

CIO News:

Avoid these architecture mistakes in your disaster recovery planning

By Linda Tucci, Senior News Writer
02 Apr 2009 | SearchCIO.com

The complexity of IT systems today accounts for 80% of systems downtime, yet companies design their disaster recovery systems in the same way -- complexly -- representing one of the biggest mistakes in disaster recovery planning.

So say different experts schooled in the disaster recovery (DR) discipline, suggesting that the secret to effective, long-term disaster recovery planning is hidden in the biggest mistakes they see companies make.

As companies bolster their DR capabilities to reduce ordinary downtime costs and stay online 24/7 -- and not to just mitigate risk -- they are employing a greater breadth of technologies, according to Forrester Research Inc., from server virtualization to host-based replication and optimized wide area networks (between remote, colocated or outsourced sites).

The problem can be getting all those pieces to work together when you need them.

Mistake No. 1: Lack of flexibility.

The default approach to disaster recovery planning is often to conjure a list of common disasters and figure out how to bring IT systems back up if they occur, said Paul Clifford, founder of Davenport Group, a storage and data recovery consulting firm based in St. Paul, Minn. But that leaves out the most critical element of a good DR plan -- flexibility.

"Probably the biggest mistake we see CIOs make is not building a plan that can adapt," Clifford said.

His remediation? "Simplify IT architecture and centralize control, so you truly have a data-centric approach. If you simplify and centralize, then you get the important element, which is flexibility," Clifford said. (He's the expert who said 80% of system downtime relates to complexity.)

Virtualization is one way to build in flexibility; so are dual-mirrored storage area networks that give you options at multiple locations.

"Once you have the data in hand, and the full-blown images for the other site, you can do whatever you need to do, whether that is bringing it up in a colo [colocation center] or sending data back to the main office, or establishing another primary location," Clifford said.

Mistake No. 2: Underestimating the time it takes to fully implement a disaster recovery plan. That is especially true for companies with remote locations, Clifford said.

"I have yet to call on the IT team that is overstaffed and overfunded," he said. "It takes a very long time to deploy these solutions in remote offices. Nothing happens in a 30-day window."

In fact, in many cases full implementation never actually occurs.

A DR plan involving multiple locations is especially demanding. CIOs need to understand the nuances of each location, from the operational and technical resources to the personnel.

"In some locations we see folks doing backup tapes on servers and relying on admin staff to swap out those tapes and get those tapes off-site. That is just fraught with peril," Clifford said.

Mistake No. 3: Staying with a provider that isn't meeting your needs or isn't as cost effective as it once was.

Stephanie Balaouras, an analyst at Cambridge, Mass.-based Forrester, notes that many companies are bringing disaster recovery back in-house for several reasons. These range from "frustrations with DR service providers" to more aggressive recovery objectives, as well as technical advances that allow CIOs to bring the capabilities in-house.

Technology is both more affordable than it used to be -- host-based replication is cheaper and more "bandwidth-efficient" than storage-based replication "and it will still help you achieve recovery objectives you can measure in hours," Balaouras said -- and a company can use it for more than just DR. IT can use a recovery site for application development and testing, for example.

"That said, there are good reasons to stay with a DR provider," Balaouras advised. Provider costs are also decreasing, with the advent of commoditized solutions and a la carte offerings.