Archiving Cries for a Holistic Architecture

2008 EMC Proven™ Professional Knowledge Sharing

Paul Kingston
Solutions Architect
EMC²
kingston_paulr@emc.com

Table of Contents

Table of Figures

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Introduction

Since computers are business tools in today's competitive marketplace, there is an enormous amount of data being created worldwide. While companies may invest in fast systems and storage to work with "living" documents or other files that change frequently, over time data that does not change consumes considerable space on this primary storage. This slows search times and requires more frequent capacity upgrades. Enter archiving; archiving takes fixed content that should never change, such as Driver's License photographs (even if you would prefer your picture looked different, the state is concerned with controlling its authenticity), medical records, and financial statements; and moves them to less expensive storage that will not need to support the same transactional rate of access as other data.

Archiving involves applications that move the data, the connectivity from the primary storage to the storage destination, and the target storage itself. EMC offers a range of products in the Backup Recovery and Archiving (BURA) space to provide feature-rich storage technology designed specifically for global archiving needs. Many industries and governments regulate how data should be stored, accessed, and retained, amongst other rules. EMC Development and Product Management attempt to stay current with global opportunities to satisfy business requirements.

This article explores several technical configurations by studying six implementations of archiving solutions. While at times examining scenarios utilizing various storage and application topologies, the focus is primarily set on sharing knowledge of EMC technologies and services. The underlying message is that managing the project and designing the architecture must follow a distinct methodology to be comprehensive and successful. This awareness is enhanced by the conceptual analysis of the application and storage scenarios in a variety of company sizes, using an assortment of technologies to meet customer needs.

Managing an Archiving Project

Managing the implementation of an archive solution requires attention to technical requirements, project timelines, and documentation. In an archiving solution that involves Professional Services, a technical team will be assigned to a project. The Project Manager must rely on these resources to detail the technical requirements provided by the customer. Additionally, the Project Manager will rely on these resources to possess the technical expertise to design and deliver the solution with attention to detail and quality.

Timelines are important in any project. It is important to begin with the solution design, rather than the customer's requests or assumptions, when setting expectations for the delivery of archiving products and implementation services. If the solution is not understood, deliverables may not be met on time, and the overall timeline may be disregarded. This will impact the customer and their confidence in the management of the engagement. Consult technical resources to understand how long the solution will take to design and implement in the customer's environment. The technical team will consider known variables based on internal information gathered from the account's documented history and conversations with the account team. Project Managers need to identify and detail risk factors to timelines as often as they occur and as early as they are foreseen. It is often useful to communicate common issues with similar projects during the initial meeting.

Timelines consist of identified deliverables and actions with target dates. We must document ownership of these action items. Any changes must be documented. The Project Manager should record the reasons for the changes, and at whose request they were made, should post-mortem analysis ever be required due to a major failure or error. Team members' effort to adapt to and account for timeline changes can aid in scoping future projects or justifying scope changes to the existing project. Document technical prerequisites at the start of an engagement. As a project moves forward, you may find that additional technical requirements will arise as typical assumptions do not apply to the customer's environment or situation.

Documenting these changes and publishing them to the team enables the Solutions Architect or other team members to identify any risks the changes may pose to the solutions' functionality. Project Managers are responsible for ensuring successful closure of engagements, and may be evaluated on cost and/or time budgets. Any tasks that are delayed jeopardize the forecasted closure of the engagement, and should be fully documented. Project Managers should communicate responsibilities to the individual and their management when tasks are delayed. All meetings should be fully documented, as this is most often when new action items are communicated and owners identified. These should be appended to the published timelines, including target dates, when possible.

There are distinct phases in engagements. At minimum, these include scoping, designing, and implementing the solution. The scoping and designing processes may overlap or be interdependent. For example, a Solutions Architect may discover while designing a project that an unexpected requirement makes it necessary to revisit the scope of the solution and possibly seek additional funding from the customer. Additionally, a high level conceptual design is often required to scope fair pricing for the engagement. The implementation of the solution, however, needs to always follow a detailed and comprehensive design and site preparation.

In this article, we will examine several case studies to demonstrate where proper management and a comprehensive design of archiving environments would have enabled customers to better meet their goals. While portions of the case studies may seem similar to actual customer implementations, the details and outcomes have been modified to most closely meet the goals of this paper, and each scenario should therefore be assumed to be fictitious. The case studies begin by analyzing a scenario in which a customer's applications are writing to non-specific storage. The intent is to demonstrate that designing an archive project to account for all the needs of a customers' data, and the technology capabilities involved in the data's lifecycle will benefit customers regardless of product choices, budget, or size of the business.

Customer Scenarios

1. Tax Returns

A small company with three employees performs accounting services and provides tax advice. A couple of types of important business data are created every year. The company electronically creates tax returns for its customers, and computer documents such as financial spreadsheets. They need to keep all of their financial data for three years to remain in compliance with local regulation. Additionally, the customer would like to keep the tax return data for a few years, so that their customers and consultants can reference previous year's returns when creating new ones. The documents related to the financial consulting business should be replaceable as customers' needs change.

Understanding the data and the customer's needs, we know that we must first provide a mechanism to segregate the data so that different retention policies can be applied. We must ensure the company's financials are retained for three years, while tax returns and consulting documents must have more flexible retention enforcement. We then see the customer's budget dictates that products such as EMC Centera® or EMC Documentum will not be good options for their environment. The data first needs to be centralized since there is no requirement and no value to having the data segregated by individual employee or workstation. This requires a separate computer investment, which is a considerable investment for a small company. Next, we need to discover what capabilities are available to protect that data for defined periods of time.

At this point, it is desirable to control the process and the access to delete data. We discover that the deletion right can be disallowed to all but a single user. We decide that this user can have a defined account on the storage device. The profile credentials can be kept secret, and embedded into a script that will perform the file deletions based on specific rules. This decreases the likelihood that data will be accidentally or maliciously deleted prematurely. We provided a second account to the user who will perform the deletes. Only this account has the right to execute the deletion script. The script deletes data based on the date it was written to the centralized storage.

The value of designing the end to end solution before deploying any of its individual components is now apparent. We know that the date we move the data from the individual workstations to centralized storage needs to be standardized to accurately reflect the data's age. The answer is another script which runs on each user's

workstation and copies financial data to a location on the centralized storage on the same day that the files are created. The tax returns and consulting documents may be copied to a separate partition of the centralized storage to enable different rights to be applied. We can run scripts to easily manage data that does not yet have a defined retention period.

Reengineering this company's solution at any point in the project may prove to be relatively painless, but will undoubtedly involve some additional time and effort than would be necessary if the Lifecycle Management of the data was fully designed before the engagement began. In this environment, the storage is basic disk storage with no Write Once Read Many (WORM) technology on top of it, so the data can be manipulated by modifying the necessary access capabilities of user accounts within the operating system of the centralized storage. However, if a company used shrink-wrap applications to perform their archiving and storage with increased features and protections, this would likely not have been the case.

One of the company's employees works from a remote location. A year after archiving to centralized storage, the company decided that it would like to keep an offsite copy of the data to protect it in case of disaster. In this environment, we could write a script to copy the data to a remote share on their network. This would cost the company little, as the script would have likely cost nearly the same if implemented initially.

If this were a larger company using Centera for archiving, the cost would be substantially different. Keeping an offsite copy of archive data is recommended for Centera deployments. Many customers, however, incorrectly assume that replication can be turned on at any point in the future, and that their data will then be copied to the remote storage. Like metadata and retention periods, replication has to be configured before the data is first written to Centera. Centera replication is a point-forward technology, and will not copy any data written to a Centera prior to replication being configured.

The pre-existing data may be copied by EMC Professional Services using data migration tools, but this service can be costly and may be unnecessary if the initial implementation was properly designed. If a replica cluster is intended to be implemented in the near-term, Centera replication can be configured to queue up data to be replicated by pausing replication and designating the Centera's local IP addresses as the replication target.

Pointing the Centera at itself is important because the application will try to connect to all of the IP addresses listed in a replication configuration each time it makes a Pool Open call. For more information on Centera Replication, see this section on page 19.

In this scenario, the company's data protection consisted only of restricted access to the centralized storage. Since hacking into a computer system is not uncommon, this protection is not robust enough for companies that can afford more. Centera offers additional security features, for more information on these, see this section on page 23.

2. Email

Another case that demonstrates the value of standardized project planning involves a medium sized company that is archiving email to optical storage. Historically, the company suffered two major outages to their production e-mail systems because of a server design that funneled the delivery of their e-mail through inadequate infrastructure. The architecture was redesigned and the company was then faced with a new problem. They were rapidly approaching the capacity of the optical platters in their jukebox, but had no written procedures on what to do with the platters or how to continue operations.

They decided to keep a copy of every platter onsite, and another copy at a remote location in case of disaster. This required them to assign two of the drives in the six drive jukebox to archiving current data, and the other four drives allocated to manually copy one disk to the other by pulling the data back through the Small Computer Systems Interface (SCSI) connections to the server and subsequently to the replica platter with OS copy commands. This project is time consuming and would have been unnecessary, should the business's objectives been understood and properly scoped before the initial installation.

Four years pass, and the company is familiar with the speed at which they can retrieve email from the archive. They determine that a separate environment must be built to handle recalls of large securities industry regulatory requests. A request arrives calling for a year and a half of several employees email, totaling hundreds of thousands of documents. The technology owners of the company find it difficult to dedicate time to assist in the implementation of this archive restore environment, but allocate resources from time to time to troubleshoot issues. Archive experts are brought in and discover

that searching the optical media for individual files is extensively time consuming. Instead, they stage entire platters on a revolving Storage Area Network (SAN) cache. Files are retrieved that match the search list while they reside in cache. This involves juggling platters for some time in and out of the jukeboxes. Additionally, the disks must be copied first from the optical media currently in production when bad disks are found amongst the duplicates recalled from the Disaster Recovery center. Bad disk errors are not uncommon, and various drivers and patches are tested to perform the recovery of failed disks. Of the three disk jukeboxes allocated to this effort, one is used to copy optical disks, one for copying data to the SAN cache, and the third for ad-hoc searches when individual messages are not retrievable from the bulk SAN copy.

Sometimes the archiving application can read neither disk in a duplicate set. In these cases, the team uses scripts to pull individual files of the disk with special utilities. These disk reconciliations are time consuming, and rapidly the team's 48 hour window draws to a close. With two hours to spare, the team has provided three dumps of sorted and zipped files to a remote administrator who prepares separate optical images to provide to the company's compliance department. It is only then discovered that none of the retrieved emails include attachments.

Further investigation reveals that the version of the software installed in the restore environment has a defect in how messages are associated with attachments. While the most recent version of the software was chosen due to its performance enhancements and advanced application modules, the customer must now revert to a known working version of the software to meet the requirements of the regulatory inquiry. Quickly, an earlier version of the software is installed on the ad-hoc search server and new searches are performed for messages with attachments. Old known issues like memory leaks are carefully managed by the application experts as a second server is added to the retrieval effort. Six hours later, the retrieval is complete.

Although the production of documents was successful, the opportunities for design revision were apparent. The restore environment should have been designed, configured and tested prior to production use, the same as any archiving implementation, once the volume of retrieval was assessed, and the decision to procure funds for the environment was approved. Because the project was managed by internal technical resources, many of the contacts needed to acquire emergency hardware and

network changes were known. Identifying escalation resources could be a challenge for an external Project Manager. For this reason, contact lists should be a data collection priority from any customer prior to the engagement of services professionals.

Creating a contact list can be important not only to the implementation, but also to the design. An architect may need to inquire about specific capabilities or settings on various components of the customer's environment. Changes to existing applications may be required to utilize features of the design, and coordination must occur directly with the owners of those application servers, when possible. The customer resources who own the technology implemented during the project must be available throughout the design and implementation. For Centera administration, it is a Best Practice to use a distribution list for the storage administrators when configuring Simple Mail Transport Protocol (SMTP) notification. Too often, a customer provides individual resources' email addresses for this setting.

3. West Coast Replication

A company replicates to a cluster in a campus datacenter, and in a chain fashion to a remote datacenter in a West Coast state.

Figure 1: Chain Replication

In this configuration, chain replication is configured to not replicate deletes to the West Coast Centera, thereby retaining the data in the event of mistaken deletion or future needs for deleted content. The remote Centera uses capacity at a greater rate than the primary and bunker site, and relies on the retention policies of the application server backups to preserve the records of the stored data following the user directed deletion.

If the Content Addresses are not preserved in this or some other method, the record locators for these deleted objects will not exist, and the stored data may be all but useless to the company. In this case, replication may be used to partially satisfy the retention requirements of the company, but at a cost. Not only must the application server databases and metadata be preserved indefinitely, the deleted data remains recoverable, although likely less accountable to the end user data management as it is not visible to the primary application.

In the event of legal action, the company may be held liable by the existence of historical data. In most cases, Information Lifecycle Management defines a point when data is no longer useful and should be destroyed. Using retention definitions on Centera enables the deletion of data once the retention period has expired, and the data becomes functionally obsolescent. For more information on Centera and the use of retention, see this section on page 21.

4. Phone Calls

One customer had recordings of phone calls that they needed to archive. Their design was created with the application vendor, and did not involve the storage vendor. The result was a design in which six servers recorded calls from distinct geographical districts. The calls were categorized as emergency or non-emergency calls and placed on EMC Symmetrix® for high availability and fast write and read times.

The emergency calls were programmatically identified and archived to Centera by a separate server. This server then held all knowledge of where the calls were and how to retrieve them. The company needed to report which district was utilizing the most storage and producing the most emergency calls, but this could not be done from the standpoint of storage. Any such reporting would have to be produced by the district's servers, extrapolated and multiplied to estimate the amount of storage being consumed

on Centera. A more robust solution would have enabled the application to archive directly from the district's servers, thereby associating metadata unique to that district with the files. Alternatively, keeping the data segregated on the Symmetrix may have enabled the application archiving the data to add identifying metadata associated with their originating file shares.

5. Disconnected Administrators

One customer had Centeras installed at several sites across the United States. The company housed a variety of data for its own customers. Some of their Centeras were configured with Compliance Edition Plus (CE+) security, making remote administration unavailable. The storage administrators had become accustomed to a process of reviewing email alerts from the Centeras. They connected using Centera Tools only if problems were noticed that they might be able to rectify, or if configuration changes were necessary.

Two of the Centeras were not emailing home to EMC due to the customer's improperly configured SMTP rules for those devices' IP addresses. The storage administrators were not receiving email from these Centeras because they were configured to email alerts and Health Reports to individuals who were no longer with the company. Additionally, these two Centeras were configured as CE+, so the administrators could not remotely update the notification settings, nor remotely monitor the health of the Centeras. Neither was EMC aware of disk drive issues that were occurring on the primary Centera in this replicated pair, until the customer called because they were unable to access their data from the archive. EMC dialed into the primary cluster and assisted the customer in failing over the applications to the replica, until the failed disks on the primary Centera could be replaced.

Some of the data was unavailable until the primary Centera could be repaired. This was because items that were in the replication queue were physically located on the failed disks. The Centera regenerated the data, but it was not replicated before the customer pointed their server to the replica cluster. We discovered that the primary Centera had been ingesting data at a greater speed than their Wide Area Network (WAN) could transfer the data to the replica site. Once the environment was repaired and the application configured to archive to the primary Centera once again, the customer made

plans to increase their network bandwidth for Centera's throughput needs. With proper planning, these network inadequacies could have been identified before implementation. For more information about network requirements, see this section on page 18.

6. Data Hosting

Business policies may require accounting for and managing infrastructure investments, including efficiency reporting and data utilization. This may include chargeback to departments or customers, as in the case of a company that stores data for customers that use their applications and services. Each of their customers has their own application environment. In some cases, these are blade server clusters and in others hosted application environments in remotely managed VMware servers.

The company needs to know the Centera capacity each of their customers use for chargeback purposes. Initially, the servers were all writing to the default pool with no distinguishing metadata, so no chargeback reporting was possible. We recommended that they send each server's data to a segregated Virtual Pool. The company was concerned about the 98 pool limit, as they thought they could exceed this number.

The company then wanted to use Centera Seek to index and report the data, but the large number of individual files was also approaching the limits of this application. The best solution was to add custom metadata from each server that could be queried and reported by a custom application using Virtual Pools to organize and segregate the data. This allowed some customers to share Virtual Pools, should the Centera's recommended limit be reached. It was understood that all of the data would have to be recalled and rewritten to obtain reliable capacity information on the data that was archived before the solution was properly designed, and that this would mean additional cost and temporary network storage allotments. For more information on capacity reporting, see this section on page 24.

Before the Engagement

A Project Manager's responsibility begins before the formal engagement of the project. To prepare for the initial customer planning meeting, the project manager should facilitate the preparation of resources involved in delivering the products and services. This includes validating that the solution has been qualified. EMC's Solution Validation Center (SVC) qualifies many designs prior to their delivery. For all implementations and major modifications to EMC storage environments, EMC Change Control (CCA) is required. There is not currently a required SVC validation of an end-to-end solution for an EmailXtender implementation archiving to a pair of replicated Centeras. This qualification rests primarily with the Solutions Architect designing the project's scope or delivery.

The Project Manager often assists in writing the Statement of Work (SOW). This limits the responsibilities of the Professional Services team to the services necessary to deliver the solution and related work. Should the customer request work beyond what is necessary to simply install a functional environment to meet their Information Lifecycle requirements, it is important that the Project Manager ensure that language is entered into the SOW to explicitly define that work and its duration.

In cases where no SOW is necessary because a sold service has pre-defined tasks that are already documented, the Project Manager working with the project's technical resources or account team should validate that the pre-defined scope has been provided to the customer. As much as possible, the customer's environment should be verified to meet the technical prerequisites of the scope of work with the presales support team, before meeting with the customer.

The purpose of the customer planning meeting is to introduce the Professional Services team to the customer resources who will be responsible for the project and the technology throughout and following the implementation. The Project Manager is responsible for collecting the customer contact information including roles, responsibilities, and locations. During this initial meeting, the Project Manager should lead the technical resources in reviewing the customer's environmental prerequisites, the team's expectations, and the solution. There is usually information gathering that the technical resources will need to perform, and as this information is being collected

throughout each meeting and subsequent correspondence, the Project Manager should triage it to the appropriate Professional Services resources and keep a record of all collected information. Action items will be created during the initial customer meeting, including the expectation of attendance at future meetings and checkpoints. With each of these assigned tasks and data collection requirements, the Project Manager can begin to formulate timelines specific to this engagement, and combine those with estimates previously completed for this standardized or customized service.

Designing the Solution

We must begin with data characteristics to design the most comprehensive and functional archiving implementation for any size budget and business. There are many questions that a Solutions Architect will answer during the design process, and we will examine some of these. In answering these questions, we will reveal Best Practices and recommendations for certain configurations, and demonstrate how effective planning can maximize the functionality of a solution and thereby enable the customer's technical owners to meet or exceed business objectives.

Archiving is, over time, necessary to increase the usability of primary-tier storage, valuable for preserving fixed content for its useful life, critical to maximizing the financial return of a company's infrastructure investments, and often required for regulatory purposes. Archiving does not provide the same functionality as some other storage processes and technologies. The Professional Services team must assess the customer's expectations and the capabilities of the existing infrastructure and purchased solution to achieve the customer's objectives. To succeed, the team must understand the expectations about the use and storage of data. Thereby, the Solutions Architect can define limitations, policies, and procedures for working with the customer's data.

The first question may be: *What do you want to archive?* Does the current data set include unsupported data types? Application files, for example, typically do not perform well on archive storage. The comparably slower performance of an Optical media Jukebox, tape, or even Centera to primary-tier storage does not provide the search times and access capabilities that applications typically require. This is seen in an environment where a customer wants to implement archiving for a Microsoft Windows Common Internet File System (CIFS) from a Celerra® to a Centera. The file system is

comprised of several department shares and user directories. The customer wants to regain space that is used by infrequently accessed files to increase the disk performance of the primary storage, reduce the time it takes to back up the file systems, and lengthen the schedule of storage upgrades to the primary array. On some of these shares, they store application executables and library components. These files need to be explicitly excluded from archiving policies to ensure that access to these programs is not delayed or altogether unsuccessful.

While not always posing them as questions, we'll discuss six major items to consider when designing an archiving solution. Once again, the information supporting these points will focus on EMC Centera and archiving products. The technical design will differ when using other products, and only professionals that are trained and knowledgeable about how to architect the implementation of archiving products should design solutions, regardless of the type of storage or applications used.

1. File Sizes

Certain features and pre-requisites of the archive storage and application require the Solutions Architect to know the expected average file size of the data subject to archiving policies. Another customer was archiving Network Attached Storage (NAS) file shares to Centera. One of the applications writing to their Celerra created backup images of end-user workstations and laptops. The customer was experiencing intermittent data unavailability when trying to access files on their Centera. The error messages were that the Network File System (NFS) server was unavailable. EMC Support found that the customer's Centera File Archiver (CFA) appliance needed to be rebooted to regain access to the Centera. The cause of the problem was traced to the applications' backup images which in some cases exceeded the maximum file size of the CFA, which is 60 GB. While the CFA is becoming deprecated, being replaced by the Rainfinity File Management Appliance, the CFA's underlying Centera Universal Access (CUA) platform, which is still supported, retains this file size limitation.

In addition to considering large file sizes, very small files can also be undesirable to an archiving process. File system archiving applications such as DiskXtender, Rainfinity FMA, FileNet and others do not perform containerization of data. The applications move the data to archive in whatever state it is in when it matches the move policies. In

DiskXtender, a minimum file size can be specified to qualify for migration. While FMA will skip files that are 8 KB or smaller on a Celerra, and 4 KB or smaller on a NetApp appliance, files as small as 10 KB will still be eligible for archiving. Archiving files this small can significantly degrade the performance of Centera, as each file is processed by a thread worker before processing the next file.

Recommended maximum file sizes for a Centera may be 100 GB, which would ultimately contain one thousand 100 MB blob slices. The Centera will first divide archived files into 100 MB slices before storing and protecting the data. The time to read files larger than 100 MB may depend on the number of blob slices. Centera Development has run tests to show that files that contain certain numbers of slices perform poorly, while increasing or decreasing the number of slices improves performance.

This aberrant behavior has been addressed, but no publication of a performance report that states the optimum or problematic number of blob slices has been seen. Such a report might indicate the optimum file sizes over 100 MB for Centera archiving, but currently available information allows us to assume that there is no known performance impact for different file sizes below the recommended maximum file size. The important thing to recognize is that with file sizes larger than 100 MB, the protection scheme should be considered closely not only for a possible performance tradeoff, but also the ultimate number of objects the files will create and that the Centera can support. The number of files is important for correctly sizing the environment.

With each file and administrative change to a Centera, multiple objects may be stored. Each Centera node can be expected to support 50 million objects, and in some environments it may be necessary to allocate additional Centera nodes because of the number of files, before capacity utilization dictates the need for expansion. Other applications such as FMA, CUA, and Centera Seek also have a maximum number of files that each node can support.

File size limitations are certainly not unique to Centera. The total size of optical platters may be 5 GB or 9 GB. Archiving applications may be unable to write a single file across multiple platters. Small files can be problematic whether writing across a network, SCSI or other connectivity. As the application processes each file, connects to the archive

storage, and then writes the small file before moving to the next file, performance is lost per thread worker between the writing of each file. Tape devices may experience shoeshining when archiving very small files. Whenever possible, multi-threaded applications should be employed to write small files to any storage. Improved write and read times and managing the total object count on Centera are good reasons to containerize. Data that is smaller than 1 MB should be containerized before being written to Centera whenever possible. Applications that perform archiving of small e-mail files, such as Symantec Enterprise Vault and EMC EmailXtender can perform containerization before writing to Centera.

2. Network Requirements

In an environment that moves data over the network, customer network architects and engineers need to be identified and familiarized with the project's intents, timelines, and resources. The understanding of the network and security requirements of a solution between application, primary and archive storage tiers and management consoles invites the networking teams to be involved in the design process. If they cannot provide the necessary configurations and contingencies, the network may prove inoperative for the solution. In production situations, these issues may result in data unavailability.

Network requirements are sometimes not investigated until the solution is being delivered. While the cost of making networking changes may not be substantially different when a solution is purchased rather than at time of implementation, the cost may potentially be budgeted. Additional cost may be incurred by both the customer and vendors if projects are delayed because networking requirements were not identified earlier on. A good example involves customers who wish to perform offsite replication with Centera, but do not have a WAN that can support throughput equal to their rate of ingest at the primary site. This will lead to a perpetual replication queue and underprotected data if deployed without increasing the bandwidth available between the sites, or tapering the rate of archival to their network limitations.

In addition to bandwidth, network security is frequently a cause of project delays. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port requirements for Centera archiving and management environment are not always

planned when a Centera expert is not engaged early to review the architecture requirements with the customer.

Network ports that are most commonly associated with Centera implementations are detailed in the table below.

Port	Inbound/Outbound	Protocol	Description
3682	Inbound	UDP TCP	Remote Management via CLI/CV
3218	Inbound	UDP TCP	SDK access
22	Both	TCP	SSH daemon
25	Outbound	TCP	SMTP for health reporting and alerting to EMC
1610	Both	UDP	SNMP GET and GETNEXT
162	Both	UDP	SNMP Traps
7069	Both	HTTP	Default Main Centera Console server port
7059	Both	HTTP	Default Centera Console shutdown port
514	Outbound	UDP	Syslog

Figure 2: Centera Network Ports

The most common problem with network security occurs when multiple companies are involved in the ownership of facilities, processes, people, and data. Some very large companies outsource components of their Information Technology Department (IT), and their end users may have to cross several networks to access their data. This may involve firewalls and Network Address Translation (NAT). A networking expert with a complete understanding of Centera's requirements is needed to prepare this environment for a future Centera implementation.

3. Centera Replication

Centera can support a variety of replication topologies. It is important to compare each alternative to the customer's environment and needs. Customers may have an existing relationship with application vendors that can cause them to primarily design the archiving solution from the perspective of the application; instead of the entire lifecycle of the data and the needs for that data after the archiving is complete. This can preclude the configuration of specific metadata, Network Data Management Protocol (NDMP) backup accessibility, and data mobility. DiskXtender, for example, can be configured to

provide a scheduled server export to a remote server. If this server is placed online at a replica site location, a replica Centera can be available for high availability and dual site access of data as recent as the last server metadata import. This is an attractive Disaster Recovery configuration when it meets the customer's Recovery Point Objective (RPO) needs. Additional backup and recovery operations can add to the point of recovery options, all the while reducing the time it takes to recover the environment as opposed to server restores from tape or even disk backup that often requires manual server and network configuration.

A Best Practice surrounding Centera Replication is to make an application's database of Content Addresses available through replication, if possible. Another option is often used with medical image PACS application environments. In this scenario, the application provides data portability to a remote application server that archives at the remote site to a standalone Centera simultaneously, or nearly so, to the archival at the primary site to another standalone Centera. Drawbacks of this design include uniquely stored objects at either site, making recovery of storage failures dependent on the regeneration of the data using the application. Technology upgrades that involve data migration from one Centera to another will not support this as one replicated environment, but rather two migration efforts will need to commence to preserve the integrity and mobility of the data and its offsite copy. Data migrations may reveal data integrity issues caused during times of multiple disk failures or other issues. In this case, a replica cluster does not exist to regenerate the data from. Instead the application must recreate and rewrite the data, or restore it from a Centera Backup and Recovery Module (CBRM) NDMP backup.

Replication designs may also include the replication of deletes. Replication topologies differ, and the business continuance and Disaster Recovery requirements of a company will determine the best design. Replication ensures recoverability should an issue occur with the primary site or storage. This functionality should not be confused with backup and restore, or facilitation of retention requirements.

4. Data Retention

When designing the archiving implementation, a Solutions Architect will ask: *How long does the data need to be kept in archive? Does business policy dictate that data needs to be deleted? Is it necessary to destroy the data so that it cannot be recovered?* Different storage media allows administrators to perform different actions concerning data retention and deletion. Tape storage, for example, is not WORM compliant. Therefore, the deletion and overwriting of data can be supported. WORM optical disk, however, allows no deletion and the media must be physically destroyed when the retention period has expired and policy requires the data's deletion. Centera allows for specifying retention periods and policies and deleting eligible data. On Centera, there are also advanced options for tailoring the design, including shredding the data.

The utilization of compliance security is often ill-planned or misunderstood. There are four different levels of compliance with Centera, and often customer's are ready to deploy a solution with a compliance mode that they don't understand, and that does not meet their needs.

Basic Edition	• No enforcement of retention	• No license required • No data shredding, only deletion
Governance Edition	• Enforces retention • Can set default retention by cluster or pool	• Allows audited delete prior to expiration of retention • Retention classes may be lengthened or shortened
Compliance Edition Plus	• Strictest retention enforcement • Remote administration disabled	• Retention classes may only be lengthened
Advanced Retention	• Increases flexibility when retention needs change • Event based retention • Min/max governor • Litigation hold	• Event Based Retention (EBR) supports retention classes • Litigation holds must be issued and released by an external application

Figure 3: Centera Compliance Options

Broker-trader's have purchased Governance Edition licensing that does not meet the stringent requirements of SEC-17a4, though Compliance Edition Plus licensing does. Hospitals have incorrectly assumed that they needed Compliance Edition Plus protection, when Governance is better suited for their records due to privacy laws requiring some hospitals to delete patient records upon request. This deletion can be performed when providing a reason for doing so, such as patient requested deletion, but only when using the Governance Edition licensing.

Basic, or no compliance, is a very common setting with Centera deployments. A customer often expects their application to prohibit deletion, which it may do, but only from that application. Other applications with the correct security context may be able to affect the data, if retention is not enforced on the Centera. Centera expects the application to set a retention period for the data, but does not expect the application to enforce that retention. Indeed, many applications can perform deletes of data before the retention period has expired. The deletion from Centera will not actually occur, but the application may lose its knowledge of the data, thereby orphaning the data. It would not be retrievable by the end-user until the application database is restored. Typically, orphan file cleanup is a manual process. DiskXtender for NAS, however, offers a configurable orphan cleanup process that enables an administrator to retain orphans for a specified period of time, and then have them automatically deleted from Centera if the stub files have not been restored to the NAS before the expiration of the rule.

Centera can also set retention on items, if the application server cannot or if the configuration of the entire solution makes this difficult, but this cannot occur with Basic security. Instead of Basic, Governance Edition is required. Governance will allow everything written to the Centera to be retained for a single defined period of time. It may be desirable for some content to have different retention periods. If this is the case, and it is still desired for the storage to set the retention, Advanced Retention is needed.

Advanced retention allows separate pools of data to receive variable retention periods applied by the Centera. Advanced Retention has other capabilities to have a retention period change according to an event, or to have minimum and maximum retentions that allow more flexibility for complex legal policies.

If no industry regulation defines what retention policies a business should adhere to, a company's IT resources are often unsure about their retention requirements. They may request an infinite retention without realizing the ramifications of never being able to delete the data should business mergers, legal action, capacity management, or other situations cause groups of this data to be undesirable.

Centera supports retention classes that can be set to any length of time and modified at any time thereafter, if retention preferences change. Retention classes can be used to maintain data for 30 days to protect against accidental deletion of the data, and later extended to 3 years to meet changing business objectives. In a non-regulated business, this retention may later be changed to 1 year to enable storage administrators to allow deletions and regain capacity availability after 1 year has commenced since the data was originally archived.

The differences between compliance options can pose challenges for customers who purchase and deploy the wrong option for their business needs. An expert must analyze, educate, and design compliance settings based on the data's importance, regulation, and expectations about when and how to delete the data. Centera's compliance capabilities are one of its major strengths. The determination about compliance settings must be made collaboratively by the customer and Centera experts. The account teams and the customer must understand that the result of these discussions may indicate the need for additional compliance licensing, which could mean additional cost before the solution can be satisfactorily deployed.

5. Centera Security

Centera provides WORM technology by prohibiting the modification of any files managed by its CentraStar operating system (OS). The Content Addressable Storage (CAS) method of storing data provides additional protection. This helps to prevent any unauthorized access to data by requiring the requesting application to know the data's record locator, which is a 27 or 53 character Content Address. This Content Address is how the Centera locates a c-clip containing the User Data. A c-clip contains a Binary Large Object (blob), which is the archived file reduced to simple binaries and therefore unreadable by external applications for most file types (text documents do not change, for example.) A c-clip also contains a Content Descriptor File, which is the metadata

describing the User Data, and the protected copies of the blob and CDF, according to the protection scheme on the Centera cluster.

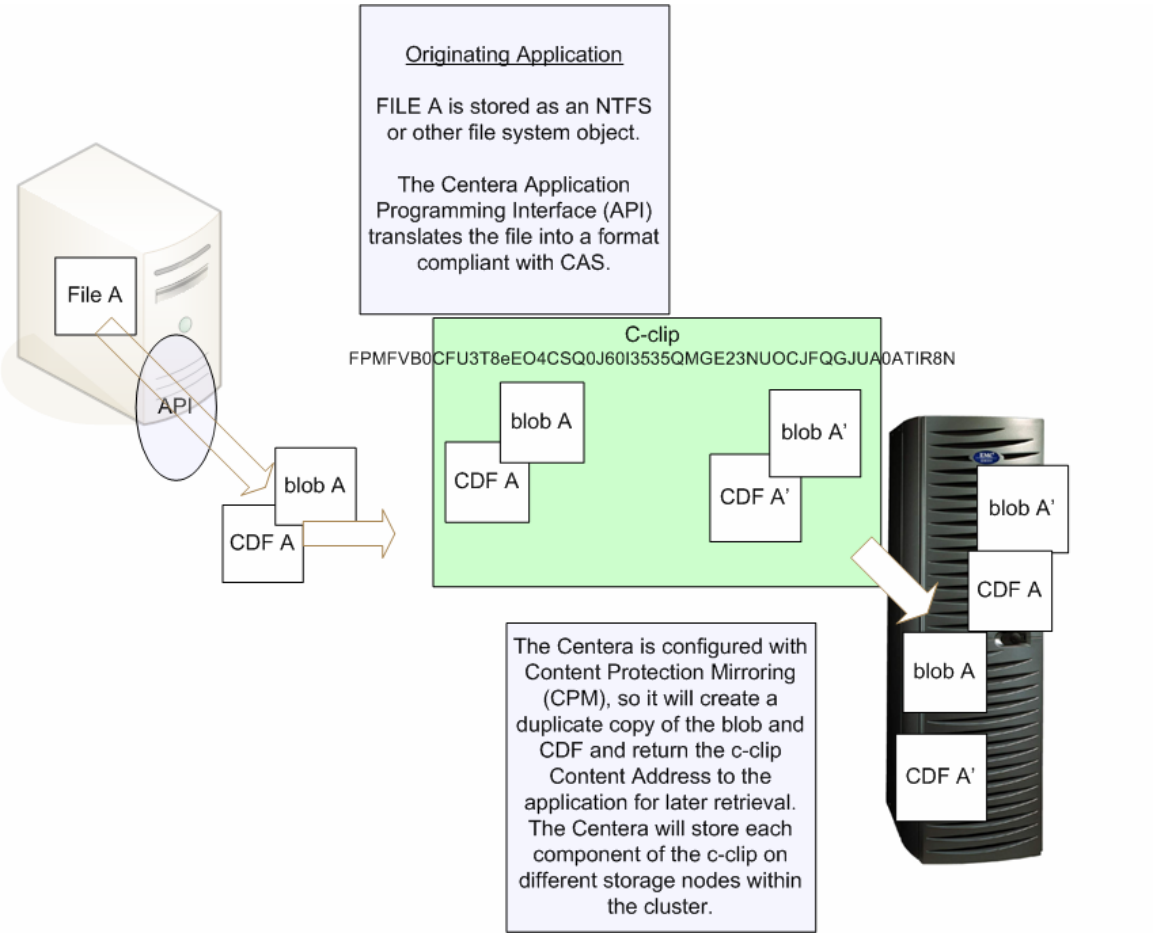

Figure 4: Archiving Files to CAS

An application or user must first authenticate to the cluster and its replication partners to write, delete, or access data on a Centera. In the past, Centera was accessible with an anonymous profile. Since CentraStar 3.1, anonymous access is disabled by default. While it can be enabled, it is a Best Practice to create an Access Profile for each application. The Access Profile can then write to a Virtual Pool. Virtual Pools provide a method of logically segregating data on a Centera. The Pool can allow the profile certain rights, and the profiles can be assigned a subset of those rights. Multiple profiles can be assigned to a single Home Pool, and an individual profile can read (only) pools other than its Home Pool, thereby extending the flexibility of rights assignments. (There is also a cluster level definition which is highest in the permissions hierarchy.)

6. Capacity, Health, and Reporting

Managing and supporting Centera includes monitoring capacity and cluster health. Operations teams may be responsible for responding to email alerts generated by the cluster. Centera Console offers a Web User Interface that graphically depicts the replication relationship, performance, and utilization of Centera. Several TCP and UDP ports are used by Centera Console, Centera Viewer and other tools, SMTP mail relaying, Secure Shell (SSH) and File Transfer Protocol (FTP) access, Centera replication, Centera Universal Access, EMC Secure Remote Gateway, Syslog, and Simple Network Management Protocol (SNMP), when applicable to the design.

Some customers have Engineering teams responsible for the configuration and design of the application infrastructure. They may not support the environment after implementation, but will make necessary configuration and procedure changes regarding the storage and environment. The Engineering team may have access to Centera Viewer for the creation of Virtual Pools and Access Profiles, Replication, and other settings, while the operation teams require only SMTP and Centera Console or EMC Control Center monitoring. The Engineering team may decide that only certain individuals and specific application servers can delete archive data, and set up Access Profiles with the necessary permissions to support this.

DiskXtender for Windows also offers similar alerting options to augment Centera's reporting, or used when the archive storage media is other than Centera. DiskXtender can alert administrators whenever Warnings or Errors occur on the server. An important configurable alert is the Extended drive free space alert. This alert is generated when the free space falls below the threshold specified in the properties of the extended drive.

The configuration of archiving applications is paramount to utilizing the features of the storage and data movement and retention policies. Prior to the design and installation of these applications, the implementation team must work in conjunction with the design of the storage, network, and archiving policies in order to provide function-rich longevity. Reporting on how much Centera capacity to use, and by whom, might be fulfilled by a couple of different mechanisms, namely Centera Console and Centera Seek with Chargeback Reporter. Both of these applications require configuration values be

present before any of the data is written to Centera, in order to receive complete and accurate usage accounting.

The data should be written to a Virtual Pool using an Access Profile created specifically for the application server performing the archiving. Administrators should take advantage of any capability the archiving application has to write additional metadata that can aid in identifying the characteristics of data that is being archived. The Centera Software Development Kit (SDK) will support server specific custom metadata configured as environment variables on any application server. Some applications can add even more specific information about data such as the originating path and file name. In addition to server-wide Environment variables, DiskXtender can add custom metadata to individual Media Groups through the Media Group properties. Administrators should learn about all of the metadata capabilities their archiving applications offer, and utilize these options to the fullest extent possible to allow reporting capabilities that can meet and exceed future business requirements.

Centera Console can historically report on cluster performance and pool capacity usage. Centera Seek indexes all of the metadata for user files on a Centera. Proper environment design can enable reporting on storage usage trends by business unit, geography, or applications over time.

Along with the management of capacity on the Centera, the IT department is typically responsible for the retention and destruction of archive data to meet with the business policies, even if there are none. In the days when optical and tape storage were the de facto standard for fixed content archiving, administrators were tasked with the offsite protection of data and when required, the physical destruction of storage media. Although with Centera, business requests to physically destroy the equipment are rare, a company's investment in the technology can be better leveraged by understanding the data replication, data retention, and data shredding features of the storage.

Centera Console offers improved monitoring of replication, and displays alerts when significant replication issues occur. Chargeback Reporter, shipped with Centera Seek, cannot natively report how much data is stored with expired retention periods. This is due to the report configuration format not supporting expressions mathematically

comparing metadata fields, in this case *creation_date* and *retention*. Seek does index this data, however, and these reports are possible with custom development.

Implementation Prerequisites

Although it may seem obvious that the physical environment needs to be prepared for the installation of new equipment or software, customers may not be fully prepared when an implementation team arrives onsite. Without a diligent site survey, change control qualification, and coordination between the customer's project resources and the Professional Services teams, requirements such as application server builds, routable network security, and management workstations may not be ready at time of implementation. Regardless of the storage technology being employed, certain environmental factors, such as the physical rack or floor space, power and connectivity will need to be defined and allocated before installation.

Data mobility requirements for the solution need to be identified and simplified, which may include Network Address Translation, firewall and routing rules, application middleware, or gateway appliances to support archiving, replication, remote access and administration, or SMTP messaging. In the case of Centera, premium support is available from EMC for the storage array's hardware and configuration, but remote access to the device is required. This can be achieved through either a modem or Virtual Private Network (VPN) appliance.

In addition to validating site preparedness, the Project Manager has other tasks that must be completed before sending resources onsite to deliver the services. All change control documentation and approval must be complete. An internal meeting should be held so that the Solutions Architect can review the detailed design with the implementation resources. Fully prepared Implementation Specialists are as important to successful services engagements as fully prepared customer sites. The physical readiness of the customer's site is only part of the required preparation. Customer resources and technology owners must be identified and scheduled during the Implementation Specialist visit, so that applicable knowledge transfer and training can be delivered.

Solution Delivery

The configuration will be tested in some fashion as part of the implementation. Testing an environment can happen in multiple ways. The solution should be tested by several applicable methods to ensure the greatest accuracy and reliability. While it may be necessary to test components of the overall solution with different tools due to the implementation of these components at different times, the architecture should be qualified theoretically and against approved and documented procedures at the start of the project. The Project Manager should oversee this qualification during the Design phase of the project, and should ensure the documented acceptance of the design is published to the customer after having gained acceptance by the internal Account Team and Technology Solutions Group engaged in the solution's delivery.

It should go without saying that the solution is best tested from creation of the data to the final target destination of the data, be that a Centera replica, optical disk, or tape storage. While the scope of work on an individual project may be focused on only a single component of the solution, such as the archiving-tier storage array, the responsibility ultimately lies with the customer to fully test the solution. The Professional Services team may not have access to all of the systems, or they may not be employed to configure preexisting installations. It is in any vendor's best interests, however, to ensure that the vendor's products will help their customer base achieve their productivity and management goals.

It should therefore be a Solutions Architect's responsibility to qualify and document the environment so that it can be proven to test theoretically against the Professional Services documented Procedures and Best Practices. The Project Manager can provide the greatest success for a working environment and an efficient delivery by insisting that this Detailed Design is published to all applicable stakeholders and that any discrepancies to the prerequisites are identified before an implementation is scheduled. Furthermore, implementation resources should be expected to verify all prerequisites including software versions, hardware specifications, and technological capabilities such as network bandwidth and settings meet the requirements in the Detailed Design before configuring any new settings, hardware or software. Should discrepancies be found at this stage, they should be documented and communicated to the Project Manager. The Project Manager should ensure that the documentation is published to determine if

additional qualification is necessary, or if postponing the work and extending the timeline is necessary before additional configuration occurs.

Project Closure

Ultimately, a Project Manager must drive project closure. This includes escalation of chronically delayed tasks or support issues, and documentation of the customer's acceptance of each project deliverable. As much as possible, the Project Manager should verify that the implementation meets each requirement in the detailed design. When knowledge transfer is included in the scope of work, the project manager should ensure that appropriate customer participation is planned. The Project Manager is often the primary point of contact in a Professional Services engagement. Awareness of customer perception is important during all phases of a project, not least at project closure. Implementing the best archiving solution will cause lasting effects on customer perception as the solution meets the long-term needs of their business.