IBM System p: Improving Server Utilization

BY CHARLIE CLER

Improving server utilization has become a hot topic in many businesses, especially with the increased focus on power and cooling consumption. This article will present a general approach for improving server utilization, as well as an in-depth analysis of System p* server capabilities for improving utilization. Last, recommendations on how to begin a server utilization improvement project and potential challenges are provided.

Business Challenges and General Concepts

Many IT operations managers are facing a number of expense pressures. Some of these include:

- Budget reductions—do more with less
- Server sprawl—running out of datacenter floor space
- Increasing power and cooling costs
- Reaching peak power load in the datacenter
- Rising systems administration costs

Improving the utilization of existing server assets can help address these challenges. Improvements in server utilization and the retirement of older servers can result in reduced demands for floor space, power and cooling as well as reduction in hardware maintenance and administration expenses.

Let's start our discussion with a simple definition for overall utilization. Overall utilization can be defined using the following equation:

Overall Utilization % = Utilized Capacity/Installed Capacity

Utilized capacity is the sum of measured CPU utilization for each of your servers. Installed capacity is the sum of the maximum performance ratings for all of your installed servers. You'll need to represent these two numbers using a unit of performance and not a percentage. Use of percentages doesn't provide distinction between different sized servers. It's recommended that you use a unit of performance like estimated transactions per minute (TPMs) or rPerfs (a System p metric). For example, using TPMs, utilized capacity for an individual server can be calculated using the following equation:

Server Utilized Capacity =
(CPU Utilization Percent) (Server TPM rating)

Referring back to the overall utilization percent definition, improvements can be made by: (1) Increasing the utilized capacity or, (2) by reducing the installed capacity. This is shown in *Figure 1, Improving utilization of installed capacity*.

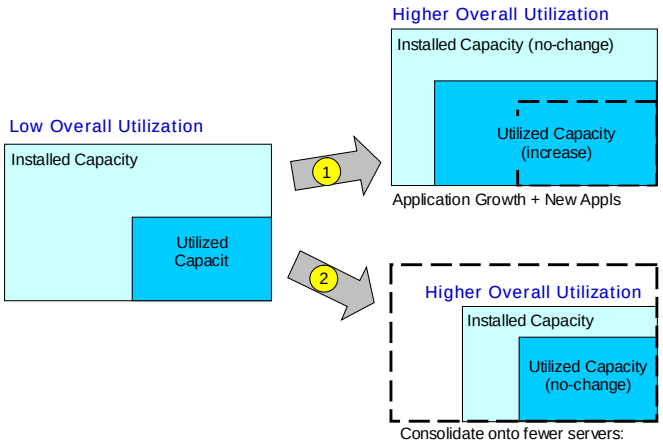

Figure 1 – Improving utilization of installed capacity

The first option in the figure shows higher utilized capacity through the growth of existing applications and/or the deployment of new applications. The effect will be an increase in overall utilization. Financial benefits might include a temporary halt or reduction in the growth rate for new server procurements and their environmental needs (power, cooling and floor space). The second option shows how consolidation onto fewer servers can be used to reduce the installed capacity, thereby improving overall utilization. You may have existing end-of-life (EOL) projects underway which provide a good starting point for removal of servers from your environment.

Partition Utilization vs. Server Utilization

System utilization has traditionally been measured at the operating system level. This measurement might represent an entire stand-alone server or a single partition running on a partitionable server. For partitionable servers, the ultimate goal is to improve the utilization of the entire server, and not just the utilization of individual parititions.

To improve the overall utilization of a partitionable server, it is recommended that you use the following two step process:

1. Reduce individual partition allocations to drive up each partition's utilization.
2. Deploy additional workload to increase overall server utilization.

This is shown graphically in *Figure 2 – Improving partition and server utilization*:

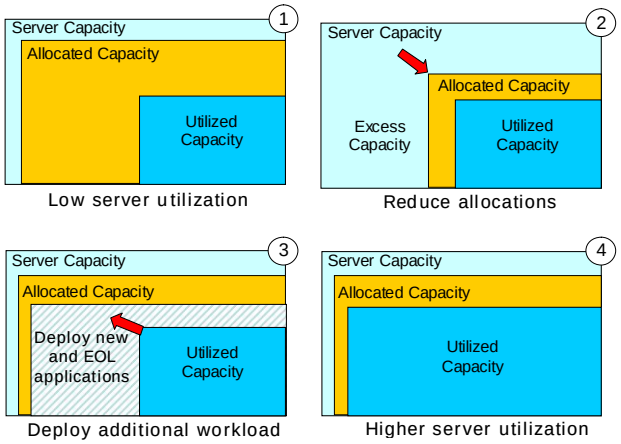

Figure 2 – Improving partition and server utilization

The allocated capacity is initially oversized in relation to the utilized capacity, resulting in low server utilization (1). The partition allocations can be reduced to align more closely with their utilization rates. Historical utilization data can be used to decide which partitions are good candidates for allocation reductions. When allocation reductions are completed (2), excess server capacity is easily identifiable and is no longer "hidden" within the individual partitions.

At this point, resource allocation for individual partitions has been optimized. Now we can focus on overall server utilization. Using the excess server capacity, new and rehosted EOL applications can be deployed (3). The result is higher server utilization (4).

For servers that do not support partitioning, you have the option to run multiple applications on a single operating system image, which is known as application stacking. In this situation, you would move directly from step #1 to step #4, skipping over the partition optimization steps.

Improving Utilization of System p Servers

Now that we have a high level plan on how to improve server utilization, let's look at specific capabilities within the System p servers that can be used to improve server utilization. Each System p processor generation, from legacy (pre-POWER4), to POWER4 and POWER5, has introduced additional CPU allocation and/or CPU sharing capabilities. Highlights for each generation include:

Generation	**Function**
Legacy	Application Stacking
POWER4	Logical Partitions (LPARS)
POWER5	CPU & I/O Virtualization (APV)

Your business may have more than one generation of System p servers installed, therefore it will be important to develop a utilization improvement plan which accommodates the generations present in your environment. Let's take a closer look at the capabilities provided by each generation.

Legacy servers encompass all System p servers introduced prior to POWER4. Legacy servers don't have partitioning capabilities and therefore run a single OS per server. Application stacking will be the primary method for improving utilization. To assist with the prioritization of processor and memory resource sharing among stacked applications, IBM provides a tool called Workload Manager (WLM). WLM was introduced with AIX* V4.3 and has been included in all subsequent releases of AIX at no additional charge.

Use of application stacking also helps to reduce the overall quantity of OS instances, which can result in lower administration costs. For this reason, application stacking should also be considered for use with POWER4 and POWER5 systems.

POWER4 and AIX V5L added logical partitioning (LPAR) capability to the System p product line. LPAR technology allows a single server to be divided into multiple partitions, each running its own OS instance (AIX, Linux* or i5/OS*). With LPAR, processors, memory and I/O resources are allocated independent of one another and without regard to physical boundaries.

Each LPAR comprises one or more CPU along with required memory, which is defined to the nearest 256-MB increment required. LPAR allows resource allocations to more closely match capacity planning specifications as compared to stand-alone servers. With LPAR, businesses also have the option to purchase enterprise class servers and then deploy LPARs of various sizes rather than having to buy a variety of different server models to satisfy various application needs. AIX V5.2 added the capability to dynamically move processors and memory between LPARs without OS reboots.

With the introduction of **POWER5** servers, IBM added a new hardware feature called Advanced POWER Virtualization

(APV). Coupled with AIX V5.3, APV provides full virtualization of CPU and I/O resources.

The shared processor pool component of APV provides CPU virtualization, which allows CPU resources to be assigned to LPARs in 1/10th CPU increments. This frees LPARs from the constraints of physical processor boundaries and permits them to share resources from the shared processor pool. The POWER hypervisor dispatches CPU cycles to the LPARs running in the shared processor. In addition, the hypervisor automatically distributes unused CPU cycles (in 1/100th CPU increments) to LPARs experiencing high utilization.

The APV hardware feature also includes the Virtual I/O Server (VIOS), which allows physical Ethernet and SCSI resources to be shared between LPARs. Sharing of I/O resources is an important feature, which allows large numbers of LPARs to be hosted on a single server with a reduced number of adapter cards.

Recommendations

As discussed earlier, the general approach for improving server utilization is to reduce excess CPU allocations at the partition level, followed by the deployment of additional workloads that will consume excess capacity. A summary of these steps, organized by processor family is shown in *Figure 3, Improvement plan, by processor family.*

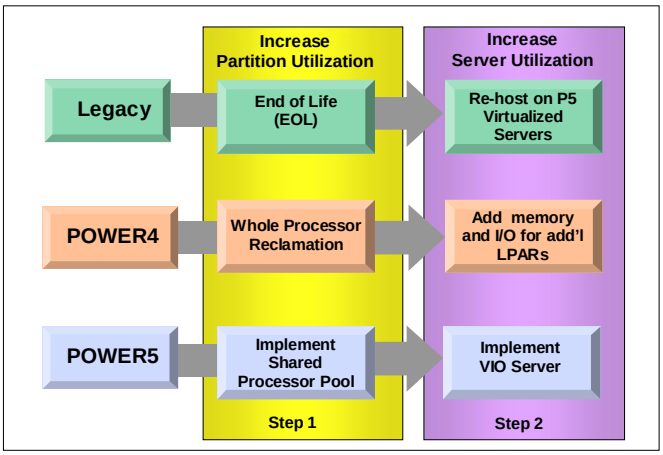

Figure 5 – Improvement plan, by processor family

Looking more closely at the legacy servers, it's recommended they be targeted for retirement under an EOL project. These servers should be fully depreciated and have higher maintenance costs. Also, their overall performance is small in comparison to POWER4 and POWER5 servers. Removing the legacy servers from service will reduce the overall installed capacity and increase the utilized capacity when applications are redeployed onto existing servers.

For POWER4 servers, partition utilization can be increased by removing excess CPUs from the partitions that currently have low utilization. Historical performance data can be used to determine which LPARs have excess capacity and to estimate how many CPUs can be removed from each of them. For LPARs running AIX V5.2 or AIX V5.3, processors can be dynamically removed without having to stop/start the LPAR. Step 2, improving server utilization, will require the deployment of new or EOL applications to consume the freed up CPU resources.

Improving partition utilization on POWER5 servers is best accomplished by configuring all AIX V5.3 partitions to use the shared processor pool. It's recommended that AIX V5.2 partitions be upgraded to AIX V5.3 so that they too, can participate in the shared processor pool. To increase POWER5 utilization at the server level, implementation of the VIOS will allow additional partitions to be deployed without the need to add more physical I/O resources.

If you have both POWER4 and POWER5 servers present in your environment, a best practice for deploying new or EOL workloads (Step 2) is to place larger applications on POWER4 and smaller applications on POWER5. This is due to: (1) The whole CPU granularity on the POWER4 isn't fine-grained and larger LPARs will likely make better use of the allocated CPUs, and (2) you can't run as many LPARs on POWER4 because each LPAR requires physical I/O because there is no VIOS capability on POWER4. The shared processor pool and the VIOS features on POWER5 allow you to run many more LPARs on single server.

As servers are driven to higher levels of utilization you will need to consider the increased demand for memory and I/O. Currently, memory cannot virtualized, therefore the total amount of memory required per server will increase as additional workloads are deployed. The required ratio of memory to CPUs has shifted from a typical 4-to-8 GB per CPU to 16-to-32 GB or more per CPU.

Key Learning Points

Improving server CPU utilization can have a significant, positive impact on your operational costs. It is important to establish metrics which allow you to track improvements to your overall server utilization. As your project gets underway, be aware that some servers, for business reasons, may be excluded from the techniques discussed here. Examples include database hot-standby servers and mission critical servers.

Last, your success will be highly dependent on management support of your project. Processor reclamation, application

stacking, and end-of-life projects will require closer communications between application development, capacity planning, and I/T operations organizations. Success will help to drive your company toward Service Level Agreements (SLAs) and away from each application group owning their own server resources.

Additional information on partitioning and Advanced Power Virtualization can be found at the System p Redbook portal, www.redbooks.com/portals/UNIX.

Charlie Cler *has been with IBM for more than 22 years, working with the System p5 servers since their beginning 17 years ago. He currently supports IBM customers in a system architect role. Charlie can be reached at cbcler@us.ibm.com.*

