Issue Date: AIX EXTRA eNewsletter
November 2004, Posted On: 11/1/2004

Understanding Virtualization

Jaqui Lynch

Virtualization is a big buzzword these days. Although logical partitioning (LPAR) and dynamic LPAR provide a form of virtualization, the introduction of POWER5 technology offers true virtualization and provides flexibility in resource allocation while minimizing constraints.

Before getting started with virtualization in the POWER5 world, several things must be in place. The first prerequisite is the combination of POWER5 hardware and AIX V5.3 or an enabled version of Linux. For this article, we'll use AIX V5.3 as a reference. Additionally, virtualization requires the Advanced Power Virtualization (APV) feature. This comes standard with the IBM eServer p5 590 and 595 servers, but is a chargeable feature on the 570, 550 and 520 models. Lastly, a Hardware Maintenance Console (HMC) is required to implement virtualization and the new Capacity on Demand (CoD) options.

So what does virtualization offer? What are the components? Firstly, virtualization is built on the basics of LPAR technology. A POWER5 system is always running under a hypervisor, even if all of the resources are assigned together. The hypervisor takes care of time slicing and dispatching for the partition workloads between processors. With the use of an HMC, partitions can be created just as they can with POWER4 systems—the minimum granularity is one processor, and I/O devices are still assigned at the slot level.

POWER5 technology takes this to the next level with the APV feature. APV provides for micro-partitioning and virtual I/O (VIO) server. In the case of micro-partitioning, a set of processors can be assigned to the shared processor pool (SPP). So, on a 16-way we could have seven processors dedicated to LPARs as we did previously and the other nine could be in the SPP or in the CoD pool. Let's assume they are in the SPP.

Processors in the SPP can be assigned to LPARs, just like regular processors. However, the granularity is different. The initial assignment is a minimum of one tenth of a processor with increments being in 1/100 of a processor. Keep in mind that these partitions still need I/O resources and dedicated memory. The I/O and networking resources can be dedicated as per previously or virtualized through the VIO server, or you can combine the two options.

When defining a partition you now have to choose between a shared-processor and a dedicated-processor partition. For a shared partition it's also necessary to determine whether the partition is capped or uncapped. This is an important decision as it affects both performance and license charges. An uncapped partition can use all of the processors in the pool if it needs to and if it's more important than the other workloads running. A capped partition can use processors up to the number of processors specified in the cap. License charging is typically based on the maximum processors that this workload could use—in the case of uncapped partitions this would be all the processors in the pool, whereas with a capped partition this would be the number of processors specified in the cap. Additionally, the number of processors allocated to a partition is referred to as the entitled capacity.

The other key component of the APV feature is the VIO server. This facility allows you to virtualize both I/O and network resources. The VIO server is a custom AIX V5.3 partition that can own I/O and network resources. In the case of I/O, disks and adapters are shared using the virtual SCSI server. This runs in the VIO server and effectively allows you to have a physical disk that has multiple logical volumes (LVs) allocated on it. You can then export each LV to a different client LPAR. The client LPAR would see these LVs as normal SCSI disks, even though they may be fibre attached to the server. This allows you to take a 146 GB disk drive and carve it into three 45 GB LVs that could be used as boot disks by three different LPARs.

Network-adapter sharing is done using a shared Ethernet adapter configured in the VIO server. Effectively, the Ethernet adapter belongs to the VIO server and is exported to the client partitions. IP addresses are mapped accordingly and the traffic is routed through the one physical adapter in the VIO server. From the VIO server traffic either flows over the real network or through the hypervisor to the actual partitions.

These are some of key components of virtualization. It's important to note that virtualization takes additional processor power and memory—it's not free. Part of any new server strategy is going to need to include an analysis of where virtualization

might work for your site. There are definite cost savings to be made by sharing resources using virtualization, but it's important to understand the workloads running and whether they need dedicated bandwidth or can work happily in a shared or virtualized environment.

Recommended Reading

- "Advanced POWER Virtualization on IBM p5 Servers: Introduction and Basic Configuration," SG24-7940
- "Advanced POWER Virtualization on IBM p5 Servers: Architecture and Performance Considerations," SG24-5768

About the Author(s):

Jaqui Lynch: Jaqui Lynch, an *eServer Magazine, IBM edition for UNIX* technical editor, is a senior systems engineer focusing on pSeries and Linux at Mainline Information Systems. During her more than 26 years in the IS industry, she's been responsible for a wide 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.