

Upgrading to Dell | EMC CX4 Series Storage Systems

Applied Technology

Abstract

This white paper presents the presales considerations for upgrading to Dell/EMC CX4 series storage systems. It explores the choice of buying a new system versus leveraging existing assets by performing a data-in-place upgrade.

Dell Inc.
April 2009

Visit www.Dell.com/emc for more information on Dell/EMC Storage.

Copyright © 2009 Dell Inc. THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.



Table of Contents

Executive Summary	3
Introduction	3
Audience.....	3
Terminology.....	3
Upgrading vs. Buying New	4
Data-in-Place Upgrades	4
Planning for the Upgrade	5
Downtime.....	5
CAP2 Conversion Readiness Check.....	6
Adequate FLARE revision and storage system health.....	6
Disk compatibility with the new FLARE revision	6
Reserved space expansion	6
Physical enclosure relocation.....	7
Power Requirements.....	7
Mixing 2 Gb/s and 4 Gb/s DAE Technologies	8
Cabling and Connectors.....	8
Host Software Upgrades	9
Special Case Upgrade Paths	9
Special Considerations for Replication Software	9
Private LUNs.....	9
MirrorView.....	10
Upgrading with Celerra Gateways Attached	10
Backup.....	12
Performing the Upgrade	12
System Scaling after the Upgrade	12
Conclusion.....	12
Appendix A: Example CX4 Conversion Readiness Check	13
Appendix B: Virtual LUN Migration Performance.....	16

Executive Summary

Customers who purchase IT solutions are concerned about the return on investment (ROI) of their purchases. Creating solutions that offer customers the ability to extend the usable life of an investment provides measurable financial benefits.

The Dell/EMC CX4 series storage systems offer a data-in-place upgrade capability: As users' performance, capacity, or functionality requirements expand, the customer can upgrade from a CX or CX3 series storage system to a CX4 series storage system with minimum impact to existing business processes. Because of the common hardware architecture, upgrades can occur while keeping the data in place. This white paper discusses the reasons for performing this hardware upgrade and the planning considerations.

Introduction

Dell/EMC CX4 series storage systems offer many advantages over the CX and CX3 series storage systems. For an overview of what the CX4 series has to offer, see [*Introduction to the Dell/EMC CX4 Series Featuring UltraFlex Technology*](http://www.dell.com/emc) at www.dell.com/emc. When deciding to invest in CX4 storage systems, a major consideration is whether to perform a data-in-place upgrade for an existing system or buy a new CX4 system and leverage the existing hardware elsewhere. This white paper discusses the advantages and disadvantages of both solutions and discusses the considerations for performing a data-in-place upgrade.

Conversion kits permit the user to expand capacity, improve performance, and increase port count within the Dell/EMC family of products. Current upgrade paths include support for converting an existing CX300/500/700 or CX3 series into the new CX4 series.

If any of the presales considerations require significant work to prepare for the upgrade, Dell professional services can be engaged to complete the necessary preparation work.

Audience

This white paper is intended for Dell customers, Dell system engineers, and members of the Dell professional services community who are interested in learning more about a data-in-place upgrade to a CX4 series storage system. It is the responsibility of the Dell presales technical resource to review these considerations when discussing an upgrade with the customer.

The upgrade procedure is performed by a Dell service professional.

Terminology

Disk-Array Enclosure (DAE) – The generic name for a 15-disk enclosure.

DAE2 – A 2 Gb/s Fibre Channel disk enclosure.

DAE2-ATA – A 2 Gb/s ATA disk enclosure.

DAE2P – A 2 Gb/s point-to-point Fibre Channel disk enclosure. DAE2, DAE2-ATA, and DAE2P disk enclosures may all be present on CX series storage systems.

DAE4P – A 4 Gb/s point-to-point Fibre Channel disk enclosure. This is the standard disk enclosure for CX3 UltraScale series storage systems, as well as CX4 storage systems. The only difference between the two is the bezel on the front of the enclosure.

DAE-OS – The first DAE in the system, containing the Dell/EMC operating environment (FLARE[®]) on the first five drives.

Standby Power Supply (SPS) – A battery that allows a clean shutdown during power failures, so that no data is lost.

Storage Processor Enclosure (SPE) – An enclosure that holds dual storage processors, power and cooling modules.

Disk Processor Enclosure (DPE) – An enclosure that holds one or two storage processors, power and cooling modules, and up to 15 disks. Specific to the CX300, CX500, AX4 storage systems.

Upgrading vs. Buying New

The main reason to upgrade an existing system is to enhance the existing investment in disks, disk enclosures, and cabinets by adding capacity, performance, and functionality. All existing equipment remains in place and only the storage processors (SPs), and potentially the standby power supplies (SPS), are exchanged. Another major benefit is that all data and configuration remain on the system (“data-in-place”). This eliminates any risks that may be incurred by requiring a data migration, and saves the time that would be required to configure a new system. When the upgrade is completed and the new system is booted, all data is available and all hosts have access to their LUNs.

One consideration for upgrading from a CX (or CX3 that has been upgraded from a CX) system is that existing back-end loops are always 2 Gb/s. This may not be a concern if the storage system is not servicing high-bandwidth applications and 2 Gb/s back-end connectivity is not the limiting factor. For example, performance in a small-block random I/O environment (most databases) does not differ greatly between systems with 2 Gb/s or 4 Gb/s back-end loops. If upgrading to a system with additional back-end loops (for example, upgrading from a CX500 with two pairs of back-end loops to a CX4-480 with four pairs of back-end loops), the new back-end loops may be run at 4 Gb/s while the existing back-end loops remain at 2 Gb/s.

For customers who want to invest in a full 4 Gb/s storage system, purchasing a new CX4 series system is a better choice than upgrading. This gives customers the performance benefits of an end-to-end 4 Gb/s storage system and the ability to repurpose their existing storage investment in another application. There are many solutions for data migration to a new system that can minimize application downtime.

To summarize, customers should perform a data-in-place upgrade if their objectives are to:

- Add capacity, performance, or functionality to their existing hardware.
- Leverage their existing hardware investment.
- Avoid configuring a new storage system.
- Avoid data migration to a new storage system.

Customers should buy a new CX4 series storage system if their objectives are to:

- Invest in an end-to-end 4 Gb/s storage system.
- Redeploy the existing system in another application.

Data-in-Place Upgrades

Dell/EMC storage systems offer the investment protection of data-in-place upgrades. These upgrades allow customers to preserve their existing disks, DAEs, and cabinets, as well as data and software configurations. A customer who has purchased a CX or CX3 series storage system can upgrade to a CX4 series system to add performance, functionality, and scalability. Table 1 shows the valid upgrade paths at the time of this paper's publication.

A data-in-place upgrade typically consists of a storage processor conversion and a software conversion/upgrade. Upon completion of the upgrade, the new system has the same identity as the old system and all data and the configuration are preserved. No host or SAN reconfiguration is necessary, but there may be suggested software updates to Dell PowerPath®, Navisphere® Agent, and Navisphere CLI to operate with the new FLARE revision.

Table 1. Currently supported conversions to CX4 series storage systems

Convert From	Convert To			
	CX4-120	CX4-240	CX4-480	CX4-960
CX300	Yes	Yes	Yes	Yes
CX3-10c	Yes	Conditionally	Conditionally	Conditionally
CX500		Yes	Yes	Yes
CX3-20		Yes	Yes	Yes
CX4-120		Conditionally	Conditionally	Conditionally
CX700			Yes	Yes
CX3-40			Yes	Yes
CX4-240			Yes	Yes
CX4-480				Yes

Conditionally supported upgrade paths must meet certain prerequisites. See the “Special case upgrade paths” section for more information.

Unsupported upgrades

In addition to the upgrade paths absent from Table 1, the following upgrades are not supported:

- A CX3-10c using SATA for the OS drives may only be upgraded to a CX4-120. There is no procedure to swap the SATA drives for the Fibre Channel drives that are required in the larger systems.
- A CX4-120 using SATA for the OS drives may not be upgraded. There is no procedure to swap the SATA drives for the Fibre Channel drives that are required in the larger systems.
- An AX model array may not be upgraded into the CX family.

Planning for the Upgrade

If deciding to upgrade, there are a number of planning considerations that should be thought out before the upgrade is purchased. A high-level data-in-place procedure is outlined in Figure 1.




Figure 1. Data-in-place upgrade flow chart

The detailed procedure for performing the upgrade is found in the Dell Professional Services’ Procedure Generator. The Dell service professional that performs the upgrade follows this procedure.

The following sections list the technical considerations that should be taken into account before the upgrade takes place.

Downtime

Because physical hardware is swapped during the conversion, the storage system must be offline during the conversion. Some of the major steps that are taken during this time are:

1. Stop all I/O to the system and power down.
2. Convert storage processors to the new model and swap SP cables as necessary.
3. Boot the array to **Utility Partition** – A recovery mode where software changes can be made.
4. Convert databases – All FLARE databases are converted to a new format.
5. Upgrade FLARE – FLARE is upgraded to the latest level that supports the new hardware.
6. Boot the system.
7. Reconnect the hosts – The conversion does not break existing zoning or LUN mappings.

It is realistic to expect 8-12 hours of downtime for the upgrade (3-6 hours for in-family upgrades). This assumes that all considerations in this paper have been addressed before starting the upgrade. If any other steps are required during the conversion (such as power reconfiguration or re-racking of components), then additional time must be planned when scheduling and utilizing the outage window.

CAP2 Conversion Readiness Check

The CAP2 tool includes a conversion readiness check (CRC). The CAP2 CRC automatically checks to see if an existing CX or CX3 series system is ready to be converted to a CX4 series system. CRC makes sure that no LUNs are in conflict with the expanded FLARE space, port usage is not in conflict, and array limits will not exceed those of the new platform. If LUNs are in conflict with the expanded FLARE space, consult the “Reserved space expansion” section below.

You should check the overall condition of the storage system using the (CAP2) tool to ensure that there are no trespassed LUNs, hardware failures, cabling issues, and so forth, as these may complicate the upgrade. Resolve any issues before attempting the upgrade.

The CAP2 CRC is *not* a complete check of all considerations; it is the responsibility of the presales account team to read this paper and identify what actions are required prior to an upgrade. For a sample of a CRC report, please see “Appendix A: Example CX4 Conversion Readiness Check” at the end of this whitepaper.

Adequate FLARE revision and storage system health

Specific changes have been made in the FLARE operating environment to enable the conversions; therefore, the FLARE revision may need to be updated before the conversion can be performed. Reference the Procedure Generator for the minimum release and patch level required for each hardware platform.

In addition to the FLARE minimum revision requirement, please ensure that all applicable field change orders (FCOs) have been carried out prior to the system upgrade. Also, it is important that any open EMC® Technical Advisories (ETAs) that are applicable to the storage system are addressed and resolved before the data in place upgrade is attempted.

Disk compatibility with the new FLARE revision

The majority of data-in-place upgrades, especially an upgrade to a CX4 from a CX or CX3 array, will result in a newer revision of FLARE running on the upgraded system. It is important to verify before the upgrade that all disks in the system are compatible with the new model and FLARE revision. The document *CX4 Series Storage Systems Disk and FLARE OE Matrix* on EMC Powerlink lists the minimum FLARE revision required for each drive.

The existing disks can be easily inventoried for this purpose by using the following Navisphere Secure CLI command:

```
C:\>naviseccli -h <SP IP address> getdisk -tla
Bus 0 Enclosure 0 Disk 0
Dell/EMC TLA Part Number:005048646
```

The output of this example is truncated, but the getdisk command will return information for all disks in the system. The part numbers returned can be matched up with those in the *CX4 Series Storage Systems Disk and FLARE OE Matrix*.

Reserved space expansion

With the CX4 series, the space reserved by FLARE on the first five drives has grown to approximately 62 GB per drive. This allows for new features, expansion of existing features, and growth in the future. Therefore, if the upgrade is from a CX or CX3 series storage system with LUNs bound on the first five drives, those LUNs are most likely located where FLARE needs to expand. Furthermore, if the system drives are 36 GB drives, you must replace them with 146 GB (or larger) drives before the conversion.

A new CX4 conversion readiness check report has been added to the CAP2 tool (see the sample in “Appendix A: Example CX4 Conversion Readiness Check”) to automatically check LUN placement. The output is a listing of all LUNs that would conflict with the new FLARE reserved space. Dell recommends using virtual LUN technology (LUN migration) to relocate these LUNs. There are some special LUNs used by replication software that may not be moved using virtual LUN technology; these are known as *private* LUNs. If you are using replication software, please see the “Special Considerations for Replication Software” section.

It is the responsibility of the presales account team to identify if LUN migration will be required. The CAP2 tool can be used to identify LUNs that need to be migrated. Use “Appendix B: Virtual LUN Migration Performance” as a guideline to estimate how long LUN migrations will take.

All conflicting LUNs identified by CAP2 must be relocated *before* the upgrade. For metaLUNs with at least one component residing in the new FLARE reserved area, the entire metaLUN must be relocated. Even if CAP2 identifies free space on the first five drives, there is no guarantee that newly bound LUNs will not conflict with the new reserved area. Therefore, all LUNs identified by CAP2 should be migrated off of the first five drives.

Physical enclosure relocation

CX4 series systems utilize SPE and SPS models that may have different rack heights than previous systems. Depending on how the current system is racked, enclosures may need to be re-racked during the upgrade. This should be planned prior to the upgrade, especially since two people may be required to lift some of the enclosures. Use the NEMA rack heights provided in Table 2 as a guideline to determine whether or not enclosures require relocation.

Table 2. Dell/EMC SPE + SPS + DAE-OS rack heights

Convert From	Convert To			
	CX4-120 (6U)	CX4-240 (6U)	CX4-480 (6U)	CX4-960 (9U)
CX300 (4U)	+ 2U	+ 2U	+ 2U	+ 5U
CX3-10c (5U)	+ 1U	+ 1U	+ 1U	+ 4U
CX500 (4U)		+ 2U	+ 2U	+ 5U
CX3-20 (5U)		+ 1U	+ 1U	+ 4U
CX4-120 (6U)		Same	Same	+ 3U
CX700 (8U)			- 2U	+ 1U
CX3-40 (5U)			+ 1U	+ 4U
CX4-240 (6U)			Same	+ 3U
CX4-480 (6U)				+ 3U

As an example, a CX3-40 SPE, SPS, and the first tray of disks (DAE-OS) require 5U of rack height. If upgrading to a CX4-960, the same components require 9U of rack height. If this space is not available, Dell recommends that you move the DAEs up in the rack to accommodate the new enclosures.

As an alternative to re-racking, the new SPE and SPS may be located either at the top of the cabinet or in an adjacent cabinet. They must be located together because of the short length of the SPS signal cable. Back-end cables are currently available in lengths up to 5 meters, so the SPE must be within that distance of the DAE-OS. Take care when creating a non-standard rack configuration, as this may cause confusion, which could lead to human error in future servicing of the system.

Existing equipment is utilized as much as possible in the conversion. For example, if upgrading from a CX300 or CX500, the SPs in the DPE are replaced with link control cards (LCCs), and this enclosure remains the DAE-OS for the upgraded system. When upgrading a CX4-120 to a CX4-240 or CX4-480, the SP CPU module is replaced in the original chassis.

Power Requirements

When performing an upgrade, the new system will most likely draw more power than the current system. Use the EMC Power Calculator (on EMC Powerlink) or the Dell Data Center Capacity Planner (Dell’s DCCP at www.dell.com/calc) to determine what the new power requirements are. It is important to note the power distribution limitations of the current cabinet. For example, a Dell 40U rack requires two 30-Amp feeds. For redundancy, each 30-Amp feed must be able to power all enclosures in the rack. For compliance with North American wiring requirements, the feeds must be loaded at no more than 80 percent of their rating. Therefore, the enclosures placed in this type of rack must not exceed 24 Amps of current draw. Ensure that the new configuration will not exceed this number if it will be in this type of rack.

To summarize:

-
- For a Dell 40U cabinet, check the power requirements for all equipment to be located in the rack and ensure that the total current draw does not exceed the total available current.
 - If the Dell/EMC array is racked in a custom enclosure, check the power distribution limitations of that enclosure.
 - If the power requirements for the new equipment exceed the capacity of the current rack, equipment must be moved to a new rack that can support the power requirements, or an additional rack must be provided to house some of the enclosures.

Mixing 2 Gb/s and 4 Gb/s DAE Technologies

Dell/EMC CX series storage systems (CX200 through CX700) were designed around 2 Gb/s Fibre Channel technologies. All enclosures and drives on the existing CX system will be supported on the new system after the upgrade. Existing 2 Gb/s enclosures will remain and 4 Gb/s enclosures can be added. When added to a 2 Gb/s back-end loop, 4 Gb/s enclosures and all disks within them will run at 2 Gb/s. If upgrading to a system with more back-end loops, 4 Gb/s enclosures with 4 Gb/s disk drives can be connected to the new loops so that the disks and loops can run at 4 Gb/s. No enclosures or disks should be added during the upgrade, but once the new system is running new enclosures and disks may be added.

When planning for the upgrade, consider whether 4 Gb/s back-end loop bandwidth is necessary and plan accordingly. It may be better to purchase a new CX4 series system if the system is intended for high-bandwidth applications and all back-end loops need to be running at 4 Gb/s to meet performance requirements.

NOTE: Resetting a loop from 4 Gb/s to 2 Gb/s does not occur during a array conversion; however, it could be an issue whenever a 2 Gb/s DAE is migrated to a different array (no data is migrated, just empty capacity). If a 2 Gb/s DAE is attached to a 4 Gb/s back-end loop, the 2 Gb/s DAE will be ignored and the loop will continue to operate 4 Gb/s. If the array is rebooted, then the loop, as well as all other 4 Gb/s DAEs and drives on the loop, will be reset to 2 Gb/s.

To summarize:

- Existing disks and enclosures will continue to operate at 2 Gb/s. 4 Gb/s disks and enclosures added to these back-end loops will operate at 2 Gb/s.
- If upgrading to a system with more back-end loops, the new loops can operate at 4 Gb/s.
- For an end-to-end 4 Gb/s solution, consider purchasing a new CX4 series system.

Cabling and Connectors

When upgrading from a CX series storage system, some new back-end cables may be needed for the upgrade. The CX series systems with DAE2 and DAE2-ATA enclosures use HSSDC connections. All back-end cables are HSSDC-to-HSSDC. DAE2P enclosures offered for the CX300/500/700 use HSSDC2 connections. These are connected to existing equipment using HSSDC2-to-HSSDC cables. The CX3 and CX4 series storage processors have SFP connectors for the back-end loops. They connect to DAE3P and DAE4P enclosures using SFP-to-HSSDC2 cables. With the addition of CX4 data-in-place upgrades, SFP-to-HSSDC cables are available to connect CX4 series SPs to pre-existing DAE2 and DAE2-ATA enclosures. Use Table 3 as a quick reference for the type of connector used by each enclosure.

Table 3. Cable connector types used by CX enclosures

Connector type	Connects To
HSSDC	DAE2, DAE2-ATA, CX300 through CX700 SP
HSSDC2	DAE2P, DAE3P, DAE4P
SFP	CX3 and CX4 series SP

Host Software Upgrades

Before performing the hardware upgrade, software updates to host software may be required for compatibility with the new system. These updates may include PowerPath, Navisphere Host Agent, Navisphere Server Utility, Navisphere CLI, admsnap, and others. For a listing of revisions of Navisphere software supported with each FLARE revision, see the “Navisphere management software requirements” table in the EMC *Navisphere Host Agent/CLI and Utilities Release Notes* on EMC Powerlink. For recommended revisions of PowerPath and admsnap refer to E-Lab™ Navigator on EMC Powerlink.

Special Case Upgrade Paths

Some upgrade paths require special considerations. It is important to obtain the latest version of CAP2, which contains pre-upgrade checks for many of these situations.

CX3-10c and CX4-120 Upgrades

- CX3-10c systems that use SATA for the OS drives may only be upgraded to a CX4-120 with SATA OS drives. CX4-120 systems that use SATA for the OS drives may not be upgraded. There is **NOT** a procedure to swap the SATA drives for the Fibre Channel drives that are required in the larger systems. The CAP2 CRC check will verify the OS drive type.
- Low-cost HBAs in any attached servers (for example, Emulex LP-101) that are supported by the CX3-10c and CX4-120, but are not supported by the rest of the Dell/EMC family, must be swapped out prior to conversion. Use E-Lab Navigator to validate the new configuration.
- The version of PowerPath bundled with the CX3-10c and CX4-120 is not supported by any other models in the Dell/EMC family. PowerPath licenses must be purchased prior to the upgrade to the new storage system.

CX3-20f to CX4 Upgrades

The CX3-20f system has six FC front-end ports per SP while the base CX4-240 has only two FC front-end ports per SP. There cannot be a loss of port count during a conversion. Therefore, an additional FC I/O module is required on each SP. A FC UltraFlex I/O kit, consisting of a pair of FC I/O modules, must be purchased separately from array upgrade kit.

Special Considerations for Replication Software

Private LUNs

Private LUNs used by EMC replication software cannot be migrated using LUN migration. These include clone private LUNs (CPLs) used by SnapView™ clones, write intent log (WIL) LUNs optionally used by MirrorView™/Synchronous mirrors, and reserved LUNs used by SnapView snapshots, SAN Copy™ incremental sessions, and MirrorView/Asynchronous mirrors. If CAP2 identifies these as LUNs requiring relocation, they need to be moved elsewhere, which may involve some disruption and reconfiguration of the replication software that uses them. CPL and WIL LUNs can be moved with no disruption to clones or synchronous mirrors. Moving LUNs from the reserved LUN pool requires destruction of all SnapView sessions, MirrorView/Asynchronous mirrors, and SAN Copy incremental sessions using those reserved LUNs. Full resynchronizations are required for SAN Copy incremental sessions and MirrorView/Asynchronous mirrors. For information about migrating private LUNs, review the specialized procedure documented in the CPG.

MirrorView

When MirrorView/Synchronous or MirrorView/Asynchronous is installed, the port used by MirrorView will not change. The logical port numbering will be carried through the conversion. This eliminates the need for a zoning change to connect the MirrorView ports. Table 4 lists the iSCSI and FC MirrorView ports in use with previous models.

Table 4. Port usage and MV ports by model

Array	Port numbers and type per SP					
	0	1	2	3	4	5
CX300	FC	FC				
CX300i	iSCSI	iSCSI				
CX3-10c	iSCSI	iSCSI/MV	FC	FC/MV		
CX3-20	FC	FC/MV				
CX3-20c	iSCSI	iSCSI	iSCSI	iSCSI/MV	FC	FC/MV
CX3-20f	FC	FC/MV	FC	FC	FC	FC
CX500	FC	FC/MV				
CX500i	iSCSI	iSCSI				
CX3-40	FC	FC/MV				
CX3-40c	iSCSI	iSCSI	iSCSI	iSCSI/MV	FC	FC/MV
CX3-40f	FC	FC/MV	FC	FC		
CX700	FC	FC	FC	FC/MV		
CX3-80	FC	FC	FC	FC/MV		

Legend:

FC = Fibre Channel port, iSCSI = iSCSI port, /MV = port used by MirrorView

All mirrors need to be administratively fractured and all SAN Copy sessions stopped before the upgrade. SnapView snapshot sessions and clone groups do not need to be altered. All persistent replication sessions and checkpoints will be preserved through the upgrade. This means, for example, that a SAN Copy incremental session does not require another initial synchronization following the upgrade.

If upgrading from a CX300 with SAN Copy/E installed, the SAN Copy/E enabler must be removed before the upgrade or incompatibility messages will be displayed after the upgrade. All CX4 series platforms support full SAN Copy.

Upgrading with Celerra® Gateways Attached

Although this is unlikely to be the case, Dell/EMC array upgrades are supported when attached to Dell/EMC N500G (Celerra) NAS gateways. This is the only Celerra NAS gateway that Dell has resold up to this time. Upgrades are supported for gateway systems only and not integrated NAS systems. Be sure to indicate that a NAS gateway is attached when generating the conversion procedure in the CPG.

The N500G uses standard Dell/EMC LUNs for both control data and user data. Celerra control LUNs are usually placed on the first five Dell/EMC drives, which may conflict with the reserved space expansion required when upgrading from a CX or CX3 series to a CX4 series. The CAP2 conversion readiness check report will report which LUNs need to be moved and Celerra CLI tools may be used to determine which of the LUNs identified for migration are NAS LUNs.

It is advisable that your trained Dell technician carry out the following activities.

The following NAS command initiated on the connected NAS device(s) will list the Dell/EMC disk volumes owned by an individual Celerra:


```
[nasadmin@eng25214 bin]$ nas_disk -list
id   inuse  sizeMB  storageID-devID  type  name           servers
1    y      11263   APM00063304671-0000  CLSTD  root_disk      1,2
2    y      11263   APM00063304671-0001  CLSTD  root_ldisk     1,2
3    y      2047    APM00063304671-0002  CLSTD  d3             1,2
4    y      2047    APM00063304671-0003  CLSTD  d4             1,2
5    y      2047    APM00063304671-0004  CLSTD  d5             1,2
6    y      2047    APM00063304671-0005  CLSTD  d6             1,2
7    y      521539  APM00063304671-0010  CLSTD  d7             1,2
8    y      521539  APM00063304671-0011  CLSTD  d8             1,2
15   n      1099246  APM00063304671-0012  CLSTD  d15
16   n      1099246  APM00063304671-0013  CLSTD  d16
17   n      549623  APM00063304671-0014  CLSTD  d17
18   n      1099246  APM00063304671-0016  CLSTD  d18
19   n      549623  APM00063304671-0015  CLSTD  d19
20   n      1099246  APM00063304671-0017  CLSTD  d20
```

The storageID-devID information must be used to determine if the LUNs to be migrated (as reported by the CAP2 report) are NAS LUNs.

The NAS control LUNs are reserved LUNs for the NAS operating system and operational database located on the back-end storage. The Celerra control LUNs currently use approximately 30 GB of Dell/EMC storage. A typical control LUNs configuration is as follows:

- Two 11 GB LUNs
- Four 2 GB LUNs

Celerra control LUNs for current NAS implementations always have IDs 1-6 in the output of the “nas_disk -list” command. The remaining IDs in the listing above are the data LUNs associated with the Celerra. After the control and data LUNs have been identified, the EMC identifier and ALU can be used to map the NAS LUNs back to the Dell/EMC array. The ALU directly corresponds to the LUN# column in the CAP2 CX4 conversion readiness check report.



If NAS LUNs are identified for migration the following considerations apply:

- LUN migration *must* be used to execute migration of NAS LUNs. The recommended priority for the LUN migration is high. This provides the best balance between application response time and the time for LUN migration to complete. Refer to “Appendix B: Virtual LUN Migration Performance” on page 16 for LUN migration rates.
- NAS control LUNs *must* be migrated to Fibre Channel drives
- NAS control LUNs *must* be moved to a single RAID group.
- NAS control and data LUNs *must* be migrated to a RAID group of the same geometry. For example, LUNs on a RAID 5 4+1 *must* be migrated to another RAID 5 4+1. Refer to EMC knowledgebase article emc160173 for additional information.
- NAS control and data LUNs *must* be migrated to LUNs with identical block counts. Refer to EMC knowledgebase article emc157859 for additional information.

-
- NAS data LUNs *should* be moved to a RAID group not already containing NAS LUNs to avoid performance conflicts with other defined NAS data LUNs.
 - NAS data LUNs *should* be migrated to target disk technology compatible with the source to avoid performance issues. For example, LUNs currently on Fibre Channel disk should be migrated to Fibre Channel disk.
 - The NAS control LUNs *should* be migrated back to the original RAID groups if there is sufficient space to accommodate NAS control LUNs after the CX, or CX3, to CX4 conversion.

Backup

Before the upgrade, a full backup of all data on the storage system is recommended. This includes retrieving SP collects from each SP, which CAP2 does when running the conversion readiness check.

Performing the Upgrade

A Dell service professional or service partner performs the data-in-place upgrade by following the Procedure Generator. Dell technical support should be contacted if there are any issues during the upgrade.

System Scaling after the Upgrade

Following an upgrade, it is possible to add additional DAEs and disks to existing back-end loops or to new loops if they are available. Disks added to existing 2 Gb/s back-end loops operate at 2 Gb/s. Disks added to new back-end loops may run at 4 Gb/s if the disk is capable of running at this speed. Once disks are added, data can be reorganized within the Dell/EMC by using metaLUNs or LUN migration to spread existing data to new drives.

Depending on the upgrade path, the new system may have more fibre and/or iSCSI ports than the previous system. These ports may be used to expand the storage system's connectivity to new SANs or direct-attached hosts.

Customers should check the current memory allocation to read and write cache and increase the allocations if possible on the new system.

The number of LUNs, number of connected hosts, or allowed number of replication sessions may also be increased. For these limits, see the *Open Systems Configuration Guide*.

Conclusion

Dell continues to bring Dell/EMC customers extensive investment protection and deployment flexibility. By providing a clear upgrade path of features, functionality, and performance—on both a software and a hardware level—the Dell/EMC product line continues to meet customer demands for *pay-as-you-grow* features and functionality. Customers who start with a CX or CX3 series storage system can upgrade to a CX4 series storage system, while keeping their data in place. Within the CX4 series, customers can start small and upgrade to a larger system for increased performance, functionality, and scalability. With consistent hardware and software architectures, customers can keep their business process in place while their applications and processes grow.

Appendix A: Example CX4 Conversion Readiness Check

The report below is an example of what CAP2 includes in the CX4 conversion readiness check report. In this example, CAP2 has identified one LUN that must be relocated from the first five drives. CAP2 identifies each LUN as either a user LUN, metaLUN, write intent log (WIL), clone private LUN (CPL), or reserved LUN (RL). These are identified in the usage column to the far right of the migration table. MetaLUNs are shown in a separate table. User LUNs and metaLUNs may be relocated using virtual LUN technology, while WIL, CPL, and RL may require other procedures as described in "Special Considerations for Replication Software." CAP2 also identifies free space to which these LUNs could potentially be migrated.

For Dell employees, see the *Release 28 Conversion Training (IMPACT) module* for more examples of CRC checks with special cases.

```

=====
Conversion Lun Check report: FAIL
-----
Summary of Luns/Metaluns which need to be migrated:
  Count      Lun Type/Description      Capacity (MB)
  -----
  0           User Metalun              0
  0           User Lun                  0
  0           Write Intent Log (WIL)    0
  1           Clone Private Lun (CPL)   128
  0           Reserved Lun (RL)         0
  -----
  1 (total)                                128 (total)

Summary of free disk space:
  Disk Count  Description                Disk Size (MB)
  -----
  22          Unbound Disk              2874236
  0           Empty RAID Group         0
  -----
  22 (total)                                2874236 (total)

Summary of free available RAID group capacity:
  Description                Capacity (MB)
  -----
  Inactive Hot Spare        273776
  System RAID groups        137784
  Non-System RAID groups    2030719
  -----
  2442279 (total)

Migration Table
Lun(s) which need to be migrated:

  Storage Group  Lun#  Lun Name  Capacity (MB)  RAID Type  RAID Group  Count of Disks  Size on Disk (MB)  Usage
  -----
  N/A            2046  LUN 2046  128            RAID5      0           5               160                SP-A CPL
  -----
  128 (total)                                5 (total)  160 (total)

Metalun(s) which need to be migrated:
None

Free Space Table
Available Disk(s):

  Bus  Enc  Slot  Drive#  Disk Size (MB)  State  RAID Group ID  RAID Type  Disk Type
  ---  ---  ---  ---     -----

```

0	5	14	89	68238	Unbound Disk	N/A	N/A	Unknown
0	6	14	104	136888	Unbound Disk	N/A	N/A	Unknown
0	7	0	105	136888	Unbound Disk	N/A	N/A	Unknown
0	7	1	106	136888	Unbound Disk	N/A	N/A	Unknown
0	7	2	107	136888	Unbound Disk	N/A	N/A	Unknown
0	7	3	108	136888	Unbound Disk	N/A	N/A	Unknown
0	7	4	109	136888	Unbound Disk	N/A	N/A	Unknown
1	4	14	194	68238	Unbound Disk	N/A	N/A	Unknown
1	6	0	210	136888	Unbound Disk	N/A	N/A	Unknown
1	6	1	211	136888	Unbound Disk	N/A	N/A	Unknown
1	6	2	212	136888	Unbound Disk	N/A	N/A	Unknown
1	6	3	213	136888	Unbound Disk	N/A	N/A	Unknown
1	6	5	215	136888	Unbound Disk	N/A	N/A	Unknown
1	6	6	216	136888	Unbound Disk	N/A	N/A	Unknown
1	6	7	217	136888	Unbound Disk	N/A	N/A	Unknown
1	6	8	218	136888	Unbound Disk	N/A	N/A	Unknown
1	6	9	219	136888	Unbound Disk	N/A	N/A	Unknown
1	6	10	220	136888	Unbound Disk	N/A	N/A	Unknown
1	6	11	221	136888	Unbound Disk	N/A	N/A	Unknown
1	6	12	222	136888	Unbound Disk	N/A	N/A	Unknown
1	6	13	223	136888	Unbound Disk	N/A	N/A	Unknown
1	6	14	224	136888	Unbound Disk	N/A	N/A	Unknown
-----				2874236(total)	-----			

Free capacity in System RAID Group(s):

RAID Type	Index	Free Capacity (MB)	Current Luns
RAID5	0	137784	2046
		137784(total)	

Free capacity in Non-System RAID Group(s):

RAID Type	Index	Free Capacity (MB)	Free Contiguous Capacity (MB)	Current Luns
RAID5	1	6448	6448	110 74 2047
RAID5	2	27056	27056	99 100 92 4 112 70
RAID5	6	12529	12529	13 16 66
RAID5	32	85595	65115	140 152 153 154 159
RAID5	10	63920	40960	41 150 151
RAID5	11	4128	4128	42 64 65 22 23 103 108 109
RAID5	12	3504	2480	43 57 58 27 104 44
RAID5	16	1056	1056	47 0 2 101
RAID5	17	65568	61440	67 51 52 61 5 84
RAID5	18	83547	83547	15 36 49 1
RAID5	20	38431	28191	96 32 33 18 24 30 29
RAID5	30	0	0	12
RAID5	25	119216	119216	82
RAID1	26	43628	43628	83
RAID5	27	648	648	86
RAID5	31	393819	393819	48
RAID5	33	187995	187995	141 155 156
RAID5	34	345691	345691	143 144 145 147 157 158
RAID5	3	37296	16816	89 90 91 93 94 6 7 75
RAID5	4	57776	57776	68 9
RAID5	5	25008	20480	10 11 8 88 97 98
RAID5	7	12320	12320	37 95 106 107 148 149
RAID5	8	432	432	38 62 63 28
RAID5	9	3504	2480	39 55 56 105 111 40
RAID5	13	2480	2480	45 77 59 69 73 102
RAID5	14	14768	10240	46 3 53 54 60
RAID5	15	19888	14768	50 76 21
RAID5	19	1567	1567	71 72 19 20

RAID5	29	47346	47346	14 25 34 35 31 26
RAID5	21	119216	119216	78
RAID5	22	119216	119216	79
RAID1	23	42604	42604	80 87
RAID1	24	43628	43628	81
RAID5	28	891	891	85
-----		2030719 (total)	1936207 (total)	-----

=====
Zero Mark Check report: WARNING

Cannot generate Zero Mark Check report due to Array not online

=====
Start Bundle Revision Check report: FAIL

The current bundle 2.19.600.5.019 must be upgraded to at least 02.19.600.5.034 before doing the conversion.

Appendix B: Virtual LUN Migration Performance

The time to complete a LUN migration operation varies depending on the user selected migration rate. Low, medium, and high rates are mostly constant, while ASAP depends on the type of disks involved, RAID layout, Dell/EMC model, available write cache, and other activities in the storage system. The test results in Table 5 can be used as a guideline to estimate how long a LUN migration operation will take to complete.

Table 5. LUN migration rate by priority

Setting	Observed rate
ASAP	Up to 650 GB/hr (depends on config)
High	12.24 GB/hr
Medium	5.76 GB/hr
Low	3.24 GB/hr

The test parameters were as follows:

- System: CX4-960
- Write cache: 10.7 GB
- Drives: 15k rpm, 146 GB