



OPENBENCH LABS

Data Center SAN Infrastructure

Analysis: Use iSCSI SAN Software to Virtualize and Consolidate Systems and Storage



Analysis: Use iSCSI SAN Software to Virtualize and Consolidate Systems and Storage

Author: Jack Fegreus, Ph.D.
Managing Director
openBench Labs
<http://www.openBench.com>
July 12, 2009

Jack Fegreus is Managing Director of openBench Labs, consults through Ridgetop Research, and is on the Board of Advisors of IT Brand Pulse. He currently contributes to InfoStor and Virtualization Strategy Magazine. He has served as Editor in Chief of Open Magazine, Data Storage, BackOffice CTO, Client/Server Today, and Digital Review. Previously Jack served as a consultant to Demax Software, CTO of Strategic Communications, and IT Director at Riley Stoker Corp. Jack holds a Ph.D. in mathematics and worked on the application of computers to symbolic logic.

Table of Contents

Executive Summary	04
VMware iSCSI SAN Scenario	06
Snapshots & Replication in a VOE	09
Desktop Storage Consolidation	13
Customer Value	15

Executive Summary

“With StarWind Enterprise Server, any 32-bit or 64-bit Windows® Server can be used to provide iSCSI target devices capable of sharing access with multiple hosts—important for Windows Server clusters and VMware environments—and having no limits on the storage capacity exported to clients or on the number of client connections.”

Whether at a small- to medium-size business (SMB) or a global enterprise, IT decision makers universally agonize over servers and networks. The root cause of most slowdowns and interruptions to computer operations, however, is found within the site's data storage configuration. That's one of the reasons driving CIOs to focus on infrastructure virtualization as a way to lower IT operating expenditure (OpEx) costs.

openBench Labs Test Briefing: StarWind Enterprise Server iSCSI SAN Software

- 1) **StarWind consolidates resources and virtualizes storage at the disk-block level:** Image-based logical devices serve logical blocks from any underlying physical storage and can be easily moved, cloned, or copied to another storage device.
- 2) **StarWind management console provides a unified single-pane-of-glass management for all StarWind servers:** With storage resource management being IT's biggest cost driver, StarWind provides the means to manage heterogeneous storage resources centrally.
- 3) **Full support for VMware ESX server including VMotion and VMware Consolidated Backup:** IT is no longer constrained to use identical hardware and can mix and match servers and storage at their disposal to assemble a highly-available storage infrastructure that covers disaster recovery needs.
- 4) **Image-based virtual disks pools support thin provisioning:** StarWind automates thin provisioning by consuming virtual disk blocks, which are assigned to a logical disk in the file header, only when data is written to the disk rather than when the volume is assigned or formatted.
- 5) **Iometer benchmark pegs sequential throughput at near-wire speed while supporting high IOPS levels:** Sequential reads and writes reached 80MB and 65MB per second respectively, while a stream of random 8KB reads (75%) and writes (25%) on sustained 3,250 IOPS.

Resource virtualization eliminates the constraints that physical limitations put on devices. By eliminating device-specific restrictions common to storage hardware packaging, IT administrators can group storage resources into generic pools, which are characterized strictly by device functionality. What's more, administrators can manage all of the members in a pool with one set of tools, which simplifies the task of providing a network of servers with highly available and extremely flexible storage.

As a software-only product for iSCSI storage virtualization, StarWind Enterprise Server provides IT with a highly cost effective means to leverage storage

virtualization and help deliver critical data management services for business continuity. With StarWind Enterprise Server, any 32-bit or 64-bit Windows® Server can be used to provide iSCSI target devices capable of sharing access with multiple hosts—very important for Windows Server clusters and VMware

environments. Most importantly, StarWind differs from many competitive offerings by including two nodes with a StarWind Enterprise Server license to support mirroring and replication. StarWind also places no limits on the storage capacity exported to clients or on the number of client connections.

Through the creation of snapshots that are compatible with Microsoft's Volume Shadow Copy Services (VSS) with unlimited rollback points along with asynchronous and synchronous replication and mirroring, StarWind extends the benefits of storage virtualization well beyond higher productivity for IT administrators and higher utilization of storage devices. CIOs can address key business application issues, such as data availability for service level agreements (SLAs), via the storage virtualization provided by a StarWind-based iSCSI SAN.

The simplicity of providing virtualization with a software iSCSI SAN solution allows StarWind to fit a wide range of IT needs. For SMB sites, where direct attached storage (DAS) is dominant, IT can leverage StarWind Enterprise Server with existing Ethernet infrastructure to adopt a SAN paradigm with minimal capital expenditure (CapEx) and OpEx costs. At large enterprise sites, IT can immediately garner a positive return on investment (ROI) by leveraging StarWind Enterprise Server to extend storage consolidation to desktop systems for dramatically higher storage utilization, on-demand storage provisioning, higher data availability, booting over the SAN, and tighter risk management.

It is the virtualization of systems via a Virtual Operating Environment (VOE); however, that now garners the lion's share of attention from CIOs. More importantly, to maximize the benefits of a VOE, such as VMware® Virtual Infrastructure, shared SAN-based storage is a prerequisite. Advanced features, such as VMware Consolidated Backup and VMotion™, are designed specifically to leverage shared storage in order to provide a virtual machine (VM) with an unobtrusive snapshot-based backup process and mobility within a VOE.

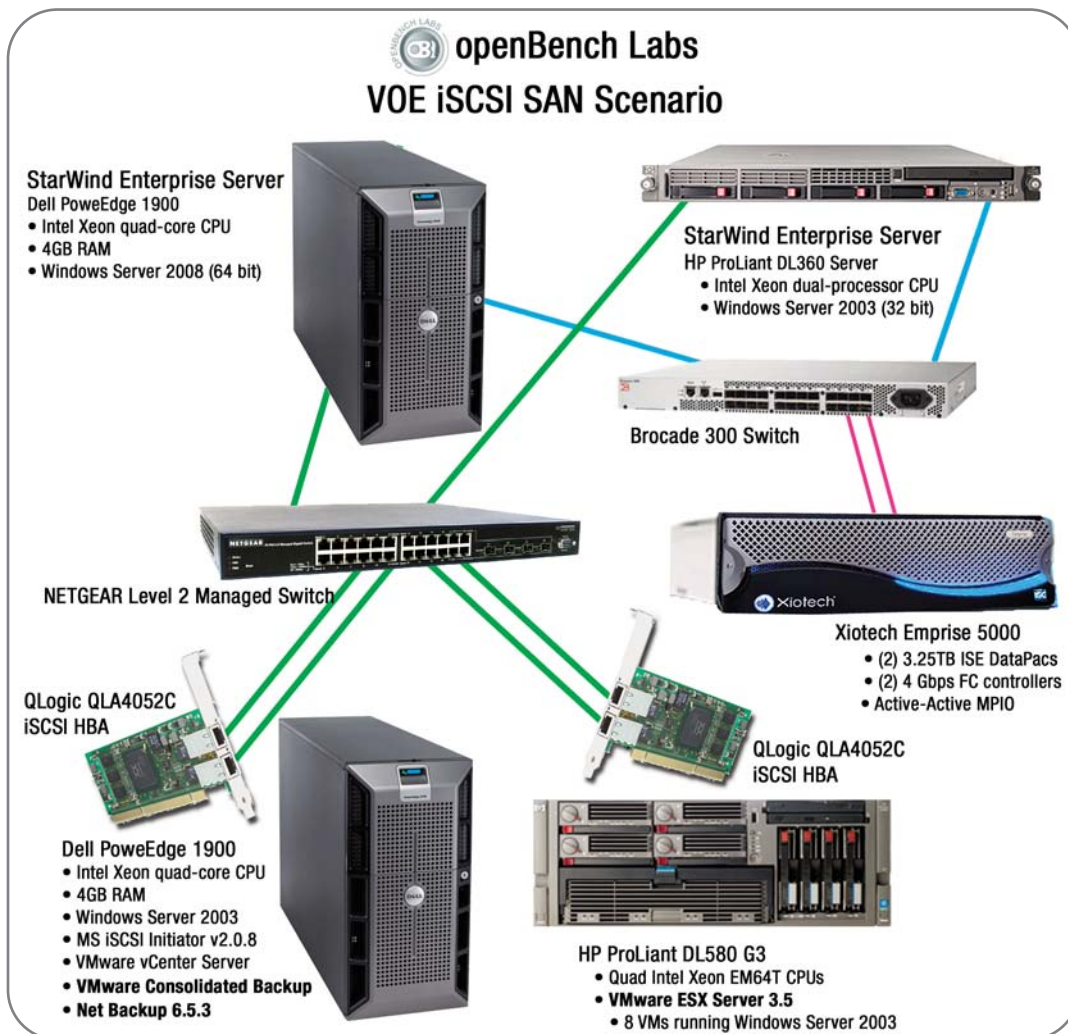
In particular, VM mobility is critical for advanced features such as load balancing and disaster recovery. The image of files stranded on a server that is not functional does not make a great poster for high availability. That explains why strong growth in server virtualization is spurring strong iSCSI adoption as the least complicated way to migrate from DAS to a SAN and harness holistic system and storage virtualization synergies. Remarkably, IDC reports revenue growth for iSCSI SAN equipment vendors grew 40 percent.

StarWind Enterprise Server further compounds the SAN value proposition with integrated software for the automation of storage management functions to ensure the high availability of data and protection against application outages. ***StarWind Enterprise Server's data management features include thin storage provisioning, data replication and migration, automated snapshots for continuous data protection, and remote mirroring for less than \$3,000.***

VMware iSCSI SAN Scenario

“VCB provides the framework needed to support both image-level and file-level backup operations; however, VCB requires that an ESX host share logical disks with a Windows-based server and that makes VCB support an excellent test of the iSCSI SAN flexibility, ease-of-use, as well as, the I/O throughput provided by StarWind Enterprise Server.”

iSCSI FROM DESKTOP TO VM



With server virtualization and storage consolidation the major drivers of iSCSI adoption, openBench Labs set up two servers running StarWind Enterprise

Server in a VMware iSCSI SAN test scenario. We focused our tests on StarWind's ability to support a Virtual Operating Environment featuring VMware ESX server hosting 8 virtual machines (VMs). With disk sharing essential for such advanced features as VMotion and VMware Consolidated Backup (VCB), the lynch pin for supporting these features with StarWind is the device sharing capability which is dubbed "clustering" in the device creation wizard.

We began by installing StarWind Enterprise Server on a pair of servers, a Dell PowerEdge® 1900 running the 64-bit version of Windows Server 2008 and an HP ProLiant DL360 running the 32-bit version of Windows Server 2003. We exported all iSCSI targets using a dedicated Gigabit Ethernet adapter over a private network. Both servers were configured with an 8Gbps FC HBA and all disk block pools for iSCSI target volumes originated on volumes imported from a Xiotech Emprise 5000 system. With two, active-active, 4Gbps FC ports, the Emprise 5000 system provided a uniform storage base for StarWind that eliminated back-end storage as a potential bottleneck.

As with any SAN device, StarWind Enterprise Server is about packaging a collection of physical disk blocks and presenting that package to a client system as a logical disk. There are several ways to do this within StarWind, which provides a complete set of tools for IT administrators to establish control of an entire iSCSI SAN infrastructure from a central point of management. The most functional and sophisticated of StarWind's storage virtualization methods provides thin storage provisioning, automated snapshot creation, and shared disk access.

VOE SUPPORT RIGORS

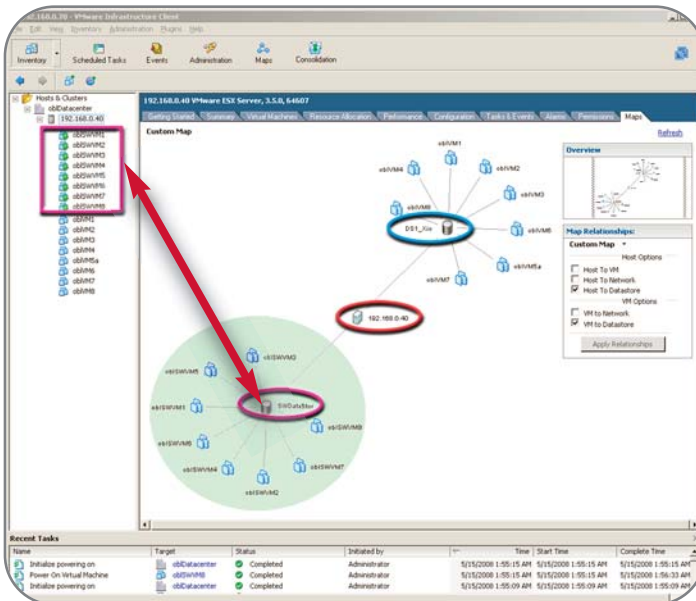
The openBench Labs VOE was hosted on a quad-processor HP DL580 ProLiant server, on which we installed a dual-port QLogic iSCSI HBA, the QLA4052, along with a 4Gbps Fibre Channel HBA. Similarly, we configured a quad-core Dell PowerEdge server, which ran Windows Server® 2003, VMware vCenter Server (a.k.a. Virtual Center), VMware Consolidated Backup (VCB), and Veritas NetBackup 6.5.3, with both a QLogic iSCSI HBA and an 8Gbps FC HBA. In this configuration, we were able to test StarWind's ability to provide shared access to iSCSI targets, which is critical for VCB.

Backup is not only a critical business support issue in a VOE, for IT it's a major technical support issue. Best IT practices call for classic file-level backups of VMs to be augmented with image-level backups that can be moved among VOE servers to enhance business continuity. VCB provides the framework needed to support both image-level and file-level backup operations; however, VCB requires that an ESX host share logical disks with a Windows-based server and that makes VCB support an excellent test of the iSCSI SAN flexibility, ease-of-use, as well as, the I/O throughput provided by StarWind Enterprise Server.

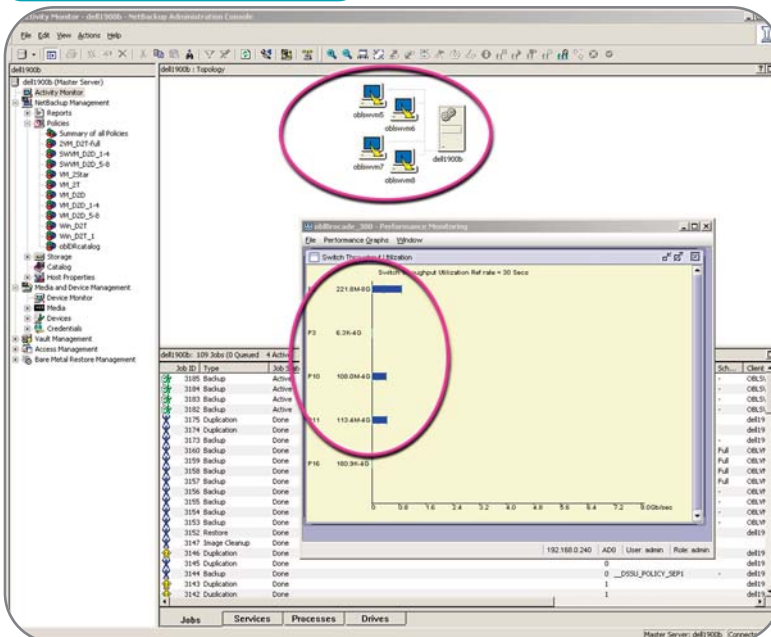
The pivotal component in a VCB package is a VLUN driver, which is used to mount and read ESX-generated snapshots of VM disks. VCB is installed on a

Windows Server host—dubbed the “proxy server”—that has SAN access to the datastore volume, which contains the appropriate VM files, on the ESX server. The proxy server must be able to mount the ESX datastore and access all vmdk files associated with the disk volumes that belong to the VM to be backed up.

ESX DATASTOR AND VM CONFIGURATION



STARWIND VCB SUPPORT



As a result, a VCB-based backup process stresses the flexibility of the SAN configuration as much as it stresses I/O throughput.

With both the proxy server and the ESX server able to access the datastore over the SAN, from the perspective of a VM and its business applications, the VCB backup window only lasts for the few seconds that the ESX Server needs to take and later remove a snapshot of the VM’s logical disk, which is represented by a vmdk file. During a backup, the proxy server instructs the ESX Server to initiate a snapshot, which freezes the vmdk file in a state that reflects the VM at the instance of the snapshot, and directs all new data writes to a special file dubbed a delta disk file.

Next, the ESX Server creates a snap ID and a block list of the frozen vmdk file and sends them to the proxy server. The Windows-based proxy server uses the VLUN driver to mount the vmdk file as a read-only drive. In this way, all data during the backup process is accessed and moved over the storage network for minimal impact on production processing.

At the end of a backup, the proxy server dismounts the vmdk file and the ESX Server removes the snapshot from the VM by consolidating the delta disk file into the vmdk file. As a

Using vCenter Server, we cloned an existing network of eight VMs, which was resident on an FC-based datastore, and deployed the new VMs on an iSCSI datastore that was exported by StarWind. We easily shared access to both the iSCSI and FC datastores with a Dell server running Windows Server 2003, VCB, and NetBackup.

With StarWind’s support for shared iSCSI targets, dubbed “clustering” for support Windows HS cluster configurations, we were easily able to implement and integrate VMware Consolidated Backup (VCB) with Veritas NetBackup (NBU) 6.5.3. During the backup process, the VCB proxy server initiated all I/O for both VCB and NBU. In the first stage of a backup, VCB read data on the shared ESX datastore, SWDataStor, and wrote a copy of the data files associated with each VM to a local directory on another the data to another iSCSI StarWind disk. With two Ethernet paths and four simultaneous VM backups in process, both reads and writes proceeded at over 100MB per second.

Snapshots & Replication in a VOE

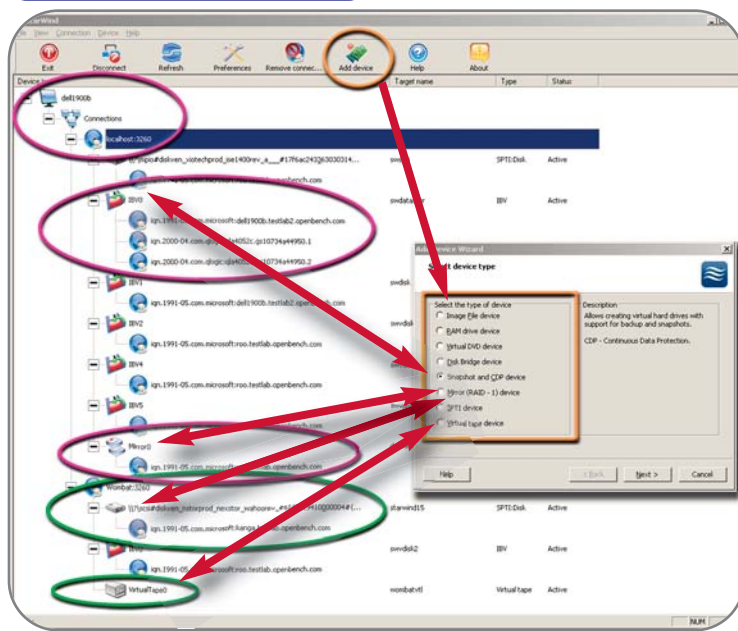
“While automated snapshots cannot replace a traditional backup process, in many scenarios that involve data corruption, StarWind’s automated snapshot and replication scheme provides a far superior Recovery Point Option (RPO) and a far better Recovery Time Option.”

FEATURE-BASED DEVICE PROVISIONING

As a full-featured, software-only, iSCSI target solution for Windows, StarWind Enterprise Server is a conceptually simple and deceptively complex application. Ostensibly, the purpose of installing StarWind Enterprise Server is to package collections of physical disk blocks and present those collections as logical disks to client systems. What complicates the StarWind equation are those robust full-featured services. The result can be a daunting initial out-of-box experience.

For IT, the business value of the StarWind solution lies in optimal resource utilization, minimal IT administration, and maximal performance. To meet that value proposition while also providing a wide array of optional features, StarWind utilizes a number of ways to package and virtualize pools of physical disk blocks.

ISCSI STORAGE PROVISIONING



set, such as a SATA drive, StarWind offers a Disk Bridge device that virtualizes the command set so that a non-SCSI drive or array can be employed.

While the menu to add a new iSCSI device is feature-centric, the eight options represent three different technologies for device virtualization. The most basic level of virtualization and simplest management task for IT involves exporting a raw physical device.

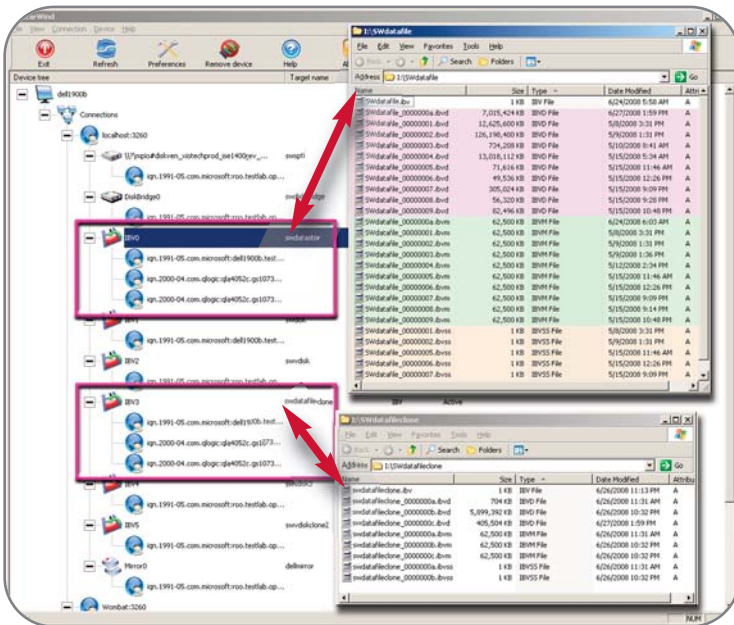
If the device uses the SCSI command set, then all StarWind needs to do is route the command and data streams. All extended functionality is provided by the device itself. This device choice is dubbed a SCSI pass-through interface (SPTI) device. If the device does not utilize the SCSI command

On any server running StarWind Enterprise Server, IT administrators can launch the StarWind management console and manage all StarWind Servers on a LAN. Provisioning a new storage device amounts to choosing a device with the appropriate support features, which range from a simple hardware SCSI pass-through interface (SPTI) device to an iSCSI disk with automated snapshots for continuous data protection (CDP). Within our VOE, all testing was done strictly using the latter family of image-based virtual disks, which StarWind designates as IBV volumes.

To do more than just virtualize the SCSI command set, StarWind needs to virtualize disk blocks and that means leveraging NTFS on the host server. The easiest way to do this is with an image file that represents the iSCSI device, which is exactly the technology found in the free version of StarWind. It is also the virtualization technology used to provide VTL and RAID-1 mirror devices.

ISCSI TARGET REPLICATION

ISCSI STORAGE PROVISIONING



The most sophisticated StarWind features involve thin provisioning, automated snapshots for continuous data protection, and replication—dubbed cloning. Providing those features requires the much more sophisticated StarWind image-based virtual disk (IBV) scheme. In this scheme, StarWind virtualizes disk blocks using separate files for disk headers (.ibvm), block data (.ibvd), and disk snapshots (.ibvss). This virtual disk image technology is what we employed in our VOE to virtualize iSCSI targets for the ESX server and the Windows VCP proxy server.

To test the advanced features of StarWind and how these features add value to a VOE, we created a new snapshot of our IBV device *SWdatafile*. *SWdatafile* was imported by our ESX server and used for the datastore, *SWDataStor*, which contained the vmdk files of our eight VMs. We then mounted that snapshot as an IBV device, *SWdatafileclone*, which was virtualized as an independent copy or *SWDataStor* on the ESX host server.

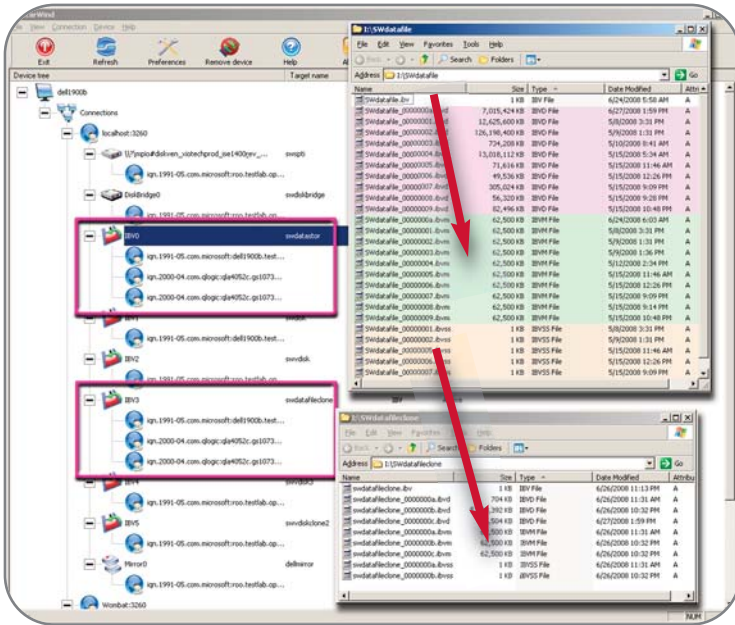
With our ESX datastore resident on a StarWind IBV device, we had encountered no issues in leveraging the advanced features that StarWind offers for IBV target devices. Since our datastore would immediately be required to contain the logical disk files of eight VMs, we needed to define a volume with sufficient room to provide for the VMs and also support future growth. To meet those needs, we created a 500GB target volume with thin provisioning. Initially this volume began with a few MBs of meta data, but it quickly swelled to consume about 155GB of real disk space as we loaded our eight VMs.

In addition to setting up thin provisioning to garner better utilization of storage, openBench Labs also configured the IBV device for the ESX datastore with support for snapshots. IT administrators can manually create a snapshot at any time. What’s more, IT can also enable automatic snapshot creation for an IBV target device.

The central issue with any automated snapshot scheme is the extent to which snapshots grow and consume resources. ***In our VOE testing, snapshot overhead for our ESX datastore typically was on the order of 100KB.*** That gives system and storage administrators a lot of leeway when deciding how to configure the

frequency for automatic snapshot creation and whether to limit the number of snapshots maintained. If a limit is placed on the number of snapshots, StarWind overwrites the oldest snapshot when the limit is reached.

ISCSI STORAGE PROVISIONING



The easiest way to mount and leverage a StarWind snapshot is through the creation of a new IBV device as a “linked clone.” *The linked clone process utilizes a snapshot of the original device to set up pointers, rather than duplicate the device’s data, in creating a point-in-time replica.* Once the replica is created, any changes to either the new device or the original are independent, which makes snapshots and replication powerful tools for disaster recovery.

We used the StarWind new device wizard to easily mount and access IBV device snapshots as a new IBV device. We employed the wizard to create a new device based on a snapshot of our SWdatafile device. We simply pointed to the snapshot we wanted to use and a location to create the initial files to support the new device dubbed SWdatafileclone.

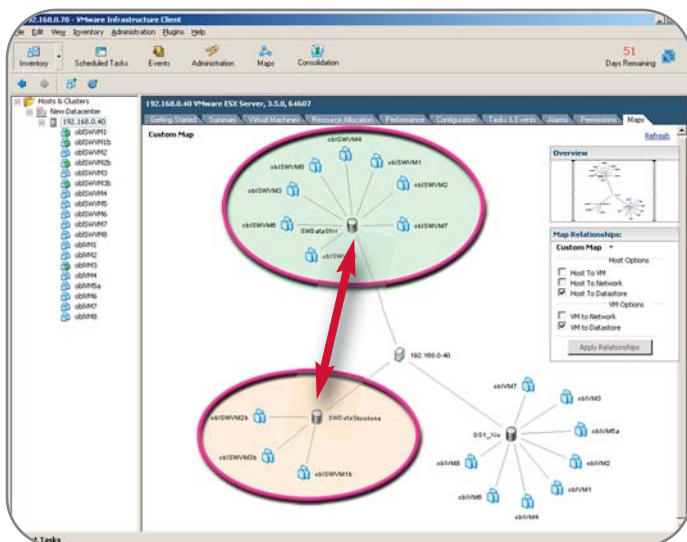
Snapshots and replication have profound implications for a VOE. Once we created the new IBV device, SWdatafileclone, we

used the VI Client to rescan the iSCSI HBA in our ESX host. The new device was found, mounted, and its VMFS datastore was imported. We dubbed the datastore SWDataStorclone. We now had two 500GB datastores and 16 VMs within our ESX host; however, we had only consumed 160GB of space on our StarWind Server.

For an ESX server hosting VMs running mission-critical applications, the StarWind process of creating automatic snapshots and mounting the snapshots as linked clones is much quicker and easier than a traditional VCB-based backup. While automated snapshots cannot replace a traditional backup process, in many scenarios that involve data corruption, StarWind’s automated snapshot and replication scheme provides a far superior Recovery Point Option (RPO) and a far better Recovery Time Option. What’s more, each of the VMs in

In addition to using the replicated datastore as a means of restoring the original VMs, we were also able to use the replicated datastore as a way to replicate the VMs resident in the datastore. The process involves the creation of a new VM by using the vmdk file for a system disk of one of the replicated VMs. Using the replicated system disk of oblSWVM1, we created oblSWVM1b. To complete the task, an administrator must launch the VM, change its LAN address, and join a domain under the new name.

ESX DATATOR AND VM CONFIGURATION



the replicated datastore can be used to easily create a new VM that can be mounted and run in its own right.

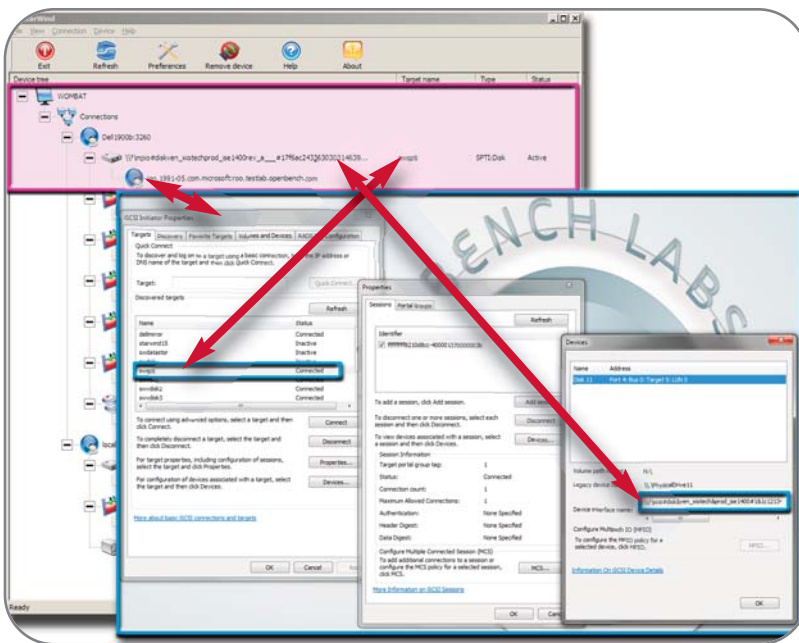
In addition to being able to clone a connected on-line disk using snapshot technology, the files representing an off-line virtual disk can simply be copied to another drive on the server or another server to create a new and fully independent virtual drive. ***As a result, a StarWind IBV device represents a virtualized iSCSI target that is entirely independent of the underlying storage hardware and the server running StarWind.***

Desktop Storage Consolidation

“Using an iSCSI SAN to provision desktop storage also provides IT with a means to deploy enterprise-class functions such a replication, volume snapshots, and continuous data protection.”

DESKTOP STORAGE CONSOLIDATION

iSCSI STORAGE PROVISIONING



One of the most promising areas to consolidate storage and improve resource utilization rates is on the desktop. The value proposition of iSCSI storage provisioning has always been its simplicity and low cost compared to an FC SAN. Nonetheless, it was the availability of the extra bandwidth provided by Gigabit Ethernet that launched significant adoption of iSCSI SANs.

The upcoming release of Windows 7 provides strong support for IT to provision systems with iSCSI storage and introduce desktop virtualization services. We were easily able to provision a workstation running Windows 7 RC with all of the iSCSI disk options available within StarWind Enterprise Server.

Now growing regulatory demands on risk management and data retention, along with the

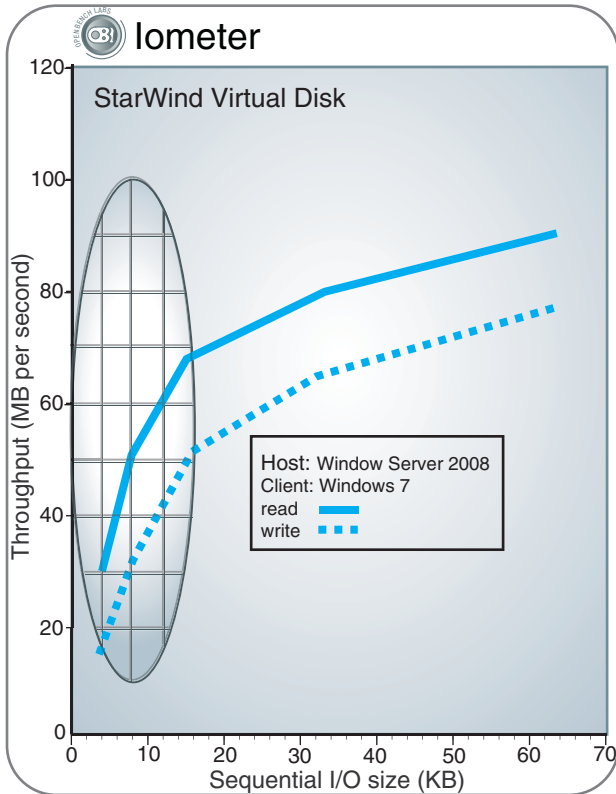
issues of Green IT are spurring the use of iSCSI SANs as a means to consolidate desktop storage resources. In particular, using an iSCSI SAN to provision desktop storage also provides IT with a means to deploy enterprise-class functions such a replication, volume snapshots, and continuous data protection. To support these aims, the new Windows 7 OS shares the Quick Connect iSCSI technology being introduced with Windows Server 2008 R2,

To assess the use of StarWind Enterprise Server as a means for IT to provision desktop systems with storage, openBench Labs utilized a Dell Precision T5400 workstation with a dual-core Xeon processor. Unlike on our server systems, we employed a dedicated 1Gbps Ethernet adapter for iSCSI traffic on the workstation, rather than an iSCSI HBA. On this system, we ran the Iometer benchmark to test the performance of a number of iSCSI virtualization technologies provided by StarWind Enterprise Server. While advanced functionality can sometimes trump performance on servers, the playing field is

much more level on the desktop.

WIRE-SPEED PERFORMANCE

iSCSI STORAGE PROVISIONING



Given the simplicity of SPTI devices, we chose to use several to act as base-level devices in our benchmarks. What’s more, we benchmarked the Fibre Channel performance of the underlying Emprise 5000 which sustained 5,300 8KB I/O requests per second on a single DataPac and provided 700MB per second sequential I/O.

As expected, sequential I/O on our Windows 7 workstation was entirely limited by the single Gigabit Ethernet connection. Equally interesting, sequential read and write throughput on all of our iSCSI targets, from a simple SPTI device to sophisticated image-based virtual disks were all statistically identical. While we measured no difference in sequential throughput, running Iometer with 8KB random access I/O requests, which were in a mix of 75% reads and 25% writes, uncovered differences that were an order of magnitude in performance.

On Windows 7 clients, sequential file I/O was statistically the same with all of the StarWind iSCSI devices. For mainstream desktop applications, 8KB read and write requests characterize nearly all I/O and we measured read and write throughput to be consistently around 50MB per second for 8KB reads and 30MB per second for 8KB writes.

Sequential access, however, stresses read-ahead caching far more than it does the ability of a drive to process I/O requests, which is needed for storage on servers for database processing. In tests of random I/O, StarWind devices that virtualized the SCSI command set significantly outperformed SPTI devices.

With an SPTI device, IOPS peaked at about 350. On the other hand, using a Disk Bridge device, which also exports a full disk target, but virtualizes the command set to enable the use of non-SCSI drives as simple targets, more closely reflected what we had measured with the underlying disk over Fibre Channel. In particular, when we connected to the same hardware, but as a Disk Bridge device, our StarWind iSCSI target sustained 3,250 8KB IOPS. Using full feature IBV devices also sustained that same 3,250 IOPS level on target devices used to provision our Windows 7 workstation.

Customer Value

“With two servers included in a StarWind Enterprise license, the ability to replicate and mirror devices eliminates single points of failure and all of the disk-related interruptions that get in the way of mission-critical business processing.”

For CIOs, the top-of-mind issue is how to reduce the cost of IT operations. With storage volume the biggest cost driver for IT, storage management is directly in the spotlight.

StarWind Enterprise Server Quick ROI

- 1) Lower the Total Cost of Storage Hardware by Eliminating Device-specific Feature License Purchases
- 2) Consolidate Site-wide Storage Resources for Higher Resource Utilization
- 3) Centralize Site-wide iSCSI Storage Provisioning via the StarWind Management Console
- 4) Unified SAN Thin Storage Provisioning Only Draws Upon Storage Capacity When Data is Written
- 4) High Availability and Disaster Recovery Supported with Automated Snapshots, Replication and RAID-1 Mirroring

At most SMB sites today, IT has multiple vendor storage arrays that all have similar functions that must be managed in different ways. From an IT operations perspective, multiple arrays with multiple management systems forces IT administrators to develop many sets of skills. Worse yet, if IT attempts to automate based on one of these proprietary functions, they may be unable to move data from one system to another. By tying data to the functions on one box, purchase options

and vendor choice is greatly reduced.

There are also substantial capital costs to be paid when the same critical management functions are repeatedly licensed for every storage array. For example, licenses for snapshot, mirror, and replication functionality on an IBM DS4200 storage array provisioned with 8TB of SATA disk storage will more than double the cost of the array. With an iSCSI SAN anchored by StarWind Enterprise Server, IT buyers can plan for new storage devices not with an eye to the bottom line, but with laser-like precision.

By building on multiple virtualization constructs, including the notion of a space of virtual disk blocks from which logical volumes are built, StarWind Enterprise Server is able to take full control over physical storage devices. In doing so, StarWind provides storage administrators with all of the tools needed to automate such critical functions as thin provisioning, creating disk snapshots, disk replication and cloning, as well as local and remote disk mirroring for High-Availability (HA) and Disaster-Recovery (DR) capabilities.

What's more, StarWind helps IT utilize storage assets more efficiently and avoid vendor lock-in. By leveraging the storage capacity of many disk arrays, IT

administrators no longer need to keep spare space available for each individual storage array. This is a particularly important advantage when pushing an iSCSI SAN out to desktop systems, which are notoriously underutilized and over provisioned.

StarWind can also play a very important role in supporting business continuity. With two servers included in a StarWind Enterprise license, the ability to replicate and mirror devices eliminates single points of failure and all of the disk-related interruptions that get in the way of mission-critical business processing. ***Replicating a disk that has automated snapshots to periodically create known-good points in time will provide a far better recovery image than any backup saveset.***