

EMC - April 2003

Centera v2.0, Compliance Edition, Compliance Edition Plus

What?

EMC recently unveiled the second generation of Centera, their Content Addressed Storage (CAS) solution. In addition to several technology updates, EMC announced new functionality aimed at helping companies in regulated industries achieve compliance. Centera Compliance Edition (herein referred to as CE) is designed for environments that must adhere to strict records retention and preservation regulations. An enhanced version of Compliance Edition, Compliance Edition Plus (herein referred to as CE Plus) is available specifically for financial services organizations that need to comply with the Securities and Exchange Act of 1934, Rule 17a-4.

Why

Over the next several years IT will see tremendous growth in Reference Information. (See ESG Research Report, Reference Information: The Next Wave.) Reference Information is defined as "digital assets retained for active reference and value." These assets are usually fixed in nature, as their content does not change over time. (EMC refers to Reference Information as Fixed Content).

The growth in Reference Information is driven in part by regulatory requirements. Regulating bodies, such as the SEC, FDA and DoD, require certain electronic records to be retained for specified periods of time. Because e-communications regulations that manage this type of information are all increasing in number and scope, regulations such as the Electronic Signatures and Global Commerce Act and HIPAA are attempting to accommodate the evolution of media and technology. Meanwhile, the Patriot Act and Sarbanes-Oxley look to police fraud and corporate malfeasance. Regardless of the regulations, there are commonalities across these different rules. Records cannot be changed, retention schedules need to be established, and information protection and security must be guaranteed. These requirements pose significant challenges to customers

because they need to safeguard this growing class of information for longer periods of time. The technology challenges manifest themselves in many different ways, including: the ability to access information quickly, the requirement to prove records' authenticity, and the ability to mitigate risk by keeping only what is required for the mandated time periods.

ESG expects these technology challenges to increase in number and complexity. As an example, we expect fines to be levied on companies who are unable to access records under examination in a timely manner. In the case of near-line or off-line media, it can take weeks or months to complete records searches stemming from audits or legal inquiries. In the eyes of regulators, elongated retrieval times raise a red flag, opening a company to further inquiry and potentially damaging exposure.

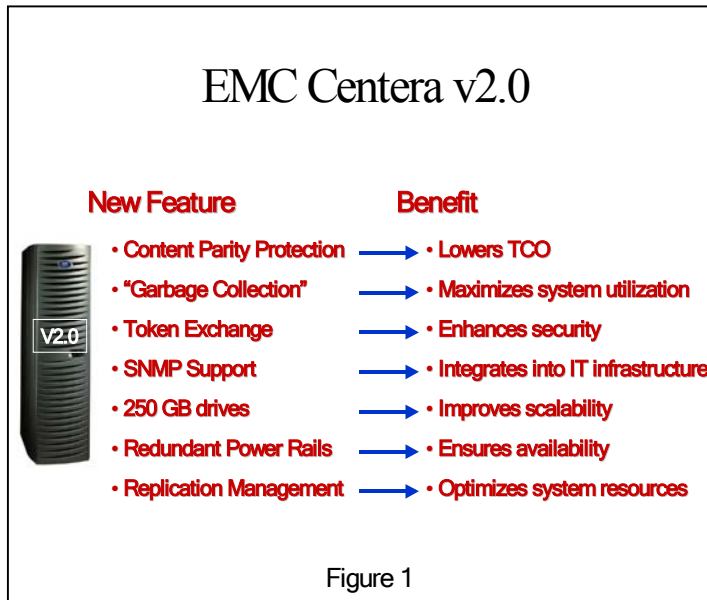


Figure 1

Making Compliance Easy

When it comes to compliance, EMC is proving its thought leadership. By providing technology that helps its customers meet their business and regulatory challenges, the company is showing that it is ahead of the curve. In addition, EMC's strong understanding of the regulations has allowed them to architect a solution that enables customers to comply with even the most stringent regulations. Based on retention policies, Centera CE enforces retention periods as set by the application. Once assigned, records cannot be deleted prior to their retention period's expiration. The inherent content-addressing technology of Centera ensures that these records cannot be overwritten. These features are crucial for regulations that mandate specific retention, preservation and disposition practices. To ensure privacy, a number of regulations, including HIPAA and DoD 5015.2 require the deletion and destruction of expired electronic content. Centera CE utilizes a

government-endorsed algorithm for shredding such objects. Once content has been destroyed, the space can be reclaimed for use in storing other information.

Centera CE Plus, containing all the features of Centera CE, is architected for financial services firms regulated by SEC rule 17a-4. This rule mandates that certain records need to be stored in a non-erasable, non-rewritable format for specific periods of time. In the past, financial services companies have used optical disk or WORM (Write Once Read Many) tape. With CE Plus, customers now have an alternative to optical storage without being subjected to the many limitations of near-line media. As with Centera CE, applications can set retention periods for each record. In CE Plus however, if an application were to fail to set a retention period, Centera defaults to “infinity” as opposed to the “zero” default in CE. For additional security, the remote management dial-in and local management consoles within Centera are disabled to ensure secure, controlled environments. ESG expects to see CE and CE Plus implemented in regulated industries as part of a total compliance solution.

What's New in v2.0?

In addition to introducing new compliance capabilities, EMC made significant technology enhancements to the base Centera product line. One key addition to Centera v2.0 is Content Parity Protection (CPP). The initial version of Centera required content to be mirrored. CPP utilizes Centera's Redundant Array of Independent Nodes (RAIN) architecture to parity-protect content. Instead of calculating parity for a set of volumes (as is the case with RAID - Redundant Array of Independent Disks), parity is calculated on the content itself. While mirroring provides faster performance and increased protection, CPP reduces the amount of disk capacity required for purchase and lowers the cost of ownership. Customers must purchase a minimum of 16 nodes to run CPP, as opposed to 8 for mirrored solutions.

Automated garbage collection, another new feature in v2.0, is designed to help customers reclaim space within Centera. Objects stored inside the system that do not have any associated object pointers are automatically deleted. The same is true for object pointers that do not have any related objects. This process helps customers maximize capacity utilization within Centera.

Centera 2.0 introduces additional security features. A software token has been developed for exchange between the application (or server) and Centera, providing an extra measure of security. Security can be

established at the DNS or application levels. Applications or servers need permission to store and retrieve data from Centera.

Further integration between Centera and the rest of the IT infrastructures is enabled with the addition of SNMP support in v2.0. Now framework applications can report on the health of a Centera system.

New functionality has been added to Centera's remote replication management, enabling operators to optimally schedule data replication procedures between Centera systems. Operators can schedule replication tasks to occur during off-peak hours to optimize network resources.

EMC also upgraded Centera hardware: 250GB ATA drives are now shipping with systems. With the new drives and the CPP features, the 8-node entry point for a Centera system is now 4TB and can scale to 23TB in a 32-node system.

Backward compatibility is always a customer concern when new versions are released. However, EMC has all the bases covered. All software enhancements, except CPP, run on v1.0 Centera hardware. (CPP requires the AC transfer switch and Y cable, which are absent from v1.0.) Also, the 250GB drives are supported in v1.0 hardware configurations, allowing customers to scale by adding capacity rather than new systems.

The Partner Factor

Centera systems are deployed as total solutions with integrated content management applications. The application controls what objects are stored. The value to customers is an integrated solution that stores digital assets for extended periods of time. Today, EMC boasts a comprehensive partner list that has integrated with the Centera API. The 120-plus partners on this list range from email archiving companies to medical imaging vendors. There are 50-plus partner applications available today. EMC has made partnering easy through an openly available Centera API and an online web connection to a Centera Lab in their Hopkinton, MA headquarters. Partners need not buy, borrow or even see Centera in order to complete integration testing—it can all be done over the Internet. With this ease of integration, ESG expects this partner list to continue growing, especially with the introduction of the Centera CE and CE Plus.

Who Cares?

In order to address business and regulatory requirements with the appropriate technologies, customers need different options to store information for different purposes. If the content being stored is Reference Information, customers need to weigh the cost of storage along with accessibility and availability requirements. Customers who need to meet compliance regulations no longer require separate solutions for their regulated and non-regulated environments. Because of Centera's flexibility, they can implement Centera with several applications, leveraging the functionality for regulatory or general business requirements.

Meeting compliance regulations is not limited to the storage of records; it means accessing those records quickly, and insuring that they haven't changed from their original form. Now, for the first time, it is possible to prove original data integrity via on-line disk storage.

Centera's pricing makes it attractive for storing any type of Reference Information. Standard Centera list prices for mirrored configurations are around .035/MB (based on useable capacity). CPP configurations should be less than .03/MB (based on useable capacity) for a CPP configuration. We expect Centera Compliance Editions to carry a 20-30% price premium. Going forward, all Centera systems will ship with the technology updates

associated with v2.0. The CE and CE Plus ship with the configuration changes as discussed above.

Bottom Line

Centera eases the pain of storing increased digital assets for longer periods of time. Whether customers are governed by regulations or by their own business processes, Centera is a formidable, cost-effective solution. The new features in v2.0 demonstrate EMC's commitment to providing a cost effective solution without sacrificing availability. Centera's CE and CE Plus systems exemplify EMC's knowledge of this space and clearly demonstrate their ability to translate that knowledge into functionally integrated solutions for regulated industries. Anyway you slice it, Centera is uniquely positioned to take advantage of, and store, waves of Reference Information both today and in the future.

Author: Brian Babineau

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Storage Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Storage Group, Inc. and is intended only for use by Subscribers or by persons who have purchased it directly from ESG. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Storage Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.