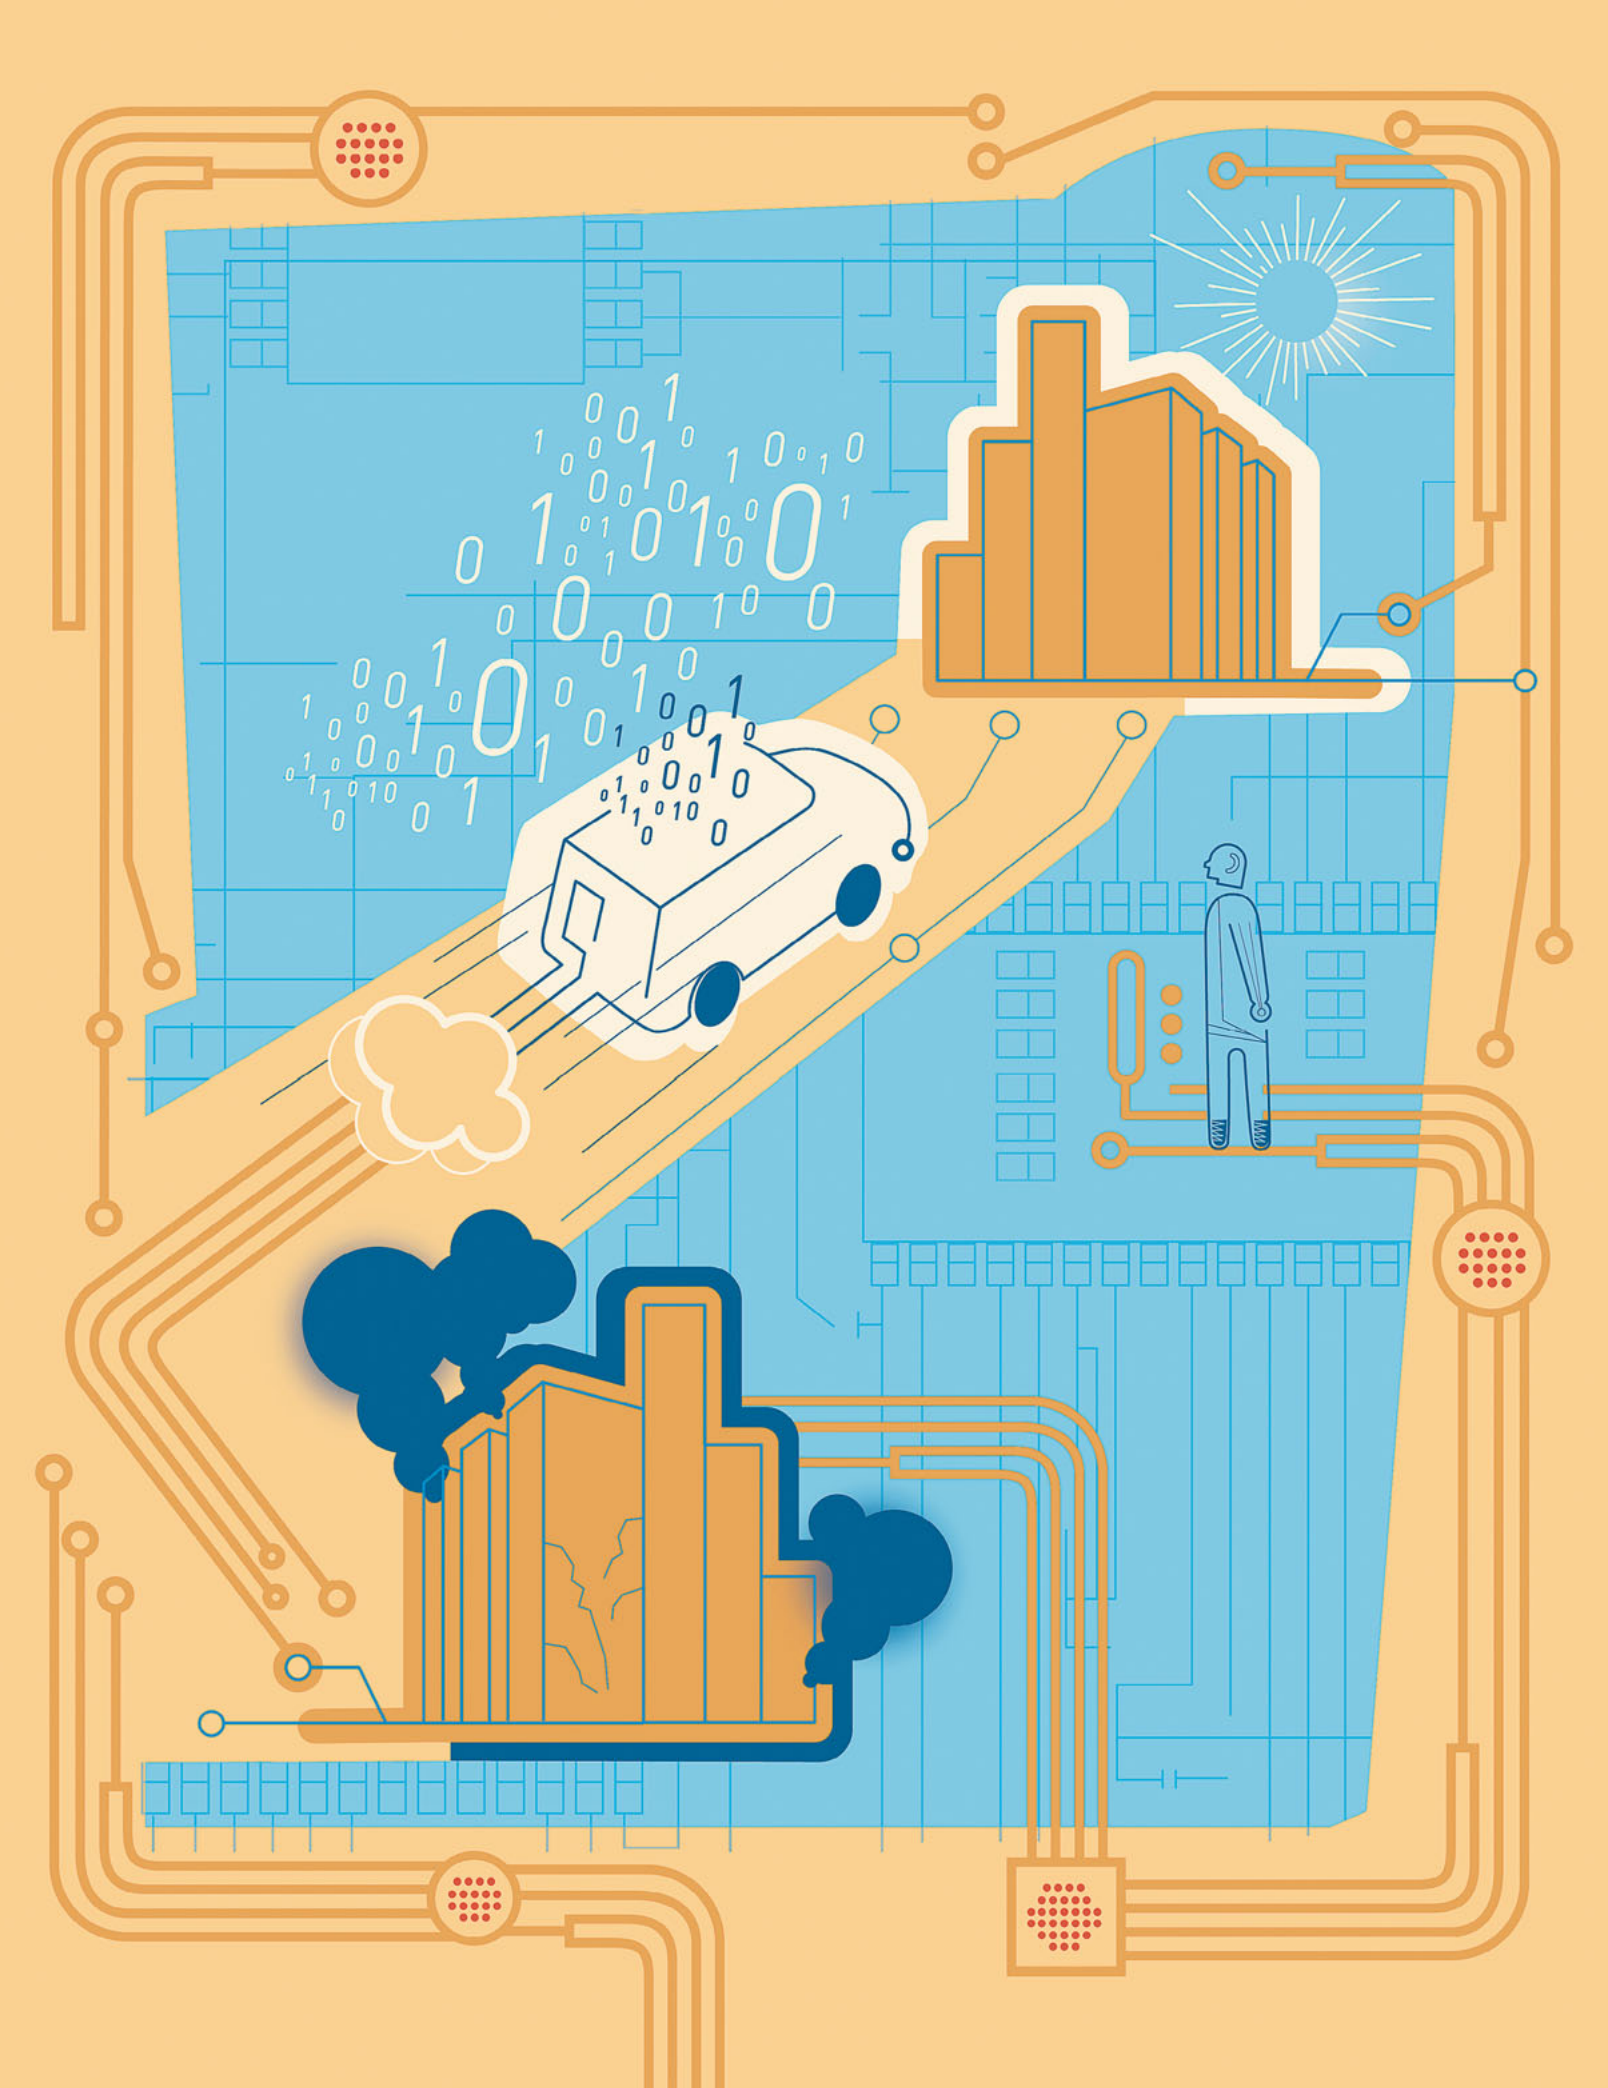

Pedal to the metal

BY SANDRA GITTLEN

Illustration by Celia Johnson

WAN acceleration lies at the heart of optimizing data replication and business continuity.

Businesses are under pressure to consolidate their assets and centralize management of applications and data. While they have seen great fiscal returns from pooling their resources, they have also seen an increased exposure in a critical area: disaster recovery.

"Due to heightened risk, fiduciary responsibility, increased competition, and regulation, upgrading disaster recovery capabilities is a top priority for enterprises," says Kevin Hohenbrink, product manager at F5 Networks.

A 2007 report by Cambridge, Mass.-based Forrester Research Inc. and commissioned by F5, found that of the 200 North American IT decision makers polled, 50% rank the ability to improve their time to recover and decrease data loss from disasters as "very critical."

Though many companies have advanced their disaster recovery plans from relying on tape backups to doing data replication and backup between data centers and major sites, Hohenbrink says there is still room for improvement.

According to the Forrester report, fewer than half of the respondents could recover from a primary data center failure in five hours or less, and only 55% of North American respondents say they would lose five hours of data or less.

Forrester concluded from these metrics that enterprises must determine "how they can optimize their existing data center recovery solutions to the point that their recovery time and recovery point capabilities are measured in minutes, not hours."

A two-pronged challenge

Hohenbrink says the success of data replication lies in the use of WAN acceleration and optimization tools.

WAN acceleration methods such as offloading work from servers—including security and TCP activities, compressing data, pattern matching (which helps companies avoid storing duplicate data), and caching—optimize and accelerate WAN performance without having to boost bandwidth.

"Enterprises that are using data replication face a two-pronged challenge: They must determine how to optimize performance of the replication solutions as well as extend these solutions to cover critical remote sites at a reasonable cost," Hohenbrink says.

Hohenbrink admits that enterprises face significant challenges. "In many instances, issues related to wide area network connectivity between sites are key limitations to achieving these

goals," he says. Among these obstacles are latency, packet loss, network congestion, actual bandwidth not matching expected bandwidth, and the expense of ever-increasing pipes.

Another critical factor is the amount of data that enterprises must protect. According to the Forrester study, "many enterprises in financial services, life sciences, and managed services measure their storage capacity in petabytes." The study finds that "because storage capacities grow dramatically each year, enterprises that lack enough bandwidth today to support their disaster recovery solutions will only find this problem compounding over time."

Assess the risk

Data replication efforts can be hindered by any or all of these problems, leaving organizations to rely on stopgap measures. "Often, they think there are only two obvious things that can be done: buy more bandwidth, or don't replicate as much data," says Charlie Cano, F5 solutions architect.

But that is a dangerous viewpoint. "Buying bandwidth is a costly endeavor, and it only addresses the low-hanging fruit," he says, adding that limiting the amount of data replication due to WAN performance problems could put recovery efforts at risk. For instance, if organizations choose to not back up remote sites, then they lose access to critical data from the field in the event of a disaster. Since more and more organizations are distributed, including law firms, construction firms, and consultancies, this could have a devastating result.

The Forrester report bears this out, showing that remote sites, which can be anything from local branches and sales offices to regional data centers and headquarters, are rarely considered as part of an enterprise's disaster recovery plan. "A majority of enterprises report that many of their remote sites are not protected from any local disaster through remote application or remote backup to a central data center or other facility," according to the study.

Joe Skorupa, research vice president for enterprise network services and infrastructure at Gartner Inc. in Stamford, Conn., says companies must undergo a risk assessment to evaluate the level of exposure they have at the data center as well as at remote sites. "You have to decide how big an exposure you're willing to accept in terms of data being out of sync. If you can't deal with any exposure, then you'll have to write data synchronously to multiple locations," he says.

Because of its real-time nature, synchronous replication has many limiting factors, however. These include shorter distances and higher costs. In reality, most companies can accept some level of loss and therefore are better suited for asynchronous data replication.

Manage expectations

Stephanie Balaouras, senior analyst at Forrester, says IT should not undertake this assessment alone. "The onus is usually on IT to execute, but really it should be a case of IT working with line-of-business owners to define data recovery and data loss requirements. IT needs to understand the line-of-business sensitivity to data loss and downtime," she says.

One way to manage expectations is to put in place a charge-back and reporting process so the line-of-business unit can see the costs of various solutions, including WAN acceleration and increased bandwidth. "If you can prove that by deploying some WAN acceleration you can improve recovery without increasing bandwidth, that is a justification that will help line-of-business owners," Balaouras says.

She advises companies to widen their scope to look at mission-critical business processes rather than "narrowing their view to a single application."

Map security needs, test regularly

The next step, according to F5's Cano, is to map out security needs. "If you're transferring business-critical, core information, you need to have encryption and authentication in place," he says. Organizations that fall under compliance mandates such as the Sarbanes-Oxley Act of 2002 or the Health Insurance Portability and Accountability Act (HIPAA) must integrate regulations into their data replication plans.

Once organizations understand their security and compliance obligations—in the data center and at remote sites—and they've put a plan in place, they must test that plan on a regular basis. Some major things to consider are how many people will need to access data and applications remotely during a disaster and what types of devices they will be using.

"You have to look at ensuring high availability and optimizing the data transmission over all types of connections," Cano says. He adds that companies should go over these plans with employees, ensuring that all licenses are up to date, that they have the appropriate capacity, and that recovery time objectives (RTOs) and recovery point objectives (RPOs) are met.

To achieve the expected RTOs and RPOs, companies may need to use a combination of synchronous, asynchronous, or even batch/scheduled replication technologies. Hohenbrink says companies must consider site separation "because recovery sites must be far enough away to escape the likeliest of local and regional threats such as natural and man-made disasters."

As they go through the process of measuring their RTOs and RPOs, companies can meld in products that help optimize the response of their wide area network. Ken Salchow, manager of technical marketing at F5, says there are solutions for every aspect of data replication security, optimization, and acceleration. For instance, the BIG-IP Global Traffic Manager provides high availability and traffic management for IP applications and services running across multiple data centers. The FirePass SSL VPN device ensures that companies have a trusted access solution in place in case of a disaster. Finally, F5's WANJet enables the delivery of LAN-like performance over the WAN by accelerating mission-critical applications.

Salchow adds: "In a disaster, most people are so focused on the networks themselves, they're not looking at the big picture, which is delivering applications and data to users. Our products help you do that."

Additional resources

F5's application optimization portal *(www.f5.com/solutions/techportals)*
F5 white paper on the myth of bandwidth and application performance *(www.f5.com/solutions/technology/bandwidth_myth_wp.html)*
F5 white paper on WANJet in data replication environments *(www.f5.com/solutions/technology/data_replication_wp.html)*
WAN optimization demo *(www.f5.com/products/WANJet)*

Sandra Gittlen is a freelance writer in Northborough, Mass.