IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Infrastructure Convergence: The Integration of Technology to Lower Cost and Improve Business Response

Sponsored by: HP

Michelle Bailey

November 2009

IDC OPINION

IDC believes that we are entering into a new technology business cycle where converged IT infrastructure will become an increasingly attractive option over standalone solutions.

After more than a decade of physical server sprawl, massive growth in storage terabytes, and the emergence of hierarchical network architectures, IT complexity has reached an all-time high. Today, IT organizations are forced to act as their own systems integrators. The ongoing management of this sprawling infrastructure has led to spiraling costs, and the time required for testing and tuning during deployment has led to a drain on available staff resources and ability to quickly realize business value.

Customers have already begun the work to reinvent their IT infrastructure and are quickly moving to more efficient and business-ready platforms. Server virtualization is now mainstream, while deduplication technologies and tiered storage are helping to lower the burden on information systems. As a result, many are rethinking their plans for 10GbE adoption and virtual I/O technologies as network performance becomes a critical gating factor for a server and storage rearchitecture. In addition, aging datacenter facilities have forced IT organizations to introduce supplemental power and cooling equipment and deploy energy management strategies. While these individual technologies are helping to lower up-front capital costs, they have not yet helped to lower management overhead, and their introduction requires significant integration investments.

With the sprawl in IT systems, many customers find themselves with a complex and inflexible set of infrastructure that limits their ability to respond to changes in the business. For many, the goal of offering "IT as a service" remains elusive, and their legacy infrastructure is an increasingly larger part of their IT maintenance budget that constrains the ability of IT departments to launch new projects and innovate.

IDC believes that the next technology cycle will have a converged architecture as a central design feature and that efficiency, high performance, and the ability to quickly rightsize the infrastructure relative to application and business demands will be critical. The movement to integrated or converged technology is not new. We have seen this trend in the server market twice in the past 30 years, and our market data indicates that we are at the very beginning of the next phase in IT. This next business cycle is "service centric" and has a very clear set of objectives: simplify and reduce complexity, maximize resource utilization, and shorten time to market to more quickly realize business value.

SITUATION OVERVIEW

IT infrastructure has become the backbone for most businesses today, supporting a wide array of applications and enabling new business initiatives. For organizations to continue to grow and evolve their business, ongoing innovation and IT expansion will be essential. Most companies today support their own technology and over many years have built a complex series of hardware and software solutions that they themselves have carefully tuned and integrated to meet the performance and availability objectives of each application. As the size of IT environments has grown, so too have the management costs associated with system deployment and maintenance, placing significant pressure on IT budgets and staffing resources.

Changing the Economics of the Datacenter

Over the past 13 years, IDC has tracked the changing landscape of the datacenter. During this time, we have observed a dramatic shift in both the types and the numbers of technologies installed. Figure 1 demonstrates the rising costs and management burdens of the datacenter over time.

Spiraling Operational Costs for the Datacenter

The number of servers installed on a worldwide basis has increased sixfold from just over 5 million in 1996 to more than 30 million in 2008. Until recently, customers had been deploying at least one physical server for each installed application (often more when taking into account test and development environments), leading to physical server sprawl. As a consequence of the huge numbers of installed servers, staffing costs on systems maintenance have risen 600% to over $120 billion annually, and the cost to power and cool installed servers has more than tripled from $2 billion to $10 billion per year.

Server Virtualization Impacts Installed Infrastructure, But Not Management Costs

Along with the growth in the number of installed servers has been a continued improvement in the capabilities of system hardware, particularly x86 server hardware. Most applications consume only a fraction of an average server's total capacity, typically 5–10% of an x86 server. Overprovisioning of server resources has been a standard tactic to ensure application performance and service levels. Server virtualization really has proven to be the "killer app" for the datacenter. By allowing IT organizations to run multiple applications per physical server, virtualization has not only improved the utilization rates of server hardware but also lowered spending requirements for new server hardware. IDC expects the server installed base to stabilize over the next five years; however, this tells only part of the story for IT departments. While the growth in physical infrastructure flattens out, the growth in the number of virtual machines explodes. IDC expects virtual machine (VM) density (or virtual machines per physical server) to rise from five on average in 2008 to more than eight by 2013.

FIGURE 1

Sprawling Infrastructure: Operational Costs Rise Dramatically

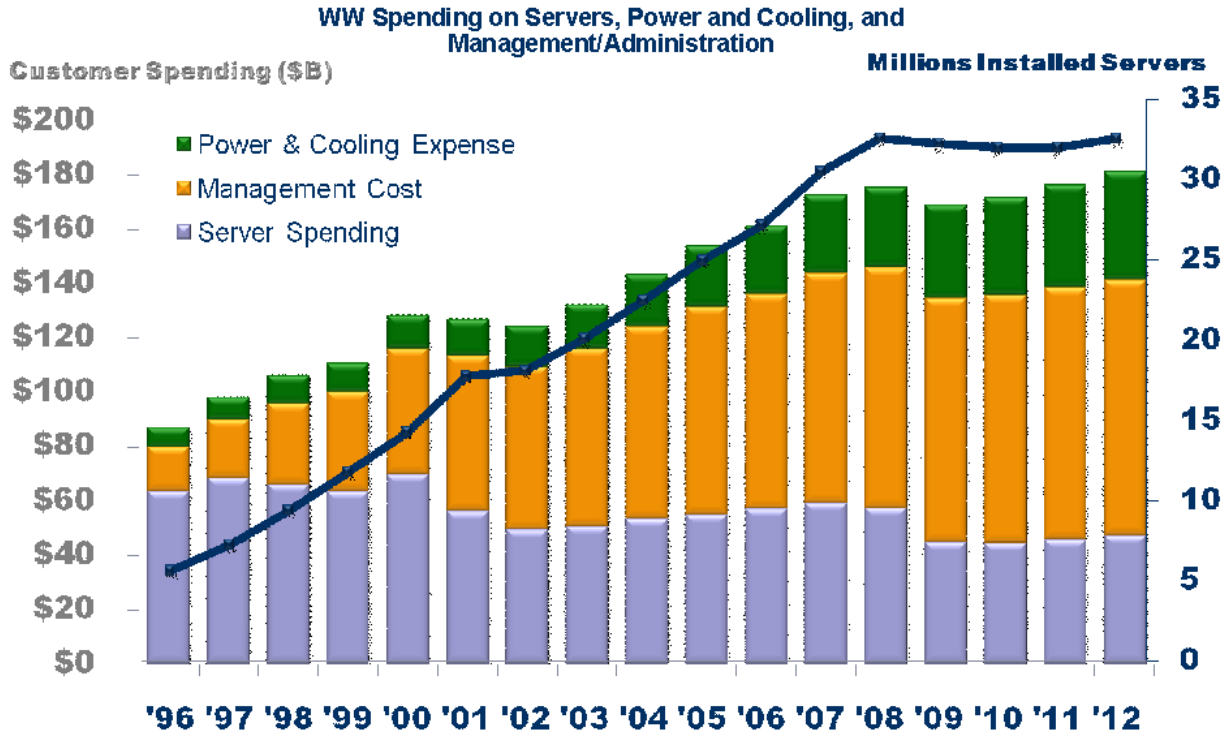

Source: IDC, 2009

Virtual Machine Sprawl Is a Rising Datacenter Cost

IDC expects that there will be more than 100 million virtual and physical servers installed by 2013 (see Figure 2). To date, most customers have underinvested in systems management and automation tools relative to the investments that have been made in virtualization adoption. This has meant that many datacenters still employ manually intensive processes, including the integration of service management frameworks, resulting in greater burdens on staffing. This manual intervention often means that systems are disconnected from business processes and the ability to adapt to change in response to the business is hampered. IDC expects that IT management costs will actually rise if customers don't significantly improve automation capabilities for their virtualized environments (see Figure 3).

FIGURE 2

New Economic Model for the Datacenter: Shifts to Automation Tools Are a Requirement

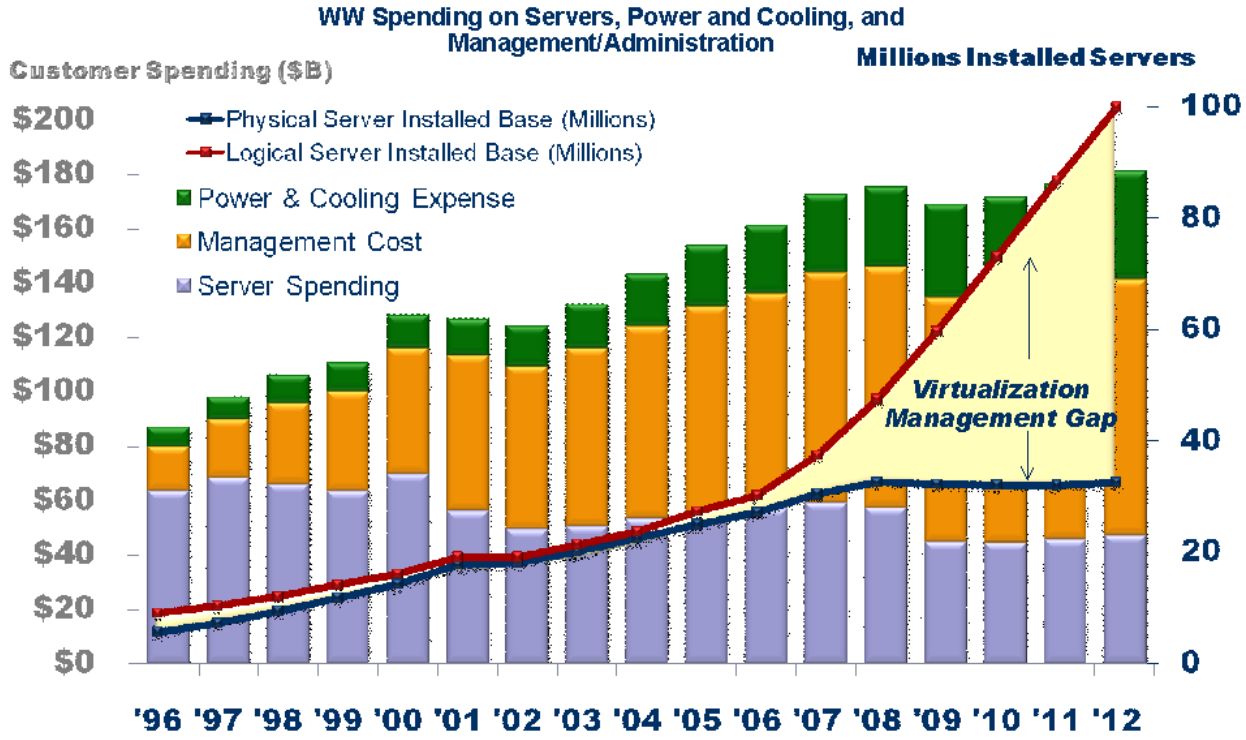

Source: IDC, 2009

FIGURE 3

New Economic Model for the Datacenter: Shifts to Automation Tools Are a Requirement

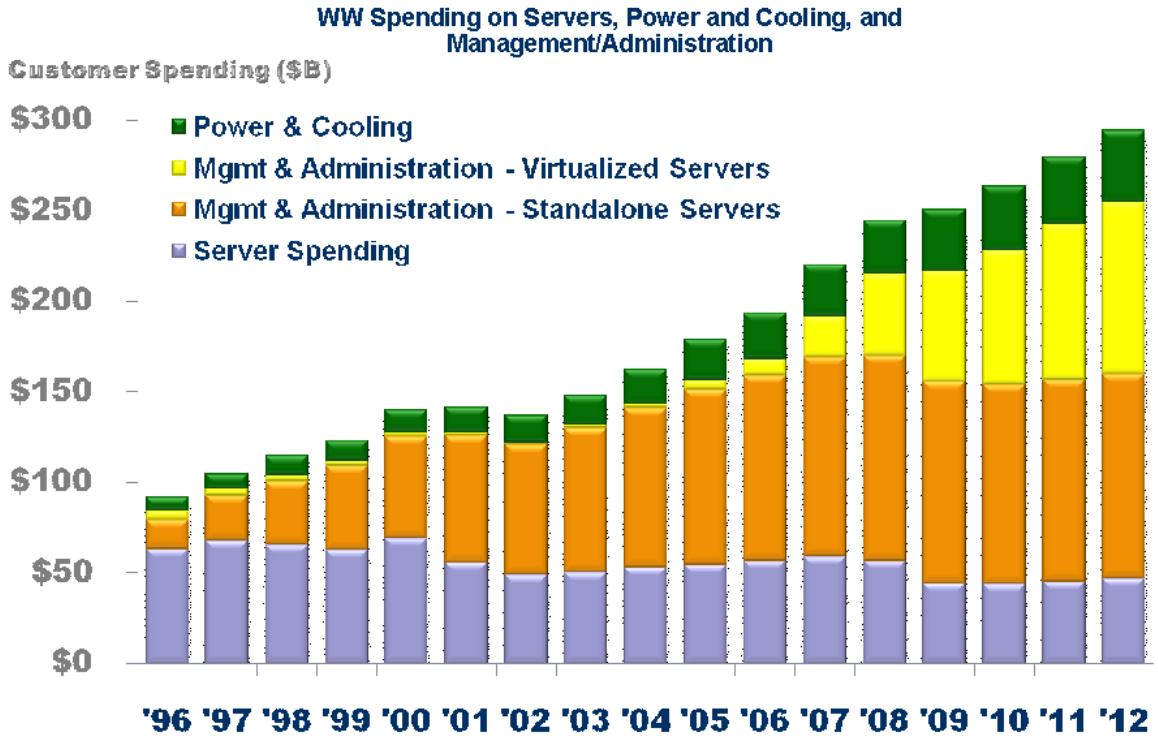

Source: IDC, 2009

Aging Datacenter Facilities Hamper Future Growth

The average age of a datacenter in the United States is 12 years, meaning that the typical datacenter was not built to support the sprawling systems environment of today. Power and cooling has become the number 1 operational problem for the datacenter. As a result, the blueprint for datacenter facilities is changing. Ultimately, IDC sees a new blueprint for the future datacenter facility that is scaled by constructing modular building blocks rather than traditional large-scale buildouts. Containerized datacenters are a new option for modularizing the datacenter and shortening time to market; customers can literally have a prepopulated datacenter delivered to their site in a shipping container within a matter of weeks and can bypass the design, build, construction phase.

COMBATING MASSIVELY MODULAR TECHNOLOGIES

Over the past 15 years, customers have traditionally sourced multiple products from a variety of suppliers as they construct their IT solutions, often acting as their own systems integrator in mixing and matching servers, storage, networking, and systems software to optimize the performance of their applications. Today, IT organizations can select technologies across a broad spectrum of solutions, from best-of-breed to low-cost or "good enough" alternatives. Until now, customers have expected their suppliers to provide them with multiple options across their technology portfolio, and as a result, the IT market has evolved to a point where technology has become "massively modular."

Battling Complexity to Facilitate IT as a Service

Thousands of hardware devices and software solutions are available with standards and APIs built to facilitate the interoperability of these components. As the number of systems and devices continues to proliferate, IT organizations are faced with the challenge of selecting appropriate technologies and making them work within their existing environment. For some, this process has reached a point where it is too complex, requiring sophisticated staffing expertise and significant lead time for deployment and testing. Traditionally, server, storage, and network systems are installed in their own "silos," which has made interoperability, efficiency, and problem resolution a challenge. As virtualization becomes the default standard for the datacenter, these barriers become a greater challenge and inhibit the ability of IT departments to offer IT as a service.

Recently, customers have dealt with these complexities by standardizing on fewer vendors and technologies and by using reference architectures and blueprints. In addition, the adoption of server virtualization has enabled customers to leverage resources across a pool of server, storage, and networking technologies. While these investments have helped to consolidate and rein in IT sprawl as well as set a foundation for a virtualized pool of resource, ongoing challenges remain integrating business policies into the application portfolio and automating a defined set of infrastructure services.

Convergence Is Necessary

IDC believes that we are entering a new business cycle in IT where customers are prepared to trade choice for both ease of installation and simplicity of management. In essence, multiple technologies converging into a single, tightly integrated system that is application aware, automated, and governed by policies simply isn't economically possible without this next phase of converged systems. Convergence not only lowers complexity for customers but also is an enabling vehicle for a shared-service model of computing, one that maximizes hardware utilization, improves availability, contains management costs, and reduces time to deployment.

THE EVOLUTION OF CONVERGED SYSTEMS: A NEW BUSINESS CYCLE FOR IT?

What Is a Converged System?

The term *converged system* refers to a new set of enterprise products that package server, storage, and networking architectures together as a single unit and utilize built-in service-oriented management tools for the purpose of driving efficiencies in time to deployment and simplifying ongoing operations. Within a converged system, each of the compute, storage, and network devices are aware of each other and are tuned for higher performance than if constructed in a purely modular architecture. While a converged system may be constructed of modular components that can be swapped in and out as scaling requires, ultimately the entire system is integrated at either the hardware layer or the software layer.

Convergence Is an Evolution, Not a Disruption

Converged systems are a natural evolution for server technology. Figure 4 demonstrates how, over the past 25 years, modular technologies have eventually become integrated during each technology cycle. Integrated technologies are then later disrupted by a new, lower-cost modular solution that is "good enough" to take on workloads from the more established integrated technology. This process typically takes 10–15 years to mature. It also demonstrates that for more demanding and complex applications to migrate from a high-performance system to a less performant technology, an integration phase is actually essential. This integration can typically be achieved by only a select set of vendors that are market leaders and have a broad set of products and solutions that can be tightly coupled.

For example, workload migration from mainframe solutions to Unix/RISC architectures initially took place for low-end applications such as IT infrastructure and application development along with smaller ERP applications and data marts, thus leaving the mainframe as the mainstay for higher-end transactional applications such as OLTP and batch processing. Over time, we have seen a gradual migration of these higher-end applications to Unix/RISC; however, vendors had to integrate the operating systems with the hardware platform as well as create a tighter relationship between software tools, middleware, and the application stack to ensure reliability and scalability of complex workloads. Creating these interdependencies is the fastest (and possibly only) method to penetrate higher-tier applications and typically can be done only by a single vendor or an extremely tight partnership among a very limited number of vendors.

FIGURE 4

Evolution of Platforms and Applications

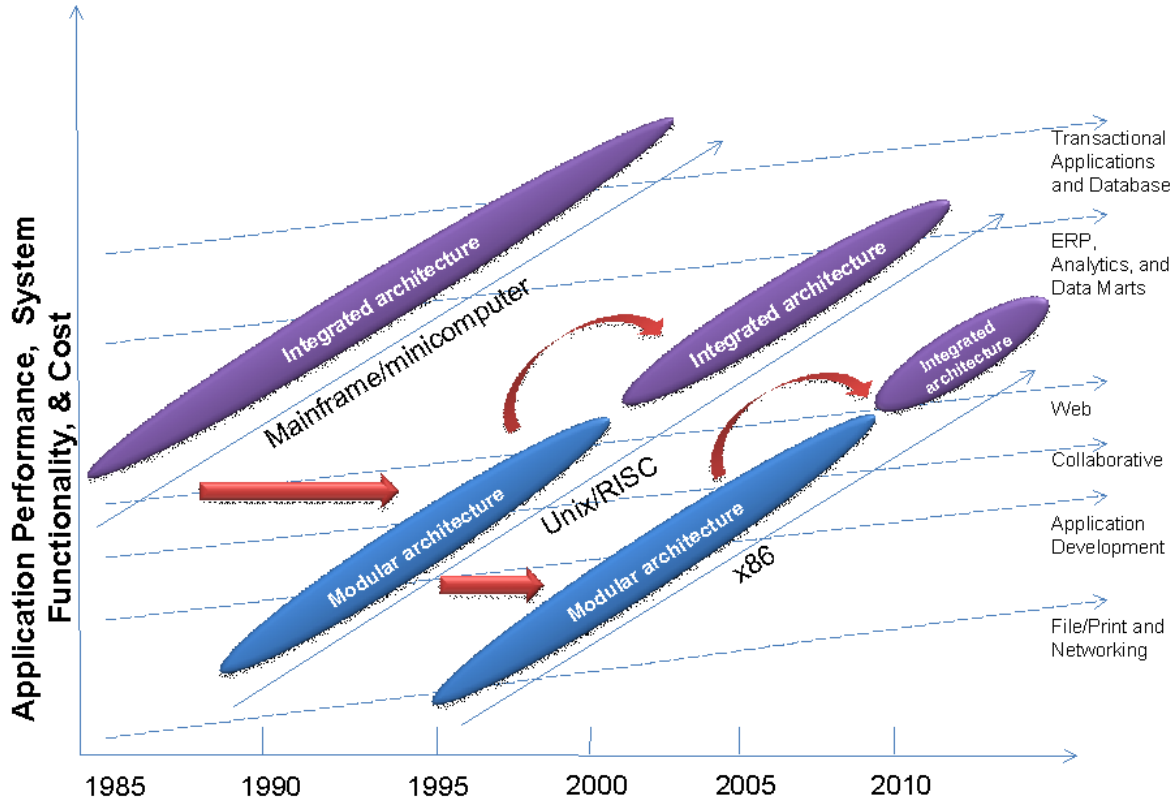

Source: IDC, 2009

IDC believes we are seeing the same technology cycle within x86 hardware. Much of the "low-hanging fruit" from Unix migration has already been taken, and in order for x86 architecture to take on higher-end Unix workloads as well as support a rapidly growing virtualized server installed base, the move to an integrated platform is actually necessary.

This next phase in IT is evolutionary, not disruptive. Newer converged architectures will need to coexist with installed technology and support datacenter strategies that have been in place for the past several years, particularly with respect to server virtualization, consolidation, and addressing power and cooling constraints. Again, this is not a disruption for clients but a next step in improving x86 performance and driving down the cost of computing.

Why Converged Infrastructure?

Ideally, IT organizations want to create a dynamic environment where infrastructure can automatically grow, shrink, and redeploy at an optimized price point or monthly

fee, as business demands change. While this has long been a vision, the realities of building in this type of flexibility are difficult to achieve. Virtualization has become the start for many companies looking to lower hardware cost and create a more dynamic and mobile infrastructure. Virtual machines can be moved from one physical server to another; however, this requires a shared storage environment and a high-performance, predictable network architecture. Building an environment where network policies, security policies, and application SLAs move with the virtual machines across servers, storage devices, and networks requires a tightly coupled set of infrastructure, management tools, policy tools, and operating environments. Without this coupling, application visibility is limited and the ability to dynamically reallocate resources is compromised.

Converged infrastructure also brings the benefits of faster deployment time (relative to standalone hardware), improved availability with preset SLAs, and lower management costs through the use of automation tools. The true benefit is not so much in the hardware but in the agility of the environment that offloads the manual tasks of staff, ensures a predictable operating environment, and shortens time to business value.

Private Clouds Versus Public Clouds

IDC believes that the short-term benefits of a converged infrastructure lie in the ability of an organization to build its own set of infrastructure that offers cloud computing–type services within the walls of its own datacenters. These "private clouds" of computing provide a service-based approach to IT delivery yet are able to manage existing security and compliance issues that are not yet fully covered with external providers offering shared infrastructure. Over time, as service providers build their shared infrastructure environments, IT organizations may choose to federate some of their applications and data to a public cloud or outsource entirely with SLAs for security, availability, and performance.

Requirements for a Converged System

What's new with converged infrastructure?

Customers should look for several core elements that are required in a converged system:

- **Built-in virtualization capabilities that are leveraged across the entire converged system to improve utilization, drive efficiencies, and enable a more dynamic infrastructure**
  - IT organizations continue to move toward virtualized solutions due to the improvements in hardware utilization and lower cost. During the next phase of server virtualization, customers are seeking to drive up VM densities to achieve maximum leverage on their installed systems. Key to this goal is rich hardware in terms of processing and memory capabilities. Virtual storage and networking are needed to optimize the broader infrastructure with the explosion of virtual servers. Storage and networks need new levels of flexibility and integration to enable virtual environments to continue to grow without creating new problems or operational costs.
- **Integrated, flexible network components to enable VM mobility**

- Ongoing server virtualization has real implications for adjacent network and storage environments. For customers to get any real benefit from server virtualization, the ability to move virtual machines from one physical system to another around the datacenter is essential. Without mobility, customers are faced with limited backup and maintenance windows and cannot invoke high-availability tools that are essential in this type of consolidation to avoid a failure of having all their "eggs in one basket." A shared storage environment such as SAN or NAS is a technology requirement for VM movement, and a reliable, high-performance network is critical to ensuring predictable performance. In addition, the ability to support a variety of network protocols, such as Ethernet, Fibre Channel, and FCoE, is critical as customers deploy a multilayer networking strategy and balance cost against performance.

Flexible storage attachment options for rightsizing terabyte costs and storage tiers

- Today's IT organizations deploy a variety of storage solutions. Robust Fibre Channel SAN and NAS solutions are increasingly being balanced against lower-cost iSCSI solutions, and alternative storage media such as SATA and Flash are balanced against traditional disk solutions for tiered architectures. A converged system requires a flexible storage attach environment for customers to continue to rightsize their storage strategies.

Inherent resiliency for uptime and application performance to support mission-critical applications

- Converged systems by design will support hundreds (if not thousands) of virtual machines and their applications. In addition, larger unvirtualized applications such as enterprise databases and transactional applications require a robust set of hardware and uptime guarantees in excess of 5-9s. Converged systems will inherently offer this type of continuity based on their underlying architecture of tightly coupled systems with automated failover capabilities and redundant system components for both virtual and bare metal applications.

Built-in automated policy management and orchestration tools for both deployment and ongoing application management

- IT organizations have long sought a panacea where their infrastructure flexes and shrinks as business demands change. This goal remains elusive and will be impossible without the ability to match hardware capacity and application performance against business objectives. At the heart of this transition is an automated policy management system that allows for business priorities to be mapped to applications. As the requirement for more resources or new deployments is triggered, automated tools perform these tasks for the systems administrator. This includes provisioning a full set of services as application requirements based on a predefined set of policies that encapsulate required hardware resources, application response time, network connections, and storage access. This step requires the integration of both IT best practices and more sophisticated management tools.

Energy management and power-capping capabilities that span the system and the datacenter facility

- Given the major challenges of today's aging datacenter facilities, the ability to measure and manage energy requirements and balance power availability against application requirements can be critical to some IT organizations that are close to their threshold of power and cooling capacity in the datacenter. Such a system not only will require the ability to balance total system power consumption against processor speeds but also must be able to tie into installed building management systems that feed information to in-row and perimeter infrastructure systems such as PUDs, UPSs, and CRACs.

Predictability and repeatability

- A converged infrastructure should take the guesswork out of systems deployment and operation. Mature technologies combined with a sophisticated management suite that automate today's manual systems administrator tasks are a table stakes requirement. Without a converged system, performance, availability, and security are only as good as the weakest link in a chain of modular systems.

A shared resource environment that enables dynamic allocation of resources as application requirements change that create infrastructure-as-a-service solutions

- While virtualization is at the heart of helping IT organizations pool their IT resources, offering IT as a service requires far more than just aggregated systems with virtual machines. Truly turning infrastructure into a service offering depends on a rich set of system management tools and hardware that is application aware and can automatically allocate and reorient resources based on application demands.

HP'S CONVERGED INFRASTRUCTURE SOLUTIONS

As a new cycle for enterprise technologies emerges, IT organizations will find that they need to think differently about not just their systems but also their operations, policies, processes, people, and strategy. This type of transformation requires not only that IT groups be able to take a holistic view of technology as servers, networks, storage, and management tools converge but also that suppliers be able to offer and bring together a broad set of technology solutions. HP can draw upon its entire technology portfolio in creating a converged infrastructure solution from the server to the datacenter facility itself.

HP has created an architecture that brings together four essential elements of its systems hardware, software, and services portfolio, as shown in Figure 5.

FIGURE 5

HP's Converged Infrastructure Solutions

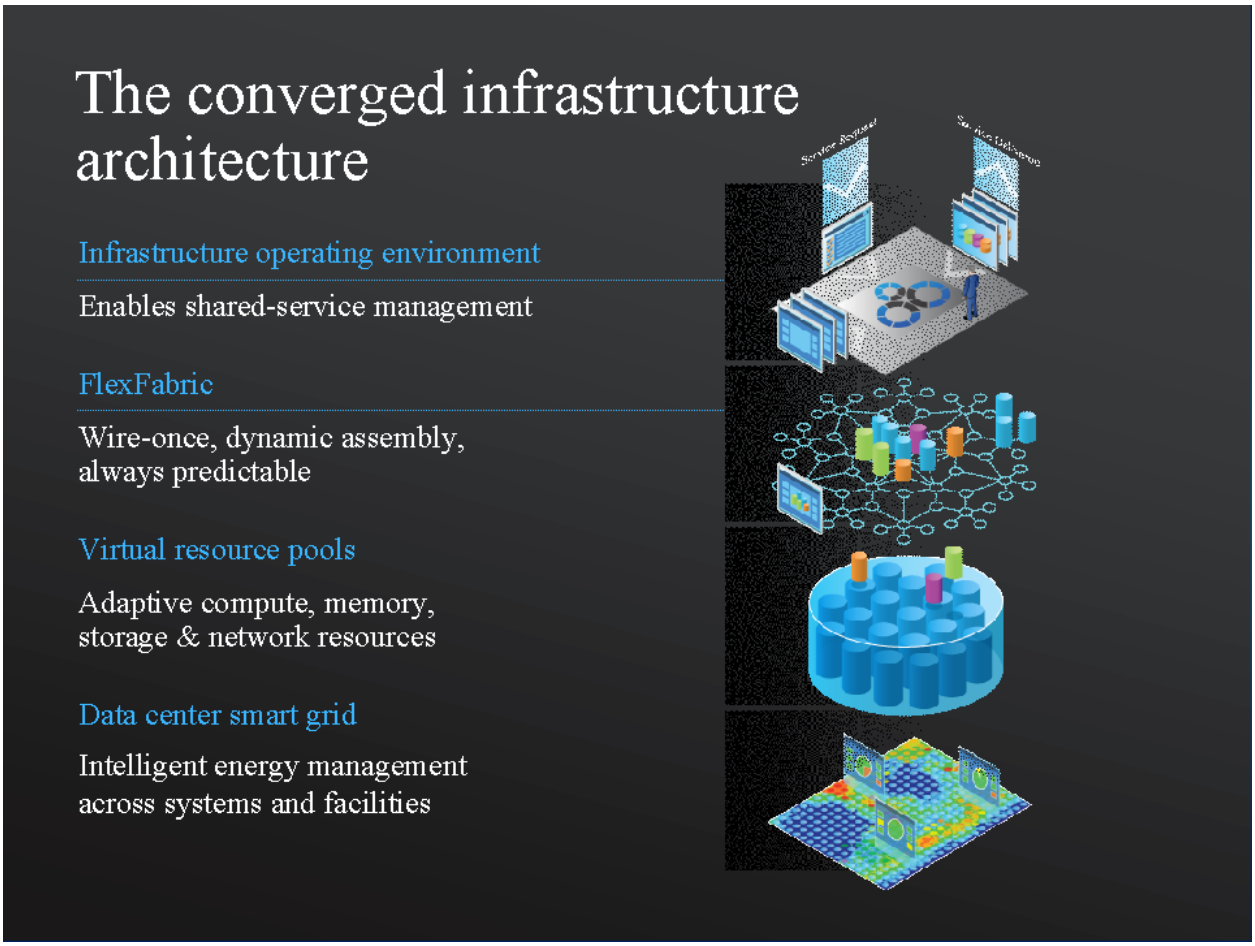

Source: HP, 2009

HP Infrastructure Operating Environment

The HP Infrastructure Operating Environment brings together a set of application management tools that help customers quickly deploy new infrastructure solutions as business demands change. These tools manage applications based on business priorities and policies for rapid provisioning of IT elements that are ultimately delivered as a service. This set of tools manages the infrastructure across the life cycle of an application for deployment, maintenance, resource provisioning, high availability, and business continuity.

HP FlexFabric

HP FlexFabric virtualizes network systems and protocols so that customers can choose network connectivity based on application demands without having to hardwire network components to single devices. HP FlexFabric consolidates Ethernet and Fibre Channel architectures and ultimately means that customers require fewer physical connections by virtualizing at the access layer and allowing for shared network connections over a

single physical connection. HP FlexFabric utilizes the HP ProCurve and HP Virtual Connect technologies, which are primarily built to support virtualized server environments and allow for predictable performance as virtual machines are moved across the network, particularly with respect to quality of service and security.

HP Virtual Resource Pools

HP Virtual Resource Pools are a combination of HP's bladed ProLiant, Integrity, StorageWorks, and ProCurve systems in a shared services architecture to create a virtualized set of server, storage, and network systems. This combined set of hardware is built for maximizing the utilization of hardware components and is the foundation for rapidly deploying new applications across a pool of resources, particularly in support of virtual machine technology.

HP Data Center Smart Grid

The HP Data Center Smart Grid optimizes customers' power and cooling environments by combining information from hundreds to thousands of sensors across the datacenter in real time. This environmental monitoring occurs at the system, rack, and facility levels to help customers not only visualize potential areas of concern but also monitor the total capacity of the datacenter. These tools help customers extend the life of their datacenters by taking action based on this large set of information through the use of power capping, in-row cooling, and variable frequency drives.

CHALLENGES

The greatest challenge for adoption of converged systems is not technology but a change in thinking. To garner full benefit from converged systems, IT organizations will have to think differently about procurement and be satisfied with a smaller set of qualified vendors. This is counter to traditional procurement that has at least a wider list of vendors and typically multiple product options. However, HP understands that customers need to evolve from within their existing IT investments and across their heterogeneous environments. This is why the HP converged infrastructure strategy and architectural approach is based on modular, standards-based components with tight, open integration with partners. Converged systems also require that the IT strategy link to the goals of the business. This type of purchase requires a commitment to a singular architecture, and the needs of the business must be served on an ongoing basis with this investment. At the same time, customers should develop a policy-based automation scheme that exploits the full benefit of mobility across the system and prioritizes application needs relative to business outcomes.

Converged technology requires more interaction between previously siloed technology groups, IT departments, and even business users. This is important in a service-centric IT infrastructure if it is to deliver better business outcomes and typically requires a more complete set of governance rules. For example, networking and server groups would have to agree on this architecture, and facilities and IT departments would have to be aligned on the goals for energy management. The advantage to an integrated technology approach is that it does help drive a consensus around application priorities and goals for IT metrics, improving communication between teams through a single management tool and framework.

CONCLUSION

As the market moves into this next phase of IT, the ability to lower complexity and improve flexibility will be key, especially as businesses recover from the economic downturn and look toward innovation as a means to drive revenue and profits. For most businesses today, IT is at the heart of bringing new products and services to market, and a flexible infrastructure that not only can support a broad set of applications but also can scale without a linear increase in management overhead will be imperative. Customers need to consider a set of solutions that they will be able to leverage for the next several years and that not only are capable but also will fit in their existing infrastructure.

During this next phase of IT, customers will need to develop a set of governance rules that map back to their application portfolio and link to their IT infrastructure, creating a more defined connection between IT and changes in the business. This process involves not only a technology assessment but also a business assessment and rationalization of the application portfolio. Customers will need to choose their suppliers carefully and should look for a vendor that can provide a holistic set of IT infrastructure, software, and services offerings to bring this future vision together.