

The **Future** of Disaster Recovery Solutions **is Here**

By Leonid Shtilman



MOST businesses secure their information infrastructure by regularly backing it up onto tape. Some have gone further, enhancing their backup strategy with expensive disk arrays and mirroring. Whether a natural disaster, a blackout, or a hard disk failure occurs, these backups would ensure the survival of information. Should human or software error corrupt a business's data (which accounts for approximately 40% of all application-related disasters), they would simply reach for a recent backup, which would help them get back on their feet.

In most cases, an organization using this approach will be able to eventually restore corrupted data from tape or disk arrays, but recovery is rarely quick enough to mitigate lost revenue and productivity. Data integrity may be compromised, and information added or modified since the last backup will probably be lost.

In other words, most organizations are satisfied with receiving information in real-time but safeguarding it much less frequently. The argument is simple: traditional backup is not enough. It may cost you downtime affecting productivity, but more importantly, it may seriously affect your company's bottom line.

RAID, TAPE BACKUP AND SNAPSHOTS

There are a number of traditional solutions available to protect data including RAID, tape backup and snapshots. If data becomes corrupt, RAID merely ensures that the resulting garbage is safe from disk failure.

Both tape backup and snapshots are "all or nothing" solutions, unable to prevent the loss of hours of valid, valuable data creation prior to an incident. Even a good replication system requires resynchronization from a backup server and is likely to replicate corruption to the replica server as well. All of the aforementioned measures have their place in an effective overall application protection strategy but none are enough on their own.

To elaborate on the most common of these methods, snapshots maintain frequent watch over your data, a complement to conventional periodic backups (typically, once a day). Snapshot-based backup solutions will generally give you added security, for the simple reason that they allow you to store the state of your data more often (typically, once every few hours). Should corruption occur, you will no longer have to spend hours restoring all of your data from yesterday's tape backup. Instead, you will be able to relatively quickly restore data from a previously stored system snapshot.

On the downside, disk activity typically needs to be suspended while snapshots are taken. Also, snapshots require tight integration with the specific application server being secured, and data integrity can only be assured if either the application itself or a system administrator invokes them in proper timing. If you're out of luck, and the most recently recorded snapshot is out of sync with your application's latest consistent state, restoring data from the snapshot may have unpredictable results.

This is better than regular tape backup, but certainly not perfect.

CONTINUOUS BACKUP GETS THE JOB DONE

Users interested in the highest data integrity and recovery speeds are probably best served by boosting their backup strategy with the latest continuous backup solutions. These may be added to your existing backup infrastructure and will monitor your application servers while capturing and recording all operations applied in a journal at all times.

Should data corruption occur, affected servers may simply be "rewound" by playing back an opposite operation or "counter-event" previously logged in the journal. Not only does this allow you to back up vast amounts of data accumulated over a long period of time but it also means that recovery will be instantaneous.

Snapshots are typically made once every few hours, and may additionally be taken to create offline backups, usually on a daily basis. If a disaster occurs, and only tape backups and/or snapshots are available, you stand to lose as many as 24 hours of updates, not including the amount of time it will require to restore operations. At the very least, you will lose all updates made during the hours that have gone by since the last snapshot or backup was made. Also note that rapid restoration from snapshots will only be possible if you restore entire volumes, and things will significantly slow down if you attempt to restore individual files, directories or databases.

Ultimately, continuous backup may be added and configured to monitor every update made to your servers, either all the time or, if you prefer, in between the two most recent snapshots. Should your data be corrupted, you will be able to choose between virtually unlimited restore points. Continuous backup is the only solution that will allow you to restore single or multiple databases to the most recent consistent state, sometimes logged just minutes before corruption occurred, so that the highest data integrity and negligible data loss, if any, are ensured.

HOW TO SELECT A CONTINUOUS BACKUP SOLUTION

When selecting a continuous backup solution as part of an overall application protection strategy, it is important to remember that a technology is not in itself a solution to a business problem, it is only a component of the solution. The failure to recognize this distinction has led to enormous losses of resources invested in IT over the last decade or so. Companies have purchased one technology after another without understanding how the technology would actually provide a solution to the problems that they faced.

In addition, it is critical to concentrate on the principle of real-world focus. Real-world focus means that any solution must be well matched to the problems of your company's IT environment, not to those that are easiest to address. Without this focus, products are likely to create as many complications and problems as they solve.

The solution should be able to seamlessly integrate with the environment including security, bandwidth management, and alarm-reporting mechanisms. There should be an intuitive interface with powerful application-specific aids. The issue of installation is very important, an advanced solution will not require you to reconfigure the servers or reboot the servers during installation.

The product should be application-aware, it should be able to recognize and record points at which databases are in a consistent state. It should

also have auto-discovery of database files to relieve administrators of the need for detailed knowledge of database storage architecture.

The continuous backup solution needs to be flexible. There should be centralized management of all scenarios enabling you to perform replication on specific data. If you're operating in a cluster environment, your solution must support clusters.

Look for a solution with a clear return on investment, typically measurable to the cost of a single recovery incident. There are several factors that can keep the cost down including a simple setup which reduces administrative costs and no unnecessary third party software investment or additional bandwidth requirements. And of course the initial installation should not require any downtime.

Perhaps most importantly, find a product that you can test before you buy it. Many continuous application availability products claim to solve your problems. Verify their claims by downloading a trial version to see for yourself.

THE BEST OF BOTH WORLDS

You might also want to consider enhancing your backup strategy with real-time replication technology, which will add a high-availability layer to your data infrastructure, along with optional automatic failover. Generating a real-time copy of your data that is continuously kept up-to-date will achieve the most important benefit, 24/7 availability for your users. They will appreciate being able to work uninterrupted and will be even happier when, should disaster strike, you'll have them up and running again in a matter of seconds. 🌐

Desktop Environments

Leonid Shtilman holds a PhD in Mathematics from the Israeli Technion, and a PhD in Engineering from Tel-Aviv University. He has over 15 years of experience in technology and management. In 1999, Shtilman founded XOsoft, a leading provider of continuous application availability software. His work in developing and bringing to market XOsoft's WANSync business continuity software suite and its continuous backup product, Data Rewinder, has helped the company grow a large global customer base.

Prior to founding XOsoft, he served as CTO of Identify Software, a developer of unique software development and performance monitoring tools, which he co-founded in 1995. Shtilman was also one of the founders of C3D Inc. (NASDAQ: CDDD), a developer of new optical media for three-dimensional data storage. His research has gained substantial support by the DOD and DOE. Previously, he also served in various positions at NASA, MIT and a number of universities.