



BY CHERYL WATSON

# Vendor Performance Ratings: Why Your Results May Differ – Part III

There are several factors that need to be considered if your installation's performance ratings differ from your vendor's claims.

**PART** I of this series (*Technical Support*, December 1997) defined various terms, examined the use of vendor claims for CPU performance and described how vendors meet their claims. Part II (January 1998) examined IBM's *Large Systems Performance Reference (LSPR)* Amdahl's and HDS's performance ratings, and why you should use these claims. This month's article examines the factors that may cause your installation's performance ratings to differ from the vendors' claims. The final article in this series will complete this discussion and conclude with some tips on what you can do when your results differ from your vendor's.

## WHY YOUR EXPERIENCE MAY DIFFER

Why wouldn't you get the same performance as the vendor out of a processor model for your workloads? There are many reasons, and I'll address the most important among them:

1. Workloads vary.
2. Your workloads don't match the vendor's.
3. You measure different things.
4. Your mix doesn't match the vendor's.
5. The workloads vary throughout the day.
6. The volume affects capacity.
7. Constraints in software affect capacity.
8. Constraints in hardware affect capacity.
9. LPAR affects capacity.
10. Dispatch priorities affect capacity.
11. Software levels affect capacity.
12. Levels of PTFs affect capacity.
13. Different facilities are invoked.
14. The amount of storage affects capacity.
15. The level of tuning affects capacity.
16. User's behavior changes.
17. The one thing that remains consistent is that you will always have change!

18. All of these reasons.

This article describes the first nine reasons and the concluding article the remaining nine reasons.

## WORKLOADS VARY

The primary reason that a single performance estimate will not work for most sites is that performance differs for each type of workload. In the newer processors, the range of this difference is getting larger with each new model.

I think the following summary of these points made in IBM's *LSPR* manual [REF001] helps provide some insight into what you might expect to see:

- a. The actual MIPS rate for a model will, in general, be highest for workloads at the batch end and lowest for workloads at the online end of the spectrum.
- b. When comparing n-way models to their corresponding uni-processor model, the actual capacity will be higher for workloads at the batch end and lowest for workloads at the online end of the spectrum.
- c. When comparing models with larger high-speed buffer caches to those with less, the capacity will be higher for workloads at the online end and lowest for workloads at the batch end of the spectrum.

One problem is that your workloads aren't necessarily designed like those that IBM uses for their *LSPR* workloads. For example, you might have some TSO users who use a

lot of SAS (close to FPC1 workloads), others who access DB2 frequently (close to DB2 workloads), and others who spend the bulk of their time in ISPF (close to TSO workloads). The number of each type of user will determine which part of the scale you're on when evaluating TSO. BMPs (IMS batch programs) may look much like IMS and yet may have many of the characteristics of CBW2. To use LSPR effectively, you must be aware of the workload mix you're executing. The workloads were described in Part II of this series.

#### YOUR WORKLOADS DON'T MATCH THE VENDOR'S

IBM has defined some very specific workloads, and while Amdahl and HDS have their own workloads for testing, we don't know what they are. You will need to determine how well your workload matches the vendor's workloads before you can tell if their estimates will be useful.

Here are a few examples where the performance of some workloads might not meet the vendor's expected performance claims:

1. IBM's TSO workload is ISPF-based and has a large amount of editing and browsing types of transactions. If your workload is primarily FOCUS or ADABAS, then your performance probably won't be the same. FOCUS and ADABAS have characteristics that are much closer to CICS and DB2 than TSO.
2. A few customers found that some batch jobs took much longer than expected when they moved from a bipolar processor to a CMOS processor. It turned out that the problem was due to the fact that the packed decimal instruction set was much slower on the CMOS 9672 models than on the bipolars. A heavy use of packed decimal instructions tends to occur in COBOL programs that use subscripts for heavy table processing and were compiled with a compiler option of 'TRUNC=BIN'. IBM didn't run into this particular combination of heavy packed decimal work because their benchmark programs used indexes rather than subscripts. (I remember teaching students that they should try to use indexes rather than subscripts back in early 1970, but programmers

and even vendors are still using subscripts!) This phenomenon has been significantly improved with some microcode changes, but it still exists in many of the IBM 9672 models and HDS Pilot models. For more information on this, see **WSC Flash 9608** and the archives from "Cheryl's List" [REF002].

---

**The primary reason that a single performance estimate will not work for most sites is that performance differs for each type of workload. In the newer processors, the range of this difference is getting larger with each new model.**

---

3. As mentioned earlier in the DB2 workload description, IBM's DB2 transactions don't cause the DB2 Sort Assist facility to be invoked. Since many applications do require a DB2 sort, your workloads could get better or worse performance when moving between processors with or without the sort assist facility.
4. One of the most common problems I've seen recently results from a much larger occurrence of work that uses floating point. SAS, for example, uses floating point for most of its work. Any installation with a large percent of SAS in their daily processing should consider the FPC1 workload as being more representative of SAS than other workloads. Since FPC1 isn't used to determine the #MIX from IBM, SAS users are often surprised by some quite low (or quite high) ITRRs on FPC1 workloads on some models.

#### YOU MEASURE DIFFERENT THINGS

In describing IBM's LSPR technique in Part II of this series, I referred to their use of "processor utilization." This is all of the captured CPU usage for the measurement

interval and includes CPU time consumed by all the system address spaces such as MVS, JES, RACF, VTAM, GRS, CONSOLE, etc., not simply the time recorded by the application in the SMF type 30 (job termination) or type 72 (workload by performance group or service class) records.

IBM can obtain all of the measurements because they run in a dedicated, standalone environment. It's much harder for an installation to obtain all of the CPU usage measurements for a specific workload. For example, if you run TSO and CICS at the same time, how much of MVS, RACF, VTAM, etc., is being used by the TSO workload and how much by the CICS workload? You simply can't tell.

So if you see a CICS ITRR (discussed in Part II) between two machines of 1.2, does that mean reduction in CPU time is just in CICS or will part of it be seen in reduced CPU time for MVS? You don't really know because IBM is really measuring multiple things at one time — that is, the SMF time of the region, MVS, VTAM, initiators, JES, etc.

#### YOUR MIX DOESN'T MATCH THE VENDOR'S

The published #MIX ITRR by IBM and the average performance estimate by Amdahl represent some mix of workloads. In IBM's case, the assumption is that there is 20 percent TSO, 20 percent CICS, 20 percent IMS, 20 percent DB2, and 20 percent traditional batch. This isn't representative of any installation that I've ever seen.

So you'll need to determine your own mix. For daytime processing, you might want to look at your peak processing period and determine the make up of the work at that time. For example, you might be moving from a 9672-R53 to a 9672-R83 and you run 50 percent CICS, 10 percent TSO, 10 percent batch, and 30 percent "other things" such as MVS, RACF, VTAM, monitors, operation's started tasks, and scheduling programs.

When using a variety of workloads, it's easiest to determine the percent of each type of work during the peak interval (that's when the capacity of the machine is the most important). Simply group MVS and supporting functions with the miscellaneous workloads and use the #MIX ITRR. Let's assume that you had some work on an eight-way 9021-982 and planned to move it to a 10-way 9672-RX5. Also assume that you were running 70 percent CICS, 10 percent TSO, and 20 percent other (MVS) during

Figure 1: IBM LSPR ITRRs

Machine	TSO	CICS	DB2	IMS	CB84	CBW2	FPC1	#MIX
9021-982	6.00	7.23	6.07	6.18	7.13	7.20	6.95	6.48
9672-R15	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
9672-R55	4.40	4.34	4.02	3.98	4.28	5.02	4.86	4.20
9672-RX5	6.48	6.61	5.92	6.11	6.76	9.09	9.61	6.36

the peak intervals. From Figure 1 (which was described in Part II), we can calculate the ITRR for CICS to be .91 (6.61/7.23), the ITRR for TSO is 1.11 (6.48/6.00), and the ITRR for #MIX is .98 (6.36/6.48). That's 70 percent at .91, 10 percent at 1.11, and 20 percent at .98 for a combined ITRR of .94.

**THE WORKLOADS VARY THROUGHOUT THE DAY**

Of course, what has just been discussed is for the typical peak processing time. What about the other times of the day? If you run busiest during daytime processing and are able to complete nightly processing in plenty of time, you can probably just use the daytime estimates.

But if you have a tight batch window at night, as many installations do, you will need to calculate a daytime ITRR and a nighttime ITRR to better determine the effect of a processor change. It would be quite possible to find a site with a mix of 70 percent online during the daytime peak hour only to find the mix has shifted to 70 percent batch in the nighttime peak hour.

As more companies are going to international processing windows, the variation between day and night processing is reduced. Even the online workloads will vary dramatically throughout the day.

**THE VOLUME AFFECTS CAPACITY**

IBM, as does Amdahl and HDS, ensures that they are running the system at close to capacity — not exceeding it and certainly not severely under-utilizing it. For IBM that means that measurements are taken at close to 100 percent utilization for batch and FPC1, and at both 70 percent and 90 percent for the online workloads.

Your results will almost certainly vary if you run at different capacities. Frankly, few sites will upgrade to a new machine and immediately run at between 70 percent and 100 percent busy. A new machine

almost always has excess capacity, and this will affect how much CPU is needed for the workload.

For some models, being underutilized will actually provide worse CPU overhead due to the management of high-speed cache and how work is dispatched to the CPUs. Other factors, such as LPAR processing, can add to several "low utilization" effects. For most models, however, being underutilized will result in less CPU time per transaction than the work will see as the system gets busier.

That means that shortly after moving to a new processor, you will tend to see very good performance. As you get more work on the system, which may be many months later, the CPU usage of the system will increase.

In almost every analysis I've made, jobs will take more CPU time when the CPU utilization is at its highest. This is often referred to as the "multi-programming" effect. If you measure the data at 50 percent CPU busy, it will always be to the vendor's favor because the machine will be able to get the work done in less time than estimated at higher utilization.

This phenomenon is seen very frequently. An installation that has been severely constrained for months (running well over 100 percent busy for long periods of time) might replace their current machine with a model that has a higher capacity so the entire workload can be processed while only running at 60 percent busy on the new processor. Prior to the upgrade, jobs were taking additional CPU time due to the high CPU utilization. After the upgrade, the jobs take less CPU time due to the low utilization. The result is that the new machine will seem faster than expected based on the vendor's estimates.

Therefore, you will need to wait until you've reached full utilization on your processor before knowing whether you have obtained the processor capacity that you had planned for.

**CONSTRAINTS IN SOFTWARE AFFECT CAPACITY**

As IBM points out in their *LSPR* manual, they ensure that no software or hardware constraints exist during their measurement period. That means that on your new machine, poorly structured and managed JES initiators or an IPS parameter set to limit the amount of work on the system could seriously affect capacity.

Unfortunately, this happens quite often when an installation upgrades to a new processor. There are several dozen parameters that should be modified when you upgrade to a larger capacity machine. If these parameters aren't modified, you could be restricting the capacity of your new machine. A simple parameter, such as the domain constraints in the IPS, could cause an increase in the amount of swapping, and therefore, overhead in the new model.

**CONSTRAINTS IN HARDWARE AFFECT CAPACITY**

IBM also eliminates hardware constraints during their testing because they don't want to consider the CPU cycles spent dealing with the constraint. For example, they don't want to spend CPU cycles in paging when the intention is to determine the speed of the CPU for a specific type of work.

You should be aware that if you have any hardware constraints, such as lack of I/O paths, poor cache hit ratios, poorly performing DASD, storage shortages, or other hardware constraints, you could be impacting the potential capacity of your machine.

**LPAR AFFECTS CAPACITY**

Perhaps the most significant reason that your workloads may not match your vendor's expectations is that all performance claims are made for a non-LPAR environment. The vendors aren't trying to hide anything, but they simply can't account for all of the variations seen in LPAR configurations.

LPAR processing, whether it's from IBM's PR/SM, Amdahl's MDF, or HDS's MLPF will take additional cycles for processing time. A small portion of LPAR processing may be displayed in the partition data available from RMF and CMF, but that is only the LPAR management time and does not include the bulk of the actual overhead. Most LPAR overhead is actually experienced by the workloads, and their CPU time (TCB or SRB) will increase in an LPAR situation.

The amount of increase is quite variable and dependent on several factors. The primary

factors are the number of LPARs on the machine, the total number of shared logical CPUs, the ratio of logical to physical CPUs, and the activity in the other LPARs. An increase in any of these four factors will cause an increase in the CPU time for your work. This CPU time has not been considered in the vendor's announced performance claims (nor can it be). The LPAR overhead could be as small as 2 percent (in a production LPAR that's given 95 percent of the machine) to 25 percent (in a grossly over-configured, multiple LPAR, multiple CPU environment).

You need to take this into consideration if you are running in any type of LPAR environment.

In Part IV of this article, I'll complete this discussion and conclude with some tips on what you can do when your results differ from your vendor's.

---

## BIBLIOGRAPHY

[REF001] *Large Systems Performance Reference*, IBM, SC28-1187

[REF002] Watson, Cheryl, Listserv archives can be found at <http://www.watsonwalker.com/archives.html>. 



---

NaSPA member Cheryl Watson is an internationally-known consultant and author who has been working with IBM mainframes since 1965. She is the president of Watson & Walker, Inc., and is the principal author of Cheryl Watson's TUNING Letter, a journal of OS/390 advice published since 1991. She will soon be releasing BoxScore, the first in her ADVISOR Series of low-cost performance analysis software. Visit her web page at [www.watsonwalker.com](http://www.watsonwalker.com).

©1998 Technical Enterprises, Inc. For reprints of this document contact [sales@naspa.net](mailto:sales@naspa.net).