

Greenplum Database Enterprise Data Clouds on EMC CLARiiON CX4-960 Building Blocks

Best Practices Planning

Abstract

Enterprise Data Clouds are a compelling new way to approach large-scale data warehousing and analytics. Greenplum Database software, in concert with the EMC[®] CLARiiON[®] CX4-960 storage system, delivers on the reliability, flexibility, and scalability promises of private cloud computing, bringing the power of self-service and unified data access to data warehousing and analytics. In this white paper, we provide a series of recommendations to help you deploy this combination to your best advantage.

March 2010

Copyright © 2010 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 other trademarks used herein are the property of their respective owners.

Part Number h7015.1

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Terminology	5
Technology overview	6
Greenplum Database	6
Greenplum Enterprise Data Cloud	7
EMC CLARiiON	8
Configuration building blocks	9
Storage building block	9
Scaling out the storage building block	10
Compute building block	11
Flexibility	11
Configuration guidelines	11
Balanced I/O configuration	11
Balancing I/O activities between the two storage processors	11
Balancing data distributions among LUNs	12
Balancing traffic across the back-end buses	12
Balancing I/O traffic among the UltraFlex I/O modules	12
Disk and RAID configuration	12
Cache memory configuration	13
Write cache configuration	13
Read cache configuration	13
Transaction log placement	13
Deploying Flash drives	13
Greenplum/CLARiiON building block and EDC	14
Conclusion	16
References	16

Executive summary

A series of concurrent business and technology trends is serving as the catalyst for a new way to manage information in large enterprises. First, each passing year sees dramatic growth in the overall amount of data that must be captured, organized, and managed. Second, while large organizations used to enjoy distinct technological and financial advantages because of their size, decreasing hardware and software costs have leveled the playing field, translating into significant new competitive pressures from all directions. Finally, the IT backlog is rising at a time of increased economic stress, reducing IT's ability to respond to the needs of the business.

To help respond to the aforementioned challenges, many organizations are turning to Enterprise Data Clouds (EDC). These are solutions made up of hardware and software components designed to leverage the power of massively parallel processing (MPP). They utilize from tens to hundreds of nodes, all coordinated to process in parallel. They scale from tens of terabytes to tens of petabytes, and their data is automatically partitioned across nodes. This scalability is made possible by employing a “shared nothing” architecture. Once the EDC is in place, business analysts are able to use self-service to quickly set up their own data warehouses with minimal support from IT. This rapid return on investment (ROI) and independence make managing change much easier, while helping to reduce friction between the two groups.

The set of new features in the EMC® CLARiiON® CX4 product family is designed to facilitate participation in an EDC. The top-end member of the family, the CX4-960 delivers unparalleled levels of performance and scalability, easily supporting multi-threaded large random read requests on a sustained basis from many complex concurrent user queries. The system supports up to 960 drives and accommodates a multitude of drive types ranging from extremely high-performance drives to extremely dense drives, offering not only good scaling in total capacity, but also storage tiering to allow for the most cost-effective way of storing and managing EDC content.

Introduction

This white paper provides specific recommendations for configuring and managing a CX4-960 system with related Greenplum software to be an effective participant in an EDC.

The guidelines and best practices presented will let the reader configure their Greenplum/EMC node in the most cost-efficient way, with an optimal price/performance balance. We start with a basic building block of servers and EMC storage that can then be replicated (almost infinitely) to scale out to your compute and storage needs. Using this fundamental building block, you can scale out your Greenplum/EMC-based EDC from tens of terabytes to tens of petabytes.

Building blocks provide a simple way to configure a balanced I/O subsystem that will perform well with Greenplum. This approach eliminates the need to perform extensive planning of where to place partitions, how many servers to purchase, or how to lay out data. The building block can fit in a single 42U rack and provide 3 GB/s of raw I/O and as much as 30 GB/s of effective I/O (using compression).

As we'll see throughout the paper, Greenplum and EMC make a particularly attractive combination, sporting a number of compelling synergies, including:

- Deep compression and partitioning using EMC's storage migration service
- Ability to scale compute servers separately from storage servers
- No need for mirroring
- Improved high availability model using LUN takeover
- Enhanced archiving using Greenplum's Gpsuspend backup utility
- Effective virtualization with server motion and load balancing

Configuring a Greenplum/CLARiiON node as part of an EDC is quite straightforward. We'll describe the following procedures and best practices:

-
1. Balancing I/O activities between the two storage processors.
 2. Creating a balanced data distribution among LUNs.
 3. Balancing traffic across the back-end buses.
 4. Distributing I/O traffic among the UltraFlex™ I/O modules.
 5. Configuring disk and RAID.
 6. Setting efficient read/write cache values.
 7. Optimizing placement of the transaction log.
 8. Deploying Flash drives.
 9. Configuring the Greenplum Database and EDC.

Audience

This white paper is intended for Greenplum practitioners and/or IT staff responsible for planning and managing the storage infrastructure for their enterprise data warehouse and analytic deployments. A basic understanding of CLARiiON storage and Greenplum Database technologies is assumed.

Terminology

Analytics: The study of operational data using statistical analysis with a goal of identifying and leveraging patterns to optimize business performance.

Business intelligence (BI): The effective use of information assets to improve the profitability, productivity, or efficiency of a business. Frequently, 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 Enterprise Data Cloud.

Cloud computing: A form of distributed computing where computational, data storage, and other assets are accessible locally yet are hosted elsewhere.

Compute nodes: Computers that are dedicated to performing complex computations; the underlying data is generally stored and managed on separate storage servers.

Data warehouse (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.

Decision support system (DSS): A set of business applications and processes that provide answers in response to different queries pertaining to the business, based on the business's information assets, to help direct or facilitate key business decisions.

Disk array enclosure (DAE): The physical enclosure with disk drive slots to support up to 15 drives to be accessed from the CLARiiON storage processors using the UltraPoint™ connection CLARiiON technology. DAEs support the ability to grow the total number of disk drives used in a CLARiiON system in a modular fashion.

Enterprise Data Cloud (EDC): A hardware and software solution designed to enable self-service provisioning of data warehouses from tens of terabytes to tens of petabytes, on tens to hundreds of nodes working together in parallel.

Flash drive (FD): A solid-state data storage device with no moving parts. These types of drives deliver significantly faster performance than traditional disk drives.

I/O cards: The flexible I/O modules that can be added to CLARiiON CX4 systems to expand the connection ports for increased front-side connections from the servers on the storage area network (SAN),

or back-end ports to provide for more I/O paths for the storage system

Logical unit number (LUN): A storage system object that can be made visible and usable as a server operating system “disk device” from the underlying storage system.

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.

Private cloud: A conglomeration of compute and storage servers that are generally dedicated to one particular organization. These computers may be hosted behind the enterprise’s firewall, or may be distributed across the Internet.

Redundant Array of Inexpensive 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 would be 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.

RAID 5 (R5): A RAID option where the actual data distributed and stored inside a set of drives is effectively protected by an additional set of parity data of the distributed content across the drives computed and stored in an additional drive. Under EMC CLARiiON implementation, the extra parity data is systematically rotated among all the drives in that RAID set to avoid any particular write hot spots when parity data adjustment has to be made against any piece of data stored inside a LUN or LUNs created from this RAID group. RAID 5 protects against loss of data, or data inaccessibility, in the event that one of the drives in the RAID set should experience a drive failure.

RAID 10 (R10): A RAID option that combines the performance-enhancing features of RAID 0 with the data integrity capabilities of RAID 1. Data is striped over mirrored drive pairs.

Scale out: A technique that increases total processing power by adding additional independent computational nodes, as opposed to augmenting a single, large computer with incremental disk, processor, or memory resources.

Self-service provisioning: A fundamental philosophy of the Enterprise Data Cloud, where business analysts are provided with the tools and technology to let them quickly construct their own data warehouses with minimal support from IT staff.

Shared nothing: A distributed computing architecture made up of a collection of independent, self-sufficient nodes. This is in contrast with a traditional central computer that hosts all information and processing in a single location.

Technology overview

Before itemizing the steps necessary to deploy Greenplum on CLARiiON storage as part of an EDC, let’s examine each of the components in more detail.

Greenplum Database

The Greenplum Database is at the heart of each node in an EDC. It’s designed for business intelligence and analytical processing, utilizing a shared nothing, massively parallel architecture to support tremendous scalability, multi-level fault tolerance, and redundancy. Since Greenplum’s philosophy is to maximize uptime while minimizing the IT burden, the database is designed for online system expansion. Typical installations range from tens of terabytes to petabytes.

Greenplum lets clusters of servers act as a database supercomputer. Although not required, you have the freedom and flexibility to partition your information using several different criteria, including:

-
- Date
 - Range
 - Value

All queries are executed using parallel processing. In spite of this power, developers and users are free to use familiar SQL Server statements, in any of the following SQL standards:

- SQL-92
- SQL-99
- SQL-2003 OLAP extensions

In addition to working with native SQL, developers are also free to employ MapReduce for high-scale data analysis. Finally, Greenplum supports access via a broad collection of industry-standard interfaces such as:

- SQL
- ODBC
- JDBC
- DBI

Taken together, these capabilities make a node running the combination of Greenplum and CLARiiON an ideal participant in an EDC.

Greenplum Enterprise Data Cloud

Enterprise Data Clouds represent an innovative new way to manage the information challenges of the 21st century. The sheer amount of information that must be stored, managed, and queried continues to grow at an accelerating rate. To make matters worse, this data is most commonly stored in multiple silos, using multiple formats.

Administering this information collection is proving to be too large a task for most backlogged IT organizations. There's little time to configure data warehouses, yet business analysts need access to this data as quickly as possible.

Existing technologies and architectural approaches have proven to be unable to address these needs. Some reasons include the following:

- OLTP-style databases simply won't scale to support modern data warehousing and analytic applications.
- Enterprise data warehouses were an attempt to create a "data mainframe," but they are expensive and rigid, two traits that are distinctly undesirable in today's cost-conscious, on-demand world. Furthermore, many organizations have found that building the single, all-encompassing data model mandated by this approach simply won't work in the real world.
- Data warehousing appliances, which are proprietary, turnkey hardware and software solutions, perpetuate rigid, fragmented silo approaches to information access.

Given that existing technology has been unable to properly address the dramatic growth and distribution of information, it's no surprise that many enterprises find themselves in the predicament illustrated in Figure 1.

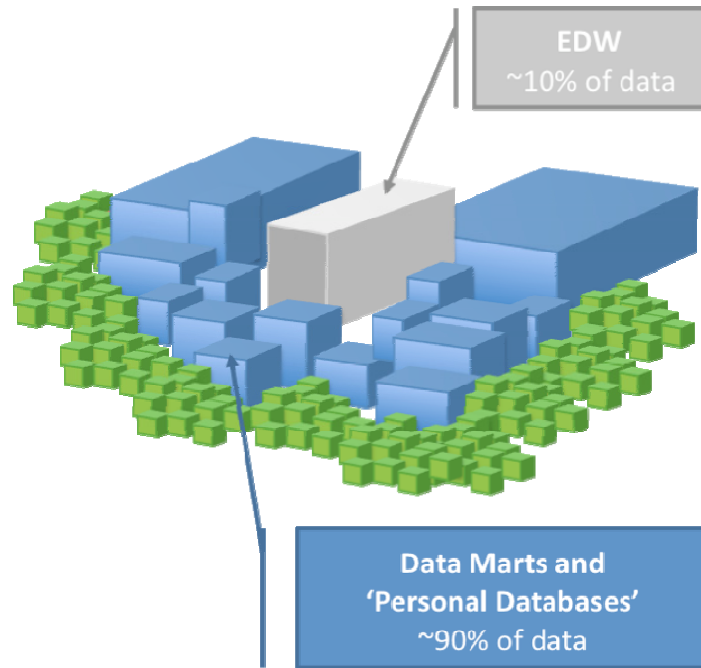


Figure 1. The majority of the organization’s data is hidden and locked away in silos

In contrast to the shortcomings of the above approaches, Enterprise Data Clouds offer a number of substantial advantages:

- **Self-service provisioning.** Business analysts are provided with a Web-based user interface to quickly produce virtual data warehouses. These warehouses can be created instantly, and combined from multiple locations. Greenplum employs scatter/gather streaming technology to let business analysts quickly load their own data. This relieves IT of many burdens, letting them focus solely on assembling pools of servers for provisioning.
- **Parallelization and expandability.** The Enterprise Data Cloud offers extreme scale and elastic expansion. Your data volumes can be dynamically expanded or reduced, depending on your needs. It also supports massively parallel analytic processing using SQL or MapReduce.
- **Scalability and performance.** Enterprise Data Clouds scale from tens of terabytes to tens of petabytes. In spite of the sheer amount of data available for access, business analysts are free to run extensive queries without having to worry about impacting production applications or other analysts.
- **Deployment flexibility.** The Enterprise Data Cloud can run on internal hardware, or using external resources hosted in the cloud. As we’ll see in this white paper, the EMC CLARiiON CX4-960 is an ideal platform for this type of application.
- **Data mart consolidation.** Data marts on existing platforms can easily be migrated to the EDC within a very brief amount of time, while preserving the organization’s investment in supporting technologies such as business intelligence.

EMC CLARiiON

The EMC CLARiiON CX4 series delivers industry-leading innovation in midrange storage with the fourth-generation CLARiiON CX storage platform. The unique combination of flexible, scalable hardware design and advanced software capabilities enables EMC CLARiiON CX4 series systems, powered by Intel Xeon processors, to meet the growing and diverse needs of today’s midsize and large enterprises. Through innovative technologies like Flash drives, UltraFlex technology, and CLARiiON Virtual Provisioning™, customers can decrease costs and energy use while optimizing availability and virtualization.

Configuration building blocks

Greenplum is an MPP database, designed to work optimally in a homogenous compute environment. By combining the processing power of several to several hundred machines, you can easily scale Greenplum to your compute and storage requirements. To this end, we will outline a building block that combines a Greenplum compute with EMC CLARiiON storage. Using this building block, you can start with as few as four servers with 21 TB of aggregate storage and grow to over a thousand machines with 5 PB of uncompressed storage.

Storage building block

The building block starts with an EMC CX4-960 configured with 71 drives. This configuration will require 21 rack units of space and about 2,300 watts of power, thus occupying half of a standard 42U rack. The detailed components are as follows:

- Two service processors
- Five Vault hard drives
- 64 data drives (600 GB or 1 TB)
- Two spare drives
- Four FlexIO modules
- Five DAEs
- One battery backup unit

The net usable space for the database depends on a number of factors, including the drive size, drive count, and RAID protection level.

The raw usable capacity (R) is defined as follows:

D = Drive count
C = Drive capacity
O = RAID overhead

$$R = D * C * O$$

Assuming RAID 10 with 1 TB drives, the equation would look as follows:

$$R = 1,000 * 64 * .5$$
$$R = 32,000 \text{ MB}$$

Once you have determined R, you need to subtract out file system overhead and sort space required by the database. For file system overhead, we will assume 10%, which is fairly conservative. For sort space, we'll assume 33% overhead, after file system formatting.

Given the above, usable space (U) for the database can be calculated as follows:

$$U = (R * 0.9) / 1.33$$

Assuming 1 TB drives, the equation works out as follows:

$$U = (32,000 * 0.9) / 1.33$$
$$U = 21,654 \text{ MB}$$

With 2 TB drives on the horizon, the usable space will soon jump to 43 TB. If you plan on using 600 GB FC-AL drives, the usable space works out to ~13 TB.

The following table of attributes itemizes the components of the building block:

Table 1. Components of the building block

Item	Space usage
DAEs (five @ 3U each)	15U
SP chassis	6U
Total	21U
Disk type	Count
Data	64
Hot spares	2
System drives	5
Total	71 disks

In a RAID 10 configuration, measured sustainable throughput for this building block is approximately 3.1 GB/s.

Scaling out the storage building block

As we will describe later, scaling out your EDC entails deploying additional CX4 units, as illustrated in Figure 2:



Figure 2. Sample scale-out configuration with four CX4 building blocks and Cisco UCS blades

There are a number of dramatic benefits to this approach, including:

- High storage capacity per unit of rack space. Even when the overhead of RAID 10 is included, this can be as much as two to 10 times greater than other offerings
- Superior virtualization with server motion and load balancing
- Simpler, proven storage management including disaster recovery and backup

-
- Improved high-availability model using LUN takeover

Compute building block

In addition to storage, the Greenplum system will require a number of commodity servers to provide compute resources for the database. A typical compute node would be configured as follows:

- Two CPU sockets and at least eight CPU cores (total)
- 4 GB to 6 GB of RAM per CPU core
- Two 10 Gb NICs (for redundancy) or four 1 Gb NICs (for redundancy and throughput)
- Two dual-port 4 GB HBAs (for redundancy)
- CentOS 5.x, RHEL 5.x, or SLES 10 SP2

In most situations, four servers per CLARiiON storage node would be sufficient. However, since every workload is unique, this number can easily be doubled or halved, depending upon your compute requirements. For the majority of workloads, four servers will be a good fit.

Flexibility

One of the most compelling aspects of the Greenplum/CLARiiON configuration is its ease of setup and maintenance. By detaching the storage nodes from the compute nodes, you're free to add as many computational resources as necessary. You may also incorporate additional storage nodes, or simply expand a single resource without impacting any other nodes. As we'll see later, configuring the Greenplum Database is handled through a software interface, thus simplifying managerial and administrative tasks.

Finally, you're able to choose the optimal connectivity option, from 10 Gb iSCSI to Fibre Channel or Fibre Channel over Ethernet (FCoE).

Configuration guidelines

Now that we've itemized all of the components that make up the Greenplum/CLARiiON building block, let's explore how to configure these elements for optimal performance. Note that performance of data warehousing and analytical workloads is primarily dependent on the ability of the I/O system to perform sequential scans at high bandwidth.

Even when these systems are used with high concurrency, the database and operating system together reorder the I/O into blocks of 512 KB reads, which is a "piecewise sequential" workload. Consequently, the performance tuning of the system focuses on achieving the maximum possible sequential bandwidth from each CLARiiON array in the system. Additionally, because the Greenplum Database is a scale-out MPP architecture, multiple CLARiiON arrays are commonly used to increase the I/O performance of the system beyond the limitations of any single CLARiiON array.

Balanced I/O configuration

MPP database systems operate at peak effectiveness and efficiency when their resources are distributed across the entire cluster. Thus, it's optimal to ensure that all of the work that you assign to the CLARiiON storage platform is balanced. The Greenplum Database will automatically adjust your data across the compute nodes. However, to achieve the same benefits with the storage system, there are a number of configuration options to consider, as we'll explore next.

Balancing I/O activities between the two storage processors

The CX4-960 employs the traditional system architectural technique of having I/O requests from servers against LUNs automatically serviced by one of the two storage system processors (SPs) that "owns" the LUN at any one instance in time.

To ensure that all LUNs used by the Greenplum engine are balanced, each SP is configured with exactly half of the LUNs. In the case of the RAID 10 building block standard, the Greenplum choice is to establish 16 LUNs from eight 4+4 R10 groups using 64 of the 71 drives in the CX4-960 unit, and associating eight LUNs with each storage processor.

Balancing data distributions among LUNs

To improve performance, streamline operations, and reduce manual administrative responsibilities, the Greenplum Database automatically distributes data across nodes and LUNs. The 16 LUNs on the CLARiiON are organized by Greenplum such that each of the four compute nodes owns four of the LUNs. By default, queries will be processed in parallel across these LUNs, thereby engaging all spindles and CPUs to compute the results.

Balancing traffic across the back-end buses

In the CLARiiON system architecture, each disk drive in a DAE is simultaneously available from both SP-A and SP-B. Each SP uses one of its back-end I/O ports to communicate with the DAE's link control card (LCC).

A back-end bus consists of a back-end I/O port from each SP connecting into the same DAE. For the Greenplum building block configuration that we're describing in the paper, there are five back-end bus ports used on each SP, which results in two back-end buses per DAE. This combination ensures that each back-end bus has the same number of drives, and is perfectly balanced inside the storage array.

Each DAE has two 4 GB connections, one from each SP, which gives it a theoretical throughput of 720 MB/s. In reality, however, you will typically experience a throughput rate of less than 600 MB/s from each DAE, due to the aggregate processing limitations of the CX4 SP.

The practical bandwidth of the pair of CX4 SPs for real-world bandwidth intensive workload is approximately 3 GB/s. The configuration using five DAEs that we've been describing is designed to fully drive the disks and DAEs to achieve that peak SP bandwidth. Additional DAEs and drives will allow the usable capacity to be boosted but will not change the available peak bandwidth on the CX4-960.

Balancing I/O traffic among the UltraFlex I/O modules

In a standard CX4 system, the FC ports in the I/O modules are partitioned evenly between front-end and back-end ports. This division helps balance I/O traffic through the CLARiiON storage processors. Should more front-end or back-end ports be needed, the CX4-960 lets you add more as necessary.

In general, it's a good idea to balance the number of front-end and back-end ports. However, this isn't the case with the Greenplum building block that we're describing in this paper. In this example, there are a total of eight front-end ports (per SP) and five back-end ports. The rationale behind this configuration is that additional back-end ports are not necessary. Because with the CLARiiON architecture, 10 total back-end ports are already able to fully support the bandwidth feed to optimally engage the CLARiiON storage processors.

Disk and RAID configuration

For the building block that we're describing in this paper, we recommend RAID 10 protection. Although RAID 10 has the largest overhead with respect to data protection, it also lets you employ slower SATA drives. Despite using SATA drives, RAID 10 still lets you realize the performance you would expect from FC-AL- or SCSI-based drives. By closing the performance gap between FC-AL and SATA drives, RAID 10 delivers superior cost-effectiveness with both 1 TB and (eventually) 2 TB disks.

Cache memory configuration

The CX4-960's kernel lets you use up to 16 GB of physical system memory, which is a key performance differentiator. Even after allowing for system-resident kernel software and drivers, there is still more than 13 GB of available storage processor memory available to be configured as data cache for use by the different LUNs. This approach lets you configure dramatically larger read and write caches. For example, you may configure more than 10 GB of mirrored write cache, which is a significant increase from the 3 GB maximum supported in the CX3-80.

In addition, the CLARiiON cache configuration supports a wide range of cache page sizes. Applications such as EDC tend to use larger database page sizes, such as 32 KB. In general, it's wise to match the cache page size to a multiple of the database page size. Selecting the maximum value (16 KB) will tend to enable the storage system to send the largest possible single disk I/O requests through the back end to the physical disk drives.

Write cache configuration

Given the massive size of most EDC deployments, there's minimal advantage to allotting significant cache for reading or prefetch operations. Thus, we recommend that all cache memory be assigned to the write cache. By following this approach, you help ensure that writes are batched as large as possible, which can dramatically enhance the rate at which the database loads. Additionally, it's quite common to experience numerous "smaller" write operations that are driven by sorting or temporary area overflow. Since you've maximized the write cache, these additional write operations will also occur as quickly as possible.

Read cache configuration

Surprisingly, given the predominant read patterns typical of EDC-style workloads, allocating storage read caching for data LUNs generally doesn't produce as much of a performance payoff as might be expected. As a matter of fact, implicit LUN data prefetching may have the counterproductive effect of wasting I/O: Data that has been prefetched may not be consumed before it is aged out of the cache. Because of the caching present on database servers, it's nearly impossible for the CLARiiON storage system to achieve any meaningful read cache rehits.

Thus, we recommend disabling data caching and prefetching on all data LUNs and at the array level. To disable LUN prefetching (which can be done dynamically), set the prefetch policy for a particular LUN to NONE using the Navisphere® Manager or via NAVISECCLI. Read cache can easily be disabled at the array level within Navisphere, eliminating the need to disable read caching for each LUN.

Transaction log placement

Best practices for most database system dictate that the transaction log is placed on its own dedicated storage – even Flash drives if available. On the other hand, Greenplum Databases don't have this requirement, since there is minimal logging. By leveraging visibility bits, Greenplum is able to avoid much of the logging that occurs in other database solutions. This is another example of Greenplum's simplicity.

Deploying Flash drives

Flash drives (FD) are one of the many storage options available with the CLARiiON product line. Without any moving parts, FDs are capable of servicing many random I/O requests per unit time, with low milliseconds I/O service time. A typical FD can sustain up to 200 MB/s of random or sequential I/O. As an added benefit, unlike traditional spinning disks, there is no penalty for random I/O when using FDs.

Although the majority of EDC workloads don't require FDs, there are certain scenarios where they add value:

- As a staging area where new table partitions can be loaded and later migrated to traditional storage. This can be useful if the most recent data is accessed more frequently or if trickle feeding is used.

-
- High-performance storage of indexed tables. When an index is used to address table data, the access pattern to the base table becomes random instead of sequential. By placing the base table on FD, indexed access achieves maximum performance. The tables themselves can be stored on normal disk in these configurations. Greenplum supports B-Tree, Bitmap, GIST, hash, and other indexing methods, although these types of indexes are not commonly employed in data warehouse environments,
 - Temporary storage area for use in doing ELT or other transformations.

Although not included by default, the Greenplum/CLARiiON building block that we're describing in this paper includes four empty drive bays that can be used for FDs. These drive bays are balanced across four DAEs and eight SP back-end storage buses, ensuring optimal performance. If you require additional FD storage space, then you'll need to conduct some additional planning. In this case, your options include adding additional DAEs or substituting existing drives with FDs. In general, deploying and configuring FDs should be considered on a case-by-case basis.

Greenplum/CLARiiON building block and EDC

Now that we've described the optimal configuration for the CX4-960 building block, let's examine how Greenplum employs this foundation as a participant in an EDC. For brevity's sake this exploration will be high level; for more detailed hands-on instructions, consult the *Greenplum Database Administrator Guide*.

If you're using a traditional hardware configuration, the Greenplum/EDC building block delivers an excellent experience with superior performance. However, as we described earlier, sophisticated virtualization support is one of the most attractive capabilities of this combination. Unlike technologies from other database providers, Greenplum's parallelization architecture lets it aggregate the performance from multiple slower virtual machines to deliver dramatically faster results. While virtual machines always introduce some degree of latency, Greenplum's aggregation approach along with its increased flexibility and superior management capabilities outweigh these drawbacks.

As the logical view in Figure 3 shows, each CX4-960 LUN can be represented with one or more mount points (also described as filespace). These filespace service VMware ESX containers, thus enabling the server and database virtualization that are fundamental components of an EDC. For optimum efficiency, there will be four filespace per LUN, thus saturating I/O bandwidth. For enhanced flexibility, you may elect to create additional filespace, but there should always be a minimum of four.

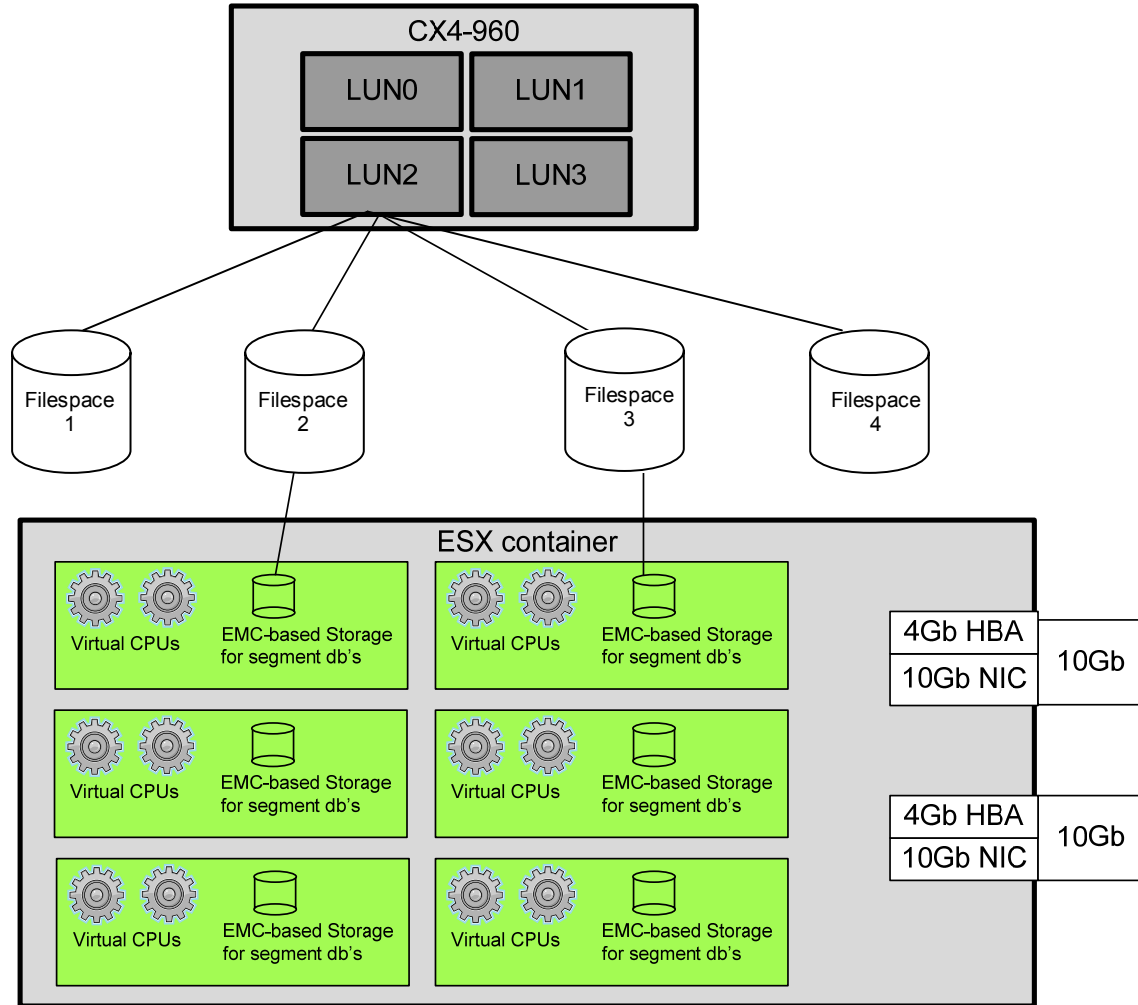


Figure 3. Virtualization and storage

This is a powerful architecture, yet Greenplum’s self-service provisioning capabilities and intuitive software hide the underlying complexity from the analysts, as illustrated in Figure 4.

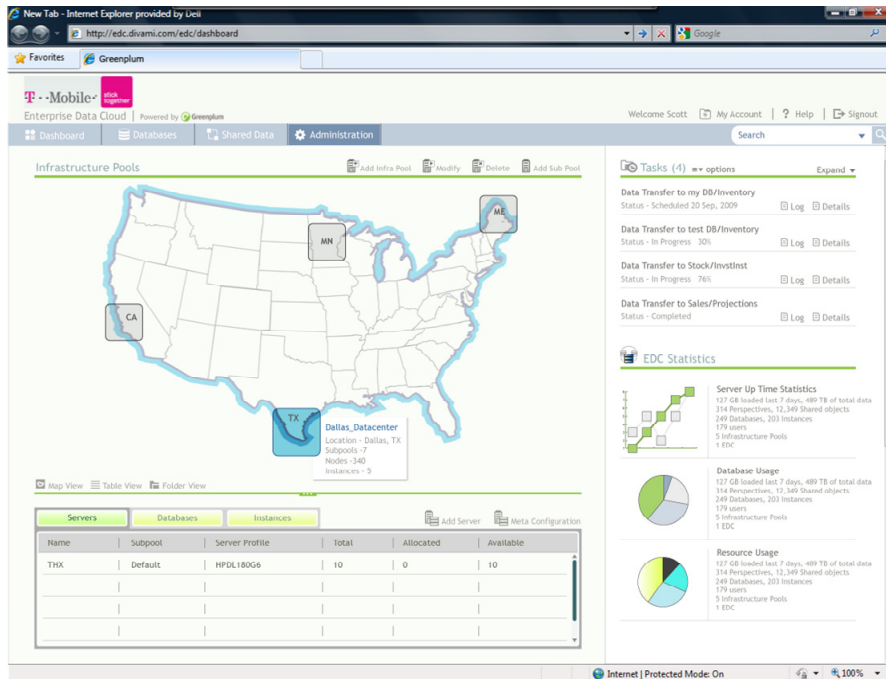


Figure 4. Self-service provisioning by analysts using Greenplum software

The combination of Greenplum and EMC enables rapid virtual provisioning of data warehouses. This speed translates to new opportunities for collaborative analytics, helping organizations realize dramatic benefits from the vast quantities of information available in an EDC.

Conclusion

From controlling costs, dealing with expanding data volumes, and responding to competitive challenges, Enterprise Data Clouds help address the challenges facing business and IT today. By bringing the power of self-service to data warehousing and analytics, an EDC helps enterprises derive more value from their information repositories, quickly and effectively.

Based on extensive experience with EDC deployments and testing of Greenplum on CLARiiON storage, the building block described in this paper provides a configuration that is balanced, thereby maximizing the available hardware. Furthermore, if you need more capacity or throughput, scaling up the solution is as simple as adding more building blocks, either at implementation time or after the platform has been deployed. The Greenplum Database supports seamless expansion, so it's easy to add storage and servers after a deployment.

References

The following EMC white papers are available on EMC.com:

- *Introduction to the EMC CLARiiON CX4 Series Featuring UltraFlex Technology*
- *Deploying EMC CLARiiON CX4-960 for Data Warehouse/Decision Support System (DSS) Workloads*

The following Greenplum documents are available on Greenplum.com:

- *Greenplum Database Administrator Guide*
- *Next Generation Data Warehousing Architecture with Greenplum and VCE*