acceligence

Why is disaster recovery always such a disaster?

Insight

August 25, 2026

Category: technology

Tags: business continuity, disaster, resilience, tabletop

The nature of disaster recovery planning has always been challenging, especially due to the different and unpredictable nature of such disasters. Hurricanes, earthquakes, and tornadoes pose very different issues than a terrorist attack, ransomware, or a raging wildfire.

It's also exceedingly difficult to anticipate logistical issues. For example, when Hurricane Sandy hit New Jersey back in 2012, many businesses thought they had covered their energy needs by keeping portable power generators in the back room. Unfortunately, many of those portable generators were gasoline-powered, and the nature of the damages forced the state to ration gasoline, and gas containers for generators were banned. Few businesses projected a power outage coupled with a generator gasoline ban.

Many also failed to project the lengthy duration of the power outages, having prepared for the typical outages of several hours to maybe one day. Sandy's outages lasted more than a week.

But the biggest threat to enterprise DR plan effectiveness today is that the nature of the IT environment has dramatically shifted, with a sharply spiked percentage of controls now existing well beyond the perimeters of the enterprise environment.

We're not mostly talking about the obvious changes, including SaaS, cloud, and GenAI/Agentic – although those are definitely factors.

The biggest change, and the latest AI versions are a key part of this, is that so many functions are now controlled by third parties.

Frank Trovato, a principal advisory director at Info-Tech Research Group, said environmental changes are a big part of the reason most disaster recovery arrangements fail in the real world.

“The exponential growth of SaaS has changed how organizations need to address DR and overall resilience,” Trovato said. “They have no control over SaaS outage recovery. They can’t failover to their own warm standby and that is a huge vulnerability. For example, if M365 has an outage, that impacts an organization’s primary means of internal and external communications, not to mention the project work maintained in MS Teams channels, SharePoint, and OneDrive. And the organization is often just waiting for the vendor to resolve the outage, with no control over the situation.”

A recent story in CIO.com talked about a potential slew of outages coming in November because of a seemingly innocuous update to DNSSEC. The cause: third-party dependencies that are beyond the radar of most enterprise CIOs.

One senior site reliability engineer at credit card giant Visa described the exposure in frighteningly specific terms.

“The risky areas are usually not the obvious managed DNS services. They are the older internal applications, hardcoded resolvers, containerized workloads, sidecar configurations, custom scripts, partner integrations, VM images, stale base images, and service-to-service dependencies that nobody has touched in a long time,” the Visa official said. “These systems can keep working quietly for years, then fail during a DNS or certificate-related change because they bypassed the normal platform standards.”

Think about that during your next DR tabletop preparations.

This means that you need to factor every reasonable partner into your exercises. But before you can do that, you must incorporate them into your thinking. Have your teams evaluate every business unit and list every function they perform.

Then make as comprehensive a list as possible of every third-party involved and list every function they perform or support.

Have some IoT functions, such as assembly line machinery that automatically downloads updates? What happens if this hypothetical disaster hits them? What happens if the mothership stops responding to machine queries? Does it proceed without the patch, or does it shut down? Don’t rely on what the vendor tells you. Test it yourself.

And don’t forget about communications. During a cyberattack, assume that all network communications – including phone calls if you are relying on VoIP – are being monitored by your attackers. By the way, buying burner phones for key personnel is a great first step. But, in an emergency, do your people even know where those burner phones are? Is there a procedure to periodically check those burner phones to make sure they still connect and have power?

Here’s the secret: Doing this to more effectively set up meaningful DR tabletops will get management support. But

your real objective is to map as many dependencies as you can. That information will address a massive number of IT and cybersecurity issues, not to mention compliance.

This procedure should also be applied to every single company that the M&A team is considering for acquisition. Without it, they are not truly understanding the risks they are about to purchase.

The DR exercise is an excuse to do what you should have done two years ago, but no one would approve of the massive amount of time it would require.

Looked at that way, your DR efforts could be a gift in disguise.