

OF A
THREE-PART
SERIES

ALWAYS ON: The Intersection of High Availability and Business Continuity

By Jon William Toigo CEO,
Toigo Partners International Founder,
The Data Management Institute

 CA XOsoft

ABSTRACT

High Availability (HA) has become synonymous with Business Continuity for a growing number of companies. High availability strategies, generally provided via a combination of server failover clustering and disk array mirroring, were once the exclusive purview of companies with both “always on” business requirements and “bottomless” bank accounts—the latter being required to fund a combination of expensive and proprietary hardware platforms and private networks of sufficient bandwidth to carry replicated data.

To those companies that could afford them, technologies such as multi-hop array mirroring and point in time disk mirror splitting provided rapid recovery of data from most interruptions. For everyone else, tape backup provided a hardware neutral and highly affordable recovery scheme—albeit, one with potentially significant recovery time delays.

Products like CA XOsoft WANSyncHA™ have changed this equation. Providing always-on continuity for common and critical application environments such as Microsoft Exchange Server or SQL Server platforms can now be accomplished reliably and affordably.

Since different applications and processes impose different recovery parameters, CA XOsoft also supports the management of a continuum of disaster recovery solutions, from tape backup, to mirroring, to failover clustering. This gives planners the ability to apply the right level of protection to the right assets, while preserving the simple manageability of a much less complex business continuity capability.



jtoigo@toigopartners.com

TABLE OF CONTENTS

Introduction	3
Historical Perspective	4
Current Situation—An Uneasy Blending High Availability with Business Continuity	5
“Assured Recovery” More Than a Tag Line	6
Platform Agnosticism	6
Integrated Processing	6
Unified Management	6
Conclusion	7

INTRODUCTION

Business Continuity has been elevated to a Front Office issue in many companies by a combination of three factors: (1) legal and regulatory mandates requiring data protection and preservation, (2) highly visible natural and manmade disaster events, and (3) commonsense concern with operational efficiency and risk reduction. As a result, a growing number of firms are today reviewing their contingency plans and re-evaluating whether current disaster recovery provisions are equal to the task of business continuity.

Given the "always-on" operational model of most firms in the Internet Age, and the growing dependency of business upon information technology infrastructure and electronic data, developing a dependable business continuity capability almost inevitably becomes a discussion of data protection and infrastructure replacement. While the importance of personnel to the recovery effort should not be minimized or understated, this paper focuses on a strategy for recovering the technological infrastructure and data assets that support key business processes as a means for supporting effective operational recovery from a disaster event.

More to the point, this paper focuses on a strategy for recovering business infrastructure and data assets using WANsyncHA™ from CA XOsoft. WANsyncHA™ is a software product that delivers the means not only to replicate infrastructure and data, but also to integrate other modalities of data protection into a common scheme of monitoring and management. The product has expanded the range of options available to planners for assuring operational resiliency and continuity. It has also enabled the centralization of resiliency services management and provided a method for overseeing the efficacy of the recovery strategy as business and technology change over time.

Thank you for taking the time to review this paper.

HISTORICAL PERSPECTIVE

For the past 30 years, disaster recovery and business continuity planning have been undertaken using concepts and methodologies derived from the earliest years of mainframe computing. Companies seeking to make their automation-dependent business processes resilient needed to choose between tape backup and disk mirroring as the foundation of their recovery strategy. The data protection method they chose was guided as much by budget as by any assessment of the criticality of the data and the business processes they supported.

Then, as now, data has been viewed as the centerpiece of the recovery effort. Absent data, it was—and is—pointless to recover hardware platforms or secure new work facilities. Without its data, the business was dead in the water.

Tape backup has become a mainstay of contemporary data protection and disaster recovery strategy, providing a familiar mechanism for making data copies. Improvements made to the technology over time have led to increased tape capacity and performance and bolstered the efficacy of tape backup, even as industry pundits declared it obsolete.

In the mainframe world, tape backup was well suited to the preponderance of applications that constituted the targets of recovery planning. Given the homogeneity of hardware and operating systems in the mainframe data center, companies needed only their backup tapes and a mainframe-equipped recovery facility (a "hot site") that provided sufficient MIPS to process their workload in order to recover.

With the advent of distributed computing, application hosting platforms moved outside of the glass house of the data center, but tape continued to have efficacy as a business continuity strategy for another reason. In the distributed computing environment, tape provided a low cost mechanism for making data "portable"—that is, for making data capable of being restored to and hosted on substantially different hardware platforms than those used in the production shop. This, in turn, enabled the cost effective definition of the recovery environment using the principle of "minimum equipment configuration"—planners did not require one-for-one hardware replication at the recovery site in order to sustain mission critical apps at subsistence levels during an emergency.

Minimum equipment configurations, which were designed to be adequate to support minimalistic workloads in the hours and days following a disaster, represented a huge cost savings from the recovery strategy maintenance perspective. Tape provided the means to unlock data from proprietary storage platforms and to re-host it on the fly on different equipment. (See Figure 1)

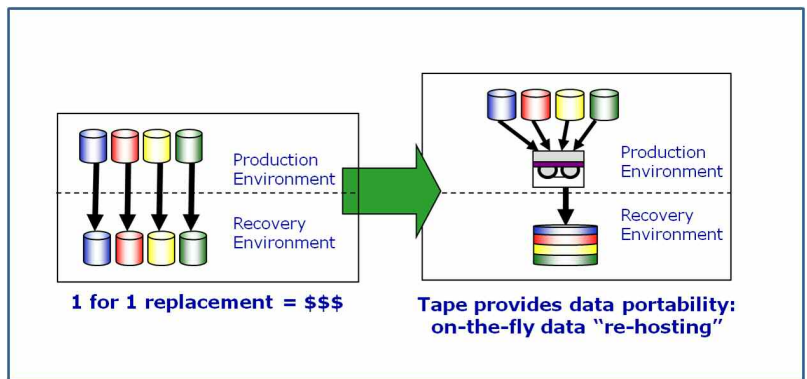


Figure 1

The only downside of tape was its restore speed. Given that a "disaster" is defined as an unplanned interruption in normalized access to mission critical data for whatever constitutes an unacceptable period of time, the essential goal of continuity planning to reduce "time to data." The shorter the duration of an interruption event, the less likely the event is to result in a full-fledged business disaster. With large and growing amounts of data that needed to be restored from tape, combined with the comparatively slow speed of tape to disk transfer, tape increasingly fell out of favor as the preferred means to recover mission critical data.

Companies with deeper pockets and a perceived need for "continuous" or "uninterrupted" operations opted for the alternative to tape: disk to disk data mirroring. Mirroring schemes usually required a company to deploy identical vendor storage gear at both sides of a network connection and to purchase specialized software for replicating changed data on an on-going basis.

In more advanced schemes, multi-hop mirroring was used, providing layers of protection against different threats. In multi-hop mirroring, data is replicated between sites as a "back end" or "behind the array" process that involved three identical storage platforms—two local and one remote. As shown in Figure 2, data written to array 1 was replicated "synchronously" to array 2, which was co-located to array 1. This replication process was made possible by the close proximity of the two arrays and provided protection against a production array failure. Once data was replicated between the co-located arrays, array 2 then initiated a second replication process, this time asynchronous, between array 2 and array 3, which was located some distance from the production environment—usually at a backup facility or hot site. Having array 2 manage this task prevented latency from being introduced into the production environment and provided additional protection against a facilities outage. (See Figure 2)

Then as now, the expense of a disk mirroring strategy was high. Cost components included wide area network connections to link together the production and the recovery facilities, the contracting or building of a recovery facility itself (which needed to be predefined and equipped with necessary hardware), synchronous and asynchronous replication software, and the expenses associated with product lock-ins to a particular vendor and its array platforms. (Typically, mirroring software works only across like storage platforms from the same array manufacturer.)

In the mainframe environment, mirroring was simplified by infrastructure homogeneity. Companies that sought always-on operations and that had the financial resources tended to choose mirroring as a strategy for their most mission critical application data. For less critical business applications and data, tape was integrated into a tiered strategy of data protection and disaster recovery.

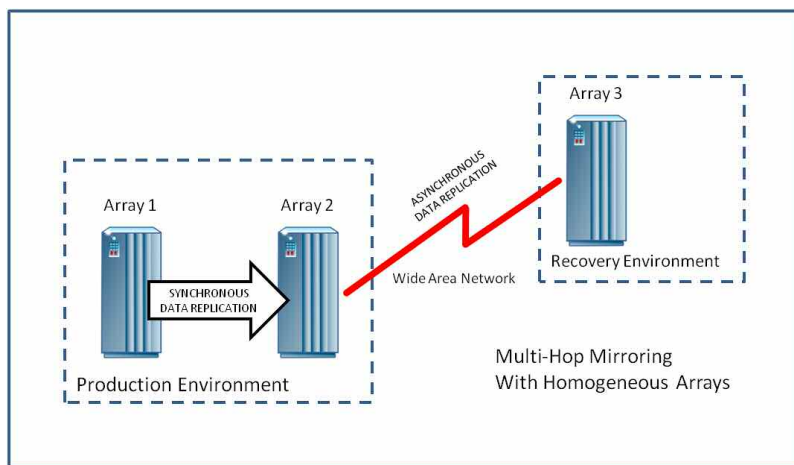


Figure 2

As mission critical applications found their way into distributed platforms, mirroring received a boost by virtue of the fact that companies could sometimes leverage their existing distributed environment to create make-shift hot sites, replicating data between servers and storage devices in branch offices, for example. Further support for this strategy was provided by the Internet, and by falling costs for broadband network connections generally. Always-on or high availability architectures began to proliferate and business continuity, in some cases, benefited from data replication and data protection processes that could be built into the infrastructure rather than being bolted-on after the fact.

Still, the range of options available to planners remained the same as it had been in the mainframe data center: tape backup or disk mirroring. Both provided the means to replicate data, but concerns began to mount that data replication itself was insufficient to recover a business process in a timely way.

CURRENT SITUATION—AN UNEASY BLENDING HIGH AVAILABILITY WITH BUSINESS CONTINUITY

It is difficult to gauge the exact point at which business process recovery usurped systems or data recovery as the new focus of business conti-

nity. It made sense on many levels.

For one, delays in business recovery were increasingly tracked down to the failure of planners to consider the broader business context of the strategies they had created for disaster recovery. Most plans tended to focus on a specific piece of hardware—a hold-over from mainframe-centric planning—rather than on the business process. In some cases, the company successfully recovered infrastructure componentry and data, but did not successfully recover the business process—which amounted to no recovery at all. In short, commencing in the late 1990s, planners began searching for the means to recover the business and not just its servers, storage, networks and data.

Other factors steering change included the advent of high availability (HA) server clustering and server virtualization. Over several years, HA architecture gained a following among those who sought to prevent equipment level, and application level, failure events altogether by creating a shadow infrastructure behind the production infrastructure. Rather than keep the shadow infrastructure in reserve, it was often used in production to co-process workload as well as to provide a failover capability for the application in question. This thinking eventually found its way into business continuity architecture, as well, extending the metaphor of HA beyond the data center and over geographical distance.

Certainly the availability of increasingly standardized computing platforms contributed some cost-efficiency to the strategy of HA, but this can be seen as cutting both ways. On the one hand, standardized servers and network switches made the on-the-fly replacement of processing and networking platforms with commodity off-the-shelf gear much more viable, which militated against the fielding and maintenance of remote shadow infrastructure. However, the opposite spin on the same cost and availability story appears to have held sway. Rather than eschewing shadow infrastructure in favor of on the fly hardware replacement, a growing number of companies are embracing always-on clustering and remote clustering paradigm—and for exactly the same reasons of low cost and standardization.

In effect, as network switches, servers and desktops (though not storage arrays) become standardized commodities, with prices falling

dramatically on a year over year basis, the idea of replicating infrastructure in its entirety and failing over between identical infrastructure has become both more attractive and economically viable.

The exception to this story, of course, is storage. Storage vendors have consistently resisted the commoditization of their products. Instead, the dominant theme in the storage market for the past decade has been for vendors to differentiate their products—by embedding more “value add” technology on proprietary controllers—rather than acquiesce to standards-based commoditization of storage wares.

This has contributed to a bifurcation of recovery requirements in most companies. Part of the recovery strategy must deal with data recovery, while the other part must deal with process recovery. From an infrastructure standpoint, companies find themselves requiring an HA failover solution for their servers and networks, and a separate mirroring or tape backup solution for their data and storage in order to recover a business process. Designing, deploying and managing such a complex and multifaceted recovery capability can be daunting and accounts, at least in part, for the failure of about 50% of companies to undertake any sort of business continuity planning until recently, when they were compelled to plan by legislative or regulatory mandate.

Fortunately, products like CA XOssoft’s WANSyncHA™ have entered the market at exactly the right moment to address the current complexity of the recovery environment.

“ASSURED RECOVERY” MORE THAN A TAG LINE

CA XOssoft adds a new dimension to the tape versus disk mirror discussion that has dominated business continuity planning since the mainframe era. The product features three key capabilities that improve the efficacy of planning activities from the standpoint of both cost and time to data. These are platform agnosticism, integrated processing, and unified management.

Platform Agnosticism

Platform agnosticism refers to the capability of CA XOssoft’s technology to operate with virtually all server, network and storage hardware platforms. Companies can develop production and shadow infrastructure as best meets their needs.

WANSync supports cross-vendor data replication: making data portable between unlike storage platforms. As a result, planners can drive cost out of their shadow infrastructure and implement minimum equipment configurations with unlike hardware. This capability was, up until now, available only in tape backup technology.

In addition to heterogeneous hardware support, also supports failover between differing server and operating system architectures, such as cluster server environments and VMware or other types of virtualized server environments.

This is especially helpful in environments where Microsoft guidance has been observed in creating production clusters of Exchange Mail or SQL Server, but failover is to an ISP or disaster recovery service provider running virtualized servers.

With CA XOssoft WANSync, the business is in complete control of the creation of appropriate business continuity strategies for server, network and storage. Failover logic is “application aware,” meaning that you can design a “scenario” that restores not only hardware platforms, but the entire application environment supporting a business process. Scenario templates are available for common applications right out of the box.

Multiple scenarios can be developed and implemented based on the requirements of the business and the technology assets they use. When the need arises to failover to a recovery environment, this can be accomplished automatically or manually, as preferred by the consumer.

Integrated Processing

Essentially, CA XOssoft enables organizations to centralize the configuration and implementation of individual recovery processes into a holistic, centrally managed, recovery capability. With its robust scripting facility, the technology can even be used to enable the management and monitoring of vendor-specific disk mirroring processes from a single console. While customization work may be required to monitor third party hardware-based mirroring processes today, this is an exciting opportunity for integration going forward.

Better yet, data mirroring between unlike storage environments can be accomplished readily using CA XOssoft’s integral data replication software. Recovery scenarios driven by this platform agnostic engine provide enormous transparency into the replication process and enable at-a-glance confirmation that the data you think you are replicating is actually available at the recovery site when you need it.

Since, as previously noted, planners may prefer to associate different business recovery objectives with different recovery methodologies, it is likely that the business continuity strategy will comprise a mixture of mirroring and tape backup processes. These can also be integrated by CA XOssoft WANSync, which supports not only native and third party mirroring-based data replication but also CA’s tape backup product, ARCserve.

Unified Management

Closely linked to the integration of the myriad processes that planners have engineered to make business process recovery viable is the unified management that CA XOssoft brings to the recovery capability overall. Using the CA XOssoft console, planners can keep track of the recovery strategy in detail. They can see what data is being replicated or backed up,

manage changes in their infrastructure configuration, and administer a host of other tasks that comprise the day to day maintenance of the business recovery capability.

By adding the Assured Recovery option, planners can also validate the failover process and perform inspections of mirroring operations without disrupting the mirroring process itself. These are two very important points that merit additional attention.

By providing a simple means for process validation, WANsync with Assured Recovery facilitates a program of active testing that does not require the allocation of labor resources required in traditional test methodology. At the discretion of the user, remote shadow infrastructure can be activated to determine whether shadow platforms are consistent with recovery requirements.

Truth be told, about 70% of CA XOsoft customers purchase the Assured Recovery option together with WANsyncHA in order to obtain a convenient capability to test the shadow environment without disrupting their production environment. Using this technology, they are afforded something that they don't get from proprietary disk mirroring solutions: visibility into data replication processes.

With most proprietary disk mirroring solutions, the only way to know whether a mirror process is capturing and replicating the right data is to break the mirror. Some vendors have proffered workarounds, such as proxy validation approaches in which two PCs are connected at either end of the link used for data replication. One leading storage vendor asserts that, by successfully failing over one PC to the other, you confirm and validate the integrity of the mirroring process overall. This strategy, however, only confirms the operation of the link, not the underlying mirroring process itself.

Truth be told, many companies have discovered that data they thought they were replicating is not present on the mirrored array when they most need it. This is not necessarily the fault of the mirroring software, but of changes in storage configurations and data locations that have not been accounted for in the replication process. Unfortunately, the inconvenience of breaking the mirror to assess the validity of the replication strategy often deters planners from checking

for problems.

With CA XOsoft Assured Recovery, this impediment is removed. Full visibility is provided into the replication process. If desired, the shadow infrastructure associated with the mirror can be started—while mirroring continues—to assess and verify that the right data will be available when it is needed for business recovery.

Unified management also means unifying the various scenarios used to failover discrete infrastructure components in order to recover a key application or business process. With CA XOsoft, planners can see the entire strategy, with all of its complex interdependencies, on a single pane of glass. This facilitates the activation of the recovery plan in total or in part based on the nature of the interruption event. It also provides the means to test the strategy modularly, which can be a great timesaver.

CONCLUSION

CA XOsoft is on the forefront of a significant change in business continuity best practice. The product is engineered to support the realities of contemporary business computing, while respecting current business sensibilities with respect to business continuity strategy program cost-efficiency and risk reduction.

The solution set provided by CA XOsoft can be leveraged to streamline and unify existing data protection and infrastructure failover strategies into a coherent and verifiable continuity capability. It can simplify testing, reduce management expense, and deliver visibility into the processes that must work if a business is to recover from disaster.

It should be noted that this product does not eliminate the need for business analysis that is the essential precursor to effective planning, nor does it eliminate the need for data management to ensure that data associated with mission critical processes is correctly identified and exposed to appropriate protective services. These are the "heavy lifting" components of sound business continuity planning.

What CA XOsoft can do is provide a way to convert the paper Business Continuity plan into a real Business Continuity capability, which is what the planner set out to build in the first place.



By Jon William Toigo CEO,
Toigo Partners International Founder,
The Data Management Institute
1538 Patricia Avenue
Dunedin, Florida 34698
727-736-5367
jtoigo@toigopartners.com