

Revivio Optimizes Time-to-Restore

Analyst: Anne MacFarland

Management Summary

We are creatures of pulse, and rhythm seems a natural way to organize out life processes. The habit of managing routines by the frequency with which they are done is as natural as using recipes for cooking and measuring ingredients by spoonfuls or pinches. Methodologies and scheduling are key modes of control. We have brought this perspective to technology in many ways beyond the clocking of processors.

Now, consider the need for data redundancy to assure data persistence and integrity, historically known as *backup*. This data replication addresses three vulnerabilities: wide-scale disaster, malicious tampering, and the far more frequent *oopses* that need an *undo*.

The problem of any backup schedule or snapshot approach to data redundancy is that it pits the scheduled frequency against the resource cost of performing the replication in an enterprise game of chicken. You end up scheduling snapshots and full back-ups as infrequently as you can get away with, and hoping for the best. Like any form of insurance, you are betting against fate, and praying to the law of averages.

Time-stamping data gives administrators a way to take frequency, and the betting or praying, out of the equation altogether. With time-stamping and modern processing power, we now have the capabilities needed to pull together data sets by time as well as by application. This moves the point of control from the *physical level* of assigned volumes or back-up tapes to the *logical level* of aggregated data pointers to a set that documents a point in time. By managing the logical views of point-in-time data sets, the administrator can quickly assemble application data as it was before an event.

Automated recovers to an aggregated *any point in time* is can be much faster than recovering from point-in-time snapshots, particularly if your budget for snapshots is limited. It also gives a better way to look into the function of an application over time, and a wider choice of data sets for modeling and predicting behavior. While not a process for deep archiving or long-term retention, with remote replication it does support disaster recovery. And it gives the flexibility of *undo* at the application level to support recovery from both casual and malicious unfortunate events.

Revivio is one of the companies offering what they call Time-Addressed Storage (TAS). If the ability to easily recover to a point in time would help you sleep at night, read on for more details.

IN THIS ISSUE

➤ Parameters of Restore	2
➤ Who is Revivio?	2
➤ Revivio Time Addressed Storage Basics	2
➤ How Revivio Restoration Works	2
➤ Conclusion	3

Parameters of Restore

System restore has a well-established litany of process. First you analyze what failed to determine where to recover from. Then you access that data, move the data so the application can recognize it, and apply transaction logs to roll the data forward to a near-current state. This process takes considerable time, even for small errors that must be undone.

Today's enterprises have new demands for recoverability.

- Applications must have the resilience to recover transparently from small glitches. There must be a pervasive ability to *undo*.
- Many enterprises need a 2-hour recovery from *any* disaster for core applications.

This is a substantial challenge that traditional back-up and restore methodologies cannot meet. *Time Addressed Storage* (TAS), for which Revivio has 14 patents pending, can instantly recreate any point in time by pulling together the requisite data sets. It is like a spare tire that can be installed without stopping your car.

Who is Revivio?

Revivio launched its *CPS-1000* product in October of 2003, concentrating on large financial services enterprises in the Northeast. Early customers began testing with a single application, or in labs. These early customers are now preparing for a broader rollout.

Customers have found that Revivio not only changes how they protect their stored data, but it changes how they run their operations. One trading company wanted to set up an external website for client companies outside the internal firewall. They didn't want to back up through the firewall, but needed more than local back-up, so they used Revivio. Now that they are familiar with the process, they see it as enabling opportunistic decentralization of process.

The body of customer experience has validated the Revivio process. It is time to examine this process more closely.

Revivio Time-Addressed Storage Basics

You may think that an application knows

only the physical location of its data, but administrators have been using redirection since the memory maps of RAID were first invented, to access redundant copies of data should the primary data become corrupted or unreadable. This virtualization of RAID was invented to address the physical fallibility of drives and media.

Revivio's TAS uses similar memory maps, but they are organized, not by location, but by time. This process of organization is a complex, high-speed block mode process, like a RAID controller with a clock. It works with database data, modern file systems, and even with Microsoft *Exchange*¹.

Revivio's CPS-1000 is a self-sufficient appliance², containing a cluster of processing blades (more may be added) connected by a proprietary high-speed interconnect³. There are no disks in the box. Storage must be provisioned to Revivio for storing the time-stamp maps and near-current data, as well as a pool to store a copy of current data. The segregation of these two pools of data (see the exhibit on next page) is key to supporting a near-instantaneous restore.

How Revivio Restoration Works

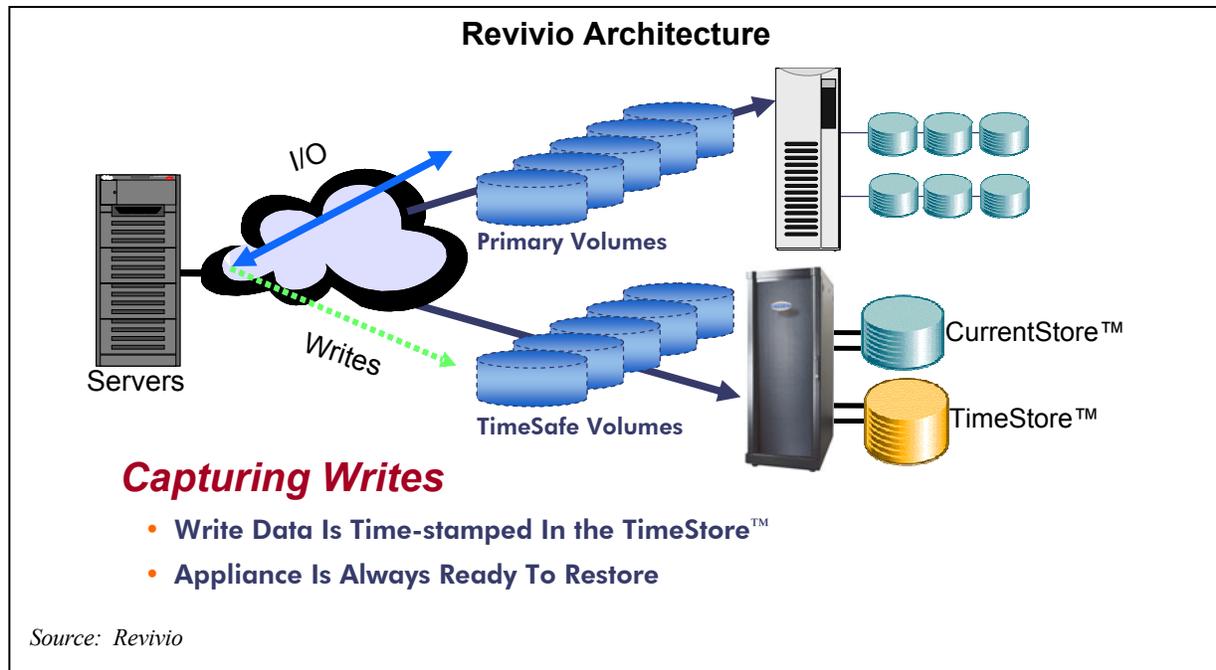
While Revivio could be used as primary storage, most customers use Revivio, like RAID, as a local replication, or *secondary storage*. The CPS-1000 receives this secondary, synchronously mirrored, *TimeSafe*, copy of the primary data volumes.⁴ It stores this data in storage directly attached to the CPS-1000, the time-stamped data in a *CurrentStore* pool of high-performance storage. There is another pool of data storage, called the *TimeStore*.

¹ Exchange has traditionally been hard to back up because it is not transactional in nature (Future editions will be less intractable as they are moving to the transactional environment of SQL). Revivio addresses the problem by protecting the files and data structures as an integrated, time-stamped data set. A virus, too, will be time-stamped, which enables the system to quickly determine a safe point in time to retreat to.

² Revivio CPS-1000 does not use host-based agents.

³ Of course, all interfaces to the box are open standard interfaces.

⁴ The remoteness of the Revivio appliance to the primary storage is that of fiber channel (40 km). This limitation is not one of signal degradation, which may be addressed in a number of ways, but of propagation delay, which is a matter of physics.



Before data is overwritten with more current data on the *CurrentStore*, the old data is copied to the *TimeStore* pool, with indexing and a time stamp. *TimeStore* contains a metadata database that allows instant creation of a memory map to a point in time. This is not a logging/journaling rebuild, but something inherently faster. The size of this *TimeStore* pool will depend on how much you write, how large the writes are, and how far back you want this instant recoverability to. This will depend on the nature of the application – and the enterprise.

When there is a corruption of data, the data in the *CurrentStore* will also be corrupted, but the *TimeStore* pool will remain pure. The Revivio patented process masks out the corruption, and lets the application recover from a second before the corrupting event.⁵

There is, of course, more to recovery than simply having a clean point in time – the data must be accessible to the application. The *TimeSafe* Volumes delivered from the application in the first place were LUN masked and preconfigured. When they are recovered from the *TimeStore* (as *TimeImage* volumes) the data can be moved to the *TimeSafe* preconfigured LUNs.

These recovery volumes (think of them as

time-shots) do not need the quiesce of snapshots or the administrative overhead of traditional business continuation volumes (BCVs). They can be used for testing and local synchronization, and can also be replicated synchronously to address the system's vulnerability to large-scale disasters.

Conclusion

Revivio's TAS gives enterprise applications a current-state resilience that high-transaction-environments desperately need. It gives application administrators an improved methodology for restore that they may desperately crave. It can give enterprise officers peace of mind, and can remove limitations on how they architect their processes. If you feel these pains, take a look at how Revivio might bring relief.



⁵ Revivio could support sub-second granularity, but most customers do not want 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

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.