IBM Technical Brief

SAP® Application Server Consolidation on IBM System z™: Opportunity, Metrics, and Case Studies

Author:
Michael R. Sheets
IBM Americas Solutions Advanced Technical Support

Version: 1.02
Date: March 16, 2009

Filename: SAP_Systemz_Sercon.pdf

TABLE OF CONTENTS

TABLES

FIGURES

EQUATIONS

1 Introduction

SAP customers never run just one SAP system – they run many. Part of the reason is that each production SAP system is part of a "flock" of SAP systems with associated non-production systems. These associated non-production SAP systems are for purposes such as development, test, training, and quality assurance. Customers may also have several of these flocks to service different geographies, lines of business, and functions such as SAP's Business Intelligence Warehouse (BW). More recently, SAP has been introducing "confederations" of production systems. For example, SAP is beginning to require certain infrastructure SAP systems such as Enterprise Portal, Solution Manager, and Master Data Management. Some of SAP's industry solutions, such as Retail, also require an associated BW system for operations. The result of all this is that the total number of SAP systems at each customer is growing over time.

This means that the total number of servers is growing as well – especially for large customers. However, typically these servers, especially the SAP Application Servers, have relatively low average utilization. If the peak utilizations happen at complementary times, these might be excellent candidates for server consolidation.

Over the past few years, the author has collected extensive utilization data from a dozen installations, both IBM and external customers, to investigate and quantify the opportunity for consolidating SAP Application Servers onto IBM's System z. The purpose of this paper is to discuss the opportunity in general, the data collection methodology, the development of appropriate metrics for quantifying consolidation opportunity, the necessary server virtualization characteristics, the results of the analysis, and conclusions.

2 Notices

without notice. Consult your local IBM business contact for information on the product or services available in your area.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

3 Trademarks

SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver and all SAP product and service names mentioned herein are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

AnyNet, CICS, DB2, DFSMSdfp, DFSMShsm, DFSMSrmm, DFSORT, DS6000, DS8000, ESCON, FICON, GDPS, HiperSockets, HyperSwap, IBM, IBM eServer, IBM logo, IMS, InfoPrint, Language Environment, Multiprise, NetView, OMEGAMON, Open Class, OS/390, Parallel Sysplex, PrintWay, RACF, RMF, S/390, Sysplex Timer, System Storage, System z, System z9, SystemPac, Tivoli, TotalStorage, VSE/ESA, VTAM, WebSphere, z/Architecture, z/OS, z/VM, z/VSE, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a trademark of Linus Torvalds in the United States and other countries.

Unicode is a trademark of Unicode, Inc.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other products may be trademarks or registered trademarks of their respective companies.

4 Feedback

Please send comments or suggestions for changes to msheets@us.ibm.com

5 Acknowledgements

The author would like to thank all the installations who provided the utilization data that is the basis of this paper. Further, the author would like to thank Jim Dilley for the opportunity and encouragement to make this study.

6 SAP Application Server Consolidation Opportunity

Most large SAP customers run a typical three-tier configuration for their production systems with multiple SAP Application Servers driving the database server(s). In the case of System z DB Servers, most often these Application Servers are UNIX systems. However, other platforms are supported as well. Typically, these Application Servers were initially similar, limited in number, and relatively easy to manage.

Over time though, many customers have found that their business volume requirements grew and they now have a larger and rather disparate mix of Application Servers. Some customers have not up-leveled their SAP applications for quite some time and are now in the process of planning some significant SAP application upgrades to take advantage of new SAP functions and features. Typically, newer releases of SAP have more function and as a result consume more resources. Overall customers see their number of application servers increasing. Even in cases where the above is not true, customers have a large number of non-production SAP systems that is most often adds to server sprawl. All of this has coincided with an increased customer focus on efficiency in cost, operations management, and reduced down time.

In addition, there have been changes in System z. The delivery early this decade of SAP support for Linux for System z (zLinux) [1] dramatically increased the number of SAP Application Server features and functions that can be run in a zLinux environment. This, in turn, opened up the industry leading System z virtualization and provisioning options, features, and functions such as PR/SM LPAR and z/VM [2] to many more SAP customers. Over time IBM has enhanced these as well as making the zLinux offering more price competitive. Customers now see more value in SAP environments not only for classic server consolidation, but other functions as well – such as using System z and z/VM provisioning to provide disaster recovery functions for the application servers.

However, sizing in a virtualized environment is much more complex than conventional dedicated server sizing. Because of this, and the novelty of running SAP in a virtualized environment, there was no experience or rules of thumb available. Now, based on quantitative data from twelve installations, we have some experience.

In this section, we will discuss sizing in a virtualized environment from a conceptual perspective. This will be done from the perspective of migrating from existing dedicated servers to virtualized servers.

6.1 Simplistic Example

Capacity planning and sizing of existing systems should take into account some kind of peak utilization. While the average utilization might be of interest, it is not prudent to install just enough capacity for average utilization. If one did so, demand would not be met a significant period of time. This would be analogous to an electric utility installing only enough generating capacity for the average load. Customers do not want to get up in the middle of the night to balance their load to meet the average.

However, in a virtualized environment, there is the hope that when one server peaks, others may be idle. That is, there are complementary peaks – and that we can get rid of excess unused idle-time (or “white space”). Thus, we would have to configure only for the peak of the sum of all servers’ utilization.

A simplistic example is shown in the accompanying three figures below. Let us assume the current running SAP system has four identical application servers. “Figure 1”, below, shows the utilization graphs of these servers over the same representative period of time. Note that while each of these Application Servers peaks at a utilization of 9, their utilization profiles are unique.

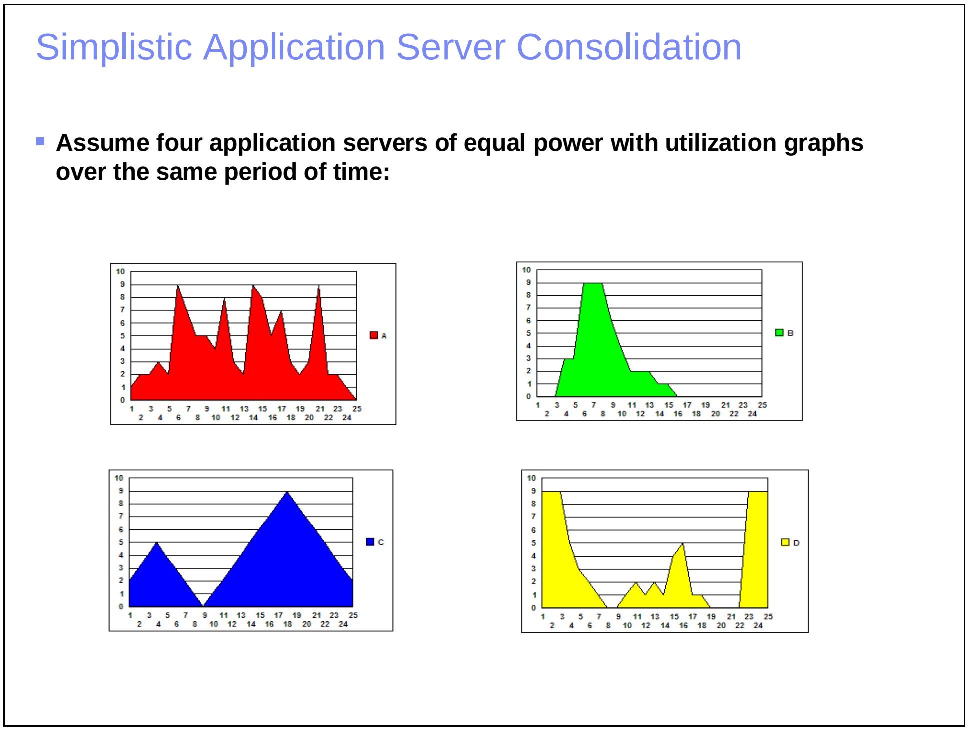

Figure 1: Simplistic Application Server Consolidation: Four Application Servers' Current Utilizations

To consolidate this same work onto one server, we would have to add the utilization of all four servers together for each time interval. This is shown in "Figure 2", below. The graph to the right in that figure shows the aggregated utilization of all four servers.

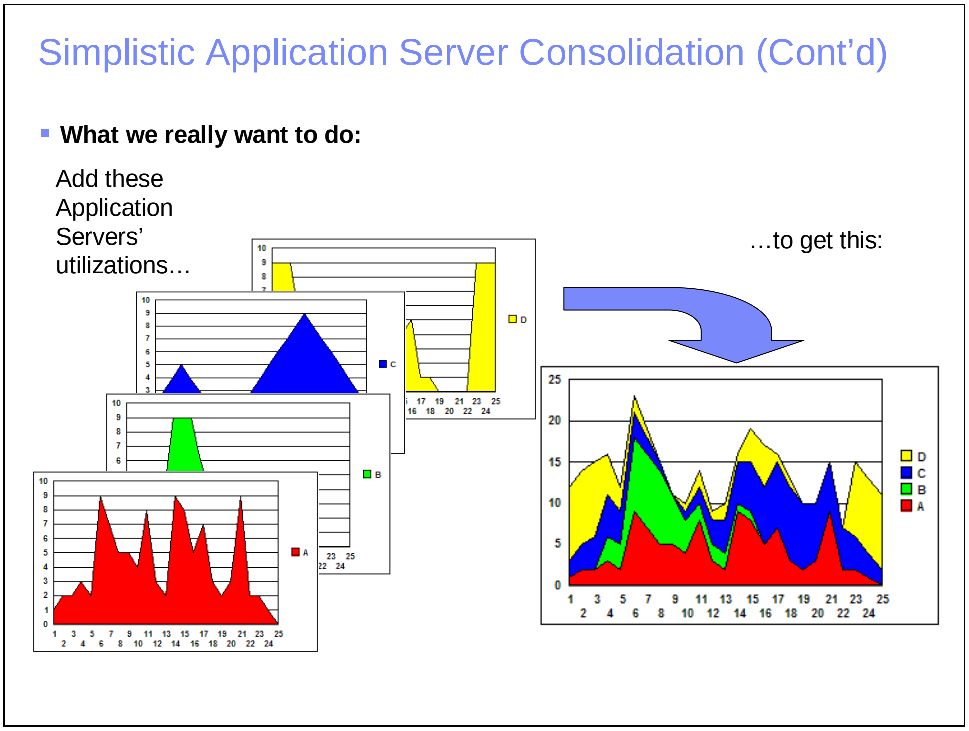

Figure 2: Simplistic Application Server Consolidation - Initial Summing of Utilizations

There is the opportunity to consolidate even further if the consolidating server has advanced workload management and virtualization characteristics **and** if the installation can differentiate their work. For example, if there is some discretionary work, it can be deferred until after the aggregate peak. In other words, the peak can be managed down if it is possible to manage the work to give better service (higher priority) to the “loved” work while deferring the discretionary work. As mentioned above, this requires a very sophisticated set of workload management and virtualization features and functions. System z is particularly strong in this realm as discussed in “System z as an Industry Leading Server Consolidation Platform” on page 13.

This “managed peak” is graphically represented in “Figure 3”, below. The graph to the right accomplishes the same work as the graph on the left. While the peak on the right is somewhat less, that peak is also broader. In terms of integral calculus, the integrals (or the areas) of the two curves are the same. Note that the graph on the right is a single color going from dark to light as the utilization goes up. This is an attempt to graphically indicate two things:

- The single color is to indicate that each application server’s work is broken into very small pieces and blended together – like “mashed potatoes”. Further, it needs to be dispatched by a very dynamic sophisticated scheduler designed to dynamically handle changing workloads.
- The darkness is an attempt to indicate the importance of the work. As you can see, we are still getting the important work done while we are allowing the discretionary work drive to high utilizations.

We will discuss more about both of these, and their implications on server requirements in “System z as an Industry Leading Server Consolidation Platform”, on Page.13

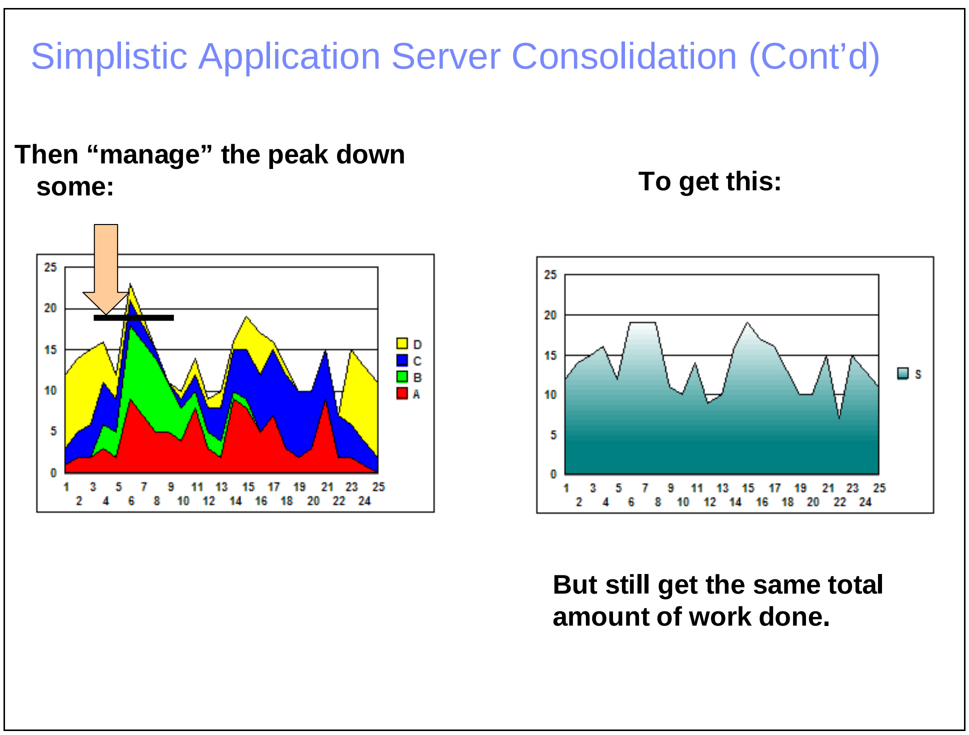

Figure 3: Simplistic Application Server Consolidation - Managed Peak

6.2 Some Metrics and Math

To go into more depth on these topics, it is necessary to define and discuss some metrics. In an attempt to keep this at a high level, we will describe the metrics and their characteristics – but we will offer no mathematical proofs.

The first metric we want to introduce is “Equation 1”, the Sum of the Peaks (SOP).

$$Sum\_of\_Peaks = \sum_{i=1}^{n} Max(Server_i\_Util)$$

Equation 1: Sum of the Peaks

In this equation, we first find the peak utilization of all time intervals for each server. Then we are adding these peaks for all n servers. This metric does reflect some degree of virtualization. Specifically, it shows the required capacity if we configure just enough capacity for each server’s maximum utilization. However, it does not reflect the benefit of advanced virtualization as shown in the earlier figures. It does not take into account when the peaks occurred or, more importantly, if the peaks are complementary. However, SOP is a relatively common and easy to obtain metric. For example, IBM’s Insight for SAP tool [3] reports the peak utilization for each server.

A metric that better represents the characteristics of advanced virtualization is “Equation 2”, the Peak of the Sums (POS).

$$Peak\_of\_Sums = \mathrm{Max}\left(\sum_{i=1}^{n} Server_i\_Util\right)$$

Equation 2: Peak of the Sums

In this equation, for each time interval, we add the utilization of all n servers. Then we find the maximum of that sum. This corresponds to the graph on the right of “Figure 2: Simplistic Application Server Consolidation - Initial Summing of Utilizations” on page 8. POS is more difficult to calculate than SOP because it requires all servers’ utilizations for common time intervals over a lengthy period of time. The practicalities of this will be discussed more in “Practicalities” on page 12. A major purpose of this paper is to show some actual observed relationships between the relatively easy to obtain SOP and the more difficult to obtain POS[1].

To mathematically implement a managed peak as shown in “Figure 3: Simplistic Application Server Consolidation - Managed Peak” on page 8, we use the following equation.

[1] The ratio of POS/SOP is bounded by 1 and 1/n where n is the number of equal capacity servers being considered for consolidation. This ratio is also defined as “Equation 4: Consolidation Factor”.

$$Managed\_Peak = \text{Percentile}\left(\sum_{i=1}^{n} Server_i\_Util\right)$$

Equation 3: Managed Peak of Sums via Percentiles

This equation is a variant of "Equation 2". It reduces the POS by selecting a specific percentile of all the sums. For example, the 50^{th} percentile would be the same as the median of the sums. That is, half of the sums are higher and half are lower.

Finally, we define a Consolidation Factor. This factor defines the fraction of the SOP capacity required after doing server consolidation using advanced virtualization. The lower this number, the greater the server consolidation opportunity.

$$Consolidation\_Factor = Peak\_of\_Sums / Sum\_of\_Peaks$$

Equation 4: Consolidation Factor

This equation also shows the relationship between the more difficult to obtain POS and the relatively easy to obtain SOP. A similar variant of the Consolidation Factor uses the Managed Peak as the numerator. Later in "Discussion" on page 17, we will show that this metric varies from installation to installation. One of the unique characteristics of this metric is that it is not so much dependent on the average utilization of the current servers but rather on the degree that the workload can be smoothed (i.e., complementary peaks).

"Table 1" below, applies these metrics to the simplistic example introduced earlier. The managed peak in this case is 52.8% of the SOP and 47.5% of the currently installed capacity. This indicates we could replace the current four servers with one server with sophisticated advanced virtualization and with only 47.5% of the current capacity. We would expect the average utilization of this workload on this new server to be 42% (= 0.38/0.9) with peaks at 100%. However, remember that this simplistic example uses arbitrary data and therefore is not at all like real installation results. In "Case Studies", on Page 15, we will show results from real installations.

	Capacity	**Ratio to Installed**	**Ratio to Sum of Peaks (Consolidation Factor)**
Installed	40	1.0	
Sum of Peaks (SOP)	36	0.900	1.0
Peak of Sums (POS)	23	0.575	0.639
Managed Peak	19	0.475	0.528
Average (Mean)	13.68	0.342	0.380

Table 1: Simplistic Application Server Consolidation – Quantitative Results

6.3 Practicalities

Discussed below are some practicalities we have observed with this analysis and typical installations.

Intervals:
To get a proper POS, we must have utilization data for all servers over the same period of time. Because we are doing capacity planning based on the observed utilization, this monitoring period of time should be representative of the installation. In addition, since we are implementing a managed peak based on percentiles of observations, we must have a complete cycle of the periodic usage. For example, if the installation has monthly cycles or peaks, then correspondingly, we should have a typical month's monitoring period.

Further, we need a consistent interval for the utilization data of all servers. If we have 15-minute intervals for one server and 5 minute intervals for all the other servers, we have to convert them all to the same 15-minute intervals. The shorter the interval, the "spikier" the data is. In addition, for the same elapsed time, more data must be collected and analyzed. Conversely, longer intervals implicitly "smooth" the utilization with resultant loss of information about the peaks. In general, we recommend intervals in the 15-minute interval range as a reasonable compromise of showing the peaks and minimizing the data collected. Also, intervals in the range of 15 minutes do not have to be perfectly coincident – rounding to the nearest interval is close enough.

Number of Servers:
In general, the larger the number of servers being consolidated into a single virtualized server, the greater the potential for a better (i.e., lower) Consolidation Factor. This is especially true as the number of SAP System IDs (SIDs) goes up. The logic is that the more disparate work you consolidate, the more likely the servers have complementary peaks and the better the Consolidation Factor.

On the other hand, as the number of servers goes up, the more difficult it is to get utilization data for all the servers for the same period of time. Probably the most important aid in this is to have automated tools that collect the utilization data. Nevertheless, even in this case, there can be data gaps in the collection of some servers due to re-boots or other problems. This must be looked into during the analysis and a decision made to include or exclude the data from servers missing data.

Utilization:
It is important that all the CPU utilization be collected and analyzed. Specifically, using standard UNIX/Linux SAR terms [4], this includes both User and System times. However, it excludes Wait and Idle times. The data from the installations shown in "Case Studies", on Page 15, came from a mixture of sources: Insight, NMON, VMSTAT, etc. Windows environments have similar considerations.

There has been some speculation that there are significant fixed CPU costs on UNIX and Windows servers – especially at low utilizations. According to this speculation, the observed utilizations should be discounted somewhat. To date, we have not seen any definitive study supporting this. It is our feeling that there are operating system costs associated with both wait/idle management and application service requests. Further, both of these costs are variable. The so-called "fixed costs" at low utilizations could be significantly reduced at high utilizations. We feel these costs could be different for each specific operating system release. Since we are recommending using a relatively high percentile of managed peak, there is significant utilization and we feel these costs, if they exist, are minimal. As a result, in general, we do not recommend adjusting the observed utilizations.

Managed Peak Percentile:
For many years, there has been "conventional wisdom" that with System z you can capacity plan for 100 percent utilization for the 90^{th} percentile POS using 15 minute or longer intervals. Because the purpose of this paper was to consider consolidation on System z, we also used the 90^{th} percentile. Later in "Discussion" on Page 17, we will show the relationship between the POS (or 100^{th} percentile) and the 90^{th} percentile POS observed in several installations.

Server Weightings:
The forgoing equations and math implicitly assume the existing servers all had the same capacity. In fact, this is rarely the case. Virtually all installations run a mix of servers with different capacities. As a result, it is necessary to weight the servers based on their capacity. This was done using proprietary SAP capacity ratings.

Most Utilized Server:
While evolving this methodology, it became apparent that it was useful to determine the most utilized server. Conventional wisdom is that server consolidation is generally not useful for servers with high utilization. In fact, there was one installation (Installation "E" in "Case Studies" on Page 15) where about half of all the total utilization of eleven different servers was attributable to one server. It may be worthwhile excluding the sub-set of heavily utilized images from server consolidation.

Upgrades:
Often in planning a consolidation, installations will want to upgrade some number of the SAP SIDs being consolidated. Different software or application levels can influence performance. As a result, it is necessary to weight the specific SIDs based on these effects. IBM sizing groups have some guidelines for this.

6.4 System z as an Industry Leading Server Consolidation Platform

All of the above formulae and machinations can be done relatively easily with spreadsheets and other mathematical tools. However, building a server that actually can deliver the characteristics implied by the mathematics is quite difficult. Further, there are other requirements in a virtualized environment we haven't touched on yet. This section intends to address some of these requirements and to demonstrate that System z is uniquely positioned to be the best server platform in a consolidation environment. It does this through a combination of hardware, microcode, PR/SM LPAR, and z/VM

- Security: If you run lots of different work on a system, it is necessary that one set of work not impact other work. For example, you cannot have rogue non-production work bringing down production. Further, you want to make sure one set of work cannot observe or monitor other work. System z's PR/SM LPAR is the only platform partitioner in the world that is certified by the Common Certification Criteria [5] as EAL5 [6]. This effectively means that from a security perspective different partitions are functionally equivalent to separate machines. This is not something new. In the early 1990s a US Government Agency stated this publicly at the Share User Group. Similarly, z/VM provides a good level of protection for work running under each z/VM through its EAL4 certification.

- Isolation: You also want to be able to manage the different kinds of work differently. There are in PR/SM LPAR, z/VM, and even z/OS complementary mechanisms for managing different kinds of work. The main purpose of these mechanisms is to allow the installation to differentiate the work. This should be done at various levels (e.g., SID, operating system image, process/thread, batch job, LPAR, virtual machine, etc.). This is necessary to facilitate managed peaks. The idea is to give excellent service to the most loved work while at the same time using discretionary work to drive to high utilizations.
- Fine granularity of sharing: Resources should be managed in small increments. For example, with advanced virtualization, CPU's or cores should be shared. Further, this sharing must be event-driven. For example, PR/SM LPAR and z/VM know when a subordinate operating system is about to go into wait and can immediately switch to another so there are no lost CPU resources.
- Memory virtualization: As the load dynamically shifts from system to system or Application Server to Application Server, it would be advantageous to also shift memory resources. z/VM provides extensive memory virtualization functions.
- Autonomically responsive: Workloads themselves can change dramatically in the short term. An advanced consolidation platform must be robust enough to automatically adapt its management mechanisms for these changes to continue to provide the desired service goals.
- Dynamic policy changes: Installation requirements and priorities can change. For example, an installation may have different priorities at night than during the day. Similarly, each SAP system may have a different month-end peak time. This requires that the instillation's specification of priorities be changed on the fly, perhaps via an automation tool, without an outage.
- Disaster recovery: Some customers want disaster recovery procedures for SAP Application Servers as well as the SAP DB Server. Advanced management and automation tools such as Tivoli Systems Automation provide full features and functions to help with this.
- Provisioning: With the complexities of modern SAP applications and the rate of business change, there is a constant requirement to add and remove systems, and SIDs. This can be done with Tivoli products as well.

System z environments are uniquely positioned to provide these functions with the depth and maturity of offerings no other platform can match.

7 Case Studies

We were able to obtain detailed SAP Application Server utilization data from twelve different installations. Essentially, we put each of them through the analysis discussed above. However, over time we did refine the analysis process. Always the goal was to quantitatively estimate their server consolidation opportunity and size them for System z using zLinux and z/VM.

7.1 Data and Spreadsheet Analysis

Some of the raw utilization data came from Insight, the rest from tools like NMOM. In general, all the analysis was done with three-dimensional spreadsheets. We used one sheet for each server with each row containing the time stamp and utilization information. A separate summary sheet was used to implement many of the considerations discussed in "Practicalities" on page 12, calculate the SOP and POS, as well as various other statistical calculations. The result was some large spreadsheets – up to 56 MB.

7.2 Results

Listed in "Table 2: Installation Observed Results", below, are the twelve separate installations for which data was collected. To protect their identity, eleven of the installations are identified by letter only. They are listed in date sequence. Each of these is a unique installation. The twelfth installation is one of IBM's largest SAP installations and is identified as "IBM". Note the different Consolidation Factors listed (i.e., 100^{th} percentile through mean). Some of these are plotted graphically for all installations in "Figure 4: Installation Consolidation Factors Using Different Percentiles" on page 17.

As this set of installations was being accumulated over time, the experiences were used as a basis for developing much of what is discussed in "Practicalities", on page 12.

Installation	**A**	**B**	**C**	**D**	**E**	**F**	**G**	**H**	**I**	**J**	**K**	**IBM**
Date	Aug 2004	Sept 2004	Nov 2004	Dec 2004	Feb 2005	Feb 2005	Mar 2005	Mar 2006	Aug 2006	Sept 2006	Jan 2007	Mar 2007
SIDs	1	1	23	1	10	80	12	8	~20	1	48+	9
Appl. Images	4	5	26	7	11	54	52	12	32	6	31	46
Comments	Already running on zLinux		No SOP avail.						All non-prods		Some non-SAP servers too	
Monitoring Period	2 days	11 days	2 days	2 days	4.5 days	13 days	5 weeks	9 days	13.5 days	6 days	13 weeks	4 weeks
Periodic Peaks?	No	Yes	No	Yes	No	No	Yes	Yes	No	No	No	Yes
Mean utilization	0.134	0.198	0.061	0.037	0.190	0.065	0.159	0.242	0.049	0.186	0.064	0.251
Consolidation Factors												
Sum of Peaks	1.000	1.000	NA	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Peak of Sums (100 %tile)	0.659	0.869	NA	0.490	0.575	0.282	0.345	0.668	0.225	0.666	0.287	0.458
99th %tile	0.595	0.728	NA	0.364	0.549	0.251	0.280	0.567	0.194	0.552	0.211	0.418
95th %tile	0.454	0.636	NA	0.321	0.501	0.224	0.255	0.519	0.143	0.472	0.178	0.384
90th %tile	0.410	0.570	NA	0.279	0.470	0.209	0.239	0.489	0.127	0.425	0.161	0.366
85th %tile	0.377	0.510	NA	0.234	0.452	0.198	0.228	0.467	0.119	0.392	0.150	0.352
50th %tile	0.122	0.109	NA	0.132	0.377	0.142	0.181	0.320	0.085	0.248	0.104	0.296
Mean	0.173	0.212	NA	0.142	0.379	0.142	0.179	0.340	0.090	0.253	0.109	0.289

Table 2: Installation Observed Results

8 Discussion

In all cases, we provided a sizing in terms of System z requirements. In some cases, we factored in projected workload increases due to additional business volumes, up leveling of software, and LPAR and z/VM effects. The y-axis in "Figure 5: IBM Installation Aggregated Utilization", below, shows an example of a sizing. The 90th percentile of the POS is about 24 IFLs on an IBM System z9. That is, the current 46 different servers, with 169 processors, can be consolidated onto one IBM System z9 using 24 z9 IFLs with z/VM.

Not all the installations shown ended up consolidating on system z. In fact, in the case of Installation B, IBM recommended not consolidating on System z. This was because of a combination of low Consolidation Factors and high average utilization. In Installation E we determined that about half of the total Application Server utilization came from one of the 11 total application images. Therefore, we recommended consolidating only the other ten images. However, the customer ended up consolidating everything on z. Some installations are still considering their options. Others are in the process of consolidating. So far, however, installations that have consolidated have validated this sizing methodology.

"Figure 4", below, shows various Consolidation Factors for the different installations. This is intended to give a feeling of the distribution skews in each installation behind the tabular data. As stated earlier, our advice is to use the POS or, when there is an extensive period of utilization data, the 90th percentile POS. This graph also allows us to compare the different installations.[2]

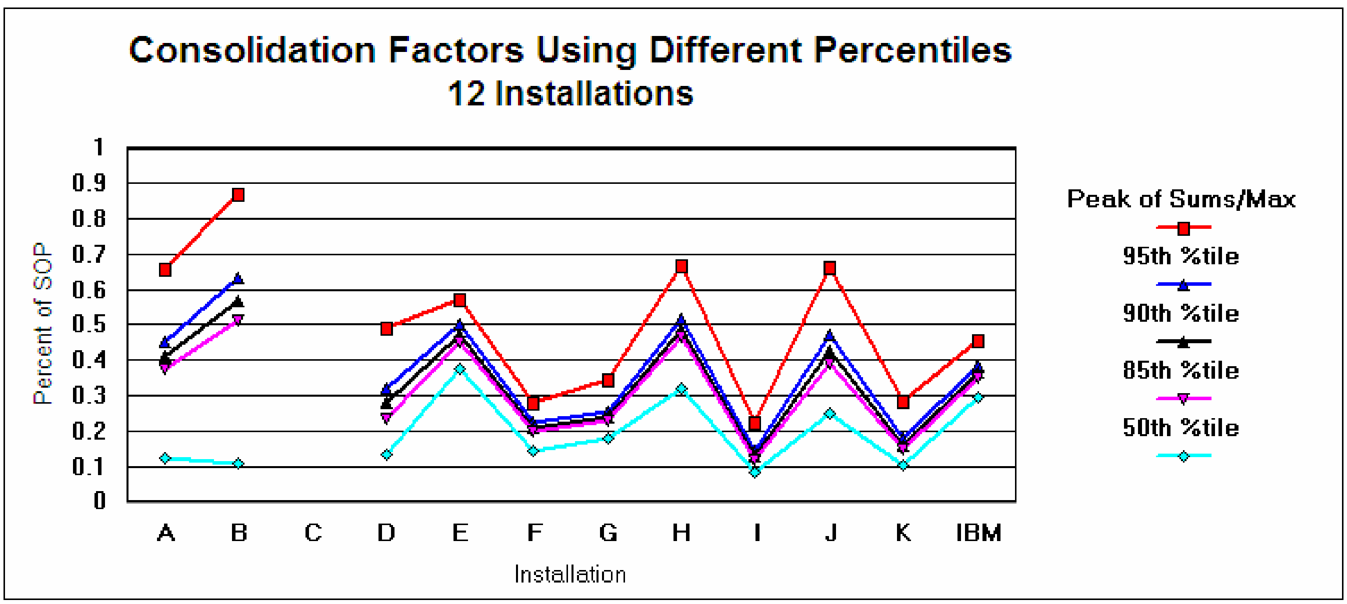

Figure 4: Installation Consolidation Factors Using Different Percentiles

8.1 Search for Gross Correlations

Despite the small sample size and the lack of controlled experiments, we tried to find immediate gross correlations to the Consolidation Factors. The number of application images was first tried – but that

[2] Installation C is not graphed because, as noted in "Table 2" above, no SOP was available.

did not work. Neither the number of SIDs nor the length of observed time seemed to have any bearing either. Next, we tried the mean utilization of the current sets of application servers. While the range was unexpectedly wide (3.7% to 25.1%[3]), there seemed to be no immediate correlation. However, this wide range did seem to indicate that different installations have different "styles" of running application servers. Some seem to make more of an effort to run at higher utilizations.

Finally, we attempted to look at periodic peaks. About 40% of the installations had periodic utilization peaks similar to classic OLTP workweek loads. This is demonstrated in "Figure 5", below[4]. You can clearly see weekend lulls – despite the fact that this represented several independent lines of business production and non-production work. However, the existence of periodic peaks seemed to have no bearing on the Consolidation Factors either.

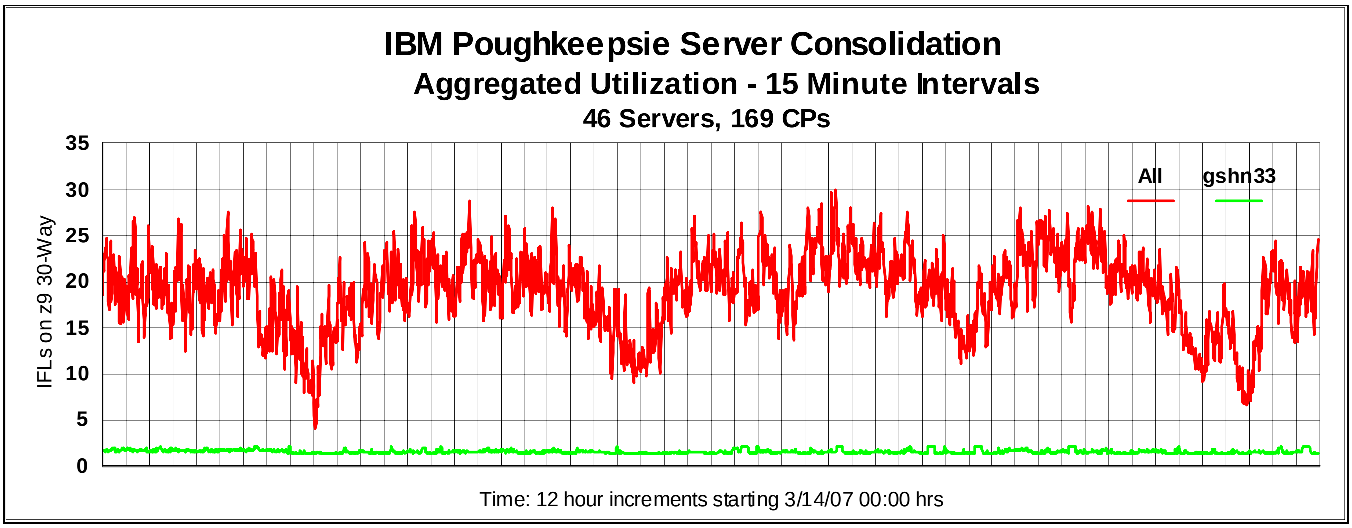

Figure 5: IBM Installation Aggregated Utilization

Finally, conventional wisdom is that non-production work, in addition to being intrinsically more discretionary, can consolidate well because their work is sporadic and unrelated. Installation I, whose analysis consisted of non-production servers only, has the lowest Consolidation Factors of any installation and seems to confirm the convention wisdom.

8.2 Slicing and Dicing Installations

One intuitive feeling is that consolidating all the servers onto one increases the consolidation opportunity when compared to consolidating onto two or more servers. Early on experience with Installation C confirmed this. This installation had two independent business units. Consolidating each business unit separately onto two servers required 18% more resources than combining them onto one server (using the 90th percentile POS). This did not take into account that CPU increment offerings, which could raise this even more.

[3] The IBM installation was surprisingly high at 25.1% – especially considering that it represented 46 servers over a four-week interval.

[4] The green plot at the bottom of the graph is the utilization of the one server (gshn33) that consumed the most CPU resources during this monitoring period.

This was looked at in more detail in Installation K and IBM. In the case of Installation K, we looked at separating production from non-production. From "Table 2: Installation Observed Results" on page 16, we can see that the 90^{th} percentile Consolidation Factor of everything was 0.161. Separately, Production was 0.184 and non-production was 0.179. We found that the 90^{th} percentile consolidation onto two servers required 13% more capacity.

In the case of IBM, we tried two options. First was to separate each of the three lines of business into three servers with production and non-production together for each line of business. Everything together had a 90th percentile Consolidation Factor of 0.366. Separately they were: 0.422, 0.453, and 0.304. Three servers required 9% more resources. Then we tried further separating each line's production from non-production – into six total servers. Further breaking the load into six servers required 18% more capacity.

9 Conclusions

The data so far show a wide range of average utilizations and Consolidation Factors. The results, metrics, and methodologies are of definite value sizing the specific cases where we had data. Further, these results can probably be used for other similar environments at that same Installation if they keep to their "style" of running the Application Servers. However, it is felt that the range of results is too wide to predict for an unknown customer.

The Consolidation Factor metrics ("Equation 4: Consolidation Factor" on Page 11) proved to be particularly powerful in the System z environment.

Finally, the data seem to confirm the conventional wisdom that consolidation into a single server saves resources.

10 References

[1] IBM Corp. 2002. *SAP on DB2 UDB forOS/390 and z/OS: Implementing Application Servers on Linux for zSeries* (SG24-6847) or from http://www.ibm.com/redbooks

[2] z/VM Information: http://www.vm.ibm.com/?

[3] Insight: http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS381

[4] AIX SAR Command: http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.cmds/doc/aixcmds5/sar.htm

[5] Common Criteria: http://en.wikipedia.org/wiki/Common_Criteria ,

[6] Evaluation Assurance Level (EAL): http://en.wikipedia.org/wiki/Evaluation_Assurance_Level