

## Actona's WAN File Services for Active Data — It's as Easy as Cache

Analyst: Anne MacFarland

### Management Summary

The trunk of the tree may be what supports its size, but all the functional stuff is happening at its edges – the roots, the leaves, flowers and fruits – and in the layers of the bark that transport the sap between the different extremities. Enterprise IT systems have evolved to have a similar structure. Data center environments support the scale of business processes. Networks of many kinds transport information from here to there. Moreover, most data is produced or captured (along with most revenue) at edges, often far from the data center.

Most of this data is far more than the blips of blocks transmitted by point of sale (POS) devices and kiosks. It is business records: correspondence, papers, proposals, transcripts, presentations, etc. All of these are critical information to every enterprise. As we have switched from the interruptiveness of meetings to less intrusive e-mail, the volume of these records has ballooned. The challenge becomes how to manage, protect and make available all of this information with little or no local support expertise at all these branch and remote locations. **Centralization of information is significantly easier to manage and makes data protection more efficient.** However, accessing data across the WAN in hundreds of branch operations, each involving perhaps fifty or more knowledge workers each generating and reusing files of various sorts is untenable.

**Information from these branch offices can be pushed to the data center for safekeeping in a variety of ways, all of which must be managed.** In traditional architectures, the information is then pushed back out to the edges for reuse when requested. A tree with such a recursive architecture would not thrive without an abundance of resources. Neither does an enterprise. Supporting branches the old way consumes a lot of resources, is expensive and scales poorly.

Actona Technologies, Inc., has addressed the problem using a different approach. It has taken the old paradigm and switched the direction of the arrow, so that active data files are cached on the inbound trip to the data center. **Actona has changed write then forward, to write-to-cache AND forward. This automatic and minimally intrusive process keeps full functionality locally while generating full business documentation at the data center.** This caching with an in-bound emphasis (but an out-bound capability) can be used to support hundreds of knowledge workers at hundreds of locations without the management and logistical complexity of previous architectures. For more details, read on.

### IN THIS ISSUE

➤ The Scope of the Problem .....	2
➤ Actona Rethinks WAN File Services .....	2
➤ The ActaStor Solution .....	3
➤ Demonstrable ROI .....	4
➤ Conclusion.....	4

## The Scope of the Problem

Large and complex enterprises can easily spawn dozens, even hundreds of branches and other kinds of distributed work units. These branches produce business, revenue, and lots and lots of data, much of it in the form of files of various sorts.<sup>1</sup> Locally-stored data is easy to access – but it also must be actively protected. Despite attempts at standardization, branches vary and so do their data protection needs. Data services must be implemented and monitored for each location. **Remote management of these distributed IT environments can compensate for a lack of local IT skills, but it does not reduce the management burden. This is a huge, expensive problem that grows with every new location that is added.**

Enterprises have attacked the problem of remote file services with network solutions – thin client options and browser-aggregated portals. When read-write access is involved, this can set up an unfortunate tradeoff between telecommunications costs and functionality/performance. Every action button involves considerable messaging and page reloads. Any limitation to functionality, performance, or perceived environmental richness generates user resistance.

It is useful to step back and think about the basic requirements of actively used enterprise data files. Such data needs to be easy to use and it needs to be protected from loss or corruption.

Usability, the aspect that matters to the client-end user, involves the following criteria:

- Full-function access to all the relevant information
- Reasonable response time
- Currency of data shared with others

Protection can be thought of as *or else demands* for the following:

- Integrity, data protected as created;
- Restorability, copy secure and available; and
- Recoverability, done quickly and seamlessly.

**The answers to the demands for data protection are most efficiently provided in a well-staffed data center.**

---

<sup>1</sup> Point-of-Sale device data and other forms of limited data capture are a separate problem, not addressed by Actona *ActaStor*.

## Actona Rethinks WAN File Services

**Actona has rethought how WAN File Services should be done, working not from existing architectures or device capabilities, but from the branch's file service requirements, discussed previously, for data protection and data access.** They have turned the traditional outward-push of branch support and fixed content management around, optimizing the inward push process instead. They use the highly-evolved routines of cache on that inbound information transfer to keep a copy of active information<sup>2</sup> local for quick access, while pushing the master copy of the data to the core, where data protection, migration, and other services can be done efficiently. Should local files become unusable, the masters can quickly be accessed over the WAN using advanced compression and protocol optimization techniques. This front-side data replication addresses the need for data resiliency, reduces the need for management at the edge (the cache is self-managing) and reduces dependency on and use of the network. But there is more.

By propagating only segment-level or block-level changes, and propagating them at file close, Actona changes this information push from a big batch headache to a small incremental pause-on-close routing – a pause which secures the data and, in doing so, secures the documentation of business process for internal and governmental review. There is great value in these small pauses.

Actona's intellectual property is all software. The solution is delivered at the branch as *ActaStor EdgeServer*, an access gateway that users see as a file server, and at the data center as *ActaStor CoreServer*, a stateless, pass-through device that can sit in front of (or on) a NAS gateway and serves edge requests. The software can run on any Intel server running *Red Hat Linux*. Actona has qualified their software on Dell, HP, and IBM servers, so enterprises will be able to reuse existing equipment – or add the software to existing Intel nodes.

***ActaStor* is targeted at very large enterprises with distributed knowledge workers, whose demanding use of files renders traditional network-based modes of support, such as thin clients and portals, astoundingly expensive.** Actona's *EdgeServer* software can

---

<sup>2</sup> Both data and metadata.

generally handle 100 or so users (depending on the ferocity of their needs). It's *CoreServer* can service up to 100 EdgeServers, and are usually installed in an N+1 active-active cluster. Typical connectivity latency is 60-150 milliseconds, although Actona's *ActaStor* can support higher latencies.

Standard benchmarks are not designed to be useful in analyzing distributed environments. So, Actona put together a weighted average of typical office productivity applications doing typical open/modify/save/close functions, to demonstrate the performance possible. (See graph, below.). ActaStor is not as fast as a file server on a LAN because of the need for global coherency, but is close enough so that most users don't notice, and gaining the benefits of storage consolidation, data protection and global access at near-LAN speed becomes an easy decision to make.

### The Acta-Stor Solution

#### Cache Architecture and Degrees of Coherence

Actona employs asynchronous caching with validate-on-open and synch-on-close. If there are new changes to the file from other sources, these will be caught by the validate-on-open, and those changes will be streamed to the edge server,

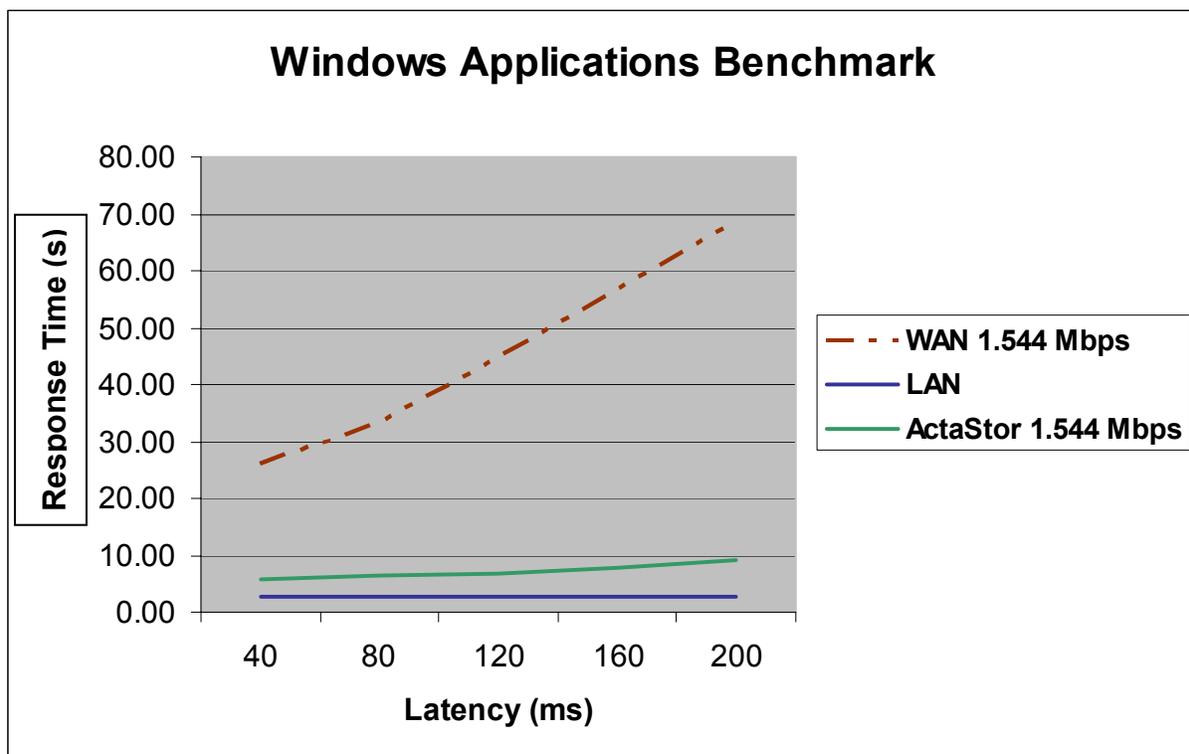
ensuring that the edge user has the latest copy of the file. The synch-on-close feature ensures that all data is propagated to the data center upon issuance of a close command. This guarantees that data will not be lost after an application is closed or if users exceed storage quotas without realizing it.

#### Coherence comes with multiple options.

The central administrator can specify that files which will only be used locally (like those in a *MyDocs* folder), receive only **local coherence**. With local coherence, the changes are sent on to the data center asynchronously as network conditions permit, and handling of CIFS commands is highly optimized. This minimizes response time and networking costs.

**Global coherence** is for documents that are shared between sites – and assures that all users sharing the document are working on the same current “page.” And, when the file is closed, Actona will notify any user waiting for write access to it.

Finally, there is **strict coherence**, where every change is propagated synchronously over the WAN. This causes perceivable pauses, but is used only occasionally, mainly for certain meta-data files used by some applications.



Source: Actona

There is always a trade-off between pervasive data freshness and system performance. Users will know what matters most - and where - and can choose the level of coherence, within their scope of authority.

### ***CIFS and Network Optimization***

Actona has developed its own native CIFS server, rather than using the translation services of an application like *SAMBA*<sup>3</sup>. This means that **authentication and authorization are fully supported, and there is no need for a privileged user access at the core** - a potential source of security problems.

Actona uses compression, but much of the network expense is not the bulk of the files themselves, but the bulk of the messaging that accompanies file activities, due to the chatty nature of CIFS. The real performance benefits are to be gained by protocol optimization. **Actona's CIFS software can sort out which messaging can be served and satisfied locally and which must go to the core server.** Typically, more than 90% of requests can be satisfied locally out of a department's own active files.

### ***Management***

Cache, in all its many forms, has a long, rich heritage and well-evolved functionalities. Both *EdgeServer* and *CoreServer* are simple, limited-use devices. Management of the Actona system has been designed to be comprehensive, web-based, and automatic. It can plug into larger management frameworks (via SNMP), such as *Tivoli Storage Manager* or HP's *OpenView*.

### ***Additional Features***

Beyond simply providing WAN File services in an optimized way, Actona seeks to reduce the need for management at remote locations. While the bulk of its intellectual property is in its revolutionary in-bound use of cache, it also retains the out-bound receptivity of traditional edge servers. Actona's edge server can support an automated, policy-based overnight push-out of patches, upgrades, and synchronized files from the *CoreServer*.

Finally, because most remote environments need them in a self-managing form, Actona's *EdgeServer* supports local print services. The goal is to keep branch operations simple, keep

data safe, and keep the solution transparent by taking as many managed nodes as possible out of edge environments.

### **Demonstrable ROI**

Actona has developed a spreadsheet ROI tool that includes hardware savings (CapEx), management savings (OpEx), and productivity gains. Users of the tool can discount factors (like productivity) that may be too soft. Actona has said that many of their customers find the ROI tool a conservative calculation. Because Actona addresses the problems of high-function branch environments, there is a multiplier effect, based on the number of workers in an office, the number of offices, the amount of data, and its growth rate. Typical ROI is under six months, and soars over time, and with wider deployment.

### **Conclusion**

**Actona has done several things right.**

1. **It has separated out the problem of active data in distributed environments** from other solutions, and targeted how to provide a good user experience and full functionality while also getting all enterprise data under active management expeditiously.
2. **It has achieved the needed enterprise coherency** as inconspicuously as possible.
3. **It has targeted a lightweight approach** that minimizes the need for operational management, using all the automation traditions of cache.

In short, Actona provides a solution for a painful problem, using hardware and networks that most enterprises already have, without any side effects. Mother Nature would be proud.

Does this sound like a dream you might have in a hammock under a tree? Think of your branch offices, unsecured data fluttering like leaves in the breeze. Think of how *ActaStor* lets distributed work groups more effectively do their work and share it with others. Think of how, if all the distributed data was automatically, incrementally, and unobtrusively trickled to a data center for protection, you could reduce risk of unprotected data hiding out in the branches. You could then use your enterprise knowledge assets more effectively. And this would be good.



<sup>3</sup> Open Source software developed by Robert Tridgell. More information can be found at [samba.org](http://samba.org).

### ***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](http://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](mailto: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***

***The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and “clipper.com”*** are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *“Navigating Information Technology Horizons”*, and *“teraproductivity”* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***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.