

OUTDATED

StarWind Virtual SAN Provides HA Storage for Hyper-V Clusters with DAS-Based Virtualized iSCSI Devices

Cost Effective HA Storage for Critical Business Applications

FEBRUARY 2015

TECHNICAL PAPER BY DR. JACK FEGREUS



Trademarks

"StarWind", "StarWind Software" and the StarWind and the StarWind Software logos are registered trademarks of StarWind Software. "StarWind LSF5" is a trademark of StarWind Software which may be registered in some jurisdictions. All other trademarks are owned by their respective owners.

Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, StarWind Software assumes no liability resulting from errors or omissions in this document, or from the use of the information contained herein. StarWind Software reserves the right to make changes in the product design without reservation and without notification to its users.

Technical Support and Services

If you have questions about installing or using this software, check this and other documents first - you will find answers to most of your questions on the [Technical Papers](#) webpage or in [StarWind Forum](#). If you need further assistance, please [contact us](#).

Copyright ©2009-2015 StarWind Software Inc.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of StarWind Software.

In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms".

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

About StarWind

StarWind is a pioneer in virtualization and a company that participated in the development of this technology from its earliest days. Now the company is among the leading vendors of software and hardware hyper-converged solutions. The company's core product is the years-proven StarWind Virtual SAN, which allows SMB and ROBO to benefit from cost-efficient hyperconverged IT infrastructure. Having earned a reputation of reliability, StarWind created a hardware product line and is actively tapping into hyperconverged and storage appliances market. In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms" following the success and popularity of StarWind HyperConverged Appliance. StarWind partners with world-known companies: Microsoft, VMware, Veeam, Intel, Dell, Mellanox, Citrix, Western Digital, etc.

Contents

Driving HA Through Virtualization	4
<i>Building HA on Virtualized DAS Devices</i>	4
<i>Layering Storage and CPU Availability for Near Fault Tolerance</i>	5
Maximizing iSCSI MPIO Performance	7
<i>Leveraging Virtual 10GbE NIC Ports</i>	7
StarWind Wizard-Based Configuration	9
<i>Hyper-V Failover Cluster on HA Storage</i>	10
<i>4-Node Virtual SAN Performance</i>	11
Application Continuity and Business Value	12

“By layering Hyper-V HA systems support on a foundation of StarWind HA-storage resources, IT is able to provide a near-FT for key business applications, for which IT is frequently required to support an SLA addressing business continuity.”

Driving HA Through Virtualization

In this analysis, openBench Labs assesses the setup of the High Availability (HA) features of a StarWind Virtual Storage Area Network (Virtual SAN) in a Microsoft® Hyper-V® environment. Using StarWind Virtual SAN software, IT is able to create virtual iSCSI storage devices using internal direct attached storage (DAS) on one or more servers. Virtual iSCSI devices can be shared among multiple Hyper-V hosts to create a cost-effective virtual infrastructure (VI) without single points of failure.

For technically savvy IT decision makers, this paper examines how to use StarWind to implement multiple, software-defined, hyper-converged, Virtual SAN that support VI high-availability (HA) features, including Hyper-V Live Migration in a Failover Cluster. Using two, three, or four StarWind servers, IT is able to configure storage platforms with differing levels of HA support, on which HA systems can be built using native VI options. By layering Hyper-V HA systems support on a foundation of StarWind HA-storage resources, IT is able to provide a near fault tolerance (FT) for key business applications, for which IT is frequently required to support a service-level agreement (SLA) addressing business continuity.

In particular, StarWind is seeking to change old HA notions that equate high-performance shared storage with an expensive complex hardware-dependent SAN. A StarWind Virtual SAN typically requires no capital expenditures and avoids disruptive internal IT struggles between storage and network groups via a wizard-driven management console.

Building HA on Virtualized DAS Devices

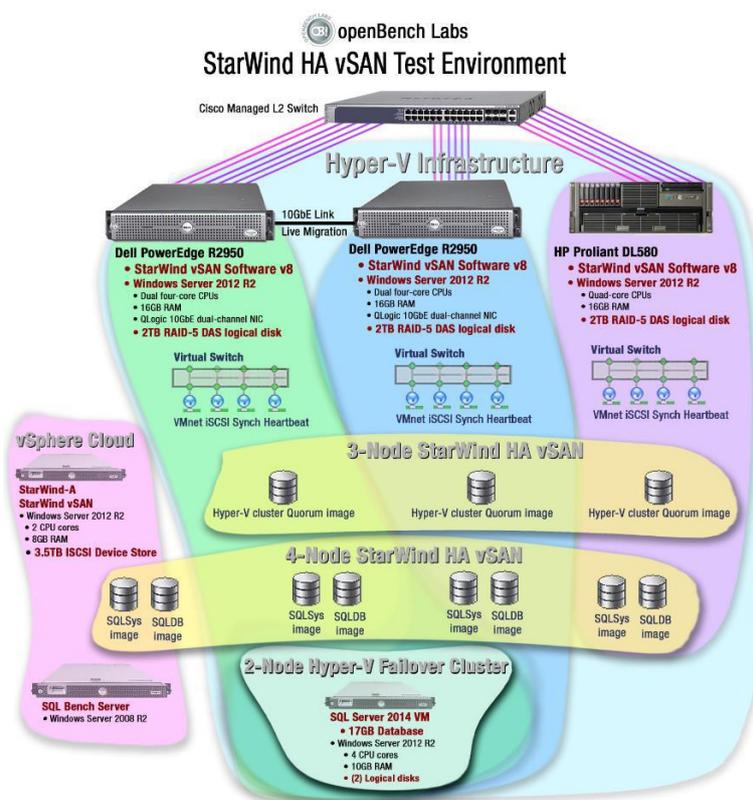
Demands on IT executives continue to center on projects with a potential for high revenue-generation. These business applications are typically data-rich, involve transaction-driven I/O, and help drive the double-digit growth of stored data. In addition, VI capabilities for enhanced availability are motivating CIOs to run business-critical applications, including on-line transaction processing (OLTP) applications on VMs. Consequently, data protection, disaster recovery (DR), high availability (HA), and fault tolerance (FT) are all becoming hot-button IT strategic planning issues for CIOs.

For robust application support, A StarWind Virtual SAN provides critical HA heartbeat and synchronization services to iSCSI targets without requiring special hardware, such as flash devices. With no expensive hardware constraints, a StarWind Virtual SAN enables IT to deploy cost-effective HA host and storage solutions for business-critical applications running on VMs. More importantly, a StarWind Virtual SAN is equally at home in a data center or at a remote office.

Building on a mature iSCSI software foundation, StarWind uses advanced virtualization techniques to represent device files as logical disks and decouples these devices from iSCSI targets. As a result, StarWind enables IT to address the needs of line of business (LoB) users by exporting logical disks with explicit support features tailored to application requirements.

In particular, IT administrators are able to invoke a wizard to create feature-rich logical devices utilizing NTFS files located on any standard storage resources, including DAS. What's more, IT administrators are able to modify and replace decoupled logical device files while leaving client iSCSI target connections live during the entire process.

Layering Storage and CPU Availability for Near Fault Tolerance



2-Node Hyper-V Cluster on a 3-Node StarWind Virtual SAN with 4-node Virtual SAN services

The StarWind iSCSI-based Virtual SAN service runs directly on the Windows Server OS of a Hyper-V host to virtualize devices for the hypervisor. In contrast, competitive products typically use Hyper-V VMs, which requires those services to pass all I/O through the hypervisor for virtual devices used by the hypervisor.

Creating a robust near-FT environment to run business-critical applications on a VM is a complex multi-dimensional infrastructure problem. In a Hyper-V, VI, failover clustering leverages shared SAN-based volumes to enhance VM host availability solely with respect to host execution. Shared data volumes are simply assumed to be constantly available.

“Given the many new roles of CSV in a failover cluster, Scaling the number of nodes in an HA Virtual SAN minimizes the impact of synchronizing a replica device in the event of a Virtual SAN device or node failure—much like a multi-spindle RAID array minimizes the impact of a disk crash.”

On the other hand, a multi-node HA StarWind Virtual SAN cluster ensures the availability of the data volumes exported to clients, which are typically Hyper-V and vSphere hosts needing specialized storage for deploying VMs. By synchronizing virtual device files that are replicated across multiple nodes having independent storage volumes, such as local DAS arrays, a HA StarWind Virtual SAN provides its iSCSI clients with a high level of data availability.

To support our Hyper-V VI environment, we used two Dell PowerEdge R2950 servers running Windows Server 2012 R2 with Hyper-V and an HP DL580 server running Windows 2012 and Hyper-V. To test our VI, we first created a 3-node HA StarWind Virtual SAN to provide a shared quorum disk for a 2-node Hyper-V Failover Cluster. By using a 3-node StarWind Virtual SAN, the quorum disk would remain available even if both of the nodes in Hyper-V cluster were to fail.

To provision the failover Hyper-V cluster with cluster shared volumes (CSV) based on HA storage, we leveraged StarWind’s ability to incorporate an active vSphere StarWind VM with physical Hyper-V servers to create a 4-node HA Virtual SAN. More importantly, in Windows Server 2012, CSV has been expanded beyond the Hyper-V workloads to include scale-out file server and SQL Server 2014. Given the many new roles of CSV in a failover cluster, scaling the number of nodes in an HA Virtual SAN minimizes the impact of synchronizing a replica device in the event of a Virtual SAN device or node failure—much like a multi-spindle RAID array minimizes the impact of a disk crash.

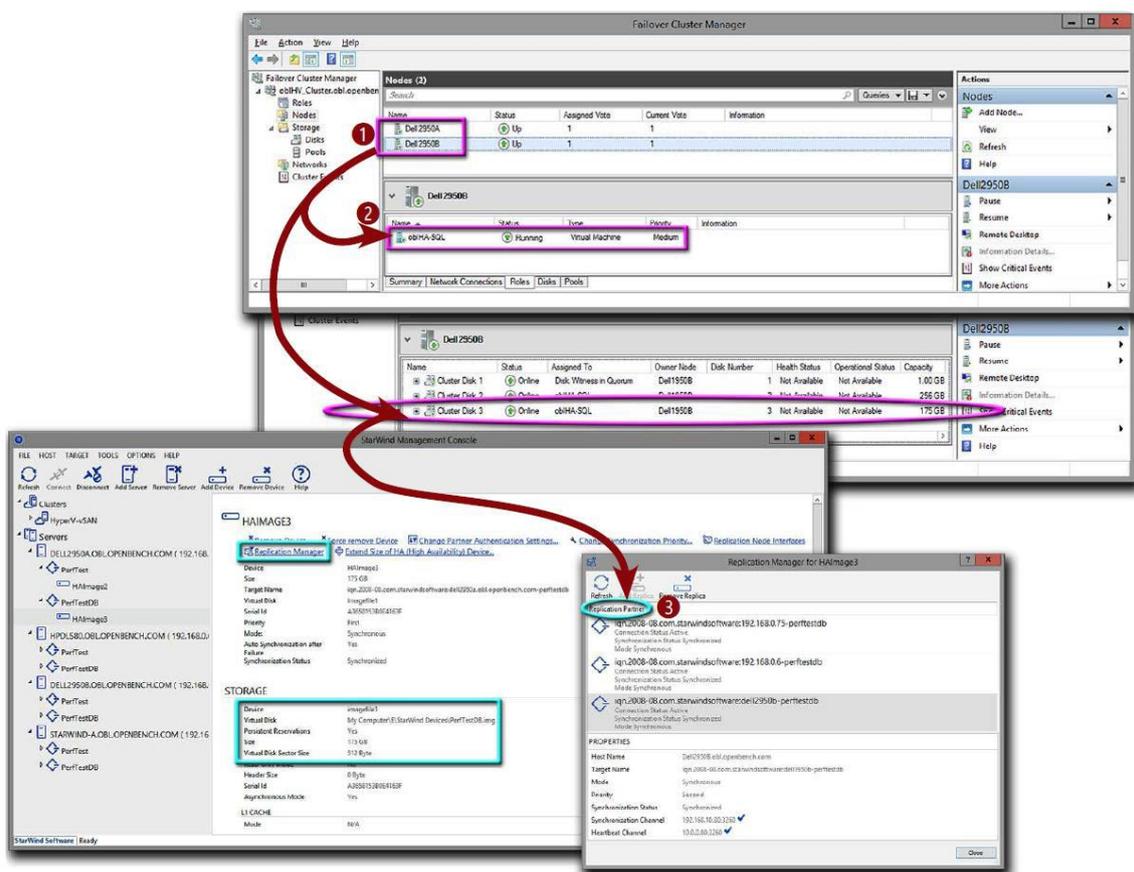
With a 4-node StarWind Virtual SAN, we were able to provide continuous processing for an I/O-intensive OLTP application in the event of either a storage or a host computing failure, while not incurring any performance penalty as resources were restored. What’s more, our configuration with separate logical devices supporting the VM’s OS and the VM’s data also demonstrate the ability to implement different levels of HA support for greater granularity in StarWind resource consumption.

Maximizing iSCSI MPIO Performance

We initiated storage provisioning for the StarWind Virtual SAN environment by configuring an array of six local SATA drives on each host as a RAID-5 volume. Using this volume, we configured a logical 50GB OS volume on each server and a logical multi TB VM storage volume to support both StarWind disk device files and local Hyper-V volumes. StarWind devices were used to provision the failover cluster with CSVs that had HA support for deploying near-FT VMs and to provision local Hyper-V hosts with volumes that featured advanced storage capabilities provided by an internal log-based file system for provisioning local VMs not requiring HA support.

By default, StarWind creates a 128MB L1 write back cache in RAM for each logical disk device and also supports an L2 cache using local flash storage devices. This memory cache can provide a significant boost for configurations using consumer-grade SATA disk drives. Given the 128MB cache on each of our enterprise SATA drives, as well as the significant amount of cache on the host's RAID controllers, we chose not provision a StarWind L1 RAM cache, which essentially would have extended host memory and CPU resources to cache data located in our existing data cache.

Leveraging Virtual 10GbE NIC Ports

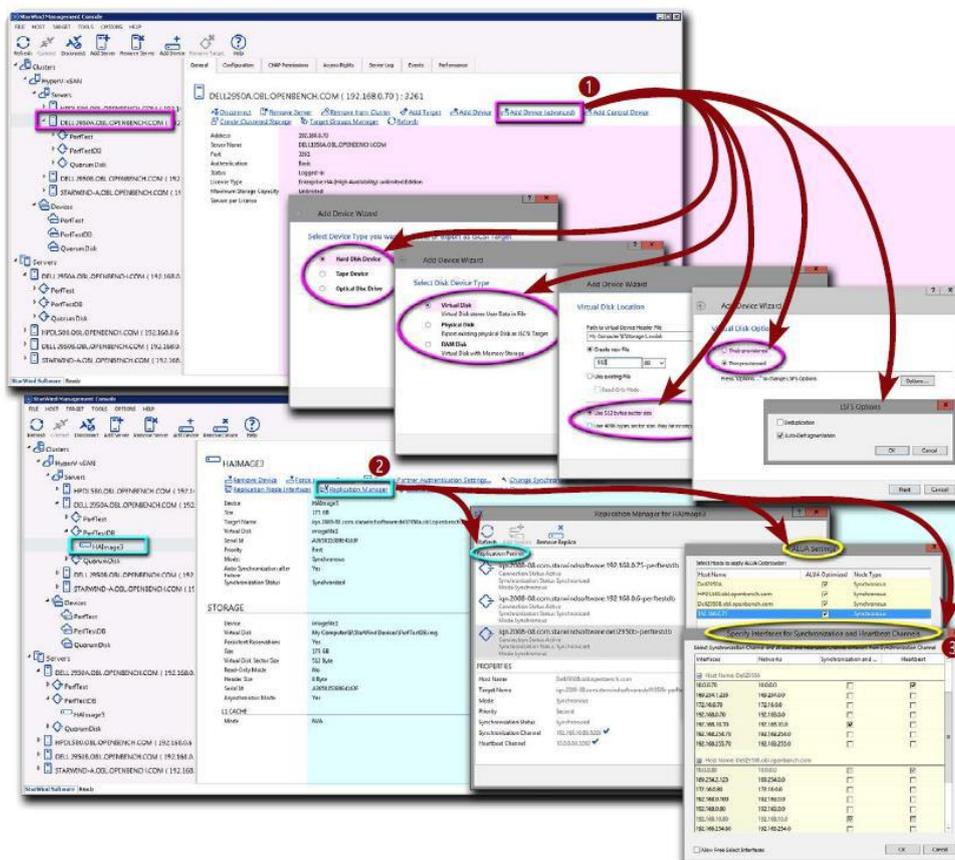


Virtual SAN Network Throughput Optimization via Virtual Switches

To maximize throughput in our Virtual SAN, we used a key Hyper-V networking feature to configure Virtual SAN ports on each host. For VM networking, Hyper-V utilizes virtual 10GbE switches. In particular, when an IT administrator creates a virtual switch^① to provide access to external systems, the virtual switch is linked to an installed physical NIC. If the physical NIC is shared between the virtual switch and the host, the host provisions a new virtual port^② based on the characteristics of the 10GbE virtual switch^③, rather than the characteristics of the host's physical NIC—a standard 1GbE NIC in our tests.

We leveraged virtual 10GbE host ports to support a dedicated 1GbE NIC for each StarWind network function: iSCSI data traffic, synchronization, and device heartbeat to simplify network management. We did not share any of the Hyper-V virtual switches used for Virtual SAN ports with any Hyper-V VMs. More importantly, we measured a significant I/O throughput boost streaming synchronization data, using a virtual 10GbE port data when restoring a failed device.

StarWind Wizard-Based Configuration

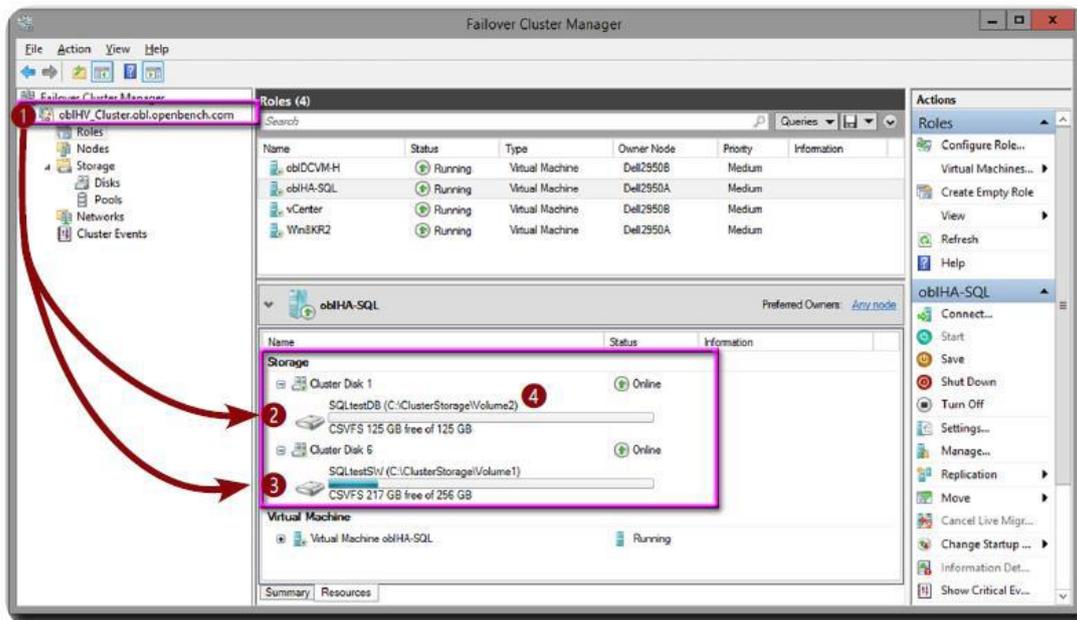


StarWind Virtual SAN Device and HA Partner Configuration

With the StarWind Virtual SAN service installed on each Hyper-V host, we utilized the Management Console to create, attach, and manage logical device files on multiple servers. Using the context-sensitive configuration wizards in the Management Console, an administrator follows a step-by-step series of options to create¹ and configure iSCSI devices and targets independently of each other. Consequently, device features, can be modified without changing active target configurations. StarWind device options include block sector size and the use of a log-based file system capable of supporting thin-provisioned and data de-duplication. For an HA SAN device, a thick-provisioned disk image is required.

Wizards to modify device properties after configuration include the Replication Manager². IT administrators use the Replication Manager to set HA characteristics including the replication partner nodes, asynchronous logical unit access (ALUA) path settings—all target paths are optimized by default—and assign the networks³ that will be used to isolate synchronization and heartbeat traffic from iSCSI client traffic. Isolation of data synchronization traffic is critical in a scaled-out Virtual SAN, in which the volume of synchronization data received from Virtual SAN partner nodes will be two to three times greater than the volume of iSCSI data received from a client, as each Virtual SAN node consolidates a full copy of all iSCSI data sent by a client host server.

Hyper-V Failover Cluster on HA Storage



Hyper-V Failover Cluster with HA Storage

To create a near-FT platform to run our OLTP business application, we configured a failover Hyper-V cluster, dubbed obIHV_Cluster¹, on two Dell R2950 servers. To facilitate Hyper-V cluster failover, we used our 3-node StarWind Virtual SAN to provision a logical quorum disk² that would remain available if either or both of the cluster hosts failed,

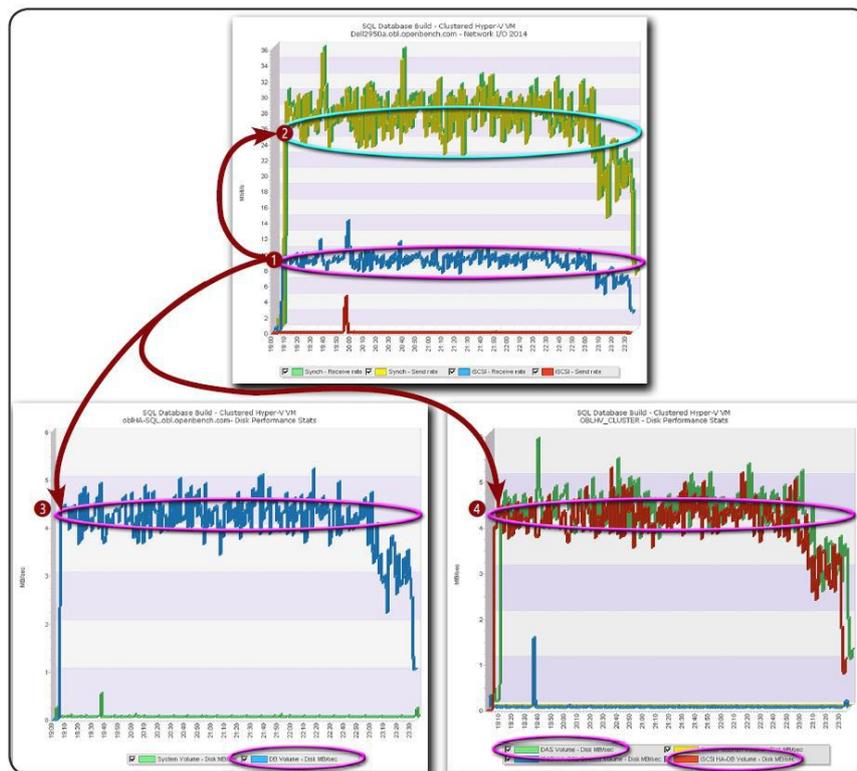
Using the Failover Cluster Manager, we added two CSV³ that were created on the 4-node HA Virtual SAN to obIHV_Cluster for the deployment of near-FT VMs. In particular, we used one CSV to both of the cluster shared volumes to provision the role of a failover cluster VM for obIHA-SQL⁴, which would run our OLTP business application. One CSV supported the system while the other supported the application database.

Our OLTP test scenario is built on a business model for a brokerage firm. In this model, customers generate transactions to research market activity, check their account, and make stock trades. The company makes market transactions to execute customer orders and updates customer accounts. To support this business model, the underlying SQL Server database has 33 tables that contain all of the unique active trading information of customers, along with the customer's detailed trading history. The database also contains a set of fact and dimension tables with common data, including company names, stock exchanges on which companies trade, client identification, broker identification, and tax codes.

To test this model, all of the queries directed at the model's database are generated on an external client running on a VM in our vSphere VI. The client program uses ten weighted transaction templates

and a random number generator to create SQL queries that spread transactions for new trading activity across users and stocks. To ensure that our queries sufficiently exercise the underlying storage subsystem rather than cache, we populate the database with 24GB of table data.

4-Node Virtual SAN Performance



4-Node Virtual SAN Database Build

With the infrastructure provisioned for the test VM, oblHA-SQL, we ran a script to populate the database table space with 24GB of application data. While populating database tables and generating indices on oblHA-SQL, each node in the 4-node Virtual SAN exhibited a nearly identical network traffic pattern:

- Each of the four Virtual SAN nodes received a steady 8-Mbps stream of iSCSI write data **1**,
- Each Virtual SAN node sent synchronization data **2** consisting of new iSCSI data to each replication partner at 24 (8x3) Mbps.
- Each Virtual SAN node received synchronization data **2** from each replication partner containing new iSCSI data at 24 Mbps.

On the oblHA-SQL VM **3**, the four 8-Mbps iSCSI Virtual SAN streams to the Virtual SAN nodes were manifested as a 4-Mbps disk I/O stream to the VM's database volume. Similarly, we monitored two 4-Mbps write streams on the logical Hyper-V cluster host. In particular, we monitored a 4-Mbps write stream to the logical CSV disk **4** supporting the oblHA-SQL database and a 4-Mbps write stream DAS volume **4** on which the StarWind logical device files were located.

“StarWind provides storage administrators with tools needed to automate key storage functions, including as thin provisioning, disk replication, and data de-duplication, in addition to providing multiple levels of HA functionality to support a near-FT environment for business continuity.”

Application Continuity and Business Value

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 IT spotlight.

HA StarWind Virtual SAN Feature Benefits	
1	StarWind Reduces Capital Expenses associated With Storage Provisioning: Starwind’s Virtual SAN service enables SATA and SAS DAS storage to be shared in an iSCSI Virtual SAN without requiring the use of Flash Storage as a Level-1 cache.
2	StarWind Reduces Capital Expenses associated With Server Provisioning: StarWind’s light-weight Virtual SAN service runs directly on Windows host servers, which enables IT to retask older servers to scale out storage HA features.
3	StarWind Reduces VI Management Overhead: By decoupling iSCSI devices and targets, StarWind enables device reconfiguration without impacting iSCSI target sessions. In particular, HA support levels can be reconfigured without involving a Hyper-V host.
4	Simplifies Virtual SAN Scale Out: By virtualizing HA support as a device-specific function, we were able to scale out multiple levels of HA support as a datastore-specific feature.

At most small-to-medium business (SMB) sites today, IT has multiple vendor storage arrays with similar functions and unique management requirements. From an IT operations perspective, multiple arrays with multiple management systems forces IT administrators to develop many sets of unrelated skills.

Worse yet, automating IT operations based on proprietary functions, may leave IT unable to move data from system to system. Tying data to the functions on one storage system simplifies management at the cost of reduced purchase options and vendor choices.

By building on multiple virtualization constructs a StarWind Virtual SAN service 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 key storage functions, including as thin provisioning, data de-duplication and disk replication, in addition to providing multiple levels of HA functionality to support business continuity.



Jack Fegreus

Jack Fegreus is the founder and Managing Director of openBench Labs and consults through Ridgetop Research. Previously he was Editor in Chief of Digital Review and Data Storage. Jack was the founding Editor of Open Magazine, BackOffice CTO, and Client/Server Today. Jack also served as the CTO of Strategic Communications and IT Director at Riley Stoker Corp. Jack holds a Ph.D. in Mathematics and worked on Lie Algebras the application of computers to symbolic logic.

openBench Labs

Web-site: www.openbench.com

Tel: (508) 329-1645

E-mail: jack.fegreus@openbench.com

Fax: (508) 329-1640

StarWind Software

US Headquarters

 +1-617-4497717
 +1-617-507-5845

EMEA and APAC

 +44 20 3769 1857 (UK)
 +49 302 1788 849 (Germany)
 +33 097 7197 857 (France)
 +7 495 975 94 39
(Russian Federation and CIS)
 1-866-790-2646

Support Portal: <http://www.starwind.com/support>

Sales: sales@starwind.com

Support Forum: <http://www.starwind.com/forums>

General Information: info@starwind.com



www.starwind.com StarWind Software, Inc. 35 Village Rd. Suite 100 Middleton, MA 01949 USA

©2015, StarWind Software Inc. All rights reserved.