

IT Sense



JON WILLIAM TOIGO

Season of Disasters

There's a tendency to think of disasters in seasonal terms. In Florida, where I'm based, June 1 is the date assigned as the beginning of hurricane season, despite the fact the really disastrous weather events seem to be coming later and later in the year—from September all the way through December. Elsewhere in the country, we hear worries of another active winter storm season, though this is also an imprecise bracketing of ice storms and blizzard events, which last year clung like ice to the roof gutters in much of the country until well into spring. And, while less weather-related than geophysical, we're now hearing that the West Coast is "overdue" for the Big One—an earthquake that might someday give Nevada its own ocean-front vistas—based on historical intervals between tectonic events.

Truth be told, most disasters—which I define in the context of business technology as "any unplanned interruption of normal access to data for whatever constitutes an unacceptable period of time"—don't fall into the category of "smoke and rubble" events with broad geographical impact. Statistically, these disasters comprise less than 5 percent of the total interruption pie chart.

In our world, disruption events tend to emanate from four primary sources: planned downtime, software glitches, hardware failures, and user errors. Trend lines are upward for the last three and downward for the first. These days, a lot of maintenance activity is being postponed (hence the downward trend in planned downtime) for lack of time and resources, which actually contributes to the upward trend in hardware- and software-related outages. Some call these types of interruption events "crises" that can, in theory, be resolved more rapidly than cataclysmic outages stemming from earthquakes, floods, hurricanes, and the like.

In truth, however, the recipe for recovering from any outage is the same. You need to re-platform the application, re-connect the application to its data, and re-connect users to the application.

Mainframers know this drill well, and their data center operations procedures manuals usually feature one of several methods for accomplishing these tasks in response to any interruption event. For those firms with deep pockets, there's usually a second data center, either owned by the company or leased from a commercial hot site provider, waiting to pick up the workload from the

primary data center should it become impaired or isolated. The less well-heeled tend to take a more winding road to recovery: backup to tape, store tapes offsite and, in the event of an emergency, rely on a combination of service bureau processing, just in time discovery of alternative space, and next-box-off-the-line delivery of new or off-lease gear.

Virtual Tape Libraries (VTLs) are being increasingly used; in crude terms, they're simply disk buffers that front-end tape libraries. While some vendors offer VTLs as special storage hardware appliances that are capable of prestaging tape volume for later writes, they also provide a means to replicate their data to a remote site. The gotcha is that usually the gear receiving the write remotely must be the same brand as the VTL appliance used in the primary site—doubling the cost of the solution and locking the consumer into a particular vendor's stuff.

A better alternative is to go with a software-based VTL. The strengths of a software-based approach are many. One benefit is that you can use your own DASD rather than paying a huge markup on otherwise commodity disk to store your tape volumes. Moreover, if you're replicating some or all of your data over a network, this approach frees you from a vendor lock-in on hardware, increasing your flexibility.

If you go the software route, be sure to select a product that 1) facilitates network-based communication without a lot of hassle; 2) gives you the flexibility to alter the size of the cache of disk; and 3) leverages all LPAR and CPU enhancements, including dynamic allocation of DASD (so you can grow and shrink your storage pool as needed) and System z Integrated Information Processor (zIIP) eligibility for the VTL application and its workload (so you can offload processing).

Also, make sure your virtual tape supports, well, tape. In my experience, disasters know no season, and for all the talk about disk-to-disk replication replacing tape, I keep hearing that Sony ad replay in my mind: There are only two kinds of disk—those that have failed and those that are about to. **Z**

JON WILLIAM TOIGO is a 25-year veteran of IT and the author of 13 books. He also is CEO and managing principal partner of Toigo Partners International, an analysis and consulting firm serving the technology consumer in seven countries. Email: jtoigo@toigopartners.com; Website: www.it-sense.org