Cheryl Watson's

TUNING Letter

2004, No. 2

A PRACTICAL JOURNAL OF z/OS TUNING AND MEASUREMENT ADVICE

Inside this issue...

This is a special issue; one that we've been researching for more than six months. The z990 series machines that debuted in September 2003 are the most powerful machines that IBM has ever produced. But not everybody is happy with them. It's taken us several months of investigation, but we think we understand the reasons behind the apparent underperformance. That's the focus of this issue; our analysis starts on page 26.

New customized LSPR workload mixes are at the heart of this problem. So before we can discuss z990 performance in depth, we need to provide an introduction and update to LSPRs (page 3). It will probably surprise most of our readers to discover their workloads probably don't match the traditional workload LSPR definitions.

Another factor is the move to faster and fewer processors. Although we wrote about this situation in 1998, it's more pronounced when moving to the faster z990s. So we've updated that article and have included it on page 40.

The 40th anniversary of the IBM S/360 was on April 7, 2004, and it brought a large number of significant announcements. We describe these announcements on page 45 in our *What's New?* section.

All of the performance results in this issue were obtained using our **BoxScore** software product. BoxScore was designed to show the differences in actual versus expected speed and capacity when moving to a new machine or environment. You'll need BoxScore or something like it to make sure that you got what you paid for.

The next issue is almost complete and will follow this one shortly.

We think this is one of the most valuable issues that we've ever produced. We hope you think so too.

This Issue:

Cheryl Watson's TUNING Letter

Management Summary

This entire TUNING Letter is devoted to the research we've been doing over the last six months regarding z990 expectations versus actual performance. The z990s became generally available in September 2003, and we started hearing rumors of underperformance in November. Then we started getting results back from customers using our BoxScore software product, which is designed to measure performance differences associated with hardware or software changes. The results didn't look good - with many customers reporting a 10 to 15% underperformance for batch and a 15-20% underperformance for CICS using the traditional methods of comparison.

At the end of 2003, IBM began to realize that most of the published LSPR ratings apply only to about 20% of the z990 customers. They have found that about 80% of their customers now run a low I/O density workload mix, which most closely matches the CB-L (commercial batch long) workload. To the casual observer, this would seem like good news because the CB-L workload tends to result in a higher MIPS rating than most other workloads. But, in fact, it's bad news: The increase in MIPS, when moving from a z900 to a z990 with the same or fewer processors, is less for CB-L than for any other workload. In one example, if you were expecting a 1000 MIPS increase in capacity based on an average MIPS rating, the use of CB-L ratings would lower that expectation to be only a 900 MIPS increase.

We have confirmed that when using these CB-L ratings most batch work falls within IBM's guidelines for acceptable performance (plus or minus 5% from the expected ratings). We have some concerns about this range of plus or minus 5% when sizing your capacity. On these very large machines, that 5% could amount to a range of 500 MIPS or more.

From our customers' BoxScore runs, we have also concluded that CICS is still not meeting even the minimal CB-L workload expectations. They are finding that CICS work is still taking from 5-10% more CPU than expected. We don't know why this occurs; IBM disagrees with us, but we stand behind our results. (They don't use CICS transaction level analysis, and we do.)

There are two main points in this issue: 1) You must approach the sizing of new machines in a different manner if you expect IBM to warranty your expectations; and 2) It's extremely important to insure that you received the added capacity that you purchased. ■

Cheryl Watson's TUNING Letter

2004 Number 2 (Vol. 14, No. 2)
Published 6 times a year
by Watson & Walker, Inc.
www.watsonwalker.com

Publisher: Tom Walker
Executive Editor: Cheryl Watson
Editor: Clark Kidd
Production Manager: Linda May

2004 SUBSCRIPTION RATES: Electronic version (CD-ROM & email) $765 per year. Printed version $545 per year, $600 outside North America. Payment may be made by a check drawn on a U.S. bank, a money order, any major credit card, or by wire transfer. Issues are mailed via first class mail. All back issues are available. Send all correspondence to: Watson & Walker, Inc., 2130 Bispham Road, Sarasota, Florida 34231, USA. Tel: 800-553-4562 or 941-924-6565. Fax: 941-924-4892. For customer service, send email to admin@watsonwalker.com.

Note: Implementation of any suggestions contained in this newsletter should be preceded by a controlled test and is the responsibility of the reader.

Focus: LSPR Update

We're concerned that many installations seem to be seeing disappointing results when they move to the new z990 processors. We think there are several reasons for this disappointment, which we'll explain in more detail starting on page 26. But for a complete understanding of what has happened, we first need to provide an update and explanation of IBM's Large System Performance Reference (LSPR) methodology that they use for processor sizing. Please review this section before continuing to the next article on *z990 Expectations*. It's critical to understanding what has happened over the last couple of years.

IBM publishes their LSPR tables and methodology on one of their Web sites [see REF001 in our Bibliography on page 24]. We've explained this LSPR methodology in our CPU Chart [REF002] and in previous TUNING Letters [REF003]. But there are several things that we've recently discovered that are very important to understanding the performance of the z990s. This article covers the following topics:

- History
- IBM's Sizing Tools
- New Benchmarking Workloads
- z900 LSPRs
- z990 LSPRs
- z890 LSPRs
- Bibliography

History

MVS LSPRs

Many years ago, IBM designed a set of seven benchmark workloads to represent the types of work being run by their customers. These workloads were TSO, CICS, DB2, IMS, CB84 (commercial batch using traditional techniques), CBW2 (commercial batch using newer data-in-memory techniques and heavy DB2), and FPC1 (scientific and engineering, and also representative of SAS work). We describe these workloads in detail in our OS/390 CPU Chart. These workloads were used over the course of several MVS releases.

IBM then ran stand-alone benchmarks for each workload on different machines to determine the ITR (Internal Throughput Rate) ratio between each machine and a "base" machine, which was given an ITR of 1. The ITR is calculated as the units of work completed divided by processor busy time in seconds. When an ITR is compared to a base machine's ITR, the result is an ITRR (ITR ratio). The higher the resulting ITRR value, the greater the capacity for that workload. A machine with an ITRR of .5 has only one quarter the capacity of a machine with an ITRR of 2.0. For our CPU Charts, we convert the ITRRs to MIPS so that they are easier to understand.

MIX - MIX Workload

At the same time, IBM also provided a mixed workload (which we have always named "MIX"). IBM calculated this MIX workload as the harmonic mean of the benchmarks for five of the seven workloads: TSO, CICS, DB2, IMS and CB84.

The harmonic mean is used instead of the arithmetic mean because these values represent rates instead of totals. The MIX workload ITR is calculated as:

MIX ITR = 1 / ((.20 / TSO_ITR) + (.20 / CICS_ITR) + (.20 / DB2_ITR) +
(.20 / IMS_ITR) + (.20 / CB84_ITR))

The value of .20 is used because it represents the percentage of each of the five types of workloads. As an example of the difference between the harmonic mean and the arithmetic mean for the z990 2084-332, the arithmetic mean would produce a total MIPS estimate of 9249 MIPS, whereas the harmonic mean would produce our published total MIPS estimate of 9059. This difference is minor and almost unnoticeable for small processors, but becomes more significant as the speed of the processor increases.

An excellent CMG 2000 paper was written about the use of harmonic mean by **Dr. Sudhir R. Nath** of Wells Fargo Service Company [REF009].

Figure 1 - OS/390 LSPRs for z900s

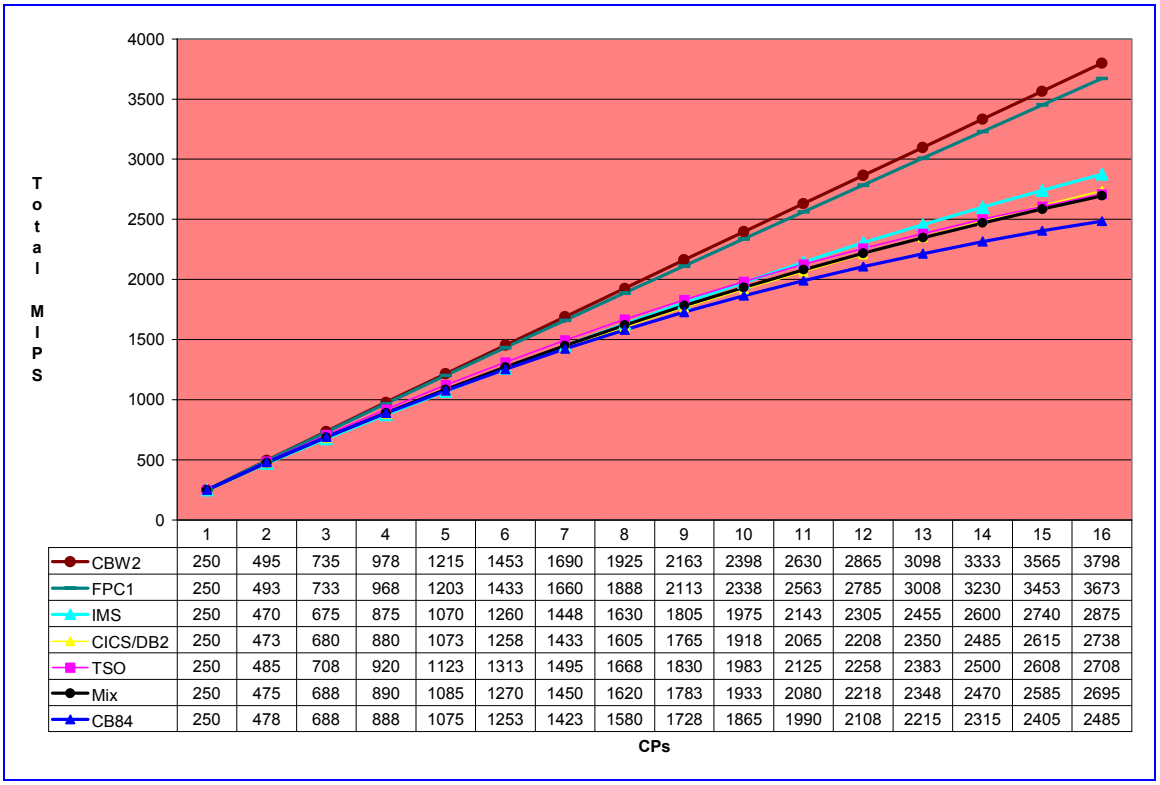

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CBW2	250	495	735	978	1215	1453	1690	1925	2163	2398	2630	2865	3098	3333	3565	3798
FPC1	250	493	733	968	1203	1433	1660	1888	2113	2338	2563	2785	3008	3230	3453	3673
IMS	250	470	675	875	1070	1260	1448	1630	1805	1975	2143	2305	2455	2600	2740	2875
CICS/DB2	250	473	680	880	1073	1258	1433	1605	1765	1918	2065	2208	2350	2485	2615	2738
TSO	250	485	708	920	1123	1313	1495	1668	1830	1983	2125	2258	2383	2500	2608	2708
Mix	250	475	688	890	1085	1270	1450	1620	1783	1933	2080	2218	2348	2470	2585	2695
CB84	250	478	688	888	1075	1253	1423	1580	1728	1865	1990	2108	2215	2315	2405	2485

Change in CICS & DB2 LSPRs

In June of 1999, the IBM LSPR group changed the benchmark jobs to more adequately reflect the workloads commonly being run in the field. In doing so, they combined the CICS and DB2 workloads into one benchmark and changed the base operating system from OS/390 R1 to OS/390 R4. IBM used the 9672-R15 processor as a base machine. The new MIX workload was the harmonic mean of the TSO, CICS/DB2, IMS and CB84 workloads. This greatly changed the meaning of MIX, because it caused TSO, IMS and CB84 to each become 25% of the workload instead of 20%, and CICS/DB2 was reduced from 40% to 25% of the workload.

Figure 1 shows an example of these workloads for the z900 models 1C1-1C9 and 110-116 using the OS/390 LSPRs after this change in 1999. We have converted the ITRRs to MIPS, but you can see the wide variation at the higher n-ways between CBW2 and CB84. Notice that the MIX is quite low in this chart because the FPC1 and CBW2 workloads are not used in the MIX.

Introduction of 64-bit LSPRs

In December of 2001, IBM produced a new set of LSPR ratings based on OS/390 R10 running on a z900 2064-1C1 base machine. What was very confusing about this set of ratings was that the TSO and CICS/DB2 workloads were run in 64-bit mode, but the other workloads were run in 31-bit mode. The MIX workload of TSO, CICS/DB2, IMS and CB84 was, therefore, a mix of incompatible (and physically impossible) benchmarks.

zSeries LSPRs

In May of 2003, IBM once again changed the workloads. These benchmarks were run using z/OS 1.4 and based on a z990 2084-301 machine. IMS, CB84 and CBW2 remained the same benchmarks, but were renamed to OLTP-T (OnLine Transaction Processing - Traditional), CB-S (Commercial Batch - Short) and CB-L (Commercial Batch - Long), respectively. The TSO and FPC1 workloads were dropped. The CICS/DB2 workload was modified to include a front-end Web access and was renamed to OLTP-W (OLTP Web-Enabled). And a new workload, WASDB (WebSphere Application and Database), was introduced. All of the workloads were run in 64-bit mode. The new MIX workload was now calculated as the harmonic mean of these five: OLTP-T, OLTP-W, CB-S, CB-L and WASDB. These are the current workloads for the zSeries and are described below.

From our perspective, there were two major flaws in these changes. According to IBM, they dropped TSO because it no longer represented a significant workload in most installations. This has not been our experience, and we would have liked to see the TSO benchmark retained. Secondly, there have been no published comparisons for these new workloads running in 31-bit mode and in 64-bit mode. This makes it impossible to determine what effect the architecture had on the results. This would have been extremely useful information and we wish it had been published and made available. Based on past experience, we believe that TSO will tend to match the MIX workload.

Another difference in the May 2003 LSPRs was the addition of PR/SM. Prior benchmarks had been run in basic mode; all the new benchmarks were run in LPAR mode. Thus, the overhead of running in LPAR mode (typically .5 to 1.0 percent) is now included in the LSPR ratings. Remember that the majority of "LPAR overhead" in most installations comes from multiple LPARs that are trying to share the same physical CPs. This overhead increases as the number of logical CPs exceeds the number of physical CPs. Because the LSPR benchmarks never use more logical CPs than available physical CPs, this larger (and more variable) overhead that many users will experience is not re-

flected in the LSPR ratings. We've written about LPAR overhead several times in our TUNING Letters. Please see REF015 and REF016 for some articles on this topic.

Here is a more detailed view of each of the new workloads (these descriptions come from our zSeries CPU Chart):

CB-L MIPS - Commercial Batch Long Job Steps

The CB-L workload is the same as the previous CBW2 workload, but has been renamed. This is a commercial batch workload that is most representative of new applications that exploit ESA functions, such as data in memory. It consists of programs written in C, COBOL, FORTRAN and PL/I. The steps include sorts, DFSMS utilities, compiles, VSAM and DB2 utilities, SQL processing, SLR processing, GDDM graphics and FORTRAN engineering routines. There is a lot of JES processing, and the workload spends about 50% of the time performing DB2 activities.

Even though this newer CB-L and the older CBW2 represent the same types of work, they are difficult to compare. If you are trying to compare CB-L MIPS with the CBW2 MIPS from an older CPU Chart, be aware that the CBW2 workloads were run in 31-bit mode on OS/390 R10, but the CB-L workloads are now run in 64-bit mode on z/OS. Therefore, you can't make a reliable comparison between the two workloads.

CB-S MIPS - Commercial Batch Short Job Steps

The CB-S workload is the same as the previous CB84 workload, but has been renamed. This represents a typical, traditional, view of commercial batch work. This workload consists of COBOL, Assembler H, and PL/I programs, along with compilers and utilities such as DFSORT. The BSAM, QSAM, BDAM, and VSAM access methods are used.

Once again, it is problematic to try and compare this CB-S with the older CB84 workload. The CB84 MIPS from our older CPU Chart represent work run in 31-bit mode on OS/390 R10. But the CB-S MIPS represent workloads run in 64-bit mode on z/OS. Therefore, you can't make a reliable comparison between the two workloads.

WASDB MIPS - WebSphere Application Server and Data Base

This is an entirely new workload that was created for the z/OS LSPRs. This workload uses WebSphere applications and a DB2 database. It contains Java classes, Java servlets, Java Server Pages (JSPs) and Enterprise Java Beans (EJBs). It emulates an online brokerage application and was written using IBM VisualAge for Java and WebSphere Studio tools.

OLTP-W MIPS - Web-enabled On-line Workload

The OLTP-W workload is similar to the previous CICS/DB2 workload, but has been changed to have a front-end Web-enabled access. This is very similar to most of the first Web-enabled applications in most installations. It has ten transaction types accessing DB2 databases. CP/SM (CICSplex/System Manager) is used to manage an MRO (Multi-Region Option) complex using a ratio of 1:3 between TORs (Terminal Owning Regions) and AORs (Application Owning Regions). These transactions are typical of an order entry and product delivery system.

The LSPR documentation provides more details on the architecture of the workload: "The J2EE (Java 2 Enterprise Edition) 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 V4.0.4. The application uses the J2EE architected Common Client Inter-

Figure 2 - z900 MIX Using Two LSPRs

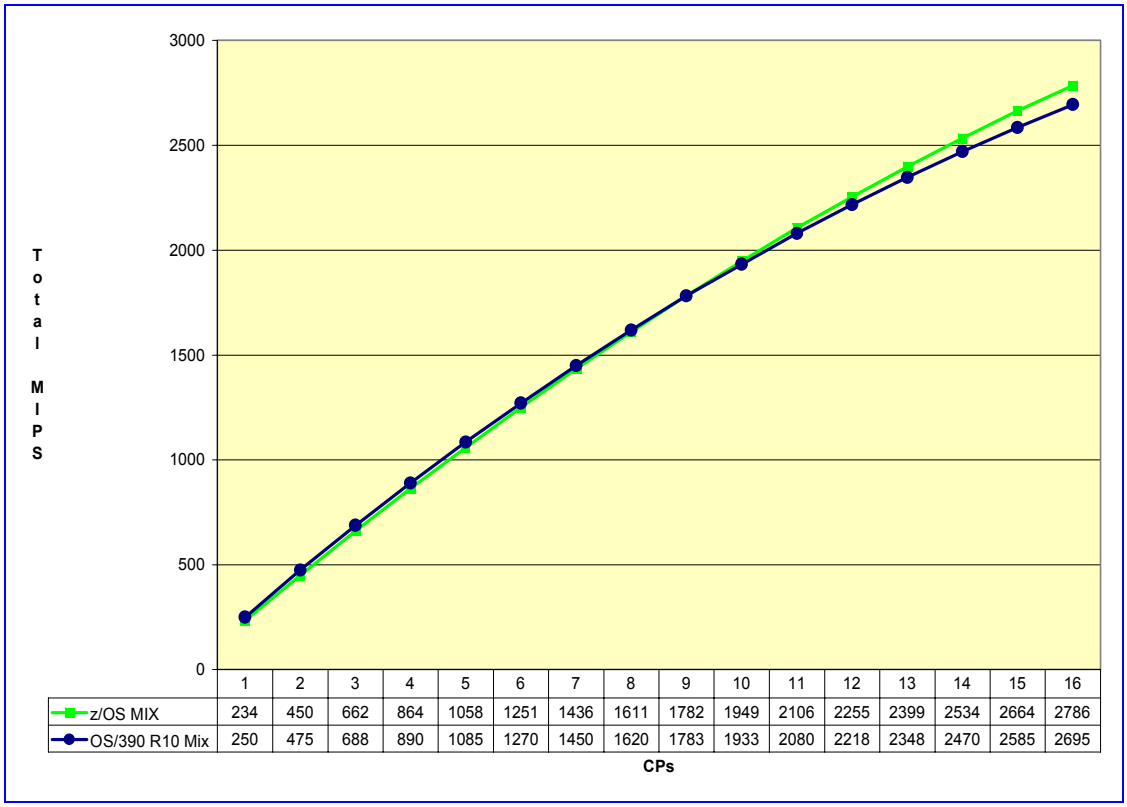

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
z/OS MIX	234	450	662	864	1058	1251	1436	1611	1782	1949	2106	2255	2399	2534	2664	2786
OS/390 R10 Mix	250	475	688	890	1085	1270	1450	1620	1783	1933	2080	2218	2348	2470	2585	2695

CPs

face (CCI). Clients access WebSphere services using the HTTP Transport Handlers. Then the appropriate servlet is run through the webcontainer, which calls EJBs in the EJB Container. Using the CTG ECI, CICS is called to invoke DB2 to access the database and obtain the information for the client."

Because the Web-enabled front-end has been added to this workload, you can't make a reliable comparison between this OLTP-W workload and the previous CICS/DB2 workload.

OLTP-T MIPS - Traditional On-Line Workload

The z/OS LSPR OLTP-T workload is the same as the previous IMS workload. It was renamed because the IMS workload produces similar performance to the CICS DLI applications and can be used for benchmarking both types of traditional online workloads. There are twelve transaction types. Enough MPRs (Message Processing Regions) are started to bring the system to the desired utilization (70% and 90%) without causing contention within an MPR. The DLI HDAM and HIDAM access methods are used with VSAM and OSAM databases. Batch Message Processing regions (BMPs) are not included.

If you are trying to compare the IMS MIPS from an older CPU Chart with the OLTP-T MIPS, be aware that in LSPRs up through OS/390 R10, the IMS workloads were run in 31-bit mode, but the OLTP-T workloads are run in 64-bit mode. Therefore, you can't make a reliable comparison between the two workloads.

Figure 3 - z900 MIX Using Two LSPRs

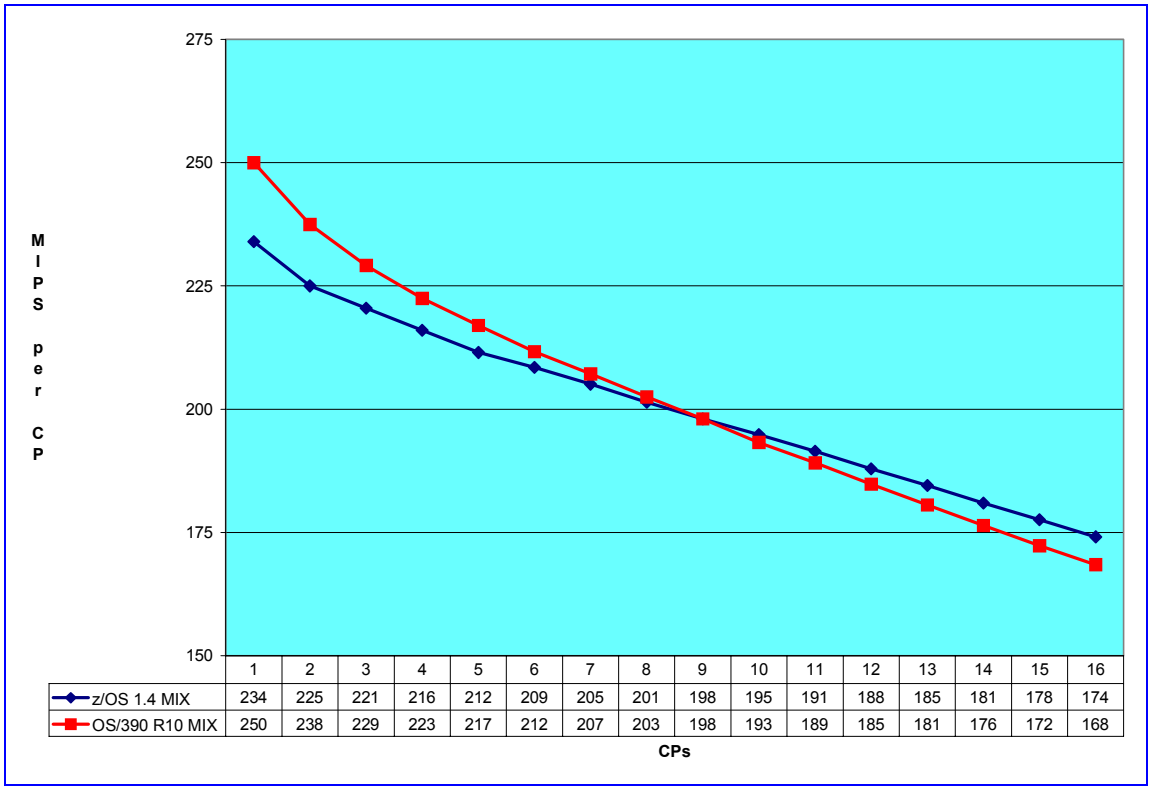

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
z/OS 1.4 MIX	234	225	221	216	212	209	205	201	198	195	191	188	185	181	178	174
OS/390 R10 MIX	250	238	229	223	217	212	207	203	198	193	189	185	181	176	172	168

Comparing Different LSPRs

IBM recommends that you not compare LSPR ratings that are based on different operating systems and different machines. We absolutely agree. That is why we have different CPU Charts that correspond to the different LSPR benchmarks. Our zSeries CPU Chart is based solely on the z/OS 1.4 benchmarks that are compared to the 2084-301. Our OS/390 CPU Chart is based on several of the benchmarks, but we indicate the LSPR tables used for each analysis and recommend that you not compare machines whose MIPS are not based on the same set of tables.

But we think it's important to understand how these LSPR versions differ because we know some installations will end up comparing them, despite the warnings. If you're going to do something dubious, you should at least understand the risk that you're taking. We've included several plots to help understand the difference in ratings. Here are some observations:

1. Figure 2 shows the OS/390 versus z/OS MIX MIPS for one of the z900 series machines (models 1C1-1C9 and 110-116). At the high end that plots 16 CPs, there is a difference of about 90 MIPS or about 3%. This chart refers to the total MIPS, or the capacity of the machine, based on two different sets of benchmarks.

2. Figure 3 shows a slightly different view of the same data. This shows MIPS per CP, which represents the speed of each CP. There are two things to notice from this plot. First, the scale of a plot can give you a very different view. Look at the 1-way. In Figure 2, the one-way capacity looked the same, simply because of the scale. In Figure 3, we see that the largest change in speed occurs in the 1-way rather than the 16-way. There is

Figure 4 - z900 Batch Workloads

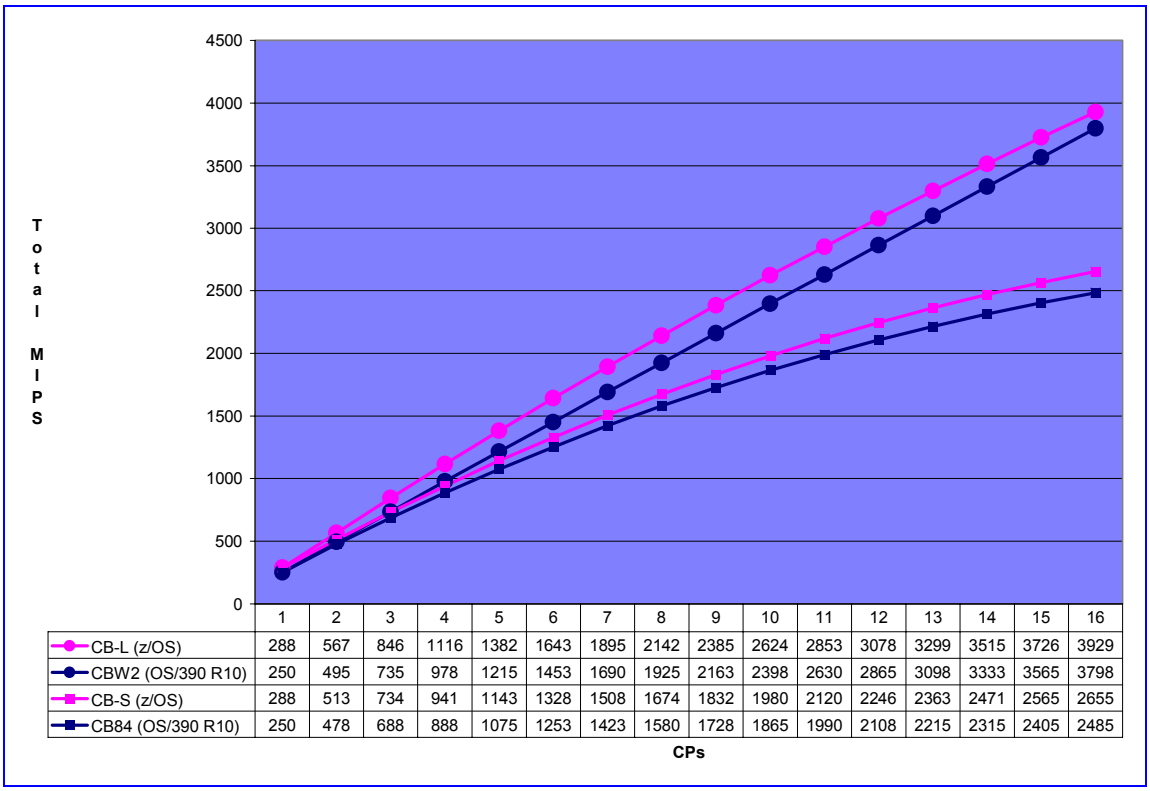

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CB-L (z/OS)	288	567	846	1116	1382	1643	1895	2142	2385	2624	2853	3078	3299	3515	3726	3929
CBW2 (OS/390 R10)	250	495	735	978	1215	1453	1690	1925	2163	2398	2630	2865	3098	3333	3565	3798
CB-S (z/OS)	288	513	734	941	1143	1328	1508	1674	1832	1980	2120	2246	2363	2471	2565	2655
CB84 (OS/390 R10)	250	478	688	888	1075	1253	1423	1580	1728	1865	1990	2108	2215	2315	2405	2485

actually a 7% difference in the speed of the 1-way compared to a 3% difference in the speed for the 16-way. The second thing we see is that the OS/390 R10 LSPRs result in higher MIPS for anything less than a 9-way, similar MIPS for a 9-way or 10-way, and lower MIPS for anything greater than a 10-way.

3. We also wanted to see how the workloads that are supposed to be fairly similar compare to one another. Figure 4 compares the associated MIPS for CB-L (z/OS 1.4) and CBW2 (OS/390 R10) workloads. See the top two lines. They are fairly close. The same is true for the older CB84 and newer CB-S (lower two lines). What's most important to note is that CB-L is used as part of the MIX workload for z/OS, but CBW2 is not used as part of the MIX workload for OS/390. This is one of the reasons that the z/OS MIX MIPS are higher than the OS/390 MIX MIPS in Figure 2.

4. Figure 5 shows a similar plot for the online workloads. Two interesting things show up here. First, the WASDB workload is the most effective at the higher CP levels. Secondly, the OLTP-W workload (that is a modification of the older CICS/DB2 workload) tracks more closely with the older IMS workload; and the OLTP-T workload (that is based on the older IMS) tracks more closely with the older CICS/DB2 workload. We have no explanation for this.

Summary

It is reasonable that IBM must make changes to their LSPR workloads to more accurately reflect the workloads being run in today's installations. Because of the changes, however, it makes it very difficult to determine whether the change in ITRR comes from the change in architecture (31-bit versus 64-bit), the host software (OS/390 to z/OS) or the base machine (z900 versus z990). In addition to

Figure 5 - z900 Online Workloads

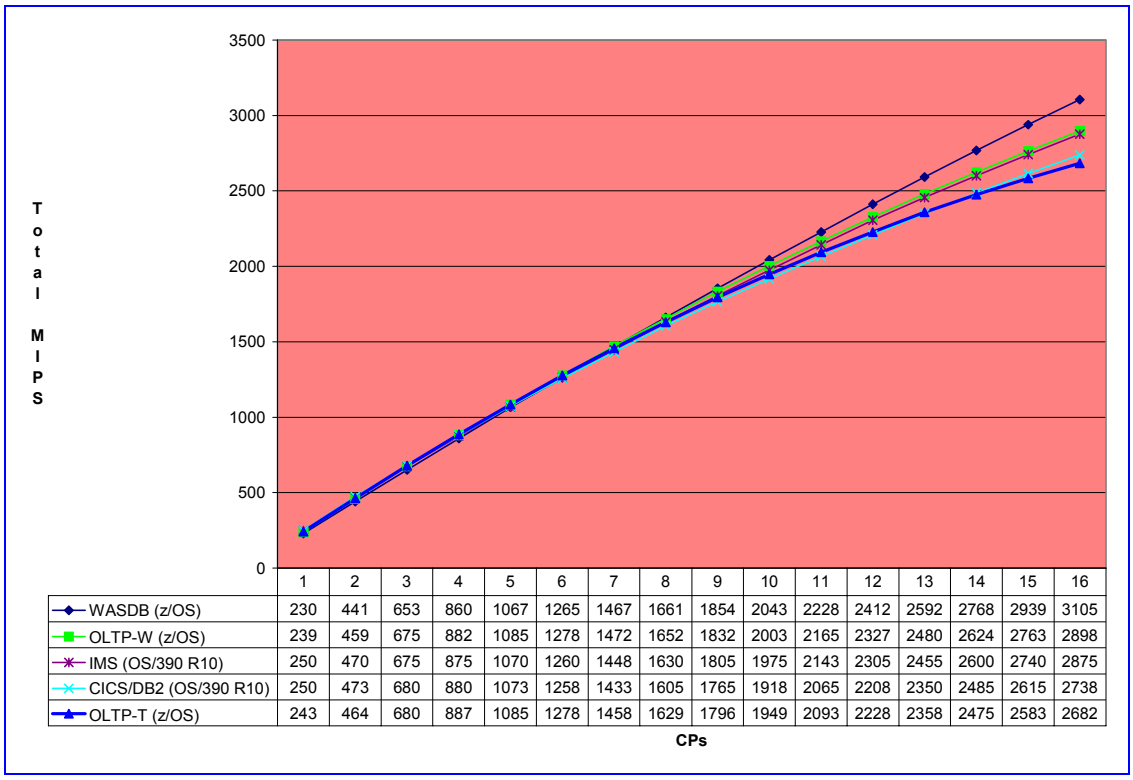

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
WASDB (z/OS)	230	441	653	860	1067	1265	1467	1661	1854	2043	2228	2412	2592	2768	2939	3105
OLTP-W (z/OS)	239	459	675	882	1085	1278	1472	1652	1832	2003	2165	2327	2480	2624	2763	2898
IMS (OS/390 R10)	250	470	675	875	1070	1260	1448	1630	1805	1975	2143	2305	2455	2600	2740	2875
CICS/DB2 (OS/390 R10)	250	473	680	880	1073	1258	1433	1605	1765	1918	2065	2208	2350	2485	2615	2738
OLTP-T (z/OS)	243	464	680	887	1085	1278	1458	1629	1796	1949	2093	2228	2358	2475	2583	2682

this, the meaning of the MIX workload changes even more because of the variability mentioned above and the differing combinations of workloads. This is one of the many reasons that IBM does not recommend using the MIX workload for capacity estimates. We don't approve of that either. We also know that most installations use only the averages. As we shall see, this also has its drawbacks. Estimates by workloads provide the only meaningful results, but there are some problems even with that technique.

IBM's Sizing Tools

IBM has several free tools to help customers plan for an upgrade. (These are in addition to the IBM fee-based consulting services.) The tools are available from your IBM representative or IBM business partner, and we don't really understand why more people don't take advantage of them. A brief description of these tools and some additional tools and services can be found in Appendix C of IBM's LSPR manual [REF004].

There is little published information about the tools other than a brief description in the LPSR manual, but you can obtain additional information from your representative. **JoAnne Brown** of IBM introduced some of the tools at session 2513 at the last SHARE [REF005] and showed the use of the zPCR tool along with some sample output. The tools are:

- **zPCR** - Processor Capacity Reference for zSeries. This is a PC-based tool that provides capacity relationships for zSeries machines. You can download data into a spreadsheet from the RMF Overview Report and the RMF CPU Activity Report.

- **PCRW** - Processor Capacity Reference (z800, z900 and earlier). This is similar to zPCR, but uses the OS/390 LSPRs instead of the z/OS LSPRs. A similar tool, PCR, can be used for older S/390 machines. Older machines running in LPAR mode can also use the LPAR/CE tool.

- **LPAR/CE** - LPAR Capacity Estimator. This is another PC-based tool that helps an installation understand the impact of changing the LPAR configuration of a machine. It also provides guidance when moving to a different machine in LPAR mode.

- **CP2000** - Performance Analysis and Capacity Planning. This is a PC-based tool, but has a mainframe component to collect the data it needs from SMF and RMF. This tool performs functions including capacity estimates, health check analysis, CPU analysis, workload analysis and many others. Graphs are provided.

- **SPSSZR** - S/390 Parallel Sysplex Quicksizer. This PC-based tool provides capacity planning recommendations for parallel sysplex environments.

- **SOFTCAP** - Software Migration Capacity Planning Aid. This tool is the only one available for download by the customer. It is PC-based and helps you understand the impact of moving between architectures, machines, operating systems or WLM goal/compat mode. [REF002]

- **BWATOOL** - Batch Workload Analysis Tool. This is an MVS batch job that processes SMF type 30 and type 70 records in order to analyze batch jobs. CPU time and elapsed time are analyzed for recommendations on the best processor matches or the critical path.

- **Unknown** - Capacity analysis software. Although we don't know the name of this software, the WSC team has products to analyze how well your new machine has met its capacity estimates. It's quite similar to our BoxScore, differing (from what we can tell) in its use of weighted means rather than average percent of change. This tool uses special customized workloads (described below) in order to estimate and confirm the new processor's capacity. We explain more about the weighted means on page 34.

When we polled the attendees at Cheryl's last SHARE "Hot Flashes" session, only four to five people out of more than 250 indicated they had ever used these tools. That's a shame! Before your next upgrade, be sure to contact your IBM representative for access to these tools.

Of course, if you want to work a bit harder, you can do much of this work on your own after reading the LSPR manual and understanding the concepts of ITRRs, harmonic means, and the techniques for capacity planning.

New Benchmarking Workloads

New Customized Workload Mixes

One of the most useful things that came out of JoAnne's SHARE session was a description of some new customized workload mixes that IBM is now using to characterize customer work: 'custom online', 'custom other' and 'low I/O density'. The workloads that we previously described (CB-L, CB-S, etc.) are now called 'primitives', and the new customized workload mixes are created using a combination of the primitives. **Walt Caprice** also mentioned these three workload mixes in his SHARE session [REF006].

When doing capacity planning or sizing with these tools, IBM typically divides the customer's workloads into two categories: online work (those workloads requiring a user response, such as CICS and IMS) and non-online work. They then use customized workload mixes for these two types of work. As an example, the custom online workload mix is the average of the WASDB, OLTP-W and OLTP-T primitive workloads. Of course, you can modify these combinations of primitives if you feel that your workloads are weighted more heavily towards specific types of work.

The first two new workload mixes have recently been added to our CPU Chart and BoxScore product:

CUSONL - Custom Online Workload Mix

The Custom Online workload is a custom workload mix designed for online applications. It is computed by taking the harmonic mean of the WASDB, OLTP-W and OLTP-T primitive workloads.

This is a specialized zSeries workload mix that you will not find in IBM's LSPR tables, but it may prove to more accurately meet your needs if you run the type of workload described.

CUSOTH - Custom Other Workload Mix

The Custom Other workload mix is designed for non-online applications. It is computed by taking the harmonic mean of the CB-L workload value and the CB-S workload value.

Again, this is a specialized zSeries workload mix that you will not find in IBM's LSPR tables, but it may prove to more accurately meet your needs if you run the type of workload described.

Low I/O Density Workload Mix

Additionally, it is IBM's belief that a large percentage (80% or so) of their customers now run in a "low I/O density" environment on z990s. IBM defines this as less than 30 DASD I/Os per second per unit of CPU usage as measured in MSUs. According to IBM, this situation results in an installation needing to use a much different customized workload mix than the two described above. The low I/O density workload mix is defined as a combination of 60% CB-L, 20% WASDB and 20% OLTP-W.

Figure 6 - RMF CPU and Workload Activity Extracts

```
                          P A R T I T I O N   D A T A   R E P O R T
...snip

 MVS PARTITION NAME                    SYSA
 IMAGE CAPACITY                          37
... snip

 --------- PARTITION DATA -----------------  -- LOGICAL PARTITION PROCESSOR DATA --
                    ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----
 NAME       S   WGT  DEF    ACT  DEF   WLM%  NUM   TYPE   EFFECTIVE       TOTAL

 SYSA       A    80    0     30  NO     0.0    2   CP    00.24.24.877  00.24.30.583
 SYSB       A     5    0      2  YES    0.0    2   CP    00.01.16.196  00.01.20.202
 SYSC       A    15    0      5  NO     0.0    2   CP    00.03.45.182  00.03.52.502
... snip

*******************************************************************************

                        W O R K L O A D   A C T I V I T Y
...snip

   REPORT BY: POLICY=STANDARD
                     Standard Policy

   TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----
   AVG    110.39   ACTUAL            1.14.263   SSCHRT 881.8   IOC    579383
   MPL    110.38   EXECUTION         1.14.206   RESP      1.6  CPU      6411K
   ENDED    2276   QUEUED                  25   CONN      1.1  MSO      4658K
   END/S    2.53   R/S AFFINITY             0   DISC      0.2  SRB    376059
   #SWAPS   2172   INELIGIBLE               8   Q+PEND    0.3  TOT     12025K
   EXCTD       0   CONVERSION               0   IOSQ      0.0  /SEC    13359
   AVG ENC  0.08   STD DEV          58.03.846
   REM ENC  0.00
   MS ENC   0.00
```

You can compute your own I/O density by looking at the RMF Workload Activity Report for the total interval and the RMF Partition Data Report (both shown in Figure 6). In this example, the DASD SSCH rate for the entire interval (Policy summary report) is 881.8 SSCHs per second. There were 30 MSUs used during the interval, so the I/O density is 881.8 / 30 = 29.4. This is only slightly under the guideline of 30. We're finding many sites with I/O density rates of 12 to 18. Please note that the IBM published MSUs for the z990 and z890 are for software pricing and are 10% less than the approximate capacity of the machine. If this were a z890 or z990, the calculation would be: 881.8 / (30 /.90) = 26.5.)

You can also find a SAS program on our Web site to calculate the I/O density for each of your systems over many intervals at **www.watsonwalker.com/lowio.txt**. The program is designed to handle data stored in an MXG, NeuMICS or ITSV (ITRM) SAS database. An example of output from this program is shown in Figure 7. Please notice how the I/O density can vary a great deal from one LPAR to the next running on the same physical machine. Also, notice that all LPARs would be considered "low I/O," because they are running I/O densities lower than IBM's limit of 30.

The concept of low I/O was introduced by IBM in 1999 when installations were upgrading from a G4 to a G5 or G6 machine. These new machines were significantly faster than the G4 and IBM determined that installations with a low I/O density appeared to perform more like CBW2 (which is very CPU intensive) than any other workload. (The CBW2 workload was renamed to CB-L for the zSeries LSPRs.) This was documented in a paper written by **Kathy Walsh** of WSC in 1999 and last updated March 2002 [REF008]. This paper only referred to non-online work being affected by the low I/O density. A similar paper was written in October 2003 by **Marty Dietch** of IBM, but it only referenced the older machines and not the z990 [REF011].

Figure 7 - I/O Density Report

```
BoxScore V1.6, Watson and Walker, Inc.
17:08 Wednesday, March 24, 2004

I/O Density - DASD I/Os per Second per MSU

                                    Average     Minimum     Maximum
                                      I/O         I/O         I/O
 System     CPU Type     MVS Level  Density     Density     Density

  SYSA      2064-110     z/OS 01.04    18.5        5.3      1475.3
  SYSB      2064-110     z/OS 01.04    18.1        4.0       458.7
  SYSC      2084-310     z/OS 01.04    29.7        4.8      3496.2
  SYSD      2084-310     z/OS 01.04    15.6        1.3      1294.0
  SYSE      2084-310     z/OS 01.04    24.1        2.8      1166.1
```

IBM now feels that the low I/O classification also relates to the z990s. The original LSPR manual [REF004] didn't mention low I/O when the z990s came out. On 15Jan2004, a brief mention of it was added to the LSPR manual [REF004]. Unfortunately, this reference was not added to the Web site, but only as part of Chapter 4 (*Using LSPR Data*) in the manual. An update was made on 2Apr2004 that contains the z890 updates and describes the three new customized workload mixes, including the low I/O content workload mix in Chapter 4.

Based on IBM's performance metrics at the back of the LSPR manual, we calculate the I/O densities for each of the IBM workloads to be:

CB-L = 28
CB-S = 62
WASDB = 9
OLTP-W = 6
OLTP-T = 52

IBM claims that this low I/O density situation occurs because of the move to significantly faster CPs. They say that this was true with the previous G4 to G5/G6 machines, and now occurs when moving from z900s to z990s. Although we are not sure we can justify the use of CB-L based on intuitive logic (and because of the I/O densities mentioned above), we agree with IBM that it much more closely approaches the actual performance that our customers are seeing. The problem for most installations occurs because this is so different from prior years where people successfully used the primitive workloads such as CB84, TSO, CICS/DB2 and IMS for both capacity planning and before/after comparisons.

Because IBM now feels the low I/O condition applies to 80% of the z990 installations, we've also added it as a third customized workload mix to our zSeries CPU Chart and BoxScore.

LOIO - Low I/O Workload Mix

The low I/O workload mix is a custom workload mix designed for a system that has a low I/O rate. It is computed in the IBM manner by taking the harmonic mean of 60% of the CB-L workload, plus 20% of the WASDB workload, plus 20% of the OLTP-W workload.

This is a specialized zSeries workload mix that you will not find in IBM's LSPR tables, but it is described in the LSPR manual. It may prove to more accurately match your environment if you run on an image that has a low I/O density.

What does this mean to you? If you're one of the 80% of installations who have a low I/O density, then your expectations should be much closer to the low I/O workload mix than to the MIX workload.

And Even Newer Workloads

Just as we were completing this newsletter, we discovered one more reference to customized workload mixes. We don't know yet how these newly defined workload mixes will be used by IBM.

The **zSeries 990 Technical Guide, SG24-6947-00**, didn't mention the low I/O concept or customized workload mixes in the original Jul2003 edition. The draft (**SG24-6947-01)** that was published 7Apr2004 contains several new customized workload mixes, including the low I/O content workload mix (section 8.8.1).

Table 1 shows the new workload mixes described in this z990 Technical Guide. We don't know whether these will eventually be documented in the LSPR manual or not.

Table 1 - New z990 Customized Workloads

Workload	CB-L	CB-S	WASDB	OLTP-W	OLTP-T
MIX	20%	20%	20%	20%	20%
TI-MIX	30%	10%		30%	30%
TD-MIX	45%	15%		20%	20%
TM-MIX	52.5%	17.5%		15%	15%
CB-MIX	75%	25%			
LoIO-MIX	60%		20%	20%	

The workloads are described as:

- TI-MIX - Transaction intensive
- TD-MIX - Transaction dominant
- TM-Mix - Transaction moderate
- CB-Mix - Commercial batch mix
- LoIO-Mix - Low I/O content

One of the changes to note is the difference in the customized commercial batch mixed workload. The zPCR tool was using a combination of 50% CB-L and 50% CB-S. The newer z990 Technical Guide uses a combination of 75% CB-L and only 25% CB-S.

Now that we've described both the primitive and customized workloads, let's look at the LSPRs for each of the newer machines - the z900s, z890s and z990s.

z900 LSPRs

The z900 processors became available in May 1999 and the original benchmarks for them were run using OS/390 R10 on a z900 2064-1C1 processor. As mentioned earlier, the MIX workload was a combination of some 31-bit workloads and some 64-bit workloads, and was even more unreliable than in the past in forecasting expected performance. More consistent z900 LSPR ratings came from running the benchmarks created for the z990 processor on the z900s. These are the CB-L, CB-S, WASDB, OLTP-W and OLTP-T workloads that we described earlier.

Analyzing the Latest z900 LSPRs

Converting the z900 ITRRs to MIPS, we see the five workloads and the MIX workload as shown in Figure 8. *(Editor's Note - We're using MIPS as the basis of our charts, but you would produce exactly the same charts if you used ITRRs. For most of our examples, we use just the CB-L and CB-S workloads because they represent the two extremes of all the workloads and are easier to view. The other workloads always fall between these two workloads, both for the z900s and z990s.)* As the number of CPs increases, the variance of the MIPS by workload increases. CB-L appears to do much better at the higher end (more CPs), and CB-S does much worse at the higher end. In fact, we can see that CB-L workloads always do much better than other workloads.

What does this mean to you? If you make the mistake of using the average (MIX) MIPS, either from us or from someone else, you could be very far off in your expected performance. The MIX workload provides lower MIPS than the CB-L, WASDB and OLTP-W workloads, but provides higher

Figure 8 - z900 Total MIPS Using z/OS LSPRs

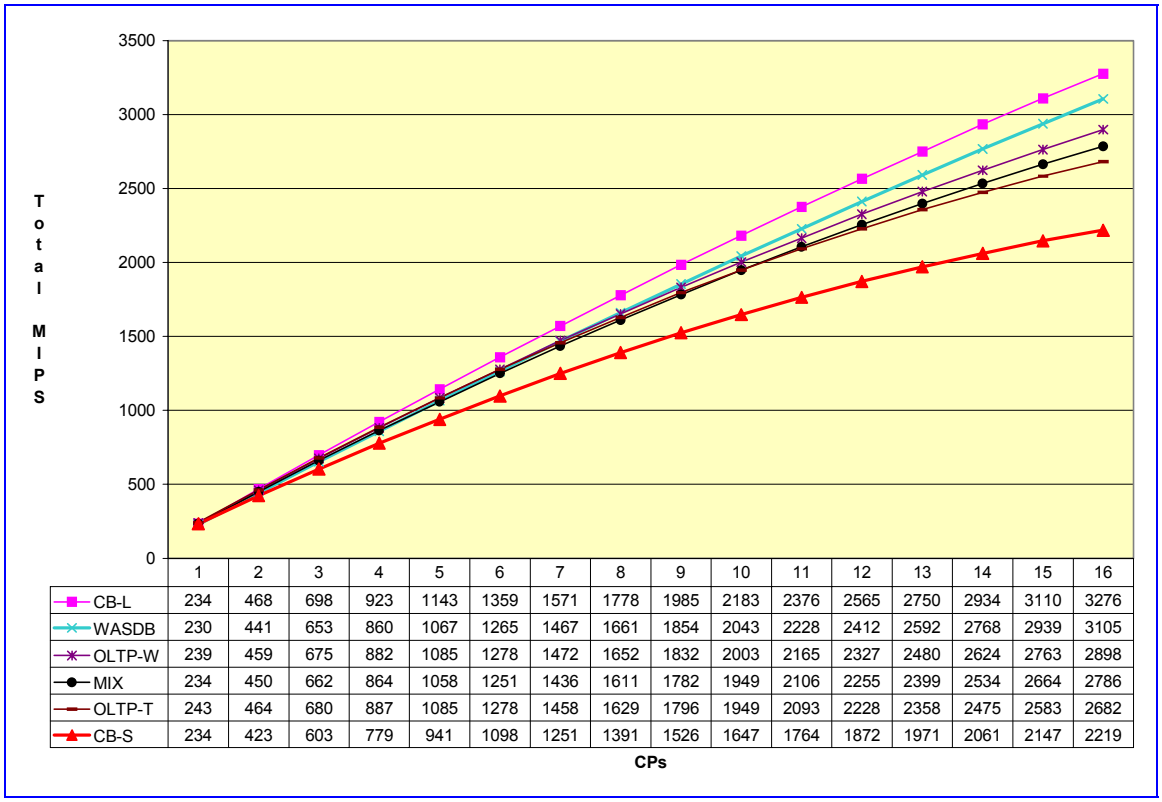

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CB-L	234	468	698	923	1143	1359	1571	1778	1985	2183	2376	2565	2750	2934	3110	3276
WASDB	230	441	653	860	1067	1265	1467	1661	1854	2043	2228	2412	2592	2768	2939	3105
OLTP-W	239	459	675	882	1085	1278	1472	1652	1832	2003	2165	2327	2480	2624	2763	2898
MIX	234	450	662	864	1058	1251	1436	1611	1782	1949	2106	2255	2399	2534	2664	2786
OLTP-T	243	464	680	887	1085	1278	1458	1629	1796	1949	2093	2228	2358	2475	2583	2682
CB-S	234	423	603	779	941	1098	1251	1391	1526	1647	1764	1872	1971	2061	2147	2219

Figure 9 - z900 MIX MIPS per CP

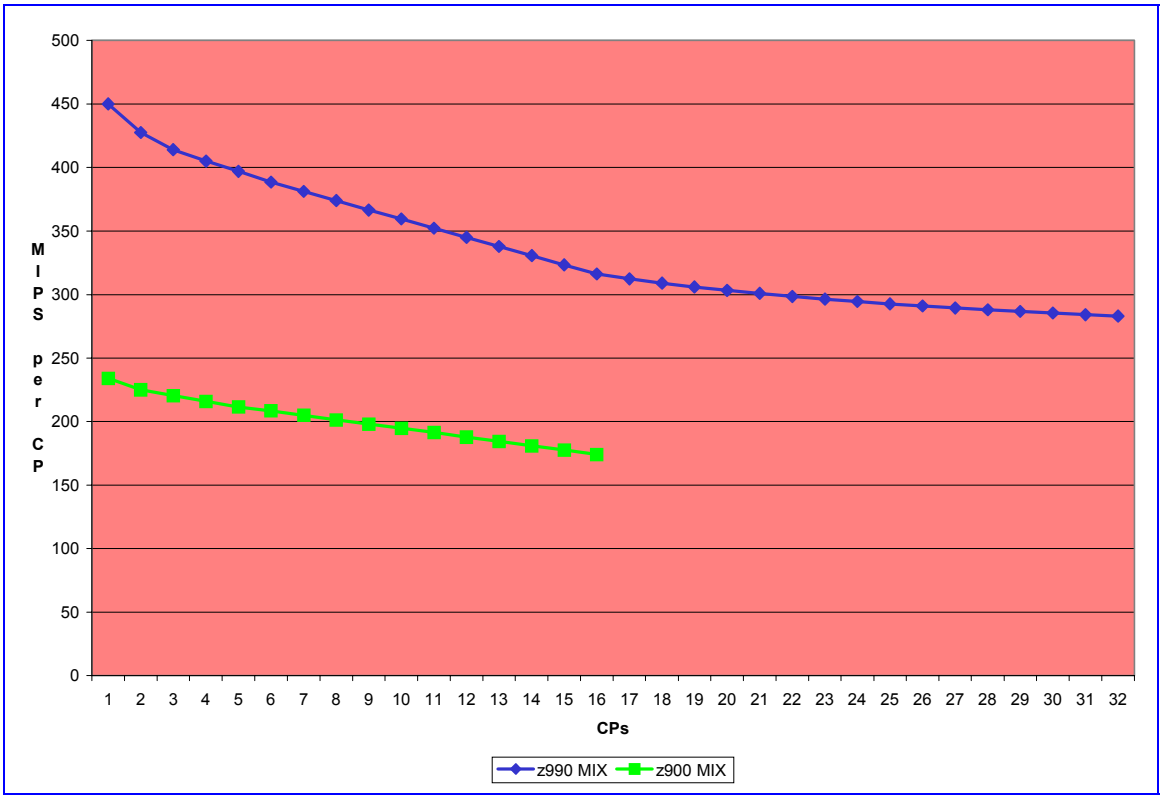

MIPS than the CB-S and OLTP-T workloads. Using MIX, your expectations could be wrong, and in fact, will almost always be wrong. The only way to properly predict the performance is by using the correct set of workload estimates. We'll cover this subject in more detail later.

Figure 8 looked at the total MIPS for each workload, but Figure 9 provides a view of the LSPR using MIPS per CP. Although the total MIPS predict the total capacity of a machine, the MIPS per CP predict the speed of each processor. A specific job usually runs on only one processor at a time, and the speed gives an indication of the change in the amount of CPU time used by the job. In Figure 9 we've graphed both the z900 processors and the z990 processors. Notice how the MIPS per CP decreases as more processors are added to the configuration. This drop is the multiprocessing effect (or 'MP' effect) due to the overhead associated with processor communication (such as SIGPs). Traditionally, this drop has been about 4-5% for the first few processors, changing to 2-3% as you add a greater number of processors.

What does this mean to you? Because of this inherent drop in CPU speed when more processors are enabled, it's important to assign the minimum number of CPs to any image. Reducing the number of processors will allow each to run at a higher speed. We'll discuss the z990 processors from this graph a bit later.

Some workloads do better with more processors, while others don't. This is easily seen in Figure 10. For now, let's concentrate on the two bottom lines on the graph. These represent the MIPS per CP for the z900 CB-S workload (bottom line) and the z900 CB-L workload (the line above it). You see that the MP effect is less noticeable for the CB-L work. In other words, performance doesn't degrade as much for CB-L work when more processors are added.

Figure 10 - CB-L and CB-S MIPS per CP

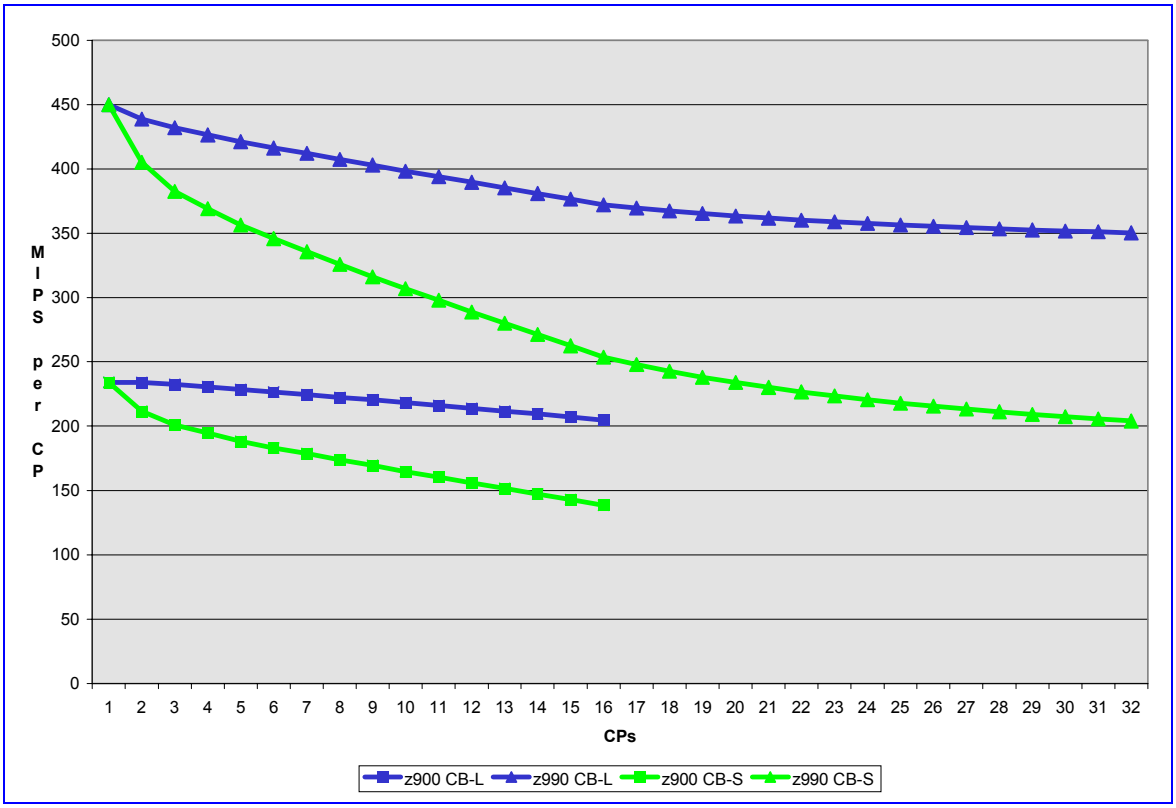

Figure 11 - z900 - All Workloads

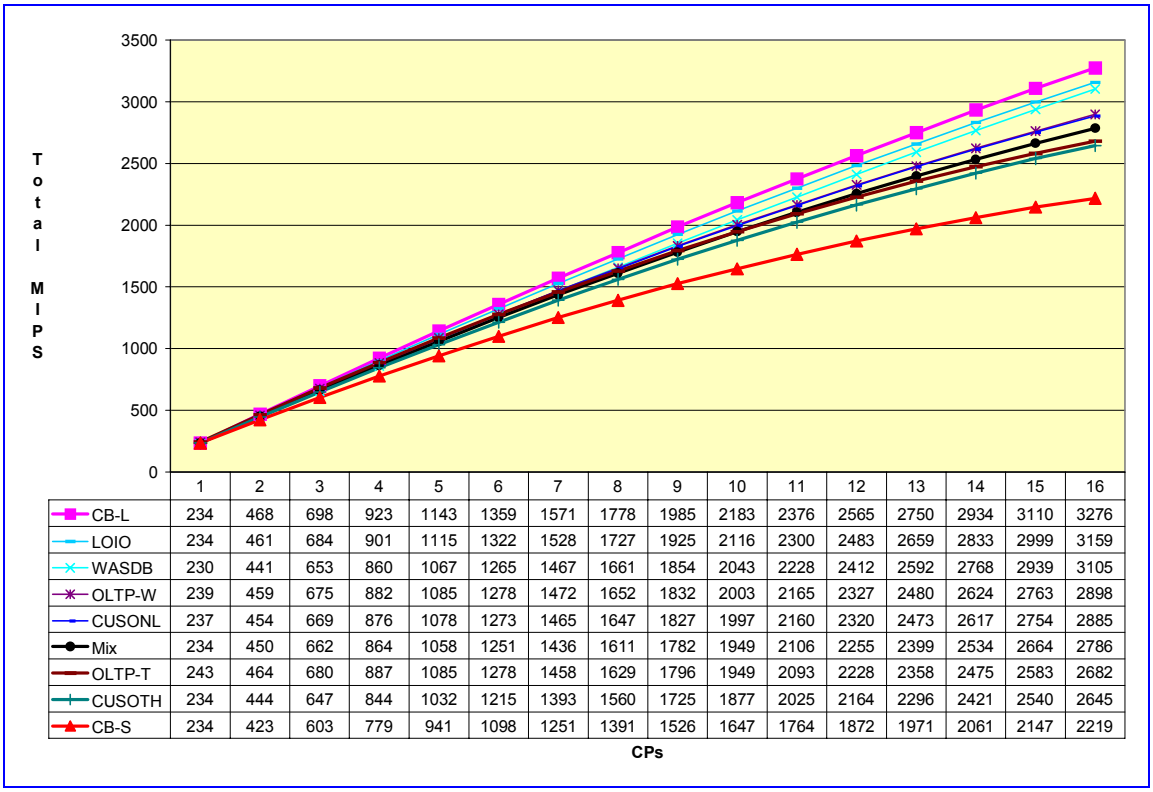

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CB-L	234	468	698	923	1143	1359	1571	1778	1985	2183	2376	2565	2750	2934	3110	3276
LOIO	234	461	684	901	1115	1322	1528	1727	1925	2116	2300	2483	2659	2833	2999	3159
WASDB	230	441	653	860	1067	1265	1467	1661	1854	2043	2228	2412	2592	2768	2939	3105
OLTP-W	239	459	675	882	1085	1278	1472	1652	1832	2003	2165	2327	2480	2624	2763	2898
CUSONL	237	454	669	876	1078	1273	1465	1647	1827	1997	2160	2320	2473	2617	2754	2885
Mix	234	450	662	864	1058	1251	1436	1611	1782	1949	2106	2255	2399	2534	2664	2786
OLTP-T	243	464	680	887	1085	1278	1458	1629	1796	1949	2093	2228	2358	2475	2583	2682
CUSOTH	234	444	647	844	1032	1215	1393	1560	1725	1877	2025	2164	2296	2421	2540	2645
CB-S	234	423	603	779	941	1098	1251	1391	1526	1647	1764	1872	1971	2061	2147	2219

Figure 12 - z990 Total MIPS (32)

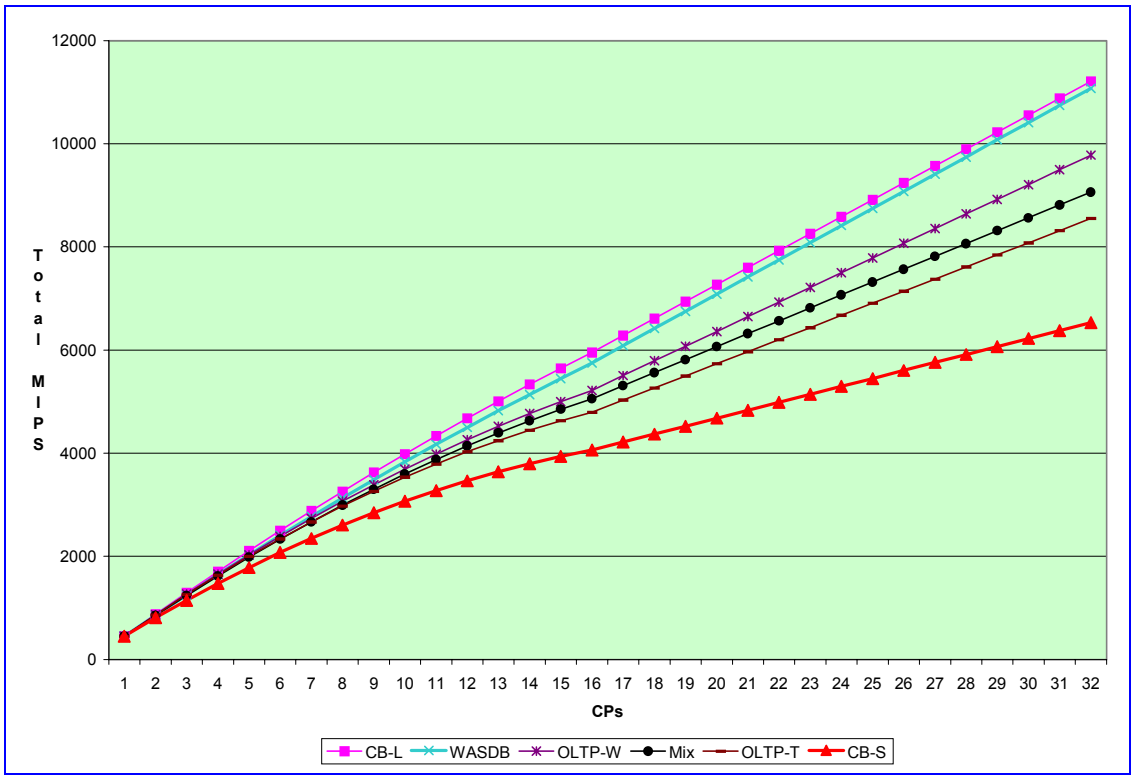

From what we've seen thus far, the CB-L workloads show improved relative performance over the CB-S workloads on the z900. The CB-L workloads consistently provide more MIPS than CB-S workloads, especially as the number of processors increase.

We used the MIX workload for our initial analysis in Figure 2. But now let's take a look at both the primitive workloads and customized workload mixes together. In Figure 11 we can make the following observations for the z900:

1. The CB-L, WASDB and LOIO workloads are somewhat similar (top three lines of the plot) and perform the best of all the workloads plotted. CB-L is the best (i.e. produces more MIPS), followed by LOIO and then WASDB.

2. The OLTP-W and CUSONL workloads are almost identical (next two lines from the top, but superimposed on one another).

3. The MIX workload is the next line down and is almost identical to the OLTP-T workload until they start to separate at about 11 CPs.

4. The OLTP-T and CUSOTH workloads are also very similar (second and third line from the bottom).

5. The CB-S workload is the worst of all those plotted (bottom line).

6. On a 16-way machine, the difference between the highest and lowest total MIPS is over 1000 MIPS (from 3276 MIPS for CB-L to 2219 MIPS for CB-S). This is over a 47% difference - quite significant!

Figure 13 - z990 Total MIPS (16)

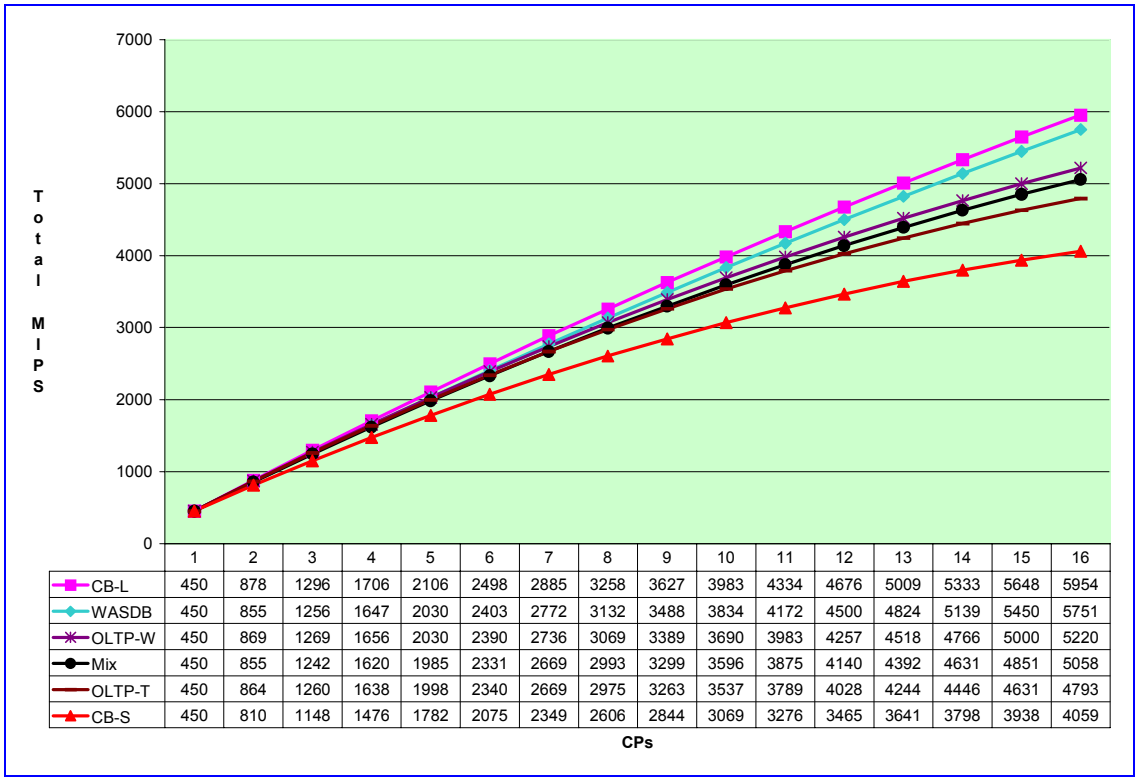

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CB-L	450	878	1296	1706	2106	2498	2885	3258	3627	3983	4334	4676	5009	5333	5648	5954
WASDB	450	855	1256	1647	2030	2403	2772	3132	3488	3834	4172	4500	4824	5139	5450	5751
OLTP-W	450	869	1269	1656	2030	2390	2736	3069	3389	3690	3983	4257	4518	4766	5000	5220
Mix	450	855	1242	1620	1985	2331	2669	2993	3299	3596	3875	4140	4392	4631	4851	5058
OLTP-T	450	864	1260	1638	1998	2340	2669	2975	3263	3537	3789	4028	4244	4446	4631	4793
CB-S	450	810	1148	1476	1782	2075	2349	2606	2844	3069	3276	3465	3641	3798	3938	4059

7. On a 16-way machine the difference between the highest and lowest MIPS per CP (Figure 10) is 66 MIPS (205 for CB-L versus 139 for CB-S).

8. In the previous section, we discussed the difference between the z900 LSPRs derived from OS/390 and z/OS. The most important point there was that the CBW2 work, which had not been part of the MIX for OS/390, is now included in CB-L, which is part of the MIX for the z/OS LSPRs.

z990 LSPRs

Now let's take a look at the LSPRs for the z990s. Figure 12 shows the MIPS by workload for the z990. Notice how there seems to be a slight wiggle around the 16-CP area? During IBM's initial testing, none of the existing operating systems supported more than 16 logical processors. As noted in the LSPR manual and LSPR Web site, IBM chose to overcome this limitation by using two LPARs for the benchmarks, running each LPAR as a 16-way. We can now understand why this variation occurs by first looking at Figure 13, which shows just the first 16 CPs. Notice that as the number of CPs increases, the range of MIPS between CB-L and CB-S increases even more.

Because the testing environment used the same number of logical CPs as physical CPs, the performance was almost equivalent to running on dedicated CPs. IBM then used a straight extrapolation between the 16-way and 32-way. This extrapolation is fairly easy to see in Figure 12, and appears to be the cause for the wiggle between the 16-way and 17-way that we see in the same figure.

Figure 14 - CB-L and CB-S Total Mips

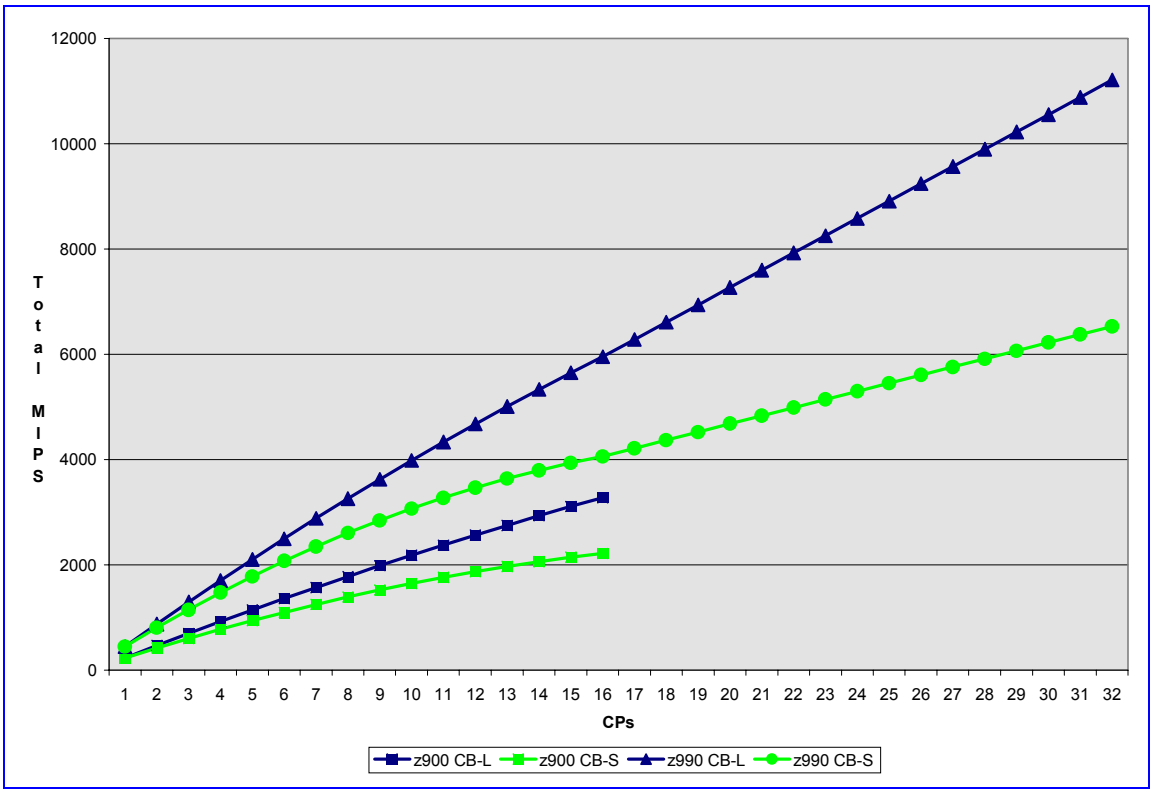

Understanding that, it doesn't explain what we see in Figure 14. The darker lines are CB-L MIPS and the lighter lines are CB-S MIPS. The top two lines are for the z990 and the bottom two lines are for the z900. We can see from this chart that CB-S work starts to drop off more rapidly than CB-L work from the 8-way to the 16-way, while CB-L work doesn't have nearly as much degradation (if any) as the MP increases. But because of what we assume is the straight line extrapolation, this degradation does not continue past 16 CPs. This makes us quite uncomfortable about the projected performance and capacity estimates for the models above the 16-way models.

Going back to Figure 9 and looking at the MIX MIPS per CP for the z990 (top line), we can see this same phenomenon. The machines above the 16-way don't have the same degradation as the first 16 CPs. This doesn't seem reasonable, and gives us even more reason to doubt the capacity claims of the larger MPs (i.e. the larger n-ways). The same observations can be made regarding Figure 10, where the CB-L MIPS per CP have almost no degradation. We just don't think this seems reasonable. We could be wrong, of course, but we suggest being **VERY** careful when planning for anything larger than a 16-way z990.

We were using the MIX workload for our first analysis in Figure 2. But let's take a look at all of the customized and primitive workloads together, as shown in Figure 15. We can make the following observations for the z990:

1. The CB-L, LOIO and WASDB workloads are similar (top three lines of the plot).

Figure 15 - z990 - All Workloads

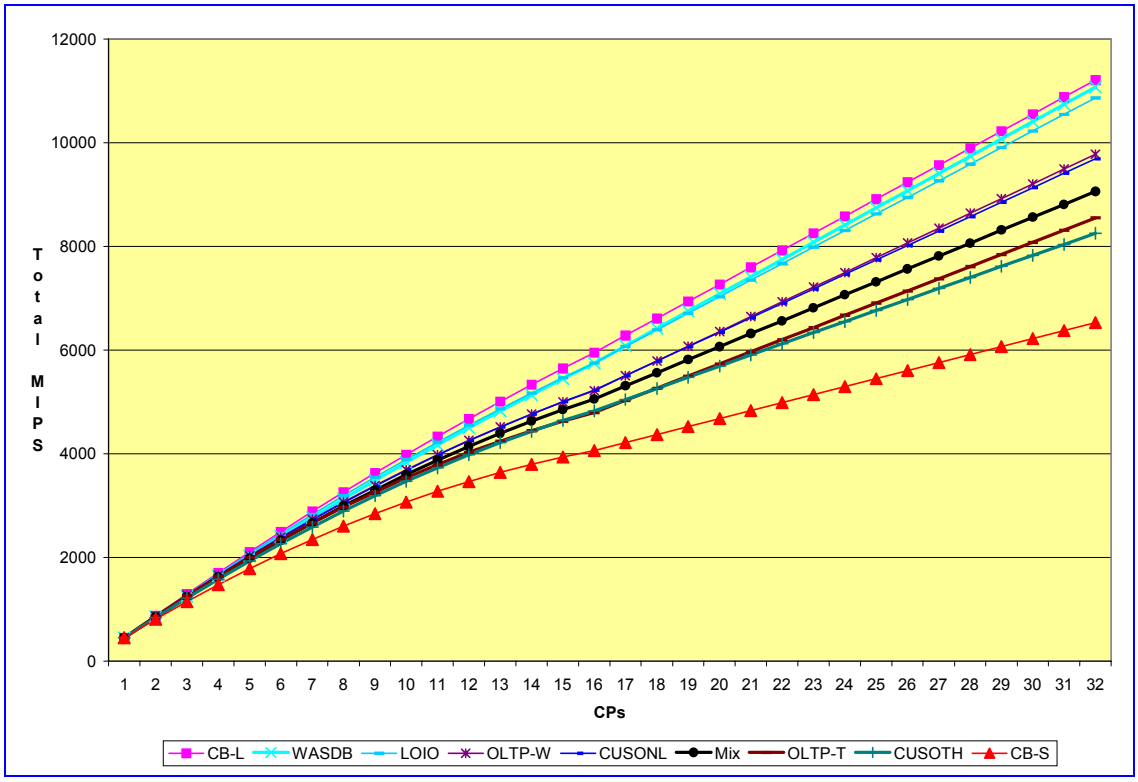

2. The OLTP-W and CUSONL workloads are also similar (next two lines from the top and overlapped).

3. The MIX is the next line down and does not seem to be representative of any of the other workloads, although it certainly represents a midpoint (this is an important concept).

4. The OLTP-T and CUSOTH workloads are quite similar (second and third lines from the bottom).

5. The CB-S workload is the worst and stands apart from all others (bottom line).

6. The main difference that we notice between the z900 and z990 workloads is that several of the z990 workloads tend to cluster together (i.e. look very similar). This means there is little difference between CB-L, LOIO and WASDB. But it's all relative, and we need to be careful of the scale of the plot. There is still a significant difference of 345 MIPS between CB-L and WASDB for the 32-way. There is also little difference between OLTP-W and CUSONL (87 MIPS for the 32-way). The same is true for OLSTP-T and CUSOTH (although why this is true is probably no more than coincidence).

Figure 16 - z890-x70

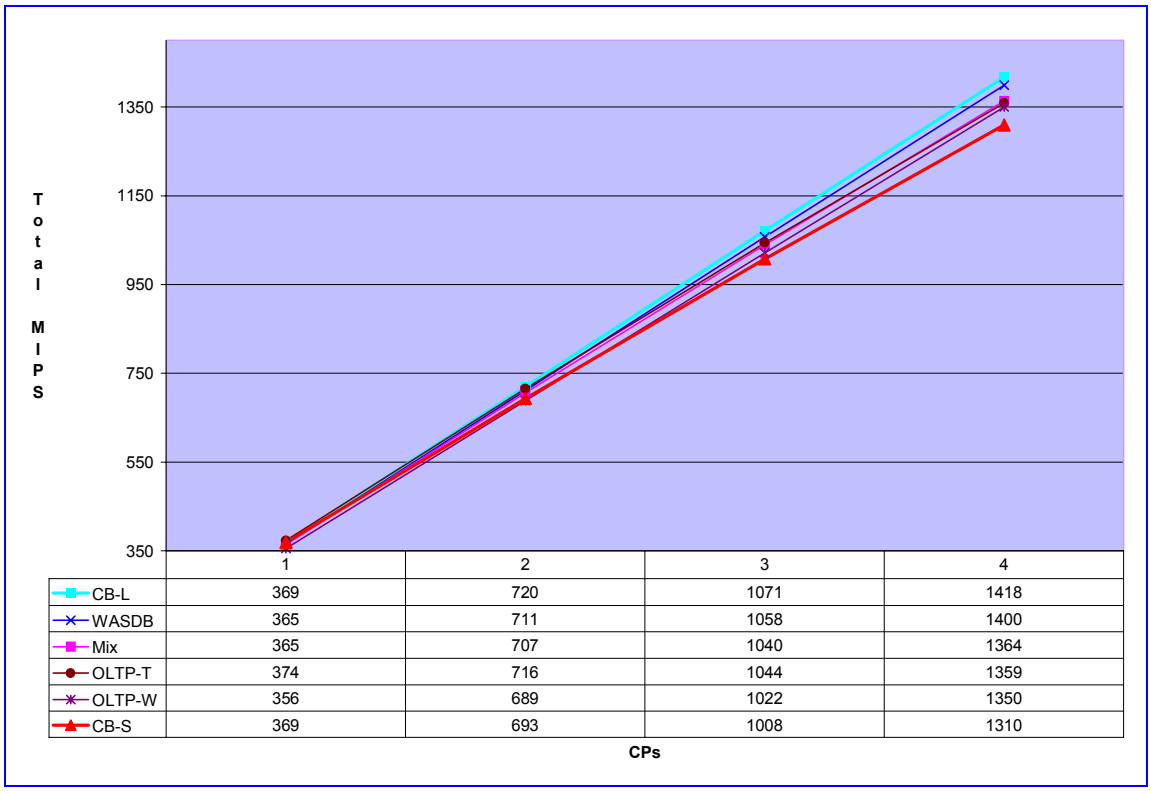

	1	2	3	4
CB-L	369	720	1071	1418
WASDB	365	711	1058	1400
Mix	365	707	1040	1364
OLTP-T	374	716	1044	1359
OLTP-W	356	689	1022	1350
CB-S	369	693	1008	1310

z890 LSPRs

The z890s, announced in April of 2004, support a maximum of four CPs. Therefore, you won't find the major differences in the various workloads that are seen on the faster processors. At the high end, however, the difference between CB-S and CB-L could still be more than 100 MIPS. Figure 16 shows the largest of the z890s, the x70 models.

There appears to be far less degradation due to MP overhead, however (at least according to the ITRRs). This can be seen more easily in Figure 17 in the MIPS per CP plot. We don't fully understand this behavior because the z890 engines are simply z990 engines modified to run slower.

Because we haven't talked to any z890 customers, we don't know whether the low I/O condition applies to the z890s. There is more about the z890s in our *What's New?* section starting on page 45.

Figure 17 - z890 & z990 MIPS Per CP

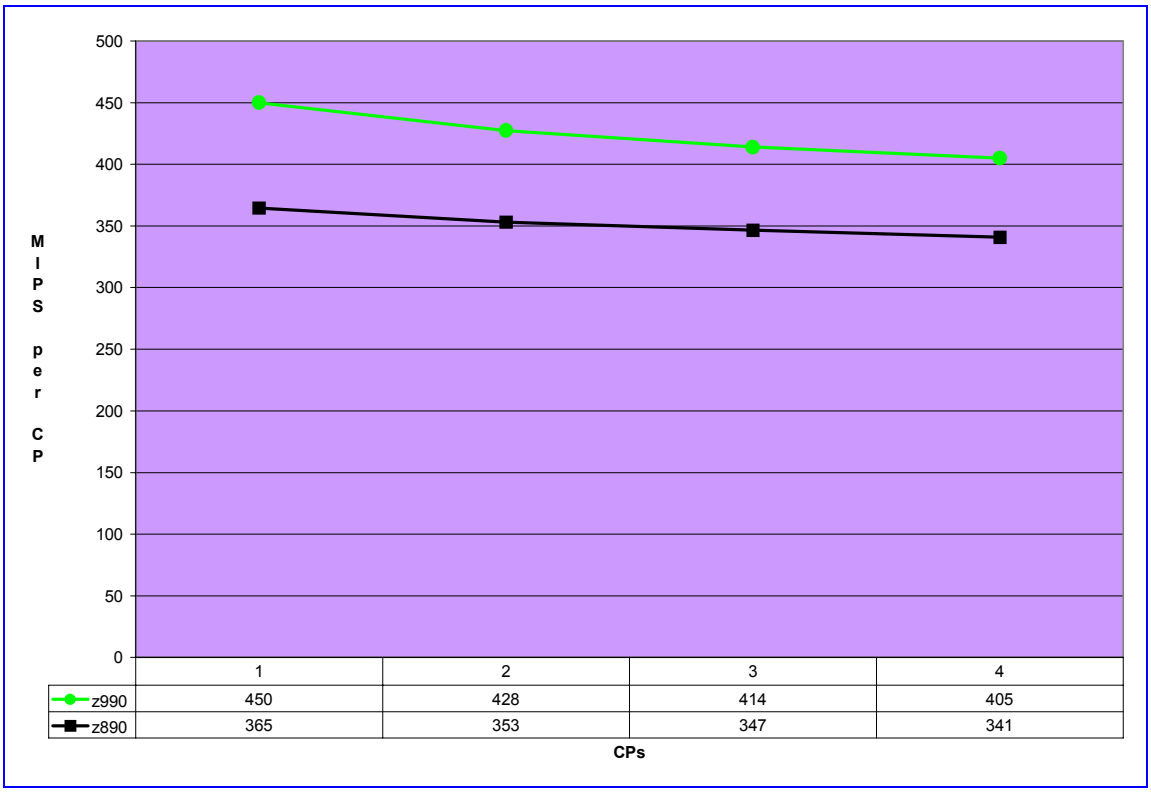

Bibliography

This is the bibliography for the *LSPR Update* article that began on page 3. Some of these items may also be referenced in the next article on *z990 Expectations*.

We would like to make an observation about several of the IBM articles and our CPU Chart. Throughout the IBM documentation, there are many references to the invalid use of MIPS from consultants 'or vendors' CPU Charts. They say repeatedly that MIPS are invalid and that the only thing that is reliable is the LSPR ITRRs. We don't believe that their warnings apply to our CPU Charts. Their main concern is that most CPU Charts, with ours being the only exception we know about, only publish a single MIPS value for a machine. This is usually based on the MIX ITRR. As we've mentioned in this article, we don't agree with that method either, and it's the reason our CPU Charts contain MIPS for all of the workloads. We even include the customized workloads for which not even IBM publishes ITRRs.

The second problem with most other MIPS charts is that they have only a single chart. But as both we and IBM recommend, the relative capacity of a machine can only be determined by looking at ratings that come from the same LSPR set of benchmarks. We do that by having a separate z/OS CPU Chart and by identifying the LSPR benchmark set in a column within the OS/390 CPU Chart. Because we base our MIPS on the LSPR ITRRs, we feel very confident that we reflect the most accurate view of the processor capacities, while still using the MIPS that are more familiar for most installations.

REF001 - IBM LSPR Web site - **www.ibm.com/servers/eserver/zseries/lspr**
REF002 - Cheryl Watson's zSeries CPU Chart, April 2004; Cheryl Watson's OS/390 CPU Chart, April 2004
REF003 - Cheryl Watson's TUNING Letters: 2003 No. 3, pages 23-25; 2001 No. 2, pages 35-36; 2000 No. 5, pages 49-50; 1997 No. 4, pages 10-26
REF004 - **IBM Large Systems Performance Reference, SC28-1187-09**, 13May2003; update on 15Jan2004; update on 2Apr2004
REF005 - SHARE Long Beach session 2513, **JoAnne Brown**, *Processor Sizing*, **www.share.org**
REF006 - SHARE Long Beach session 2514, **Walt Caprice**, *WSC Short Stories and Tall Tales - z990 Performance Considerations*, **www.share.org**
REF007 - WSC Presentation PRS135, **Kathy Walsh,** *Information on Capacity Planning for S/390 G5 and G6 Processors using LSPR data*, 28Mar2002, **www.ibm.com/support/techdocs**
REF008 - **IBM eServer zSeries 990 Technical Guide SG24-6947-01** draft, 7Apr2004 (but no mention of low I/O in SG24-6947-00, Jul2003), **www.redbooks.ibm.com**
REF009 - CMG 2000 Proceedings, **Dr. Sudhir R. Nath** of Wells Fargo Services Company, *Harmonic Mean Analyses Of CPU Speeds*, **www.cmg.org** (Members Only section of the Web site)
REF010 - **Gregory V. Caliri** of BMC Software, *Large Scale Processor Reference (LSPR) And Its Implications,* **www.bmc.com/offers/performance/whitepapers/docs/2003/Large_Scale_Processor_Implications.pdf**
REF011 - Oct/Nov2003 - z/Journal, **Marty Deitch**, *IBM's LSPR Benchmark Results: The Truth About Mainframe MIPS,* **www.zjournal.com/PDF/deitchoct.pdf**
REF012 - BoxScore, **www.watsonwalker.com/boxscore.html**
REF013 - Cheryl Watson's TUNING Letters: 1998 No. 6, page 28; 1998 No. 4, page 37
REF014 - Cheryl Watson's TUNING Letter 2002 No. 4, page 20
REF015 - Cheryl Watson's TUNING Letters: 1998 No. 6; 2002 No. 3, pages 28-33; 2003 No. 5, pages 37-41; other references:
February 1991, FOCUS: PR/SM
March 1991, Reducing I/O Elongation
May 1991, Amdahl's MDF
December 1991, Dedicated and Shared LPARs
July 1992, Defining a Sysprog LPAR
Sept/Oct 1992, PR/SM Overhead
May/June 1993, PR/SM Changes (SP 4.3)
July/Aug 1993, FOCUS: LPAR Update (EMIF, MDF, SP 4.3)
1997, No. 4, LSPRs and You
1998, No. 1 & 2, Configuration Changes
REF016 - **WSC White Paper WP100258** - *Performance Considerations When Moving to Fewer Faster CPUs,* **www.ibm.com/support/techdocs** ■

Focus: z990 Expectations

The z990 (T-REX) processor has been in the field since September 2003 and we've been getting quite a bit of feedback on its performance. The results are varied, but overall we feel that customers are not seeing the performance they expected. In this article, we'll first cover the results from our z990 customers (using both our BoxScore product and other tools). Then we'll explain the causes for these missed expectations, both real and perceived. And finally, we'll present our recommendations. When we mention LSPRs, it will be helpful if you've already read our *LSPR Update* beginning on page 3. That article is an important prerequisite to understanding what is happening on the z990s.

- BoxScore Results
- Reasons for Underperformance
- z900 to z990 Migration
- More BoxScore Results
- Our Conclusions and Recommendations

BoxScore Results

In early November of 2003, we started getting feedback from our BoxScore customers who had installed z990s. As you probably know, our software product, BoxScore, is designed to determine whether you are getting the capacity you expect after a hardware or software upgrade. [Please see REF012 for additional information on BoxScore.]

So far, we've only seen moves from z900s to z990s. **With one exception, all BoxScore sites are seeing between 8% and 12% fewer MIPS than expected for batch processing using the traditional primitive workloads. Four CICS installations were experiencing between 10% and 15% fewer MIPS than expected.** That one exception experienced performance that was 4% better than expected for their batch work on their new z990. Their CICS work, however, failed to meet expectations by about 10%.

At least two of the installations went back to IBM with the BoxScore results and were able to get hundreds of thousands of dollars back in credits. (Note: the customers told us these were credits due to the underperformance of the z990. IBM says that credits haven't been given due to underperformance but because of improper sizing. All of the installations reporting underperformance asked that their names not be used.)

Our long-time readers will remember that we've previously seen underperformance in BoxScore results during moves to new processors. You might recall the problem with CMOS machines and COBOL programs that used subscripting instead of indexing. [REF013] Or you might remember the problems with the change to high-speed cache processing on the z900s that caused increases in run times for SAS and other products. [REF014] In both previous cases, customers were able to identify from the BoxScore reports the specific programs that were not meeting expectations and determine the commonalities. But thus far, we have not identified any specific application or program type that seems to be causing the variance from the expected z990 results.

Figure 18 - BoxScore Percent of Change Plot

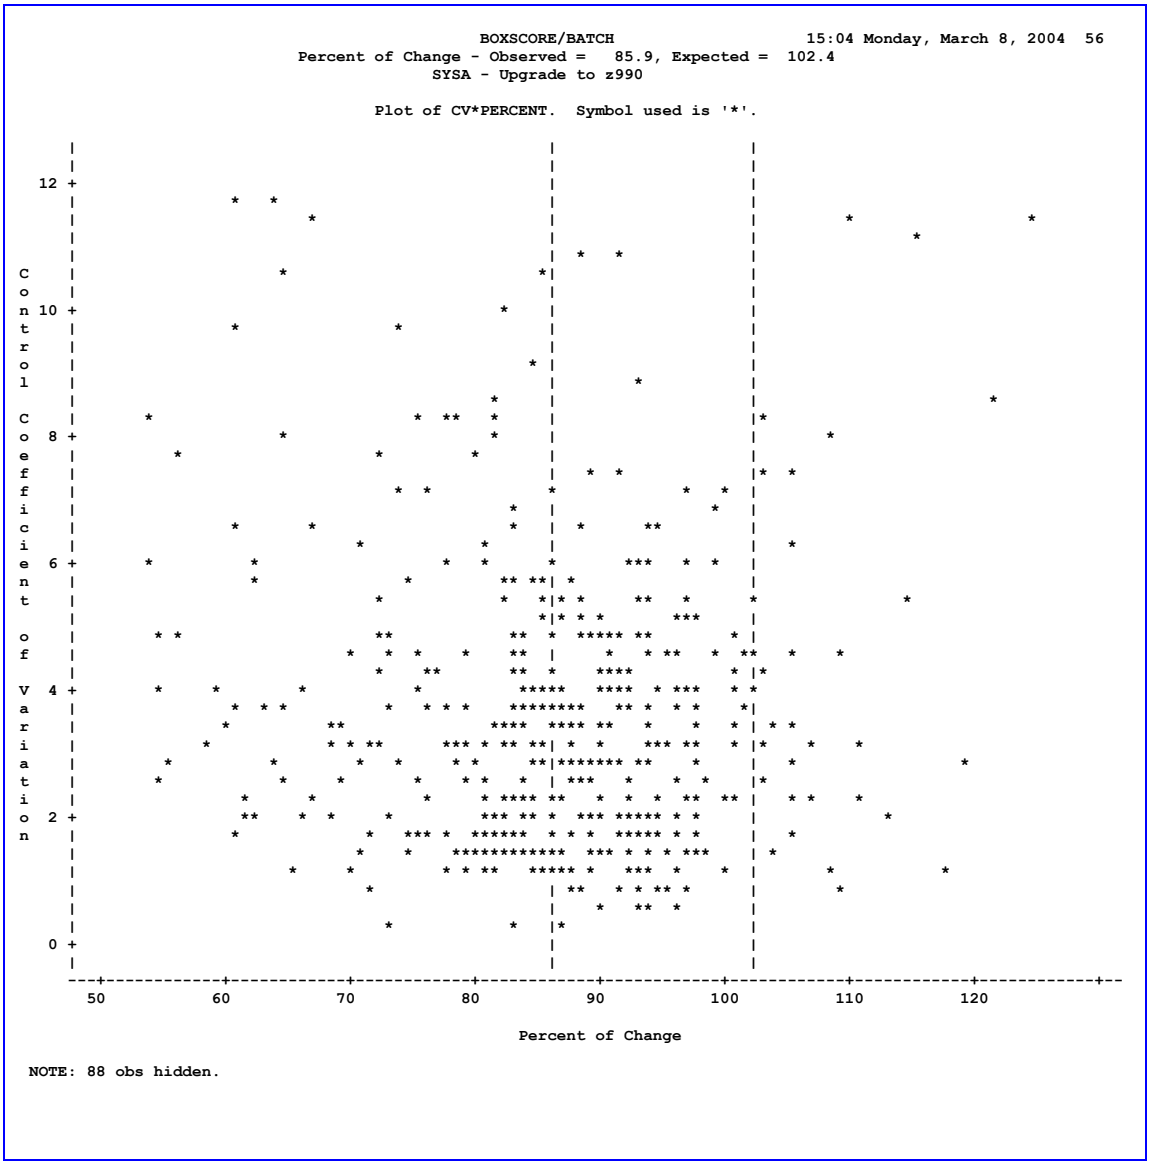

Refer to Figure 18, which shows a BoxScore Percent of Change Plot. This is an example of a move from a z900 2064-113 to a z990 2084-310, which should have provided double the MIPS per CPU. The plot for this change shows an increase of 85.9% in effective MIPS per CPU (as seen by the middle vertical line of '|'), but the expected change was 102.4% (the rightmost vertical line of '|') so the customer was very disappointed. The points on the far left are those jobs performing the worst, while the points on the right are those doing the best. As you can see, however, the work is quite clustered and evenly distributed around this 85.9% increase. This BoxScore run was made using the CB-S LSPRs, because the installation had successfully used CB-S LSPRs in the past for moving batch between two machines.

Another BoxScore report is shown in Figure 19. This reports that from an LPAR point of view, the batch work experienced an 11.1% reduction from the expected result, in both capacity and speed. Looking at it from the CEC point of view, the batch work experienced an 8.1% reduction from the expected result.

Figure 19 - BoxScore Summary Report

```
                                        BOXSCORE/BATCH        15:04 Monday, March 8, 2004  49
V1R6 (c) Watson & Walker, Inc.          Summary - CPU per I/O

                        +-------------------------------------------------------+
                        +  The work analyzed during this period experienced a   +
                        +       46.2% decrease in CPU time per I/O              +
                        +    between the two environments analyzed.             +
                        +-------------------------------------------------------+

+-BoxScore/BATCH: -11.1%  +----Capacity (MIPS/LPAR)---+ +--% Delta--+ +--Speed (MIPS/Logical CPU)-+
+                         +   Expected   +  Observed  + +           + +   Expected   +  Observed  +
+ From an LPAR view,      +--------------+------------+ +-----------+ +--------------+------------+
+   STUDY                 +              +            + +           + +              +            +
+   had 11.1% less        + Max 2884.5   +   2520.3   + +  -12.6%   + + Max  412.1   +    360.0   +
+   speed and capacity    +              +            + +           + +              +            +
+   than expected from    + Avg 2349.0   +   2087.1   + +  -11.1%   + + Avg  335.6   +    298.2   +
+   published performance +              +            + +           + +              +            +
+   estimates.            + Min 2349.0   +   1724.3   + +  -26.6%   + + Min  335.6   +    246.3   +
+                         +              +            + +           + +              +            +
+-------------------------+--------------+------------+ +-----------+ +--------------+------------+

+-BoxScore/BATCH: -8.1%   +----Capacity (MIPS/CEC)----+ +--% Delta--+ +-Speed (MIPS/Physical CPU)-+
+                         +   Expected   +  Observed  + +           + +   Expected   +  Observed  +
+ From a CEC view,        +--------------+------------+ +-----------+ +--------------+------------+
+   STUDY                 +              +            + +           + +              +            +
+   had 8.1% less         + Max 3982.5   +   3404.0   + +  -14.5%   + + Max  398.3   +    340.4   +
+   speed and capacity    +              +            + +           + +              +            +
+   than expected from    + Avg 3069.0   +   2818.9   + +   -8.1%   + + Avg  306.9   +    281.9   +
+   published performance +              +            + +           + +              +            +
+   estimates.            + Min 3069.0   +   2328.9   + +  -24.1%   + + Min  306.9   +    232.9   +
+                         +              +            + +           + +              +            +
+-------------------------+--------------+------------+ +-----------+ +--------------+------------+
```

Figure 20 shows an LPAR view of the migration. (By the way, we use "control" to identify the old machine or situation and "study" for the new.) We've underlined some of the items that we want to mention. The first underlined item shows that this move was from a z900 2064-113 processor to a z990 2084-310 processor. The next two underlined items show that the LPAR had 11 and 7 logical CPs assigned out of the 13 and 10 physical CPs on those boxes. The ratio of logical to physical CPs went from 1.6 to 1.9, which is a little greater, but not enough to cause significant LPAR overhead. The next three underlined items show that we expected an ITRR of 1.56, but saw an ITRR of 1.43, representing an underperformance of 8.1%. Please note that all of these results were calculated based on the CB-S workload. These results are typical of most of our customers; the results for CICS were usually much worse.

Starting in November, because all of our BoxScore customers were seeing these types of results, we started investigating the reasons for the apparent underperformance. The articles in this newsletter are the result of the last six months of investigation, analysis of z990 migrations, conversations with our customers and discussions with IBM. We think we understand what is happening, and we will try to explain it to you in the remainder of this article.

Figure 20 - BoxScore Identification and CPU Comparison Reports

BOXSCORE/BATCH - Dropping Outliers 15:04 Monday, Jan 8, 2004 51
V1R6 (c) Watson & Walker, Inc. CPU Comparison

Item	Control	Study	Delta	% Delta	ITRR	Comments
System Identification:						
System	SYSA	SYSA				
Model-Version	2064-113	2084-310	*****			
Common name for processor	2064-113	2084-310	*****			
Manufacturer	IBM	IBM				
MVS release	z/OS 01.04	z/OS 01.04				
Architecture mode	64-bit	64-bit				
Central storage	12288MB	12288MB	0MB	0.0%		
Number of logical CPUs	11.0	7.0	-4.0	-36.4%		WWCB061-I # of log. CPU
Number of physical CPUs	13.0	10.0	-3.0	-23.1%		WWCB024-I # of phys. CP
LPAR status	SHR	SHR				WWCB026-I LPAR used in
LPAR weight (avg)	660.0	410.0	-250.0	-37.9%		WWCB097-I IRD decreased
Number of active LPARs	6.0	6.6	0.6	10.0%		WWCB098-I # LPARs inccr
Total number of LPs in CEC	21.0	19.3	-1.7	-8.1%		
LPAR LPs to CP ratio	1.6	1.9	0.3	19.5%		WWCB127-W LP to CP rati
Weight of other LPARs (avg)	340.0	586.0	246.0	72.4%		
Percent of CEC this LPAR	66.0%	41.2%	-24.8%	-37.6%		WWCB131-W % of CEC decr
Total CPU busy	801.9%	533.4%	-268.5%	-33.5%		WWCB068-W CPU busy is lo
Avg CPU busy	72.9%	76.2%	3.3%	4.5%		
Max CPU busy	96.2%	98.4%	2.2%	2.3%		
Min CPU busy	13.0%	13.7%	0.7%	5.4%		
. . .						
Speed of one CPU (physical):						
Expected SU/second	8724.10	17003.18	8279.08	94.9%	1.95	
Expected avg MIPS/CPU	151.6	306.9	155.3	102.4%	2.02	WWCB028-I Expected faste
Expected max MIPS/CPU	211.5	398.3	186.8	88.3%	1.88	WWCB096-I Weight increas
Expected min MIPS/CPU	151.6	306.9	155.3	102.4%	2.02	WWCB098-I # LPARs increa
Observed MIPS/CPU	151.6	281.9	130.3	85.9%	1.86	WWCB030-I CPU is faster
Change from predicted avg			-25.0	-8.1%		
Machine capacity (physical):						
Expected avg MIPS	1971.0	3069.0	1098.0	55.7%	1.56	WWCB032-I Expect more ca
Expected max MIPS	2749.5	3982.5	1233.0	44.8%	1.45	WWCB127-W LP to CP ratio
Expected min MIPS	1971.0	3069.0	1098.0	55.7%	1.56	WWCB129-W % of CEC incre
Observed MIPS	1971.0	2818.9	848.0	43.0%	1.43	WWCB034-I More capacity
Change from predicted avg			-250.1	-8.1%		

Reasons for Underperformance

While investigating the situations seen by our BoxScore installations, we believe we've discovered three main causes for the underperformance. Although we think the third reason is the most significant, we'll cover the other two first and then turn our attention to the third reason. (We also believe there is still another reason for underperformance for CICS, but we haven't been able to track it down. We'll cover that in a later issue as we acquire more information.)

- LPAR Configuration

  Because the z990 is twice as fast as the z900, most installations are making a move to faster but fewer processors. If the installation does not change the number of LPARs, having fewer CPs can lead to an unacceptable ratio of logical CPs to physical CPs. When there are too many logical CPs to physical CPs (more than 3 to 1), the amount of LPAR overhead can be quite high. One site, for example, went from a 12-way z900 to an 8-way z990 with 30 logical CPs. Although this was barely acceptable on a z900 (with a ratio of 2.5 to 1), it was simply too much contention on the z990 (with a ratio of 3.7 to 1) and caused additional overhead. In our example in Figure 20, the change was from 1.6:1 to 1.9:1, so the increase in LPAR overhead was minimal.

It should be possible to reduce the number of logical CPs to an acceptable number once you're aware of the problem. This step of planning LPARs should be done before you make a final decision on a processor.

- Moving to Faster CPs

Because of the significant difference in the speed of the CPs between the two machines (the z990 is rated at 450 MIPS per CP versus the z900 at 234 MIPS per CP), you will probably run into the fewer but faster phenomenon that we described in our TUNING Letter 1996, No. 6. The original issue is on our Web site under "Sample Issues," but we've updated the article for the current machines and included it in this issue starting on page 40.

The major problem that we see with the faster CPs is that higher importance online systems will tend to dominate the system at the expense of lower importance work. Lower importance work might then take more elapsed time, and therefore more CPU time than expected. Some people have tried using restricted resource groups for peak period times, but the problem continues. When looking at the speed and capacity of the higher importance work, they should tend to get better than expected performance. BoxScore customers, however, are seeing the opposite condition where CICS is not performing as expected.

- Low I/O Density

We feel that the primary reason for dissatisfaction is that IBM did not prepare customers for the low I/O expectations (described starting on page 12). It is obvious that IBM only identified the low I/O situation after the z990s were already in the field and complaints started showing up. The first public notice of the low I/O condition for z990s was at the SHARE conference in Long Beach in February 2003, five months after the first z990 became generally available. Unfortunately, probably less than a hundred people saw those excellent presentations from **Walt Caprice** [REF006] and **JoAnne Brown** [REF005]. In January 2004, IBM included a slight mention of this issue in their LSPR manual [REF004] and again in a Redbook draft of the z990 Technical Guide [REF008]. There is still no mention of it on the LSPR Web site. We feel that most people have never seen these references, and that IBM should do more to make them aware of it.

Because IBM says that this situation applies to 80% of the z990 customers, we think that this should be highlighted and much better understood. To understand why the low I/O density situation is causing this dissatisfaction, we'll address the LSPRs again in the next section.

z900 to z990 Migration

As you saw in our *LSPR Update* starting on page 3, the CB-L workload appears to get the best performance of all workloads, while the CB-S workload shows the worst performance. CB-L results in the highest MIPS in comparison to other workloads, and the least degradation as the number of processors increases. As we've already seen (Figure 15), CB-L is very similar to both the WASDB workload and the LOIO (low I/O) workload mix. The LOIO workload is the one that, according to IBM, applies to 80% of their customers. Everybody should be thrilled that they seem to match the performance of CB-L.

Figure 21 - z900 8-way and 12-way to z990s Expected % of Change

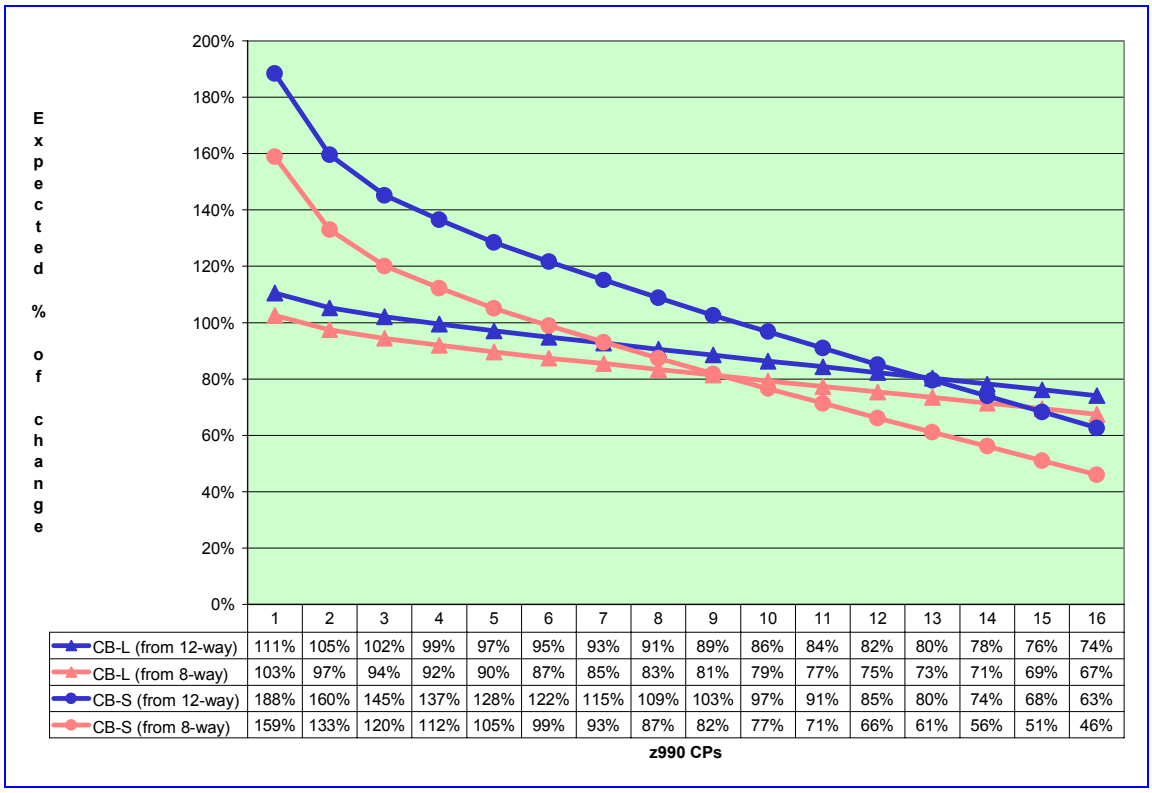

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
CB-L (from 12-way)	111%	105%	102%	99%	97%	95%	93%	91%	89%	86%	84%	82%	80%	78%	76%	74%
CB-L (from 8-way)	103%	97%	94%	92%	90%	87%	85%	83%	81%	79%	77%	75%	73%	71%	69%	67%
CB-S (from 12-way)	188%	160%	145%	137%	128%	122%	115%	109%	103%	97%	91%	85%	80%	74%	68%	63%
CB-S (from 8-way)	159%	133%	120%	112%	105%	99%	93%	87%	82%	77%	71%	66%	61%	56%	51%	46%

But there have been complaints when upgrading from a z900 to a z990 processor. Why is this? We think that the reason is evident in Figure 21. The darker colored lines show the expected percent of change in MIPS per CP when going from a z900 12-way to each of the possible z990 processor configurations from a 1-way up to a 16-way. The expected percent of improvement for CB-L work is significantly less than the expected percent of improvement for CB-S work until you get to a 13-way z990. The lighter colored lines show the expected percent for change from a z900 8-way to each of the z990 processors. In both cases, we can see from where the CB-L and CB-S lines cross that until you move to a machine with slightly more processors, the CB-L work is expected to perform the worst (improve the least) and the CB-S work is expected to perform the best. According to many of our BoxScore customers, they are certainly seeing the smaller performance boost rather than the larger.

To get more specific with this figure, let's assume that you're moving from a 12-way z900 (about 2255 MIPS using a MIX workload with 214 MIPS per CP) to an 8-way z990 (about 2293 MIPS with 407 MIPS per CP). The average total MIPS are about the same, but the workloads are very different. As we see from Figure 21, the CB-L workload for this move is expected to see a 91% improvement in MIPS per CP, while the CB-S workload is expected to see a 109% improvement. That's an 18% difference in MIPS per CP for a relatively common type of move. In this example, we're specifically talking about the following when going from a 12-way z900 to an 8-way z990 (these are MIPS per CP):

z900 (CB-L = 214) to z990 (CB-L = 407) for a gain of 193 MIPS or 91% improvement
z900 (CB-S = 156) to z990 (CB-S = 326) for a gain of 170 MIPS or 109% improvement

Because the z990s are twice as fast as the z900s, most installations will be moving from a larger number of CPs on a z900 (such as the 12-way in Figure 21) to fewer CPs on the z990. But because of IBM's claim that 80% of their customers have a low I/O density (and fall into the LOIO workload mix), then 80% of their customers are seeing the lowest possible improvement when moving from a z900 to a z990.

The net of this is that the majority of z990 customers will be getting the minimum performance improvement expected by the LSPRs, which is nowhere close to the MIX workload. This situation may not be apparent to some installations for the following reasons:

1. Most installations simply don't confirm the speed and capacity of new machines. We think that is a poor way to manage an installation, but it's true. Unless performance to the user greatly degrades or bills sharply increase, nobody will notice a 8-12% underperformance (unless they have BoxScore or a similar product). One installation that saw no z990 underperformance also admitted they were upgrading to a configuration with a lot of unused capacity, and that they had not done any measurements after the migration. This is certainly one way to manage capacity if you have the dollars to throw at it.

2. Installations that comprise the 20% minority with high I/O density can successfully use the published LSPR primitive workloads for their expectations. We have to say, however, that none of our BoxScore customers have been found to be in that category. All had low I/O densities.

3. The installation has increased the capacity by a large amount. There are two reasons for not noticing an underperformance here:

    a. There is enough capacity so that all of the work is easily handled and throughput is increased and response times are decreased. If an installation doesn't bother about billing or throughput measurement, they wouldn't notice a small increase in CPU time.

    b. They are moving to a z990 with more processors than their older z900. In this situation, the low I/O workload does provide more capacity than the other workloads.

Very few installations even know that they should be using the low I/O workload for estimating their capacity. They will probably see the underperformance based on their planning that used other LSPR workloads for sizing.

Even when they've heard about low I/O, the apparent performance of low I/O appears to produce the highest MIPS of any workload so people think it's a good thing. Very few people look at the comparisons to see that the low I/O workload mix also results in the smallest expected growth between the z900 and the z990.

Figure 22 - Low I/O Results

Item	Control	Study	Delta	% Delta	ITRR
Speed of one CPU (physical):					
Expected SU/second	8724.10	17003.18	8279.08	94.9%	1.95
Expected avg MIPS/CPU	204.9	389.4	184.5	90.0%	1.90
Expected max MIPS/CPU	211.5	398.3	186.8	88.3%	1.88
Expected min MIPS/CPU	151.6	306.9	155.3	102.4%	2.02
Observed MIPS/CPU	204.9	381.3	176.3	86.1%	1.86
Change from predicted avg			-8.2	-2.1%	
Machine capacity (physical):					
Expected avg MIPS	2664.0	3894.3	1230.3	46.2%	1.46
Expected max MIPS	2749.5	3982.5	1233.0	44.8%	1.45
Expected min MIPS	1971.0	3069.0	1098.0	55.7%	1.56
Observed MIPS	2664.0	3812.6	1148.6	43.1%	1.43
Change from predicted avg			-81.6	-2.1%	

More BoxScore Results

After IBM discovered that the low I/O situation was appearing on the z990s, we enhanced BoxScore to add three new customized workload mixes: CUSONL, CUSOTH and LOIO. We sent this new code to several z990 users who reran the BoxScore reports. In most cases, when using the low I/O workload (60% CB-L, 20% WASDB and 20% OLTP-W), the batch work seemed to match expectations fairly closely, although the CICS work still did not meet expectations. So here are more results.

Using Low I/O for Expectations

Figure 22 is a rerun of the same system shown in Figure 20. Notice that the revised expectation changed from 1.56 to 1.46 because the low I/O workload expectation was used. And now the results show a 2.1% underperformance rather than an 8.1% underperformance using the CB-S workloads. We are seeing similar results for most of the BoxScore analyses of batch workloads.

IBM believes that if you come within 5% of expectations, they (IBM) have met their commitment, and the majority of our customers are coming within 5% of the low I/O expectations *for batch work*.

CICS Expectations

But we still cannot account for the underperformance that we see with CICS work. BoxScore uses the SMF 110 records or similar data that collect CPU usage by transaction. Figure 23 shows one of the worst results for CICS (most sites see between 8% and 15% underperformance). This report was based on about 8.5% of eight million transactions using the OLTP-W LSPR estimates. They saw a 47.9% improvement rather than the expected 102.6% improvement.

IBM disagrees with our CICS results. They don't use the transaction level data themselves, but instead use SMF type 30 records for the CICS regions and type 72 records for the service classes. We found that data to be too variable and too suspect for analysis. In BoxScore, we know exactly which transactions are being analyzed. Using type 30 and type 72 records, you have a combination of transaction types and I/O activity. We strongly believe that the technique used by BoxScore provides the most accurate analysis of CICS workloads.

Figure 23 - BoxScore CICS Results Using OLTP-W

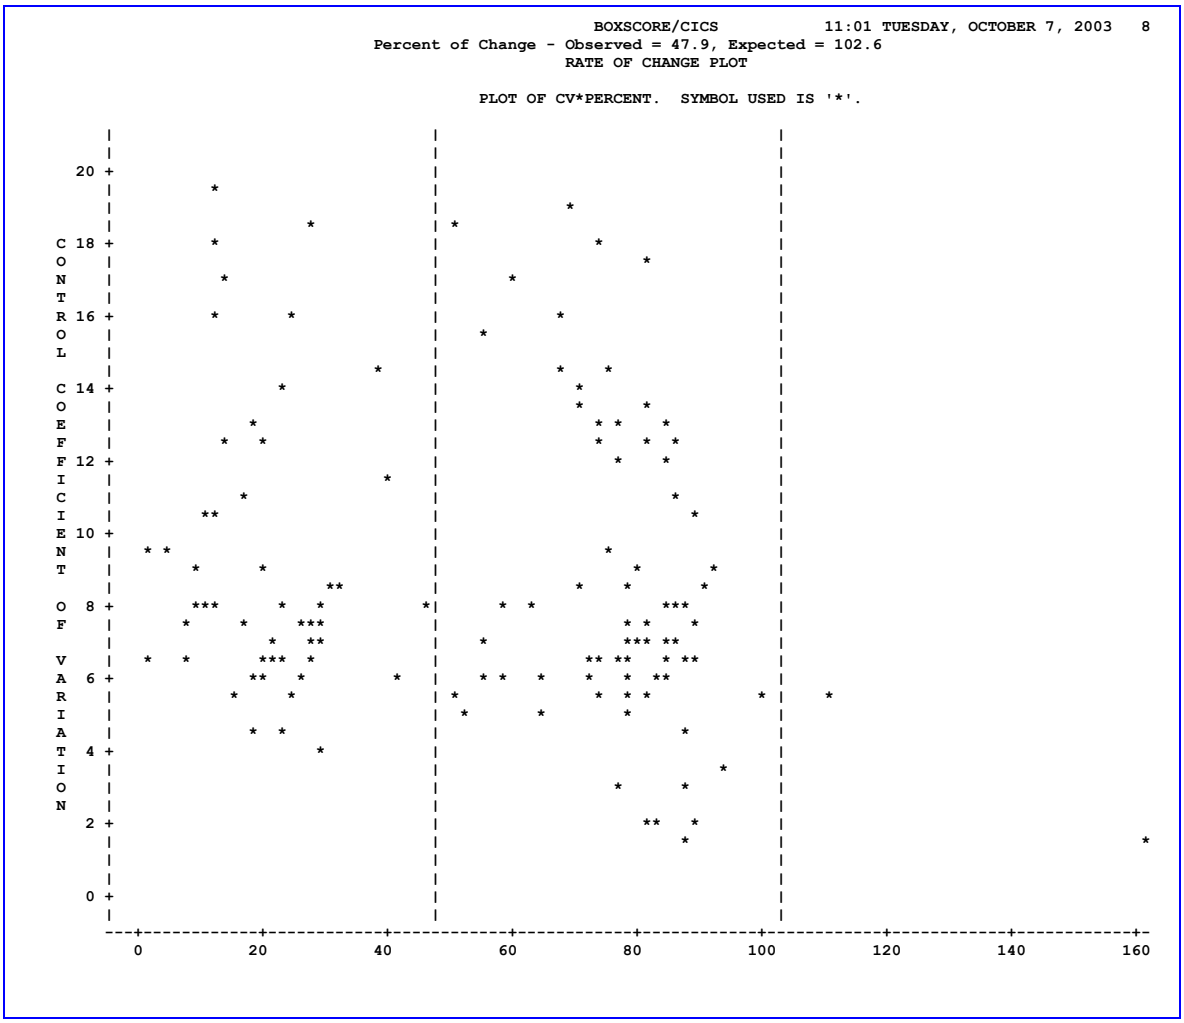

These results are especially important if you do any type of billing based on CICS transaction data.

A difference of 15% from the expected results can produce very high bills for users.

Weighted Average

BoxScore uses the arithmetic average of the changes seen by all of the stable job steps or transactions in order to determine the overall change. Because we are using a 5% to 10% sample of all job steps or transactions (that is the percentage that is usually found to be stable), we feel that our results provide a good view of the rest of the work.

IBM, on the other hand, gives more weight to the steps that use more CPU time. We could agree with that method if we were using all of the steps, but we aren't. We think that a single large step could skew the average too much. There is nothing to indicate that the unstable steps (those not analyzed) represent the same type of CPU usage.

Here is a very simple example with two steps that shows the difference between our calculation and IBM's:

	Control Steps	Control CPU/IO	Study Steps	Study CPU/IO	Study Avg CPU	W & W CPU Ratio	Weighted CPU Ratio
Step1	10	.40	5	.20	110.00	2.00	220.00
Step2	10	1.00	10	.80	20.00	1.25	25.00
Total/Avg	20	1.40	15	1.00	130.00	**1.63**	**1.88**

The columns are:

Control Steps - number of steps found in the Control (before) period for this step name
Control CPU/IO - the average CPU per I/O computed for all of the Control steps
Study Steps - number of steps found in the Study (after) period for this step name
Study CPU/IO - the average CPU per I/O computed for all of the Study steps
Study Avg CPU - average CPU time per step on the new machine
W & W CPU Ratio - the ratio of the control CPU/IO to the study CPU/IO (e.g. .40/.20=2.00). The total for the run is the arithmetic average of these values (1.63).
Weighted CPU Ratio - The CPU Ratio multiplied by the average CPU used during the Study period (e.g. 2.00 * 110.00=220.00). The average for the run is the sum of these divided by the total average CPU time (245.00 / 130.00 = 1.88).

So that our BoxScore customers could see what to expect when IBM analyzes their data, we enhanced BoxScore to add some lines showing the weighted average as well as the BoxScore results. Figure 24 shows one example. This is a batch analysis for a move from a z900 2064-113 to a z990 2084-310, and the customer used the MIX workload rather than low I/O mix. The expected ITRR was 1.95; the BoxScore analysis provided a 1.86 ITRR; and the weighted analysis came up with 1.90.

Figure 24 - BoxScore Addition of Weighted Average

Item	Control	Study	Delta	% Delta	ITRR
Speed of one CPU (physical):					
Expected SU/second	8724.10	17003.18	8279.08	94.9%	1.95
Expected avg MIPS/CPU	184.5	359.5	175.0	94.9%	1.95
Expected max MIPS/CPU	211.5	398.3	186.8	88.3%	1.88
Expected min MIPS/CPU	151.6	306.9	155.3	102.4%	2.02
Observed MIPS/CPU	184.5	343.0	158.5	85.9%	1.86
Change from expected avg			-16.5	-4.6%	0.95
Weighted MIPS/CPU	184.5	349.9	165.4	89.6%	1.90
Change from weighted avg			-9.7	-2.7%	0.97
Machine capacity (physical):					
Expected avg MIPS	2398.5	3595.5	1197.0	49.9%	1.50
Expected max MIPS	2749.5	3982.5	1233.0	44.8%	1.45
Expected min MIPS	1971.0	3069.0	1098.0	55.7%	1.56
Observed MIPS	2398.5	3430.4	1031.9	43.0%	1.43
Change from expected avg			-165.1	-4.6%	0.95
Weighted MIPS	2398.5	3498.7	1100.2	45.9%	1.46
Change from weighted avg			-96.8	-2.7%	0.97

Another addition to the BoxScore product to help understand weighted averages was a change to our Inventory Report. This report lists every step that is used in the analysis. It can now be sorted on the weighted ratio (among other things) so that you can remove outliers if you wish. Figure 25 shows an example of that report. We have even less confidence in the weighted average method after looking

at several of our BoxScore Inventory reports. For most situations, the program that shows up with the most CPU time is the sort utility. Let's explore this in more detail.

Storage Impact

One of the other things we've noticed from the BoxScore runs (especially after looking at the weighted Inventory reports) is that the most significant CPU user is often the sort program (DFSORT, SyncSort, etc.). When most installations move to a z990 from a z900, they significantly increase the storage on the machine. It's not unusual to see an increase between 50% and 200% in central storage. From our years of BoxScore analysis, we've seen that sorts and copy jobs are the ones most affected by a change in storage. If you're trying to measure the speed of the machine, you will want to remove sorts from the analysis or vary storage offline until measurements are complete. On the other hand, if you want to see the total effect on the customers due to both the CPU speed and the storage, then you definitely want to include them in the analysis.

Figure 25 - BoxScore Inventory Report

BOXSCORE/BATCH 16:18 Sunday, April 18, 2004 11
Matched Job/Step/Program Inventory

Weight	CPU Ratio	Job	Stepname	Program	Control Records	Control CPU	Control CPU per Study CV	I/O	Study Records	Study CPU	Study CPU per CV	I/O
618.79	1.09555	JOB0500	STEP0900	DFSRRC00	4	0:42:29.5300	5.6	0.0138	5	0:47:04.1100	3.9	0.0126
256.02	0.97672	JOB0100	STEP0200	DFSRRC00	4	0:17:05.6000	1.9	0.0205	5	0:21:50.6000	1.5	0.0210
227.46	0.95538	JOB0700	STEP0300	DFSRRC00	4	0:17:04.4900	6.4	0.0115	5	0:19:50.4400	10.8	0.0121
226.80	0.95114	JOB0700	STEP0100	DFSRRC00	4	0:17:01.4300	5.6	0.0115	5	0:19:52.2600	11.7	0.0121
152.30	0.87159	JOB0700	STEP0200	DFSRRC00	4	0:12:28.1200	2.6	0.0010	5	0:14:33.7000	1.8	0.0012
122.40	0.68148	JOB0400	STEP0650	IKJEFT01	5	0:12:59.7600	5.7	0.1150	5	0:14:58.0300	10.8	0.1687
109.46	1.02272	JOB0200	CA	CAPUMAIN	3	0:05:16.8600	1.3	0.0007	5	0:08:55.1600	7.6	0.0007
107.74	0.88852	JOB0700	STEPUTIL	DFSRRC00	16	0:27:45.8300	2.6	0.0061	15	0:30:18.9100	4.9	0.0069
107.29	0.88524	JOB0900	STEP0100	DFSRRC00	3	0:05:19.8600	2.8	0.0080	5	0:10:06.0100	4.8	0.0090

Range of Performance

As previously noted, our research discovered that IBM maintains acceptable performance is achieved if the machine performs within 5% of expectations. But we think that this is too great a range with today's machines. As an example, look at Figure 24, where BoxScore reports that the machine underperformed by 4.6%. With this machine, that's 165 MIPS lost. Using IBM's 5% variance, you could lose up to 180 MIPS and still be within expectations. When sizing for this machine, you would need to be able to handle a 360 MIPS variance of capacity. We feel this is unacceptable for many installations, and could make a difference in whether there's room for one or more applications.

Summary

Our current BoxScore customers are not getting the expected performance out of their z990 processors. Although some of this is due to tuning (such as LPARs and goals), the majority of it is because of expectations based on traditional workloads. The low I/O concept came as a surprise to both IBM and their customers. If you are a low I/O candidate, and perform sizing based on the low I/O workload, you will probably come within 5% of those expectations. If you can stand 5% underperformance for batch work, then you'll be fine. In our opinion, however, CICS work is definitely not meeting any type of expectations in many of the installations (even using low I/O). We hope to have more information on CICS soon. As we mentioned earlier, IBM does not agree with our CICS assessment, but we stand by it.

Observations

Many of the old rules have changed, and if you aren't prepared for these changes it will leave you in the difficult position of having a machine with less capacity than you expected. Here are some observations that we think are important to anyone planning their next upgrade:

1. Most of the LSPRs on IBM's Web site are only applicable to 20% of the installations according to their recent conclusions regarding low I/O density. When using the low I/O workload mix, you can expect to see the least amount of increase in speed and capacity as compared to other workloads when moving from a z900 to a z990. The difference in expected capacity between the low I/O mix and the MIX workload varies, but is usually about 8-12% for a typical migration.

2. ITRRs for the low I/O workload mix are not published any place other than in **Cheryl Watson's zSeries CPU Chart** as MIPS, but they may apply to 80% of the customers. Although BoxScore has been enhanced to support its use, we have the following reservations about low I/O methodology:

    a. The documentation for low I/O has not been disseminated to most customers. But according to IBM, your IBM representative or business partner should be aware of the situation. There is a brief mention of low I/O in the LSPR manual, but no reference on the LSPR Web site. This is very inadequate for a concept that is so crucial to explaining performance.

    b. There are no published ITRRs for the customized workload mixes in the LSPR manual. IBM says that the majority of customers should use these or come up with their own workload mixes, but no customized ITRRs are available other than through the IBM sizing tools (that are only available through your marketing representative). To properly determine ITRRs for the low I/O workload mix, the customer would need to compute the harmonic mean of three workload ITRRs (CB-L, WASDB and OLTP-W). *Editor's Note - or you can use our April 2004 z/OS CPU Chart that has already done the calculations and provides these workloads as relative MIPS.*

    c. It is obvious that the low I/O methodology was identified only after the first z990s were out in the field and some reports of underperformance appeared. These first customers attempting to size their new z990 processors did not have this information available to them. We believe that, unless people read this newsletter or see one of the SHARE presentations, they will still not know about it. If they use the zPCR tool, their actual performance should come quite close to the results of that tool. Otherwise, they are buying machines that will probably perform below their expectations.

    d. If this low I/O workload mix truly represents 80% of their customers, then IBM should make a concerted effort to notify prospective customers of this fact, including a major update to the LSPR Web site.

3. It is very important that you use IBM's sizing tools when contemplating an upgrade. Several of the tools are free, although they require some effort on your part. This is the best way to

keep from being surprised, and you can't beat the price. zPCR should be a requirement for every upgrade.

4. The software MSUs for the z990 are approximately 10% below the assumed capacity of the machine, but appear to be more accurate for capacity planning (for the majority of installations that are defined as having a low I/O density). Although this often happens to be about the same amount that many people see as underperformance, we believe that this is simply a coincidence and do not agree with people who say that IBM knew about this when setting the MSUs. The software pricing reduction was announced in August 2003 and the low I/O condition wasn't identified until November or December of that same year. Because many software vendors use MSUs for pricing (and MSUs are based on the capacity of the MIX workload), it's important for software vendors to realize that the lower MSUs for the z990 tend to track more closely to the actual capacity.

5. Don't assume that because you match a workload that produces a higher MIPS rating (such as CB-L), that it will result in more capacity. As we saw in Figure 21, it's all relative.

6. As always, it's important to never count on a single average MIPS rating, as shown in most MIPS charts. We have long believed that you should be using workload estimates, and that's why we publish them.

7. People are not used to having to accept a 5% reduction in capacity when it's over a hundred MIPS. You will need to add this to your plan.

Our Recommendations

Here are some of the recommendations that we'd like to offer those installations planning to install a z990:

1. When increasing your capacity and moving to a much faster processor, you can see some extremely large increases in CPU usage due to latent demand on the previous system. This is often seen when installations implement a large capacity increase and then find themselves out of capacity on the first day. This latent demand is difficult to estimate and determine. (BoxScore is not affected by this because it's looking at the impact on each job using CPU per I/O, and not at the increased number of jobs.)

2. As IBM has noted, there can certainly be LPAR configuration problems and increased overhead for some installations. This is an old problem, but is exaggerated by the large size of these new processors and might certainly be a factor to consider when moving to a z990. Our article on page 29 discusses these considerations. Please review that material and correct the LPAR situation if your configuration is now using, or will be using, too many logical CPs for the physical CPs. IRD may help this situation, but it will take some work. See our TUNING Letter 2003, No. 5 for our article on IRD.

3. Be very careful when sizing the z990 for your particular workloads. If you had been using the CB84 LSPR workload for batch, understand now that the newer CB-L (similar to the older CBW2) workload may be closer to what you will see. **Never, ever, use average MIPS**

or MSUs to size a new machine! You should also determine your I/O density. If you have a CPU-intensive system, then you should use only our low I/O workload MIPS for capacity estimates.

4. If you continue to use average MIPS, plan on using more MIPS than you expect, and take this into account when pricing future software from your ISVs. This recommendation is based on the majority of our BoxScore results. **All but one BoxScore run shows stable job steps taking between 8% and 12% more CPU time than expected when going from the z900 to the z990 and using average MIPS for sizing. This is even worse for CICS transactions.**

5. Plan on tuning your new z990 system after it's installed. This has been a requirement for most installations. For WLM, you may need to change velocities and importance levels in order to achieve the performance you had before.

6. Be sure to get a performance warranty from IBM before purchasing or leasing a z990 (or any machine, for that matter). Make sure it includes penalties that allow you to either reject the machine if it doesn't perform or force enough of a price reduction to pay for the increased ISV software costs associated with adding an engine or two. Obviously, we'd also suggest getting software (such as our BoxScore product or something similar), to provide your own results instead of relying on IBM's analysis. That's why we wrote BoxScore. According to our IBM contacts, they do not feel that having these warranties or measurement software is necessary. They say that IBM is determined to satisfy their customer whether a performance warranty is in place or not. And they also have their own tools to perform the same type of analysis. Our opinion is that having independent advice and measurement software is always a good idea.

7. The greatest danger may loom for the large number of installations that are moving to a z990 because of the improved price/performance. In most of these cases, the installation will attempt to keep the capacity the same, while achieving price reductions in maintenance and software costs. This usually means that the move is to faster but fewer CPs and the work may run at the lowest end of IBM's expectations. They might think that they're moving from 2000 MIPS to 2000 MIPS, but might end up with only 1800 MIPS due to the low I/O condition. In this situation, performance problems can, and usually will, occur. ■

Processor Upgrades

This is a reprint and update of an article from our 1998, No. 6 TUNING Letter about moving to faster processors. We thought it was appropriate to update this article because most people moving to z890s and z990s will be experiencing this type of move. This article covers typical upgrade problems, such as moving to the same number of CPs (but that are faster), or moving to faster, but fewer, CPs.

Typical Upgrade Problems

We get ***lots*** of email from people who tell us heart-breaking stories about how things went terribly wrong with their processor upgrades. Here's a list of how to avoid the biggest mistakes when upgrading a processor:

Get a Performance Guarantee

Insist on getting a performance guarantee or warrantee from the vendor. If you don't and the machine fails to meet your expectations, then you have no recourse. Be careful of the penalty, however, because it is easy for the vendor to simply upgrade your machine if the capacity is less than expected. That is not really good enough. Most sites cannot afford to accept a CPU upgrade because of the increased software cost. We think the only valid penalty clause is one that allows you to return the machine after a period of time and would allow you to bring in a different machine without penalty.

*You simply **cannot** use CPU busy as an indicator of either capacity or speed after a move*

Although we recommend that you get a performance guarantee, we also realize that it is very difficult to confirm capacity. The guarantees should be related to the performance of your critical workloads. Performance guarantees are normally based on capacity projections, but most people notice what's occurring with a specific workload (and that is normally determined by speed, not capacity).

You can't easily make judgments about the speed of a new machine without really understanding the underlying relationships. In most hardware performance guarantees, the vendor includes a restriction stating that the workload must remain the same. This is almost impossible to do. You are either moving to a faster processor where the work will change characteristics (or the users will change their behavior), or you're changing the number of CPs that will affect the number of concurrent users.

Remove Limiting Parameters

Don't let your current parameters limit the capacity of a new machine. Most sites forget to update their parameters when upgrading and may artificially limit the amount of new work on the machine. The most common parameters that are overlooked are:

1. Resource group maximum in goal mode (limits the amount of CPU used by some service classes).

2. VTAM LUs for TSO users (limits the number of TSO users).

3. Storage limits or artificial limits, such as transaction class limits (can limit the number of transactions).

These (and many more) parameters can be found in our two-part series in our 1998, No. 1 and No. 2 TUNING Letter issues.

Review Software Costs

Complete an analysis of the software costs associated with any hardware upgrade. In many cases, the software costs far exceed the hardware costs, yet many sites still forget to do a complete analysis until they've already committed to a specific processor. You should also consider some of the benefits of two processors connected with a coupling facility to qualify for both PSLC discounts on software used on all machines and a reduction in costs for software that only needs to be run on one machine. We're aware of several sites that found that buying two machines plus the coupling facility resulted in a significant reduction in the total outlay for the upgrade.

The increase in software costs is causing many installations to find replacements for products that are marketed by a select group of intractable software gougers (pardon us, we meant to say "companies"). However you choose to deal with the software issue, it cannot be ignored.

Expect Latent Demand

As we've recommended in the past, make it your business to understand latent demand in order to understand why a processor upgrade seems to take more capacity. One of the most common questions we get relates to a condition where the CPU busy is much higher than expected after a processor upgrade.

As an example, a 1000 MIPS machine is currently running at 100% busy during peak period. The site upgrades to a 1500 MIPS machine where they expect to be running at 67% busy with the same workload. We can guarantee this will never happen. We would expect, instead, for the new machine to be running at 100% during peak period too. The reason is latent demand. This is work that is currently being limited by the current capacity and finally has a chance to run (for example, a faster processor allows data entry clerks to enter data faster and process more transactions). It could also be work that is normally run during off-peak hours and finally has a chance to run during peak period (test batch jobs, for example).

In the first example, you can actually measure the increase in work on the system by looking at the increase in transactions. In the second example, the total activity for the day may not increase, but the peak period will increase. For this reason, any analysis of capacity should include both a peak period transaction rate and a daily total transaction rate. You can either use transaction rates from your subsystem, such as CICS, or you can use I/O rates as an indicator of the amount of work you've processed.

You simply cannot use CPU busy as an indicator of either capacity or speed after a move.

Avoid Increased Expectations

Here's a topic bound to cause disagreements in any installation. When you upgrade your processor, you will most often improve the response times and turnaround times for your users. They'll love you!

But what if that increased capacity was purchased because of planned increases in specific workloads or new applications? What happens to the user's response time when that new work comes into the system? It will increase again and the users will hate you!

This is a common scenario and one that we think is avoidable. We believe that you should not give all of the new capacity to users if you plan to take it away again. The easiest method is to place the image in an LPAR, give it a weight that is needed to handle today's workload, and cap the LPAR so the users can't get more. That will leave some idle CPU capacity until the new work appears on the system when you can increase the weight and, therefore, the capacity of the image.

The primary argument against this is that you are throwing away CPU time that users could have access to. Our experience is that once you "loan" your users that new capacity, you will never be able to take it back again without complaints. By that time, too, they might have changed their behavior and require the new level of resources.

We'll get off our soapbox now, but if we could convince installations to use this technique more often, we know that the number of processor upgrades and the number of complaints relating to them would be reduced.

Summary

We've listed the primary reasons for complaints after a processor upgrade. You may also run into some minor problems, such as increased expectations or a misunderstanding of LPAR configuration. If you understand these issues, however, we think you'll be in a much better situation to evaluate what might be happening to your system after an upgrade.

Moving To Faster CPUs

This first situation implies that you are moving to faster CPs with either the same number of CPs or a larger number. This is the most common scenario today, because all of the newer processors (z900, z890 and z990) are faster than ever before.

This is the easiest move of all, and seldom causes any problems except those we listed above. The biggest problem with faster CPs is simply the increased cost of the software, so be sure to determine how expensive the upgrade will be for both hardware and software before you commit to the upgrade.

The previous two topics we just covered (*Expect Latent Demand* and *Avoid Increased Expectations*) are most likely to occur when moving to faster CPs and not reducing their number.

Moving To Fewer But Faster CPs

If you are moving to fewer CPs, with each CP being faster than the older processors and providing more capacity, you might think that there is nothing to worry about. Actually, there are several things to worry about. First, take into consideration all of the items we mentioned in the previous section about moving to faster CPs.

Then you need to look at your workloads. Some workloads really run better with more CPs instead of faster CPs. Take the example of moving from a 2064-102 (2-way z900) to a 2084-301 (1-way z990). The MIPS for the two-way z900 are 427 MIPS or 214 MIPS per CP, while the MIPS for the z990 are 450. This move represents a 5% increase in CPU capacity, so it should present no problem.

But this is moving from a two-way to a one-way that's faster. If all the applications are small applications that don't want to dominate the CPU, there will be no problem. If one of the applications is CPU-intensive and would take more than one of the z900 engines if it could, then you might have problems when you reduce the number of CPs. If the application is high priority, such as CICS, and it had been limited on that 214 MIPS CP, it would take more CPU on the z990 than it had on the z900. The other applications could suffer.

This phenomenon occurs all of the time, such as when moving from a 2064-104 (4-way of 810 MIPS total) to a 2084-302 (2-way of 855 MIPS total). In many installations, that type of upgrade won't work. All it would take in this last example is to have two high priority applications that are CPU hogs, and all other work would suffer. The four-way would tend to limit these two applications to no more than 203 MIPS each (810 / 4), but the two-way would allow them to gobble up 428 MIPS (855 / 2) if they weren't constrained.

Many installations have been hurt during upgrades like this. One installation combined two machines with a total of twelve processors to a single machine with 30% more capacity but only five processors. It was a disaster! They happened to have four very large online systems that dominated the new machine. On the twelve processors, the online systems were restricted by the speed of a single CPU. They thought the machine would last for a year and it was out of capacity by 10 a.m. on the first Monday morning!

In order to see if an application is likely to take more CPU, look to see how much it's using on the current machine. If a single workload (especially a high-priority workload) is using a full CP's worth of processor, then it will most likely take more CPU resources on a faster processor. Maybe that's what you want, but you'd better plan for it. A quick way to see how much a workload is taking is to look at an RMF or CMF report during the peak periods for that workload.

The example in Figure 26 shows a sample RMF report by service class period. "AVG" and "APPL%" are the two fields of importance for sizing CPs. "AVG" is the average number of concurrent address spaces. "APPL%" is the percent of a single CP that's needed for the workload. Thus, a value of 250% indicates that the workload needs 2.5 CPs. In our example, this work is taking 33.4% of a single CP.

If a high priority single workload takes over 80% of a single CP, it's an indication to you that it will probably take more CPU on a faster processor.

Figure 26 - RMF Workload Activity Report

```
                                W O R K L O A D   A C T I V I T Y
                                                                                                  PAGE    1
      OS/390                    SYSPLEX SYSB              DATE 04/18/2002          INTERVAL 05.34.429   MODE = GOAL
      REL. 02.10.00             RPT VERSION 02.10.00      TIME 14.53.25

                                      POLICY ACTIVATION DATE/TIME 04/18/2002 14.53.03

------------------------------------------------------------------------------------------------------------- SERVICE CLASS PERIOD(S)

REPORT BY: POLICY=STANDARD   WORKLOAD=ONLINE     SERVICE CLASS=CICSA        RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=3

TRANSACTIONS     TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--  ---SERVICE----   --SERVICE RATES--  PAGE-IN RATES    ----STORAGE----
AVG       3.21   ACTUAL              16.541   SSCHRT  11.8  IOC          0   ABSRPTN      1156  SINGLE     0.0   AVG          0.00
MPL       3.21   EXECUTION           16.539   RESP     1.2  CPU       1243K  TRX SERV     1156  BLOCK      0.0   TOTAL        0.00
ENDED       59   QUEUED                  2   CONN     0.7  MSO          0   TCB         111.5  SHARED     0.0   CENTRAL      0.00
END/S     0.18   R/S AFFINITY            0   DISC     0.1  SRB          0   SRB           0.0  HSP        0.0   EXPAND       0.00
#SWAPS       0   INELIGIBLE              0   Q+PEND   0.4  TOT       1243K  RCT           0.0  HSP MISS   0.0
EXCTD        0   CONVERSION              0   IOSQ     0.0  /SEC      3716  IIT           0.0  EXP SNGL   0.0   SHARED       0.00
AVG ENC   3.21   STD DEV             12.674                               HST           0.0  EXP BLK    0.0
REM ENC   0.00                                                            APPL %       33.4  EXP SHR    0.0
MS ENC    0.00
```

In Figure 27, we've extracted some data from a few key workloads. Notice that PRDCIC is currently using 97.4% of a single CPU. Because CICS is primarily a single tasking address space (even in CICS TS, over 80% of the work is still done under a single TCB), we would suspect that this CICS region needs more CPU than it's getting. If you move it to a faster CP, it will take more resources. If you reduce the number of CPs but make them faster, CICS may take more resources than anticipated and the lower priority workloads may suffer. Of course, the CICS users would be thrilled.

In the same figure, PRDIMS is also taking almost an entire CP. The AVG field indicates that there are five address spaces in this service class. In order to see if any IMS region will take more CPU in the new machine, you'll need to analyze each of the five address spaces (perhaps use report classes to see if the IMS regions are using equivalent resources, or whether one or more of them uses a disproportional amount).

Figure 27 - Extracts from RMF/CMF reports

Svc Class/ Period	Workload	APPL%	AVG # of AS
PRDBAT	Batch	66.7	65.00
PRDCIC	CICS	97.4	1.00
PRDIMS	IMS	98.4	5.00
HISTC	Hi Pri STC	22.5	1.00
PRDTSO/1	TSO 1st period	41.5	2.75
PRDTSO/2	TSO 2nd	13.5	1.08

When sizing for fewer CPs, look at your primary, high priority applications to see if they will tend to take more CPU resources if available. If so, be prepared to deal with the situation by somehow restricting their access to the CPU (e.g. by lowering their dispatch priority). ■

What's New?

*This section provides a description of new announcements and new products. Announcement letters can be found at **www.ibm.com/news**. Click on Announcements and search for the announcement number.*

On April 7th, IBM celebrated the 40th anniversary of the System/360 by announcing several major additions to the mainframe field. IBM has a separate Web page for this anniversary (**www.ibm.com/servers/eserver/zseries/40years/**) where you can read about the history of the mainframe and download a special PDF file commemorating this event. If you click on 'Read mainframe history' on the right, you'll have access to what occurred each decade. I (Cheryl) especially liked the article written by the Clipper Group (upper right). I started in mainframes the year after this announcement, so I've been able to follow the evolution of the industry. It's been quite a ride, and gets better every year!

This set of announcements is one of the largest in several years and will affect all mainframe users. We've highlighted several of the announcements below:

- z890 - New Low End to Mid-Range Processors (announcement #104-117)
- zAAP - zSeries Application Assist Processor (#104-117 and #104-118)
- "Baby" Shark (#104-119)
- z990 Enhancements (#104-118)
- z/OS 1.6 Additions (interspersed in other announcements)

IBM hopes this 40th anniversary announcement will show the evolution that has occurred and will continue to occur with the mainframe. Enhancements are being made to both the hardware and software that continue to process the majority of the world's commercial data. IMS is 36 years old, CICS is 35 and DB2 is 21 years old. These subsystems are continually being enhanced to support the new and larger volumes of data associated with today's applications. New subsystems such as WebSphere can work with the legacy subsystems (IMS, CICS and DB2) to modernize traditional applications. A new version of WebSphere (V5R1) was announced on April 20th.

z890 - New Low End to Mid-Range Processors

Overview

We think that the heart of the April 7th announcement is the introduction of a new zSeries model, the z890, which provides 28 levels of capacity within this single model. We're most excited about this announcement because it provides a growth path for the small and mid-sized market. The smallest model is only 27 MIPS (32% the capacity of the smallest z800) and the largest is 1,364 MIPS (2.2 times the capacity of the largest z800)[1]. Our only disappointment is that this announcement didn't occur last year before some smaller sites moved off the mainframe. This offering allows small users

[1] IBM indicates that this machine ranges from 26 MIPS to 1365 MIPS. If you use their published ITRRs, you would get 17 to 1364, so they must have more precise ITRRs which they use internally.

to start with a zSeries machine for about US $200,000, while providing growth for the current z800 machines.

If you move from a 9672 to a z890 of the same capacity, the price/performance for the maintenance and software can produce a 40% savings in the first five years using EWLC[2]. Even with a 50% boost in capacity, the z890 will still cost less than the older 9672, which won't even support the latest (and only supported) software releases. In addition to these cost savings, the z890 can also run z/OS.e, which was previously available only to z800 users and can provide even greater savings. z/OS.e can provide pricing at 10% of PSLC prices for qualifying installations. See our TUNING Letter 2002, No. 3, page 49 for a discussion of z/OS.e.

Even if you're a z990 customer, you might want to read how this one is designed because it may foreshadow the design of future machines.

The z890 is very interesting under the covers. It's essentially a z990 4-way processor, which can be modified to look like 28 different machines of varying speeds. There are seven speed ratings, which IBM accomplished by slowing down the speed of the z990 chip. Because you can order one to four processors at one of the seven speeds, you have up to 28 options for ordering. All processors must be ordered with the same speed rating. Note that you will always get all of the processors on delivery, but you must order a feature code in order to activate them. The varying speeds are accomplished by micro-code, which can essentially slow down the processing speed.

You get the reliability and robustness of the z990 chip, but in the form of a smaller processor with lower software pricing. Of course, because it's built on the z990 chip, you get all of the standard facilities of the z990, such as 64-bit architecture, support for IFLs and ICFs, on/off capacity on demand, PCI Crypto and OSA-Express. The actual Multi-Chip Module (MCM) contains 5 CPs, the fifth being used as a SAP (System Assist Processor). Others can be configured as standard CPs, IFLs (Integrated Facility for Linux), ICFs (Internal Coupling Facilities), and zAAPs (new zSeries Application Assist Processors, which we describe on page 48).

Table 2 - z890 Models and MIPS

Speed	Feature	MIPS	Feature	MIPS	Feature	MIPS	Feature	MIPS
1	110	27	210	50	310	72	410	99
2	120	45	220	90	320	131	**420**	**171**
3	130	90	**230**	**171**	330	262	430	3295
4	140	108	240	212	340	311	440	410
5	**150**	**171**	250	333	350	491	450	644
6	160	212	260	405	360	599	460	783
7	170	365	270	707	370	1040	470	1364

The model number is 2064-A04, with feature codes as shown in Table 2. The first digit of the feature code is the number of CPs and the second digit is the speed level (1 to 7, with 7 being the full-speed model). For increased capacity, you can move to the right on the chart, down the chart, or in a diagonal direction in order to increase the capacity. Depending on your workload, you may prefer to

[2] EWLC (Entry Workload License Charge) first became available for z800 machines to provide WLC (including sub-capacity pricing) for stand-alone z800s running either z/OS or z/OS.e. EWLC is now also available for the z890s.

have more processors (move right), faster processors (move down), or more and faster processors (move diagonally down and right). The options will keep you on your toes trying to figure out the software charges for each of the options. As an example, if you want 171 MIPS, you can choose the 150 (one 171 MIPS CP), the 230 (two 85.5 MIPS CPs) or the 420 (four 43 MIPS CPs). Which will work better for your workloads? These moves can all be made dynamically.

MSUs

One of the more important things to understand about the z890 announcement is that the MSUs are different from those used with the z800s and 9672s. For this understanding, we need to go back to the z990 announcement. IBM originally announced MSUs for the z990s that were similar to other MSUs relating to capacity. There were about 6 MIPS per MSU. This was handy because you could look at an MSU rating and easily see the average MIPS capacity of the machine. So a machine rated at 300 MSUs was about 1800 MIPS.

In order to improve the price/performance of the z990, IBM announced MSU ratings for the z990 models that were about 10% lower than the capacity MSUs. These MSUs were then used for pricing. We discuss the difference between the hardware (capacity) and software (pricing) MSUs in our zSeries CPU Chart (September 2003 and April 2004). This reduction came as a great benefit to installations that were trying to justify the move to a z990. Almost all software vendors followed suit and also used the IBM software MSUs for pricing. So, z990 users saw an effective 10% reduction in software costs. Using software MSUs, the ratio is about 6.6 MIPS per MSU.

The good news is that IBM has continued to provide these reduced MSUs for the z890. Just be aware that they no longer publish the hardware (capacity) MSUs. This means that a 100 MSU z800 has about 600 MIPS, but a 100 MSU z890 has about 660 MIPS. IBM has always said that you should not be using MSUs (or average MIPS) for capacity and sizing of a new machine, but unfortunately people still do. We strongly agree that you should use IBM's LSPR ITRRs[3] (or our CPU Chart, which is based on LSPR ITRRs) for sizing. But if your management insists on using MSUs, then they need to be aware that the meaning of an MSU has now changed. Just remember that the older 9672s, z800s and z900s still use the original MSU ratings (6 MIPS per MSU). The z890 uses the new MSU ratings (6.6 MIPS per MSU). And for the z990, IBM originally published both (so you needed to clarify which you're using). Also on April 7th, IBM replaced the z990 hardware MSU ratings with software MSU ratings on their LSPR Web site. For our subscribers, we continue to publish both hardware and software MSUs, although we have to estimate some of the values that aren't published. On the CPU Chart, any values we estimate are shown in italics.

Here's an example of the way in which this change in the meaning of MSUs could lead to a mistake. Assume that you are running a 100 MSU z800 and you want to move to a z890 and increase capacity by 25%. You might think that you'd need a 125 MSU z890, but you'd only really need a 113 MSU z890, which would cost less, especially when you consider all of the software costs. Again, we recommend that you use IBM's free sizing facilities (see page 10) or use workload-based LSPRs or our CPU Chart, and not MSUs.

[3] ITRR - Internal Throughput Rate Ratio

Miscellaneous Support

The minimum amount of memory orderable for the z890 is 8 GBs, with 8GB increments up to 64 GBs. Support is provided for up to thirty LPARs[4]. Crypto processors are optional. Two Logical Channel SubSystems (LCSS) are available, with 16 hipersockets (four times the number on the z800). Another feature is the OSA-Integrated Console Controller (OSA-ICC), which provides a system console and operations support for multiple LPARs at a much lower cost. OSA-Express 2 support is provided with up to 40 ports[4]. The z890 can have up to 40 FICON channels[4] and 420 ESCON channels[4]. There's a new cryptographic function on every processor unit (CPACF) that is enabled via a feature code. The PCIXCC cryptographic adapter is designed to meet FIPS 140-2 level 4 certification. And the PCICA feature offers high-scale performance for SSL transactions.

What This Means to You

The z890 is an absolute requirement for the small to mid-range non-zSeries sites that need to move to supported hardware. The price incentives make it extremely attractive, and the design means that growth can be dynamic between 27 MIPS and 1364 MIPS, with much smaller increments than on the z990. IBM says that there are at least 2000 z800s in the field, and many of these can now move to the z890 instead of the z990 for their next upgrade. The z890, of course, can easily be upgraded to a z990-A08 (the smallest of the z990 models). As we'll see below, the z890 is also an attractive option as a stand-alone coupling facility, a platform for z/VM and Linux, and as a replacement platform for VSE.

For larger installations, the z890 is small enough that you might consider it for a small development machine. You can also use it for your disaster recovery configuration. Another option is to move to a z890 when the smallest z990 (one-way) is not compatible with your workloads that run better with multiple CPs. As we said in the beginning, this design might well be the design for future generations, which means that you need to look more closely at the differences in the number and speed of your processors when considering an upgrade. On these newer machines, you can upgrade in a variety of ways.

A good Redbook on the z890 was published 21May2004: **SG24-6310-00 - IBM eServer zSeries 890 Technical Introduction**. Also, please see our discussion of z890 LSPRs on page 23.

zAAP - zSeries Application Assist Processor

Overview

zAAP provides the ability to assign one or more of the CPs to process only Java work. This is done by ordering a feature code to activate the CP for this special processing. The purpose of this facility is to remove the work done by the zAAP processor from software charges, while providing needed processing power for new Java work. It's thought that Java work tends to take two to three times the resources of traditional techniques. Some IBM-Main contributors have provided examples of Java programs that take 5 to 10 times more processing time when compared to a similar COBOL program.

[4] Feature code 110 has more restrictions than the other feature codes: A maximum of 15 LPARs, 24 OSA-Express ports, 32 FICON channels and 240 ESCON channels.

Although this concept of adding an assist processor isn't new, it adds a level of processing and charge back complexity that will prove to be interesting. IBM has already provided similar offloading facilities by allowing the customer to specify CPs to be used as IFLs or ICFs, but charge back has almost never been an issue for these uses.

The zAAP processors will be available June 30, 2004 for z890s and z990s, although software support won't be available until z/OS 1.6 (which is scheduled for September 2004). zAAP, like its counterparts, will always run at the full speed of the processor, not scaled back. For the z890s that are run at the lower speed levels, a zAAP processor could greatly improve performance. A 2086-110 CP would run at 27 MIPS, but a corresponding zAAP on the same machine would run at 365 MIPS. On a z990-301, however, the standard CP would run at 450 MIPS, but so would the zAAP. In addition, there is between 2% and 3% overhead incurred by moving work to the zAAP. For planning purposes, IBM suggests that you use a 5% overhead factor to allow for all situations.

There is special pricing for this, which is US $125,000 per zAAP processor[5]. This is less than the price for the full-speed standard CP.

The Java Virtual Machine (JVM) at the SDK 1.4 level will be modified to signal the start of Java code. That will be intercepted by MVS and continuing execution will possibly be moved to a zAAP CP. We say 'possibly' because the ability to run Java on a zAAP can be controlled, so that a zAAP is always used or sometimes used. An IEAOPTxx parmlib parameter called IFACROSSOVER controls this decision. If this option is set to 'YES' or omitted, then Java work can be dispatched on a zAAP or can crossover and be run on a standard CP. If the option is set to 'NO', then Java work can only be dispatched on the zAAP.

If IFACROSSOVER=YES, then another parameter will control the actual dispatching technique. Option IFAHONORPRIORITY=YES tells z/OS to dispatch the next highest priority work on the standard CPs regardless of whether it is Java work. This is the default and the way that z/OS works today. If IFAHONORPRIORITY=NO is specified, then z/OS will only dispatch zAAP-eligible work if nothing else is ready to dispatch on the standard processors.

Each installation will need to determine whether they have enough Java work to justify the purchase and use of an entire CP just for that workload. There are mechanisms to help you determine how much you could utilize a zAAP (once you are using Java SDK 1.4 and have APAR **PQ86689** applied). See the WSC White Paper, **WP100417** (*z/OS Performance: Capacity Planning Considerations for zAAP Processors*, 7Apr2004), for techniques on making this estimate. The heaviest users of Java on z/OS are expected to be WebSphere and DB2, although new subsystems will start to take advantage of the new zAAP facility.

You will need z/OS 1.6 and JVM/SDK 1.4.1 to use the zAAP. Some subsystems that will exploit zAAPs are:

- WebSphere 5.1
- CICS/TS 2.3
- DB2 V8
- IMS V8
- WebSphere WBI for z/OS

[5] This may vary in other countries.

We have a lot of questions about this new facility, and we will provide an update when more details are available. As an example, there are new fields in the SMF type 30 and type 72 records to show the amount of CPU time spent on the zAAP and the amount of time the work was eligible to run on the zAAP. CPU time, of course, won't be added to the normal CPU times already recorded in these SMF records. That means a new charge back method will be needed in order to charge the users of this new CP for its use and to justify its expense. But it appears that the amount of time spent in the zAAP will not always be the same. And remember that zAAP CPU time represents a chip running at full power. Consistent and reliable charge back for Java processing will be difficult. Watch for our article on SMF measurements in the next TUNING Letter for more information.

What This Means to You

Because most people are starting to run more mainframe Java applications, such as WebSphere and DB2, the use of the zAAP facility will be closely evaluated by many z890 and z990 installations. If you are doing much Java processing, there is a strong possibility that you can add processing power for Java without increasing the costs of your other software that happens to be running on the same machine.

There will be tools available, such as the SMF field that indicates the amount of time eligible for zAAP processing, to help you determine whether it will be useful. Charge back issues (especially consistency) will be complex, at best. But the large software cost savings may make it worthwhile.

"Baby" Shark

Overview

The TotalStorage Enterprise Storage Server (ESS), Model 750, was announced, and can be thought of as a "Baby" Shark. Although it is sized ideally for the z800 and z890 market, it provides excellent price/performance in a smaller Shark model than was previously available. You can order the 750 with between 1.1 and 4.6 terabytes of storage and 8 GBs of cache. This is compared to the model 800, which can have from 520 GBs to 55.9 terabytes of storage, and from 8 GB to 64 GBs of cache.

Although the 750 is offered at a lower cost per terabyte than its larger brother the 800, it still provides all of the 24x7 support and performance options we have come to expect. These include RAID5 or RAID10, PAV, priority I/O queuing, multiple allegiance, redundant failoverable hardware, up to 64 paths to disk, and support for FlashCopy, PPRC and standby capacity on demand.

The price for the model 750 starts at about US $125,000, but can come packaged with the z890 at an even lower cost. This means that you can start with an ESS at a 40% reduction from larger models.

z990 Enhancements

The following facilities will become available on the z990 on May 28, 2004:

- zSeries Application Assist Processor (zAAP) - discussed on page 48, but will require z/OS 1.6
- Security Enhancements, including
  - EAL5 Certification (which has already been given to z800 and z900)
  - Cryptography extensions
- Parallel Sysplex clustering enhancements
- Expanded z/Architecture
  - 4 Logical Channel Subsystems
  - Up to 1024 channels
- Improved Networking and Connectivity features
  - OSA-Integrated Console Controller
  - FICON Performance Improvement
- On/Off Capacity on Demand enhancements
  - IFLs, ICFs and zAAPs
  - Increased flexibility with CBU and On/Off COD together

What This Means to You

Large installations will be happy to hear about the extension to four logical channel subsystems and 1024 channels. And Java sites (WebSphere and DB2, especially) can possibly make a significant reduction in software costs with the addition of a zAAP.

z/OS 1.6 Additions

Overview

z/OS 1.6 was previewed on February 10th, and we described it on page 35 of our TUNING Letter 2004, No. 1. With more than 70% of the MVS customers currently running z/OS 1.4, this will be the next major release to which people will migrate. z/OS 1.6 becomes available in September 2004. Although there wasn't a separate announcement regarding z/OS 1.6 on April 7th, there were references to z/OS 1.6 enhancements in other announcements:

- Support for zAAPs on z890s and z990s
- Support for up to 24 processors in a single LPAR (maximum of CPs and zAAPs; but they have a statement of direction to support more than 24 processors in a later release of z/OS and/or z/VM)
- Support for four Logical Channel Subsystems (LCSS) with support back to z/OS 1.4
- Support for the OSA-ICC (Integrated Console Controller) with support back to z/OS 1.3
- z/OS 1.6 with optional RACF feature is under evaluation for Common Criteria certification at Controlled Access Protection Profile (CAPP) at EAL3+ and Labeled Security Protection Profile (LSPP) at EAL3+

What This Means to You

The major addition here is the support for the z990 enhancements and the features of the z890. Most of these changes are significant only to larger installations.

Other Announcements

z/VSE 3.1

On April 7th (announcement #204-055), IBM announced z/VSE 3.1 that will not only run on a z890 machine, but will exploit it. The availability date of this VSE version hasn't been announced yet, but it signals that IBM has not left their VSE customers without growth potential. This is an important move for the VSE customers. z/VSE can run on the z890 only in 31-bit mode. IBM sees an extension of VSE by allowing it to communicate with IFLs on the same machine running Linux. Support is added in z/VSE 3.1 for Fibre Channel Protocol (FCP) channel-attached SCSI disks.

z/VM 5.1

Announcement #204-057 provides the details regarding this new release of z/VM, which becomes available September 24, 2004. This is an important release because it requires an architectural level set, which means that it can only run on zSeries machines. There are several enhancements that VM customers have been waiting for:

- Requirement for ECKD™ disks for Linux servers removed - SCSI-only environments are now supported
- Improved cryptographic performance with PCIXCC support
- z/VM security manager (RACF) support for authorization control of virtual server access to Guest LANs and virtual switches
- Support for the OSA-Express Integrated Console Controller

In addition to new features, z/VM 5.1 has a new pricing metric - value units - which is based on the number of engines that will run z/VM 5.1. The new pricing can result in a 50% reduction for z/VM. The ranges are listed in the announcement letter, but we'll give an example here:

1 - 3 CPs are charged at 10 value units per CP
4 - 6 CPs are charged at 9 value units per CP

So one CP is 10 value units and four CPs are 39 value units. The price per value unit is $2,250. So a one-way processor for z/VM 5.1 costs $22,500 (it previously cost $45,000) for OTC (one time charge). Annual maintenance is $563 per value unit.

IBM Tools

IBM now has over 150 database tools for IMS, DB2 and CICS. They say that 19 out of the 24 DB2 tools exploit DB2 V8 today, and the rest will do so by June 2004. The Tools group is very excited because of the inroads they've made recently: over 2000 customer winbacks (customers who had bought other tools and moved back to IBM tools) and over 9000 migrations.

IBM now provides fixed price migration services to assist customers in most of their migrations from products provided by BMC, Computer Associates, Compuware and others. More complex migrations can be negotiated.

Mainframe Charter

There were also some additions to the Mainframe Charter community arena. IBM will be providing a z990 to Marist College in order to promote development of Linux applications and skills on the mainframe. We discussed the Mainframe Charter in our TUNING Letter 2003, No. 4, page 49.

There have also been some updates to the IBM Scholars Program for supporting zSeries in universities. See **www-306.ibm.com/software/info/university/products/zseries** for more information.

Control Block Updates

The z890 and z990s with microcode updates contain changes to the STSI and Node Descriptor control blocks. This information primarily applies to software vendors, but you should be aware of the changes in case you need to obtain updates from your software vendor. IBM has already notified ISVs of these changes, and they should be documented in a new Principles of Operations manual by the end of May 2004.

According to IBM:

> *With the delivery of the zSeries z890 and the updated z990 driver supporting new functions, both announced on April 7, 2004, there are two changes related to the reporting of zSeries model names.*
>
> *The Store System Information Instruction now includes SYSIB 1.1.1 words 25-28 which contain the z990 model in the form A08, B16, C24 or D32 and the z890 model in the form A04. SYSIB 1.1.1 words 16-19 continue to contain the z990 model in the form of 3xx, where xx represents the number of CPs, and will contain the z890 Capacity Settings in the form xy0, where x represents the number of CPs and y represents the level of CP capacity.*
>
> *The model field of the Node Descriptor now contains the z990 model in the form A08, B16, C24 or D32 and will contain the z890 model in the form A04. Prior to the offerings announced on April 7th, the z990 Node Descriptor contained the z990 model in the form 3xx.*

Here's one example. The z990 2084-A08 model has feature codes of 301 to 308 to indicate the number of CPs. Likewise the 2084-B16 uses feature codes 309 to 316. If you issue the 'D M=CPU' command from a z/OS console, you will see the following:

```
Message IEE174I
CPC ND = 002084.A08.IBM.02.000000012345 (Node Descriptor)
CPC SI = 2084.308.IBM.02.0000000000012345 (STSI instruction)
```

The Node Descriptor on the original 2084s contained a 308 instead of A08.

ECMB Update

Brian Currah from BDC Computer Services, Inc. provided some additional information about the ECMBs (Extended Channel Measurement Blocks) we mentioned in the last issue (see TUNING Letter 2004, No. 1, page 10):

- As we noted, the move to ECMBs will provide a modest amount of virtual storage constraint relief, because the previous allocation requirement for above-the-line SQA storage will be eliminated by obtaining storage within a common data space instead. But another benefit occurs because of the reduced demand for the real storage needed to back those pages. Because the CMB storage is allocated in one contiguous pool at IPL time, multiple real storage frames will be needed to back this storage, and they will be unavailable to other applications for the life of the IPL. This problem would have been even worse if the ECMBs had not moved, because their size was doubled. But ECMB storage is not allocated in a pool, and it is not allocated at all until the device referenced by the ECMB is accessed for the first time.

- We stated that the IOSCMB macro will return results in either a CMB or an ECMB, but didn't indicate that the user of the macro can specify which format is desired. Regardless of whether the data is stored in a CMB or an ECMB, the macro will convert the data and return the output in the format requested. This means that existing monitors can be upgraded to use the IOSCMB macro but can still get the data back in the familiar CMB format they expect. Similarly, if you wish to upgrade your monitors to use the new ECMB format, you may do so even if the data is not yet stored in that format within the operating system. You just need to be running the level of the operating system that supports the IOSCMB macro.

- One of the SHARE presentations we attended indicated you could use bits within the CMCTCPMFMODE flag to determine whether ECMBs were being used. Brian indicates those bits were actually added by APAR **OW38548** (04Jun1999) and relate to FICON channel measurements. The correct indicator that you should be testing is CMCTECMBMODE (bit 0) within a new field called CMCTECMBFLAGS.

We always appreciate these updates from our users based on their own experiences with some of these new facilities. Many times they will pass along information that will help the rest of us to save time and avoid problems. Thanks to Brian for clarifying this new facility.

Change in Throughput

Editor's Note: This section has been included because of its comments about the z990, but isn't recent information.

Although we have been discussing the capacity and speed of the z990, there are also some additional considerations. As an example, response times have improved dramatically in some installations when moving to a z990. This often occurs when increased capacity is installed. It could also be due to the scaling and improvements made in several areas of capacity. Figure 28 shows an interesting diagram from an IBM presentation. In this diagram, you can see the relationship between the z990, z900 and older G5s and G6s for the maximum I/O bandwidth, cycle time, gigabytes (GB) of storage and number of CPUs.

Our concentration in this article has been on the speed and capacity of the machine, which is fairly closely related to the cycle time. But you can see that for large installations, the potential for the higher bandwidth, greater number of CPs and larger storage could significantly affect the amount of work that can be processed by a z990. Things like larger storage sizes can also improve throughput by decreasing the amount of I/O required on the machine. ■

Figure 28 - z990 Throughput

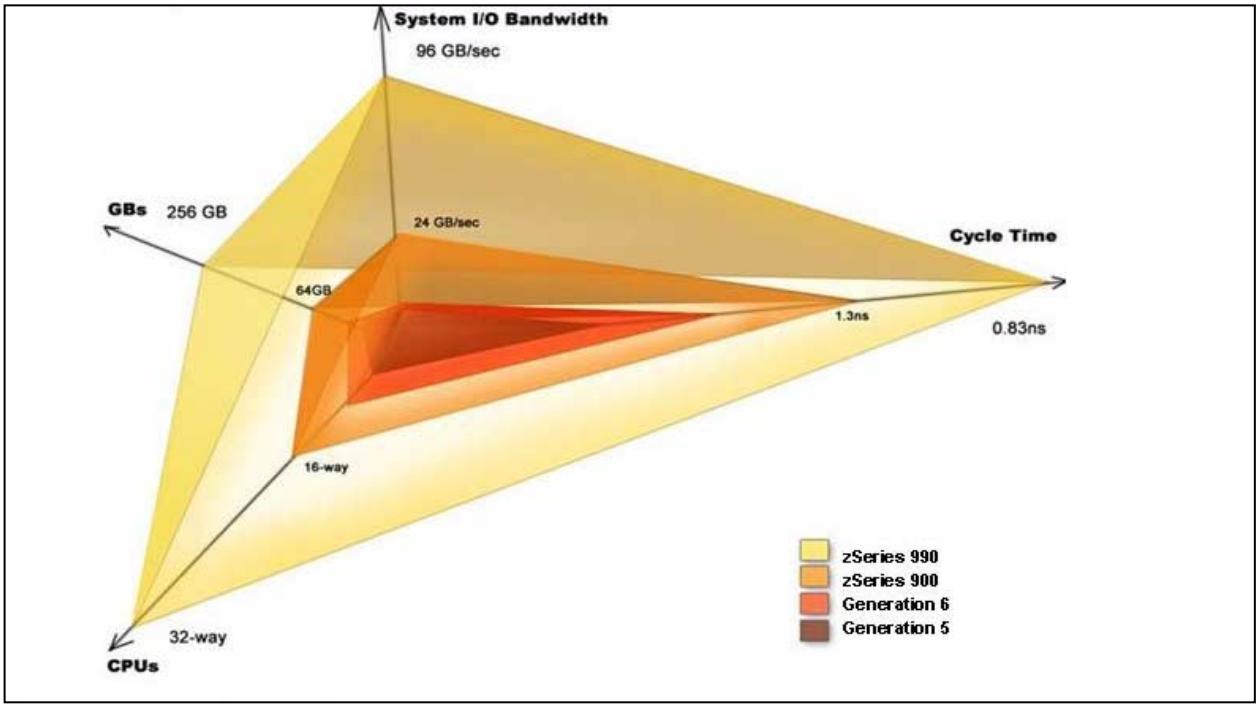

Cheryl's List

Here's a summary of the last transmissions sent to subscribers of our free electronic Cheryl's List. We've eliminated any sections printed in a previous newsletter. Past issues of Cheryl's List can be obtained in full at ***www.watsonwalker.com/archives.html****. See the last paragraph for instructions on how to sign up for the list.*

Cheryl's List #87 - 31 March 2004

1. Cheryl Watson's TUNING Letter 2004, No 1

The fifty-page 2004, No. 1 TUNING Letter was emailed to electronic subscribers today. The print issues will be mailed next week. *The rest of this item was published in TUNING Letter 2004, No. 1 in the Management Summary section.*

Cheryl's List #88 - 7 April 2004

1. 40th Anniversary of S/360 and IBM Announcements
2. Special Offering on Video Tape Performance Classes
3. Typo Correction

1. 40th Anniversary of S/360 and IBM Announcements

Today, April 7th, IBM celebrated the 40th anniversary of the System/360 by announcing several major additions to the mainframe field. IBM has a separate Web page for this anniversary (**www.ibm.com/servers/eserver/zseries/40years/**) where you can read about the history of the mainframe and download a special PDF file commemorating this event. If you click on 'Read mainframe history' on the right, you'll have access to what occurred each decade. I especially liked the article written by the Clipper Group (upper right). I started in mainframes the year after this announcement, so I've been able to follow the evolution of the industry. It's been quite a ride, and gets better every year!

This set of announcements is one of the largest in several years and will affect all mainframe users. Our next TUNING Letter will be out in a few weeks and has many pages describing these new announcements and their impact on installations of all sizes.

We're extremely excited about the heart of the announcement: the latest line of small to mid-range processors. Called the z890, this single machine can be configured into one of 28 different capacities, designated by a feature code. You can configure the machine as a one-way to a four-way with up to seven speed ratings ranging from 26 MIPS to 1365 MIPS. Starting at about $200,000 (US), this machine fills the need for a low-end box and yet provides a machine that can be upgraded to more than twice the capacity of the largest z800. The good news is that IBM is providing the same

MSU reduction that they provided on the z990 (this reduces software costs by about 10%). On the other hand, IBM is no longer providing the MSUs that relate to capacity, so you need to be careful and avoid using the MSUs for capacity purposes. A 100 MSU z890 has about 10% more MIPS than a 100 MSU z800. We'll cover this in much more detail in our upcoming TUNING Letter, along with our new CPU Chart.

The rest of this item was expanded on in much greater detail in our What's New? section starting on page 45.

2. Special Offering on Video Tape Performance Classes

We stopped teaching classes in 1999, but our videos have gone on teaching new performance analysts and capacity planners ever since then. Our week-long classes with video tapes and course materials have until recently been offered on our Web site for between $1995 and $2195. We are now pleased to announce that we will be offering them for $1500 each.

The basics presented in these videos haven't changed since Cheryl taught them. There have been some new measurement fields and some new facilities, such as IRD, but the concepts, the logic, and the reports simply haven't changed that much. If you have some new system programmers, performance analysts or capacity planners, these courses can provide their training for a fraction of the cost of other alternatives. Learn how WLM really works, what sysplex, shamplex and parallel sysplex are all about, how data sharing works, how to interpret RMF and CMF reports, which SMF records provide the most valuable insights, and how best to analyze your system performance.

The four classes available are:

Advanced OS/390 Performance & Capacity Planning (filmed May 1999)
OS/390, Parallel Sysplex & Workload Manager (filmed October 1998)
Advanced MVS Performance and Measurement (filmed in late 1995)
Exploiting MVS/ESA Facilities (filmed in late 1995)

See **www.watsonwalker.com/classintro.html** for a description of the courses.

3. Typo Correction

Zion Botzer from Leumi Card pointed out an IBM error found on page 5 of our TUNING Letter 2004, No. 1. The title for APAR **OA06168** is incorrect on the IBM APAR database. The title is "Various RMF Problems Caused by **OA05167**." The actual APAR that caused the problem is **OA05197**, *not* OS05167.

==========

SUBSCRIBE TO CHERYL'S LIST

If you would like to obtain an email subscription to Cheryl's List, just go to **www.watsonwalker.com** and fill out the form under "Cheryl's List." That signs you up. Remember, it's a one-way list, from us to you. If you make a "reply," it will come just to us, not to the other members of the list. ■

ADVICE YOU CAN TRUST!

In 1965 Cheryl Watson began studying mainframe measurement and performance, and since 1989 she and Watson & Walker have been dedicated to providing education and expert advice for analysts everywhere.

Through our publication, **Cheryl Watson's TUNING Letter**, you can stay abreast of the latest changes, improvements, techniques and prices, saving you and your company time and money.

And through our **BoxScore** software, you can use Cheryl's experience to monitor changes in processor speed and gauge the effect on CPU usage after hardware or software changes.

Since January 1991 Cheryl Watson has filled thousands of pages with her practical and impartial advice for z/OS tuning and measurement. Over ten thousand analysts in more than forty countries rely on her to keep themselves up-to-date and their systems finely tuned.

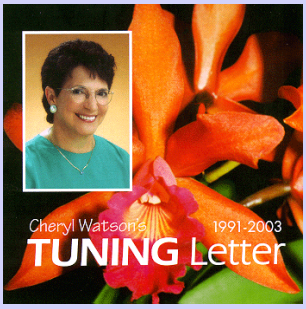

Six times a year readers benefit from timely and in-depth articles on WLM, Parallel Sysplex, LPAR Management, WebSphere and UNIX, as well as the latest information on processor upgrades, software pricing, critical APARs and much, much more. Readers also receive our CPU Chart containing all zSeries processors with information such as MSUs and MIPS by workload.

And there are archives! Electronic subscribers who receive the **TUNING Letter** via email also receive our searchable CD-ROM of all past issues. Updated annually, it is one of the most comprehensive bodies of performance advice ever published.

Don't delay. Take advantage of this valuable resource today!

Cheryl Watson's **ADVISOR** Series

BoxScore

BoxScore is the first in **Cheryl Watson's ADVISOR Series** of software for the mainframe.

After making a change in software or hardware, you might ask:

Did I get the CPU I paid for?
Did that software upgrade cost me CPU?
What's my LPAR overhead?
Did my tuning efforts save any CPU?
Which jobs were affected the most?
How much is data sharing really costing me?
Is the new release of CICS costing me more?
Have ANY of my jobs significantly changed their CPU usage?

BOXSCORE HAS THE ANSWERS!

It not only shows you what happened, but, through Cheryl's advisory messages, tells why it happened and what steps you might want to take next.

Visit us at
www.watsonwalker.com
or call 1-800-553-4562