

How Much of Your Midrange Computing Power Is Actually “Usable”?

Bradford Camp, Safeway, Inc.
Christopher Lynn, Safeway, Inc.

When examining the CPU utilization of midrange servers in an enterprise, at first glance, they often appear to be significantly under-utilized. Is that really the case? This paper discusses how high availability, load balancing, multiple standards (platforms, operating systems, middleware services, etc.), and application incompatibility affect the amount of computing capability which is actually usable. Can you explain to executive management why 25-30% peak CPU utilization is reasonable if you engage in a lot of high availability, load balanced, and multi-platform solutions? Would you like to be able to?

Computing Power – The Under-Utilization Gap

There are many facets of capacity planning but organizations often focus primarily on the performance and resource utilization levels of individual servers. Although those activities can be very valuable, we have also branched out into what we refer to as Macro Capacity Planning. Our definition of Macro Capacity Planning includes examinations of resource utilization at the enterprise level. This enables us to increase resource usage efficiencies and price performance by identifying under-utilized resources. Typically, these efforts result in summary capacity reports covering each of the technology categories we utilize (for example, mainframe, midrange servers, Tandem, etc.).

Recently, as we examined each of these categories we noticed that the gap between available midrange CPU resources and actual CPU utilization was wider than we believed it should be. As a result of this observation, we decided to take a closer look at this environment to understand and, if possible, narrow this gap.

We have defined the following terms as they are used in this paper.

Term	Definition
Available	Physical resources on the server
Usable	Resources not reserved for HA purposes or unusable due to incompatibilities
Unusable	Available, but not Usable
Unutilized	Usable, but not used

By defining the computing power of each midrange server in our enterprise based on one standard (we chose the SpecIntRate2000 Rating) and identifying additional factors affecting the amount of usable computing power, we determined that we had not been adequately reporting the amount of truly usable but unutilized midrange computing resources.

How Much Power Is Really “Usable”?

In order to evaluate our midrange CPU utilization gap and to be consistent, we needed a standard computing power identifier, or computing capability, for each server in the enterprise.

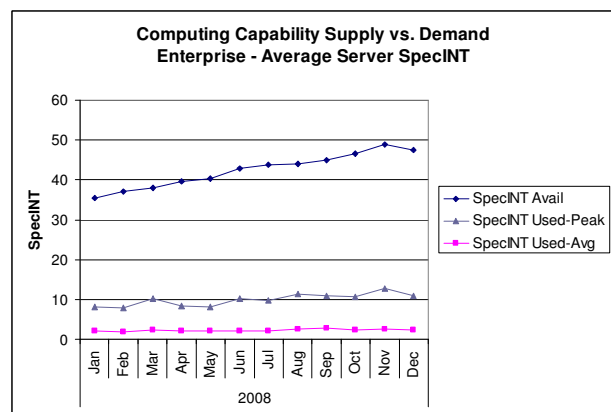


Figure 1 – Computing Capability Supply & Demand

The most readily available option was to use the Spec Rating (Integer CPU2000Rate, hereafter referred to as SpecINT) as reported by our primary metric gathering tool. This approach allowed us to define our enterprise-wide, per server, and per

“average” server computing power rather than simply looking at CPU utilization percentages.

This is important because a 90% idle 10 SpecINT server does not represent the same amount of “waste” as a 90% idle 200 SpecINT server. We wanted to be able to identify unused computing capability, not just a percentage of existing CPUs or of some un-quantified whole.

Next, we calculated the enterprise SpecINT used value by multiplying the CPU utilization rate of each server by its SpecINT. Finally we rolled all of that information into a comparison of SpecINT Available vs. SpecINT Used for an “average” server (total SpecINT in the midrange environment divided by the total number of servers). The gap between these two values, SpecINT Available and SpecINT Used, was much larger than we believed it should be (see Figure 1 above).

So, what might explain the size of this gap? We had been introducing more virtual environments and migrating existing applications to them. We had also been engaged in consolidation efforts to reduce the number of servers deployed and to more efficiently use our large shared environments. Why were these efforts not more successful at shrinking the gap?

As we closely examined these server environments it occurred to us that there may be other reasons why a low proportion of CPU resource was being used. First, we considered the type of environment (production, quality assurance, or development) these servers supported. Recognizing that production was our largest and most critical, we focused our attention on that group of servers. Since we were evaluating CPU usage at an enterprise level, another consideration which directed us toward production was the more erratic use of QA and development environments which tend to distort the overall CPU usage picture, while production servers tend to have more consistent usage patterns.

Another consideration was whether to focus on average or peak CPU utilization. We collect and store both metrics, but, since we were examining

the gap between total available and actual used, we decided to focus our attentions on peak. The primary motivator for that approach was the realization that these midrange systems need to be able to address peak requirements, so they really should be sized with the highest utilization time periods in mind.

As we explored possible factors influencing this gap we began to consider that the amount of computing power we believed to be available may not actually be usable. With this in mind we turned our focus to the following three key factors:

- Sustained Peak Maximum CPU Level (SPMCL)
- High Availability (HA)
 - Redundant Solutions
 - Load Balancing (LB)
- Incompatibility
 - Multiple Technology Standards
 - Application Incompatibility

Figure 2 shows the amount of computing power which becomes unusable due to these three factors.

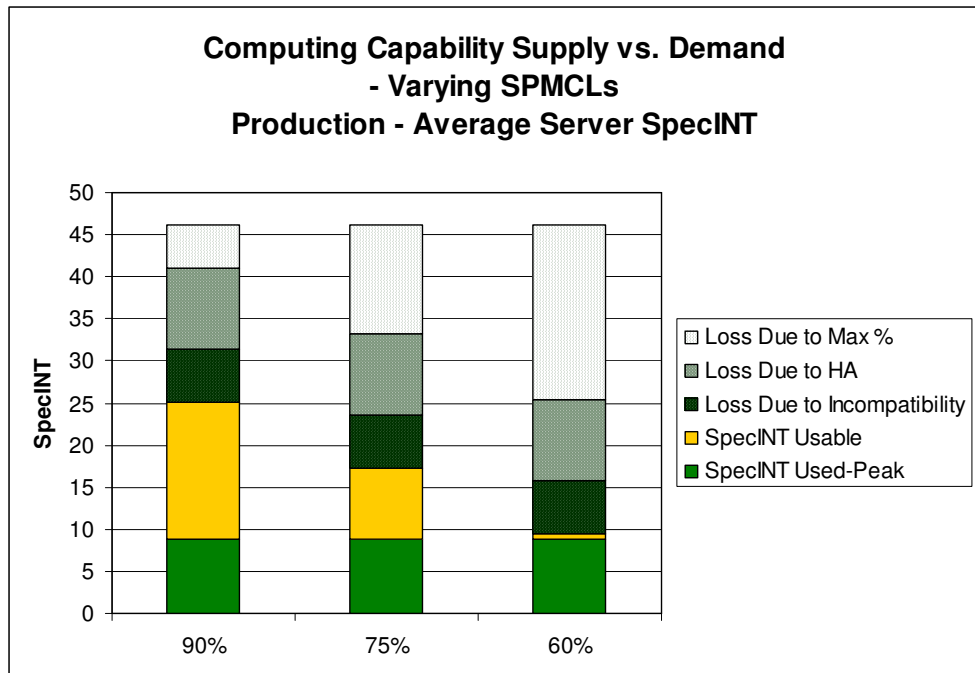


Figure 2 – Usable Power – Varying SPMCLs

Following is a detailed description of these factors.

SPMCL – As a general rule of thumb, in our enterprise, midrange servers should not be required to sustain CPU utilizations exceeding 90% for long periods of time (over one hour). Your requirements may be different. With those differing requirements in mind, Figure 2 illustrates the

Practical SpecINT Usable for an Average Server based on different definitions of acceptable SPMCLs; 90%, 75% and 60%. The “Loss Due to Max %” sections in the graph represent the unusable computing resources, based on different SPMCLs, before taking into consideration any other factors. The workloads that you run should be taken into consideration as you determine the “right” SPMCL(s) for your organization.

High Availability (including Clusters and Load Balancing) - One of the primary explanations for the overall increase in computing power in our environment is our requirement for High Availability solutions. This increase in computing power due to HA also contributes significantly to the wide CPU utilization gap. As you examine Figure 2 above, notice the “Loss Due to HA” values. These values represent, in large measure, the computing resources on hand to address potential system failures. By design, unless system failures occur, these resources will never be utilized. And when they are used, the failed resources will not be used, which means they will never be used at the same time, and what is not used appears to be wasted from a capacity perspective.

Incompatibility – The final major factor contributing to the decrease in usable resources is called incompatibility (see “Loss Due to Incompatibility” in Figure 2). Generally speaking, this means that there are some applications that can not share computing resources with other applications due to incompatible operating systems, middleware versions, etc., or due to an application’s inability to coexist (play nice) with other applications. These situations often result in computing resources being available, but going unused.

Taking the points listed above into consideration, Figure 3 shows more accurately the practical amount of usable midrange computing power of the production environments.

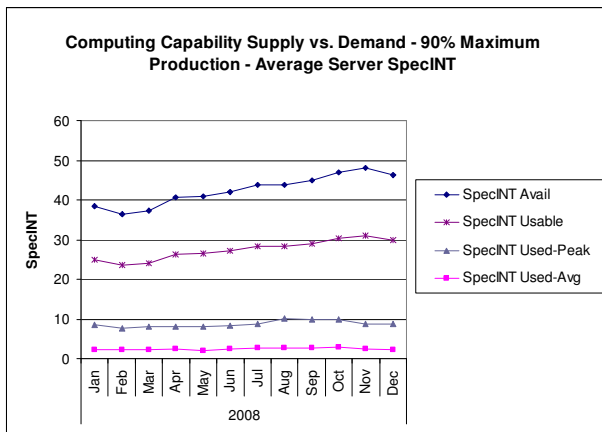


Figure 3 – Computing Power Supply & Demand – Usable

Quantifying the Key Factors

You may be thinking to yourself, “The above concepts and graphs seem to make some sense, but where did the numbers come from?”

As we examined our environment we were able to identify the systems utilizing HA solutions, and based on our familiarity with the systems along with feedback from other IT specialists in the enterprise, we were able to estimate the number of servers impacted by incompatibilities (servers ineligible for consolidation). The spreadsheet in Appendix A contains the formulas we used to quantify the impact of these two key factors. A detailed description of the spreadsheet follows.

The spreadsheet is divided into three basic sections:

- High Availability
- Incompatibility
- Totals and Summaries.

The High Availability section includes two subtypes; “Clusters” and “Load Balanced Groups”. The Incompatibility section also includes two subtypes; “Operating Systems, Middleware, etc.”, and “Don’t Play Nice.” Within each of the subtypes there are additional rows representing specific configurations and the subsequent columns show the following values:

- *Always Unused % SpecINT*
- *SPMCL*
- *Practical Peak CPU %*
- *Unusable Available % SpecINT*
- *% of Production Enterprise*

Always Unused % SpecINT – This is the percentage of computing power that is not usable because of the HA design, or a form of incompatibility. For HA environments, the assumption is that in the case of a server failure, the remaining servers in the cluster or load balanced group need to serve 100% of the workload, including work usually handled by the failed server. The values for the Incompatibility rows are estimates based on observed unused computing power of a sampling of the servers running incompatible applications.

SPMCL – This represents the CPU utilization level that your organization believes midrange servers should not exceed for sustained periods of time (over one hour). You may have different SPMCL levels for differing server workload types (for example, your web server thresholds may be different than your batch server thresholds). In that case, you may add additional rows with unique

SPMCL values for servers with different workload types.

Practical Peak CPU % - The actual amount of usable computing power after considering HA and Incompatibility as well as the *SPMCL*.

Unusable Available % SpecINT – The cost (or amount of “waste”) attributable to the three key factors.

% of Production Enterprise – This represents the portion of the production enterprise affected by this type of HA, or Incompatibility. This value is estimated, however, it is based on information gathered in our enterprise regarding the key factor types. We used server lists containing cluster information, load balancer configuration lists, our CMDB, and first-hand knowledge of application incompatibility issues within our IT organization to determine these values. You may have additional sources available to determine even more accurate estimates.

Using the values in the *Unusable Available % SpecINT* and *% of Production Enterprise* columns along with server count and computing power information in your environment, you can calculate the SpecINT Usable values for your enterprise.

Conclusions

You may be “wasting” less midrange computing power than you think you are. You can more accurately report on the computing power which is actually usable in practice if you are willing to:

- Define your Sustained Peak Maximum CPU Level(s) (SPMCLs)
- Gather some information (if you don’t already have it) about your clustered and load balanced environments
- Identify environments which are significantly under-utilized because the applications running on them can not run in your standard environments or do not “play nice” with other applications

Once you have this information you are ready to calculate the amount of truly usable midrange server computing power in your enterprise.

After you know how much is unusable due to HA designs and incompatibilities, your IT organization will be in a better position to make more informed, cost effective decisions.

You will be able to recommend specific HA designs based on business requirements while taking into consideration the cost of unusable resources associated with that design. You may also be able to move incompatible workloads to less powerful environments (once you learn about the incompatibilities) in an effort to “waste” fewer resources.

Now, instead of telling management that 80% of the midrange resources are being wasted, you can inform them that 60% of those “wasted” resources actually exist to ensure availability and performance.

Appendix A - Unusable midrange computing power calculation spreadsheet

A sample spreadsheet, including formulas, can be included as an attachment to this document.

Unused Computing Power Factors:	Type	Config	Always Unused % SpecINT	Sustained Peak Maximum CPU Level	Practical Peak CPU %	Unusable Available % SpecINT	% of Production Enterprise
High Availability	Clusters Active/Passive	1/1	50%	90%	45%	55%	8%
		2/1	33%	90%	60%	40%	3%
		3/1	25%	90%	68%	33%	8%
		4/1	20%	90%	72%	28%	10%
	Load Balanced Groups	2 member	50%	90%	45%	55%	8%
		3 member	33%	90%	60%	40%	5%
4 member		25%	90%	68%	33%	4%	
Total						46%	
Incompatibility	Operating Systems, Middleware, etc.	Varies	50%	90%	45%	55%	5%
	Don't Play Nice	Varies	75%	90%	23%	78%	12%
Total						17%	
Unused Total							63%
All Other Servers				90%	90%		37%
					# Prod Svrs:		2000

Legend
Estimate
Significant
Key Variable