Hewlett Packard
Enterprise

Best practices for HPE StoreEasy 1000

Setup and configuration

Technical white paper

Contents

Introduction

As you work with your large number of users, you need to accommodate their demands to store and access growing volumes of files such as business documents, images, audio files, and videos. You need efficient file storage that is secure, yet easy to manage.

Technologies used to address these challenges

The HPE StoreEasy 1000 Storage family is designed for small businesses, branch offices, and workgroup environments. The HPE StoreEasy 1000 servers offer various storage density solutions from 1U to 2U form factors. If you are running object storage, large data set management content delivery, or other data-intensive workloads on general purpose servers, the HPE StoreEasy 1650 E Storage server allows you to save valuable data center space through its unique density-optimized 2U form factor, which holds up to 28 LFF or 50 SFF hot plug drives.

The HPE StoreEasy Storage family caters to your unique storage needs. They integrate easily into new and existing environments offering your storage administrators and IT generalists a straightforward, consistent management experience.

Objective of this white paper

The objective of this white paper is to provide best practices for managing your HPE StoreEasy 1000 Storage server. The topics covered in this white paper include network configuration, storage configurations, data protection, security, monitoring, and performance tuning.

Note
Not all of Microsoft®'s documentation referenced have been updated to include Windows Server® 2016 as of the creation of the white paper, but that the information is still applicable.

The StoreEasy 1000 system has various performance options available. This table depicts best practices for increasing performance response and throughput.

Table 1. Performance

Increase performance	For more information
Increased throughput by aggregating network adapter ports	NIC teaming
Lower CPU utilization using Server Message Block (SMB) direct support of remote memory direct access (RMDA)-capable network adapters offloading network-related functions from the CPU	SMB 3.1.1 Multichannel
Tuning operating system default parameters to improve I/O performance	Tuning the operating system
Increasing number of users supported on the StoreEasy platform	Tuning the operating system

Best practices for fault tolerance

To increase fault tolerance and high availability of data, there are multiple tools and applications that can enable these capabilities. This table depicts best practices for increasing fault tolerance.

Table 2. Fault tolerance

Increase fault tolerance	For more information
Configure multiple network ports that provide redundant network routes utilizing network interface card (NIC) teaming	NIC teaming
Configure your storage pool and RAID configuration to protect from hard drive failures in your data storage	StoreEasy 1000 disk configuration

Best practices for networking configuration

Implementing fast and efficient network configurations will provide the following advantages:

- Increased throughput by aggregating network adapter ports
- Lower latency by decreasing the network response time
- Lower CPU utilization through the use of SMB direct, including support of RMDA-capable network adapters offloading network-related functions from the CPU

Network Configuration Tool

The StoreEasy Network Configuration Tool (NCT) enables the configuration of the network interfaces on HPE StoreEasy 1000 Storage. The use of NCT for network configuration is considered a best practice and Hewlett Packard Enterprise encourages the use of this tool to configure network interfaces.

- NCT analyzes all available network interfaces of the StoreEasy system allowing the selection of different network configurations and implements these changes, which reduces the chances of incorrect configuration.
- NCT also validates the configuration after changes are implemented to verify that the changes in the network environment are compatible with the network environment. This validation can help with troubleshooting errors in the networking environment.
- Use NCT to perform the following tasks:
  - Change the network interface settings and configure network teams
  - Configure VLAN assignments and assign VLAN IDs
  - Configure the IP address for the selected interface
  - Confirm network settings and diagnose environmental network issues using the network validation system

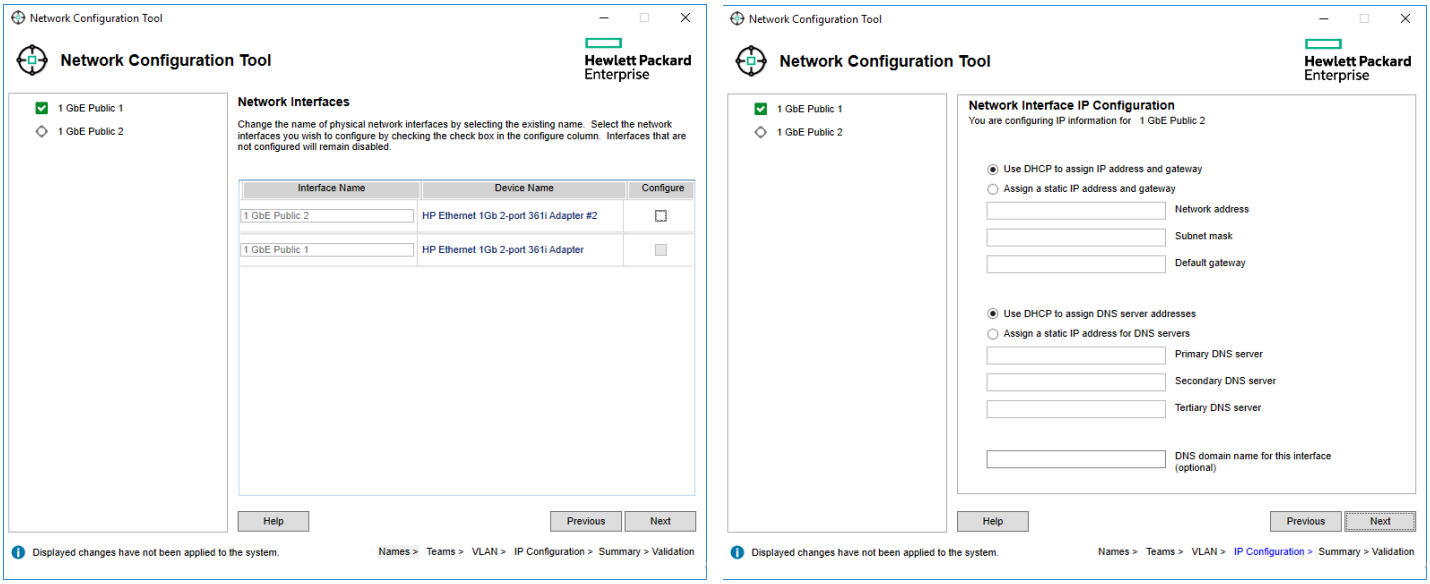

Figure 1. Network Configuration Tool

NIC teaming

NIC teaming, also known as load balancing and failover (LBFO), enables bandwidth aggregation and traffic failover in the event of a network component failure. Implementing a network configuration that provides high availability of network connections to HPE StoreEasy 1000 requires that network ports on different network adapters be members of NIC teams. Beginning with Windows Server 2012, the ability to implement NIC teaming is included in the operating system; therefore, Hewlett Packard Enterprise no longer has a separate application to enable this capability.

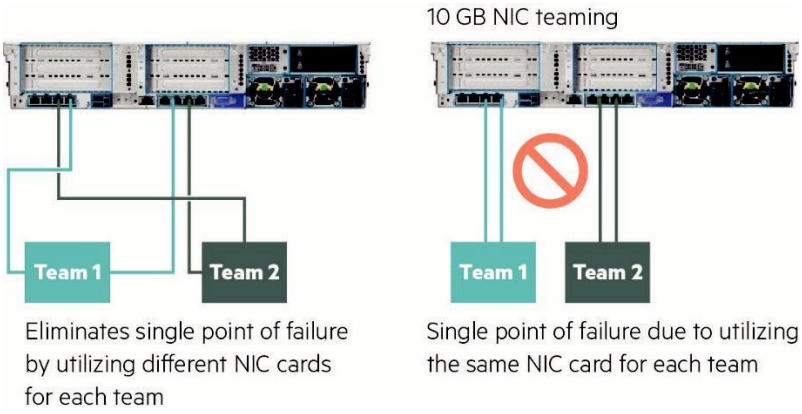

Figure 2. 10 GB NIC teaming

SMB 3.1.1 Multichannel

Simplified SMB Multichannel and Multi-NIC Cluster Networks is a new feature in Windows Server 2016 that enables the use of multiple NICs on the same cluster network subnet, and automatically enables SMB Multichannel.

Simplified SMB Multichannel and Multi-NIC Cluster Networks provide the following benefits:

- Failover Clustering automatically recognizes all NICs on nodes that are using the same switch/same subnet—no additional configuration needed.
- SMB Multichannel is enabled automatically.
- Networks that only have IPv6 Link Local (fe80) IP Addresses resources are recognized on cluster-only (private) networks.
- A single IP Address resource is configured on each Cluster Access Point (CAP) Network Name (NN) by default.
- Cluster validation no longer issues warning messages when multiple NICs are found on the same subnet.

Network implementation best practices

When configuring the networking on the HPE StoreEasy 1000 node, it is a best practice to:

- Remove network bottlenecks by implementing multiple network segments utilizing multiple network ports or by implementing NIC teaming
- Disable unused network ports
- Offload CPU processor load by modifying the RSS Profile for 10GbE network interfaces from NUMAStatic (default) to NUMA. Implement NUMA by issuing this PowerShell command, '`Set-NetAdapterRss -Name "<10Gb Adapters>" -Profile NUMA`'

Best practices for storage configurations

Built on Microsoft Windows® Storage Server 2016, you can provide users with a highly available solution so that their data is ready when they need it. Refer to the HPE StoreEasy 1000 Administrator Guide and Quick Start Guide. After the initial configuration has been completed, the rest of the storage configuration steps can be accomplished through the Windows 2016 Server Manager.

Storage Quality of Service

Storage Quality of Service (QoS) in Windows Server 2016 provides a way to centrally monitor and manage storage performance for virtual machines using Hyper-V and the Scale-Out File Server roles. The feature automatically improves storage resource fairness between multiple virtual machines using the same file server cluster and allows policy-based minimum and maximum performance goals to be configured in units of normalized IOPS.

Storage QoS in Windows Server 2016 can be used to accomplish the following:

- **Mitigate noisy neighbor issues.** By default, Storage QoS ensures that a single virtual machine cannot consume all storage resources and starve other virtual machines of storage bandwidth.
- **Monitor end-to-end storage performance.** As soon as virtual machines stored on a Scale-Out File Server are started, their performance is monitored. Performance details of all running virtual machines and the configuration of the Scale-Out File Server cluster can be viewed from a single location.
- **Manage Storage I/O per workload business needs.** Storage QoS policies define performance minimums and maximums for virtual machines and ensure that they are met. This provides consistent performance to virtual machines, even in dense and overprovisioned environments. If policies cannot be met, alerts are available to track when VMs are out of policy or have invalid policies assigned.

Storage QoS requirements

Storage QoS supports two deployment scenarios:

- **Hyper-V using a Scale-Out File Server.** This scenario requires both of the following:
  - Storage cluster that is a Scale-Out File Server cluster
  - Compute cluster that has least one server with the Hyper-V role enabled

For Storage QoS, the Failover Cluster is required on Storage servers, but the compute servers are not required to be in a failover cluster. All servers (used for both Storage and Compute) must be running Windows Server 2016.

- **Hyper-V using Cluster Shared Volumes.** This scenario requires both of the following:
  - Compute cluster with the Hyper-V role enabled
  - Hyper-V using Cluster Shared Volumes (CSV) for storage

Failover Cluster is required. All servers must be running the same version of Windows Server 2016.

Data deduplication

Data deduplication, often called Dedup for short, is a feature of Windows Server 2016 that can help reduce the impact of redundant data on storage costs. When enabled, data deduplication optimizes free space on a volume by examining the data on the volume by looking for duplicated portions on the volume. Duplicated portions of the volume's data set are stored once and are (optionally) compressed for additional savings. Data deduplication optimizes redundancies without compromising data fidelity or integrity.

Data deduplication helps storage administrators reduce costs that are associated with duplicated data. Large data sets often have a lot of duplication, which increases the costs of storing the data. For example:

- User file shares may have many copies of the same or similar files.
- Virtualization guests might be almost identical from VM-to-VM.
- Backup snapshots might have minor differences from day to day.

Data deduplication can be used in:

- General purpose file servers
- Virtualized Desktop Infrastructure (VDI) deployments
- Backup targets, such as virtualized backup applications
- Other workloads

Configuring the storage system

The initial configuration task (ICT) presented when the system is initially started up is provided for the purpose of walking a user through those steps that should be taken to properly implement the system. The ICT window launches automatically at logon for any user who is a member of the local administrators group. Use ICT to perform the task of implementing:

- System settings
- Networking
- Notifications
- Storage configuration
- Protect this server

Initial configuration best practices

- Review and implement the steps identified in the Quick Start Guide that comes with your StoreEasy 1000 system.
- Once the HPE StoreEasy 1000 Storage is connected to your network and to external storage enclosures (if present), powered up, and logged on, complete all the ICT tasks stepwise for the system to be configured.
- The ICT application should only be used for initial configuration and once all tasks are completed, select the "Do not show this window at the next logon" check box in the ICT window.

StoreEasy 1000 disk configuration

It is recommended that the StoreEasy Pool Manager be used to configure your disks instead of using the Smart Array Configuration Utility. Pool Manager analyzes the configuration of your StoreEasy Storage system and guides you to choose a disk configuration that reduces the chance of data loss or unavailability.

When choosing a RAID level, there is basically a tradeoff among storage space efficiency, availability of data in the event of drive failure, and read/write file access performance. Select the appropriate RAID level based on your requirements. For example, if some files in your document folder are more critical than other files, choose a higher RAID level, such as RAID 1+0, which increases the chances of data retention in case of multiple drive failures.

Using the StoreEasy Pool Manager, you are presented with various possible configurations, which include several configurations that automatically add spare drive(s).

For detailed information on configuration steps, refer to HPE StoreEasy 1000 Administrator Guide.

Pools should follow these best practices

- Selection of pool configurations with a greater number of pools provides better availability. A disk failure would impact only the respective volume rather than all volumes.
- Storage pools should contain one logical unit number (LUN); multiple LUNs in a pool may suffer performance degradation. When multiple LUNs are contained within a storage pool, any operations performed on a LUN will affect the other LUNs in that storage pool until the specific operation has completed.
- Midline (MDL) or Nearline SAS drives of 2 TB or greater should use RAID 6; the added redundancy of RAID 6 will protect against data loss in the event of a second drive failure with large MDL SAS drives.
- When using RAID 5 with an MDL SAS drive, it is recommended that you assign a spare drive to the RAID 5 group per enclosure.
- RAID groups should not contain drives from multiple enclosures. This also applies to spare drive assignment to the RAID group.
- It is recommended that you use RAID 1 when configuring virtual machine (VM), with one RAID 1 LUN per VM.
- Consider using RAID 10 (RAID 1+0) for applications that have a very high write ratio (over 50% of the access rate).
- RAID 50 (RAID 5+0) should be considered for FC or serial-attached SCSI (SAS) providing the highest performance-to-capacity ratio.
- All drives in a single enclosure should be the same type of drive (all SAS hard drives, or all SATA hard drives or all Solid State Drives).
- To take advantage of the speed of the 12 GB SAS backplane, use hard drives that are rated for 12 GB.

Tuning the operating system

The Windows Storage Server 2016 operating system default settings should be tuned based upon your environment and use case.

Tuning of SMB can increase file server performance in many cases. To tune your server properly, you should run a performance test over a period of time to capture peak and steady state workload of the server. The following registry settings listed should specifically reviewed to see where you could possibly remove bottlenecks.

- **Max Threads Per Queue**—The default setting is 20; increase this value based upon your server hardware. For the StoreEasy 1650, StoreEasy 1650 E, and StoreEasy 1850, it is recommended to increase this value up to 64. One indication that you need to increase this value is if the SMB2 work queues are growing very large by reviewing the performance counter for "Server Work Queues\Queue Length\SMB2 NonBlocking."
- **Additional Critical Worker Threads**—The default setting is 0. By increasing the number of threads the file system cache uses for read-ahead and write-behind requests will improve I/O performance, especially on systems with multiple processors. If you have multiple processors, recommend changing this up to 64. Even if you do not have multiple processors, you may see value of increasing this value. You can determine if you need to increase the value by looking at the performance counter for "Cache\Dirty Pages" and if this value is growing to consume a large portion (over ~25%) of memory or if the system is doing a lot of synchronous I/Os, then increase the value.
- **SMB2 Credits minimum and maximum settings**—The defaults are 512 and 8192, respectively. Some clients, especially Windows 7 clients, can benefit by increasing these parameters when copying files over a high-bandwidth, high-latency link. Monitor the "SMB Client Shares/Credit Stalls/Sec" to see if there are any issues with credits and adjust appropriately.

As a best practice, it is recommended that you periodically run a performance report to look at the bottlenecks. In reviewing the performance results, you may be able to increase throughput and or increase the number of simultaneous users on your StoreEasy 1000 system by increasing the number of processors, the amount of memory installed per processor, or by implementing 10GbE networking.

Best practices for data protection

As file storage data grows or as file storage consolidation from other storage mediums are leveraged into the StoreEasy solutions that have been implemented, it becomes increasingly more important to be able to recover or restore the data.

Snapshots

Volume Shadow Copy Service

The ability to perform accurate and efficient backup and restore operations requires close coordination of backup applications, business applications, and the storage management hardware and hardware-related software. Volume Shadow Copy Service (snapshots) provides this coordination of actions that are required to create a consistent shadow (also known as snapshot or point-in-time copy) of the data that is to be backed up. Here are some typical VSS scenarios:

- Back up application data and system state information, including archiving data to another hard disk drive, to tape, or to other removable media
- Efficiently perform disk-to-disk backups

Require fast recovery from data loss by restoring data to the original LUN or to an entirely new LUN that replaces an original LUN that failed.

Windows features and applications that use VSS include Windows Server Backup, Shadow Copies, System Center Data Protection Manager, and System Restore.

Plan your backup strategy carefully when backing up virtual volumes utilizing VSS because of the fact that VSS does not support creating a Shadow Copy of a virtual volume and the host volume in the same snapshot set. VSS does support creating snapshots of volumes on a virtual hard disk (VHD) if backup of the virtual volume is necessary.

When configuring VSS to make Shadow Copies of Shared Folders, the following best practices apply:

- Locate the snapshot copy of a volume on a different volume so that each volume has its own snapshot volume.
- The Diff Area storage is the storage space that Shadow Copies of Shared Folders allocates on a volume for maintaining the snapshots of the contents of shared folders. Instead, the storage volume and the original volume need to be the same volume, or they need to be on separate physical disks.
- Adjust the Shadow Copy schedule to fit the work patterns of your clients.
- Do not enable Shadow Copies on volumes that use mount points. You should explicitly include the mounted volume in the schedule for Shadow Copy creation (for previous versions of a file to be available, the volume must have a drive letter assigned).
- Perform regular backups of your file server; snapshots of shared folders are not a replacement for performing regular backups.
- Do not schedule copies to occur more often than once per hour. The default schedule for creating Shadow Copies is at 7:00 a.m., Monday through Friday. If you decide that you need copies to be created more often, verify that you have allotted enough storage space and that you do not create copies so often that server performance degrades. There is also an upper limit of 64 copies per volume that can be stored before the oldest copy is deleted. If Shadow Copies are created too often, this limit might be reached very quickly, and older copies could be lost at a rapid rate.
- Before deleting a volume that is being Shadow Copied, delete the scheduled task for creating Shadow Copies. If the snapshots are no longer needed, then the snapshot volume can also be deleted.
- Use an allocation unit size of 16 KB or larger when formatting a source volume on which Shadow Copies of Shared Folders will be enabled if you plan to defragment the source volume on which Shadow Copies of Shared Folders is enabled. If you do not, the number of changes caused by defragmentation can cause previous versions of files to be deleted. Also, if you require NT File System (NTFS) file compression on the source volume, you cannot use an allocation unit size larger than 4 KB. In this case, when you defragment a volume that is very fragmented, you may lose older Shadow Copies faster than expected.

After backing up a volume that contains Shadow Copies, do not restore the volume to a different volume on the same computer. Doing this would cause multiple snapshots with the same snapshot ID on the system and will cause unpredictable results (including data loss) when performing a Shadow Copy revert.

Replication

Windows DFS Replication and Namespaces

DFS Namespaces and DFS Replication are role services in the file and storage services role. DFS Namespaces (DFS-N) enables you to group shared folders that are located on the same or different servers into one or more logically structured namespaces. Windows DFS Replication (DFS-R) enables efficient replication of folders across multiple servers and multiple locations. There was new functionality implemented in Windows 2012 R2, which includes the ability to support data deduplication volumes. The following list of the current capabilities were included for DFS-R in the Windows 2012 R2 release and that are still available in Windows Server 2016:

- New Windows PowerShell commandlets for performing the majority of administrative tasks for DFS-R. Administrators can use the extensive Windows PowerShell commandlets to perform common administrative tasks, and optionally automate them by using Windows PowerShell scripts. These tasks include operational actions such as creating, modifying, and removing replication settings. New functionality is also included with the commandlets, such as the ability to clone DFS Replication databases and restore preserved files.
- New Windows Management Infrastructure (WMI)-based methods for managing DFS Replication were in implemented beginning with Windows Server 2012 R2 includes new Windows Management Infrastructure (sometimes referred to as WMI v2) provider functionality, which provides programmatic access to manage DFS Replication. The new WMI v1 namespace is still available for backwards compatibility.
- The ability to perform database cloning for initial sync was introduced, providing the ability to bypass initial replications when creating new replicated folders, replacing servers, or recovering from a database. Now the “Export-DfsrClone” commandlet allows you to export the DFS Replication database and volume configuration .xml file settings for a given volume from the local computer to clone that database on another computer. Running the commandlet triggers the export in the DFS Replication service and then waits for the service to complete the operation. During the export, DFS Replication stores file metadata in the database as part of the validation process. After you press the data and copy the exported database and .xml file to the destination server, you use “Import-DfsrClone” to import the database to a volume and validate the files in the file system. Any files that perfectly match don’t require expensive inter-server metadata and file synchronization, which leads to dramatic performance improvements during the initial sync.

- Added the ability to support the rebuilding of corrupt databases without unexpected data loss caused by non-authoritative initial sync. Beginning in Windows Server 2012 R2, when DFS Replication detects database corruption, it rebuilds the database by using local file and update sequence number (USN) change journal information, and then marks each file with a "Normal" replicated state. DFS Replication then contacts its partner servers and merges the changes, which allows the last writer to save the most recent changes as if this was normal ongoing replication.
- The ability to provide the option to disable cross-file remote differential compression (RDC) between servers. In Windows Server 2012 R2 or Windows Server 2016, DFS Replication allows you to choose whether to use cross-file RDC on a per-connection basis between partners. Disabling cross-file RDC might increase performance at the cost of higher bandwidth usage.
- Includes the ability to configure variable file staging sizes on individual servers. DFS Replication now allows you to configure the staging minimum size from as little as 256 KB to as large as 512 TB. When you are not using RDC or staging, files are no longer compressed or copied to the staging folder, which can increase performance at the cost of much higher bandwidth usage.
- Provides the capability to restore files from the "ConflictAndDeleted" and "PreExisting" folders. DFS Replication now enables you to inventory and retrieve the conflicted, deleted, and pre-existing files by using the "Get-DfsrPreservedFiles" and "Restore-DfsrPreservedFiles" commandlets. You can restore these files and folders to their previous location or to a new location. You can choose to move or copy the files, and you can keep all versions of a file or only the latest version.
- Updated capability to enable automatic recovery after a loss of power or an unexpected stoppage of the DFS Replication service. In Windows Server 2012 and Windows Server 2008 R2, the default behavior when an unexpected shutdown happened required you to re-enable replication manually by using a WMI method. Beginning with Windows Server 2012 R2, it defaults to triggering the automatic recovery process. You must opt out of this behavior by using the registry value. In addition, if the only replicated folder on a volume is the built-in "SYSVOL" folder of a domain controller, it automatically triggers recovery regardless of the registry setting.
- In Windows Server 2012 and earlier operating systems, disabling a membership immediately deleted the "DfsrPrivate" folder for that membership, including the "Staging," "ConflictAndDeleted," and "PreExisting" folders. After these folders are deleted, you can't easily recover data from them without reverting to a backup. Beginning with Windows 2012 R2, DFS Replication now leaves the DfsrPrivate folder untouched when you disable a membership. You can recover conflicted, deleted, and pre-existing files from that location if the membership is not re-enabled (Enabling the membership deletes the content of all private folders).

The following list provides a set of scalability guidelines for DFS-R that have been tested by Microsoft on Windows Server 2012 R2 or Windows Server 2016:

- Size of all replicated files on a server: 100 terabytes
- Number of replicated files on a volume: 70 million
- Maximum file size: 250 gigabytes

The following is the published tested DFS Replication limits; this information is from the blog on Understanding DFS Replication "limits":

- Each server can be a member of up to 256 replication groups.
- Each replication group can contain up to 256 replicated folders.
- Each server can have up to 256 connections (for example, 128 incoming connections and 128 outgoing connections).
- On each server, the result of the following formula should be kept to 1024 or fewer: (number of replicated folders in replication groupx x number of simultaneously replicating connections in replication groupx) + (number of replicated folders in replication groupy x number of simultaneously replicating connections in replication groupy) + (number of replicated folders in replication groupn x number of simultaneously replicating connections in replication groupn)
- A replication group can be arbitrarily large, scaling to several thousands of members. However, each member can be connected to, at most, 256 partners. (See our blog post for more about this 256-member recommendation.)
- A volume can contain up to 8 million replicated files, and a server can contain up to 1 terabyte of replicated files. (See Microsoft's blog post for more about this 1 TB recommendation.)

DFS-R relies on Active Directory Domain Services for configuration; it will only work in a domain. DFS-R will not work with workgroups. Extend the Active Directory Domains Services (AD DS) schema to include Windows Server 2016, Windows Server 2012 R2, Windows 2012, Windows Server 2008 R2, Windows Server 2008 schemas. Ensure that all servers in a replication group are located in the same forest. You cannot enable replication across servers in different forests.

- DFS-R can replicate numerous folders between servers. Ensure that each of the replicated folders has a unique root path and that they do not overlap. For example, D:\Sales and D:\Accounting can be the root paths for two replicated folders, but D:\Sales and D:\Sales\Reports cannot be the root paths for two replicated folders.
- For DFS-R, locate any folders that you want to replicate on volumes formatted with the NTFS. DFS-R does not support the Resilient File System (ReFS) or the file allocation table (FAT) file system. DFS Replication also does not support replicating content stored on Cluster Shared Volumes.
- DFS-R does not explicitly require time synchronization between servers. However, DFS Replication does require that the server clocks match closely. The server clocks must be set within five minutes of each other (by default) for Kerberos authentication to function properly. For example, DFS Replication uses time stamps to determine which file takes precedence in the event of a conflict. Accurate times are also important for garbage collection, schedules, and other features.
- Do not configure file system policies on replicated folders. The file system policy reapplies NTFS permissions at every “Group Policy” refresh interval.
- DFS-R cannot be used to replicate mailboxes hosted on Microsoft Exchange Server.
- DFS-R can safely replicate Microsoft Outlook personal folder files (.pst) and Microsoft Access files only if they are stored for archival purposes and are not accessed across the network utilizing a client such as Outlook or Access.

Backup and restore

It is always a best practice to have a backup solution that backs up your data from the storage arrays. This will provide assurance that you will be able to recover your data in the event of a catastrophic failure or due to file corruption due to user error, application problems, or database issues.

When planning and implementing your backup and restore process, take note of the following:

- For a backup to succeed, in other words, the node must be running.
- Before putting a system into production, test your backup and recovery process.
- When you perform a backup (using Windows Server Backup or other backup software), choose options that will allow you to perform a system recovery from your backup. For more information, see help or other documentation for your backup software.
- When backing up data on disks, notice which disks are online on that node at that time. Only disks that are online and owned by the node at the time of the backup are backed up.
- When restoring data to a disk, notice which disks are online on that node at that time. Data can be written only to disks that are online and owned by the node when the backup is being restored.
- The HPE StoreEasy 1000 server operating system should be backed up before and after any significant changes in configuration. Perform system state backups of each of your HPE StoreEasy 1000 servers. System state backups, also called authoritative system restore (ASR) capable backups.
- Review backup logs routinely to ensure that data is being properly backed up.
- Monitor network reports for errors that may inhibit or cause congestion on the network, which can create backup errors.
- It is recommended to have at least 1.5 times in storage capacity of the items that you want to back up.
- Rotate backup media between onsite and offsite locations on regular intervals. When backing up shared volumes, refer to the information in this link, technet.microsoft.com/library/jj612868.aspx.

Antivirus

Antivirus products can interfere with the performance of a server.

Antivirus software that is not cluster-aware may cause problems with Cluster Services. For more details, visit this link: support.microsoft.com/kb/250355.

It is recommended that you disable scanning of DFS sources. To review further details about this, go to support.microsoft.com/kb/822158. Antivirus software may create issues when dealing with virtual machines; this link will provide further details—support.microsoft.com/kb/961804.

Best practices for security

User access

User access should be managed through Active Directory and not local accounts. Limit user access to shares, folders, or files by implementing domain-level group permissions.

Create a different user for each system administrator that will use the system. Alternatively, configure the system to use Active Directory and make sure all users use their own accounts to log into the system.

When scripting, use the lowest privilege level required. If a script requires only read access to the system, use a "Browse" account. If a script does not need to remove objects, use a "Create" account.

Encryption

Windows BitLocker

Windows BitLocker Drive Encryption is a data protection feature of the operating system. Having BitLocker integrated with the operating system addresses the threats of data theft or exposure from lost, stolen, or inappropriately decommissioned computers. When installing the BitLocker optional component on a server, you will also need to install the "Enhanced Storage" feature. On servers, there is also an additional BitLocker feature that can be installed—BitLocker Network Unlock.

The recommended practice for BitLocker configuration on an operating system drive is to implement BitLocker on a computer with a Trusted Platform Module (TPM) version 1.2 or 2.0 and a Trusted Computing Group (TCG)-compliant BIOS or Unified Extensible Firmware Interface (UEFI) firmware implementation, plus a personal identification number (PIN). By requiring a PIN that was set by the user in addition to the TPM validation, a malicious user that has physical access to the computer cannot simply start the computer.

Encrypting File System

Windows Encrypting File System (EFS) includes the ability to encrypt data directly on volumes that use the NTFS so that no other user can access the data. Encryption attributes can be managed through folder or file objects' properties dialog box. EFS encryption keys are stored in Active Directory objects on the domain controller with the user account that encrypted the files. Some EFS guidelines are:

- Teach users to export their certificates and private keys to removable media and store the media securely when it is not in use. For the greatest possible security, the private key must be removed from the computer whenever the computer is not in use. This protects against attackers who physically obtain the computer and try to access the private key. When the encrypted files must be accessed, the private key can easily be imported from the removable media.
- Encrypt the "My Documents" folder for all users (User_profile\My Documents). This makes sure that the personal folder, where most documents are stored, is encrypted by default.
- Teach users to never encrypt individual files but to encrypt folders. Programs work on files in various ways. Encrypting files consistently at the folder level makes sure that files are not unexpectedly decrypted.
- The private keys that are associated with recovery certificates are extremely sensitive. These keys must be generated either on a computer that is physically secured, or their certificates must be exported to a .pfx file, protected with a strong password, and saved on a disk that is stored in a physically secure location.
- Recovery agent certificates must be assigned to special recovery agent accounts that are not used for any other purpose.
- Do not destroy recovery certificates or private keys when recovery agents are changed (Agents are changed periodically). Keep them all, until all files that may have been encrypted with them are updated.

- Designate two or more recovery agent accounts per organizational unit (OU), depending on the size of the OU. Designate two or more computers for recovery, one for each designated recovery agent account. Grant permissions to appropriate administrators to use the recovery agent accounts. It is a good idea to have two recovery agent accounts to provide redundancy for file recovery. Having two computers that hold these keys provides more redundancy to allow recovery of lost data.
- Implement a recovery agent archive program to make sure that encrypted files can be recovered by using recovery keys. Recovery certificates and private keys must be exported and stored in a controlled and secure manner. Ideally, as with all secure data, archives must be stored in a controlled access vault, and you must have two archives: a master and a backup. The master is kept onsite, while the backup is located in a secure, offsite location.
- Avoid using print spool files in your print server architecture, or make sure that print spool files are generated in an encrypted folder.

The Encrypting File System does take some CPU overhead every time a user encrypts and decrypts a file. Plan your server usage wisely. Load balance your servers when there are many clients using Encrypting File System.

HPE StoreEasy 1000 data encryption

HPE StoreEasy 1000 data encryption prevents data exposure that might result from the loss of physical control of disk drives. This solution uses self-encrypting drive (SED) technology to encrypt all data on the physical drives and prevent unauthorized access to data-at-rest (DAR). When encryption is enabled, the SED will lock when power is removed, and it will not be unlocked until the matching key from the HPE StoreEasy 1000 systems is used to unlock it.

Keep the encryption key file and password safe. If you lose the encryption key and the HPE StoreEasy 1000 systems are still functioning, you can always perform another backup of the encryption key file. However, should you lose the encryption key file or the password, and should the HPE StoreEasy 1000 system then fail, the HPE StoreEasy 1000 system will be unable to restore access to data. Ensure that backup copies of the latest encryption key file are retained and that the password is known.

The importance of keeping the encryption key file and password safe cannot be overstated. Hewlett Packard Enterprise does not have access to the encryption key or password.

Different arrays need separate backups, although the same password can be applied.

The SED datastore provides an open interface for authentication key management. The SED datastore tracks the serial number of the array that owns each SED, which disallows SEDs from being used in other systems.

Best practices for system monitoring

HPE System Management Homepage

System Management Homepage (SMH) is a web-based interface that consolidates and simplifies single system management for HPE servers. The SMH is the primary tool for identifying and troubleshooting hardware issues in the storage system. You may select this option to diagnose a suspected hardware problem. Go to the SMH main page and open the “Overall System Health Status” and the “Component Status Summary” sections to review the status of the storage system hardware.

Windows performance monitoring

It is recommended to have monitoring running on the node. This monitoring should be set up so all the relevant counters are logging information to a binary file. Circular logging should also be used to overwrite previous data. This also ensures the captured data file does not grow too large and therefore does not eat up disk space. This information is useful when a performance issue occurs; rather than waiting for performance data to be gathered once an issue has been raised, the performance data is already in hand.

HPE Notification Tool

The HPE Notification Tool eliminates the necessity to manually login into the StoreEasy server and check various system parameters like free space, unallocated space, used space, etc. The Notification Tool will provide email alert based on how you have configured the tool. Reference the HPE StoreEasy 1000 Administrator Guide for details on configuring the Notification Tool.

Summary

The HPE StoreEasy 1000 offers a tightly integrated, converged solution when paired with HPE storage. By following guidelines and best practices as outlined in this technical white paper, you can increase the efficiency and performance of your HPE StoreEasy 1000 converged solution.

Resources

Implementing Microsoft's NIC teaming

SMB Multichannel

Performance Tuning Guidelines for Windows Server 2012 R2

Failover Clustering in Windows Server 2016

Double Take Availability

Storage Replica overview

What's new in Failover Clustering in Windows Server 2016

Storage QoS overview

Simplified SMB Multichannel and Multi-NIC Cluster Networks

Learn more at

hpe.com/us/en/product-catalog/storage/file-and-object-storage/pip.file-and-object-storage.5335825.html

Sign up for updates

4AA6-3618ENW, January 2017, Rev. 1

Hewlett Packard Enterprise