

EMC Centera

Content Addressed Storage



Centera Content Addressed Storage Product Description Guide

Chapter 1: Introduction

Market Overview	2
-----------------------	---

Chapter 2: The Value of Centera

Product Overview	5
<i>Accessing Centera</i>	5
<i>C-Clip Functionality</i>	6
Centera Architecture	7
<i>Continuous Content Availability and Authenticity</i>	9
Centera Management	10

Chapter 3: Centera's Application Programming Interface (API)

Basic API Functions	13
---------------------------	----

Chapter 4: EMC Global Services

EMC Professional Services	15
<i>Operations Management</i>	15
<i>Information Storage Integration</i>	15
<i>Business Continuity</i>	16
EMC Customer Service	16
<i>Remote Support</i>	17
<i>Escalation</i>	17
<i>Change Control</i>	17
<i>Installation Support</i>	18
<i>Post-sale Warranty and Coverage</i>	18
<i>EMC Technical Support and Contact Information</i>	18
Global Technical Training	18
<i>EMC Educational Services</i>	18
<i>EMC Proven Professional Certification Program</i>	19
<i>EMC e-Learning</i>	19

Chapter 1

Introduction



Surveys of Global 2000 enterprises indicate that the fastest growing type of data is so-called “fixed content.” Fixed content comprises any form of digitized information assets retained for reference and value: documents, e-mail, check images, streaming video/audio, X-rays, final-form CAD drawings, etc.

This rapid accumulation of fixed content demands a new category of storage designed for the secure, online storage and retrieval of such information for years and years—Content Addressed Storage (CAS). Rather than access a data object by its file name at a physical location, a CAS device uses a content address to store and retrieve the object. Because content often accumulates without any upper limit, a CAS repository must be extremely easy to scale to a petabyte (1,024 terabytes) while maintaining sub-second access performance.

EMC Centera™ Content Addressed Storage is a purpose-built, hardware and software disk-based storage product that fulfills the challenge of securely holding hundreds of terabytes of business content online for months, years, or decades. Centera complements EMC’s traditional CLARiiON®, and Symmetrix®, block access storage devices that are optimal for transaction-oriented, update-intensive, data store solutions. Similarly, EMC Celerra™ file servers excel at providing file-based access for collaborative applications.

Traditional disk storage systems with block or file access schemes are well suited to the storage of tens of terabytes, typical of transactional and collaborative applications. But today, organizations also need to keep large, rapidly growing repositories of fixed content information online, and this content needs to be accessed by a variety of network attached clients. It becomes a challenge to manage the logistics of data placement and capacity scaling while also assuring authenticity of the content over its life. Centera, with its CentraStar™ software, and advanced modular disk-based hardware architecture, handles the management of the physical storage resources transparently to the user and application. Content management applications can easily interface to Centera facilities to permit painless scaling to a petabyte.

Market Overview

While transactional data stores and applications based on relational database technologies are growing at an annual rate of between 30 to 50 percent, the content explosion is increasing storage capacity requirements in some markets by as much as 100 percent or more. Customers need new ways to manage these increasing amounts of fixed content, which is typically unchanged once it is created, and which can create additional business value when accessible online.

The Content Addressed Storage market cuts across multiple vertical industry segments such as:

- **Healthcare and Life Sciences:** The effectiveness of patient care depends on effective sharing of fixed content information (X-rays, MRI, drug trial results). Centera allows users to easily store, retrieve, and preserve content so regulatory requirements are adequately satisfied. Faster access to information helps improve the quality of care.
- **Financial Services:** Centera addresses two major needs: 1) stringent regulations that mandate content integrity, and 2) cost-effective online access as a way to repurpose information for new service offerings and revenue streams.
- **Enterprise Content Management:** Across the enterprise, from e-mail messages and attachments to Web transactions, fixed content is growing exponentially. Centera can help users manage this growth—without disruption.

- **Film, Broadcast and Media:** The creation, reuse, and sale of video, film, and audio assets are key revenue sources that must be protected. Centera can be used as a digital asset repository to address the issues of long-term retention and content authenticity for copyright protection.
- **Applications for all industries:** Centera addresses significant challenges presented by the growth of fixed content in Federal and State governments, telecommunications, discrete manufacturing, oil and gas, and more.

In each of these market segments, content must be preserved intact for years, if not decades. Particularly, in regulated industries, such as financial services, life sciences, healthcare, and government, data must not be altered or erased within the legally defined retention periods. In the past, data has often ended up on tape or optical disk. These are sub-optimal solutions based on inadequate access speed and unverified reliability and integrity. The solution needed is EMC Centera, Content Addressed Storage (CAS).

The following examples illustrate the need for Centera CAS solutions:

Medical imaging: Over 400 million patient studies were completed in the U.S. in 2000. Each study is composed of one or a series of images, which range in size from about 15MB for standard digital X-rays to over 1GB for oncology studies. As X-rays are created in the radiology department of the hospital, they are stored online for immediate use by attending physicians for a period of approximately 60-90 days. At the point where the patients are cured or discharged, the access needs for their particular X-rays drop off dramatically. However, HIPAA requirements stipulate that these studies must be kept in their full glossy image formats for a minimum of seven years. Beyond 90 days, hospitals may back up images to tape or send them to an offsite archive service for long-term retention. The cost of restoring or retrieving an image when in long-term storage could be 5-10 times more expensive than leaving the image online, and recovery times can often be measured in hours or days. Medical image solution providers offer hospitals the capability of viewing medical studies such as X-rays online with sufficient response times and resolution to allow rapid assessment of patient situations. Centera is the optimum and most cost-effective storage solution to facilitate long-term storage and immediate access of medical images online, within a hospital or at a clinician's office.

Financial Services: Check images, each with a size of about 25KB, are created at the bank, and they can be stored internally or sent to a check imaging service provider who may process 50 to 90 million check images a month. Typically, check images are actively processed in transactional systems for about five days. For the next 60 days, check images may be requested by regional banks or individual consumers for verification purposes, at a rate of about one-half percent of the total check pool (250,000 to 450,000 requests for check look-ups). Beyond 60 days, access requirements drop dramatically to as few as one for every 10,000 checks. In this case, the check images would be stored on Centera starting at day 60 and held there indefinitely. It should also be noted that check imaging is not the only financial services application requiring the content storage facilities of Centera. Customer transactions initiated by e-mail, contracts, and security transaction records also need to be kept online for as long as 30 years.

Research shows that customers want higher reliability and faster access to fixed content than current solutions (tape or optical) provide. Many of the technologies currently addressing the content market do not provide immediate, online access to the information. The customer accepts a trade-off, giving up functionality (typically response time) in order to achieve a lower initial acquisition price. In some tape and optical-based applications, multiple concurrent users accessing content is cumbersome to implement and creates bottlenecks.

Although tape and optical solutions may appear less expensive due to lower initial acquisition costs, in the long run they cost more as they require increased focus on manual media management (movement of tapes, conversion to new formats, etc.). And they are more difficult to manage through migrations and changes in technology. Likewise, when dealing with large quantities of fixed content, low-end NAS or SAN alternatives also create management challenges because they do not have the content addressing and self-healing/self-managing features of Centera.

Chapter 2

The Value of Centera

The Centera approach to storing fixed content offers the end user many benefits:

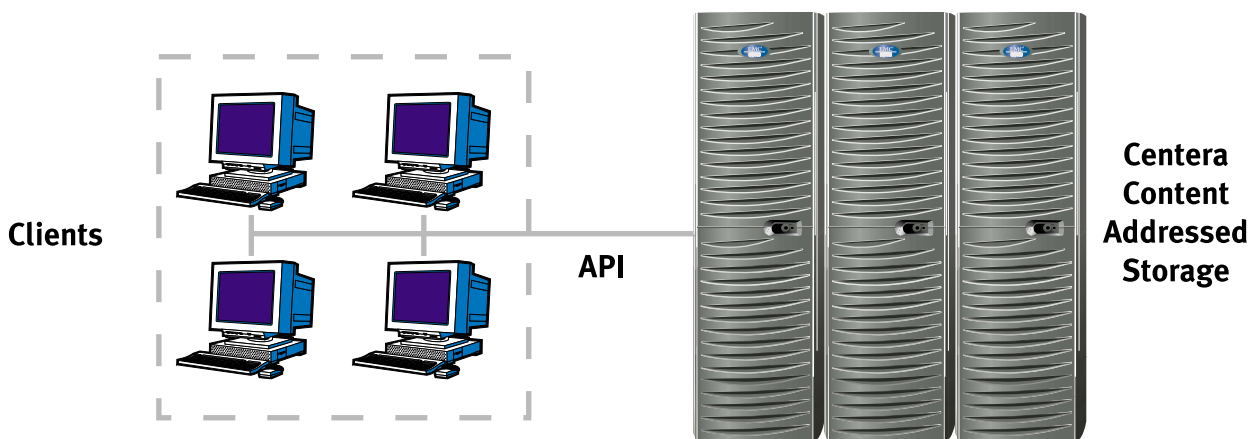
- **Ease of Management:** Centera technology greatly simplifies system planning and management of hundreds of terabytes of content storage. With no RAID types to choose, LUNs to bind, or file systems to create, customer applications are released from the requirement to compensate for, or manage, the complexity of traditional storage topologies.
- **Content Authentication:** Any object presented to the system is stored in such a way that it is unchangeable and authenticated, which is transparent to the end-user application.
- **Non-eraseability:** Data objects cannot be deleted prior to the expiration of their particular retention periods.
- **Efficient Replication:** As you'll see, Centera uses a unique content-derived address to ensure that only one copy of content is stored (and a replica for reliability) no matter how many times it is used. This can significantly reduce the amount of information a customer stores and is a key factor in lowering the TCO for storing content.
- **Scalability without Reconfiguration:** Centera, based on redundant arrays of independent nodes (RAIN), is architected to be highly scalable (terabyte to perabytes).
- **Easily Installed and Upgraded:** A significant effect of this breakthrough technology is that systems are able to be installed or upgraded in an hour or less.

Product Overview

Accessing Centera

Centera introduces a new data access paradigm known as Content Addressing. Instead of requiring an application to store and retrieve static content in the traditional address location-based approach, Centera presents a flat address scheme to the archive program. When a reference data object is initially stored in the Centera repository, the application is given a "claim check" that is uniquely derived from the object's content.

Figure 1

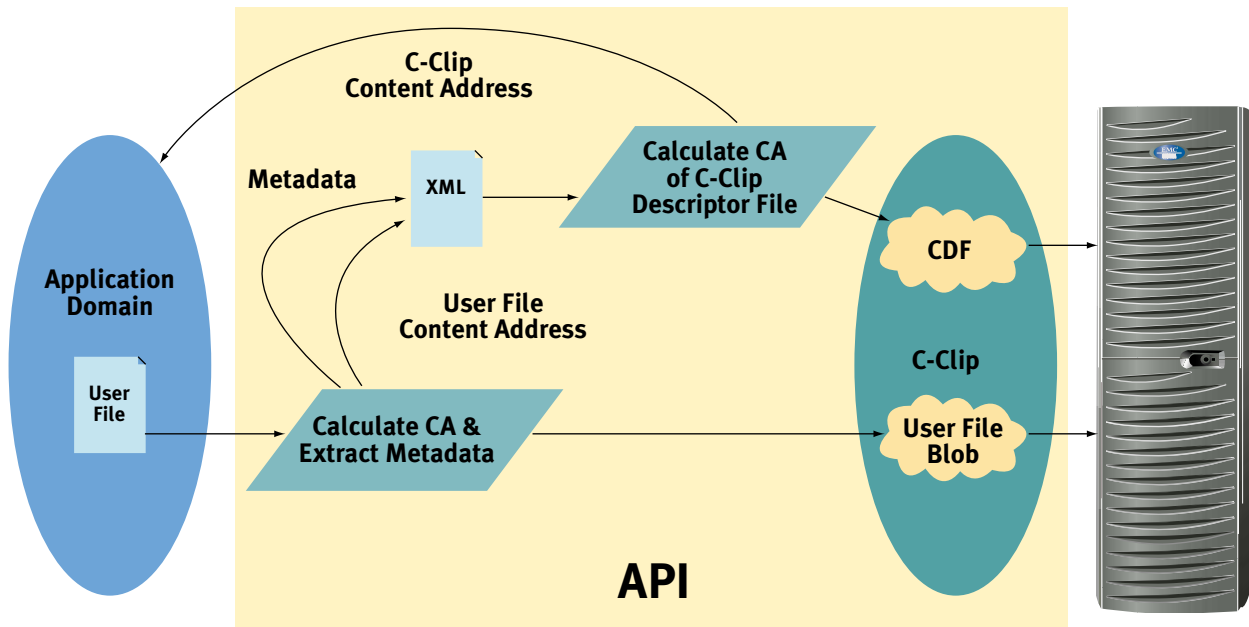


Subsequent access of that data object is made by simply giving the claim check that uniquely identifies the object back to the repository and the data object is returned. Content addressing greatly simplifies the storage resource management tasks, especially when handling hundreds of terabytes of static objects.

Centera stores a data object, also known as a BLOB (Binary Large Object) synchronously as shown in Figure 2.

1. An application delivers a data object to the Centera Application Programming Interface (API), which calculates a 128-bit claim check, i.e., the Content Address (CA), from the object's binary representation.
2. The Centera repository then stores the BLOB and a mirror copy.
3. This CA, which is unique for the object's content, and metadata about the object (e.g., file name, creation date, etc.) are then inserted into an XML file, called a C-Clip™ Descriptor File (CDF), which in turn has its content address calculated. Note: The C-Clip is the union of the CDF and its content object.
4. This C-Clip CA is only returned to the application once two copies of the CDF and two copies of the BLOB have been safely stored in the repository.

Figure 2



Future requests for the retrieval of the data object occur when the application submits the C-Clip's CA for that object to the repository via the API. There is no centralized directory in Centera and no pathnames or URLs are used. Only the C-Clip's CA is used as a reference. Where the data is physically stored in Centera is transparent to the application.

C-Clip Functionality

The CA is essentially a fingerprint that assures the authenticity of the stored object (i.e., the User File BLOB). If an object is retrieved and altered by one bit, when the modified object is subsequently stored, the Centera API will produce a new CDF with a new content address for the altered BLOB. The original object is still unchanged and accessible by its original content address. This Write Once Read Many (WORM) attribute assures a level of versioning integrity that file servers can't provide. Additionally, Centera features an operational mode where data cannot be deleted prior to the expiration of a defined retention period. These non-rewritable and non-erasable properties of Centera give it WORM attributes without sacrificing online performance.

Another benefit of the C-Clip access method is that if multiple clients store the same BLOB in the Centera repository, only one instance of the object (plus a mirror copy) will be stored while each

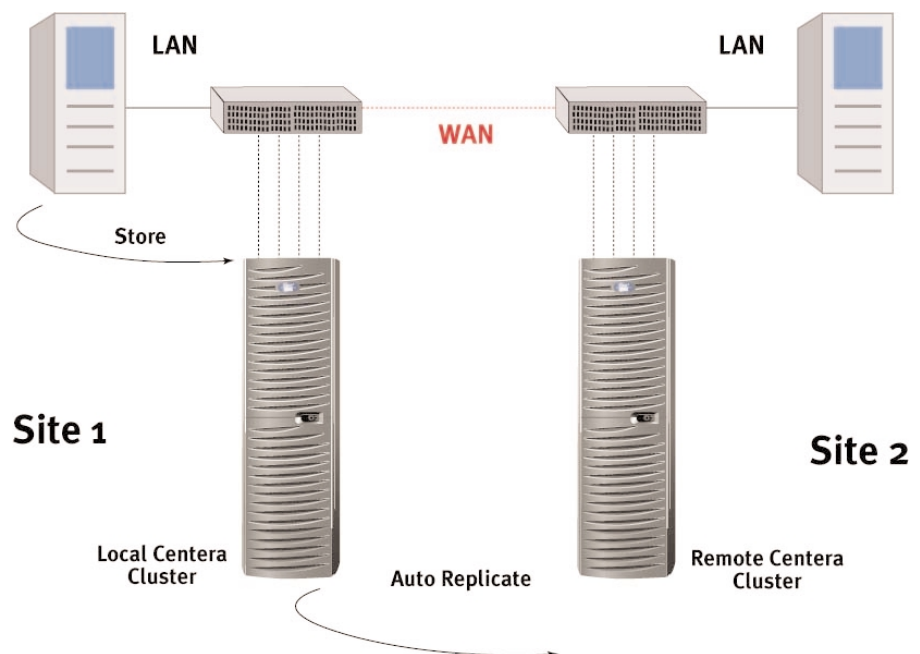
referencing client gets its own C-Clip CA pointing to the common object. This is very useful in optimizing space consumption in applications such as e-mail attachment archiving.

The C-Clip method ensures that application developers, users, and storage managers do not need to think about where data is physically or logically located. The C-Clip's CA is a globally unique identifier permitting a data object to be retrieved from anywhere, exactly as it was stored, irrespective of the relative locations of the content and the user.

Because Centera uses the location independent addressing scheme described earlier, the result is data mobility that facilitates a simple yet robust disaster recovery topology. When a data object is initially stored in the local Centera, the object can also be asynchronously and automatically replicated to the remote site over a wide area network (WAN) such that the object resides both locally and remotely. Replication is offered as a CentraStar software option to be configured on both local and remote Centera systems.

If a disaster should occur locally, when the local site is rebuilt the remote site will be able to re-populate the local Centera as needed. Note that this replication facility can be operated bi-directionally in an active-active topology.

Figure 3



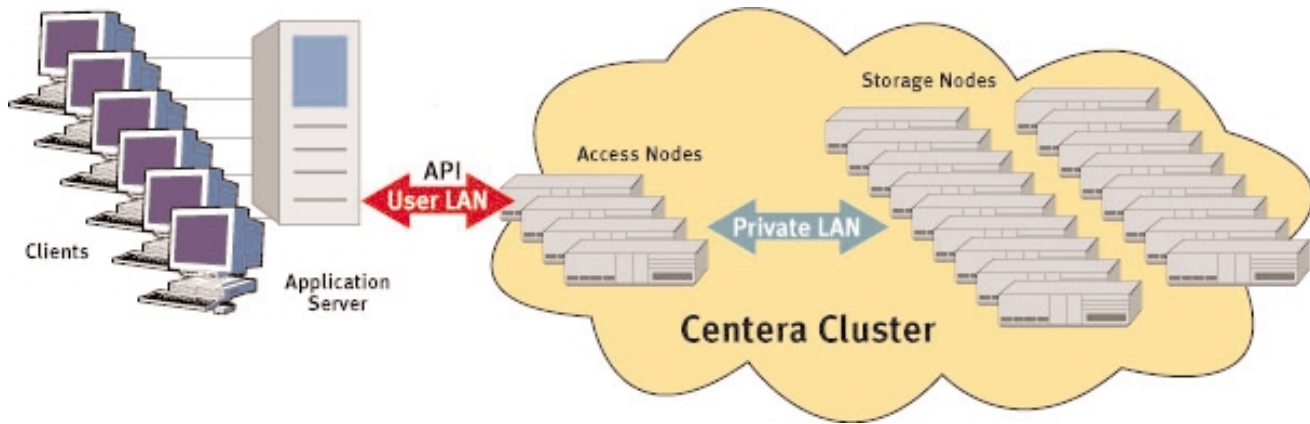
Users access EMC Centera through content-based software applications that incorporate Centera's powerful API. Examples of the type of third-party software that use the Centera API include content/document management, medical imaging, e-mail archiving, and a wide variety of vertical applications that benefit when large amounts of content objects are accessible online and shareable by a large number of users.

Centera Architecture

The Centera architecture presents a “no-single-point-of-failure” platform that is highly scalable and implements non-disruptive servicing. Centera is built upon a Redundant Array of Independent Nodes (RAIN) that are deployed in one or more six-foot NEMA standard 19-inch racks. A single rack can hold 8, 16, 24, or 32 nodes to provide 7.6TB, 15.4TB, 23.0TB, or 30.8TB of raw capacity, respectively. Each node contains processing power, 960GB of raw storage capacity, and is interconnected with all other nodes in the cluster via a private LAN. Each node executes an instance of Centera software in one of two operational modes:

- Storage Node: the node facilitates long-term storage of BLOBs and CDFs
- Access Node: the node is the conduit for interaction between the application server and the storage nodes.

Figure 4



The throughput needs of the application (e.g., medical imaging) will determine how many front-end (or access) nodes must be configured at the time of installation. Each access node is connected to the application server infrastructure via a 100 megabit-per-second Ethernet cable. Given that Centera will be connected to the application server via multiple access nodes, the application will enjoy not only scalable bandwidth with low access latency, but also high availability.

Note that as access nodes do not provide long-term storage for content, the “usable” protected storage of a Centera cluster is determined solely by the number of storage nodes. For example, if a 32-node cluster has four access nodes and 28 storage nodes, the usable capacity will be 20TB.

The usable capacity of a specific deployment is a function of the number of storage nodes and the user-selected means of protection: either Content Protection Parity (CPP) or Content Protection Mirroring (CPM). Mirroring protection places two physical copies of a data object on two different nodes within the cluster. (CPM mirrors data objects as opposed to mirroring a disk or a node.) Therefore, CPM utilizes 50 percent of the raw capacity of the storage nodes to facilitate usable data object storage. CPP employs a more space-efficient parity algorithm. Any data object greater in size than 20KB will be broken into six separate data fragments, and a seventh parity fragment will automatically be created through an “exclusive OR” calculation on the six data fragments. Each of the seven fragments are then stored on separate storage nodes. If a node or disk were to fail, the missing fragment can be recreated on the fly by using the same exclusive OR process. From a usable capacity perspective, CPP utilizes approximately 75 percent of the raw storage capacity, depending upon the mix of sizes of data objects being processed.

Content Protection Parity, like CPM, executes at the data object level and not the disk or node level. Therefore, when a device fault occurs, CentraStar software can employ the power of many storage nodes in parallel to regenerate the missing fragments of the impacted data objects. In contrast, in a RAID (Redundant Array of Independent Disks) device, only a handful of peer disks of a failed disk would be heavily exercised simultaneously to recreate the missing data.

CPP is available for all Centera cabinet configurations, except the eight-node entry system as two of the eight nodes will be configured as access nodes, leaving only six storage nodes (CPP requires

seven storage nodes). Also, it is important to note that in applications involving average object sizes < 1MB, mirroring will offer higher object/hour throughput rates for store and retrieve operations as compared to parity protection.

Centera cabinets can be ordered in four configurations with resulting usable capacities shown below for either CPM or CPP:

# Nodes	# Access Nodes	Raw TB	Protected TB		Usable TB	
			Mirrored	Parity	Mirrored	Parity
8	2	7.7	3.8	na/	2.9	n/a
16	4	15.4	7.7	11.5	5.8	8.6
24	4	23.0	11.5	17.3	9.6	14.4
32	4	30.7	15.4	23.0	13.4	20.2

- (1) Usable capacity = protected capacity without access nodes
 (2) Parity assumes approximately 75% of raw capacity is utilized

This architecture gives Centera its extreme scalability for capacity and performance as well as its radically unique simplicity of managing physical storage resources. Capacity scaling is accomplished merely by non-disruptively adding nodes in 8TB (raw) chunks. As capacity is scaled, Centera's performance is also scaled to handle the management of data across the incremental nodes. Additionally, the number of access nodes can also be augmented to facilitate more bandwidth to the application server.

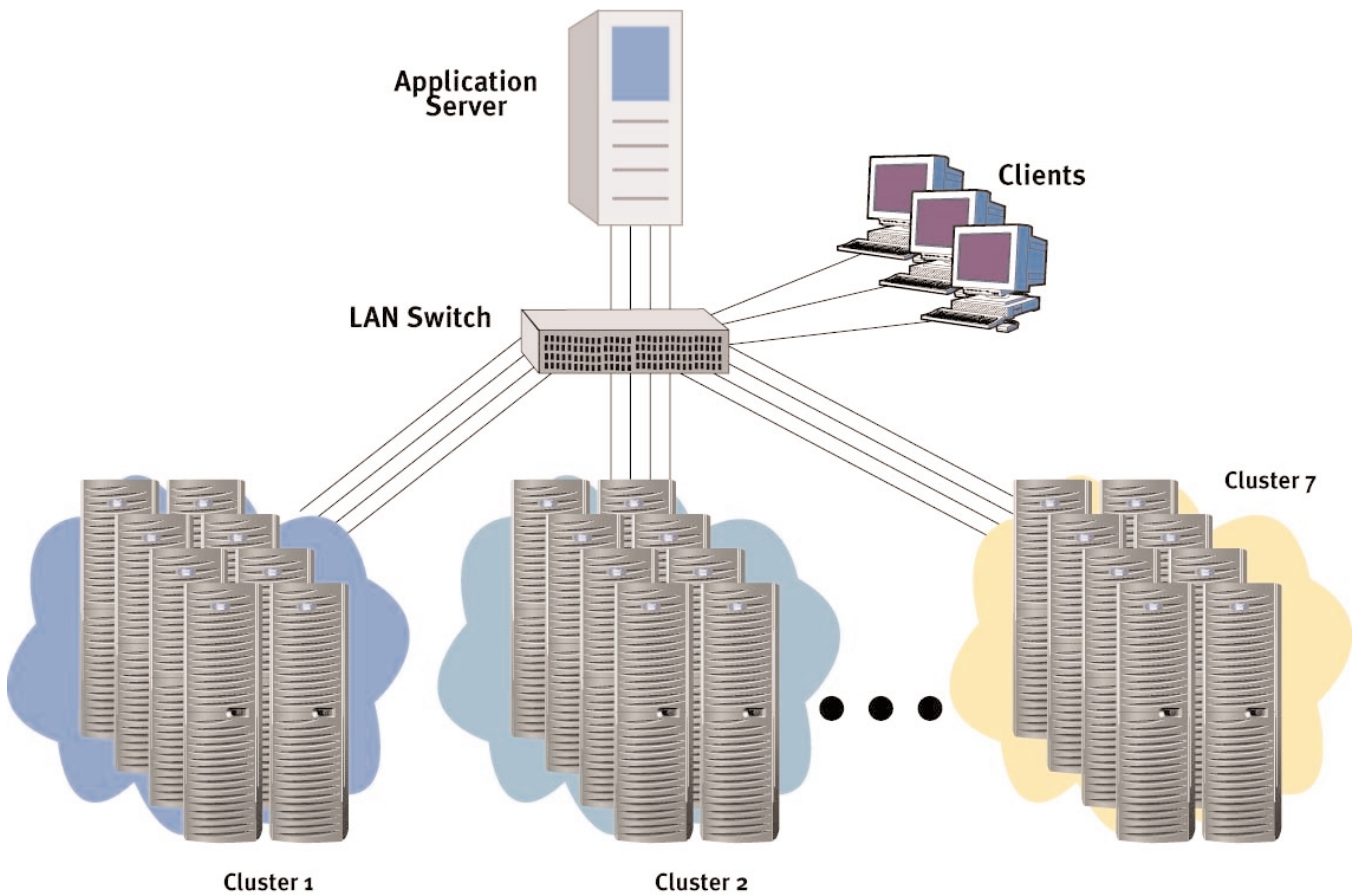
Perhaps the most significant benefit of Centera's architectural implementation is that the addition of nodes requires no time-consuming, complex management effort on the part of either the application (or developer) or the system administrator. This benefit permits a single administrator to handle hundreds of terabytes of content storage—far greater capacity than can be managed in traditional storage products.

Multiple Centera racks can be configured as a single "cluster" delivering up to 180TB of usable storage. The Centera repository connects directly to one or more Windows or UNIX application servers via multiple TCP/IP LAN connections. If 180TB is insufficient, the client application API can reference several clusters, as depicted in Figure 5, to access up to one petabyte of data. Building such a network of clusters is simply a matter of non-disruptively connecting new clusters to the application server's LAN infrastructure as older clusters are filled.

Continuous Content Availability and Authenticity

Centera features complete redundancy to protect against component faults that will inevitably occur. Every data object is mirrored. Should there be a disk drive failure, Centera will "self-heal" by detecting the fault and generating new mirror copies of content objects. Failed components can be replaced non-disruptively.

Figure 5



The Centera cabinet can be powered by two different AC power sources to permit the application server to have continuous ability to store and retrieve objects even if one AC power source is lost. In environments where AC power loss occurs momentarily and with some frequency, the customer is advised to power Centera with third-party uninterruptible power supplies (UPS) to ride through such outages with no loss of function.

The management of storage resources has been a primary cost factor, especially when content repositories grow to hundreds of terabytes. Centera makes quantum leaps in ease of management through its introduction of the concept of organic management (i.e., self-management), its RAIN architecture which is dynamically expandable when more storage is required, and with self-healing that occurs after hardware failures. When new storage units are added to the cluster and powered-on, they are automatically “auto-discovered” and join the cluster.

Using CentraStar™, Centera’s operating software, when drives fail, they are “fenced off” from the cluster by the system and affected objects are regenerated to ensure that a fully redundant mirror copy of the content is always available. In the event of hardware failures, the user data remains available at all times. No power down or restore is required. Additionally, data integrity checking is run in the background. This task continuously recalculates content addresses of all the objects and compares the calculations to the content addresses originally stored in the CDF. The existence of protected copies of objects in the cluster is also verified.

Centera Management

Managing Centera operations over its multi-year life is extremely simple as the Centera software, CentraStar, does the work of handling the logistics of storing and retrieving data objects. CentraStar delivers storage, retrieval, and network-awareness intelligence. It provides the Centera system with the capabilities of self-management, auto-configuration, self-healing, non-disruptive

maintenance and upgrades, and content replication. The software technology that determines and assigns a unique address to every stored object is also part of CentraStar.

Centera's Content Addressing scheme permits easy capacity scaling of the repository while also enabling stored objects to be mobile across a user's content infrastructure. The Content Address, a unique identifier to a data object, assures the authenticity of the referenced information, no matter where in space and time the object is accessed. Retrieval requires no knowledge of the storage environment or physical location of the objects. Those relationships are maintained by the CentraStar operating environment.

The system administrator doesn't need to worry about volume creation and management or file system structure and maintenance. All the system administrator needs to do is monitor Centera's available capacity and its object store/retrieval performance. The system administrator uses a GUI, Centra Viewer, that provides a simple means of displaying capacity utilization and operational performance. Centra Viewer runs on any Windows PC that is LAN-connected to Centera's access nodes. Centra Viewer also permits the system administrator to change any site-specific information such as the Centera's public IP addresses and subnet masks and end-user contact name and e-mail address.

The system administrator can also enable the use of Simple Network Management Protocol (SNMP) to alert an enterprise network management system as to any faults that might occur within normal Centera operation. A "green" status signifies all is well with Centera. A "yellow" status means a self-healing fault has occurred, such as a disk or node fault, but because regeneration is underway, the Centera is fully functional. A "red" indication would mean that a fault requiring immediate attention has occurred, such as a loss of a network device that could not be self-healed.

As mentioned previously, if more capacity is needed, it is simply a matter of either installing more storage nodes within a cabinet if there is sufficient rack space, or adding another cabinet of storage to the existing cabinet(s) via two gigabit Ethernet LAN cables—true "plug-n-play" storage. Centera Content Addressed Storage represents a new and radically simpler means of storing hundreds of terabytes of fixed content online, providing fast access time at essentially zero management cost.

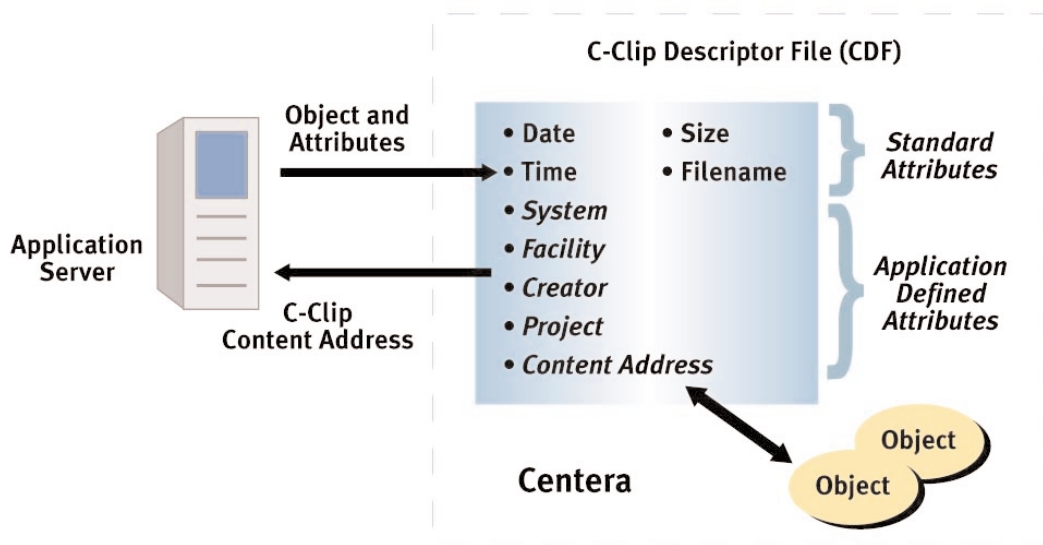
Increasingly, system administrators require the ability to integrate a Centera into a secured network environment. Centra Viewer provides access to a pull-down CLI menu that permits a system administrator to optionally assign a login and password for an application community hosted on a server that will be verified every time that application community accesses the cluster to store or retrieve data. If the login and password don't verify, then access is denied as the access attempt is judged to be illegitimate. The system administrator can also define a set of authorized functions (store, retrieve, delete, query, etc.) that a specific application community is permitted. In this way, the end-user system administrator can facilitate secure access to the Centera repository, especially in a deployment where multiple application communities are involved.

Chapter 3

Centera's Application Programming Interface (API)

As was described previously, application server access to the Centera repository is facilitated by use of a simple, robust set of API commands combined with a powerful “claim-check” mechanism, the C-Clip. A C-Clip Descriptor File (CDF) is an XML description for the referenced data object. It contains not only the content address pointer(s) to a particular data object(s), but also standard system-generated metadata and any application-specific metadata associated with the object (e.g., a C-Clip pointing to an engineering document may include system name, facility location, creator, project name, etc.).

Figure 6



A content-based application will typically use a database to connect “owners” of content to their specific collections of images, documents, etc. (e.g., X-rays, check images, scanned mortgage contracts, and more). The database would be indexed by account number or medical record number and would contain fields pointing to the various data objects. It is very easy to adapt such an application's database to Centera's API by simply using the C-Clip content address as the pointer to specific objects. When a patient's X-ray is stored, its C-Clip CA is then inserted into the database as a field in his/her medical record database. When the physician needs to review the patient's X-ray, the application simply uses the C-Clip CA claim check entry in the patient database to retrieve the X-ray.

A very powerful feature of the C-Clip Descriptor File is that it can store metadata generated by the application for each stored object by use of XML tags. Each tag will have a name designated by the application and can have one or more attributes. As an example, suppose Centera is used to store all the recorded music of a musical artist in MP3 format. Each song would be a separate MP3 file that Centera will store as a BLOB; each BLOB will have its own unique content address. The song title could be stored in the CDF as part of the application metadata.

Suppose this music library spanned many albums, such that a number of highly popular songs would have shown up on a “greatest hits” album as well as the original album. Centera's content address scheme would ensure that such songs would only be stored once, but could be referenced by both an original album CDF and a greatest hits album CDF.

In applications where data must be retained in its original form for a specified number of years, typically driven by regulations, a retention period value can be specified in the metadata of each object. Centera software will enforce the retention period by preventing the data object's deletion prior to the expiration of the retention period.

Another example demonstrating the storage efficiency of Centera is the archiving of e-mail attachments. If 25 people received a multi-megabyte presentation, Centera would store only one instance (but mirrored for continuous availability) of the presentation with 25 CDFs pointing to the attachment object linking each addressee with the common attachment object.

Basic API Functions

There are five basic API functions: Store (Write), Retrieve (Read), Exists, Delete, and Query. The Store and Retrieve functions are self-explanatory and are arguably the most commonly used API functions. Exists is used by the application to check to see if a particular object is already stored in a Centera repository. If an object is already stored, there is value in not consuming time and network cycles by transmitting the BLOB a second time over the LAN connection between application server and Centera, especially if the BLOB is very large. The second attempt to store an object should simply cause a second CDF to be generated for the new reference to the already stored BLOB. Delete is an important function in at least two situations. In one instance, an application may need to deterministically remove a data object that must be expunged. Or, the end user might need to delete objects that meet some criteria (e.g., objects older than seven years) to reclaim the capacity for storing new content.

As mentioned previously, in regulated environments where data items have legally mandated retention periods, Centera API permits compliant applications to store each data object with a specified time to delete (generally multiple years). All attempts to delete such protected objects will fail until the retention period expires giving Centera non-erasability.

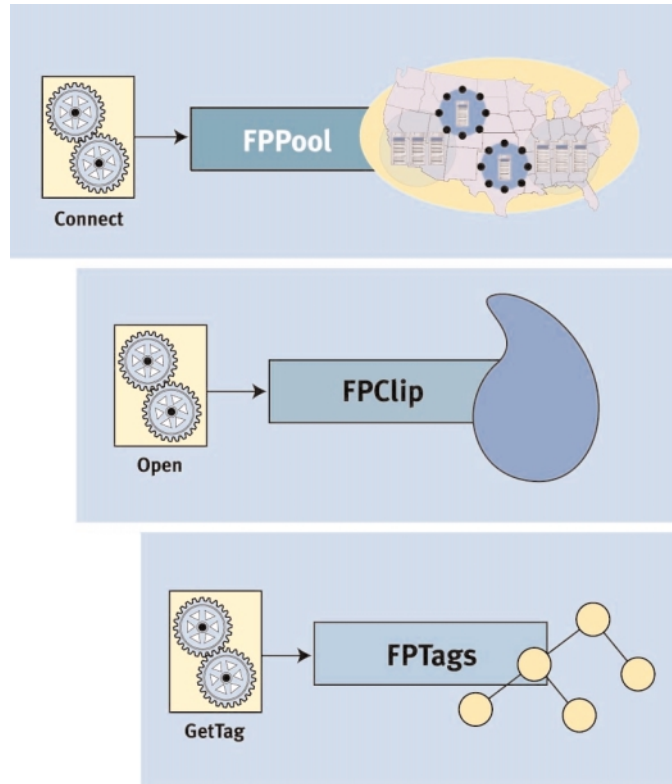
Once all CDF pointers have been deleted, the space occupied by the data object will be reclaimed by a background process known as "garbage collection." Optionally, for privacy or security reasons, the deleted data may be overwritten ("shredded") as the space is reclaimed.

Query permits the application to determine what data objects have been stored in the Centera over a specified time interval.

EMC offers a Software Developer's Kit for Centera that allows the application developer to select one of two basic approaches. The fastest and simplest API adoption approach is to use the high-level toolbox functions, referenced above (which are aggregates of lower-level APIs) that permit an application to be quickly and easily modified to execute basic store and retrieve functions. The toolbox permits C-Clip storage (or retrieval) to be accomplished in one single ANSI C function rather than via multiple API function calls.

The second, more flexible and powerful approach is to use the low-level API functions themselves. This approach permits the application to interact with Centera at the pool, clip, and tag level.

Figure 7



Functions at the pool level include opening a connection to Centera cluster (the “FPPool”) to permit storage or retrieval of an object. When storing an object, the application would then create a new, empty C-Clip structure (“FPClip”), followed by the opening and creation of a tag (“FPTag”). The object would then be stored in the repository via an FPWrite function, at which point Centera would return the object’s C-Clip CA to the application. The application then terminates the object write by closing the C-Clip and then closing the pool.

This is a very simplistic overview of a write transaction. There are a number of access modes available to the application developer, as well as a number of “returns” that the application must be prepared to handle. The low-level API functions are implemented in both ANSI C and Java and the Centera API is thread-safe.

Centera’s API is offered as a DLL that is supported on Windows NT and Windows 2000, and a library file for Solaris, AIX, HP-UX, IRIX, and Linux. For more information, please refer to the Centera Programmer’s Reference Guide.

Chapter 4

EMC Global Services

To help you realize the full potential of your networked storage infrastructure, EMC Global Services puts an industry-leading network of interlinking services at your disposal. This network provides a continuum of “best-in-class” services that support you throughout your complete information life-cycle—as you acquire, as you deploy, and as you grow your infrastructure.

Key Global Services components include:

- Professional Services: surrounds you with strategic planning, implementation, and operations management services
- Customer Service: available around-the-clock to optimize the efficiencies of your E-Infostructure
- Global Technical Training: educates and empowers you with the skills and knowledge needed to self-sufficiently manage and leverage your E-Infostructure

These components, while independent in areas of expertise, complement each other so customers achieve competitive advantage. Your business success through your EMC E-Infostructure is our sole concern, and we do whatever it takes to attain 100 percent customer satisfaction.

Professional Services

As a company, EMC has focused, information storage experience, unparalleled in the industry. Within Global Services, the Professional Services organization represents more than 1,200 Professional Services experts, boasting an average of 20 years of industry-related expertise overall. EMC has categorized our areas of focus to areas we know are important to you. Each solution set includes multiple offerings ranging from consulting, planning, and design services, to consolidation, migration, and operations management.

Operations Management Consulting

EMC Professional Services analyzes a customer’s current testing processes and methodologies. Based on the customer capabilities, Professional Services recommends improvements in the areas of human resources, process, and technology. EMC Professional Services consultants can also be engaged to go onsite at the customer’s location to augment in-house staff. These expert consultants facilitate knowledge transfer to the customer’s staff until the staff is properly trained and able to manage the environment on their own.

An Operations Management engagement helps you address the following business challenges:

- Meeting the demands of instituting effective storage management practices
- Understanding storage management is instrumental to leverage advanced software features
- Maintaining control of data while receiving the assistance needed to manage it effectively
- Creating storage management self-sufficiency, the most cost effective management model

An Operations Management Consulting engagement consists of the development of a storage management strategy. This includes defining the best approach to managing infrastructure operations, developing processes and metrics, and the selection, implementation, and integration of storage management products.

Operations Management Consulting offerings include:

- Operations Management Planning
- Operations Management Design and Implementation
- Operations Management Support

Information Storage Integration During an Information Storage Integration engagement, Professional Services presents you with a comprehensive storage systems strategy, based on the unique information gleaned from an in-

depth consultation. From there, the conceptual storage architecture and design is identified, and a plan for the installation and integration of storage infrastructure products is created. Information Storage Integration services assists you in building a storage architecture foundation to meet current and future business requirements.

As you continue to focus on decreasing your storage costs, you will typically experience increasing data requirements driven by new applications and increasingly complex storage architectures. Members of the EMC Professional Services team are experts in managing and protecting information storage, and they can give you the knowledge and assistance you need to meet those challenges.

Information Storage Integration offerings include:

- Storage Infrastructure Strategy and Planning
- Storage Infrastructure Design and Implementation

Information Storage Consolidation

EMC Professional Services provides all services, including project management, for a full data center migration project. This engagement uses an EMC best practice methodology called “Eccelerate,” which includes risk analysis, planning and design, implementation, and management. There are two offerings Professional Services will deliver in an Information Storage Consolidation engagement:

- Information Storage Consolidation Planning develops and delivers a strategy and architectural approach that addresses consolidation needs at a variety of levels. The strategies range from specific application consolidations to enterprise wide multi-environment situations. The Professional Services consultants delivering this offering review the current storage architecture and technology and identify current and future storage requirements for scalability and opportunities to leverage operational efficiencies. An information storage consolidation strategy and architecture is developed to support both customer business objectives and IT requirements. Technical components are recommended, and high-level deployment and migration plans are developed. Overall cost and ROI are calculated and a formal business case is proposed to support the consolidation effort.
- The Data Migration offering provides the detailed planning and physical migration of data from one storage environment to another. This effort could involve a small- to medium-data set migration that is moving data to a new storage base or a full data reorganization in a data center.

Business Continuity

The Business Continuity services that EMC Professional Services consultants offer include the development of a comprehensive plan for data availability—the availability of storage assets as part of a business continuance initiative. Included in the plan is a strategy and tactical details on the implementation of data replication devices and hierarchical storage technologies. Customers receive a proposal based on the unique needs of their businesses that reveals the financial impacts of not having an effective, business continuity plan in place.

EMC understands the challenges customers face delivering realtime data, and we also acknowledge the need to reduce cycle time for backups. EMC provides comprehensive business continuity solutions that meet multiple levels of availability and business continuance requirements. EMC Professional Services has the experience and expertise to help determine the right solution for each customer. Moreover, we’ll implement that solution quickly and cost-effectively.

Customer Service

EMC Customer Service (CS) provides optimum availability of customers’ information. EMC CS delivers service excellence through a number of key attributes and functions, including:

- A global organization with 5,500 field and support center staff
- A proactive infrastructure based on resolving potential issues before they are problems

- 24x7 coverage
- Technical expertise in hardware, software, and multi-vendor environments
- Advanced remote support
- Comprehensive e-services
- Best practices change control

A large part of EMC CS success is its investment center philosophy. EMC CS is not measured on a profit and loss basis as most organizations. Rather, it is measured purely on the customer's level of satisfaction. This structure creates a culture of unyielding service and customer commitment. All CS staff, from the support center to the field-based customer engineer and regional service manager, has one common mission—do whatever it takes to meet customer expectations.

Field staff provide onsite ongoing maintenance and problem resolution as necessary. An EMC customer engineer (CE) will be sent to your site with parts or equipment to remedy any issues. While the CE is en route to your site, the support center continues to analyze the system remotely to pinpoint the cause of the issue, providing additional information to the CE for use upon arrival at the site.

Additionally, EMC has self-help, e-services capabilities via the EMC Powerlink website. This robust, Web support engine searches the EMC Knowledgebase for known solutions to frequently asked questions and allows customers to open, update, track, and resolve service cases online via the Web.

Remote Support

EMC's sophisticated remote support capabilities provide the support you need, before you even know you need it. We recognize that constant information availability is critical to your business success. With our highly advanced remote support, your EMC system is never alone. Through automatic status notification, the EMC system permits our service experts to monitor your installation 24-hours a day, 365-days a year. And by dialing back into the EMC system, we take action quickly, analyzing events and abnormalities in your system, and resolving the majority of issues before they have a chance to affect your business. This proactive and pre-emptive service approach is unmatched in the industry.

Escalation

EMC's Global Technical Support organization has an automatic escalation policy to ensure that every service call receives prompt attention. The escalation and notification procedure is an integral part of the EMC customer service philosophy. It outlines escalation contacts and guidelines, and translates error codes to priority levels. This provides EMC support staff with a clear and concise plan of action with definitive milestones to resolve service issues and ensure rapid response and resolution. Escalation is based on the impact an issue has on the customer, and not on a specific technical or non-technical problem.

Examples of EMC Severity Levels:

- Severity Level One—Severe Customer Business Impact: Customer or workgroup cannot perform normal job functions
- Severity Level Two—Moderate Customer Business Impact: Customer or workgroup performance of job function is degraded
- Severity Level Three—Limited Customer Business Impact: Customer or workgroup performance of job function is largely unaffected

Change Control

EMC's industry-leading change control process enables customers to take full advantage of the outstanding connectivity, flexibility and upgrade capability engineered into every EMC system. Every change made to your IT environment introduces risk. EMC's change control process mitigates that risk so you continually realize the full value of your E-Infostructure. Our dedicated change control team meticulously plans and orchestrates changes to your EMC solution. From

standard microcode upgrades to massive data center relocations, our disciplined process ensures that your customized changes are implemented correctly and timely the first time.

Installation Support

EMC specialists install and configure the Centera system according to a customer's needs in a process that typically takes less than two hours. Centera installation is a matter of unpacking the cabinet, moving it to its operating site, cabling power, modem, and LAN connections, and then booting the unit. The user-supplied network connection and site-specific information is then entered into the Centera. The final step is to verify that the customer can store data to, and retrieve it from, the Centera.

Post-sale Warranty and Product Support

Coverage of the Centera system includes a two year Scheduled Service Warranty on hardware, as well as an upgrade option to Premium Response Service. Scheduled Service provides for free installation and configuration and maintenance visits every six months to replace faulted capacity. The Premium Response Service option extends the coverage of the Scheduled Service Warranty by ensuring that an EMC specialist will make service visits whenever any component fault is observed. Coverage of installed Centera systems is provided on a 24-hour, seven-day-a-week basis.

Both the Scheduled Service Warranty and the Premium Response Service upgrade can be extended beyond their two-year terms. Centera software warranty is 90 days and can be extended in one-year terms. Software warranty and post-warranty maintenance include software upgrades for the purchased software.

EMC Technical Support and Contact Information

EMC has multiple globally dispersed support centers, including dedicated software assistance centers and a Microsoft Competency Center, which are available 24x7. The support centers are also connected via an intelligent call-management system enabling EMC to leverage its worldwide pool of resources and provide customers fast and effective problem resolution. Use the following numbers to contact EMC and obtain technical support:

US: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Other locations: 1 + (508) 497-7901, or contact the nearest EMC office

Global Technical Training

The EMC Global Technical Training department has developed a comprehensive customer education program. The educational offering is a blended solution of e-learning and instructor-led courses. Students can attend courses at one of EMC's training facilities, or if a particular customer has a large group of employees that require training, instructors will provide training at the customer site.

Developing in-house talent is beneficial because it creates a self-sufficient staff and reduces reliance on service calls for problem mitigation and resolution. Although focused on providing students with the knowledge and skills needed to operate and maintain the E-Infostructure, the courses also teach students how to use EMC products at optimum efficiency and to leverage system capabilities in the overall IT environment.

Educational Services

Both e-learning and traditional instruction are available. Our Web-based program offers access to training whenever it is convenient for you.

***The EMC Proven Professional
Certification Program***

The EMC Proven Professional Certification Program is aligned with other IT industry certification programs, notably Microsoft and Cisco. Students can achieve an Associates or a Masters level of certification in the Proven Professional program. Four tracks are offered, based on IT job roles: Operator, Builder, Architect, and Instructor.

- Operator: manage data center operations
- Builder: implement and integrate data centers
- Architect: design enterprise storage networking solutions
- Instructor: knowledge transfer of networked storage infrastructure

E-learning

EMC e-learning incorporates online learning into the suite of training, education, and certification solutions available to customers, partners and employees.



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000

In North America
1-866-464-7381

EMC², EMC, CLARiiON, and Symmetrix are registered trademarks and C-Clip, Celerra, Centera, ContraStar, and where information lives are trademarks of EMC Corporation. Other trademarks are the property of their respective owners.

© 2003 EMC Corporation. All rights reserved.
Published in the USA. 3/03

C938.3

Product Description Guide