THE CLIPPER GROUP
Navigator™

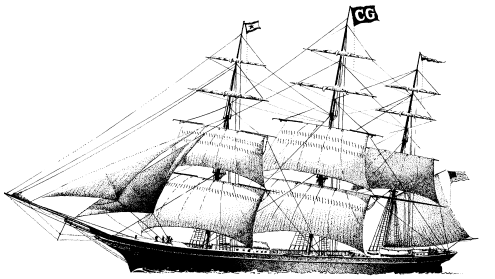

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2003020 May 14, 2003

IBM's Soups Up Its pSeries Servers and Offers Temporary, Extra Power On Demand

Analyst: David Reine

Management Summary

You are in the car dealer's showroom looking for the perfect vehicle: one that is simultaneously economical, sufficient, and powerful.

- *Economical* – able to deliver a lot of miles per gallon of gas
- *Sufficient* - able to carry you, your expected passengers and/or cargo
- *Powerful* - able to carry that extra load or to accelerate quickly when you need to pass

You want it to be efficient, effective and, yet, somehow able to keep up with your unexpected needs. What you want sounds impossible, like a combination of a Cooper Mini and a BMW wagon – able to do more, work when needed, but at a reasonable price.

Today, you want the same exceptional value from your enterprise servers, especially in an era of widespread consolidation from smaller servers. **You need to handle your new and continually changing mix of workloads (the *sufficiency* test), to do this at the lowest total cost of ownership (the *economics* test), and to be able to have the capacity to expand to meet your unexpected processing needs (the *reserve power* test).**

In the past, you would overprovision by adding extra processors to SMP servers or by having spare rack-mount servers standing by. Of course, you probably had to pay dearly for this standby power, to meet a future demand to process an extreme amount of work, such as an abnormally high number of transactions via your web portal or queries to your data warehouse. And you probably do not mind paying for a reasonable amount of headroom, but you know that there is a tsunami of demand waiting around the corner. **You want to be prepared, but you do not want to pay unreasonably for this protection (think *insurance*) against what might happen some time in the future.**

According to IBM: *you can have your cake and eat it, too*. Their recent mid-range to high-end pSeries announcement offers models with many new bells and whistles, but their ability to provision spare capacity for unexpected use at a reasonable cost gives enterprises something new to consider. Read on to learn more about IBM's new *eServer pSeries 650, 670* and *690* models, and its new capability for On/Off Capacity on Demand.

IN THIS ISSUE

On-Demand Resources for Enterprise Requirements

So you want to have your cake and eat it too? **The question is: how do you provision for the variety of variable (i.e., unpredictable) needs for IT services.** It is not about providing enough capacity, because you can accomplish this by extensive overprovisioning, with very high capital requirements and increased staff costs. We seem more focused, today, on the staff costs of provisioning and administration; but ask your CFO if he is concerned about ballooning capital costs for IT. **So, what you really want is enough reserve power, when you need it, in a vehicle that does not have ridiculously-high administrative burdens.** This all boils down to four requirements:

- The need for *more power*;
- The need for *reserve power*;
- The need for administrative simplicity; and
- All wrapped in an *economical, flexible financial offering*.

More Power

Processing power is usually a function of the speed of the processor, the amount of memory, the sophistication of the server, and the number of processors. The fastest systems usually come with a higher cost per processor, although fewer may be needed. Regardless, **when you need more processing power, you are speaking about adding processors to a server or more servers to a cluster or complex**.

More Reserve Power

Likewise, reserve power usually translates into more processors and memory. The question is: how much more do you really need? *10%? 40%? 100%?* Now you understand why overprovisioning is so attractive. **It is easier to overprovision than to accurately provision and - until recently - less painful.**

Administrative Simplicity

If you have to stop what the enterprise is doing to add capacity or if the addition requires complex adjustments or staff-involved workload balancing, then adding capacity on a just-in-time basis is not attractive. **You need the convenience of adding capacity with a few mouse clicks.**

Economical and Flexible Financing

In the end, this is all about money – when, and how, to spend it. **Overprovisioning has severe economic consequences: you buy more than you need at today's prices and pay more for software, etc., rather than spending less up front for hardware and software, and buying it later for less, because hardware prices are always declining.**

This is really all about paying a premium for insurance, where the coverage is having enough power available when you need it, on a just-in-time basis. The amount that you pay for the coverage needs to be aligned with the benefits that you might receive and the risks that you hope to mitigate. You do not want to overpay for the benefit (e.g., the cost of the processors) for the risk that you are avoiding (e.g., the business that might be turned away). This is a delicate balancing act; of course, as you might have experienced in trying to decide how much disability or life insurance you want to buy

IBM Responds with New *pSeries* Models

Yes, this is about more hardware. In this case, it is about UNIX servers. IBM has upgraded is mid-range to high-end *pSeries* offerings. As with any upgrade, they go faster, have more capacity, perform more work more efficiently, and cost less on a per transaction basis. That makes the *eServer pSeries 650, 670 and 690 family* interesting in its own right, and worthy of consideration. However, this issue of **The Clipper Group Navigator** is about *reserve power*, and that is where the new pSeries models offer something new.

Innovative Approach

Previously, you could order extra processors for pSeries, i.e., more than you thought that you might need. You had two choices:

1. You could pay for them up front and have them idle until needed, or
2. You could pay a lesser fee to have them installed, but not activated

until you need them. Once activated, you would buy them at the current price, and increase the capacity[1]. IBM calls this *Capacity Upgrade on Demand (CUoD)*.

Buying too much capacity is not the right answer for a serious, but occasional spike in demand. Above a reasonable buffer, it is just too expensive. Adding permanently to the available capacity is also not the right answer for meeting an unexpected spike in demand, because it has a high purchase price and usually increases software license charges.

What you really want to do is borrow some power from a friendly neighbor, and give it back when you are done. You might create a mutual-need arrangement with another enterprise, but that may be messier than you want to consider today.[2] **What you really want to do is to rent the additional hardware that you need for a day or two, without long-term, asset-based ramifications.**

On/Off Capacity on Demand

With its new pSeries offerings, IBM is making this possible. **You can start with a small number of processors activated to meet your expected needs, including a buffer for some unanticipated needs, and have extra processors and memory installed for real-time activation, when the need arises.** The extra processors can be activated under an option IBM calls *On/Off Capacity on Demand (OOCoD).*[3] *The extra memory can only be activated under the CUoD option, i.e., once activated, it is purchased and stays installed.*

Exhibit 1 on the next page shows the three new pSeries models and their maximum number of processors. It also shows the minimum number of processors that you must buy at the outset, and the processors that may be installed and held in reserve, if you pay an up-front premium of about 20% for each inactive processor. The premium is a prepayment for the option to activate the spare processors, dynamically configurable when you need them.

As indicated in Exhibit 1, the new pSeries models range from

- Two active and up to 6 inactive CPUs @ 1.45 GHz, to
- Up to 16 active and up to 16 inactive CPUs @ 1.7 GHz.

This ensures significant scalability and investment protection within a single pSeries server. Combined with a potentially large reserve of inactive memory (as shown in Exhibit 1) and *Dynamic Logical Partitioning (DLPAR)* functionality, you can direct the power of the pSeries to the application partition in need when management requires it.

Risk Management

Just like renting a power tool at the local rental center, there is a charge for each day that you use it. The premium is a separate payment for a guarantee that a tool will be available, when and where needed. Moreover, like the rented power tool, if you keep it for too long, you would have been wiser to buy it in the first place. Therefore, you have to carefully consider your needs, the duration of your needs, and the terms and conditions of your rental.

IBM has now enhanced its Capacity Upgrade-On-Demand (CUoD) functionality for the mid- to high-end pSeries Servers to enable a unique on-off capability (OOCoD). IBM is delivering the capability to install inactive processors at only 20% of their list price and then to activate these reserve *Power4* and *Power4+* CPUs, for short, variable lengths of time in order to assign required resources to the applications in need when they need them. These additional CPUs may be activated in two-CPU blocks for 30-days, with the user having the authority, and responsibility, to turn that processing power off in 24-hour increments.

---

[1] And costs related to the increased capacity, such as software charges.

[2] It is the concept behind resource sharing in a grid environment, with either arranged sharing among friendly enterprises or among strangers linked into a grid pool.

[3] For further information about On-Off Capacity on Demand, see **The Clipper Group Captain's Log** dated May 8, 2003, and entitled *On-Off Capacity on Demand – Does This Make Sense for Your IT Infrastructure*? at http://www.clipper.com/research/TCG2003019.

pSeries Model #	Processor Speed	# of Active Power4 Processors	# of Inactive Power4 Processors	Total # of Power4 Processors	Maximum Inactive Memory
650	1.45 GHz	2	6	8	32 GB
670	1.5 GHz	8	8	16	64 GB
690	1.7 GHz	16	16	32	128 GB

Exhibit 1 – Active and Inactive Processors and Memory by Model

Reserve Power Requires Forethought

These 30-day blocks are available at 33% of the announced fee for the standard CuoD purchase upgrade. If enterprise planning indicates that more than three 30-day blocks will be required, then it may indicate a decision to purchase the upgrade rather that renting.

Combining OOCoD with Dynamic Logical Partitioning, IBM delivers a *virtual server*, the capability to assign additional CPU cycles and memory blocks to any new or existing mission-critical application. This is a new way to meet unexpected increased demand on your servers, whether driven by customers on the Internet or by your marketing department's data mining application.

As indicated earlier, software charges are an important issue. Traditionally, when you increase the capacity of a server, you must pay additional software license fees as if you were going to have the capacity turned on forever. Traditional licenses do not have a provision for reducing capacity. If you turn it on, you pay for full fee. Even worse, many software licenses do not distinguish between active and inactive processors, requiring you to pay for software based upon the total number of processors supported by the platform.

In addition, you have to factor the maintenance costs for inactive processors into your consideration equation. The warranty on inactive processors begins on installation, and 20% of the usual maintenance will be charged for inactive processors after the warranty expires.

IBM is pushing hard for software vendors to allow temporary rental to correspond with its hardware rental offering. IBM's major software offerings (e.g., *Lotus, Tivoli, DB2, and WebSphere*) will have a daily rental price, charged only when reserve processors are active, in addition to their one-time license charge, based on processors allocated to each application. This eases resource activation.

New Offerings from IBM

In fact, IBM has announced a series of market programs that enable IT to activate additional pairs of processors while minimizing user risk.

- You can activate processors on a permanent basis for a fixed cost to satisfy the long-term requirements of consolidation or user growth.
- You may also activate CPUs on a short-term basis in blocks of CPU-days. These blocks may be broken into as many increments as the user needs, with a granularity of 1-day.
- If you are not sure if this is the correct strategy for your enterprise, IBM has a "Try & Buy" Program, available for a 30-day period.
- You may also extend your first activation through September 30, 2003, at no additional charge.

Conclusion

You need to think about the reserve power that you need and how you plan to provision for it. With its new pSeries offerings, IBM is delivering more powerful and more dynamically configurable servers, with flexible reserve power options. **On-demand provisioning may be the beginning of new and powerful procurement alternatives to meet an enterprise's on-demand requirements.**

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is a Senior Contributing Analyst for the The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and a MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 37.***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.