



In Touch and Always On®



RealTime

Back-up and Recovery Solutions for Episys™

White Paper

November 2003



Strategic Technology Solutions Delivers Mainlines' RealTime™ Solutions

Episys Data Recovery and System Restoral in Seconds

Situation Analysis

It is no longer acceptable for Credit Unions to draft Business Continuity plans, then put them on a shelf to gather dust. Your organizations must reexamine their vulnerabilities and make verifiable plans to address very real risks—especially when it comes to your members' information. Here's what most Credit Unions face:

97% of data or application failures are due to

- system malfunction (44%),
- human error (32%),
- application error (14%), and
- virus (7%).

All too often, these failures result in data corruption. So what good is backup without recovery?

There are recovery tools available, but what is their real cost to the organization? Hours, even days of downtime?

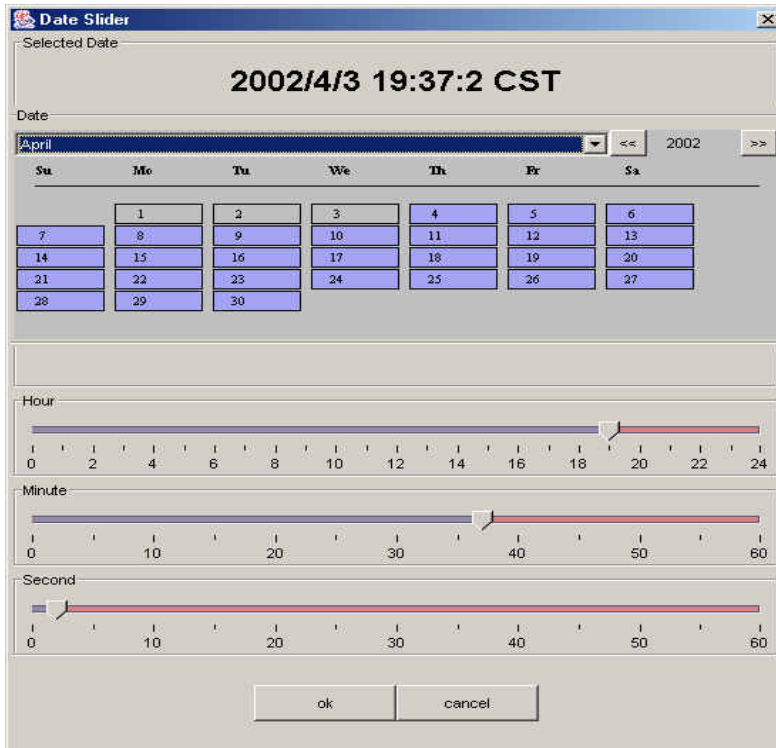
In a recently published article, "Ten Best Practices For Enterprise Backup/Recovery", Gartner™ emphasizes the issue of *recovery*. [Mark Nicolett and Raymond Paquet, ZDNet Tech Update, 5/2/02.] "Traditional solutions continue to focus on backup; data protection, speed of data movement, storage and management of the backup data sets. Recovery has largely been an afterthought. Depending on the type of failure, it has until now been difficult if not impossible to determine the amount of time and effort required for not just the recovery of the data itself, but the ability of the enterprise to resume normal operations."

Fortunately, it's no longer necessary to measure recovery time in hours or days. Recovery can and should be accomplished in minutes.

RealTime™, a revolutionary data restoration system, was engineered from its inception for immediate data recovery—just 20 minutes to recover 1 terabyte to 1 petabyte databases! Not only does **RealTime** deliver a continuous, non-intrusive backup of hot applications, it allows the orderly and almost instantaneous recovery of operations. The same Gartner article suggests that organizations should "Define recovery SLAs ... based on recovery speed, recovery point (data currency), and backup data retention." **RealTime** makes this a reality. By bringing new capture and recovery methodology to market, it is now possible to guarantee not only the backup, but predictable and timely operations recovery as well.

RealTime™ Overview

RealTime's design changes the backup restoration paradigm from a multiple-hour point-in-time backup to a continuous backup integrated with storage. The results are hot, non-intrusive backup without impact on production applications, and rapid (in minutes), verifiable restoration that is globally consistent across volumes.



(Figure 1: *RealTime* Time Slider)

How *RealTime* restoration differs

Today's standard backup and restoration products offer reliable, accurate backup, though they are complex and can take from three to seventeen hours to restore data from the most recent backup to repair the results of a system crash or data corruption. They may include mirroring or replication to protect against downtime in a system crash, but do not protect against data corruption.

How do *RealTime's* restore processes differ? In addition to offering reliable, accurate backup, *RealTime's* unique architecture captures data continuously and journals it in such a way as to allow immediate, point-in-time recovery when needed. In addition, *RealTime* maintains historical/chronological information with its innovative data capture process, which provides for data corruption protection. Although currently available competitive tools claim "rapid recovery," they require a complete restoration or reconstruction of the application data, in addition to applying transaction logs to the point of recovery. This takes hours. In addition, they are still subject to data loss and corruption.

The *RealTime* Edge: Pick a point in time through the Time Slider screen, and *RealTime* will restore data to that point in a matter of minutes.

With *RealTime*, users can restore and verify a 1-terabyte database in under 20 minutes.

By comparison, *RealTime* restores data by undoing its uniquely queued data writes in reverse order, literally backing out data until it has moved through the system crash point. This technology gives *RealTime* the ability to restore a 1-terabyte database—even a 1-petabyte database—in fewer than 20 minutes.

Key *RealTime* Features

In addition to *RealTime's* data capture and immediate recovery ability, a wide range of business coninuanance enhancements have been engineered into its design.

- **Continuous data capture.** RealTime captures data continuously, without impact to the application. In the process, it tags the data to allow for point-in-time recovery.
- **Cache and transport of data without network or system latency.** Using log files to hold captured data, RealTime is able to transport data across the network without impact.
- **Maintain forward and reverse incremental data.** Using patent-pending technology, RealTime is able to generate both forward and reverse block incremental images to facilitate a unique instant data reconstruction process.
- **Automate recovery and backup with a single interface.** The RealTime Administrator is a user-friendly GUI that allows for easy maintenance and optimization of RealTime's complete range of backup, recovery, and verification processes.
- **Data analysis through VirtualCopies.** RealTime's VirtualCopies feature lets the user make a writable copy of the data, without requiring significant additional storage space. Once the VirtualCopy is created, the user can perform data mining, financial analysis, testing, and disaster recovery. These tasks can be run without impact on the performance of the production host and without the cost of additional mirrored disk space.
- **Verify and rehearse backup procedures.** By creating a virtual copy on the backup system, users can “practice” a backup and restore situation on the virtual copy and verify the process easily and routinely, without any impact to the production system.
- **Protects continuously against data corruption.** Unlike replication and snapshot backup systems that will replicate data corruption into the copy, RealTime allows users to move incrementally through data instances from point of awareness to the point prior to the corruption instance.

RealTime Recovery

RealTime's patent-pending data capture and journaling process is at the heart of its approach to recovery. *RealTime's* automated recovery process does not require restoring full volumes of data or applying application logs to reconstruct the application data. Using this radically new approach an administrator can restore one terabyte of data or greater in less than twenty minutes.

The administrator simply selects a target point in time to recover using *RealTime's* Time Slider screen (Figure 1), and the *RealTime* Recovery Agent will restore data to that point in a matter of minutes. This automated recovery minimizes operator error and enables recovery by someone without specific knowledge of database or application configurations.

The Time Slider screen simplifies recovery by providing the administrator a single interface, and by automating restoration to any point-in-time recovery. This provides administrators the means to recover from system crashes, data corruption, and any type of data failure in minutes, rather than hours or days.

Unlike other tools that require the administrator to start over if analysis reveals that the selected target time was incorrect, the *RealTime* administrator can use the Time Slider screen to select a new target point-in-time to make adjustments in mere minutes. Recovery methodologies include full or partial restoration, and disaster failover:

- Full recovery can be accomplished using RealTime's automated process to restore a virtualized image on the backup server for recovery point validation. The recovery sequence is then applied to the production server to synchronize the production data within minutes.
- Partial recovery follows the same automated procedure using a virtual image to establish the restore point on the backup server. Affected tables and records are extracted and then inserted into the production application. This allows the production database to remain operational during recovery.
- System failover is accomplished by launching the application on the backup server using the mirrored copy of the application data. Additionally, RealTime can correct the replicated copy to remove any corruption that occurred during the primary failure, a capability no other product provides.

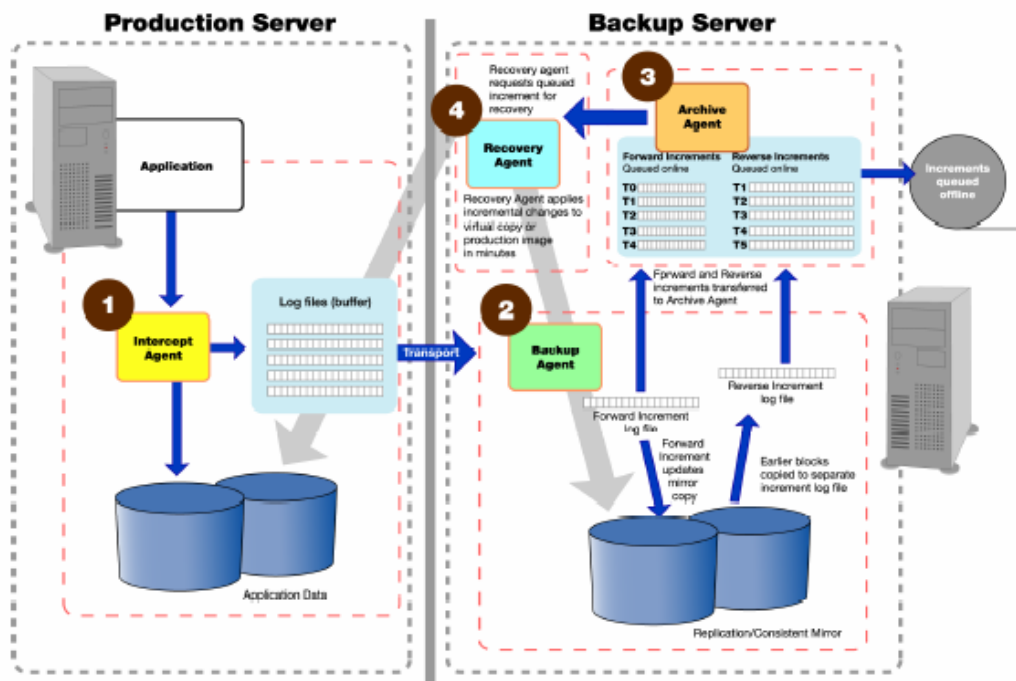
RealTime Architecture

RealTime runs across two systems: a production/application server, and a backup server. It comprises four principal components: Intercept Agent, Backup Agent, Archive Agent, and Recover Agent. These agents govern *RealTime*'s data capture, journaling, transport, archiving and restoration. The Intercept Agent, Backup Agent and Archive Agent perform all data preparation and structuring to facilitate immediate recovery.

- The Intercept Agent provides continuous backup, capturing incremental block level data changes synchronously in transit to storage, packaging them into log files for asynchronous transport to the backup server.
- The Backup Agent manages RealTime's unique data journaling process, maintaining a full replication of the storage devices associated with a protected application. Before applying changes to the replicated copy of data, the Backup Agent first makes a copy of the former data into a separate log file on the backup server.
- The Archive Agent – with its revolutionary recovery method - records log files as forward and reverse increments. A journal of forward and reverse increments is maintained on line and queued for immediate recovery.
- The Recovery Agent determines the required forward and reverse increments to be used in recovery. Then, efficiently, it applies the reverse increments to a virtual image of the application data or directly to production storage to complete the recovery in minutes.

The RealTime Data Protection Recovery Process

1. The Intercept Agent simultaneously writes data and meta data to a sequential log file while it updates the primary storage. The intercept agent ensures that all transactions and their sequences are catalogued. This is crucial to making sure that all data structures can be restored to a consistent state across application storage devices during recovery. Log files are configurable in size and number, and act as a fault-tolerant buffer during communication disruption.
2. As the log file buffers are filled, or a scheduled time period elapses, the catalogued information is transferred to the Backup Agent, usually on a separate backup server. The backup agent uses the log file to update and maintain a consistent mirror. Before overwriting the mirrored data with the changed blocks the Backup Agent makes a copy of the earlier block into a separate log file, creating a reverse incremental.
3. The log file with the forward and reverse increments of before and after block differences are handed off as paired data sets to the Archive Agent. A journal of these log files is maintained on line and queued for immediate recovery. When integrated with a media management tool like Veritas Netbackup or Tivoli TSM, the Archive Agent migrates the forward and reverse increments to the backup system.
4. The Recovery Agent queries the Archive Agent for the required reverse increments. The Recovery Agent applies the reverse incremental data to a virtual image of the application data or directly to production storage. The recovery agent sequentially applies the individual block updates across all volumes to bring data to a consistent state at the precise time selected.



(Figure 2: RealTime Process)

Key Technical Differentiators

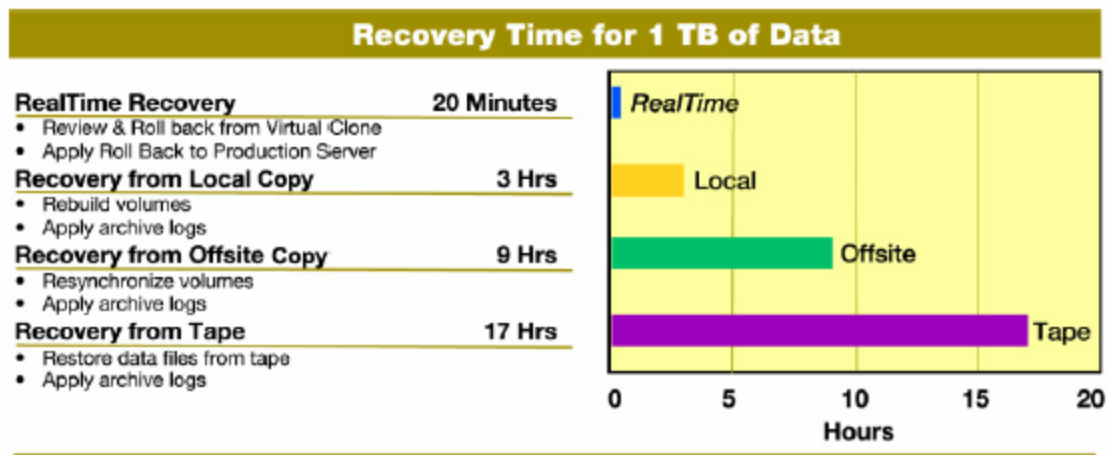
Backup products traditionally capture data at the file layer, making it difficult to support openfiles or active databases. Replication products normally capture at the physical layer and are unable to maintain consistency across unlinked devices such as database tables on different volumes. Backup, snapshot and replication products require integration with applications using agents to shut down or queue the application to assure data is consistent on storage before generating a copy—and require additional application journaling processes to reconstruct data between cycles.

RealTime uses an agent placed in the I/O stream at a point in the file stack that is:

- below the file systems and/or raw volumes, and
- above the logical and/or physical volumes.

RealTime captures data as it is transported to the logical or physical storage device, the only point in the I/O path that can ensure consistent write across all storage devices. This capability is critical because capturing write activity at this layer allows *RealTime* to be application and storage independent, allowing it to support any application or storage device. Capturing data here facilitates management of open files or applications without impact or special application agents.

RealTime's historical journals of forward and reverse increments facilitate recovery to ANY point-in-time, independent of storage or application. The management and recovery using incremental data eliminates the need to recover entire volumes and reduces significantly network bandwidth requirements for data protection and recovery.



(Figure 3: Recovery Time Comparison Chart)

Summary

RealTime has significant technical, operational, and cost advantages over current replication and backup products. Its unique design enables the single most important feature: its ability to restore vast amounts of data accurately and verifiably with unprecedented speed.

Traditional backup products provide full recovery of aged data, but there is a constant risk of losing data between scheduled backups. Using *RealTime* eliminates that risk. Backup tools as a group also have a significant impact on server operations, forcing backup processes to be scheduled during non-business periods. Large databases must have numerous partial backups performed over several days, in order to achieve full backups, as a result of the time impact.

With its ability to “time slide” through the historical record of duplicated data rapidly and accurately, *RealTime* offers cutting edge data restoration and replication tools that separate it from its competitors.

About Strategic Technology Solutions

Strategic Technology Solutions (STS) is a premiere provider of contact/call center, telephony and communications disaster recovery solutions. STS is partnered with Mainline Information Systems to deliver RealTime Host Processing Back-up and Recovery solutions to Credit Unions nationwide. Please contact a STS Representative to find out more about the Mainline | *RealTime* solutions which can benefit Credit Unions member data. For additional information go to www.stsolutions.net or call us at 1.866.STS.4CUs.