DOCULABS

*Doculabs MarketFocus Paper*

# The Business Value of EMC Centera: An IT Perspective

As the proliferation of digital information continues to grow at an exponential rate, organizations are struggling with how to cost effectively harness the business value of this information. Much of this digital information is "unstructured" content that needs to be retained for long periods of time and for which speed of access and assured content authenticity are paramount. Examples of such unstructured content include documents of all types, email, medical images and patient care information, product design documents, testing results, etc.

In order for organizations to use their digital content as an asset to drive their business (as opposed to a liability), they require robust business applications to manage the information, as well as a robust storage infrastructure. With this in mind, many organizations are leveraging EMC Corporation's Centera solution (and its Content Addressed Storage approach) to create highly available repositories of digital content. Centera provides them with quick access to any needed information, enforcement of their retention management policies, and a lower total cost of ownership than other archiving technologies.

EMC commissioned Doculabs to interview a number of Centera customers, specifically IT management, and to document the business benefits they realized. This paper provides validation from those customers and insight into how those organizations are leveraging their Centera-based archives. Throughout this paper we present Centera customer statements or Doculabs' analysis that show the far reaching gains that can be derived from these benefits.

Overall, the IT management we interviewed said that their Centera deployments:

- Cost significantly less in terms of dollars and resources compared to tape or optical alternatives
- Provide far superior response times compared with their previous systems – enabling retrieval in seconds instead of minutes, hours, or even days.
- Enforce their information retention and disposition policies within storage (which customers view as imperative for their day-to-day business, governance, and regulatory requirements)
- Provide better manageability than other disk-based solutions, contributing to a lower total cost of ownership

The customers correlated these and numerous other benefits into realized business value such as improved service levels; the ability to repurpose information for additional uses; reductions in downtime; improved disaster recovery capabilities; and improved compliance.

# What's Inside

This paper includes the following major sections:

### *About Doculabs*

Provides a brief overview of Doculabs, an independent consulting firm that advises organizations on their strategies for enterprise content management (ECM).

### *Scope of this Paper*

Describes the objective of this paper, which is to highlight the business benefits that EMC customers have derived from their use of Centera and its approach to content addressed storage.

### *Content Storage Management Challenges*

Highlights some of the problems or challenges associated with storage management of unstructured content.

### *Customer Perspectives and Doculabs' Opinion on Centera's Benefits*

Highlights the key business benefits that EMC Centera can deliver in order to address the content storage management challenges described above. Details are based on feedback from a set of Centera customers and Doculabs' industry experience.

### *An Overview of EMC Centera's Approach to Content Storage Management*

Provides Doculabs' review of Centera's components, capabilities, administration, and architecture.

## About Doculabs

Doculabs, Inc. is a technology consulting firm backed by research and extensive client experience. Our services lower the business risk of technology decisions through client-specific recommendations, objective analysis, and in-depth research. Founded in 1993, Chicago-based Doculabs provides consulting services that are based on our fundamental belief that in order to protect a client's long-term interest, technology advisors should not be implementers.

Doculabs' specializes in providing consulting services to assist our clients with their content management strategies. For example, we work with our clients to assess their current state for content management, define their requirements and identify their gaps, develop business cases for content management, devise conceptual architectures, recommend appropriate technology solutions, and develop deployment road maps for moving forward with their content management strategies and implementations.

Hundreds of leading organizations have turned to Doculabs for assistance with their strategies for content management. Doculabs' consulting services are completely objective because the firm does not sell software or integration services. For more than 10 years, our consulting methodology has provided customers facing mission-critical challenges with the information and advice they need to make confident and well informed decisions.

## Scope of this Paper

*Doculabs interviewed numerous EMC customers to identify the types of benefits and the impact that Centera delivered.*

To assess the benefits and value of Centera in customer deployments, Doculabs interviewed a number of customers with Centera in production. These customers were large companies in a variety of industries with high-volume content storage requirements, including firms in the financial services, manufacturing, medical and entertainment sectors.

For each customer interview, we asked a number of questions about the customer's experiences with storage management, both before and after the Centera solutions were deployed. The objective was to identify the impact of Centera, in order to help illustrate or generalize the types of benefits that most customers can expect from their Centera investments.

We focused our questions in a number of key categories, including storage capacity, response times, efforts and costs associated with maintenance, media management, technology migration and disaster recovery.

## Content Storage Management Challenges

Recent studies show that as much as 70% of new content created in organizations today is considered unchanging information – including scanned images of paper documents, email, finalized electronic documents, report output, and other information that needs to be securely archived in final form and remain accessible over time. The long-term uses can be for day-to-day business, governance, compliance or regulatory requirements. As the proliferation of digital content today continues to grow at an exponential rate, organizations are struggling with how to best store, manage and retrieve this information.

*Traditional long-term storage approaches for content (such as tape and optical storage) have numerous limitations that impact the bottom line, including constant management, slow retrieval, questionable long-term accessibility, questionable content authenticity, reliability risks, and high maintenance costs.*

Historically, organizations have used older technologies such as tape and optical storage in combination with any file systems based disk for their long-term content storage needs. These approaches have numerous limitations, including:

- **Slow retrieval**.
  Retrieval is typically extremely slow from near-line and offline storage solutions such as tape and optical, especially for organizations that have amassed collections of thousands of tapes and disks. Thus, it may require hours – and in some cases days or even weeks – from the time a business user first requests a document to the time the restored document is delivered to the user.
- **Questionable long-term accessibility.**
  With older technologies, there is always the risk of technology obsolescence, as file formats, operating environments, or storage platforms become incompatible with newer systems or may not be supported by the original vendors. Such obsolescence requires migration to newer storage solutions, which can result in data loss.
- **Questionable content authenticity**.
  With solutions such as tape, it can be difficult to know if the correct version of a document has been retrieved. Given the *chronological* process of archiving content (i.e. making a backup tape of a server each night), there can be multiple versions and copies of the same document, stored over and over again, across different tapes or disks technologies over time. Even unintelligent arrays of hard disks face similar challenges, as they do not prevent content redundancy and do not provide assured content authenticity and security, which are critical for long-term archival and compliance.
- **Reliability risks**.
  With systems such as tape or optical storage, there is always the potential for mechanical failure. In addition, there is potential for media degradation or damage over time.
- **High cost of maintenance**.
  Maintaining tape and optical storage systems requires maintenance of hardware devices and mechanics as well as the media itself. Specialized skills, manual handling, and physical storage space all add to maintenance costs.

- **Cost of technology upgrades and migrations due to obsolescence**. Tape and optical storage solutions have a history of obsolescence of technology and drive formats. As a result, organizations are forced to periodically replace equipment and migrate data, as newer equipment may not be backward-compatible.

- **Content disposition challenges**. Organizations must effectively manage the destruction or disposition of information in accordance with their records management programs. This is a challenge with solutions such as optical, in which a piece of content that needs to be destroyed is stored on the same platter as other content that still needs to be retained. Accomplishing the disposition requires a time-consuming conversion process in which the platter is read, the unwanted data is purged, the data to be retained is written to a new platter, and the original platter is physically destroyed. The situation is similar when disposing from tape, and the process involves a series of manual steps.

EMC designed its Centera solution to address these types of limitations. The following section provides perspective from Doculabs and a number of EMC customers on the business value that Centera provides.

# Customer Perspectives and Doculabs Opinion on Centera's Benefits

This section highlights the key benefits that Centera can provide, based on Doculabs' perspective and interviews with a number of Centera customers.

## *Summary of Overall Findings*

During the interviews, it became clear that the differences in customers' business and technical environments were greater than expected, yet the benefits of Centera were similar in many areas. For example, we found that:

- **Centera provided dramatic improvement in the retrieval time to deliver information to those who need it.** The customers we interviewed said that they measured their retrievals from Centera in *seconds*, as opposed to *minutes or hours* from optical media, and *days or weeks* from tape backups. Many interviewees reported that retrievals were "nearly instantaneous". They viewed this level of performance coupled with assured content authenticity as critical to improving overall business process efficiency.
- **Centera offers improved storage management while reducing the cost and effort involved.** 80% of the interviewees cited the improvements in manageability, maintenance cost, and effort as one of the top three reasons for choosing Centera. Universally, all of the interviewees reported that managing Centera represented only a small fraction of the time spent managing all of the storage solutions they used. In addition, Centera eliminated costs associated with media, media management, library management, and equipment maintenance typical with tape or optical storage. The bottom line for these customers was a lower total cost of ownership than their previous storage approaches.
- **Centera provides secure, long-term archival of business-critical information.** Centera provides an efficient, secure, tamper-proof records storage platform. The system ensures long-term online accessibility to content, minimizing concerns about information loss or the technology obsolescence associated with alternative storage technologies. An interviewee from an entertainment company noted that growth in information discovery requests and litigation made it important for them to select a system like Centera to ensure data accessibility through metadata searching, and to protect against loss or corruption

Numerous other benefits of Centera were cited as well, including high scalability and high reliability. In addition, the customers we interviewed considered Centera to be a critical element in their organization's overall strategies for compliance, records management, content management, and information lifecycle management (ILM).

With its wealth of benefits, Doculabs believes that in many cases Centera is preferable to alternative archive approaches such as tape, optical and traditional magnetic disk storage, and that any organization with long-term archival needs for large volumes of information should consider Centera.

## *Details*

This section reviews in greater detail some of the key findings from the customer interviews. We asked questions about a wide array of value-added areas that customers might have considered, and asked each customer to identify their top three drivers for selecting Centera.

The following chart shows the frequency of the top three drivers across the customers Doculabs interviewed. For each driver, the chart shows the percentage of interviewees that said the driver was one of their top three priorities when they evaluated Centera.

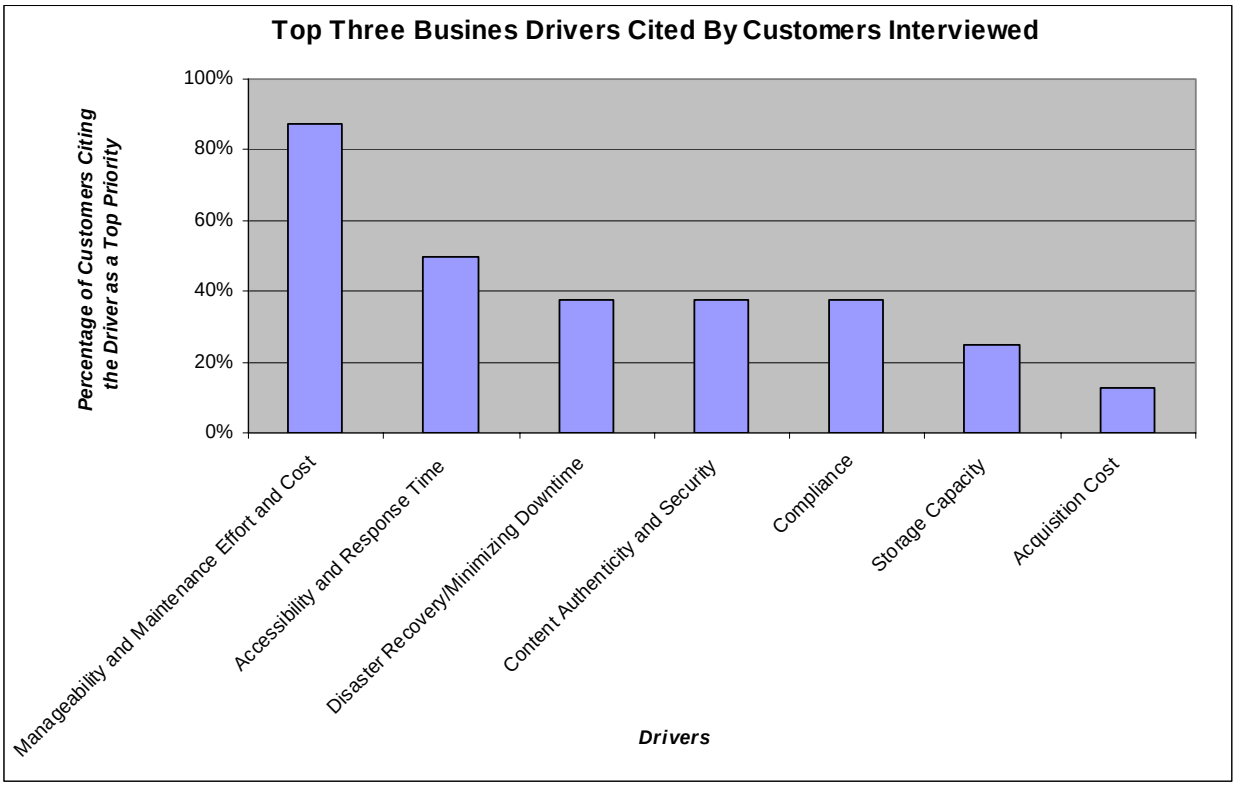


As the figure shows, the two drivers that interviewees cited most frequently were manageability and maintenance effort and cost and accessibility and response time. The following subsections provide additional details on the benefits and returns that Centera is providing customers in these areas, along with details on other benefits of Centera that the interviewees shared with Doculabs.

(Important note: due to competitive concerns or company policies, we have honored the interviewees' requests not to cite their organizations by name.)

### *Manageability and Maintenance Cost and Effort*

80% of those companies interviewed cited the **improvements in manageability, maintenance cost, and effort** as one of the top three reasons for choosing Centera as their information archive.

**Customer Quote:**
***"Once you set [Centera] up and set up your rules there is zero FTE time required to manage it."***

The organizations we interviewed were looking not just at the cost of the storage technology itself, but the total cost of ownership associated with the storage solution. In this area, customers wanted systems that were easy to maintain and required minimal personnel for maintenance. In addition, their objectives were to eliminate time and costs associated with their disk-based file system, tape, and optical storage solutions – items such as backup management, the cost of media, and the effort required for tape handling and media management.

In Doculabs' experience, organizations frequently underestimate this maintenance effort, leading to higher cost of operations. But the customers we interviewed universally cited the benefits that Centera's manageability provides for lowering total cost of ownership. For example:

- An interviewee from a financial services firm said that with Centera, "you just don't have to touch it; it simply runs itself," and estimated that the company had eliminated more than $10,000 in monthly maintenance costs associated with its old tape storage approach by moving to Centera.

**Customer Quote:**
***"We realized an 18-month payback in time alone compared with maintaining our optical systems."***

- An interviewee from a major insurance company said that Centera nearly eliminated the 120 hours per month they used to spend on optical storage maintenance, translating into an "18-month payback in time alone compared with maintaining our optical systems".
- An interviewee from a health care services provider commented that "when we went to Centera, we did not dedicate any new employees to manage it. Too many organizations fail to take this factor into account."
- An interviewee from a large radiological services organization said that only half of an FTE spends time managing the Centera environment.

In addition to the maintenance time and cost savings that Centera provided the customers we interviewed, many of them commented on the specific maintenance features and inherent capabilities of Centera that they found especially valuable. Specifically, customers commented on Centera's:

- "Self-healing" features to continuously monitor the system to detect and automatically repair errors
- The ability to easily add new capacity to the storage environment, as the system auto-discovers and configures the new capacity as it is installed
- The ability to upgrade firmware without having to take the system offline which minimized total application downtime

### *Accessibility and Response Time*

50% of the companies interviewed cited accessibility and response time as one of the top three reasons for choosing Centera. Based on our research and customer experience, Doculabs has observed this area as critically important to application end-users, as user frustration from slow response times leads to abandonment of the system and lower adoption rates – which reduces the value of the technology investment.

**Customer Quote:**
***"Retrieval is very fast for our doctors. From tape, it might be 45 minutes to an hour, depending on image size. With Centera, it is nearly instantaneous."***

Response time is defined as the length of time it takes from the first request for information from a business user to the successful delivery of that information. In most of the organizations we work with, poor response times are viewed as service and performance issues by users, and are a leading source of user frustration with their current systems (and with IT). In this area, the customers we interviewed cited the dramatic improvements in response time that Centera provided in comparison with alternatives such as tape and optical storage.

On average, the customers we interviewed all said that Centera provided them with consistent retrieval times of anywhere from sub-second response times to no more than three to five seconds, as opposed to minutes or even hours from previous storage solutions. (Note: reported response times represent the total round-trip time for a retrieval, i.e. from the time information is requested until it is presented in the application.) Doculabs recognizes that, in addition to Centera, network latency and application performance contribute to the overall response times observed by the companies interviewed.

**Customer Quote:**
***"Response time used to be 27 seconds on average [from optical]. Now it is down to 1.5 seconds or less."***

Overall, the customers we interviews said that in comparison with previous solutions, Centera's retrieval speed coupled with assured content authenticity delivered ROI (return on investment) by improving and accelerating business processes – including business-critical processes such as patient evaluations, claims processing, and customer service.

For example:

- An interviewee from an insurance firm said that "response time used to be 27 seconds" on average from their old optical system, but with Centera, "response time is down to 1.5 seconds or less." The firm currently uses Centera to manage approximately 130 million images, and is adding more than 1 million new documents every month.
- An interviewee from a health care services provider said, "retrieval is very fast for our doctors. From tape, it might be 45 minutes to an hour, depending on image size. With Centera, it is nearly instantaneous."
- An interviewee from a credit union said that user response times from optical used to be 45 seconds on average, and had dropped to three to five seconds on average since moving to Centera.

Most other interviewees noted similar improvements, stating that access from Centera is so fast that it was deemed unnecessary to track the cost associated with each Centera retrieval.

**Customer Quote:**
*"Centera helped mitigate risks associated with data loss or technology obsolescence" that previously existed with alternate storage approaches.*

Several customers cited the need for information accessibility over time, and said that "Centera helped mitigate risks associated with data loss or technology obsolescence" that previously existed with alternative storage approaches. In addition, customers stated that Centera's long-term accessibility supports compliance and facilitates information discovery requests (such as for audits or litigation). Since Centera information is readily accessible, it negates the need to go to optical storage or tape backups or to perform lengthy restores in order to find information. Some customer examples in this area include the following:

**Customer Quote:**
*"Centera improved internal and regulatory compliance while also improving the speed of discovery."*

- An interviewee from an entertainment company noted that growth in information discovery requests and litigation made it important for them to select a system like Centera to ensure data accessibility through metadata searching, and to protect against loss or corruption.
- An interviewee from a large banking group said that Centera improved internal and regulatory compliance while also improving the speed of discovery

Several customers noted that they are using Centera for email archiving in order to offload their mail servers and reduce the size of user inboxes. They cited the importance of fast retrieval time on archived email as a key factor in their decision to use Centera, as slower response time would cause issues for their internal users. Another benefit of using Centera for email archiving is reduced management costs. One of the firms interviewed had previously commented that their administration costs of Centera were magnitudes less than the SAN administration costs for email.

### *Other Benefits Cited*

Apart from the benefits described above, the customers we interviewed cited numerous additional benefits that their Centera solutions delivered. The following table identifies a number of these Centera benefits, and provides customer examples of the impact of each benefit.

| Area of Benefit | Customer Example |
|---|---|
| Elimination of the time and cost associated with doing restores from tape backups (which can take more than a month in some cases, such as with tape), reducing the cost of discovery | There are numerous well-publicized examples of high cost of restoring data from tape for information discovery purposes – especially for email discovery. For example, an energy company that was required by a plaintiff to recover email from certain employees estimated that it would cost $6.2 million and take six months to recover the email from 93 backup tapes. With Centera's content addressed storage approach and metadata searching, such discovery costs can be minimized or eliminated. |
| Reduced downtime | An interviewee from a credit union reported that Centera is extremely reliable in terms of uptime, and that most upgrades or capacity additions can be done without taking the system offline – reducing costs associated with system downtime compared with tape. An interviewee from a major insurance firm previously commented that downtime costs his organization $5 million per hour and that Centera has nearly eliminated storage downtime in comparison with previous solutions. |
| A solid platform for storing multiple terabytes of data, with the ability to easily add capacity as needed | An interviewee from a major insurance firm, who now has 51 terabytes of total capacity installed, said "I love upgrades on the fly, with no disruption to the business customers." |
| Security and content authenticity to facilitate governance and compliance | "Centera fits very well with the security demands of health care, and it makes compliance easier," said an interviewee from a health care services provider. |
| Ability to integrate with enterprise content management systems | A major health insurance firm is using an enterprise content management system to capture millions of images (including claims documents, reports, and statements), with all images and documents archived to Centera for managed storage. |
| The ability to easily replicate or mirror Centera archives, minimizing concerns about archival and backup for disaster recovery and business continuity | Nearly all of the interviewees said that their Centera environments were mirrored or replicated to satellite sites with identical setups for backup and disaster recovery purposes, reducing costs associated with downtime or data loss in case of disaster. |
| The ability to improve customer service by providing fast access to the information their users need to do their jobs | A major insurance firm stated that Centera's fast response times and 24x7 availability – which was previously impossible with optical – has contributed to improved user satisfaction and customer responsiveness. |

## An Overview of EMC Centera's Approach to Content Storage Management

EMC developed its Centera storage solution to address the needs of secure, high-volume, long-term storage of content. First introduced in 2002, Centera is an efficient, secure, tamper-proof, magnetic-disk based storage platform that provides users with fast, online access and assured information authenticity making it ideal for day-to-day records management purposes.

The system can be configured to add write-once/read-many (WORM)-like capabilities, providing in one solution all of the desired features for long-term, archival of information for governance and regulatory requirements.

Centera's approach is designed to overcome the many limitations associated with tape or optical systems. For example, beyond faster retrieval, Centera provides reduced risk of technology obsolescence, reduced risk of mechanical failure or media degradation, and less expensive maintenance. In addition, Centera is designed to interface with any application from virtually any platform, allowing organization to further leverage their existing investments in systems.

Organizations use Centera for a variety of content storage and archival applications. Examples include storage for:

- Content managed by an enterprise content management (ECM) system
- General records management application
- Financial records (such as statements, check images, contracts, and regulated correspondence)
- E-mail archival/compliance applications
- Medical records (such as medical reports, forms, and X-rays)
- High volumes of "fixed" content such as images, enterprise reports or other computer-generated output (COLD/ERM)
- Internal governance, Sarbanes-Oxley and other regulatory compliance applications
- Digital asset management applications (such as for graphics, voice, and video streams)
- Engineering and CAD/CAM files

## *How Centera Works*

Centera uses an approach called Content Addressed Storage (CAS) in which each content object has a "digital fingerprint". The digital fingerprint is derived from the content itself. As a unique identifier, it enables content to be managed at the object level, ensuring content authenticity and making information location- and application-independent.

When a request to store content is made, a separate metadata file about the content is also stored. The content object's content address goes into a Centera metadata file (called a C-Clip Descriptor File or CDF) specific to that user's use of the content. A Centera metadata file contains:

- A time stamp
- User information
- Content address for the data object that was stored
- Optionally stored metadata about the object (typically a subset of the metadata originally stored by the user's application)
- Retention period for which the content must be kept
- Business logic to be read when the content is accessed (optional).

Once the Centera metadata file is created, a Content Address is calculated and returned to the application. The content itself is only stored once, which greatly reduces storage. Subsequent requests to store a specific piece of content results in the same content address being calculated, and Centera returns a pointer to the already-stored content. The result is that many users can request that a piece of content be stored, they can store it at different times, annotate it with different metadata and request different retention periods – yet only one copy of the content is kept.

The following figure illustrates the operation of Centera content addressed storage (CAS) and its digital fingerprint approach (in this case integrated with a content management system).

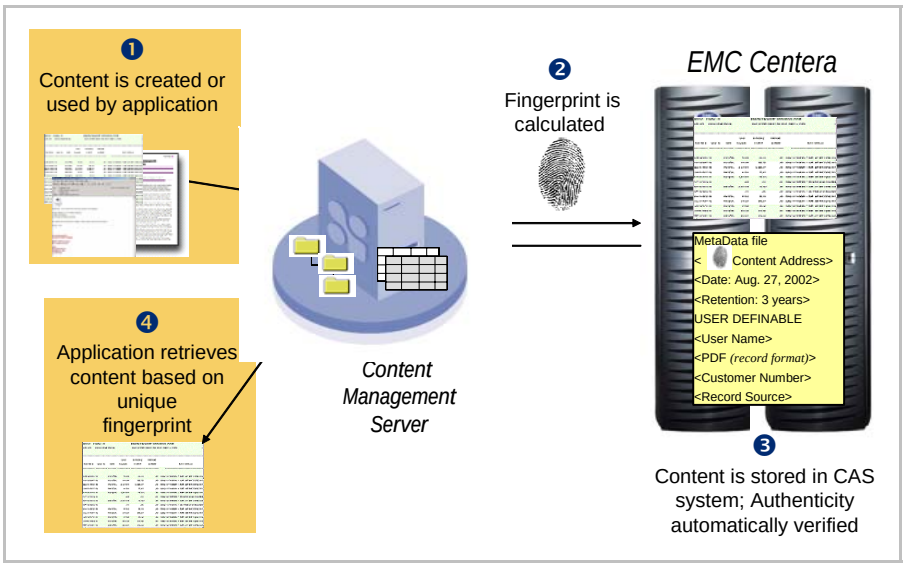


*Content objects are assigned a digital fingerprint for archival in Centera; Centera handles authentication for retrievals. (Source: EMC)*

### *Benefits of Centera's Metadata*

Centera's metadata is useful for a number of reasons. For example, Centera's metadata:

- **Ensures that information retention and disposition policies are enforced in storage.** Today it is unacceptable for content to be stored on unintelligent arrays from which it might be accidentally deleted, when an application has specified that the content be retained or when the storage administrator know that it has to be retained. With Centera, an application's retention period for a content item is enforced intrinsically in storage, and deletion is not allowed before the retention period expires.
- **Yields positive impacts on application performance.** Applications can use Centera metadata to do queries about content, without having to open the content objects themselves.
- **Enables cross-application search.** Centera offers a search capability within the archive that can traverse all metadata, irrespective of application. This is important for stakeholders who are concerned about an organization's "total" archiving needs and management – such as individuals from Legal, Corporate Records, Archives, and financial or executive ranks that have a fiduciary responsibility. Such decision-makers take a broader view of the entire archiving environment, and are highly concerned with information access and standardized policy management across the archive.
- **Facilitates chargeback reporting.** Centera can leverage the stored metadata to provide sophisticated and customizable capacity utilization reporting. This capability is particularly valuable to allocate the cost and use of the information infrastructure to the departments and applications that use it.

### *Architecture and Administration*

Centera's integrated hardware and software system supports virtually any application from any platform including: Microsoft Windows NT, Windows 2000, Windows 2003 Server, and Windows XP Professional – as well as Solaris, Linux, HP-UX, UNIX, iSeries, and mainframe platforms.

Centera is designed to be scalable. Centera's architecture is based on redundant arrays of independent nodes (RAIN), which offers petabyte scalability. Additional storage is added to the Centera environment by adding nodes which are automatically recognized by the system, making both storage and performance scalability easy with little or no administrative or application downtime during an upgrade.

The Centera architecture offers a number of advanced features such as "self-healing" to continuously monitor the system to detect and repair errors, self-management to automatically rectify errors, content mirroring or parity protection to maintain duplicate copies and enable automatic recovery within a single Centera cluster, and long-distance replication between Centera clusters for disaster recovery. One customer reported that they replicate from the main Centera environment to a backup Centera environment at their disaster recovery site more than 100 miles away.

Such features and architecture allow a typical storage administrator to manage orders-of-magnitude more content than with alternative approaches.

Centera provides an open API that is supported by nearly 220 ISV applications and software that allows non-integrated application to leverage Centera via NFS, CIFS, FTP or HHTP protocols and mainframe tape and optical emulation. With all these options, organizations can fully leverage their existing technology investments with Centera as their unified, multi-application archive.

*For more information about EMC, visit the website www.emc.com. For more information about Doculabs, visit the web site www.doculabs.com.*