



Evaluating a Business Continuity Solution

Ten Factors to Consider Before Buying



White Paper

Table of Contents

The Need for a Comprehensive Solution 1

 The Evaluation Metrics 2

 Recovery Time Characteristics 3

 1. Recovery Time Objective (RTO) 3

 2. Recovery Time Granularity (RTG) 3

 Recovered Data Characteristics 4

 3. Recovery Point Objective (RPO)..... 4

 4. Recovery Object Granularity (ROG) 4

 5. Recovery Event Granularity (REG) 5

 6. Recovery Consistency Characteristics (RCC)..... 6

 Recovery Scalability Characteristics 8

 7. Recovery Service Scalability (RSS) 8

 8. Recovery Service Resiliency (RSR)..... 8

 9. Recovery Location Scope (RLS)..... 8

 10. Business continuity Cost (RMC)..... 9

 Applying the Business continuity Metrics 10

The Asempra Business Continuity Server 10

 Instantaneous Availability..... 11

 Guaranteed Data Useability 11

 Disaster Recovery 12

Backup Windows Eliminated 12

True, Continuous Real-time Data Protection..... 13

Business continuity Scorecard..... 13

Summary 13

About Asempra 14

The Need for a Comprehensive Solution

Continuous access to Windows data and applications is critical. With downtime costs ranging from thousands to millions of dollars per hour in lost productivity, opportunity costs and hard currency, it is imperative to keep Windows-based ecosystems available at all times.

Higher service level requirements, fewer resources, and ever tightening budgets have made it increasingly difficult for IT departments to meet the challenges of real-time business continuity, data protection, disaster recovery, compliance and governance. Today's global, continuous business operations and regulatory business climates are complex and costly. The use of multiple tools that deliver only partial solutions is not acceptable. Trapped between requirements and incomplete tools, IT professionals exert great efforts - sometimes unsuccessfully - to recover data manually, to protect business Information assets, and to meet the requirements of the firm.

In a recent survey of disaster planning professionals, "Fast Recovery from Disaster" was named the #1 criterion for selecting a business continuity solution.

Contingency Planning
Association

A business can be disrupted by something as mundane as a missing file, a temporary power failure, or as extraordinary as a major natural disaster. The first step in business continuity is to protect (i.e., to be able to recover) critical business data. To handle this diverse set of failures, IT systems must be able to access and recover items as granular as a single email or file to something as global as the entire data center, locally, remotely and immediately. At the same time, it is imperative that IT recover data to a usable and consistent state so that when that data is linked to active applications the applications can utilize the recovered data. It is also essential that a business have the ability to recover its data to any point-in-time or to a point prior to a failure, without losing any critical information.

The need for a complete, high performance easy-to-use solution for protecting business-critical Exchange, SQL and Windows File Server data is clear.

This white paper describes ten parameters against which solutions should be assessed when evaluating application-aware, real-time continuous data protection technologies for Windows environments. The following sections focus on the

metrics necessary to evaluate and objectively measure business continuity solutions.

The Evaluation Metrics

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

InfoStor,
Dec., 2006

In order to evaluate a business continuity solution, one must have properly defined metrics. Data recovery service level agreements (SLAs) are traditionally measured by recovery time objectives (RTO) and recovery point objectives (RPO). RTO defines the time required to recover a set unit of missing data, and RPO defines the potential data loss – the time gap between the most recent application-consistent recovery point and the physical failure point. RTO and RPO may be good objectives for setting SLAs with regard to data recovery, but they are not sufficient for measuring a business continuity solution. For example, a snapshot tool may recover a server's data in minutes; however, a snapshot tool does not have the ability to recover a granular object. When one needs to locate a lost object from snapshots, the process is manual and the RTO could be many hours. In this case, RTO has nothing to do with the tool per se, inasmuch as it is entirely dependent on a manual process.

While a data replication tool is capable of delivering zero or near zero RPO when a server fails, it is not capable of recovering business data if the data is corrupted, and the corrupted data is replicated. As a result, IT needs more comprehensive metrics to properly evaluate a business continuity solution. There are ten core metrics that fall into three categories – Recovery Time Characteristics, Recovered Data Characteristics, and Recovery Scalability Characteristics. The following sections explore these metrics in detail.

Recovery Time Characteristics

1. Recovery Time Objective (RTO)

Time-to-Restore was named as the #1 data protection problem in a survey of IT management

Peripheral Concepts,
Dec., 2006

When applying RTO as a measuring parameter for a business continuity solution, it defines how quickly the solution is capable of recovering the data and application it is designed to protect. A block-level Continuous Data Protection (CDP) and storage snapshot tool can recover a volume or a database in hours - or less if the solution is capable of provisioning its secondary storage volume as the primary storage. A file-based CDP solution can require minutes to recover a file, and hours or even more than a day to recover a file system. The RTO of most recovery solutions depends on whether or not a verification process is needed prior to the recovery and the size of the data set to be recovered. The solution that can break established boundaries to provide instant recovery regardless of data set size would greatly reduce or eliminate business down time.

2. Recovery Time Granularity (RTG)

"...Asempra's Business Continuity Server delivers granular and real-time recovery, serving up the application and data immediately, and with guaranteed data usability."

Lauren Whitehouse,
Analyst,
Enterprise Strategy
Group (ESG)

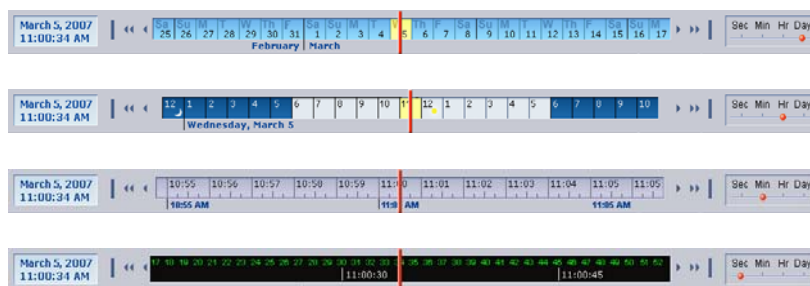
RTG determines the time spacing for selecting a recovery point. This is an important parameter for recovering from a logical failure. Unlike RPO, which determines the last recovery point prior to a physical failure, RTG defines a logical recovery point selection.

A data replication solution may have zero RPO for recovering from physical failure. However, when this same replication tool encounters a logical failure (such as partial data corruption that is not detected for a period of time), the replication tool is not able to recover the data. As a result, the RTG would be undetermined. This situation requires a more sophisticated solution with the ability to detect and provide a recovery point in the past prior to - but also as close as possible to - the logical failure.

The RTG of different CDP solutions, such as block or file journaling, can be completely different, depending primarily on how a particular CDP solution keeps track of data history.

There are two classes of CDP: real-time CDP solutions store their protected data in a time-based continuous data store, while near-CDP tools collect data into snapshots, which are then

stored in the protected storage medium on a predetermined period.



A sophisticated CDP solution will have the ability to support Recovery Time Granularities measured in Days, Hours, Minutes or Seconds

When a physical failure occurs, the RTG for true real-time CDP solutions reduces recovery time objectives (RTO) to seconds; whereas the RTG for near CDP solutions could be hours, depending on the frequency at which the snapshots are taken. By making recovery times independent of the data set size and virtualizing the data set, applications can be restarted within seconds or minutes, unlike other technologies that require a complete recovery of the data set first.

Recovered Data Characteristics

3. Recovery Point Objective (RPO)

As a measuring parameter, a Recovery Point Objective defines the minimum time gap between the last physical failure and the point-in-time where data can be recovered. Obviously, the smaller the time gap, the less data are lost. Since CDP (block or file journaling tools) and data replication tools continuously protect the changed data, their RPO capability is within a second. In contrast, the RPO of snapshot tools is in the range of minutes to hours, depending on the snapshot duration.

4. Recovery Object Granularity (ROG)¹

Recovery Object Granularity measures the level of object granularity a business continuity solution is capable of recovering. To illustrate, object granularity may be a storage volume, a file system, a database table, a transaction, a mailbox, a message, etc. Many storage snapshot tools and block-level CDP tools available today are capable of recovering data of ROG in volume only. To recover an individual file or message, a

¹ David Freund. "Backup is dead. Long live backup!" InfoStor, August, 2004

manual process must take over, sometimes taking days, to find individual recovery objects.

Again, even though some block-level CDP tools claim the capability of recovering a database, the recovery process may be manual and labor intensive, in which case the ROG should not be defined as the database. Rather it should be a volume instead.

The ROG of file-based journaling solutions ranges from files to directories, so while the boundaries differ from snapshot and block-based tools, the underlying limitations are the same: ROG is gross, and the capability to recover down to an individual object like an email or file needs to be done manually, meaning at great cost and slowly.

5. Recovery Event Granularity (REG)

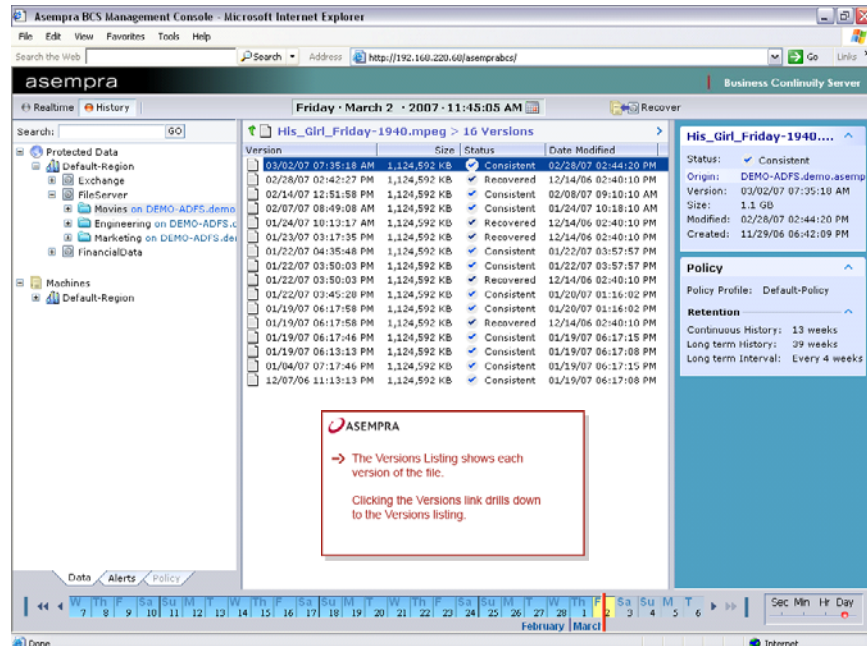
Recovery Event Granularity measures the capability of a business continuity solution to track events and to recover a failed application or missing data to a specific event. Most traditional data protection tools are not designed to explicitly track any events. The only implicit event most of these tools have awareness of is the time at which the data are backed up.

"The Asempra Business Continuity Server simplifies the process by providing real-time backup, with instant recovery and guaranteed data integrity."

Steve Duplessie,
Senior Analyst &
Co-Founder
Enterprise Strategy
Group (ESG)

A true real-time business continuity solution will be purpose-built to allow for highly efficient journaling and sequencing of data changes along with their associated metadata and events during data protection. The data object management processes will generate new object versions as consistent events arrive. These object versions will be managed in a traverse-able timeline. Object metadata will be indexed according to configuration, and index entries will be time-stamped as well to provide search-ability.

In a true real-time business continuity solution – one that supports robust ROG - a data object can be of any granularity, and may have hierarchy. For example, a file is an object, and a database is also an object. A version of a file is generated when the file is closed after a sequence of updates. A version of a database is essentially a collection of files which are at consistent state.



"Asempra's Business Continuity Server simplifies the backup and recovery process by providing real-time backup, with instant recovery and guaranteed data integrity. Asempra has created a solution that's reliable and easy to use, with a price SMBs can afford."

Dianne McAdam,
Director, Enterprise
Information Assurance,
Clipper Group

Metadata-, time-, and event-indexing capabilities are the enabler to track real-time continuous object history, locate missing information, and deliver object recovery of different granularities.

6. Recovery Consistency Characteristics (RCC)

Recovery Consistency Characteristics define the usability of recovered data by the associated application. The RCC of a business continuity solution depend not only on how data is captured and stored, but also on the data type being protected.

When a volume-based hot snapshot is taken on a non-journal file system, the recovered file system is likely to be inconsistent or corrupted, because there are probably incomplete updates to the files, or the file system structure may be in the process of being modified.

When a hot snapshot is taken on a journal file system, the directory structure can be repaired during recovery but the file system, as a whole, is not consistent, due to the fact that journal file systems only journal the file system structure, they do not journal file content. While a journal file system is capable of self recovering its directory structure using its journal, the content of the active files cannot be consistently restored. Since a file system is actively modified by many applications simultaneously

during runtime, the only time that a snapshot can be taken with consistency is when the file system is shutdown - an ironic failure to meet business continuity requirements if using the business continuity tool requires taking the data offline. Since block-level CDP or file journaling solutions are similar to taking hot snapshots continuously, the RCC of these solutions is inconsistent when applied to the file system – an equally ironic case where the data is protected continuously, and known to be corrupt.

“The importance of data places increasing demands on every company’s backup and recovery infrastructure. This poses particular challenges for IT administrators at mid-market companies, who are forced to manage increasing amounts of data with limited budgets and very little staff. Asempra’s Business Continuity Server delivers the enterprise robustness mid-market customers need, and does it at an affordable price.”

Mike Karp,
Senior
Storage
Analyst
Enterprise
Management
Associates

Unlike journal file systems, databases typically journal their own content; they also have built-in crash recovery to repair their content upon failures. During runtime, the on-disk image of a database is usually in a crash consistent state at best, with the state of its log ahead of the binary updates. Unless a database is placed in a quiescent mode, a snapshot of a database can only capture a crash consistent image, an image wherein every open file is known to be compromised.

A database typically and periodically flushes its memory to the persistent storage to complete all the binary updates; this is known as a checkpoint event. At the time point of a checkpoint event, a database is in a strongly consistent state, with its log and binary synchronized with one another.

Since most databases store their contents in file systems today, the RCC of block-level CDP tools is crash consistent at best. Block-level CDP tools must be able to preserve write-order across multiple volumes if a database file spans multiple volumes. Without this, the RCC of these solutions has no consistency. A real-time continuous protection solution if capable of tracking database checkpoint delivers strongly consistent RCC.

A true real-time continuous protection solution avoids recovering inconsistent or corrupted data through the use of comprehensive 3-D journaling (real-time data, metadata, and event) and continuous object store indexing techniques. These elements are combined to preserve a true data history of the application with consistency marking to ensure application recoverability across multiple dimensions of time (second, minutes, hours, and days), object granularity, and application events can have strong consistent RCC only if it tracks consistency events and recovers each individual object in a file system to their own consistency point.

Recovery Scalability Characteristics

7. Recovery Service Scalability (RSS)

Recovery Service Scalability is important in evaluating a business continuity solution. A business continuity solution must be able to scale with the applications and the data it protects. RSS is measured by service (number of applications or data sets the solution is capable of protecting) and capacity (the maximum size of the data it can store). Since CDP captures and processes data in real-time, a true real-time CDP solution must be able to keep up with the applications it protects. The ideal solution should be able to scale easily both in service and capacity through simple addition of processors and storage hardware without major reconfiguration and downtime. Since most CDP solutions today are not based on a GRID architecture, their RSS is usually a cause of concern.

8. Recovery Service Resiliency (RSR)

Recovery Service Resiliency defines how well a business continuity solution tolerates failures. A business continuity service must not cause an application to fail. It must be more reliable than the application it protects. When a business continuity service fails, the data service must fail over to another business continuity instance, such that an application would be continuously protected. A resilient recovery service should not corrupt its protected data; it must be able to self recover from its own failure. The self recovery should not be destructive, and there should be no impact to the applications it protects. A business continuity solution must be secured, such that individuals without the proper authorization cannot freely configure its policy. Unlike an application, a business continuity solution should not allow any individual to alter its protected data. Data history can only be purged by policies.

9. Recovery Location Scope (RLS)

Recovery Location Scope defines where the protected data must be presented when recovery takes place. Most data business continuity solutions, by design, require that the protected data be presented locally before it can be recovered

back to the primary storage. The RLS of replication tools is also LAN because the recovery location must be where the replicated data resides. Internet-based data management services, on the other hand, protect and recover data over the Internet. As a result, their RLS is WAN. As businesses become more global, and government regulatory requirements for business continuity become more stringent, it is increasingly important and valuable for a business continuity solution to be able to support both LAN and WAN RLS.

10. Business Continuity Cost (BCC)

"Small and medium sized enterprises are constantly grappling with how to cost effectively keep their applications available and their data protected on a continuous basis. Asempra's Business Continuity Server™ provides small and mid-market companies with the instantaneous application availability and real-time continuous data protection they need to keep their business up and running."

Sonia Lelii, Storage
Research Analyst
Aberdeen Group

Business Continuity Cost defines the cost efficiency of a business continuity solution. Data services such as backup, snapshot, replication, hierarchical data management, information lifecycle management, and archive are traditionally separate tools with very different architectures. This is simple because some of the tools are schedule-based with different interval requirements, while others are real-time, either synchronous or asynchronous. Typically, the tools that manage data history are not real-time, while the ones that do not manage data history are real-time. As tools that manage history move towards capturing data in real-time, they become more available and include other real-time services.

For better BCC, some of the existing CDP solutions already combine backup and replication, and others provide both backup and hierarchical data management. Without service consolidation, IT administrators have to manage several different tools manually. Consolidation of data services makes it easier for IT administrators to manage their data. In many cases, this is also a more cost-effective way to operate than through a loose collection of piece-meal tools.

From the above metrics, it is apparent that although all CDP solutions (block-level, file journaling, and Asempra's Business Continuity Server) have one common characteristic (the ability to capture data in real-time), their resemblance stops there. Most solutions have significantly different recovery characteristics. The recovery characteristic depends on the data type captured by the CDP product, and the way the protected data and metadata are stored. In short, CDP is not the answer to all recovery problems; it is the specific technology capability behind each of these, so-called, CDP products that makes the difference in an IT environment. Based on the needs

of an IT environment, one CDP product may work considerably better than another.

Given these metrics, it is possible to create a business continuity solution that scores high in all measurable parameters.

Applying the Business Continuity Metrics

Evaluation metrics are only meaningful if you can validate them with real world examples. Therefore, it makes sense to present a practical example comparing the business continuity solutions that exist in the market today. The business continuity market has offered multiple data protection and recovery tools for years, but recently there have been a handful of emerging products and companies in this space that are altering the landscape.

These products need to be included in the measurement of business continuity solutions in order to be comprehensive in our evaluation. And while several of these new products have taken a new storage approach which continually protects data at the block-level, one company has taken a much more innovative approach to solving the problem of business continuity.

“For business-critical applications that require rapid data recovery and zero loss of data ... Asempra’s Business Continuity Server delivers granular and real-time recovery, serving up the application and data immediately, and with guaranteed data usability.”

Lauren Whitehouse,
Analyst,
Enterprise Strategy
Group (ESG)

The Asempra Business Continuity Server

Asempra’s application-aware, real-time continuous data protection technology provides guaranteed application and data availability for Windows environments. Using Asempra’s patent-pending Virtual On-Demand Recovery™ technology, an application’s data is available for use within minutes, even seconds, of recovery. The data recovered is guaranteed to be completely usable on the first recovery, eliminating the need for multiple, costly manual recoveries that often result from other technologies and products.

With the Business Continuity Server, simple point and click global to granular recovery provides recovery flexibility that ranges from individual objects, such as a single e-mail or file, all the way to a complete data center. Asempra dramatically reduces the cost and complexity of mid-market IT data protection by consolidating the needs of backup, business continuity, disaster recovery, business continuity, compliance

and governance into just one solution. In order to protect data intelligently and eliminate data loss, the Business Continuity Server transparently and continuously protects application data as transaction events occur in real-time. Backup windows are eliminated and recovery point objectives (RPO) approach zero data loss.

Instantaneous Availability

Virtual On-Demand Recovery™ reduces recovery time objectives (RTO) to seconds. By making recovery times independent of the data set size and virtualizing the data set, applications can be restarted within seconds or minutes. Unlike other technologies that require a complete recovery of the data set first, the BCS allows you to recover objects instantly at a granular level – giving you the ability to recover what you want, when you want it. The Business Continuity Server recovers the data your application needs on demand, while it's running.

Guaranteed Data Useability

With traditional data protection products, data recovery is not only painful and time-consuming; it may deliver data that is corrupted and unusable. The Business Continuity Server's sophisticated event capturing and indexing technologies guarantee data is completely usable on the first recovery.

The Business Continuity Server avoids recovering inconsistent or corrupted data through the use of its patented comprehensive Real-Time Event Journaling™ and three-dimensional continuous object store indexing techniques (real-time data, metadata, and event).

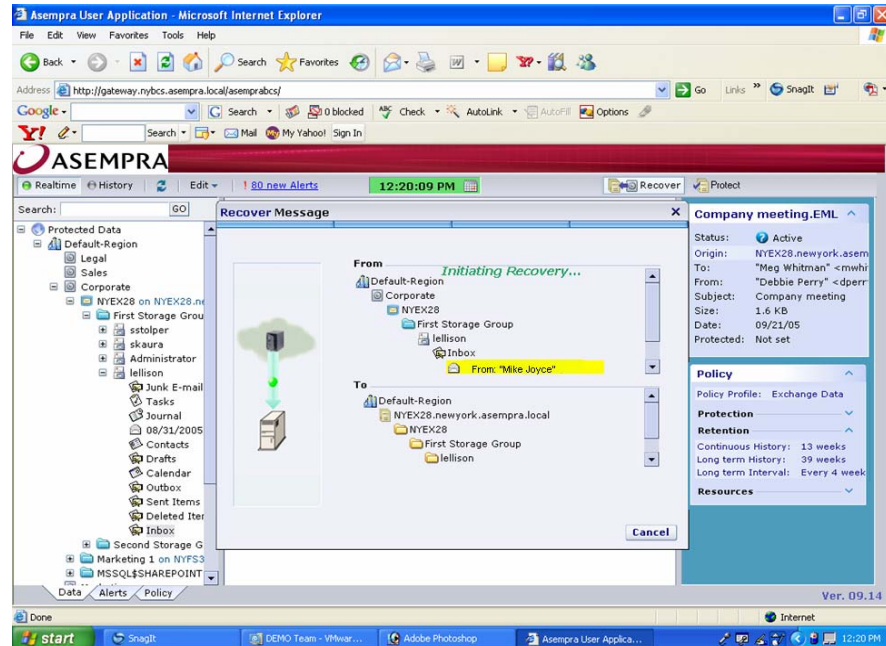
These elements are combined by the Business Continuity Server to preserve a true data history of the application with consistency marking to ensure application recoverability across multiple dimensions of time (second, minutes, hours, and days), object granularity, and application events.

Additionally, the Business Continuity Server's extensive searching and versioning capabilities can be used to locate and recover an individual object, such as a file, to help meet the needs of compliance and governance.

"Asempra's CDP is transaction-aware and application-specific. Its technology communicates directly with Exchange, SQL Server and Windows file systems... Before a transaction is copied, it checks the integrity of all of the data prior to forwarding it to the recovery server..."

Marc Staimer, President
and CDS
Dragon Slayer
Consulting

Recover Data in Seconds



Easily Recover Objects from a Single File or e-mail up to a Complete Data Center

Disaster Recovery

The BCS provides disaster recovery (DR) by protecting data over existing IP LAN and WAN networks. It applies patent pending delta reduction, temporary file reduction, and database stream prioritization techniques to minimize both network bandwidth and storage resource utilization.

The BCS also manages application server failover. For example, in the event of an Exchange failure, the affected servers are recovered to standby servers, and the BCS updates the Active Directory (AD) and Domain Name Service (DNS) entries so that users are routed to the new Exchange server. Users are back online in minutes or even seconds.

Backup Windows Eliminated

In addition to backing up production servers continuously, the BCS offers data export capability using the Common Internet File System (CIFS) protocol. This feature effectively eliminates the backup window on production servers, and improves production server performance by relieving it from having to perform hot or cold tape archiving.

True, Continuous Real-time Data Protection

The BCS enables instantaneous application availability and data recovery for Windows-based data in minutes or seconds, from any location, at any point-in-time with guaranteed data usability and reliability. The BCS provides real-time CDP, near CDP, disaster recovery, snapshot, business continuity, compliance and governance support in a single easy-to-use solution.

Business Continuity Scorecard

Because multiple tools are often unsuccessfully cobbled together to attempt to solve data protection issues, the management complexity and cost of data protection and recovery is enormous. This leaves IT professionals spending countless hours trying to integrate disparate tools and manually recover data in an effort to simulate a real-time infrastructure required to support their business. With such a myriad of protection and recovery tools to choose from, it makes sense to develop metrics that enable IT management to evaluate solutions that best fit their environment.

Summary

In most industries today, the service level agreements for data protection and recovery have moved to a point where there is no time for backup windows, no tolerance for data loss, and very little margin for recovery downtime. With the increased business demands for the disaster recovery of mission- or business-critical data and new compliance requirements, it is plainly apparent that legacy tools for data protection and recovery are ill-equipped to handle today's requirements.

Realizing that traditional data protection and recovery tools no longer suffice and that new technologies have emerged to address increased business expectations, the time has come to put in place the objective metrics needed to properly evaluate the array of business continuity solutions.

The ten metrics of business continuity enable IT management to apply thoughtful consideration to their own internal business requirements against the products they are evaluating. When looking at recovery time, recovery data, and recovery scalability characteristics, the Asempra Business Continuity Server

provides a compelling business continuity solution that will exceed service level expectations.

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

About Asempra

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, 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 while supporting today's data protection, disaster recovery, business continuity, and compliance and governance needs. Asempra is headquartered at 640 West California Avenue, Suite 110, Sunnyvale, California 94086. For more information, please call 408.215.5800 or visit www.asempratechnologies.com