

Issue Date: AIX EXTRA eNewsletter
September 2004, Posted On: 9/1/2004

AIX EXTRA: Planning for POWER5

Jaqui Lynch

The introduction of POWER5 technology and the new IBM eServer p5 systems has caused a major change in the thought processes involved in planning in the pSeries world. Previously, sizing a system correctly involved pre-planning for resource consumption in terms of memory, CPU, DASD and adapters. POWER4 technology added the new dynamic partitioning features, and now POWER5 technology adds true virtualization, providing even greater flexibility in resource allocation that minimizes constraints but also requires other types of planning. Additionally, all new eServer p5 systems are simultaneous multi-threading (SMT)-capable.

The newly announced eServer p5 systems are the follow-ons for the pSeries 615, 630 and 650 models. There are several key points to note when planning for the new systems. The first is that they require a minimum of AIX v5.2 (ML4) or an enabled version of Linux (SLES 9 or RHAS 3 for POWER). Secondly, as with the POWER4 systems, a Hardware Maintenance Console (HMC) is required to run the system in partitioned mode. However, the HMC can't support both POWER4 and POWER5 systems simultaneously, as different microcode is needed on the HMC. An older HMC can be upgraded to support POWER5 technology, however, there's no going back to support POWER4 technology once this is done.

The HMC is primarily used for controlling LPARs. It's also necessary to implement virtualization and the new capacity on demand (COD) options. The HMC acts as a service focal point. Since its primary functions are management and service, a failure event affecting only the HMC doesn't bring down the complex or LPARs. Changes to resources handled by the HMC can't be addressed until the HMC is restored to service.

Minimum requirements for the partitions have also changed. The eServer p5 systems still require each partition to have exclusive access to its assigned resources when virtualization isn't implemented. When using virtualization and AIX v5.3 (or the correct Linux version), the minimum assigned resources for an LPAR (at AIX v5.3) are now 1/10th of one processor, 128 MB of memory, one boot disk, an adapter to access the disk, an Ethernet adapter to access the HMC and methodologies to run diagnostics and for installation (e.g., NIM for installs). Resources are assigned at the slot level, which means that all disks attached to an adapter are assigned as one group to one LPAR. However, if virtualization is implemented then disks and adapters can be virtual rather than real. This allows for sharing Ethernet and other adapters between LPARs. Virtualization requires the purchase of IBM's Advanced Virtualization Feature and enables sub-processing sharing, virtualization of I/O and Ethernet sharing through a hosting partition, and provides the Partition Load Manager.

When planning for memory, it should be noted that the eServer p5 servers use some of the memory to retain critical information. The amount that must be reserved depends on several factors, including the number of I/O drawers, number of partitions and how maximum memory for those partitions is set. Numbers for the eServer p5 systems have yet to be published. However, they should be somewhat similar to the POWER4 numbers. The key point here is to remember that not all of the memory will be available due to the overhead of page tables, the Translation Control Entry (TCE) and the Hypervisor.

Several additional areas in the eServer p5 world now require planning. Clearly the whole virtual environment needs addressing. Decisions must be made as to processor and memory allocations and whether processors will be in a shared pool or dedicated. This has an impact on workload license charges so these issues need to be solved upfront. Additionally, the utilization of virtual I/O server affects the number of adapters that are needed and careful planning is needed to help ensure that performance can be maintained. Two new Redbook drafts go into detail on the servers and planning issues, and serve as a good starting point:

- IBM eServer p5 Virtualization Performance Considerations (SG24-5768)
- Introduction to Advanced POWER Virtualization on IBM eServer p5 Servers (SG24-7940)

One last item to note is a change in the connectivity to the monitor switches. POWER4 systems have a PS/2 keyboard and mouse ports, whereas POWER5 systems have only USB ports. So it's important to ensure that any attached monitors are correctly configured to support any attached POWER5 systems.

As you can see, with the new eServer p5 systems, planning is now even more critical. To get the best value from these systems, it's important to implement a resource plan and maintain planning worksheets and checklists to help ensure that resources aren't over-committed and that unnecessary overhead is avoided. Additionally, it's critical to review the HMCs and see where older systems can be consolidated so that an HMC can be freed (or purchased) to support only eServer p5 systems.