

EMC Centera Content Addressed Storage System



Innovations for the Content Wave

- The Big Picture**
- Purpose-built magnetic-disk-based records storage that overcomes current archival media limitations and facilitates compliance with records retention and preservation requirements
 - Self-management and self-configuration capabilities allow you to manage more content with fewer staff at lower cost
 - Future-proof architecture—Leverages technology to maintain long-term retention, assured content integrity
 - Self-healing—Redundant Array of Independent Nodes (RAIN) architecture, no single-points-of-failure, continuous data integrity checking/repairing, non-disruptive repair, phone-home capability, dual AC power sourcing
-

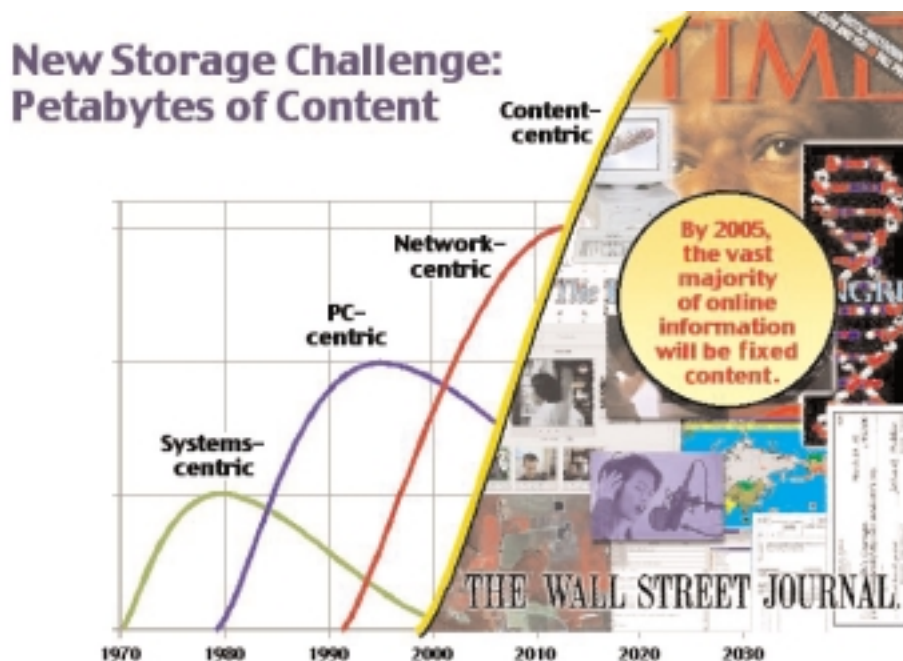
Fixed content information demands a new category of storage: Content Addressed Storage (CAS)

EMC Centera™ is a networked storage system specifically designed to store and provide fast, easy access to fixed content (information in its final form). It is the first solution to offer online availability with long-term retention and assured integrity for this fastest-growing category of information. Integrated with either an in-house-developed application or one from an expanding group of industry-focused EMC partners, Centera is for businesses that store large and expanding amounts of fixed content. Centera greatly simplifies the task of managing, sharing, and protecting large-scale fixed content repositories, enabling businesses to use this traditionally offline information to cost-effectively support new sources of revenue generation, expanded business models, and increased service levels to users and customers. Centera features WORM (write once, read many) attributes of non-rewritability and non-eraseability with disk performance and TCO (Total Cost of Ownership) superior to current archiving approaches.

The definition of fixed content is information that is not intended to change. It takes many forms, such as critical business, legal, and reference documents; X-rays; e-mail attachments; check images; broadcast content; and satellite imagery; among many others. Unlike databases or files, which change and are updated constantly, the value of fixed content stems from the combined attributes of authenticity, long life, scale, and accessibility. Existing storage architectures were not optimized for this new combination of requirements.

Once relegated to storage archives or file cabinets, fixed content is being driven online, fueled by regulatory requirements, digitization across virtually all industries, and the desire to leverage this content into new services and revenue streams. Just as the growth of applications such as Computer Aided Design (CAD) and the explosion of the Web drove the emergence of NAS to share file-based information, the need to manage and access large quantities of fixed content has given rise to a new category of networked storage—Content Addressed Storage (CAS).

More information will be created in the next two years, most of it in fixed content form, than in the entire history of humanity. Centera is ideally suited for simple, scalable, and secure storage and retrieval of this tidal wave of fixed content information.



Source: David Forrester, *Wave of Power*, and EMC.

Centera: The world's first CAS solution intended for fixed content management

Content Addressed Storage offers a significant new approach to managing information, perfectly suited to the needs of fixed content. Content addressing eliminates the need for applications to understand and manage the physical location of information on storage media. Instead, addresses are calculated based on the content itself, and serve as a unique claim check that applications can use to find and retrieve stored objects. This claim check not only simplifies the task of managing huge numbers of objects, it ensures absolute authenticity of the content itself.

EMC Centera is the world's first CAS solution designed from the outset to store and provide simple, scalable, and secure access to fixed content. Centera's CentraStar™ software operating environment employs an innovative content addressing system to simplify management, ensure content uniqueness, and deliver the scalability needed for terabyte to petabyte-level fixed content requirements. And the Centera system accomplishes this while ensuring dramatically lower overall management costs.

The Centera approach to storing fixed content offers valuable benefits:

- **Ease of Management:** Centera technology greatly simplifies system planning and management of hundreds of terabytes of content storage. With no RAID types to choose, LUNs to bind, or file systems to create, applications are released from the requirement to compensate for or manage the complexity of traditional storage topologies.
- **Speed of Access:** Centera enables shared, networked, fast access to a single copy of fixed content at sub-second speeds, enhancing the value and usability of information previously stored in less-accessible forms.
- **Content Authentication:** Utilizing breakthrough C-Clip™ technology, any object presented to the system is stored in such a way that it is unchangeable and authenticated, transparent to the end-user application.
- **Non-Eraseability:** Configurable retention settings ensure data is not erased prior to the expiration of its defined retention period.
- **Record-Level Content Management:** Content addressing allows retention protection and disposition of each individual data object, as opposed to managing retention/disposition at the platter or tape media level.
- **Efficient Storage Utilization:** Centera uses a unique content-derived address that permits only one protected copy of content to be stored no matter how many times it is used. This significantly reduces the total number of copies of information stored and is a key factor in lowering the cost of storing and managing content.
- **Scalability without Reconfiguration:** The architecture of the Centera system, based on redundant arrays of independent nodes (RAIN), is designed to be highly scalable to hold petabytes of content. Centera auto-discovers and configures new capacity as it is installed.
- **Self-Healing:** Centera continuously monitors the integrity of stored objects to detect/repair soft errors and automatically reconfigure the system and regenerate objects if necessary. It also reports these incidents through EMC's remote monitoring system so remedial action can be taken at an appropriate time.



Centera's Redundant Array of Independent Nodes (RAIN) architecture and intelligent operating environment provide non-disruptive terabyte-to-petabyte scalability; self-managing, self-healing, and auto-reconfiguration proficiency; and guaranteed authenticity of all fixed content objects.

-
- **Business Continuity Protection:** Centera can be configured to maintain duplicated copies of fixed content at a remote site, eliminating the possibility of a site disaster destroying all copies of information.
 - **Easy Installation and Non-Disruptive Upgrades:** A significant effect of Centera's breakthrough technology is that systems can be installed in under an hour and upgrades managed without disrupting content access.

Centera RAIN architecture—A Redundant Array of Independent Nodes

Centera is built upon a no-single-points-of-failure Redundant Array of Independent Nodes (RAIN) architecture that is deployed in one or more six-foot NEMA standard 19-inch cabinets. Each cabinet can be configured with 7.68 to 40.96 raw (3.58TB to 19.10TB mirrored or 11.92TB to 31.79TB parity-protected) storage capacity*. Capacity is added in increments of 8 nodes, configured as storage nodes or access nodes. Each storage node contains processing power, 1.28TB raw (up to 0.60TB mirrored or up to 0.99TB parity-protected) storage capacity, and is interconnected with all other nodes in the cluster via a private LAN*. Each node executes its own instance of the CentraStar operating environment. Access nodes are configured in pairs: each access node provides 1Gb per second Ethernet connection. The throughput and capacity needs of the application determine the most appropriate configuration.

Available and supported today by EMC and a broad range of industry-focused partners

Centera systems are available direct from EMC for customers with in-house application development capabilities and are endorsed by a large and increasing number of EMC partners integrating solutions across a variety of industries. Each partner adds its own application and industry-specific value to Centera, tailoring the system for individual markets and requirements. A complete list of Centera partners, their solutions, and information on how to become an EMC Centera partner is available on EMC's website at www.EMC.com/centerapartners.

* Usable capacity calculated for an average file size of 250KB. Usable capacity will depend on the file size and the number of storage nodes. Minimum allowed configuration for parity protection is 16 nodes.



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381

EMC, EMC², and where information lives are registered trademarks and C-Clip, Centera, and CentraStar are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

© 2003 EMC Corporation. All rights reserved. Published in the USA. 12/03

Data Sheet
C931.4