

## WHITEPAPER

# Key Considerations for Leveraging Virtualization and Keeping Your Applications Available

Today's IT organizations are faced with the daunting task of optimizing all aspects of their departments, including people, processes and technology. Optimizing and streamlining server utilization through virtualization represents one particularly exciting example. We found that one of the most popular usage models for virtualization is to drive down server procurements in development, test and production environments. When this model is followed, future server purchases are avoided; instead, new workloads are established on existing systems.

Since its inception, the term “virtualization” has been defined as the pooling of various IT resources in a way that masks the physical nature and boundaries of those resources from the users (see sidebar next page). It promises to unlock underutilized server capacity through partitioning, enabling the ability to run multiple operating systems on one server. The most obvious benefits of virtualization are not just the ability to achieve higher server and overall computing infrastructure utilization, but also improved responsiveness and flexibility—since virtual resources can be moved or modified dynamically to reflect changing business needs.

Disaster recovery represents an innovative application of virtualization technology. Traditionally, disaster recovery has been defined as the process of regaining access to the data, hardware and software necessary to resume critical business operations after a natural or human-caused disaster. In many instances, this is achieved through the use of redundant infrastructures and systems

Are you trying to keep a lid on your power and cooling costs and think virtualization alone is the answer?

housed in remote locations, away from primary data centers. While an effective strategy, many IT organizations simply cannot afford to buy and maintain “two of everything.”

Through the virtual partitioning of hardware to support multiple operating systems, virtualization purports to do away with the “two of everything” approach by delivering a level of redundancy that can be used for disaster recovery purposes—either on the same hardware or on a different piece of equipment within the same data center. Known as virtualized DR, this practice enables local availability within a single data center by facilitating failover from one partition to another partition in the same server, or in another server located in the same data center.

In this way, virtualization promises to bring the same benefits to disaster recovery as it does to the data center—namely, improved responsiveness and flexibility. For this reason, virtualization is rapidly gaining momentum as a means of supporting business critical applications (those requiring 99.9 percent >

availability), according to IDC Research. The firm also predicts that by 2010, applications requiring upwards of 99.999 percent availability or more (also known as fault tolerant) will also be addressable by virtualization.<sup>1</sup>

However, virtualization is not a panacea, particularly considering the heightened availability implications of virtualized servers and data centers, and virtualized DR. For example, the notion of local availability delivered on a single server or in a single data center directly contradicts a key tenet of the Seven Tiers of Disaster Recovery (see sidebar on page 4), which states unequivocally that a DR strategy with no off-site data may mean, quite possibly, no recovery.

Virtualization and virtualized DR are complex initiatives requiring careful, ongoing consideration of issues pertaining to people, processes and technology. This white paper explores the key challenges that organizations must consider as they undertake virtualization and virtualized DR initiatives, and the availability implications of each.

#### **These challenges include:**

- Virtual machine sprawl
- Increased complexity for server administration and management
- Change management and architecting the infrastructure migration from a physical to virtual environment
- Balancing investment versus higher utilization rates
- Need for new skill sets and processes

### **Virtual machine sprawl**

One of the first assumptions of a virtualization implementation is that it will allow an organization to cut down on the number of servers it needs to support its applications. While the number of physical servers may decrease, each one—by virtue of its supporting multiple applications—becomes significantly more critical, which in turn requires close consideration of disaster recovery strategies to support each virtual machine.

## **O** RIGINS OF VIRTUALIZATION

Virtualization has received significant attention over the past year, but the term has been widely used since the 1960s or earlier, and has been applied to many different aspects and scopes of computing—from entire computer systems to individual capabilities or components.

The term virtual machine dates from the IBM M44/44X experimental paging system, which in the mid-1960s introduced a new process for paging memory allocation, or the process of managing program access to virtual memory pages that did not currently reside in random access memory (RAM). When a page had to be loaded and all existing pages in RAM were in use, one of the existing pages had to be swapped with the requested new page. The paging system determined the page to swap by choosing one based on an algorithm that determined the best choice: the page least likely to be needed within a short time. IBM M44/44X was purely a research system, designed and operated at IBM's Thomas J. Watson Research Center at Yorktown Heights, New York.

#### **This challenge grows more onerous when one considers:**

- ***The tendency for virtual machines to multiply rapidly:*** One of the great promises of virtualization is the ability to place multiple heterogeneous operating systems on one server, thus enabling a wide range of operating system choice and freedom. However, because more than 50 percent of all virtual machines are running production level applications<sup>2</sup>—and because production level applications require superior availability and security—the instances themselves cannot be consolidated. The number of virtual machines

may grow rapidly, and without the proper back-up and redundancy, an organization's level of risk rises with the addition of each and every new virtual machine.

- *Virtual machines may be more vulnerable to individual "choke points"*: One of the first issues often overlooked when contemplating a virtualization solution is the high power consumption and high heat output of a single server hosting multiple virtualized servers. The reason for this is that virtualized servers cause the hardware CPUs to run at a higher utilization rate, requiring greater power and resulting in greater heat output, which in turn requires advanced (and often expensive) cooling mechanisms. A single malfunction of a cooling mechanism could be a recipe for a crash with costly consequences. In addition, virtualized servers require power redundancy (or the ability to withstand a single circuit failure). Without it, one accidental trip over a plug could bring down dozens of applications.

Availability implications aside, virtual machines may place other stresses on data centers. The number of virtualized servers and amount of data center space needed increases in direct correlation to the number of instances being supported, and many data centers are simply not equipped to deliver the space needed as well as deal with the new power and cooling requirements.

## Increased complexity for server administration and management

Many organizations view server virtualization as a way to lower costs, and believe reducing staff levels and lowering server support ratios will be a part of that. However, recent research shows that this is not always the case, with initial installations of virtualization products actually increasing the need for support because the environment is now more complex.

Virtualization often drives improved total cost of ownership and better use of hardware resources. However, in looking at potential staff savings, hardware support issues like installation and configuration represent just a small percentage of support staff time, which also includes software installation, configuration management, patching, capacity and performance monitoring, problem analysis and resolution and making labor reduction minimal.

Think virtualization can help you lower IT headcount? Think again. Going virtual may actually be more complex and more difficult to manage than the data center designs of the past.

Virtualization can serve the double effect of making these routine software-related tasks more complex, requiring higher investments of time and attention from IT staff. Coupled with greater operating system choice enabled through virtualization—which often drives up the number of instances, middleware and applications in need of support—the result may be a distracted and stretched IT staff that fails to pay due attention to important disaster recovery considerations.

In sum, consolidation through virtualization may reduce the physical elements of support, but it does not reduce—and may even increase—the software side of support. Before embarking on any virtualization initiative, organizations need to consider the ideal manner in which to control these support costs while obtaining the requisite staff and expertise. Looking beyond the need for expertise in installing virtualization products, organizations need resources who understand the disaster recovery needs of virtualized machines and environments.

## Change management & architecting the infrastructure migration from a physical to virtual environment

Virtualization places greater emphasis on a comprehensive understanding of interdependencies—that is, how a change to a particular hardware element may impact another infrastructure element and the applications and operating system instances that element supports.

With virtualization, the task of taking and maintaining inventory of servers, their application workload roles or identities (including systems and application software tied to each specific server model) becomes more difficult and critical, while the interdependencies multiply and become more complex.

In a virtualized environment, there is a higher likelihood of a failure to recognize interdependencies—and as a result, a higher risk for unplanned downtime which can affect multiple applications and become much more costly. Disaster recovery strategies can mitigate these risks by backing up and ensuring availability for virtual machines. As one respondent in an *SQL Magazine* survey noted, “Interdependence oversights fall into the ‘oops, I forgot’ class of problems and are a reminder that high availability is equal parts technology and human policies and procedures.”<sup>3</sup>

Patch management provides an example. In a non-virtualized environment, an organization may have to plan for application downtime around a Microsoft® Windows® Update patch being installed on the host server. In a virtualized environment, patching and rebooting a single host server creates a much more significant impact, because numerous operating systems and applications may be running on the server.

Virtualization also requires IT administrators to understand what applications are co-dependent and which must reside on the same server. For example, if application A is dependent on data inputs from application B, the two applications need to reside on the same server. If they are housed on separate servers, downtime for the server hosting application B may result in data discrepancies between the two applications, resulting in business inefficiencies or worse, costly errors.

## SEVEN TIERS OF DISASTER RECOVERY

The Seven Tiers of Disaster Recovery were originally defined by SHARE Inc., an independent, volunteer-run IBM mainframe user group, to help identify the various methods of recovering mission-critical computer systems as required to support business continuity.

Although the original published concept dates back to the 1990s, DR specialists today continue to use the seven tiers to illustrate continuity capabilities and costs at a very high level. The definitions for the various tiers have been updated as technology has evolved in support of today’s business requirements.

The first of the seven tiers, tier 0, refers to organizations with no off-site data. According to the tiers, no off-site data is equivalent to having no business continuity plan. There is no saved information, no documentation, no back-up hardware and no contingency plan.

The time necessary to recover in this instance is unpredictable. In fact, it may not be possible to recover at all.

The constantly changing nature of business applications and the IT infrastructures supporting them also requires the regular capture and storage of virtualized IT blueprints, also known as images. Images are often set up in-house and undocumented, and the challenge of re-staging an image can be so great that in the event of a disaster, an IT department may have to start from scratch—thereby hurting recovery point and recovery time objectives. A common mistake is to believe that virtualized disaster recovery environments existing within a single data center do not require remote back-up. These single location environments are vulnerable to the same threats of damage and destruction as non-virtualized environments, and in fact, the re-staging of images requires even greater consideration of remote back-up.

Automated provisioning of physical servers can drastically reduce physical server deployment time, enabling faster, more agile response to dynamic workloads and IT needs. But in a virtualized environment, physical server provisioning is not enough, and automated provisioning of virtual servers once the physical host server is up and running also needs to be considered.

# RISK MANAGEMENT

In business, the term operational risk management (ORM) refers to the oversight of many forms of day-to-day operational risks including the risk of loss resulting from inadequate or failed internal processes, people and systems, or external events.

The Basel Committee on Banking Supervision (BCBS) identifies seven categories of operational risk including losses arising from disruption of business or system failures, such as computer hardware and software failures.

In June 2004, the BCBS presented the second Basel Accord, also known as Basel II, which provided revised standards for measuring and ensuring the adequacy of a bank's capital to cover risk, with the ultimate goal of promoting greater stability for the world's financial system. What began as a trickle in the banking industry soon spread across sectors as diverse as insurance, asset and wealth management, energy, healthcare and government organizations.

The ability to accurately assess and manage operational risk allowed organizations not only to increase transparency and protect shareholders' interests, but also to achieve competitive advantage—through more productive use of capital resources.

Provisioning for a virtual server can be significantly more complex than provisioning a single physical server. The heterogeneous mix of virtual servers running on a physical server has to be planned. Physical server resources, such as CPU and memory, have to be calculated and properly allocated to each virtual server to provide a balanced physical server workload.

Did you know that virtual machines may be more vulnerable to individual “choke points”?

Improperly provisioned virtual servers can result in decreased server consolidation ratios and higher total cost of ownership. The complexity grows if an organization has a heterogeneous environment featuring multiple virtualization platforms.

In sum, virtualization and virtualized DR require IT staffs to consider and implement changes to their infrastructures in a more rigorous, exhaustive and well thought-out manner than ever before. At a higher level, the overall move from a physical to a virtual environment—and educating IT staffs on the full range of consequences and implications—can take months and require outside support. Organizations that fail to devote the adequate time, training and support resources may find themselves learning some hard and expensive lessons during either the normal course of business or in the event of a disaster.

## Balancing investment versus higher utilization rates

Looking beyond added staff and data center maintenance costs, virtualization often requires significant investments in the form of one-time, up-front expenses (for example, prototyping and new hardware purchases) as well as ongoing costs such as software licensing. Organizations must consider and understand the true cost scope in order to determine if higher server utilization rates (and the corresponding promise of lower costs) merit the level of investment and potential for higher operational risk (see sidebar).

Organizations that forge ahead must evaluate how to address these needs in the most cost-efficient manner while upholding the need for enhanced disaster recovery protection. In addition, because the virtualization market may consolidate over the next few years, organizations must ensure retention of unique, vendor-specific processes and procedures.

**Prototyping:** IT departments often begin virtualization projects by running production-like pilots to uncover any technical glitches or previously unknown requirements.

Some server vendors may lend a box for free, or at a significantly reduced price, but limit the deal to the piloting stage of the project. Start-up help from a vendor's professional services may not be included in the pre-sale. Or if it is, this help may be provided by the vendor's account executive or systems engineer—not necessarily the vendor's virtualization specialists.

#### **Preparing for a server consolidation negotiation:**

Before enlisting a professional consulting firm to assist with server consolidation, organizations must comprehensively audit their existing systems and applications. These internal audits can be both time-consuming and expensive, but research indicates that organizations that fail to carry them out are unable to clearly indicate requirements and expectations to vendors.

**Understanding the nature of your applications:** Virtualization is not appropriate for every type of application, and organizations should have a strong grasp of their existing and future application needs before investing heavily in virtualization. For example, applications that are CPU or I/O intensive are not ideal for virtualization, because of the need for fully dedicated resources and unshared high capacity.

**Software licensing:** Virtualization may result in drastic changes to software pricing schemes, with software vendors moving away from more traditional CPU pricing models to usage-based utility pricing. Organizations have benefited from CPU pricing, since it is relatively straightforward and allows employees, customers and partners to use software freely, without careful monitoring. The move toward “usage-metering” could mean escalating prices, while multi-core processing associated with virtualization may result in organizations paying for more than they're getting. Prices tied to hardware “are increasingly poor proxies for software value,” according to Amy Konary, program director, software pricing, licensing and delivery, IDC Research.<sup>4</sup>

## **Need for new skill sets and processes**

According to IDC Research, lack of technical expertise is a top-of-mind hurdle for organizations considering virtualization.<sup>5</sup> It calls for new IT capabilities including enabling applications to dynamically acquire more resources to accommodate peak performance; CPU capacity prioritization; and storage I/O and network traffic prioritization. This represents a significantly higher level of strategic planning than is customary for many IT departments, and a whole new level of skill sets that may not be inherent in existing staffs, no matter the size.

Another important skill set to consider is security. Virtualization requires a different information security skill set, because in virtualized environments the security must be provided at the hypervisor level versus the virtual machine level. The term hypervisor refers to the virtualization platform sitting between hardware and operating systems; its role is to manage hardware partitions and securely provide high performance access to the correct resources.

Compounding the problem is the fact that virtualization, as with any emerging technology, may be a prime target of new security threats—for instance, Gartner predicted that through 2009, 60 percent of virtual machines in production will be less secure than their non-virtualized counterparts.<sup>6</sup>

Prior to a virtualization or virtualized DR deployment, organizations need to implement comprehensive training programs. Virtualization software vendors may provide training on their products, but it may not encompass all of the knowledge areas required including hardware, firmware, middleware and software. Anyone working in the virtualized environment should possess a comprehensive understanding of all of these areas. In addition, change management, incident management and escalation, and hardware and software upgrade processes must be tested and documented for the new virtualized environment.

#### **Have you thought about:**

- Prototyping?
- Preparing for a server consolidation negotiation?
- Understanding the nature of your applications?
- Big changes in software licensing?

Organizations should also consider testing virtualization in a lab environment to gain experience and build internal processes. After a proof of concept is completed, low risk/low visibility applications should be deployed first, and as experience and comfort levels grow, the organization can start migrating systems with more risk and visibility to the virtualized environment.

As virtualization becomes more pervasive within organizations, the need to develop a cohesive business case will eventually arise. IT managers are not trained in this and often have a great deal of difficulty creating business cases that map to company objectives. In today's corporate setting, where institutional resistance is often fierce, IT managers may need to devote significant time and attention to formulating strong, persuasive ROI cases in order to secure top-level backing and support.

## Conclusion

Virtualization promises to extend to disaster recovery the same benefits it brings to the data center—increased capacity utilization, speed, agility and flexibility. Just as disaster recovery may benefit from virtualization, virtualization may benefit from disaster recovery, and organizations should link virtualization plans to disaster recovery plans in order to avoid inherent risks.

For example, by linking virtualization to disaster recovery, organizations can ensure proper availability, security and protection for these servers, helping to minimize operational risk (see sidebar on page 5). In addition, disaster recovery techniques can provide a means of storing and protecting entire virtualized environments through virtualization infrastructure images so they can be readily accessible in the event of a disaster. Organizations that rely solely on virtualized DR and local availability within a single data center for disaster protection—with

**Did you know virtual machines have a tendency to multiply rapidly?**

no off-site back-up—face tremendous risks, since the majority of disasters involve threats that force the shut-down of entire facilities. Standards and guidance for geographic separation, set forth by the National Institute of Standards and Testing, call for alternate data centers to be sufficiently far removed from their primary counterparts so as to not be negatively impacted by the same event.

As the Information Availability expert, SunGard Availability Services is ideally poised to mitigate these risks. SunGard has helped more than 10,000 customers worldwide develop and implement comprehensive strategies that focus on people, processes and technology—the three critical elements to help ensure that people always have access to the information they need.

SunGard is now bringing the same approach to virtualization, and more specifically, virtualized disaster recovery, supporting organizations as they identify opportunities for virtualization as part of larger disaster recovery strategies and link virtualization initiatives to disaster recovery planning—ultimately mitigating risk and optimizing infrastructure performance. For example:

### People

Today, SunGard's virtualization experts are helping organizations develop and implement virtualization strategies through staff virtualization training, assistance with IT audits, prototyping and testing, data center consolidations and moves (both planning and execution), strategy development and finally, change management and roll-out. SunGard is also incorporating policies and procedures for preserving virtualized machines and environments into customers' disaster recovery plans, and testing these as part of mock disaster run-throughs.

### Processes

SunGard is taking a leadership position in developing tools and techniques to support virtualized disaster recovery. One example is automated provisioning, which

**Do you know where your IT blueprints are? What about those undocumented images? Downtime could be in your near future.**

will enable SunGard to rapidly build and deploy customers' virtual images on SunGard systems. Automated provisioning will allow event-driven activations in addition to disaster-driven declarations—meaning that certain events such as an overloaded Web server will automatically trigger traffic failover to a SunGard-housed system.

Another example is virtual image capture and storage. For organizations supporting virtualized environments on their own premises, virtual image snapshots can be securely and reliably stored on SunGard systems. In the event of a disaster or business disruption, the organization can easily access the images through a portal, which also highlights critical interdependencies. Or, in the event of damage to a customer's primary data center, SunGard may activate the images in order to provision, build and activate a duplicate infrastructure for the organization on SunGard's own systems.

### Technology

SunGard manages nearly four million square feet of data center space worldwide, featuring infrastructure that can be optimized to deliver virtualized environments as managed services, as well as to provide back-up support to customers' own on-site implementations through:

- Experience with virtualization software and managing virtualized environments and complex interdependencies.
- Construction design and technological advancements (for example, cooling and power generation).
- Remote, secure locations, providing added assurance in the event of a disaster or disruption at the primary data center.
- Support for a wide range of hardware, operating systems and software (enabling a vendor-neutral, multi-platform approach to server consolidation projects in order to optimize value).
- Advanced data recovery techniques—including vaulting or replication—to support traditional or virtualized storage environments.

When combined with end-user recovery, these techniques allow organizations to bring RTOs and RPOs to less than four hours.

- SunGard is also an expert in datacenter moves and consolidations, having planned, managed and executed thousands of these events for customers—while maintaining operational continuity throughout. Many organizations leverage datacenter moves and consolidations as opportunities to simultaneously implement virtualization, and only SunGard offers the level of industry experience and depth of knowledge needed to concurrently maximize these initiatives.

## The Most Commonly Overlooked Considerations of Virtualization and DR

### Did You Know?

- Virtual machines tend to be more critical than their non-virtual counterparts
- They require greater (not less) consideration of remote back-up and support
- Virtual images also require a mechanism for bullet-proof back-up
- Virtual environments drive up the number of system and application interdependencies
- Downtime in a virtual environment has the potential to bring down many more applications than a non-virtual environment
- Virtualization tends to increase—not reduce—the need for administration and support
- Many data centers are simply not designed with virtualization in mind
- Lack of technical expertise is a top-of-mind hurdle
- Virtualization is not free, it often requires significant investments in the form of one-time, up-front expenses as well as ongoing costs
- Building a strong business case is key

SunGard is well positioned to tie all of these components together: expertise in virtualization software and environments, staff training and planning, automated provisioning, virtual image support accessible through a portal, data center preparedness and extensive support for all types of software, hardware and operating systems, combined with AdvancedRecovery<sup>SM</sup> solutions and more than 25 years of disaster recovery experience.

If implemented and managed properly, virtual machines and virtualized DR infrastructures can dramatically shrink recovery point and recovery time objectives. However, if not properly managed, virtualization can introduce significant availability risks which can hamper business resiliency and overall competitive position. By comprehensively addressing the challenges raised in this white paper, organizations can position themselves to maximize virtualization initiatives, reap the full benefits of virtualized DR and ultimately protect and strengthen their businesses.

## Authors

**This white paper is based on the collective experience of SunGard Availability Services and was written by:**

**Jim Champion, VCP, MCP**

Lead Consultant

**Jim Grogan, MS, CISM**

Vice President, Consulting Product Development

**Mark Stoecklein, MBA, PMP, MCSE**

Senior Director, Technical Solutions Group

**Jack Wade, MCSE**

Practice Manager

**Managing Editor**

**Valeria Maltoni**

Director, Marketing Communications

## About SunGard Availability Services

SunGard Availability Services is the pioneer and leading provider of information availability services, helping to ensure that 10,000 customers in North America and Europe have access to their business-critical information systems. With four million square feet of operational space, SunGard offers a complete range of information availability services for more than 30 technology platforms, from 48-hour disaster recovery hotsites to always-on, high-availability infrastructure, and electronic vaulting services. SunGard also provides technology and systems management services for application and data center outsourcing, as well as business continuity consulting services and planning software.

**For more information about SunGard Availability Services, visit [www.availability.sungard.com](http://www.availability.sungard.com) or call 1-800-468-7483.**

## References

1. John Humphreys, IDC. Virtualization 2.0: "The Next Phase in Customer Adoption." Doc #204904. December 2006.
2. John Humphreys, IDC. Virtualization 2.0: "The Next Phase in Customer Adoption." Doc #204904. December 2006.
3. Brian Moran, SQL Server Magazine. "Top Causes for SQL Server Downtime." SQL Server Magazine. 28 August 2003.
4. Amy Konary, IDC, quoted in: "In Depth: Software Vendors Try New Pricing Schemes for a Virtualized World." InternetWeek. 31 July 2006.
5. John Humphreys, IDC. Virtualization 2.0: "The Next Phase in Customer Adoption." Doc #204904. December 2006.
6. Neil MacDonald, Gartner. "Security Considerations and Best Practices for Securing Virtual Machines." 6 March 2007.

**SUNGARD**<sup>®</sup> | Keeping People  
Availability Services | and Information  
*Connected.*<sup>®</sup>

680 East Swedesford Road | Wayne, PA 19087  
800-468-7483 | [www.availability.sungard.com](http://www.availability.sungard.com)