02	2	0.39
2098-K03	3	0.56
2098-K04	4	0.72
2098-K05	5	0.87
2098-L01	1	0.25
2098-L02	2	0.47
2098-L03	3	0.68
2098-L04	4	0.87
2098-L05	5	1.04
2098-M01	1	0.29
2098-M02	2	0.55
2098-M03	3	0.79
2098-M04	4	1.02
2098-M05	5	1.23

Part 5 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

Model	# CP	Linux WASDB/L
System z10 BC continued		
2098-N01	1	0.34
2098-N02	2	0.64
2098-N03	3	0.93
2098-N04	4	1.19
2098-N05	5	1.43
2098-O01	1	0.39
2098-O02	2	0.74
2098-O03	3	1.06
2098-O04	4	1.36
2098-O05	5	1.64
2098-P01	1	0.44
2098-P02	2	0.84
2098-P03	3	1.21
2098-P04	4	1.55
2098-P05	5	1.87
2098-Q01	1	0.49
2098-Q02	2	0.94
2098-Q03	3	1.35
2098-Q04	4	1.73
2098-Q05	5	2.09
2098-R01	1	0.55
2098-R02	2	1.05
2098-R03	3	1.51
2098-R04	4	1.93
2098-R05	5	2.32
2098-S01	1	0.61
2098-S02	2	1.17
2098-S03	3	1.68
2098-S04	4	2.16
2098-S05	5	2.60
2098-T01	1	0.69
2098-T02	2	1.31
2098-T03	3	1.88
2098-T04	4	2.41
2098-T05	5	2.91
2098-U01	1	0.76
2098-U02	2	1.46
2098-U03	3	2.10
2098-U04	4	2.69
2098-U05	5	3.24
2098-V01	1	0.86
2098-V02	2	1.64
2098-V03	3	2.35
2098-V04	4	3.01
2098-V05	5	3.63
2098-W01	1	0.95
2098-W02	2	1.82
2098-W03	3	2.62
2098-W04	4	3.35
2098-W05	5	4.04
2098-X01	1	1.06
2098-X02	2	2.03
2098-X03	3	2.92
2098-X04	4	3.74
2098-X05	5	4.50
2098-Y01	1	1.20
2098-Y02	2	2.30
2098-Y03	3	3.31
2098-Y04	4	4.24
2098-Y05	5	5.11
2098-Z01	1	1.31
2098-Z02	2	2.50
2098-Z03	3	3.60
2098-Z04	4	4.61
2098-Z05	5	5.56

	Part 6 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			Linux						
	Model	# CP	WASDB/L						
zSeries 900(1XX Series)									
	2064-101	1W	0.39						
	2064-102	2W	0.77						
	2064-103	3W	1.14						
	2064-104	4W	1.50						
	2064-105	5W	1.84						
	2064-106	6W	2.16						
	2064-107	7W	2.47						
	2064-108	8W	2.77						
	2064-109	9W	3.07						
	2064-1C1	1W	0.40						
	2064-1C2	2W	0.80						
	2064-1C3	3W	1.18						
	2064-1C4	4W	1.57						
	2064-1C5	5W	1.90						
	2064-1C6	6W	2.24						
	2064-1C7	7W	2.56						
	2064-1C8	8W	2.87						
	2064-1C9	9W	3.18						
	2064-110	10W	3.49						
	2064-111	11W	3.79						
	2064-112	12W	4.10						
	2064-113	13W	4.41						
	2064-114	14W	4.71						
	2064-115	15W	5.02						
	2064-116	16W	5.33						
zSeries 900 Turbo (2XX Series)									
	2064-2C1	1W	0.49						
	2064-2C2	2W	0.98						
	2064-2C3	3W	1.43						
	2064-2C4	4W	1.90						
	2064-2C5	5W	2.31						
	2064-2C6	6W	2.72						
	2064-2C7	7W	3.11						
	2064-2C8	8W	3.49						
	2064-2C9	9W	3.87						
	2064-210	10W	4.24						
	2064-211	11W	4.61						
	2064-212	12W	4.99						
	2064-213	13W	5.37						
	2064-214	14W	5.74						
	2064-215	15W	6.11						
	2064-216	16W	6.48						
zSeries 990									
	2084-301	1W	0.75						
	2084-302	2W	1.49						
	2084-303	3W	2.18						
	2084-304	4W	2.89						
	2084-305	5W	3.55						
	2084-306	6W	4.18						
	2084-307	7W	4.77						
	2084-308	8W	5.35						
	2084-309	9W	5.93						
	2084-310	10W	6.51						
	2084-311	11W	7.08						
	2084-312	12W	7.65						
	2084-313	13W	8.23						
	2084-314	14W	8.80						
	2084-315	15W	9.37						
	2084-316	16W	9.94						

Part 7 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

Model	# CP	Linux WASDB/L
System z9 EC		
2094-401	1W	0.34
2094-402	2W	0.67
2094-403	3W	0.98
2094-404	4W	1.29
2094-405	5W	1.61
2094-406	6W	1.89
2094-407	7W	2.16
2094-408	8W	2.43
2094-501	1W	0.67
2094-502	2W	1.31
2094-503	3W	1.93
2094-504	4W	2.53
2094-505	5W	3.16
2094-506	6W	3.71
2094-507	7W	4.24
2094-508	8W	4.76
2094-601	1W	0.81
2094-602	2W	1.59
2094-603	3W	2.35
2094-604	4W	3.08
2094-605	5W	3.84
2094-606	6W	4.52
2094-607	7W	5.17
2094-608	8W	5.80
2094-701	1W	1.00
2094-702	2W	1.97
2094-703	3W	2.91
2094-704	4W	3.82
2094-705	5W	4.73
2094-706	6W	5.57
2094-707	7W	6.36
2094-708	8W	7.14
2094-709	9W	7.91
2094-710	10W	8.68
2094-711	11W	9.44
2094-712	12W	10.20
2094-713	13W	10.97
2094-714	14W	11.73
2094-715	15W	12.50
2094-716	16W	13.25

Part 8 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

Model	# CP	Linux WASDB/L
System z10 EC		
2097-401	1W	0.40
2097-402	2W	0.77
2097-403	3W	1.14
2097-404	4W	1.50
2097-405	5W	1.85
2097-406	6W	2.19
2097-407	7W	2.53
2097-408	8W	2.86
2097-409	9W	3.19
2097-410	10W	3.51
2097-411	11W	3.82
2097-412	12W	4.13
2097-501	1W	0.86
2097-502	2W	1.66
2097-503	3W	2.44
2097-504	4W	3.20
2097-505	5W	3.94
2097-506	6W	4.67
2097-507	7W	5.38
2097-508	8W	6.07
2097-509	9W	6.75
2097-510	10W	7.42
2097-511	11W	8.07
2097-512	12W	8.71
2097-601	1W	1.16
2097-602	2W	2.24
2097-603	3W	3.29
2097-604	4W	4.31
2097-605	5W	5.30
2097-606	6W	6.27
2097-607	7W	7.22
2097-608	8W	8.14
2097-609	9W	9.04
2097-610	10W	9.91
2097-611	11W	10.77
2097-612	12W	11.60
2097-701	1W	1.67
2097-702	2W	3.22
2097-703	3W	4.72
2097-704	4W	6.17
2097-705	5W	7.57
2097-706	6W	8.93
2097-707	7W	10.25
2097-708	8W	11.53
2097-709	9W	12.77
2097-710	10W	13.98
2097-711	11W	15.15
2097-712	12W	16.28
2097-713	13W	17.38
2097-714	14W	18.45
2097-715	15W	19.51
2097-716	16W	20.53

Table 10. ITR Ratio Table - z/VM 5.4

The tables in this section present LSPR capacity data in the form of ITR ratios relative to the 2094-701. The data represents z/VM V5-R4.0 (z/Architecture) workload running in 64-bit real.

Part 1 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM e(logo)Server zSeries 800		
2066-0E1	1	0.06
2066-0A1	1	0.12
2066-0B1	1	0.18
2066-0C1	1	0.22
2066-0X2	2	0.24
2066-001	1	0.30
2066-0A2	2	0.39
2066-002	2	0.53
2066-003	3	0.75
2066-004	4	0.98
IBM e(logo)Server zSeries 890		
2086-110	1	0.04
2086-210	2	0.08
2086-310	3	0.11
2086-410	4	0.15
2086-120	1	0.07
2086-220	2	0.14
2086-320	3	0.20
2086-420	4	0.27
2086-130	1	0.14
2086-230	2	0.26
2086-330	3	0.39
2086-430	4	0.52
2086-140	1	0.18
2086-240	2	0.33
2086-340	3	0.49
2086-440	4	0.65
2086-150	1	0.28
2086-250	2	0.52
2086-350	3	0.77
2086-450	4	1.03
2086-160	1	0.34
2086-260	2	0.63
2086-360	3	0.94
2086-460	4	1.25
2086-170	1	0.59
2086-270	2	1.11
2086-370	3	1.65
2086-470	4	2.20

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 2 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z9 BC R07		
2096-A01	1	0.04
2096-A02	2	0.08
2096-A03	3	0.12
2096-B01	1	0.06
2096-B02	2	0.12
2096-B03	3	0.18
2096-C01	1	0.08
2096-C02	2	0.15
2096-C03	3	0.22
2096-D01	1	0.10
2096-D02	2	0.19
2096-D03	3	0.28
2096-E01	1	0.12
2096-E02	2	0.23
2096-F01	1	0.15
2096-F02	2	0.28
2096-G01	1	0.19
2096-H01	1	0.22
2096-I01	1	0.26
2096-J01	1	0.30
IBM System z9 BC S07		
2096-K04	4	0.36
2096-L03	3	0.33
2096-L04	4	0.43
2096-M03	3	0.41
2096-M04	4	0.54
2096-N02	2	0.36
2096-N03	3	0.52
2096-N04	4	0.68
2096-O02	2	0.42
2096-O03	3	0.62
2096-O04	4	0.81
2096-P02	2	0.49
2096-P03	3	0.71
2096-P04	4	0.93
2096-Q02	2	0.56
2096-Q03	3	0.81
2096-Q04	4	1.07
2096-R01	1	0.33
2096-R02	2	0.63
2096-R03	3	0.91
2096-R04	4	1.20
2096-S01	1	0.37
2096-S02	2	0.70
2096-S03	3	1.02
2096-S04	4	1.34
2096-T01	1	0.42
2096-T02	2	0.79
2096-T03	3	1.15
2096-T04	4	1.51
2096-U01	1	0.46
2096-U02	2	0.88
2096-U03	3	1.28
2096-U04	4	1.68
2096-V01	1	0.52
2096-V02	2	0.99
2096-V03	3	1.44
2096-V04	4	1.89
2096-W01	1	0.59
2096-W02	2	1.11
2096-W03	3	1.62
2096-W04	4	2.12
2096-X01	1	0.66
2096-X02	2	1.24
2096-X03	3	1.81
2096-X04	4	2.38
2096-Y01	1	0.73
2096-Y02	2	1.38
2096-Y03	3	2.02
2096-Y04	4	2.64
2096-Z01	1	0.83
2096-Z02	2	1.57
2096-Z03	3	2.30
2096-Z04	4	3.01

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 3 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z10 BC		
2098-A01	1	0.04
2098-A02	2	0.06
2098-A03	3	0.09
2098-A04	4	0.11
2098-A05	5	0.14
2098-B01	1	0.04
2098-B02	2	0.07
2098-B03	3	0.11
2098-B04	4	0.14
2098-B05	5	0.16
2098-C01	1	0.06
2098-C02	2	0.10
2098-C03	3	0.14
2098-C04	4	0.18
2098-C05	5	0.22
2098-D01	1	0.07
2098-D02	2	0.12
2098-D03	3	0.17
2098-D04	4	0.22
2098-D05	5	0.26
2098-E01	1	0.08
2098-E02	2	0.14
2098-E03	3	0.20
2098-E04	4	0.25
2098-E05	5	0.30
2098-F01	1	0.09
2098-F02	2	0.15
2098-F03	3	0.22
2098-F04	4	0.28
2098-F05	5	0.34
2098-G01	1	0.10
2098-G02	2	0.18
2098-G03	3	0.26
2098-G04	4	0.33
2098-G05	5	0.40
2098-H01	1	0.12
2098-H02	2	0.20
2098-H03	3	0.29
2098-H04	4	0.37
2098-H05	5	0.45
2098-I01	1	0.13
2098-I02	2	0.23
2098-I03	3	0.32
2098-I04	4	0.42
2098-I05	5	0.50
2098-J01	1	0.15
2098-J02	2	0.26
2098-J03	3	0.36
2098-J04	4	0.46
2098-J05	5	0.56
2098-K01	1	0.16
2098-K02	2	0.29
2098-K03	3	0.40
2098-K04	4	0.52
2098-K05	5	0.63
2098-L01	1	0.19
2098-L02	2	0.34
2098-L03	3	0.48
2098-L04	4	0.61
2098-L05	5	0.74
2098-M01	1	0.22
2098-M02	2	0.39
2098-M03	3	0.55
2098-M04	4	0.70
2098-M05	5	0.85

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 4 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z10 BC (continued)		
2098-N01	1	0.25
2098-N02	2	0.44
2098-N03	3	0.63
2098-N04	4	0.80
2098-N05	5	0.97
2098-O01	1	0.28
2098-O02	2	0.50
2098-O03	3	0.70
2098-O04	4	0.90
2098-O05	5	1.09
2098-P01	1	0.32
2098-P02	2	0.55
2098-P03	3	0.78
2098-P04	4	1.00
2098-P05	5	1.21
2098-Q01	1	0.35
2098-Q02	2	0.62
2098-Q03	3	0.88
2098-Q04	4	1.12
2098-Q05	5	1.36
2098-R01	1	0.40
2098-R02	2	0.70
2098-R03	3	0.98
2098-R04	4	1.26
2098-R05	5	1.53
2098-S01	1	0.45
2098-S02	2	0.79
2098-S03	3	1.11
2098-S04	4	1.42
2098-S05	5	1.72
2098-T01	1	0.51
2098-T02	2	0.89
2098-T03	3	1.25
2098-T04	4	1.60
2098-T05	5	1.94
2098-U01	1	0.57
2098-U02	2	1.00
2098-U03	3	1.41
2098-U04	4	1.80
2098-U05	5	2.19
2098-V01	1	0.65
2098-V02	2	1.13
2098-V03	3	1.60
2098-V04	4	2.05
2098-V05	5	2.48
2098-W01	1	0.73
2098-W02	2	1.28
2098-W03	3	1.80
2098-W04	4	2.30
2098-W05	5	2.79
2098-X01	1	0.82
2098-X02	2	1.44
2098-X03	3	2.03
2098-X04	4	2.60
2098-X05	5	3.16
2098-Y01	1	0.95
2098-Y02	2	1.67
2098-Y03	3	2.35
2098-Y04	4	3.01
2098-Y05	5	3.65
2098-Z01	1	1.05
2098-Z02	2	1.84
2098-Z03	3	2.59
2098-Z04	4	3.32
2098-Z05	5	4.02

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 5 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM e(logo)Server zSeries 900 (1XX Series)		
2064-101	1	0.38
2064-102	2	0.70
2064-103	3	1.01
2064-104	4	1.32
2064-105	5	1.61
2064-106	6	1.88
2064-107	7	2.14
2064-108	8	2.38
2064-109	9	2.60
2064-1C1	1	0.39
2064-1C2	2	0.73
2064-1C3	3	1.07
2064-1C4	4	1.41
2064-1C5	5	1.75
2064-1C6	6	2.07
2064-1C7	7	2.38
2064-1C8	8	2.68
2064-1C9	9	2.98
2064-110	10	3.27
2064-111	11	3.55
2064-112	12	3.83
2064-113	13	4.10
2064-114	14	4.36
2064-115	15	4.61
2064-116	16	4.86
IBM e(logo)Server zSeries 900 (2XX Series)		
2064-2C1	1	0.48
2064-2C2	2	0.88
2064-2C3	3	1.29
2064-2C4	4	1.70
2064-2C5	5	2.10
2064-2C6	6	2.48
2064-2C7	7	2.86
2064-2C8	8	3.22
2064-2C9	9	3.57
2064-210	10	3.91
2064-211	11	4.25
2064-212	12	4.58
2064-213	13	4.90
2064-214	14	5.21
2064-215	15	5.51
2064-216	16	5.80

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 6 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM e(logo)Server zSeries 990		
2084-301	1	0.75
2084-302	2	1.41
2084-303	3	2.06
2084-304	4	2.71
2084-305	5	3.32
2084-306	6	3.90
2084-307	7	4.45
2084-308	8	4.98
2084-309	9	5.49
2084-310	10	6.00
2084-311	11	6.49
2084-312	12	6.96
2084-313	13	7.43
2084-314	14	7.88
2084-315	15	8.33
2084-316	16	8.76

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 7 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z9 EC		
2094-401	1	0.34
2094-402	2	0.65
2094-403	3	0.95
2094-404	4	1.25
2094-405	5	1.53
2094-406	6	1.80
2094-407	7	2.05
2094-408	8	2.29
2094-501	1	0.67
2094-502	2	1.26
2094-503	3	1.85
2094-504	4	2.42
2094-505	5	2.97
2094-506	6	3.49
2094-507	7	3.98
2094-508	8	4.45
2094-601	1	0.81
2094-602	2	1.53
2094-603	3	2.24
2094-604	4	2.93
2094-605	5	3.60
2094-606	6	4.23
2094-607	7	4.82
2094-608	8	5.39
2094-701	1	1.00
2094-702	2	1.89
2094-703	3	2.76
2094-704	4	3.62
2094-705	5	4.44
2094-706	6	5.21
2094-707	7	5.95
2094-708	8	6.65
2094-709	9	7.33
2094-710	10	8.00
2094-711	11	8.65
2094-712	12	9.29
2094-713	13	9.91
2094-714	14	10.51
2094-715	15	11.11
2094-716	16	11.68

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 8 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z10 EC		
2097-401	1	0.36
2097-402	2	0.67
2097-403	3	0.96
2097-404	4	1.24
2097-405	5	1.52
2097-406	6	1.79
2097-407	7	2.06
2097-408	8	2.33
2097-409	9	2.60
2097-410	10	2.86
2097-411	11	3.12
2097-412	12	3.38
2097-501	1	0.77
2097-502	2	1.43
2097-503	3	2.06
2097-504	4	2.66
2097-505	5	3.25
2097-506	6	3.83
2097-507	7	4.40
2097-508	8	4.95
2097-509	9	5.50
2097-510	10	6.03
2097-511	11	6.55
2097-512	12	7.07
2097-601	1	1.04
2097-602	2	1.93
2097-603	3	2.78
2097-604	4	3.59
2097-605	5	4.38
2097-606	6	5.15
2097-607	7	5.90
2097-608	8	6.63
2097-609	9	7.35
2097-610	10	8.05
2097-611	11	8.73
2097-612	12	9.39
2097-701	1	1.50
2097-702	2	2.78
2097-703	3	3.99
2097-704	4	5.14
2097-705	5	6.25
2097-706	6	7.34
2097-707	7	8.39
2097-708	8	9.40
2097-709	9	10.39
2097-710	10	11.35
2097-711	11	12.28
2097-712	12	13.18
2097-713	13	14.06
2097-714	14	14.91
2097-715	15	15.74
2097-716	16	16.56
2097-717	17	17.35
2097-718	18	18.13
2097-719	19	18.88
2097-720	20	19.62
2097-721	21	20.34
2097-722	22	21.04
2097-723	23	21.73
2097-724	24	22.40
2097-725	25	23.05
2097-726	26	23.69
2097-727	27	24.32
2097-728	28	24.94
2097-729	29	25.56
2097-730	30	26.16
2097-731	31	26.76
2097-732	32	27.35

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 9 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z196		
2817-401	1	0.42
2817-402	2	0.78
2817-403	3	1.14
2817-404	4	1.49
2817-405	5	1.84
2817-406	6	2.17
2817-407	7	2.51
2817-408	8	2.84
2817-409	9	3.16
2817-410	10	3.48
2817-411	11	3.80
2817-412	12	4.11
2817-413	13	4.41
2817-414	14	4.71
2817-415	15	5.01
2817-501	1	1.02
2817-502	2	1.90
2817-503	3	2.76
2817-504	4	3.61
2817-505	5	4.43
2817-506	6	5.24
2817-507	7	6.04
2817-508	8	6.82
2817-509	9	7.58
2817-510	10	8.33
2817-511	11	9.06
2817-512	12	9.78
2817-513	13	10.48
2817-514	14	11.17
2817-515	15	11.84
2817-601	1	1.33
2817-602	2	2.47
2817-603	3	3.58
2817-604	4	4.67
2817-605	5	5.73
2817-606	6	6.77
2817-607	7	7.79
2817-608	8	8.78
2817-609	9	9.74
2817-610	10	10.69
2817-611	11	11.61
2817-612	12	12.52
2817-613	13	13.40
2817-614	14	14.26
2817-615	15	15.10

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Part 10 of 10

All capacity numbers are relative to the IBM 2094-701. The data represents a zVM workload running on System z. Information presented in this table is current as of May 2010.

Processor	#CP	zVM 5.4 High*/LV
IBM System z196 (continued)		
2817-701	1	2.08
2817-702	2	3.84
2817-703	3	5.56
2817-704	4	7.23
2817-705	5	8.86
2817-706	6	10.44
2817-707	7	11.97
2817-708	8	13.46
2817-709	9	14.91
2817-710	10	16.32
2817-711	11	17.69
2817-712	12	19.03
2817-713	13	20.32
2817-714	14	21.58
2817-715	15	22.80
2817-716	16	23.99
2817-717	17	25.16
2817-718	18	26.32
2817-719	19	27.46
2817-720	20	28.58
2817-721	21	29.69
2817-722	22	30.78
2817-723	23	31.86
2817-724	24	32.92
2817-725	25	33.97
2817-726	26	35.00
2817-727	27	36.02
2817-728	28	37.02
2817-729	29	38.01
2817-730	30	38.99
2817-731	31	39.95
2817-732	32	40.90

*** High RNI (Relative Nest Intensity) describes a workload category representing heavy use of the memory hierarchy - see workload description for more information**

Table 11. ITR Ratio Table - z/VM 5.2

The tables in this section present LSPR capacity data in the form of ITR ratios relative to the 2094-701. The data represents z/VM V5-R2.0 (z/Architecture) workload running in 64-bit real.

	Part 1 of 8		
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008		
			z/VM
	Model	# CP	WASDB/LVm
zSeries 800			
	2066-0E1	1W	0.06
	2066-0A1	1W	0.13
	2066-0B1	1W	0.18
	2066-0X2	2W	0.26
	2066-0C1	1W	0.22
	2066-0A2	2W	0.42
	2066-001	1W	0.29
	2066-002	2W	0.56
	2066-003	3W	0.82
	2066-004	4W	1.06
zSeries 890			
	2086-110	1W	0.04
	2086-210	2W	0.07
	2086-310	3W	0.11
	2086-410	4W	0.14
	2086-120	1W	0.07
	2086-220	2W	0.14
	2086-320	3W	0.21
	2086-420	4W	0.27
	2086-130	1W	0.14
	2086-230	2W	0.27
	2086-330	3W	0.40
	2086-430	4W	0.52
	2086-140	1W	0.18
	2086-240	2W	0.35
	2086-340	3W	0.52
	2086-440	4W	0.67
	2086-150	1W	0.28
	2086-250	2W	0.56
	2086-350	3W	0.81
	2086-450	4W	1.06
	2086-160	1W	0.35
	2086-260	2W	0.68
	2086-360	3W	0.99
	2086-460	4W	1.30
	2086-170	1W	0.61
	2086-270	2W	1.19
	2086-370	3W	1.75
	2086-470	4W	2.28

	Part 2 of 8		
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008		
			z/VM
	Model	# CP	WASDB/LVm
System z9 BC R07			
	2096-A01	1W	0.04
	2096-A02	2W	0.09
	2096-A03	3W	0.13
	2096-B01	1W	0.06
	2096-B02	2W	0.12
	2096-B03	3W	0.18
	2096-C01	1W	0.08
	2096-C02	2W	0.15
	2096-C03	3W	0.22
	2096-D01	1W	0.10
	2096-D02	2W	0.20
	2096-D03	3W	0.29
	2096-E01	1W	0.12
	2096-E02	2W	0.23
	2096-F01	1W	0.15
	2096-F02	2W	0.29
	2096-G01	1W	0.19
	2096-H01	1W	0.22
	2096-I01	1W	0.26
	2096-J01	1W	0.30

	Part 3 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			z/VM						
	Model	# CP	WASDB/LVm						
System z9 BC S07									
	2096-K04	4W	0.37						
	2096-L03	3W	0.34						
	2096-L04	4W	0.45						
	2096-M03	3W	0.43						
	2096-M04	4W	0.56						
	2096-N02	2W	0.37						
	2096-N03	3W	0.54						
	2096-N04	4W	0.70						
	2096-O02	2W	0.44						
	2096-O03	3W	0.64						
	2096-O04	4W	0.83						
	2096-P02	2W	0.50						
	2096-P03	3W	0.74						
	2096-P04	4W	0.96						
	2096-Q02	2W	0.58						
	2096-Q03	3W	0.85						
	2096-Q04	4W	1.10						
	2096-R01	1W	0.33						
	2096-R02	2W	0.65						
	2096-R03	3W	0.95						
	2096-R04	4W	1.24						
	2096-S01	1W	0.37						
	2096-S02	2W	0.72						
	2096-S03	3W	1.06						
	2096-S04	4W	1.38						
	2096-T01	1W	0.42						
	2096-T02	2W	0.81						
	2096-T03	3W	1.19						
	2096-T04	4W	1.55						
	2096-U01	1W	0.46						
	2096-U02	2W	0.91						
	2096-U03	3W	1.33						
	2096-U04	4W	1.73						
	2096-V01	1W	0.52						
	2096-V02	2W	1.02						
	2096-V03	3W	1.49						
	2096-V04	4W	1.94						
	2096-W01	1W	0.58						
	2096-W02	2W	1.14						
	2096-W03	3W	1.67						
	2096-W04	4W	2.17						
	2096-X01	1W	0.65						
	2096-X02	2W	1.27						
	2096-X03	3W	1.87						
	2096-X04	4W	2.43						
	2096-Y01	1W	0.72						
	2096-Y02	2W	1.41						
	2096-Y03	3W	2.07						
	2096-Y04	4W	2.70						
	2096-Z01	1W	0.82						
	2096-Z02	2W	1.61						
	2096-Z03	3W	2.36						
	2096-Z04	4W	3.07						

	Part 4 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			z/VM						
	Model	# CP	WASDB/LVm						
System z10 BC									
	2098-A01	1	0.04						
	2098-A02	2	0.07						
	2098-A03	3	0.10						
	2098-A04	4	0.12						
	2098-A05	5	0.15						
	2098-B01	1	0.05						
	2098-B02	2	0.09						
	2098-B03	3	0.12						
	2098-B04	4	0.15						
	2098-B05	5	0.18						
	2098-C01	1	0.07						
	2098-C02	2	0.13						
	2098-C03	3	0.18						
	2098-C04	4	0.22						
	2098-C05	5	0.25						
	2098-D01	1	0.09						
	2098-D02	2	0.16						
	2098-D03	3	0.21						
	2098-D04	4	0.26						
	2098-D05	5	0.31						
	2098-E01	1	0.10						
	2098-E02	2	0.18						
	2098-E03	3	0.25						
	2098-E04	4	0.30						
	2098-E05	5	0.35						
	2098-F01	1	0.11						
	2098-F02	2	0.20						
	2098-F03	3	0.27						
	2098-F04	4	0.33						
	2098-F05	5	0.39						
	2098-G01	1	0.13						
	2098-G02	2	0.24						
	2098-G03	3	0.32						
	2098-G04	4	0.40						
	2098-G05	5	0.47						
	2098-H01	1	0.15						
	2098-H02	2	0.26						
	2098-H03	3	0.36						
	2098-H04	4	0.44						
	2098-H05	5	0.52						
	2098-I01	1	0.16						
	2098-I02	2	0.30						
	2098-I03	3	0.41						
	2098-I04	4	0.50						
	2098-I05	5	0.58						
	2098-J01	1	0.18						
	2098-J02	2	0.33						
	2098-J03	3	0.45						
	2098-J04	4	0.55						
	2098-J05	5	0.65						
	2098-K01	1	0.20						
	2098-K02	2	0.37						
	2098-K03	3	0.51						
	2098-K04	4	0.62						
	2098-K05	5	0.73						
	2098-L01	1	0.24						
	2098-L02	2	0.44						
	2098-L03	3	0.60						
	2098-L04	4	0.73						
	2098-L05	5	0.86						
	2098-M01	1	0.28						
	2098-M02	2	0.50						
	2098-M03	3	0.69						
	2098-M04	4	0.85						
	2098-M05	5	0.99						

	Part 5 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			z/VM						
	Model	# CP	WASDB/LVm						
System z10 BC continued									
	2098-N01	1	0.32						
	2098-N02	2	0.58						
	2098-N03	3	0.79						
	2098-N04	4	0.97						
	2098-N05	5	1.14						
	2098-O01	1	0.36						
	2098-O02	2	0.65						
	2098-O03	3	0.89						
	2098-O04	4	1.08						
	2098-O05	5	1.27						
	2098-P01	1	0.40						
	2098-P02	2	0.72						
	2098-P03	3	0.99						
	2098-P04	4	1.21						
	2098-P05	5	1.42						
	2098-Q01	1	0.45						
	2098-Q02	2	0.81						
	2098-Q03	3	1.11						
	2098-Q04	4	1.36						
	2098-Q05	5	1.59						
	2098-R01	1	0.50						
	2098-R02	2	0.90						
	2098-R03	3	1.24						
	2098-R04	4	1.52						
	2098-R05	5	1.78						
	2098-S01	1	0.56						
	2098-S02	2	1.01						
	2098-S03	3	1.39						
	2098-S04	4	1.70						
	2098-S05	5	2.00						
	2098-T01	1	0.63						
	2098-T02	2	1.14						
	2098-T03	3	1.56						
	2098-T04	4	1.91						
	2098-T05	5	2.24						
	2098-U01	1	0.71						
	2098-U02	2	1.28						
	2098-U03	3	1.75						
	2098-U04	4	2.14						
	2098-U05	5	2.51						
	2098-V01	1	0.80						
	2098-V02	2	1.44						
	2098-V03	3	1.97						
	2098-V04	4	2.42						
	2098-V05	5	2.83						
	2098-W01	1	0.89						
	2098-W02	2	1.61						
	2098-W03	3	2.21						
	2098-W04	4	2.71						
	2098-W05	5	3.18						
	2098-X01	1	1.00						
	2098-X02	2	1.81						
	2098-X03	3	2.48						
	2098-X04	4	3.04						
	2098-X05	5	3.56						
	2098-Y01	1	1.14						
	2098-Y02	2	2.07						
	2098-Y03	3	2.84						
	2098-Y04	4	3.48						
	2098-Y05	5	4.08						
	2098-Z01	1	1.26						
	2098-Z02	2	2.27						
	2098-Z03	3	3.12						
	2098-Z04	4	3.82						
	2098-Z05	5	4.48						

	Part 6 of 8		
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008		
			z/VM
	Model	# CP	WASDB/LVm
zSeries 900 (1XX Series)			
	2064-101	1W	0.37
	2064-102	2W	0.73
	2064-103	3W	1.07
	2064-104	4W	1.39
	2064-105	5W	1.71
	2064-106	6W	2.01
	2064-107	7W	2.29
	2064-108	8W	2.56
	2064-109	9W	2.80
	2064-1C1	1W	0.38
	2064-1C2	2W	0.75
	2064-1C3	3W	1.11
	2064-1C4	4W	1.45
	2064-1C5	5W	1.78
	2064-1C6	6W	2.09
	2064-1C7	7W	2.39
	2064-1C8	8W	2.68
	2064-1C9	9W	2.97
	2064-110	10W	3.25
	2064-111	11W	3.52
	2064-112	12W	3.79
	2064-113	13W	4.05
	2064-114	14W	4.31
	2064-115	15W	4.57
	2064-116	16W	4.81
zSeries 900 (2XX Series)			
	2064-2C1	1W	0.47
	2064-2C2	2W	0.91
	2064-2C3	3W	1.34
	2064-2C4	4W	1.75
	2064-2C5	5W	2.14
	2064-2C6	6W	2.52
	2064-2C7	7W	2.88
	2064-2C8	8W	3.22
	2064-2C9	9W	3.56
	2064-210	10W	3.88
	2064-211	11W	4.20
	2064-212	12W	4.51
	2064-213	13W	4.82
	2064-214	14W	5.11
	2064-215	15W	5.40
	2064-216	16W	5.68
zSeries 990			
	2084-301	1W	0.75
	2084-302	2W	1.47
	2084-303	3W	2.16
	2084-304	4W	2.81
	2084-305	5W	3.44
	2084-306	6W	4.03
	2084-307	7W	4.60
	2084-308	8W	5.14
	2084-309	9W	5.67
	2084-310	10W	6.18
	2084-311	11W	6.67
	2084-312	12W	7.16
	2084-313	13W	7.63
	2084-314	14W	8.08
	2084-315	15W	8.52
	2084-316	16W	8.94

	Part 7 of 8								
	All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008								
			z/VM						
	Model	# CP	WASDB/LVm						
System z9 EC									
	2094-401	1W	0.35						
	2094-402	2W	0.67						
	2094-403	3W	0.99						
	2094-404	4W	1.28						
	2094-405	5W	1.57						
	2094-406	6W	1.84						
	2094-407	7W	2.09						
	2094-408	8W	2.34						
	2094-501	1W	0.67						
	2094-502	2W	1.31						
	2094-503	3W	1.91						
	2094-504	4W	2.49						
	2094-505	5W	3.04						
	2094-506	6W	3.57						
	2094-507	7W	4.06						
	2094-508	8W	4.54						
	2094-601	1W	0.81						
	2094-602	2W	1.59						
	2094-603	3W	2.32						
	2094-604	4W	3.02						
	2094-605	5W	3.69						
	2094-606	6W	4.33						
	2094-607	7W	4.93						
	2094-608	8W	5.51						
	2094-701	1W	1.00						
	2094-702	2W	1.95						
	2094-703	3W	2.86						
	2094-704	4W	3.72						
	2094-705	5W	4.54						
	2094-706	6W	5.32						
	2094-707	7W	6.07						
	2094-708	8W	6.78						
	2094-709	9W	7.47						
	2094-710	10W	8.13						
	2094-711	11W	8.78						
	2094-712	12W	9.41						
	2094-713	13W	10.02						
	2094-714	14W	10.62						
	2094-715	15W	11.20						
	2094-716	16W	11.74						

Part 8 of 8

All capacity numbers are relative to the IBM 2094-701. Information presented in this table is current as of October, 2008

		z/VM
Model	# CP	WASDB/LVm
System z10 EC		
2097-401	1W	0.39
2097-402	2W	0.72
2097-403	3W	1.04
2097-404	4W	1.34
2097-405	5W	1.62
2097-406	6W	1.89
2097-407	7W	2.15
2097-408	8W	2.40
2097-409	9W	2.64
2097-410	10W	2.86
2097-411	11W	3.08
2097-412	12W	3.29
2097-501	1W	0.83
2097-502	2W	1.56
2097-503	3W	2.24
2097-504	4W	2.87
2097-505	5W	3.47
2097-506	6W	4.04
2097-507	7W	4.58
2097-508	8W	5.09
2097-509	9W	5.58
2097-510	10W	6.04
2097-511	11W	6.48
2097-512	12W	6.89
2097-601	1W	1.13
2097-602	2W	2.11
2097-603	3W	3.02
2097-604	4W	3.86
2097-605	5W	4.66
2097-606	6W	5.42
2097-607	7W	6.13
2097-608	8W	6.81
2097-609	9W	7.45
2097-610	10W	8.05
2097-611	11W	8.62
2097-612	12W	9.16
2097-701	1W	1.63
2097-702	2W	3.03
2097-703	3W	4.33
2097-704	4W	5.53
2097-705	5W	6.65
2097-706	6W	7.70
2097-707	7W	8.69
2097-708	8W	9.62
2097-709	9W	10.48
2097-710	10W	11.29
2097-711	11W	12.05
2097-712	12W	12.77
2097-713	13W	13.44
2097-714	14W	14.07
2097-715	15W	14.68
2097-716	16W	15.26

Appendix B. IBM Capacity Planning Tools

A number of tools have been developed by IBM to assist in understanding the effects on capacity with ***IBM System z*** processors when:

- Considering an upgrade to a new processor
- Implementing logical partitioning on a processor, or changing the partition configuration on a current processor.
- Migrating to a Parallel Sysplex environment
- Upgrading to more current versions of z/OS, CICS, or IMS
- Converting to 64-bit addressing (z/OS only)

All of these tools are available within IBM for use by your IBM or IBM Business Partner representative, who can assist you in assessing various aspects of capacity planning. Most of these tools must remain with the IBM representative who can work directly with you. Any output generated from these tools can be freely disseminated when capacity planning help is being provided.

Both the **zPCR** tool and the **SoftCap** tool are available directly to customers via the website noted in the abstracts on the following pages.

The tools referenced are developed by IBM's *Capacity Planning Support* (CPS) team, a part of IBM's Advanced Technical Skills, in Gaithersburg, Maryland.

The following pages include an abstract for each of these tools:

1. **zPCR** Processor Capacity Reference for IBM System z
2. **zCP3000** Performance Analysis and Capacity Planning
3. **zPSG** Processor Selection Guide for IBM System z
4. **SoftCap** Software Migration Capacity Planning Aid
5. **BWATOOL** Batch Workload Analysis Tool
6. **zMCAT** Migration Capacity Analysis Tool for IBM System z
7. **zTPM** Tivoli Performance Modeler for IBM System z
8. **zVM-Planner** z/VM Planner for Linux Guests on IBM System z Processors

zPCR Processor Capacity Reference

zPCR is a PC-based productivity tool under Windows, designed to provide capacity planning insight for ***IBM System z*** processors running various workload environments under z/OS, z/VM, and Linux. Capacity results are based on IBM's **LSPR** data. In addition to the current z/OS-1.11 workload data, Linux and z/VM are provided.

Capacity is presented relative to a Reference-CPU, which may be assigned as any 1-way processor with any capacity-scaling factor and metric. Function in **zPCR** includes:

1. **LSPR Processor Capacity Tables**: Displays processor capacity ratios for up to 10 SCP/workload environments. Specific workloads and the order presented is user controlled. Capacity tables provided include::
   - **Multi-image:** Each processor represents a partition configuration typical for the model. All partitions are assumed to be running the same z/OS LSPR workload. (General Purpose CPs only).
   - **Single-image**: Each processor represents a single shared partition with all CPs assigned. z/OS, z/VM, and Linux LSPR workloads can be displayed for General Purpose CPs; z/VM and Linux LSPR workload can be displayed on IFL CPs.
2. **LPAR Configuration Capacity Planning:** Any specific LPAR configuration can be defined. Capacity projections are generated for each partition as well as the LPAR host as a whole. The LPAR host processor can be configured with General purpose CPs, zAAPs, zIIPs, IFLs, and ICFs where valid. Partitions are defined, specifying type (General Purpose, IFL, or ICF), SCP/workload, and LP configuration (dedicated/shared with number of logical CPs), and weight/capping assignments zAAPs and zIIPs are always associated a General Purpose partition. Partition configurations can be created directly from MVS RMF data or from previous zPCR studies.

   In Advanced-mode, a current LPAR configuration can be cloned, and the result modified to represent an alternate host. Capacity comparisons between the alternate and the current can then be made, for the overall processor or by partition. Up to 5 alternate configurations can be created.

Results are presented in tables and graphs which can be easily captured for notes, presentations, or handouts. A complete study can be saved for future reference. A User's Guide, QuickStart Guide, integrated online help, and relevant documentation are included.

IBM customers can obtain **zPCR** via the Internet at:

www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381

zCP3000 Performance Analysis and Capacity Planning

zCP3000 is a PC-based productivity tool that runs under Windows. It is designed to do performance analysis and capacity planning functions for ***IBM System z*** processors, running various SCP and workload environments. It can also be used to graphically analyze logically partitioned processors and DASD configurations. Input normally comes from the customer's system logs via a separate tool:

1. From ***z/OS SMF***, using **CP2KEXTR**
2. From ***VM Monitor***, using **CP3KVMXT**
3. From ***VSE CPUMON***, using **VSE2EDF**

Some of **zCP3000**'s features include:

- Enterprise Analysis
- LPAR Analysis
- Workload Analysis
- Sysplex Analysis
- Coupling Facility Analysis
- CF Structure Analysis
- Performance Index Analysis
- zAAP/zIIP Potential Analysis
- zIIP DRDA Analysis
- ZIIP IPSec Analysis
- Capacity Report with graphs and "SmartText"
- Alternate Processor Analysis
- zTPM (simulation) Interface both passing configuration information to and integrating data from zTPM into zCP3000 Report document.
- CF Link Type/Distance Analysis
- Consolidation/Decentralization Analysis (Quick Migration Mode)
- Growth Analysis
- Performance Variable Analysis
- Health Check Analysis
- Metrics Report
- ESCON/FICON Aggregation Analysis
- Enterprise DASD Analysis
- Data Set Analysis
- Print Processor List
- RIOC Adjustment
- Sysplex Aggregation Report
- SYSID DASD Analysis
- Optional CSV output
- Capture Ratios
- Merge workloads
- SYSID Storage Analysis
- Tape Analysis

zPSG Processor Selection Guide

zPSG is a PC-based productivity tool under Windows. It is designed to provide sizing approximations for ***IBM System z*** processors intended to host a new application, planned to be implemented using popular, commercially available software products. Current application support includes:

- **WebSphere Application Server** on z/OS or Linux
- **WebSphere Portal Server** on z/OS or Linux
- **WebSphere Process Server** on z/OS or Linux
- **WebSphere Message Broker** on z/OS or Linux
- **WebSphere MQ** on z/OS or Linux
- **Apache Webserver** on Linux
- **DB2 Data Warehouse** on z/OS
- **DB2 Transaction** on z/OS

Additional software products will be added to **zPSG** depending on requirements and the availability of capacity planning data.

For each application, you will characterize the average transaction or the average user, selecting features of the software than will be exploited, and the frequency of their use. The application is sized to a single System z processor within the tool (the specific processor model is not surfaced). Then using an LSPR workload category deemed representative of the application, capacity projections are developed for all System z processors. Capacity is given in terms of the expected utilization or in terms of the expected transaction rate that can be supported at a given utilization (SDP or Saturation Design Point). Projections are available for any/all of the processors that are included in the associated LSPR table. A summary report is also available, documenting the capacity estimate in terms of MIPS, estimated zAAP/zIIP eligibility, sizing inputs, and assumptions. Results are presented in tables and graphs which can be captured for notes, presentations, or handouts. Studies can be saved for future reference.

zPSG is installed as a stand-alone tool.

SoftCap Software Migration Capacity Planning Aid

SoftCap is a PC-based productivity tool under Windows, designed to assess the effect on ***IBM System z*** processor capacity, when planning to upgrade to a more current release of z.OS (or OS/390) or a major subsystem (CICS or IMS). **SoftCap** assumes that hardware configuration remains constant while the software version changes. The capacity implication of an upgrade for the following software components can be assessed independently or in any combination.

- **MVS** OS/390 v2r10 through z/OS V1R11
- **CICS** CICS/MVS 2.0 through CICS/TS 4.1
- **IMS** IMS/ESA 3.1 through IMS/ESA 11.0

Software Upgrades

For **z/OS**, input required by **SoftCap** includes the current version/release and the utilization represented by each of the following components: **BATCH**, **TSO**, **WEB**, **CICS**, **DB2**, **IMS**, and **SYSTEM**. If ***Compatibility Mode*** is supported for the current release, migration to ***Goal Mode*** may be specified. The target OS version/release must also be specified (beginning with z/OS-1.8, the processor's N-way configuration is taken into account).

For both **CICS** and **IMS** software upgrades, the current and planned version/release, and a high-level description of the subsystem's implementation is required.

Results are provided for each software component, showing the net change in capacity and the effective change in processor utilization that can be expected. If upgrading multiple components, results are generated showing the combined effect on capacity.

Migration to 64-bit Addressing Mode

If the host is a **zSeries** processor currently running OS/390 V2R10, **SoftCap** can also assess the effect on capacity when migrating from 31-bit to 64-bit addressing mode (z/OS must run in 64-bit mode, while OS/390 V2R10 can run in either mode). This projection is based on current 31-bit storage contention metrics and workload mix that you provide..

IBM customers can obtain **SoftCap** via the Internet at:

www.ibm.com/support/docview.wss?uid=tss1prs268

BWATOOL Batch Workload Analysis Tool

BWATOOL is a productivity tool that runs as a z/OS batch job on ***IBM System z*** hardware. It is designed to analyze client SMF type 30 and 70 data, producing a report showing how long batch jobs run on the currently installed processor. Both CPU time and elapsed time are reported. Similar results can then be projected for any ***IBM System z*** processor model. Basic questions that can be answered by BWATOOL include:

1. What jobs are good candidates for running on any given processor?
2. How much would jobs benefit from running on a faster processor?
3. For jobs within a critical path (batch window), what overall change in elapsed time might occur with a new processor?

CPU time projections are based solely on processor speed, using IBM LSPR data for z/OS batch-oriented workloads. Elapsed time projections are based on processor speed, queue time, and I/O time.

Various parameters are available to filter the SMF data, to specify the current host model and target, to define the type of batch involved (CPU intensive, I/O intensive, or somewhere between), and control the various reports.

Five types of reports are available:

1. Job report
2. Job step report
3. Job time line report
4. Critical path report
5. Job information report.

Some of the reports include graphs.

A complete User's Guide describes how one can use the tool and interpret the results.

zMCAT Migration Capacity Analysis Tool

Capacity expectation for a new ***IBM System z*** processor model, relative to a currently installed processor, is usually determined by comparing the capacity of benchmark workloads such as those carried in LSPR data. Applications such as the ***Processor Capacity Reference for IBM System z (zPCR)*** enable the comparison of processors for a variety of workloads and workload mixes. But how close is the benchmark workload to your customer's real production workload? Is there a way to get a better idea of the actual change in throughput seen on the production system?

The ***IBM System z Migration Capacity Analysis Tool*** (**zMCAT**) is intended to do just that. It can be used to compare the performance of production workloads before and after migration of the system image to a new processor, even if the number of engines on the processor has changed. Workloads for which performance is to be analyzed must be carefully chosen because the power comparison may vary considerably due to differing use of system services, I/O rate, instruction mix, storage reference patterns, etc. This is why "your mileage may vary" from an internal throughput ratio (ITRR) based on LSPR benchmark data.

zMCAT is a workstation-based tool that enables interactive filtering of workloads using a variety of parameters, for example, excluding workloads that have too small a sample, or with too much variance. The upgrade analysis can focus on either batch jobs or the intervals of the online regions. Input to zMCAT originates in the SMF type 30 records that are processed by **ZMCATX**, the ***zMCAT Extractor*** utility. Multiple days or weeks in both the before and after period can be collected all at once and then transmitted to the workstation for analysis. The data are then used to establish consistent comparisons. The primary result of running a zMCAT analysis is the calculation of the processor speedup. As user filtering progresses, this value is kept current and displayed to the user. While the estimated impact of the upgrade is derived directly from the measurements of the workloads from the customer's environment, zMCAT also allows for the studied judgment of the analyst through the intervention of workload selection.

zTPM Tivoli Performance Modeler

zTPM (Tivoli Performance Modeler) is a productivity tool designed to let you build a model of a z/OS based ***IBM System z*** processor, and then run various "what if scenarios". **zTPM** uses simulation techniques to let you model the impact of changes on individual workload performance. **zTPM** runs on a Windows based PC. **zTPM** uses RMF or CMF reports as input. Based on these reports, **zTPM** can create summary charts showing LPAR as well as workload utilization. An automated Build function lets you build a model that represents the system for any reporting interval. Once the model is built, you can make changes to see the impact on workload performance. These can be changes to workload volumes or changes to the hardware configuration. You can model the impact of combining system images from multiple processors onto a single processor, or model the impact of workload growth over several predefined time intervals. You can also model the impact of changing LPAR definitions. **zTPM** also allows you to estimate the impact of latent demand when replacing a processor running at high utilization. A Wizard feature makes it simple to define multiple scenarios on a single screen. Once the scenarios are defined, **zTPM** will build and execute each modeling run to give you the results. **zTPM** also includes a 123 and Excel spreadsheet. These spreadsheets contain several buttons which simplify the process of consolidating the results in a single place, and charting the results in a user friendly format. Coincident with the announcement of the z10 family of processors, **zTPM** has been updated to support **HiperDispatch for z10 processors.**

An external version of **zTPM** is also available for sale to customers from the IBM Software Division as the ***Tivoli Performance Modeler for z/OS*** (5698-A18).

zVM-Planner z/VM Planner for Linux Guests

zVM-Planner is a PC-based productivity tool under Windows, designed to provide capacity planning insight for ***IBM System z*** processors running various Linux application environments as guests under z/VM. Capacity results are based on analysis of a variety of benchmarks, both Linux native and Linux under z/VM. The tool is generic concerning software release levels, generally applying to z/VM v5.1 and later.

zVM-Planner input consists primarily of VM guest definitions and capacity requirements for each intended Linux application (a variety of Linux applications are supported). The expected capacity requirement for each Linux guest (a required input) can usually be obtained using the companion **zPSG** sizing tool or Techline assistance. The combined guest capacity requirement is then determined for optimally complementary peaks and for totally concurrent peaks. The degree of peak concurrency is user-selectable between these values. The resulting capacity requirement is combined with that of VM to support the entire complement of guests. Another companion tool, **zPCR**, can then be used to identify a processor model and partition configuration that can accommodate the VM image. All capacity values must be relative to a Reference-CPU setting that is common for all tools involved.

For scenarios where new Linux guests are to be added to an existing VM image, a **zVM-Planner** model of the existing VM guest configuration should first be built. The new Linux guests can then be added to determine to overall VM capacity requirement.

For scenarios where a VM image is to be added to an existing host processor, a **zPCR** model of the existing LPAR configuration should first be built. A partition that can deliver the capacity required by the VM image can then be added. **zPCR** can help assess any processor upgrade that may be necessary to accommodate the VM image.

Several guest metrics are available to help balance how the overall capacity will be distributed, including the number of VCPs (virtual CPs), and Share and Capping assignments.

Results are presented in tables and graphs that can be captured for documentation purposes. **zVM-Planner** studies can be saved for future reference. Both a User's Guide and integrated context sensitive help are included.