

# So-Called ‘Small Disasters’ Can Equal Big Trouble

By **RON LEVINE**

Forget about storms, earthquakes, and fiery explosions that are the images that first come to mind when we hear the word “disaster.” The chances of one of these catastrophic events actually hitting and destroying your individual IT operation are infinitesimal.

What should be worrisome are the far more common minor mishaps that occur every day and yet are capable of bringing businesses to their knees. Everyday occurrences like coffee spills, power surges, communication outages, or an employee mistakenly deleting or overwriting a key file trigger most IT disasters.

Why? Because it’s these seemingly small mundane events that attack a business’ most valuable asset – its stored data.

Headline-generating disasters that make the 6 o’clock news can wipe out complete information systems and network environments, but they seldom occur at any individual company and rarely strike the same site twice. On the other hand, recurring small disasters (usually never heard about outside the company’s IT department) continually cause havoc and financial losses to unprepared IT dependent businesses. These more common types of IT disasters usually interrupt processing operations because they block access to, cause corruption of, or destroy critical data.

## Data Disaster Contingency Planning

Since organizations today depend so heavily on access to and availability of stored information, it is imperative this information be fully protected and easily recoverable, as anything that affects these critical files can quickly change a small incident into a major company-wide disaster. A well prepared IT administration (one with a data disaster contingency plan in place and tested) can fully protect itself from damages caused by most types of data interruption or destroying mishaps. Data replication techniques provide the cornerstone of this protection.

High-risk incidences such as the introduction of a virus destroying or corrupting files, a crashed disk on a server, a malfunctioning or blocked SAN or NAS, a storage unit upgrade which goes amuck, a communications glitch causing lost customer transaction data, or an employee’s careless mistake – all can be guarded against with a solid data replication solution.

The typical everyday IT emergencies don’t make very suspenseful movie plots, but these seemingly mundane incidences have the potential of placing an enterprise into Chapter 11. A data replication component is an essential subset of the company’s overall disaster contingency plan. And if history is a guide, chances are it will be this component of the overall plan that is most likely triggered in response to an event, thus averting a potential disaster.

Basically, data replication is the process of creating a duplicate copy of the data at a secondary location. The copied data then can be made available to an application in the event of a loss of the primary data source. The choices include local data replication, for example RAID 1 mirroring or distance-based data replication, a replication technology that can move the data over distance to a remote site. While both provide instant online availability of the copied data, only distance-based data replication guards against a catastrophic loss of the entire data center.

After deciding the local vs. distance-based data replication issue, planning should focus on the level of data replication required by assessing data protection and recovery needs; each business (and even individual site needs within a business) has unique requirements. After completing the needs assessment, the proper options for meeting those needs can be selected and installed. The final step is testing the data recovery solution with a live run-through.

## Data Recovery Needs

Each business is different, but almost all rely on stored files processing. One chief technology officer at a storage manage-

ment company begins the needs assessment phase by answering the following questions in order to help determine the level of data disaster recovery necessary at client sites:

- Does the computer system (and its stored data) directly generate income for the company?
- Do day-to-day business functions rely on access to critical data?
- Is there an impact on customer service if the data cannot be accessed and timely information and answers provided?
- What would be the financial impact on the company if data access is interrupted for a day, a week, a month?
- Is the company involved in emergency services that rely on data availability?
- Is the computer storing quality assurance data? What quantity of the manufactured products would be rendered useless if that data were unavailable?
- Is there critical research being done which would be compromised or destroyed if the storage sub-system malfunctioned?

The responses to these questions provide answers to the type and speed of data recovery necessary to protect the company.

## Choosing The Right Option

With answered questions in hand, the executive begins working with the client to develop what he calls an “application continuity curve” (ACC) to select the right level of data replication protection and recovery needed. The client plots the ACC, which illustrates the relevant importance to the company of each type of stored data/files.

For example, the customer may start out with a normal backup and recovery solution for low priority data that would not be part of the disaster recovery effort (such as MS Word’s My Documents files, appointment schedules, e-mail messages). Then, move onto mid-level data (e.g., marketing information, proposals, profit projections) that would be recovered more quickly as part of the disaster recovery capability. And finally select top priority data the company can’t operate without (customer lists, order processing database, human resource files, tax data) for continual access and high-availability solutions.

The curve assigns relative weight to data that is to be protected/accessed in the event of a disaster. This pictorial representation clearly shows the direct correlation between the value of the data, and the investment that can be justified to protect that data. For example, the My Documents directory on a laptop will not justify the same level of investment as an order processing application. As the value of the data increases and its recovery time decreases,

the cost of recovery increases. With the curve completed, the client can decide how long they can remain “down” within each scenario and select the recovery options that are most effective and cost justified.

### Data Replication

Data replication allows companies to have an exact duplicate of their data on a secondary device, usually at a remote location. The replication technique can be either hardware- or software-based and the data can be replicated at the volume, file, block, or byte level. Which one is right for you is dependent on your specific IT environment. Some factors to consider are:

- Installed hardware;
- Available WAN bandwidth;
- Data type (files vs. database);
- Budget;
- SLA (Service Level Agreement) requirements.

Each approach brings different benefits and consequences. An independent storage services company can work with you to provide a balanced review of the options and help you choose which one is best for your organization.

Once the decision to replicate a particular data element has been made, make sure that the replication data path can support the amount of traffic that will be generated moving data from the primary (source) location to the secondary (target) location. The importance of this step cannot be overstated. For example, database applications will generate more traffic than file

serving applications and, therefore, will need more bandwidth availability.

Overhead associated with the replication process itself is another extremely important consideration. Check with your supplier to make sure that there is sufficient processing power and memory in the storage device and/or server to replicate the data without affecting overall subsystem performance.

Finally, decide whether to operate the replicated site in a hot, warm, or cold state. A hot state is used for the most critical of applications – those requiring contiguous operation and zero downtime. With this option, the data is replicated and the application generating the replicated data is in an online and active state. This option is the most costly, but also most invisible to customers.

A warm state option is employed for applications that are important but not considered critical. A copy of the application is loaded but not running in an executable fashion. This level of recovery works for applications that have a few minutes to a couple of hours of downtime leeway. It is less costly to implement than the hot state recovery technique.

A cold state data replication function is preferred by businesses that can sustain a defined extended period of downtime. It is the least costly option as the replication site is operated as a data repository only, without any applications loaded at the target location.

When it comes to safeguarding your data and implementing a data replication

technique, there is no “one size fits all” solution. That is why it is prudent for most businesses to employ the services of an experienced storage consultant to address the data protection and recovery aspects of the company’s overall business contingency plan.

### Completing The Data Disaster Recovery Plan

After installing the physical data replication and recovery solution, spend some time on the easily overlooked plan particulars. Write down the names and numbers of every person responsible for each aspect of data administration in the company and ensure each knows what they are to do in the event of a data disaster. Make it specific as who will notify who, who has keys to the buildings, the computer rooms, and data storage vaults, who the back-up will be if some key person is unavailable, etc. Make sure names and phone numbers are reviewed and updated regularly.

A major part of an effective data disaster recovery plan, and one most likely to be overlooked, is the testing or rehearsal. Plan for a disaster well in advance, and include the recovery solution vendor(s), every employee, and all who might be involved in a true data disaster situation (don’t forget the communications vendor).

Arrange a mock run, skipping no steps from original data disaster discovery through recovery to online business as usual. If using an automated recovery solution, ensure that it all functions as expected. Have everyone report to any off-site facilities that they would go to if the emergency were real, and ask that some pose as customers to make sure that the backup data all works correctly.

Review the outcome of the test to determine what went well and what needs improvement. Once changes are made, hold another rehearsal. During one test, pose the scenario in which the disaster occurs during a normal workday. During another, have the disaster occur off-hours or during a holiday when many of the utilities and other contacts are closed.

The idea is to be prepared!



Ron Levine is a computer information systems instructor at Santa Barbara City College and frequently writes for *Disaster Recovery Journal*.

■ To comment on this article, go to 1504-16 at [www.drj.com/feedback](http://www.drj.com/feedback).



**Disaster Recovery 2000**

**Fifteen years serving the business community with Disaster Recovery solutions.**  
**Priced for your budget! No more excuses! Full Business Resumption!**

<b>Emergency Response Plan</b> <b>IS Recovery Plan</b> <b>Administration Recovery Plan</b> <b>Easy to Install</b> <b>Easy to Use</b> <b>Easy to Maintain</b>	<b>Menu Driven</b> <b>Relational Database</b> <b>Data Gathering Kit</b> <b>Installation Guide</b> <b>Recovery Team Designations</b> <b>Team Member Job Descriptions</b>
---	--

**Much More!**  
**DEMO DISK AVAILABLE UPON REQUEST**

**Disaster Recovery 2000**  
**155 Tara Hills**  
**Whittier, North Carolina 28789**  
**(828) 497-7273**  
**Disasterrecovery@earthlink.net**