

# SQL Server 2008 Data Warehouse Deployments with EMC CLARiiON CX4 Storage Systems

*Applied Technology*

---

## **Abstract**

This white paper describes some of the challenges presented by data warehouse deployments such as huge amounts of data, high-speed streaming I/O, and changing requirements. The infrastructure must provide high availability, modular growth, and state-of-the-art maintenance and monitoring facilities. Microsoft SQL Server 2008 software and the EMC® CLARiiON® CX4 midrange storage systems provide a solid foundation on which to build a successful data warehouse.

September 2008

---

---

Copyright © 2008 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 H5702

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	5
Terminology .....	5
<b>The data warehouse challenge .....</b>	<b>6</b>
<b>SQL Server 2008 manages the data warehouse.....</b>	<b>7</b>
Partitioning for scalability, availability, and performance .....	8
<b>CLARiiON storage systems manage the data .....</b>	<b>9</b>
Aging data with CLARiiON virtual LUN technology .....	10
<b>SQL Server 2008 features for DW.....</b>	<b>10</b>
Star join query optimization.....	10
Partitioned table parallelism.....	10
Partition-aligned indexed views .....	11
Enhanced data compression .....	11
MERGE statement .....	11
<b>CX4 UltraFlex features for DW .....</b>	<b>11</b>
UltraFlex I/O modules .....	11
Expanded storage caching support .....	12
<b>Conclusion .....</b>	<b>13</b>
<b>References .....</b>	<b>13</b>

---

## Executive summary

The successful enterprises of today recognize the value of business information. The modern enterprise is dependent on electronic data for sustaining and advancing the business. In many cases, electronic data has become a key business asset, directly responsible for business revenue generation.

When the term *data warehouse* was first coined, the emphasis was on capturing and retaining as much business data as possible, leveraging the economics of electronic data storage technology advances. As the volume of captured data grows exponentially, the challenges are shifting toward cataloging and managing the massive amounts of information collected each day.

But the real value of a data warehouse lies not in the amount of business data that can be captured, retained, and managed. Rather, the value comes from processing the raw data into information that can be used to guide the business direction, create competitive advantages, and identify marketable products and services to increase revenue and profits.

Microsoft's SQL Server platform integrates enterprise-class data management with versatile analytical tools. The EMC® CLARiiON® family of storage systems adds scalability, reliability, flexibility, and manageability to the mix. Together, SQL Server 2008 and CLARiiON CX4 storage systems provide a solid foundation on which to build the sophisticated management decision support system that today's enterprises depend on.

## Introduction

A *data warehouse* is a central repository for all or significant parts of the data that is collected by an enterprise's business systems. Data from various online transaction processing (OLTP) applications and other sources is selectively extracted and organized on the data warehouse database for use by analytical applications and user queries. *Data warehousing* is the process of designing, building, and maintaining a data warehouse system.

A data warehouse has many of the same challenges as a traditional warehouse for storing inventory:

- The desired data must be available when needed and structured so that applications can locate and access it quickly.
- The infrastructure must be able to grow over time and to adapt to changing business needs.
- The aging and eventual obsolescence of data must be managed.

The ultimate value of the data warehouse is not in the amount of data stored and managed; rather, the value is derived from *using* the data stored to enhance and improve the effectiveness of the business. Typical use cases may be the business reports that track, monitor, and report the health and progress of the business. More companies are leveraging accumulated historical business operation data to predict market trends, identify competitors' strengths and weaknesses, help shape new strategic direction, gain a competitive business edge, or shape a new market opportunity.

Microsoft SQL Server 2008 is a comprehensive database platform providing enterprise-class data management with integrated business intelligence tools. As a core component of the Microsoft Windows Server System, it dramatically reduces the time required to bring data warehousing applications to market. The SQL Server 2008 database engine provides more secure, reliable storage for both relational and structured data in combination with the best in analysis, reporting, integration, and notification. A SQL Server data warehouse is a reliable, secure, and productive data platform for line-of-business and analytical applications. It provides an end-to-end development environment that includes many new technologies that empower developers and significantly increase their productivity. The comprehensive analytical, integration, and data migration capabilities put critical, timely information into the hands of decision makers, empowering them to make better decisions faster.

The EMC CLARiiON CX4 series with breakthrough UltraFlex™ technology delivers advanced software functionality, best-in-class performance, and unparalleled flexibility providing investment protection. Unmatched scalability is delivered on this platform with up to 960 drives and 4,096 LUNs with the CX4-

---

960 model, accommodating a multitude of drive types ranging from extremely high-performance drives to extremely dense drives. The result is not only good scaling in total capacity, but also high transfer rates, highly available data, and storage tiering to allow for the most cost-effective manner of storing and managing SQL Server database solutions. For additional information on the CLARiiON CX4 series with UltraFlex technology, refer to the [EMC.com website](http://EMC.com).

All CLARiiON storage systems are delivered with software that lets you discover, monitor, provision, and report on multiple CLARiiON arrays from a simple browser interface. EMC offers other advanced software capabilities, including powerful management tools, full and incremental replication software, and data mobility tools.

## **Audience**

The intended audience for this white paper is decision-makers tasked with planning, deploying, or maintaining a data warehouse, including database administrators, storage administrators, and IT professionals.

## **Terminology**

**AS:** SQL Server Analysis Services.

**ATA:** Advanced Technology Attachment. A standard for connecting hard drives to computer systems over a network. Also known as IDE (Integrated Drive Electronics).

**business intelligence (BI):** Analyzing available information to help the enterprise make business decisions.

**data mining:** Identifying patterns and establishing relationships in raw business data: connected events, sequenced events, predictions about future events based on past events, and so on.

**clone:** A full binary copy of a storage element.

**FC:** Fibre Channel. A technology for transmitting data between storage devices and computer devices at data rates of up to 4 Gb/s (and 10 Gb/s in the near future) on a dedicated network (optical fiber, or for short runs, coaxial cable or twisted pair).

**HA:** High Availability. A system or component that is continuously operational for a long length of time.

**LUN:** Logical Unit Number. A unique identifier used to identify logical storage objects in a storage system.

**mirroring:** A technique in which data is written to two duplicate disks simultaneously. If one of the disk drives fails, the system can instantly switch to the other disk without any loss of data or service.

**RAID:** Redundant Array of Independent Disks. A way of storing the same data in several different places on multiple hard disks. I/O operations can overlap in a balanced way, thereby improving performance. Since multiple disks increases the mean time between failures, storing data redundantly also increases fault tolerance. A RAID storage system appears to the operating system to be a single logical hard disk device.

**RS:** SQL Server Reporting Services.

**sliding window:** A technique for modifying the range of data available in a partitioned table over time, adding new data in a new partition and rolling older partitions out to less costly storage options.

**snapshot:** A virtual point-in-time image of a storage element.

**SSIS:** SQL Server Integration Services.

**table partitioning:** Breaking a single database table into sections stored in multiple files.

**virtual LUN technology:** A technique whereby data stored on a LUN is bitwise “migrated” to another LUN within the same storage system, without restricting access to the data.


# The data warehouse challenge

The content of a data warehouse database changes infrequently compared to that of a transactional database. Used primarily for querying, analysis, data mining, and reporting, the data is typically loaded, cleansed, and updated on a periodic basis based on business rules and requirements. By the very nature of the architecture and charter of data warehouses, their storage requirements increase over time, frequently reaching the multi-terabyte or even petabyte range. By maintaining a data warehouse separate from the transactional database, the analysis applications, which may make heavy demands on the system, do not impact the response time for the mission-critical online transactions. Typical uses of the data warehouse include business intelligence, data mining, and report generation.

The ideal data warehouse infrastructure will:

- *Be available when needed.* If a warehouse door is shut, it does not matter how many useful parts and products are in the warehouse. When the storage components holding the data are inaccessible, business applications depending on the warehouse cannot be run until access is restored.
- *Adapt to changes in structure and scope.* The scope of the data warehousing project will increase over time. Furthermore, reorganizations, product introductions, new pricing schemes, new customers, changes in production systems, and so on, will require changes in the warehouse data, database, and applications. A modular infrastructure with the ability to grow in both storage capacity and access speed is required.
- *Incorporate a range of storage options in a modular, scalable architecture.* Data warehouses tend to accumulate data and continually grow in size over time. However, the value of the data tends to decrease over time. A storage system that can store and protect all the data, matching the value of the data to the appropriate tier of storage to maximize price and performance, is a major plus.
- *Process large sequential I/O operations efficiently.* The data in data warehouses tends to be write-once, read-often, and is accessed in bulk rather than as individual records. New data is typically written into the database in batches independent of analysis applications.
- *Include a rich set of system management tools.* The storage system administrator must be able to identify performance bottlenecks; add, remove, and retarget storage objects; and back up and restore data in a cost-effective, reliable manner.

The columns in Figure 1 represent the challenges presented in a data warehouse deployment as described in the paragraphs above. The rows in the table list some of the features of CLARiiON storage systems and SQL Server 2008 databases that offer solutions to these challenges.

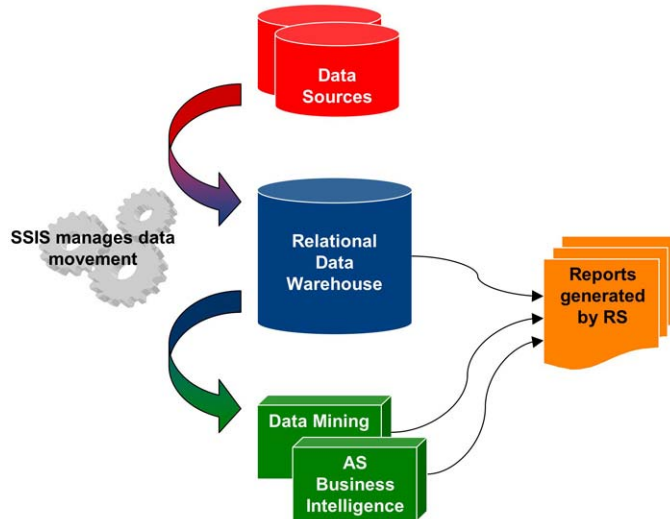


**Figure 1. Meeting the data warehouse challenge: SQL Server 2008 and EMC CLARiiON**

---

## SQL Server 2008 manages the data warehouse

SQL Server 2008 includes the tools you need to create a data warehouse and its associated analytical applications. As shown in Figure 2, data from various sources can be integrated and written to create the relational data warehouse using SQL Server Integration Services (SSIS), where Analysis Services (AS) and Reporting Services (RS) can operate on it and users access the results.



**Figure 2. SQL Server system architecture**

- The *SSIS* pipeline architecture allows data to be retrieved from various sources, manipulated in memory, and stored in the destination database without being written to disk in intermediate locations. This has tremendous performance benefits. The kinds of manipulations that can be performed in the pipeline include data validation, surrogate key lookups, derivation of new columns, aggregation, sorting, and more. SQL Server 2008 Integration Services introduces improved transformation pipeline performance and improved lookup performance.

Other aspects include:

The processing power of multi-core processors can be better utilized by the parallel processing improvements made in SSIS. Several threads are now able to work together to accomplish the task that used to be done by a single thread in SQL Server 2005 Integration Services.

SSIS allows the cache to be populated using a separate pipeline in the same package or a different pipeline in the same package or a different package. In addition the cache can now be saved to virtual memory or permanent file store so that multiple Lookup components can now share the same cache. A miss-cache feature was also introduced that improves Lookup component performance.

- *Analysis Services* takes data from the relational warehouse and prepares it for fast access by users. *Analysis Services* has an aggregation engine that helps it respond rapidly to queries at various levels of detail. The query language facilitates the expression of complex queries by client tools and web clients. SQL Server 2008 *Analysis Services* has made improvements with query speed using the new block computation, writeback partitions with MOLAP, and enhance backup optimized for larger databases.

Other aspects include:

Multidimensional expression query processing improvements in SQL Server 2008 is achieved by doing work for the non-null values in the cube space.

Prior to SQL Server 2008, AS required writeback partitions to have ROLAP storage, causing performance problems for cell writeback since ROLAP partitions must query the data source on demand to retrieve their data. AS in SQL Server 2008 allows the retrieval of writeback data from the compressed MOLAP format, which is faster than querying the data source.

---

The backup storage subsystem in Analysis Services has been rewritten to use a new format which greatly increases the performance and scalability of backup tasks.

- *Reporting Services* facilitates easy creation and retrieval of reports in many formats, including standard ones like HTML, Excel, and PDF. Reports can be defined to run using either a relational database or Analysis Services as the data source, and can be cached for rapid retrieval once they are generated. SQL Server 2008 Reporting Services has had major improvements to increase scalability.

The Reporting Services reporting engine has been improved so that it could render much larger reports than in the past. Moreover, Reporting Services does not run inside Internet Information Services (IIS) anymore. RS memory settings can now be configured to run more effectively with other services running on the same server hardware.

- The *SQL Server relational engine* provides reliable and scalable long-term storage for information in a data warehouse. In particular, table partitioning features provide a natural way to manage the data.

## ***Partitioning for scalability, availability, and performance***

Partitioning is a method of splitting up a large data set into smaller, more manageable chunks. All management operations on very large tables can be performed at a more granular level when the table is partitioned. Instead of actively managing a multi-terabyte database, a group of partitions, each of which is probably less than 100 GB, is managed.

Other benefits are as follows:

- *Partitioning tables can increase database availability.* For example, loading and deleting data can be done by partition. Backup can be limited to only the partitions that have been modified. An individual partition can be restored without bringing the table down. Thus, query activity to the rest of the table continues uninterrupted.
- *Partitioned tables are easier to maintain.* Indices can be reorganized, optimized, and rebuilt by partition, which is both faster and less intrusive than doing the same operation on the whole table.
- *Table partitioning can improve query performance.* Partitions limit the amount of data to be referenced, and, therefore, improve query performance. They also facilitate the use of multiple processors performing the query in parallel on distinct partitions.
- *Table partitions leverage tiered storage.* The declining business value of date-partitioned data can be matched to multi-tiered storage systems in a straightforward way using a *sliding window* technique, as shown in Figure 3.
- *Table partitions facilitate storage load balancing.* By evenly distributing the partitions across the available disk volumes, both I/O and lock contention are minimized.

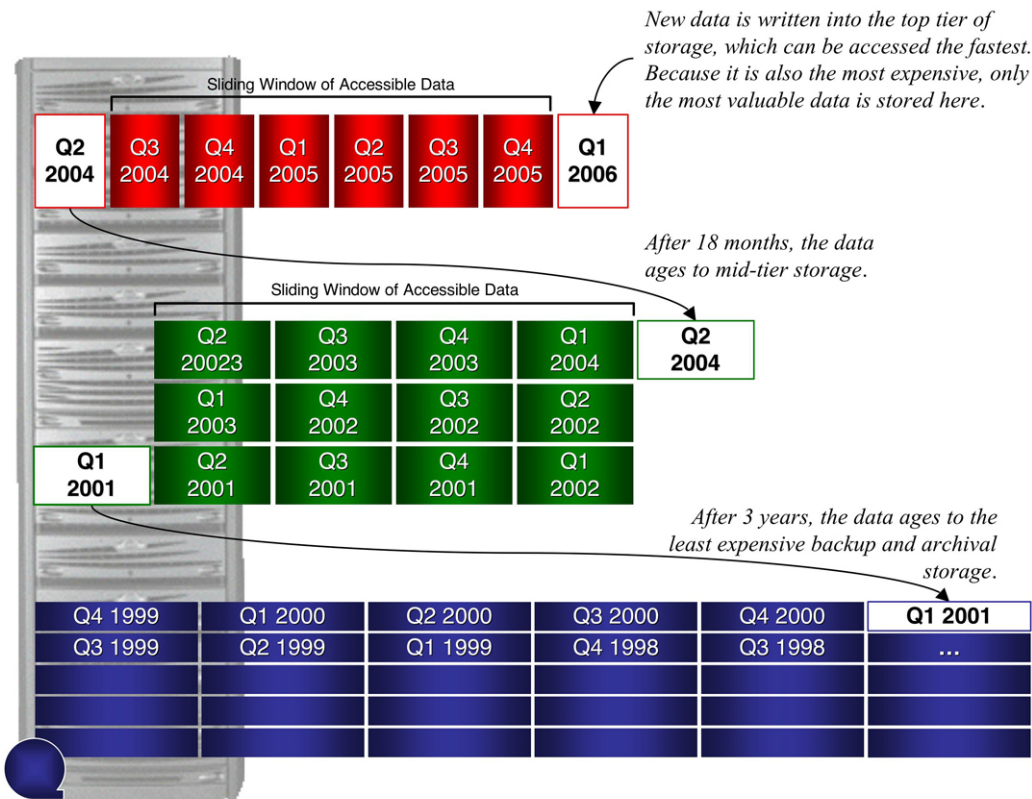


Figure 3. Sliding window aging scheme

## CLARiiON storage systems manage the data

The CLARiiON networked storage systems deliver industry-leading performance, availability, and scalability for data warehouse deployments. The ability to scale to many hundreds of terabytes in a single system in a modular, pay-as-you-grow model delivers strong investment protection. CLARiiON storage systems offer the following benefits:

- *Available when needed.* A time-proven HA architecture guarantees continual availability of warehouse databases stored in the CLARiiON system even with unanticipated storage component hardware failure; hot swap component support allows for failed components to be replaced and repaired without shutting the storage system down.
- *Adaptable to changes in structure and scope.* CLARiiON has a modular architecture to allow for capacity and corresponding performance scaling by adding storage enclosures on the fly to the storage system.
- *Multi-tiered storage.* CLARiiON provides integrated support for both high-performance FC drives and cost-effective, lower performance ATA drives within the same storage system. The integrated support nature allows data store and access in a uniform manner and protocol from the database software.
- *Efficient sequential I/O.* The use of RAID 5 or RAID 6 striping across multiple back-end disk spindles in CLARiiON storage systems produces high bandwidth aggregate data streaming, and is a good match for data warehouses, which typically read data more often than they write. This is especially true with the SQL Server database engine, which optimizes for sequential reading of table scans even with multiple concurrent read streams.

- 
- *Powerful management tools.* EMC's Navisphere® Manager provides an easy-to-use browser-based interface for configuring, monitoring, and administering CLARiiON storage systems. It displays the ever-growing storage infrastructure: where the data resides and how it is being used. Bottlenecks are quickly identified, allowing adjustments to be made to the storage infrastructure to contend with growth and change. Data replication and migration capabilities, like virtual LUN technology, SnapView™, and SAN Copy™, safely and efficiently back up, restore, and age the warehouse data.

### ***Aging data with CLARiiON virtual LUN technology***

CLARiiON virtual LUN technology makes it possible to transparently manage the data lifecycle. Using virtual LUN technology, logical volumes of data (for example, a table partition) can be moved between storage tiers, from Fibre Channel disks to ATA disks, for example, without disrupting applications. Virtual LUN technology can also help improve capacity utilization by moving data to underutilized spindles in the array.

## **SQL Server 2008 features for DW**

Although SQL Server 2005 had many features that made DW deployments feasible, SQL Server 2008 introduces enhancements and new features that make it a much more powerful relational data warehouse for decision support. The capabilities that extend SQL Server 2008 for DW deployments include the following and are discussed in the next sections:

- Star join query optimization
- Partitioned table parallelism (PTP)
- Partition-aligned indexed views
- Enhanced data compression
- MERGE statement

### ***Star join query optimization***

While most OLTP databases are designed using normalization rules for the database schema, in contrast, relational data warehouses are designed using dimensional modeling usually in the form of a star or snowflake schema. With dimensional modeling, the fact table contains the primary data while the dimension table contains descriptions about the data in the fact table. Due to this separation, the dimensional tables tend to be an order of magnitude smaller than the fact tables.

Since most star join queries process a large portion of the fact table, SQL Server 2008 has made star join query optimizations based on the volume of fact table data consumed by the query. Using a technology based on bitmap filters, SQL Server 2008 eliminates the processing of unnecessary rows from the fact table when evaluating queries. Because this enhancement is made on the SQL Server database engine, application modifications are not required to benefit from this optimization.

### ***Partitioned table parallelism***

SQL Server 2005 introduced the capability of partitioning large datasets into smaller, more manageable chunks. In addition multiple file groups can be designated for the partitions to help balance I/O for better performance. The processing of partitioned data also provided an inherent boost in performance since the worker thread could act on each partition simultaneously. There was a limitation that only one thread could be allocated to each partition in SQL Server 2005.

SQL Server 2008 introduces partitioned table parallelism to improve the performance of partitioned tables. Depending on how many processors are available in the server, SQL Server can now allocate multiple threads per partition to dramatically improve performance of queries that act on partitioned tables. Not only will the processing power of the server hardware be better utilized, the storage back end may also deliver more data since there are more requests from the server. The increased threads per partition will

---

drive more requests down to the storage system keeping it busy. This is crucial for the CX4 platform with its faster processing power, larger RAM configuration, and increased number of drives supported.

### ***Partition-aligned indexed views***

In SQL Server 2005, indexed views (or summary aggregates) defined on a partitioned table must be dropped when switching in or out a partition. Indexed views are used extensively for data warehousing to improve query performance but the need to re-create the indexed views can be an expensive task when new partitions of quarterly or monthly data was rolled in. In SQL Server 2008, there is no need to rebuild the aggregates because SQL Server automatically moves the matching partitions for partition-aligned indexed views.

### ***Enhanced data compression***

SQL Server 2005 SP2 introduced a new variable length storage option to help dramatically reduce the size of the database. The vardecimal storage format used for decimal and numeric data not only reduces the storage needs of the database but has the potential for increased performance and efficiency. With compressed data, each I/O operation on the disks results in reading up more data when compared to uncompressed data. This can help alleviate the storage bottleneck that may be caused by the scanning of large fact tables common in data warehouse queries. In addition, the data is also stored in the compressed format in the buffer pool. This means more data can be stored in the buffer, which increases the chances of a requested page being found in the buffer.

SQL Server 2008 enhances the native compression capabilities by adding row compression and page compression. Row-based compression extends the variable length storage format to all fixed-length storage formats available in SQL Server. Page compression reduces the data redundancy of columns in one or more rows of a database page. The use of data compression increases the processing power needed because of the compression and decompression step that is needed before the data can be processed. While queries are retrieving data during table scans, the server processor is usually waiting for the data to be delivered and has extra processing power needed to do the decompression. Data load is usually a processor-intensive task and the compression impact usually manifests itself during this phase.

### ***MERGE statement***

The MERGE statement allows multiple database manipulation language (DML) operations on a table or view in a single Transact-SQL statement. The MERGE statement helps facilitate efficient INSERT and UPDATE operations for slowly changing dimensions (SCD) and to maintain the fact table for scenarios commonly found in data warehouse deployments. Although running separate INSERT, UPDATE, and DELETE statements can achieve the same result as the MERGE statement, the MERGE performs better because it only requires a single pass of the data.

## **CX4 UltraFlex features for DW**

The fourth-generation EMC CLARiiON CX4 series with UltraFlex technology is based on extensive technological innovation, delivering a midrange storage solution with unmatched application performance, new levels of scalability, reliability and flexibility, and affordability. Two new CX4 system features in particular enhance data warehouse deployments for SQL Server 2008.

- UltraFlex I/O modules
- Expanded read/write caching support

### ***UltraFlex I/O modules***

Previous CLARiiON storage systems were architected with a fixed number of usable I/O ports for supporting front-end I/O from different servers, using FC or iSCSI protocols, as well as a fixed number of

---

usable I/O ports for the storage processors to push I/O through the system back-end I/O buses to the physical drives. The new CX4 systems are designed to have the I/O ports separated out into I/O modules that can be plugged in to the storage system processors. Depending on the particular system type, up to six such I/O modules can be plugged in to each of the storage processors.

For the FC type modules, which come with four ports at 4 Gb/s speed, these can be flexibly adapted to be used as front- or back-end connection ports. The iSCSI modules come with two 1 Gb/s ports for front-side iSCSI connections.

This new architecture provides for great flexibility in terms of configuring the storage system for more front-side or back-end physical I/O paths to optimize and adapt to the changing I/O requirements that are frequently a challenge in many rapidly growing and evolving SQL Server 2008 data warehouses.

By default, all CX4 series storage systems ship with two I/O port modules per storage processor. The first I/O module has four Fibre Channel GBIC slots that are split evenly between front- and back-end ports and a second I/O module that provides two 1 GIGE ports for iSCSI connectivity. This versatile configuration allows the CX4 series to be deployed in any environment because both iSCSI and FC connectivity comes built in. Depending on the specific CX4 storage system model, varying numbers of I/O modules can be added to the storage system to scale the performance of either protocol to meet the workload requirement.

Today, demanding high performance data warehouses rely on 4 Gb FC to transport data from servers to the storage back ends. As new technologies such as 10 GIGE and 8 Gb FC gain wider adoption, I/O modules based on these technologies will allow customers to run their data warehouse on either protocol. By isolating the connection ports in this modular fashion, new UltraFlex I/O modules based on next generation connection technology offering bandwidth improvements can readily be introduced and added to any existing data warehouse deployment by merely adding or replacing I/O modules, without affecting any data already existing and stored inside the storage system.

## ***Expanded storage caching support***

The new system architecture supports up to 16 GB of memory in each of the CX4-960 storage processors for a total of 32 GB. The CX4's write cache has been enhanced to take advantage of the larger RAM available on the next generation platform. Write-intensive workloads are now able to use more than 3 GB of cache to increase response times and throughput. In the context of data warehousing, the load phase and TEMPDB are two examples that may benefit from the larger write cache. The online LOGs may also benefit but bulk loads with SQL Server should be done with minimal logging enabled.

Because the writes from data loads tend to be large bursts of write I/O down to the storage system, having a larger write cache will help to absorb this burst of I/O. With predominant I/O tending to be read focused, parity RAID, such as RAID 5 and RAID 6, is a common choice for data warehouse deployments as this helps to minimize the physical storage cost to support data RAID protection. RAID 5 or RAID 6 data writes are generally the most efficient when these can be done as storage LUN stripe boundary aligned, full stripe, writes. With many database files created and manipulated within NTFS file systems, it is not always possible for write I/O requests from SQL Server to be done on an ideally aligned storage LUN stripe boundary. The expanded write cache will help significantly in buffering up the incoming new data, coalescing contiguous data chunks, and allowing the new data to be destaged to the back-end drives as stripe aligned, full stripe, writes, which greatly enhances the data load I/O process.

Additional storage processor memory can also be dedicated to provide for read caching. For the LUNs holding the database files, the extra read cache will help with providing more headroom for doing predictive LUN data read ahead by the storage system to improve typical data table scan type I/O efficiency. The additional cache memory will also enhance the I/O latency for I/O going to the TEMP space file systems, where transient data is being created and reused quickly (such as intermediate data sorting).

---

## Conclusion

The reliability, scalability, and performance features of the EMC CLARiiON CX4 storage system form a solid foundation on which to build a data warehouse. The EMC storage replication features effectively address many of the key operational challenges of managing and operating an ever-growing warehouse database or databases, such as backup, replication across usage scopes, business repurposing, and disaster protection. In particular, key CLARiiON features such as virtual LUN technology and consistent storage replication support can be fully leveraged to reduce operational and support cost without compromising practical performance, taking full advantage of typical data warehouse operational characteristics. The UltraFlex technology and large write cache capabilities found in the fourth-generation CX4 series protects your investment by allowing the storage system to scale with the ever-increasing capacity and performance requirements of modern data warehouses. The combination of the CX4 storage system and new SQL Server 2008 data warehouse product suite is an ideal choice for any new and serious enterprise data warehouse deployment.

## References

- *EMC CLARiiON: Meeting the Data Warehouse Challenge* white paper
- *Deploying EMC CLARiiON CX4-960 for Data Warehouse/Decision Support System (DSS) Workloads* white paper
- *Introduction to the EMC CLARiiON CX4 Series Featuring UltraFlex Technology* white paper
- [Microsoft SQL Server page](#) on Microsoft.com
- [TechNet webcast: SQL Server 2005 Series \(Part 6 of 10\): Managing Large Databases using Partitioning](#)