

Microsoft SMB File Sharing Best Practices Guide

Tintri VMstore, Microsoft SMB 3.0 Protocol,
and VMware 6.x

Author: Neil Glick
Version 1.0 06/15/2016

Contents

Executive Summary	1
Consolidated List of Practices.....	1
Benefits of Native Microsoft Windows 2012R2 File Sharing	1
SMB 3.0	2
Data Deduplication	2
DFS Namespaces and DFS Replication (DFS-R)	2
File Classification Infrastructure (FCI) and File Server Resource Management Tools (FSRM)	2
NTFS and Folder Security	2
Shadow Copy of Shared Folders	3
Disk Quotas	3
Deployment Architecture for a Windows File Server	3
Microsoft Windows 2012R2 File Share Virtual Hardware Configurations.....	4
Network Best Practices for SMB and the Tintri VMstore	5
Subnets	6
Jumbo Frames	6
Creating the Microsoft Windows 2012R2 File Server	6
Utilizing the Microsoft Windows 2012R2 File Server	7
Host Side Compression	8
Host Side Quotas	8
Protecting the File Share	9
Tintri Snapshots and Replication	9
Microsoft Shadow Copies of Shared Folders	10
File Server High Availability	10
VMware vMotion	10
VMware vSphere High Availability	11
VMware Distributed Resource Scheduler (DRS)	12
VMware Site Recovery Manager (SRM)	12
Conclusion	12

Executive Summary

This paper demonstrates best practices for creating a Microsoft SMB file sharing solution with Tintri VMstores. Microsoft SMB file sharing solutions can be architected in a number of ways. This paper explores some of these scenarios but limits the scope to only the best practices that apply to Tintri VMstores.

The Tintri VMstore is purpose-built storage solution for virtual machines. IT administrators with working knowledge of virtualization can easily deploy Tintri storage. When deploying Tintri storage, there are no prerequisite operations such as LUN provisioning, HBA compatibility checks, or FC LUN zoning operations. From a VMware administrator point of view, the entire Tintri VMstore is presented as a single datastore.

Tintri VMstore delivers extreme performance, VM density, and a wide variety of powerful data management features, which are seamlessly integrated with vSphere. These examples of data management functionality include snapshots, clones, instant bottleneck visualization, and automatic virtual disk alignment. Tintri VMstore extends and simplifies the management of virtual machines (VMs) through intrinsic VM-awareness that reaches from the top of the computing stack, all the way down into the storage system.

Consolidated List of Practices

The table below includes the recommended practices in this document. Click the text on any of the recommendations to jump to the section that corresponds to each recommendation for additional information.

DO: Use a 10GbE network for data traffic.

DO: Use sockets and minimize usage of cores

DO: Use Paravirtual SCSI controller type and place the file share on a separate SCSI controller than the boot disk.

DO: Use different subnets to separate the VMstore Admin, Data and Replication Networks.

DO: Ensure that all servers and switches are configured with the same jumbo frame settings.

DO NOT: Enable Microsoft deduplication on a Tintri all flash array.

DO: Use Tintri protection utilities in conjunction to Shadow Copies of Shared Folders to ensure data longevity and integrity.

DO: Use GUID Partition Table (GPT) and not Master Boot Record (MBR) as the partition style for the file share disk.

DO: Give appropriate permissions to users and groups who will use the new file share.

DO NOT: Enable host side compression unless a T600 or T500 is used and there are enough resources to accomplish the extra overhead.

DO: Use Tintri VM-consistent snapshots to ensure proper backup of user's data.

DO: Enable VMware vSphere HA in the virtualized file share environment.

DO: Enable VMware vSphere DRS as long as resources(memory and CPU) are available for VM's to migrate to.

Benefits of Native Microsoft Windows 2012R2 File Sharing

The native file sharing role in Microsoft Windows 2012R2 offers many benefits:

- Native support for the SMB 3.0 protocol
- Data Deduplication
- DFS Namespaces and DFS Replication (DFS-R)
- File Classification Infrastructure and File Server Resource Management Tools (FSRM)
- NTFS and Folder Security
- Shadow Copies of Shared Folders
- Disk Quotas

SMB 3.0

Microsoft Server Message Block (SMB) protocol is a network file sharing protocol standard supported by Microsoft Windows. It is a client-server protocol that consists of a set of data packets containing a request by the client and a response sent by the server. SMB 3.0 is the latest version of this protocol. It was introduced with Windows Server 2012 and Windows 8 and contains features not available in SMB 2.x. Additional details of the SMB 3.0 protocol can be found at www.microsoft.com.

Highlights of the SMB 3.0 protocol includes:

- SMB Transparent Failover
- SMB Scale Out
- SMB Multichannel
- SMB Direct
- SMB Encryption
- VSS for SMB file shares
- SMB Directory Leasing
- SMB PowerShell

Data Deduplication

The Microsoft data deduplication role can be used with the file share to reduce storage capacity usage. Resource usage is kept low by throttling the CPU and memory that the process consumes. Be aware that Microsoft data deduplication is post-process and can introduce more IO on the underlying storage infrastructure.

When deploying file shares with the Tintri VMstore it is not recommended to enable Microsoft deduplication when using Tintri VMstore all flash arrays since inline data deduplication and compression is performed on the array as data is written to disk.

When using Tintri hybrid arrays enabling host side deduplication should be investigated on a case by case basis. Since the T600 and T500 arrays do not compress or deduplicate, host side deduplication can be enabled if enough resources exist to perform the actions. For the T800 series data is not deduplicated on the hard drives, but is compressed. If the data is highly compressible, deduplication may not be needed.

DO NOT: Enable Microsoft deduplication on a Tintri all flash array.

DFS Namespaces and DFS Replication (DFS-R)

DFS Namespaces enable shared folders located on different file shares to be grouped together into one or more logical namespaces. Each namespace appears as a single file share to users. This allows logical grouping of resources that are located on multiple servers at different sites.

DFS Replication (DFS-R) replicates individual folders, as well as folders in the DFS namespace, across servers and sites. DFS Replication uses a compression algorithm called remote differential compression (RDC), which detects changes in files and allows DFS-R to replicate only changed blocks.

File Classification Infrastructure (FCI) and File Server Resource Management Tools (FSRM)

File Classification Infrastructure (FCI) automatically classifies files stored in the file share through the File Server Resource Management Tools (FSRM). This allows for file expiration, custom tasks and reporting based on the files business value. In today's file sharing infrastructures, data sprawl is a real and difficult problem to tackle. FCI and FSRM can help eliminate some of the manual tasks administrators face when classifying the importance of files stored on their file shares.

NTFS and Folder Security

The Microsoft NTFS file system is a log based file system that keeps track of changes to the files on the file share. In the event of an error CHKDSK needs to only roll back transactions to the last commit point to recover consistency in the file system. This alleviates the need to run a disk repair utility on a partition running NTFS in the event of an error.

NTFS offers folder security for the file share server. Two sets of permissions are available when using NTFS; the share permissions set on the folder itself and NTFS permissions set on the folder which can also be set on individual files. In the event of a permission conflict, the more restrictive of the two are applied.

Shadow Copy of Shared Folders

Microsoft Shadow Copy provide point-in-time copies of files located on the file share. Users can view previous versions or recover files accidentally deleted. It is recommended to use the integrated protection utilities located on the Tintri VMstore to provide synergistic backup and recovery capabilities for the file share for the following reasons:

- Shadow Copies of Shared Folders are not transportable
- When storage limits are reached, the oldest shadow copy will be deleted to make room for newly created copies.
- A limit of 64 shadow copies per volume can be stored. When 64 is copies are created, the oldest copy will be deleted and cannot be retrieved.
- Shadow Copies of Shared Folders can only be set on a per-volume basis. Specific folders and files in a volume cannot be selected or not selected.

DO: Use Tintri protection utilities in conjunction to Shadow Copies of Shared Folders to ensure data longevity and integrity.

Disk Quotas

Disk quotas are a useful mechanism in Windows file sharing and can be enabled directly to a disk or can be applied to specific users. Each method limits the amount of disk space that users can utilize on a specified disk. Directory quotas can be set, but are enabled in FSRM and are out of scope for this paper. More information regarding disk quotas can be found at www.microsoft.com and in the host side quotas section of this paper.

Deployment Architecture for a Windows File Server

This guide provides best practices for hosting an SMB file sharing solution on the Tintri VMstore in a VMware ESXi virtualized environment. The file sharing server documented in this paper uses a 2TB VMDK to provide the storage for a Microsoft Windows 2012R2 virtual machine. The Tintri VMstore can utilize vDisks up to 64TB. This machine is configured to run the Microsoft File Server Role. Client connectivity to the file server was provided through a 10GbE network architecture which is recommended, but not required.

DO: Use a 10GbE network for data traffic.

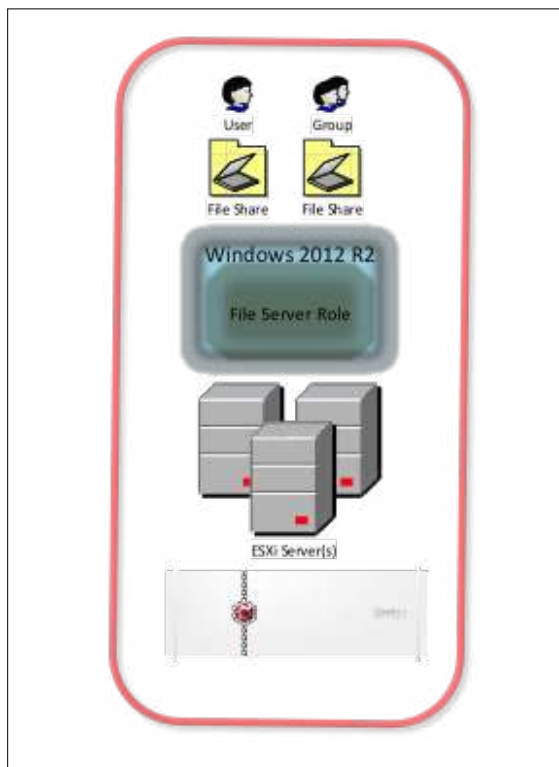


Figure 1 – Deployment Architecture for a Windows File Server

Microsoft Windows 2012R2 File Share Virtual Hardware Configurations

Microsoft Windows 2012R2 can be set up in a variety of ways, but the following virtual hardware configuration was selected to provide the best file share performance for the scenarios tested. Alternative configurations can be selected to fit the environment and purpose the file share will provide.

Four GB of memory and six virtual CPUs were selected to run the file server. The amount of RAM and number of CPUs will vary greatly from each environment depending on workload requirements put on the server.

Modern operating systems allow for a greater number of virtual sockets to achieve the appropriate number of CPUs. If possible, limit the number of cores and increase the number of virtual sockets as needed. This not only allows for the hot add of virtual CPUs, but enables the optimization of applications for Non-Uniform Memory Access (NUMA). More information regarding NUMA can be found at www.intel.com. Whichever method of CPU and RAM layout is chosen, ensure appropriate licensing is followed.

DO: Use sockets and minimize usage of cores

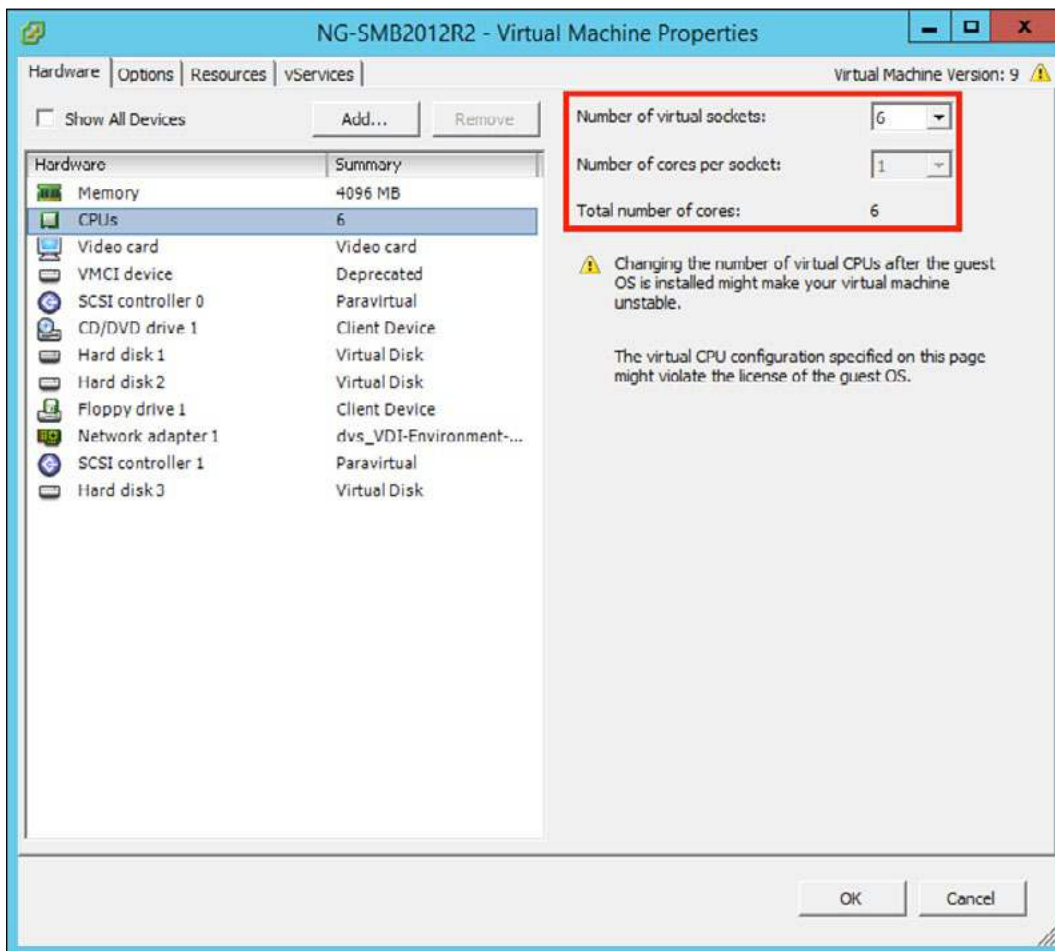


Figure 2 – Memory and CPU Overview

The Two TB VMDK can be seen in Figure 3. A Paravirtual SCSI controller was chosen to optimize performance and lower CPU usage. Paravirtual SCSI controllers can be selected with ESXi 4.x and later. The new hard disk was added to a new SCSI controller separate from the boot disk, SCSI (1:0) Hard disk. This method can improve performance and helps to avoid data traffic congestion when multiple disks are added to a virtual machine. For more information see the VMware vSphere Virtual Machine Administration guide Update 1 for ESXi 6.0 and vCenter Server 6.0.

DO: Use Paravirtual SCSI controller type and place the file share on a separate SCSI controller than the boot disk.

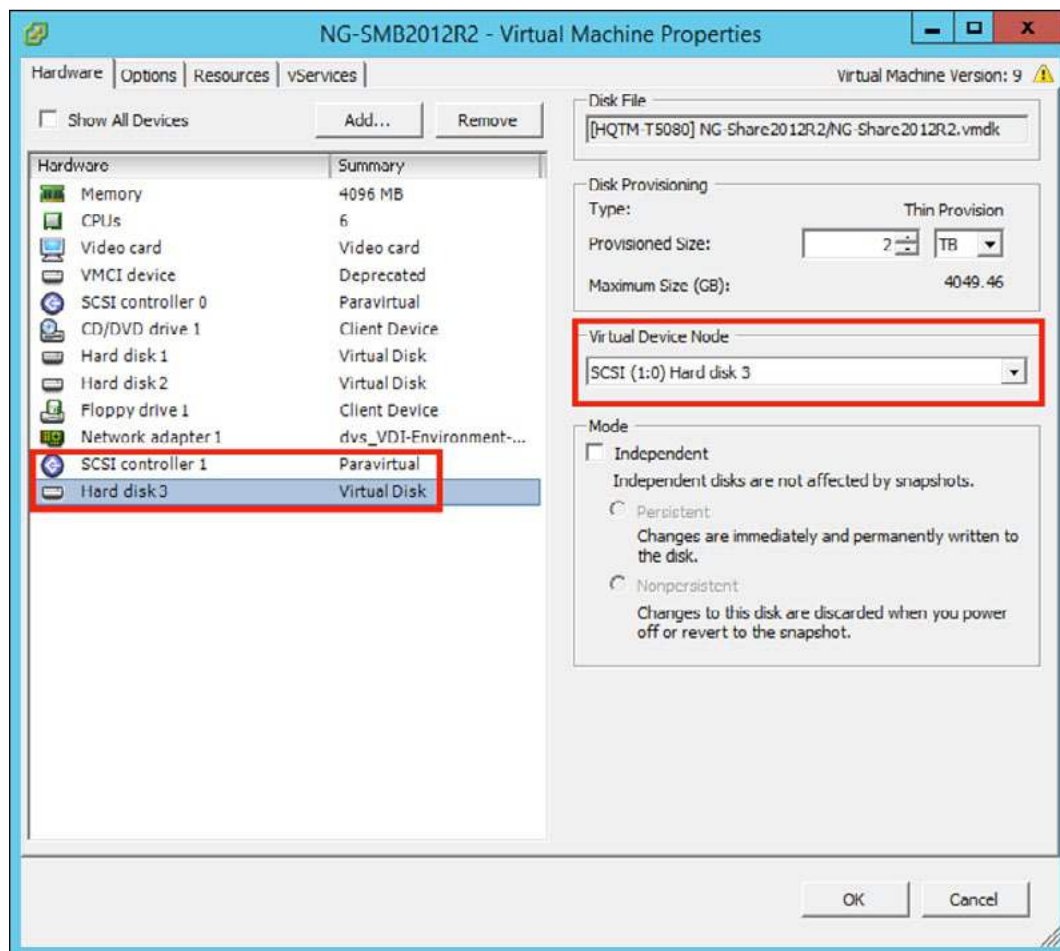


Figure 3 – SCSI Controller and Virtual Device Node

Network Best Practices for SMB and the Tintri VMstore

The Tintri VMstore provides shared highly available storage for SMB file shares. It has multiple hardware redundancies to ensure high availability. Each VMstore has two physical controllers deployed in an active/standby configuration. Each controller has separate Ethernet ports which are used to create the following networks:

Admin Network	Dual 1GbE ports for VMstore management traffic
Data Network	Dual 10GbE ports for data access*
Replication Network	Optional separate dual 1GbE or 10GbE ports for replication

* Note: the Tintri VMstore model T820 comes standard with 1GbE network ports for the data network. Optional 10GbE network cards are available. All other VMstore models come with 10GbE standard.

If a controller fails or if a network interface in the active controller fails, then the standby controller will seamlessly take over the management and data networks. The virtual machines on the ESXi hosts will continue to run without disruption. This seamless failover is also used during Tintri OS upgrades where failover between controllers is performed manually.

Subnets

The VMstore Admin Network, Data Network and Replication Network should use different subnets to isolate the network traffic.

DO: Use different subnets to separate the VMstore Admin, Data and Replication Networks.

Jumbo Frames

The Tintri VMstore supports Ethernet jumbo frames which can be employed to improve the performance of large data transfers over the Data Network. If jumbo frames are deployed then the VMstore, network switches, and the ESXi Hosts must all be configured with the same jumbo frame settings. Erroneous mismatches between jumbo frame settings can result in poor performance and network connectivity issues.

DO: Ensure that all servers and switches are configured with the same jumbo frame settings.

Creating the Microsoft Windows 2012R2 File Server

Creating the new Microsoft Windows 2012R2 file server simply involves creating a virtual machine in vCenter and placing it on the Tintri VMstore. To attach the file share VMDK storage to the file server follow the steps in Figure 2 and 3 located in Microsoft Windows 2012R2 File Share Virtual Hardware Configurations section pictured earlier in this paper.

The size of the VMDK will depend on the total storage needed to support the file server. A separate volume is not created since Tintri storage is presented to vCenter as a single datastore. To learn how to easily add storage to vCenter, see the Tintri VMstore System Administration Manual located at support.tintri.com.

Both a large single VMDK and multiple smaller disks striped together were tested using Vdbench. Vdbench is a command line performance utility that generates disk I/O. More information regarding Vdbench can be found at <http://www.oracle.com/technetwork/server-storage/vdbench-downloads-1901681.html>.

Each method performed well and no performance gain was realized using either method. It should be mentioned that Microsoft volumes larger than 2TB cannot be hot expanded through vCenter while the virtual machine is powered on.

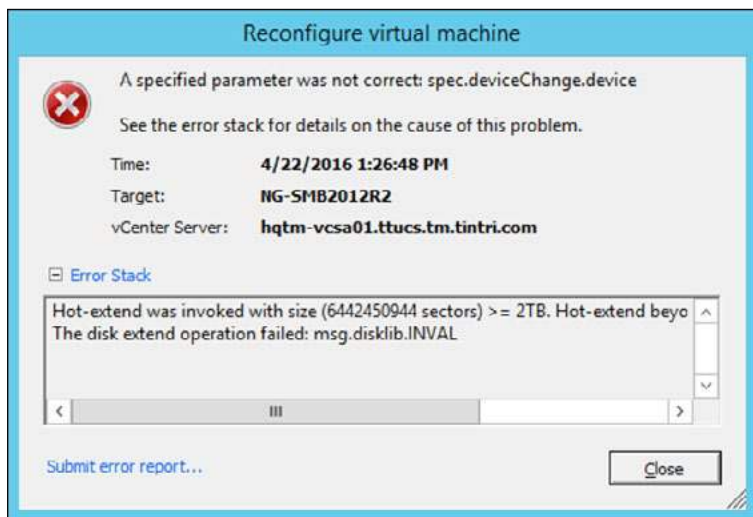


Figure 4 – Expanding a VMDK beyond 2TB While Virtual Machine is powered on.

Many choices exist when configuring storage on the Microsoft file share; basic, dynamic, simple, spanned and striped volumes. Deciding which method to implement is out of scope for this paper and should be considered for present and future growth. For further information see www.microsoft.com.

After the VMDK has been added to the virtual machine it will appear in Microsoft Disk Management as a new offline disk. Disks 2TB and larger can only be initialized as GUID Partition Table (GPT) format and will not give the option to select Master Boot Record (MBR). GPT is a new standard and is replacing MBR. Unless there is a specific reason, use GPT and not MBR.

DO: Use GUID Partition Table (GPT) and not Master Boot Record (MBR) as the partition style for the file share disk.

After the new disk has been set online and initialized a new volume can be created. Select the type of volume that best suits your file share. When creating the new volume an allocation unit size (AUS) will need to be selected in the Format Partition screen. The choices range from default to 64K in size. Testing was done on various AUS and no significant performance benefits were realized. Traditionally if smaller files will be stored on the file share a smaller AUS can be selected and vice versa. See www.microsoft.com for more information on what AUS works best.

Utilizing the Microsoft Windows 2012R2 File Server

Once the new drive has been created, it's time to share it and set properties on the new drive. Assign the appropriate users and the level of access each user or group will need. Figure 5 shows full control given to everyone, this is for example purposes and appropriate permissions will need to be given to users.

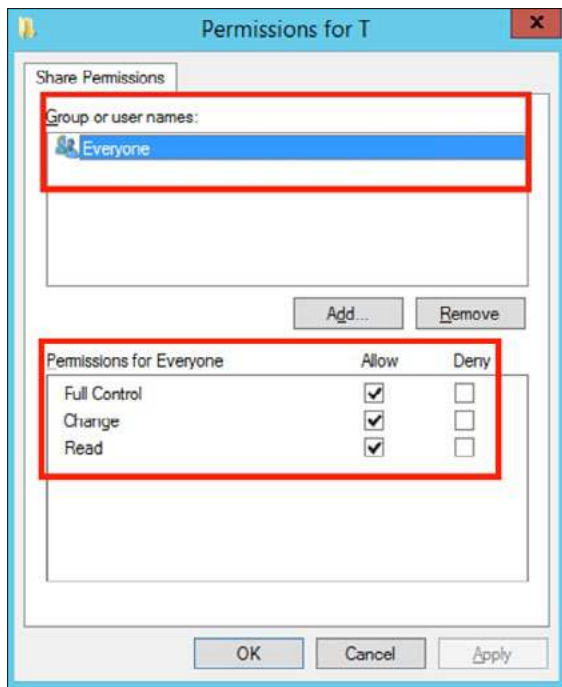


Figure 5 – Setting permissions on the new file share.

DO: Give appropriate permissions to users and groups who will use the new file share.

Host Side Compression

Microsoft Windows provides compression for the new file share, but similar to data deduplication, it is not recommended to enable Microsoft compression when using Tintri all flash and T800 arrays since compression is performed on the array as data is written to disk. Since the T600 and T500 arrays do not compress, host side compression can be enabled if enough resources exist to perform the actions.

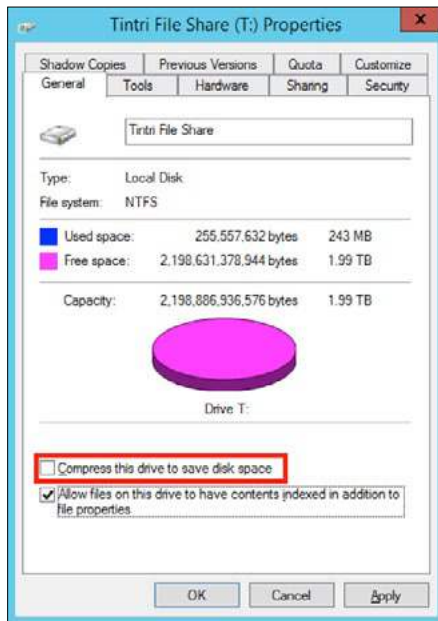


Figure 6 – Setting host side compression.

DO NOT: Enable host side compression unless a T600 or T500 is used and there are enough resources to accomplish the extra overhead.

Host Side Quotas

Microsoft File Shares allow for host side quotas to be set that are very useful in a file share environment. Quotas help prevent the abuse of a shared resource and gives the administrator power to assign specific amounts of drive space to all or specific users.

In Figure 7, a disk quota has been set to allow 20MB of data written to the file share. In the example below, a warning will be triggered at 10KB, but will allow users to continue to write data until they hit the 20MB limit. This quota applies to all users unless a per-user quota is created.

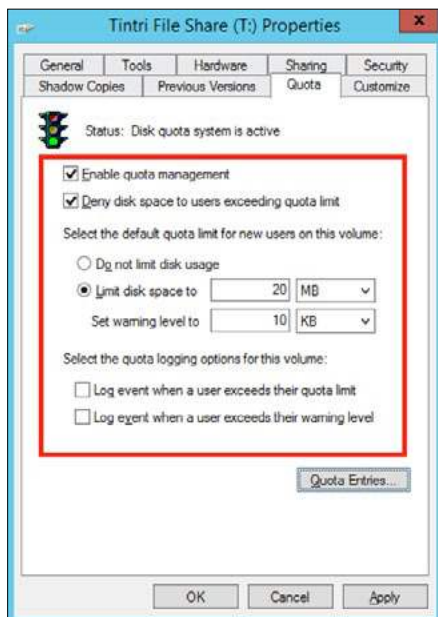
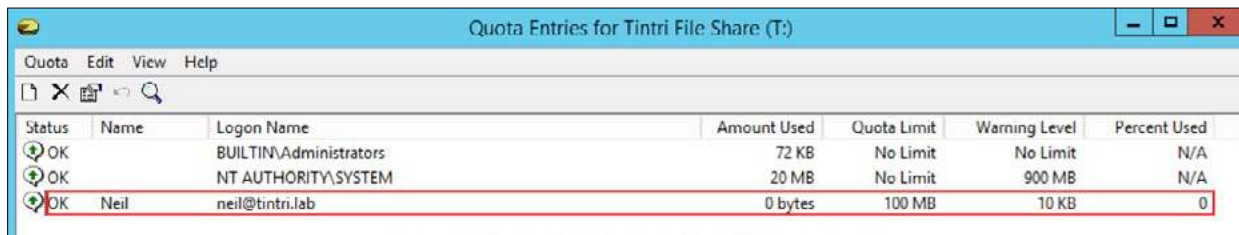


Figure 7 – Setting host side disk quotas.

In Figure 8, a per-user quota has been set on the share for a user named “Neil”. This user can write 100MB before no longer being able to add data to the share. Per-user quotas take priority over disk quotas and allow for further granularity.



Status	Name	Logon Name	Amount Used	Quota Limit	Warning Level	Percent Used
OK		BUILTIN\Administrators	72 KB	No Limit	No Limit	N/A
OK		NT AUTHORITY\SYSTEM	20 MB	No Limit	900 MB	N/A
OK	Neil	neil@tintri.lab	0 bytes	100 MB	10 KB	0

Figure 8 – Setting host side per-user quotas for individuals.

Protecting the File Share

Creating and using the file share is only part of the equation. Protecting user’s data is paramount and is incredibly simple to setup and utilize with Tintri snapshots and replication along with Microsoft Windows Shadow Copy.

Tintri Snapshots and Replication

Creating a snapshot schedule with Tintri is extremely simple and flexible. In the VMstore right click the virtual machine to protect and select “protect”. As seen in in Figure 9, Hourly, Daily, Weekly, Monthly and Quarterly snapshots can be selected as well as the amount and days of snapshots to keep.

Either crash or VM consistency can be chosen, but it is a best practice to utilize VM-consistent for backup purposes. Crash-consistent snapshots take a point in time copy while the virtual machine is active and no attempt is made to flush in-memory I/O. VM-consistent snapshots work in conjunction with VMware Snapshots to quiesce the virtual machine before the snapshot is taken. See the Tintri VMstore System Administration Manual, located at support.tintri.com, for more information.

DO: Use Tintri VM-consistent snapshots to ensure proper backup of user’s data.

If a second Tintri VMstore is available it is recommended to protect the file share by replicating the snapshots to a secondary VMstore in the event of a VMstore or site failure.

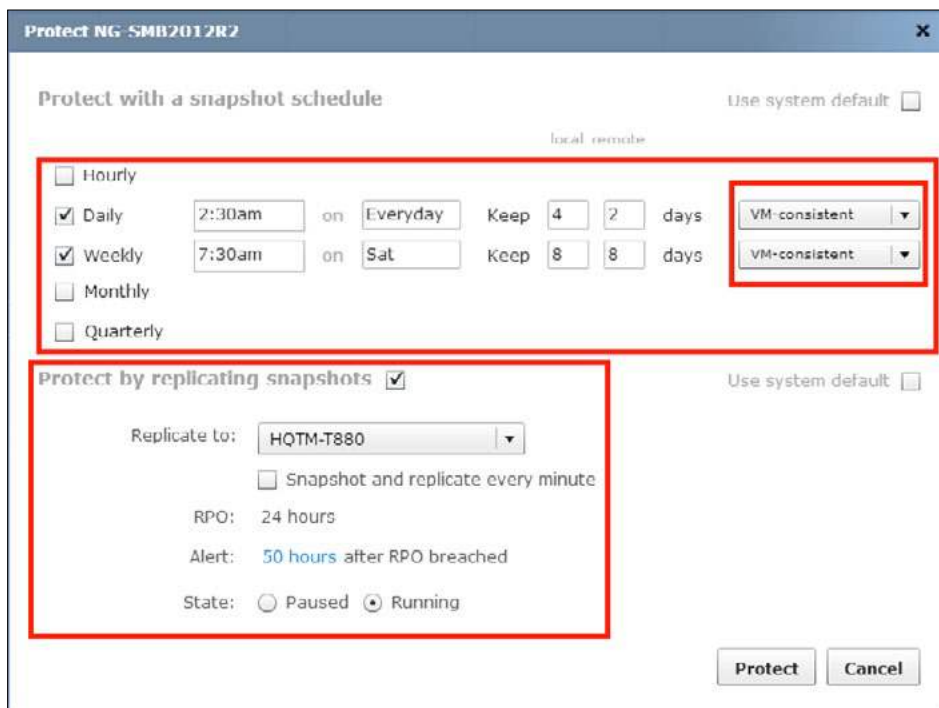


Figure 9 – Creating snapshot and replication schedules.

Microsoft Shadow Copies of Shared Folders

Microsoft Shadow Copies of Shared Folders in conjunction with Tintri snapshots and replication will help protect the file share from accidental deletions and enable users to restore files without contacting backup administrators. Shadow Copies of Shared Folders are set up on the machine hosting the share and provide point in time copies of files with the capability of restoring single files to entire folders.




Figure 10 – Microsoft Shadow Copies of Shared Folders.

The amount of space, location of the backups and schedule are all configurable, however the copies are not transportable, when allocated space for the copies has been reached the oldest shadow copy will be deleted to make room for new copies, a limit of 64 copies per volume can be kept, the copies are read-only and Shadow Copies for Shared folders can only be set per volume basis, as stated earlier in this document.

Restoring Shadow Copies for Shared Folders can occur in two ways, the volume can be explored for the desired file or the entire volume can be restored to a certain point in time. New files and file permissions on existing files will not be changed, but if updates have been made to files, a volume restoration will overwrite any changes made to files. For more information see www.microsoft.com.

File Server High Availability

Building upon protection of the file share, high availability not only ensures data protection, but data availability. Deploying VMware in a highly available cluster enables the use of VMotion, High Availability and Distributed Resource Scheduler (DRS).

VMware vMotion

VMware vMotion allows for the live migration of running virtual machines from one physical ESXi server to another, with no downtime. vMotion retains the VM's active memory, network identity and connections. Migrating VM's from a failing, over provisioned, or scheduled downtime of an ESXi server ensures simple and highly available behavior from the file share.

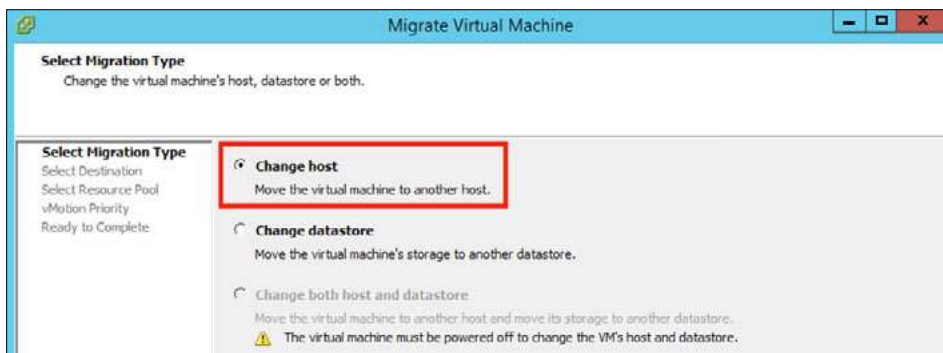


Figure 11 – Performing a VMware vMotion.

If there is a requirement to migrate data off one storage array and onto another, the Change datastore option can be selected and will utilize VMware Storage vMotion if the feature is available. Storage vMotion performs a live migration of a VM's disk files without disruption. Storage vMotion is beneficial, but care must be taken to understand the impact to the underlying storage's snapshots and protection capabilities.

VMware vSphere High Availability

VMware vSphere High Availability (HA) provides business continuance in a VMware cluster and allows VM's to automatically restart on alternate ESXi hosts in the cluster if a physical server goes down.

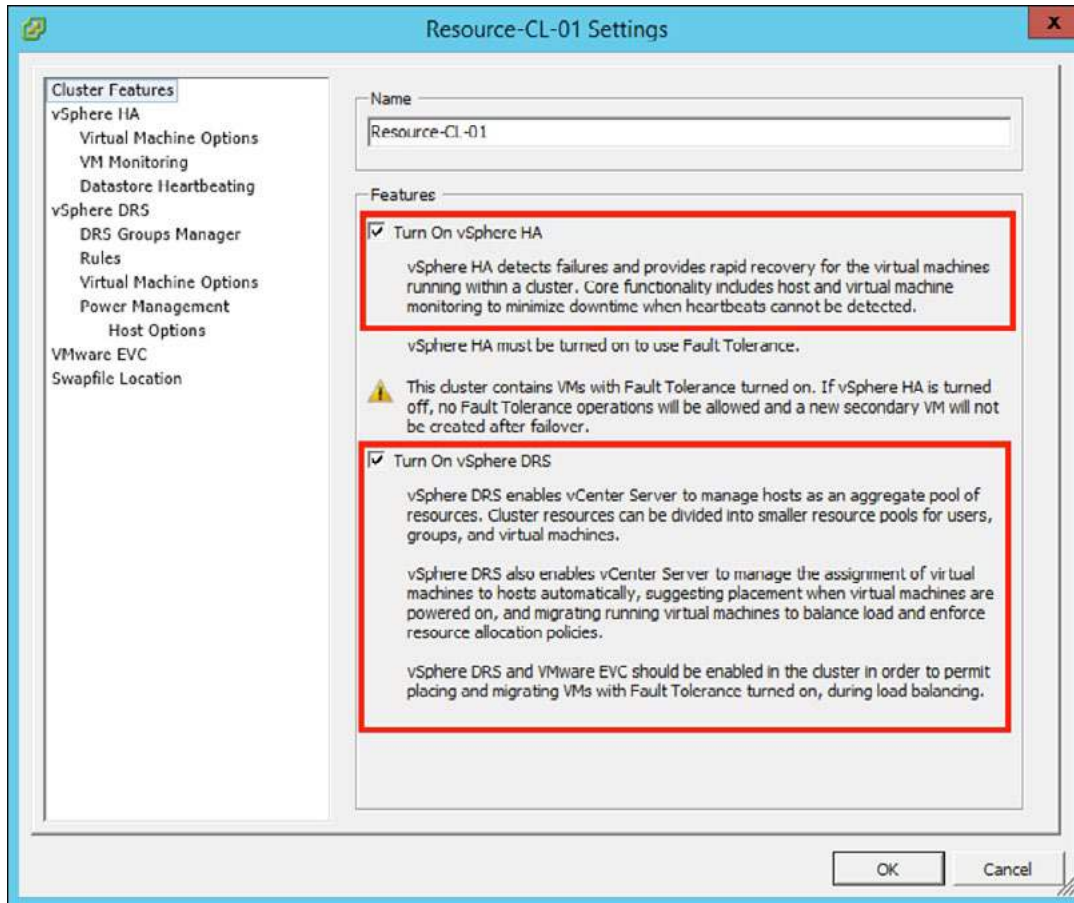


Figure 12 – Enabling vSphere HA and DRS.

VMware vSphere HA should be enabled if a cluster of ESXi hosts is available to the file share infrastructure. In the event hardware failure, it will ensure the recovery of the file share.

DO: Enable VMware vSphere HA in the virtualized file share environment.

VMware Distributed Resource Scheduler (DRS)

VMware DRS monitors CPU and memory of the ESXi servers in a VMware cluster. As resources are consumed, DRS can move virtual machines to other ESXi hosts to balance out resource consumption. The automation level is set depending on the amount of administrator interaction that is required. VMware DRS will ensure the file share is always on a healthy ESXi host with enough resources to provide users with the performance and accessibility they require.

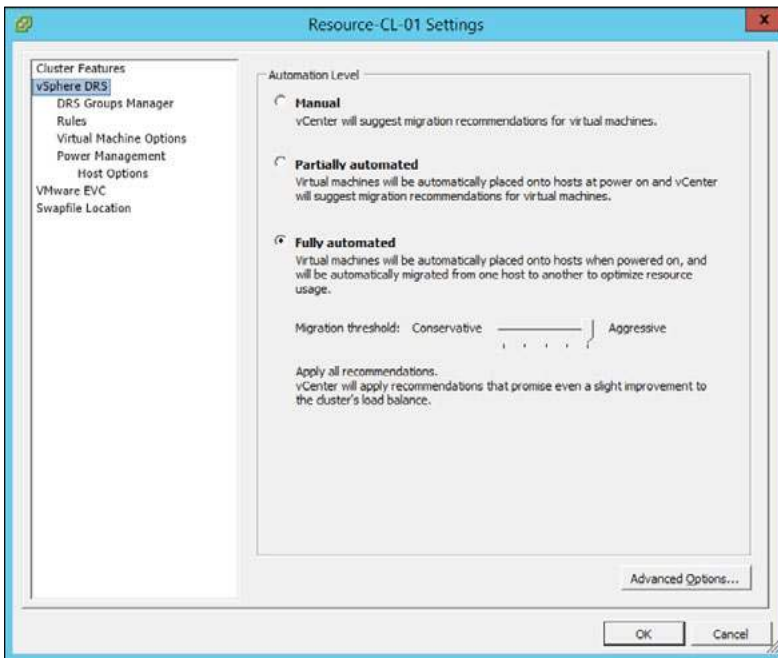


Figure 13 – Setting automation level in vSphere DRS.

DO: Enable VMware vSphere DRS as long as resources (memory and CPU) are available for VM's to migrate to.

VMware Site Recovery Manager (SRM)

If file share availability must survive a site outage, VMware SRM can be used. SRM allows for the complete automation of a disaster recovery of one vCenter to another. SRM allows for non-disruptive failover testing, automated orchestration workflows, recovery of network and security settings as well as custom automation that is necessary for the disaster recovery. If more information is required, see www.vmware.com.

Conclusion

Running SMB file shares using Microsoft Windows 2012R2 and VMware vSphere on Tintri is not only simple and reliable, but easy to setup and administer.

For more information see:

www.microsoft.com

www.vmware.com

www.tintri.com