
Introduction to the CX300

Abstract

This white paper introduces the architecture and functionality of the CX300 storage system. It compares the features and capabilities of the CX300 with other EMC[®] CLARiiON[®] storage systems. The document lists the CX300's software capabilities.

Published 6/13/2006

Copyright © 2004, 2006 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 C1090

Table of Contents

Executive Summary	4
Intended Audience	4
Introduction	4
Where the CX300 Fits into the EMC CLARiiON Storage Family	5
The CX300 Major Subassemblies	6
The CX300 Storage Processor	7
The CX300 Drive Enclosures: The DPE2 and DAE2	8
The CX300 Drive Canister	9
The Standby Power Supply	9
Cabling the CX300 Storage System	10
Software on the CX300	10
FLARE Operating Environment	10
Navisphere Management Suites	11
MetaLUNs	11
Access Logix (Included in Navisphere Management Suite)	11
SnapView Snapshots (Optional Application)	12
SnapView Clones (Optional Application)	12
SAN Copy (Optional Application: Cannot Be Loaded on CX300)	14
PowerPath	14
Non-Disruptive Upgrade	14
Upgrading to a CX500/CX700 or CX3 Series	15
Summary	15
Appendix	16

Executive Summary

This white paper describes the architecture and functionality of the CLARiiON® CX300 storage system. It compares the features and capabilities of CX300 with other CLARiiON storage systems. The CX300 storage processors (SPs) and SP subcomponents are described. In addition, major system subassemblies and configurations are described and illustrated. The paper provides details on the software capabilities of the CX300 and also introduces new features.

Intended Audience

This white paper is intended for System Engineers, EMC® partners, members of the EMC and partners sales and professional services communities, or anyone who desires an in-depth understanding of the CLARiiON CX300 storage system.

Introduction

The CX300 is the entry point into EMC's seventh-generation, full Fibre Channel CLARiiON CX series storage systems. The CX300 consists of the following modular components:

The DPE2 (2 Gb disk processor enclosure) — Houses the storage processors (SPs) and the first 15 drives.

A DAE2 (2 Gb disk array enclosure) — Houses up to 15 drives and can be added to, to scale to a maximum of 45 drives, yielding an array total of 60 drives.

A standby power supply (SPS) — protects the system's mirrored write cache.

The following are CX300 data storage-system features:

Supports both direct-attach and shared storage (SAN) environments

Hot-swappable storage processors with up to 1 GB of memory per storage processor

FC-AL and FC-SW front-end support

RAID level 0, 1, 1/0, 3, 5 and individual disk support, along with global hot sparing

NDU (non-disruptive upgrade) capability

Five-drive minimum to 60-drive maximum system configuration

Supports data-in-place upgrade to the CX500, CX700, and CX3 systems.

Supports data-in-place upgrade from prior generation storage systems like CX200LC and CX200. No conversions from CX300 to CX400 and CX600 are supported.

Supports Windows, Solaris, NetWare, Linux, HP-UX, and AIX ¹ operating systems

The CX300 supports the following software packages:

Base (FLARE™) code

Navisphere® Management Suites (include Access Logix™ technology and Navisphere Agents)

CLARAlert®

PowerPath® (Full-Featured Suite, PowerPath Base, or Utility Kit PowerPath)

SnapView® (snapshots and clones)

The CX300 will not be offered as a single storage processor configuration (CX200LC) as compared to the CX200.

Where the CX300 Fits into the EMC CLARiiON Storage Family

The CX300 storage system is an intelligent storage product in the CLARiiON product family. It has the functionality to provide disaster tolerance, data integrity, high availability, and configurability. Additionally, it offers excellent price/performance benefit along with scalable capacity and upgradability.

See Table 1 for a comparison of many of the CX300 features with those of the CLARiiON CX200 storage system.

Table 1. Side-by-Side Comparison of the CX300 with the CX200 Storage Systems

Components/Connectivity	CX200	CX300
Processor architecture per SP	1 x 800 MHz	1 x 800 MHz
Physical memory per SP	512 MB	1 GB
Max drives per storage system	30	60
Min drives per storage system	4 (5 with write cache)	4 (5 with write cache)
Max IOPs	40,000	50,000
Max bandwidth (cached/sustained)	200/200 MB/s	680/360 MB/s
Max front-end FC ports per SP (2 Gb)	1	2
Max back-end FC ports per SP (2 Gb)	1	1
Max initiators per storage system	30	128
Max H/A hosts per storage system	15	64
Max drives per RAID group	16	16
Max LUNs per RAID group	128	128
Max LUNs per storage system	256	512
Minimum configuration	4 U	4 U
Software Support		
SnapView snapshots	No	Yes
SnapView clones	No	Yes
MirrorView™	No	No
SAN Copy™	No	No

The CX300 Major Subassemblies

The minimal CX300 assembly consists of a single 3U disk processor enclosure (DPE2) and a 1U standby power supply (SPS), for a total of 4U. A fully configured CX300 consists of the DPE2 and SPS, plus three 3U disk array enclosures (DAE2s), for a total of 13U. Figure 1 shows minimum and maximum configurations of the CX300 assembly.

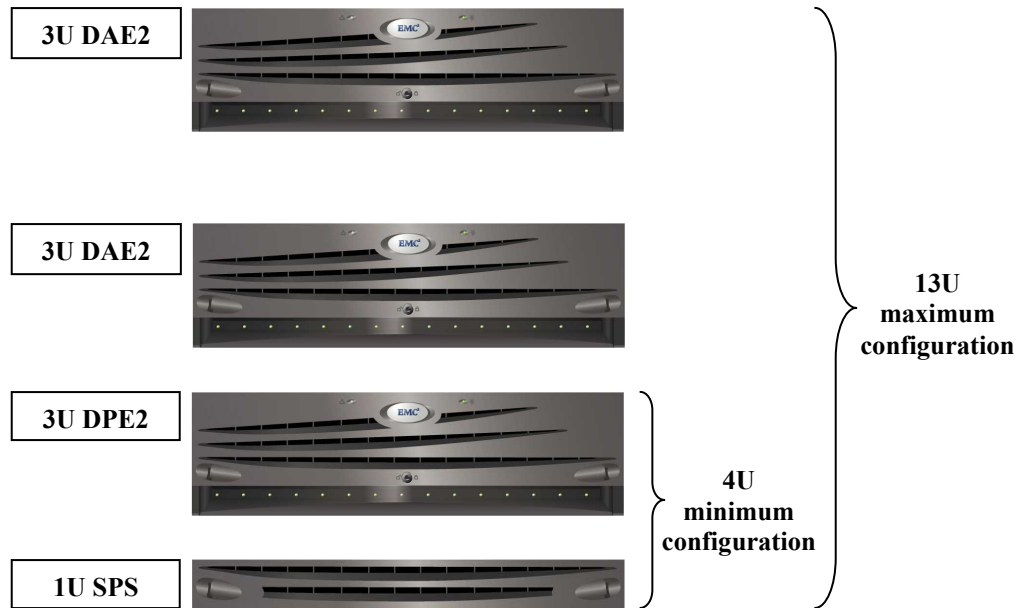


Figure 1. Minimum and Maximum CX300 Configurations

The SPS, DPE2, and DAE2 assemblies are rackmounted for efficiency, which in turn allows deployment flexibility. The CLARiiON 40U cabinet—with industry-standard 19-inch rack—houses different configurations of multiple CX300 systems, as well as FC switches (1U and/or 2U). The 40U cabinet comes standard with two power distribution units (PDUs). Power supplies on the left use the left PDU; power supplies on the right use the right PDU. Additionally, all CLARiiON storage systems can be installed in non-EMC industry-standard cabinets.

The CX300 Storage Processor

Each CX300 SP is powered by a single 800 MHz Intel Pentium III Tualatin-LP microprocessor, and has 1 GB of system memory also known as cache. The DIMMs used on the CX300 SP are not compatible with the CX500 SP or the CX700 SP. There are two 2 Gb Fibre Channel communication messaging interface (CMI) channels between the SPs. All FC protocol chips reside on a 64-bit, 100 MHz PCI-X bus. The LAN and other non-Fibre Channel I/O reside on a separate 64-bit, 33 MHz PCI bus.



Figure 2. CX300 Storage Processor (Back View)

Figure 2 shows the back view of the CX300 storage processor. LEDs indicate LAN connectivity, SP power, boot progress, and fault status, as well as power supply and blower status.

Each CX300 SP has two small form-factor (SFF-) or Lucent connector (LC-) style, 2 Gb optical front-end ports. These are unique ports and can be independently set to run at 1 Gb or 2 Gb speeds. Each SP has two HSSDC-style 2 Gb copper back-end ports, which with the peer SP form two redundant FC back-end loops for disk connectivity and capacity expansion.

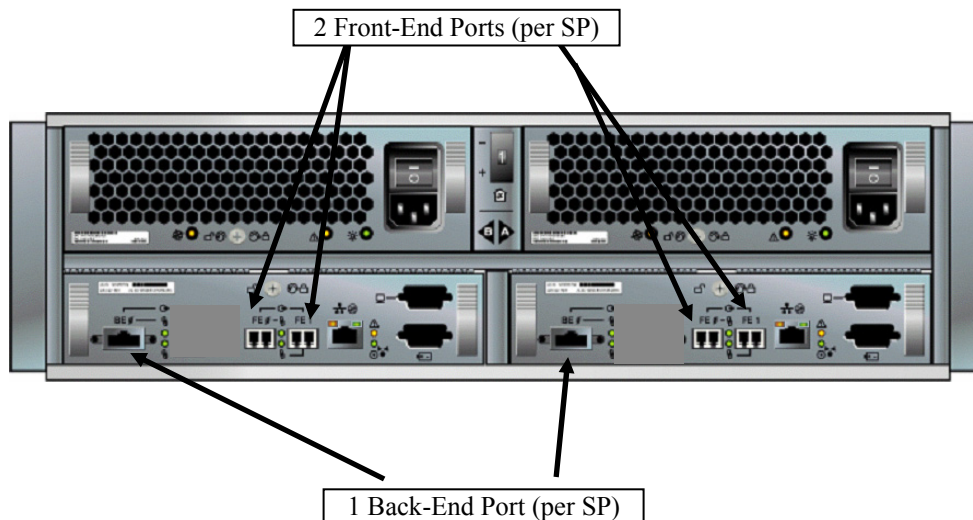


Figure 3. CX300 Storage Processor Front-End and Back-End Ports

The SPs are enclosed within the DPE2 (discussed in the next section). The DPE2 chassis is the base unit of the storage system.

The CX300 Drive Enclosures: The DPE2 and DAE2

The DPE2 and DAE2 use the same enclosure; the only difference is that the DPE2 has two CX300 SPs, while the DAE2 has two LCCs (link control cards). The first five drives of the DPE2 also contain storage system software (refer to “FLARE Operating Environment” on page 10). The enclosure has room for up to 15 drives. Drive 0 is at the furthest left side of the enclosure and drive 14 is at the furthest right side of the enclosure.

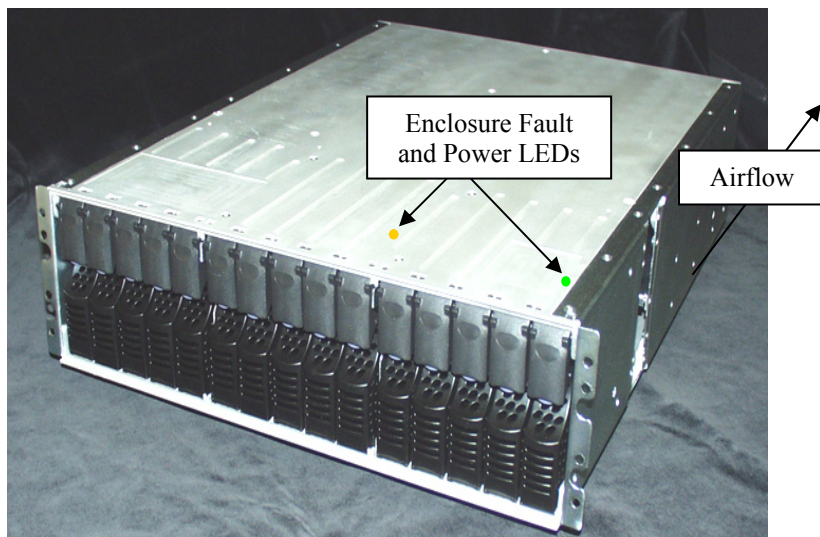


Figure 4. Front View of the CX300 DPE2/DAE2

Each CX300 SP has two 2 Gb optical front-end ports. These are shared ports, and the storage system can be set to run at 1 Gb or 2 Gb. Each SP has one HSSDC 2 Gb copper back-end port, which—with the peer SP—form one FC back-end loop for disk connectivity. A single microprocessor powers each CX300 SP. This processing power is the basis of performance, stability, and reliability in the CX300 storage system. Each SP has 1 GB of system cache. There are two 2 Gb Fibre Channel CMI channels between the SPs.

Enclosure fault and power status LEDs, as well as disk activity and fault LEDs, are visible from the front of the enclosure. Other fault indicators and unit ID switches are located at the rear of the enclosure. See Figure 4 for a front view of a CX300 DPE2 (or DAE2) enclosure.

Table 2. Components Housed within DPE2 and DAE2 Enclosures

DPE2	DAE2
2 CX300 SPs	2 LCCs
2 power supply/blower assemblies	2 power supply/blower assemblies
15 FC disk drives	15 FC or ATA drives disk drives

The CLARiiON with ATA implementation is at the disk-array enclosure (DAE2) level. With the exception of the LCCs and disk modules, all of the other DAE2 assemblies are identical to those of a standard Fibre Channel DAE2 assembly. The newly designed bridge controller cards (BCCs), for the ATA drives, convert serial ATA to serial Fibre Channel architecture. For more information on the ATA technology, refer to the *Introduction to CLARiiON with ATA Disk Drives and Enclosures* white paper.

The CX300 Drive Canister

Figure 5 shows a close-up view of the CX300 disk drive canister.

The canister holds 1-inch, 12-volt-only disk drives. For additional reliability, the drives are designed with a positive latching mechanism included on the canister itself to ensure a more complete insertion when inserting drives into the CX300 chassis. Each drive canister has drive-ready and fault LEDs. Fibre Channel drives supported at the initial release of the CX300 storage system include the 300 GB 10K rpm, 146 GB 10K rpm, 73 GB 10K/15K rpm, and the 36 GB 10/15K rpm disk drives. These are full 2 Gb FC assemblies.

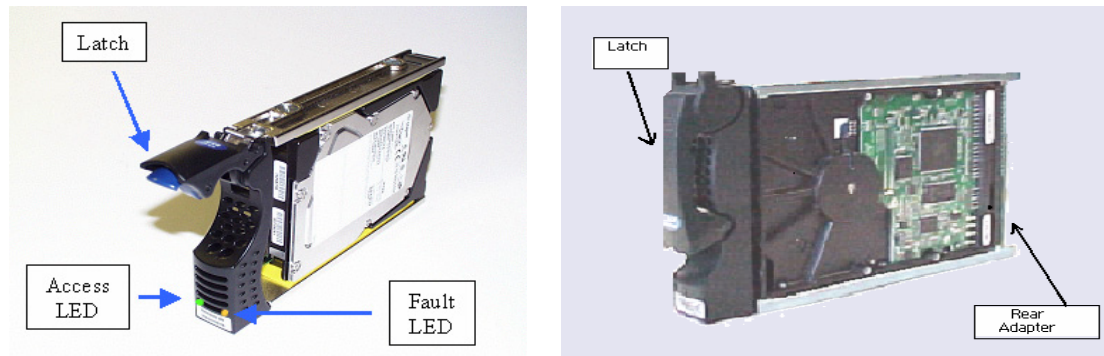


Figure 5. Two Gb Canister for Fibre Channel **Figure 6. Two Gb Canister for ATA drives**

The CX300 also supports the 500 GB SATA 7200 rpm ATA drives. Figure 6 shows the modified ATA disk carrier. The modified carrier allows you to convert FC to ATA and also converts standard single-port access for ATA to dual-ported access for more highly available and reliable disk access.

The Standby Power Supply

The CX300 DPE2 uses the same SPS as the CX500 and CX700 storage systems (Figure 7). The CX300 only requires a single SPS by default. In the case of the CX300 and the CX500 systems, only one of the SPS output connectors is used. This connector powers up one of the power supplies in the DPE2. The second SPS output is used only by the CX700 to power its first DAE2, which contains its boot drives.

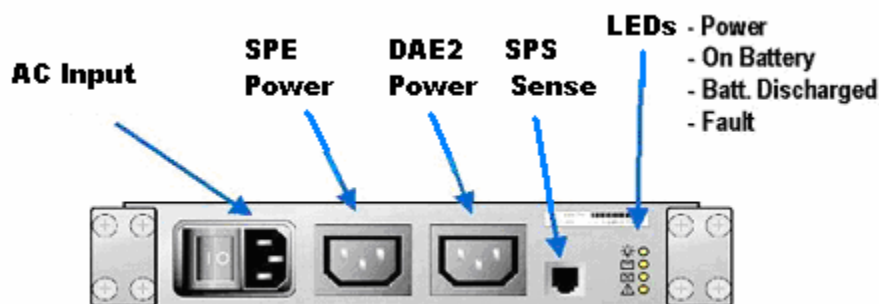


Figure 7. View of CX300 Standby Power Supply (SPS)

If using two SPSs for the CX300, it is important to cable each SPS so it connects completely to either the A side or the B side.

Looking at the system from the back:

- ◆ SPS on the right — Power-out and sense cables should connect to power supply A and SP A.
- ◆ SPS on the left — Power-out and sense cables should connect to power supply B and SP B.

If an SPS is cabled with the SPS sense cable going to SP A and the power-out cable going to power supply B (or vice versa), error conditions such as `Invalid SPS config` will be displayed when the SPSs are tested or are charging.

Cabling the CX300 Storage System

Figure 8 shows a rear view of a 60-drive CX300 system. Each SP has one back-end loop (BE#0). Each disk enclosure that connects to the same back-end loop must have a unique disk enclosure address (for instance: EA#0, EA#1, EA#2, and EA#3). The CX300 storage system is cabled in a similar manner to the CX200 storage system.

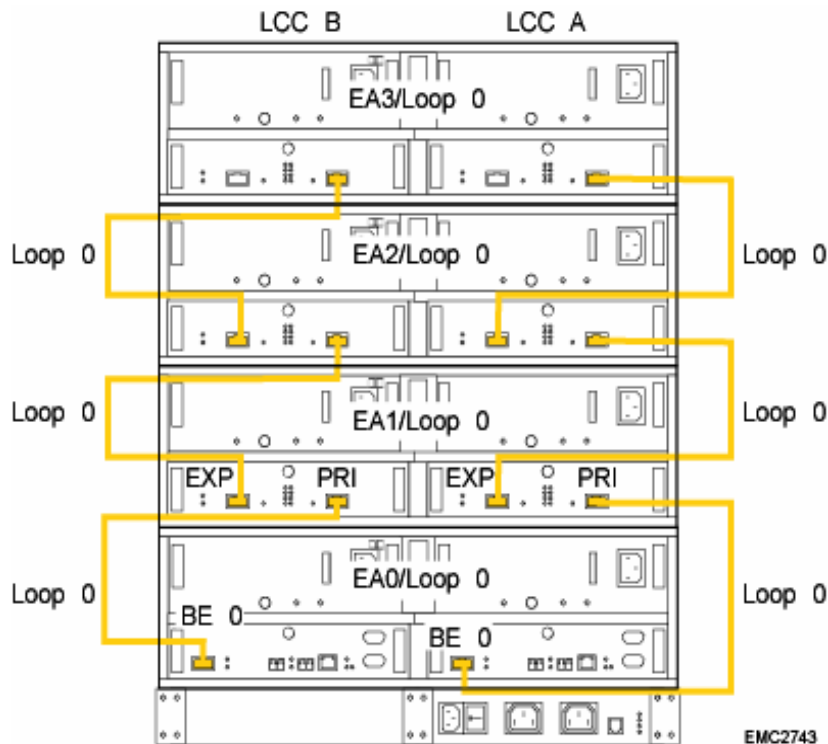


Figure 8. CX300 Cabling Diagram

Software on the CX300

The CX300 is designed to focus on intelligent storage management functions that can be operated in both direct-attach and SAN environments. The CX300 is capable of incorporating a full range of intelligent storage functions, such as logical unit number (LUN) masking and storage replication applications.

FLARE Operating Environment

All CX300 drives have an area reserved for configuration information. Additionally, the first five disks in the DPE2 have reserved areas that are used for software images and write cache destage, called the *vault* area. To provide space to support these features, user space on the first five drives starts at 6.4 GB and starts on all the other drives at 34 MB, as shown in Figure 9.

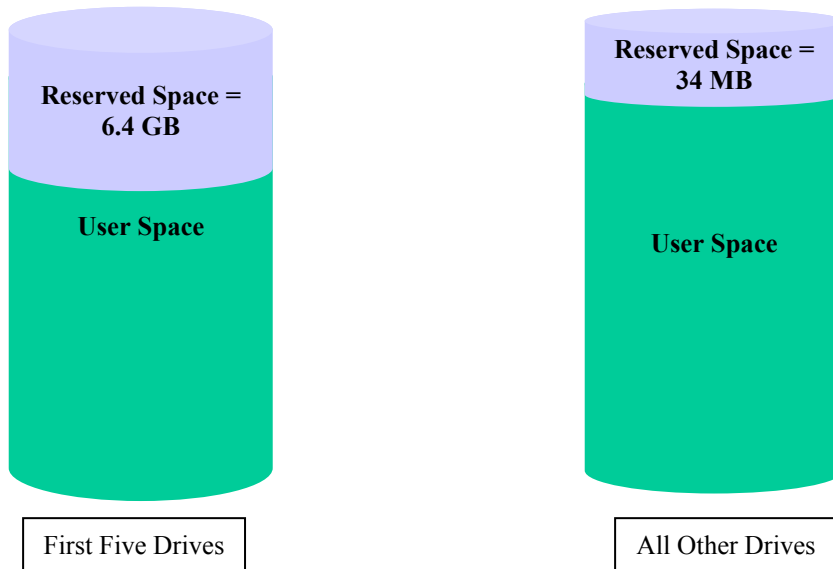


Figure 9. CX300 Disk Layout Showing the Private Space Usage (Picture Not to Scale)

This layout is also common to the CX700 and CX500 storage systems.

Navisphere Management Suites

Customers can manage their CX300 with Navisphere Manager. Navisphere allows support for multiple storage systems on the same subnet. Remote management is supported. Navisphere Manager also supports NAS and/or legacy storage systems, including support for advanced storage applications like SnapView.

MetaLUNs

With Navisphere, the CX300 storage system can utilize the LUN expansion feature, in which the base LUN can be expanded in size to form an entity called the *metaLUN*. The metaLUN is a combination of more than one LUN, joined together by using either the striping or concatenation method. If the operating system supports LUN expansion, with the metaLUN feature a LUN of a larger size will now be visible to the host. For more detail, see Table 3 for the supported metaLUN features on the CX300.

Table 3. MetaLUN Features Supported on the CX300

Feature	CX300 Support
Maximum number of metaLUNs per storage system	256
Maximum LUN size (base or metaLUN)	2 TB
Maximum number of combinations (striping/concatenation) per base LUN	16

For more detail on metaLUNs, refer to the *EMC CLARiiON MetaLUNs, Concepts, Operations, and Management* white paper.

Access Logix (Included in Navisphere Management Suite)

Access Logix includes enabling functions between the storage processors, host bus adapters (HBAs), FC fabric switches, and Navisphere, for heterogeneous host connectivity (refer to the *EMC Support Matrix* for the latest list of supported platforms). Access Logix provides users an easy-to-manage facility of *storage groups* into which they can group their hosts and associated LUNs.

The Core (FLARE) software for the CLARiiON storage system that performs all the functions includes both non-Access Logix and Access Logix code. Users must install the AccessLogixOption on the storage system to enable the Access Logix capability.

Table 4. Access Logix Features Supported on the CX300

Feature	CX300 Support
Max storage groups per storage system	128
Max LUNs per storage group (initiator)	256

SnapView Snapshots (Optional Application)

SnapView allows users to create an instantaneous virtual copy of a LUN. The virtual copy is then accessible from a different host. This enables offloading of the backup processing from the production host. Third-party backup applications integrated with the SnapView APIs or Navisphere CLI can take full advantage of this functionality. SnapView can also be used to run decision-support applications on frozen versions of databases, while the data on the production database remains unaffected. In case of source corruption, contents of the snapshot (point-in-time copy) can be rolled back. See Table 5 for the supported SnapView snapshot features on the CX300.

Table 5. SnapView Snapshots Supported on the CX300

Feature	CX300 Support
Snapshots per source LUN	8
Snapshots per storage system	100
Snapshot sessions per source LUN	8
Snapshot sessions per storage system	25
Reserved LUNs for snapshots per storage system	25

SnapView Version 2.0 supports persistence. Persistence is maintained by storing the frozen point-in-time view of the source LUN on the Save LUN as updates are made to the source LUN by the host. The point-in-time view remains available after trespasses, NDU operations, and SP reboots.

SnapView Clones (Optional Application)

SnapView clones allow users to create a full copy of a LUN, equal in size to the source LUN. Once created, this full copy can then be fractured from the source LUN and be mounted on a different host for processing backup applications without affecting the production data.

A source can have up to eight clones and will be contained within a clone group. A clone group consists of all the clones that are created for a particular source LUN. In case of source corruption, contents of the clone can be reverse-synchronized to the source LUN, thus avoiding downtime of the production database. See

Table 6 for supported clone features on the CX300.

Table 6. SnapView Clones Supported on the CX300

Feature	CX300 Support
Clones per source LUN	8
Clone images per storage system (includes sources in total count)	50
Number of clones in a clone group	8
Number of clone groups (source LUNs) per storage system	25

SAN Copy (Optional Application: Cannot Be Loaded on CX300)

SAN Copy is a generalized copy facility that enables a CLARiiON storage system to copy data within a single storage system or from one storage system to another. The CX300 cannot host the SAN Copy application, but can participate in a SAN Copy session hosted by a supported CLARiiON storage system.

With SAN Copy, the CX300 can participate in the SAN Copy data transfer operation by serving as either a source or a destination. Since the SAN Copy application is not running on the CX300, the software cannot be used to copy LUNs within the CX300. SAN Copy requires that Access Logix software be installed on the storage system.

SAN Copy Version 2.0 allows users to maintain incremental tracking of copies, so that only data that has changed on the source is copied to the destination. As noted above, since the CX300 cannot host the SAN Copy application, it can only serve as a destination system in an incremental SAN Copy session.

PowerPath

EMC PowerPath is host-resident software that works with both CLARiiON and Symmetrix® storage systems to deliver intelligent I/O path management. Using PowerPath, administrators can improve the server's ability to manage heavy storage loads through continuous and intelligent I/O balancing. PowerPath automatically configures multiple paths and dynamically tunes for performance as workload changes. PowerPath adds to the high-availability capabilities of the CLARiiON and Symmetrix storage systems by automatically detecting and recovering from server-to-storage path failures.

In addition to full-featured PowerPath, the CX300 also supports both PowerPath Base. PowerPath Base provides an active/passive path management but only provides failover support with no load balancing or multipathing. If customers want the ability to load balance and /or have the multipathing capability, then they need to upgrade to the full-featured version of PowerPath.

Non-Disruptive Upgrade

The non-disruptive upgrade (NDU) feature of the CX300 is the implementation of a transparent upgrade mechanism within the CX series storage system product line. The NDU feature, when coupled with PowerPath failover software, allows upgrades of the storage-system FLARE software or other software upgrades, without stopping host access to data. During the NDU process, one SP at a time is upgraded and rebooted during the upgrade. While one SP is upgrading and rebooting, PowerPath failover software redirects I/O through the alternate SP as needed.

The NDU function is initiated through either the Navisphere Web-based browser or NaviCLI. NDU upgrades the storage system through several basic steps using the Navisphere management software. Those steps are:

1. Select the storage system to be upgraded through the NDU process.
2. Select the software to be used as the upgrade package.
3. Initiate the upgrade process.

Upgrading to a CX500/CX700 or CX3 Series

As the customer's environment grows, the CX300 storage processor can be upgraded to a CX500, CX700, or CX3 system. Since hardware and cabling must be reconfigured during the upgrade, the storage system is taken offline. However, there is no data loss or data movement required. The upgrade path is one-way: that is, a CX700, CX500, or CX3 cannot be downgraded to a CX300. Older CX storage systems can also be upgraded to current models using a fully supported data-in-place upgrade procedure.

Summary

The CX300 is a full-featured, full 2 Gb data storage system. The CX300 shares with the CX500 and CX700 common hardware components and storage software functions including: management, LUN masking, and path failover. This storage system has four front-end FC ports and one back-end FC port, up to 1 GB of cache memory per storage system, and an Ethernet port for managing the storage system. The CX300 also can be upgraded to a CX500, CX700, or CX3 system.. Functional improvements in the CX300 include:

- Enhanced product density with a fully functional 15-drive storage system in only 4U of rack space
- Enhanced path failover capabilities with the implementation of PowerPath
- Flexibility to choose between high-performance Fibre Channel DAE2s and high-capacity ATA DAE2s
- Investment protection with data-in-place upgrades to CLARiiON family members (CX500, CX700, or CX3 systems)

Appendix

The following information explains the effect on write cache availability for various component failures on the CX300.

Table 7. Component Failure and Write Caching on the CX300 (Single SPS)

Component Failure and Write Caching	
1st SPS (battery)	Write Disabled
2nd SPS (battery)	N/A
1st Vault Drive	Write Enabled (Selectable)
User induced/NDU	Write Disabled
Fan [*]	Write Enabled
Power Supply	Write Disabled
Storage Processor	Write Disabled
A/C Power	Write Disabled

Table 8. Component Failure and Write Caching on the CX300 (Dual SPSs)

Component Failure and Write Caching	
1st SPS (battery)	Write Enabled
2nd SPS (battery)	Write Disabled
1st Vault Drive	Write Enabled (Selectable)
User induced/NDU	Write Disabled
Fan [*]	Write Enabled
Power Supply	Write Disabled
Storage Processor	Write Disabled
A/C Power	Write Disabled

*** Single fan failure note:**

There are two fans in each of two DPE fan modules, for a total of four fans. Write cache is disabled only after both fans in a single fan module fail, two fans out of the four total available fail, or a complete single fan module is pulled. If a single fan in the fan module fails, the other existing fan will speed up to maintain proper air flow through the SPE. Once the fan module is pulled for replacement, two minutes are allowed before the array shuts down for thermal reasons. These two minutes allows for fan module replacement.