

Dell Compellent FS8600: CIFS File Server Consolidation Guide

A Dell Compellent Technical Tip



Compellent

THIS TECHNICAL TIP IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Trademarks used in this text: Dell™, the DELL™ logo, and Compellent™ are trademarks of Dell Inc.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

June 2012 Rev. A

Contents

General Syntax	5
Document Revision	5
Preface	5
Customer Support.....	5
Introduction	6
Core concerns	6
Active Directory.....	6
Active Directory Domains.....	6
Active Directory Domain Controllers.....	6
Active Directory Security Identifiers.....	7
Active Directory Domain Trusts.....	7
Hypothetical Example	7
Permissions	8
Access Control Lists & Access Control Entries.....	8
Inheritance.....	8
Migrating with permissions	9
Shares.....	9
Share names, DNS and networking	9
Shares and subshares.....	10
Antivirus.....	10
NAS Volumes.....	10
Snapshots	10
Replication.....	10
Initial file ingestion.....	11
Initial ingestion and Storage Center tiering	11
Optimizing performance	11
Ingestion from production.....	11
Ingestion from backups	11
Migration tools	12
Robocopy.....	12
Xcals	12

Tables

Table 1. Conventions5

Table 2. Revision History5

Figures

Figure 1. Domains before and after migration8

General Syntax

Table 1. Conventions

Item	Convention
Console or file trace. (Commands are shown with system prompt '#' in console traces.)	Monospace Font
Other user input (file names and paths, field names, keys)	Monospace Font
Website addresses	http://www.compellent.com
Email addresses	name@dell.com

Document Revision

Table 2. Revision History

Date	Revision	Description
June 12, 2012	A	Initial draft.

Preface

This document is intended to be a resource used in the planning for the consolidation of multiple CIFS file shares and CIFS servers on to a Dell Compellent FS8600 NAS platform.

It is not intended to be a walkthrough of a migration, or a direct template for migration planning. It is similarly not intended to be the authoritative document for all FS8600 CIFS related subjects.

Customer Support

For support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

Introduction

As unstructured file data continues to grow at challenging rates, IT file server infrastructure supporting business and mission critical systems can easily become strained and fragmented. This can manifest in a variety of painful forms such as overall system availability, lost storage efficiencies, complicated or fickle backup architectures and rising management costs.

The Dell Compellent FS8600 is uniquely suited to overcoming the challenges presented by explosive file data growth and file sprawl. With the combination of FluidFS's Scale-Out Single Namespace and Storage Center's Data Progression and Dynamic Capacity, IT teams can reduce management overhead, and leverage the platform's unique storage efficiencies and cost of ownership attributes while providing class leading availability and data protection.

It should be noted that the FS8600 is a member of the Dell FluidFS NAS product family, and is in many ways conceptually very similar to the FluidFS products available on the Dell PowerVault and Dell EqualLogic product lines. While many of the concepts, particularly front end NAS concepts are similar or even identical, some subjects such as management of the platform and interaction with backend storage may be radically different, and the relevant resources for each FluidFS product should be leveraged for the appropriate product.

Core concerns

File server migrations and consolidation efforts can be daunting challenges, with unforeseen roadblocks and complex management of interdependent systems. An understanding of the core concerns and concepts involved, along with proper planning can limit the scope of effort required to consolidate established file server and file share sprawl onto an FS8600.

Active Directory

As Active Directory continues to have a dominant foot print in data center and office IT infrastructure, understanding the topology and specific details of any Active Directory systems that may be involved in file server migrations and consolidations is crucial. Outlining all the elements and concepts of Active Directory falls outside the scope of this document, however the following sections outline the subjects directly related to CIFS migrations and consolidation onto an FS8600.

Active Directory Domains

Domains are one of the fundamental building blocks of Active Directory. Domains represent the primary mechanism for outlining or limiting the scope of pools of resources that should be logically managed together.

Active Directory Domain Controllers

Domain controllers (DC or DCs) are logical server hosts that have been specifically designated to be an authority and resource manager for a specific domain. More than one domain controller can be present within a domain, and through various replication technologies the configuration for the domain is kept consistent across all domain controllers in the domain, and in some cases, with resources outside the domain.

When a FS8600 is said to “bind” or “join” to an Active Directory domain, it does so by communicating with the domain controllers in the environment and incorporating itself into the Domain topology. As

long as the FS8600 continues to be a member of that domain, it will use the domain controllers in that domain as the point of visibility into the rest of the domain topology.

It should be noted here that in order to join an FS8600 to a domain, an Active Directory user with sufficient privileges must be used to complete the operations required to bind the FS8600 to that domain. If requesting or provisioning a user with the needed privileges requires an administrative escalation, (e.g. filing a ticket with a corporate I.T. team) that should be noted and the time required to complete that escalation included in the planning timeline.

Active Directory Security Identifiers

Active directory creates a unique security identifier (SID) for each user, group, machine or other object created within an active directory topology. When an attribute or parameter has to be set for any of these objects, the attribute references the SID, which is translated into a human readable name by the Active Directory services of the domain that owns the SID. Any time a SID must be interacted with, a DC for the associated domain must be visible to identify and own that SID.

Active Directory Domain Trusts

Domain trusts are the mechanisms through which a logically isolated domain establishes a logical relationship with another domain with the intent of sharing or negotiating access to resources across domains.

For the purposes of this document, the primary concern of domain trusts are user management and visibility, and the available points of visibility for a host responsible for migrating file and permission data from the source file share to the FS8600.

Hypothetical Example

In conjunction with other efforts to unify various development teams, it is decided that various shares and domains will be consolidated onto an FS8600, and a new domain created with the intent of consolidating established domains over time. For the first phase of a staggered migration, it is determined that the following shares will be migrated:

[\\nas01.corp.local\appsteam](#)

[\\nas02.corp.local\appsusers](#)

[\\files.apps.corp.local\appsrepo](#)

[\\test3.devapps.corp.local\devappsrepo](#)

[\\util.devops.corp.local\source](#)

These shares will be migrated to the new FS8600 and new domain to create the following shares:

[\\fs8600.systems.corp.local\appsteam](#)

[\\fs8600.systems.corp.local\appsusers](#)

[\\fs8600.systems.corp.local\appsrepo](#)

[\\fs8600.systems.corp.local\devappsrepo](#)

[\\fs8600.systems.corp.local\source](#)

The complex system of trusts between domains will be simplified by replacing the system of trusts between the various domains to a two way trust between the “systems” domain to the “apps” and “devapps” domains, a single one way trust from the “corp” to “system” domains, and collapsing the “devops” domain into the new “systems” domain. In this case, the FS8600 needs to be bound to the “systems” domain.

Figure 1. Domains before and after migration



For the migration of each share in each domain, a host must be involved with visibility to the DCs for that domain in order to properly migrate the SIDs and other attributes from the source share to the FS8600 share.

Permissions

While very closely related to Active Directory, particularly SIDs, file level permissions can present a separate challenge in file share migrations.

Access Control Lists & Access Control Entries

An Access Control List (ACL) is a list of Access Control Entries (ACEs) that make up the permissions structure of each file or folder, which make up the building blocks of permissions schemes in a CIFS environment. An ACE is a set of permissions and other attributes directly associated with a single SID, and ACEs from multiple domains can make up an ACL for a file or folder.

An explicit ACE is one that has been administratively configured for a specific file, and is ordered in the ACL before any inherited ACE. An inherited ACE is one that is applied by the rules of folder and file inheritance by other ACLs and ACEs in the parent file structure. An inherited ACE is ordered below all explicit ACEs in an ACL.

When a file is migrated, it is important that the ACEs follow the file in order, as access to a file or folder is dependent on the order of the ACEs in that as the ACL is traversed, as the first “allowed” or “denied” entry will allow or deny access regardless of other ACEs that may exist in the ACL.

Inheritance

Inheritance is the process by which certain permission attributes are automatically applied to files based on the folder they are written to. Inheritance is actually an attribute that makes up an ACE on a parent folder, that is obeyed and inherited as sub folders and files are created below the parent folder. Similarly, there is an attribute that can be applied to an ACE to block inheritance.

Because the inheritance attributes make up part of the ACE, they should be maintained as they are migrated from source to destination share. It should be noted that inheritance policies on the share

itself or root folder for the share may override the inheritance policies of the first sub-folder, so this should be checked as well before a migration begins.

Migrating with permissions

Similar to requiring that tools used to move files be SID aware, the tools should also be permissions aware, in that they can walk and maintain the ACEs in a ACL (as they correspond to SIDs) and copy those properly with the file to the destination share.

If the tools available, or the migration path available, do not allow for the proper migration of permissions, there are tools available to “re-apply” a permission tree to a file set after they have been copied.

Shares

For some platforms, shares represent isolated containers that must remain largely static over their lifetime. For FluidFS and the FS8600, shares are views for CIFS clients into various assigned parts of the overall single namespace of the file system, as isolated by the NAS Volumes. This provides significant flexibility in creating a file set structure optimal for file system consolidation.

Share names, DNS and networking

One of the decisions that must be made early on in the planning phase is if DNS can be changed to reflect share movements across hosts, or if IP addresses can move with the share, or if neither of these is possible. The overall effort that may be required in reconfiguring clients to reflect new IPs, DNS or share names should be evaluated against the effort required to reconfigure all resources dependent on that name.

In the case of being able to alter DNS, the overall Universal Naming Convention (UNC) used address the share by various clients may not have to be updated. For example in the case of hypothetical example used earlier, if DNS for \\files.apps.corp.local\\appsrepo can be updated to point from the old IP of the host to the new IP pool of the FS8600, any client already configured to target that name will be able to follow migration without being reconfigured. In these sorts of migrations, it is generally considered best practice to reduce the Time To Live (TTL) property for the associated DNS A or CNAME records prior to the migration. An example:

- After a share is disabled or marked read only on the original host and the file data migration completed, the DNS can be re-configured to point to the IP pool of the FS8600.

If it is possible to move the IP of the original host to the FS8600, this offers another route for maintaining established client configurations. An example:

- After a share is disabled or marked read only on the original host and the file data migration completed, the IP for the share/host can be removed from the source host/share and configured on the destination FS8600.

If it is not possible to change DNS or move the related IP, then a new IP pool and new DNS entries must be assigned to the FS8600, and clients must be reconfigured to target the new DNS entries or IP addresses.

It should be noted that the FS8600 is fully VLAN aware, which can be used to overcome some obstacles in many consolidation efforts, where using the earlier example test3.devapps.corp.local may expect untagged frames from VLAN238 and util.devops.corp.local may be dependent on VLAN686. In this case,

both VLANs can be left tagged to the FS8600 Ethernet ports at the switch level, and the FluidFS platform is left responsible for understanding and managing the VLAN tagged frames.

Please note to consult the **Dell Compellent FS8600 Best Practices** document for further information about DNS, IP addresses and VLANs with the FS8600 platform.

Shares and subshares

Because shares represent points of visibility into the file system as isolated by NAS Volumes, the creation of shares below shares (subshares) can provide much needed control granularity. These can be created at arbitrary points within the file system, as needed, for either migrations or ongoing use cases. It should be noted that subshares can introduce a level of management complexity with regards to permissions and particularly inheritance, and should be used only as needed as opposed to a default assumption.

Antivirus

Each share or subshare can be configured to enable Antivirus scanning, via the ICAP protocol to a pre-determined Antivirus host. For more information regarding Antivirus support, please consult the **Dell FluidFS NAS Solutions Administrator's Guide**.

NAS Volumes

NAS Volumes represent the primary abstraction and isolation mechanism on top of the single namespace, single instance FluidFS file system. The FS8600 features Snapshots and Replication are NAS Volume Centric; therefore planning should include considerations towards a NAS Volume scheme. For more information about NAS Volumes, please consult the Dell FluidFS NAS Solutions Administrator's Guide and Dell Compellent FS8600 Best Practices documents.

Snapshots

FluidFS Snapshots are an elegant, efficient and high performance data protection technology built into the core FluidFS architecture. Because of their performance, reliability and ease of use, they can provide the primary data protection mechanism for any data set profile. One of the primary benefits of consolidation of file shares to a unified FluidFS platform is the collapsing of administrative overhead required to manage the backups of the distributed or unorganized file share hosts.

The FS8600 also provides two methods for user recovery of files from snapshots. Snapshot versions of files can be access through the "View Previous Versions" Windows tool. They can also be access through a hidden ".snapshots" folder that FluidFS places in every share. System Administrators, as well as users, can access these areas to recover file data preserved in an existing Snapshot.

Replication

Many data sets may be required to be backed up to an offsite location or otherwise be stored in more than one system or location. With the replication available through the FluidFS platform, this data can be replicated from an active FS8600 in one site, to another FS8600 either in same location or in a separate geographic location entirely. This can significantly simplify the backup systems involved in meeting data protection and data retention requirements.

For environments that intend to use a replication in such a way as to "stand up" on the replication destination, all Active Directory resources available in the primary site must be available in the site with the replication target.

Initial file ingestion

Initial ingestion and Storage Center tiering

Any planning for the initial seeding of the FS8600 should be aware of the size of the data set to be imported, and the total available capacity in the Storage Centers Tier 1 storage. In order to avoid any overall Storage Center performance degradation, imports should either be planned in such a way as to not overrun Tier 1 storage, or the Storage Center volumes assigned to the FS8600 should have their storage profiles altered to a lower tier. Initial replication between FS8600s should also be considered in the planning, as initial replication of large data sets can similarly threaten to overrun Tier 1 storage.

Optimizing performance

For data sets that can be broken up during initial ingestion so that data can be migrated from multiple hosts, there are both migration and life-cycle performance benefits from performing that migration from multiple hosts. For performance during the migration, the extra hosts can improve overall throughput of data from the source to the FS8600. Migrating data from multiple hosts also takes best advantage of the overall FluidFS architecture, ensuring that data is best placed evenly around the entire cluster in order to provide best performance for the data set going forward.

Ingestion from production

Often the simplest method for migrating file data to a new platform is to simply move or copy it from its production platform. While this may be the simplest, it is also likely to impact the current platform and the production file data set.

If the current platform has access to the new FS8600 and provides the functionality for mounting other shares on other hosts (e.g. Windows Storage Server platforms), it may be easiest to simply mount the intended destination FS8600 share directly on the current platform, and copy file data directly from the current platform to the FS8600. This minimizes the amount network “hops” required to seed the FS8600. It is important to remember that the FS8600 must be bound to Domain Controllers with the required visibility to translate the SIDs and other Active Directory file information associated with the data set.

If it is not possible to mount the share from the current platform, a host with visibility to both the source platform and the destination FS8600 can be used to migrate the file data to the destination FS8600. In this case, the host being used to perform the migration, as well as the FS8600 must have visibility to the Domain Controllers responsible for the domains involved.

Ingestion from backups

In environments where the performance impact or network exposure required are the limiting factors in migrating file data, a common strategy is to seed and populate from the backup system associated with the current platform. Because this introduces no performance impact to the current system, and requires no network exposure, it is often used for massive file sets or complicated migration strategies.

A strategy that builds on this is to use a backup system as the mechanism responsible to initially seed the destination FS8600, and then use host level tools to keep that seeded data in sync or at least in close proximity to the state of the current platform.

The inclusion of another system in the strategy can add complexity and introduce new challenges, such as performance contention for resources on the backup system. Another area of concern is that the

permissions scheme of the current platform may be proprietary, and items such as file permissions ACEs or other attributes may not be stored in a translate-able way. Because the backups are taken from the perspective of the current platform, this may mean that file data attributes could not be understood by the destination FS8600. A final example of this introduced complexity is that due to the ways many enterprise backup systems catalog or maintain data on medium such as tape drives or libraries, file attributes such as access and write time stamps may be inaccurate when the data is seeded to the new FS8600.

The seeding of data to the FS8600 can be performed using share mounts of the FS8600 to the backup host. As the FS8600 only supports NDMP restores from data that was initially backed up via NDMP from an FS8600, NDMP backup systems must be able to export NDMP data through open files.

Migration tools

There are a number of tools and utilities available to assist in the various tasks required for file data migrations or consolidations. The list below is not a complete list of available tools, nor is it a list of tools supported by Dell Compellent. It is simply a list of tools widely used by IT organizations for file data migrations and are commonly understood.

Robocopy

Robocopy, and a number of similarly purposed tools, is specifically designed for moving large file sets of complex data to new destinations. It has earned a reputation for reliability and performance, with modern versions of *robocopy* including support for ACLs and ACEs, as well as features such as multithreading for improved performance on multi-CPU systems across large file sets. There are a number of ways to automate *robocopy* activity, either through built in tools such as *robocopy*'s "run hours" flag or through tools such as Microsoft Windows's "Scheduled Tasks" automation tool. Using this automation, file sets can be kept in sync or in a close proximity state for a period of time running up to a migration window.

Xcacs

Xcacs, and other similarly purposed tools, provide the functionality to harvest or save permissions data and schemes across large file sets, and re-apply them to a similar or copied file set. For environments where the tool used to migrate the file data itself is unable to migrate the file permissions, *Xcacs* can be used to re-apply the permissions after the file data migration. It can also re-apply permissions to data seeded from a backup system.