Using VMware vSphere and IBM Storwize V7000 as a complete virtualization solution

A dynamic and responsive infrastructure for ever-changing business needs and requirements

Rawley Burbridge

IBM Systems and Technology Group ISV Enablement
January 2011

Table of contents

Abstract

The intent of this paper is to examine the virtualized environment that IBM Storwize V7000 and VMware vSphere provide. VMware vSphere when combined with the IBM Storwize V7000 provide a complete server and storage virtualization solution in which the infrastructure is dynamic and responsive to ever-changing business needs and requirements.

Executive summary

Server virtualization with VMware vSphere has been successful in helping customers utilize hardware more efficiently, increase application performance and availability, and decrease management costs. Similar to how VMware vSphere abstracts server workloads from the underlying hardware to achieve server virtualization; IBM® Storwize V7000 abstracts volumes from the physical storage resulting in a virtualized storage environment.

This paper examines an environment using storage provided by IBM Storwize V7000 and server virtualization with VMware vSphere. An insight is provided on the solution that these two products combine to create and how it increases the flexibility, responsiveness, and availability of application workloads. The paper also discusses and demonstrates real-world scenarios. The paper helps the reader to understand the reason for opting IBM Storwize V7000 as the storage choice for VMware vSphere deployments.

Purpose of the paper

The intent of this paper is to examine the virtualized environment that IBM Storwize V7000 and VMware vSphere provide. Real-world scenarios and how product functions and features apply to those scenarios are also discussed. The paper provides an overview of VMware vSphere and IBM Storwize V7000 technologies and also discusses and demonstrates the key points valuable to a virtualized environment.

This paper helps the reader to understand why storage virtualization is important and the value that storage virtualization can add to a VMware vSphere environment.

Although best practices are provided in this paper it is not a best practice reference guide. Best practices and deployment guidelines are provided in *Deploying IBM Storwize V7000 in VMware Environments*, which can be found in the IBM TechDocs repository.

IBM TechDocs: **ibm.com**/support/techdocs/atsmastr.nsf/Web/TechDocs

Additionally some performance aspects are discussed in this paper with the primary purpose of exploring the performance-tuning capabilities of the IBM Storwize V7000.The Enterprise Strategy Group has prepared a report titled *Real-world Mixed Workload Performance in VMware Environments* which examines the performance characteristics of the Storwize V7000 and virtualized workloads.

This paper can be accessed from the IBM ISV Storage Solutions website.

ISV Solutions: **ibm.com**/systems/storage/solutions/isv/isv_vmware.html

This paper should not be used for hardware certification. For information on official hardware certification refer to the VMware Hardware Compatibility List and IBM System Storage Interoperability Center (SSIC) websites.

VMware Hardware Compatibility List:
http://www.vmware.com/resources/compatibility/search.php?ie=UTF-8&q=vmware%20hcl

IBM SSIC: **ibm.com**/systems/support/storage/ssic/interoperability.wss

Solution overview

The core function of virtualization is abstraction. Both VMware vSphere and IBM Storwize V7000 abstract workloads from the underlying hardware and then add value through features and functionality.

VMware vSphere

VMware vSphere is a virtualization platform capable of transforming a traditional datacenter and industry standard hardware into a shared mainframe-like environment. Hardware resources can be pooled together to run varying workloads and applications with varying service-level needs and performance requirements. VMware vSphere is the enabling technology to building a private or public cloud infrastructure.

The components of VMware vSphere fall into three categories: Infrastructure services, application services, and the VMware vCenter Server. Figure 1 shows a representation of the VMware vSphere platform.

Figure 1: VMware vSphere platform

Infrastructure services

Infrastructure services perform the virtualization of server hardware, storage, and network resources. The services within the infrastructure services category are the foundation of the VMware vSphere platform.

vCompute

vCompute is the title given to the services that virtualize server resources. These services also aggregate server hardware into logical pools which allows management of resources instead of managing megahertz and megabytes.

ESX and ESXi

VMware ESX and ESXi are the hypervisor choices available for the vSphere platform. They provide a production proven high-performance virtualization layer built for running multiple virtual machines on a single physical server. Using advanced memory management and processor scheduling techniques, VMware vSphere ESX and ESXi hosts are able to achieve high consolidation ratios.

Dynamic Resource Scheduler (DRS) and Distributed Power Management (DPM)

DRS is a service that allows aggregation of resources from multiple VMware vSphere hosts and then performs allocation of the resources based on virtual machine and workload needs. Production workloads can be guaranteed resources while combined with nonproduction workloads. DPM works in conjunction with DRS to manage the energy consumption of vSphere environment. When workload demands decrease, for example, during off-peak hours, workloads are consolidated and vSphere hosts are put into energy-saving states based on need.

vStorage

vStorage services interact with the storage resources used by a VMware vSphere environment and help to simplify the management, increase efficiency, and improve the utilization of the storage.

Virtual Machine File System (VMFS)

VMFS is a purpose built clustered file system for running virtualized workloads in a shared environment. VMFS takes away the complexities usually found in a clustered environment and enables easy setup and scaling of VMware vSphere environments.

Thin provisioning

Virtual machine disk space consumption is allocated on demand as space is needed when thin provisioning is used. Thin provisioning can result in storage space use reduction as virtual machines are most often over provisioned. It also allows the IT administrator to

allocate storage for the environment when it is actually needed, resulting in a more efficient use of storage.

Storage I/O control

A new feature released with VMware vSphere 4.1, storage I/O control, prioritizes access to storage resources based on established workload needs. Storage I/O control can ensure that production workloads are guaranteed storage resources, and unreliable workloads do not monopolize and impact other workloads.

vNetwork

Virtualization introduces additional complexities to networking that the vNetwork services attempt to address.

Distributed switch

The vNetwork distributed switch is a service which allows management of a virtualized switch for an entire VMware vSphere host cluster instead of managing virtual switches for each host independently. Management is simplified and access policies are enforced at a centralized location for a VMware vSphere host cluster.

Network I/O Control

10 GbE is becoming an adopted technology by more customers. As I/O is consolidated onto fewer paths with 10 GbE, it is important to prioritize access to network resources. Similar to how storage I/O control prioritizes access to storage resources, Network I/O Control prioritizes access to network resources to ensure workload service levels.

Application services

The components categorized as application services address availability, security, and scalability concerns for all applications running on the vSphere platform, regardless of complexity of the application.

Availability

A key component of virtualization is the encapsulation of workloads. A server, the services, and applications running on it, regardless of the operating system or applications, are encapsulated into a container which is managed as a single unit. Because of this encapsulation, VMware vSphere availability services address planned and unplanned downtime of the workload agnostically.

vMotion

vMotion is the live migration of running virtual machines between VMware vSphere hosts with no downtime or service disruption to applications. vMotion eliminates downtime for planned maintenance tasks.

Storage vMotion

Virtual machine disks can be non-disruptively migrated between VMFS volumes with storage vMotion. The feature helps to avoid application downtime needed for planned storage maintenance. During migrations, virtual machine disks can also be converted to different disk types. Converting virtual machine disks to be thinly provisioned during a migration is a common use case.

High availability (HA)

VMware HA is a vSphere host clustering solution which automates application restarts in an unplanned downtime event. Operating system clustering is often costly and complex. HA provides a clustering solution which is operating system agnostic, is easy to setup, and can support the most simple or complex applications.

Fault Tolerance

Fault Tolerance is an innovative technology which provides continuous availability for a virtual machine or application. Fault Tolerance works by creating a live shadow copy of a virtual machine on another host in the vSphere cluster. The primary and secondary virtual machines operate in lockstep so that even in the event of an unplanned downtime, services are not interrupted.

Security

The VMware vSphere platform includes components that enforce security policies and network segmentation in a shared environment.

vShield zones

vShield zones is a hypervisor aware firewall that is easily configurable, yet a powerful tool for enforcing security policies in a virtualized environment. You can configure security for an application, regardless of whether that application consists of a single virtual machine or a group of virtual machines.

Scalability

VMware vSphere meets the requirements of the most demanding application workloads and maintains top performance through automated load balancing with VMware DRS and hot add enables processor, memory, network, and storage resources to be nondisruptively added to virtual machines.

IBM Storwize V7000

The IBM Storwize V7000 enables customers to improve application flexibility, responsiveness, and availability while reducing storage usage and complexity through storage virtualization and the following features:

- Metro Mirror and Global Mirror perform synchronous and asynchronous data replication between IBM Storwize V7000 systems at varying distances to protect data and keep services online in downtime situations.
- IBM Tivoli® Storage FlashCopy® Manager creates instant volume copies for data protection and flexibility.
- Virtualization enables volume migration and mirroring to any storage managed by the IBM Storwize V7000 system.
- IBM Easy Tier™, thin provisioning, and an intuitive easy-to-use graphical user interface increase efficiency for the storage and storage administrator.

IBM Storwize V7000 is a modular storage system built from the IBM SAN Volume Controller technology base and utilizing Redundant Array of Independent Disks (RAID) technology from the IBM System Storage® DS8000® family to deliver a virtualized, enterprise-ready, mid-range storage solution. The subsequent sections describe the IBM Storwize V7000 hardware and concepts.

Storwize V7000 hardware and concepts

A Storwize V7000 system is made up of 2U drive enclosures. One control enclosure with the option to add up to nine expansion controllers is required for additional capacity. All enclosures have the option of 12 drive slots (as shown in Figure 2) or 24 drive slots (as shown in Figure 3).

Figure 2: Storwize V7000 12 disk enclosure

Figure 3: Storwize V7000 24 disk enclosure

Within the control enclosure are two node canisters, also known as an I/O group (which form an active-active processing unit) act as the management point for the storage system and service all I/O to host systems.

Storwize V7000 may have between one to nine expansion enclosures which are connected to the control enclosure through redundant serial-attached SCSI (SAS) connections and form the systems internal storage capacity. The control and expansion controllers support a wide range of drives:

2.5 inch drive support

- 300 GB solid-state drive
- 300 GB 10K RPM SAS disk
- 450 GB 10K RPM SAS disk
- 600 GB 10K RPM SAS disk

3.5 inch drive support

- 300 GB 15K RPM SAS disk
- 450 GB 15K RPM SAS disk
- 600 GB 15K RPM SAS disk
- 1 TB 7,200 RPM near-line SAS disk
- 2 TB 7,200 RPM near-line SAS disk

Drive types can be intermixed between enclosures. The maximum number of internal drives supported by the Storwize V7000 system is 240. The maximum number of raw internal capacity currently supported is 240 TB. Figure 4 shows a Storwize V7000 system with mixed enclosure and drive types.

Figure 4: Storwize V7000 with 10 enclosures

In addition to managing internal storage, one major feature of the Storwize V7000 is the ability to manage storage provided by external storage systems. The Storwize V7000 acts as the virtualization layer between the host and external storage system. Up to 32 PB of external storage can be managed by a single Storwize V7000 system, and as the storage is virtualized, volumes can be nondisruptively moved between external and internal storage capacity. Figure 5 shows a representation of this technology.

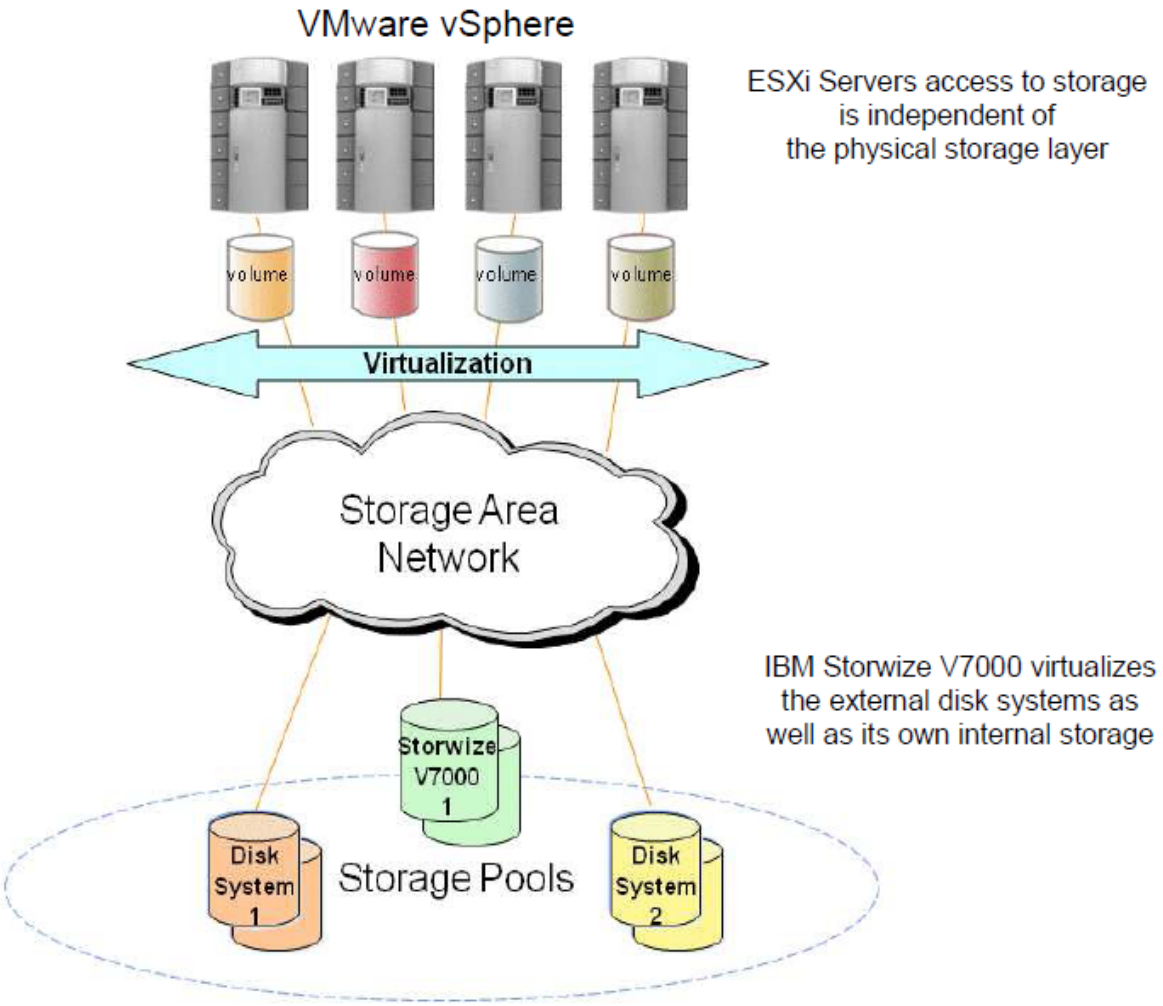

Figure 5: IBM Storwize V7000 virtualizing external and internal storage

It is important to understand the terminology and storage layers involved with storage virtualization in the Storwize V7000 system.

- **Managed disks (MDisks)**: RAID arrays that are presented to the Storwize V7000 from internal and external storage are called MDisks.
- **Storage pool:** The Storwize V7000 takes a single MDisk or aggregates multiple MDisks and creates a storage pool.
- **Volume**: A volume is a storage unit that is presented and accessed by host systems. Volumes are created from the storage capacity present in storage pools.

Figure 6 displays a representation of the storage virtualization layers.

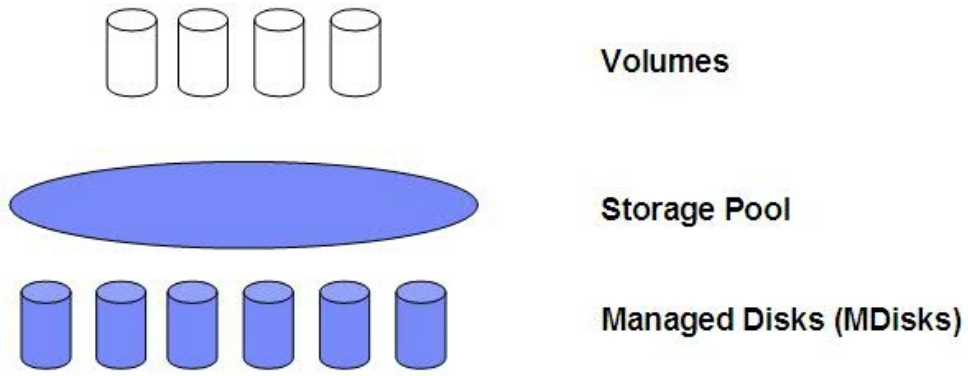

Figure 6: Storage virtualization layers

Solution flexibility

Large IT infrastructure is often complex, costly, and slow to adapt. Changes in technology are cost prohibitive to purchase and implement. Inefficiencies exist which result in loss of productivity and money. However, the virtualization infrastructure provided by a VMware vSphere and IBM Storwize V7000 solution eliminates many of these problems.

Management efficiencies

VMware vCenter server and the IBM Storwize V7000 management GUI provide centralized points of management for the entire virtualized server and storage infrastructures.

VMware vCenter Server

VMware vCenter server is the management foundation of the vSphere platform. vCenter manages all host and virtual machine resources but also features a plug-in architecture for extension of services, such as application monitoring, backups, and capacity planning. Figure 7 displays the vCenter home screen and the administrative options that are available.

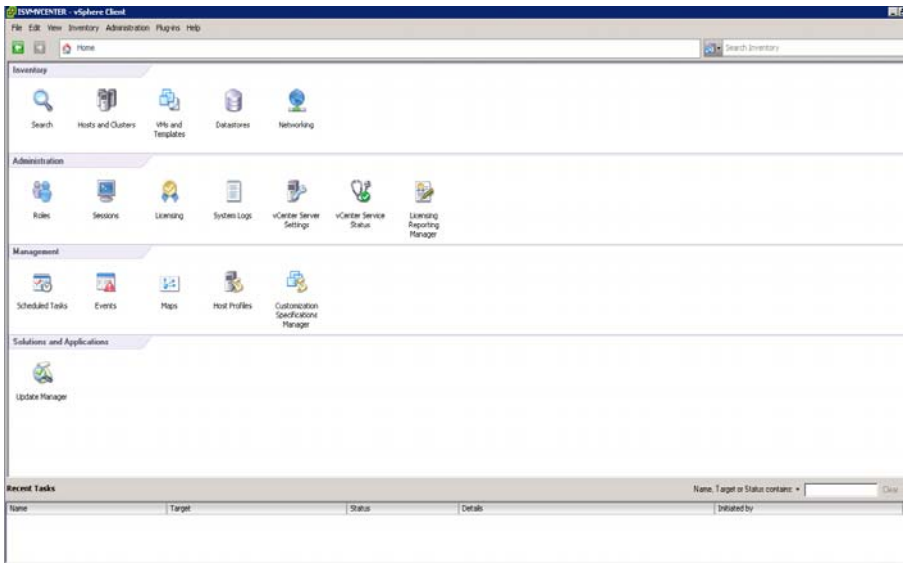

Figure 7: VMware vCenter home screen

The vSphere 4.1 release introduced scalability enhancements to vCenter server to meet the most demanding enterprise workloads. Support for 1,000 hosts and 15,000 virtual machines per vCenter instance allow administrators to manage their entire infrastructure from one interface.

IBM Storwize V7000 management

The IBM Storwize V7000 system includes an embedded web-based management GUI. Initial setup is performed by generating a configuration file with a provided workstation application, copying it to a USB key, and placing the USB key into the Storwize V7000 controller.

Many features in the management GUI for the IBM XIV® Storage System are incorporated into the Storwize V7000 GUI. Management tasks are simplified with wizards and preset configurations. The system includes a **Getting Started** page which contains a guided workflow for configuring the system.

Figure 8 shows the **Getting Started** page and the suggested tasks.

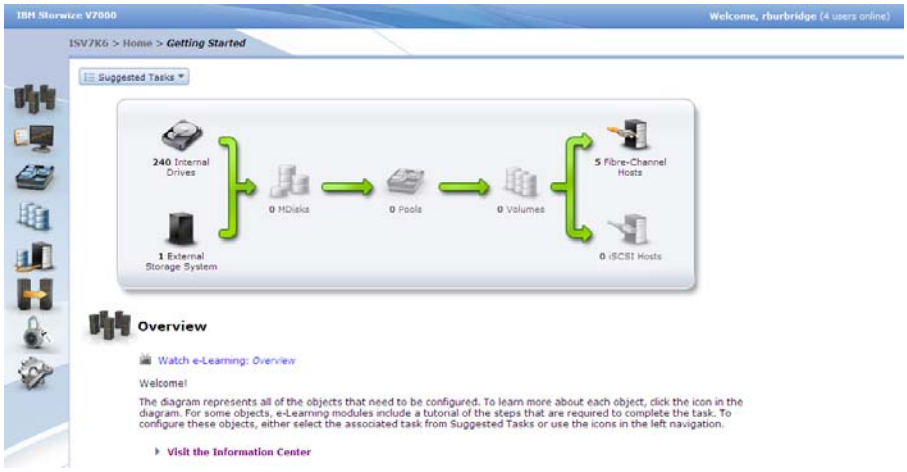

Figure 8: Storwize V7000 Getting Started page

With support to manage 240 TB of internal storage and 32 PB of external storage, Storwize V7000 provides administrators the flexibility to manage their entire storage portfolio as virtualized capacity and from a single management interface.

External storage virtualization

The ability to virtualize external storage systems is a pivotal feature built into Storwize V7000. Much of the Storwize V7000 codebase has been inherited from the IBM SAN Volume Controller, which has been virtualizing storage since the year 2003. The Storwize V7000 benefits from the experience gained over those years.

Many advantages are gained by virtualizing external storage systems. Some of them include: Better performance through the Storwize V7000 caching, centralized management of multiple storage systems, and volume abstraction from the back-end storage. The Storwize V7000 can also migrate nonvirtualized workloads from existing storage systems to virtualized workloads on any storage management by the Storwize V7000 system.

Storage vMotion can nondisruptively migrate virtual machine files between VMFS volumes. A viable option for migrating to new storage within VMware vSphere is to utilize this feature. New storage is provisioned and configured on the vSphere hosts and virtual machines are migrated over. The process can be time consuming as each virtual machine must be migrated, and Storage vMotion is not supported with virtual machines that have physical raw device mappings, as in the case of virtual machines configured in a Microsoft® cluster.

The migration process with the Storwize V7000 system is simple, and although a short downtime window is needed, the entire vSphere environment can be migrated immediately, including clustered virtual machines. The steps for importing an existing storage system that contains a vSphere environment are outlined in the following sections.

Prepare

The Storwize V7000 acts as a host initiator when communicating with external storage systems. The Storwize V7000 has to be zoned to external storage system and an access group has to be set up for

it. Figure 9 shows the host **ISV7K6** setup with no assigned storage. The logical drives used by VMware vSphere are mapped to the **ISV-3650M3-c01** host group.

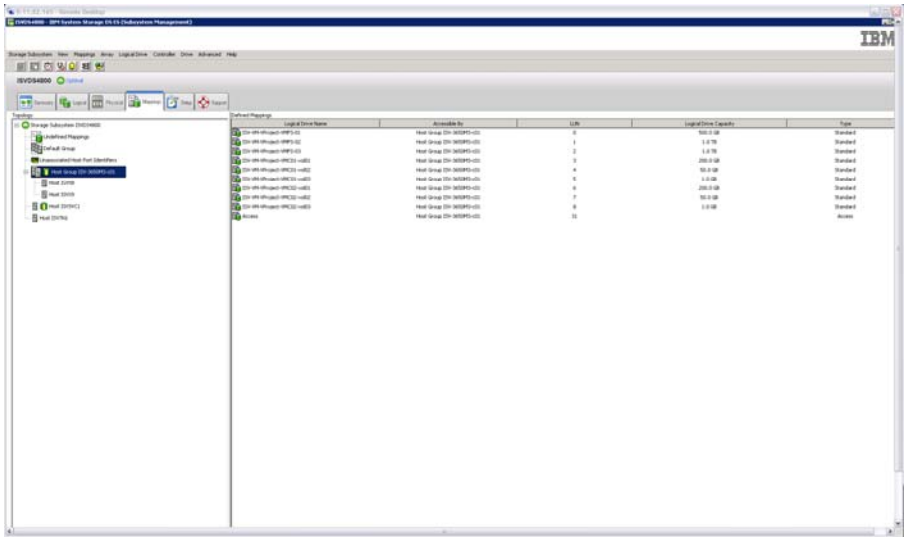

Figure 9: IBM System Storage DS4800 hosts

SAN fabric zones should also be created for the vSphere hosts and Storwize V7000. Single initiator host zones should be used. For more detailed information, refer to the *Deploying IBM Storwize V7000 in VMware Environments* white paper.

Import external storage

When the storage system and zoning is complete, a brief downtime is required for transitioning storage access from the external storage system to the Storwize V7000 system. After taking all of the virtual machines offline, unregister them from the vCenter server inventory. As storage might be briefly unmapped from the vSphere hosts, this ensures that no virtual machines are orphaned.

Soon after the virtual machines are offline and unregistered, it is time to remap the storage within the external storage system. The DS4800 system uses IBM System Storage DS Storage Manager for management. The procedure for remapping the storage on the DS4800 system is to change the logical drive mapping from the **ISV-3650M3-c01** host group to the **ISV7K6** host. With this change, the logical drives are removed from the vSphere hosts and mapped to the Storwize V7000 system.

Figure 10 shows this change being made for the logical drive **ISV-VM-VProject-VMFS-01**.

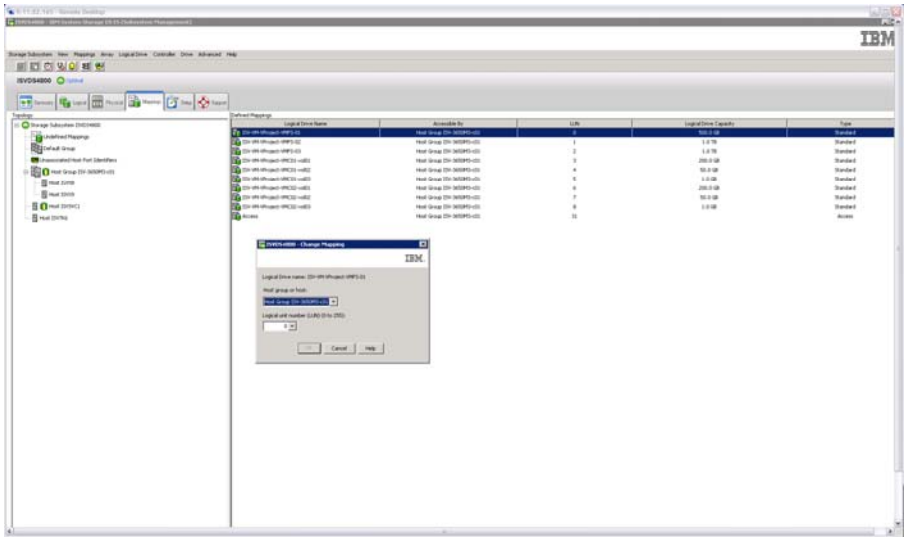

Figure 10: Changing logical drive mapping

This step should be repeated for all logical drives.

After remapping all the logical drives to the Storwize V7000 system; you need to issue a command to the vSphere hosts to rescan all host bus adapters (HBAs). This ensures that the logical drives are removed.

Figure 11 shows that all virtual machines and logical drives have been removed from the vSphere hosts.

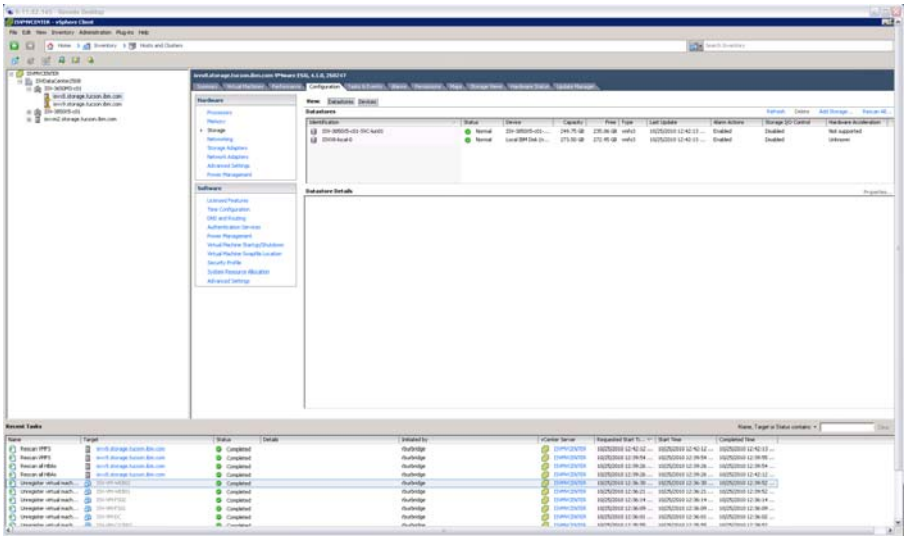

Figure 11: Rescan storage on vSphere hosts

Now that you have removed the storage from the hosts and remapped to the Storwize V7000, you need to detect the storage on the Storwize V7000 and remap to the vSphere hosts.

The **Detect MDisks** button on the external storage view within the Storwize V7000 management GUI is used to detect the storage. Figure 12 shows the discovery process.

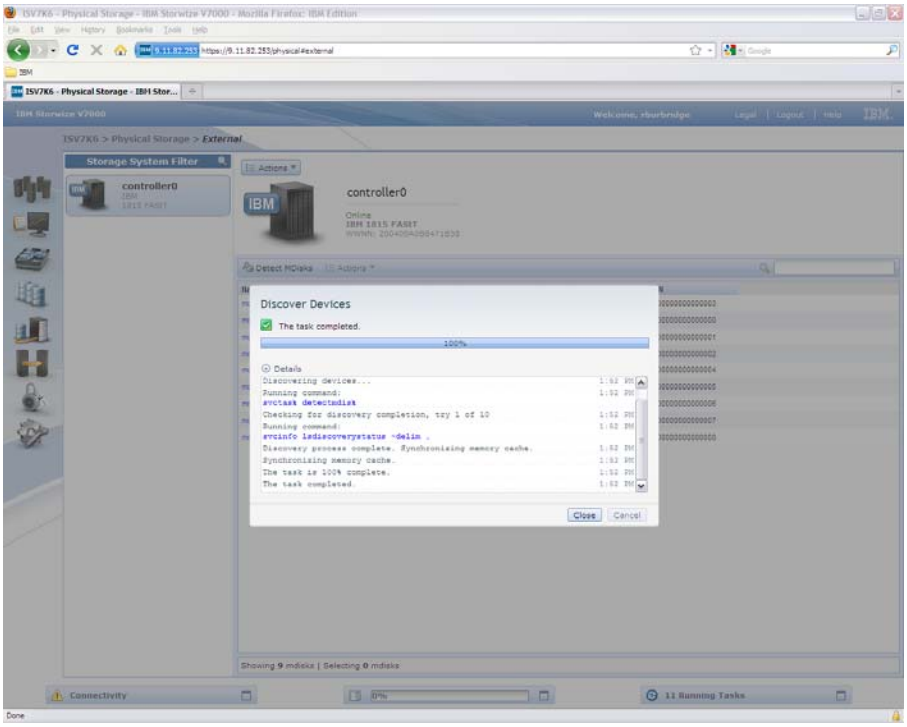

Figure 12: Detecting MDisks

The author found it beneficial to rename the newly discovered MDisks at this point as the storage is being provided by the external storage system. MDisks are renamed to match the logical drive names found in the DS Storage Manager GUI for the DS4800 system. The renamed MDisks are show in Figure 13.

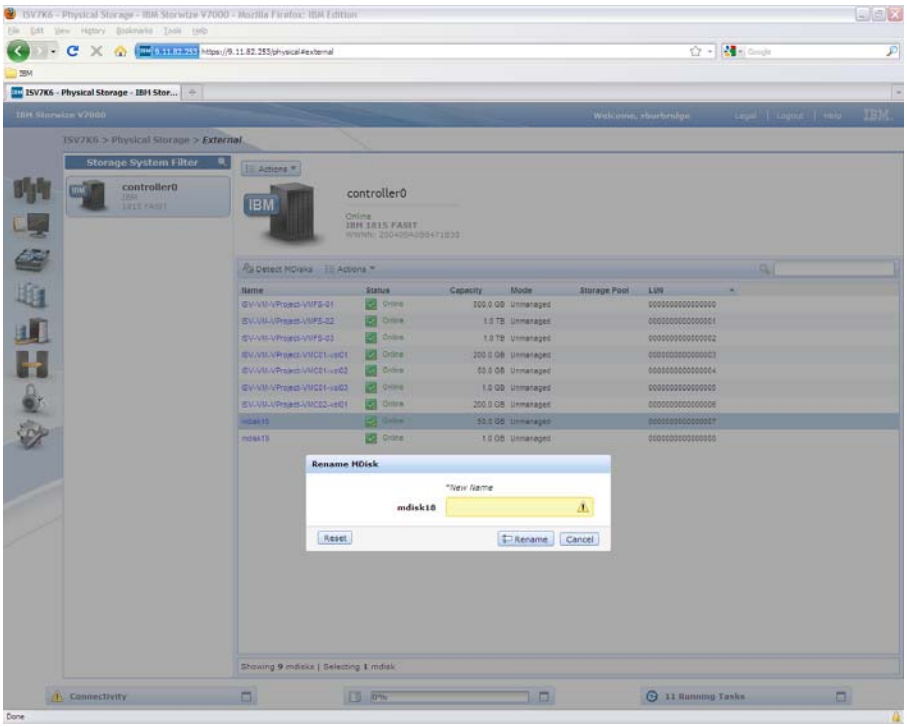

Figure 13: Renaming MDisks

You can access the Import wizard from the Actions menu. This wizard helps you import the data stored on the MDisks to image-mode volumes, which are volumes that map directly to the MDisk. Additionally, the wizard can migrate the MDisks to an existing storage pool. You can generate a new storage pool by not selecting an existing storage pool. Figure 14 and Figure 15 show the import wizard and image-mode volumes.

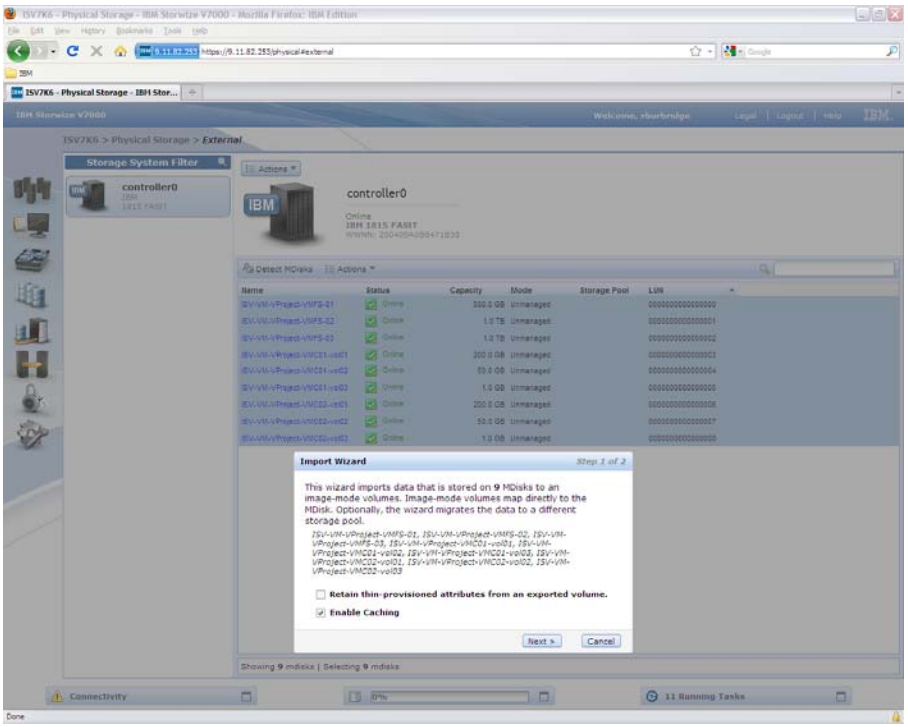

Figure 14: Import Wizard

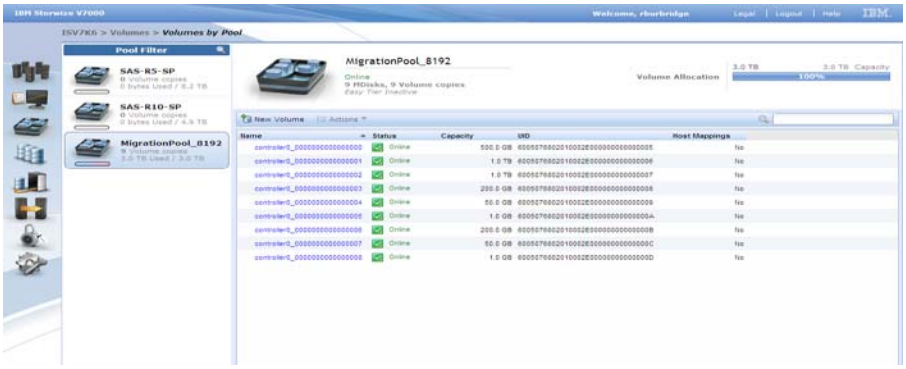

Figure 15: Image-mode volumes in storage pool

You can now map the storage back to the vSphere hosts. Mappings can be easily created and modified with the Storwize V7000. Figure 16 shows the mappings created for each volume and vSphere host. It is also possible to migrate the image-mode volumes to virtualized volumes.

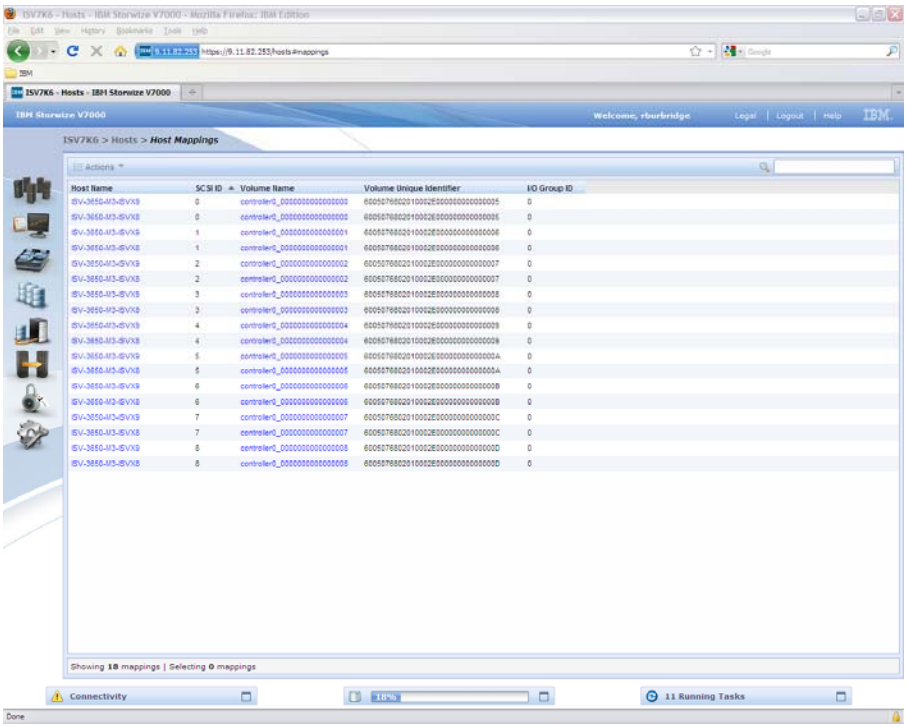

Figure 16: Host mappings

The following steps must be completed within the vCenter server to finalize the storage changes:

1. Run a rescan operation on the vSphere hosts to detect the volumes being presented from the Storwize V7000.
2. Use the **Add Storage** wizard to reattach the existing VMFS volumes.
3. Register the virtual machines to the vCenter server inventory.
4. Switch on the virtual machines and restore services.

The preceding steps outline how to transition an external storage system and vSphere environment to a Storwize V7000 system. After the steps are completed the transitioned workloads gain the benefit of all of the previously mentioned features of the Storwize V7000. The first benefit to examine is thin provisioning.

Thin provisioning

Thin provisioning is the allocation of storage space based on the on the minimum space required at a given time. With conventional storage provisioning, space is allocated for what the requirement might be and not what the current requirement is. Thin provisioning results in a reduction of storage space usage compared to conventional provisioning.

Thin provisioning with the Storwize V7000

The Storwize V7000 system includes thin provisioning of volumes for no extra charge and is supported for all host and external storage platforms that the Storwize V7000 system supports. New

volumes can be created as thin-provisioned and existing volumes can be migrated to thin-provisioned volumes. Thin-provisioned volumes on the Storwize V7000 system have two main attributes:

- **Real capacity**: Real capacity is the amount of storage space actually being used by a volume.
    - **Virtual capacity**: Virtual capacity is the amount of space allocated to the volume, and is also the amount of space available with the host.

The image-mode volumes from the imported vSphere environment were converted to generic or thick volumes, meaning storage consumption is equal to the allocation, and the **ISV-3650M3-VMFS-01** volume shows this in Figure 17.

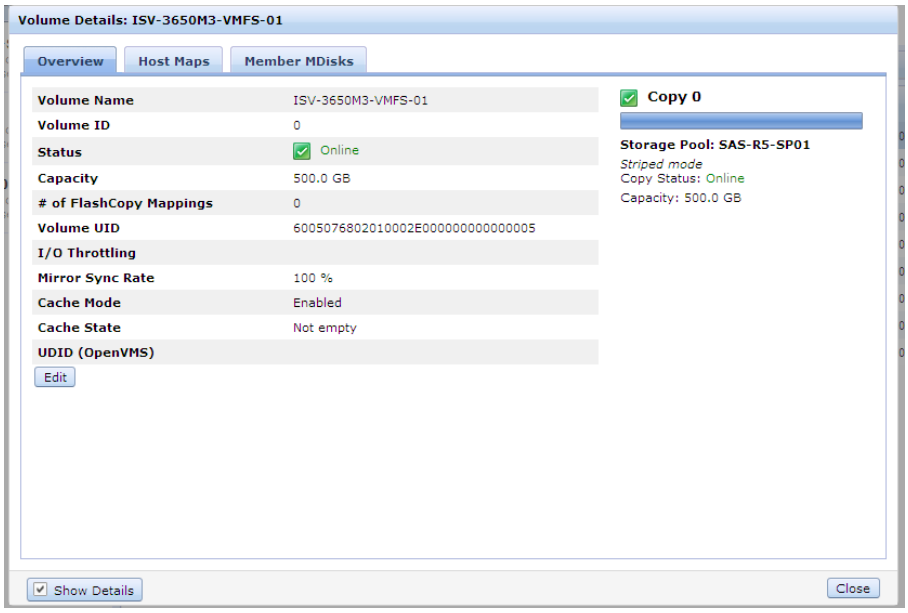

Figure 17: Volume details

The Storwize V7000 can easily convert the volume to be thin provisioned by adding a mirrored copy. A mirrored copy is a secondary copy of the data that makes up a volume. The host system interacts with only one volume; however the Storwize V7000 system maintains two copies of the data. In this example, the primary copy resides on the **SAS-R5-SP01** storage pool and is a thick volume. Figure 18 and Figure 19 show the steps needed to add a mirrored copy, which is thin provisioned.

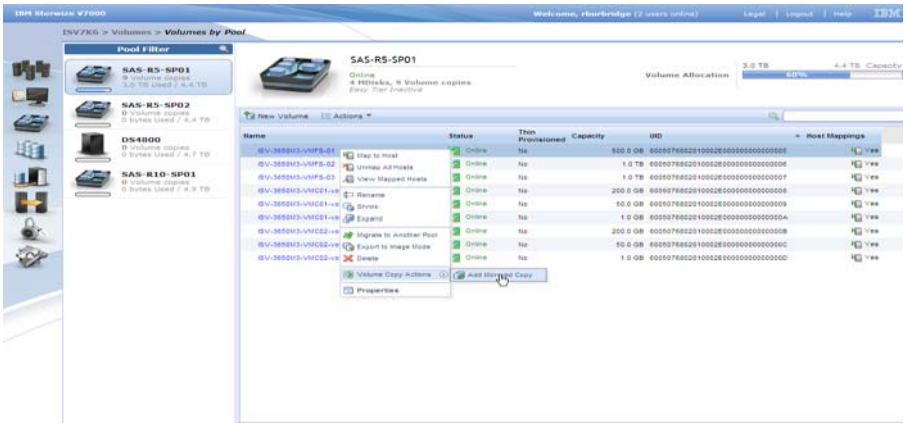

Figure 18: Adding mirrored copy – step 1

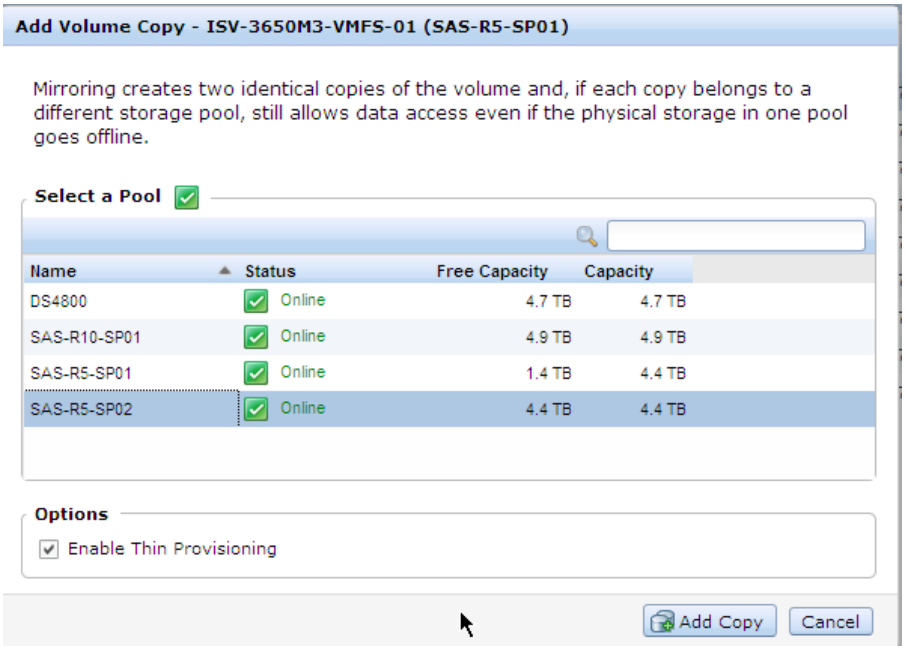

Figure 19: Adding mirrored copy – step 2

After the synchronization of data from the primary to the secondary volume copies is completed, you can find the real capacity of the thin-provisioned copy, as shown in Figure 20.

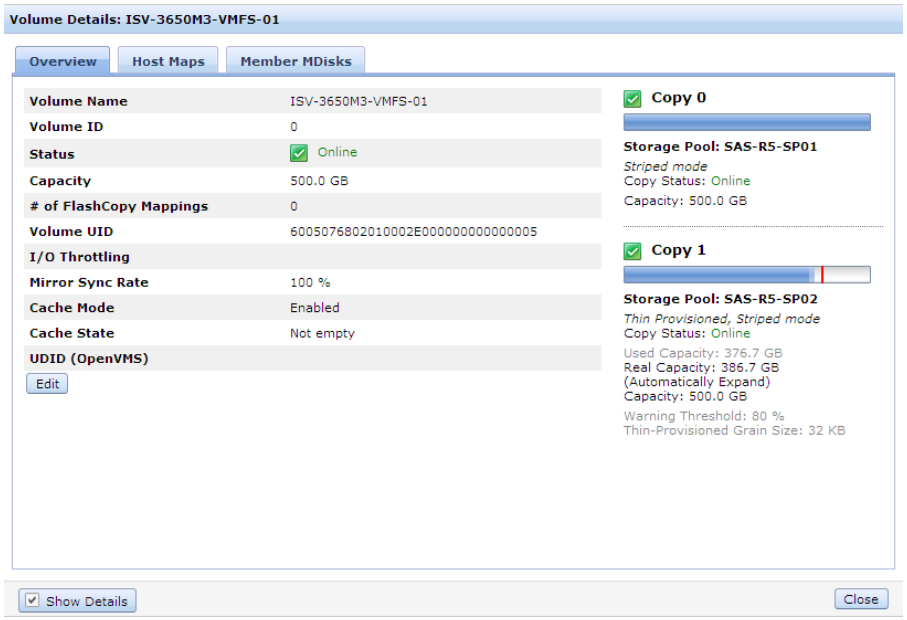

Figure 20: Volume properties

Thin provisioning with VMware vSphere

Thin provisioning in VMware vSphere is simple to implement and the benefits are realized almost immediately. Virtual machine disks can be created as thin-provisioned or existing virtual machines can be converted to thin-provisioned disks during a storage vMotion operation. Virtual machines have two main attributes that relate to storage provisioning:

- Provisioned space: Provisioned space is the amount of storage space that is allocated or provisioned for the potential use of a virtual machine.
- Used space: Used space is the amount of storage space that the virtual machine is actually using.

Figure 21 shows the VMFS usage after virtual machines were converted to thin-provisioned disks.

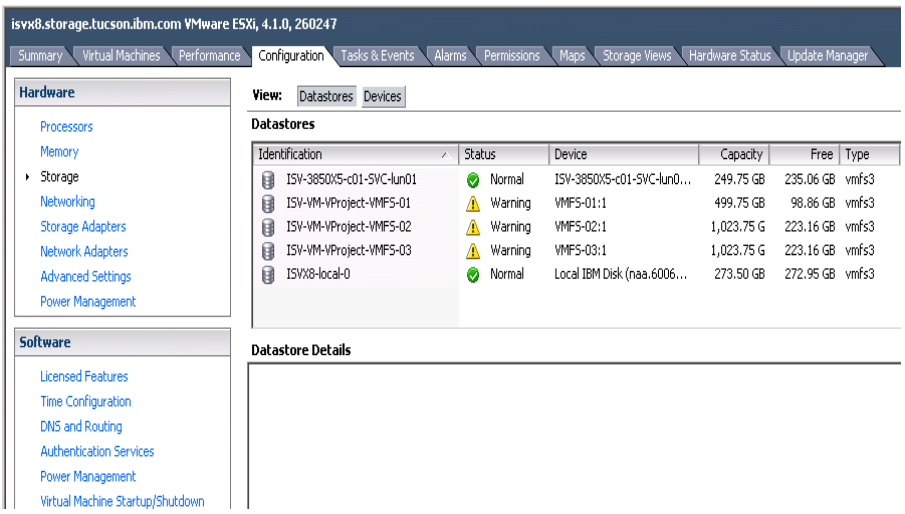

Figure 21: VMFS usage

VMFS volumes can be over-provisioned, which is when the sum of all virtual machine provisioned space exceeds the size of the VMFS volume. Whenever storage space is over-provisioned, you need to create and follow monitoring and remediation plans. vCenter includes the ability to set alarms that trigger when certain values are met.

Thin on thin

The last option for thin provisioning with the Storwize V7000 and VMware vSphere is to place thinly-provisioned virtual machines on thinly-provisioned VMFS volumes. This configuration provides a very high-storage efficiency as only the space used by virtual machines is consumed within the VMFS volume and the Storwize V7000 volume. The risk is that storage might become over-allocated and impact the systems.

A best practice when running in a thin on thin configuration is to only over-provision at one management level. Do not let virtual machine provisioned space exceed the capacity of a VMFS volume. This ensures that if all virtual machines consumed their allocated capacity, the VMFS volume might not become full and impact services. Real capacity might need to be closely monitored on the Storwize V7000 system to ensure that real capacity of volumes do not exceed the capacity of a storage pool.

As long as proper monitoring and alerting is setup, the benefits of thin provisioned can be realized with minimal risk.

Solution responsiveness

The demands of applications drive the IT infrastructure businesses must implement to meet the demands.. Applications are generally dynamic: they require more capacity or more performance, and the infrastructure need to respond accordingly. VMware vSphere and Storwize V7000 create an infrastructure solution that can quickly respond to changes in application demands.

Workload migration

Storage virtualization has many advantages over traditional storage management. The Storwize V7000 system has the ability to migrate volumes between storage pools, which allow the infrastructure to respond to changes in workload demands.

Volume migration to meet performance demands

A change in performance requirements is common with many application workloads. Increases in processor and memory capacity can be easily introduced with VMware vSphere. Hot-add even allows for those changes to be made without system downtime. However, it is much harder to implement changes in storage infrastructure that meet the changes in performance demands. New storage needs to be provisioned and workloads need to be taken offline to transition to the new storage.

With the Storwize V7000 system, volumes are virtualized and they can be nondisruptively moved between storage pools. Volumes inherit the performance characteristics of the storage pool where they reside. For example, a storage pool with a single MDisk consisting of eight disk spindles might not perform as well as a storage pool with four MDisks consisting of 64 disk spindles.

The graph in Figure 22 shows the effect of a manual migration on a workload. In this example, a virtual machine running a simulated Microsoft SQL workload profile in Iometer (SQL1) is impacted by the introduction of another workload on the same storage pool (SQL2).

- At 5:30 PM, the SQL1 workload starts and the response time is 10.5 ms/op.
- At 6:30 PM, the SQL2 workload starts and the response time increases to 20 ms/op.
- At 8:00 PM, a volume migration is initiated for SQL1. During the migration, the duration of the response time decreases.
- At 10:30 PM, the volume migration completes and the response time is 8 ms/op.

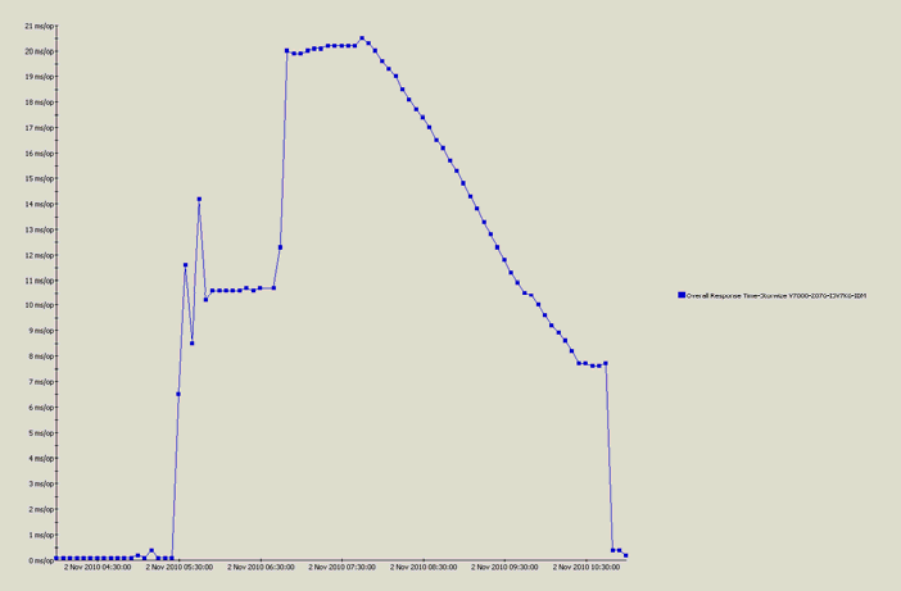

Figure 22: The effect of manual migration on response time

Automated performance tuning with Easy Tier

The Storwize V7000 has a classification of a storage pool called a hybrid pool. A storage pool is considered as a hybrid pool when it contains a mixture of solid-state drives and standard spinning disks. The main advantage of a hybrid storage pool is that IBM Easy Tier can be enabled.

Easy Tier enables effective use of solid-state drive storage by monitoring the I/O characteristics of a virtualized volume or storage pool and migrating the frequently-accessed portions of data to the higher-performing solid-state drives. Easy Tier reduces the overall cost and investment of solid-state drives by effectively managing their usage.

Workloads that experience changes in performance requirements benefit from Easy Tier without any management interaction. Figure 23 shows the impact that Easy Tier can have on a workload. The previously mentioned SQL1 and SQL2 Iometer workloads were allowed to run with their volumes in a hybrid storage pool. As data was migrated to the solid-state drives, response time decreased.

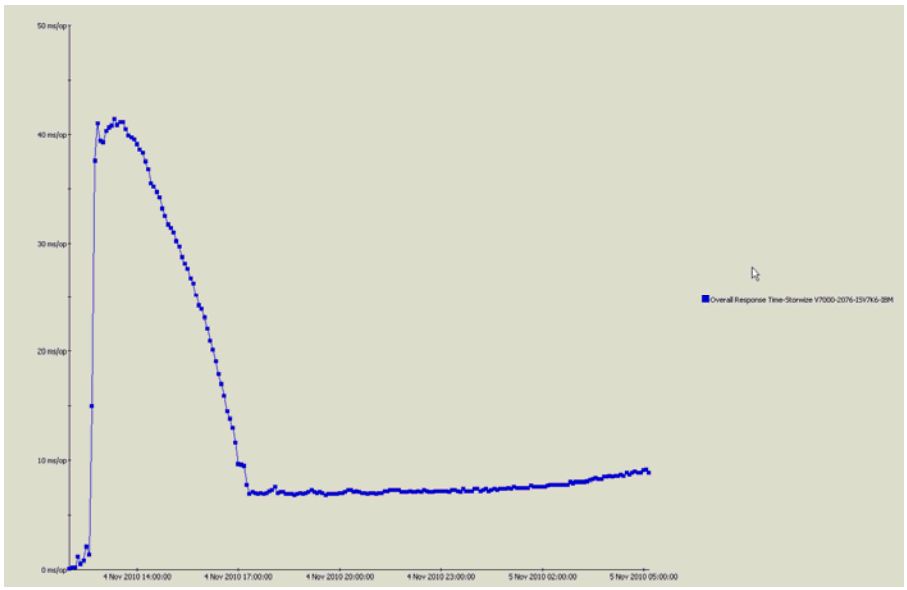

Figure 23: Impact of Easy Tier on workload

Responding to storage space demands

Another type of workload demand that VMware vSphere and Storwize V7000 are able to react to quickly is changes in storage space. Data requirements are ever increasing, and therefore, a solution that can quickly respond to those demands is valuable to customers.

Both, VMware vSphere and Storwize V7000 support nondisruptive expansion of storage. There are three locations that can and may need to be expanded when more storage space is needed.

- Storwize V7000 volume: Expanding the Storwize V7000 volume might impact a VMFS volume, or in the case of a raw device mapping, might impact a logical drive on a virtual machine. Figure 24 shows the Storwize V7000 Expand Volume screen.
- VMware VMFS volume: You can expand a VMFS volume if there is free space available on the volume. You can also add an extent, which is spanning together multiple volumes. Figure 25 shows the window used when expanding a VMFS volume.
- Virtual machine disks: A .VMDK file can be expanded if there is free space available on a VMFS volume. Figure 26 shows the window used when expanding a .VMDK file.

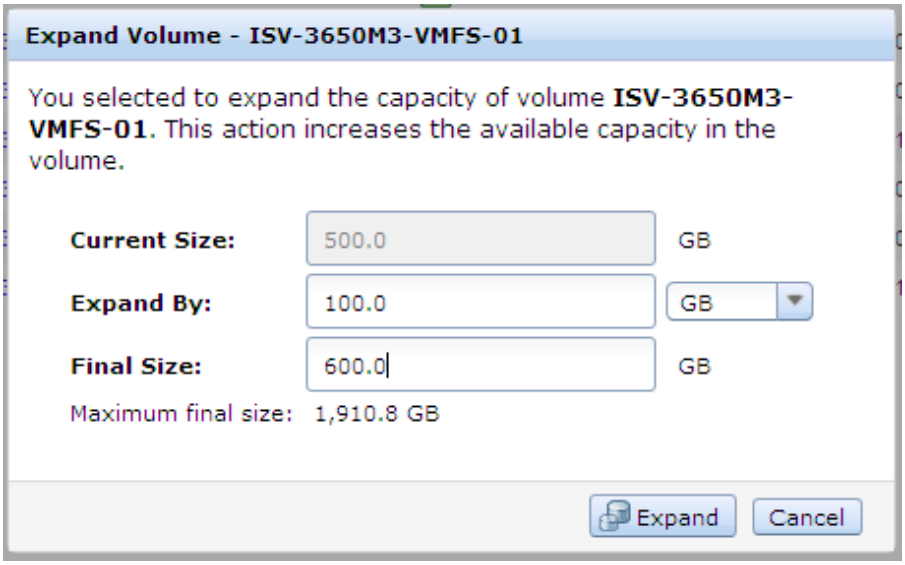

Figure 24: Expanding Storwize V7000 volume

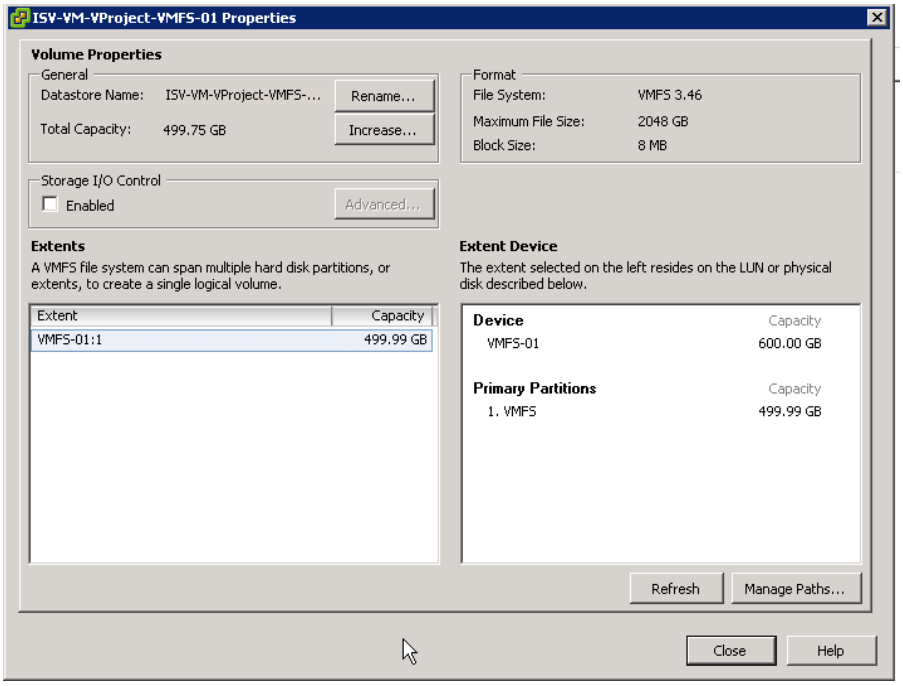

Figure 25: Expanding VMFS volume

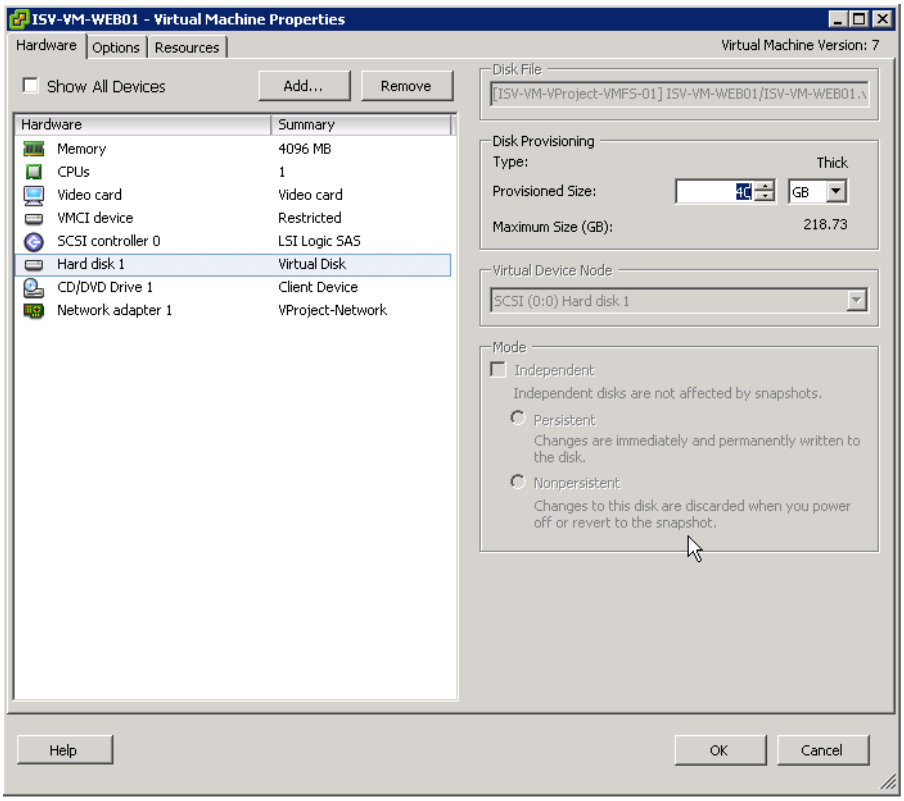

Figure 26: Expanding a VMDK file

In the example discussed in this section, the logical volume within the virtual machine operating system must also be expanded before the free space can be utilized. In Microsoft Windows® operating systems, this change is made in Disk Manager. Figure 27 shows the Disk Manager screen.

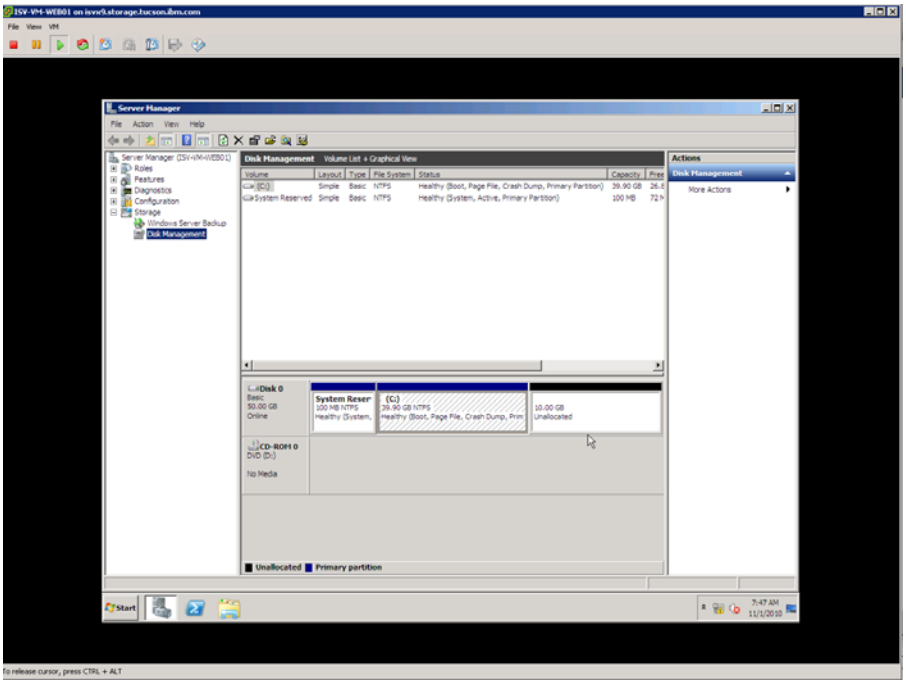

Figure 27: Disk Manager

FlashCopy services

The FlashCopy function enables point-in-time full volume copies of data, with the copies immediately available for read or write access. The FlashCopy function creates a copy of a source volume on a target volume; and this relationship is called a FlashCopy mapping. Mappings can exist from a source volume to multiple target volumes.

FlashCopy has many use cases, one of which enables the VMware vSphere and Storwize V7000 solution to be responsive to changes in business demands.

Cloning virtualized environments with FlashCopy

The ability to clone virtual machines is one major advantage that virtual machines have over physical servers. New systems can be quickly provisioned, or existing systems can be cloned for testing. VMware vSphere has cloning features built into the product which enables cloning of virtual machines. However, by using the IBM FlashCopy function built into Storwize V7000, entire virtual environments can be quickly cloned.

Creating FlashCopy mappings is an easy process that you can achieve with a few mouse clicks in the Storwize V7000 GUI. Figure 28 shows the FlashCopy menu and options available.

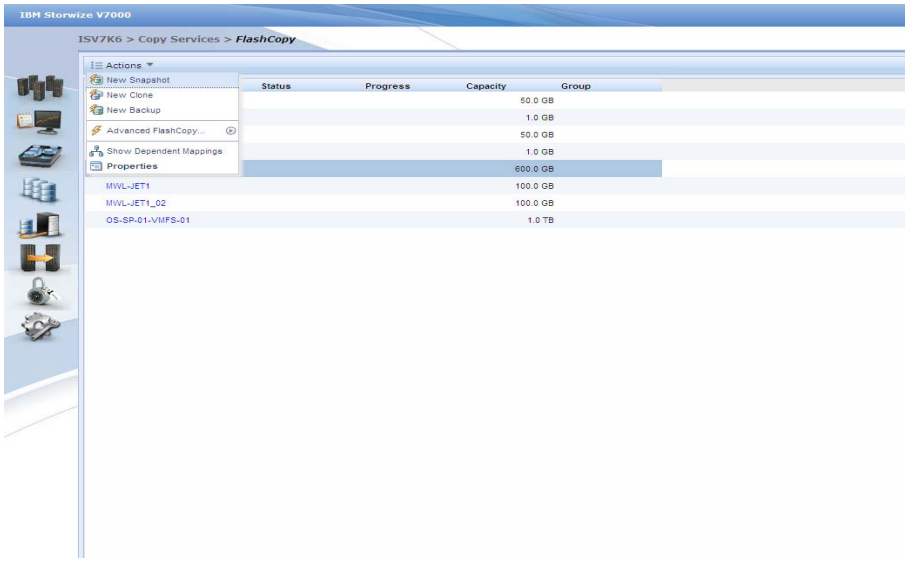

Figure 28: FlashCopy menu

If the intention is to create application-consistent clones of the virtual machines, you need to create virtual machine snapshots prior to creating the IBM FlashCopy of the VMFS volume. After creating IBM FlashCopy, you need to map the target volume to a vSphere host, and after doing this, the copies are immediately available for read and write access.

After mapping the copy to the vSphere host, you need to add it as an existing VMFS volume. As it is a copy of an existing VMFS volume, you need to mount it with a new VMFS signature. that the option to perform this is shown in Figure 29.

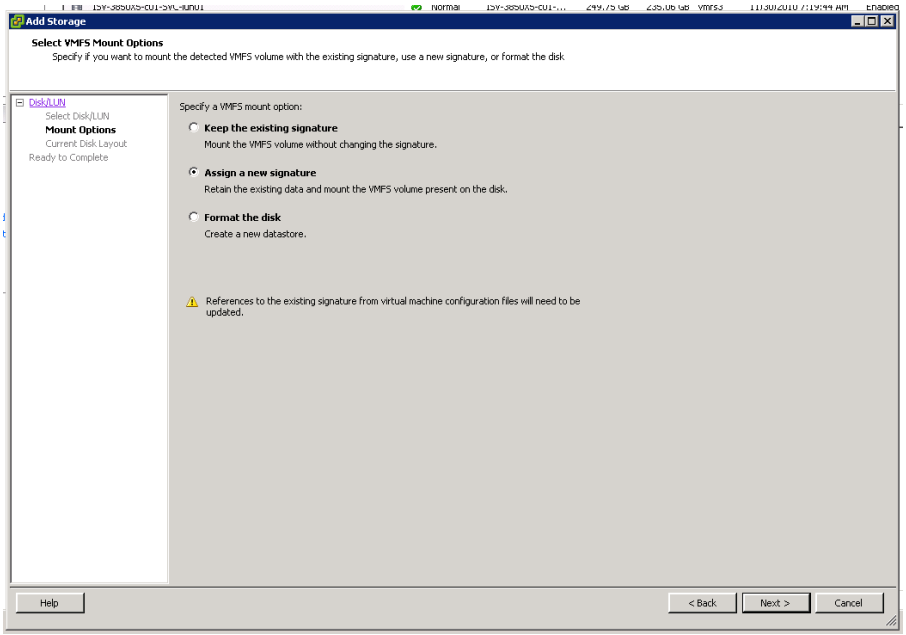

Figure 29: VMFS mount options

After mounting the copy as a new VMFS volume, register the virtual machines in vCenter and switch them on. Note that you need to rename the cloned virtual machines when you add them to the vCenter inventory so that they do not conflict with the original virtual machines. Figure 30 shows a virtual machine from the cloned VMFS volume being registered to vCenter.

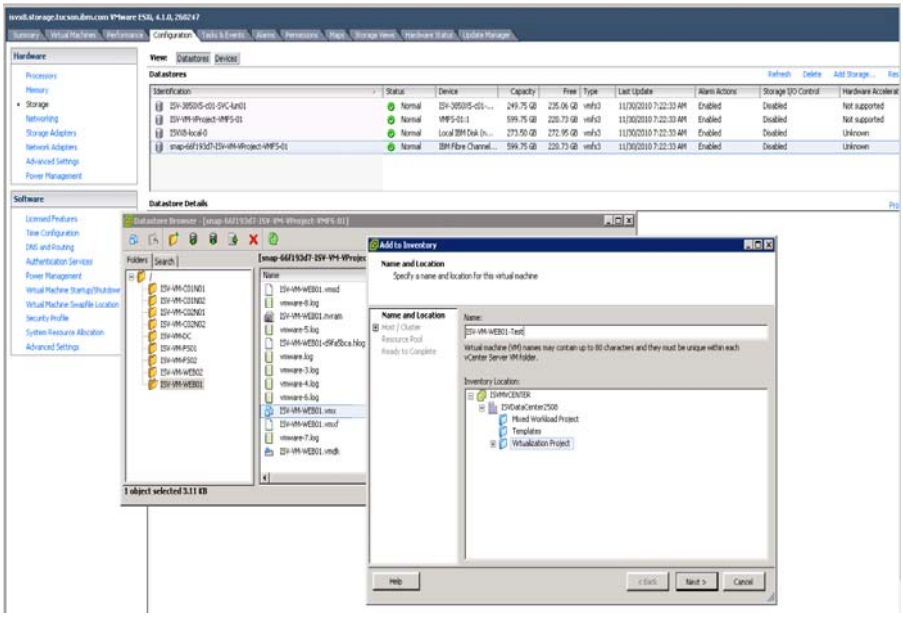

Figure 30: Registering virtual machine clones

Virtual machines on a cloned VMFS volume are identical to the original systems. This functionality allows the entire environment to be quickly cloned and used for testing changes or patches. However, you need to, take precautions to ensure that cloned systems are properly isolated from the original systems. A test virtual network with no access outside the host is recommended to ensure isolation is maintained.

Solution availability

The solution provided by VMware vSphere and Storwize V7000 provides the availability to meet nearly any application business requirement. Features such as volume mirroring, Metro Mirror, and Global Mirror, VMware Fault Tolerance, and VMware Site Recovery Manager can reduce instances of unplanned downtime and bring services back online during disaster recovery scenarios.

Volume mirroring and VMware Fault Tolerance

Volume mirroring is a feature of the Storwize V7000 system which allows for a single volume to have data-consistent primary and secondary copies on storage managed by the Storwize V7000 system. The primary and secondary copies of data can be placed in separate storage pools, and if available on separate external storage systems. By doing so, you can protect a volume from a storage system failure.

VMware Fault Tolerance is a permanent feature available in VMware vSphere that operates at the virtual machine layer, allowing the feature to be operating system and application agnostic. Management of VMware Fault Tolerance is very simple, as it is just a check box that you need to enable, and failover is automatic. VMware Fault Tolerance works by creating a secondary virtual machine, identical to the

primary virtual machine, on a separate VMware vSphere host. The virtual machines operate in a virtual lockstep, ensuring that failover can be instant with no downtime.

VMware Fault Tolerance protects against vSphere host failures. However, as the primary and secondary virtual machines both reference the same VMDK files, there is no protection against a storage failure. Storwize V7000 volume mirroring can compliment VMware Fault Tolerance by providing high availability of the VMFS volumes that Fault Tolerance-enabled virtual machines are using.

VMware Fault Tolerance is enabled through the vSphere client, as shown in Figure 31. After enabling the feature, the virtual machine icon takes on a blue hue and an additional Fault Tolerance window displays statistics on the Fault Tolerance relationship, as shown in Figure 32.

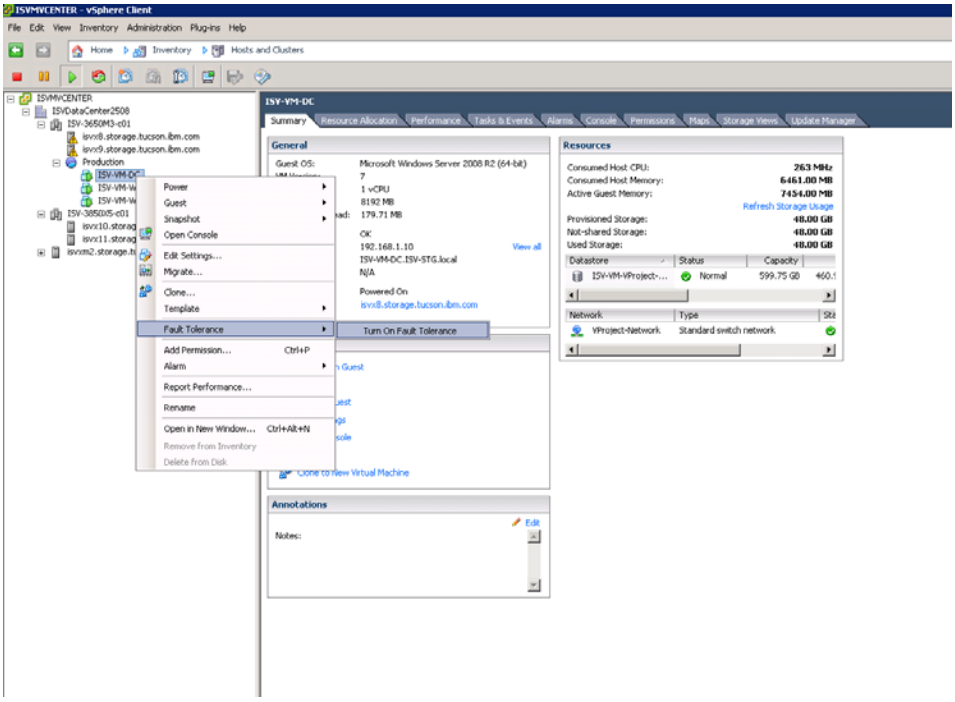

Figure 31: Enabling VMware Fault Tolerance

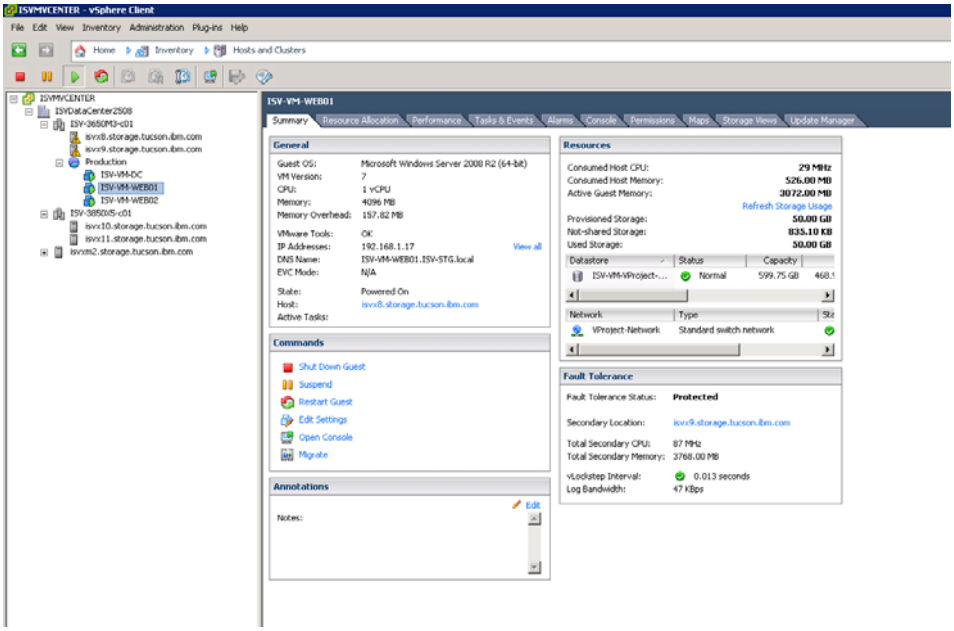

Figure 32: VMware Fault Tolerance statistics

The steps for creating a volume mirror were outlined in the “Thin provisioning with the Storwize V7000”.

IBM Storwize V7000 Metro Mirror and Global Mirror with VMware Site Recovery Manager

Metro Mirror and Global Mirror are copy services built into the Storwize V7000 system and they mirror data between volumes on different Storwize V7000 systems and at varying distances.

Metro Mirror provides a synchronous-copy process between a primary and secondary volume, across storage pools, and Storwize V7000 systems. Typically, host applications write data to the primary volume and the updates are copied to the secondary volume. As Metro Mirror is a synchronous process, it requires adequate bandwidth and low latency.

Global Mirror provides an asynchronous-copy process between a primary and secondary volume, across storage pools, and Storwize V7000 systems. Global Mirror works over greater distance and in scenarios where bandwidth might be very limited.

VMware Site Recovery Manager is a business continuity and disaster recovery solution that helps to plan, test, and run a scheduled migration or failover of data center services from one site to another. Site Recovery Manager is an extension of VMware vCenter that enables integration with the Storwize V7000 system and management of replication services.

Another paper, *Implementing disaster recovery solutions with IBM Storwize V7000 and VMware Site Recovery Manager*, provides more details on the solution that Storwize V7000 and VMware Site Recover Manager creates, as well as details on how to configure and manage the solution.

The link to the paper is: **ibm.com**/support/techdocs/atsmastr.nsf/WebIndex/WP101765

Summary

A fully-virtualized solution offers significant advantages for companies that are looking at virtualization to increase flexibility and responsiveness of IT infrastructure, and provides the following advantages:

- Migrating existing storage and workloads to a virtualized storage environment
- Thin provisioning of data for efficient use of storage space
- Responding to changing business demands through automated and manual data migration
- Mirroring and replication of data for increased availability

This paper explained the VMware vSphere and Storwize V7000 solutions and also the features that they include. When combined, the two products can create an environment where server and storage resources are fully virtualized.

Appendix A: Materials List

Storage	
IBM Storwize V7000	10 drive trays, 216 300 GB 10 K SAS HDD, 24 solid-state drive
Servers	
2 IBM System x3650 M3	Processor type: Intel Xeon 5670 (6 cores per socket @2.93 GHz) Total processor cores: 24 Total RAM: 288 GB DDR-2
Host bus adapters	
QLogic 8 Gb Fibre Channel (FC) Dual-port HBA for IBM System X	
Fibre Channel switches	
2 IBM 2498-B24, each with 24 8 Gbps FC ports	
Virtualization software and guest operating systems	
Server virtualization	VMware vSphere ESXi 4.1 Installable
Guest OS	Windows Server 2008 R2 Enterprise Edition

Appendix B: Resources

These websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
  ibm.com/partnerworld/systems

- IBM Publications Center
  www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

- IBM Storwize V7000
  ibm.com/storage/storwizev7000

- IBM Redbooks®
  ibm.com/redbooks

- IBM System Storage Interoperation Center (SSIC)
  ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes

- VMware
  http://www.vmware.com
  - VMware vSphere 4.1 Documentation
    http://www.vmware.com/support/pubs/vs_pages/vsp_pubs_esx41_vc41.html
  - Hardware Compatibility Guide -
    http://www.vmware.com/resources/compatibility

  - VMware Fiber Channel SAN Configuration Guide-
    http://www.vmware.com/pdf/vsphere4/r41/vsp_41_san_cfg.pdf

- IBM TechDocs Library
  ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

DIA Data Integrity Assurance®, Storewiz®, Storwize®, and the Storwize® logo are trademarks or registered trademarks of Storwize, Inc., an IBM Company.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.