G00219322

Storage Infrastructure Issues to Consider in a Virtual Server Environment

Published: 3 October 2012

Analyst(s): Valdis Filks

IT managers who manage or need to improve the productivity of server-virtualized environments must consider several storage options and protocols. We highlight what can and cannot be achieved with today's storage technologies and solutions.

Key Findings

- IT organizations engaged in server virtualization projects must review and modernize their storage infrastructure.
- Storage virtualization, as a point solution, will not optimize storage costs and improve storage administration productivity when used without storage management and planning tools.
- Backup and recovery (B&R), business continuity plans and disaster recovery testing for geographically distributed virtualized servers and storage need to be implemented and periodically tested to verify that applications can be recovered within the expected service levels.
- Storage protocols offer a choice of features that can provide different performance, cost and availability requirements for virtualized servers.

Recommendations

- To avoid bottlenecks and contention, determine storage performance and availability requirements for applications moving to or running on virtualized server environments.
- Implement multipathing, highly available scalable and flexible storage infrastructures that can accommodate dynamic storage requirements.
- Factor in the extra purchases and maintenance costs of new storage solutions that may be required in a server virtualization project.
- Leverage input/output (I/O) virtualization to reduce the number of I/O adapters and top-of-rack (ToR) switches to simplify cabling within the rack, and factor in extra purchase, support and maintenance costs.

Table of Contents

Analysis

Storage and Server Virtualization Are Fundamentally Different

Server virtualization has improved server resource utilization and changed server management in many areas, but has not reduced storage costs or improved storage utilization, disaster recovery, or general storage and data management tasks. Many organizations that expect to repeat the server physical-to-virtual (P2V) changeover savings for storage are assuming that applying the same virtualization process to storage will repeat the successes that server virtualization projects provide. This assumption is flawed because storage has already been virtualized and consolidated within shared storage area network (SAN) arrays.

Server virtualization enables the sharing of resources within the server domain, but the benefits are not as large in an already consolidated storage environment. For over a decade, standard storage arrays have offered the equivalent consolidation features of server virtualization, such as a single point of management, multitenancy, partitioning, the sharing/pooling of storage resources, controllers, cache, disks, internal buses and external interconnections that are shared among many servers and applications. Snapshots and replication have also been used to virtualize, copy and move data. Only in the largest and most disparate of environments is a separate storage virtualization layer required to combine, manage and migrate data among many different heterogeneous storage arrays.

To avoid escalating storage costs and improve storage management efficiency, IT managers need to understand how server virtualization solutions use various storage systems and technologies, as well as how the server and storage devices process, store data and behave when resource-constrained. These differences are greatest when storage needs to be physically moved and when extra capacity needs to be added. The differences and limits between server and storage virtualization are exacerbated because of the order of magnitude difference between the response time of server memory and physical hard-disk drive (HDD) storage.

Motion Tools Optimize Virtual Servers: Storage Is Constant

Server virtualization removes the virtual host dependencies from the physical hardware, making the server CPU and RAM resources elastic and variable. When a virtual host requires more server resources, it can be moved to another physical server within minutes. These virtual host movements are often not accompanied by the movement of storage. In most cases, the storage stays on the same device, while the virtual host moves location. Storage resources remain fixed and cannot be changed or moved with the same ease among storage systems. Server virtualization has not increased the flexibility of storage resources, even though it has added an extra layer of abstraction. This is because most data centers have consolidated and centralized their storage into one or two storage arrays that are shared among many virtual hosts. When virtual hosts are not moved between data centers, there is little reason or value to move the virtual hosts' storage between storage arrays. Server virtualization techniques solve or alleviate server resource contention, but do not resolve storage capacity or performance contention problems within the storage subsystem.

Storage Can't Rely on Another Layer: The Lender of Last Resort

In today's computer systems, the lender of last resort is the disk, because a piece of data in a CPU cache, RAM memory or any other component can only be permanently stored on disk (or tape). It has nowhere else to go. Server virtualization does not dynamically enable increases in storage capacity or storage performance, or create a new storage area. IT departments need to manage their server virtualization project expectations, taking into account that the storage management workload will not decrease or become simpler. IT departments must not rely on the expected promises of virtualization that resources will always be found somewhere within the consolidated environment.

Once IT departments reach their physical storage limit, there is no virtual storage area that can be exploited. However, storage-array-based techniques, such as thin provisioning, compression and primary deduplication, can temporarily solve the storage capacity problem, as long as the data is suitable for these techniques. Primary deduplication of active application data on disk systems is becoming available in solid-state drive (SSD) and hybrid SSD and HDD systems, such as IceWeb, Nimble, Nimbus, Oracle Zettabyte File System (ZFS), Pure Storage, Tegile Systems and Whiptail. IT departments should also be aware that when all primary data has been thin-provisioned, deduplicated or compressed at the source, they will again reach the compressed storage capacity limits of a storage system.

Adding disks or removing data from the array by a technique such as archiving can increase the limits. IT departments must anticipate these limits to prevent the service or application from failing

This research note is restricted to the personal use of mserale@villanova.edu

or becoming unavailable in real time. Once this resource is exhausted, there is nowhere else to go, such as swap or page files that servers use for overflow.

The Performance Differential Between Storage and CPU Affects Everything

Virtual hosts can be moved among physical servers in less than a second, because the limits are electronic CPU, RAM and network speeds. Data movement is limited by the electromechanical performance limitations of HDDs. Even if IT departments use SSDs that are 100 times faster than HDDs, the time to switch is determined by replication latency and bandwidth limitations if the data has not been previously replicated or copied. To move data or switch from one HDD or SSD system to another in or among data centers via snapshots or remounting replicated copies among physically separate arrays takes more time than it takes for a server to stop and start executing instructions. Therefore, the time it takes the virtual host migration to move data can be at least tens of seconds if a logical snapshot is used. Often it is minutes, and sometimes hours, if the data needs to be physically moved. Due to the much longer length of time that it takes to move data, a system or application outage is often required.

The bottom line is that servers and disk storage systems have fundamentally different characteristics, and consequently need to be managed and treated differently. From a systemic perspective, the slowest or most-unreliable component or fault domain on a system determines the overall performance and availability of that system.

Storage Considerations in a Server-Virtualized Infrastructure

Storage Utilization

While storage utilization issues are not specific to server-virtualized infrastructures, maximizing storage capacity utilization is one of the primary methods used to reduce storage costs and achieve a high ROI from the installed storage infrastructure. Feedback from Gartner clients shows that high storage utilization rates can approach the 70% to 80% level, whereas organizations with severe storage management problems can have storage utilization levels as low as 10% to 20%. However, storage utilization rates can depend on the computer platform, and all systems cannot be measured in a similar fashion.

As a guiding principle of good storage management practices, Gartner advises organizations to strive for higher storage utilization rates. Organizations should monitor and record storage utilization levels before server virtualization projects begin and after they're completed to ensure that storage efficiency does not deteriorate after a server virtualization project, and that storage costs do not increase.

Action Items:

- Aim to have storage utilization levels around the 70% to 80% level, with thin-provisioned arrays.
- Always leave buffer space and spare capacity; 25% is a good rule as an average.

This research note is restricted to the personal use of mserale@villanova.edu

- Use data reduction (e.g., compression and deduplication) and thin-provisioning technologies to improve storage utilization rates and to contain management costs.
- Continually monitor storage utilization rates, and add this as a key performance indicator to system management tools and dashboards.
- Record before-and-after values for storage utilization and capacity requirements when implementing server virtualization projects.

Reduction Technologies

Most server virtualization solutions, by definition, reduce the physical number of servers, but not the logical number of server images and, therefore, run many different OSs within the virtualized server. When the same OS is used by the virtual hosts, approximately 70% of the data on the boot or root disks is the same, except for the specific configuration data files. This characteristic makes deduplication of boot or root disks a very good candidate for primary deduplication.

Action Items:

- Implement deduplication for boot and/or root disks or volumes within virtualized server infrastructures.
- Verify and test that deduplication for primary volumes does not cause an unacceptable performance overhead. If it does, then use deduplication only where the application performance requirements are not negatively affected.
- Test, verify and implement other data reduction techniques, such as compression, to reduce primary storage requirements.

Example Vendors/Solutions:

- Dell, EMC VNX, IceWeb, NetApp, Nexenta, Nimble, Nimbus, Oracle ZFS, Pure Storage, Tegile and Whiptail

Storage Management and Configuration

The ability to manage systems at a higher level can improve operational productivity if storage array virtualization features are exploited. However, the additional virtualization layers in server virtualization systems can obscure underlying problems and block application access to array features, especially capacity information from a storage management perspective. Server virtualization solutions enable organizations to implement a virtual storage pool on a server that shares and distributes storage to all hosts within those virtualized servers. When this is implemented, storage administrators can lose accountability of storage usage in the server virtualization layer. Luckily, most storage resource management products can interrogate the virtualized storage layer in hypervisors, and storage can be monitored and managed.

Action Items:

- Provision storage to virtualized servers, instead of to hypervisors, to enable access to array management features.
- Implement storage management tools that show and maintain the visibility of storage usage and the relationships among applications, virtual hosts, physical servers and storage devices in virtualized server environments.
- Use storage capacity planning tools that can forecast and estimate storage growth requirements and upgrade cycles.
- Implement storage monitoring tools that can display the end-to-end topology among the components in the data path to help problem determination tasks.

Example Vendors/Solutions:

- Aptare, EMC (Ionix ControlCenter), Hitachi (Storage Command Suite), HP (Storage Essentials), IBM Tivoli (Storage Productivity Center), NetApp (OnCommand), Quest (Storage Horizon), Storage Fusion, Symantec (CommandCentral), SolarWinds (Profiler)

Performance

Shared-system environments require comprehensive and detailed IT management disciplines, as individual applications have an increased capability of negatively affecting each other. Therefore, storage performance considerations must be planned, and changes need to be modeled to avoid bottlenecks and contention. If this is not done, it can lead to a deterioration of application and service performance levels. Storage I/O-intensive applications should not be virtualized or consolidated onto the same system unless there is sufficient capacity and throughput available. Restoring, backing up and migrating data requires large amounts of I/O bandwidth; therefore, architects of server virtualization systems should design infrastructures that can support these processes and the software products that require the virtualized servers to perform them.

During data movement, high amounts of storage bandwidth usage can cause performance bottlenecks on any shared components of the infrastructure, making application performance unpredictable. Due to the increased complexity of monitoring and the larger number of components and relationships within a virtualized server environment, storage performance tools become extremely important. IT departments need to consider storage performance when designing server-virtualized environments, and need to purchase storage performance monitoring tools to solve performance problems within the server-virtualized infrastructure.

Action Items:

- Monitor storage and I/O, as they are more sensitive to latency delays than other external connections, such as Internet Protocol (IP) networks.
- Create near-real-time or real-time storage performance monitoring systems that enable exception reporting and alerts when storage performance issues occur.
- Include storage performance metrics — e.g., I/O operations per second (IOPS) and bandwidth (MBps) within your dashboard reporting systems.

- Use storage resource and performance trending products to plan and forecast system performance requirements and upgrades.

Example Vendors/Solutions:

- Brocade (Network Advisor), Cisco (Data Center Network Manager [DCNM]), NetApp OnCommand, Veeam Software (Veeam Monitor), Virtual Instruments (NetWisdom and VirtualWisdom)

Backup and Recovery Considerations

As hosts move into multiple logical servers, they drive a change in the B&R process, and in the implementation and design of the B&R infrastructure. This change provides IT departments with an opportunity to modernize the B&R infrastructure, and to reduce costs while improving the functionality and responsiveness of the B&R system. Where the traditional B&R architecture has one agent on each host, it is possible to reduce this agent proliferation and have one host or agent per hypervisor instance or physical system to manage all virtual hosts on that system. New techniques, such as client-side deduplication, disk image backups or snapshots, can be implemented to perform backups.

Action Items:

- Review the B&R infrastructure, and modernize the system as required, potentially investigating virtual machine (VM) backup point solutions or exploiting recent VM protection options from the incumbent backup application.
- Consolidate and reduce the number of B&R products and processes.
- Renegotiate B&R purchase, support and management costs.
- Investigate new recovery techniques and APIs from the hypervisor and backup application vendors.

Protocol and Storage Connection Considerations

Gartner client inquiries show that virtual server infrastructures are more latency sensitive than bandwidth sensitive. This has caused a migration to FC from file-based protocols, especially when there are a high number of virtual hosts per physical server. Due to the nature of virtualized hosts, all interfaces need to be virtualized to connect or map storage to a virtual host. From a storage perspective, the virtual infrastructure requires virtual FC worldwide networks (WWNs) and virtual worldwide port numbers (vWWPNs). For network-attached storage (NAS) and IP-based solutions, it requires virtual IP addresses and virtual Media Access Control (MAC) addresses (virtual network interface cards [vNICs]). While the numbers of physical interfaces may decrease, logically, they do not decrease in number. There is no reduction in the number of entities that need to be tracked and configured. Interfaces are virtualized, but still need to be managed, reported on and displayed in graphical topology and reporting tools.

This research note is restricted to the personal use of mserale@villanova.edu

Server virtualization increases the size of the fault domain by having many systems within one physical server. Regardless of the protocol used, at least two physical storage paths or connections always should be used to obtain increased availability. Automated failure and active-to-active multipathing should also be enabled.

Storage protocols are grouped into two main categories:

- Block-based — ATA Over Ethernet (AoE), FC, Internet Small Computer System Interface (iSCSI), Serial Attached SCSI (SAS), FCoE
- File-based — Network File System (NFS), Common Internet File System (CIFS)

Each category has different price and performance characteristics and considerations when selecting storage for virtualized servers. We briefly outline the issues concerning the main protocols in use today.

AoE

AoE can be used on all Ethernet switches and networks. Most customers create a virtual LAN (VLAN) in their data center Ethernet network to segregate this from other network traffic. Gartner recommends that at least 1 Gigabit Ethernet (GbE) is used, but 10 GbE is preferred. AoE operates at Layer 2 of the Open Source Initiative (OSI) model and is not a routable protocol. Very little configuration is required, other than the configuration for enabling jumbo frames. Although AoE is a niche protocol, customer feedback from management and performance perspectives is positive, as relatively little management is required — less than for FC, NFS and iSCSI. Performance increases automatically as additional Ethernet/AoE connections are installed. Drivers for some systems can be problematic to obtain. Prospective customers should determine availability and compatibility of drivers before purchasing storage that is connected using AoE. Presently, Coraid is the only vendor using this protocol.

FC

FC still has the lowest overhead and latency of all the mainstream protocols. Reliability and availability characteristics are inherent in the deterministic design of FC, and multipathing is simple, mature and well-supported. Host bus adapters (HBAs) can be tuned to change queue depth, timeouts, I/O tagging and logical unit number (LUN) limits. It is recommended that Peripheral Component Interconnect Express (PCIe) slots, where available, are used for the FC HBA connections. This is because FC is more latency-sensitive than iSCSI or file-based protocols, and 8 Gbps and 16 Gbps FC requires the higher performance provided by PCIe ports. Certain restrictions apply to FC in virtual server environments. Virtualized FC WWNs must be used and N-Port ID virtualization (NPIV) must be supported from the hypervisor, the physical HBA, SAN switch or SAN storage, if directly attached. When using NPIV, all storage or virtual-to-virtual (V2V) or P2V server/host moves must have the storage, HBAs and SAN switch ports in the same FC fabric zone. This can limit the flexibility of V2V or P2V server movement using solutions such as VMotion.

This research note is restricted to the personal use of mserale@villanova.edu

iSCSI

The iSCSI is a mature and low-cost protocol, as it can use standard Ethernet network interface cards (NICs) and TCP/IP networks. For the majority of applications, the performance of iSCSI is sufficient. However, iSCSI is mainly used in small or midsize business (SMB) and small data centers. TCP offload engines (TOEs) are no longer required today, as the drivers in the OSs are more efficient. TOEs are only required in the highest-performing latency-sensitive systems. When this level of performance is required, then use FC.

Mapping of iSCSI storage is slightly simpler than with FC, as the iSCSI storage is bound to the virtual hosts via the virtual IP and MAC addresses (vNIC). As virtual hosts move within and among physical servers, these addresses do not change and follow the virtual host. There are often requirements to have all iSCSI storage within the same subnet. Because iSCSI is routable, work-arounds are theoretically possible. Therefore, hypervisors can be limited when moving virtual hosts to another machine in a different subnet, because hosts are identified to one another by IP addresses. Similar to FC zones, this may limit V2V and P2V virtual host switching.

Block-based protocols FC and iSCSI have adapters that offload CPU resources into the HBA and provide the highest levels of connectivity performance, but at a cost. Traditional NICs have used CPU cycles to run iSCSI and TCP/IP stacks, but newer models provide both TCP and iSCSI offload.

File-Based NAS Protocols

The NFS NAS protocol is widely implemented with Unix systems, and widely used with VMware. Under most workloads and in most cases, the performance of NFS is equivalent to FC, but may require extra tuning in very-high-performance environments. NFS is easier to configure and manage than FC, but is limited by the virtual IP and subnet restrictions. Most virtualized server solutions can provide multipathing at the TCP/IP layer, which enables redundancy by moving IP addresses among ports. Similarly, NAS appliances implement path availability by using multiple connectivity and redundancy options between dual and multiple NAS heads.

CIFS is the NAS protocol associated with Windows. Because applications running in VMs are often running in Windows, it is not unusual for a NAS array to support VMware ESX with NFS, while the applications within the VMs are simultaneously using CIFS.

Block-based and file-based protocols can be configured to provide storage directly to a large virtual server, shared-hypervisor storage pool, from which all the virtual hosts are allocated storage. This reduces configuration time to the detriment of flexibility, as storage cannot be directly tuned or allocated specifically to each virtual host. IT operational departments need to determine their requirements according to performance and availability demands, whether to allocate storage directly to each virtual host or to a central hypervisor-managed pool that provides the next level of granularity to the individual virtual hosts.

Action Items:

- Connect storage to the servers with PCIe ports to avoid internal system bottlenecks, improve storage response times and enable hot swap ability/replacement and performance.

This research note is restricted to the personal use of mserale@villanova.edu

- To avoid performance problems, use the higher-performance 8 Gbps or 16 Gbps FC and 10GbE protocols for storage connections in virtualized server environments.
- Determine your storage change and usage rates, and chose your storage protocol based on ease of management, performance and component costs.
- Provide directly mapped or indirectly mapped storage to virtual hosts, as dictated by application performance, availability and storage administrative overhead.
- Consider mixing GbE iSCSI or NAS support for low-performance servers with FC for high-performance servers. This reduces interconnection costs and is supported by many arrays.

New Storage Technologies for Virtualized Environments

As server virtualization becomes more popular, old problems are solved and new problems appear. To resolve the new storage management problems caused by server virtualization, a new industry and class of products appeared hoping to benefit from the move toward server virtualization, but few products survived. The promise of new opportunities and markets created by server virtualization is difficult to ignore. Many new vendors have taken advantage of this trend, but the cost benefits, value proposition, advantages and disadvantages of these products need to be considered.

Software Storage Virtualization Layers

The installation of extra drivers into OSs that intercept I/O operations has been available for a long time. This methodology has again become popular due to the introduction and increased popularity of server virtualization solutions. Software layers — e.g., I/O operations drivers that are installed in the hypervisor, or I/O operations redirection techniques within the storage network — enable storage to be intercepted and managed closely to the server, without the need for any hardware device. This enables many features, such as replication and snapshot, which are traditionally performed within the storage array or appliance to be performed within the software stack. Additionally, it can allow the individual virtualized server to share storage at the hypervisor layer.

These new products and software layers are transparent to most storage resource management, server management or SAN management tools, and are difficult to monitor and combine into a capacity planning system. Although these solutions reduce configuration complexity, they introduce an extra layer of software to manage and a level of complexity to consider when problems occur. The solutions tend to increase, not decrease, storage requirements, while increasing the use of virtualized server resources in real time. Adoption of these technologies has been slow and mainly within midsize data centers, but has not expanded to larger environments because of software-increased change control and licensing overheads. During the last two years, Double-Take Software was acquired by Vision Solutions, and AutoVirt ceased to trade.

Action Items:

- Ensure that the software virtualization solutions are mature and reliable. Perform error-injection testing in test or proof-of-concept environments to determine their ability to recover from path or device errors.

- Determine whether hypervisor or virtualization software patches and upgrades can be performed independently of the storage virtualization software components and patches.
- Determine the software dependencies within these configurations. Can patches be installed? Can different versions of the storage virtualization software exist and work simultaneously?
- Determine whether future software upgrades will be affected by an extra level of software in the storage stack.
- Determine and evaluate the level of instrumentation, visibility and problem determination tools that are provided with each solution. These will be required during daily configuration and monitoring tasks, and in performance and problem diagnosis.

Example Vendors:

- DataCore Software, HP, Virsto, Vision Solutions

I/O Virtualization and Top-of-Rack Switches

A change in the market is that extra networking hardware layers are being created to manage storage within server virtualization environments. Many new vendors have jumped into this market, but adoption and success have not been automatic. Aprius was acquired by Fusion-io, Xsigo was acquired by Oracle, and Virtensys was acquired by Micron. The future of their products is in doubt. When physical hosts are consolidated into virtualized environments and the number of physical servers is reduced, the number of external storage and network connections within a rack may actually grow. In many cases, a 1 rack unit (RU) or 2RU server may require as many as six or eight network ports, or even more, to support the following requirements:

- Two ports of 8 Gbps FC for SAN attachment
- Two ports of 1 Gbps Ethernet for LAN attachment
- One port of 1 Gbps Ethernet for each — hypervisor out-of-band management and backup

Additionally, cable management for a rack containing 20 or more servers becomes complex, increasing costs and reducing reliability due to human error. This drives many companies to consider I/O virtualization strategies that replace discrete Ethernet FC adapters with a pair of high-performance and low-latency connections, such as InfiniBand, PCIe bus extensions or converged I/O adapters. In this scenario, the physical NICs and HBAs are replaced with virtual adapters within a ToR switch. The number of logical connections is maintained, but the number of physical adapters and the associated cabling is reduced.

A ToR switch can be deployed using a variety of technological approaches, including:

- IP SAN (iSCSI) and IP LAN over a two-port 1 Gbps or 10 Gbps Ethernet NIC
- FC SAN and IP LAN over a two-port 10 Gbps FCoE converged network adapter (CNA)
- IP or FC SAN and IP LAN over a two-port InfiniBand adapter

- IP or FC SAN and IP LAN over a two-port PCIe extender/adapter

In nearly all cases, the server links connect to ToR switches that break out the SAN and LAN traffic for connection to the respective backbones. These switches may be transparent to most storage resource management tools; therefore, the relationship between the applications and storage cannot be reconciled or managed. Similarly, storage performance and problem determination can become more complex. However, they are being increasingly integrated into server-oriented virtualization and system management frameworks. The most mature products are iSCSI- and InfiniBand-based, although all major server and networking vendors are devoting significant resources to FCoE-based solutions.

Actions Items:

- Use IO virtualization and ToR architectures to consolidate and simplify server SAN and LAN access.
- Ensure that the performance requirements of all the combined virtualized servers within a rack can be provided by the selected ToR switches. Areas that need to be considered are server-to-server and server-to-core network bandwidth, as well as the overhead of any protocol conversion functions (e.g., FCoE to FC).
- Select switches that support active upgrades and hot field-replaceable units so you won't need to shut down servers during a scheduled maintenance or when upgrading the ToR switch. Alternatively, you can select dual, multipath switch configurations that can be restarted and upgraded in turn.
- Check to see whether the consolidation switch is compatible with standard high-volume low-cost adapters, or whether special higher-cost adapters are required.
- Verify that traditional storage fan-out architectures cannot be used with existing FC, FCoE and iSCSI solutions before investigating ToR I/O virtualization switches.
- Verify that the total cost of ownership of the new ToR I/O virtualization switch infrastructure is less than that of traditional fan-out architecture, as described above.

Example Vendors:

- Blade Network Technologies, Brocade, Cisco, HP, Micron, Oracle

Business Continuity and Disaster Recovery

Most server virtualization solutions do not improve the availability of virtualized hosts from a business continuity perspective, where the disaster recovery sites are geographically distributed over metropolitan-area networks (MANs) or longer distances. Within a data center, when a virtual host is switched between one physical server and another, but the data and storage do not move or need to be reconfigured, the process of moving the virtual host is easy and can be accomplished in a subsecond time frame. In this case, the switch is limited by the speed at which the virtual host can be transferred from one physical server to another and the time to reconfigure the network

This research note is restricted to the personal use of mserale@villanova.edu

definitions. Storage is not transferred or moved and, therefore, does not slow down the switchover time.

In a geographically dispersed disaster recovery scenario the time to move or switch applications is predominantly determined by the time to move the storage. This is because the data or storage also needs to be moved from one site to another, and the secondary and remote copy of data has to be reconfigured to become the primary storage in the disaster recovery site and then checked for consistency. In this situation, the recovery of the virtual hosts is determined by the much longer time to switch and reconfigure the storage between the disaster recovery sites, not the time to switch and move the virtual hosts, which can be done much faster. The reason is that for the host, mainly state information needs to be transferred.

While tools exist to manage and automate the switchover, the fundamental storage replication functionality and time delays or restrictions are applicable to all servers. Therefore, processes such as file system checking are not removed within server virtualization solutions if a system was not shut down gracefully. In a MAN or wider business continuity virtualized server infrastructure, total switchover and failover time can take several minutes or longer, as the fundamental process of storage replication and switchover has not been solved or improved by the server virtualization solutions. Take this into account when designing business continuity plans for virtualized servers.

Action Items:

- Test disaster recovery processes at least every quarter.
- Determine the availability and data integrity requirements of the application running on the virtual hosts, and implement synchronous or asynchronous replication as necessary.

Example vendors:

- All vendors that offer storage arrays, appliances and software replication solutions that perform synchronous or asynchronous replication. Software products that monitor and can simulate disaster recovery testing are Continuity Software and VirtualSharp.

Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"Prioritize Software Features When Buying SSD Arrays"

"Storage Management Tools in a Virtualized Server Environment"

"The Future of Storage Management"

"Increases in Disk Capacity Are Affecting RAID Recovery; Start to Purchase New RAID Technologies"

"How Much and What Type of Disk Storage Do IT Departments Need?"

"Innovations in Storage Technologies Are Not Enough to Reduce Storage Costs"

"Application Requirements Must Drive Storage Purchasing Decisions"

"Recommendations for SAN Fabric Dashboards"

"Recommendations for a Storage Array Dashboard"

"When Migrating Storage, Use the Tools in Server Virtualization Products"

This research note is restricted to the personal use of mserale@villanova.edu

GARTNER HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters
AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit http://www.gartner.com/technology/about.jsp