THE CLIPPER GROUP
Navigator™

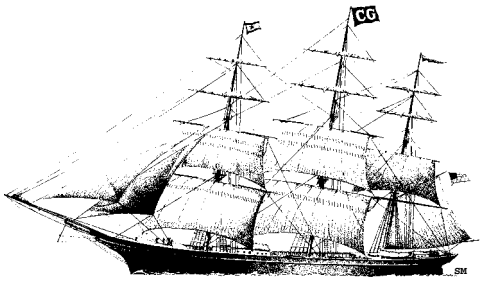

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2004091 November 10, 2004

Update — Archivas Liberates Fixed Data for Use by Multiple Applications[1]

Analyst: Anne MacFarland

Management Summary

Archivas' Founder and CTO Andres Rodriguez was CTO at the New York Times when its digital archives were established. Over the course of setting up the massive repository (10-15 TB on line, ten times that on tape), it was brought home, again and again, that **IT storage infrastructures were designed around drive fallibility and optimized for device management, not for information use.** Moreover, these traditional infrastructures assume that information will be used through the gateway of a single application. This makes the use of information by people using different applications an inherently complex procedure. Nevertheless, as business processes are consolidated and analyzed seven ways from Sunday, this is exactly what businesses want to do.

Rodiguez founded Archivas, Inc. of Waltham, MA, and set out to design the software to support a proper fixed-content repository of information that would be useful for decades, with the following specifications: The repository should free the information from the clutches of the application that generated or captured it, but the repository should harvest all the relevant metadata from the application and store it with policies governing use of the data as an object. It should isolate *management of the information* from *management of the hardware* on which it was stored, to enhance and optimize the former, and automate the latter. It would have to scale to thousands of Terabytes, which meant it should withstand multiple points of failure. Archivas also wanted sub-second response time to meet the needs of business analysis, which dictated rampant decentralization of intelligence and parallelization of functionality. They wanted the repository to meet the retention requirements of government regulations. Of course, the archive had to be able to be replicated remotely. Of course, it had to use standard, open protocols and interfaces.

The Archivas Cluster, *ArC,* is a solution for enterprises that need the many-to-many data access capability described above, but there is more to the story. **Archivas' approach also addresses a set of problems that have been a serious worry for everyone charged with long-term retention of electronic data.** The need to maintain obsolescent applications and the outdated hardware they required to access data will inevitably become an impossible burden over time. This takes the Archivas solution from meeting the immediate needs of compliance with government regulations to a larger scale satisfaction of the requirements of any long-term data retention.

If you need to access data from multiple applications, or if your use of data will exceed the lifespan of the application that generated or captured it, *read on* for more information on Archivas' data liberation technique.

IN THIS ISSUE

[1] This is an update to **The Clipper Group Navigator** published on June 18, 2004, entitled *Archivas Liberates Fixed Data for Use by Multiple Applications*, available at http://www.clipper.com/research/TCG2004051.pdf, and contains additional information.

The Archivas Process

When information is imported into the *ArC* repository, policies are set, and the metadata is harvested from the application to give the archive object self-sufficiency. The retention period that is set cannot be curtailed. ArC has an internal, inviolate clock mechanism to prevent evasion of the retention mandate, so Archivas is a good candidate for the software portion of an enterprise compliance solution.

Archivas' early focus is on government and media customers, whose multiple users using multiple applications need to access the same fixed content. A good example is NASA, an early customer whose customer base for Hubble images uses a variety of applications to analyze them.

The Archivas Architecture

The ArC repository, like many fixed content repositories, is a cluster of intelligent but relatively limited-capacity nodes. (See *Hardware Supported* box on the next page for what they support.) The ratio of storage capacity processor gives an expedient *find* experience akin to using a library of well-indexed books of at most a few hundred pages – or a multi-volume encyclopedia.

However, a cluster of intelligent elements is common in technology. We must look more closely at what functionality is put where to see how Archivas differs from other clustered repositories.

Objects

The nodes contain *archive objects*, aggregations of the information, together with the metadata relevant to it that has been gleaned from the application, and the policies (retention, migration, authentication, and protection) that govern its use. Such enrichment of data into objects facilitates a distributed approach to access control that allows the archive to scale very large. It also will allow the authentication and usage policies to be of object-level granularity, though this capability is not in the early release currently shipping[2].

These objects are replicated to other nodes to protect the integrity of the data in a larger scale version of RAID, called RAIN. At first release, Archivas offers RAIN-1 and RAIN-5, and RAIN N+K, which multiply protects data against multiple sources of failure, in Version 1.1, planned for September 2005.

RAID (or RAIN) N+K allows policies to determine the number of redundant instances of data and the number of parities that support rebuild. Because data is fixed and synchronization is not an issue, this approach is feasible. Multiple instances of data can also improve retrieval performance for large amounts of data.

Objects are inherently extensible. As enterprise processes evolve, the attributes (access hooks) of the objects can be augmented as a background process.

Nodes

The node is the basic functional unit of the Archivas ArC Cluster. Felicitously, each node runs on a Lintel server with embedded storage, something that is easy to visualize. But do not be misled. The node is a software construct, not a hardware element, and is not limited by hardware preconceptions. The physical manifestation could be a wide variety of sizes and media, once it has qualified with the software. Moreover, the hardware can often be what you already have.

Each node contains a full stack of functionality down to the Linux OS that runs on the processor. The most important element is the Archivas *Fixed Content File System* (FCFS), which, together with the distributed cluster-wide database[3], forms a potent data management layer. The FCFS manages and groups the objects by policy. The cluster database provides the relational information to provide views (or symbolic links) to repository users. This enhances the general usefulness of the repository beyond the attributes endowed to a data object when it was imported. It allows an object to be viewed as a member of various collections. This gives the enterprise tools to evolve the usefulness of the repository over time, as needs (and organizational relationships) change.

The Node is supported by two protocols – a

[2] It will be in the GA release due out April 2005.

[3] This linking is made possible by the nature of fixed content.

Hardware Supported

Archivas has chosen to standardize on Intel servers running *Linux* with embedded disk, because it is the cheapest way to get a node. LinuxNetworx, Penguin Computing, SCC, Rackable, Dell, and IBM *xSeries* hardware have been certified, with more in the offing. RedHat Linux has been used. The ArC uses simple storage without RAID. The density of internal server storage can be up to 1TB per 1U rack-mount. Release 1.0 will support 1U, 2U, and 4U nodes. In theory, one could vary node performance by endowing them with direct attached FC disks or tape.

Archivas can support heterogeneous hardware within a cluster.

messaging protocol for control flow and TCP/IP for data flow. Other elements that support the file system include a request manager that administers access to the objects and an administrative engine that configures the system. Beneath that, running on the operating system, is the core management of storage, metadata and policy. Because the data is fixed, there is no need to manage write-access or synchronization.

Cluster Functionality

The Cluster manages the recognition and configuration of newly added nodes using cluster wide policies. It also balances the cluster for activity and capacity, as dictated by the policies set on the objects when they were imported. The cluster synchronizes the distributed database as more objects are added. Administrators do not manipulate files. Human administration is reduced to adding new and replacing failed nodes. Rebuilds and data migration are automatic. The nodes have the intelligence to make failover and quiesce transparent to the end user.

There is a GUI (and, for those who prefer, a CLI), by which administrators can monitor events and their automated remediation, as well as information about nodes, policies, and logs. This monitoring can also give information on how the ArC is used. The cluster can also feed metrics via SNMP up to system management products, which see the ArC as a device. The cluster appears to the application as a simple target – a giant directory or file structure, depending on their expectations. There is no complex data access process. It is just there, like a transparently vast and surprisingly potent instance of NAS.

Archivas Partner Strategy

Archivas is keeping its focus on the repository by partnering for auxiliary functions. Many of the data retrieval needs of certain industries are rapidly evolving (health care is a good example). Archivas has partnered with content management and industry-specific content generating applications to provide the taxonomies to support mapping and indexing. E-mail indexing, for instance, is done at the client side and is limited by what the e-mail application supports. Current partners are iLumin, Connected, Hummingbird, Witness, and Merlin. Archivas partners with Veritas for remote replication of the ArC repository, and with Verity for search. Archivas will partner with hardware vendors to present an appliance-style solution. This is a good start for a company with five early customers, an early release in August, and a general release scheduled for April 2005. ArC Software is a one-time charge of 1¢/MB of data under management, plus an annual software maintenance fee. This simplicity of pricing makes Archivas a good partner for a variety of archiving solutions.

Conclusion

The value of a repository is in its breadth of contents and in its usefulness. Its role is to expose information to users in a way that is simple to use, yet secure. Archivas is architected to do just that, and to grow with time. If a repository of fixed information is a key to your long-term business application support plans, think on what Archivas has done, and how they have done it.

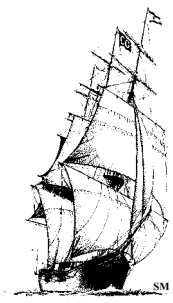

About The Clipper Group, Inc.

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 Author

Anne MacFarland **is Director of Enterprise Architectures and Infrastructure Solutions for 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 Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 28. (Please dial "1-28" 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.