WHITE PAPER

Improving Systems Management and Availability with x86 Virtualization

February 2009

PREPARED FOR

IBM

TABLE OF CONTENTS

Executive Summary

Virtualization is clearly having a major impact across the IT industry. It is generating real excitement with users, and it has proven its ability to deliver some fundamental business benefits in a variety of real-world environments. Virtualization is certainly not new technology; it has been used on mainframes and other platforms for decades. However, what has been generating all the excitement recently is the growing maturity of virtualization technology on industry-standard x86-based hardware.

Much of the interest in deploying virtualization on x86 hardware has been driven by the desire to lower costs through improved utilization of computing resources and

IDEAS RECOMMENDATIONS FOR USERS

- **Focus on immediate management priorities with virtualization.**
  - Leverage provisioning and patch management tools and procedures
  - Manage the relationship between physical and virtual resources
  - Optimize remote management capabilities
- **Pursue strategic management goals over the longer term.**
  - Power management
  - Billing
  - Automation
- **Choose the appropriate virtualization platform based on availability needs. Match workloads against the maturity of x86 versus that of other platforms (UNIX, mainframes), and evaluate the tradeoffs between x86 server virtualization platforms.**
  - Determine appropriate service levels for workloads virtualized on x86
  - Weigh the acceptable levels of downtime against the investments required to maintain uptime
  - Categorize potential causes of downtime and appropriate protections
  - Identify hardware-driven versus software-driven availability issues
- **Evaluate the x86 virtualization platform as an integrated solution.**
  - Ensure that the platform can be supported as a complete system, rather than as a collection of components

As a result of management and availability benefits, virtualization can significantly improve the value of x86 systems in enterprise data center environments.

optimized use of resources such as power and cooling. However, when coupled tightly with x86 systems that are optimized for enterprise usage, virtualization also has the potential to deliver significant value in other ways. In particular, virtualization can simplify the management of systems and infrastructures, and it can help improve the availability of services. As a result of management and availability benefits, virtualization can significantly improve the value of x86 systems in enterprise data center environments. To achieve the benefits of x86-based virtualization, however, organizations must deploy x86 virtualization technology in the form of complete systems that tightly couple several components, including:

- x86 server hardware
- Key third-party software for virtualization
- Management tools that are optimized for both virtual and physical components

Because virtualization can fundamentally redefine the way that workloads are deployed, administrators need to be aware of the impact that the deployment of virtualization will have on management and availability. Virtualization has the potential to deliver substantial advantages in these areas, but it may also introduce some challenges. The hardware platform on which virtualization is deployed can play a significant role in determining the extent to which users achieve the possible benefits while overcoming potential obstacles. This paper reviews some of the ways that deploying virtualization on leading-edge x86 servers such as IBM's System x™ and BladeCenter™ platforms, as well as using IBM's management software, can maximize the manageability and availability benefits of adopting virtualization.

IBM has optimized its X-Architecture™ servers – particularly the newest generation – for virtualization deployment. IBM also works closely with third-party providers of virtualization technology to ensure their products work on IBM hardware. Further, IBM adds value with its own system management options that are optimized for virtualization, in particular IBM Systems Director, which can be used to coordinate and manage both virtual and physical resources in the data center with a single set of tools. IBM makes sure that these components are all highly integrated, allowing users to deploy x86-based virtual infrastructure that functions as a complete system.

IBM System x and BladeCenter servers, along with IBM Systems Director and IBM Tivoli® management software, are a powerful complement to the leading x86-based virtual machine (VM) platforms. IBM's fourth generation of System x servers, based on its eX4 technology, offer some of the strongest scale-up capabilities in the industry, with SMP and memory ranges that are matched by no other platform. The impressive scale-up range in System x maximizes the number of virtual machines that can be consolidated on a single system, which makes the best use of resources such as power, cooling, and space. To protect against the greater cost of failure (in the event of unplanned downtime), the eX4 architecture stands out with a number of resiliency features, including redundant I/O, Memory ProteXion, and memory scrubbing and mirroring. Hot-swap memory modules help to avoid the cost of moving large numbers of VMs to other hosts during planned downtime.

IBM Systems Director 6.1 provides a single front-end for controlling both virtual and physical layers of infrastructure.

IBM also adds considerable value for managing virtual infrastructures on System x and BladeCenter. IBM Systems Director 6.1 provides a single front-end for controlling both virtual and physical layers of infrastructure. Having visibility over both layers via the same interface is an important advantage for administrators struggling to deal with the complex interaction between physical and virtual resources. Further, the IBM Systems Director Active Energy Manager plug-in helps users optimize their power consumption by coordinating the power draw of virtual machine hosts with the workloads that are deployed on them. Tivoli Provisioning Manager addresses one of the thorniest management problems associated with virtualization by helping to standardize and automate the installation and maintenance of software on large numbers of virtual machines. Finally, on IBM's BladeCenter platform, administrators can use BladeCenter Open Fabric Manager to virtualize system components that are outside the scope of virtual machines (i.e., the LAN and SAN connections needed to connect virtualized workloads with the outside world).

Background

IT managers making technology deployment decisions today have several key trends to consider. First, the growth in demand for computing resources shows no signs of slowing down. The impact of the economic recession on trading volumes may be uncertain, but companies are likely to increase the computing resources that they dedicate to analyzing risk. Other compute-intensive workloads, such as data mining, are also being driven deeper into organizations as more levels of management participate in identifying new business opportunities from transaction data.

Further, the adoption of service-oriented architectures (SOAs) may result in the need for more scalable services as they become shared by multiple business units. And as the web becomes ever more integrated into business and consumer activities, vast new computing resources are needed to host web applications with boundless, Internet-scale user populations. Organizations need to maintain performance and availability to handle these growing workloads. But at the same time, they increasingly have to confront hard limits in terms of real resources such power, cooling, and space, as they struggle to maintain data centers that were built over 10 years ago under far more modest operating conditions.

From a technology standpoint, the processors from industry-standard suppliers such as Intel and AMD are rapidly increasing their capabilities by introducing refinements such as 64-bit processing, multicore designs, and virtualization hardware. Many organizations, especially on the larger side, perceive opportunities for simplifying their infrastructure by reducing the number of platforms they have to support. In some cases, these organizations have a long-term vision of exploiting the economics of industry-standard technology as much as possible, by consolidating workloads on hardware with x86 processors from Intel and AMD, and operating systems such as Windows and Linux.

As virtualization capabilities on the x86 platform mature, virtualization is becoming a standard part of x86-based IT infrastructure at multiple levels, including processors, server hardware, and operating systems.

Virtualization Benefits

As virtualization capabilities on the x86 platform mature, virtualization is becoming a standard part of x86-based IT infrastructure at multiple levels, including processors, server hardware, and operating systems. Virtualization has already proven its ability to deliver several fundamental benefits, including the following:

- **Consolidation and improved resource utilization.** Consolidating servers with virtualization enables administrators to reduce the number of physical machines that they have to acquire and manage. The reduced hardware footprint results in lower maintenance costs, cooling requirements, and power consumption.
- **Legacy application support.** Virtualization enables administrators to migrate legacy applications to new hardware without disturbing their environment. If an application depends on a particular OS version, migrating that application to new hardware will not be possible if that hardware does not support the OS version. Virtualization can be used to extend the life of such applications. The OS required for the legacy application can be hosted in a virtual machine running on the new platform until that application can be replaced or rewritten.
- **Improved test and development processes.** Virtualization simplifies and improves the quality of testing and development. IT managers can rapidly allocate resources as needed to support test processes. If necessary, the resources required to perform testing can be assigned by tapping into production resources temporarily. Production systems can be copied easily for testing purposes, and developers can easily test applications on different operating systems. Virtual machines also help to isolate bugs during the testing process by neutralizing hardware variability.

In addition to these basic use cases, virtualization also has the potential to deliver significant value in other ways when it is coupled tightly with x86 systems that are optimized for enterprise usage. In particular, virtualization can help to improve the manageability and availability of systems and infrastructures.

IBM recently introduced a model for efficient use of IT called the New Enterprise Data Center (NEDC), which helps users take advantage of IT to drive business innovation. The model defines paths for organizations to better manage costs, improve operational performance and resiliency, and respond quickly to business needs. Both virtualization and industry-standard x86 systems play a central role in the New Enterprise Data Center approach, and for IBM to deliver the benefits of NEDC on the System x and BladeCenter platforms, it must deliver x86-based solutions as a tightly integrated stack that has been thoroughly tested and supported for use in virtual infrastructure.

The following sections examine how users can benefit from virtualization to improve their systems management procedures and systems availability using a combination of System x and BladeCenter servers; IBM Systems Director and Tivoli management software; and third-party virtualization software platforms.

Indeed, virtualization makes new classes of management possible, such as resource pooling, in which a number of physical servers are treated as a single logical set of computing resources that can be assigned to virtualized workloads on demand.

Virtualization and Manageability

Virtual machines inherently decouple workloads from details about the hardware on which they are deployed. As a result, administrators gain significant flexibility in matching applications with their required computing resources. Moreover, because virtual machines run the same operating systems as physical hardware, they do not introduce any major operational changes. As a result, virtual machines represent a powerful and intuitive mechanism for assigning resources to workloads with a new degree of freedom.

With the right management tools, virtualization can significantly reduce the time required to deploy new systems. Rather than set up new physical systems when users require new computing resources, administrators can maintain a pool of virtual resources that they draw upon. When servers are implemented as virtual machines, they can be created almost instantly through either manual commands or scripts. With physical servers, the end-to-end process of installing a new system might take a month or longer, if one includes the time required to spec the needed system; push the request through an organization's requisition process; order the hardware and have it delivered from the vendor; configure the system; and set it up in a server room. With virtualization, that time can be reduced to hours or less by loading preconfigured software images to a new VM on an existing server. This process allows organizations to respond to business opportunities more rapidly than before.

One side effect of the ease of deployment, though, is that it could significantly increase the number of virtual servers that have to be managed. As users become more comfortable with the ease of deploying virtual systems, they may respond by increasing the rate at which new VMs are created. Hence, administrators may soon have to contend with "virtual machine sprawl," in which managing all the software needed to host workloads on large numbers of virtual machines becomes a significant burden.

The deployment of virtualization in itself does not result in radical changes to system management tools or procedures. From an operational standpoint, administrators generally treat virtual machines the same way as physical servers. Since virtual machines run standard operating systems, most management tools will continue to function as before. However, administrators are likely to add new tools specifically to manage virtual machine hosts, as well as the virtual infrastructure that is deployed on top of virtual machines.

Indeed, virtualization makes new classes of management possible, such as resource pooling, in which a number of physical servers are treated as a single logical set of computing resources that can be assigned to virtualized workloads on demand. The relative ease of controlling virtualized workloads simplifies the process of automating certain tasks, such as capturing the state of an application and migrating it to another host, either in response to failure conditions, or for load-balancing purposes. Power consumption may also be controlled in this manner, by

In IBM's management architecture, IBM Systems Director captures operational data and events at the level of hardware and specific virtualization platforms, and feeds this data at a higher level of abstraction to Tivoli, which can be used to manage advanced functions such as failover, disaster recovery, and maintenance of Service Level Agreements (SLAs).

migrating off lightly utilized hosts so that they can be shut down to reduce overall power draw.

Of course, at the same time as organizations adapt to new virtualization capabilities, they must continue to address traditional management concerns related to issues such as user identities and security; service level agreements; asset management; and governance (i.e., regulatory compliance and strategic accountability). Some of these operational aspects may be impacted by the adoption of virtualization.

Key Issues with Virtualization Management

Some of the key technical issues related to virtualization management include the following:

- **Provisioning and patch management –** Installing, maintaining, and patching the software on virtual machines as their number increases.
- **Visibility over resources at multiple tiers –** Understanding the relationship between physical and virtual resources in order to properly maintain service levels in virtualized workloads.
- **Remote management –** Maintaining secure access to virtual machines, which are not physically connected to any peripherals and thus can x only be managed remotely.
- **Downtime –** Minimizing downtime due to server hardware and I/O via the ability to reconfigure hardware and I/O with minimum disruption.
- **Energy management –** Taking advantage of consolidation and load balancing to optimize power consumption in virtual machine hosts.
- **Billing and metering –** Tracking consumption of shared virtual computing resources and, where appropriate, allocating costs to users on a chargeback basis.
- **Heterogeneous VM platform and operating system support –** With competition accelerating in the x86 virtual machine market, it may be necessary to support multiple virtual machine hosts, while the virtual machines themselves may be running multiple classes of operating systems on the same hosts.
- **Automation and resource pooling –** Treating multiple servers as a single pool of resources that can be tapped dynamically and automatically in response to changing workload conditions or downtime events.
- **Security management –** Ensuring applications and data remain protected, even as they are accessed through purely virtual paths.
- **Managing an integrated environment –** Being able to treat the following combination as a complete system in which all the components work together seamlessly: hardware based on industry-standard processors and operating systems; native system management software; and third-party virtualization software.

The combination of IBM Systems Director and Tivoli can thus be used to address most of the technical issues with virtualized infrastructures on an end-to-end basis, spanning multiple server and virtualization platforms.

IBM's Approach to Virtualization Management with System x and BladeCenter

Systems management involves an extraordinarily broad range of challenges. The introduction of virtualization may bring on yet additional management requirements, even as some management activities are simplified by the use of virtualized resources. In general, IBM addresses management through two classes of products (see Table 1):

- **IBM Systems Director** is optimized for hardware-specific management tasks on various server platforms. IBM Systems Director can be used to manage environments with a mix of physical and virtual systems using a single, consistent interface that covers a broad range of operational details. In addition to controlling low-level functions in IBM servers, IBM Systems Director can be used to manage IBM POWER® processor-based systems, BladeCenter, System x, IBM System z®, and IBM System Storage™ servers via the same interface. Systems Director 6.1 can also manage hardware from other suppliers, and it can manage multiple hypervisors and virtual machine platforms at a detailed level.
- **IBM Tivoli** is a systems management platform that provides integrated visibility, control, and automation across multiple business units and heterogeneous server platforms. Tivoli is optimized for aligning IT operations with business at a high level, while enabling governance and control of automated business processes independently of specific hardware attributes.

TABLE 1. Positioning IBM Systems Director vs. IBM Tivoli Capabilities

System x/BladeCenter Management	Enterprise IT Management
» **IBM Systems Director 6.1** converges management of physical hardware and virtual machine platforms	» **IBM Tivoli** is optimized for comprehensive service-level management » Integrates with IBM Systems Director to provide a total solution
Benefits	**Benefits**
» Enables detailed control of IBM systems » Ensures that virtual and physical layers can be managed with a single interface » Combines management of multiple virtualization environments based on x86, UNIX, or mainframes (simplification helps to avoid admin-driven errors and improve standardization which can lower costs) » Can be used to manage both IBM and non-IBM x86 platforms	» Controls and automates heterogeneous IT infrastructures » Enables integrated monitoring and event management across the organization » Provides billing and chargeback capability on both physical and virtual machine usage, enabling inter-departmental sharing and compliance tracking

IBM has optimized System x and BladeCenter servers for remote management at multiple levels.

In IBM's management architecture, IBM Systems Director captures operational data and events at the level of hardware and specific virtualization platforms, and feeds this data at a higher level of abstraction to Tivoli, which can be used to manage advanced functions such as failover, disaster recovery (DR), and maintenance of Service Level Agreements (SLAs). The combination of IBM Systems Director and Tivoli can thus be used to address most of the technical issues with virtualized infrastructures on an end-to-end basis, spanning multiple server and virtualization platforms.

Provisioning and Patch Management

One management challenge that arises quickly with broad virtualization deployments is *provisioning* virtual machines (i.e., installing and maintaining the necessary software on virtual machines so that they can perform a particular workload). The software needed for a workload typically includes an operating system, middleware, application software, and virtual infrastructure, all of which have to be patched on occasion. The burden of maintaining software in virtual machines increases significantly as the number of VMs grows. Such growth sometimes results from "virtual machine sprawl," which arises as a result of the relative ease with which VMs can be created.

To meet the challenge of provisioning software in virtualized environments, administrators require the ability to create "master images" of virtual machines using class-based templates that describe generically what software is needed on a particular VM. Once the template is defined, administrators can rapidly instantiate a virtual machine based on that class, which automatically installs all of the software itemized in the template. As a result, administrators can rapidly set up large numbers of VMs with relatively few commands.

IBM Tivoli Provisioning Manager is a provisioning platform that can be used to manage software changes by keeping server software compliant with administrator-defined policies. Tivoli Provisioning Manager is built on an SOA to help maximize the alignment between IT infrastructures and business operations. It provides a toolkit that can be used to build customized automation packages that tailor servers to particular environmental requirements, or IT/business processes. IBM Tivoli Provisioning Manager works with both virtual machines and physical servers. On System x and BladeCenter, Tivoli Provisioning Manager can thus be used to automate the deployment of VMware virtual machines.

Maintaining Visibility over Physical and Virtual Resources

The deployment of virtualization can introduce yet another layer of complexity to management burdens that may already be challenging. Virtualization management tools are often optimized for the control of virtual resources, and the hypervisors that manage those resources. However, on occasion, administrators may need to gain insights into the behavior of other components of the system in order to maintain overall service levels. In particular, the behavior of hardware on which

To assist in optimizing power usage for System x and BladeCenter servers, IBM Systems Director Active Energy Manager can be used to measure, monitor, and manage System x and BladeCenter energy components.

virtual machines are hosted can have a major impact on the resources that are required by a virtual machine.

Virtualization management tools that are integrated with hardware management tools can therefore provide a more complete picture of the state of a virtualized workload, and may also provide greater control over its behavior. When service levels are not being met, administrators need visibility over the relationships between physical and virtual resources to help them understand which real resources are being consumed by virtualized workloads. Understanding the behavior of virtualized workloads thus requires visibility of both physical and virtual resources, preferably from a single interface.

IBM Systems Director 6.1 provides a single interface for controlling multiple levels of virtualized and physical infrastructure. In addition to providing tools for managing the hardware-related details of System x and BladeCenter servers, IBM Systems Director integrates with popular VM management software such as VMware vCenter Server (formerly called VMware VirtualCenter). This integration allows administrators to coordinate the use of virtual and physical resources from a single interface.

Remote Management

Networks are becoming an increasingly critical component of IT infrastructures. The importance of remote management continues to grow as administrators try to minimize visiting the physical location of servers to perform maintenance. Because virtual machines have no physical connection with servers or peripherals, they are inherently managed remotely, further increasing the importance of management tools that work well over a network. Moreover, virtual machine deployments are frequently coupled with the installation of blade servers and initiatives such as disaster recovery, each of which imposes its own requirements for remote manageability.

IBM has optimized System x and BladeCenter servers for remote management at multiple levels. At the lowest level of System x and BladeCenter hardware, peripherals such as the Remote Supervisor Adapter II and Advanced Management Module, respectively, provide a means for administrators to maintain servers “out-of-band” (i.e., even if they are disconnected from the primary network because the operating system or hypervisor has not been booted). IBM Systems Director is designed to handle both in-band and out-of-band functions from the same interface, allowing administrators to standardize on a single tool for all kinds of management activities. Multisystem management is facilitated with Tivoli Provisioning Manager for OS Deployment (TPMfOSD), which helps to automate the maintenance of the Basic Input Output System (BIOS), the lowest level of software in server hardware. IBM also provides a Management Module Web GUI, which is optimized for ease of use, and is thus ideal for delegating some administrative tasks to personnel with limited IT training.

. . . tools such as IBM Tivoli Intelligent Orchestrator can be used to automate a variety of management processes related to servers, operating systems, storage, middleware, applications, and network devices.

Billing and Metering

The deployment of virtualization often results in infrastructure that is shared by different business units or departments, and users may demand accounting of resource consumption for chargeback purposes to make sure they are getting their "fair share" of virtualized resources. IBM Tivoli Usage and Accounting Manager can be used to accurately track resources in shared computing infrastructures. The tool helps administrators measure the costs of shared resources, allowing costs to be allocated to departments or users. IBM Tivoli Usage and Accounting Manager Collects detailed IT resource usage data for multiple virtual machine and operating system platforms, including Microsoft Windows, Linux, and VMware.

Energy Management

Energy optimization has become a key goal for many organizations. In some cases, organizations are under corporate mandates to become more "green." In other cases, organizations are starting to hit hard limits in the power and cooling resources available to data centers that were designed to host systems with far lower consumption. Virtualization can help organizations take better advantage of available power resources by enabling workloads to be placed optimally on servers in a way that minimizes the draw on power. To achieve this benefit, it is necessary to have tools that can precisely measure the power that is being consumed by particular virtual machine hosts, and then balance virtual machines across the hosts to maximize the utilization of computing resources.

To assist in optimizing power usage for System x and BladeCenter servers, IBM Systems Director Active Energy Manager can be used to measure, monitor, and manage System x and BladeCenter energy components. The tool performs power and thermal trending analysis, and enables administrators to define policies for power capping and putting servers into power-saving mode. The data produced by IBM Systems Director Active Energy Manager can also be exploited by Tivoli Monitoring and Tivoli Usage and Accounting Manager, allowing power-based chargeback policies to be established.

Heterogeneous VM Platform and Operating System Support

Competition in the x86 virtualization market is heating up, and customers have a choice of several viable x86-based server virtualization platforms. Most of the leading x86-based virtual machine platforms offer tools for managing various aspects of virtualization, which x86 system providers must support in order for their customers to productively deploy the platform.

System x and BladeCenter have first-tier support for leading x86 and IBM POWER operating system platforms such as Windows, Linux, and AIX®. Further, IBM has close relationships with key virtual machine software suppliers such as VMware, Virtual Iron, Microsoft (Virtualization Server and Hyper-V), Red Hat, and Novell. Some System x and BladeCenter servers even offer an embedded a hypervisor in a flash memory module so that servers can quickly and easily boot up as virtual machine hosts (currently, VMware's ESXi hypervisor is supported in flash memory).

IBM System x and BladeCenter servers have several features and options that help to protect the security of virtualized workloads.

IBM Systems Director 6.1 provides "single pane-of-glass" management for multiple virtualization platforms, including VMware, Microsoft, Xen, IBM PowerVM™, and IBM z/VM® virtualization infrastructure, which has the potential to promote standardization and lower administrative costs.

Automation and Resource Pooling

As users extend the scope of virtualization from single servers to multiple systems throughout a data center or organization, they may become interested in automating some virtualization management tasks. Coupling basic virtualization functions with higher-level management tools that are specifically enabled to automate the management of multiple virtual systems creates a foundation for "virtual infrastructure." Such virtual infrastructure allows multiple virtualized systems to be treated as a flexible pool of resources that can be tapped dynamically in response to changing workload conditions or downtime events.

IBM helps to enable virtual infrastructure by providing tools for automating virtualization management at multiple levels. On IBM's BladeCenter platform, administrators can use BladeCenter Open Fabric Manager to automatically move an entire virtual machine environment, including the hypervisor and all of its VMs, from one blade to another, based on event-driven policies. IBM Systems Director 6.1 can trigger the migration of individual virtual machines from one host to another based on low-level hardware events, working with the native management systems of particular virtual machine platforms to perform the migration. Finally, at the level of enterprise IT infrastructure, tools such as IBM Tivoli Intelligent Orchestrator can be used to automate a variety of management processes related to servers, operating systems, storage, middleware, applications, and network devices.

Security Management

Virtualization can affect security in several ways. On the one hand, relying on virtualization to host critical workloads potentially introduces new security concerns, since it may introduce new paths through which breaches can occur. On the other hand, virtual machines make it possible to control access directly at the level of the hypervisor – for example, by keeping data behind closed data center doors, rather than on desktop drives in exposed cubicles – potentially providing a powerful way to monitor and regulate traffic to and from virtual machines.

IBM System x and BladeCenter servers have several features and options that help to protect the security of virtualized workloads. First, the ability to boot compact, embedded hypervisors directly from flash memory significantly reduces the attack surface for potential incursions. At a higher level, the IBM Tivoli portfolio includes a broad range of tools to maintain enterprise-level security, addressing identity and access management; data and information security; and application security. IBM's Virtual Infrastructure Access service helps businesses use virtualization to secure applications and data. Finally, looking forward, IBM is releasing core software technology developed in its research labs to help users maintain security as they virtualize broad segments of their infrastructure, such as entire data centers. As

Management tools play a key role in helping administrators maintain uptime.

part of this effort, IBM is developing secure hypervisor (sHype) technology. sHype is a policy-driven security architecture that enables administrators to separate virtual resources into different "domains" that can coexist in the same physical infrastructure without concerns that users of one domain will access resources from another. IBM is also working on extending the industry's Trusted Computing initiative into the virtual server realm.[1]

Virtualization and Availability

Virtualization can be used to fundamentally improve the overall reliability of a computing infrastructure. Virtualization enables fewer physical servers to be deployed, which reduces the footprint for potential hardware failures that result in unplanned downtime. Moreover, the servers that are deployed can be configured with high availability (HA) features (such as redundancy and hot-plug components) to reduce downtime. The ability to migrate virtual machines from one host to another with little or no interruption to processing provides yet another means to reduce planned downtime. Such migration allows workloads to be temporarily moved so that hardware maintenance can be performed on the hosts with minimal disruption. When coupled with HA clustering functions, virtualization can be used to restart workloads on a backup host in the wake of a primary host failure – dramatically simplifying the implementation of DR procedures. Traditional HA and DR implementation requires applications and their dependencies to be adapted so that they can be restarted on backup systems – which is a notoriously complex and error-prone process. With virtualization, the entire workload can easily be relaunched simply by restarting the virtual machine on which it is hosted.

Key Dependencies for Maximizing Availability When Adopting Virtualization on x86

To fully achieve the availability improvements that virtualization promises, administrators must address a number of deployment requirements (see Table 2). First, the scalability and reliability of the hardware platform itself becomes much more critical, since individual servers will typically host multiple consolidated workloads rather than single instances of workloads, which increases the impact of failure. Next, storage management plays a critical role in maintaining uptime for workloads. While virtualization simplifies the task of ensuring that workloads have the necessary computational resources, the virtualization platforms and storage subsystems must also be integrated to ensure that virtualized applications have reliable access to data.

Management tools play a key role in helping administrators maintain uptime. Further, x86 servers must fully support the software-based mechanisms in the leading x86-based virtual machine platforms for availability and disaster recovery.

[1] With Trusted Computing, hardware is equipped with specialized Trusted Platform Modules (TPMs) that provide a "root" level of trust on which other software-based levels of trust can be layered.

The IBM System x and BladeCenter platforms are fully optimized to meet both scale-up and scale-out requirements.

Finally, the combination of server hardware and virtualization software should be offered and supported as a complete solution, whereby the supplier takes responsibility for helping users maintain uptime of the entire environment, rather than just individual components.

The following sections examine how IBM addresses these requirements with its System x and BladeCenter servers.

TABLE 2. Impact of Hardware Platform on Availability of Virtualized Workloads

Requirement	Importance to Availability
Server Hardware Scalability	» Provide sufficient computing resources within a single server footprint » Optimize the use of real resources (space, power, and cooling) » Deliver resources for the most demanding workloads when needed » Support both scale-up (consolidation) and scale-out (virtual infrastructure) deployment
Server Hardware Reliability	» Minimize planned and unplanned downtime to sustain multiple workloads that depend on a single system
Storage Integration	» Ensure that virtualized workloads have continuous access to data
Management Integration	» Provide a single point of management for multiple tiers (hardware, virtualization, software)
Leverage Industry Ecosystem	» Fully exploit the availability capabilities of leading virtualization software platforms (live migration, DR functions)
Solutions Focus	» Optimize for availability in complete systems (hardware + virtualization software + management functions)

Scalability Impact on Availability

The first consideration for maintaining availability with virtualization is to make sure that virtualized workloads always have sufficient resources to provide the necessary service levels, even under variable demand. This problem must be addressed in two ways. First, to maximize consolidation ratios and make optimal use of real resources such as space, power, and cooling, “scale up” Symmetric Multiprocessing (SMP) configurations should be exploited to achieve as much performance as possible within a single server footprint. Also, to take full advantage of virtual infrastructure for meeting performance requirements, in which service levels are maintained by spreading workloads across multiple hosts (so that no single host becomes a performance bottleneck), the chosen server platform

For users who want to scale out using virtual infrastructure, IBM's BladeCenter platform provides an ideal form factor in which to deploy multiple virtual machine hosts with an optimal footprint.

must also be able to "scale out" (i.e., offer options that allow multiple servers to be deployed and managed in unison with maximum efficiency).

The IBM System x and BladeCenter platforms are fully optimized to meet both scale-up and scale-out requirements (see Table 3). The eX4 architecture underlying IBM's newest generation of System x enterprise servers is based on a unique chipset developed by IBM that enables an extraordinary SMP scale-up range in individual servers. System x3850 M2 and x3950 M2 servers can have as few as 4 sockets or as many as 16 sockets. When these sockets are populated with the newest x64 processors containing 6 cores per socket, a single System x server can have as many as 96 cores. Moreover, using the System x ScaleXpander kit, this performance range can be extended incrementally, whereby systems may install only a few processors at first, and then increase their processing capacity by adding additional chassis or "nodes" with more sockets. IBM calls it "pay-as-you-grow" expandability.

TABLE 3. Impact of Hardware Scalability on Availability of Virtualized Workloads

	Scale Up	Scale Out
Requirements	» Maintain service levels for most demanding virtualized workloads » Quickly deliver resources when needed	» Ensure virtualized workloads have continuous access to data
System x850 M2 and x3950 M2 / BladeCenter LS42 Solutions	» Unique eX4 chipset » Scales from 4 to16 sockets (96 cores) with a modular design (ScaleXpander Kit) » Scales from 7 to 28 PCIe slots (including 8 hot-plug) » Proven performance with record TPC-C results » BladeCenter LS42 blade scales from 2 to 4 sockets via modular Multiprocessor Expansion Unit)	» IBM BladeCenter is optimal for deploying VMs, and widely used for deploying virtual infrastructure » BladeCenter Open Fabric Manager minimizes downtime and complexity associated with reconfiguring LANs and SANs
Key System x850 M2 and x3950 M2 Memory Functions	» The largest memory capacity of x86 servers (1 TB) » Xcelerated Memory Technology » Huge snoop filter (320 MB) » Energy-smart DDR2 memory	

System x and BladeCenter provide a variety of features to minimize both planned and unplanned downtime resulting from hardware functions.

The eX4 chipset also enables the largest memory range in the x86 industry, allowing a single server to hold as much as 1 Terabyte (TB) of memory. Support for large memory capacity is particularly critical with virtualization deployments. Depending on the workload that is being virtualized, virtual machines often share CPU cycles much more readily than memory, potentially demanding large amounts of memory to be installed, even with relatively few processors. The eX4 chipset also builds in a variety of optimizations to boost the performance of accessing memory, including Xcelerated Memory Technology to allow memory to be accessed at higher frequencies; large-capacity snoop filters to search for memory locations more efficiently; and support for fast, energy-smart double data rate (DDR) memory modules.

For users who want to scale out using virtual infrastructure, IBM's BladeCenter platform provides an ideal form factor in which to deploy multiple virtual machine hosts with an optimal footprint. BladeCenter Open Fabric Manager simplifies the task of load balancing large numbers of virtual machine hosts by neutralizing the need to reconfigure SAN and LAN connections when hosts have to be relocated from one blade to another. In addition, the BladeCenter LS42 blade server is scalable from 2 sockets to 4, and from 8 DIMM sockets to 16, using the snap-on Multiprocessor Expansion Unit.

Hardware Reliability Benefits

As administrators plan for how to maintain the availability of virtualized workloads, they generally have to consider two classes of outages: planned and unplanned. One of the management advantages of virtual machines is that they allow workloads to be migrated from one host to another relatively easily and with minimal downtime. Indeed, in some ways, the use of such "live migration" functions reduces the importance of keeping individual servers running at all times. However, functions such as live migration do not necessary protect workloads against unplanned downtime, in which a host ceases to operate due to component failure or operator error (live migration requires both the source and destination hosts to be fully operable at the time of migration). Moreover, even in cases of planned downtime – when large numbers of virtual machines are hosted on a single server – the process of migrating all the VMs to another host may become fairly time-consuming. In these cases, it may be preferable to avoid shutting down the virtual machine host, even while maintenance is performed.

Hence, a server platform used to host virtual machines should offer resiliency functions that allow it to continue running even when certain components fail. The platform should also offer online maintenance procedures that allow the hardware to be reconfigured without interrupting processing. System x and BladeCenter provide a variety of features to minimize both planned and unplanned downtime resulting from hardware functions (see Table 4). The eX4 architecture has a number of resiliency features to protect servers against component failure in I/O, power, cooling or memory. Dual PCIe controllers with redundant I/O links allow processing to continue even if an I/O link fails. If a link fails, the I/O traffic will automatically be rerouted through the other link. Similarly, the various BladeCenter

IBM BladeCenter Open Fabric Manager allows administrators to easily change connections between servers, LANs, and SANs without adjusting MAC addresses or WWNs.

chassis include hot-swap and redundant power supplies, cooling, I/O switches, and a midplane with redundant power and signal connections to blades. I/O switch modules offer failover to other switches. BladeCenter S provides hot-swap internal RAID-5 storage and optional integrated SAN. The BladeCenter servers themselves offer online hot-spare memory protection and even failover to other blades.

System x eX4 architecture has extensive protection against memory failure. For example, with memory scrubbing, an automatic daily test of all system memory is performed to detect and report any developing memory errors that could cause a server outage. Standard ECC memory can correct only single-bit memory errors. Patented IBM Chipkill technology detects and corrects multi-bit errors. It works with the Memory ProteXion feature (also called redundant bit steering) that enables the server to sense when a multi-chip error has occurred, so that it can route data around the failed chip. For even greater protection, the memory mirroring feature, maintains two copies of memory in separate banks, and if a memory module fails, the system can still access memory from the second bank.

To minimize planned downtime, IBM's eX4 architecture supports hot-plug I/O slots, which allow I/O devices to be added, removed, or replaced without shutting down servers. Further, IBM's eX4 architecture also supports hot-swap memory modules, which enables the memory of host hardware to be reconfigured without interrupting their operation. The base x3850 M2 or x3950 M2 server contains two hot-plug Active PCIe slots. By using the ScaleXpander Kit, you can increase the number of hot-plug slots available to your VMs from two to four, six, or even eight.

In some cases, bringing down virtual machine host hardware for maintenance purposes may be unavoidable. While virtual machines can be migrated easily from one host to another, the virtual machine hypervisor itself often needs to be migrated from one server to another – an operation that presents its own challenges. Configuring virtual machine hosts that are integrated into SAN and/or LAN infrastructures often requires server management teams, storage management teams, and network management teams to cooperate. When LAN and SAN connections in virtual machine hosts have to be reconfigured, coordinating the efforts of these teams can present a significant logistical challenge.

This coordination burden can be simplified by applying some elements of virtualization to the server hardware itself, in particular, by allowing the MAC addresses associated with particular LANs and the World Wide Names (WWNs) associated with SANs to be dynamically reassigned to server hardware. IBM BladeCenter Open Fabric Manager allows administrators to easily change connections between servers, LANs, and SANs without adjusting MAC addresses or WWNs. This dynamic reconfiguration allows server administrators to relocate an entire virtual machine host, including the hypervisor and all its VMs, without the involvement of network or storage management teams. As a result, server administrators can be productive without having to involve LAN- and SAN-management personnel.

The Predictive Failure Analysis capability in System x and BladeCenter helps to identify malfunctioning components – including processors, memory, HDDs, power supplies, fans/blowers, and voltage regulator modules – before they actually fail, and issues alert events to the IBM Systems Director interface.

TABLE 4. Impact of Hardware Reliability on Availability of Virtualized Workloads

	Avoid Unplanned Downtime	Avoid Planned Downtime
Issue	» Virtual infrastructure provides only limited protection against hardware failures and catastrophic SW failures	» Virtual infrastructure operations (i.e., live migration) can be time consuming with a large number of VMs
BladeCenter Solution	» Hot-swap/redundant I/O » Hot-swap/redundant power and cooling » Chipkill and online hot-spare protection » Predictive Failure Analysis and light path diagnostics – Blade server and I/O failover » Integrated RAID-5/SAN in BladeCenter S	» BladeCenter Open Fabric Manager can migrate VM host images without reconfiguring LAN/SAN infrastructure » VMware ESXi hypervisor embedded on an internal flash key means almost instant virtualization, saving installation time
System x3850 M2 and x3950 M2 Solutions	» Hot-swap/redundant I/O » Hot-swap/redundant power and cooling » Memory ProteXion and Chipkill protection » Memory scrubbing and mirroring » Predictive Failure Analysis and light path diagnostics – Hot-plug I/O	» Hot-swap PCIe peripherals – Hot-swap memory » VMware ESXi hypervisor embedded on an internal flash key means almost instant virtualization, saving installation time

Managing Availability of Virtualized Workloads

Management tools play a key role in helping administrators maintain uptime. With virtualization, it is especially important for administrators to have visibility over both the physical and virtual parts of the infrastructure so that they can precisely identify the root cause of any failure conditions. As discussed earlier, IBM addresses management with two classes of products: IBM Systems Director, which is optimized for hardware-specific management tasks on various server platforms; and IBM Tivoli, a set of tools that are optimized for management across multiple business units and heterogeneous server platforms.

IBM Systems Director is the management interface responsible for controlling all hardware-related functions of System x and BladeCenter servers. In addition to

IBM's offerings will be of great value to users who want to achieve more from virtualization than its basic use case of consolidation.

being optimized for controlling System x hardware, IBM Systems Director integrates with popular VM management software such as VMware vCenter Server (formerly called VMware VirtualCenter). This integration enables low-level events in System x and BladeCenter hardware to trigger high-level virtualization events such as virtual machine migration, which can be used to avoid downtime in the event that an impending failure is detected. As a result, administrators can spend the majority of their time using a single interface. The ensuing familiarity helps to avoid downtime caused by operator error.

System x and BladeCenter servers also provide other notable management features to help maintain server uptime. As discussed earlier, the Remote Supervisor Adapter II and Advanced Management Module allow administrators to diagnose server problems even if their operating system or virtual machine hypervisor does not boot. Another feature called light-path diagnostics provides visual pointers inside the server chassis to specific components that need to be replaced. The Predictive Failure Analysis capability in System x and BladeCenter helps to identify malfunctioning components – including processors, memory, HDDs, power supplies, fans/blowers, and voltage regulator modules – before they actually fail, and issues alert events to the IBM Systems Director interface. These events could trigger the migration of virtual machines to other hosts, as described above, or on the BladeCenter platform, the entire virtual machine host image can be migrated to another blade server using BladeCenter Open Fabric Manager.

Since improving availability can mean different things to different organizations, users sometimes need help understanding what level of protection is appropriate for their particular requirements. In response, IBM has put together its High Availability Center of Competence (HACoC) to assist users in planning HA and disaster-tolerant solutions. The HACoC initiative works with customers to develop customized solutions for secure and resilient computing environments. These solutions are designed to help customers meet their goals of maintaining continuity and reacting to changes with flexibility and agility, while reducing operational costs associated with downtime or security breaches.

The IDEAS Bottom Line

Virtualization can simplify management tasks that are tedious and time-consuming, and help to reduce both planned and unplanned downtime. To achieve manageability and availability benefits with virtualization on x86 servers, customers must carefully consider not only which virtual machine software to adopt, but also the hardware platform on which to deploy their virtual machines, and the software to use for managing their virtualized environment. IBM's offerings will be of great value to users who want to achieve more from virtualization than its basic use case of consolidation. Using virtualization to consolidate servers has clearly proven its ability to lower acquisition costs, and some maintenance costs. However, in many organizations, the high operating costs typically associated with managing systems and maintaining service levels offer much more dramatic opportunities for savings.

Virtualization can simplify systems management and increase systems availability, potentially lowering their costs. In order to capitalize on these benefits, though, virtualization must be developed, delivered, and supported as a complete, integrated system. One of the challenges of using industry-standard technology to drive critical initiatives such as virtualization is ensuring that layers from different suppliers properly fit together. With industry-standard systems, the processors, server hardware, operating systems, and virtual machine platforms are all developed independently from each other. If users encounter "finger-pointing" from the suppliers of these components when they try to manage their virtualization initiatives, they will have difficulty achieving the benefits they set out to gain from virtualization in the first place. Therefore, the hardware platform used to host virtual machines must be supported as a unit, rather than as a collection of components.

IBM has a great deal of experience participating in the industry-standard ecosystem, and in delivering x86-based virtualization solutions, it has carefully evaluated where third-party technology is most appropriate, and where to add its own value. IBM relies on third parties to optimize the processor technology and virtual machine platforms, while it innovates with the eX4 chipset built around the processors; redundancy features in server hardware; management software; and I/O infrastructure virtualization technology. With System x and BladeCenter, IBM delivers and supports all these components as a complete system, ensuring that customers can focus on achieving the operational benefits of virtualization, rather than spending most of their time trying to get the components to work together.

Americas
Ideas International, Inc.
800 Westchester Avenue
Suite N337
Rye Brook, NY 10573-1354
USA
Tel + 1 914 937 4302
Fax +1 914 937 2485

Asia/Pacific and Worldwide Headquarters
Ideas International Limited
Level 3
20 George Street
Hornsby, NSW, 2077
Australia
Tel +61 2 9472 7777
Fax +61 2 9472 7788

Europe, Middle East, Africa
Ideas International Europe
The Courtyard
Lombard Street
Abingdon, Oxon OX14 5BJ
United Kingdom
Tel + 44 (0) 1235 462 890
Fax + 44 (0) 1235 462 891

www.ideasinternational.com