



xiotech
White Paper

Best Practices for Deploying Highly Resilient, Cloud-Ready, Scale-Out Storage Infrastructures Using VMware vSphere 4 and Xiotech ISE

Xiotech's Emprise 5000 and VMware vSphere 4 bring together the concepts of consolidation, high availability and fault tolerance for both storage and servers.

This document contains virtualization concepts and scalable configurations using the Emprise 5000 and VMware vSphere.

It is intended for IT Professionals managing virtual servers and storage.

Chad Bittner
Technical Marketing
Xiotech Corporation

Rev 1.0

Table of Contents

- Introduction..... 2**
- Xiotech Emprise 5000 Overview.....3**
 - Lowest Capital and Operating Costs.....4*
- VMware vSphere™ 4 Overview.....6**
- Why Xiotech Emprise 5000 and VMware vSphere 4?..... 7**
- Deploying a Virtualized Environment.....7**
- Configuring the Emprise 5000.....8**
- Configuring the VMware vSphere 4 Cluster.....11**
- VMware vSphere 4 Backup and Replication Methodologies.....15**
 - Backup Methodologies.....15*
 - Replication Methodologies.....16*
- Conclusion..... 17**

Introduction

In today's world, every company is attempting to do more while minimizing investments in new technology. This is known as the "do more with less" principal. Virtualization is the vehicle that many companies have adopted as the leading strategy for achieving this goal. Virtualization has many forms, but the primary forms used by companies are Server Virtualization and Storage Virtualization.

Server Virtualization can be simply explained as running multiple Guest Oses or Virtual Machines on a single physical host, allowing for more effective use of physical CPU and memory resources. In reality, server virtualization is a very complex balancing act. All server virtualization products have varying ways to help make this balancing act automated and transparent so that administrators can easily implement the server virtualization software.

Storage Virtualization can be simply explained as a pool of storage that is provisioned in many ways depending on the application. Like server virtualization, storage virtualization allows administrators to purchase storage in large blocks and then provision out just what is needed for a specification application. This allows for centralization of the storage and easier management, as opposed to the traditional individual server based storage.

VMware vSphere 4 and the Emprise 5000 are leading technologies in the virtualization market and one of the easiest to implement. As virtualization technologies continue to penetrate the market, choosing the right solution will be an integral part of the IT budget.

Xiotech Emprise 5000 Overview

It is very important to understand the storage involvement in a virtualized infrastructure. Regardless if you are talking about server, switch or storage virtualization, the storage is the foundation and core of it all.

Xiotech's Emprise 5000 system is a revolutionary concept in data storage. It is built on patented Intelligent Storage Element® (ISE) technology and provides a perfectly balanced building block of virtualized storage with performance, reliability, and scalability that delivers the lowest total cost of ownership of any SAN or SAN-like storage system on the market.

Key Emprise 5000 characteristics:

- **Flexible storage building block** – attach one directly to a server or many via a switch
- **More IOPS**, important for database and Web serving applications, key for virtualization
- **Greater throughput** for applications like data warehousing and decision support
- **Dense capacity** – use 30% less of your data center for the same amount of storage
- **Linear scaling** of both performance and capacity – multiply your IOPS, MBps, Gbps, and TB simply by plugging additional systems into a switch
- **Better scaling for all applications** at all times with less trouble
- **More than 100x the reliability** of traditional disk drives – never replace a drive again
- **Minimal attention needed from experts** and tuning from administrators
- **Lowest capital and operating costs in the industry**
- **Reduced power and cooling costs** through automated shut-down and power-up capabilities of PowerNAP™

Back to Basics: Don't Pay for What You Don't Need

Typical SAN systems come with high-level functionality you may not need, and all that functionality adds to the price tag. It also consumes performance and introduces a lot of complexity, which often requires high-priced administrators.

Emprise 5000, on the other hand, is a flexible, foundational building block of storage. It is a complete, self-enclosed virtualized storage solution, which you can configure to meet your unique needs or optimize for a specific application. It provides superior capabilities in the dimensions that matter most: performance, reliability, cost, and ease of use.

Start Small, Grow Big: You can start with a single Emprise 5000 system direct-attached to a server. It is easy to deploy and requires little ongoing administration. As your needs grow, simply add a switch and plug in more Emprise 5000 systems. You can grow your storage network virtually without limits, without ever losing your investment.

Use It Right Where You Need It: Emprise 5000 offers an easy-to-deploy, cost-effective solution for a variety of applications.

Get Everything You Pay For

Typical SAN performance is limited by bottlenecks in the system—from SAN controllers that slow data movement to back-end switches that add another point of contact. Scalability also is restricted by how many drives you can effectively put behind a set of controllers.

Emprise 5000 changes this paradigm. It eliminates the storage controllers on the front end and switches on the back, so data is free to flow at screaming speeds. And because storage is decoupled from the servers, you can add servers solely for processing power—and gain the full capabilities of each.

Keep Your Data Safe

With typical SANs, disk drive failure is inevitable, expensive, and puts your data at risk. When a disk fails, a human must pull the failed drive and insert a replacement—of the right kind. This takes time and effort, and mistakes can be devastating. What's more, data must be rebuilt on the new drive. This process can take hours or days, introduces further risk to your data, and can cause system performance to plummet.

Emprise 5000 with its ISE technology provides reliability that other storage systems can't touch—extending reliability by orders of magnitude beyond that of a single drive.

Never Touch Another Disk: Instead of individual disk drives, Emprise 5000 utilizes sealed *DataPacs*. Preventive diagnostics keep the disk drives within the DataPacs running longer, and self-healing capabilities repair those that do fail—in-place, without human intervention, and without impacting normal operation or performance.

Lowest Capital and Operating Costs

Switch-connected Emprise 5000 systems **cost less to buy and much less to maintain** than any other SAN-like solution.

Better for Your Data Center: Emprise 5000 uses 30% less space than traditional storage devices to provide the same capacity, significantly reducing your square footage costs. Emprise 5000 also reduces power and cooling costs.

No Bottlenecks: Storage controllers, back-end switches, and high-end features of typical SAN systems not only introduce bottlenecks, they increase up-front costs while adding complexity—which often leads to expensive vendor services.

All the Performance, Not All the Disk: Many storage vendors must load their systems with a large number of disk drives in order to come close to the performance of a single Emprise 5000 system. This drives up the costs of their systems, while introducing more opportunities for disk-related failures. Emprise 5000 delivers the performance you need without the extra disk.

VMware vSphere™ 4 Overview

Leverage the power of virtualization to transform your IT infrastructure into private clouds built on VMware vSphere—enabling IT to increase control through service-level automation. VMware vSphere, the industry’s most reliable virtualization platform, dramatically reduces capital and operating costs and maximizes IT efficiency—with the freedom to choose any application, OS, or hardware.

At a high level, VMware vSphere 4 vCompute includes:

VMware vCenter Server: Provides a central point of control for virtualization management, essential for administering infrastructure and application services, with deep visibility into every aspect of virtual infrastructure, automation of day to day operational tasks and the scalability to manage large datacenter environments.

VMware ESX and VMware ESXi Server: Provide a robust, production-proven, high performance virtualization layer that abstract server hardware resources and allow their sharing by multiple virtual machines. Unique memory management and advanced scheduling capabilities of VMware ESX and ESXi allow for the highest consolidation ratios and the best application performance, in many cases, even better than physical servers.

VMware DRS: Aggregates compute resources across many clusters and dynamically allocates them to virtual machines based on business priorities, reducing management complexity through automation.

VMware DPM: Included with VMware DRS, automates energy efficiency in DRS clusters by continuously optimizing the power consumed in a DRS cluster.



Fig. 1. VMware vSphere 4 (VMware)

Why Xiotech Emprise 5000 and VMware vSphere 4?

Xiotech and VMware together provide easy to use, cost effective, storage and server virtualization solutions for business's today. By focusing on automation, ease of use, high availability, and improved infrastructure scalability, our customers realize greater flexibility in building a dynamic technology infrastructure. The Emprise 5000's innovative approach to storage compliments VMware vSphere's approach to virtualization. The advanced capabilities of the ISE related to performance, reliability and modularity are an architectural match for vSphere and its advanced resource management.

Xiotech is a Select Member of the VMware Technology Alliance Program and long time partner of VMware. By combining our technologies, customers have been able to deploy a virtualized solution that not only encompasses virtualized servers but storage as well, while minimizing their total cost of ownership (TCO) and maximizing their return on investment (ROI).

Deploying a Virtualized Environment

The configuration represented in this document is two or more vSphere servers connected to one or more Emprise 5000's via an FC switch. The functionality and configuration of the Emprise 5000 is identical in both single and multiple ESX/ESXi servers, but the vSphere capabilities between the two are dramatically different. In order to use the advance vSphere functionality of HA, DRS (VMotion), and fault tolerance a minimum of a two node cluster with a vCenter server and proper licensing is required. See figure 2.




Fig 2 Two Node VMware vSphere Cluster

Configuring the Emprise 5000

This section describes configuring storage and connecting it to servers using the Emprise 5000 Web based interface. To access the ISE interface, browse via a web browser to one of the IP Addresses assigned during the installation process. Once the required fibre channel zoning is in place, the creation of storage LUNs and the presentation to the vSphere servers is a very simple process requiring very little time and expertise.

First, create the shared storage for the VMware vSphere HA/DRS cluster. In order for virtual machines to be protected by VMware HA or load balanced using VMotion and VMware DRS, they must be placed on a shared VDisk.

In order to create the VDisk the follow items need to be entered.

Name – this is the friendly name for identification in the Emprise 5000. In this example, we called the VDisk “V1”

Cache Policy – Select Write-Back as the default cache policy

Capacity/Redundancy – Select the Pool, Raid type and enter the capacity of the VDisk to be created for the Datastore. RAID 1 is the default and recommended selection.

The screenshot shows the 'Storage Configuration' interface with a 'Create a new volume' section. The 'Name' field contains 'V1'. The 'Cache Policy' dropdown is set to 'Write-Back'. Below, there are two storage pools: 'Pool 1: Performance' and 'Pool 2: Performance'. 'Pool 1' shows 'RAID 1: 336 GB' selected. A 'Capacity in Gigabytes' field contains '300' with a range of '1 GB - 2047 GB'. 'Pool 2' shows 'RAID 1: 36 GB' and 'RAID 5: 57 GB' options.

Fig. 3. Creating VDIs for Booting VMware ESX Server Hosts

Figure 3 shows the creation of a VDisk to be used as a datastore within the ESX environment.

Xiotech recommends that you create datastores that are 300GB to 1TB in size for the storage of Virtual Machines. Raw Device Maps (RDM) can also be used. However, they should be used only when necessary for isolation of data or when creating a Microsoft Cluster with a physical and virtual node. RDMs and VMFS Datastores are almost equal in performance except when under very high I/O loads.

Note: When creating Datastores on the Emprise 5000 they should be created in a manner that load balances the DataPacs on the Emprise 5000. For example, when creating two datastores, they should be created so that one datastore exists on Pool1 and the other exists on Pool 2.

Next, create the host entries for each ESX Server in the cluster. When creating a host entry, enter the name of the server. This is simply a friendly name that will be used in the interface to identify the server. Next select the Operating system type of Linux as this is the closest match to the ESX server operating system. The final entry is the WorldWide Names (WWN) of the Fibre channel ports in the ESX server. Select all of the WWNs that are associated with the server and press the 'Add' to configure them for that host. Then press the 'Submit' button at the bottom of the page.

Detailed Host Information

E5K-ESX4i-1 Information

Name:

Operating System Type:

Select 1 or more available HBA Ports to be contained in this Host.

Available HBA Ports:

Selected HBA Ports:
2100001B3206B2CA
2101001B3226B2CA

Fig. 4. Host entry creation and WWN assignment

Figure 4 shows the first server in the ESX cluster. This physical host has a dual port QLogic QLE2462 FC HBA installed and zoned to see the Emprise 5000.

Note: When zoning the hosts to the Emprise 5000, a zone should be created for each FC port in the server and zoned to a single controller on the Emprise 5000. This allows for multipathing and link resiliency in the event that a cable, FC port, FC switch, switch port, or controller goes offline. Reference Xiotech's support matrix for further information on zoning requirements. <http://xiotech.com/Support-Center.aspx>

Lastly, the VDisks need to be assigned to the hosts. This step allows the servers to access the storage and is configured under the VDisk. Simply select the servers you want to access the VDisk and click the 'Add' button. Once the servers are in the 'Selected Hosts' box, simply press the 'Save Changes' button at the bottom of the page.

Detailed Volume Information

V1 Information

Name: V1

Status: **Operational**

Detailed Status: [None]

Type: Primary Volume

Created: Mon Nov 30 06:34:41 2009

Cache Policy: Write-Back

Pool Identifier: Pool 1

Redundancy Level: RAID 5

Capacity (GB): 300

Global ID (GUID): 6000C508805F7000000E000200000000

Local ID (LDID): 16

Snapshots:

Select 0 or more available Hosts for volume mapping.

Available Hosts:

Selected Hosts: EDIT

E5K-ESX4i-1 : lun(0)

E5K-ESX4i-2 : lun(0)

Add >>

<< Remove

Fig. 5. Assigned VDisks to ESX hosts

Figure 5 shows the first VDisk being assigned to the two servers in the ESX cluster.

Configuring the VMware vSphere 4 Cluster

Next configure the shared storage for the templates and virtual machines.

Using the Virtual Infrastructure client, log into the VMware vCenter server and create a Datacenter by right-clicking on Hosts & Clusters and selecting New Datacenter.

Right-click on the new Datacenter and New Cluster. Select both **VMware HA** and **VMware DRS**. Now add the VMware ESX hosts to the cluster.

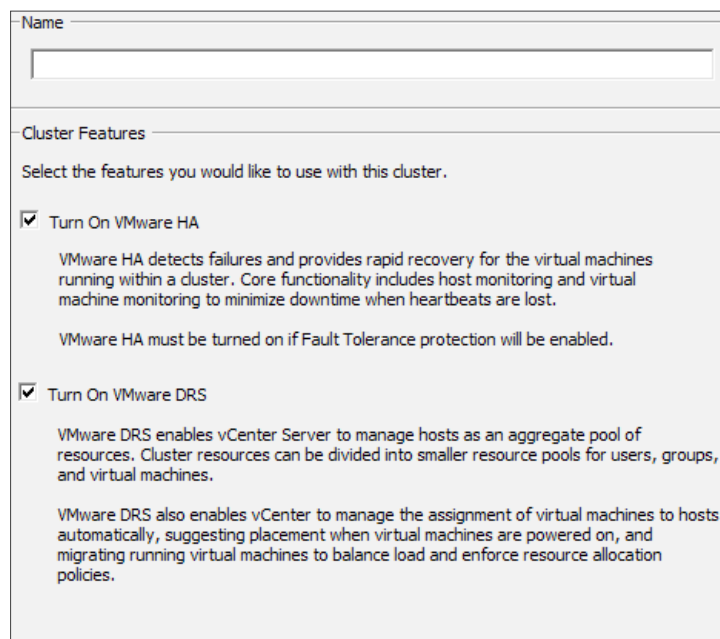


Fig. 6. Create Cluster

Next configure each host to use the shared storage.

First, select an ESX Server host and click the **Configuration** tab.

Select **Storage Adapters**.

Verify that the host can see the storage.

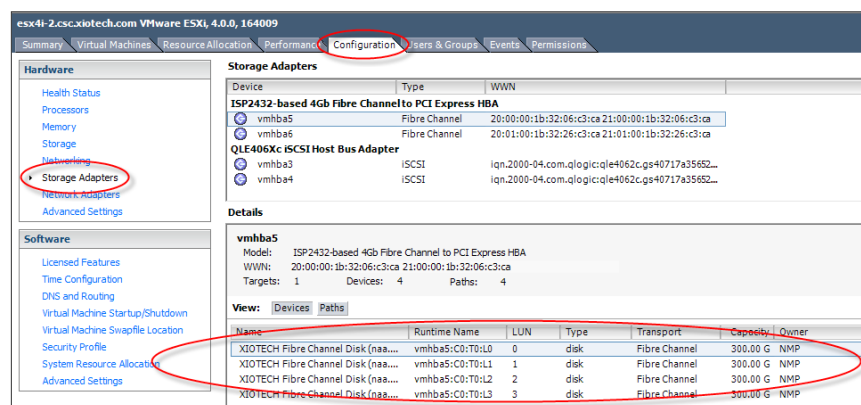


Fig. 7. Configure Hosts to See Storage

Next step is to create a datastore, which is storage that has been configured for VMs.

From the **Configuration** tab, select **Storage (SCSI, SAN, and NFS)** to display the available Datastores.

Select **Add Storage**.




Fig. 8. View Available Datastores

Select the type of storage you are adding to the ESX Server host. In this case the default of **Disk/LUN** is correct.

Click **Next**.

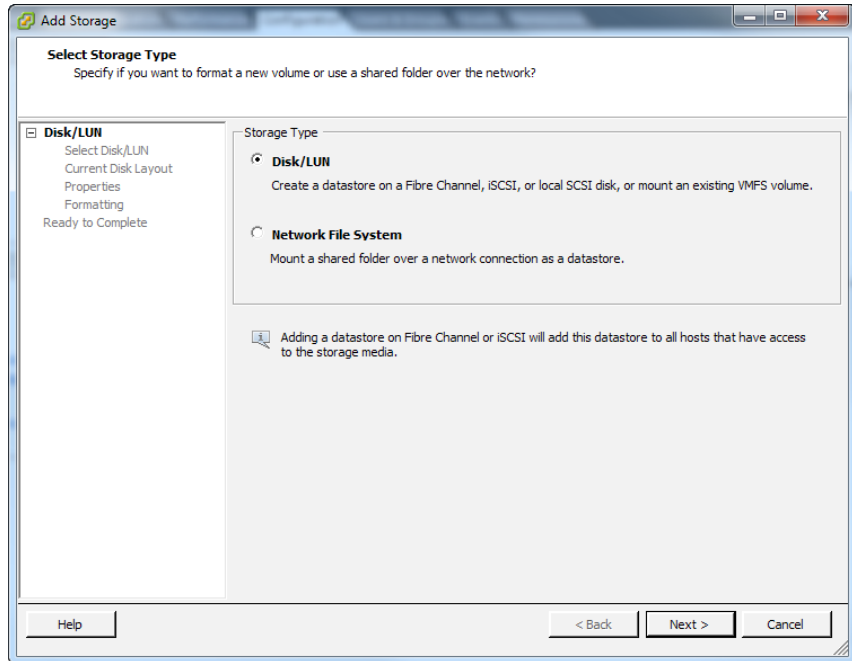


Fig. 9. Select the Storage Type to Add

From the list of available Disks/LUNs, select the LUN you wish to configure.

Click **Next**.




Fig. 10. Select the LUN to Configure

Verify the current disk layout of the selected LUN is correct.

Click **Next**.




Fig. 11. Verify Layout of the LUN

Enter a name for the datastore. This will be used to identify the datastore when you create the VMs.

In Fig. 12, the datastore name is **SharedVM**.




Fig. 12. Name the Datastore

Format the disk/LUN using the VMFS-3 file system.

VMware's VMFS-3 file system is designed to accommodate VMware's use of very small files for configuration and large files for virtual disks.

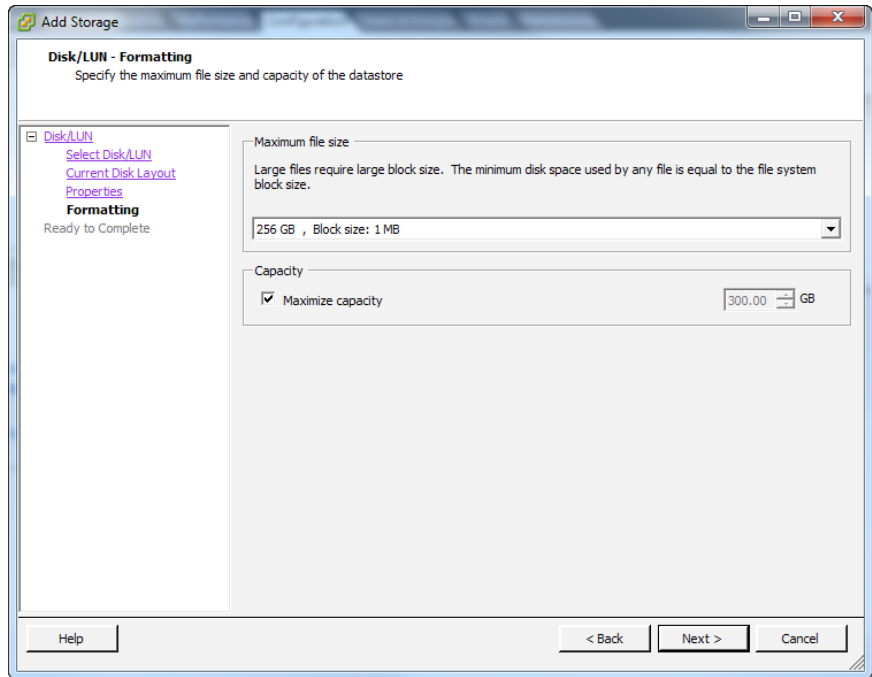


Fig. 13. Format the Disk/LUN

Datastores are now available for you to use when creating VMs.

Reference VMware's [Guest Operating System Installation Guide](#) for information on creating VMs.

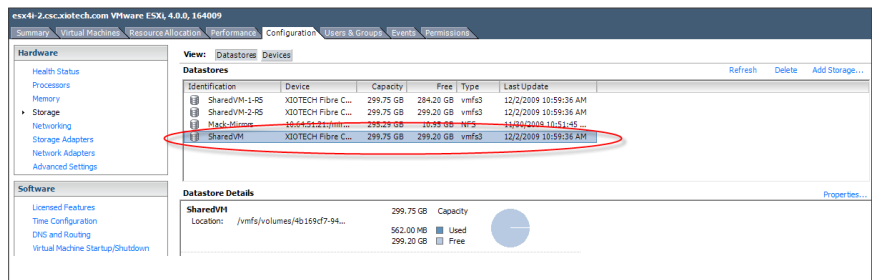


Fig. 14. View Available Datastores

VMware vSphere 4 Backup and Replication Methodologies

Data protection is not an option, it is a business requirement. Loss of data or uptime can equal loss of revenue. Organizations depend on data protection methods for daily recovery operations; therefore, protecting vSphere from potential outages is a high priority. However, when an outage does occur, it is the time it takes to recover the data and the recovery point that matter most.

A backup and replication strategy is measured by two indicators: recovery point objective (RPO) and recovery time objective (RTO). RPO is the point in time (prior to the outage) to which systems and data will be restored. RTO is the period of time after an outage at which the systems and data will be available again. The lower the RPO and RTO, the better the reliability.

Recovering from an outage with a traditional backup/restore is typically time and resource intensive, with data only as current as the last backup. If there is an outage and the last backup was performed the night before, the recovery point is the previous day. All the data at the time of the backup is recoverable. However, if the outage takes place late in the day following the last backup, nearly an entire day of changed data is gone... lost forever. There is also the possibility that the restore may not work at all, so having an effective methodology for data recovery is very important and should not be taken lightly.

Backup Methodologies

VMware has a number of methods for the backup and recovery of VMs in the VMware ESX Server environment.

- **Traditional Backup:** Deployment of a backup client to every VM that requires backup services.
- **VMware Data Recovery:** Provides a centralized management interface to enable backup and recovery of your VMs directly through VMware vCenter Server.
 - The entire inventory of VMs are automatically identified through VMware vCenter Server
 - Uses an intuitive wizard driven workflow to create, configure and schedule backup jobs
 - Multiple restore points for each VM are displayed to easily select a specific point in time copy to restore
 - Automatically monitors VMs that are moved by VMware HA, VMware VMotion and VMware DRS, so scheduled backups continue uninterrupted

- **VMware Consolidated Backup:** provides an easy to use, centralized backup facility for VMware vSphere. Backing up your virtual machine contents from a centralized Microsoft Windows 2003 proxy server rather than directly from VMware ESX lets you:
 - **Reduce the load** on your VMware ESX server during backups
 - **Improve the manageability** of your IT resources by using a single agent running on the proxy server rather than having an agent on every virtual machine
 - **Eliminate backup traffic** on the local area network by utilizing Fibre Channel tape devices for virtual machine backups

Replication Methodologies

VMware supports two types of replication, Array/Hardware based and software based replication.

- **Array/Hardware Replication:** Provides replication using either the native replication in a storage array or an appliance that is in-line between the physical servers and the storage array. This method of replication allows for a per Datastore or RDM level of replication control.
 - VMware Site Recovery requires array based replication in order to provide the automated fail over and Disaster Recovery plan testing.
- **Software Replication:** Provides replication using either agents installed in the Guest OS of a virtual machine or an application that interfaces directly with the VMware environment. This method of replication allows for a per virtual machine level of replication control. Here are some examples of software replication vendors:
 - **Vizioncore** vReplicator is a virtual machine replication solution that enables disaster recovery and business continuity by helping organizations meet recovery time objectives (RTO) and recovery point objectives (RPO) while offering a low total cost of ownership.
 - **Double-Take®** Availability provides affordable data protection, ensures minimal data loss and enables immediate recovery from any system outage. Double-Take Availability continuously captures changes as they happen and replicates those changes to one or more servers at any location, locally or globally.
 - **Veeam** Backup and Replication offers a major breakthrough in full and incremental backup speeds. Using technological advancements within VMware vSphere 4 and the VMware vStorage APIs for data protection, Veeam customers experienced the following benefits derived from native support of thin provisioned disks and changed block tracking.

Conclusion

Virtualizing an IT infrastructure can be a bit overwhelming with all the media exposure it is receiving today. A way to minimize this is to speak with virtualization professionals that have years of industry experience and hands-on knowledge with virtualized products. Please contact Xiotech with any questions you may have on this solution and they will be happy to discuss with you how the Emprise 5000 and VMware vSphere 4 can help meet your business needs.