

Information Integrity, Access, and Security

An Evaluation of EMC Celerra

I. Executive Summary

It is the opinion of Kahn Consulting, Inc. (“KCI”) that EMC Corporation’s (“EMC”) Celerra product line (“Celerra”) offers capabilities that are essential to meet the growing need that today’s organization have to better manage and control their information assets. By protecting the integrity of information; controlling access to information; enforcing retention periods; offering business continuance and disaster recovery capabilities; protecting system security; and ensuring the proper destruction of information, Celerra can help to ensure that an organization’s critical information is stored and managed in a manner that meets compliance and business requirements.

II. Evaluation Overview

Kahn Consulting, Inc. (“KCI”) was engaged by EMC Corporation (“EMC”) to evaluate the company’s Celerra product line.¹ The primary focus of this evaluation is those Celerra capabilities that address the integrity, accessibility, security, and privacy of information. In conducting this evaluation, KCI has assessed Celerra’s capabilities using criteria derived from broad compliance requirements and best practices related to information management.

Evaluation Background

Whether because of new laws and regulations; greater scrutiny from regulators, boards, partners, and customers; or simply because of the exponential growth of digital business information, organizations worldwide today face new requirements and expectations regarding the way that business information should be stored and managed. As a result, today’s organizations should be evaluating products and services that can help them meet higher standards for information management while controlling budgets and meeting performance and service level goals.

Many organizations are now adopting an approach to information management that seeks to ensure that the time and resources spent managing a given piece of information is commensurate with the value of that information. Often referred to as “information lifecycle management,” this management strategy correlates the capabilities of their storage systems to the value - and the nature of - the information being stored. In such a model, information that has the lowest value to an organization (i.e., as indicated by its frequency of access, business function, and so on) is stored on

Not a legal opinion or legal advice. For all questions regarding compliance with specific laws and regulations seek legal counsel.

WHERE LAW & TECHNOLOGY MEET

KAHN
CONSULTING INC.

July 2006

PO BOX 1045 • HIGHLAND PARK IL • 60035
PHONE: 847.266.0722 • FAX: 847.266.0734 • EMAIL: INFO@KAHNCONSULTINGINC.COM

systems that have the lowest cost - and typically also the lowest performance and most limited set of management capabilities.

However, as organizations come to rely on digital information for an increasing number of business processes with increasing value and risk associated with them, the volume of information that requires significant management control capabilities also increases.

For example, in a situation where employees only rarely use the email system for business activities with high value or risk, it may be appropriate to store and manage email data on storage systems with limited controls related to data integrity, accessibility, protection, and confidentiality. However, in the situation where most of today's organizations operate - i.e., one where the email system is routinely used for negotiating contracts, corresponding with regulatory, handling customer complaints, and other high value/high risk business activities,² such controls become critical.

Information that may have previously been viewed as ancillary to the business process, or of limited value, today may be the lifeblood of an organization. Whether found on employee desktop computers, in the enterprise email system, in file shares, or other locations, the unstructured information found across the enterprise is typically of greater value today than ever before. Consequently, the systems used to store and manage this information should offer capabilities that address this increasing value and that enable organizations to proper control and manage their information assets.

While those evaluating storage systems often exclusively focus on performance, integration, and similar factors, KCI recommends that organizations increasingly consider compliance requirements and best practices in their evaluation process.

WHERE LAW & TECHNOLOGY MEET



III. About Celerra

“Celerra” refers to an EMC product line of Network Attached Storage (“NAS”) devices. A “Celerra” device is comprised of integrated hardware and software components that are designed to provide file storage capabilities for large-scale environments. Celerra products are available in a number of variants that are engineered to provide specific levels of performance and storage capacity at specific price points. In this evaluation, the baseline capabilities of the Celerra product line were examined - and not any capability that might be specific to a particular product variant.

Network Attached Storage

Network Attached Storage (“NAS”) generally refers to a group of software, hardware, protocols, and other components that enable the storage and retrieval of digital information from storage devices over a network.

NAS devices can offer several advantages over storage devices that are directly attached to a computer (e.g., a hard drive in a desktop computer), including:

- Centralized management and administration of storage devices
- Economies of scale resulting from device consolidation, resulting in more storage capacity across the enterprise at a lower cost
- Sharing of information in heterogeneous computing environments (i.e., such as users on Windows and Unix computers)
- Centralized backup, archive, and recovery
- Greater enterprise control over the retention and management of information assets

A typical enterprise use of NAS devices is for providing file-sharing capabilities that extend desktop storage capacity and facilitate collaboration. In addition, organizations often leverage high performance/high capacity NAS devices to support enterprise applications such as email.

Celerra is designed to offer these and other benefits and capabilities, as explored in more detail below.

Celerra Architecture

A Celerra controller device is a rack-mounted enclosure comprised of one or more Celerra “blades” and one or more Celerra “Control Stations” that, taken together, control the operations of one or more storage systems (such as a disk array). The function of each of these components is as follows:

- Each blade is a self-contained computer containing an independent file system running within EMC’s Data Access in Real Time (“DART”) operating system. Each blade is connected to a storage system, and controls requests to store and retrieve information on and from the storage system, and to perform other operations.
- Each Control Station is a self-contained computer that provides administrative access to the Celerra blades for the purposes of monitoring, configuration, software installation, and other purposes.

WHERE LAW & TECHNOLOGY MEET

KAHN
CONSULTING INC.

- Each Celerra controller device can connect through a variety of interfaces to a disk array or other storage system used to store information written to the Celerra. A disk array (such as EMC's CLARiON) is a device that contains multiple disk drives and supporting hardware and software that enable the device to provide high storage performance and various storage capabilities, such as RAID ("Redundant Array of Independent Disks").

Celerra File Level Retention

Celerra implements a capability called "File Level Retention" that is designed to provide "Write Once, Read-Many" protection for files. This Celerra capability is designed to protect files and directories from deletion, alteration, renaming, or overwriting during a designated "retention period." Celerra File Level Retention is designed to provide organizations with the ability to protect the integrity of stored information and to enforce retention periods.

This Celerra capability is built upon native functionality found in the file system protocols that Celerra supports, namely Common Internet File System ("CIFS") and Network File System ("NFS").

Within Celerra, WORM files have one of three states:

- 1) **Clean.** The beginning state of a file before it is designated as WORM. It can be deleted, altered, renamed, or overwritten in the same way as any other file.
- 2) **Worm.** The state of a file that has been set to read-only within the file system and cannot be deleted, altered, renamed, or overwritten until it has reached its retention period.
- 3) **Expired.** The state of a file that has reached the end of its retention period, at which point its retention period can be extended (and therefore the file reverts to WORM state), or the file can once again be deleted, altered, renamed, or overwritten.

This functionality can be accessed in Celerra by a variety of third party applications so that the information that they store in Celerra can be designated as WORM.

WHERE LAW & TECHNOLOGY MEET

KAHN
CONSULTING INC.

PO BOX 1045 • HIGHLAND PARK IL • 60035
PHONE: 847.266.0722 • FAX: 847.266.0734 • EMAIL: INFO@KAHNCONSULTINGINC.COM

IV. Celerra Capabilities

This part of the Evaluation is divided into sections that describe the integrity, accessibility, security, and privacy capabilities that are desired in storage systems; explain why each capability is desired; and evaluate Celerra's capabilities against desired capabilities.

Protecting Information Integrity

Desired Capability. Information should be protected from unauthorized alteration. A system that protects information from unauthorized alteration can minimize the likelihood that the integrity of the information will be diminished.

Information Management Principle. Information has integrity if it can be demonstrated that it has not been altered and remains accurate since it was created or archived. Business best practices and many laws and regulations require digital information to have integrity. In addition, there may be cases where it is desirable to prevent any alteration to a file once it has been created, including alteration by the user, system, or administrator that created or has access to the file.

Celerra Capabilities. Celerra CWORM works to prevent the unauthorized alteration of electronic information by enabling files to be designated as WORM, and subsequently protected from deletion, alteration, renaming, or being overwritten for a pre-designated period of time.

Celerra's capabilities in this regard rely in large part upon the native functionality of the CIFS and NFS file system. In addition, no additional or proprietary encryption or similar algorithm is employed as part of the CWORM functionality. The CWORM capability is in large part dependent upon security protections generally built into Celerra and the related file system protocols, and also on adherence to quality information security practices as it relates to management and administration the Celerra system.

Enforcing Retention Periods

Capability. A storage system should offer the ability to code and enforce retention periods for the files stored within the system.

Information Management Principle. Laws, regulations, standards, and practices often require organizations to retain specific types of information for specified periods of time. Organizations retaining electronic information require the ability to designate and enforce retention periods for that information.

Celerra Capabilities. Celerra enables retention period coding using two primary methods.

- 1) **DART OS.** Using standard file system protocols, the Celera DART operating system provides the ability to declare a retention period for a file at the time the file is written to the Celerra. Once a file is declared in this way, the file is write-protected and cannot be deleted until the retention period expires. A limitation of this approach is that, in Celerra, entire file systems can be deleted by authorized Celerra administrators. So, while this function may be useful for basic file-level retention periods, it should not be relied upon for the enforcement of legally-mandated retention periods.
- 2) **Celerra file-level retention.** Using Celerra "File-Level Retention" capabilities, a file can be designated as WORM ("write once, read many,") and cannot be deleted or

WHERE LAW & TECHNOLOGY MEET



altered by users or administrators until the end of its retention period. In addition, entire file systems can be designated as WORM. Although individual WORM files, paths to those files, and individual directories cannot be deleted by users or administrators, administrators can delete entire WORM file systems. The deletion of a WORM file system is an event for which an audit log³ is created by Celerra.

Protecting Privacy and Controlling Access

Desired Capability. Storage systems should have the capability to prevent users from accessing information they are not authorized to access.

Information Management Principle. Various privacy-oriented statutes, regulations, and standards require organizations to protect the confidentiality of information. This may be a obligation related to protecting an external parties' information (such as a customer or partner), or may be related to protecting the confidentiality of internal information regarding employees, strategies, and so on.

Celerra Capabilities. Celerra works to protect the confidentiality of files and directories stored within the system through capabilities such as:

- 1) **File extension filtering.** Celerra uses a file extension filtering mechanism, in combination with file access control lists, to control access to files. Using this capability, Celerra can prevent users or groups from accessing particular file types, and control the types of files that can be stored in a given repository. This type of control could be useful to organizations seeking to control the use of file shares and shared storage system by helping to technologically enforce policy controls regarding access to, and use of, such systems.
- 2) **Accessed Based Enumeration.** In a Windows environment, Celerra provides the capability to prevent unauthorized users from not only accessing directories and files they are not authorized to access, but also from viewing the names of those files and directories. This capability is important in file sharing environments where directories and files contain descriptive names that could reveal confidential information.

Disaster Recovery and Business Continuity

Desired Capability. Storage systems should have the capability to replicate information stored within the system to redundant systems so that in the event of a system failure, disaster or other event resulting in the loss of data, data can be recovered in a manner that minimizes business impact and loss.

Information Management Principle. Standard disaster recovery techniques require that information be stored in at least two physically separate locations. This is also a requirement of some regulations. Data that does not exist in two or more places can be permanently lost if the device or facility housing the data is damaged destroyed, or otherwise made unavailable. Thus, there is a need for organizations to copy important data to different physical locations for backup, disaster recovery, and business continuity purposes.

Celerra Capabilities. Celerra provides capabilities designed to support disaster recovery and business continuance operations, including:

WHERE LAW & TECHNOLOGY MEET

KAHN
CONSULTING INC.

- 1) **Replication.** The Celerra “Replicator” functionality create a copy of a Celerra file system and can keep this read-only copy synchronized with the content of the original file system. This replication function can be used to replicated data to another blade in the same Celerra cabinet (to help protect against data corruption, hardware failure, etc.) or to a remote Celerra system (for disaster recovery purposes).
- 2) **Redundant hardware components.** To help protect against the loss of data and system downtime, Celerra provides redundancy for many key hardware components, including power supplies, connection to storage systems, and network connections.
- 3) **Backup.** The Celerra DART operating system support sophisticated data backup techniques that, among other things, allow information to be archived to tape and other media in a manner that minimizes impact on system performance.
- 4) **Notifications.** Celerra can be configured to automatically provide a variety of detailed notification to system administrators when Celerra components are failing or when the system otherwise requires the administrator’s attention in order to prevent potential system downtime and/or data loss.

Managing Data for Efficient Access

Desired Capability. Storage Systems should provide the capability to efficiently move information from one storage system to another as the value of the information changes, and/or the organizations management and access needs change.

Information Management Principle. The value of a given piece of information typically changes through its lifecycle. For example, transactional information that starts out as high-value, frequently accessed, and frequently changing may ultimately need to be retained as a fixed content business record that should not change, is rarely accessed, but must be retained for many years. Storage systems should provide the flexibility to support organizations’ needs in this area.

Celerra Capability. Through the use of EMC’s “FileMover” software, Celerra enables the automatic movement of data files from higher speed, higher cost storage to slower and less expensive storage environments (and back again) based on file size, frequency of access, and other configurable factors. In addition, Celerra can be configured to move data from Celerra to EMC’s Centera product line. Centera is designed to provide fine-grained control over retention periods and other capabilities that support the long-term retention of business records.

Proper Destruction of Information

Desired Capability. Storage systems should provide the capability to properly destroy information once it is no longer needed.

Information Management Principle. Destruction is the final lifecycle stage of most information. In the digital world, it can be difficult and expensive to ensure that electronic information is properly destroyed. This can lead to situations where “deleted” files are recovered or recreated in the course of litigation, for example. In addition, the requirement to properly destroy certain types of private information is a requirement of existing and emerging privacy laws and regulations in the US and abroad, including the Federal Trade Commission rules regarding the proper disposal of consumer information.⁴

WHERE LAW & TECHNOLOGY MEET

KAHN
CONSULTING INC.

Celerra Capabilities. Celerra can be configured to use EMC's "Certified Data Erasure" service to overwrite and digitally "shred" information in a manner that conforms with the US Department of Defense 5220.22-M (i.e., DoD 5015.2) standard for permanently deleting digital information.

System Protection

Desired Capability. Storage systems should provide information security controls and capabilities that protect the system and its contents from unauthorized access.

Information Management Principle. The integrity of information stored within a storage system is largely dependent upon the security of that system. As such, organizations wishing to rely on electronic information should ensure that the systems they use to store and manage that information provide adequate security controls.

Celerra Capabilities. In Celerra, no direct access is provided to the DART operating system. Rather, administrative access is provided through a Celerra Control Station. This has the effect of separating administrative access from data access, thereby making it more difficult for unauthorized users to gain access to data stored and managed by Celerra.

WHERE LAW & TECHNOLOGY MEET



V. About Kahn Consulting

Kahn Consulting, Inc. (KCI) is a consulting firm specializing in the legal, compliance, and policy issues of information technology and information lifecycle management. Through a range of services including information and records management program development; electronic records and email policy development; Information Management Compliance audits; product assessments; legal and compliance research; and education and training, KCI helps its clients address today's critical issues in an ever-changing regulatory and technological environment. Based in Chicago, KCI provides its services to Fortune 500 companies and government agencies in North America and around the world. Kahn has advised a wide range of clients, including International Paper, Dole Foods, Sun Life Financial, Time Warner Cable, Kodak, McDonalds Corp., Hewlett-Packard, United Health Group, the Federal Reserve Banks, Ameritech/SBC Communications, Prudential Financial, Motorola, Altria Group, Starbucks, Mutual of Omaha, Sony Corporation, and the Environmental Protection Agency. More information about KCI, its services and its clients can be found online at: www.KahnConsultingInc.com.

V. Endnotes

¹ In undertaking this engagement, KCI exclusively relied upon information supplied by EMC through internal and external documentation, and interviews with EMC representatives. KCI does not conduct independent laboratory testing of information technology products, and as such, did not evaluate Celerra in a laboratory setting or otherwise field-test any EMC products.

² See, for example, the AIIM/Kahn Consulting 2005 Email Survey that found that a majority of organizations used email for such activities. Available at: <http://www.aiim.org/article-pr.asp?ID=29428>

³ As with any sensitive system, organizations using the WORM capabilities should ensure that administrators with the authority to delete file systems and perform other significant activities in a Celerra system have proper training and security clearance.

⁴ "Disposal of Consumer Report Information and Records," 16 CFR Part 682.

Entire contents © 2006 Kahn Consulting, Inc. ("KCI"). Reproduction of this publication in any form without prior written permission is forbidden. KCI and EMC shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice. All rights reserved. www.KahnConsultingInc.com info@KahnConsultingInc.com 847-266-0722

WHERE LAW & TECHNOLOGY MEET

