Technical Report

Windows File Services Best Practices with NetApp Storage Systems

Brahmanna Chowdary Kodavali, Reena Gupta, NetApp
August 2011 | TR-3771

Abstract

This document describes Windows® File Services best practices and recommendations for using a NetApp® storage system in a Windows file sharing environment. It also provides basic guidelines to build a storage infrastructure for Windows File Services. It describes the typical deployment-related best practices used by system administrators and architects for integration with Microsoft® Active Directory® and services, securing and optimizing NetApp storage systems using CIFS or SMB 2.0 protocol. The scope of this guide is limited to Data ONTAP® 7G / 7-Mode releases.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

The Microsoft Common Internet File System (CIFS) protocol is natively integrated into NetApp Data ONTAP. As a result, Windows 2008 R2, Windows 2008, Windows 2003, Windows 2000, Windows 7, Windows Vista®, and Windows XP computers do not require additional client software to access data on NetApp storage systems. In a Windows file sharing environment, NetApp storage systems appear on the network as native file servers. NetApp storage systems running Data ONTAP 7.2.4 (7G) and later or 8.0.1 (7-Mode) include support for the Microsoft Windows 2008 R2 Active Directory (AD) and Windows 7 clients. NetApp systems can be installed into Windows 2008 R2, Windows 2008, 2003, or 2000 "mixed-mode" or "native-mode" AD domains. Beginning with Data ONTAP 7.3.1 (7G) and 8.0.1 (7-Mode), the new version of the CIFS protocol, SMB 2.0, is also supported.

2 Audience

This document is targeted to the technical audience—system administrators, architects, system engineers, application vendors, and so on—who design Windows file services solutions, who deploy and implement these solutions, and who maintain and administer the NetApp storage system in that environment.

3 Infrastructure Layout

Customers' environments can be divided according to their company size or number of employees into three categories: small environments, medium-sized environments, and large enterprise environments. The following sections describe these environments and offer guidelines for building a tiered storage infrastructure.

Small Environments

These environments:

- May be classified under the tier-3 storage model
- May have tier-1 availability requirements, such as ISP or ASP, in the business of hosting IT services
- May not have any remote offices

Medium-Sized Environments

These environments:

- May be classified under the tier-1, -2, or -3 storage models
- May have some remote offices; use low-end storage platforms for the remote offices

Large Enterprise Environments

These environments:

- May be classified under the tier-1, -2, or -3 storage models
- May have multiple remote offices and sites; use low-end storage platforms for the remote offices

Are highly recommended to use a tiered storage model in their data centers

3.1 Tiered Storage Infrastructure Architecture

Based on the different business criteria and SLA requirements, the storage infrastructure could be classified into three or more tiers. Table 1 offers some guidelines for building a tiered storage environment in your data center.

Table 1) Tiered storage architecture guidelines.

Business Criteria	Tier 1	Tier 2	Tier 3
Availability	99.999% (~5 minutes downtime per year). No single point of failure, maximize uptime, true HA solution, failover trunks, multipath disks, server failover, storage failover. Enterprise-level network switches or network directors.	99.99% (~52 minutes downtime per year). Redundancy internal to servers, networking, and storage: for example, storage failover, redundant power supplies. Servers or switches cannot be clustered at complete functional unit level. Simpler configurations, network switches from standard vendors.	99.9% (~8 hours downtime per year). Storage failover, server, or switches cannot be clustered. Simpler configurations. Network switches from standard vendors. Low- to mid-priced commodity equipment—fixable but with significant downtime.
Accessibility and Security	Private networks (dedicated network ports or VLANs for storage data, management of traffic, SnapMirror®). Disabled telnet, accessible through SSH only, restricted logon from certain servers or from certain administrators, role-based access control, secured hosting environment. Data security as per the application requirement. Data encryption, centralized logging (syslog), and NTP time services.	Shared network (shared network ports or VLANs for data and backup, separate ports for management of traffic), SSH access, and restricted logon from certain servers or from certain administrators and business application team members. Data security as per the application requirement. Centralized logging (syslog) and NTP time services.	Shared network (shared network ports for data and backup, separate ports for management of traffic), SSH access, and restricted logon from certain servers or from certain administrators and business application team members. Data security as per the application requirement. Centralized or localized logging (syslog) and NTP time services are usually Internet-based.
Scalability and Growth	Scale up and scale out. Consider growth for 18 months minimum.	Scale up. Consider growth for 18 months minimum.	Scale up. Consider growth for 18 months minimum.
Cost	Depends on the business application and the budget (low volume, high cost).	Depends on the business application and the budget (moderate volume, medium cost).	Depends on the business application and the budget (high volume, lowest cost).
Manageability	Enterprise management tools.	Standard management tools.	Standard management tools.
Data Protection	Hourly Snapshot™ copies (local backup), SnapMirror every hour to local and remote sites or MetroCluster™ with SyncMirror®.	Hourly Snapshot copies (local backup), SnapMirror every hour, 4 hours to DR site.	Hourly Snapshot copies (local backup), SnapMirror every 8 hours or daily SnapVault® to secondary storage (D2D backup) at local site.

Business Criteria	Tier 1	Tier 2	Tier 3
Recoverability	Separate DR site with hot-standby servers. Automated site failover.	Separate DR site with cold-standby servers. Manual site failover.	Tapes at the remote site. No failover. Manual recovery at remote site at time of disaster.
Production, Staging, Development, and Test Environments	Dedicated storage for production and staging environments. Primary reasons—change control, fault isolation, performance, and availability. Shared storage for dev and test.	Shared storage for production and dev/test. Dev/test are isolated from production and could be virtualized. Staging not required.	Shared storage for production and dev/test. Dev/test could be virtualized. Also could be used for archive data. Staging not required.
Disk Drive Types	Serial-attached SCSI (SAS), Fibre Channel (FC), serial ATA (SATA) disks.	Fibre Channel, serial ATA disks.	Fibre Channel, serial ATA disks.
Interoperability	Physical equipment, API, and protocol interoperability standards are high. Comprehensive testing is done during proof-of-concept phases and through QA testing. Heterogeneous yet stable environment needs careful attention to interoperability.	Fewer vendors; some interoperability testing and qualification required for critical systems. Mostly rely on vendor product compatibility matrixes.	Possibly one or two vendors, relatively simple configuration. Heavily rely on vendor product compatibility matrixes. May not be any prequalifications. Possibly standardize on fewer commonly used vendors.
Data Cloning Requirements	Cloning for dev/test environment.	Cloning for dev/test environment.	Cloning for dev/test environment.
Operational Profiles	DR events are relatively frequently rehearsed and tested, with well-defined procedures and dedicated operational staff.	DR procedures are well-defined and regularly but not necessarily frequently tested. Operational staff may have secondary responsibilities in more technical areas.	DR procedures are usually defined in a basic sense, but may not include regular testing. Knowledge about DR may rest with one or two individuals.
Application Examples	SAP, CRM, finance apps, Agile (sales), external Web site.	AutoSupport™, intranet Web sites, FTP, Exchange e-mails.	Home directories, corporate or group shares, HR apps, monitoring apps, nonproduction apps, archiving.

4 Storage Configurations

Any storage system must be prepared for initial setup and storage provisioning before it can be used to store user data or application data. When used properly in production deployments, a core set of storage best-practice guidelines that apply to all NetApp storage systems helps customers realize maximum data availability and integrity along with optimal performance.

4.1 Hardware Configurations

It's very important to lay out the correct hardware configuration in the beginning, keeping future growth in mind to avoid disruption to the production environment later on. NetApp recommends the following initial best practices for different types of environments, depending on the customer's requirements.

General Recommendations

- Use the sizing tools to choose the right storage platform and the number of disks right from the beginning.
- Use multipath for the disk drives in order to have good storage resiliency.
- Use a cluster in active-active mode rather than active-passive mode. This provides more throughput and storage efficiency from the NetApp storage system.
- Use a private network to separate all the storage traffic from other network traffic.
- Use the correct switch port settings along with right flow control, and the MTU size on the switch port and the network cards should match up.
- Make sure that the network cards are set for full duplex (for 100BaseT cards; gigabit and larger are full duplex only).
- Make sure that all of the components in your Fibre Channel loop are the same speed. For example, if you need 4GB Fibre Channel speed, then the FC port on the NetApp system, the DS14MK4 shelf, the SFP module in the ESH module, and the disk drives should all be capable of 4GB FC speed.
- Do not mix different RPM disk drives within the same aggregate.

Recommendations for Better Performance

- Based on the applications requirements of different departments, use high-speed FC disks for critical data and low-speed SATA disks for less important data. For example, use 300GB 15K RPM FC disk drives for a high-performance requirement and use low-RPM SATA disk drives for backup and archiving.
- Use 4Gb FC backplane speed to the FC shelves.
- Use 10GbE (single port, not dual ports, unless they run out of PCI slots) network interface with 9000 MTU size, depending on the workload.
- Use jumbo frames for 10GbE as a rule; end-to-end support for them is very important.
- Based on the workload, be sure to have an adequate number of Performance Acceleration Module (PAM) cards on the appropriate platforms.

4.2 Software Configurations

RAID Configurations

Use NetApp RAID-DP® (RAID Double Parity) for increased storage resiliency. Create the RAID group with the default size of 16 for FCAL drives and 14 for ATA drives (including 2 parity disks). This provides an optimal level of performance with increased resiliency against multiple drive failures. RAID groups are the basic building blocks for creating an aggregate and volumes. RAID group size (the number of disks in a RAID group) can vary from 2 to 28 drives. The optimal size for a RAID group is based on a number of factors: the time taken for reconstruction in case of a drive failure, the willingness to dedicate additional drives for parity, and the increased usable storage available when fewer parity drives are used. For more information on RAID groups, refer to the "NetApp System Administration Guide" for the version of Data ONTAP installed on your storage system.

Recommendations

- Use the default RAID group size when creating aggregates or traditional volumes.

- Allow Data ONTAP to select disks automatically when creating aggregates or volumes.
- The number of RAID groups in a particular aggregate is largely irrelevant from a performance perspective. However, the number of RAID groups does change the availability and average RAID reconstruct times for the aggregate.

Aggregate Creation

Create the largest aggregates possible. Try to maximize the number of disks in any aggregate, especially during the creation of the aggregate and when you have a good idea about the amount of data to be stored on this aggregate. This maximizes the performance (random read throughput) available to the aggregate and allows all the storage provisioning benefits of the flexible volumes to be realized. Independent Snapshot copy schedules are available at the aggregate level apart from the ones at the FlexVol® level. Data ONTAP 7G allows up to 100 aggregates (including traditional volumes) on a single storage system.

Recommendations

- Create the largest aggregates possible with the maximum number of disks supported in them.
- Avoid creating small aggregates, because they can become disk-bound even for sequential workloads.
- Allow Data ONTAP to choose disks and adapters during aggregate creation. Because Data ONTAP automatically spreads aggregates across disks adapters, let it choose the member disks of an aggregate.
- Increase the aggregates in increments of RAID group sizes, if necessary. Do not add one or two disks at a time.
- Refer to the System Configuration Guides available on the NetApp Support site for the maximum size of an aggregate.

Hot Spares

With NetApp's self-healing RAID software, disk failures automatically trigger parity reconstructions of affected data onto a hot-standby disk. Note, however, that a hot spare disk must be available to Data ONTAP for this self-healing process to begin. Therefore, at a minimum, resiliency planning should include keeping at least one hot spare disk for each type of disk drive present in the storage system. Disk drive differences are FC, SAS, SATA disk drive types, disk size, and rotational speed (RPM).

Recommendations

Here are some general recommendations to increase storage resiliency:

- Maintain two hot spares for each type of disk drive in the storage system to take advantage of the Maintenance Center.
- For active-active configurations, hot spares must be owned by the right storage controller; with SyncMirror, hot spares must be in the right pool.
- NetApp recommends using two spares per disk type for up to 100 disk drives. For each additional 84 disks above that, another hot standby disk should be allocated to the spare pool.

Flexible Volumes

Administrators can implement multiple flexible volumes to host various types of CIFS data. Data ONTAP 7G introduces flexible volume technology, a breakthrough technology in which volumes are logical data containers that can be sized, resized, managed, and moved independently from the underlying physical storage. This enhances the storage administrator's ability to address a variety of data management requirements while preserving the familiar semantics of volumes and the current set of volume-specific data management and space allocation capabilities. Like qtrees, FlexVol volumes offer the flexibility for

dividing system data into more granular volumes, setting file security types and quotas, and setting CIFS opportunistic locking. In addition, independent Snapshot copy schedules can be configured per FlexVol volume (unlike a qtree). The use of FlexClone® is available at the FlexVol level only, and FlexVol volumes can be mirrored (synchronously and asynchronously) by using volume SnapMirror for disaster recovery. This feature is implemented in 7-Mode as well.

For a single storage system, Data ONTAP 7G/7-Mode allows up to 200 FlexVol volumes for FAS2000 and FAS200 series platforms, and 500 FlexVol volumes for all other platforms.

Recommendations

- Prefer FlexVol volumes over traditional volumes.
- Create volumes based on needs for different Snapshot copy schedules, security requirements, or quota requirements.
- Make sure that all files and folders are created with Unicode.
- Convert any existing non-Unicode folders to Unicode.
- To resolve any volume language warnings, if the volume contains only NDMP Backup (NetBackup™ DSU) or SnapVault® qtrees, change `vol lang` to `C.UTF-8`.

Root Volume

The root volume can exist as a traditional RAID 4 or RAID-DP volume or as a FlexVol volume that is part of a larger hosting aggregate. In practice, having the root volume on a FlexVol volume makes a bigger difference with smaller-capacity storage systems than with very large ones, in which dedicating two disks for the root volume has little effect.

Recommendations

- Create a root volume distinct from any data volumes. No user data should be stored on the root volume.
- Create a RAID-DP root volume for increased resiliency.
- For small storage systems in which cost concerns outweigh resiliency, a FlexVol-based root volume on a regular aggregate may be more appropriate.

Snap Reserve

Snap reserve specifies a set percentage of the disk space for Snapshot copies. By default, snap reserve is 20% for a volume and 5% for an aggregate. NetApp recommends keeping the default snap reserve settings initially and later increasing or decreasing them based on the Snapshot copy disk consumption.

You should adjust the snap reserve to be slightly more than your Snapshot copies consume at their peak. It might be necessary to monitor the system for some time to determine the average Snapshot copy size. You should consider the following change rates for CIFS:

- Daily change rate between 0.1% and 1%
- Weekly change rate between 1% and 7%
- Monthly change rate between 5% and 20%

The snap reserve can be changed at any time. Be sure not to raise the snap reserve to more than the free space on the volume, or client machines may abruptly run out of storage space.

NetApp recommends that you observe the amount of snap reserve being consumed by the Snapshot copies frequently. Do not allow the amount of space consumed to exceed the Snapshot copy reserve. If you exceed your Snapshot copy reserve, consider increasing the percentage of the snap reserve, or delete Snapshot copies until the amount of space consumed is less than 100%.

Snapshot Schedule

The default Snapshot copy schedule automatically creates one nightly Snapshot copy Monday through Saturday at midnight and four "hourly" Snapshot copies at 8 a.m., noon, 4 p.m., and 8 p.m. Data ONTAP retains the two most recent nightly Snapshot copies and the six most recent hourly Snapshot copies, and deletes the oldest nightly and hourly copies when new copies are created. The following command shows the default Snapshot copy schedule in Data ONTAP:

```
snap  sched  volume_name  0  2  6@8,12,16,20
```

The Snapshot copy schedule can be changed based on the follow selection criteria:

- If users rarely lose files or typically notice lost files right away, use the default Snapshot copy schedule.
- If users commonly lose files or do not typically notice lost files right away, delete the copies less often than you would with the default schedule. NetApp recommends keeping two weekly Snapshot copies, six nightly copies, and eight hourly copies:

```
snap sched  vol1  2  6  8@8,12,16,20
```

- Depending on the applications, more weekly Snapshot copies might be required.
- Based on the RPO and RTO requirements, Snapshot copies can be scheduled more frequently. The RPO and RTO parameters are closely associated with recovery.
- You can create different Snapshot schedules for different volumes on a storage system. On a very active volume, schedule Snapshot copies every hour and keep them for just a few hours, or turn Snapshot off. For example, the following schedule creates a copy every hour and keeps the last three:

```
snap sched vol2 0 0 3
```

- This schedule does not consume much disk space, and it lets users recover files in recent Snapshot copies as long as they notice the loss within a couple of hours.
- Use hourly Snapshot copies for users' home directories and group shares especially. This adds a real business benefit for users who mistakenly delete or overwrite files, and it consumes very little space.
- When you create a new volume, it inherits the Snapshot schedule from the root volume. After you use the volume for a while, check how much disk space the Snapshot copies consume and how often users must recover lost files. Then adjust the schedule as necessary.
- If you are using volume SnapMirror, keep in mind that the destination volume should not have a Snapshot schedule set, because it carries all the source Snapshot copies as a result of volume SnapMirror.
- The secondary storage systems that are used for backup or archiving require longer-term retention than the primary storage systems; you might need more daily and some weekly Snapshot copies. You could reduce or stagger the frequency of hourly copies for the secondary storage systems, because they could conflict with the SnapMirror or SnapVault schedule.
- When monitoring Snapshot schedules, keep in mind the maximum number of Snapshot copies possible in Data ONTAP (255 for a single volume).

Qtrees

NetApp has traditionally recommended the use of qtrees for user and project file data. Qtrees enable you to partition volumes into smaller segments that can be managed individually. A qtree is a special type of directory on the NetApp storage system that enables you to specify a security style and to set limits on disk space consumption. In Data ONTAP, qtrees offer the most flexibility for dividing data volumes into granular chunks, setting file security types and quotas, applying storage object security by using Storage-Level Access Guard, and setting CIFS opportunistic locking and qtree SnapMirror for disaster recovery and SnapVault for backup.

Recommendations

Use qtrees to put the CIFS shares on; this helps in setting up the SnapVault relationships. As a general practice, use qtrees when you need different:

- Security styles
- Storage object security
- Oplock settings
- Quota limits
- Backup and restore operations

Note: Don't use a qtree for each user or each group directory. This would result in thousands of backup relations and lead to complications if SnapVault backup needs to be used.

Data ONTAP 7G/7-Mode allows up to 4,995 qtrees per FlexVol volume or traditional volume.

Security on Qtrees, Volumes, and FlexVol Volumes

The security setting on a qtree, volume, or FlexVol volume destination for resident or migrated data for Windows file services should be set to NTFS. Even though UNIX® clients will access data, NTFS should be specified because it is the dominant security style for this type of environment. If UNIX (also known as nonnative) access to the Windows data is required, implement user mapping on the NetApp system to grant proper access. For information about user mapping, see information about the `/etc/usermap.cfg` file in Data ONTAP admin guides.

Security style can be set to UNIX if it is a pure or heavy UNIX file serving environment running over NFS v3 or NFSv4. NetApp recommends that you limit or restrict the use of mixed security style qtrees, volumes, and FlexVol volumes. Very few situations require the mixed security style, and use of this style can cause additional administrative overhead in dealing with the management of two sets of permissions styles in one qtree.

Quotas

Even though FlexVol volumes can be easily resized, NetApp recommends keeping disk space consumption under control, especially for home directories, project shares, and corporate shares volumes. You can set user quota, group quota, or tree quotas in Data ONTAP. For more granular controls and reporting, NetApp recommends using our partner software from NTP or Northern Parklife. Quotas can also be configured to track disk usage.

Recommendations

- Use quotas to control disk space usage.
- Always set a default quota on the volume.
- Quotas must be applied to all volumes that contain home directories. You can define default or explicit user quotas for home directories.
- You can also use quotas to measure the disk space used in order to charge back.
- For more information on Data ONTAP quota implementation and partner solutions, refer to www.netapp.com/us/library/technical-reports/tr-3425.html.

Deduplication

If the solution design includes space saving as one of the objectives, then apply deduplication on the primary and secondary storage as necessary. For data that does not change frequently, but that is mostly the same in all Snapshot copies, deduplication maximizes the use of disk space, which means cost savings. For example, in environments like home directories, project shares and corporate share data, logs, and archive data that don't change frequently, deduplication technology can save 20%-50%. For

more information on NetApp deduplication technology, see TR-3505, "NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide."

Backup and Restore

- NetApp recommends using Snapshot copies to create images of volumes for backup and recovery purposes.
- Back up NetApp system configuration, especially before any upgrades; keep a copy of the `/etc` folder.
- For the secondary storage, choose the right NetApp storage platform, considering the CPU utilization for all the mirroring and tape backup operations.
- Not all the applications need tape backup; often disk-to-disk backup is sufficient.
- Backup and retention policies can be set according to the requirements of your business. For example, Compliance=7 years, Archive=5 years, Business=3 years, Scratch=2 weeks.
- Use SnapVault for most of the CIFS shares backups to overcome backup issues in large-file-count environments, to restore issues for media errors, long durations, and so on.

Automatic E-mail Notification to Netapp

1. Enable and configure the **AutoSupport** feature in Data ONTAP.

Capacity Planning

- Keep the space utilization in all the volumes and aggregates under 80%.
- Do not let used space in any volume (traditional or FlexVol) exceed 90% for an extended period of time.
- Use compression utilities like tar to store data efficiently wherever possible.
- Use NetApp deduplication technology to achieve space savings.
- Use 10k RPM FC disk drives or SATA disk drives unless it's necessary to use 15k RPM FC disk drives in SAN or higher-performance environments.

Growth Considerations

Consider at least one year's growth when provisioning space for any type of data storage.

Daily Operations

- Monitor NetApp storage system health regularly.
- Install and use a monitoring tool such as Operations Manager.
- Regularly check the NetApp system's syslog message (`/etc/messages`) files. You can also set up alerts in Operations Manager to be triggered from certain types of warnings.
- Check the system CPU usage during peak load times.
- Run the `netdiag` command to check the health of network interfaces and network.

5 Integration with Microsoft Windows

The following sections describe some typical best practices employed by Windows system administrators when deploying a NetApp storage system in a Microsoft Windows environment.

5.1 Windows Workgroup Mode Authentication

This method of authentication is the second-most-used mode and is a step up from the UNIX mode authentication choice because passwords are not sent in clear text. This method uses NT LAN Manager

(NTLM) authentication without a domain controller against a small list of users who are specified on the NetApp storage system. This list is kept in the local user database and is relevant only to that storage system. The obvious limitation here is that authentication is done locally on the system itself and therefore there is no centralized authentication across the entire environment.

From a security standpoint, workgroup mode is better because it supports SIDs and NTLM authentication (although locally on the system). This mode also works fine in conjunction with a domain controller, meaning that you can have local user and group accounts defined on the system and simultaneously be part of a domain. If your domain goes down, you still retain some level of access to the system for emergency use through a local user account, such as Administrator:

```
net use * \\toaster\c$/user:toaster\administrator <passwd>
```

where you use the system name instead of the domain name for a user account to gain quick access.

Limitations with Workgroups

- You can create a maximum of 96 local user accounts.
- You cannot use User Manager to manage local user accounts on your storage system.
- To add local users to the system, you must use the `useradmin` command in Data ONTAP. You can view the local users only through the Microsoft Computer Management MMC snap-in, but you can manage the local groups on the storage system.

Recommendations

- Always create a local administrator CIFS account on the storage system, so that you always have some level of administrator access to the system.
- Choose a workgroup name that is synonymous with the rest of your Windows clients. Typically, the name “workgroup” is left alone because this is the default and makes administration easier.

Use Cases

- Generally recommended for customers that:
  - Are small shops with just a handful of workstations—don’t want to deal with domain administration
  - Don’t have the budget to implement a domain—have fewer than 96 users
- Applications that you must authenticate but for which you may not want to use a domain account include:
  - SnapDrive® for Windows
  - SnapManager® for Exchange
  - Internet Information Services (IIS)

5.2 Active Directory Domain Mode Authentication

NetApp highly recommends using a domain-style mode of authentication. The choice of using a Windows 2008 R2, Windows 2008, Windows 2003, or Windows 2000 Active Directory domain depends on what your organization’s requirements are and what your existing domain structure looks like.

Active Directory stores information about network components. It allows clients to find objects within its *namespace*. The term namespace (also known as *console tree*) refers to the area in which a network component can be located. Active Directory provides a namespace for resolving the names of network objects to the objects themselves. Active Directory can resolve a wide range of objects, including users, systems, and services on a network. For more detailed information, see http://technet.microsoft.com/en-us/library/bb742424.aspx.

NetApp systems support the Lightweight Directory Access Protocol (LDAP) and the Kerberos authentication protocol to communicate with Active Directory and to authenticate access to the NetApp

system's resources. Functional levels determine the available Active Directory Domain Services (AD DS) domain or forest capabilities. They also determine which Windows Server® operating systems you can run on domain controllers in the domain or forest. However, functional levels do not affect which operating systems you can run on workstations and member servers that are joined to the domain or forest. There are two types of functional levels in Active Directory: domain and forest.

Domain functional levels. Domain functionality activates features that affect the whole domain and one of the following domains only. These levels are distinguished by the version of the Windows Server operating system that is permitted on the domain controllers present in the domain. With each successive level increase, the domain functionality activates features of the previous domain level.

Forest functional levels. Forest functionality activates features across all the domains in your forest. NetApp storage systems can join and participate in any of the following domain or forest function levels of Active Directory.

Table 2) Forest functional levels

Domain or Forest Function Level	Supported Domain Controllers
Windows 2000 mixed	Windows NT® 4.0
Windows 2000 native	Windows 2000 Windows Server 2003
Windows Server 2003	Windows Server 2003 Windows Server 2008
Windows Server 2008	Windows Server 2008 Windows Server 2008 R2
Windows Server 2008 R2	Windows Server 2008 R2

Windows 2000 Domains

Windows 2000 Server introduced two Active Directory domain functional levels, mixed and native, to support different deployment scenarios. There may be some deployments still using Windows 2000 domains in mixed or native mode. However, the Windows 2000 Servers have certain limitations, and they are now considered as legacy systems.

Limitations of Windows 2000 Servers

- No support for Universal Groups
- No support for Group Policies (only System Policies)
- Active Directory database is limited to 40MB (24,000 accounts versus 3 to 10 million for AD)

In both Windows 2000 Server mixed and Windows 2000 Server native, domain styles support legacy clients. NetApp storage systems can still interoperate in either **mixed-mode** or **native-mode** Active Directory domains and adhere to the environment in which they are installed.

Recommendations

- Consider implementing a Windows 2000 native and later functional level. Windows 2000 mixed-function level should be considered only as a stop gap until you can migrate your entire legacy domain.
- Use Windows 2000 mixed domain only if you have a mixture of Windows 2000 and later domain controllers as well as Windows NT 4.0 backup domain controllers.

Windows 2003 Domains

In Windows Server 2003, the terms mixed and native have been superseded by *Raise Function Level*. Windows 2003 introduced two additional modes, Windows Server 2003 interim and Windows Server 2003 (also known as Windows Server 2003 native).

Windows 2003 interim mode is much like Windows 2000 mixed, but it offers a few improvements, such as replication. Interim mode is an interim solution that provides compatibility with NT domains until they can be upgraded to Windows 2003. It has the same limitations and caveats as Windows 2000 mixed mode.

Recommendations

Use Windows 2003 interim mode when:

- Migrating from Windows NT domains without going to Windows 2000 first
- Upgrading the first Windows NT domain to a new forest
- You only have Windows NT 4.0 and Windows 2003 domain controllers in your environment (no Windows 2000 domain controllers)
- You have Windows NT 4.0 groups with more than 5,000 members (Windows 2000 Server does not allow you to create groups with more than 5,000 members)

Windows Server 2003 (native) mode offers some additional enhancements, including all the improvements found in Windows 2000 Server native mode. One of these enhancements is the ability to rename domains and domain controllers easily. Windows Server 2003 mode can support Windows 2003 as well as Windows 2008 domain controllers.

Use Windows 2003 Server (native) mode when:

- You are building a new domain with Windows Server 2003 or Windows 2003R2.
- All domain controllers in your organization's existing domain have been upgraded to Windows Server 2003. If you're currently running a native 2000 domain, you would need to upgrade all Windows 2000 domain controllers to Windows Server 2003.
- You want the same features (plus some additional) that a native-mode 2003 domain provides.
- You have a mix of Windows 2003, Windows 2003 R2, and Windows 2008 domain controllers.

Note: Both the Windows Server 2003 interim and Windows 2000 Server native domain styles support legacy clients. NetApp storage systems can interoperate in either **interim-mode** or **native-mode** Active Directory domains and adhere to the environment in which they are installed.

Windows 2008 and 2008 R2 Domains

Windows 2008 Active Directory domain is a native-style domain mode that currently supports only Windows 2008 domain controllers. This type of domain can support Windows Vista clients and legacy clients such as Windows XP, Windows 2000, and so on. Windows 2008 and 2008 R2 domains include all of the Active Directory Domain Services (AD DS) features that are available at the Windows Server 2003 domain functional level and some more enhancements. These include Distributed File System Replication support for Windows Server 2003 System Volume, AES 128 and AES 256 support for the Kerberos protocol, and support for Read Only Domain Controller. Windows 2008 Servers support SMB 2.0 and 2008 R2 supports SMB 2.1, which is the next version of the CIFS protocol. NetApp storage supports up to SMB 2.0 only.

Recommendations

Use Windows 2008 domain mode when the following requirements are proposed:

- Building a new domain with Windows Server 2008
- All the domain controllers in your environment are Windows Server 2008

- Implementing Read Only Domain Controllers for remote offices
- Wanting to use restartable Active Directory Domain Services (AD DS)
- Implementing auditing for AD DS with a new audit policy subcategory (Directory Service Changes) to log old and new values when changes are made to AD DS objects and their attributes
- Implementing fine-grained password policies, that is, multiple password policies within a single domain
- Using the new AD database mounting tool (Dsamain.exe) to improve the recovery processes for Active Directory data
- Using the improved installation wizards and management snap-ins for AD DS

For more information about the enhancements in Windows Server 2008 from Windows Server 2003, refer to http://technet.microsoft.com/en-us/library/cc753208.aspx.

For more information about the enhancements in Windows Server 2008 to 2008 R2, refer to http://technet.microsoft.com/en-us/library/dd391932(WS.10).aspx.

2008 R2 Domain Features

Windows 2008 R2 introduces some additional features for the domain:

Active Directory Recycle Bin: With the Active Directory Recycle Bin, we can undo an accidental deletion of an Active Directory object. Accidental object deletion causes business downtime. Deleted users cannot log on or access corporate resources. Active Directory Recycle Bin works for both AD DS and Active Directory Lightweight Directory Services (AD LDS) objects. This feature is enabled in AD DS at the Windows Server 2008 R2 forest functional level.

Active Directory Module for Windows PowerShell: The Active Directory module for Windows PowerShell provides command-line scripting for administrative, configuration, and diagnostic tasks, with a consistent vocabulary and syntax. The Active Directory module enables end-to-end manageability with Exchange Server, Group Policy, and other services.

Offline Domain Join: Offline domain join makes provisioning of computers easier in a data center. It provides the ability to preprovision computer accounts in the domain to prepare operating system images for mass deployment. Computers are joined to the domain when they first start.

Which Domain Mode Should I Choose?

Your primary long-term objective should be to get your organization's Active Directory to Windows 2003 native or Windows 2008. Although NetApp storage systems can interoperate with any of the Active Directory domain modes, NetApp recommends at least raising the domain or forest function levels to Windows 2003, Windows 2003 R2, Windows 2008, or Windows 2008 R2. This is especially important because the legacy operating systems from Microsoft are nearing end of support.

To choose a domain mode for any new domain controller, you must check which mode your existing domain is in. Follow these steps:

1. Open the Active Directory Domains and Trusts console.
2. Right-click the domain in the left pane and select Properties.
3. Look at the domain functional level and the forest functional level on the General tab of the domain's property page to see the mode in which your existing domain is running.

Guidelines for Raising Domain and Forest Functional Levels

To raise the domain or forest functional level, follow these guidelines:

- You must be a member of the Domain Admins group to raise the domain functional level.
- You must be a member of the Enterprise Admins group to raise the forest functional level.

- You can raise the domain functional level on the primary domain controller (PDC) emulator operations master only. The AD DS administrative tools that you use to raise the domain functional level (the Active Directory Domains and Trusts snap-in and the Active Directory Users and Computers snap-in) automatically target the PDC emulator when you raise the domain functional level.
- You can raise the forest functional level on the schema operations master only. Active Directory Domains and Trusts automatically target the schema operations master when you raise the forest functional level.
- You can raise the functional level of a domain only if all domain controllers in the domain run the version or versions of Windows that the new functional level supports.
- You can raise the functional level of a forest only if all domain controllers in the forest run the version or versions of Windows that the new functional level supports.
- You cannot set the domain functional level to a value that is lower than the forest functional level.
- You cannot reverse the operation of raising the domain and forest functional levels. If you must revert to a lower functional level, you must rebuild the domain or forest, or restore it from a backup copy.

For more information about Active Directory Domain Services Domain and Forest functional levels, see http://technet.microsoft.com/en-us/library/cc754918.aspx.

5.3 UNIX Mode Authentication

Although this method of authentication does not exactly qualify as a Windows domain type and is the least preferred method, it's worth discussing this legacy mode.

This method is used to authenticate against the system's local `/etc/passwd` or the organization's NIS or LDAP infrastructure, or both. Security is the most obvious limitation in this method, because passwords are sent in clear text. This option also requires a registry change on the Windows client to allow clear-text passwords.

This authentication mode works by mapping the Windows user name to a UNIX user name and checking the supplied password against a UNIX hash. Based on order password lookup in the `/etc/nsswitch.conf` file, user mapping is done against the users from the `/etc/passwd` file or NIS or LDAP.

Another important limitation of this style of authentication is that there are no security identifiers (SIDs), so again it's not very secure. Because of the lack of security and limitations, NetApp does not recommend this mode of authentication.

This option is most useful in UNIX shops with a handful of Windows machines and no Windows domain established.

5.4 Kerberos Authentication

Data ONTAP includes native implementations of the NTLM, NTLMv2, and Kerberos protocols and therefore provides full support for the Active Directory and legacy authentication methods. The security style of authentication that Data ONTAP uses depends on the client and what they can negotiate with. This is true for both mixed and native domains of functional-level Windows 2000 and later. For example:

- Standalone computers running Windows 2000 or later that are not part of an Active Directory domain use NTLM only.
- Computers running Windows 2000 or later that are part of an Active Directory domain default to using Kerberos first, then NTLM.

With Kerberos authentication, Windows clients contact the KDC service that runs on Windows 2000 and later domain controllers and do a TGT (Ticket Granting Ticket) and TGS (Ticket Granting Service) exchange with KDC. Clients then pass the authenticator and encrypted session ticket to the NetApp storage system, from which a CIFS credential is constructed to create a session ID for SMB traffic.

Here are some things to keep in mind regarding Kerberos:

- The system (Windows Server or client) must establish its own authenticated connection to a domain controller (DC).
- The system must also contact Kerberos KDC (which may not be the same box as the DC) to authenticate as a client to the DC.
- The system has the information needed to decrypt the client's Kerberos ticket, but unfortunately there is only a numeric SID, which requires contacting the Local Security Authority (LSA) on a DC to convert that to a string form for mapping to get a UNIX credential. This is a fairly fast procedure, and the SID-to-name information is cached locally. This is one reason to make sure that SID caching is enabled on the system (the default).

5.5 Domain Controller Discovery

A NetApp system can be joined to and operate in any Active Directory mode. It attempts to automatically sense what type of domain exists on the network. It first searches for an Active Directory domain controller or LDAP server by querying the DNS server. This is the same method used by Windows 2008, Windows 7, Windows Vista, Windows XP, and Windows 2003 computers. The storage system attempts to search for domain controllers or LDAP servers under the following conditions:

- The storage system has been started or rebooted.
- A CIFS `resetdc` command has been issued.
- Four hours have elapsed since the last search.

When deployed in a Microsoft Windows Active Directory environment, NetApp systems perform the following discovery process to find and connect to domain controllers:

1. Verify the cached server address bias ("last connection" cache).
2. Verify the domain controller priority groups:
    a. Preferred: Domain controllers defined in the cifs prefdc list
    b. Favored: Domain controllers that are members of the same AD site or that share the same subnet as storage system sorted by fastest response time, or else in random order
    c. Other: Domain controllers that are not members of the same AD site sorted by fastest response time, or else in random order
3. Query directory SRV records in DNS

This discovery process runs completely through all of these steps regardless of any successful connections found. All addresses are discovered at once, categorized, prioritized, and cached. From this list, Data ONTAP selects the optimal DC to be used.

Recommendations

To improve DC discovery and connections:

- Use Microsoft Sites (described in section 5.7) to make sure that the domain controllers selected are physically as close to NetApp systems as possible.
- Make sure that a domain controller is relatively close to the NetApp system (on the same LAN). Placing NetApp systems remotely (over a WAN) might have performance implications on authentication requests from clients and client logon times. A good practice is to place a domain controller near any users or servers (including NetApp systems) so that the users can still log on even if the WAN connection fails.
- Use a list of preferred DCs (`CIFS prefdc`). This is not always required, but might be useful for troublesome domains where the closest DC may not be chosen for various reasons and you want to specifically designate a DC. There is only one `prefdc` list, used to identify preferred addresses for

DC connections. Most customers use this command to control which addresses they prefer to have Data ONTAP refer to for DC connections.

Note: Even if you designate a DC from the `prefdc` list, the NetApp system still selects DCs that have better response times. This is so a server you specify doesn't fail, causing the NetApp system to wait for it to come back up.

For more information on how NetApp systems interact with Microsoft Active Directory, see TR-3367, NetApp Storage Systems in a Microsoft Windows Environment.

5.6 Using DNS in Active Directory

Active Directory and the Domain Name System (DNS) are tightly integrated and share the same hierarchical structure. Therefore AD relies on DNS to resolve names and services to IP addresses. In addition, Active Directory clients and client tools use DNS to locate domain controllers for administration and logon.

Like all Active Directory systems, NetApp storage systems use DNS to identify AD-defined sites and domain controllers and to locate special network services such as those that run on domain controllers, Kerberos, and KPASSWD services.

Recommendations

- You must have a DNS server installed and configured for Active Directory.
- An Active Directory compatible DNS server must be online and configured properly to install a NetApp storage system into an AD domain.
- Microsoft recommends that you use Microsoft DNS Server, supplied with Windows 2008/2003 Server, as your DNS server. However, Microsoft DNS is not required. Any non-Microsoft DNS server, such as a Berkeley Internet Name Domain (BIND) server, can be used. However, you should verify that the version being used supports SRV records, or update it to a version that does.
- NetApp recommends using a Microsoft DNS server to avoid the administration of two different types of DNS servers in your environment.
- The DNS domain name does not need to match the name of the Active Directory domain of which the NetApp system is a member. For example, an organization's DNS name may be *hq.princeton.com,* but the Active Directory domain name might be *marketing.princeton.com.*
- In some cases, for smaller environments you can use the same Active Directory host (DC) for your DNS server as well unless you have a designated Microsoft DNS server or set up a specialized BIND DNS server (with secure update patches).
- NetApp recommends using Dynamic DNS updates. If these updates are not available, you can manually add the "A" record on the DNS server (described in section 5.6.4).
- Make sure that DNS options are configured properly on the NetApp storage system as well (described in section 5.6.5).

Note: If DNS is not enabled, or if it is not configured correctly, Data ONTAP cannot find the service records it needs to locate DCs, KDCs, LDAP servers, and KPASSWD servers, so it cannot join the AD domain.

DNS Server Type

The DNS server that you use:

- Must support the SRV RR (RFC 2782)
- Should support the dynamic update protocol (RFC 2136)

BIND version 8.1.2 or later (a popular DNS server implementation) supports both the SRV RR and dynamic update (version 8.1.1 does support dynamic updates, but it has flaws that were fixed in 8.1.2.). If

you are using a version of BIND that does not support dynamic update, you must manually add records to the DNS server.

Verifying DNS Configuration

Open the DNS Management Console on the MS-DNS server. Verify that you have a DNS domain with the same name as your corresponding Active Directory domain. It must contain the four SRV record folders (child domains): `_msdcs/ _sites/, _tcp/, and _udp/`. These must exist and should look similar to those in Figure 1. Notice that all four folders are present for the *southfield.lab* domain.

Figure 1) DNS database records.

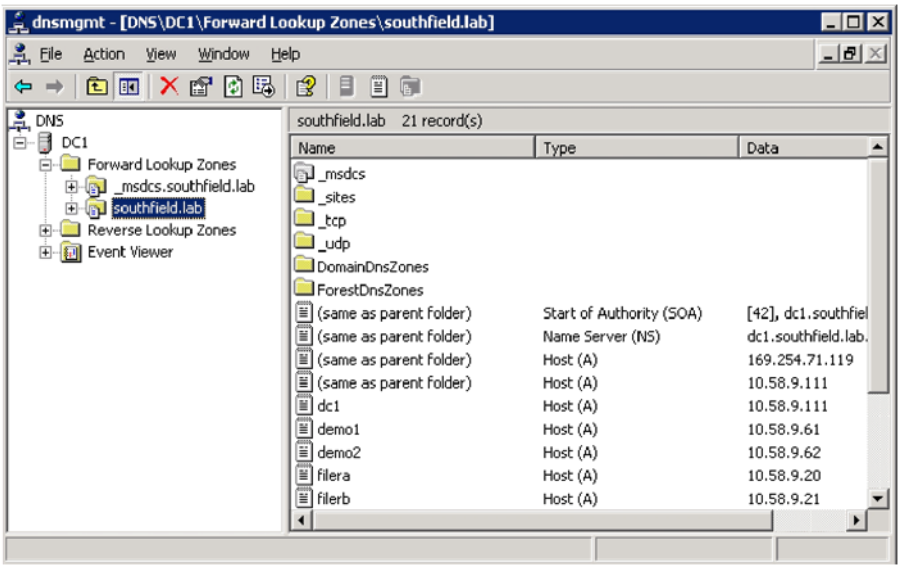

If they don't look like this, your AD functions might have broken. One sign of this is a long log-on time to the DC. The Preparing Network Connections window remains on the screen for quite a while and many AD operations return errors when you try to perform them. This can happen if you did not manually configure your DNS server and instead let the DCPROMO process do it for you.

Another reason for the lack of SRV records (and of all other records, for that matter) is that you did configure the DNS server manually, but you made a mistake, either with the computer suffix name or with the IP address of the DNS server.

Recommendations

To try to fix the problems, first see if the zone is configured to accept dynamic updates:

1. Right-click in the zone you created and click Properties.
2. On the General tab, under Dynamic Update, select Nonsecure and Secure from the drop-down list, and click OK to accept the change.
3. Restart the NETLOGON service to force the SRV registration:

   To stop and start netlogon, run the following command at the command prompt:

```
net  stop  netlogon & net  start  netlogon
```

4. Return to the MS DNS console, click in your zone, and refresh it (F5). If all is correct, you should see the four SRV record folders.

If the four SRV records are still not present:

- Double-check the spelling of the zone in the DNS server. It should be exactly the same as the AD domain name.
- Check the computer's suffix. You can't change the computer's suffix after the AD is installed, but if you have a spelling mistake it's better to remove the AD now, before you have any users, groups, and other objects in place. You can then repair the mistake and rerun DCPROMO.

Note: Make sure that an "A" record exists in DNS that corresponds to the AD domain name. It generally does exist, and the customer would probably know about any deviation.
If there is no "A" record for the domain in DNS, `cifs setup` gives up on an Active Directory domain joining, and it assumes that the domain is an NT4 domain. This is not what you want.

Dynamic DNS Updates

Using dynamic DNS updates can prevent errors and save time when sending new or changed DNS information to the primary master DNS server for your storage system's zone. Dynamic DNS allows your storage system to automatically send information to the DNS servers as soon as the information changes on the system.

NetApp recommends enabling the dynamic DNS updates. Otherwise, you must manually add the DNS information (DNS name and IP addresses) to the DNS server when a new storage system is brought online or when existing DNS information changes. This process is slow and error-prone. In a disaster-recovery situation, manual configuration can result in long downtimes. Data ONTAP supports a maximum of 64 Dynamic Domain Name Server (DDNS) aliases.

Recommendations

The following conditions apply when you are using dynamic DNS updates:

- By default, dynamic DNS updates are not enabled in Data ONTAP.
- Dynamic DNS updates are supported on UNIX and Windows systems.
- On Windows DNS servers, secure dynamic DNS updates can be used to prevent malicious updates on the DNS servers. Kerberos is used to authenticate updates. Even if secure dynamic DNS updates are enabled, your storage system initially tries to send updates in clear text. If the DNS server is configured to accept only secure updates, the updates sent in clear text are rejected. Upon rejection, the storage system sends secure DNS updates.
- For secure dynamic DNS updates, your storage system must have CIFS running and must be using Windows Domain authentication.
- Dynamic DNS updates can be sent for the following:
  - Physical interfaces
  - VIF and VLAN interfaces
  - vFiler® units
- You cannot set TTL values for individual vFiler units. All vFiler units inherit the TTL value that is set for vFiler0, which is the default vFiler unit and is the same as the physical storage system.
- DHCP addresses cannot be dynamically updated.
- In a takeover situation, the hosting storage system is responsible for sending DNS updates for IP addresses for which it is responding.
- For both manual and auto configured global IPv6 unicast addresses, the dynamic DNS update is sent after Duplicate Address Detection is performed. For IPv6 addresses of any other type and scope, your storage system does not send any dynamic DNS update.
- For DDNS updates to function properly, you must configure a reverse lookup zone also on the DNS server.

Dynamic DNS on your storage system can be enabled by setting the following option entry:

```
options dns.update.enable {off|on|secure}
```

"A" Record for NetApp Storage System in DNS

If the DNS server in your environment doesn't support the dynamic updates, then you must create an "A" record on the DNS server manually to access the storage system by name.

1. Open the DNS management console on the MS-DNS server.
2. Right-click the desired AD domain under Forward Lookup Zones (in this case, *southfield.lab*) and select New Host (A).
3. Enter the NetBIOS name of the NetApp system (see section 5.9) and its IP address.
4. Click Create Associated Pointer (PTR) Record. This is useful for `nslookup` and other troubleshooting tools.

Figure 2) Creating "A" DNS record for storage system.

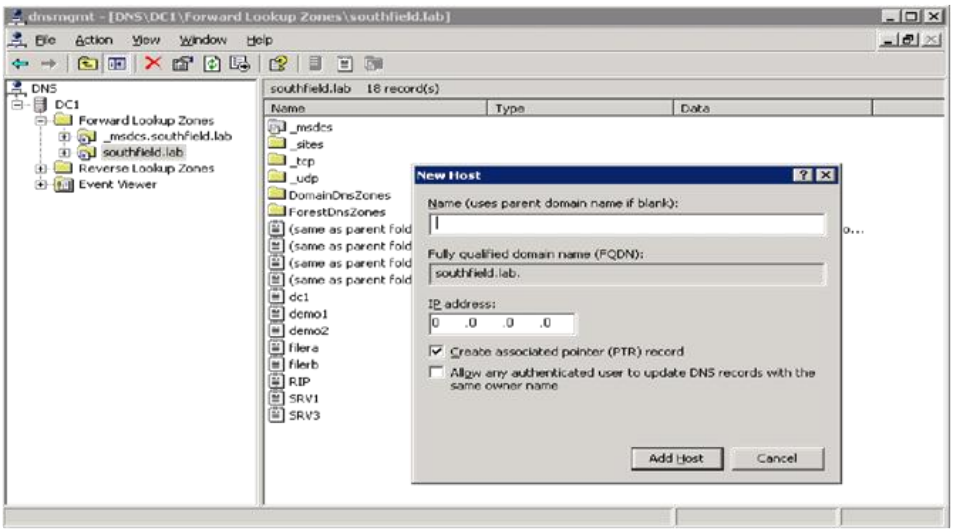

Configure DNS on the NetApp Storage System

The following steps enable DNS without having to use the NetApp system's initial `setup` command (not to be confused with `cifs setup`).

1. Create or confirm the system's `/etc/resolv.conf` file with nameserver entries:
   `nameserver <dns-server-ip-address>`

```
nameserver <dns-server-ip-address> nameserver <dns-server-ip-address>
```

2. Set the following options on the NetApp system's command line:

```
toaster> options dns.enable on
toaster>  options  dns.domainname  <DNS  domain  name>
```

3. Edit the NetApp system's `/etc/rc` file with the following entries so the DNS options are persistent across system reboots:

```
options dns.domainname <DNS domain name> options dns.enable on
```

4. Make sure that the NetApp system's `/etc/nsswitch` file has the hosts entry that looks like this:

```
hosts: files dns nis
```

Note: The name service switch configuration (`nsswitch`) file contains the preferred order in which name services are contacted for name resolution by the storage systems, until the name is successfully resolved. In most cases you shouldn't have to edit this file, but it should resemble the

example in step 4 for the 'hosts' entry for DNS resolution to occur after checking the local `/etc/hosts` file.

5.7 Microsoft Site Awareness

Active Directory sites are used to logically represent an underlying physical network. A site is a collection of networks connected at LAN speed. Slower and less reliable wide area networks (WANs) are used between sites (locations) that are too far apart to be connected by LAN.

Sites allow clients and servers to locate domain controllers that are physically close, so that they can be accessed efficiently. Sites also provide information that is useful for efficient and timely replication between DCs. Without a proper site configuration, it is possible for a server to connect to and use a domain controller that is halfway around the world. A defined Microsoft site is not required for Active Directory; however, Microsoft recommends it as a best practice. If a site has been defined, site-specific DNS queries are used to discover where the optimal (closest) DCs, LDAP servers, KDC, and Kerberos services are located.

Data ONTAP automatically detects a defined Active Directory site and allows the NetApp system to identify its own site in that environment. A server's site depends on what subnet it is on. A NetApp device determines its own site membership by making LDAP queries based on its IP address and subnet mask.

5.8 Time Synchronization in Active Directory and Kerberos Domain

In order for Kerberos authentication to work in the Active Directory environment, the storage system's time must be synchronized with the Windows domain time. You must configure the time daemon on the NetApp storage system, and make sure of close synchronization with the domain before running the cifs setup on the NetApp storage system. The system won't join the Active Directory domain if the clock settings are more than 5 minutes apart. Second, if the time synchronization is not enabled, and if the system's time drifts by more than 5 minutes from the domain's time, the storage system's connection to the domain controllers fails and therefore client authentication attempts will fail too.

NetApp systems support two time service protocols:

- TIME protocol (also known as rdate) is specified in the RFC 868 standard. This standard allows time services to be provided on TCP or UDP port 37. The storage system uses only UDP port 37.
- Simple network time protocol (NTP) is specified in the RFC 2030 standard and is provided only on UDP port 123.

When your storage system has the `timed.enable` option set to on and a remote protocol (rdate or ntp) is specified, the storage system synchronizes to a network time server.

If the `timed.enable` option is set to off, your storage system is unable to synchronize with the network time server using NTP. The `rdate` time protocol can still be used by manually issuing the `rdate` command from your storage system console.

Recommendations

- Establish a timed service in your organization and have all domain controllers and storage systems point to the same timed server.
- Configure the time daemon on the NetApp storage system:
  1. Specify a protocol by entering the following command:
     `options timed.proto ntp`
  2. Specify up to five time servers by entering the following command:
     `options timed.servers <FQDN of windows domain>`
     This way a domain controller is automatically used as a time server.
  3. Turn on the time daemon by entering the following command:
     `options timed.enable on`

4. Set the timed window for adding a random offset within five minutes of the actual time update/verification:
   `options timed.window 5m`
   This way not all the systems are talking to the time server at exactly the same time every hour.

- Set the `timed.enable` option to `on` in a cluster configuration.

Note: If no internal time server is available in the network, there is another popular external NTP time source called "Pool." For more information, see www.pool.ntp.org.

5.9 NetApp System's Netbios (Windows) Name

The storage system's NetBIOS name is the name by which it is known on the network and by which it is referred to by Windows clients. Select a name that follows the same naming convention you use in your Windows environment today for file servers. This name can be the same as the system's existing name.

Recommendations

- For ease of administration, NetApp recommends that you keep the same name of the storage system for both UNIX and Windows environments.
- The name can be no longer than 15 characters. If you are installing a cluster, the host name must be unique for each storage system in the cluster.
- Some customers use nbaliases and DNS aliases after they migrate several file servers to one NetApp system. The following considerations must be taken into account in such situations:
  - Remove all old computer accounts; otherwise the alias does not work correctly.
  - Use `ADSIEdit` to add domain name suffixes (all aliases) to the computer object of the system. Otherwise Kerberos does not work for the aliases, and the authentication for the storage system fails back to NTLM.

Note: When you create a name for the storage system in an Active Directory domain, the NetBIOS name you select is appended with the DNS name.

5.10 Storage System Description in Active Directory

The NetApp storage system adds the description `Network Appliance Filer` to its Description field on the computer account. This aids in determining whether the machine account you're looking at in the `Users and Computers` MMC is an actual Windows Server or a NetApp storage system. It also aids in determining which version of Data ONTAP the NetApp system is running. You should leave the default description unless you must change it to suit your organization's needs.

Why This Is Important

This helps customer who wants to do a domain query for computer objects that have operating systems with Data ONTAP in them. If there are no specific organizational units for systems (lumped in with all servers), it can be difficult for administrators to know whether they are looking at a native Microsoft server or NetApp NAS storage systems.

After running the `cifs setup` command or on the restart after the Data ONTAP update, Data ONTAP attempts to update its operating system name and version in Active Directory. This information can be found under the Operating System tab of the Properties dialog box for the storage system's AD computer account object.

The NetApp system uses the LDAP interface to Active Directory to update its account information. The NetApp system authenticates to LDAP by logging in to LDAP using its own AD account and password. Consequently the system account needs permission to modify itself to update its OS name and version values. In some cases the update fails due to the default security settings on the system's AD computer account.

Recommendations

To allow the OS version update, the `SELF` permission entry should be adjusted on the storage system's AD object. Since Active Directory security settings and customer security needs differ from installation to installation, here are three examples of how to enable the OS update feature of Data ONTAP, from the most general to the most restrictive:

- Give SELF: Full Control on the storage system's AD object.
- Give SELF: Write on the storage system's AD object.
- Give SELF: Specific Property Permission for OS Updates on the storage system's AD object.

Note: There is no way to force the storage system to change its OS description in Active Directory except by:

- Terminating `cifs`, deleting the system account from AD, and rerunning `cifs setup`: OR
- Upgrading Data ONTAP

Note: Active Directory changes need time to propagate, so it is possible to make the security changes just described on one DC but not have them available at other DCs for some time, depending on AD propagation delays. In such cases it might appear that the security setting change has no effect if an upgraded system's first DC contact has not yet received the changes.

5.11 Storage System Computer Account Creation

Similar to users who require a valid account before being allowed to access a networked resource, workstations, servers, and other devices participating in an Active Directory domain must have a computer account. This provides a means for authenticating and auditing computer access to the network and access control, security, and management of domain resources. A NetApp storage system (like a Windows computer) can belong to only one domain and can have only one computer account defined for it in Active Directory.

Permissions Required to Create a Computer Account in the AD Domain

- By default, members of the Account Operators group can create computer accounts in the Computers container and in new organizational units.
- By default, Authenticated Users in a domain are assigned the Add Workstations to a Domain User right and can create up to 10 computer accounts in the domain.
- There are two additional ways to give a user or group permission to add a computer to the domain: Use a Group Policy object to assign the Add Computer User permission; or, on the organizational unit, assign the user or group the Create Computer Objects permission.

Computer accounts can be placed into one of two general places in your Active Directory domain:

- The default is a Windows Active Directory Computers built-in container object.
- You can specify an organizational unit (OU) for your AD layout.

Note: The default location for the storage system computer account (if no OU is defined) is the built-in Computers container object.

Note: If you are using Group Policy objects, do not place the NetApp storage object in any of the default Windows OUs, because Group Policy objects cannot be assigned to the default OUs.

There are two methods by which a system computer account can be created. Select one of the following methods for creating the appliance system account.

Create a Storage System Computer Account in AD Before Running CIFS SETUP

If your security structure does not allow you to assign the setup program the necessary permissions to create the storage system domain account, or if you intend to use Windows NT4-style authentication, you must create the storage system domain account before running `cifs setup`.

At a minimum, the following permissions are required at the OU level to enable the storage administrator to add the computer account through the `cifs setup` process without precreating the account:

- Change Password
- Write Public Information

Do Not Precreate a Storage System Computer Account

Allow `cifs setup` to create the system account automatically *during* the join process. Before adding a storage system to a Windows Active Directory domain, organizational unit (OU), or other Active Directory container object, you must make sure that the storage system administrator account has sufficient privileges and permissions to add a Windows Active Directory server to that domain or object.

The following permissions are required at the OU level to enable the storage administrator to add the computer account through the `cifs setup` process without precreating the account:

- Change Password
- Write Public Information
- Create Computer Objects

Recommendations

Do not precreate the storage system computer account in the Active Directory domain unless your security structure doesn't allow storage administrators to have sufficient privileges to create the computer objects in Active Directory.

5.12 Adding the Storage System to a Domain

The NetApp system can be added as a member server to:

- A Windows NT4 domain
- A Windows 2000 domain (mixed and native)
- A Windows 2003 domain
- A Windows 2008 or 2008 R2 domain

The process of doing this requires the administrator to run `cifs setup` and to supply answers to the questions asked. The following sections describe some of the requirements for `cifs setup`, as well as a checklist for adding the storage system to an Active Directory domain.

Domain Administrator Privileges Required

Domain\Administrator or equivalent rights are required when joining a storage system to an Active Directory domain. Data ONTAP does not use the same Kerberos calls that a Windows system uses. Microsoft publishes a private set of calls for its operating systems to use and another, "public," set for other vendors. This behavior shows up only when you attempt to add a storage system to a domain as a user who is not a member of the domain administrators group. If your organization's security structure does not allow you to use `Domain\Administrator`, there are workarounds to give specific rights and privileges to specific domain users or groups that your organization has authorized for joining Windows Servers to the domain. These additional rights and privileges are defined in section 5.11.

Removal of Old System Account

If you are moving the system from an old domain to a new domain, remove your system from the old domain before running `cifs setup`. To do this:

1. Run `cifs terminate` to stop cifs if it's already running.
2. When prompted during `cifs setup` whether you want to delete the account information, respond `Yes`.

5.13 CIFS Installation Checklist

The following checklist outlines the steps to install a NetApp storage system into an Active Directory domain. For detailed steps and information and guidance on joining the system to an Active Directory domain, see the Data ONTAP administrators' guides.

1. Determine mixed or native mode Active Directory domain style to understand your environment.
2. Accept the CIFS license.
3. Check the CIFS support matrix on the NetApp Support site.
4. Get the details of the Active Directory domain name and Host name/ NetBIOS name for NetApp storage.
5. Determine the storage system account location in AD (for example, Computers or other OU).
6. Precreate the storage system computer account if necessary; otherwise, skip this step.
7. Configure time services on the storage system.
8. Get the domain administrator account and password or other account with the privileges defined in section 5.11.
9. Make sure that DNS is configured properly for Active Directory support.
10. Determine the DNS server type (Microsoft or UNIX BIND).
11. Get the DNS domain name and DNS name servers' IP addresses.
12. Configure DNS on the storage system.
13. Make sure that there is an "A" record for the storage system in DNS; otherwise create an "A" record for the storage system in DNS.
14. Determine whether Microsoft sites are defined (the storage system determines this automatically). Check sites by using the `cifs domaininfo` command.
15. Remove the old storage system account from Active Directory (if necessary).
16. Confirm DNS, AD domain, and Microsoft site information by using the `dns info` command.
17. Run `cifs setup` on the system and answer all the prompts as appropriate.
18. Verify AD join by using the `cifs sessions` command.
19. Verify AD join by using the `cifs domaininfo` command.

5.14 Verifying Successful CIFS Installation

```
cifs sessions
```

To verify that cifs setup has successfully joined the storage system to your Active Directory domain, use the `cifs sessions` CLI command to validate:

- A successful registration into a Windows 2000 domain, including the domain name
- A list of WINS servers (if defined)
- A currently selected domain controller for Kerberos authentication

Example:

```
Netapp> cifs sessions
Server Registers as 'netapp' in Windows 2000 domain 'SOUTHFIELD' Filer is using en for
DOS users
WINS Server:  10.58.9.133
Selected domain controller \\DC1 for authentication
=======================================

PC  IP  (PC  Name)  (user)  #shares  #files
```

Note: CIFS sessions do not show smb and smb2 sessions separately. To list specific session information we must add the –p switch with smb or smb2 as parameter.

`cifs domaininfo`

To verify that `cifs setup` has successfully joined the storage system to your Active Directory domain, use the `cifs domaininfo <AD domain name>` CLI command to validate:

- The DNS is set up correctly on the storage system
- The Windows domain and domain type
- The storage system's Microsoft site information (if defined)
- The list of connected and available domain controllers (DCs)
- The connected AD LDAP server

Example:

```
netapp> cifs domaininfo southfield.lab NetBios Domain: SOUTHFIELD
Windows 2000 Domain Name: southfield.lab Type: Windows 2000
Filer AD Site:  labsite
```

The output indicates that the storage system can see the Windows 2000 domain `southfield.lab` and has detected a site `labsite`. This is what normal output should look like.

Errors

An error message indicating that the domain you specified could not be found in DNS probably indicates a DNS misconfiguration on the system (section 5.6). At this point, do not continue with `cifs setup` because it will fail. Instead, focus on resolving the DNS issue.

Recommendations

- Make sure that you have a domain administrator account or an account with similar privileges and password information, or arrange to have the domain administrator available to enter the password during cifs setup.
- Follow the CIFS installation checklist in section 5.13.
- Answer all the `cifs setup` prompts appropriately.
- Do not configure WINS if it is not used in your environment to access the storage system as a file server.
- If you have licensed other protocols as well, select Multiprotocol Filer.
- Let `cifs setup` create the default `/etc/passwd` and `/etc/group` files for the multiprotocol environment.
- If you want to change the CIFS server name, you can.
- Select the appropriate authentication mode, usually the Active Directory domain authentication.
- Always enter the fully qualified domain name (FQDN) of the Active Directory domain.
- Configure the time service through this process, if it was not configured earlier.

- Use the appropriate admin account with sufficient privileges to join the AD domain.
- Create a local administrator account.

6 Trusts Between Domains

A NetApp storage system can belong to only one Active Directory domain at a time (unless you are using MultiStore®). Access from other domains is accomplished through trusts that are put in place explicitly between domains or by virtue of automatic transitive trusts in Active Directory.

A trust is essentially a mechanism that allows resources in one domain to be accessible to authenticated users from another domain. Domain trusts in Windows NT 4.0 are one way, not transitive, which requires numerous multiple-trust relationships. Active Directory introduced the concept of transitive trusts as automatic two-way trusts that exist between domains in an Active Directory forest. The transitive trust relationships between parent and child domains are automatically established whenever new domains are created in the domain tree. These trusts connect resources between domains, and trusts flow from one domain to the other, as illustrated in Figure 3. In other words, if Domain A trusts Domain B and Domain B trusts Domain C, Domain A trusts Domain C. This new design in Active Directory greatly simplifies the trust relationships between Windows domains because it negates the need for multiple explicit trusts between each domain.

Figure 3) Transitive trusts.

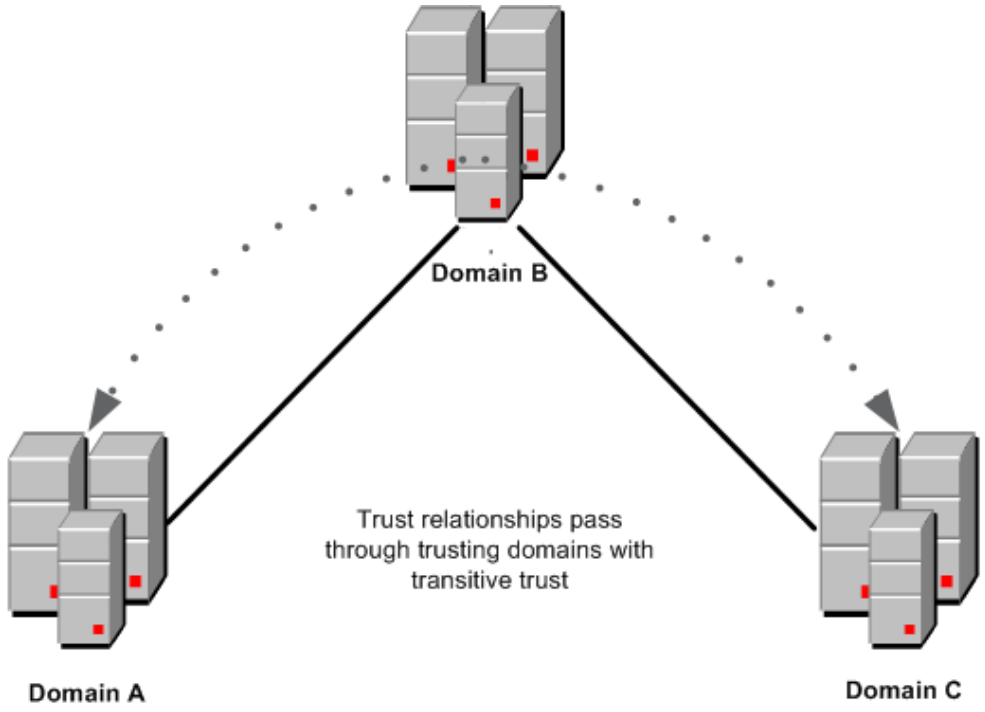

Cross-forest trusts in Active Directory are essentially two-way transitive trusts that exist between two disparate Active Directory forests.

For more information about managing the domain and forest trusts in a Windows 2008 domain, refer to Managing Domain and Forest Trusts.

A user's access to the NetApp storage system depends on whether the user's domain is trusted by the storage system's domain. Establish a trust relationship between the user's account domain and the storage system's domain if not part of any transitive trust.

Recommendations

The following best practices increase availability, make for trouble-free operations, or ease administration when you use them to administer domain and forest trusts.

- Optimize authentication speed in multidomain forests.
- When your forest contains domain trees with many child domains and you observe noticeable user authentication delays between the child domains, you can optimize the user authentication process between the child domains by creating shortcut trusts to midlevel domains in the domain tree hierarchy. For more information, see “When to create a shortcut trust” in “Understanding When to Create a Shortcut Trust” (http://go.microsoft.com/fwlink/?LinkID=107061).
- Keep a current list of trust relationships for future reference.
- You can use the Nltest.exe tool to display and record a list of these trusts. For more information, see Nltest Overview (http://go.microsoft.com/fwlink/?LinkID=93567).
- Perform regular backups of domain controllers to preserve all trust relationships within a particular domain.
- If you do not want any users from nontrusted domains to log in and access shares on the storage system, or if you want remote nontrusted users to be prompted for credentials, you can set the `cifs.guest_account` option on the storage system to null.
  This option enables a user to get access to the system, provided that either the system uses a domain controller for authentication and the user is not in a trusted domain, or the system uses the `/etc/passwd` file or the NIS password database for authentication and the user has no entry in the `/etc/passwd` file or the NIS password database. If this option is set to the name of an account in the password database, users logging into the system are assigned to the guest account if their names are not listed in the password database (when using `/etc/passwd` or NIS) or if the user is not from a trusted domain (when using a domain controller). The configured user name is used for the UNIX user ID, group ID, and group set of the specified account. If the option is a null string, guest access is disabled. The default value for this option is a null string.
- Consider MultiStore if a single NetApp storage system needs to join multiple AD/DNS/NIS domains.
- Place the Global Catalog servers of one domain close to the NetApp systems in another domain.

7 CIFS Shared Folders

NetApp storage systems allow you create multiple CIFS shares for corporate data, application data, or any other data. Creating CIFS shares on a NetApp system is a straightforward process that can be accomplished by any of the following methods:

- Using NetApp System Manager wizards
- Using the Data ONTAP command line
- Using the FilerView® administration tool
- Using the Shared Folders MMC snap-in in pre-Windows 2008 and Share and Storage Management in Windows 2008 and later

Access to the shares can be controlled through Share permissions, but to restrict the shares’ enumeration we have few options available on Data ONTAP.

To disable enumeration of a share by browsers:

```
cifs shares -change sharename -nobrowse
```

To control anonymous CIFS share lookups:

```
options cifs.restrict_anonymous <Value>
```

- 0 – No restriction
- 1 – Restrict enumeration of users and shares
- 2 – Fully restrict anonymous share lookups

Permissions and Shared Folders

- Assign appropriate permissions to limit root volume access.
- Assign permissions on shares appropriately. Sometimes the permissions at the share level are defined as less restrictive and permissions at the file/folder level (NTFS ACLs) are defined as more restrictive.
- Set NTFS permissions (ACLs) on directories to prevent unauthorized users.
- Assigning permissions to groups simplifies management of shared resources, because you can then add users to or remove them from groups without having to reassign permissions. To deny all access to a shared resource, deny the Full Control permission.
- Assign the most restrictive permissions that allow users to perform required tasks. For example, if users need only to read information in a folder and they will never delete, create, or change files, assign the Read permission.
- Organize resources so that objects with the same security requirements are located in the same folder. For example, if users require the Read permission for several application folders, store the application folders in the same parent folder. Then share the parent folder, rather than share each individual application folder. Note that if you must change the location of an application, you may need to reinstall it.
- When you share applications, organize all shared applications in one folder. This simplifies administration, because there is only one location for installing and upgrading software.
- Use centralized data folders.

Recommendations

- Carefully examine all the shares and share names in the existing Windows file serving environment.
- Create identically named shares to match your existing environment.
- Make sure that the share names you create are unique to a file server, especially when you migrate data from existing file servers to NetApp storage systems.
- Assign permissions to groups, not to user accounts. This simplifies the management of shared resources, because you can add users to or remove them from groups without having to reassign permissions.
- Assign the most restrictive permissions that still allow users to perform required tasks. For example, if users only need to read information in a folder and they will never delete, create, or change files, assign the Read permission.
- Organize resources so that objects with the same security requirements are located in the same folder. For example, if users require Read permission for several application folders, store the application folders in the same parent folder. Then share the parent folder, rather than share each individual application folder. Note that if you must change the location of an application, you may need to reinstall it.
- Organizing all shared applications in one folder simplifies administration, because there is only one location for installing and upgrading software.
- In most cases, do not change the default permission (Read) for the Everyone group. The Everyone group includes anyone who has access to network resources, including the Guest account. In most cases, do not change this default unless you want users to be able to make changes to the files and objects in the shared resource.
- Grant access to users by using domain user accounts.
- Using centralized data folders, you can easily manage resources and back up data.

- Use intuitive, short labels for shared resources, so that the shared resources can be easily recognized and accessed by users and all client operating systems.
- The recommended value for Anonymous lookup is 1 or 2, depending on the security policies.

8 DFS Integration

Distributed File System (DFS) is Microsoft's implementation of a global namespace. A namespace is a virtual view of the shared folders in an organization. A DFS namespace enables you to group shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of subfolders. This structure increases availability and automatically connects users to shared folders in the same Active Directory Domain Services site, when available, instead of routing them over WAN connections.

The major components of DFS namespace are namespace server, namespace root, folders, and folder targets. Namespace server hosts a namespace; it could be a member server or a domain controller. Namespace root is the starting point of the namespace. Folders build the namespace hierarchy. Folders can optionally have folder targets. When users browse a folder with targets in the namespace, the client computer receives a referral that directs the client computer to one of the folder targets. A folder target is a UNC path of a shared folder or another namespace that is associated with a folder in a namespace.

For more information about DFS configuration, refer to "DFS Step-by-Step Guide for Windows Server 2008".

There are two types of namespaces; you can choose one of these based on the following guidelines:

- Use standalone namespace:
  - If you do not use Active Directory Domain Services (AD DS)
  - If you want to increase the availability of the namespace by using Microsoft's failover cluster
- Use domain-based namespace:
  - If you want to make sure of the availability of the namespace by using multiple namespace servers
  - If you want to hide the name of the namespace server from users (choosing a domain-based namespace makes it easier to replace the namespace server or to migrate the namespace to another server)

In addition to the namespace type, you must choose the namespace mode, either Windows 2000 Server mode or Windows Server 2008 mode.

NetApp systems can be part of the DFS namespace as "leaf nodes," and the CIFS shares can be the targets of DFS links. A NetApp system cannot act as a DFS root and does not currently support DFS-R.

Recommendations

- Choose Microsoft DFS in Windows Server 2008 mode, because it is much more enhanced than the previous generations of DFS. It supports Access Based Enumeration, provides increased scalability, and supports unlimited folders with targets.
  The option to enable Access Based enumeration for shares on the controller is
  `cifs shares –change sharename –accessbasedenum`
- You can use DFS to stitch the shares from multiple NetApp aggregates to make up a single large global namespace as a workaround for the 16TB size limit of aggregates.
- For more DFS best practices information, refer to "Best Practices for Distributed File System."

9 Home Directories

The Data ONTAP Auto Home Directory feature enables you to create users' home directories on the NetApp storage system and automatically offer a dynamic share for each user's home directory, without creating an individual CIFS share for each. Users can then connect to their own home directory matching with their user name without seeing other users' home directories. From a CIFS client perspective, the home directory works the same way as any other share to which the user can connect.

Data ONTAP offers the share to the user with a matching name. The user name for matching can be a Windows user name, a domain name followed by a Windows user name, or a UNIX user name. Home directory names are not case sensitive.

This feature is especially useful in large enterprises to spread the users into different volumes (to control volume space) and to lower the number of CIFS shares. Think about systems with 20,000 CIFS shares and their impact during a cluster takeover.

When Data ONTAP tries to locate the directories named after the users, it searches only the paths that you specify. These paths are called home directory paths, and they can exist in different volumes. You can specify multiple home directory paths in `/etc/cifs_homedir.cfg`. Data ONTAP searches the paths, in the order you specify, until the directory that matches the user name is found.

The following differences exist between a home directory and other shares:

- You cannot change the share-level ACL and the comment for a home directory.
- The `cifs shares` command does not display the home directories.
- The format of specifying the home directory using the Universal Naming Convention is sometimes different than that for specifying other shares.

9.1 How Data ONTAP Matches a Home Directory with a User

You can specify the naming style of home directories to determine how Data ONTAP matches a directory with a user. Here are the naming styles from which you can choose:

- **Windows name.** Data ONTAP searches for the directory whose name matches the user's Windows name.
- **Hidden name.** If the naming style is hidden, users connect to their home directories by using their Windows user name with a dollar sign appended to it (name$), and Data ONTAP searches for a directory that matches the Windows user name (name).
- **Windows domain name and Windows name.** If users from different domains have the same user name, they must be differentiated by using the domain name. In this naming style, Data ONTAP searches for a directory in the home directory path that matches the domain name. Then it searches the domain directory for the home directory that matches the user name.
  Example: To create a directory for `engineering\jdoe` and a directory for `marketing\jdoe`, you create the two directories in the home directory paths. The directories have the same names as the domain names (engineering and marketing). Then you create user home directories in these domain directories.
- **Mapped UNIX name.** If the naming style is UNIX, Data ONTAP searches for the directory that matches the user's mapped UNIX name from `/etc/usermap.cfg`.

If you do not specify a home directory naming style, Data ONTAP uses the user's Windows name for directory matching. This is the same style used by versions of Data ONTAP earlier than version 6.0.

9.2 Syntax for Specifying a Home Directory by Using a UNC Name

The convention for specifying a home directory when using UNC depends on the home directory naming style specified by the `cifs.home_dir_namestyle` option.

Table 3 lists UNC names, with examples, for each name style value.

Table 3) Home directory name style syntax.

Value of cifs.home_dir_namestyle	UNC Name
Ntname or ""	\\toaster\Windows_NT_name Example: **\\toaster\jdoe**
Hidden	\\ toaster \Windows_NT_name$ Example: **\\toaster\jdoe$**
domain	\\ toaster \~domain~Windows_NT_name Example: **\\toaster\~engineering~jdoe**
mapped	\\ toaster \~mapped_name Example: **\\toaster\~jdoe**

If `cifs.home_dir_namestyle` is domain but the UNC name in the access request does not specify a domain name, Data ONTAP assumes the domain to be the domain under which the request is sent. If you omit the domain name in the access request, you can also leave out the tilde (~) before the user name.

Example: A user named jdoe is logged in as engineering\jdoe from a PC in the engineering domain. When he tries to access his home directory by using his user name in the marketing domain, he can enter either of the following commands to request access:

```
net use * \\toaster\~jdoe /user:marketing\jdoe
net use * \\toaster\jdoe /user:marketing\jdoe
```

Recommendations

- If a CIFS share for every user's home directory is required, then NetApp strongly recommends using the NetApp CIFS Homedir feature instead.
- You can also add the CIFS home directory paths to users' profiles in the Active Directory to automatically map to a drive letter when they log on.
- Use one or more qtrees or separate volumes to contain homedirs only.
- Create an `/etc/cifs_homedir.cfg` file to contain the homedir paths. You can use the `cifs homedir` command to display the current list of directory paths.
- If you want Data ONTAP to match the home directories by the Windows user names, keep the `options cifs.home_dir_namestyle` as `ntname`.
- If the naming style is hidden, users must enter their user name with a dollar sign appended to it (for example, name$) to attach to their home directory.
- Create folders in the homedir paths matching the exact user names and apply suitable permissions and ownerships of the users.
- Connecting to your own CIFS home directory by using either the "cifs.homedir" or "tilde" (~) share alias can be useful when you are writing scripts.

  Example:

```
net use * \\toaster\cifs.homedir net use * \\toaster\~
```

  You can enable wide symbolic links from a share if you want to allow CIFS clients to follow absolute symbolic links to destinations outside the share or storage system. By default, this feature is disabled. To enable the wide symbolic links from a share, use the following CLI command:

```
cifs shares -change sharename -widelink
```

10 Types of Security Groups (Local and Global)

Groups are used to organize individual user or computer accounts. Mainly they are used for security purposes. Best practices dictate that most of your security administration should be done using groups. Groups simplify administration by allowing you to assign appropriate permissions to a group of users rather than to each user account individually. When you add a user to an existing group, the user automatically gains the rights and permissions already assigned to that group. NetApp also recommends that you give each group a name that describes the group's function or purpose (for example, naming a group MrktInfo if the people in it are given access to marketing information).

The groups of greatest interest are the security groups, which are mainly used to manage user and computer access to shared resources and to filter Group Policy settings. For instance, if you want to grant users permissions to a network share, you would create a group, grant the group appropriate permissions to the share, and then add the users (or other groups) as members of that group. Windows 2000 and later operating systems provide the following types of domain security groups; the security group called "Special Groups" was introduced in Windows 2008.

- Built-in local groups
- Domain local groups
- Global groups
- Universal groups
- Special groups

10.1 Built-In (Nondomain) Local Groups

These groups are not the same as domain local groups. A local group is a collection of user accounts on a computer. A local group has only machine-wide scope; that is, it can be used to grant resource permissions only on the machine (NetApp system) where it exists. (Note: Local groups created on a domain controller are available on every domain controller in that domain and can be used to grant resource permissions on any domain controller in that domain.) You can add local users, global users, and global groups to local groups. However, you cannot add local users and groups to a global group.

Built-In Local Groups on a NetApp System

Membership: You can define a local group on the NetApp storage system that consists of users or global groups from any trusted domains. You should not create your own local groups on member servers or the storage system, because that would make future migrations and server consolidations more difficult when transferring the proper security (SIDs) of the local group.

You can add the admin users from the domain to a built-in administrators group of the NetApp system for specific administration purposes or for applications that require nondomain authentication (SnapDrive, SnapManager, and so on). When the storage system joins an Active Directory domain, then the domain's primary local groups like Domain Admins and Backup Operators groups are automatically added to the respective local groups of the storage systems. This is done by design, but it is not necessary.

Permissions: Members of a local group can be given access to files and resources on that system. Membership in certain well-known local groups confers special privileges on the storage system. For example, members of BUILTIN\Power Users can manipulate shares, but they have no other administrative capabilities.

Local users and groups are an important security feature because you can limit the ability of users and groups to perform certain actions by assigning them the rights and permissions. For instance, you may want to authorize a user to perform certain actions on a computer, such as backing up files and folders.

You could place a global group of users with backup duties in the default local Backup Operators group on the system. This would give that set of users the right to back up and restore files on the local system regardless of the permissions on the individual files.

Recommendations

The following are guidelines for creating local user groups and their limitations compared to global user groups:

- Use local groups on computers that do not belong to a domain (for example, a workgroup).
- Use local groups only on the computer on which they are created; this makes them very inflexible.
- Although local groups are available on member servers and systems, do not use local groups on computers or systems that are part of a domain. Doing so prevents you from centralizing group administration, which is a primary reason for choosing a domain structure.
- Local groups do not appear in the Active Directory service, and you must administer them separately for each computer.
- You can assign permissions to local groups to access only the resources of the computer on which you create the local groups.
- You cannot create local groups on domain controllers, because domain controllers cannot have a security database that is independent of the database in Active Directory.
- Local groups cannot belong to any other group.
- Do not create any local groups on your own; instead, use the default ones.

Note: Whenever the SID of the storage system changes (migration, `cifs setup` reinvoke), all local group ACLs on the file systems are affected. The ACLs can be changed by using the subinacl.exe tool from Microsoft; or the local groups can be imported in the new storage system by using the following command: `useradmin domainuser load`.

10.2 Domain Local Groups

Do not confuse this type of group with regular “built-in local groups” (those that reside locally on a Windows Server or NetApp system).

Membership: These groups can reside only inside a single domain (cannot cross domains) and only on domain controllers. However, the members in the group can be from any domain, although this is not used much because domain local groups are not visible outside their own domain. Membership of the domain local group is controlled by the administrators where the resource is located, not where the users are, which puts it in line with how administration is typically done.

Permissions: A domain local group has domain-wide scope; that is, it can be used to grant resource permissions on any Windows 2000 machine within the domain in which it exists (but not beyond its domain). Since a domain local group is associated with an access token built when a member of that group authenticates to a resource in that domain, any unnecessary network traffic (carrying of membership information) is avoided. If, instead, you assigned global group permission to access the resource, the global group can end up in a user's token anywhere in the forest, causing unnecessary network traffic.

Recommendations

Use domain local groups when you want to grant permissions to resources that exist within a specific domain. Domain local groups are useful in a multiple domain where you add a global group to a domain local group so you can access resources in another domain. Visit this link for more information about parent and child domains.

10.3 Global Groups

Groups with global scope help you manage directory objects that require daily maintenance, such as user and computer accounts. These are most often used to organize users who share similar network access requirements.

Membership: Global groups can have members from within their own domain only, and they are visible to all trusted domains as well. They can include only user and other group accounts in the same domain. Global groups cannot contain local groups or other global groups and they are not assigned to local resources.

Permissions: These groups can grant permissions to objects located in many domains. Resources are assigned by placing global groups within local groups on Windows NT workstations or standalone servers.

The benefit of using global groups is that you can, on the domain level, assign users to a global group and add the entire group to a local group already on a local computer. In other words, an administrator can change the domain user's global group (for example, when a new hire comes in) without having to reset any permissions on a local workstation or server.

The appealing thing about global groups is that they can be nested to allow overlapping access or to scale for very large group structures. The most convenient way to grant access to global groups is by making the global group a member of a resource group that is granted access permissions to a set of related project resources. This can simplify administration immensely. Using nested groups in Windows 2000 or 2003 servers alleviates the old 5,000-member limit, and, more importantly, it allows you to apply "most restrictive" and "most inclusive" nesting strategies to make administering resources easier. An example of this is a situation in which all members of an organization need access to the same resource, such as employee information. Not all members have authority to view all employee information, such as salary grade, yet most need access to office phone numbers and e-mail addresses. Simply create two global groups with different access rights and members and then nest them. The effect is a blend of the two that gives you the desired security and an easier administration model. There is no limit to the levels of nesting that can be applied, but tracing permissions can become problematic and some applications, such as backup, can drill down only so far.

Recommendations

- NetApp highly recommends that you implement global groups to contain users in a particular domain, especially if there is only one Active Directory domain. For environments with multiple domains, you would still employ global groups, but you would include them in universal groups.
- Groups created in Active Directory should generally be global groups, especially for user accounts.
- Global groups or domain local groups are also listed in the global catalog, but their individual members are not listed. Using these groups thus reduces the size of the global catalog and reduces the replication traffic needed to keep the global catalog up to date. Therefore you should use global groups or domain local groups if the group membership changes frequently.

10.4 Universal Groups

Universal groups are similar to global groups, but with a larger scope.

Membership: Universal groups can contain user and group accounts (universal or global groups) from any trusted domain in the forest. Universal security groups are not available in mixed mode, only in native mode.

Permissions: Universal groups can be granted permissions to related resources in any domain, including domains in other forests with which a trust relationship exists.

By default, the NetApp system includes membership in nested groups and membership in universal groups from other domains in the forest. This can be controlled by using the following option:

```
filer>  options  cifs.universal_nested_groups.enable
```

When `cifs.universal_nested_groups.enable` is off, the system does not include membership in nested groups or membership in universal groups from other domains in the forest. The default is on. This option is pertinent to all NFS clients accessing a file or directory with Windows security; it does not affect CIFS clients. Changes to the option take effect immediately, affecting all new NFS connections thereafter; it is not necessary to restart the CIFS. This option will be deprecated in a future release, when the NetApp system will always include these memberships.

Caution: All group memberships are fetched from Active Directory only when the user and the system are in the same domain tree, or when the user's domain tree has a two-way transitive trust with the system's domain tree.

Recommendations

- Although it is possible to create universal groups in the Active Directory with any domain structure, it is generally not required or recommended for single domain structures. Instead, you should employ global groups because they use fewer resources.
- Use universal groups in a multidomain environment. These groups can help you represent and consolidate groups that span domains. It helps to build groups that perform a common function across an enterprise.
- Although few organizations choose to implement this level of complexity, you can add user accounts to global groups. You nest these groups within universal groups, and then make the universal group a member of a domain local (or machine local) group that has access permissions to resources. With this strategy, any membership changes in the global groups do not affect the universal groups.
- Designate widely used groups that seldom change as universal groups. Universal groups and all of their members are listed in the global catalog. Whenever one member of a universal group changes, the entire group membership must be replicated to all global catalogs in the domain tree or forest. Therefore, if you use universal groups, use them in situations in which the membership of the group does not change frequently.

10.5 Special Groups

Special Groups is a new feature in Windows Vista and in Windows Server 2008. The Special Groups feature lets the administrator find out when a member of a certain group logs on to the computer. This feature also enables the administrator to set a list of group security identifiers (SIDs) in the registry. An audit event is logged in the Security log if the following conditions are true:

- Any of the group SIDs is added to an access token when a group member logs on.

Note: An access token contains the security information for a logon session. Also, the token identifies the user, the user's groups, and the user's rights.

- In the audit policy settings, the Special Logon feature is enabled.

For more information about the Special Groups feature, refer to this Microsoft article: "Description of the Special Groups Feature in Windows Vista and in Windows Server 2008."

10.6 Security Group Recommendations

- Microsoft recommends the following procedure for granting permissions across multiple domains using the security groups:
- Create a global group in each domain and add the appropriate users as members. Create a universal group and grant the appropriate permissions.
- Add the global groups as members of the universal group.
- If you have only a single domain structure, simply use global groups (may be domain local) and not necessarily universal groups.

- In general, object permissions should be assigned to domain local groups, whereas user and computer accounts should be placed in global groups. Then these global groups can be (not necessarily) placed or nested in domain local groups to gain access to network resources. Note, however, that in order to place the global group in a domain local group, you must be running a native-mode domain.
- If you are going to use domain local groups, use the built-in defaults and create new ones if your environment requires it. The default ones should satisfy most of your requirements. One example is a global group of users who have higher security requirements. You would add this global group to the domain local Administrators group.

11 Security

Security is the key factor to protect data from unauthorized access. This can be achieved through setting different options at different levels. We can broadly categorize them as Communication-Level Security, Storage-Level Security, and File-/Folder-Level Security. Each category adds an additional layer to the total data security.

11.1 Communication Security

Communication security plays a key role in the data transfer between machines. This is to determine the challenge/response authentication protocol to be used while negotiating and what kind of security we can apply on the data packets transferred between the machines.

Authentication Protocol NTLM, NTLMv2, or Kerberos

LMCompatibilityLevel

NTLMv2 or Kerberos is used as an authentication protocol in Windows 2000 and later operating systems. In case of the following scenarios, NTLM is used as the authentication protocol.

- Authenticating against a Windows NT 4.0 domain controller
- The client is authenticating to a server using an IP Address.
- When the remote system is in a workgroup not in the domain

Recommendations

For secure communication, set the LMCompatibilityLevel option on both server and client, depending on the organization's requirement. To configure the setting on the controller, use the following options entry: `options cifs.LMCompatibilityLevel <Value>`, where the value is one of the following:

1 = Accept LM, NTLM, NTLMv2 session security, NTLMv2, Kerberos (default)

2 = Accept NTLM, NTLMv2 session security, NTLMv2, Kerberos

3 = Accept NTLMv2 session security, NTLMv2, Kerberos

4 = Accept NTLMv2, Kerberos

5 = Accept Kerberos only

For LMCompatibilityLevel values on Windows, refer to http://technet.microsoft.com/en-us/library/cc960646.aspx.

The default values of LMCompatibilityLevel on Windows are listed in Table 4.

Table 4) Default LMCompatibilityLevel values for Windows.

Operating System	Default Value
Windows XP	0
Windows 2003	2
Windows Vista/2008	3
Windows 7/2008 R2	3

SMB Signing

SMB Signing is a feature through which communications using SMB can be digitally signed at the packet level. Digitally signing the packets enables the recipient of the packets to confirm their point of origination and their authenticity. This security mechanism in the SMB protocol helps avoid issues like tampering of packets and "man-in-the-middle" attacks.

To set SMB signing on the controller, use the following options entry:

```
options smb.signing.enable on
```

LDAP Signing and Sealing

Signing LDAP traffic verifies that the packaged data comes from a known source and that it has not been tampered with.

To enable LDAP signing, use the following options entry

```
options ldap.security.level <Value> where:
```

0 = No signing, SASL bind is used.

1 = Use LDAP signing (Default).

2 = Use LDAP signing and sealing. LDAP queries and responses are encrypted.

11.2 Storage-Level Security

Establishing a storage-level policy is an important facet of security. Monitoring the creation or modification of objects gives you a way to track potential security problems, helps to promote user accountability, and provides evidence in case of a security breach.

Auditing

Establishing an audit policy is an important facet of security. Monitoring the creation or modification of objects gives you a way to track potential security problems, helps to promote user accountability, and provides evidence in case of a security breach.

Data ONTAP supports auditing of logon, logoff, system events, and file access events similar to Windows. When you configure Data ONTAP for auditing, the event log file and the settings for all options persist across a reboot or when CIFS is terminated or restarted. Data ONTAP can audit file access events for both CIFS and NFS access. You can view the auditing events logs in a standard (EVT) format from a Microsoft Event Viewer. Backing up and dumping of log files can be controlled based on log size or periodicity.

For more information about the configuration of CIFS and NFS auditing in Data ONTAP, refer to NetApp technical report TR-3595, "File Access Auditing on NetApp Controller."

Recommendations

- **Create an audit plan before implementing the auditing settings.**
  - Decide what type of information you want to gain by collecting audit events. If you are interested in intrusion detection—tracking the attempts of users to gain access to areas for which they are not authorized—you can collect failure audits. If you are interested in forensics—using the audit log to determine exactly what happens in your organization—you can collect a combination of success and failure audits.
  - Consider the resources that you have available for collecting and reviewing an audit log. Audit events take up space on your storage system, and they take up your time and the time of people in your organization. Don't audit events that don't really interest you.
  - Specify the categories of events that you want to audit. Examples of event categories are user logon, user logoff, and account management. The event categories that you select constitute your audit policy.
- **Being specific is the key to reducing performance impact.**
  - Define the specific object to audit: specific users, specific access events, specific types of success or failure.
  - Determine which objects you want to monitor access to and what type of access you want to monitor; for example, any open request to a particular file from a specific user and whether you want to audit both success and failure for this access event.
  - For optimal system performance, minimize the number of entries in the SACL for an object. One entry in a SACL that contains 1,000 users does not degrade system performance as much as 1,000 separate entries.
  - To reduce the volume of events that is generated and to maximize the effectiveness of each event, audit only the actions that really interest you. For example, if you are interested in users reading a file, don't audit Full Control.
- **Set an appropriate size and recycling time interval for the security log.**

  It is important to configure the size of the security log appropriately, based on the number of events that your auditing policy settings generate. You can view the security log with Event Viewer.
  - Configure for automatic saving of the event logs. Data ONTAP allows you save the audit log files automatically, based on the size of the file or a time interval.

    `options cifs.audit.autosave.onsize.threshold Nsuffix` where *`N`* is the value of the size threshold.

    `Suffix` is the unit of measure percentage (%), kilobytes (k), megabytes (m), or gigabytes (g).

    `options cifs.audit.autosave.ontime.interval Nsuffix` where *`N`* is the value of the time interval.

    `suffix` is the unit of measure—seconds (s), minutes (m), hours (h), or days (d).
  - You can also configure different file extensions for saving the event logs.

    `options cifs.audit.autosave.file.extension` *`counter`*

    Examples: `eventlog.evt, eventlog1.evt, eventlog2.evt` and so on.

    `options cifs.audit.autosave.file.extension` *`timestamp`*

    Format: `base name of event file YYYYMMDDHHMMSS.evt`
  - Data ONTAP allows you to keep a maximum of 1,000 event log files with the automatic saving option, but you can keep on archiving the event log files to some other location or to a different file system for long-term retention.

Storage-Level Access Guard

Beginning in Data ONTAP 7.2.2, an additional layer of security, called Storage-Level Access Guard, was introduced on the storage object level. With this feature, storage administrators can set security (permissions and auditing) on volumes and qtrees by using the `fsecurity` console command. This security provides a not-to-exceed security and/or auditing capability that applies to all objects in a given volume or qtree. This new security level cannot be set or modified by normal CIFS or NFS administrative clients, because it is invisible to a Windows client or a UNIX host. This prevents such clients from overriding policies that should be managed at the storage level, not the protocol level.

Access to a file or directory in Data ONTAP is determined by the combined effect of both the native permissions applied to files and/or directories and the Storage-Level Access Guard permissions set on qtrees and/or volumes. For CIFS/NFS client access, three levels of security checks are performed to determine effective permissions. The checks are performed in this order:

20. Storage-Level Access Guard permissions
21. CIFS share or NFS export-level permissions
22. NTFS file/folder access control lists (ACLs) or UNIX-mode bits

Note: All accesses must pass all levels of security checks.

For details about configuring this feature, refer to TR-3596, "Storage-Level Access Guard Quick Start Guide."

Recommendations

Storage-Level Access Guard provides additional security at the storage level. It helps in separating the role of the storage administrator from that of the system administrator. This security cannot be revoked by any users or administrators from their desktops. This feature makes it most commonly used in the following scenarios:

- Intellectual property protection by auditing and controlling all user access at the storage level
- Storage for financial services companies, including both banking and trading groups
- Government services with separate file storage for individual departments
- Universities, to protect student files

Make sure that your organization defines storage administrator accounts differently from Windows or UNIX administrator accounts.

Be sure to select the Storage option from the secedit tool for configuring the Storage-Level Access Guard.

11.3 File-Level Security

This is a more granular level of security that defines the permissions at the folder or file level, either by setting the ACLs on the files and folders or by blocking the unauthorized enumeration of the folders or shares.

Fsecurity

Beginning with Data ONTAP 7.2.2, storage administrators have `fsecurity`, a Data ONTAP console command, by which administrators can apply NTFS security (permissions and auditing) over large directories. This tool significantly reduces the time to apply permissions on large directories because security settings are managed locally on the storage system, not from remote clients. In addition, storage administrators can set security on many files and directories at once using the same command. Examples of this usage include:

- File storage for large enterprise environments such as home directories

- Data migration
- Changing of Windows domain for NAS

For more details on configuring security using the `fsecurity` CLI tool, refer to TR-3597, "Bulk Security Quick Start Guide."

Recommendations

- Use the `fsecurity` command locally on the NetApp system console to apply permissions on large directories.
- Use the `fsecurity` command to display the effective permissions on the NetApp system console, especially in the case of mixed-mode permissions.
- Use the `fsecurity` command to reapply ACLs after data migration.
- Use the `fsecurity` command to set and display Storage-Level Access Guard.
- Be sure to select the correct Apply To options when you apply permissions to "this folder, subfolders and files" and so on.

Access-Based Enumeration

Data ONTAP 7.2 and later releases provide storage system support for access-based enumeration, a shared-resource security feature introduced in Microsoft Windows Server 2003 Service Pack 1 as an add-on and then as a built-in feature in Windows 2008 and later. This feature enables administrators to control the display of files and folders according to a user's access rights.

Conventional share properties enable you to specify which users (individually or in groups) have permission to view or modify shared resources. However, they don't allow you to control whether shared folders or files are visible to users who don't have permission to access them. This could pose problems if the names of shared folders or files describe sensitive information, such as the names of customers or new products under development.

Access-based enumeration (ABE) extends share properties to include the enumeration of shared resources. When ABE is enabled on a CIFS share, users who do not have permission to access a shared folder or file underneath it (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment. ABE therefore enables you to filter the display of shared resources based on user access rights.

End users see only the files and folders for which they are responsible, rather than spending time looking through lists of inaccessible folders and files. Administrators can be more productive because they don't have to help less-skilled users navigate through dense shared folders. In the Data ONTAP implementation of ABE, almost no performance effect is observed.

ABE for a CIFS share on a NetApp storage system can be managed by the CIFS shares option:

```
[-accessbasedenum | -noaccessbasedenum]
```

ABE can also be set by the abecmd.exe CLI from a Windows system for a CIFS share on a NetApp system:

```
abecmd [/enable | /disable] [/server <servername>] {/all | <sharename>}
```

Recommendations

- Use the ABE feature to protect or hide the sensitive information in the shared folders. This is very useful for the small file sharing environment, especially for a workgroup environment.
- Assign the user permissions carefully in order to give users access to only those folders they need.

- You should use this feature at the top-level home directory share, so that if users connecting to their home directories are trying to browse the home dir share to see other users' directories, they won't be able to see them.
- The Data ONTAP ABE feature is very useful with migration from Novell, where ABE was already in use. In such cases, NetApp recommends enabling ABE on NetApp file shares.

12 Group Policy Objects (GPOs)

Group Policy is a technology that enables administrators to precisely define the configuration of the users' computing environment. GPOs contain settings that are applied to users, computers, and other organizational units (OUs) to define what a system looks like and how it behaves for a defined group of users or computers. These can include configuration settings for such things as security and application deployment, for example. GPOs are generally used to control both user- and machine-based configuration settings.

GPOs maintain the configuration of possibly thousands of workstations in an environment in a consistent manner. By providing a well-managed desktop environment through group policies, Windows administrators can ease the resolution and elimination of change and configuration management issues.

Group policies should be used throughout Windows 2000 and 2003 to define user and computer configuration settings such as scripts, software policies, security settings, application deployment, user settings, and document options.

Recommendations

- Make use of users' and computers' GPOs, especially in large environments.
- Apply GPOs at the OU level in your Active Directory structure. This makes it much easier to delegate administration and have more granular control over which policies get applied where.
- Use groups to filter the policies' effects and scope. Instead of creating multiple OUs or sub-OUs with differing group policies, it's best to apply one group policy to an OU that applies to a group of users. For instance, if you have an OU that contains all users and you want to create a group policy for just the group "engineers," then you would create the policy to apply only to the engineers. This allows you to have better control of group policy scope within groups.
- You cannot apply group policy settings to the default Users and Computers containers. These are containers, not OUs, so don't be confused.
- The Default Domain Policy and the Default Domain Controllers Policy should not be changed for most environments. However, you should change the following settings to match the security requirements of your organization:
  - Default Domain Policy

    Password Policy settings

    Account Lockout Policy settings − Kerberos Policy settings
  - Default Domain Controllers Policy

    User Rights Assignment settings

12.1 GPO Support in Data ONTAP

Data ONTAP supports Microsoft GPOs and has the ability to accept GPOs directly from Active Directory by using Microsoft LDAP GPO APIs. Most GPOs mandated in a Microsoft environment to control users and computers are not valid or necessary for a NetApp storage system. Consider the following points:

- NetApp systems don't run applications.
- NetApp systems don't offer logon sessions for users other than the specifically defined administrators.

- NetApp systems are impervious to viruses.
- There is no concept of a Windows registry on a NetApp system that can be compromised by a casual user.
- Locking down a NetApp system is done using permission set to telnet in or administer by using a Web GUI. The reasons for locking down a Windows Server don't apply to a NetApp system in most cases.
- Any GPOs for users and computers used in an AD domain are transparent to the NetApp storage system, as described in the next two sections.

GPOs Set on Users or Groups

NetApp systems are not in the path for GPOs on domain users, so GPOs set on domain users look transparent to the system.

GPOs Set on Computers

NetApp systems are not in the path for GPOs on computers, so GPOs set on computers look transparent to a system. Typically, administrators use Computers containers to control users' desktop or laptop settings through GPOs. Many of these GPOs are for managing Windows tasks—for example, to restrain the installation of software apps on a user's desktop and to update patches. None of these apply to a system.

The following GPOs are currently supported for NetApp storage systems:

- Start-up and shut-down scripts
- GPO refresh time interval for the computer
- File system security settings
- Restricted group security
- Event log support
- Auditing support
- User rights assignment
- GPO refresh time interval random offset

For more information on these supported GPOs and their descriptions, refer to NetApp TR-3367, "NetApp Storage Systems in Microsoft Windows Environment."

Recommendations

- You must explicitly associate the NetApp system to an OU to apply NetApp specific GPOs; this is not done by default.
- For troubleshooting the GPOs on the NetApp systems, you can turn the GPO trace on:

```
options cifs.gpo.trace.enable on
```

13 Windows Client Features

13.1 Client-Side Caching

NetApp storage systems support the Microsoft Offline Folders feature, or client-side caching, which allows files to be cached for offline use on Windows 7, Vista, XP, 2000, and 2003 clients. You can also specify whether Windows user documents and programs are automatically cached on a share or whether the files must be manually selected for caching. Manual caching is enabled by default for new shares. The

folders that are made available offline are synchronized to the Windows 2000 local disk. Synchronization occurs when network connectivity to a specific storage system share is restored.

On NetApp systems, use the following CIFS shares options to manage client-side caching, while creating a CIFS share:

```
[-no_caching  |  -  auto_document_caching  |  -auto_program_caching]
```

When you change share properties from the command line, you can specify the following:

- −no_caching to disable client-side caching for the share
- -manual_caching to enable manual selection of files to be cached on the share
- -auto_document_caching to enable user documents to be automatically cached on the share
- -auto_program_caching to enable programs and user documents to be automatically cached on the share

To enable the Offline Folders option on a Windows client, in Windows Explorer right-click the folder, select Properties, and click the Offline Files tab. To force this feature on a specific file or folder, right-click the selected network drive or subfolder and select Always Available Offline. For more information, refer to Offline Files for Windows Vista and What's new in Offline Files for Windows 7.

Recommendations

- Use client-side caching (offline folders) to cache large executables on clients (for example, the CATIA V5 CAD application).
- Use client-side caching to give mobile users access to their data even when they aren't connected to the network.
- Always enable the "Synchronize all offline files before logging off" Group Policy setting. This means that offline files are fully synchronized and that all the files in the users' redirected folders are available when they are working offline. If this setting is not enabled, the system performs only a quick synchronization, and, as a result, only recently used files are cached.
- In general, accept the Default Settings for Folder Redirection. If you are storing roaming profiles on the server on which Offline Files is enabled, Redirected Folders enables Offline Files to be set to synchronize at logon and logoff.

13.2 Accessing Shadow Copies of a Shared Folder

Snapshot technology has been an integral part of NetApp storage system solutions since 1992. Users can view Snapshot copies created on the storage system by using the Microsoft Volume Shadow Copy Service (VSS) client application. Figure 4 shows how to access shadow copies of a shared folder.

Figure 4) Accessing shadow copies of a shared folder.

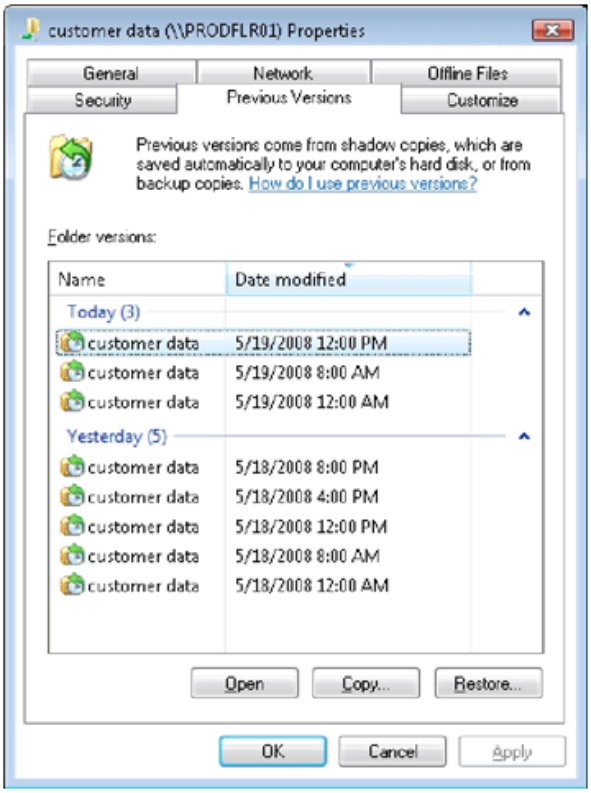

You can control the view of the Previous Versions tab if you don't want users to see Data ONTAP Snapshot copies by using the option `cifs.show_snapshot`. By default, this option is set to False. Set it to True to enable access to Snapshot copies through the Previous Versions tab.

Recommendations

- Use the Previous Version tab to give access to users' files in Data ONTAP Snapshot copies, so that users can retrieve their own files without the help of an administrator.
- You should not disable the `cifs.show_snapshot` option unless the environment doesn't allow users to look at the copies for security reasons.
- Some businesses prefer to disable the Restore button, forcing users to use the Copy-To button and to choose a restore destination manually. For more information, refer to the MS Knowledge Base Article ID: 888603.

13.3 Folder Redirection

NetApp storage systems support Microsoft folder redirection, one of the key components of Microsoft IntelliMirror technology. This option is intended mostly for organizations that have already deployed home directories and that want to maintain compatibility with their existing home directory environment. This option can also be used to redirect user-specific profile folders to an alternate location. Documents, Desktop, and Start Menu are examples of folders that you can redirect. Folder redirection enables administrators to divide user data from profile data.

Figure 5 shows how to specify a target for folder redirection on Windows Vista to a share on a storage system.

Figure 5) Specifying a target for redirecting My Documents on Windows Vista.

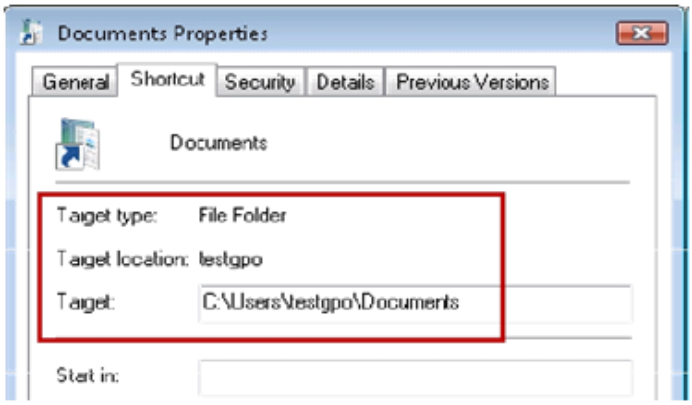

Folder redirection can also be set through a GPO configuration on the Windows Server. For more configuration details, refer to Managing Roaming User Data Deployment Guide.

Recommendations

- Don't use the redirect to home directory option unless you have already deployed home directories in your organization.
- Use Offline Folder settings on the NetApp file share where the users' data is stored. This is especially useful for mobile users; in particular, redirected folders of any type should be coupled with Offline Files.
- You can use folder redirection in Citrix environments also, where users log on from different servers in the network and need to store data centrally on the CIFS server.
- For additional best practices for using folder redirection in the terminal services environment, refer to Terminal Server Group Policy Best Practice.
- Because redirected folders can contain personal information, security must be considered for the CIFS shares used for redirected folders:
  - Restrict the share to users who need access. Create a security group for users who have redirected folders on a particular share, and limit access to only those users.
  - When creating the share, hide the share by putting a $ after the share name. This hides the share from casual browsers; the share is not visible in My Network Places.
  - Give users only the minimum number of permissions they need.
  - When redirecting to a user's home directory, make sure that correct permissions are set. When redirecting to the home directory, the default security checks are not made; ownership and the existing directory security and any existing permissions are not changed. It is assumed that the permissions on the user's home directory are set appropriately.

14 Roaming Profiles

A user profile describes the desktop computing configuration for a specific user, including the user's environment and preference settings. A profile is created the first time that a user logs on to a Windows client computer. It is a group of settings and files that defines the environment that the system loads when a user logs on.

User profiles can be stored on the local hard disk drive, or IntelliMirror can be set so that the data roams with the user wherever he or she logs on. Roaming user profiles are stored on a CIFS share on the server. This profile is downloaded every time a user logs on to any computer on the network, and any changes made to a roaming user profile are synchronized with the server copy upon logoff.

The roaming profile can be configured through the Active Directory Users and Computers MMC snap-in. Figure 6 shows how to create a roaming profile on a storage system by using the Active Directory Users and Computers MMC.

Figure 6) Using the Active Directory MMC to manage users.

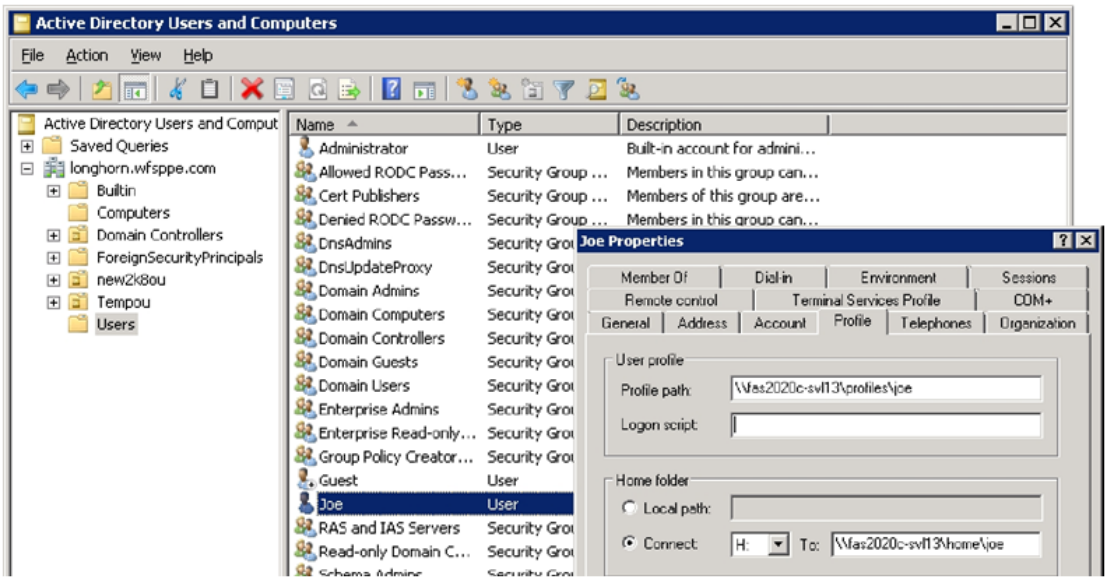

Recommendations

- **Turn off the fast logon enhancement** in Windows XP, because it takes two logons on each machine for profile changes to be registered when users change from a local to a roaming profile. This is because the user always logs on with cached credentials; therefore it takes one logon for the network to notice that the user has begun roaming and the second logon to apply these settings.

  For better results, enable the “Always wait for the network at computer startup and logon” setting, located at Computer Configuration\Administrative Templates\System\Logon.
- **Redirect the location of the My Documents folder outside of the user’s roaming profile.** To decrease initial logon time to a new computer, NetApp recommends that you redirect the location of the My Documents folder outside of the user’s roaming profile. The preferred way to do this is with Folder Redirection. If you don't have Active Directory enabled, you can do this with a logon script or instruct the user to do it manually.
- **Don't use Offline Folders on roaming profile shares.** Be sure to turn off Offline Folders for shares where roaming user profiles are stored, or you might experience synchronization problems because both Offline Folders and Roaming Profiles try to synchronize the files in a user’s profile.

Note: This does not affect using Offline Folders with redirected folders such as My Documents.

- **Do not use Encrypted File System (EFS) on files in a roaming profile.** The Encrypted File System is not compatible with files in roaming profiles. If you encrypt profile folders or files using EFS, the user’s profile cannot roam.
- **Do not redirect** `root\Application Data` **and** `root\Local Settings\Application Data` folders to a network share as part of a roaming profile. Some applications, like Google Maps, can put lots of I/O on a user profile if these two folders are redirected to a network share.
- **When using roaming profiles, consider security:**
  - Restrict the share to users who need access. Create a security group for users who have redirected folders on a particular share, and limit access to only those users.
  - When creating the share, hide the share by putting a $ after the share name. This hides the share from casual browsers; the share is not visible in My Network Places.

- Give users only the minimum number of permissions they need.

15 Citrix Environments

Citrix XenApp is a Windows application delivery system that manages applications in the data center and delivers them as an on-demand service to users anywhere who are using any device.

The integration of NetApp storage systems into Citrix server-based computing environments extends the Citrix model of high-performance, reliable user service combined with solution scalability and management simplicity.

Recommendations

- Consider implementing a server farm consisting of multiple identically configured servers. The benefits of using server farms are:
  - Because of the uniformity of the environment, users can connect to any of the servers at any time without being tied down to a specific server.
  - It's easier to scale up in a server farm when additional capacity is needed.
  - It allows performing maintenance on a particular server or the entire farm on a server-by-server basis with no disruptions to users.
- Make use of roaming profiles and folder redirection to minimize roaming profile corruption, shorten the login times for users, and improve performance in general.

For more information about user profile best practices in a Citrix XenApp (Presentation Server) environment, refer to User Profiles Best Practices for Meta Frame Presentation Server.

16 Multiprotocol

NetApp storage systems are designed to support multiprotocol natively. This provides a bridge between the UNIX file security styles and Windows file security style by mapping the UNIX identities and Windows identities. Many customers have environments in which the data must be accessed both from UNIX clients through NFS and Windows clients through CIFS. Because the UNIX security model and the Windows security model do not directly correlate, a single security style's credentials cannot be used for accessing the data with a nonnative security style.

Both models use the concept of a user for authentication and authorization. NetApp provides the mechanism of mapping the user from one security style to another security style. When a UNIX user tries to access data that uses NTFS security style, UNIX user credentials are mapped to a Windows (CIFS) user's security credentials and provide authentication. The same applies to a CIFS user accessing the data that has a non-NTFS security style.

For more information on the multiprotocol functionality and configuration, refer to TR-3490, "NetApp Storage System Multiprotocol User Guide."

17 SMB 2.0 Protocol

SMB 2.0 is a next-generation NAS protocol for Windows. The protocol has been redesigned to facilitate the next-generation NAS server requirements, especially for wireless networks and remote office deployments. The SMB 2.0 protocol is much more resilient to network interruptions. It is designed to scale and perform better, as well as to provide more security, as compared to the CIFS (SMB 1.0) protocol. There are opportunities for application vendors to use their applications to get maximum benefits from SMB 2.0 features.

Data ONTAP supports SMB 2.0 in coexistence with CIFS (SMB 1.0) starting from version 7.3.1. Microsoft supports SMB 2.0 in Windows Vista and later operating systems.

For more information about all the SMB 2.0 features, configurations, and benefits, refer to NetApp TR-3740, "SMB 2.0 – Next Generation CIFS Protocol in Data ONTAP."

Recommendations

All of the best practices described for the CIFS protocol are also valid for the SMB 2.0 protocol. However, there are some guidelines that are specific to SMB 2.0:

- Use SMB 2.0 if there are higher performance, scalability, security, and resiliency requirements, or if federal regulations demand using SMB 2.0 as a standard protocol.
- To use SMB 2.0 in your environment, your clients must have Windows Vista SP1 at a minimum. Windows Vista SP1 has the full implementation of the SMB 2.0 protocol; Vista RTM had partial implementations of some SMB 2.0 features.
- SMB 2.0 performance benefits are most visible on WAN links as compared to LAN, so use it accordingly.
- Windows Vista clients are mostly auto-tuned to get the maximum benefits of the SMB 2.0 protocol, so generally no additional performance tuning is required on the client side. However, in order to achieve the maximum performance benefit for SMB 2.0, the following best practices are very helpful.
- If you are using any applications over SMB 2.0, then wherever applicable, tune your applications to:
  - Leverage the larger block size to send data
  - Send requests for more concurrent blocks
- Use a Gigabit Ethernet or better network for high bandwidths.
- Use the best hardware on both the client and server side, because a powerful configuration yields more performance.
- Turn SMB signing off if it's not needed.

18 Recommendations for Optimal SMB Performance

- Use SMB 2.0 over SMB 1.0 wherever possible because SMB 2.0 provides better performance and more security than SMB 1.0.
- TCP Window Size: Set the value to 64K or more if SMB 2.0 is the protocol used. To set TCP window size on the controller:

```
options cifs.tcp_window_size 64240
```

- Setting TCP window size to 2MB gives very good performance when combined with Windows 7 clients.
- Max outstanding requests per session (Max Mpx): Keep the Max mpx more than the 50 (the default). This setting improves performance for parallel reads and writes. The values that can be set are 126, 253, or 1124.

```
options cifs.max_mpx <value>
```

- Negotiated Buffer Size: This setting controls the maximum negotiated read/write buffer size that can be used. The default is 32k and the size can also be set to 64k.

```
options cifs.neg_buf_size <value>
```

19 Data Migration

Data migration is one of the key requirements in enterprise data management. Data migration is usually required when the systems that hold the user or enterprise data need to be refreshed or there is a need to

move the data to a secondary storage device. This is one of the key steps when consolidating Windows file servers to NetApp storage.

Some of the industry's leading data migration tools are Robocopy, Secure Copy, and F5 ARX (Acopia).

20 Data ONTAP 8.0.1 7-Mode Features

Data ONTAP 7-Mode has maintained parity with all Data ONTAP 7G features while incorporating the additional performance enhancements due to the introduction of CIFS Waffinity.

20.1 CIFS Waffinity

Prior to Data ONTAP 7.2, WAFL® is completely single threaded. Data ONTAP 7.2 introduced Waffinity, which allows threads to run on multiple CPUs. CIFS support for Waffinity was introduced with 8.0. With this feature, CIFS-related WAFL messages are processed under the Waffinity framework.

Without CIFS Waffinity, all CIFS-related message processing happens in one common I/O domain that is single threaded, along with WAFL and SnapMirror. CIFS Waffinity attempts to move some of this processing out of this domain to a different domain that runs threads on multiple CPUs. This enables us to process multiple CIFS messages in parallel, thereby improving the CIFS performance.

CIFS Waffinity is supported only on platforms running Data ONTAP 8.0 or later on multicore processors.

Note: The following features are not available in 7-Mode:

- IP version 6 (IPv6)
- SnapLock® Compliance and Enterprise versions
- IPsec

21 Conclusion

File sharing in Windows File Services involves many factors. It's very important to follow all best practices at the time of deployment. Employing best practices can save a great deal of time and effort in ongoing management and maintenance. If the solution is deployed correctly and efficiently, it can also save on total cost of ownership and can lead to stable performance in the environment.

References

- Storage Subsystem Resiliency Guide
  http://media.netapp.com/documents/tr-3437.pdf
- Best Practices for Secure Configurations of Data ONTAP 7G
  http://media.netapp.com/documents/tr-3649.pdf
- Anti Virus Scanning Best Practices Guide
  http://media.netapp.com/documents/tr-3107.pdf
- SMB 2.0 – Next Generation CIFS Protocol in Data ONTAP
  http://media.netapp.com/documents/tr-3740.pdf
- NetApp Storage Systems in Microsoft Windows Environment
  http://media.netapp.com/documents/tr-3367.pdf
- File Access Auditing on NetApp Controller
  http://media.netapp.com/documents/tr-3595.pdf
- Storage-Level Access Guard Quick Start Guide
  http://media.netapp.com/documents/tr-3596.pdf

- Bulk Security Quick Start Guide
  http://media.netapp.com/documents/tr-3597.pdf
- NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
  http://media.netapp.com/documents/tr-3505.pdf
- Active Directory and Domain Authentication
  http://technet.microsoft.com/en-us/library/bb742424.aspx
- Unified Windows and UNIX Authentication Using Microsoft Active Directory Kerberos
  http://media.netapp.com/documents/tr-3457.pdf
- DFS Step-by-Step Guide for Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc732863.aspx
- Best Practices for Distributed File System
  http://technet.microsoft.com/en-us/library/cc736324.aspx

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

NetApp
www.netapp.com