

WHITE PAPER

Storage System Availability: No Longer Just a High-End Requirement

Sponsored by: EMC

Brad Nisbet
May 2007

IN THIS WHITE PAPER

This white paper examines how many organizations are addressing high-availability requirements with respect to their storage environments. It explores the evolution of high availability, as the need has migrated from large, high-end IT environments to more pervasive deployments among midsize organizations looking to provide increased availability among a much broader set of business needs and applications. The paper then focuses on deployments of EMC's CLARiiON CX3-UltraScale series midrange storage systems and the associated benefits customers have experienced with the high-availability characteristics of the CLARiiON solutions.

SITUATION OVERVIEW

Introduction

High-End High Availability

Historically, the notion of high availability was considered to be a characteristic of only sophisticated, high-end IT environments deployed in the largest organizations. High availability, which was once solely devoted to mainframe environments, first began to spill over into very large, open storage systems. These systems were designed to emulate the high performance and exceptional reliability of their mainframe cousins while creating opportunities for users to deploy applications on open operating systems such as Unix, Windows, and Linux.

In the past, high availability was devoted to larger, centralized pools of applications and associated storage due to the dependency on these core applications to provide business continuance, constant data availability, and ultimately a successful business. The reigning architecture consisted of business-critical applications on the front lines combined with backup and archive environments behind the scenes (typically in the form of tape). The focus on high availability was almost entirely aimed at these core business applications, typically running either on mainframe, high-end open storage systems or on a select few midrange systems capable of delivering high availability.

Today's IT environments are entirely different. They are becoming more sophisticated and complex at every level. Larger organizations have expanded their datacenters to include multiple tiers of applications to accommodate a much broader set of internal and external customer requirements. In doing so, administrators seek to expand the corresponding server and storage environments to align appropriately with these

varying applications, resulting in multiple tiers of storage. For many applications, this means looking to a new class of smaller yet scalable, modular, and more nimble storage arrays that allow for this flexibility and ultimately reduced cost.

In the datacenters of large organizations, however, the requirement of high availability remains imperative for a growing set of applications such as Exchange, business functions supported by applications such as SQL and other databases, and directory file services. The tolerance for unexpected downtime among all applications is extremely low. The desire to accommodate a growing set of second-tier applications with modular storage systems is a successful strategy only if the reliability and high-availability characteristics provided on these midrange systems are similar to those found on high-end systems.

As storage systems have evolved, suppliers have leveraged the success of existing performance and reliability characteristics of higher-end systems. These experiences and philosophies have influenced the development of availability characteristics in an expanding line of midrange offerings. This, in turn, has allowed large organizations to invest in expanding their architectures to a broader range of solutions without having to sacrifice high availability among second-tier applications.

Is High Availability Important to Midsize Organizations?

Large organizations are no longer alone in their quest for constant uptime. Midsize and smaller organizations are now looking for the same benefits that were once reserved exclusively for high-end products deployed in larger organizations that could invest the time, resources, and money toward establishing a highly available IT environment. For most businesses today, the high expectations to provide constant uptime are pervasive among a growing set of internal and external users.

Organizations or subsets of companies that rely on external access of data, such as Internet or other transaction-oriented interaction or exchange of business data, place increased importance on uptime. For these businesses, the correlation between downtime and loss of business is clear.

Midsize businesses whose IT organizations focus on providing SAN storage to support applications for internal departments are also increasing their need for high availability. Escalating requirements for service level agreements (SLAs) have led to an increased focus on uptime and the associated high availability of storage systems needed to achieve satisfactory results. While applications such as email or other business processes are internal, associated downtime can cause decreased productivity and lead directly to loss of business.

In short, customers among both large and midsize organizations have come to expect high availability as an inherent characteristic in midrange storage systems. This expectation has driven many customers to explore options based not only on typical metrics around performance, scalability, and capacity characteristics but also on the high-availability/reliability aspect of any given storage system.

High Availability on EMC's CLARiiON CX3 Midrange Storage Systems

One such company that is addressing the high-availability needs among users of midrange storage systems is EMC. The company's CLARiiON midrange storage systems have been shipping to customers since 1992. In May 2006, the company began shipping its latest generation of CLARiiON systems, the CX3-UltraScale series, which now includes the CX3-10, CX3-20, CX3-40, and CX3-80.

EMC has always focused on providing availability features on its CLARiiON products. However, with the latest CX3 line, the company has made significant investments in a number of key areas to advance the levels of high availability for these storage arrays. Significant steps forward in product hardware architecture are complemented by advancements in software functionality. EMC has also significantly enhanced processes and serviceability with respect to CX3 deployments and has established best practices to leverage lessons learned along the way. A significant portion of these enhancements comes from the increased level of direct monitoring and outage reporting and the coordinated process to improve quality and availability. All of these investments are made in an effort to increase the uptime of CX3 products and enhance the overall customer experience by providing the required and expected availability.

Customers who are interested in high availability should be asking storage suppliers about specific product hardware and software features that address high-availability in storage arrays. In addition, customers should ask their suppliers how availability is measured in order to ascertain if these features have a meaningful impact.

High-Availability Product Advancements for CX3

EMC has responded to this customer challenge by making significant advancements in CX3 hardware architecture and software functionality that contribute directly to increased levels of availability, including:

- ☒ UltraScale Architecture, which is an architecture that, along with full data path protection throughout the storage system and disk rebuild logging, allows for increased availability among drives, controllers, and other system components
- ☒ UltraPoint Technology, which is the combination of a point-to-point hardware design within the drive enclosures and unique software capabilities of the FLARE operating environment that allows for granular fault detection and isolation at the drive level
- ☒ Other enhancements to protect against disk drive data loss, including global and proactive hot sparing, sector-based checksum of all user data, and automatic background sniffing of disk media
- ☒ Redundant hardware components and online parts replacement
- ☒ Enhanced error correction capabilities in the FLARE operating system
- ☒ Ability to save write-cached data to nonvolatile memory upon power failure
- ☒ Other software-related features such as online updates and upgrades, online capacity expansion and LUN migration, and SnapView/MirrorView, which provide local and remote array-based data replication, respectively

Monitoring of Reliability and High-Availability Metrics for CX3

In addition, EMC has taken steps to monitor the reliability and uptime of the CX3 products across a significant portion of its installed base. In an effort to provide remote serviceability and event alerting to as large a population as possible, the company monitors all EMC-branded CX3 arrays that are serviced either directly by EMC or through enabled partners for system availability and downtime. This is accomplished by means of direct remote monitoring ("dial home") by EMC or via other automated or manual email or telephone notification when system outages occur. About 80% of all EMC-branded CX3s are deployed with remote monitoring, providing a significant portion of the CX3 population with direct, real-time event monitoring.

This monitoring of CX3 in all its forms serves as a key element of EMC's Total Customer Experience (TCE) program. In addition to providing enhanced levels of service and support, teams across engineering, global services, and customer quality combine efforts to monitor and continually improve the reliability of and event resolution times for CX3 products.

EMC evaluates the availability of all CX3 systems that are serviced by EMC for total possible runtime, downtime, and uptime. Runtime hours are the total number of hours (in a given time period) possible for a population of a particular array model to be up and running properly. EMC tracks the total hours of unplanned downtime, which may be caused by a variety of factors, including:

- ☒ Any hardware failure leading to an unplanned outage caused by failure of storage processors, disk drives, disk enclosures, power and cooling components, etc.
- ☒ Any software failure leading to an unplanned outage caused by the FLARE operating environment and/or array-based replication products such as MirrorView, SnapView, PowerPath, etc.
- ☒ Any process-induced failure, resulting from an EMC process, such as outages caused by an incorrect software configuration or misconfiguration in the SAN switching environment, etc.

EMC monitors all CX3 systems serviced by EMC and includes all EMC and non-EMC unplanned outages. Customer-planned downtime and outages resulting from other unplanned events such as natural disasters or application-specific misbehavior are monitored but are not included in calculations made to determine availability. The total amount of downtime is then calculated over the measurable time period and subtracted from the total runtime, yielding the total uptime for the defined population of arrays. Availability then can be calculated by producing the ratio of total uptime to total runtime. IDC has found EMC's methodology to be rigid and acceptable for measuring CX3 availability in terms of percentage of uptime over a given time period.

This sound method for monitoring and analysis has led EMC to determine that its CX3-UltraScale series of storage arrays has achieved an impressive 99.999% (5-9s) availability between May 8, 2006 (when CX3 products first began shipping) and April 28, 2007.

With the customer expectations and requirements around uptime and availability, 5-9s is meaningful. In a given year of product deployment, 5-9s translates to a very small amount of downtime — as little as just over 5 minutes for a single array, compared with 4-9s availability, which translates to as much as 52 minutes of annual downtime. Described in different terms, for every 1 million hours of runtime, 4-9s availability means 100 outage hours, compared with 10 outage hours achieved with 5-9s availability.

Equally important, EMC has developed a process by which it can better understand and improve quality, service, and uptime from understanding the outages it monitors. Upon recording an outage event, an expansive group of EMC teams engages in a process to analyze the root cause of the problem while working in parallel to develop an immediate resolution. For this improvement process, EMC considers all types of outages due to either hardware, software, process, or application issues and all other indeterminate outages.

The resolutions are entered into a central knowledge base, which facilitates the ability to broadly leverage the material for widespread distribution to the installed base, if needed. This process also arms EMC with an extremely advanced understanding of availability improvements to enhance future generations of the CLARiiON line.

Overall, the process EMC has established for analyzing availability is very impressive and consistent with methodologies found in other areas of the IT industry for characterizing system availability. The comprehensive monitoring and scrutiny of its own installed CX3 base are essential to understanding system characteristics. More importantly, EMC's process will help to address system- or customer-related issues as quickly as possible so that appropriate corrective actions can be implemented. In this way, the company can maintain high levels of customer satisfaction with its CX3 product.

EMC Customer Experiences with CLARiiON CX3

IDC spoke with several users who have deployed EMC's CX3 products primarily to provide networked storage to support a wide variety of applications. In all cases, the reliability and associated high availability of the CX3 series proved to be attractive, if not required, characteristics of the deployed SAN solution.

Large Software Company Standardizes on CLARiiON CX3-80 to Provide Reliable Storage Utility

A very large software company is consolidating islands of isolated storage that have proliferated throughout the company over the past five years. This growing, centralized storage utility, which is managed by four administrators, consists of about 3PB of SAN storage today and is expected to grow by about 1PB per year.

The storage utility supports a variety of critical applications such as email, SQL databases, several development projects, and numerous customer-facing support databases and online services. As such, the SAN environment has grown over the past several years to comprise about 200 storage arrays from multiple vendors, including a mix of EMC's CLARiiON CX700 and newer CX3 series products. The administrators plan to standardize future SAN growth on EMC's CLARiiON CX3-80 storage arrays, citing the importance of these applications to the overall business and the absolute requirement for high availability.

Consistent with administrators in a growing set of companies, the storage utility's administrators described system uptime as a critical aspect to achieve the required SLAs delivered to the departments. This is true particularly for customer-facing applications where downtime can result in direct loss of business, such as real-time online software product activation.

The group operations manager was specifically attracted to the CX3-80's ability to provide zero downtime to accommodate firmware updates for controllers, disks, and the switching environment. According to the manager, "We've had zero downtime with the EMC products. We've never had to bring a device down." In addition, the manager stated the importance of CLARiiON's integrated Virtual LUN technology to accommodate movement of data within the CX3 storage environment without having to bring devices down. "One very large benefit of EMC [is that] it can move LUNs from disk to disk in the background without affecting the application. Not all storage arrays can do this."

Overall, the high availability of EMC's CX3-80 is playing a crucial role in this software company's need to provide critical uptime for a variety of internal and external applications.

Specialty Contractor Deploys CLARiiON CX3-20 to Support Its Most Critical Applications

A leading specialty contracting firm focuses on commercial and specialty construction such as clean rooms, acoustical ceilings, commercial drywall, cold storage facilities, and mechanical insulation. A corporate datacenter currently supports about 700 employees located at 50 branch offices throughout the United States and also supports an additional 2,500 to 3,000 employees among 40 job sites.

The company's Exchange email service and several SQL databases housed at the central datacenter are the lifeblood of its business. These applications provide real-time communication and support for estimators to make bids and secure contracts. In addition, the databases support access to CAD drawings and other critical information necessary to close business or track and resolve issues at ongoing jobs. "If our systems become unavailable even for a short period of time, we risk losing multimillion-dollar contracts. We run our whole business on these SQL databases. We can't live without these applications," said the company's systems administrator.

When investigating storage array options to expand the company's growing 6TB SAN, the contracting firm considered solutions from several SAN storage vendors. Ultimately, the company chose EMC's CX3-20 storage arrays, concluding that the high availability achieved on EMC's CX3 products stood out based on EMC's reputation for product reliability and ability to provide service and support.

Major Northeast University Supports External Users

A major university in the Northeast provides Exchange email and file-sharing directories to approximately 5,000 faculty and staff, over 19,000 enrolled students (over 7,000 on campus), and thousands of transient students and alumni. Growing the centralized SAN environment from 25TB two years ago to over 110TB today has been an incredible challenge for the storage administrator, who explains the expectation he observes

around availability and uptime: "Certainly everybody expects everything all the time now. The killer applications in today's world are email and instant messaging. Being up and being available all the time [are] key to every enterprise."

Encouraged by his management to look for ways to work smarter, and not harder, and having been impressed with the reliability of the 37TB of storage deployed on an EMC CX700 since 2005, the storage administrator chose to deploy an additional 70TB on two EMC CX3-80 arrays to accommodate necessary SAN expansion in June 2006. The university has plans to deploy additional CX3 arrays within the next 12 months to migrate existing administrative applications off of an older, non-EMC high-end storage array.

Along with matching the criterion for high availability, the administrator found that the CX3 arrays also provided a low-cost solution with the benefit of 4Gbps performance, which is not readily available on all midrange SAN arrays. Overall, the administrator has been extremely pleased with the EMC arrays: "Did it meet our expectations? Definitely!"

CHALLENGES/OPPORTUNITIES

EMC's midrange line of CLARiiON CX3 storage arrays is helping customers address sophisticated network storage needs with highly available solutions. As EMC strives to improve the process by which it monitors CX3 availability and to improve product quality, it should consider the following:

- ☒ The current CX3 hardware and software capabilities provide an extremely high level of availability. As EMC strives for more sophisticated software offerings by partnering with third-party suppliers, the company needs to maintain strict levels of testing and qualification to ensure that these third-party solutions will not compromise the resulting high availability. This will be an increasing challenge for all storage suppliers as partnerships expand to provide overall customer solutions.
- ☒ Customers have found EMC's ability to move data between LUNs within a single system, without downtime, to be extremely useful and impressive. EMC should consider advancing this functionality beyond a single unit with existing and future virtualization technologies. In doing so, the company must make sure to expand the established levels of high availability beyond single arrays.
- ☒ Availability of midrange storage systems across all suppliers in the market will continue to evolve. Over the long term, high availability will continue to move toward becoming a "checklist" item, if not an assumed requirement. Although this level of availability may take some time to become pervasive in the midrange market, EMC has already established a meaningful focus in this area. Moving forward, the company will need to continue the evolution of its total storage solution availability and differentiate its offerings from those of its competitors as solutions become more sophisticated.
- ☒ The notion of high availability will continue to migrate its way downward toward smaller systems aimed at smaller organizations. The low tolerance for downtime and the desire for constant uptime are already pervasive. There is an opportunity for EMC to leverage its long-term monitoring and knowledge base of high-availability expertise and bring increased uptime into an expanded set of products in the low end such as its AX product line.

CONCLUSION

High-availability requirements are no longer exclusive to high-end IT environments and solutions. Indeed, increased uptime to avoid direct loss of business or to provide ultimate internal or external customer satisfaction is becoming a common expectation among midtier solutions as well.

With this increased focus on business continuance, availability and reliability are gaining equal footing in the selection process along with performance and scalability. Customers should continue to ask for increased visibility into storage vendors' positions regarding availability/reliability as part of their RFP process. Prudent customers will look for advancements in hardware, software, processes, and best practices that will advance the level of storage system availability. In addition, there will be a growing importance on the metrics and standards used to measure availability, and storage suppliers that have developed such practices will be well-poised to accommodate a growing population of demanding customers.

With a legacy of focusing on reliability and customer satisfaction, EMC has been able to leverage a vast amount of experience and knowledge to intensify the focus on availability in its midrange CLARiiON CX3 systems. Businesses among both large and medium-sized organizations are expressing the benefits derived from increased uptime and high availability provided by EMC's CX3 products.

EMC's advancements in and improvements to hardware, software, processes, and implementation of best practices all serve to significantly accelerate the evolution of high availability within its midrange products. The comprehensive and sound methodology employed by EMC to monitor and advance the high-availability characteristics ensures commitment to customers for future enhancements. The resulting 99.999% availability calculated by EMC should serve as an indicator of the company's commitment to providing highly available arrays to a broader set of applications among a wide variety of organizations.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2007 IDC. Reproduction without written permission is completely forbidden.