

BACKUP AND RECOVERY OF THE EMC GREENPLUM DATA COMPUTING APPLIANCE

EMC Data Domain DD890, Network File System, EMC Data Domain Boost

- Data warehouse protection
- Major storage savings with data deduplication
- Simple and fast restore

EMC Solutions Group

Abstract

This white paper provides insight into how EMC® Data Domain® deduplication storage systems effectively deal with the data growth, retention requirements, and recovery service levels that are essential to businesses. Data Domain's industry-leading technology provides a powerful backup, archiving, and disaster recovery solution that can scale with the most demanding data center requirements. This white paper explores the various practices and considerations for backing up EMC Greenplum™ Data Computing Appliance data to Data Domain systems and how to effectively exploit Data Domain's leading-edge technology.

March 2012



Copyright © 2012 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All trademarks used herein are the property of their respective owners.

Part Number H8038.6

Table of contents

Executive summary	5
Business case.....	5
Solution overview	6
Key results	7
Introduction	8
Purpose	8
Scope	8
Audience.....	8
Terminology	8
Overview of components	11
Introduction to the components	11
EMC Greenplum DCA.....	11
EMC Data Domain DD890 deduplication storage system	12
Configuration	13
Overview of the DCA and Data Domain solution	13
Physical environment.....	13
Hardware resources	14
Software resources	14
Data Domain DD890 system features, integration, and administration	15
Introduction to the DD890 system.....	15
System features	15
Integration	15
Administration	15
Scalable deduplication storage.....	15
Easy integration	15
Integration into an existing Greenplum DCA	16
Operational simplicity.....	16
Data integrity	16
Data compression	16
SISL	17
Multipath and load-balancing configuration.....	17
DD Boost.....	17
Design considerations	17
Data Domain Enterprise Manager.....	18
Data Domain file system	18
Creating and working with shares on the Data Domain system	19

Installing and configuring the Data Domain DD890 system with the DCA	20
Overview.....	20
Installing the Data Domain DD890 system	20
Configuring the Data Domain DD890 system.....	20
Data Domain and DCA connectivity	20
NFS – Configuring the DD890 and DCA.....	21
NFS – Backing up to Data Domain from the DCA.....	25
NFS – Restoring data on the DCA using Data Domain	26
DD Boost – Configuring the DD890 and DCA	28
DD Boost - Backing up to Data Domain from the DCA	31
DD Boost – Restoring data on the DCA from Data Domain	32
Backup schedule	34
NFS test results	35
Test objectives.....	35
Test scenarios.....	35
Test 1 results	35
Test 2 results	36
Test 3 results	40
Test 4 results	42
DD Boost test results.....	43
Test objectives.....	43
Test scenarios.....	43
Test 1 results	43
Test 2 results	44
Test 3 results	48
Test 4 results	50
Conclusion	51
Summary	51
Findings.....	51
References	53
White papers	53
Product documentation.....	53
Supporting information	54
Interconnect 1 – Converting ports 18 and 19 from a LAG to switch ports	54
Interconnect 2 – Converting ports 18 and 19 from a LAG to switch ports	55

Executive summary

Business case

A data warehouse enables organizations to perform educated, strategic, and reactive business and operational decisions that contribute to the overall success of an organization, both functionally and economically.

To ensure that an organization can make the right decisions in the quickest amount of time, its data warehouse must be reliable, secure, high performing, and extremely flexible to support the growing amounts of data, concurrency of many users, and increasing complexity of analysis.

Therefore, an organization's data needs to be loaded into the system so that it can be queried intelligently and quickly. The data also needs to be backed up, recovered, and restored efficiently and cost-effectively.

In today's business environments, data warehousing and business intelligence (DW/BI) systems are:

- Growing exponentially over time and also as the business grows, often leading to many terabytes of storage requiring protection.
- Frequently becoming unmanageable, lacking the ability to back up or provide effective disaster recovery (DR).

With the enormous amounts of storage increasingly required for backing up data, businesses are finding it ever more difficult to achieve storage efficiencies.

EMC has created a purpose-built data warehouse solution, the EMC® Greenplum™ Data Computing Appliance (DCA). The DCA addresses essential business requirements and ensures predictable functional, performance, and scalability results. This eliminates the guesswork and unpredictability of a highly customized, in-house solution. Many DCA product options are available, this solution uses the full-rack DCA.

Data protection for a data warehouse is a critical IT discipline and businesses have historically chosen simple approaches, such as periodically creating a full backup to tape. With the rapid growth in the volume of data stored, a simple, periodic full backup to tape or to a non-deduplicating, disk-based storage system is no longer economical. In addition, tape or non-deduplicated disk-based backups do not provide the recoverability or reliability of next-generation backup solutions like EMC Data Domain® deduplication storage systems.

The DCA, combined with Data Domain systems, provides a total solution for data warehousing deployment that addresses all these key challenges. This white paper illustrates how to backup data from the DCA to a Data Domain system, and how to subsequently recover and restore the data to the DCA. It is strongly recommended that you read this paper in conjunction with the white paper *EMC Greenplum Data Computing Appliance: Architecture, Performance and Functions - A Detailed Review*.

Solution overview

Data warehouse environments today demand a more comprehensive strategy for data protection, security, and high availability than ever before, based on numerous causes of data loss. Data recovery options must align with application and business requirements to yield the highest availability. Organizations with traditional backup and recovery systems face many challenges, for example:

- Efficient use of both infrastructure and people to support the business
 - Improve IT efficiency—save hours of staff time and boost user productivity
 - Reduce operational and infrastructure costs by eliminating the requirements of tape and reducing data center footprint requirements to a single rack
 - Reduce the need for massive backup storage
 - Correct and reduce costs—match infrastructure costs with changing information value through efficient, cost-effective storage
- Simplicity
 - Ease of integration
 - Ease of management
- Faster backup and restores
 - Minimize backup windows to reduce the impact on application and system performance
 - Meet more aggressive recovery time objective (RTO) service-level agreements (SLAs) set by business owners by maximizing end-to-end recovery performance

The DCA provides an end-to-end DW/BI solution packaged within a manageable, consolidated, self-contained data warehouse appliance that can be easily integrated into an existing data center. This white paper describes a backup and recovery environment for the DCA using EMC's Data Domain deduplication storage system.

The backup and recovery solution was implemented using Greenplum's **gpcrondump** backup utility. The utility was deployed using the following two methods:

- Over a network file system protocol (NFS)
- Using EMC's Data Domain Boost

Both methods back up the data to a Data Domain deduplication storage system.

The **gpcrondump** backup utility automates the parallel backup of distributed Greenplum databases across multiple Segment Servers. For NFS, each of the Segment Servers has a directory mapped to a Data Domain system. For Data Domain Boost, data is backed up to a dedicated storage unit on the Data Domain system.

New data protection solutions led by Data Domain's architecture improve data recoverability, while increasing the probability for a business to survive most types of outages. This is increasingly critical since most businesses can no longer survive without their IT function. A backup solution that provides an easy and fast method of data recovery is a necessary requirement rather than an optional one.

EMC has ensured that leading-edge technology is available to support the backup and recovery of its DCA. This white paper demonstrates how the Data Domain system achieves this in the most simple, efficient, and cost-effective way by providing greatly increased storage savings.

Key results

With data deduplication, organizations can reduce the amount of data that needs to be stored over an extended period of time. This offers cost savings in terms of the numbers of disks or tapes required for backup. In addition, organizations can fundamentally change the way they protect their backup and archive data. Only unique, new data is written to the disk, eliminating the restore penalties associated with incremental backups.

This solution describes a backup and recovery environment for the DCA using the Data Domain DD890 system. The test results show that this solution meets the business challenges faced by many organizations today through:

- Operational ease and efficiency through the utilization of MPP architecture to efficiently back up, in parallel, across the network to the Data Domain system.
- Major space saving advantages using Data Domain inline deduplication—in this case, test results with DD Boost demonstrate savings of 34.5x on repetitive nightly backup and an aggregate 11.3x cumulative storage reduction over a week of uncompressed backups to the Data Domain system. One of the key benefits of Data Domain deduplication storage systems is a reduction in the need for excessive amounts of backup storage.
- Fast restore times for returning uncompressed and deduplicated data to the DCA. The test results with DD Boost demonstrate that the combination of **gpcrondump** backups and Data Domain deduplication technology provides an average backup of 13.08 TB/hour and a restore of 5.90 TB/hour. This enables the backup and recovery of a Greenplum full-rack DCA (36 TB uncompressed database) in 2.75 and 6.1 hours respectively.

The Data Domain DD890 system in this solution provides an effective disk-based backup target that significantly minimizes storage while providing long-term retention.

Introduction

Purpose

The purpose of this white paper is to illustrate how Data Domain's industry-leading technology provides a powerful backup, archiving, and disaster recovery solution that can scale with the most demanding data center requirements, such as data warehousing. This white paper also explores the various practices and considerations for backing up DCA data to the Data Domain system and how to effectively exploit this technology.

Scope

The scope of this white paper is to document:

- DCA features and benefits (at a high level)
- Data Domain DD890 system features, as utilized in this solution
- Installing and configuring the DD890 system for the DCA
- The solution test objectives, scenarios, and results

Audience

This white paper is intended for:

- Field personnel who are tasked with implementing a backup and recovery solution for the DCA or a data warehouse
- Customers, including IT planners, storage architects, and database administrators involved in evaluating, acquiring, managing, operating, or designing a backup and recovery solution for the DCA or a data warehouse
- EMC staff and partners, for guidance and for the development of proposals

Terminology

Table 1 defines key terms used in this document.

Table 1. Terminology

Term	Definition
Addressable capacity	The amount of physical space available on a Data Domain system to store deduplicated and compressed backup images.
Business intelligence (BI)	The effective use of information assets to improve the profitability, productivity, or efficiency of a business. IT professionals use this term to refer to the business applications and tools that enable such information usage. The source of information is frequently the data warehouse.
Cumulative compression factor	The ratio of the logical storage size to the physically stored space.
Data Domain	EMC deduplication storage system.

Term	Definition
Data Domain Boost (DD Boost)	Extends the backup optimization benefits of Data Domain deduplication storage solutions by distributing parts of the deduplication process to the backup server or application client. DD Boost dramatically increases throughput speeds, minimizes backup LAN load, and improves backup server utilization.
Data warehousing (DW)	The process of organizing and managing information assets of an enterprise. IT professionals often refer to the physically stored data content in some databases managed by database management software as the data warehouse. They refer to applications that manipulate the data stored in such databases as DW applications.
Deduplication	<p>Deduplication is similar to data compression, but it looks for redundancy of very large sequences of bytes across very large comparison windows. Long sequences (greater than 8 KB) are compared to the history of other such sequences and, where possible, the first uniquely stored version of a sequence is referenced rather than stored again. In a storage system, this is all hidden from users and applications, so the whole file is readable after being written.</p> <p>Eliminating redundant data can significantly shrink storage requirements and improve bandwidth efficiency.</p>
Local compression	<p>Standard lossless compression algorithms. The algorithms available on a Data Domain system include lz (Lempel-Ziv), gz, and gzfast:</p> <ul style="list-style-type: none"> • lz – The default algorithm that gives the best throughput. Data Domain recommends the lz option. • gzfast – A zip-style compression that uses less space for compressed data, but more CPU cycles (twice as much as lz). gzfast is the recommended alternative for sites that want more compression at the cost of lower performance. • gz – A zip-style compression that uses the least amount of space for data storage (10 percent to 20 percent less than lz on average; however, some datasets get much higher compression). This also uses the most CPU cycles (up to five times as much as lz). The gz compression type is commonly used for near-line storage applications in which performance requirements are low.
Logical storage size	The total size of all backup images in all pools on a Data Domain system. This total size includes all pools mapped to a Data Domain system instance, which can include primary disk pools and clone storage pools.

Term	Definition
Massively Parallel Processing (MPP)	A type of distributed computing architecture where tens to hundreds of processors team up to work concurrently to solve large computational problems.
NFS	Network file system.
Redundant Array of Independent Disks (RAID)	A method of organizing and storing data distributed over a set of physical disks, which logically appear to be one single storage disk device to any server host and operating system performing I/O to access and manipulate the stored data. Frequently, redundant data is distributed and stored inside this set of physical disks to protect against loss of data access, should one of the drives in the set fail.
Scale out	A technique that increases the total processing power of a system by adding additional independent computational nodes, as opposed to augmenting a single, large computer with incremental disk, processor, or memory resources.
Shared-Nothing Architecture	A distributed computing architecture made up of a collection of independent, self-sufficient nodes. This is in contrast to a traditional central computer that hosts all information and processing in a single location.
Storage Unit	The Data Domain system exposes pre-made disk volumes called storage units to a DD Boost-enabled media server. Multiple media servers, each with the Data Domain DD Boost plug-in, can use the same storage unit on a Data Domain system as a storage server.

Overview of components

Introduction to the components

This section identifies and briefly describes the components deployed in the solution environment. The components used are:

- EMC Greenplum DCA
- EMC Data Domain DD890 deduplication storage system

EMC Greenplum DCA

The DCA is a purpose-built, highly scalable, parallel DW appliance that architecturally integrates database, compute, storage, and network into an enterprise-class, easy-to-implement system. The DCA brings in the power of MPP architecture, delivers the fastest data loading capacity and the best price/performance ratio in the industry without the complexity and constraints of proprietary hardware.

The DCA is offered in multiple-rack appliance configurations to achieve the maximum flexibility and scalability for organizations faced with terabyte to petabyte scale data opportunities.

Many DCA product options are available. This solution uses the full-rack DCA. Table 2 briefly describes the main components of DCA.

Table 2. Main components of DCA

Item	Description
Greenplum Database™	Greenplum Database is an MPP database server, based on PostgreSQL open-source technology. It is explicitly designed to support business intelligence (BI) applications and large, multi-terabyte data warehouses.
Greenplum Database system	An associated set of Segment Instances and a Master Instance running on an array, which can be composed of one or more hosts.
Master Servers	The servers responsible for the automatic parallelization of queries.
Segment Servers	The servers that perform the real work of processing and analyzing the data.
Interconnect Bus	The Interconnect Bus provides high-speed communication between Master and Segment Servers.
Admin Switch	The Admin Switch provides the management interface between the servers and additional racks.

For more information about the DCA, refer to the following white paper *EMC Greenplum Data Computing Appliance: Architecture, Performance and Functions - A Detailed Review*.

**EMC Data Domain
DD890
deduplication
storage system**

EMC Data Domain DD890 deduplication storage systems dramatically reduce the amount of disk storage needed to retain and protect enterprise data. By identifying redundant data as it is being stored, Data Domain provides a storage footprint that is up to 30 times smaller, on average, than the original dataset. Backup data can then be efficiently replicated and retrieved over existing networks for streamlined disaster recovery and consolidated tape operations. This allows Data Domain appliances to integrate seamlessly into database architectures, maintaining existing backup strategies with no changes to scripts, backup processes, or system architecture.

The Data Domain appliance is the industry's fastest, most cost-effective and scalable single-controller deduplication storage solution for disk-based backup and network-efficient DR.

The Data Domain Stream-Informed Segment Layout (SISL™) scaling architecture enables the fast-inline deduplication throughput of the DD890. A CPU-centric approach to deduplication delivers a high throughput while minimizing the number of disk spindles required.

Configuration

Overview of the DCA and Data Domain solution

This section illustrates the architectural layout of the DCA and Data Domain solution. It also provides details of the hardware and software resources that were used in the solution.

Physical environment

Figure 1 illustrates the architectural layout of the solution.




Figure 1. Solution architecture layout

Hardware resources

Table 3 and Table 4 list the hardware used to validate this solution.

Table 3. Data Domain DD890 system

Equipment	Quantity	Specification
DD890	1	96 GB memory 4 x internal SATA drives 2 x 1 GbE network interfaces 1 x dual-port 10 GbE optical network interface card
Expansion shelves (32 TB)	6	ES20 disk shelves with 16 x 2 TB SATA disks

Table 4. Backup network

Equipment	Ports	Specification
DCA Interconnect Bus*	2	10 GbE

*The existing Greenplum DCA Interconnect Bus is designed to accommodate the backup network port requirements.

Software resources Table 5 lists the software used to validate this solution.

Table 5. Software resources

Software	Version	Comment
Data Domain OS	5.1.0.4-259167	Data Domain operating system
EMC Greenplum Database	4.2.0.0	Minimum version to support DD Boost
Fabric OS	6.3.2a1	CEE/FCoE network

Data Domain DD890 system features, integration, and administration

Introduction to the DD890 system

This section describes how Data Domain systems ensure data integrity and provide multiple levels of data compression, reliable restorations, and multipath configurations. The Data Domain Data Invulnerability Architecture protects against data loss from hardware and software failures.

This section also provides details on DD890 system integration and administration.

System features

- Scalable deduplication storage
- Easy integration
- Operational simplicity
- Data integrity
- Data compression
- SISL
- Multipath and load-balancing configuration
- DD Boost

Integration

- Design considerations
- Data Domain Enterprise Manager
- Data Domain file system

Administration

- Working with shares
- Creating shares on the Data Domain system

Scalable deduplication storage

A single DD890 system achieves a single-stream throughput of up to 1.44 TB/hour—performance that is imperative for protecting large, business-critical databases in the data center. The DD890 provides an aggregate throughput of up to 14.7 TB/hour using multiple backup policies. Data Domain systems store each unique data segment only once and save significant physical storage by substituting small references for each identical redundant segment. The DD890 offers up to 30 times data reduction for enterprise recovery images, enabling cost-effective retention on disk for high-speed recoveries.

Easy integration

The Data Domain DD890 is qualified with all leading enterprise backup software and archiving applications and easily integrates into existing software infrastructures without change for either data center or distributed office data protection.

Data Domain systems integrate easily into existing data centers. All Data Domain systems can be configured as storage destinations for leading backup and archiving

applications using NFS, common internet file system (CIFS), Data Domain Boost, or virtual tape library (VTL) protocols. Consult the compatibility matrices for information about the applications that work with the different configurations. Multiple backup servers can share one Data Domain system.

Integration into an existing Greenplum DCA

The DCA architecture is designed to help you easily integrate Data Domain systems in a nondisruptive, seamless manner. You can reserve Port 19 in each of the Interconnect switches for Data Domain connectivity. All that is required is to connect the Data Domain system directly into the DCA environment and start the configuration steps for NFS or DD Boost. Refer to the Data Domain and DCA connectivity section for more options.

Operational simplicity

Data Domain systems are simple to install and manage. Connect an appliance to the backup server either as a file server via Ethernet or as a VTL via Fibre Channel. Data Domain Boost for EMC NetWorker® and Symantec OpenStorage are also supported; all three interfaces can be used simultaneously.

Data integrity

When writing to disk, the DD890 creates and stores checksums and self-describing metadata for all data received. After writing the data to disk, the DD890 recomputes and verifies the checksums and metadata. An append-only write policy guards against overwriting valid data.

After a backup is completed, a validation process looks at what was written to disk to check that all file segments are logically correct within the file system and that the data is the same on the disk as it was before being written to disk. In the background, the Online Verify operation continuously checks that the data on the disks is correct and unchanged since the earlier validation process.

The back-end storage is set up in a double parity RAID 6 configuration (two parity drives). Additionally, hot spares are configured within the system. Each parity stripe has block checksums to ensure that the data is correct. The checksums are constantly used during the online verify operation and when data is read from the Data Domain system. With double parity, the system can fix simultaneous errors on up to two disks.

To keep data synchronized during a hardware or power failure, the Data Domain system uses non-volatile RAM (NVRAM) to track outstanding I/O operations. An NVRAM card with fully-charged batteries (the typical state) can retain data for a minimum of 48 hours.

When reading data back on a restore operation, the Data Domain operating system (DD OS) uses multiple layers of consistency checks to verify that restored data is correct.

Data compression

The DD OS stores only unique data. Through Global Compression™, a Data Domain system pools redundant data from each backup image. The storage of unique data is invisible to backup software, which sees the entire virtual file system. DD OS data compression is independent of a data format. Data can be structured, such as databases, or unstructured, such as text files. Data can be from file systems or raw volumes.

Typical compression ratios are 20:1 on average over many weeks. This assumes weekly full and daily incremental backups. A backup that includes many duplicate or similar files (files copied several times with minor changes) benefits the most from compression. Depending on backup volume, size, retention period, and rate of change, the amount of compression can vary.

The best compression happens with backup volume sizes of at least 10 mebibytes (MiB—a unit of data storage that is exactly 1,048,576 bytes, the base 2 equivalent of MB). To take full advantage of multiple Data Domain systems, a site that has more than one Data Domain system should consistently back up the same client system or set of data to the same Data Domain system. For example, if a full backup of all sales data goes to Data Domain system A, the incremental backups and future full backups for sales data should also go to Data Domain system A.

A Data Domain system compresses data at two levels:

- **Global compression**—compares received data to data already stored on disks. Duplicate data does not need to be stored again, while new data is locally compressed before being written to disk.
- **Local compression**—a Data Domain system uses a local compression algorithm developed specifically to maximize throughput as data is written to disk. The default algorithm (**lz**) allows shorter backup windows for backup jobs but uses more space. Local compression options provide a trade-off between performance and space usage.

For more information on how to change compression, refer to the *DD OS 5.1 Administration Guide*. Changing the algorithm immediately affects any new data written to the system. Any data already stored on the system will be recompressed during the next cleaning run, which may take much longer to run than usual.

SISL

Data Domain SISL enables high throughput, inline deduplication. SISL identifies 99 percent of the duplicate segments in RAM, inline, before storing to disk. In addition, it stores related segments and fingerprints together, so large groups can be read at once. With these patented techniques, Data Domain can utilize the full capacity of large SATA disks for data protection and minimize the number of disks needed to deliver high throughput. In the long term, SISL allows dramatic Data Domain system performance improvements as CPU speeds increase.

Multipath and load-balancing configuration

Data Domain systems that have at least two 10 GbE ports can support multipath configuration and load balancing. In a multipath configuration on the DD890 system, each of the two 10 GbE ports on the system is connected to a separate port on the backup server.

DD Boost

EMC Data Domain Boost significantly increases performance by distributing parts of the deduplication process to the backup server, simplifies disaster recovery procedures, and serves as a solid foundation for additional integration between backup applications and Data Domain systems.

Design considerations

Retention of data, frequency, rate of change, and backup policies influence the decision when determining the amount of storage required in the Data Domain

system. For this solution, the initial capacity was chosen to accommodate a simulated 10 weeks of backup of the DCA.

Data Domain Enterprise Manager

All Data Domain systems run the DD OS, which includes Data Domain Enterprise Manager, a simple web-based rich Internet application for managing Data Domain systems. DD Enterprise Manager provides both a graphical user interface (GUI) and a command line interface (CLI) for configuration management and monitoring all system operations. The web-based GUI, available through Ethernet connections, can manage up to 20 Data Domain systems (depending on the model) at any location. DD Enterprise Manager provides a single, consolidated management interface that allows for the configuration and operation of many system features and settings.

DD Enterprise Manager also provides real-time graphs and tables that enable users to monitor the status of system hardware components and configured features. Additionally, a command set that performs all system functions is available to users through the CLI. Commands configure system settings and provide displays of system hardware status, feature configuration, and operation.

The CLI is available through a serial console when a keyboard and monitor are directly attached to the Data Domain system, or remotely through an Ethernet connection using SSH or Telnet. For more information on Data Domain Enterprise Manager, refer to the *DD OS 5.1 Administration Guide*.

Data Domain file system

Data Domain systems are designed to be a highly reliable “storage of last resort” to provide longer-term onsite retention of backups. As new backups are added to the system, old backups are aged out. Such removals are normally done under the control of backup software (on the backup server) based on the configured retention period. This process is similar to configuring tape retention policies in which older backups are retired and the tapes are reused for new backups.

When backup software removes an old backup from a Data Domain system, the space on the Data Domain system becomes available only after the Data Domain system cleans the retired disk space. A good way to manage space on a Data Domain system is to retain as many online backups as possible, with some empty space (about 20 percent of the total space available) to comfortably accommodate backups until the next scheduled cleaning run.

Space utilization on a Data Domain system is primarily affected by:

- The backup policy and redundancy in the data
- The size, redundancy, and rate of change of the backup data
- The retention period specified in the backup software

High levels of compression result when backing up datasets with many duplicates and retaining them for long periods of time.

The file system supports the following interfaces:

- NFS
- CIFS

- Data Domain Boost
- VTL

For more information on the file system, refer to the *DD OS 5.1 Administration Guide*.

Creating and working with shares on the Data Domain system

When creating shares, assign client access to each directory separately and remove access from each directory separately.

Note If replication is to be implemented, a single destination Data Domain system can receive backups from both CIFS clients and NFS clients as long as separate directories are used for each. Do not mix CIFS and NFS data in the same directory.

Installing and configuring the Data Domain DD890 system with the DCA

Overview

Data Domain DD890 integrates easily into existing DCA deployments. The DD890 storage can be exported as NFS shares which can easily be mounted on the DCA Master and Segment Servers and utilized as a backup target. With Greenplum database support for DD Boost, the DD890 also integrates easily in this case. For this solution, both configurations were tested as outlined in this section.

Installing the Data Domain DD890 system

Install the Data Domain hardware, including a system console, as described in the *Data Domain Installation and Setup Guide*, which is shipped with the Data Domain system. The *Installation and Setup Guide* provides instructions for installing the Data Domain system, connecting it to an administrative console, and powering it on. After completing the installation and powering on the system, refer to the *DD OS 5.1 Initial Configuration Guide* for additional information.

Configuring the Data Domain DD890 system

When the installation is complete, the Data Domain Configuration Wizard starts automatically. The Configuration Wizard performs an “initial” configuration of the Data Domain system, configuring only what is needed for the most basic system setup, including licenses, network, filesystem, CIFS, and NFS. For more information about using the Configuration Wizard and changing or updating the configuration, refer to the *DD OS 5.1 Initial Configuration Guide*.

Data Domain and DCA connectivity

Two options are available when connecting the Data Domain system to the DCA:

- Direct connection via Port 19 on both the DCA Interconnect switches.
- Alternatively, Ports 18 and 19 can be configured as a two-port link aggregation on both the DCA Interconnect switches. Two separate switches can then be connected, one to each Interconnect switch, to increase the connectivity options.

Note To convert ports 18 and 19 from a two-port link aggregation to switch ports, refer to the section entitled Supporting information.

For this solution, the connection was made using two Multi-Mode Fibre (MMF) cables through optical Small Form-Factor Pluggable (SFP) devices, as shown in Table 6.

Table 6. DCA and Data Domain connectivity

DCA Interconnect Bus	Data Domain system interface
i-sw-1, Port 19	Slot 4, Port 1/eth4a
i-sw-2, Port 19	Slot 4, Port 2/eth4b

The hardware used for this connectivity option is listed in Table 7.

Table 7. Hardware used for DCA and Data Domain connectivity

Hardware	Part number	Quantity
Data Domain – 10GbE card	C-10G-L2PO	1
10 GbE SFP	BRSFP-10GSW1P	2
OM3 multimode optical cable	CTX-OM3-10M	2

A connectivity kit that consists of the cables and required SFPs is also available under part number DCA1-10GBEXT.

Depending on distance, connectivity can be provided through either MMF cables and optical SFP, or twin-ax cables. For more information about options and internal connectivity, refer to the white paper *EMC Greenplum Data Computing Appliance: Architecture, Performance and Functions - A Detailed Review*.

NFS – Configuring the DD890 and DCA

For NFS Backups, the Data Domain system must be configured to enable the DCA to access the NFS directories that are used as a target for the **gpcrondump** utility to store the backups.

The following example shows how the system was configured using two 10 GbE interfaces with the IP addresses 172.28.8.209 and 172.28.12.209, respectively:

On the DD890:

```
sysadmin@dd_890_1# net show settings
port      enabled  DHCP    IP address    netmask        type  additional setting
-----  -
eth0a     yes      no      172.28.4.209  255.255.248.0  n/a
eth0b     no       n/a     n/a           n/a            n/a
eth4a     yes      no      172.28.8.209  255.255.252.0  n/a
eth4b     yes      no      172.28.12.209 255.255.252.0  n/a
eth5a     no       n/a     n/a           n/a            n/a
eth5b     no       n/a     n/a           n/a            n/a
-----  -
* Value from DHCP
```

It is very important also to configure the hosts files on both the DCA and the Data Domain DD890 system. The following example shows how the Data Domain host entries look on the Master and Segment Servers.

On the DCA:

```
[gpadmin@mdw ~]$ grep "Data Domain" /etc/hosts
### Data Domain - Backup
172.28.8.209    dd_889_1_1    # Data Domain interface eth4a, Network 172.28.8.0
172.28.12.209 dd_890_1_2    # Data Domain interface eth4b, Network 172.28.12.0
```

To add the hosts to the Data Domain DD890 system, type:

```
net hosts add <ipaddr> <host-list>
```

The following example adds both 10 GbE interfaces of the host mdw:

On the DCA:

```
sysadmin@dd_890_1# net hosts add 172.28.8.204 mdw1-1
Added "mdw1-1" -> "172.28.8.204" mapping to hosts list.
sysadmin@dd_890_1# net hosts add 172.28.12.204 mdw1-2
Added "mdw1-2" -> "172.28.12.204" mapping to hosts list
```

Repeat with both 10 GbE interfaces for all Master and Segment Servers.

Perform the following steps to configure NFS:

1. Add the NFS share **/backup** and set the access level to the DCA network.
The network used in this case is **172.28.8.0/22** and **172.28.12.0/22**.

On the DD890:

```
sysadmin@dd_890_1# nfs add /backup 172.28.8.0/22,172.28.12.0/22
```

NFS export for **/backup** is added.

2. Ensure that the network was added correctly as shown in the following example:

On the DD890:

```
sysadmin@dd_890_1# nfs show clients
path          client          options
-----
/backup       172.28.8.0/22  (rw,no_root_squash,no_all_squash,secure)
/backup       172.28.12.0/22 (rw,no_root_squash,no_all_squash,secure)
```

3. Enable NFS on the Data Domain system as follows:

On the DD890:

```
sysadmin@dd_890_1# nfs enable
```

NFS server is enabled.

4. On the DCA Master Server, mount the **/backup** NFS as shown in the following example:

On the DCA:

- Create a temporary mount point for the **/backup**.

```
[root@mdw ~]# mkdir /backup_tmp
[root@mdw ~]# ls -la /backup_tmp
total 12
drwxr-xr-x  2 root root 4096 Aug 18 13:57 .
drwxr-xr-x 30 root root 4096 Aug 18 13:57 ..
```

- Mount the **/backup** directory on the Master Server (all one line):

On the DCA:

```
[root@mdw ~]# mount -t nfs -o
hard,intr,nfsvers=3,tcp,rsize=32768,wsize=32768
dd_890_1_1:/backup /backup_tmp/
```

5. Create the DCA directory and all the server subdirectories.

Note In this case, use the user **gpadmin** to avoid any issues with permissions while running the **gpcrondump** utility.

In the following example, the directory **DCA-01** and the subdirectories **Master01**, **Master02**, and **Seg01** to **Seg16** were created using **mkdir** as follows (all one line):

On the DCA:

```
[gpadmin@mdw ~]$ mkdir /backup_tmp/DCA-01;mkdir
/backup_tmp/DCA-01/Master{01,02};mkdir /backup_tmp/DCA-
01/Seg{01,02,03,04,05,06,07,08,09,10,11,12,13,14,15,16}
```

After all the directories are created, unmount the **/backup_tmp** directory as follows:

On the DCA:

```
[root@mdw ~]# umount /backup_tmp
```

The **gpcrondump** utility creates a directory named **db_dumps** when the backup is started for the first time.

Each server should use a different mount point for this purpose.

An NFS share was created for each server on the Data Domain and the corresponding mount point created on each server was **/backup/DCA-01**. You can use Table 8 as a reference to understand the mapping of the directories in this solution.

Table 8. NFS shares and mount points

DCA Server	Host Name	NFS Share on DD890	Mount Point on DCA
Master Server 01	mdw	/backup/DCA-01/Master01	/backup/DCA-01
Master Server 02	smdw	/backup/DCA-01/Master02	/backup/DCA-01
Segment Server 01	sdw1	/backup/DCA-01/Seg01	/backup/DCA-01
Segment Server 02	sdw2	/backup/DCA-01/Seg02	/backup/DCA-01
Segment Server 03	sdw3	/backup/DCA-01/Seg03	/backup/DCA-01
Segment Server 04	sdw4	/backup/DCA-01/Seg04	/backup/DCA-01
Segment Server 05	sdw5	/backup/DCA-01/Seg05	/backup/DCA-01
Segment Server 06	sdw6	/backup/DCA-01/Seg06	/backup/DCA-01
Segment Server 07	sdw7	/backup/DCA-01/Seg07	/backup/DCA-01
Segment Server 08	sdw8	/backup/DCA-01/Seg08	/backup/DCA-01
Segment Server 09	sdw9	/backup/DCA-01/Seg09	/backup/DCA-01
Segment Server 10	sdw10	/backup/DCA-01/Seg10	/backup/DCA-01
Segment Server 11	sdw11	/backup/DCA-01/Seg11	/backup/DCA-01
Segment Server 12	sdw12	/backup/DCA-01/Seg12	/backup/DCA-01

DCA Server	Host Name	NFS Share on DD890	Mount Point on DCA
Segment Server 13	sdw13	/backup/DCA-01/Seg13	/backup/DCA-01
Segment Server 14	sdw14	/backup/DCA-01/Seg14	/backup/DCA-01
Segment Server 15	sdw15	/backup/DCA-01/Seg15	/backup/DCA-01
Segment Server 16	sdw16	/backup/DCA-01/Seg16	/backup/DCA-01

Note The mount point can be created on each DCA using the **gpssh** utility. The **/backup** directory is created by the root user but **/backup/DCA-01** must be created by the **gpadmin** user.

The following example describes one way of creating the mount points. For more options, refer to the *DD OS 5.1 Administration Guide*.

On the DCA:

```
[root@mdw ~]# mkdir /backup ; mkdir /backup/DCA-01
[root@mdw ~]# gpssh -h smdw -v -e mkdir /backup ; mkdir
/backup/DCA-01
[root@mdw ~]# for NODE in `cat ~/shosts`; do gpssh -h $NODE -v -e
mkdir /backup; done;
[root@mdw ~]# su - gpadmin
[gpadmin@mdw ~]# for NODE in `cat ~/shosts`; do gpssh -h $NODE -v
-e mkdir /backup/DCA-01; done;
```

After the mount point is created, mount the Data Domain NFS share according to Table 8. The following command provides an example:

On the DCA:

```
[gpadmin@mdw ~] $ gpssh -h sdw33 -v -e mount -t nfs -o
rw,hard,tcp,rsize=131072,wsiz=131072 dd_890_1_1:/backup/DCA-
01/Seg01 /backup/DCA-01
```

To get the best balance across both Data Domain 10 GbE interfaces (eth4a and eth4b), servers mdw and sdw1 to sdw8 should use host **dd_890_1_1** (network 172.28.8.0, interface eth4a), while smdw and sdw9 to sdw16 should use host **dd_890_1_2** (network 172.28.12.0, interface eth4b).

Check that the directories are mounted correctly as follows:

On the DCA:

```
[gpadmin@mdw ~]$ for NODE in `cat ~/shosts`; do gpssh -h $NODE -v
-e df -h /backup/DCA-01 ; done;

[Reset ...]
[INFO] login sdw33
[sdw33] Filesystem                Size      Used Avail Use% Mounted on
[sdw33] dd_890_1_1:/backup/DCA-01/Seg01
[sdw33]                          130T      4.9T   125T    4% /backup/DCA-01
[INFO] completed successfully

[Cleanup...]
```

```
[Reset ...]
[INFO] login sdw34
[sdw34] Filesystem                Size  Used Avail Use% Mounted on
[sdw34] dd_890_1_1:/backup/DCA-01/Seg02
[sdw34]                          130T  4.9T  125T   4% /backup/DCA-01
[INFO] completed successfully
```

[Cleanup...]

[some text removed]

```
[Reset ...]
[INFO] login sdw48
[sdw48] Filesystem                Size  Used Avail Use% Mounted on
[sdw48] dd_890_1_2:/backup/DCA-01/Seg16
[sdw48]                          130T  4.9T  125T   4% /backup/DCA-01
[INFO] completed successfully
```

[Cleanup...]

If all the NFS shares are mounted successfully, the Data Domain system is ready to receive any backups run by the **gpcrondump** utility.

NFS – Backing up to Data Domain from the DCA

To automate routine NFS backups, Greenplum provides the **gpcrondump** utility, which can be called directly or from a crontab entry. The **gpcrondump** utility is a wrapper for the **gp_dump** command. **gpcrondump** also enables the backup of objects other than databases and data, such as database roles and server configuration files.

By default, **gpcrondump** creates data dump files in the Master and each Segment instance's data directory in **<data_directory>/db_dumps/YYYYMMDD**. By default the segment data dump files are compressed using **gzip**.

For NFS backups, the **gpcrondump** utility with the **-u** parameter is required. Without the **-u** parameter, the backups are stored locally in the DCA array.

For the NFS solution, the intention was to send the backups to the Data Domain system so the directory used was the NFS share created in the DD890 **/backup/DCA-01** as follows:

On the DCA:

```
[gpadmin@mdw db_dumps]$ pwd
/backup/DCA-01/db_dumps
[gpadmin@mdw db_dumps]$ ls -la
total 12
drwxrwxr-x 7 gpadmin gpadmin 382 Aug 20 15:01 .
drwxrwxr-x 4 gpadmin gpadmin 158 Aug 16 17:10 ..
drwx----- 2 gpadmin gpadmin 1342 Aug 16 19:37 20100816
drwx----- 2 gpadmin gpadmin 2608 Aug 17 19:19 20100817
drwx----- 2 gpadmin gpadmin 1563 Aug 19 23:57 20100819
drwx----- 2 gpadmin gpadmin 3328 Aug 20 15:21 20100820
```

All the NFS backups were generated by running the following command as **gpadmin** user:

For gpcrondump uncompressed backups:

```
gpcrondump -x tpch -g -G -u /backup/DCA-01 -b -a -z
```

Log files for the Greenplum Database management utilities are written to ~/gpAdminLogs by default. The naming convention for management log files is:

```
<script_name>_<date>.log  
e.g. gpcrondump_<date>.log
```

Table 9 lists the gpcrondump options used in NFS.

Table 9. NFS – gpcrondump options used

Parameter	Description
-x	Database name
-g	Copy config files
-G	Dump global objects
-u	Backup target directory
-a	Do not prompt
-b	Bypass disk space check
-z	Do not use compression

Note The **gpcrondump** compressed options --rsyncable are also available, but may not be suitable for environments with a high rate of random data changes. In order to gain the maximum benefit from EMC Data Domain deduplication technology, the **gpcrondump** uncompressed option **-z** is recommended.

It is also possible to automate parallel backups with **gpcrondump** by adding an entry to the crontab of the Master Server:

```
SHELL=/bin/bash  
GPHOME=/usr/local/greenplum-db- 4.2.0.1  
MASTER_DATA_DIRECTORY=/data/gpdb_master/Data-1  
PATH=$PATH:$GPHOME/bin  
01 0 * * * gpadmin gpcrondump -x tpch -c -g -G -a -q >>  
gpcrondump_Backup.log
```

For more information, refer to the *Greenplum Database 4.2 Administrator Guide*.

NFS – Restoring data on the DCA using Data Domain

For restoring a database, Greenplum provides the **gpdrestore** utility, which is a wrapper for the **gp_restore** command.

The prerequisites for restoring data using **gpdrestore** are:

- Backup files are created by a **gpcrondump** operation.
- The Greenplum Database system is up and running.
- The Greenplum Database system has the exact same number of primary segment instances as the system that was backed up.

- The database being restored to is already created in the system.

Note `gpdrestore` is the powerful and relatively new wrapper for the `gp_restore` command and is still under active qualification. This solution uses the simpler `gp_restore` utility. This is adequate for most customer needs, as shown in the following example.

The NFS shares on the Data Domain should be available and mounted on both Master Servers and on each Segment Server for the restore to proceed.

For example, to restore the `tpch` database backup generated by the `gpcrondump` on **20/08/2010**:

1. Ensure that the Data Domain DD890 can be reached from the Master Server:

On the DCA:

```
[gpadmin@mdw ~]$ ping dd_890_1_1
PING dd_890_1_1 (172.28.8.209) 56(84) bytes of data.
64 bytes from dd_890_1_1 (172.28.8.209): icmp_seq=1 ttl=64
time=0.105 ms
64 bytes from dd_890_1_1 (172.28.8.209): icmp_seq=2 ttl=64
time=0.103 ms
[gpadmin@mdw ~]$ ping dd_890_1_2
PING dd_890_1_2 (172.28.12.209) 56(84) bytes of data.
64 bytes from dd_890_1_2 (172.28.12.209): icmp_seq=1 ttl=64
time=0.755 ms
64 bytes from dd_890_1_2 (172.28.12.209): icmp_seq=2 ttl=64
time=0.109 ms
```

2. Ensure that the database being restored to has already been created in the system:

On the DCA:

```
[gpadmin@mdw ~]$ createdb tpch
```

This command fails if the database already exists.

3. Ensure that the database contains no objects as follows:

On the DCA:

```
[gpadmin@mdw ~]$ psql -d tpch
psql (8.2.14)
Type "help" for help.
```

```
tpch=# \d
No relations found.
```

4. Run the `gp_restore` utility:

On the DCA:

```
[gpadmin@mdw ~]$ gp_restore -d tpch --gp-d=/backup/DCA-01/db_dumps/20100820 --gp-k=20100820120142
```

Log files for the Greenplum Database management utilities are written to `~/gpAdminLogs` by default. The naming convention for management log files is:

```
<script_name>_<date>.log
e.g. gpdbrestore_<date>.log
```

Table 10 lists the gp_restore options used in NFS.

Table 10. NFS – gp_restore options used

Parameter	Description
-d	The database name
--gp-d	The backup set path
--gp-k	The timestamp (key) of the backup set to be restored

For more options, see the *Greenplum Database 4.2 Administrator Guide*.

DD Boost – Configuring the DD890 and DCA

For DD Boost backups, the Data Domain system must be configured to enable the DCA to access the Data Domain storage unit that is used as the target for the **gpcrondump** utility to store the backups.

The following example shows how the system was configured using two 10 GbE interfaces with the IP addresses 172.28.8.209 and 172.28.12.209, respectively:

On the DD890:

```
sysadmin@dd_890_1# net show settings
port      enabled  DHCP    IP address  netmask      type  additional setting
-----
eth0a     yes     no      172.28.4.209  255.255.248.0  n/a
eth0b     no      n/a     n/a          n/a           n/a
eth4a     yes     no      172.28.8.209  255.255.252.0  n/a
eth4b     yes     no      172.28.12.209 255.255.252.0  n/a
eth5a     no      n/a     n/a          n/a           n/a
eth5b     no      n/a     n/a          n/a           n/a
-----
* Value from DHCP
```

Also, it is very important to configure the host files on both the DCA and on the Data Domain DD890 system. The following example shows how the Data Domain host entries look on the Master and Segment Servers:

On the DCA:

```
[gpadmin@mdw ~]$ grep "Data Domain" /etc/hosts
### Data Domain - Backup
172.28.8.209    dd_889_1_1    # Data Domain interface eth4a,
Network 172.28.8.0
172.28.12.209 dd_890_1_2    # Data Domain interface eth4b,
Network 172.28.12.0
```

To add the hosts to the Data Domain DD890 system, type:

```
net hosts add <ipaddr> <host-list>
```

The following example adds both 10 GbE interfaces of the host mdw:

On the DD890:

```
sysadmin@dd_890_1# net hosts add 172.28.8.204 mdw1-1
Added "mdw1-1" -> "172.28.8.204" mapping to hosts list.
sysadmin@dd_890_1# net hosts add 172.28.12.204 mdw1-2
Added "mdw1-2" -> "172.28.12.204" mapping to hosts list
```

Repeat these instructions on both 10 GbE interfaces for all Master and Segment Servers.

Perform the following steps to configure DD Boost:

1. Add DD Boost license.

On the DD890:

```
sysadmin@dd_890_1# license show
Feature licenses:
##  License Key          Feature
--  -----
1   XXXX-YYYYY-XXXX-YYYY  DDBOOST
--  -----
```

If necessary, use the following command to add the license:

On the DD890:

```
sysadmin@dd_890_1# license add <license-code>
```

Where <license-code> is the DD Boost license key.

2. Enable NFS on the Data Domain system as follows:

Note: The access list for the DD Boost clients is an NFS export, so NFS must be enabled.

On the DD890:

```
sysadmin@dd_890_1# nfs enable
NFS server is enabled
```

3. Enable DD Boost access to the DCA network.

The network used in this case is 172.28.8.0/22 and 172.28.12.0/22.

On the DD890:

```
sysadmin@dd_890_1# ddboost access add clients 172.28.8.0/22
172.28.8.0/22 : Added
sysadmin@dd_890_1# ddboost access add clients 172.28.12.0/22
172.28.12.0/22 : Added
```

Ensure that the network was added correctly as shown in the following example:

On the DD890:

```
sysadmin@dd_890_1# ddboost access show
DD Boost access allowed from the following clients:
```

```
172.28.12.0/22
172.28.8.0/22
```

4. DD Boost user.

By default, the DD Boost user is `sysadmin`. If you have a requirement to create a new user for DD Boost, refer to the *DD OS 5.1 Administration Guide*. This solution uses the `sysadmin` user.

5. Setup the DD Boost credentials.

Create a credentials file

On the DCA as the `gpadmin` user, create a configuration file named `.ddconfig` in the `/home/gpadmin` directory of the Master Server (`mdw`). The file should have three lines using the parameters shown in Table 11.

Table 11. Parameters for the credentials file

Parameter	Description
Host name	IP address of DD target interface (10 GbE)
User name	Data Domain Boost User Name
Password	Data Domain Boost Login Password

The following three lines show the content of a sample `.ddconfig` file:

```
172.28.8.209
sysadmin
pa55w0rd#
```

Copy the credentials file to the Segment Servers

To get the best balance across both Data Domain 10 GbE interfaces (`eth4a` and `eth4b`), servers `mdw` and `sdw1` to `sdw8` should use the `172.28.8.0` network (interface `eth4a`), while `smdw` and `sdw9` to `sdw16` should use the `172.28.12.0` network (interface `eth4b`).

e.g.

```
To backup via eth4a
172.28.8.209
sysadmin
PA55WORD#
```

```
To backup via eth4b
172.28.12.209
sysadmin
PA55WORD#
```

Use the `gpscp` command to distribute the appropriate DD Boost credentials file from the Master Server to the secondary Master and Segment Servers.

For example:

```
$ gpscp -f hostfile /home/gpadmin/.ddconfig_
=:/home/gpadmin/.ddconfig
```

6. Enable distributed segment processing:

On the DD890:

```
sysadmin@dd_890_1# ddboost option show
Option                               Value
```

```
-----  
distributed-segment-processing  enabled  
-----
```

If disabled, enable as follows:

On the DD890:

```
sysadmin@dd_890_1# ddbboost option set distributed-segment-  
processing enabled  
DD Boost option "distributed-segment-processing" set to  
enabled.
```

Note Ensure that the interface groups on the DD890 are disabled or at least not used for the DCA interfaces.

On the DD890:

```
sysadmin@dd_890_1# ddbboost ifgroup show config  
The 'ifgroup' list is empty.  
sysadmin@dd_890_1# ddbboost ifgroup status  
Status of ifgroup is "disabled".
```

7. Enable DD Boost.

On the DD890:

```
sysadmin@dd_890_1# ddbboost enable  
DD Boost enabled.
```

Alternatively, as with all other DD Boost options, you can also enable DD Boost using the Enterprise Manager GUI with the **Data Management > DD Boost** view option.

The Data Domain system is now ready to receive DD Boost backups run by the **gpcrondump** utility.

DD Boost - Backing up to Data Domain from the DCA

To automate routine DD Boost backups, Greenplum provides the **gpcrondump** utility that can be called directly or from a crontab entry. The **gpcrondump** utility is a wrapper for the **gp_dump** command. **gpcrondump** also enables the backup of objects other than databases and data, such as database roles and server configuration files.

By default, **gpcrondump** creates data dump files in the Master and each Segment instance's data directory in **<data_directory>/db_dumps/YYYYMMDD**. The Segment data dump files are compressed using **gzip**.

For DD Boost backups, **gpcrondump** with the **--ddbboost** parameter is required. Without the **--ddbboost** parameter, the backups are stored locally in the DCA array.

For the DD Boost solution, the deduplication is performed on the Segment Servers before sending the data to the storage unit on the Data Domain system. This reduces network traffic and speeds up the backup.

Note A storage unit called GPDB is automatically created on the Data Domain system during the initial DD Boost backup.

All the DD Boost backups were generated by running the following command as **gpadmin** user:

For **gpcrondump** uncompressed backups with DD Boost:

```
gpcrondump -x tpch --ddboost -z -a
```

Log files for the Greenplum Database management utilities are written to `~/gpAdminLogs` by default. The naming convention for management log files is:

```
<script_name>_<date>.log  
e.g. gpcrondump_<date>.log
```

Table 12 lists the **gpcrondump** options used in DD Boost.

Table 12. DD Boost – gpcrondump options used

Parameter	Description
-x	Database name
--ddboost	Use Data Domain DD Boost for this backup
-z	Do not use compression
-a	Do not prompt

Note The **gpcrondump** compressed options `--rsyncable` are also available, but may not be suitable for environments with a high rate of random data changes. In order to gain the maximum benefit from EMC Data Domain deduplication technology, the **gpcrondump** uncompressed option `-z` is recommended.

It is also possible to automate parallel backups with **gpcrondump** by adding an entry to the crontab of the Master Server:

```
SHELL=/bin/bash  
GPHOME=/usr/local/greenplum-db- 4.2.0.1  
MASTER_DATA_DIRECTORY=/data/gpdb_master/Data-1  
PATH=$PATH:$GPHOME/bin  
01 0 * * * gpadmin gpcrondump -x tpch -c -g -G -a -q >>  
gpcrondump_Backup.log
```

For more information, refer to the *Greenplum Database 4.2 Administrator Guide*.

DD Boost – Restoring data on the DCA from Data Domain

For restoring a database using DD Boost, Greenplum provides the **gpdbrestore** utility, which is a wrapper for the **gp_restore** command.

The prerequisites for restoring data with DD Boost using **gpdbrestore** are:

- Backup files were created by a **gpcrondump –ddboost** operation.
- The Greenplum database system is up and running.
- The Greenplum database system has the exact same number of primary and segment instances as the system that was backed up.
- The database to which the data is being restored is already created in the system.

For example, to restore the **tpch** database backup generated by the **gpcrondump** on **14/12/2011**:

1. Ensure that the DD890 can be reached from the Master Server:

On the DCA:

```
[gpadmin@mdw ~]$ ping dd_890_1_1
PING dd_890_1_1 (172.28.8.209) 56(84) bytes of data.
64 bytes from dd_890_1_1 (172.28.8.209): icmp_seq=1 ttl=64
time=0.105 ms
64 bytes from dd_890_1_1 (172.28.8.209): icmp_seq=2 ttl=64
time=0.103 ms
[gpadmin@mdw ~]$ ping dd_890_1_2
PING dd_890_1_2 (172.28.12.209) 56(84) bytes of data.
64 bytes from dd_890_1_2 (172.28.12.209): icmp_seq=1 ttl=64
time=0.755 ms
64 bytes from dd_890_1_2 (172.28.12.209): icmp_seq=2 ttl=64
time=0.109 ms
```

2. Make sure the database to which the data is being restored has already been created in the system:

On the DCA:

```
[gpadmin@mdw ~]$ createdb tpch
```

Note This command fails if the database already exists.

3. Ensure that the database contains no objects as follows:

On the DCA:

```
[gpadmin@mdw ~]$ psql -d tpch
psql (8.2.14)
Type "help" for help.
tpch=# \d
No relations found.
```

4. Run the **gpdbrestore** utility:

On the DCA:

```
[gpadmin@mdw ~]$ gpdbrestore -t 20111214224212 -ddboost -a
```

Log files for the Greenplum Database management utilities are written to `~/gpAdminLogs` by default. The naming convention for management log files is:

```
<script_name>_<date>.log
e.g. gpdbrestore_<date>.log
```

Table 13 lists the **gpdbrestore** options used in DD Boost.

Table 13. DD Boost – gpdbrestore options used

Parameter	Description
-t	Timestamp key

Parameter	Description
--ddboost	Use Data Domain DD Boost for this restore
-a	Do not prompt

For more options, see the *Greenplum Database 4.2 Administrator Guide*.

Backup schedule

Data Domain's SISL scaling architecture ensures balanced backup and restore speeds.

The backup schedule for this solution uses full backups every day. The advantage of this schedule is that when a recovery is required, just a single restore is required, as incremental or differential restores are not needed. This greatly improves the RTO.

Data Domain systems support both full and incremental backups. However, the **gpcrondump** command can perform a full backup only. Therefore, EMC recommends that you perform a full backup to Data Domain since the recovery and restore process is simple and fast. This is because only a single restore is required, regardless of where in the schedule the data restore is required.

Note With Data Domain systems, it is also possible to do a hot backup while running query and ingest loads. For more information, refer to the *Greenplum Database 4.2 Administrator Guide*.

NFS test results

Test objectives

The test objectives were to validate the success of data backup and restore over NFS by:

- Comparing Data Domain compression rates
- Quantifying the Data Domain deduplication ratio using increased data loads
- Quantifying the impact of running read queries and write queries on the database while running backups
- Quantifying the successful restoration of data with the time to complete the restores

Test scenarios

Table 14 outlines the data backup and restore test scenarios in this solution.

Table 14. NFS backup and restore test scenarios

Test	Description
1	Perform full backup to Data Domain via NFS with Greenplum backup compression off, and with Data Domain using lz , gzfast , and gz local compression options, to determine the best compression option to use.
2	Perform full backup to Data Domain via NFS with Greenplum backup compression off and Data Domain lz compression over a simulated week of full backups.
3	Perform full backup to Data Domain via NFS with Greenplum backup compression off and Data Domain lz compression with: <ul style="list-style-type: none">• Query load running• Ingest load running
4	Perform a restore of the Greenplum database via NFS from a Greenplum uncompressed backup.

Test 1 results

In Test 1, a full backup to Data Domain via NFS with Greenplum compression off and with Data Domain using **lz**, **gzfast** and **gz** was performed.

Figure 2 shows the comparison of lz, gzfast, and gz types of NFS test results.




Figure 2. NFS - Comparison of lz, gzfast and gz

Note that the Data Domain ingest rate refers to the speed at which the DD890 was able to accept incoming backup data.

For the NFS solution, the choice was to use a lighter compression type. lz was selected as it gives the best performance. The test was performed on a 2.59 TB database, where an incremental data load of 5 percent was applied since the previous backup.

Test 2 results

In Test 2, a full backup to Data Domain with Greenplum backup compression off and Data Domain lz compression was performed.

Figure 3 illustrates the 5 percent daily increase in the source database size over seven days on the DCA with an uncompressed backup of an uncompressed database. To simulate the real-life expectancies and data growth of a data warehouse, an incremental data load of 5 percent was applied between backups.

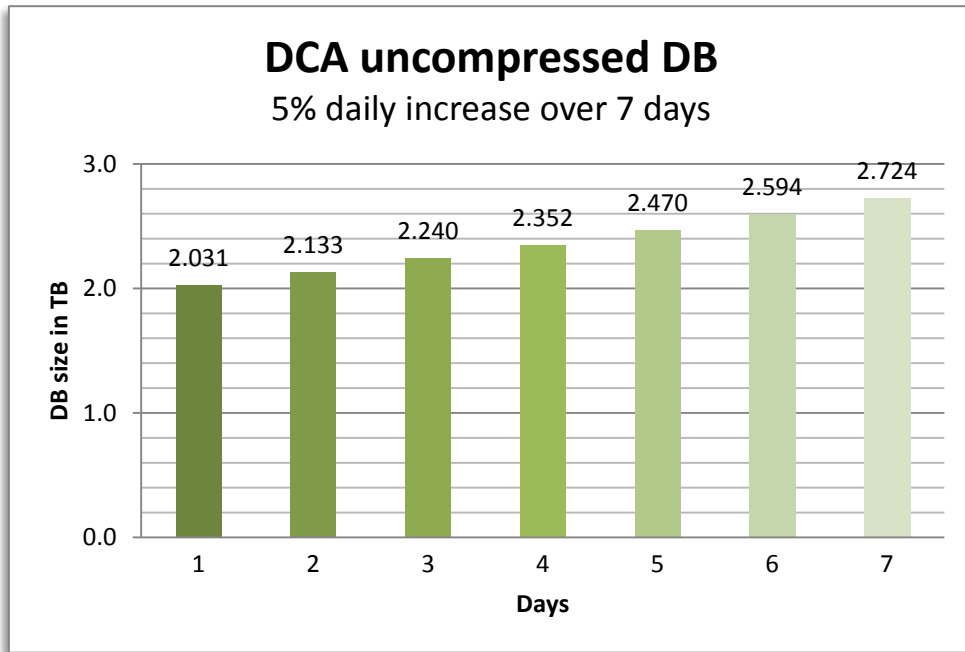


Figure 3. NFS - Five percent daily increase over a simulated seven days with an uncompressed database

Figure 4 illustrates the backup duration in minutes on the DCA with an uncompressed backup of an uncompressed database, based on the 5 percent daily increase. The first backup takes more time to complete because it is the first time the data is being written to the Data Domain system.

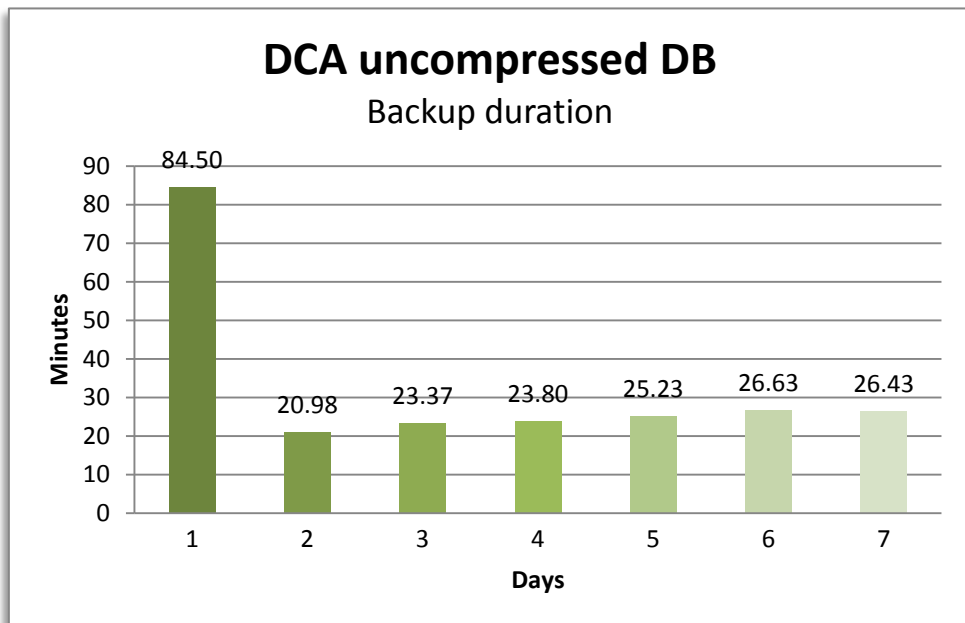


Figure 4. NFS - Backup duration with gpcrondump uncompressed backup

Figure 5 illustrates the effective backup rate in TB/hour with an uncompressed backup of an uncompressed database. As seen previously in the backup duration test, the throughput is slower on the first backup. Using **gpcrondump** uncompressed backups with Data Domain's deduplication technology results in an average backup speed of 5.6 TB/hour for backups 2 to 7.




Figure 5. NFS -Data Domain ingest rate with gpcrondump uncompressed backup

Due to the strength of Data Domain deduplication technology, there was an average 20.9x storage saving for each nightly backup. Shortly after the nightly database backup began, the savings were significant. For example, on the second day, a full backup of 1971.5 GiB used 99.9 GiB, resulting in a 19.7x storage reduction.

Figure 6 illustrates the tremendous incremental savings that can be achieved on a daily basis. Over time, the savings are even greater.

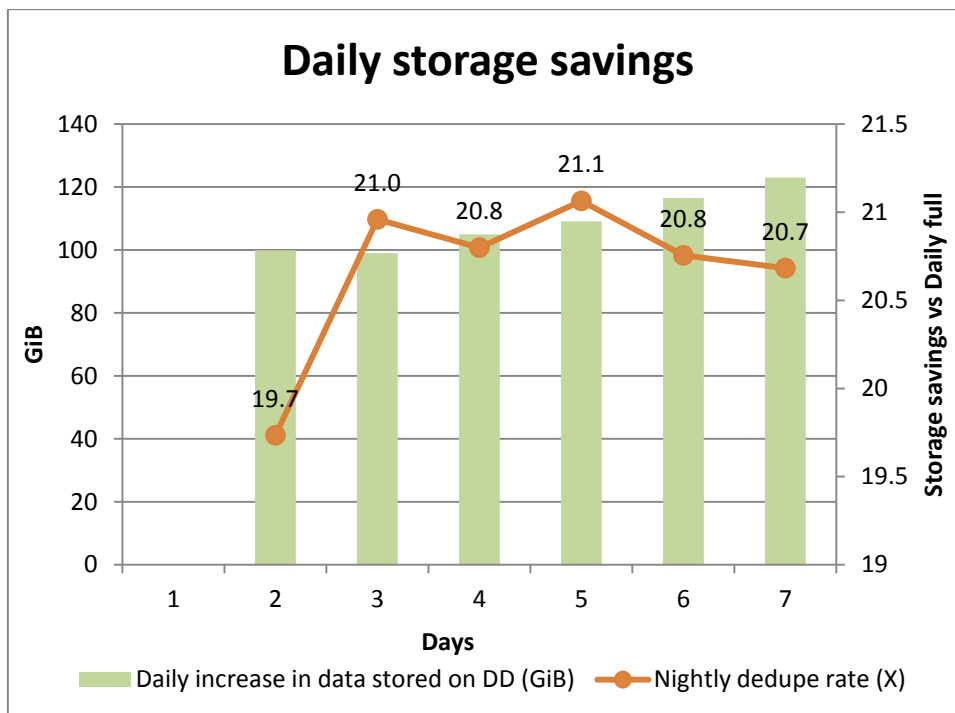


Figure 6. NFS - Storage saving after seven days

Note During all the backup tests, the DCA was idle.

Figure 7 illustrates the cumulative effect of this storage saving over a seven-day backup cycle. After seven days of running the backup, 15363.2 GiB of data was backed up; however, only 1620.5 GiB of storage was needed on the Data Domain deduplication storage system, resulting in a 9.5x storage savings. This 89.5 percent storage saving represents a significant saving in backup infrastructure and facility costs. Regular backups of large data warehouses are far more sustainable and much less costly than before.

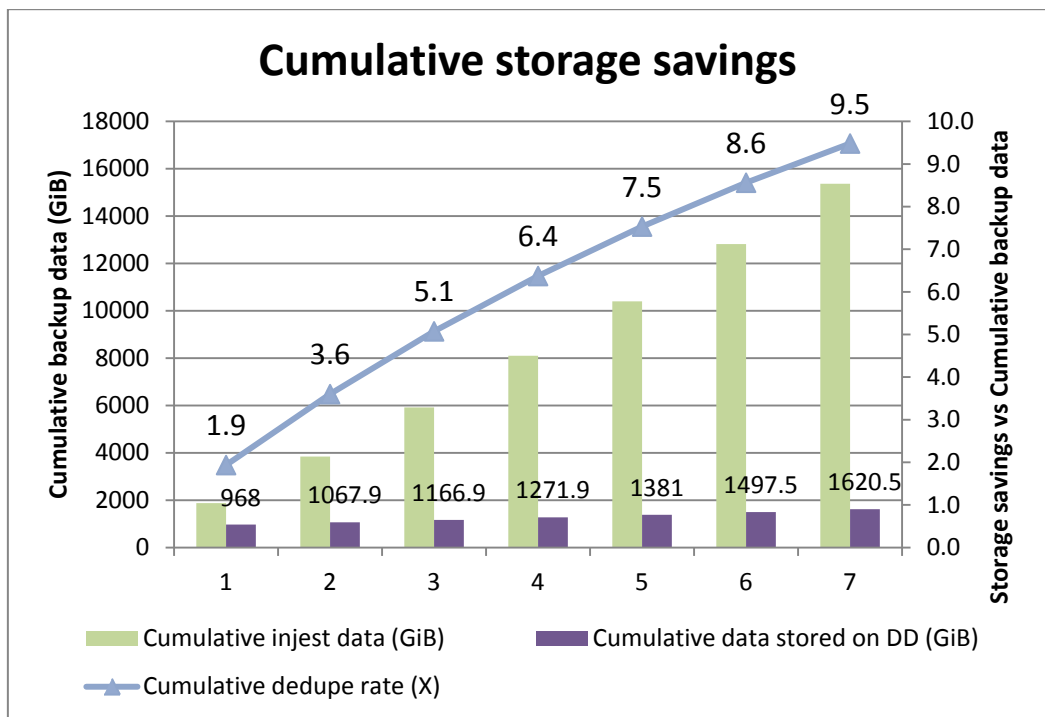


Figure 7. NFS - Cumulative storage savings over seven days

Note There is roughly a 5.5 percent difference between the daily database sizes reported using Greenplum utilities in Figure 3 versus the actual data saved to the Data Domain system as noted in Figure 7. The reason for this difference is that **gpcrondump** extracts the actual data from the Greenplum database for backing up to the Data Domain system. The 5.5 percent is additional Greenplum overhead that is not required for a successful backup or restore.

Test 3 results

In Test 3, a full backup to the Data Domain system was performed with a query load and a ingest load running against the DCA. No incremental data was applied since the previous backup. The query and ingest loads were not performed at the same time. The intention was to demonstrate the impact while a backup job was running.

Figure 8 illustrates the DCA database size before and after the ingest load. The amount of data backed up was 2.03 TB and the database went from 2.03 TB to 2.29 TB while the backup was performed.




Figure 8. NFS - Database size before and after the ingest load

Figure 9 compares the DCA backup duration in minutes with no load, query load, and ingest load running. As illustrated, under a concurrent query load, the NFS backup performance is somewhat degraded. NFS backup performance under a full ingest load, which is a write-intensive process, is less affected, with no operational issues.

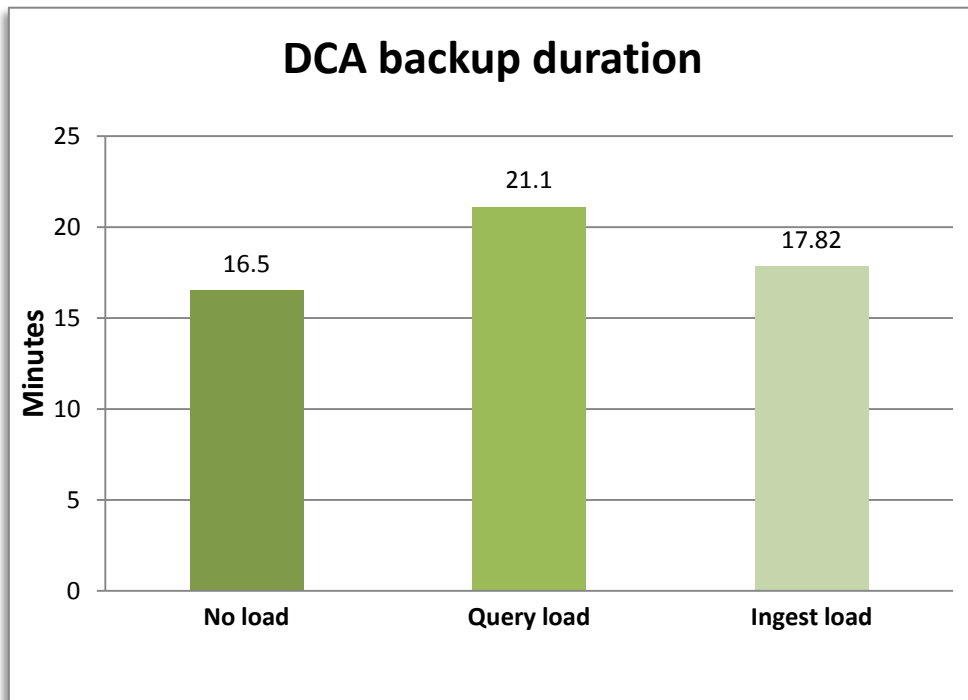


Figure 9. NFS - Backup duration under varying conditions

Figure 10 compares the Data Domain average ingest rates with no load, query load, and ingest load running.

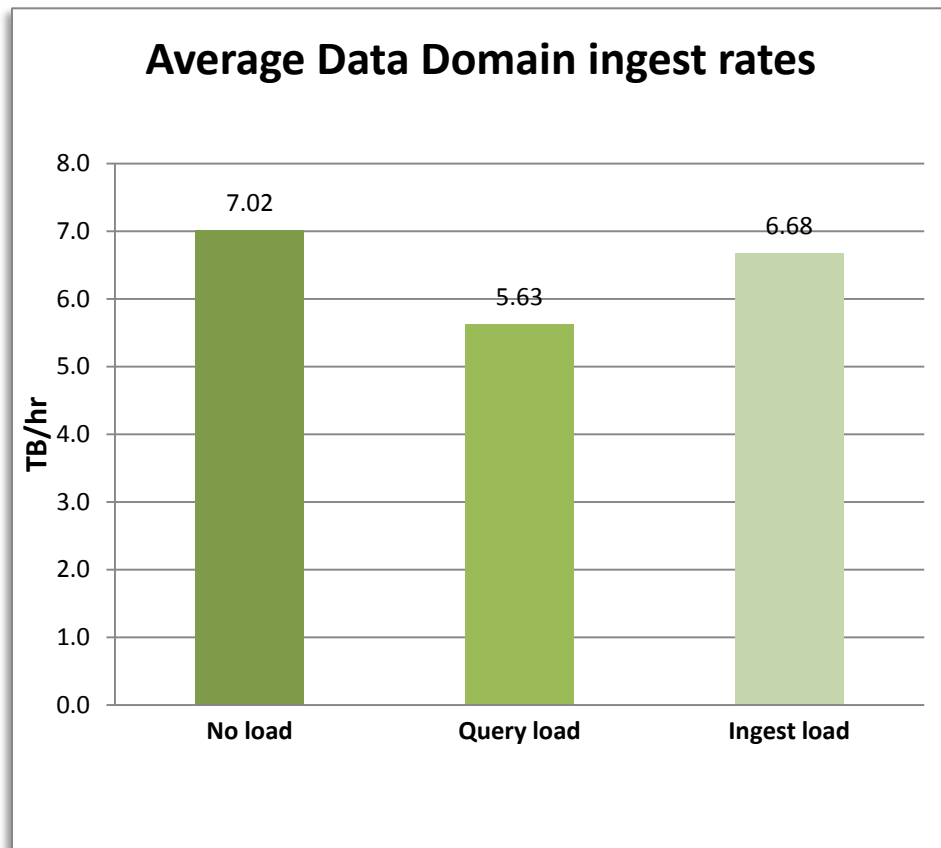


Figure 10. NFS – Average Data Domain ingest rate under varying conditions

Test 4 results

In Test 4, a restore of the Greenplum database was performed using the Greenplum uncompressed backup set from previous tests.

The Test 4 DCA restore with uncompressed data produced the following results:

- The restore of a 2.72 TB database from an uncompressed backup was achieved in 32.67 minutes.
- The average restore throughput of a Greenplum uncompressed backup was 5.00 TB/hour.

DD Boost test results

Test objectives

The test objectives were to validate the success of a DD Boost data backup and restore by:

- Comparing Data Domain compression rates
- Quantifying the Data Domain deduplication ratio using increased data loads
- Quantifying the impact of running read queries and write queries on the database while running backups
- Quantifying the successful restoration of data with the time to complete the restores

Test scenarios

Table 15 outlines the data backup and restore test scenarios in this solution.

Table 15. DD Boost backup and restore test scenarios

Test	Description
1	Perform a full backup to Data Domain using DD Boost with Greenplum backup compression off and with Data Domain using lz , gzfast , and gz local compression options to determine the best compression option.
2	Perform a full backup to Data Domain using DD Boost with Greenplum backup compression off and Data Domain lz compression over a simulated week of full backups.
3	Perform a full backup to Data Domain using DD Boost with Greenplum backup compression off and Data Domain lz compression with: <ul style="list-style-type: none">• Query load running• Ingest load running
4	Restore the Greenplum database using DD Boost from a Greenplum uncompressed backup.

Test 1 results

Test 1 comprised of a full backup to Data Domain using DD Boost, with Greenplum compression off and with Data Domain using **lz**, **gzfast**, and **gz**.

Figure 11 shows the **lz**, **gzfast**, and **gz** DD Boost test results.

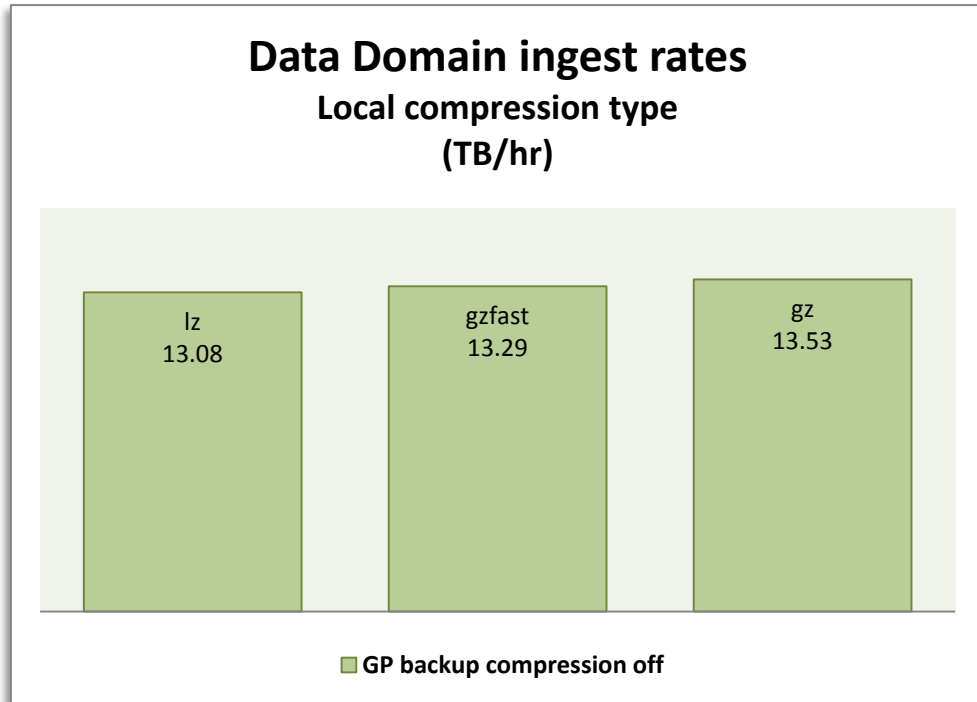


Figure 11. DD Boost - Comparison of lz, gzfast and gz

The Data Domain ingest rate refers to the speed at which the DD890 was able to accept incoming backup data.

For the DD Boost solution, **gz** provides the best backup and restore performance as well as increased storage savings. Because this is a global setting, the choice was to use the default and lighter compression option **lz**. In most cases, the DD890 is also used for other backup jobs and using **gz** can have a significant effect on their performance. The test was performed on a 2.59 TB database, where an incremental data load of five percent was applied since the previous backup.

Test 2 results

Test 2 comprised of a full backup to Data Domain, with Greenplum compression off using Data Domain **lz** compression.

Figure 12 illustrates the five percent daily increase in the source database size over seven days on the DCA with an uncompressed backup of an uncompressed database. To simulate the real-life expectancies and data growth of a data warehouse, an incremental data load of five percent between backups was applied.




Figure 12. DD Boost – Five percent daily increase over a simulated seven days with an uncompressed database

Figure 13 illustrates the backup duration in minutes on the DCA with an uncompressed backup of an uncompressed database, based on the 5 percent daily increase. The first backup takes more time to complete because it is the first time the data is being written to the Data Domain system.

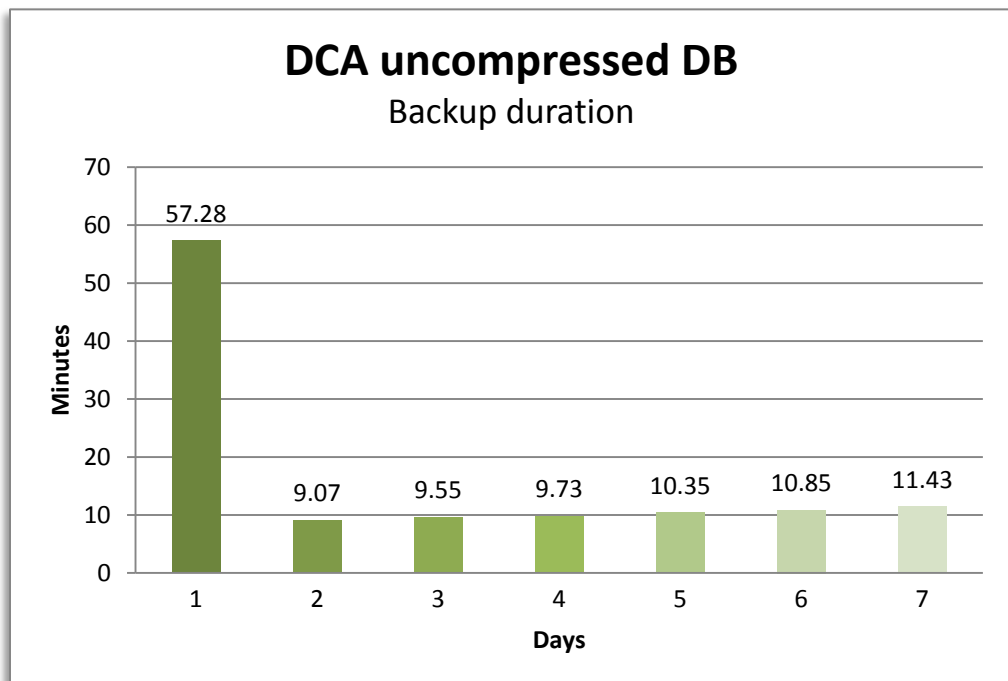


Figure 13. DD Boost - Backup duration with gpccrondump uncompressed backup

Figure 14 illustrates the effective backup rate in TB/hour with an uncompressed backup of an uncompressed database. As seen previously in the backup duration test, the throughput is slower on the first backup. The combination of **gpcrondump** uncompressed backups and Data Domain's deduplication technology results in an average backup speed of 13.08 TB/hour for backups 2 to 7.




Figure 14. DD Boost – Data Domain ingest rate with gpcrondump uncompressed backup

Due to the strength of Data Domain's deduplication technology, there was an average 34.5x storage saving for each nightly backup. Shortly after the nightly database backup began, the savings were significant. For example, on the second day, a full backup of the 1971.5 GiB used 59 GiB, resulting in a 33.4x storage reduction.

Figure 15 illustrates the tremendous incremental savings that can be achieved on a daily basis. Over time, the savings are even greater.




Figure 15. DD Boost - Storage saving after seven days

Note During all the backup tests, the DCA was idle.

Figure 16 illustrates the cumulative effect of this storage saving over a seven-day backup cycle. After seven days of running the backup, 15363.2 GiB of data was backed up; however, only 1354.9 GiB of storage was needed on the Data Domain deduplication storage system, resulting in a 11.3x storage saving. This 91.2 percent storage saving represents a significant saving in backup infrastructure and facility costs. Regular backups of large data warehouses are far more sustainable and much less costly than before.




Figure 16. DD Boost - Cumulative storage savings over seven days

There is roughly a 5.5 percent difference between the daily database sizes reported using Greenplum utilities in Figure 12 versus the actual data saved to the Data Domain system as noted in Figure 16. This is because **gpcrondump** extracts the actual data from the Greenplum database for backing up data to the Data Domain system. The 5.5 percent is additional Greenplum overhead that is not required for a successful backup or restore.

Test 3 results

Test 3 comprised of a full backup to the Data Domain system with a query load and an ingest load running against the DCA. No incremental data was applied since the previous backup. The query and ingest loads were not performed at the same time. The intention was to demonstrate the impact while a backup job was running.

Figure 17 illustrates the DCA database size before and after the ingest load. The amount of data backed up was 2.03 TB and the database went from 2.03 TB to 2.42 TB while the backup was performed.

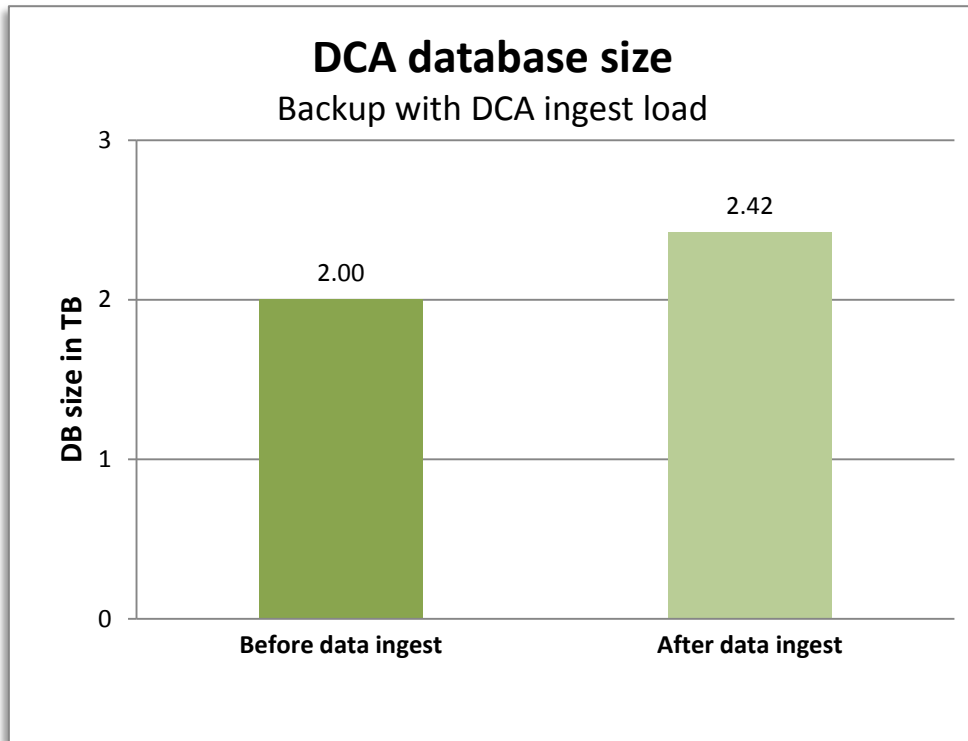


Figure 17. DD Boost - Database size before and after the ingest load

Figure 18 compares the DCA backup duration in minutes with no load, query load, and ingest load running. As illustrated, under a concurrent query load, DD Boost backup performance is degraded. DD Boost backup performance under a full ingest load, which is a write-intensive process, is less affected, with no operational issues.

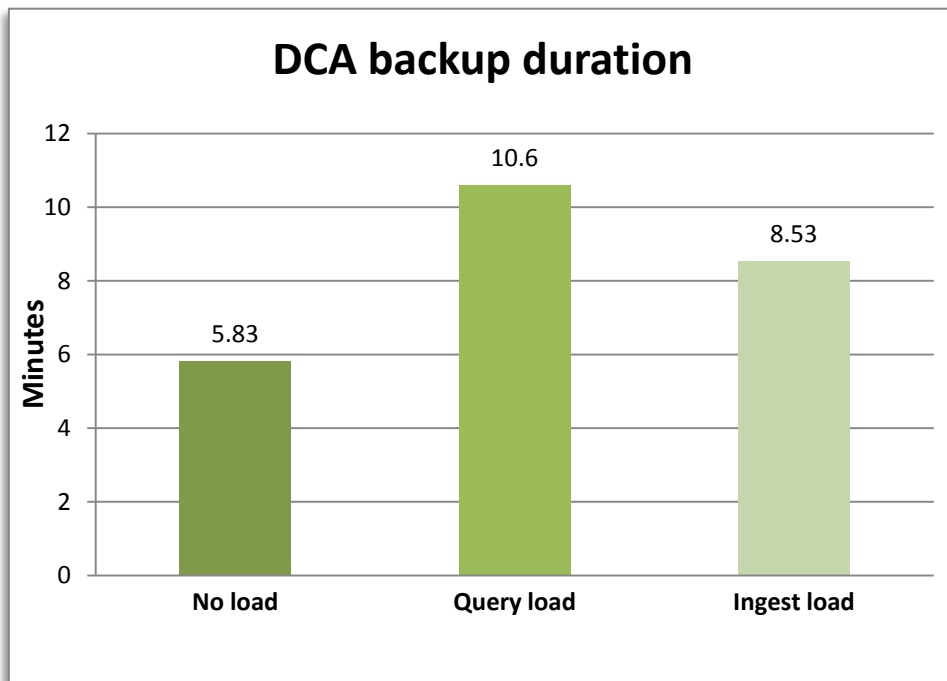


Figure 18. DD Boost - Backup duration under varying conditions

Figure 19 compares the Data Domain average ingest rates with no load, query load, and ingest load running.

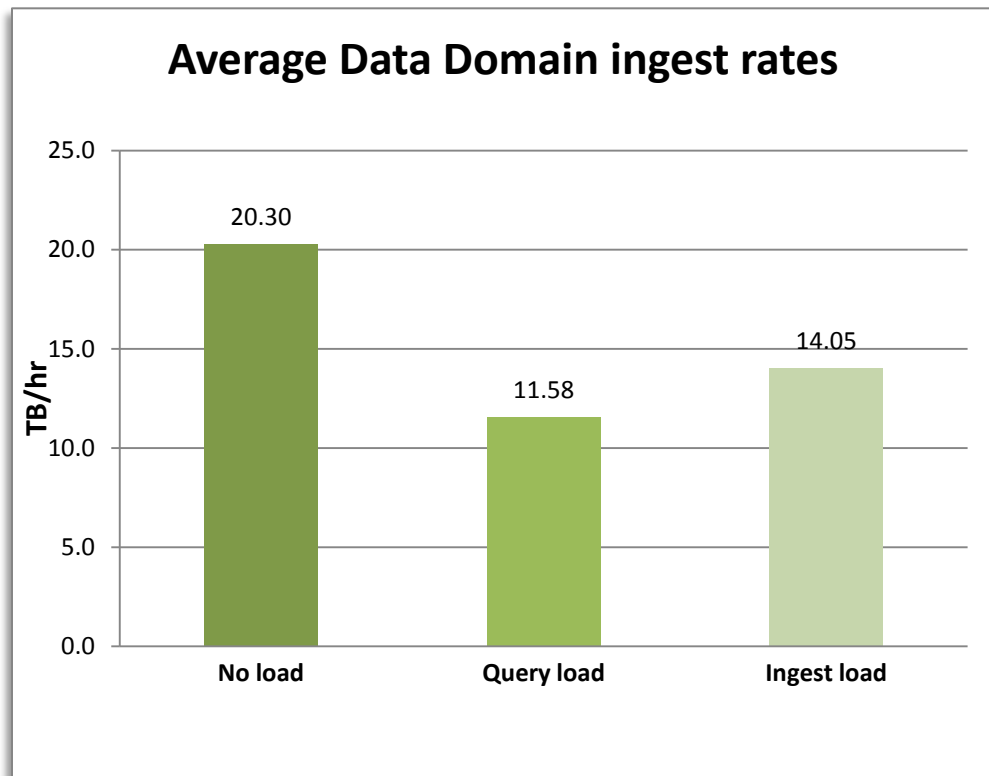


Figure 19. DD Boost - Average Data Domain ingest rate under varying conditions

Test 4 results

In Test 4, a restore of the Greenplum database was performed using the Greenplum uncompressed backup set from previous tests.

The Test 4 DCA restore with uncompressed data produced the following results:

- The restore of a 2.72 TB database from an uncompressed backup was achieved in 27.72 minutes.
- The average restore throughput of a Greenplum uncompressed backup was 5.9 TB/hour.

Conclusion

Summary

EMC provides leading-edge technology to support the backup and recovery of the Greenplum DCA. This white paper demonstrates how a Data Domain deduplication storage system achieves this in the most simple, efficient, and cost-effective manner.

With this backup and recovery solution for the DCA, organizations can achieve:

- Faster backup and restores with minimized backup windows and maximized end-to-end recovery performance
- Efficient use of both infrastructure and people to support the business, leading to reduced operational and infrastructure costs
- Operational simplicity through ease of integration and management

Findings

The EMC Data Domain DD890 deduplication storage system provides cost-effective, long-term, onsite retention and protection of business-critical data and enables WAN-efficient networked disaster recovery.

With this solution, companies can expect to achieve:

- Operational ease and efficiency through the utilization of MPP architecture to efficiently back up, in parallel, across the network to the Data Domain system.
- Major space saving advantages using Data Domain inline deduplication—in this case, test results on the DD890 with DD Boost demonstrate a saving of 34.5x on nightly backup, and 11.3x cumulative saving over a week of uncompressed backups to the Data Domain system. One of the key benefits of Data Domain deduplication storage systems is to reduce the need for excessive amounts of backup storage.
- Fast restore times for returning uncompressed and deduplicated data to the DCA. The test results with DD Boost demonstrate that the combination of **gpcrondump** and EMC Data Domain deduplication technology results in an effective backup of 13.08 TB/hour and a restore of 5.90 TB/hour. This enables the backup and recovery of a Greenplum full-rack DCA (36 TB uncompressed database) in under 2.75 and 6.1 hours respectively.

Other observations from the test cycle:

- The **gpcrondump** compressed backup option `--rsyncable`, especially for NFS backups, can provide greater backup and restore throughput. However, this is only suitable for static environments. In order to gain the maximum benefit of EMC Data Domain deduplication technology, the **gpcrondump** uncompressed option `-z` is recommended.
- Because DD Boost is more CPU intensive on the DCA Segment Servers, you can see a significant drop in backup throughput when running a query load or an ingest load in parallel.
- For DD Boost backups, if the Data Domain system is dedicated to DCA backups, you should consider Data Domain **gz** compression.

- The **gpcrondump** uncompressed DD Boost backups have a significant performance improvement over **gpcrondump** uncompressed NFS backups.
- The rate of change directly affects the average Data Domain ingest rate. For this solution, the NFS backup performance can vary from 1.37 TB/hour to 7.02 TB/hour. For DD Boost, the backup performance can vary from 1.83 TB/hour to 20.30 TB/hour.

As the key findings show, the Data Domain DD890 has been used in this solution to provide an effective disk-based backup target that significantly minimizes storage usage while providing long-term retention. Retention of many backups can be managed more easily and cost effectively when using the Data Domain deduplication storage system.

Note This white paper provides one example of the results from using a Data Domain deduplication storage system. However, customer environments differ considerably in terms of backup policies, backup windows, and rate of data change, and data retention and deduplication ratios can vary. Therefore each environment must be reviewed individually to size the Data Domain system to address each user's particular requirements. Interested customers should contact their EMC BRS representative for more information.

References

White papers

For additional information, see the white paper listed below.

- *EMC Greenplum Data Computing Appliance: Architecture, Performance and Functions - A Detailed Review*

Product documentation

For additional information, see the product documents listed below.

- *DD OS 5.1 Administration Guide*
- *DD OS 5.1 Initial Configuration Guide*
- *Greenplum Database 4.2 Administrator Guide*
- *Data Domain Installation and Setup Guide*

Supporting information

Interconnect 1 – Converting ports 18 and 19 from a LAG to switch ports

After logging into the switch, run the following commands:

```
i-sw-1:admin> cmlsh  
i-sw-1#show run
```

Ports 18 and 19 will look like this:

```
interface TenGigabitEthernet 0/18  
  channel-group 2 mode active type brocade  
  no shutdown  
  lacp timeout short  
!  
interface TenGigabitEthernet 0/19  
  channel-group 2 mode active type brocade  
  no shutdown  
  lacp timeout short
```

To setup ports 18 and 19 as switch ports, run the following commands:

```
i-sw-1#config t  
Enter configuration commands, one per line. End with CNTL/Z.  
i-sw-1(config)#no interface Port-channel 2  
i-sw-1(config)#interface TenGigabitEthernet 0/18  
i-sw-1(config-if-te-0/18)#switchport  
i-sw-1(config-if-te-0/18)#switchport mode access  
i-sw-1(config-if-te-0/18)#vlan classifier activate group 2 vlan 199  
i-sw-1(config-if-te-0/18)#no shutdown  
i-sw-1(config-if-te-0/18)#spanning-tree edgeport  
i-sw-1(config-if-te-0/18)#spanning-tree edgeport bpdu-guard  
i-sw-1(config-if-te-0/18)#interface TenGigabitEthernet 0/19  
i-sw-1(config-if-te-0/19)#switchport  
i-sw-1(config-if-te-0/19)#switchport mode access  
i-sw-1(config-if-te-0/19)#vlan classifier activate group 2 vlan 199  
i-sw-1(config-if-te-0/19)#no shutdown  
i-sw-1(config-if-te-0/19)#spanning-tree edgeport  
i-sw-1(config-if-te-0/19)#spanning-tree edgeport bpdu-guard  
i-sw-1(config-if-te-0/19)#exit  
i-sw-1(config)#exit  
i-sw-1#write mem  
Overwrite the startup config file (y/n): y  
Building configuration...  
i-sw-1#show run
```

Ports 18 and 19 will now look like this:

```
interface TenGigabitEthernet 0/18  
  switchport  
  switchport mode access  
  vlan classifier activate group 2 vlan 199  
  no shutdown  
  spanning-tree edgeport  
  spanning-tree edgeport bpdu-guard  
!  
interface TenGigabitEthernet 0/19  
  switchport  
  switchport mode access
```

```
vlan classifier activate group 2 vlan 199
no shutdown
spanning-tree edgeport
spanning-tree edgeport bpdu-guard
```

Plug the sfp into port 19; then connect the fibre cable from port 19 to Data Domain eth4a, for example.

Interconnect 2 – Converting ports 18 and 19 from a LAG to switch ports

After logging into the switch, run the following commands:

```
i-sw-2:admin> cmlsh
i-sw-2#show run
```

Ports 18 and 19 will look like this:

```
interface TenGigabitEthernet 0/18
  channel-group 2 mode active type brocade
  no shutdown
  lacp timeout short
!
interface TenGigabitEthernet 0/19
  channel-group 2 mode active type brocade
  no shutdown
  lacp timeout short
```

To setup ports 18 and 19 as switch ports, run the following commands:

```
i-sw-2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
i-sw-2(config)#no interface Port-channel 2
i-sw-2(config)#interface TenGigabitEthernet 0/18
i-sw-2(conf-if-te-0/18)#switchport
i-sw-2(conf-if-te-0/18)#switchport mode access
i-sw-2(conf-if-te-0/18)#vlan classifier activate group 2 vlan 299
i-sw-2(conf-if-te-0/18)#no shutdown
i-sw-2(conf-if-te-0/18)#spanning-tree edgeport
i-sw-2(conf-if-te-0/18)#spanning-tree edgeport bpdu-guard
i-sw-2(conf-if-te-0/18)#interface TenGigabitEthernet 0/19
i-sw-2(conf-if-te-0/19)#switchport
i-sw-2(conf-if-te-0/19)#switchport mode access
i-sw-2(conf-if-te-0/19)#vlan classifier activate group 2 vlan 299
i-sw-2(conf-if-te-0/19)#no shutdown
i-sw-2(conf-if-te-0/19)#spanning-tree edgeport
i-sw-2(conf-if-te-0/19)#spanning-tree edgeport bpdu-guard
i-sw-2(conf-if-te-0/19)#exit
i-sw-2(config)#exit
i-sw-2#write mem
Overwrite the startup config file (y/n): y
Building configuration...
i-sw-2#show run
```

Ports 18 and 19 will now look like this:

```
interface TenGigabitEthernet 0/18
  switchport
  switchport mode access
  vlan classifier activate group 2 vlan 299
  no shutdown
```

```
spanning-tree edgeport
spanning-tree edgeport bpdu-guard
!
interface TenGigabitEthernet 0/19
switchport
switchport mode access
vlan classifier activate group 2 vlan 299
no shutdown
spanning-tree edgeport
spanning-tree edgeport bpdu-guard
```

Plug the sfp into port 19; then connect the fibre cable from port 19 to Data Domain eth4b, for example.