

Realtime
publishers

The Shortcut Guide to Balancing Storage Costs and Performance with Hybrid Storage

sponsored by



Dan Sullivan

Chapter 2: Hybrid Storage for Virtualization	15
Mixed-Workload Server Virtualization.....	16
Varying Demand for Resources on Virtualized Servers	16
Demands on Systems Administrators in a Virtual Environment.....	18
<i>Optimizing Storage</i>	18
Sizing IOPS to Meet a Mixed Set of Application-Specific Demands.....	21
Provisioning Adequate Storage.....	22
<i>Quality of Service Controls for Storage Performance</i>	23
Implementing Fast Disaster Recovery Mechanisms	25
Desktop Virtualization	27
Challenges to Implementing Desktop Virtualization	27
Storage Requirements for Desktop Virtualization.....	29
Benefits of Hybrid Storage for Server and Desktop Virtualization.....	30

Copyright Statement

© 2014 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Chapter 2: Hybrid Storage for Virtualization

IT departments across the globe are under constant pressure to deliver consistent high-quality services and respond to changing demands of dynamic business environments all while cutting or controlling costs. A common approach to meeting these challenges is virtualization: running multiple versions of a virtual resource on a single actual resource. Server virtualization platforms such as VMware vSphere and Microsoft Hyper-V allow multiple virtual machine instances to run on a single physical server. This approach to cutting costs is enabled by the dramatic increase in performance of multi-core CPUs.

If CPU performance were the only critical element to overall virtualization performance, the work life of systems and application administrators would be much easier than it is in reality. The problem is that storage performance is also a key factor to overall application performance in a virtualized environment. Although there have been steady and significant improvements in the capacity of storage systems based on hard disk drives, the same cannot be said for the performance of storage systems.

Storage systems often become a bottleneck in virtualized environments, which can lead to a number of common problems, including:

- Inconsistent application performance,
- Degraded application performance,
- Inability to meet service level agreements (SLAs),
- Over-provisioning for disk capacity to address worst-case scenarios, which leads to
- High storage costs that undermine the cost benefits of server virtualization.

In addition to these problems familiar to systems administrators working with mixed-workload virtualized environments, there are performance challenges that surface in virtualized desktop infrastructure (VDI) environments. For example, when a large number of users begin working with their desktops at the beginning of a workweek or after system maintenance, there is a spike in IOPS as virtual desktops are started. These spikes are known as “boot storms” and can delay users as they wait for their slice of limited storage throughput.

In this chapter, the focus shifts to the benefits of hybrid storage in virtualized environments. In particular, this chapter addresses:

- Mixed-workload server virtualization
- VDI
- Benefits of hybrid storage for mixed-workload server virtualization and VDI

Hybrid storage systems take a significant step toward closing the performance gap between CPU performance, enabling virtualization and limited storage performance. This setup prior to the development of hybrid storage, created a rate-limiting bottleneck in servers and desktop virtualization performance.

Mixed-Workload Server Virtualization

Prior to the advent of server virtualization, it was a common practice to dedicate one or more physical servers to a single application. For example, systems administrators might dedicate a single server to Microsoft Exchange, another to Microsoft SQL Server, and another to Microsoft SharePoint. This practice streamlined systems administration but led to underutilization of CPU resources. The Microsoft Exchange server might be running at more than 90 percent CPU utilization while the Microsoft SharePoint server was running well below 20 percent CPU utilization and the database server was running at 40 percent utilization. Obviously, both the collaboration application and the database could be accommodated on a single server.

Varying Demand for Resources on Virtualized Servers

With server virtualization, multiple operating systems (OSs), each running different applications, run on the same physical server. Systems administrators could carefully craft the combination of applications running on a single physical server. For example, a single physical server could host a collaboration application such as Microsoft SharePoint that has peak utilization during the business day, with another server running extraction, transformation, and load (ETL) operations for a data warehouse that runs primarily outside of normal business hours. Not all use cases provide this kind of well-defined segregation of peak demand periods.

As Figure 2.1 shows, when peak demand periods of virtual machines running on the same physical server are not at the same time, there is little contention for limited resources, such as IOPS.




Figure 2.1: Ideally, in mixed-workload environments, the peak demand periods for each virtual machine running on a server do not overlap.

In reality, mixed workloads are more varied in their demand. The problem becomes more pronounced as more virtual machines are run on the same physical server.




Figure 2.2: Mixed workloads can have unpredictable periods of high- and low-demand for resources, including IOPS.

Mixed workloads and varying demands create challenges for administrators who are trying to optimize the utilization of all resources: compute, storage and network.

Demands on Systems Administrators in a Virtual Environment

Administrators are faced with multiple challenges when managing virtualized environments:

- Optimizing storage
- Sizing IOPS to meet a mixed set of application-specific demands
- Provisioning adequate storage capacity
- Implementing fast disaster recovery mechanisms

These demands reflect the need to control costs while maintaining SLAs. Server virtualization is a complex technical issue, but how one implements and manages such an environment is ultimately driven by business concerns.

Optimizing Storage

When optimizing storage is a top priority in virtualized environments, shared storage is often used. In fact, the vast majority of virtualized environments depend on shared storage. This choice is logical for a number of reasons.

Benefits of Shared Storage

Shared storage decouples storage from individual servers. There is less need for direct-attached storage (DAS) in each physical server. This setup helps avoid the same kind of underutilization that plagued CPUs prior to virtualization. For example, an IT department may have a standard set of server configurations it uses for all applications. Some applications may require less storage than others. In a standard configuration, all servers have the same amount of storage regardless of whether the applications running on that server need all the storage. Capacity that is underutilized in one server cannot be shared and is therefore wasted. This situation is analogous to all servers having the same set of CPUs even though some applications underutilize the CPUs.

In addition to more efficient use of storage capacity, shared storage simplifies management. Storage is treated as a single logical resource that can be used by multiple servers. Servers are allocated storage according to their needs. As needs change, more or less storage can be allocated as necessary.

Perhaps most important, shared storage enables advanced virtual machine management features in virtualization platforms, such as automatic machine migration. This feature allows for a virtual machine to be moved from one server to another without disrupting operations. This functionality could be needed, for example, if there is a hardware failure in one server or the virtualization platform moves a virtual machine to better balance load across all servers. With shared storage, the virtual machine has access to the persistent storage it is using regardless of which physical server it is running on (see Figure 2.3).




Figure 2.3: Virtual machine migration depends on shared storage (Server image source: RRZEIcons (Own work) [CC-BY-SA-3.0 (<http://creativecommons.org/licenses/by-sa/3.0>)], via Wikimedia Commons).

Although there are definite advantages to using shared storage with respect to optimizing storage, there are drawbacks. These drawbacks are not a sufficient reason to avoid shared storage, but they do encourage the use of hybrid storage that combines the performance benefits of flash devices with the capacity benefits of hard disk drives.

Disadvantages of Shared Storage

The primary disadvantage of shared storage in a virtualized environment is that it can lead to storage resource contention. There may be periods of time when two or more virtual machines are performing large numbers of reads and writes to shared storage. These periods can lead to longer queues of I/O operations and delays in finishing those operations.

This reality can be particularly problematic when a single server is hosting VMs with a combination of business-critical and non-critical applications. A non-critical application, such as a backup, can perform a continuous stream of I/O operations at the same time a large number of users are running SQL queries against a database on a virtual machine hosted on the same server as the virtual machine running the backup. Consider a possible scenario for this configuration of virtual servers:

1. A weekly full backup that is automatically scheduled to run at this time begins executing.
2. The backup process generates a large number of write operations.
3. On the other virtual machine hosting a relational database, a number of users run reports. These reports generate a large number of SQL SELECT queries.
4. The SQL SELECT queries cause index reads and then reads of data blocks.
5. At the same time, an application using the database begins executing a maintenance procedure that issues a number of UPDATE queries that require reads to select particular rows and writes to update those rows.

In this simple example of just two virtual machines on a single server, it is apparent that I/O operations are generated faster than hard disk drive-based storage can complete the operations.

In addition, write operations can be particularly problematic. Applications have to wait for acknowledgement of writes completing, which can be slowed by the need to replicate data and calculate and write parity bits.

It is tempting to think that systems administrators can assess different combinations of workloads and configure virtual machines on servers to minimize this kind of contention. This optimized setup is possible, at least in theory, but there are still issues:

- Systems administrators may not have detailed knowledge of I/O operation patterns of all applications
- Applications' I/O read and write operations might be highly variable, and predicting loads over time may not be a practical option
- Business requirements change frequently; systems administrators would have to rebalance the mix of virtual machines to keep up with changes in application workloads related to changes in business requirements
- To realize the benefits of automatic virtual machine migration, such as optimal use of CPUs and automatic failover, systems administrators should enable dynamic resource management features of the virtualization platform

Fortunately, systems administrators are no longer faced with choosing between the benefits of automatic resource allocation by virtualization and predictable performance. The combination of hybrid storage, which provides performance levels to meet demanding business applications, coupled with policy management that allocates IOPS based on application priority, provide the opportunity to realize consistent performance without sacrificing advanced virtualization management features.

In addition to optimizing storage, systems administrators have to manage IOPS. Like optimizing storage, this demand is well addressed by hybrid storage systems.

Sizing IOPS to Meet a Mixed Set of Application-Specific Demands

Sizing storage systems to meet IOPS requirements is often a problem with a moving target. It is hard to imagine a realistic, enterprise application that has a consistent load of I/O operations and the same mix of reads and writes for extended periods of time.

A typical scenario is that servers hosting virtual machines are subject to randomized patterns of I/O operations. This situation is known as the I/O blender effect (see Figure 2.4). This state is caused when virtual machines each send their I/O requests to the hypervisor and the hypervisor interleaves the I/O requests from different virtual machines.




Figure 2.4: Hypervisors on virtualized servers interleave I/O requests from virtual machines, creating the I/O blender effect.

Consider a SQL query that does a full table scan on a large table. If the database is running on a dedicated server, a series of sequential reads are sent to the I/O controller. Assuming the database table storage has been optimized, the full table scan can be accomplished using a series of sequential reads of data blocks on a disk.

Now assume the database server is running in a virtual machine on a server hosting multiple virtual machines. The other virtual machines send I/O requests to the hypervisor, which in turn, queues some of the database I/O operations and then some of the other virtual machine I/O requests, followed by another set of database I/O operations. By the time the second set of database operations is processed, the disks, which are continuously rotating, may no longer be in a position to read a database data block. The database now has to wait until the disk is in position to read the next needed data block, thus increasing the latency of the I/O operation.

One approach to dealing with the I/O blender effect is to overprovision. By adding more disk capacity to meet peak performance demand, systems administrators can ensure there are sufficient IOPS to counter the impact of I/O blender effects. The problem with overprovisioning, of course, is the additional costs as well as surplus capacity and footprint.

Flash devices do not have moving parts such as spinning platters of disks. The latency of random reads is not worse than sequential reads, so the I/O blender effect is less of an issue when employing flash devices. As with over-provisioning, with flash devices, administrators have to consider the cost of deploying flash. For this reason, hybrid storage is an effective way to balance performance and cost considerations. Another systems administration consideration in virtual server environments is provisioning sufficient storage.

Provisioning Adequate Storage

When provisioning storage, systems administrators have to consider both the size of the storage system and its performance (see Figure 2.5). A particularly difficult challenge is addressing the need for different levels of performance. Mission-critical applications should not suffer performance problems because a low-priority job is performing a burst of I/O operations on shared storage. This type of resource contention leads some to overprovision and plan for peak possible demand.

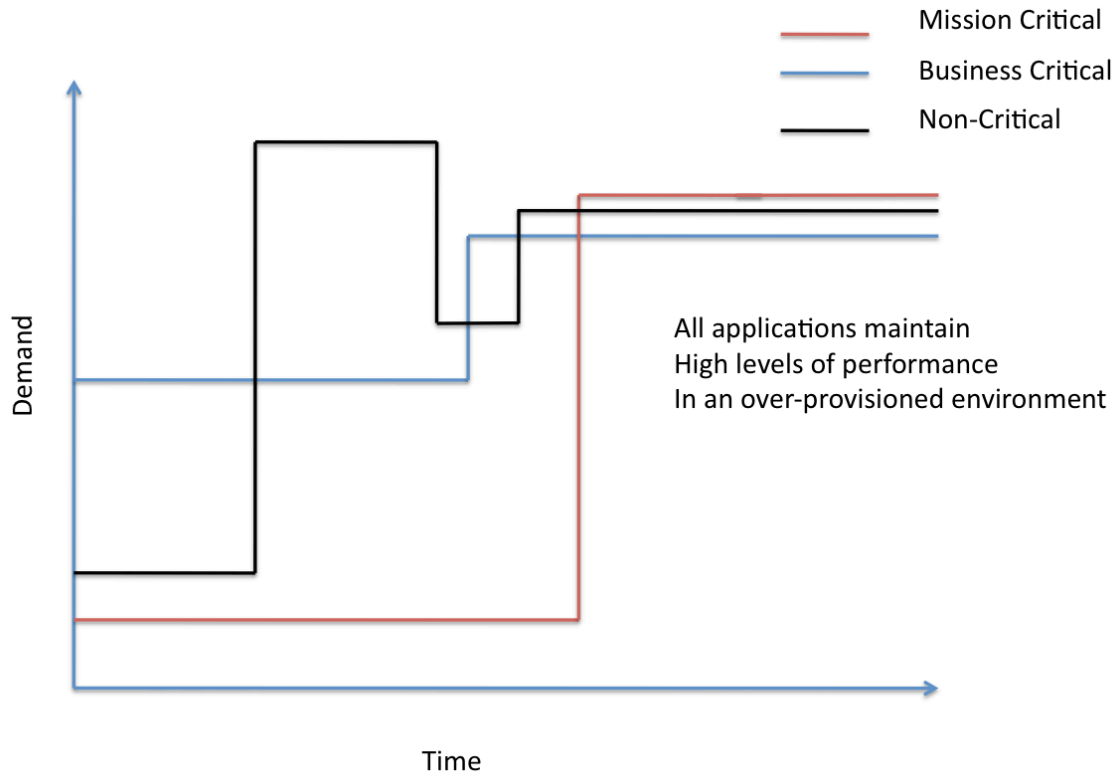


Figure 2.5: In an over-provisioned storage environment, all applications maintain high performance but at additional costs. Non-critical applications could run at lower storage performance without adversely effecting business operations.

Quality of Service Controls for Storage Performance

A better approach, at least from a cost control perspective, is to manage I/O operations according to quality of service needs. This practice is well-established in networking but it is less widely implemented in storage management. The basic idea is that applications should be categorized according to their importance to the business mission and resources should be allocated accordingly. One categorization approach is:

- Mission-critical
- Business-critical
- Non-critical

Mission-critical applications should receive the highest priority in the allocation of storage resources; business-critical are mid-level priority, and non-critical applications receive the lowest priority, as Figure 2.6 illustrates.

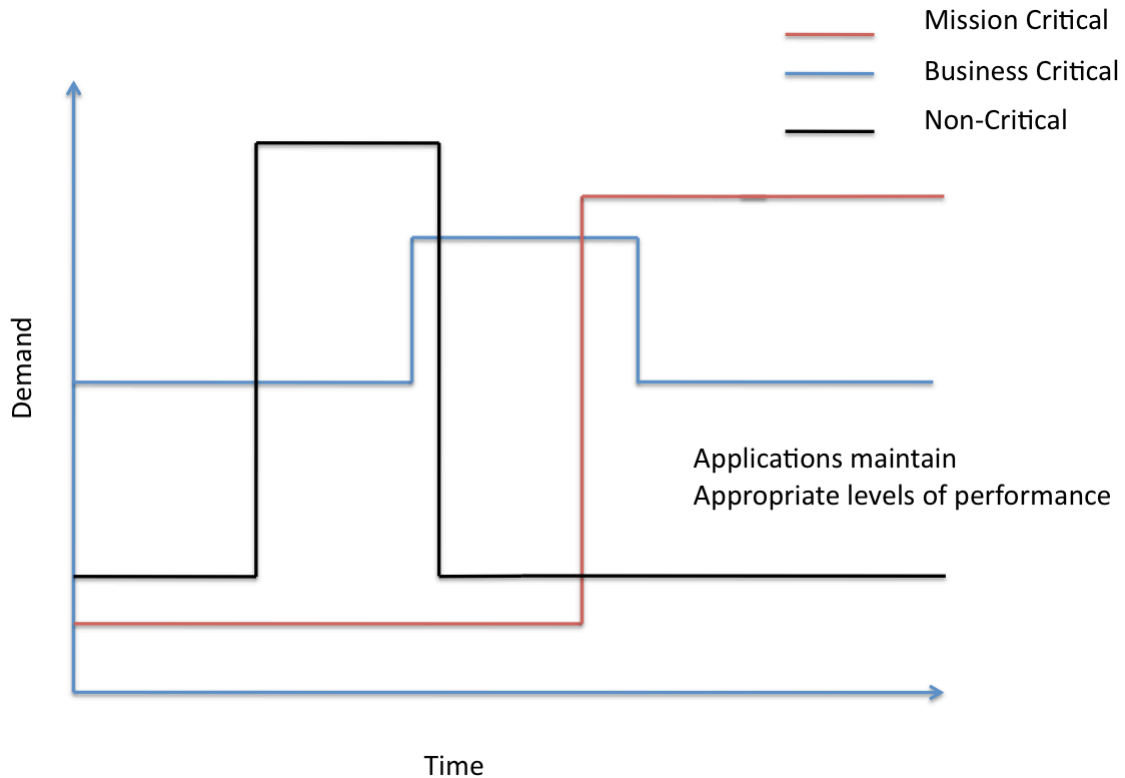


Figure 2.6: When quality of service controls are in place for storage performance, applications maintain appropriate levels of storage performance.

Provisioning storage with quality of service controls helps to minimize uncertainty about application performance. Each application can be assigned to an appropriate category, ensuring that it will maintain high levels of storage performance if needed. It also allows systems administrators to designate non-critical applications that can function with lower levels of application performance.

Note on Determining Criticality of Applications

It should be noted that designating some applications as mission-critical and others as non-critical does not imply some applications are not important. Virus scans and backups, for instances, are important operations that must be performed. In that sense, they are critical to business operations; however, from the perspective of storage performance, it is less important that applications receive priority over other applications. Virus scans and backups must run and complete successfully, but if they take longer than expected in some cases, it would not have the same level of adverse impact as a mission-critical application suffering poor storage performance.

The goal of provisioning sufficient storage in a virtual environment entails having sufficient quantity of storage and adequate storage performance levels for the mix of applications running in the environment. Hybrids alone can help ensure adequate quantities of storage and improve overall performance of storage systems, but quality of service control is an additional feature not available in all hybrid storage systems. When evaluating hybrid storage systems, consider the cost advantages of enabling quality of control services. There is a limited amount of flash in a hybrid storage system, so it must be managed to ensure the limited resource is used in an optimal way. These controls are a key feature that help avoid the need to over-provision storage systems in order to maintain adequate performance levels of mission- and business-critical applications.

Another feature of storage systems in virtual environments that should be considered is the speed with which business services can recover from disaster.

Implementing Fast Disaster Recovery Mechanisms

Planning for disaster recovery is a top concern and top challenge with regards to enterprise virtualization strategies. An Enterprise Strategy Group study from 2011 found that disaster recovery is considered as important of a challenge as is sizing true capacity and IOPS requirements (see Figure 2.7).

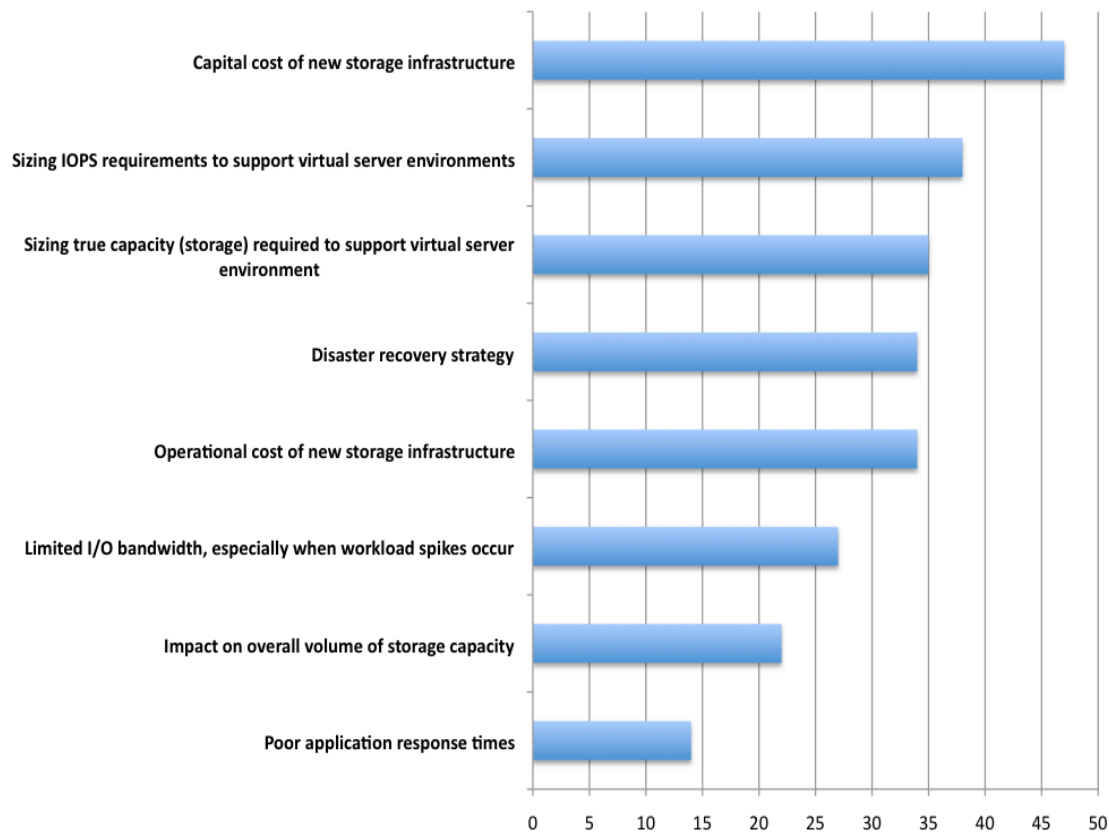


Figure 2.7: Storage challenges encountered with server virtualization (Data source: Enterprise Strategy Group, 2011).

In a virtualized server environment with a consolidation ratio, there are significantly more virtual machines than physical servers. This setup can potentially create the need to replicate a large number of virtual machines. The speed improvements of I/O operations on hybrid storage over hard disk drive storage systems can help reduce the time required to replicate virtual machine images.

There are several demands on systems administrators responsible for managing storage in a virtual server environment, including optimizing storage, sizing IOPS to meet application-specific needs, provisioning adequate storage, and implementing fast disaster recovery mechanisms. A key element common to the solution of all of these challenges is hybrid storage.

Hybrid storage enables optimal storage capacity and IOPS sizing and, when coupled with quality of service controls, can help minimize the need for over-provisioning. Many of the issues considered for mixed application virtual environments also apply to a more specialized form of virtualization: desktop virtualization. VDI environments, however, have some unique characteristics worth considering with regards to storage.

Desktop Virtualization

Desktop virtualization is technology for separating tightly coupled OSs and applications with underlying end-user hardware. The software components of a desktop are virtualized to run on servers in a data center rather than on local hardware. There are many advantages of virtual desktops:

- Centralized management of desktop images
- Data is saved and backed-up centrally
- Potentially improved security
- Access to desktops from multiple devices
- Potentially reduced maintenance costs

Along with these benefits of desktop virtualization come a number of challenges for systems administrators.

Challenges to Implementing Desktop Virtualization

Perhaps the most significant challenge to desktop virtualization in many businesses is user acceptance. The benefits of desktop virtualization, particularly the ability to access a persistent virtual desktop from any device, can be undermined by poor performance. If virtual desktops require long periods of time to boot or application performance is sluggish, users may opt not to use the virtualized desktops that IT departments spend significant time and resources to implement.

Consider the common problem of boot storms, which occur when a large number of users start their virtual desktops at the same time. Many of us work on similar schedules, so it is not surprising that multiple users will need to start their desktops around the same time (see Figure 2.8).




Figure 2.8: During a boot storm, multiple virtual desktops boot at the same time putting increased demand for IOPS on shared storage servers.

One solution to the problem of boot storm is over-provisioning. Administrators could configure shared storage systems to meet the IOPS and throughput requirements of peak demand, but this approach has cost implications, as discussed earlier. As with server virtualization, the cost savings that might be realized by implementing virtual desktop environments can be quickly undermined with the cost of additional storage hardware.

One of the problems with provisioning for peak demand with hard disk-based storage is that the cost of achieving the necessary IOPS can be quite high. Disk-based storage systems might provide in the range of 4,000 to 6,000 IOPS, but boot storms could require 10,000s of IOPS to maintain adequate performance.

Hybrid storage is once again a cost-effective solution to this problem. Flash devices provide significantly higher IOPS than do hard disk drives, so virtual desktop users can receive the overall device performance they expect and IT departments can deliver that level of service without unnecessarily high capital expenditures.

Storage Requirements for Desktop Virtualization

To successfully implement VDI, IT professionals have to address three key requirements:

- Storage system performance
- Cost
- Desktop image management

VDI has similar demands as mixed-workload virtualization in some cases and may utilize the same shared storage as mixed-workload virtualized servers. It is essential to provide sufficient IOPS and throughput to simulate a near equivalent experience of using a physical desktop. Users may tolerate some variation in response time and even some level of poorer performance, but to ensure adoption, users should not have to compromise performance for the other benefits of a virtual desktop.

The cost factor of shared storage is important as well. Trying to achieve sufficient storage performance using only hard drive-based storage systems can drive up the cost of VDI in unexpected ways. For example, hard disk drive-only storage arrays may not be able to provide sufficient IOPS to meet peak demands when all virtual desktop users are using the same storage systems. In such cases, IT departments may turn to using multiple storage arrays and dividing virtual desktop users across multiple storage systems. This setup has obvious and significant cost implications for delivering VDI.

The methods of desktop image management will vary, to some degree, according to how VDI is deployed. In some cases, users are assigned desktops from a pool of identical desktops. Users do not have the ability to change these desktops, so they can be assigned any desktop from the pool without altering their user experience. Other VDI implementations provide for persistent desktops. In this case, users can customize their desktops and save those changes. Copies of individual desktops are maintained for each user. Systems administrators have to plan their storage capacity according to the types of desktops (pooled or persistent) that they plan to deploy.

Desktop virtualization faces many of the same challenges as mixed-workload virtual environments as well as some particular to VDI. In both cases, hybrid storage can deliver the performance required to meet business and user requirements while maintaining favorable cost profiles when compared with hard disk-only storage solutions.

Benefits of Hybrid Storage for Server and Desktop Virtualization

The preceding discussion of server virtualization and desktop virtualization presents a series of challenges that systems administrators must overcome to deliver cost-effective and adequate storage performance. These include:

- Sufficient IOPS
- Quality of service
- Balancing sufficient IOPS with capacity

CPUs have steadily increased in performance at a rate faster than the speed of disk drives. Although the capacity of hard drives has grown steadily, the rate of growth in IOPS has not kept a similar pace. In order to provide sufficient IOPS for demanding applications, a combination of flash devices and hard disk drive technologies are offering an alternative to storage systems based solely on hard drive technology.

Although flash devices can increase the IOPS performance and throughput of a storage device, systems administrators may still be tempted to over-provision in order to ensure adequate performance of mission-critical and business-critical applications during peak demand periods. This problem is better addressed by using hybrid storage systems that support quality of service controls. With such controls, applications can be categorized as mission-critical, business-critical, and non-critical. During times of high demand for storage resources, those resources can be allocated to the most important applications and thus reduce the need to over-provision.

In addition to having adequate performance in a virtualized environment, administrators must consider capacity. This scenario is where hard disk drive technology excels. The cost per gigabyte of hard disk storage has decreased steadily over the past few decades. By combining flash devices with hard disk drives, hybrid storage systems can deliver a balanced storage solution that provides both the performance and capacity needed for today's enterprise applications.