



EMC Centera Content-Addressed Storage System

Automated archiving made simple, affordable, and secure

Content-addressed storage (CAS) for archiving

Fixed content has 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, and accessibility. Existing storage architectures are 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).

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, but is also in fact a digital fingerprint of the content, ensuring absolute authenticity.

EMC Centera: archiving made simple

EMC[®] Centera[®] is an IP network storage system specifically designed to store and provide fast, easy access to fixed content. It is the first solution to offer online availability with long-term retention and assured integrity for this fastest-growing category of information. Whether integrated with an in-house-developed application or an application from a continually growing group of industry-focused EMC partners, EMC Centera is the optimal information archive for businesses and organizations that require a simplified solution to expanding amounts of fixed content. EMC Centera greatly simplifies the task of managing, sharing, and protecting all sizes of content repositories. It cost-effectively puts information online in support of new sources of revenue generation, expanded business models, and increased service levels to users and customers. EMC Centera features WORM (write once, read many) attributes of non-rewriteability and non-eraseability with disk performance and TCO (total cost of ownership) superior to current archiving approaches.

EMC Centera is the world's first CAS solution, designed from the outset to provide simple, affordable, and secure access to fixed content. EMC 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 archive requirements. And the EMC Centera system accomplishes this while dramatically lowering overall management costs.

The Big Picture

- Purpose-built archiving storage platform that provides online access and assured content authenticity
- Facilitates governance and compliance needs for retention and preservation
- 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

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

The EMC Centera approach to storing fixed content offers the following valuable benefits:

- **Ease of Management:** EMC Centera technology greatly simplifies system planning and management of terabytes to petabytes of content storage. EMC Centera does not use typical file systems for storing content. 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:** EMC 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, authenticated, and transparent to the end-user application.
- **Virtual Pools:** Content-partitioning capability without the use of standard filesystems. The secure digital content in each of the pools or in the entire EMC Centera can be easily managed, searched, and retrieved with authorized access by application or department.
- **Non-Eraseability:** Configurable retention settings ensure information 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 content object, as opposed to managing retention/disposition at the platter or tape media level.
- **Retention and Disposition Management Capabilities:**
 - **Retention Classes**—Enables retention policy management for an entire class of content.
 - **Audited Deletes**—Predominantly for European privacy laws, Audited Deletes enables the deletion of a piece of content that is still under the enforcement of a retention period through a limited and tightly audited channel (not available with Compliance Edition Plus).
 - **Configurable Default Retention Period**—Allows an application owner/storage administrator to specify a default retention period in the event an application cannot or does not assign one (not available with Compliance Edition Plus).
- **Advanced Retention Management:**
 - **Event-based Retention (EBR)**—Allows applications to set an “undetermined” retention period when the content is archived, then set a specific retention period or class at the time of the “event.”
 - **Litigation Hold**—Enables an application to put a “hold” on an object, preventing deletion until the hold is removed.
 - **Min/Max Governor**—Enables the system administrator to enforce minimum/maximum retention classes within a virtual pool for more effective policy management.
- **Efficient Storage Utilization:** EMC Centera’s unique content-derived addresses permit 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 EMC Centera system, based on redundant arrays of independent nodes (RAIN), is designed to be highly scalable to hold petabytes of content. EMC Centera auto-discovers and configures new capacity as it is installed.
- **Application-independent Archive Virtualization:** Enables multiple independent clusters to be virtualized and presented as one unified archive. Such virtualized archives accommodate varying software and hardware revisions across the clusters and enable greater archive scaling despite space, power, and distance constraints of an individual data center.

- **Self-Healing:** EMC 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.
- **Business Continuity and Disaster Recovery Protection:** EMC Centera can be configured to replicate and maintain copies of 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 EMC Centera's breakthrough technology is that systems can be installed in under an hour and upgrades managed without disrupting content access.

EMC Centera Architecture—Redundant Array of Independent Nodes (RAIN)

EMC 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 standard NEMA 19-inch cabinets. Each cabinet can be configured with 16 to 256 TB raw (7.7 TB to 125.7 TB mirrored or 26.4 TB to 215.5 TB parity-protected) storage capacity* and cabinets can be linked together for continued scalability. Capacity is added in increments of two nodes, configured as storage nodes and/or access nodes. Each storage node contains processing power: 4.0 TB raw (up to 1.92 TB mirrored or 3.3 TB parity-protected) or 8.0 TB raw (up to 3.92 TB mirrored or 6.73 TB 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 1 GB per second Ethernet connection. The throughput and capacity needs of the application determine the most appropriate configuration.

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

Available and supported today by EMC and a broad range of independent software vendors (ISVs)

EMC Centera systems are available direct from EMC 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 EMC Centera, tailoring the system for individual markets and requirements. A complete list of EMC Centera partners, their solutions, and information on how to become an EMC Centera partner is available on EMC's website at www.EMC.com/centera. EMC Centera also supports non-integrated applications through the use of EMC Centera Universal Access software.

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



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

Take the next step

Discover how EMC Centera content-addressed storage can enable you to better store, protect, optimize, and leverage your fixed-content data for your enterprise. To learn more, contact your local EMC sales representative or authorized value-added reseller, call us at 1-866-464-7381, or visit our website at www.EMC.com.