FROST & SULLIVAN GROWTH CONSULTING

Solutions for the Airline and Aerospace Industries in Meeting the Legislated Demands of Record Retention and Accessibility

A Frost & Sullivan White Paper Sponsored by EMC
Analyst and Author: Jarad Carleton

In the world of aerospace, manufacturers and air carriers alike stand out as highly visible global businesses with operations on every populated continent in the world. They are also industries that have been subject to heightened levels of regulatory control combined with the threat of tangible financial loss when fines are assessed for regulatory noncompliance. When these companies are brought to court and/or fined for noncompliance, tangible financial losses are easy to measure. However, intangible financial losses will continue for months or years as a bad corporate image pervades the very market the manufacturer wishes to target. In the case of government contractors, noncompliance can result in losing all government contract business. Participants in the aerospace and airline industries will recognize that similar situations have occurred in the past. What is important to know, however, is that future occurrences of similar scenarios can be minimized or avoided entirely with the application of proper technology.

This white paper describes the data retention issues facing the global aerospace and airline industries. Laws enacted to regulate discrete aspects of both industries require companies to be responsible for maintaining product data as well as maintenance records for several years to decades as well as ensuring that the retained data is readily available on demand around the clock. EMC Corporation's Content Addressed Storage (CAS) solution, Centera, is examined in the context of the global aerospace and airline industries and how it can cost-effectively help meet legislated data retention and availability requirements for both industries. The introduction of Centera from EMC to the aerospace and airline industries represents a technological and cost-effective leap forward for long-term storage of fixed content information within a regulatory environment. Fixed content is unchanging digital information assets retained for active reference and long-term business value.

The information presented in this publication is based primarily on interviews and therefore is subject to fluctuation. Frost & Sullivan takes no responsibility for any incorrect information supplied to us by manufacturers or end-users.

Frost & Sullivan
1040 East Brokaw Road
San Jose, CA 95131-2309, USA

Record Retention Trends, & Opportunities

Record retention has always been necessary for aerospace manufacturers and airlines, however as laws change, demands placed upon record retention have increased in scope and decreased in average time allowed to access data. This data takes multiple forms, but can be segmented into two categories:

1. Dynamic data – information that is actively being edited and changed.
2. Fixed content – unchanging digital information assets retained for active reference and long-term business value.

This latter form represents more than 75 percent of information retained by manufacturers.

In the 1990s fixed content was often stored in a variety of content repositories, a small amount on disk-based storage and the majority on offline optical disk or magnetic tape. The result was often several unrelated storage systems accompanied with the challenges of maintaining them. The problem reached its peak when disparate enterprise content management (ECM) systems emerged to manage different types of fixed content ranging from documents to Web pages and digital media. Having different content repositories, for each type of ECM system placed undue strain on IT personnel and budgets and made the shortcomings of multiple storage repositories painfully obvious.

The opportunity for the storage industry was to develop a cost-effective, online, networked repository for use in the enterprise environment by applications and databases alike. The requirements for a fixed content storage solution are much different than for traditional transaction-based storage:

1. Access times need to be much faster than tape or optical solutions if fixed content is to be readily available to users, yet access time doesn't need to reach the speeds required to support engineers designing products. Therefore, an optimum storage system could leverage the latent knowledge of Internet Protocol (IP), which is cost-effective, and that all IT departments possess in order to network the storage system to business-critical applications.
2. Since the fixed content is retained for long periods of time, the storage solution must:
    a. Ensure content authenticity.
    b. Require minimum human intervention.
    c. Take advantage of technology evolution without requiring information migration.

U.S. Air Carrier Maintenance Recordkeeping System Requirements

The U.S. Federal Aviation Administration (FAA) requires that all U.S. registered aircraft operated by an air carrier have a valid standard certificate of airworthiness. This certificate is valid as long as alterations, maintenance, and preventative maintenance are performed according to FAA guidelines. Furthermore, Title 14 of the Code of Federal Regulations (CFR) outlines requirements for an air carrier's storage, retention, and retrieval of accurate

and complete aircraft maintenance records, which are the primary means used by the FAA to determine if required maintenance has been performed.

A summary of the type of maintenance and repair records that must be maintained by U.S. air carriers, some of which have multiyear retention requirements include:

- The total time-in-service of the airframe, each engine, and each propeller
- Current status of each life limited part of the each airframe, engine, propeller, and appliance
- A listing of the time of the aircraft's last overhaul, the parts overhauled, and the time-in-service remaining until the next scheduled overhaul
- The current inspection status of the aircraft
- The current status of an airworthiness directive for the aircraft
- Listing of the current major alterations of each airframe, engine, propeller, and appliance
- All records necessary to show that all requirements for the issuance of an airworthiness release have been met

The importance of maintenance record accuracy and retention has been highlighted by the U.S. Congress, which has made the falsification, alteration, mutilation, and lack of recordkeeping a criminal act subject to substantial fines, imprisonment, or both. Although FAA regulations require a copy of the aircraft maintenance log to be kept in hardcopy format in the aircraft itself, OMB Circular A-130 encouraged the use of electronic signatures to authenticate aircraft maintenance records and pushed the use of electronic record keeping systems for airline safety reviews and audits by the FAA and corporate safety compliance personnel.

The importance of maintenance data requires authorized personnel to have immediate access to all records for an aircraft in question negating the use of tape or hard copy storage as a viable option since it can take hours to days to access. Conversely, maintaining fixed content, such as historical maintenance data, on traditional SAN or NAS-based storage systems, which are optimal for maintaining dynamically changing data in the engineering and business operation environments and where data is retrieved virtually instantaneously, is not a particularly cost-effective solution for an increasingly cost conscious industry. The best solution will meet an air carrier's need for a reliable, online, fault-tolerant, and cost-effective storage solution that addresses fixed content historical maintenance data and ensures that the information has not been altered.

Aircraft Manufacturer Records Retention

Air carriers are not the only companies that are subject to multiyear data retention. Aerospace manufacturers have been required to store all data related to a product ranging from engineering schematics to wind tunnel and flight test data as part of the airworthiness certification process. Specifically, aircraft manufacturers holding a technical standard order (TSO) are required to have the following records available at factory locations at all times:

- A current and complete technical data file for each type or model article, including design drawings and specifications
- Complete and current inspection records showing that all tests and inspections have been properly completed and documented

All of these records are required to be kept as long as the manufacturer continues making the aircraft or component part and for a minimum of 2 years. This recordkeeping requirement forces manufacturers to maintain large amounts of data for periods ranging from 2 years to multiple decades. Furthermore, the manufacturer must be able to provide a FAA inspector access to the technical data files on a product when required; illustrating the need for readily accessible data through the use of disk based online storage.

Export Regulations and the Paperwork Reduction Act Necessitate Online Storage

One of the benefits of the 1995 amendment of the Paperwork Reduction Act of 1980 and OMB circular A-130 is that they encourage the use of online record keeping systems. While aerospace companies have to comply with the export compliance program outlined by the Office of Defense Trade Controls (DTC), compliance entails not only a minimum 5-year record retention, but also a thorough description of record systems concerning U.S. origin products. As is the case with many government agencies, the DTC office fails to outline what type of online record keeping systems are appropriate, leaving the door open for unintentional record keeping violations.

Unfortunately, violations of the Arms Export Control Act (AECA) and the International Traffic in Arms Regulation (ITAR), both of which fall under the jurisdiction of the DTC office can be severe for both the manufacturer and individuals within a company. Simply put, the best way to avoid record keeping violations is to ensure that any online storage system used complies with DoD 5015.2-STD, a Department of Defense design criteria standard for electronic records management applications.

Data Storage & the Information Lifecycle in Aircraft Design and Product Development

Data storage regulations for the global aerospace industry are broad and most require product information to be retained after the product has been manufactured and sold to end-users. The need to retain data however, begins when a new aerospace product concept is conceived and lasts through the manufacturing and sales process as well as through the lifespan of the product and beyond. Fortunately with modern online storage technology, aerospace manufacturers have several storage options to choose from for different phases in the lifespan of every product manufactured.

The information lifecycle at many aerospace companies today typically begins with keeping data in the product creation phase online with high performance storage solutions. Immediately after the product design is released, and Bill of Materials and other manufacturing records are transferred to Supply Chain Management, Enterprise Resource Planning, and other business systems, the product data is typically stored offline on optical disk, magnetic tape, and in some instances in hard copy form in offsite warehouses. The problem with this model of product data storage is that service level requirements for the data in question don't immediately drop to levels low enough to justify the use of offline

storage. Components manufacturers, legal teams, and other groups with a need to access aerospace product data are subject to delays in data access until the optical disk or magnetic tape can be taken out of storage and be temporarily put back online when requested. The end effect is the creation of a large service level gap that manufacturers aren't able to adequately address in a cost-effective manner.

As networked storage solutions become ubiquitous among aerospace companies, the

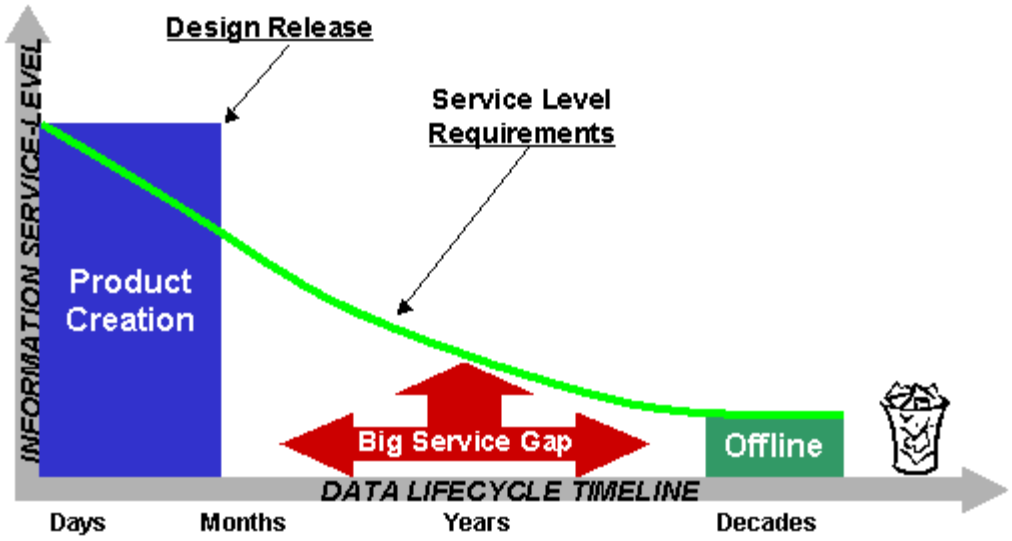

information lifecycle for product design and development will address the changing needs of manufacturers.

Today, the product data lifecycle timeline has evolved into four phases for each aerospace product and component part design. Each of the four phases has unique service level parameters, which diminish over time as data moves from one phase to the next. Information lifecycle phases in a modern discrete manufacturing environment are:

- *Product Creation* –Defined typically as a period of days or weeks or even years wherein CAD/CAM work-in-process as well as test and design simulations are conducted. Information in this phase of the information lifecycle is highly dynamic and can change on an hourly basis.
- *Availability & Sharing* – Defined typically as a period of months at which time the product design has been released to manufacturing engineering, business operations, partners, and suppliers. Information is moderately dynamic in this phase as it is shared with additional integration teams and component suppliers, for reuse or to conduct further engineering tests, resulting in altered data. A concern and often the subject of horror stories is the sharing of design documents whose versions were in question. For this reason data leaving the product creation phase at design release may be treated immediately as "fixed content" where information integrity and authenticity for these design files can be strictly maintained, avoiding costly revision control mistakes.
- *Fixed Content/Archiving* – Here the product information is no longer subject to changes and becomes fixed content. Multiple users accessing the same fixed content characterize this phase. Historical designs, parts failure reports, customer communications, and e-mails (retained for legal purposes) regarding aerospace products make up much of this

information. Note: If any change needs to be made to a document stored in this phase, the optimum storage solution will keep the original document and create a new version with a unique identifier, such as the example described in the previous paragraph.

- *Offline* –Here information is accessed infrequently, if at all and regulatory requirements no longer encourage manufacturers to keep the information readily available.

As stated earlier, a different type of storage system is needed for the lifecycle of information. To date, an aerospace manufacturer had two options characterized as "fast and expensive" or "cheap and slow." If the storage industry could provide a cost-effective, online solution that ensured information authenticity, then many of the challenges that storage technology placed on IT in the past could be eliminated. In this scenario, all fixed content could be put online, help increase end-user service levels, and be repurposed for other valuable business uses. This is the promise of the newest category of storage, Content Addressed Storage, for fixed content.

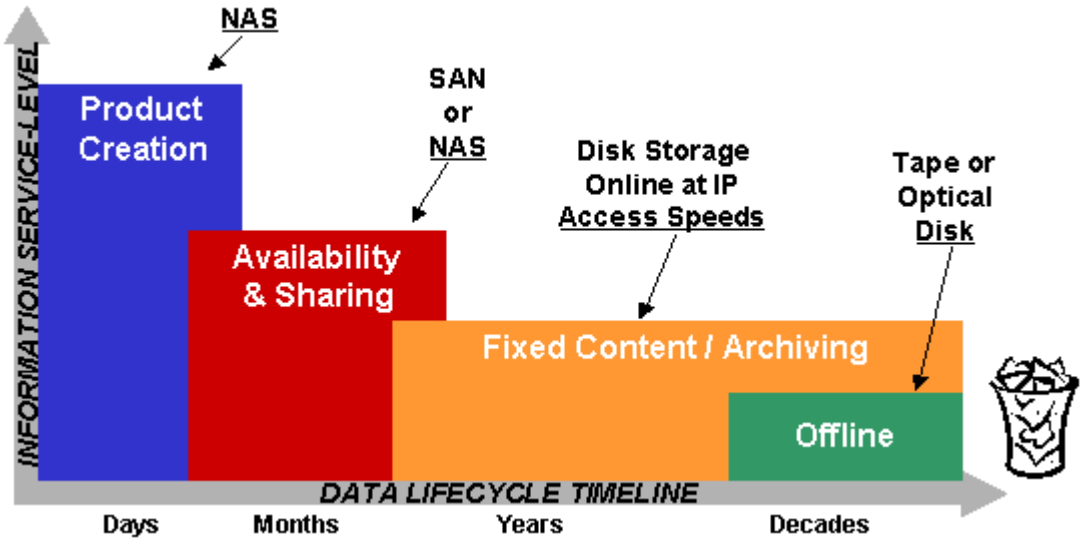

Content Addressed Storage (CAS) and EMC Centera

Content Addressed Storage addresses storage based on a Content Address (CA) not the information's physical or logical placement in the storage solution. EMC Centera™ is the industry's first implementation of Content Addressed Storage (CAS). Centera stores information objects based on a 128-bit globally unique address that is derived from the object's binary representation. With a Content Address derived from the content itself, Centera eliminates the storage of multiple copies of identical information, regardless of how many requests to store a piece of content are made. For business continuance purposes, Centera stores the content and protects it using content mirroring or content parity protection within the same Centera array. The addressing and encryption functions are similar to a public key infrastructure (PKI) ensuring, security, authenticity and nonrepudiation.

When content is stored in Centera, the unique content address for the stored object and the metadata describing the object are inserted into an XML file. The XML file is referred to as a

C-Clip™ Descriptor File (CDF), which also has a unique content address calculated for it. The CDF is then protected (in the same way as the object itself) and is the mechanism used by an application to retrieve an information object. Retrieval of content within Centera is based entirely upon the content addresses rather than through the use of a centralized directory, pathnames, or URLs. Using a content address to access fixed content makes the management of physical and logical location of the information unnecessary, which results in a dramatic reduction in system/storage management. In the event that fixed content is altered and stored again, Centera computes a different content address (because the content is different) and stores it in the array. Original fixed content is not overwritten, ensuring an intact audit trail and assurance that fixed content remains in its original state.

As an integral part of maintaining data integrity and audit trail in the event of a hardware failure in one part of the array, Centera will self-heal by detecting the fault and generating a new copy of the content objects. As this process takes place, the affected disk drive or storage node is isolated from the rest of the system until it can be replaced. Lastly, due to the fact that applications don't have knowledge of the physical placement of fixed content within Centera, components can be replaced and Centera software upgraded without disruption, demonstrating a solution architected to easily scale up to one petabyte and beyond.

Centera can and has been configured to help meet the most stringent requirements of regulated environments. Specifically, Centera enforces application-based retention periods within its microcode. Manufacturers have the ability to lengthen the retention periods, but cannot shorten them.

Just as important as data retention in a regulatory environment however, is the importance of automatically deleting unwanted information using U.S. Department of Defense data destruction standards outlined in DoD 5015.2-STD, overwriting information multiple times with random characters, complimentary values, ones, and zeros, when data reaches its expiration date. This not only frees IT personnel from doing a low level maintenance task, but also creates a standard data destruction policy within an organization eliminating potential legal liabilities, and negates any ability to recapture deleted information using disk management tools. Also essential to any data management policy, is the ability to immediately suspend data destruction in event of litigation. Centera's central policy management allows immediate extension of data retention periods as needed for any eventuality.

EMC Storage Solutions Potential for Aerospace/Air Carrier Record keeping

As discussed earlier, air carrier maintenance records, manufacturer product information, and export regulations require the following features:

- Data reliability
- Authentication/Nonrepudiation
- Access on demand
- Cost effectiveness
- Centralized control

Data reliability and nonrepudiation are two particularly important aspects of a storage solution in the airline and aerospace industries. This is because of the regulations

surrounding maintenance data retention for airlines and product data retention for aircraft manufacturers holding technical standard orders (TSO's) for different aircraft models.

Specifically, manufacturers that hold a TSO are required to keep all data for each aircraft model including design drawings and specifications. Records must also be retained for all product tests and inspections pertaining to an aircraft model as well as its component parts such as propellers. Finally, these records must be available to the FAA or other authorized government agencies on demand.

EMC Centera addresses all of these critical points and has the capacity to meet stringent DoD 5015.2 records retention requirements with Centera Compliance Edition™, meeting the storage needs of all sectors of aerospace and related industries.

Data reliability within Centera is assured as a result of its RAIN architecture that eliminates all single points of failure within the platform and enables non-disruptive servicing of the system. Centera itself is composed of independent nodes with one terabyte of raw storage capacity, and is interconnected to all other nodes in the cluster via CentraStar™ software (Centera's operating environment) and a private LAN.

Authentication of content is handled by Centera's use of content addressing, which stores fixed content within the cluster using a globally unique address derived from an object's binary representation that is subsequently time/date stamped. The system operates similar to a public key infrastructure (PKI) that establishes encryption algorithms ensuring the integrity of the data so that the end user knows it has not been tampered with.

Centera easily handles access to fixed content on demand via LAN connectivity to application servers. Because each node within the cluster operates as a storage node or an access node, performance can be scaled to meet demand by non-disruptively adding additional front-end access nodes to augment bandwidth to application servers. This enables Centera to work within a growing database-driven environment that is common in the aerospace and airline industries. Centera makes this possible by allowing database fields to interact with Centera's API to a content address as a pointer to specific objects in the cluster. Thus, when a database request is placed regarding a component part of an aircraft or missile, the application will use the Centera content address to retrieve the information.

Centera's contribution to modern cost-conscious manufacturing and air carrier environments is exemplified in two strikingly obvious ways. The first is the manner by which Centera ensures one copy and one replica of fixed content is stored on the system regardless of the number of times it is used. Thus, when Centera is used to store e-mail in a regulatory environment, it will only keep the original and protect that information via content mirroring or content parity protection, even when that document is attached to an email. Furthermore, depending upon the needs of the organization, Centera can replicate its content items to another Centera in a different geographic location if catastrophic natural disasters are a material business concern.

As airlines and aerospace companies move to implement 100 percent electronic record keeping systems, the efficient storing, protecting, and replication of information within Centera will substantially lower the TCO for its users.

Centralized control of one to many Centera clusters also plays an integral role in lowering the TCO of long-term online storage. Key to lowering TCO for long-term storage of fixed content is Centera's ability to provide a single repository for multiple applications. This translates into measurable savings in cost-per-megabyte of storage for IT budgets as maintenance and training costs significantly decrease. Elimination of multiple repositories also dramatically decreases potential mistakes that lead to financial and criminal liability during litigation. In these situations, when an IT team has to manage several repositories, data could be unintentionally destroyed before a technician is able to modify the retention parameters on one or more repositories. Manufacturing environments utilizing Centera are able to extend retention periods for specific information objects or categories of objects as controlled by the users application, effectively minimizing outside liability potential.

Conclusion

Content Addressed Storage, as implemented in EMC Centera, offers aerospace manufacturers as well as air carriers an organization-wide, managed storage solution for fixed content originating from and accessed by multiple enterprise-class business applications. Centera's ability to meet stringent regulatory requirements can enable manufacturers to avoid fines related to improper data retention and subsequent intangible financial losses from negative perceptions of the manufacturer resulting from publicity surrounding regulatory noncompliance. In addition, it provides a cost-effective networked storage solution to address the need to provide fixed content to governmental bodies, and customers, on demand, 24 hours a day. These points taken in conjunction with Centera's ability to eliminate multiple application repositories thereby reduce IT management and maintenance costs. This increases long-term storage system ROI and makes Centera an optimal solution for fixed content information in aerospace and all its related industries.

About Frost & Sullivan

Based in San Jose, California, Frost & Sullivan is a global leader in strategic growth consulting. This white paper is part of Frost & Sullivan's ongoing strategic research into the Information Technology industries. Frost & Sullivan regularly publishes strategic analyses of the major markets for products that encompass storage, management, and security of data. Frost & Sullivan also provides custom growth consulting to a variety of national and international companies.