IBM @server

SAP Solutions on IBM eServer iSeries ...
a virtualization and consolidation study

A real-world SAP business scenario demonstrating the power of server virtualization and application consolidation on the **IBM eServer i5 Solution Edition for *mySAP*^{TM} *ERP*.**

IBM eServer i5 Solution Edition for *mySAP*^{TM} *ERP*

Executive Overview

As the cost of managing systems continues to increase, consolidation of systems and applications is becoming an important consideration for many businesses, especially those that use SAP applications. Successful consolidation requires a server that can handle multiple operating systems and application environments simultaneously and affordably.

With the announcement of the new POWER5™ processor-based IBM eServer® i5 servers you can run applications supported across multiple operating systems/environments: IBM i5/OS® , IBM AIX 5L™ , 64-bit Linux® (Linux on POWER), Microsoft® Windows® and 32-bit Linux® (Integrated xSeries Adapters or Integrated xSeries Servers are required for Windows and 32-bit Linux).

The IBM eServer i5 Solution Edition for *mySAP ERP* is specifically configured and priced to provide a platform that can be used for the consolidation of systems and applications, reducing overall systems management costs and saving money. This study describes how IBM used logical partitioning on the IBM eServer i5 Solution Edition for *mySAP ERP* to demonstrate the benefits of application consolidation and server virtualization. Using logical partitioning, multiple SAP and other applications can run simultaneously in their own separate partitions on the same server footprint. As your workloads change over time, server resources can easily be shifted as needed to other partitions, helping to ensure you get the most business value possible by virtualizing your server resources.

Multiple mySAP Business Suite components are common in any SAP environment.   Customers usually have Development, Test and Production systems for each mySAP component, resulting in multiple SAP systems across the landscape.   On most other hardware platforms, customers usually put one SAP system on each footprint. Each of these servers must be sized such that they each can handle the peak workload for that system, and it also means that large amounts of computing resource will usually be underutilized. Customers may be faced with even more complexities than this in the SAP environment. Some of the SAP systems may have priority over others, and some may be more subject to unpredictable workloads than others. Accounting for these factors can make it difficult to decide how much computing resource to allocate to each partition or may even make the customer feel a need to oversize servers.   Wouldn't it be nice if your server could accommodate all of your SAP systems in a single footprint, and move resources to the systems that need them, when they need them?

By leveraging logical partitioning, iSeries customers historically have required fewer servers to support their SAP environment. Logical partitioning allows iSeries customers to separate different SAP applications from each other, or to separate Production from Development and Test systems. Although a consolidated server is typically larger, it reduces the overall number of server footprints in the environment and has clear advantages when it comes to sharing hardware

2

resources across the partitions.   Of course, the iSeries still stands out as a platform capable of running multiple SAP applications within a single partition due to the superior work management capabilities of the i5/OS. This special ability is by no means diminished when discussing the virtualization capabilities of the iSeries.

Server management has also been simplified through hardware virtualization using POWER Hypervisor™ firmware.   This is the technology that makes it possible to dynamically share resources between partitions. This capability can help to simplify resource allocation and server sizing decisions for iSeries customers. New IBM Virtualization Engine™ Console software presents a common monitoring interface for heterogeneous environments, easing consolidation of multiple operating systems and thus streamlining operations. In addition, alerts can be set at various thresholds to inform system administrators or users when a problem occurs.

In this whitepaper, we describe how IBM used new virtualization features of the eServer i5 server to maximize the usage of processor resources when running multiple SAP applications. The team also included workload driver and monitoring applications to complete the test suite.   The results of this study using an IBM eServer i5 Solution Edition for *mySAP ERP* shows how an eServer i5 server can dynamically share processors across partitions, allowing the server to handle a sudden increase in application workload in one partition by harnessing underutilized and available processors in another partition without operator intervention.

In summary, the results of this study demonstrate how the IBM eServer i5 Solution Edition for *mySAP ERP* can provide excellent server and application consolidation and virtualization opportunities to iSeries customers.

Introduction and Technical Overview

The POWER5 architecture provides new server consolidation, performance and virtualization opportunities for customers using eServer i5 servers. These include greater processor performance and the ability to run AIX 5L and Linux operating systems in separate partitions side by side with IBM i5/OS, allowing a broader range of applications to run on IBM eServer i5 servers.

IBM POWER Hypervisor firmware virtualizes a physical system so that each application appears to run on its own logical system. The following pages demonstrate how SAP and other applications can run across partitions on a single IBM eServer i5 Solution Edition for *mySAP ERP*, Model 550, without compromising performance or manageability. Also, by using the shared processor pool, processor resources can be balanced across partitions to minimize the impact of sudden increases in application workload.

In addition, we will show how the IBM Virtualization Engine Console software creates a common interface to monitor disparate workloads, taking advantage of existing products such as Management Central and IBM Director Multiplatform Agent.

Server Configuration

In our test environment on the Model 550, we created three logical partitions.
- Partition 1 - SAP Business Warehouse(BW) 3.5 on i5/OS V5R3M0
- Partition 2 - SAP R/3 Enterprise 4.70 on i5/OS V5R3M0
- Partition 3 - Client driver on Linux, hosted from Partition 2

The first and second partitions contained i5/OS and DB2, each running a SAP application in a 2-tier configuration.   The third partition contained Linux, running SAP benchmarking tools to simulate client activity for the R/3 and BW workloads.

Linux was hosted from the second partition using Virtual I/O to share the eServer i5 disk storage with the workload running on that partition.   The client simulation tools use minimal disk

resources, so the fact that the same disks were hosting Linux as well as servicing the R/3 workload had a negligible effect on the test results.

Partitions 1 and 2 used direct-attach disk drives formatted to use RAID-5, giving customers maximum flexibility in setting up their partitions. In this study, all disks were connected in RAID-5 sets of either 10 or 15 disks.

All partitions were connected to a Virtual Ethernet that appeared to the partitions as a 1Gbps Ethernet adapter. The Virtual Ethernet efficiently used the system bus to connect partitions, achieving a maximum throughput of approximately 450MB/sec.   This fast Virtual Ethernet throughput between partitions has been particularly useful for SAP on iSeries customers since it is common to transfer huge amounts of data from the R/3 system to BW.   All partitions were also connected through 100MB Ethernet adapters to an external network for administration, configuration, monitoring and control.

The Model 550 used in this study was initially configured with dedicated processors as shown in Figure 1.

SAP BW 3.5	SAP R/3 Enterprise 4.70	Client simulation tools
i5/OS ⌐DB2	i5/OS ⌐DB2	Linux
1 CPU, 8GB memory	2 CPUs, 8GB memory	1 CPU, 4GB memory
89 35GB disks	53 35GB disks (Linux hosted from shared i5/OS disk)	
POWER Hypervisor		
4-way POWER5 16GB memory, 142 35GB disk arms (4.9 TB)		

Figure 1. Logical Partitioning Schematic

Workloads

SAP R/3 Enterprise 4.70 application software is used by many diverse companies. The application workload used for this study was the Sales and Distribution script commonly used for SAP benchmarking. This script included order, delivery, posting and invoicing transactions based on the typical mix used at real customer locations. This workload on an i5/OS partition was driven by SAP tools and scripts running on the Linux partition. The tools simulated SAP GUI users and recorded average response times and throughput metrics, as well as processor utilization.

SAP BW 3.5 is a data warehousing application that provides decision-support tools used for query, reporting and multidimensional analysis. The primary BW workload used for this study was the multi-user query phase of the standard BW benchmark. This workload on an i5/OS partition was also driven by tools and scripts running on the Linux partition. The tools simulated

SAP BW users and recorded average response times and throughput metrics, as well as processor utilization.

In a later part of the study, a second workload was added to the BW system. This workload was taken from the data load phase of the standard BW benchmark. Here, an SAP batch work process was used to load an additional set of data into the BW system.

Processors and partitions

The IBM Hardware Management Console is used to configure partitions. Partitions can use one of two modes of processors on the iSeries server. A partition can use dedicated processors or it can use the shared processor pool.   Both modes may be active on a server at the same time, with some partitions using the shared processor pool while other partitions use dedicated processors.

Dedicated processors allows you to assign a whole processor to a partition. You can assign dedicated processors to a partition as long as processors are available or unassigned. Figure 2 shows how we partitioned our Model 550.
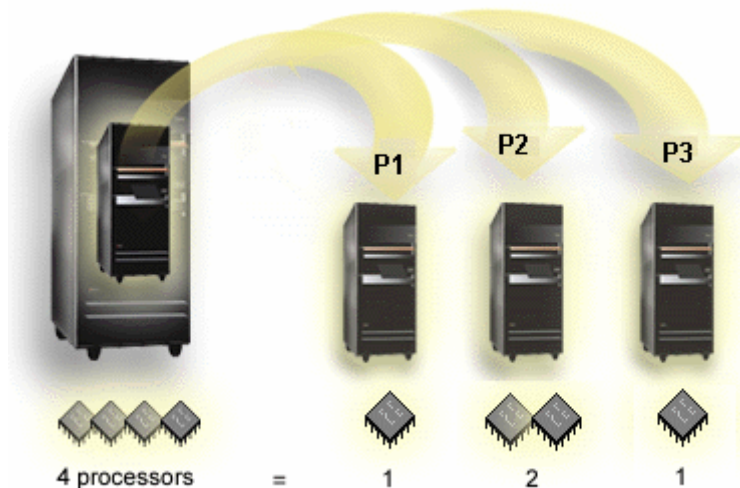

Figure 2.   Partitioning with dedicated processors.

The shared processing pool allows you to assign whole or partial processors to a partition. The physical processors are held in the shared processing pool and are shared among the partitions. Virtual processors are the whole number of concurrent operations which the operating system can utilize. The processing power can be conceptualized as being spread equally across these virtual processors. Selecting the optimal number of virtual processors is dependent on the workload in the partition: some benefit from greater concurrence, others require greater power.

For example, our eServer i5 server with 4 processors in the shared pool provides 4.0 processing units. In one of our tests, the three logical partitions were set to distribute the processing power in the following way: partition 1 has 1.0 processing units and 3 virtual processors, partition 2 has 2.0 processing units and 2 virtual processors, and partition 3 has 1.0 processing units and 1 virtual processor. The sum of the 3 logical partition's processing units is equal to the total number of processing units in the shared pool. The total number of virtual processors is equal to or greater than the number of processing units.   Figure 3 shows this example configuration.
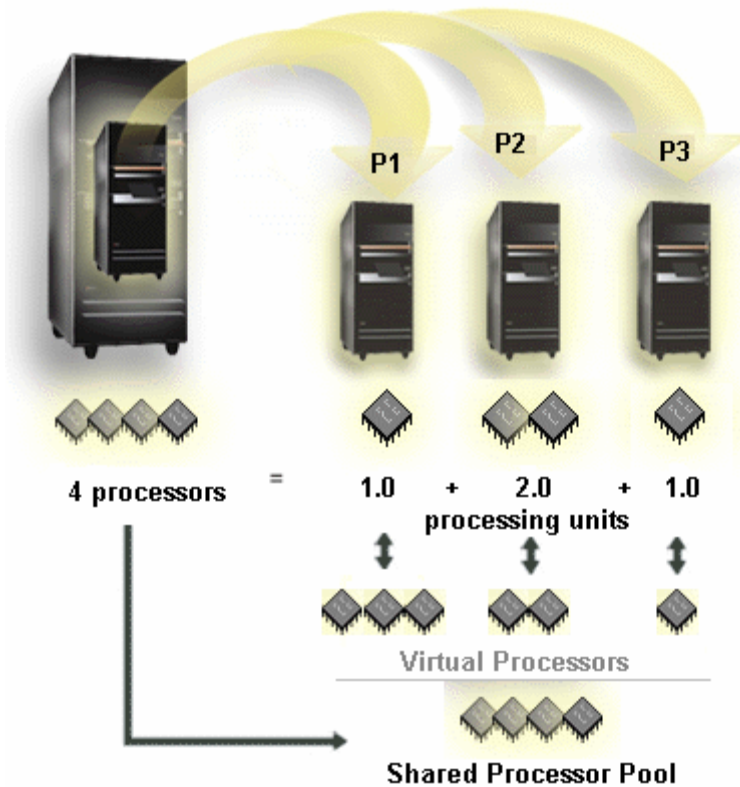
Figure 3. Partitioning using the Shared Processor Pool and Virtual Processors

Test Procedure

Comprehensively testing the dynamic partitioning (LPAR) capabilities of the IBM eServer i5 Solution Edition for *mySAP ERP*, Model 550, required several steps:

I.   The IBM team established workload baselines with partitions using dedicated processors. Each workload was run individually, and then all of the workloads were run simultaneously.

II. The team changed the partition definition to use the shared processor pool across all the partitions. Several tests were performed, each with a different number of virtual processors for the BW partition, and the results were compared the results to those of step I.

III. The team significantly increased the workload on one partition by introducing an additional batch job. The batch job was run with a dedicated processor to get a baseline. Next, it was run simultaneously with the other workloads. Then it was tested simultaneously with the other workloads using the shared pool and different numbers of virtual processors.

The benchmark client tools provided a constant simulated client transaction load to both the R/3 and BW systems. The tests were constructed so as to maintain a high use interval of approximately fifteen minutes. Using accompanying monitoring tools, the number of transactions achieved and the average response time within a specific time interval could be observed.   The batch job used in part III was constructed such that the load would also run in approximately this same period of time, and would be started at the beginning of the high use interval during each run.

6

I. Results for tests running with dedicated processor partitions

The R/3 and BW workloads were tuned to drive the processor utilization of each of those partitions to between sixty and seventy percent. The client simulation tools used to drive the R/3 and BW workloads consumed fewer than ten percent of the processor allocated to the Linux partition.

The IBM team established individual baselines for the R/3 and BW application workloads with dedicated processor resources. The team then ran the R/3 and BW application workloads simultaneously, still using dedicated processor resources.

The team found that the results differ minimally when running one at a time as compared to running them simultaneously, confirming that logical partitioning in its simplest form does not introduce significant performance degradation.

II. Results for tests using the shared processor pool

As a simply stated concept, a partition using virtual processors and the shared processor pool can provide a performance benefit because it is allowed to make more concurrent processing requests than usual.   When using dedicated processors, a partition would only make same number of concurrent requests as could be serviced by those real processors.   When those real processors are not busy 100 percent of the time, there are opportunities for improvement through virtualization.

The shared processing pool allows you to assign whole or partial processors to a partition. The physical processors are held in the shared processing pool and are shared among the logical partitions. Each physical processor provides one processing unit of resource to the shared pool. When more virtual processors than physical processors are configured, more concurrent requests can be made and serviced across the partitions.

Using the IBM Hardware Management Console to configure these partitions to all use shared processors, the R/3 partition was initially given two processing units, the BW partition one processing unit, and the Linux partition one processing unit. The partitions were initially assigned the same number of virtual processors as processing units.   The BW partition was assigned additional virtual processors over successive trials. To easily observe the effects of resource sharing, the BW partition was uncapped, allowing it to accept resources from the other partitions; while the R/3 and Linux partitions remained capped, able to share resource without being able to accept it.
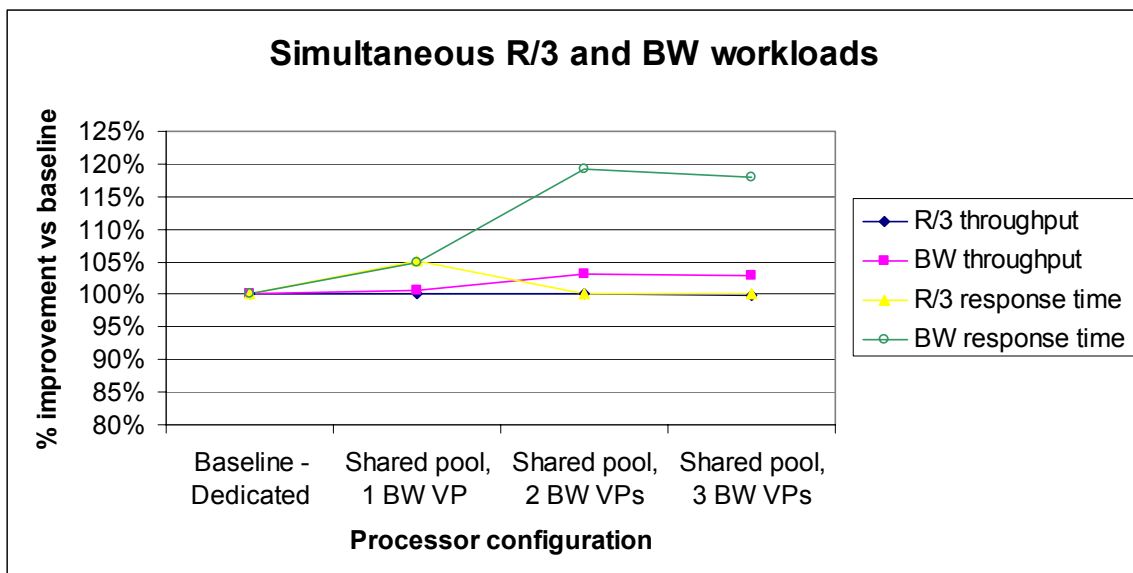

Simultaneous R/3 and BW workloads

Figure 4. Comparison of partitions using dedicated processors vs. shared processor pools

As shown in Figure 4, the average response time in the BW partition improved significantly when shared pool configurations with two and three virtual processors were used. With one virtual processor there was only a small improvement from the dedicated resource results because the OS in the BW partition was not making any more requests to the Hypervisor than before.   More virtual processors caused the OS in the BW partition to make more concurrent requests to the Hypervisor than before. Although the entitlement in processing units had not changed, these additional requests were responsible for the increase in performance because there were now four pool processors at any given point in time that might be able to service the extra requests instead of just the one that was available in a dedicated resource environment.

The R/3 partition throughput metrics remained similar to those measured using dedicated processors, but the response times were down about five percent. This is due to the fact that the R/3 partition kept its two processing units, but was now sharing processors with the other two partitions. At some points in time, requests from the BW and Linux partitions could be occupying all of the shared processors due to the increased BW virtual processors, so a request from the R/3 partition might be made to wait slightly.   But overall, the R/3 partition throughput remained constant because the allocated processing units were constant.

Although not shown in Figure 4, a nominal increase in processor utilization in each partition was noted when using the shared pool, possibly due to minor level three cache contentions between processors. This increase is more than compensated for by the overall increase in real utilization of the server.

In summary, we demonstrated in this section that using the shared processor pool and virtual processors can improve the overall server utilization across partitions.

III. Results for tests after introducing an additional workload

To investigate additional benefits of shared resource partitions, we introduced a BW data load running in a batch work process to cause a sudden increase in the processor resource demands within the BW partition. This additional workload was designed so that when it was run simultaneously with the BW query workload used in the previous tests, together they would demand more processor resources than the BW partition could provide with only a single dedicated processor.

Normally one would not schedule a BW data load during the day while people are running BW reports. However, imagine that the CEO requests a new report that requires inclusion of more up to date information than is already in BW. A new BW data load would have to happen while normal BW business is going on. In some cases this might mean that all BW work would stop so the load and report could go on, but that is not necessary because of dynamic resource sharing. Of course, even the CEO wouldn't want the R/3 production work to be significantly impacted to run his report, so we will make sure that the R/3 throughput doesn't suffer, but dynamic resource sharing can ensure that we use as much resource as possible from the other partitions to assist the overcommitted BW partition.

The BW data load ran in a SAP batch work process job that had a lower priority as compared to dialog work process jobs servicing the BW query clients. The BW data load took eleven minutes and about sixty percent of the processor when running without any other workload in the BW partition.   The BW query load from the prior tests also consumed about sixty percent of the processor when running alone.

With both of these BW workloads running with only one dedicated processor, the duration of the batch job was expected to increase, and the response time of the queries was also expected to rise. Efficiencies could be realized when the test was run in the shared pool configuration, with increased numbers of virtual processors to consume underutilized resources from the shared pool.   Figure 5 shows the effects of adding the additional batch job to the BW partition.
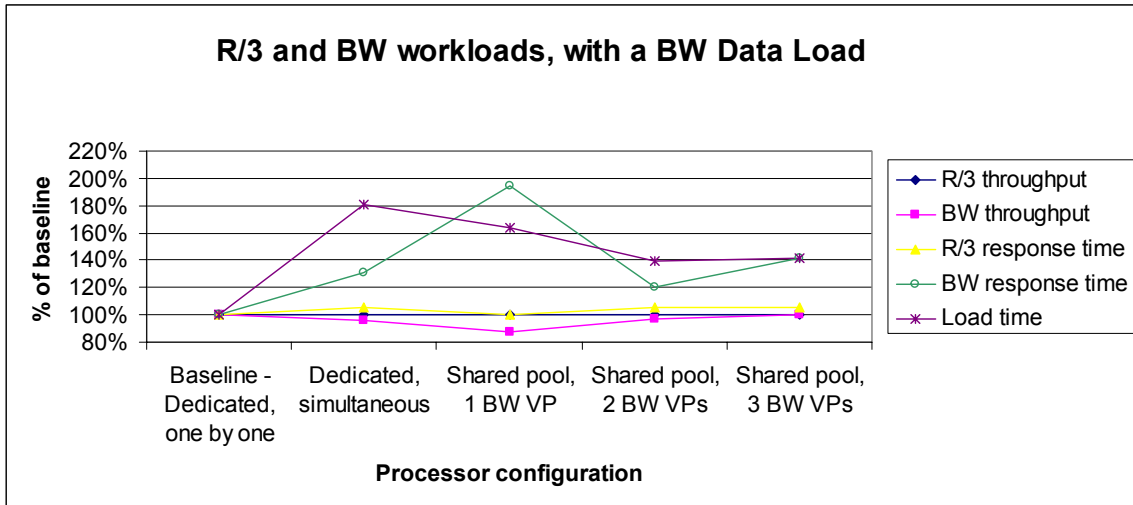
R/3 and BW workloads, with a BW Data Load

Legend:
- R/3 throughput
- BW throughput
- R/3 response time
- BW response time
- Load time

Y-axis: % of baseline (80%, 100%, 120%, 140%, 160%, 180%, 200%, 220%)

X-axis (Processor configuration): Baseline - Dedicated, one by one; Dedicated, simultaneous; Shared pool, 1 BW VP; Shared pool, 2 BW VPs; Shared pool, 3 BW VPs

Figure 5. Comparison of partitions using dedicated processors vs. shared processor pools with additional workload

We initially configured the SAP BW partition with one dedicated processor, and ran each workload individually to establish a baseline.

Next, we ran the workloads simultaneously with dedicated processors in each partition. The SAP work process jobs that handled "interactive" query requests were set to run at a higher priority than the batch job doing the data load. The increase in BW data load time between the baseline measurement and the second measurement shows that this job clearly was unable to perform as well as it had when there was no competing query workload. As expected, the BW response time increased compared to when there was no batch workload, but overall throughput remained nearly the same because of the job priority settings.

Next, we changed all the partitions from using dedicated processors to using the shared processor pool. With the shared pool, the time for the BW data load was reduced significantly. There was almost no disruption to the response time and throughput in the R/3 partition in all cases.

The data also shows that although using one virtual processor under these conditions did help the batch workload, it was disruptive for the BW query workload response time. Since two and three virtual processors yielded better results, it is clear that the shared pool of four processors had the resources to service more requests. The cause of the disruption when using only one virtual processor was that the BW partition was simply not making enough concurrent requests to take advantage of the available processor resource in the pool. This time though, there were far more processing demands on the partition than before because of the batch job. One must take some care in using the shared processor pool to make sure the benefits are realized.

It was possible for our team to estimate how many virtual processors would be most likely to produce the best results by using the performance data from the tests performed with dedicated processors. The R/3 system workload generated about seventy percent utilization of two dedicated processors. The SAP client workload driver generated almost ten percent utilization of one processor. So, of 4.0 processor units available in the shared pool on the 4-way Model 550, 1.0 were already entitled to the BW partition, about 1.4 units (0.7 times 2) would be used by the R/3 partition to run the R/3 workload, about 0.1 would be used by the SAP client driver in the Linux partition, an additional 1.5 units were potentially available to the uncapped BW partition. This can be thought of as an additional 1.5 processors that could be dispatched to by the i5/OS running on the BW partition. From this calculation, we were able to estimate that two or three virtual processors would give the most benefit; the one BW was already entitled to, plus another

one or two to use up the 1.5 units of the shared pool capacity.   The tests confirmed our estimates.

While one test showed that three virtual processors was an acceptable choice, it also showed a slight increase in BW response time compared to two virtual processors. If the number of processing units allocated to a partition remains constant, the length of time each request can be serviced must decrease as the number of virtual processors increases. While the BW partition made more concurrent requests to the Hypervisor because of the additional virtual processor, each request could be serviced for a shorter time because a constant number of processing units were still allocated to that partition. It is important to also consider the number of processing units allocated to a partition when increasing the number of virtual processors in it.

In this study, the processor utilization rates of over 60 percent are higher than usual for normal customer situations in the field, so it is possible that customers may see benefits by using more virtual processors than shown here.

The cause of a workload peak in a partition could be caused by many things in an SAP environment. If it was an R/3 development or test system, it could be something as simple as a large report recompile or importing a transport.   By introducing a sudden increase in a partition workload, we demonstrated how underutilized resources from the shared pool can be used to smooth out workload spikes and improve overall performance.

Virtualization Engine Console

Business intelligence workloads like BW have a long history of producing fluctuating demands on servers. Nightly loads may place predictable demands on the system, but daytime use by business users can vary widely from one day to the next or even from one hour to the next. Also, business intelligence workloads have markedly different tuning parameters than most transaction-based applications. Because of these factors—widely varied utilization and unique tuning characteristics—business intelligence workloads typically run on dedicated, stand-alone servers.

With the POWER Hypervisor, the eServer i5 provides makes server consolidation between unlike workloads possible. Logical partitioning provides a way to isolate workloads with specific tuning parameters from workloads with unpredictable fluctuations.

But the challenge of server consolidation very quickly becomes systems management. This is compounded when the servers run different operating systems. Virtualization Engine Console software is an IBM Licensed Program Product (LPP) that provides a single view of all the OS's and applications supported in this benchmark. This product played a key role in monitoring the three logical partitions under test and generated performance-based alerts when exceptions occurred.

The Virtualization Engine Console is integrated with existing systems management tools, like iSeries Access for Windows. Administrators create and start monitors on the systems management tools they are most comfortable with, and then the Console provides a single view for all the started monitors. For this benchmark, the team used Management Central, which is part of eServer iSeries Access for Windows, and IBM Director Multiplatform Agent, which is part of IBM Virtualization Engine systems technologies.

In traditional iSeries customer sites, Management Central is often used to manage single and multiple i5/OS partitions. IBM Director, which is the systems management tools for xSeries servers, helps customers who run Linux or Windows manage their different environments. We also could have used Web SM, which is the systems management tool for AIX and Linux. All of these tools integrate with the Virtualization Engine Console.

IBM Virtualization Engine Console software used a standardized WebSphere portal to present information to the user. Figure 6 shows a sample screen capture.
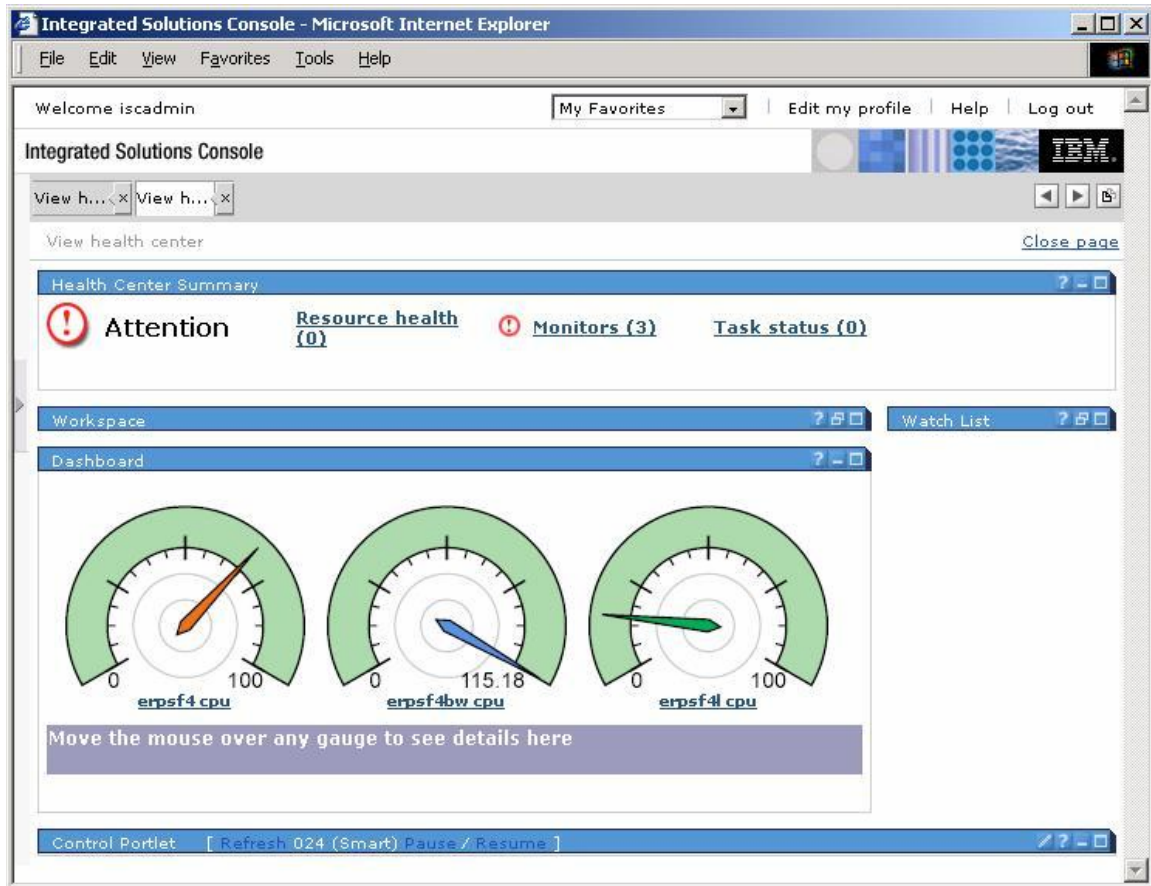
Figure 6. Virtualization Engine Console Sample Screen Capture

In the Console, three processor monitors were running, one for each partition on our Model 550. Even though these monitors were created on different operating systems from different tools, they all display the same way on the Console, simplifying the systems management of the Model 550.

As shown in the example, administrators can use the console to create alerts that warn administrators if a threshold has been exceeded, and to take a default action if no administrator responds in a timely fashion. In this example, an alert was triggered when the BW partition started using more than 100% of its entitled processor allocation. This occurred when we purposely added the batch job we described earlier so that additional processor resources would be used. In addition to processor or disk utilization monitoring at the partition level, the Console can also monitor specific applications for error conditions.

IBM Virtualization Engine Console provides a critical, cross system view of partitions in a server consolidation environment.

Conclusions

This study achieved four major goals:

1. Demonstrated the capabilities of the eServer i5 server as a consolidation platform for multiple applications and operating systems
2. Showed how hardware and software applications can be virtualized to make a single server appear as though it were multiple servers

11

3. Demonstrated management of heterogeneous operating systems and application environments
4. Demonstrated the ability of the eServer i5 server to dynamically adjust virtualized resources to maximize physical resource utilization across the server.

In summary, the results of this study demonstrate how the eServer i5 Solution Edition *for mySAP ERP* can provide excellent server and application consolidation and virtualization opportunities to iSeries customers.

Other References

As another demonstration of the versatility of these new server and software products, IBM has also developed a new benchmark—the Virtualization Grand Slam—using popular commercial applications such as PeopleSoft Enterprise One, PeopleSoft Enterprise Learning Management, Trade3, Linux file serving (Samba) and a Linux firewall.   That benchmark is documented at http://www.ibm.com/servers/eserver/iseries/hardware/virtualizationgrandslam/.

iSeries System Performance Team. "The IBM eServer i5 Three-in-One Benchmark." May 2004. http://www.**ibm.com**/servers/eserver/iseries/hardware/threeinone

Harris et al. "Logical Partitions on IBM PowerPC: A Guide to Working with LPAR on POWER5 for IBM eServer i5 Servers." IBM Redbook, August 2004. http://publib-b.boulder.**ibm.com**/redpieces/abstracts/sg248000.html

Disclaimer

It is important to understand that all of these are synthetic tests using simulated application workloads and customers should validate their own application performance prior to using this technique in a production environment. This was only a functional demonstration and these measurements were not meant as a substitute for an official benchmark.   The SAP systems were not highly tuned for performance, and the benchmark tools were used only as a mechanism to provide a consistent workload.

Trademarks and Disclaimers