



Deploying the F5 ARX with Data Domain Storage Systems

Table of Contents

Deploying the F5 ARX with Data Domain Storage Systems

Prerequisites and configuration notes	I-1
Revision history	I-3
Configuration example	I-3
Configuring the F5 ARX device	I-4
Verifying the Namespace protocol support setting	I-4
Verifying CIFS volume configuration options	I-6
Adding the Data Domain system as Tier 2 storage	I-8
Adding policies for data migration between Tier 1 and Tier 2	I-13
Appendix A: Verifying data compression is occurring	I-19

Deploying the F5 ARX with Data Domain Storage Systems

Welcome to the F5 - Data Domain Storage Systems deployment guide. This guide provides step by step procedures on deploying the Adaptive Resource Switch (ARX) with the Data Domain system.

The F5 ARX file virtualization platform decouples file access from [the](#) physical file location within Network Attached Storage (NAS) environments. The ARX platform automates file migration to the appropriate tier of storage without impacting data access, thus minimizing backup and recovery windows.

Data Domain storage systems are typically deployed as nearline storage for tape backup replacement. Data Domain systems leverage deduplication and compression technologies to reduce the data footprint. With the F5 ARX, the Data Domain system can be deployed as Tier-x nearline storage within a single global namespace. Because ARX's place rules provide seamless online automated file migration among multiple tiers, the combination of the F5 ARX and Data Domain storage systems further optimizes and reduces the cost and management of tiered storage solutions.

More information regarding Data Domain systems can be found on [the](#) Data Domain website <http://www.datadomain.com/>.

Prerequisites and configuration notes

The following are prerequisites and configuration notes for this deployment:

- ◆ This document is based on ARX version 5.1.0 and Data Domain OS 4.7.3.0.
- ◆ This document demonstrates the deployment of the Data Domain system as Tier 2 storage and it is assumed that the ARX has been preconfigured with a working Tier 1 storage.
- ◆ The following Data Domain systems are compatible with F5 ARX:
 - DD530, DD510, DD4xx, and DD120 support 150 CIFS connections
 - DD565, DD360, and DD560g support 300 CIFS connections
 - DD580, DD880, DD880g, DD690, support up to 600 connections
- ◆ The following features are not supported:
 - Data Domain Retention Lock feature
 - ARX Virtual Snapshots comprised of Data Domain snapshots
 - ARX Direct Volumes for NFS
- ◆ This document is written with the assumption that basic networking, namespace, managed volume, Tier 1 shares, and virtual service are configured and virtual service is up and running with Tier 1 storage attached to it. Refer to the product documentation on how to configure these objects.

- ◆ Refer to the Data Domain release notes for information on configuring the Data Domain DDR device to give the F5 ARX proxyuser backup Operator Privileges. Use the **cifs option set F5 <Domain Name>\<Proxy User>** Data Domain CLI Command.
- ◆ For CIFS, you must disable Named Streams (see *Verifying CIFS volume configuration options*, on page 7 for more information). No support for Compressed Files and Sparse Files.
- ◆ Although the Data Domain OS supports both CIFS and NFS version 3 (TCP only), ARX deployments with Data Domain configured for multi-protocol access are not supported.
- ◆ Hosting ARX metadata is not supported on the Data Domain system. The Data Domain system is not intended to be deployed as a primary storage and therefore is not required to host ARX metadata.
- ◆ ARX v5.1.0 adds support for CIFS multiplexing. This feature multiplexes multiple client side CIFS connections onto a fewer number of Server side connections. This new ARX feature specifically addresses back-end filesystem CIFS/TCP connection limitations/constraints. The F5 ARX CIFS Connection Multiplexing feature facilitates proxy multiplexing thousands of client connection requests though the ARX into a couple of hundred to the backend filesystem.
- ◆ ARXv5.1.0 adds support for Configuring the File System Name. CIFS client applications can query a managed volume for the name of its file system. By default, the volume chooses a name based on the CIFS capabilities of its back-end shares; for example, it advertises “NTFS” if and only if all of its back-end shares can support NTFS-related features. Some client applications require specific settings in this string in order to function properly. On the advice of F5 Support, you can use the CIFS **file-system-name** command to manually set the file system name for this volume. For more information see the ARX CLI Storage and CLI Reference guide; pay particular attention to the External Filer section.

Revision history

Revision	Description
1.0	Deployment Guide published
1.1	Removed CIFS support for Compressed Files and Sparse Files from ARX configuration, as they are not supported by Data Domain.
1.2	Updated the deployment guide to add support for Data Domain OS v4.6.3.1
1.3	Added support for ARX v5.1.0 and DD v4.7.3.0
1.4	Added the CIFS Connection Limit setting to the external filer procedure (Step 7d of <i>Adding the Data Domain system as Tier 2 storage</i> , on page 9) and the command line procedure (<i>To add the Data Domain system as Tier 2 storage from the command line</i> , on page 12). Also updated Figure 7, on page 10.
1.5	Made changes to the Data Domain configuration section based on feedback from EMC, updated tiering steps, and screen captures. Added support for ARX v5.1.5

Configuration example

In the following diagram, we show basic connectivity between clients, ARX and the Data Domain system. A client attempts to retrieve a file from a file share. The ARX proxies the request, and transparently retrieves the file from the server where the file is located (regardless of whether it is on Tier 1 or Tier 2). We configure a policy on the ARX that periodically checks the last time files were modified, and migrates the file to the appropriate Tier if the

conditions of the policy are met. In our example, the ARX policy is checking for a last modified time of less than (or more than) one day. If the policy matches, then the ARX moves the file between the Tiers.




Figure 1 Logical configuration example

Configuring the F5 ARX device

In this section, we configure the ARX device. The following procedures contain both Configuration utility graphical user interface (GUI) and command line (CLI) steps; choose the most familiar format.

To configure the ARX for the Data Domain system, you must complete the following procedures:

- *Verifying the Namespace protocol support setting*, on page 5
- *Verifying CIFS volume configuration options*, on page 7
- *Adding the Data Domain system as Tier 2 storage*, on page 9
- *Adding policies for data migration between Tier 1 and Tier 2*, on page 14

After completing these procedures, see *Appendix A: Verifying data compression is occurring*, on page 20.

Verifying the Namespace protocol support setting

The current version of the Data Domain system does not support multi-protocol configurations. Use the following procedure to verify that the namespace protocol is **CIFS** or **NFS v3-TCP**.

◆ Important

Changing protocol setting on the Namespace may require disabling or destaging all volumes. For information on disabling or destaging all volumes, see chapter 5 of the [ARX CLI Maintenance Guide on Ask F5](#).

To verify the namespace protocol from the Configuration utility

1. From the navigation pane, click **Namespace**.
2. From the list, click the name of the namespace to which the Data Domain system will be added, and click the **Edit** button. Or select the **Add** button to create a new namespace.
3. In the Protocol section, make sure there is a check in either the **CIFS** or **NFS v3-TCP** box. Click the applicable box. If you made any changes, see the Important note preceding this procedure.
4. Click the **OK** button (see Figure 2).

Edit Namespace

Namespace Name	my_namespace_nfs
Namespace Description	
Protocol(s)	<input type="checkbox"/> NFSv2 <input type="checkbox"/> NFSv3-UDP <input checked="" type="checkbox"/> NFSv3-TCP <input type="checkbox"/> CIFS
Metadata Cache Size	512 MB

NFS Properties:

NFS Character encoding	ISO-8859-1
-------------------------------	------------

OK Cancel

Figure 2 Changing the Namespace protocol

To verify namespace protocol from the command line

1. From the command line, use the following syntax:

show namespace <namespace-name>

In our example, we type

show namespace my_namespace_nfs

The following shows an example of the output from this command.

```

arx1a# show namespace my_namespace_nfs
Namespace "my_namespace_nfs" Configuration
Description
Metadata Cache Size: 512 MB

Domain Information
-----

Supported Protocols
-----
nfsv3-tcp

Participating Switches
-----
None

Volumes
-----
/data

Volume freespace: Not available (automatic)
Metadata size: Not available
Metadata free space: Not available
State: Enabled (No enabled shares)

Host Switch: arx1a
Instance: none
VPU: 1 (domain 2)
Files: Unknown used, 10M free

Metadata shares:

Filer      Backend Path    Contains Metadata  Status
-----
nfs-filer1 /meta1          No                  Unknown
  
```

Figure 3 Example output from the show namespace command

-
2. To change the protocol setting from the command line (for example, to remove support for NFSv3 UDP), use the following syntax:

```
global
  namespace <namespace-name>
    no protocol nfs3
  exit
exit
```

 **Important**

Changing protocol setting on namespace may require disabling or destaging all volumes. For information on disabling or destaging all volumes, see chapter 5 of the [ARX CLI Maintenance Guide on Ask F5](#).

Verifying CIFS volume configuration options

If you are using the CIFS protocol, you must disable Name Streams and Unicode-On-Disk on the CIFS volume configuration.

To verify CIFS volume configuration options from the Configuration utility

1. From the Navigation pane, click **Managed Volumes**.
2. Click the name of the target volume.
3. Click the **Edit** button.
4. In the **CIFS Characteristics** section, make sure the **Persistent ACLs** and **Unicode-on-Disk** boxes are checked. The Data Domain system does not support Compressed Files, Sparse Files or Named Streams (see Figure 4).
5. Click the **OK** button.

Edit Managed Volume

Enable Volume	<input checked="" type="checkbox"/> Enable Volume
Volume Name	/data
Volume Description	
Metadata File Server	filer1
Metadata Protocol	CIFS
Metadata Path	meta1
Shadow Volume	<input type="checkbox"/> Shadow Volume
Snapshot Directory Name	~snapshot
Snapshot Directory Display	None
	<input type="checkbox"/> Mark snapshot directory as hidden.
Privileged Snapshot Access	<input type="checkbox"/> Allow access to snapshots by privileged users only.
Snapshot Consistency	<input type="checkbox"/> Make snapshots consistent.
VSS Mode	Windows XP
Auto Reserve Files	<input checked="" type="checkbox"/> Auto Reserve Files
Maximum Files	4000000
CIFS Characteristics	<input type="checkbox"/> Auto-detect CIFS characteristics. <input checked="" type="checkbox"/> Compressed Files <input checked="" type="checkbox"/> Sparse Files <input checked="" type="checkbox"/> Named Streams <input checked="" type="checkbox"/> Unicode-On-Disk <input checked="" type="checkbox"/> Persistent ACLs

Figure 4 Configuring the CIFS characteristics

To verify or change the volume configuration options from the command line

1. From the command line, use the following syntax:

```
show namespace <namespace-name> volume <volume-name>
```

For example

```
show namespace my_namespace_cifs volume /data
```

The following figure shows an example of the output.

```
Volumes
-----
/data
  CIFS : compressed files: no; named streams: no; persistent ACLs:
yes
      sparse files: no; Unicode on disk: yes; case sensitive: no

      Volume freespace: 6.0GB (automatic)
      Auto Sync Files: Enabled
      Metadata size: 252k
      Metadata free space: 6.0GB
      Oplock support: Enabled
      Notify-change mode: Normal
      CIFS path cache: Not Enabled
      State: Enabled

      Host Switch: arx1a
      Instance: 1
      VPU: 1 (domain 1)
      Files: 1 used, 9.9M free

      Metadata shares:

Filer      Backend Path      Contains Metadata      Status
-----
filer1     meta1              Yes                     Online
```

Figure 5 Viewing the volume configuration options from the command line

-
2. To change the volume configuration settings from the command line, type the following command:

```
global
    namespace <namespace-name>
        volume <volume-name>
            no named-streams
end
```

Adding the Data Domain system as Tier 2 storage

With the F5 ARX, the Data Domain system can be deployed as Tier-x near-line storage within a single global namespace. In a typical environment, the Data Domain system is deployed as Tier 2 or Tier 3 storage. In the following procedure, we add the Data Domain system as Tier 2 storage.

To add the Data Domain system as a Tier 2 storage from the Configuration utility

1. From the navigation pane, click **Common Operations**.
2. Click the **Add Share** button.
The Add Share wizard opens.
3. In the Share name box, type a name for this share. In our example, we type **ddr_share**.
4. From the **Namespace** list, select the appropriate Namespace. In our example, we select **my_namespace_cifs**.
5. From the **Volume** list, select the appropriate volume. In our example, we select **/data**.
6. Click the **Next** button.

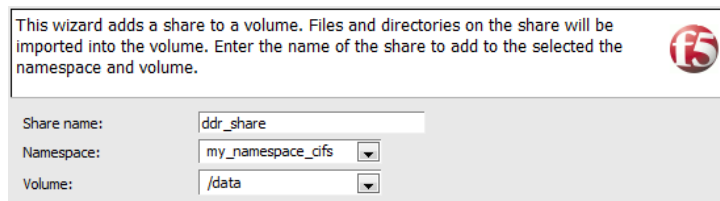


Figure 6 The Add Share wizard

7. If the Data Domain system information has not been added as a filer to ARX, click the **Add** button at the top of the page. If it has already been added, skip to the next step.
 - a) On the new file server page, in the **Name** box, type a name for this file server. In our example, we type **ddr_filer**.

- b) In the **Primary IP address** box, type the primary IP address for the file server.
- c) In the **CIFS port** box, type the appropriate port. In our example, we leave the port at the default: **445**.
- d) In the **CIFS Connection Limit** box, type the appropriate maximum number of connections. In our example, we select **600**.
Note: Refer to *Prerequisites and configuration notes*, on page 1 for a list of maximum connections by Data Domain platform.
- e) Optionally you can enable snapshot support. A snapshot is an exact copy of a managed volume at a single point-in-time. You can create regularly-scheduled snapshots in a managed volume, and you can limit the CIFS clients who can access those snapshots. To enable snapshot support for the Data Domain System, check the **This file server supports snapshots** box, from the **File Server Type** box, select **Data Domain**, and then provide management proxy user credentials for the Data Domain System. For more information on snapshot support refer to the ARX CLI Storage Guide
- f) Click the **Save** button. You return to the File Server page of the wizard.

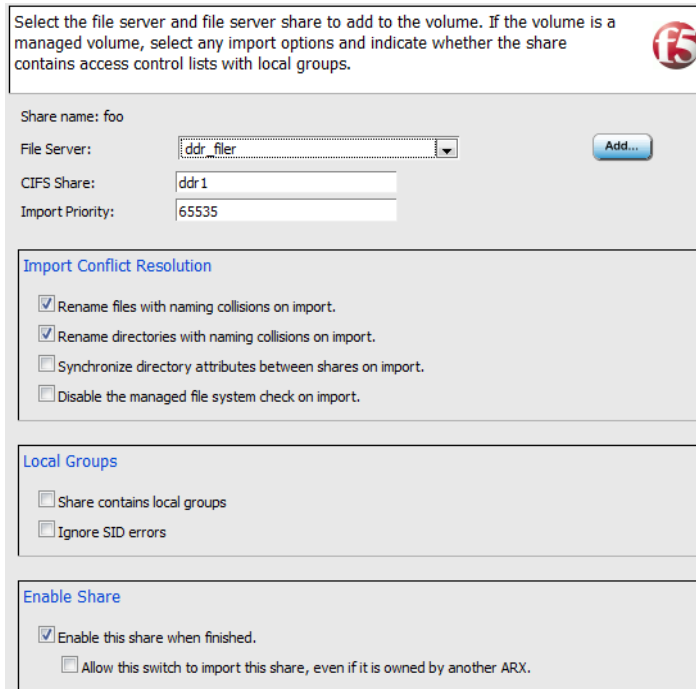
Add File Server

Name	ddr_filer
Primary IP Address	10.10.73.4
Secondary IP Address	<input type="text"/> <input type="button" value="Remove"/>
	<input type="button" value="Add"/>
CIFS Port	445
CIFS Connection Limit	0 <small>The value of 0 means unlimited. The CIFS connection limit for file servers sharing this IP address will be updated to match.</small>
Supports Snapshots	<input checked="" type="checkbox"/> This file server supports snapshots.
File Server Type	Data Domain
Management IP Address	10.10.73.4
Management Protocol	SSH
Management Proxy User	dd_user
NFS TCP Connections	1
Kerberos Service Principal Name (optional)	<input type="text"/>

Figure 7 Configuring a new file server

- 8. From the **File Server** list, select the name of the Data Domain system.
In our example, we type **ddr_filer**.
- 9. In the **CIFS Share** box, type the name of the CIFS share.
In our example, we type **ddr1**.

10. Leave other settings at the default level (see Figure 8).



Select the file server and file server share to add to the volume. If the volume is a managed volume, select any import options and indicate whether the share contains access control lists with local groups.

Share name: foo

File Server: ddr_filer Add...

CIFS Share: ddr1

Import Priority: 65535

Import Conflict Resolution

- Rename files with naming collisions on import.
- Rename directories with naming collisions on import.
- Synchronize directory attributes between shares on import.
- Disable the managed file system check on import.

Local Groups

- Share contains local groups
- Ignore SID errors

Enable Share

- Enable this share when finished.
- Allow this switch to import this share, even if it is owned by another ARX.

Figure 8 Adding the Data Domain to the ARX

11. Click the **Next** button. The summary page opens.
12. Review the changes, and then click the **Finish** button.

We now verify that Data Domain system storage had been added as to the ARX.

To verify the Data Domain system has been added from the Configuration utility

1. From the navigation pane, click **Managed Volumes**.
2. From the **Volume** list, select the appropriate volume. In our example, we select **/data**.
3. Click the **Shares** tab. You should see the newly created Share (**ddr_share** in our example), with a green status ball.

To add the Data Domain system as Tier 2 storage from the command line

1. Use the following command to add the Data Domain system as a new filer.

```
global
  external-filer ddr_filer
  cifs connection-limit 600
  cifs-port 445
  ip address 10.10.73.4
end
```

2. To add the Data Domain system share as a new share, use the following command:

```
global
  namespace <target-namespace>
    volume <target-volume>
      share ddr_share
    no import rename-directories
    no import rename-files
    filer ddr_filer cifs ddr1
  enable
end
```

◆ Note

*If you would like to copy and paste multiple lines of code from the command line at once, you may need to issue the command **no terminal confirmation** prior to copying and pasting the code from the command line*

To verify the Data Domain system has been added from the command line

1. Use the following syntax to show the new share:

```
show namespace <target-namespace> volume <target-volume>
```

In our example, we type

```
show namespace my_namespace_cifs volume /data
```

The following figure shows an example of the output.

```

Namespace "my_namespace_cifs" Configuration
Description
Metadata Cache Size: 512 MB
Proxy User: acopia

Domain Information
-----
CIFS Authentication method: Kerberos

Supported Protocols
-----
cifs

Participating Switches
-----
arxla (vpu 1) [Current Switch]

Volumes
-----
/data
CIFS : compressed files: no; named streams: no; persistent ACLs:
yes
sparse files: no; Unicode on disk: yes; case sensitive: no

Volume freespace: 6.2GB (automatic)
Auto Sync Files: Enabled
Metadata size: 252k
Metadata free space: 5.9GB
Oplock support: Enabled
Notify-change mode: Normal
CIFS path cache: Not Enabled
State: Enabled
Host Switch: arxla
Instance: 1
VPU: 1 (domain 1)
Files: 5 used, 9.9M free

Metadata shares:

-----
Filer      Backend Path      Contains Metadata  Status
-----
filer1     meta1              Yes                 Online

Share ddr_share
Filer      ddx_filer [10.10.73.4]
CIFS Share ddx1
Features   cifs-acls cifs-case-blind
CIFS Maximum Request Size 16644
SID Translation No
Ignore SID errors No
Status Online
Import Rename Files No
Import Rename Directories No
Free space on storage 291MB (305,770,496 B)
Transitions 33
Last Transition Wed Mar 17 20:05:25 2010
Share my_share_cifs
Filer      filer1 [10.10.73.1]
CIFS Share data1
Features   cifs-acls cifs-case-blind
CIFS Maximum Request Size 16644
SID Translation No
Ignore SID errors No
Status Online
Volume Root Backing Yes
Import Rename Files No
Import Rename Directories No
Free space on storage 5.9GB (6,427,648,000 B)
Transitions 13
Last Transition Wed Mar 17 20:05:10 2010

```

Figure 9 Viewing the `ddr_share` and its status


Adding policies for data migration between Tier 1 and Tier 2

The final procedure in this configuration is to define two policies, or place rules, on the ARX for placing data on the appropriate tier (Tier 1 or Tier 2). Typically, the place rule policies are defined so data that is modified more frequently is placed on Tier 1, while data which is modified less frequently (such as archives) is placed or moved to Tier 2 storage.

In the following example, policies are defined so data that has a *last modified time* older than one day are moved from Tier 1 to Tier 2, while data that has a *last modified time* newer than one day are moved from Tier 2 to Tier 1.

To add policy for data migration between Tier 1 and Tier 2 from the Configuration utility

1. From the navigation pane, expand the **Policy** menu, and then click **Place Rules**.
2. Click the **Tiered Storage** button. The Tiered Storage wizard opens.
3. In **Policy Name** box, type a name for this policy. In our example, we type **tier1-to-tier2**.
4. From the **Namespace** list, select the relevant namespace. In our example, we select **my-namespace-cifs**.
5. From the **Managed Volume** list, select the relevant volume. In our example, we select **/data**.
6. In the **Number of tiers** box select **2**.
7. Click the **Next** button.
8. Select the tier 1 share from the list and click the **Next** button. In our example, we select **my_share_cifs[share]**.
9. Select the tier 2 share from the list and click the **Next** button. In our example, we type **ddr_share[share]**.
10. The next step is to either select or create a fileset. If you have already defined a fileset, continue to the next step. If you have not defined a fileset, click the **Add** button.
 - a) In the **Fileset name** box, type a name for the Fileset. In our example we type **fileset_older_1day**.
 - b) From the **Fileset Type** list, select **age**.
 - c) From the **File Attribute Criteria** list, select **Last Modified**.
 - d) In the **File Age Criteria** section, in the **Older Than** box, type **1** and select **days** from the list. Leave the **Newer than** box at **0**.
 - e) Click **Save**. You return to the Fileset Migration wizard (see Figure 10).

Create a new fileset. 

Fileset Name:

Fileset Type:

File Attribute Criteria:

File Age Criteria:

Newer than

Older than

Use evaluation interval.

Evaluate Every

Starting At: / / - :

Month Day Year Hr Min

Figure 10 *Creating a new Fileset*

11. From the **Fileset** list, select the name of the Fileset you created. In our example, we select **fileset_older_1day**.
12. The policy needs a schedule on which to run. If a schedule already exists choose it from the list. To add a new schedule, click **Add** next to the **Schedule** field. Specify a **Schedule Name** and an interval at which to run. Select **Save** to continue. Click the **Next** button.
13. Review the summary, and then click the **Finish** button.

The next step is to verify the policy we just created.

To verify policy for data migration between Tier 1 and Tier 2 from the Configuration utility

1. From the navigation pane, expand the **Policy** menu, and then click **Place Rules**.
2. Click the name of the Place Rule you just created. In our example, we click **tier1-to-tier2**.
3. From the Place Rule Details page, you can view the statistics (such as **Total Files Migrated**, **Number of Directories Scanned** and **Number of files Scanned** to verify files are being migrated. In the following example, you can see 4 files have been migrated.

Place Rule Details

Namespace: Volume: Rule:

Status		Queue Statistics	
File Migration Status	Complete	First-time Migrates	0
Volume Scan Status	Complete	Requeued Migrates	0
Inline Status	Enabled	Queued Directory Promotes	0
Configuration		Cumulative Statistics	
Rule	tier1-to-tier2	Total Files Migrated	4
From	fileset_older_1day (fileset) files only	Total Directories Promoted	0
Source	my_share_cifs (share)	Total Failed Migrations	0
Target	ddx_share (share)	Total Failed Directory Promotes	0
Reports	tier1-to-tier2, Verbose, Delete Empty Reports	Total Retried Migrations	0
Migration Limit	n	Total Canceled Migrations	0

Figure 11 Viewing the Place Rule statistics to verify the policy is functioning

To add policy for data migration between Tier 1 and Tier 2 from the command line

1. Use the following syntax to create two filesets, one for moving file from tier1 to tier2 storage and another one for moving file in reverse direction:

```
global
  policy-simple-age-fileset fileset_older_1day
    select-files older-than 1 days
    last modified
    every 1 hours
    exit
  policy-simple-age-fileset fileset_newer_1day
    select-files newer-than 1 days
    last modified
    every 1 hours
    exit
exit
```

2. The following command creates the Place Rule that uses the Filesets you just made:

```
global
  namespace my_namespace_cifs
    volume /data
      place-rule tier1-to-tier2
        report tier1-to-tier2 verbose delete-empty
        source share my_share_cifs
        from fileset fileset_older_1day
        target share ddx_share
```

```
    enable
    exit
place-rule tier2-to-tier1
    report tier2-to-tier1 verbose delete-empty
    source share ddr_share
    from fileset fileset_newer_1day
    target share my_share_cifs
    enable
    exit
exit
exit
```

To verify the policy for data migration between Tier 1 and Tier 2 from the command line

1. Use the following syntax to show policy detail information:

```
show policy <name_space> <volume> <rule-name>
```

In our example, we type

```
show policy my_namespace_cifs /data tier1-to-tier2
```

The following figure shows an example of the output.

```

arx1a> show policy my_namespace_cifs /data tier1-to-tier2

Namespace:      my_namespace_cifs
Volume:        /data

Place Rule:     tier1-to-tier2

Configuration:
  From fileset:      fileset_older_1day
(files only)
  Source share:     my_share_cifs
  Target share:     ddx_share
  Report:           tier1-to-tier2,
Verbose, Delete Empty Reports
  Migrate limit:    0
  Volume Scan:     Enabled
  Inline Notifications: Enabled
  Promote Directories: Disabled

State:          Enabled

Status:
  Volume Scan Status: Complete
  File Migration Status: Complete
  New File Placement Status: Enabled

Cumulative Statistics:
  Total Files Migrated: 4
  Total Directories Promoted: 0
  Total Failed Migrations: 0
  Total Failed Directory Promotes: 0
  Total Retried Migrations: 0
  Total Canceled Migrations: 0
  Total Hard Links Skipped: 0
  Total Files Placed Inline: 0
  Total File Renames Processed Inline: 0
  Total Directories Placed Inline: 0
  Total Directory Renames Processed Inline: 0
  Number of Inline Overflow Errors: 0
  Number of Scans Performed: 4

Queue Statistics:
  First-time Migrates: 0
  Requeued Migrates: 0
  Queued Directory Promotes: 0

Last Scan Statistics:
  Scan Started: Wed Mar 17 16:09:10
2010
  Scan Completed: Wed Mar 17 16:09:10
2010
  Elapsed Time: 00:00:00
  Scan Report:
tier1-to-tier2_20100317135607.rpt
  Number of Files Scanned: 0
  Number of Directories Scanned: 1
  Number of Files in Fileset: 0
  Number of Files Migrated: 0
  Size of Files Migrated: 0 (0 on source)
  Number of Directories Promoted: 0
  Number of Failed Migrations: 0
  Number of Failed Directory Promotes: 0

```

Figure 12 Place Rule statistics

This completes the configuration. The ARX should now migrate files between each storage tier based on the defined file-set. For the example in this document, ARX migrates files with a Last Modified time of older than one day from Tier 1 storage to the Data Domain system, and migrate files

with Last Modified time of newer than one day from the Data Domain system to Tier 1 storage. Newly created files are also placed on Tier 1 storage.

For more information on configuring the F5 ARX, refer to the documentation, available on [Ask F5](#).

Appendix A: Verifying data compression is occurring

To verify that the data compression is occurring with a given deployment from the web interface, log on to Data Domain device and observe the summary statistics on the welcome screen. See the following example:

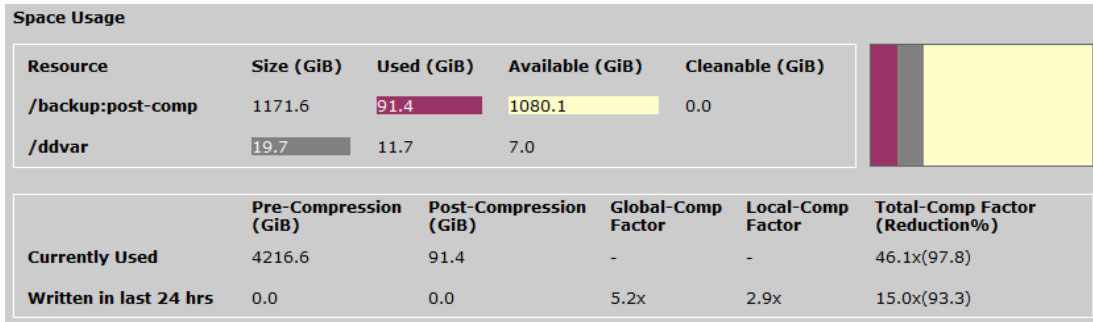


Figure 13 File System Compression Summary showing the reduction percentage

To verify that the data compression is occurring with a given deployment by command line, log on to Data Domain device console and type the following command: `filesys show compression`

The following figure shows an example of the output from this command.

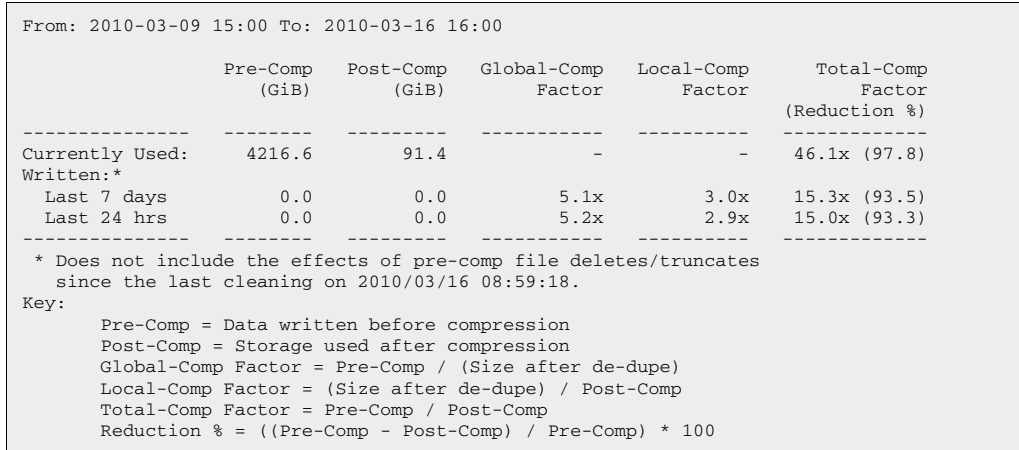


Figure 14 Verifying data compression from the command line

The statistics above show the comparison between **Pre-Comp** (data written before compression) and **Post-Comp** (storage used after compression). The Data Domain system applies compression in 2 phases: Global Compression and Local Compression.

Global Compression is the phase 1 de-duplicated data statistics. This occurs on data ingestion. As data arrives the Global Compression algorithm will identify repeating data patterns and create a set of references and checksums representing the source data that are written to disk.

As a second phase, Data Domain performs local compression, which analyzes the files with LZ, GZFAST, or GZ compression algorithms.