



Instantaneous Data Protection For the Real-time Business

Something is About to Happen to Your Data.

Are you Ready?



WHITE PAPER

Table of Contents

Highlights	3
Introduction	3
Current Data Protection Technologies	4
Tape	4
Disk	5
Snapshot	6
Replication	6
Growing Business Risks :.....	8
The Wall You Are About to Hit	8
The Bottom Line	10
The Solution? An Integrated Recovery Paltform	11
Real-time Business Meets Real-time Infrastructure	13
About Asempra Technologies, Inc.	14

Highlights

This paper discusses the short comings of existing data protection offerings and introduces a new application availability and data recovery solution for Windows environments, designed to provide real-time data protection, near-time CDP, snapshot, replication and seamless backup integration in a single easy-to-use solution.

**Intel Loses Emails in AMD Case:
Chip maker claims backup procedure backfired during legal discovery process**

**Red Herring
March 07**

Introduction

- **Recovering data at three in the morning only to discover that the data is corrupt.**
- **Struggling to find a point-in-time where the restored data is actually application consistent in order to recover a corrupted database.**
- **Recovering an entire messaging database and then searching through it for hours to find the single piece of information that actually needs to be recovered.**

Unfortunately, experiences such as these are depressingly familiar to an increasing number of IT professionals, who are trusted with the management, protection, and recovery of an exponentially growing body of data. Their task has been made ever more daunting by the increasing deployment of global web-based applications which has turned IT environments into a 24-hour operation – downtime is simply unacceptable.

Horror stories about loss of critical business data exacerbated by recovery failure are legion today, even though IT departments have a plethora of backup, archive, snapshot, and storage mirroring tools at their disposal. To date, businesses that are equipped with the most sophisticated tools still cannot claim 100% recoverability without data loss – and recovery time for these failures is still higher than anyone in IT likes to admit.

So how do you address this situation? IT cannot continue supporting a real-time enterprise with an array of un-integrated backup and recovery tools that are increasing the business risk of additional data loss, recovery downtime, and new capital and operational costs.

Current Data Protection Technologies

On average, 40% of unplanned downtime is caused by operation failures, typically people and process issues related to infrastructure changes...

– Gartner Group

Designed nearly 40 years ago during the emergence of open systems and distributed computing, today's data protection tools take scheduled or manual point-in-time images of critical business data – architected to take a snapshot of your data for a world that once accepted planned downtime. While business has since evolved into a 24x7 world, the state of data protection and recovery hasn't kept pace. Let's take a look at the tools of the trade...

Currently, data protection tools can be divided into four categories:

1. Tape-based backup tools
2. Disk-based backup tools
3. Snapshot tools, and
4. Replication tools.

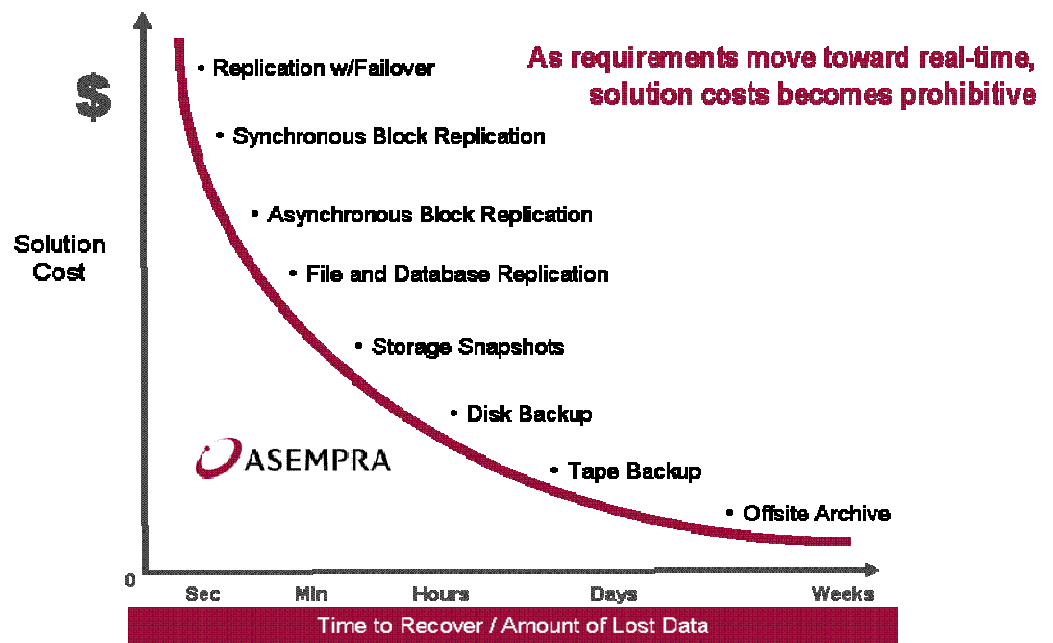
Tape

Tape-based backup tools are usually file system- and database-aware. During backup, this solution requires that the file system be shut down, or the database be put in a quiescent mode, in order to ensure data consistency during the full or incremental copy to tape. The most typical IT process is a nightly file-incremental tape backup and a weekly full backup. This would enable a recovery point objective (RPO) of no worse than 24 hours, assuming the data is actually recoverable. (RPO is a definition of how much data loss is acceptable in a particular application environment.)

Tape-based backup tools usually include media management capabilities which create a catalog that tracks the tape location of all backed up data objects. However, these tools do not own or manage the data on the tapes, they really only manage the catalog.

Disk

Disk-based backup tools are conceptually similar to tape-based backup tools; in fact many of the tape-oriented tools just extended their current product to include a disk-based backup option. The goal of this feature was to reduce backup windows or increase recovery speed utilizing the faster media characteristics of disk storage. Or put differently, to reduce the recovery time objectives (RTO) of the associated application environment. Periodically, these tools copy modified files onto a secondary storage array. Like the tape-based tools, they also create a catalog of backed up objects. On the debit side, the storage consumption of these tools can be quite high.



Snapshot

These tools take a picture of a storage volume at a particular point in time. With snapshot, per file consistency may not be achievable if the file system is not shutdown because at any moment in time a file may be partially updated.

Snapshot tools are frequently used for creating backup copies of a database. During recovery, a hot snapshot of a database is in crash consistent state and it needs to be verified before the database can be fully restored. IT sometimes uses database snapshots mounted on a backup server to perform off-host tape backup in an effort to eliminate the backup window on the primary server. Volume-based snapshots also cannot create a consistent image of a database unless the database service is shutdown or be placed in a hot backup mode.

Although snapshot tools can be integrated into primary data storage, they do not possess any object cataloging functionality. Therefore, it is not possible to perform an object search using such tools. In order to recover a fine grained object (like a file or mailbox), a complete data set (or entire database) must first be fully recovered before the desired data object can be extracted manually.

The most common practice is to perform 4 to 6 snapshots a day, and in some cases, hourly snapshots are taken.

Replication

Replication tools come in three types – block-based, file-based, or application log-based. Each type consumes different amounts of network bandwidth and has different server impact characteristics.

Some replication tools perform synchronous replication, others perform asynchronous replication. Application log shipping replication usually is more bandwidth efficient and has better recovery consistency characteristics. Unlike backup tools, replication tools do not generate object catalogs, nor do they

The cost of downtime ranges from tens of thousands of dollars hourly to as high as \$6.4M per hour.

InfoStor,
Dec., 2006

maintain data history. Both block- and file-based replication tools can generate a crash consistent remote image when failure occurs on the primary server. It takes several steps and significant manual effort on the part of the IT administrator, but the replicated server at the secondary site can recover from the crash consistent image. Most replication tools handle failover well but do not necessary have failback capability. The cost for deploying host based replication is high since remote servers and storage are required.

Growing Business Risks: The Wall You Are About to Hit

Data failures can be caused by device failures, human errors, software errors, and environmental failures. The table below catalogs and summarizes such failures, identifies their consequences, and assesses the risks that are beyond the protective capabilities of the current generation of data protection tools.

Common data loss scenarios and the consequences of using the wrong solution

Device Failure	Consequences	Risks
Storage or server failures.	Loss of data sets or inconsistent data sets.	<p>Snapshot copies may have been wiped out in a storage failure. If not, the administrator would have to spend hours determining if the snapshot copies have application consistent data. Data replication can partially address this failure by enabling immediate switchover to the standby server. However, replication does nothing to address the possibility that the data set is inconsistent and would require repair. In such instances, data loss is likely. Tape backup is always a last resort.</p> <p>RPO = seconds with replication otherwise hours at best</p> <p>RTO = minutes with a replication tool hours with a snapshot tool days with tape backup tool</p>

Human Error	Consequences	Risks
Accidental or malicious deletion or modification of data objects.	Incorrect or lost data objects.	<p>The lost data object may be located in a tape- or disk based backup catalog. Although with tape-based backup, the recoverability of the data object is unknown, owing to the fact that tape-based backup tools do not own the data. One could also manually search through snapshots for the corrupted or missing objects. Data replication does not protect against this type of failure, as such errors would be replicated to a secondary data set.</p> <p>RPO = hours at best</p> <p>RTO = hours to days with snapshot tool days to weeks with tape backup tool</p>

Software Error	Consequences	Risks
Incorrect algorithms or viruses.	Incorrect data objects, inconsistent data sets, loss of data objects or data sets.	<p>Data may be recovered from snapshots, disk-based, or tape-based backup. Data replication does not protect against this type of failure. Alternatively, IT may manually restore the current data using backup images. Recovering from this type of failure is extremely time-consuming. Manual examination of data objects is often required.</p> <p>RPO = hours at best</p> <p>RTO = days to weeks with snapshot or backup tool</p>

Environmental Failure	Consequences	Risks
Power failures, theft, or natural disasters.	Inconsistent data set, or total data loss.	<p>Data replication is usually the solution to manage these types of failures. Tape backup can be the last resort, but could take days to find the complete data set.</p> <p>RPO = seconds with sync replication and hours to days with snapshot or backup tools</p> <p>RTO = minutes with replication and hours to days with snapshots or tape backup tools</p>

"Given the current regulatory and legal landscape, companies need to manage email as business records just like they do any other form of communication. However, the growing volume of email creates challenges for message recovery, search and legal discovery. Asemptra's Business Continuity Server allows IT managers to recover an email, an information store, or a complete Exchange server quickly and with a minimum of effort."

Michael Osterman,
President
Osterman Research

As one can see from the table above, there is not a single data protection tool in existence today that is capable of addressing all common data center failures or the RPO and RTO requirements placed on IT. Existing data protection tools are ad hoc remedies rather than true solutions that address the failure issues at the architecture level.

The Bottom Line

Tape-based backup tools own a catalog but not the data. Snapshot tools have the data but not the catalog. Disk-based backup tools make inefficient use of storage; they do not necessarily eliminate data lost. Replication tools do not maintain data history, although they do perform their function in real time, unlike all other tools.

As a result, IT departments often resort to more than one solution to protect their company data, even though they still cannot guarantee that such data can be recovered with consistency and integrity.

The quest to achieve today's RPO and RTO requirements is an exercise in frustration. In spite of deploying a multitude of

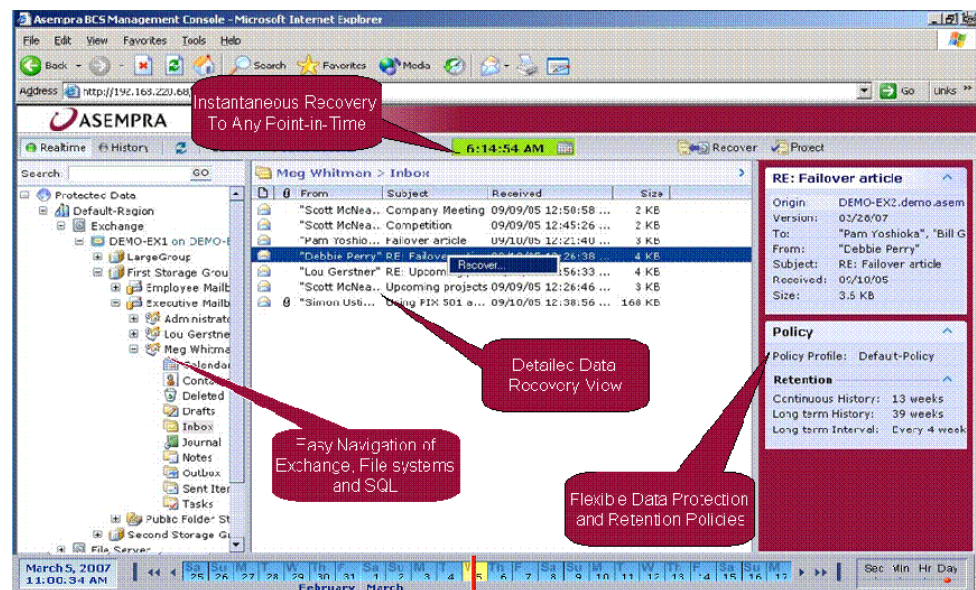
protection tools, unexpected simultaneous (or sequential) failures often force IT administrators to resort to exceedingly arduous and time-consuming manual methods to recover lost data.

The Solution? An Integrated Data Protection and Recovery Platform

Over the last two decades, operating systems and applications have changed dramatically at the architectural level to address changing business needs. In contrast, the development of data protection tools has lagged considerably behind. The available tools have remained essentially standalone products that require manual integration during day-to-day operations and recovery, the woeful inadequacy of which has become only more salient as the gap between business needs and technology development widened.

In order to successfully protect data in today's sophisticated application environment, as well as to survive tomorrow's challenges, one must create an integrated recovery platform that is capable of handling all types of failures, whether they occur singly or concurrently.

A complete, easy to use interface to a data protection solution that provides instantaneous data availability, assured data useability and granular to global point and click recovery



"Data protection is crucial to every business owner. No business can afford the loss or corruption of key data. Yet existing data protection approaches that attempt to ensure business continuity may not protect all the data all the time. For Windows environments, Asempra's Business Continuity Server 2.5 corrects those deficiencies. Asempra provides a rapid recovery experience for applications and data that is measured in seconds or minutes (rather than hours or worse). What's there not to like?"

David Hill,
Storage Analyst,
Mesabi Group

Asempra's Business Continuity Server provides transaction-aware continuous data protection and application-driven data recovery which reduces recovery time for any data set size to seconds or minutes. Transaction-aware data protection is provided through a real-time event and data capturing engine, the BCS Continuous Protection Driver (CPD) which sits on the host server of the application being protected. The CPD continually captures transaction events (I/O-writes along with event and application metadata) coming from a file-server, email, or database server and asynchronously transfers that data in real-time to a BCS cluster. The I/O writes are reduced to byte-level deltas in order to dramatically reduce the network and storage required to save the continuous application history.

The BCS cluster creates and maintains a data structure called the Continuous Object Store which contains these byte-level transaction events and indexes the events across the dimension of time in order to preserve a true application history. The retention of this application history, along with its associated storage and bandwidth parameters, are all automatically managed via the policy engine built into the BCS.

The BCS provides application-driven data recovery which reduces recovery time to seconds or minutes due to a patented innovation called Virtual On-Demand Recovery (VODR). This innovation is possible because the BCS Continuous Protection Driver senses the application I/O requests and recovers that data on-demand so that the application can start running immediately after a recovery is started. Recovery of the data that the application is not currently using is happening in the background while this on-demand recovery is taking place. This enables the application to be available within seconds or minutes instead of waiting for a full recovery to complete which could take hours or days by legacy data protection tools.

Summary: Real-time Business Meets Real-time Infrastructure

The management complexity and cost of solving data protection and recovery issues today is rooted in the fact that it takes multiple tools to deliver a solution that still doesn't meet the requirements of today's data center. This leaves IT professionals spending countless hours trying to integrate disparate tools and manually recovering data in an attempt to simulate a real-time infrastructure that is required to support their business operations.

The Business Continuity Server delivers these results for today's real-time business.

About Asempra Technologies, Inc.

Asempra Technologies
640 West California Avenue
Sunnyvale, CA 94086
Phone: 408.215.5800
Fax: 408.215.5802
www.asempratech.com

Copyright © 2005 - 2007
Asempra Technologies. All
rights reserved. Asempra,
the Asempra logo,
Asempra Business
Continuity Server and
Virtual On-Demand
Recovery are trademarks
or registered trademarks of
Asempra Technologies in
the United States and other
countries. Other names
may be trademarks of their
respective owners.

Asempra Technologies
makes no warranties,
expressed or implied, by
operation of law or
otherwise, relating to this
document, the products or
the computer software
programs described
herein. The information
contained in this document
is subject to change
without notice. Asempra
Technologies assumes no
responsibility for any
errors that may appear.

All other trade names,
trademarks, registered
trademarks and service
marks used and mentioned
in this document are the
rightful property of their
respective owners.

Asempra Technologies is a leading provider of instantaneous application and data availability solutions for Windows. Named "One of the Top 10 Startups to Watch" by ByteandSwitch, Asempra's Business Continuity Server™ enables application availability and data recovery for Windows-based application data in minutes (even seconds), from any location, at any point-in-time. Incorporating real-time CDP, near CDP, snapshot, replication, and seamless backup integration, the BCS provides data protection, disaster recovery, business continuity, and compliance and governance support in a single easy-to-use solution. Tightly integrated with Microsoft Exchange, SQL and File Server platforms, the BCS allows companies to leverage existing infrastructure and reduce management complexity. For more information, please call



Solution Centre Ltd, Vickers House,
Priestley Road, Basingstoke, RG24 9NP, UK
Tel: 01256 818600
Fax: 01256 819600
Email: sales@solutioncentre.co.uk
Web: www.solutioncentre.co.uk