TANEJA GROUP
TECHNOLOGY ANALYSTS

P R O D U C T P R O F I L E

VMware Infrastructure 3 – The Marriage of Server Virtualization with Business Continuance

July 2006

Server virtualization, as the name implies, has been categorized by the industry as a largely server-related phenomenon. However, Taneja Group sees this as short sighted. Server virtualization and the infrastructure transformation it begets to a virtualized, consolidated, centrally-managed infrastructure has far reaching ramifications to the areas of data protection, disaster recovery, and high availability. Specifically, server virtualization dramatically lowers deployment complexity for employing technologies like clustering, backup, and replication. The boundaries between servers and storage are blurring and the repercussions of technological innovations in these two industries must be viewed holistically – not in isolation.

In the following profile, we describe what server virtualization is and then examine the ramification of the technology on providing business continuance services for mission-critical applications. We specifically concentrate on the areas of data protection, disaster recovery, and high availability and how a virtualized infrastructure improves the deployment of those technologies. Further, we spotlight VMware and its latest product release VMware Infrastructure 3 that features many business continuance enhancements that simplify the deployment of virtualized infrastructure.

The State of the Market

Over the past six years, server virtualization has become a technological change agent for revamping server architectures and wringing cost out of the data center. It is difficult to find a CIO of a Global 2000 corporation who is not embarking on a server consolidation project or hasn't earmarked dollars for server consolidation initiatives in their environment. Routinely through CIO councils and other forums, we hear that server consolidation continues to be a top spending priority for large companies. This is a clear trend in the market and one that we anticipate to continue unabated for some time. Clearly, server virtualization as a technology has moved out of the development and test and desktop arenas – early market footholds– to the heart of the data center.

Three years ago, it might have been controversial to say that server virtualization would become part of the core data center, but today (2006), it is a safe bet to assume that server virtualization will become a mainstream part of the x86 computing landscape. The only question left to answer is what types of applications will it be used with (all applications or some subset of applications) and what classes of service level it will support (applications whose recovery is measured in second or hours or days).

What is Server Virtualization?

Server virtualization is a layer of software that sits between the operating system and the bare metal hardware. This layer of software, often referred to as a hypervisor, hosts multiple virtual machine instances. Each virtual machine contains an entire operating environment – applications, operating system, and virtualized CPU, memory, disk I/O interface, and network I/O interfaces. An application running in a virtual machine thinks it is running on the bare metal hardware when in fact it can be sharing the same physical hardware with 20 other applications or more. Moreover, this design allows multiple types of operating systems (e.g. Linux, Windows NT, Windows 2000, Sun Solaris, etc.) to be co-mingled and share the same physical hardware server.

The role of the hypervisor is to slice the bare metal resources of the server and parcel them out for consumption in each of the resident virtual machines. At its core, server virtualization decouples the physical hardware from the software environment which enables a number of efficiencies, one of which is the ability to run multiple, isolated operating environments on the same physical server.

Business Continuance (BC)

It goes almost without saying, we live in a 24x7x365 world where IT systems are expected to be available globally and where an outage can dramatically impact the financial bottom line or brand of a company. Business continuance is a topic on the mind of any CIO, particular in light of the tragedies of September 11th and Hurricane Katrina. Business continuance as Taneja Group defines the term refers to the technologies and IT processes that are used in the data center to ensure that applications and data are available and safeguarded from any type of failure imaginable – whether that failure be the loss of an entire data center, a software failure, or a hardware failure at the server, storage or network level.

Broadly speaking, there are three main technologies that fall under the category of business continuance. They are:

1. Backup for Data Protection
2. Clustering for High Availability
3. Replication for Site to Site Disaster Recovery

Business continuance is a critical part of the IT planning and infrastructure world without server virtualization in the mix. However, server virtualization, as we have discussed, is driving customers to consolidate their applications onto higher and higher utilized server platforms. Ironically, server consolidation taken to its natural end exacerbates the need for business continuance – a failure of a single server can impact multiple applications. Therefore, business continuance in a virtualized world becomes more important, not less.

The Problem with BC Today: Complexity

In talking to end users, one thing becomes clear. There is no question that IT administrators want to protect and safeguard their applications more than they currently do, but business continuance in today's

environment is a complicated equation that weighs risk of an outage and importance of the application against cost to administer and procure the technologies. BC solutions, like clustering and replication in particular, cost a lot of upfront money, but also are taxing in terms of administrative overhead on the IT organization.

Much of this complexity stems from the subtle differences in OS, application, and hardware settings that exist across different generations of hardware and the data center as a whole. For example, administrators must ensure that two servers that participate in a high availability cluster be mirror images of themselves. Much of an administrator's time is spent ensuring that the OS settings, application configuration, registry keys, hardware bios, etc. stay synchronized. The penalty for failure is high. If two servers fall out of synch, the failover of an application may not complete and thus an outage ensues despite Herculean efforts to safeguard the system.

Site to site disaster recovery using replication technologies suffers similar complexity problems. In site to site recovery scenarios, two operations must occur flawlessly for a successful recovery to occur. First, the data must be copied from site A to site B. Second, the application and operating environment must be duplicated and the application restarted using the replicated/copied data set. Similar to the case of clustering, there is a myriad of complexity to ensure that the physical environment in site A looks the same as the physical environment in site B.

Consider the case of backup and restore operations. Typically, a backup application will backup the data of an application on a nightly basis. Backups can be full copies of the data or incremental backups (the delta of what has changed since the last backup). After a disaster, an IT administrator must do a bare metal restore of an application environment. He must not only restore the data from tape or other media, but also rebuild the application and operating system environment. This rebuilding usually entails re-installing the OS on the target server and manually configuring the server environment to mirror the previous backup environment. Duplicating this environment is very difficult and fraught with error if the physical hardware is different from the previous version. Once again, the complexity of recreating the operating environment wreaks havoc on the administrator at the absolutely wrong moment – when a disaster has struck and he or she must get the system up and running as fast as possible.

The Benefits of Virtualization

Server virtualization enables IT to fully utilize server capacity by stacking and load balancing applications across the server infrastructure. Fundamentally, server virtualization breaks the 1:1 relationship between the operating environment (application and operating system) and the physical hardware. In turn, IT administrators can drive server utilization up from 5-15% on average to 60-75% in a virtualized environment. However, for the purposes of this paper, we will focus the discussion on the benefits that server

virtualization brings to the deployment of business continuance services.

Server Virtualization and BC

Much of the complexity associated with clustering, replication, and backup can be eliminated in a virtualized data center. Server virtualization provides a proven mechanism for storing the operating environment as a collection of virtual machine files and restoring the operating environment on different physical hardware using copies of these files. By breaking the dependence between the operating environment and the physical machine and by encapsulating the operating environments into a handful of files, virtual machines can be saved, moved, and restored with relative ease. This process can be automated and removes the complexity of bare metal operating environment configuration that plagues the use of data protection, replication, and clustering technologies in a physical world.

Virtualization and Clustering

Virtualization simplifies the deployment and hardware costs of clustering. Instead of clustering physical servers where the hardware must be a mirror image, virtualized servers running multiple virtual machines are clustered together. The server virtualization software monitors the health of each virtual machine and if there is a failure, restarts the virtual machine on another virtualized server in the cluster. By using virtual machines, the complexity of keeping the operating environment consistent between two physical servers is removed. In addition, clustering and high availability can be easily extended to all virtualized applications, regardless of type. In traditional clustering, specific scripts often must be written to monitor and failover an application. In a virtualized environment, there are two approaches to high availability –clustering software running in the virtual machine or actually clustering the virtual machines. The most common approach undertaken by most server virtualization vendors is to run the clustering software, like Microsoft Cluster Service (MSCS), in primary virtual machine and communicate with a heartbeat to a secondary or standby virtual machine. The clustering software, like any other application, thinks that it is running on the bare metal hardware, and monitors the health of the application within the virtual machine (as well as the health of the virtual machine itself). On failover, the application is restarted through a set of customized scripts in the secondary virtual machine.

In the case of clustered virtual machines, the server virtualization software monitors the health of the virtual machine through a heartbeat or ping mechanism. If a failure occurs, it orchestrates the failover process by restarting the entire virtual machine on a secondary server running the hypervisor. This approach has the benefit that no customized scripts need be created for each application and no additional clusterware is used. It can be rapidly applied to any application or operating environment. Only VMware supports this second approach today.

Virtualization and Replication

Replication and server virtualization also work well together. Server virtualization provides the mechanism for encapsulating and persisting the operating environment of a virtual machine to disk. The replication software, whether it is storage-based (e.g. EMC SRDF, Hitachi TrueCopy, or Network Appliance SnapMirror) or host-based (Legato Replistor or NSI Double-Take), can replicate the virtual machine files or the data within a virtual machine to the target data center. Instead of worrying about recreating the exact physical hardware environment, the persisted virtual machine can be restarted on any physical hardware running the server virtualization software. The manual process of bare metal restoration of the operating environment is eliminated and the disaster recovery process can be fully automated without any manual operations steps.

Virtualization and Backup

With virtualization, standard backup software can now be configured to not only backup the application data, but also backup snapshots of the virtual machine. These virtual machine snapshots can be restarted on any physical server running server virtualization software - obviating the painful and manual bare metal restore process.

VMware Infrastructure 3

With the launch of VMware Infrastructure 3 (VI 3), VMware has taken server virtualization to a new level. VMware Infrastructure 3 is an integrated suite of virtualization, business continuity and management tools for supporting data center applications. The suite is composed of VMware ESX Server (an enterprise class server virtualization layer), VMware VirtualCenter (the control point for managing and allocating virtual machines to pools of physical hardware), VMware Virtual Machine File System (a file sharing layer to share virtual machine data on Direct Attached Storage (DAS) or an IP or Fibre Channel SAN), VMware Virtual SMP (software that allows a virtual machine to scale its workload across multiple physical processors), VMware VMotion (software that allows a running virtual machine to be moved from one physical server to another), and VMware Distributed Resource Scheduler (a policy engine that balances virtual machines across resources).

In addition to providing significant product updates to the core virtualization and management products, VI 3 introduces two new products specifically aimed at business continuance – VMware HA (virtual machine heartbeat and health monitoring software that allows any application running in a virtual machine to inherit the availability benefits of the clustered ESX Server) and VMware Consolidated Backup (software that supports a centralized backup facility of all virtual machines). These products in concert with enhancements to VirtualCenter to manage the distributed infrastructure allow VMware to tell a more integrated, simplified story around providing high availability and disaster recovery services to a virtualized infrastructure.

Taneja Group Opinion

The virtualized world needs business continuance features just as much as the traditional physical topologies of today's data center. Server virtualization dramatically simplifies and in some cases improves the efficiency of traditional data protection, disaster recovery and availability technologies.

Taneja Group believes that applying server virtualization to the areas of clustering, backup and replication should increase the penetration and adoption of these technologies and extend their usage to a larger class of applications. Historically, technologies like clustering have only applied to 10% of applications in the data center. Deployment complexity and cost prevented clustering from ever breaking out of the higher SLA level application niche that it has occupied. Despite the desire of the lines of business for greater availability and resiliency in the data center, clustering has not been a mainstream phenomenon.

Replication adoption is similarly low. The cost and complexity has limited the use of replication to only the most high end customers and data centers. This has changed somewhat over the course of the past five years as disasters like 9/11 and Hurricane Katrina have raised awareness and low-cost, IP replication schemes have become more commonplace. However, site to site disaster recovery is still mostly reserved for only the top most tiers of service levels.

Only the adoption of backup is ubiquitous across the enterprise and has reached critical mass across all applications and service classes. However, backup is not exempt from complexity. Taneja Group through its research has first hand knowledge of the disruptive forces at work in data protection. Whole new product categories like Data Protection Management (DPM) have been spawned to manage the complexity that arises from managing backup on a wide scale across the data center.

Moreover, products from Bare Metal Restore (BMR) vendors, like Unitrends and Veritas/Symantec, attempt to address the issues of bare metal restore without leveraging server virtualization technology. These products represent a viable substitute for solving the backup problem, but neither DPM nor BMR addresses the issues across the full spectrum of data protection, disaster recovery, and high availability like VMware does.

Server virtualization can greatly simplify and automate the bare metal restore process that is part of any tape media-based disaster recovery scenario. Moreover, VMware's announcement of VMware Consolidated Backup is significant because it offers a better, more efficient model for conducting backup and reducing the number of agents that must be installed, configured, and managed in a virtual infrastructure environment.

With VMware Consolidated Backup (VCB) and SAN storage, a single backup proxy server can backup multiple physical servers hosting tens to hundreds of virtual machines.

Each virtual machine is analogous to an independent server that would require a backup agent in order to be protected in the physical world. VMware supports the use of traditional backup agents running in virtual machines as well. However, VCB simplifies the backup process and removes the need for having backup agents running in each virtual machine. Instead, VCB quiesces the virtual machines and takes a snapshot of the virtual machine images containing the operating system, application configuration, data, and system state that is stored on the shared SAN. Using VMware Virtual Machine File System, multiple servers can share data on the SAN and get access to these persisted copies. A backup server then backs up the snapshot of each virtual machine on the SAN and offloads it to tape or other media. This reduces the need for administering backup agents and reduces backup licensing costs.

Taneja Group believes that VMware has a great opportunity to change how servers and storage are managed, share information, and interoperate as part of a major computing shift to x86 based computing elements in the data center. VMware's vision of the "virtualized" data center where virtual machines can be moved seamlessly without disruption of service and balanced across a pool of shared CPU, memory, network, and disk capacity is compelling. Add to that mix, a new product launch with key additions like VMware HA, VMware Consolidation Backup, and VMware DRS and you have a potent product combination that moves VMware closer to its aggressive vision of a virtualized data center.

VCB and VMware HA are a good start in the areas of integrating virtualization with business continuance. However, they are more in line with improving the applicability and broad adoption of these tried and true approaches to these problems Taneja Group believes that there is a fundamental opportunity for VMware to change how high availability and DR are done on a virtualized platform. Specifically, we see the extension and application of the VMotion technology to the area of unplanned downtime as being a killer app that would change the game and redefine what people think of when they think "clustering." We realize that this would be a major engineering undertaking, but the benefits of instantaneous availability and disaster recovery of applications across sites would be a powerful differentiator and further hasten adoption of virtualization in the most rigorous enterprise environments.

Two worlds are colliding – server virtualization and storage technologies. Interestingly, the acquisition of VMware by EMC positions the combined company at this intersection. The combined entity provides a unique integrated footprint in the market across both server and storage infrastructure. Clearly, EMC's product portfolio in data protection, replication, and backup could be more closely integrated with VMware's server virtualization. There is a symbiotic relationship that can give the combined company a more tightly integrated story than any other vendor in the market. However, VMware will have to walk a tightrope with integrating with EMC storage portfolio while still continuing to be an open ecosystem player – something that has benefited the company immensely to date.

The virtualized data center is not a reality yet, but VMware is executing on a strategy that Taneja Group believes will shape the data center for the next 5 to 10 years and change how applications, OS, clustering, replication, and backup are deployed and utilized. As the other major server virtualization vendors, XenSource and Microsoft, struggle to get their products out in the market and close the feature/function gap with VMware's ESX offering, we see that VMware has a significant opportunity to build an insurmountable lead in defining the virtualized data center. VMware Infrastructure 3 puts the other players on notice that they will need to raise their game in order to compete.