

# The Value of Virtualized Consolidation



*Effortlessly balance workloads, easily meet business demands with IBM POWER7 Systems*

---

## Table of Contents

- 1 More data, more applications, less utilization**
  - 1 Smarter systems to meet workload demands**
  - 2 The advantages of virtualized consolidation**
  - 4 More performance, less consumption**
  - 4 Smarter systems: delivering capacity and scalability when you need it**
  - 5 Highly available all the time**
  - 6 Avoiding complexity and risk with Power Systems**
  - 7 Smarter systems for smarter business**
  - 7 For more information**
- 

## More data, more applications, less utilization

We live on a planet where data is exploding at an annual compound rate of 57% and nearly 6 terabytes of information are being exchanged over the Internet every second. Ironically, the average server in the average business runs at less than 10% utilization. Scratch the surface and you'll find that there are only two kinds of systems—those that can adapt to the needs of a smarter planet, and those that can't.

On a smarter planet, new application workloads are emerging. And you need to integrate applications and drive more decisions in real-time. Pivotal to optimizing these systems is a deep understanding of the business needs and the workloads that run on them.

This fact is influencing the way that IBM is designing and building systems, including how they will be configured going forward. A general purpose approach to systems design will no longer achieve the performance and business outcomes required on a smarter planet.

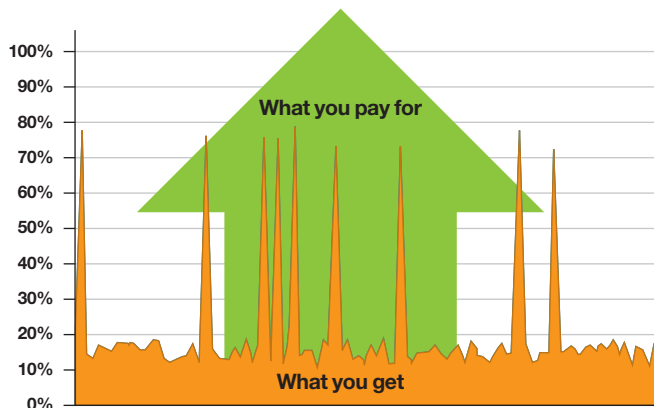
## Smarter systems to meet workload demands

If we take a deeper look at the workloads that run our world, it is apparent that not all workloads are created equally. Your business has various workloads. And different workloads place unique demands on systems – requiring different approaches based on core application function, data management, usage patterns, service levels, and integration with other applications and systems. Workloads run more efficiently when matched with the right computing resources; one size does not fit all.

For example – supporting an SAP application requires a different balance of transaction processing and database activity than business intelligence and analytics workloads. Similarly, enterprise collaboration applications require a different balance of IT system resources than the applications used to handle the complex, high volume business processes of an organization or those that manage a data center and keep it running efficiently.



The new POWER7-based enterprise class systems from IBM are not simply servers, they're fully integrated systems with the ability to run hundreds of virtual servers, helping you drive up to 90% utilization and get the most out of your investment. These next-generation systems integrate massive parallel processing, workload optimization features and analytics capabilities to optimize for the complex workloads and dynamic computing models of an increasingly data-driven world.



## The advantages of virtualized consolidation

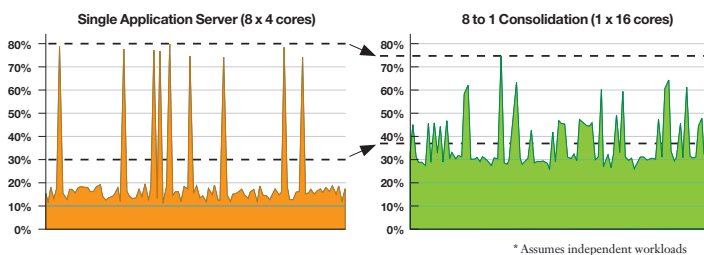
IT managers have to grapple with the issue of server sprawl and the accompanying high cost of management combined with the rapid growth of energy use in the datacenter. One solution that can address this problem is server consolidation. The challenge is to not just consolidate to fewer servers, but to get the most out of these systems by taking advantage of virtualization and other technology to dynamically share processing resources. Doing this decreases the wasted space that you are paying for, but not using. (See *Figure 1*)

Think about the last plane ride you took. How many seats were empty? How about the hotel you stayed at or the restaurant where you ate dinner. How many empty rooms or tables were there? These businesses are constantly looking for ways to fill that unused capacity – off-season hotel rates, two-for-one dinner specials, last minute plane fares. These businesses know that increased efficiency improves the bottom line, yet many IT organizations continue to accept high waste – meaning idle compute power – as a matter of course. This means paying higher software costs, higher maintenance costs, burning more energy, and taking more space in the datacenter. Even worse, it means more complexity and risk in terms of daily operations. By contrast, sharing resources through intelligent virtualization enables higher utilization which, in turn, reduces cost.

Intelligent resource sharing means several things: First, you must be able to divide processors, memory and I/O bandwidth into small chunks. (Power Systems divide processors logically into 1/100th of a processor for purposes of incremental adjustments.) Second, you must be able to dynamically adjust the assigned processing resources as workload demands both rise and fall with the business. (PowerVM does this automatically.) Third, you must have a sophisticated way of telling the system how to manage on your behalf. (Again, PowerVM and Systems Director allow for advanced workload priority and energy policy settings.)

Resource sharing drives down the wasted capacity in your system by using it for other workloads. Mainframe computers have used this concept for years and, now the performance and capability of today's UNIX systems can take advantage of resource sharing as well. Resource sharing works best when multiple workloads are placed together on the same system. In fact, the more workloads that come together, the more diverse their personalities, the better it works.

A simple example is illustrated below. Eight similar workloads can be accommodated on eight separate systems requiring a total of 32 processor cores. However, when those eight same workloads are loaded onto a single larger system and allowed to share resources, the utilization increases and the total number of cores required drops in half. This is the magic of resource sharing: if done effectively with the right systems and virtualization software, you decrease the overall amount of required system resources. In this case, not only did the number of cores decrease, but all of the costs that are associated with the system, such as maintenance, software charges, and energy usage. In most cases the required floor space will also decrease. By gathering as many workloads as possible on a single system, you can drive up the utilization and drive down your costs, making your data center more efficient and your business more productive.



By sharing resources across the system, organizations can:

- Double or triple server utilization rates, reduce total cost of ownership, and make better use of IT assets.
- Improve business responsiveness and operational speed by dynamically re-allocating resources to applications as needed, to better match changing business needs or handle unexpected changes in demand.
- Simplify IT infrastructure management by making workloads independent of hardware resources, enabling companies to make business-driven policies to deliver resources based on time, cost and service-level requirements.



## Smarter systems: IBM Power processor-based systems go to University

The University of Pittsburgh Medical Center (UPMC), an \$8 billion global health enterprise, depends on IBM Power Systems servers and AIX to operate many of its business critical databases and applications. But three years ago, UPMC was faced with a capacity problem – it was running out of CPU capacity faster than anticipated due to accelerated business growth. The solution was to aggressively adopt micro-partitioning and CPU sharing.

Within three months of identifying the capacity problem, UPMC had uncapped 90% of its LPARs and reclaimed 50% of its processors. All of UPMC's CPUs are in shared-processor pools. By minimizing CPU entitlements, yet being aggressive with virtual CPU settings, UPMC has achieved at or near 80% CPU utilization across many of its Power servers.

Processor virtualization has given UPMC the ability to provide capacity very quickly and efficiently at no additional cost.

Explains Paul Sikora, vice president of IT Transformation for UPMC: "The virtualized infrastructure flexes to meet processing peaks; the staff can respond to the demands of UPMC faster. We are more productive, more agile, and more reliable, at a lower cost point."

### More performance, less consumption

When viewing the performance of an application, much time is spent looking at the performance of the hardware that enables the application to run. POWER7 excels on many industry standard benchmarks, but what often can affect application performance even more is the ability to provide system resources to workloads in need. This is where virtualized consolidation provides surprising results.

Whether the total system size, the I/O bandwidth, the memory available or the size of the partitions, all systems have limits and a workload can only run within the resources assigned. If resources run out, the workload runs with what is available – even if that degrades performance. Think of this like a governor on a car: you need to get somewhere quickly but the regulator is set to 40 miles per hour (mph). No matter how hard you step on the gas pedal, the car cannot go faster than 40 mph. But, remove the governor and the car goes faster, up to the maximum speed the car is capable of achieving. Get in a faster car and you can go even faster.

Executing a program in a computer system is a lot like this car. The server has a governor that limits it from achieving its highest potential. It may be the size of the virtual partition that is assigned or supported. Or it may be the I/O bandwidth or memory that limits the ability to drive higher utilization. Or, it could be the system size itself. Eventually, the application runs out of resource.

That's why it makes sense to consolidate all of your workloads onto enterprise class Power Systems servers. Unused resources are available to any and all workloads on the system. When resource is available from anywhere in the system (whether it's processor, memory, or I/O bandwidth) it becomes available to other workloads, instantly and automatically – without user intervention, without moving workloads. Resources scale dynamically up and down, all day long. The system handles the adjustments. You manage your business.

One particular Power system client discovered this when they completed a consolidation of 20 Sun servers. A job that had taken 28 hours previously now took only 3 hours. Yes, the Power Systems were faster but the previous environment had capped the workload at the capacity of a single server. It had a governor in place based on the size of the system. By consolidating to a dynamic virtualized environment on a larger system, the workload was free to run unconstrained.

This is the great performance advantage of shared resource systems. You can improve service as system resources accommodate new workloads and respond to rapidly changing business needs – automatically. In addition, this leverages existing resources and does not require installation of additional hardware, saving you time and money.

### Smarter systems: delivering capacity and scalability when you need it

One of the reasons businesses have compute power sitting idle is that companies often buy systems based on “what-ifs.” What if we need more capacity? What if our business grows or has spikes in demand? What if the application was sized incorrectly? What if a new level of software requires more horsepower? All of these what-ifs lead to extra capacity which is multiplied with each server purchased, each time.

As we have seen, resource sharing helps reduce this overhead. But IBM Capacity on Demand<sup>1</sup> can provide an extra layer of protection. Inactive resources, that are not visible and not purchased, can be quickly activated and made available for permanent or temporary growth, trial usage, or unanticipated short term spikes in demand. Best of all, there is no disruption to applications or users. Additional processor and memory can be added as needed and assigned where required.

Capacity on Demand makes a significant difference in a company's financials, as well. Comparable to – and often less than – purchasing capacity from public cloud providers, capacity on demand from IBM offers:

- More granular price points. Capacity on Demand-enabled systems are sold in increments of one processor core and one GB of memory. This means that you can choose exactly the size that fits your business, both during your initial purchase and on subsequent upgrades.
- Utility-like pricing that can be billed quarterly. Usage is billed when activated, not up front and has a post-pay option that allows user to defer billing to end of quarter, which lets you more closely align your billing with actual usage
- Low temporary usage pricing. One additional POWER7 processor core with 8 GB of memory can be activated for as little as \$21 per day making this an attractive option for in-house cloud environments. This low-cost daily rate includes hardware maintenance and software charges for STG licensed software.
- Non-disruptive upgrades. Power systems with CoD resources provide instant upgrade capability that does not require system or application outage. Inactive processors and memory can easily be activated and assigned to existing workloads with out long waits or interruption to service. Automatic activations are even available for utility-like operations.

Enterprise systems also offer generation-to-generation physical upgrades<sup>2</sup> that make the purchase of newer technology less expensive. For example, clients who have POWER6 595 systems installed can upgrade to POWER7 based 795 systems for less than what a new Power 780 would cost. Live Partition Migration smoothes that migration with the ability to move workloads to alternative systems while physical upgrades are being performed. Clients can upgrade to POWER7, save money, and never take their application down.

## Highly available all the time

What happens when your system is down? Productivity stops for most of the company, while IT productivity is diverted from activities that could be focused on innovation to troubleshooting and repair. Business resiliency is crucial – it's not just a buzz word. And enterprise-class Power Systems servers can help make downtime a thing of the past, and help you ensure the resiliency of your business. In fact, according to a study by the research group IDC, enterprise class Power Systems clients reported downtime reductions of 60% - 94%, compared to their previous non-consolidated environments. In addition, this same study reported a 50% reduction of time related to IT maintenance.

The data shows that consolidation reduces downtime. This comes from fewer servers, fewer cables and connections and fewer things to manage. The system balances workloads and resources. Communications between workloads does not get unnecessarily exposed to the external world. Best of all, the most mature and sophisticated technology is available to all of your applications.

The best RAS features in the industry (outside of mainframes) are reserved for IBM Power System enterprise class. For example, Active Memory Mirroring, which reduces potential downtime by maintaining duplicate copies of the system hypervisor in memory. This ensures that hypervisor functions are always available and not subject to memory failures. In addition, the memory used on enterprise class Power Systems is more reliable, having implemented additional redundant capabilities not available in the rest of the Power Systems product family.

This may not seem like a deal breaker until you consider that the success of your business depends on how available your solutions are. With enterprise class Power Systems, you've got resiliency without downtime all while simplifying your server room.

**Table 1. Some RAS features are only available on enterprise class**

RAS Item	Power Blades	Power 750	Power 780	Power 795
Redundant/Hot Swap Fans & Blowers	●	●	●	●
Concurrent Firmware Update	●	●	●	●
Redundant/Hot Swap Power Supplies	●	○	●	●
Processor Instruction Retry & Alternate Processor Recovery	●	●	●	●
Storage Keys	●	●	●	●
PowerVM™/Live Partition Mobility/Live Application Mobility	○	○	○	○
Redundant VIOS support	–	○	○	○
Dual disk controllers (split backplane)	–	○	●	●
Hot Swap DASD/Media/PCI Adapters	–	●	●	●
Dual HMC hardware management	–	●	●	●
Redundant Service Processors with Dynamic Failover	–	–	●* New!	●
Redundant System Clocks with Dynamic Failover	–	–	●* New!	●
Redundant/Hot Swap Power Regulators	–	–	●	●
Dynamic Processor Sparing	–	–	○	○
Dynamic Memory Sparing	–	–	○	○
Hot-node Add and Repair/Hot-memory Add for all nodes**	●	–	●* New!	●*
Enterprise Memory	–	–	●	●
Hot GX Adapter Add and Repair	–	–	● New!	● New!
Mid-plane connection for inter-nodal communication	–	–	–	●
Active Memory Mirroring for Hypervisor	–	–	–	● New!

● Standard    ○ Optional    – Not available

\* Requires two or more nodes

\*\* Planned for 2H10 on 780, 1H11 on 795

## Avoiding complexity and risk with Power Systems

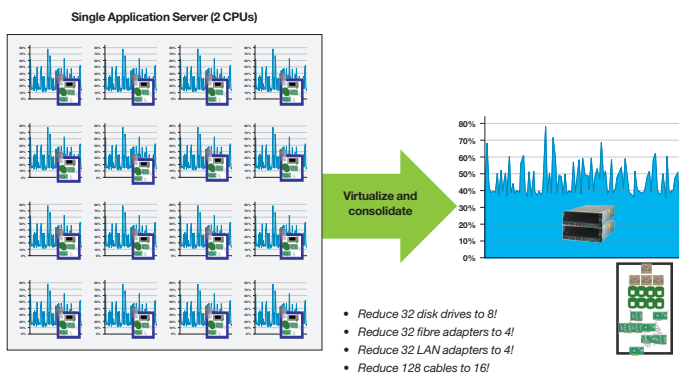
Complexity is one of the biggest issues facing business today – from globalization to collaboration, from volatile markets to uneven growth, companies are struggling to make sense of a complex world. The data center is one place where complexity can be easily managed. For instance, it is easier to secure and administer one enterprise server with virtualization than a collection of smaller servers. Think of the economies.

There are:

- Fewer operating systems to install security patches on
- Fewer systems to monitor for intrusion and compliance
- Fewer systems to synchronize security policies on
- Reduced number of administrators and administrative touch points
- Less a chance of the domino effect of a single system with weak security failing and being used for escalation of privileges and node hopping

Simplicity also comes from having fewer servers, adapters and cables to manage. Fewer connections mean fewer points of failure and less maintenance. This translates to less downtime and leads to a lower cost of operation. All of this is enabled by that first decision to consolidate multiple workloads on fewer systems with advanced virtualization software from IBM Power Systems.

POWER7 systems are helping to redefine computing to address the challenges of a smarter planet. From large scale business analytics, databases, and high volume transaction processing to consolidation that exponentially lowers costs, Enterprise class Power Systems have proven capabilities to simplify server rooms while positioning companies for future growth.



*Fewer servers, fewer cables and connections and fewer things to manage means much less downtime.*

## Smarter systems for smarter business

IBM Power Systems servers are the world wide leader in UNIX™ system sales. Why? IBM clients value the performance, virtualization, cost-savings, reliability, energy efficiency and proven roadmap provided by Power Systems servers.

Maybe you've wondered about IBM POWER7 systems, but weren't sure you needed an enterprise-class version. The roadmap for the future could not be clearer. Systems continue to grow in size and capacity. Virtualization continues to grow in sophistication and acceptance. Power Systems are the leaders in integrated, virtualized consolidation.

Other businesses like yours have chosen this path and more are choosing it every day. Surveys show that virtualized consolidation leads to lower costs, improves performance and reduces downtime, all while simplifying operations. Why not trust your business to the leaders. Go with scalable solutions that simplify complexity while still giving you all the capacity you need today – and in the future.

## For more information

For more information and to discover further how IBM POWER7 platforms deliver superior performance, using massive parallel processing, workload optimization features and analytics capabilities to optimize for the complex workloads, please contact your IBM representative or IBM Business Partner, or visit the following Web sites:

[ibm.com/systems/power/](http://ibm.com/systems/power/)

[ibm.com/aix](http://ibm.com/aix)

[ibm.com/systems/i/advantages/index.html](http://ibm.com/systems/i/advantages/index.html)

[ibm.com/linux/power](http://ibm.com/linux/power)

[ibm.com/systems/p/solutions](http://ibm.com/systems/p/solutions)



---

© Copyright IBM Corporation 2010

IBM Corporation  
Route 100  
Somers, NY 10589  
USA

Produced in the United States  
February 2010  
All Rights Reserved

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features or services discussed in this publication in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

IBM, the IBM logo, [ibm.com](http://ibm.com) and Power are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. A full list of U.S. trademarks owned by IBM may be found at: [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml).

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

- 1 Capacity on Demand is exclusively available on the Power 770, 780 and 795 systems.
- 2 Physical upgrades are exclusively available on the Power 770, 780, and 795 systems.



Please Recycle