



Issue Date: Open Systems edition
April | May 2006

The Search for the Magic Elixir

Jaqui Lynch

Finding the perfect mix for utilizing virtualization capabilities

Server consolidation is becoming more critical to many companies as a way to better manage datacenter costs and expenses. The relationship between the number of servers and the costs associated with providing support for those servers isn't linear - it can rapidly become an exponential expense. In this article, I'll define server consolidation, explain how virtualization helps with this effort and illustrate how the latest IBM* announcements provide even greater benefits in this area.

The Server-Consolidation Formula

Much of the interest in server consolidation is due to the expected savings in hardware and software maintenance; heating, cooling and power; manageability; and software licensing.

As servers age, several things happen. First, the hardware maintenance costs rise, making the box more expensive to maintain. Typically, a new server is purchased with a three-year maintenance contract. Once that period expires, a new contract must be purchased - and that's not inexpensive. However, when a new server is purchased, it comes with a minimum warranty, and the cost to upgrade it to a three-year warranty can be less expensive than extending current warranties on older servers. Additionally, many of the new servers, particularly the System p5* Quad-Core Module (QCM) models, are now shipping with a three-year warranty by default rather than a one-year warranty, making the upgrade to a 24-7 three-year warranty less expensive. Secondly, the boxes go out of marketing. This means that you can no longer buy new parts or cards for those servers; they must be purchased in the used market. While this is cheaper, some companies don't like to purchase used parts.

Software-licensing cost reductions are attainable because of increased processor speed. As an example, let's look at an 8-processor Oracle LPAR running at 1.45 GHz on a p650 server. This is rated at 18.67 rPerf if running on AIX 5L v5.3 and 14.36 rPerf on AIX 5L v5.2. This equates to a four-processor 1.65 GHz System p5* 550 - a savings of four processors of Oracle licenses. But, that assumes the LPAR was running at 100 percent and we needed all those rPerfs. If the LPAR is really running at a 50-percent processor utilization on AIX 5L v5.3, then we only needed 9.4 rPerf - and this is a 2-way 1.65 GHz System p5 550 running at 100 percent. If we had several similar workloads whose peaks didn't coincide, we could take advantage of consolidating those workloads onto separate LPARs on the same server, and they could share the processors. Thus, we could buy a smaller server and still provide the necessary performance - all at a lesser cost.

Additionally, a significant portion of the cost of running a datacenter is the cost of the people to install and manage the systems. In a consolidated environment, it's possible to take advantage of some of the features of virtualization to dynamically manage performance and resources. Add to this a properly configured Network Installation Manager (NIM) server and a Tivoli* Storage Manager (TSM) backup environment, and the systems need significantly less human intervention. This can provide considerable savings and allow systems administrators to concentrate on other areas.

Depending on your datacenter, savings in heating, cooling and power may be of interest as well. In a complex server environment, these savings can be significant.

The Virtualization Additive

Since the introduction of POWER5* technology, virtualization has offered many great options to aid in consolidations. Not only are the servers faster and larger with more available memory, but they also include options such as Virtual I/O (VIO) server, shared Ethernet, virtual Ethernet, virtual SCSI and Micro-Partitioning* technology. Additionally, simultaneous multithreading (SMT) adds a performance boost to most workloads. However, the key to taking advantage of these features is implementing AIX 5L v5.3 (or a supporting Linux* OS), taking advantage of the new features in POWER5 technology.

POWER5 servers also provide the ability to simultaneously run multiple disparate OSs on the same server in different LPARs. For example, a server can run AIX 5L v5.3, AIX 5L v5.2, SLES9 and RHES4 all in their own LPARs. And all of these,

except AIX 5L v5.2, are supported using the VIO server facilities including the Integrated Virtualization Manager (IVM).

What it All Boils Down To

The bulk of the latest IBM announcements centered around price changes and new servers with the new QCM chips. The low-end servers were refreshed to POWER5+* technology, and some newer, faster CPUs were added. Additionally, there was a refresh of VIOs to better support IVM. New adapters were announced that run at the faster backplane speed of 266 MHz used in the POWER5+ backplanes. Some of the POWER5 low-end servers will be withdrawn from marketing by the time you read this.

So what does all of this mean for server consolidation?

IBM repositioned the System p5 570 to an enterprise-class machine, so it's now treated in a similar fashion to the System p5 575, System p5 590 and System p5 595. In the meantime, IBM unveiled the 560Q, a box based on the model 570 using the QCM chips. The 560Q is available with 4, 8 or 16 processors but doesn't include any I/O drawers or Capacity on Demand (CoD) features. The 16-way configuration comes with 12 PCI-X slots, so it's a nice box to use for consolidation where you can take advantage of using VIO server to own the resources and thus, keep a handle on the number of required cards.

Important Steps

To best take advantage of a server- consolidation play, it's important to migrate as many of the LPARs as possible to AIX 5L v5.3 or one of the supported Linux OSs. Without doing so, most of the POWER5 advantages disappear, with the exception of savings brought about by faster processors and general technology-refresh savings. AIX 5L v5.2 can't take advantage of VIO server, shared Ethernet, etc. Keep in mind any migration to POWER5 technology from older technologies requires the addition of a new Hardware Management Console (HMC) or, where supported, the use of IVM.

If using IVM or VIO server, it's important to remember to add sufficient resources to support those environments. They need processor and memory resources. Additionally, there are some catches with using IVM instead of an HMC. IVM must own every resource in the box, so every LPAR must use VIO server to get its resources. This means AIX 5L v5.2 isn't supported on the server.

Finally, when planning a server consolidation, it's important to size the box for more than just the workloads being moved. Typically, a NIM LPAR and some type of backup LPAR are added at the same time and must be sized accordingly. Additionally, there's a memory overhead in the server, as there was in the POWER4* servers. This is used by the POWER Hypervisor for hardware page tables and translation control entries. Typically, an additional 8 percent should be added to calculated consolidated memory to account for this overhead.

The Experiment

Let's illustrate a specific server-consolidation scenario. [Table 1](#) shows the setups I'll assume in the datacenter. The example scenario is similar to typical environments that are currently being installed and will serve as a base for this case study. As you can see, this solution is around the 71.4 rPerf mark and has significant power and cooling requirements.

Again, remember that any implementation using AIX 5L v5.2 can't take advantage of the new features and facilities provided in the new hardware. Partitions in AIX 5L v5.2 are required to have dedicated full processors, and they can't be served by a VIO server. In this example you'd want to migrate as many of the LPARs as possible to AIX 5L v5.3.

Let's assume we can migrate everything to AIX 5L v5.3 and take advantage of the Micro-Partitioning technology. Let's also assume we have the following policies in place (per LPAR):

- 2 x fiber cards
- 2 x internal boot disks
- 2 x 10/100/1000 Ethernet cards
- 4 GB memory per processor if database
- 2 GB memory per processor for other

In the unconsolidated environment in our example, we'd have the following:

- Disks 14
- Fiber cards 28
- Ethernet cards 28
- Memory 88 GB + 8 percent overhead = 96 GB
- Processors 32 x 1.65 GHz

- rPerf 71.4

If we moved this to a purely AIX 5L v5.3 environment on the new System p5 570 with 1.9 GHz processors, it would require the following:

- Disks 14
- Fiber cards 28
- Ethernet cards 28
- Memory 88 GB + 8 percent overhead = 96 GB
- Processors 14 x 1.9 GHz
- rPerf 75.88

The full-blown 16-way is rated at 85.20, so there's some room for growth. Judicious Micro-Partitioning use should result in significant savings in licensing costs as each processor is rated at more than twice the speed of the earlier ones, in terms of rPerf (i.e., a 2-way 570 1.9 GHz is 12.27 rPerf versus the 2-way p630 1.45 GHz at 4.41).

Given the architecture of the 570, we'd have the ability to boot eight LPARs of our total 14 from the eight three-packs in the CECs. That said, we'd need additional I/O drawers (each can support two LPARs) to support the remaining LPARs. Also, the 570 only provides a total of 24 PCI-X slots, so an additional 32 slots in I/O drawers are needed. A 7311-D20 I/O drawer provides seven PCI-X slots and two six packs for disk. We'd need a total of five times the I/O drawers to support the necessary cards.

[Table 2](#) lists the environmentals of the new configuration. These are estimations and your mileage will vary, but this gives you an idea of some of the potential savings. If we now go ahead and decide to add a VIO server and use it to service some of our LPARs (e.g., provide boot disks, etc.), we can reduce resources (see [Table 3](#)).

This configuration saves a total of 20 PCI slots and allows us to go down to two times the I/O drawers as we now only need a total of 36 slots. This is a savings in rack space, cooling, heating and power, not to mention the upfront funds to purchase the drawers and their associated maintenance.

Ingredients for Success

Clearly, some significant savings can be made by judiciously implementing virtualization on the newer, faster POWER5+ systems. The examples I included don't take into account modifying the sizing such that LPARs are made smaller and allowed to grow in the shared processor pool as needed. If you're considering a server-consolidation exercise, I strongly recommend that you move any LPARs you can to AIX 5L v5.3 (or Linux) and take advantage of the benefits provided by virtualization in this environment. It's highly possible that the savings will actually pay for the purchase.

Table 1

Server	LPAR	CPUs	Memory	rPerf	AIX Version	KVA	Watts	BTU/hr
p630 1.45 GHz 4 way	LPAR1	2	4 GB	4.41	5.2	0.78	750	2450
	LPAR2	2	4 GB	4.41	5.2			
p630 1.45 GHz 4 way	LPAR3	2	4 GB	4.41	5.2	0.78	750	2450
	LPAR4	2	4 GB	4.41	5.2			
p630 1.45 GHz 4 way	LPAR5	2	8 GB	4.41	5.2	0.78	750	2450
	LPAR6	2	8 GB	4.41	5.2			
p630 1.45 GHz 4 way	LPAR7	2	4 GB	4.41	5.2	0.78	750	2450
	LPAR8	2	4 GB	4.41	5.2			
p650 1.45 GHz 8 way	LPAR9	2	4 GB	4.47	5.3	1.7	1600	5461
	LPAR10	2	4 GB	4.47	5.3			
	LPAR11	2	4 GB	4.47	5.3			

	LPAR12	2	4 GB	4.47	5.2			
p650 1.45 GHz 8 way	LPAR13	4	16 GB	9.12	5.3	1.7	1600	5461
	LPAR14	4	16 GB	9.12	5.2			
7 x I/O drawers						2.52	2380	8127
TOTAL		32	88	71.4		9.04	8580	28849
rPerf taken from rsperfs021406.doc from IBM								

Table 2

Servers	KVA	Watts	BTU/HR
OLD	9.04	8,580	28,849
570 16 way 1.9g	5.472	5,200	17,748
5x7311-D20	1.8	1,700	5,805
TOTAL	7.272	6,900	23,553
SAVINGS	1.768	1,680	5,296

Table 3

LPAR	Fiber	Ethernet	Disks	Uses	Disk	Fiber	Ethernet
	OLD	OLD		VIO?	NEW	NEW	NEW
LPAR1	2	2	2	y	0	0	0
LPAR2	2	2	2	y	0	0	0
LPAR3	2	2	2	y	0	0	0
LPAR4	2	2	2	y	0	0	0
LPAR5	2	2	2	y	0	0	0
LPAR6	2	2	2	y	0	0	0
LPAR7	2	2	2	y	0	0	0
LPAR8	2	2	2	y	0	0	0
LPAR9	2	2	2	n	2	2	2
LPAR10	2	2	2	n	2	2	2
LPAR11	2	2	2	n	2	2	2
LPAR12	2	2	2	n	2	2	2
LPAR13	2	2	2	n	2	2	2
LPAR14	2	2	2	n	2	2	2
VIOS (uses 300 GB)	0	0	0	y	4	6	6
Total	28	28	28		16	18	18
Savings over old	12	10	10				

References

For more information, check out the following resources:

- The System p5 Information Center (http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/index.htm)
- IBM System p5 and ^ p5 Facts and Features (www-03.ibm.com/systems/p/hardware/factsfeatures.html)

Jaqui Lynch, a technical editor for *IBM Systems Magazine*, Open Systems edition and senior systems engineer focusing on System p5 and Linux at Mainline Information Systems, has worked in the IS industry for more than 26 years. She's been responsible for a variety of projects and OSs across multiple vendor platforms, including mainframes, UNIX* systems, midrange systems and personal workstations. Jaqui can be reached at jaqui.lynch@mainline.com.