THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 May 20, 2002

Retrieving the Needle in the Haystack — EMC's *Centera* Manages by Content

Analysts: Anne MacFarland and Mike Kahn

Management Summary

We all know about the growth of disk storage, doubling or tripling every year. Some of the data stored is managed very well by database management systems and the like, but in no way can the recent growth of enterprise storage be attributed to the collections of structured (orderly) data, even when accounting for large databases that are being created and possibly replicated several time each day. This leads us to conclude that, in the 25-year "rush" to the paperless office of the future, almost every enterprise has accumulated massive amounts of disorderly electronic data that are occupying more than half of our storage real estate. From the initial tidy hay bales of well-structured programs and databases, enterprise electronic information has become a landscape of densely-packed and, in many cases, disorderly haystacks, containing harder-to-index and typically large data types. Think of all of the images (photos/maps/drawings), recordings (music/videos/movies), data streams (seismic/astronomical/scientific), and work objects (documents/spreadsheets/presentations/programs) that are proliferating, by not only creation and revision, but also distribution and duplication. For good reasons (legal requirements/business reference) and bad (sloppy habits), most have great longevity.

- *How many thousands/millions/billions of these objects are stored in your enterprise?*
- *How many are replicated dozens/hundreds/thousands of times? (Think of email attachments.)*
- *How do you know whether the "right one" is being retrieved for a given business use?*

This is the backdrop for a serious business discussion about the (potentially ridiculous) cost of continuing to do business as usual.

For several centuries, we have referred to objects that are stored and retrieved as "files". For the last several decades, we have tended to store each of the disorderly objects as a file on electronic media. (Think of all of the different file types: doc, xls, ppt, pdf, html, jpg, mp3, etc.) Well, our files are out of control and driving storage costs out of control. For evidence, look at the growth of your network attached storage.

Most storage mechanisms, such as disk subsystems and storage networks, have been optimized to speedily and reliably retrieve, manipulate and store chunks (blocks) of data under the control of software (operating systems, file systems, DBMSs, etc.) that manage the chunks being stored electronically. These storage mechanisms are designed to the physical characteristics of the electronic media being used. (Think about cylinders, tracks and physical offsets, and the virtual caches that represent that physicality.) **While there are many applications that manage the storage and retrieval of these disorderly objects (such as the processing of check images, photo IDs, engineering drawings), most every application has had to rely on a file system to map objects into files on a disk or tape – until now.**

With EMC's introduction of its *Centera*, we have what might be the first *content addressed storage system* from what once might have been called a disk vendor. With its available API and its low-cost hardware infrastructure, Centera may be the beginning of a new class of object-oriented storage. Read on to find out why.

IN THIS ISSUE

- Why Reference Data is Untidy..... 2
- The Solution in CAS 2
- Where CAS is useful..... 2
- The Bright Future of CAS..... 3
- Conclusion 3

Why Reference Data is Untidy

There are two components to finding a needle in any haystack:

- The size of the haystack, as compared to the needle, and
- How to discern a particular needle in the hay.

The size problem is obvious. Finding one item in a million or billion requires either a lot of table-based indexing or a better way to index. The problem of finding information relates to the limitations of location-based data organization using file systems.

Back in the days of paper files, the context of reference material, particularly in libraries, was given by its call number and place on a shelf organized by subject. You could (in most cases) go to the shelf and see what was there. While plain text search capabilities may exist, they require a dedicated infrastructure to be timely, and much electronic reference material lacks this subject-oriented accessibility. The information in the "blobs" of e-mail allotments or x-rays is only what is afforded by the file systems in which they are stored. These file systems are locationally-based directories of "what is where" with no inherent contextual component to their organization. The underlying file system attributes (name, location, data type, creator, date created and date last modified) work well to control and migrate data where staleness counts, but work poorly when content is the key. Some of us groom our file structures to suit our reference needs, but these systems are sporadically implemented and idiosyncratic.

Content Addressed Storage as a Solution

With its new *Centera* storage solution, EMC introduces the concept of *Content Addressed Storage* (CAS), which stores data as objects, not as files, or in allocations of blocks or LUNs, Each object has attributes, including a unique hash-algorithm fingerprint. Centera stores data only once (and mirrors once to another node, for redundancy) to service everyone. In addition, the object orientation of the solution will allow distributed forms of indexing and authentication. Centera will soon have the ability to purge data according to criteria established by the customer.

The trade-off is slower data access – Web-page speed, or seconds rather than in microseconds. This slower access, typical of object-oriented approaches, is acceptable in non-transactional environments, and will give a better TCO for these unchanging objects that enterprises must, or choose to, maintain.

Where CAS is Useful

As a groundbreaking product, Centera is of interest to enterprises and application developers who are focused primarily on fixed digital assets, such as content management and distribution. Over 30 software vendors and systems integrators have already signed up as Centera partners.

Content Addressed Storage has particular relevance to a number of vertical markets. The medical community relies on static information as much as it does the dynamic information from sensors and the skill of its doctors and nurses. It needs to be able to assure the authenticity of the data, and, for regulatory reasons, must document who has accessed it. The exploding volume of patient-related digital assets, and the decentralized nature of many healthcare provision systems both stress traditional storage architectures. A digital object-oriented storage platform can fill many of the patient record and medical imaging needs of medical institutions, and may enable a solution for patients who travel and suddenly need access to their x-rays.

Similarly, **financial institutions rely on the availability of fixed authentic reference assets (customer contracts) as well as real time transaction and trading information.** Check images are a significant pain point at present. Insurance policy claim files (now bristling with digital images) are another kind of data that must be accessed at multiple locations. These present a problem of unwieldiness to traditional storage architectures – one which content addressed storage can manage well.

Pervasive Uses of CAS

The bulk of e-mail archives, increasingly essential to project documentation, can be minimized by the ability of CAS to identify duplicate objects, and to store the data only once. Moreover, at the termination of a project, or other appropriate time, all the relevant files

Centera Hardware

Centera is a cabinet containing 32 nodes. Each node contains an 850 MHz Pentium III processor with 256 GB RAM and 4 disk drives of ATA (EIDE) storage, and three Base T 10/100 connections. The disk drives are currently 160 GB: capacities will increase as storage density improves. These nodes can be used as storage nodes or as access nodes, in which case the disk drives are used for internal cache. The nodes are configured as a Redundant Array of Independent Nodes (RAIN). Depending on an enterprise's data access needs, a cabinet will have two or more access nodes. The access nodes will use clustering to load balance within the cabinet.

The cabinet also contains two switches (for redundancy). Each module is connected to each switch, and also directly to the IP LAN. All components are auto-configuring and self-monitoring. If the performance of a module degrades, the unit automatically fails over to another module, recovering the data from its mirrored location on other disks. As long as there is adequate storage capacity in the cabinet, there is no need for immediate repairs and, since each cabinet holds 10 TB (mirrored), this may be quite a while.

The use of commercial rack-mount server nodes enables a refreshing new approach to service. Unless storage capacity thresholds are breached, the replacement of non-functional modules can be done at semi-annual "groomings" that are included as part of the basic service. The replacements will, over time, include larger capacity drives, but this mixing of disk size is is supported, and benefits the customer.

Asynchronous replication between paired Centeras at different sites could allow each site to be the remote disaster recovery site for the other.

Sixteen cabinets can be configured together as a single cluster, which appears as a single "pool" of 160 TB of storage of protected capacity to the application server. Up to seven clusters can be managed as a domain, for a total of a petabyte of usable (i.e., mirrored) storage.

The Centera plugs into a network like any other IP-based device. Application servers will use the Centera API to send data to Centera, which returns the Content Address, which is then stored in the application server. If the data is moved or recovered from the mirrored copy, its Content Address is unchanged and retrieval is assured.

can be easily archived or purged.

Document and video processing and other applications dependent on data sharing will be enabled without involved file system arbitration. With object-based access rights, intellectual property attributions can be maintained.

For all enterprises, CAS offers a storage method for permanent records and proprietary information that optimizes both access and guaranteed authenticity.

The Bright Future of CAS

An object-oriented system for data access gives applications using it new opportunities. Like XML as contrasted to HTML, there is an opportunity to implement more hooks. An object's attributes could include elements to facilitate business intelligence and knowledge management applications. Retention attributes could address retention scheduling dilemmas posed by the limitations of traditional, file system-based architectures. The fingerprint attribute of an object becomes critical to establishing an audit trail of distributed business processes, giving provenance to data from diverse sources. Additional attributes could specify the ways data is presented to different applications.

Conclusion

As enterprises fully embrace information systems, there is a need to address the needs of the information that helps one think as contrasted with the needs of the information that helps one act. With this new platform and product, EMC continues to meet the emerging needs of the enterprise. With CAS, enterprises can forget where in the array their data is and just use it.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Authors

Anne MacFarland **is Director of Enterprise Systems Research with The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at AnneM@clipper.com or at 781-235-0085 Ext. 28*** *(dial 128 when you hear the automated attendant)***.**

Mike Kahn **is Chairman and a cofounder of The Clipper Group.** Mr. Kahn is a thirty-year veteran of the computer industry. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis, and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation — Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull — Director of International Marketing and Support; at Honeywell Information Systems — Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. — a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan where he earned B.S.E. and M.S.E. degrees in industrial engineering.

- ***Reach Mike Kahn via e-mail at MikeKahn@clipper.com*** **or via phone at** ***781-235-0085 Ext. 21*** *(dial 121 when you hear the automated attendant)***.**

Regarding Trademarks and Service Marks

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.