

Impact Report

CAS inventor mounts second charge into fixed-content storage with startup Caringo

Analyst: [Henry Baltazar](#)

Sector: [Storage & Systems »»](#)

Date: 13 Nov 2006

451 Report Folder: [File report »»](#) [View my folder »»](#)

Content-addressed storage (CAS) specialist Caringo is seeking to disrupt the archive market with an inexpensive software offering that starts at \$1,000 per node. Founded by CAS pioneer Paul Carpentier, the startup is trying to build market momentum before accepting VC funding. Caringo only has a handful of partners and customers at this point, but with the edition of its CIFS and NFS file service gateway, its CAS clusters will be able to readily work with many standard applications.

Impact assessment

The message

Caringo is promising a tenfold performance improvement at one-tenth the price of competing products.

Competitive landscape

EMC is still the leader in the CAS market with its FilePool-derived Centera system, but Caringo must also fend off HP, Sun, IBM and NetApp, as well as startups like Nexsan Technologies, Bycast, Permabit and Archivas.

The 451 Assessment

The CAS market is currently dominated by proprietary hardware, but Caringo's software has the potential to change that with its low cost and heterogeneous hardware support. No matter how innovative Caringo's fixed-content storage cluster may be, its technology will only become a factor if quality ISVs can be convinced to support its CAS software. The upcoming file system gateway, which adds CIFS and NFS support, should broaden Caringo's appeal and attract more partnerships. The lack of economical CAS options and Caringo's pedigree give the fledgling company a fighting chance in an archive market that is not kind to startups. So far, no other company or technology has managed to mount a real challenge to EMC's supremacy in the CAS market; Caringo aims to change that.

Context

Founded in March 2005 by industry veterans Carpentier, Jonathan Ring and Mark Goros, Caringo believes it has unique insight into the flaws inherent in current CAS offerings and untapped market opportunities ignored by competitors. Carpentier, Caringo's CTO, was the cofounder of FilePool –a company EMC acquired in 2001. FilePool's technology went on to form the core of EMC's Centera, the first significant product in the CAS market and the current dominant platform for storing 'fixed' content such as document archives, email, records and digital images. Ring, Caringo's president, was vice president of engineering at Siebel Systems, while CEO Goros is a veteran of Oracle and Sybase. Caringo currently has a headcount of 17 but aims to substantially expand the CAS market while challenging entrenched competitors. The company claims its software running on off-the-shelf hardware will deliver 10 times the performance of current CAS products at one-tenth the price.

Caringo is privately funded and is ramping up its partner base before turning to VC funding. The CAS specialist currently has 12 partners and six customers. Caringo's goal is to sell through ISVs and OEMs, but it will initially sell direct to gain traction and assess user deployments. The vendor is targeting vertical markets including medical imaging, video surveillance, government and education; but it could take its product to more mainstream IT shops given its aggressive pricing and the universal applicability of fixed-content storage. Caringo lists Capricorn Technologies, Crossroads Systems, Moonwalk Universal and Storsys as members of its partner program.

With the company's upcoming CAStor File System Gateway (FSG) module, Caringo says customers will be able to connect applications to Caringo storage clusters using standard CIFS and NFS file server protocols. This feature will appeal to customers that either do not want to or cannot rewrite their applications to work with proprietary CAS APIs.

Products

Caringo's CAStor software converts off-the-shelf hardware – x86-based servers with 1GB of RAM and two or more SATA drives – into a scalable storage cluster. CAStor is installed using a bootable USB key that holds the system software and the OS (Linux-based) for the storage node. System requirements for running the software are quite modest, which should keep implementation and upgrade costs down. The CAStor software is priced at \$500 a disk, allowing SMB IT managers to implement a two-disk CAS node for \$1,000. The standard CAStor cluster ingests data using an HTTP 1.1 that Caringo calls Simple Content Storage Protocol (SCSP). The vendor claims SCSP will allow a wide range of devices and computers to communicate with its CAS, in contrast to the proprietary APIs used by other CAS vendors.

CAStor's Zero File System (ZFS) assigns each file a unique ID that it will hold for the duration of its life. Caringo's Content Integrity Seal allows IT managers to update hashing algorithms without bringing the cluster offline. The company touts this as an extremely important feature that ensures data stored on a Caringo cluster will be protected years from now, even if the initial hashing algorithm used becomes vulnerable in the future. Caringo has applied for a patent for its hash algorithm update technology.

As data is sent to the cluster, the CAStor nodes negotiate with each other to determine the best place to write the data. This intelligence ensures that the storage and access load is spread out throughout the cluster without requiring administrator input. CAStor also ensures copies of data are dispersed across multiple nodes to ensure data is not lost if a node goes down. One important limitation CAStor has in comparison to other CAS products is that it does not track and remove duplicate data throughout the cluster. Caringo will be adding single instance de-duplication in a future release in 2007. The vendor claims that the single-instance storage functionality in CAS products today causes performance degradation. Caringo's single-instance storage will be able to be applied on an as-needed basis to reduce the performance impact.

CAStor customers can add storage capacity and processing power to match their requirements. CAStor has the ability to retire older nodes and replace them with newer systems as hard drive and processor technology continues to improve.

CAStor FSG, which is in beta now and is expected to ship before 2007, has a low starting price of \$2,500. The Caringo FSG appears as a CIFS/NFS share on the network allowing users or applications to write data to it for preservation. Archive policies controlling the number of copies to maintain and the length of retention periods are set on a share level. MySQL and Oracle databases are supported, and the FSG has the ability to support Active Directory and Kerberos authentication methods. Caringo plans to support the eXtensible Access Method (XAM) standard when it is released.

Competition

EMC's Centera is the product to beat in the CAS space. Centera is entrenched in the market, and it has many ISVs that have configured their applications to take advantage of its strengths. Hewlett-Packard has two CAS offerings. The HP StorageWorks Reference Information Storage System (RISS) is somewhat similar to Centera and is the fruit of its acquisition of Persist Technologies in 2003. HP's second CAS product is the vertical HP Medical Archiving Solution, which uses technology from Bycast. Both HP and EMC have large sales and marketing arms, which will aggressively fight off all CAS competitors. Also, IBM resells Bycast and has its own Data Retention 550 product for fixed-content storage.

Other competitors to follow in the space are Sun Microsystems with its IntelliStore appliance and upcoming Honeycomb archive, and Nexsan Technologies' Assureon appliance. Other startups to watch are Archivias (which partners with Hitachi) and Permabit.

With the XAM standard still in development, many CAS implementations still rely on proprietary APIs to move data from applications to the archive. EMC and other vendors have added file service gateways but at many times the price of Caringo's \$2,500 gateway.

Caringo allows IT managers to run their CAS software on whatever hardware they choose, and with the FSG and its HTTP support it also does not require a proprietary API to absorb application data.

One factor that will make Caringo look favorable in the eyes of ISVs is the fact that it, unlike EMC or HP, is not a potential competitor in the application space. For example, an email archive software vendor would be more inclined to integrate with Caringo instead of HP's RISS, which already offers integrated email archive capabilities for its CAS.

Partnerships will ultimately determine Caringo's success, and it needs to prove to IT managers, ISVs and integrators that it will be around for the long haul. Unlike server architectures and OSs, which can be replaced relatively easily, CAS products must literally pass the test of time since the information entrusted to them needs to be preserved for several years.

SWOT analysis

Strengths

Inexpensive and hardware agnostic, Caringo's technology can fill obvious holes in the CAS market, and it has enough innovation to make competitors uneasy. Its management team also includes the original inventor of CAS.

Weaknesses

Caringo is a fledgling company with few partners, two weaknesses that stand out in a market where stability and longevity are essential.

Opportunities

The fixed-content storage needs of SMBs are currently not being met. Since it is a software-only play, server and storage hardware vendors should be receptive to Caringo's technology.

Threats

The fixed-content market has no shortage of competition, both from giants like EMC and from startups.