InformationWeek :: analytics

December 2008

Perfect Storm: New Techs Aid Disaster Recovery Plans

Analytics Report

High-speed metro Ethernet, collocation data centers, server virtualization, advanced data replication and application recovery applications—all these enable even small shops to get back online fast, without breaking the bank. What are you waiting for?

By Howard Marks

TABLE OF CONTENTS

TABLE OF CONTENTS

TABLE OF CONTENTS

Howard Marks is the founder and chief scientist at Networks Are Our Lives, a storage and networking consultancy based in Hoboken, N.J. In over 25 years of consulting, Marks has designed and implemented storage systems, networks, management systems, and Internet strategies at organizations including American Express, J.P. Morgan, Borden Foods, U.S. Tobacco, BBDO Worldwide, Foxwoods Resort Casino, and the State University of New York at Purchase.

He has been an a frequent contributor to *Network Computing* and *InformationWeek* since 1999 and a speaker at industry conferences including Comnet, PC Expo, Interop, and Microsoft's TechEd since 1990. He is the author of *Networking Windows* and co-author of *Windows NT Unleashed* (Sams).

Executive Summary

If there's been one silver lining to the scope and visibility of the disasters that have devastated various regions in the past few years, it's that senior managers at most organizations are now more cognizant of the need for business continuity and disaster recovery planning and preparedness. At the same time, the introduction of new technologies and the maturation of others have made effective disaster recovery products accessible to a wider swath of organizations than ever before.

Once, if a midsize company lost all its systems in a fire, it might never have recovered. But now, relatively new technologies, like server virtualization, metropolitan Ethernet services, continuous data protection (CDP), and inexpensive disk storage, offer new solutions to what just a few years ago seemed like intractable problems.

Of course, IT pros realize that it's generally more mundane matters than earthquake, flood, or other biblical-level pestilence that incite organizations to activate their disaster plans. And, strategies for disaster recovery planning (DRP) must reach beyond the IT department to include emergency notification, human resources, public relations, and other non-IT disciplines.

In this report, we'll help IT with its contribution to DRP by zeroing in on the process of matching disaster recovery products with an organization's business continuity objectives, and providing guidance for companies to find a path to comprehensive protection.

Research Synopsis

Survey Name: *InformationWeek* Business Continuity Planning Survey
Survey Date: January 2007
Region: North America
Number of Respondents: 560 total; 125 at companies with 500 to 4,999 employees

Purpose:
To ascertain the current status of IT disaster preparedness within a business organization, what methods are used, and the organization's intent to improve preparedness.

Methodology:

InformationWeek surveyed 560 North American organizations with business continuity/disaster recovery systems planned or in place to determine current continuity strategies and their intention to implement new strategies over the next two years. We then filtered the data to focus on 125 users in organizations with 500 to 4,999 employees. This eliminated small businesses that can't afford DR sites and large enterprise users for whom gold-plated systems are more appropriate.

The survey was conducted online, and respondents were recruited via an e-mail invitation sent to qualified *InformationWeek* readers containing an embedded link to the survey.

Times Are Changing

IT disaster recovery planning (DRP) has been practiced for decades by *Fortune* 500 companies that can afford products designed and priced to protect the mainframes and other high-end systems typical of big business IT. Smaller organizations, however, lack the deep pockets to develop such plans, so they've either ignored the problem or sent backups off site for protection.

Then came the game-changing disasters of the 21st century, notably 9/11 and Hurricane Katrina. Suddenly, CEOs of all-size companies wondered, "What if?" For too many, the answers they got back weren't reassuring. Rising importance of regulatory compliance only added to the stress level.

The good news is that 79% of our respondents already have business continuity/disaster recovery plans in place, and the remaining 21% have the issue on their radar. But has IT done everything possible to keep potential data loss to a minimum, whether the emergency is a power outage, sabotage by a malicious insider, or natural disaster?

PRIORITIZATION IS KEY

Proper strategy for protecting mission-critical applications requires balancing the possibility/frequency and costs of an outage against the costs of backup servers, DR sites, bandwidth between the primary and DR sites, and manpower to manage it all.

Most businesses today are totally reliant on IT to accomplish their missions, so failures of key IT systems typically have substantial costs in lost business and productivity. How much any

Figure 1

Business Continuity/Disaster Recovery Prevail

Does your organization have a business continuity/disaster recovery system in place?

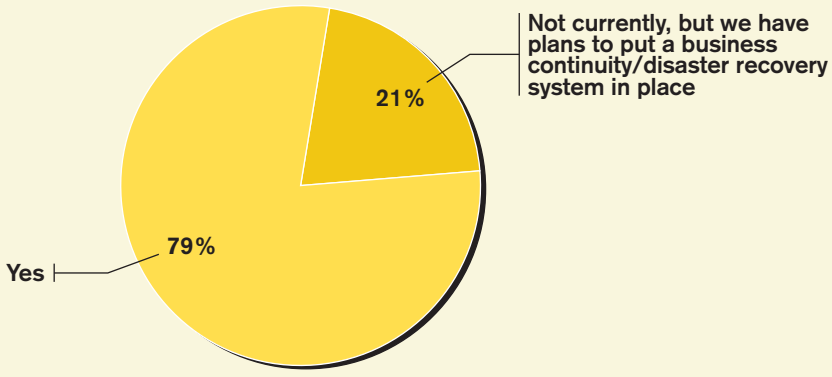

Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

given outage will affect your organization depends on how you do business, but total revenue divided by 2,000 is a rough first approximation of the hourly cost of downtime. Take a law firm with 200 lawyers, each billing $300 an hour; each hour of downtime costs $90,000 in billings.

This rough approximation overstates the cost of short outages and understates the cost as outages get longer. Most organizations will recover from a 15-minute outage with little cost. Customers placing new orders can be asked to call back, and workers can make up the lost time. But if the outage runs into days, not only will employees not be able to make up the work they couldn't accomplish during the outage without overtime—or at all—but many customers will go

Figure 2

Business Continuity Implementation Decisions A Collaborative Process

How are each of the following individuals involved in the decision to implement a business continuity plan?

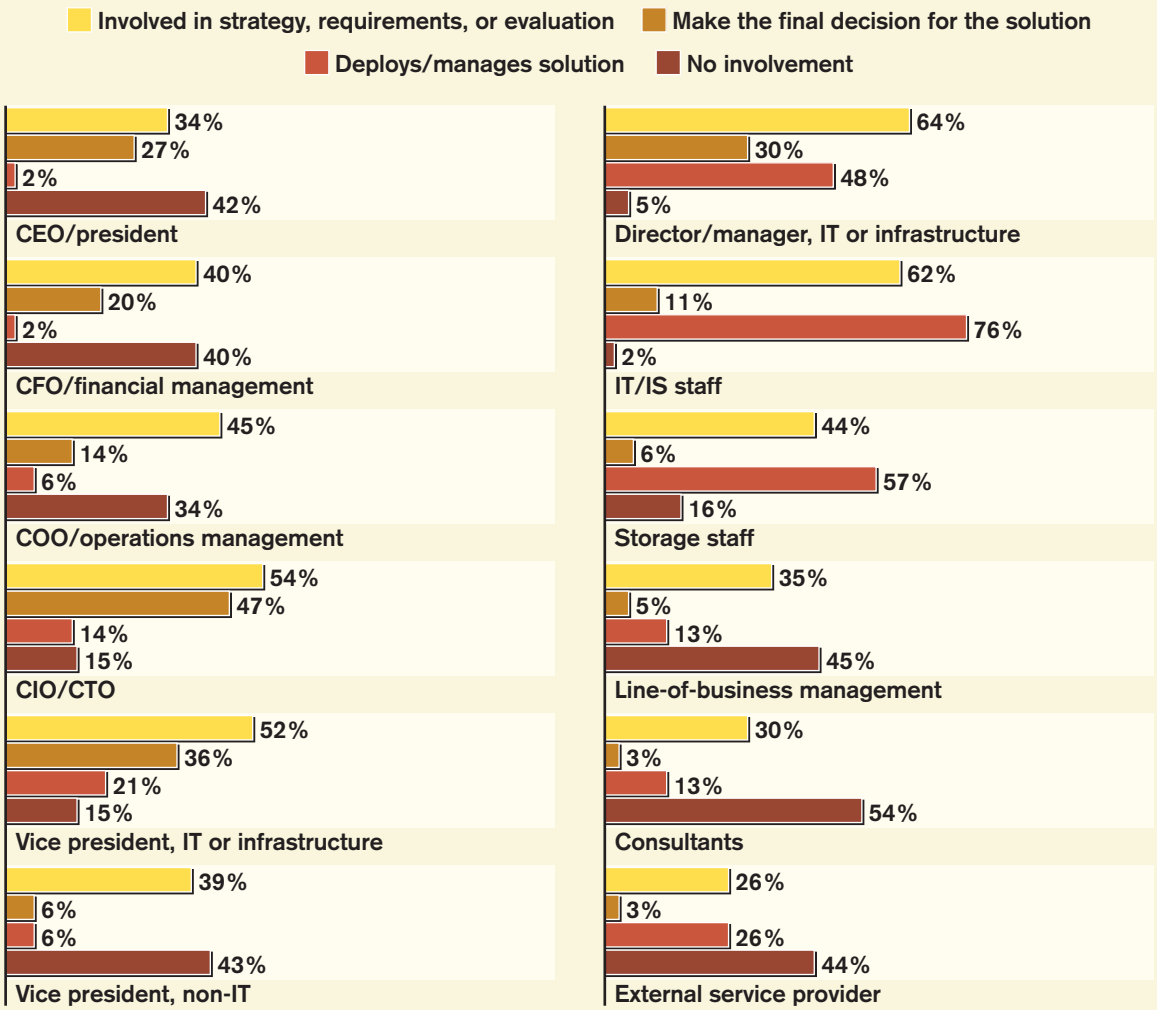

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

elsewhere, costing not just current but future sales if they continue to do business with your competitors after you've recovered.

Clearly, senior management must assess risk and decide what level of funding makes sense for a DR site and other systems to ensure that mission-critical data and applications are safe; our survey respondents indicate that while everyone has input into the DR plan, senior management has the final decision-making power.

Assessing risk and making informed decisions requires prioritization, and the best way to get a clear picture of priorities is to survey your employees. Users know what applications they run and how important they are to the organization. Some IT groups maintain lists of users, their applications, and other IT services like DNS or Active Directory, but more typically, this valuable information is locked up in someone's head.

It's also important to ask the right questions. While employees may say all of the applications they use are "important," it's better to ask what they need to get the job done on a daily basis and what they can live without for a few days or longer.

Impact Assessment: Server Virtualization For Disaster Recovery

	Benefit	Risk
IT Organization	○●●●● Smart use of server virtualization enables the IT organization to provide fast recovery for more applications with a fixed hardware budget. Utility costs are also decreased.	○○○●● Packing too many recovery servers on a single virtualization host may reduce performance during an emergency. Enterprise-class virtualization software is expensive, and app licensing can be hard to decipher.
Business Organization	○●●●● Virtualization enables businesses to have more servers standing ready, resulting in more comprehensive application availability in a crisis. Virtualized DR servers may also fill in during maintenance windows.	○○○●● Over-provisioning DR servers—even virtualized ones—can siphon funds away from other business initiatives. As in the main data center, guard against VM sprawl.
Business Competitiveness	○○●●● A company that recovers more quickly than its competitors gains a significant advantage in a regional disaster, and virtualization provides a recovery-speed advantage.	○○○○● There's no real competitive disadvantage to DR, whether virtualized or otherwise. The killer is in *not* planning for disaster, and virtualization can make that process less intimidating.

Bottom Line

Using virtualized servers for DR enables organizations to provide fast recovery for more servers and applications, reducing the reliance on complex backup/restore processes while slashing costs associated with multiple 1U dedicated servers, as well as for DR facilities including utilities and rack space. For small companies, virtualization is the magic sauce behind cloud-based DR offerings that can bring comprehensive recovery plans into financial reach.

Note: Number of dots indicates level of benefit and risk; one dot equals low benefit/risk, and five dots equal high benefit/risk.

Bear in mind that user need may or may not be synonymous with executive priorities when it comes to the bottom line. For example, the advertising department may declare Photoshop and QuarkXPress to be mission-critical applications, however, senior management may decide that the organization can survive for the first few days of a disaster without developing new ads.

Once an inventory is collected and applications are prioritized, recovery objectives for each application tier are determined. A law firm, for example, may develop an inventory that looks something like this:

Application	Tier	Real Cost Threshold	Mission-Endangered Threshold	Max Data Loss
E-mail (mailtone)	1	Less than 1 hour	8 hours	Less than 1 minute
E-mail (contents)	1	2 hours	24 hours	none
Time and billing	2	2 days	14 days	1 hour
Document management	1	Less than 1 hour	8 hours	15 minutes
Microsoft Office	1	Less than 1 hour	8 hours	15 minutes
Payroll	2	2 days	14 days	1 day
HR	1	1 hour	Days to weeks	1 day

As you can see, the e-mail application has been separated into two parts: mailtone and contents. Mailtone is the ability to send and receive new messages. It is a higher priority, because while contents may take time to restore from the backup, users can get a lot of value from a working e-mail system.

THE KEY MAN ASSUMPTION

One mistake is to assume that key IT personnel will be available when a DR plan is activated. For instance, if the plan says the Exchange servers will be restored only after Active Directory is brought back online using an application like MailMagic 5.6.4, but only Karen the Exchange administrator has ever seen MailMagic, and she's MIA, you have a problem.

Your disaster recovery plan should be detailed enough to instruct technically adept administrators who don't have specialized knowledge of your applications or environment. A good rule of thumb is that directions for recovering any Windows server or application should be detailed enough for any Microsoft Certified Systems Engineer to follow.

Only 29% of the respondents in our survey have actively planned for recovery without key personnel. More than half believe they'd muddle through, but getting replacements up to speed would extend the time until full recovery.

Figure 3

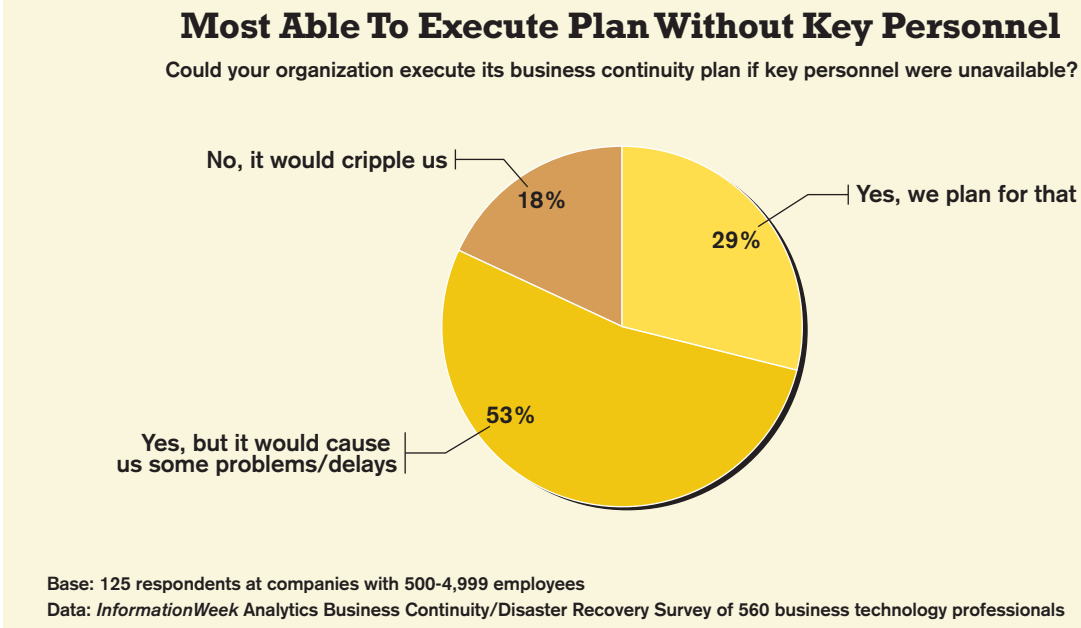

SETTING DISASTER RECOVERY OBJECTIVES

It's necessary to balance the ease and speed of recovery with the cost and degree of disruption a disaster recovery plan creates. This starts with understanding the objective measurements of system protection and recovery.

DISTANCE MATTERS

The first metric is physical distance. If your recovery site is too close, it could be destroyed by the same disaster. On the other hand, putting your DR site too far away will make maintaining the systems more expensive and difficult. In addition, longer communications lines are not just more expensive; the unavoidable latency caused by long distances can mean problems, from performance loss to application failure.

LOCATION, LOCATION, LOCATION

The second metric is deciding whether your business is synonymous with its site. Hotels are their locations: If the building burns to the ground, it won't matter much that all your applications are up and running at a DR site 300 miles away. On the other extreme, large, geographically dispersed enterprises need to have multiple DR sites, at great distances from one another, to support global operations.

Smaller firms hit the middle ground. Most will have to settle for a single DR site, and they will need to balance the safety distance provides with the funds available. Where you put your DR site depends on the type of disasters common in your neck of the woods.

Figure 4

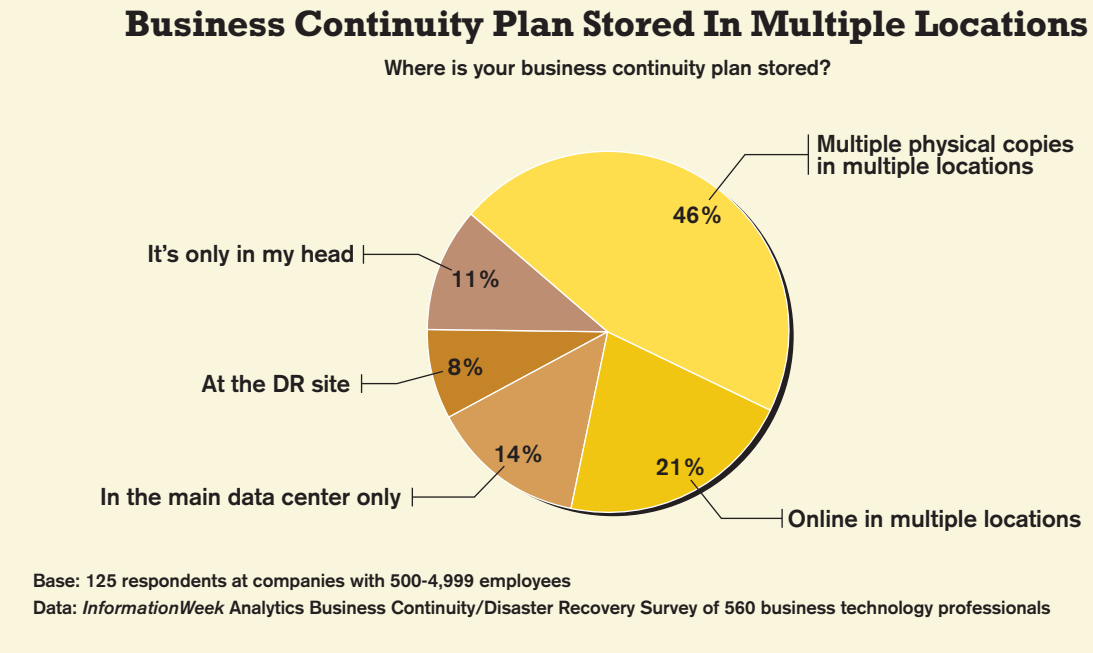

Organizations in coastal areas should look at DR sites inland, on high ground. In California, where you have the trifecta of earthquakes, landslides, and wildfires, place your DR site well outside the usual radius of such activities. For less disaster-prone locales, a DR site 30 miles to 40 miles away, ideally in an area served by a different utility company, is far enough.

TYPES OF DR SITES

When we use the term "DR site," specialized services from companies like SunGard and IBM come to mind. In addition to cold data center space and standby mainframes, these sites offer "shared server" capacity for a monthly fee and provide user workspaces with desks and PCs, as well as some technical help, when you need it. SunGard manages virtual servers for $1,000 to $1,200 per month per server, plus the cost of network bandwidth between your data center and the SunGard site.

A less expensive alternative is to set up your backup servers at a co-location data center. A co-location facility will provide fewer services but will rent you a full rack with power and a redundant Internet connection for $1,000 to $4,000 per month, depending on location, power, and bandwidth usage. An SME looking to replicate five to 10 critical servers would likely spend under $2,000 a month. You will need to provide gear, including servers, storage, firewalls, and IP KVMs, and set up terminal servers since your users aren't going to be working at the co-location center.

Given that a typical co-location data center will be significantly more resilient than what most organizations can build at their home offices—just try convincing your high-rise-office landlord

Figure 5

Sizeable Distance Separates Primary, Disaster Recovery Sites

What is the distance between your primary and disaster recovery sites?

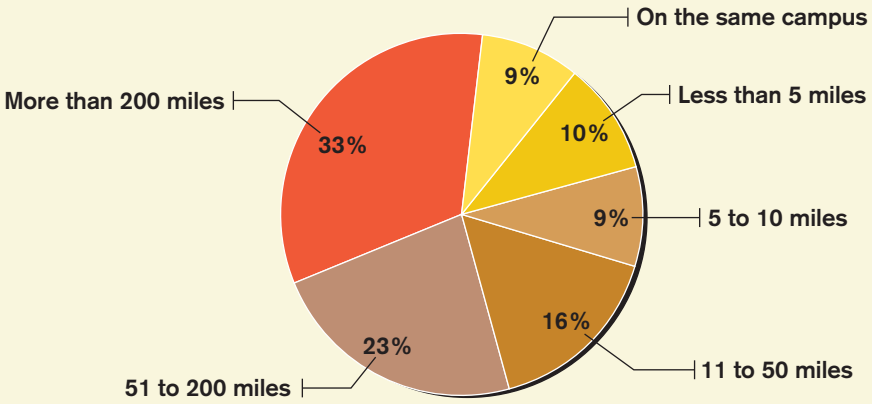

Base: 107 respondents with disaster recovery sites at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

that you need a backup generator in the penthouse—some organizations are choosing to put their primary servers in co-location centers and run backup/recovery systems locally. Remote server management cards, IP KVM switches, and especially virtual server management tools like Virtual Center can make managing these servers remotely much easier.

Using co-location servers as a primary resource makes especially good sense for organizations in areas like California and the Gulf Coast that are especially susceptible to regional-scale natural disasters. A New Orleans law firm could use terminal services or VMware VDI to run servers inland in Baton Rouge, or even Memphis, putting its primary systems in a much safer location. Should a disaster occur that takes out the co-location data center, IT personnel will be on hand at the corporate headquarters to bring recovery systems on line rapidly.

Organizations with two or more geographically diverse sites can reduce out-of-pocket costs for rent and bandwidth by using these offices as each other's DR locations. Connection costs can be cut if metro Ethernet service is available. If a metro Ethernet carrier services your building, look for a co-location center that's on its fiber optic network directly. This can save you 40% or more over using the local telco's standard services.

HAVE DATA, WILL TRAVEL

A disaster recovery or business continuity plan has to address how data gets from the primary data center to the DR site, and how quickly. The recovery time objective (RTO) defines how quickly applications must be up and running. For applications that are protected by convention-

Figure 6

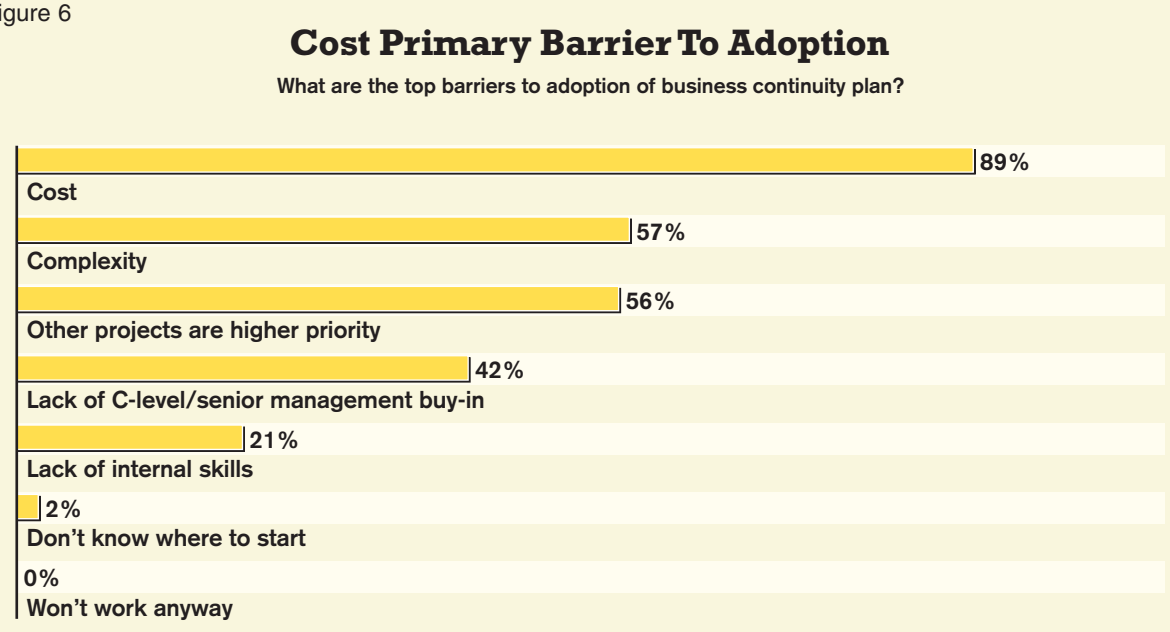

al backup-and-restore operations, it's a matter of hours—at best—so that method should be reserved for applications where the cost of anything faster is greater than the cost of not having that application. Traditional periodic backup and restore can take 12 hours or more, and recovery can take days. There's nothing sexy about shipping tapes across the country, but it can be the most cost-effective solution, especially given a large quantity of data.

If an overnight courier like FedEx takes 24 hours to deliver a box of LTO 3 tapes containing 8 TB of data, that's an effective data transfer rate of 640 Mbps at a cost of $3,000 a month in FedEx bills, plus around $1,000 per year for tapes. However, while shipping tapes is cost effective, most companies can't afford a recovery point objective (RPO) and RTO of two or more days. To meet more stringent recovery schedules, organizations will need to send data from their primary data centers to their DR site over a WAN link.

New technologies, such as data deduplication and WAN optimization appliances, can greatly reduce the bandwidth needed to synchronize data stores between sites. The actual reduction in required bandwidth is highly dependent on the data being replicated and can range from as little as 5% to as much as 95%, so testing is required. Vendors such as Cisco, Data Domain, Hewlett-Packard, Riverbed, and others offer a range of products with a variety of capabilities for conserving WAN bandwidth.

When determining your RTO, remember to consider the practical and human elements. While it's tempting to use whatever technology is necessary to minimize recovery time, it's easy to

Figure 7

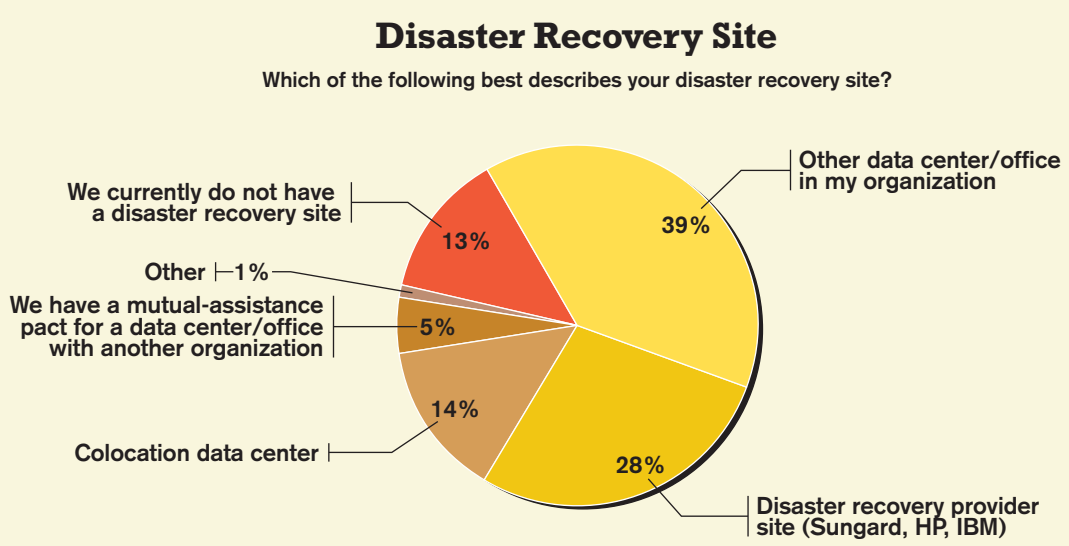

spend $10,000 to avoid losing $5,000. Don't waste time building a gold-plated disaster recovery plan that senior management won't pay for.

Even more important than speed of recovery is the point in time at which data is restored; this is known as the recovery point objective (RPO). The nature of the data being restored determines the importance of time. For instance, an attorney writing a brief is an example of re-cre-

Figure 8

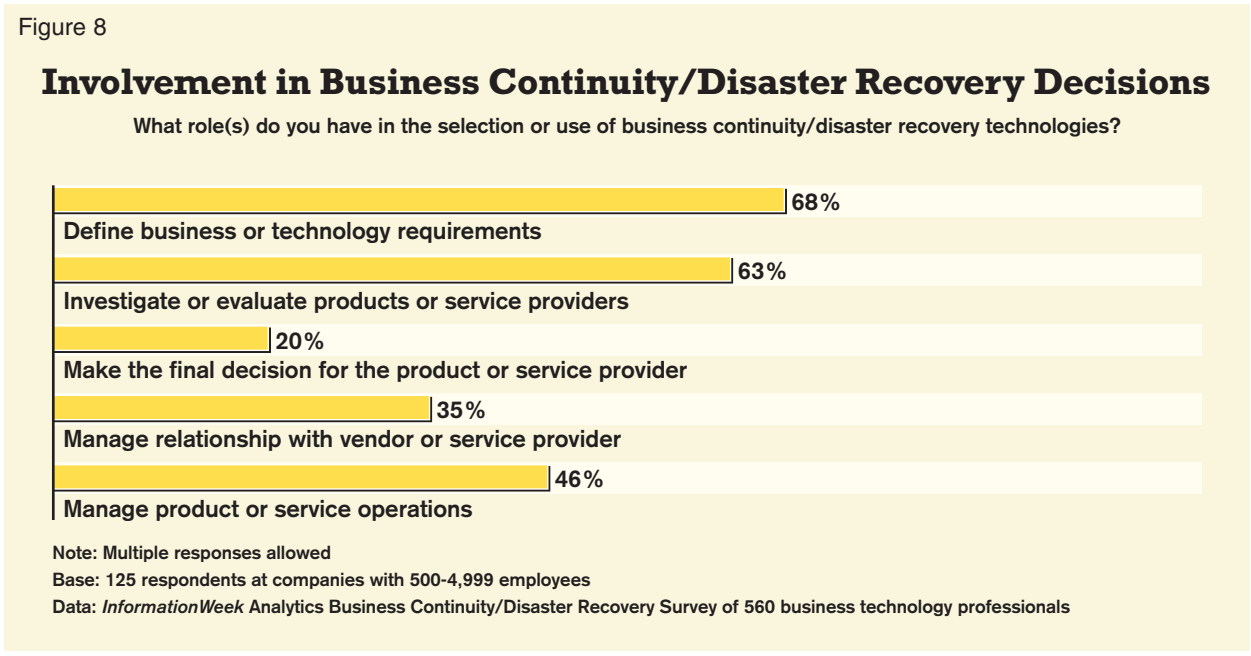

Figure 9

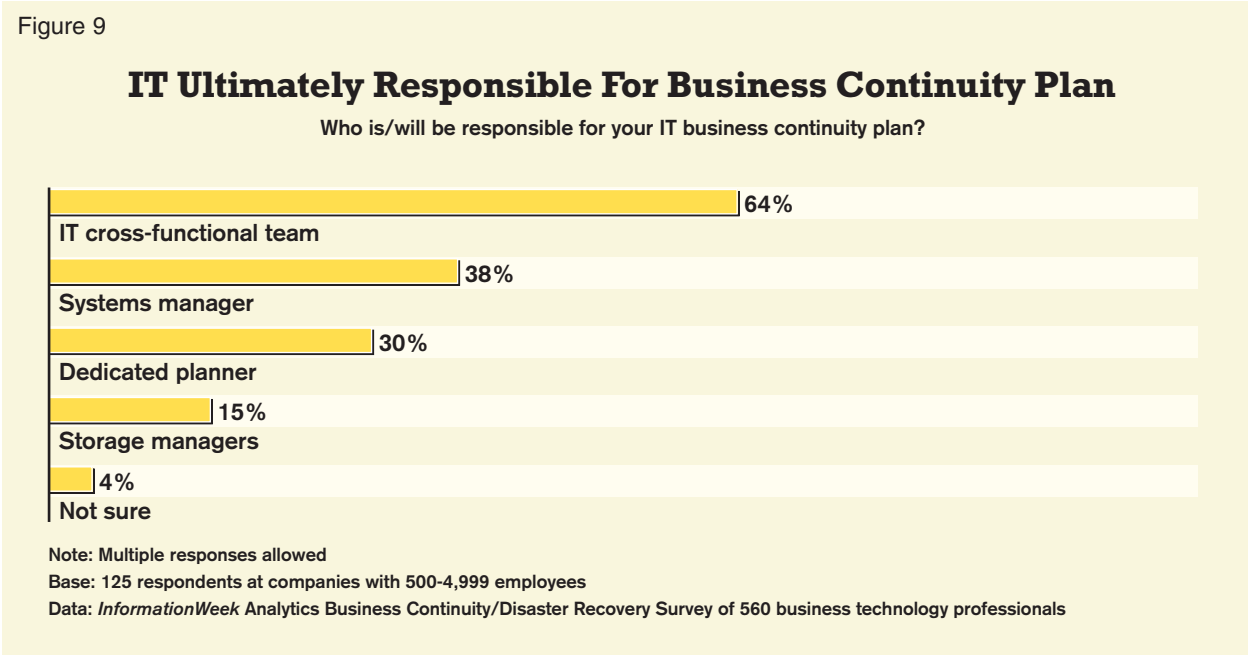

atable data—she can probably rewrite the lost paragraphs of the brief. However, an online store that accepts data from customers over the Internet is a different story. In that case, if the database is restored to a point only 10 minutes old, orders that occurred in the intervening time will be lost and the organization will have no way to know the contents, or which customers were affected. So the document management server where our lawyer friend saves her briefs may have a recovery point objective of one hour or so, but the Web store would have an RPO of just seconds.

REPLICATION OPTIONS

To reduce RPOs from hours to minutes, organizations need to replicate write requests in real time to a duplicate store at their DR sites. There are various replication products available, each designed to provide the best balance of cost, bandwidth requirement, RPO, and ease of management for a given set of uses

The first major difference among replication products is when they replicate the data. There are three commonly used replication methods: synchronous; asynchronous; and snapshot, also known as "point-in time."

Synchronous replication systems duplicate each write request, sending one to the primary data store and the other to a secondary store. They wait for both stores to acknowledge that they've received and written each request before sending an acknowledgement back to the original application or operating system.

Figure 10

Recovery Uptime For Mission-Critical Applications

If your primary data center were destroyed by fire or other disaster, how long would it take to bring your mission-critical applications back online?

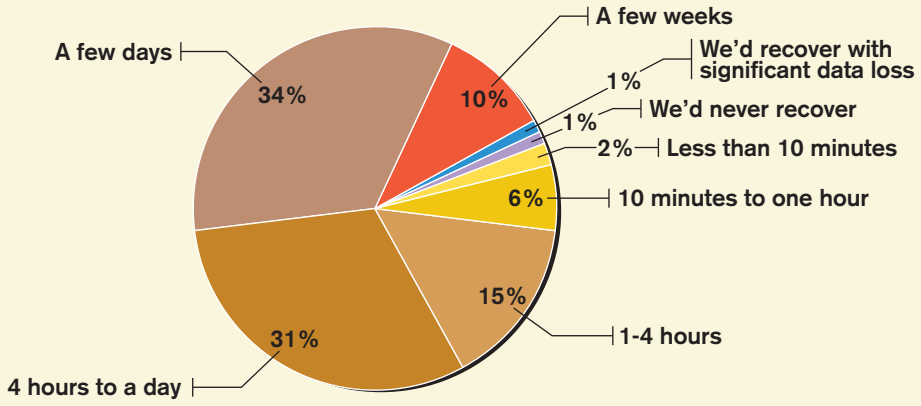

Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

While a disaster may occur in the middle of a transaction, forcing that transaction to be backed out, synchronous replication keeps the secondary storage pool fully updated at all times for an RPO of essentially 0 seconds.

While great at minimizing RPO, synchronous replication introduces some delay as data is sent to the secondary store and the acknowledgement is sent back. Because it takes about 1 ms for a packet to travel 100 km, sending a write request to your DR site 60 miles away adds 2 ms to every write operation from network latency alone. Add in transmission time—a 4 KB block takes about .7 ms on a T-3 line—and your application could slow to a crawl if your DR site is far away or you have limited bandwidth. As a result, synchronous replication is used primarily by large organizations over dedicated fiber links.

Asynchronous replication keeps latency from becoming an issue by acknowledging disk writes as soon as data is written to the primary repository. As with synchronous replication, disk writes are sent to the secondary data store, so there may be some delay. Should the amount of data transmitted temporarily exceed the available bandwidth, most asynchronous replication systems will buffer writes in memory or in a journal file on disk.

Packets can be lost or arrive out of order across a TCP/IP network, so asynchronous replication needs to have mechanisms in place to ensure that changes are applied to data in the order they were intended. Just imagine what would happen if the on-hand quantity of widgets was modified by a series of transactions, but those updates were applied out of order. Rather than show-

Figure 11

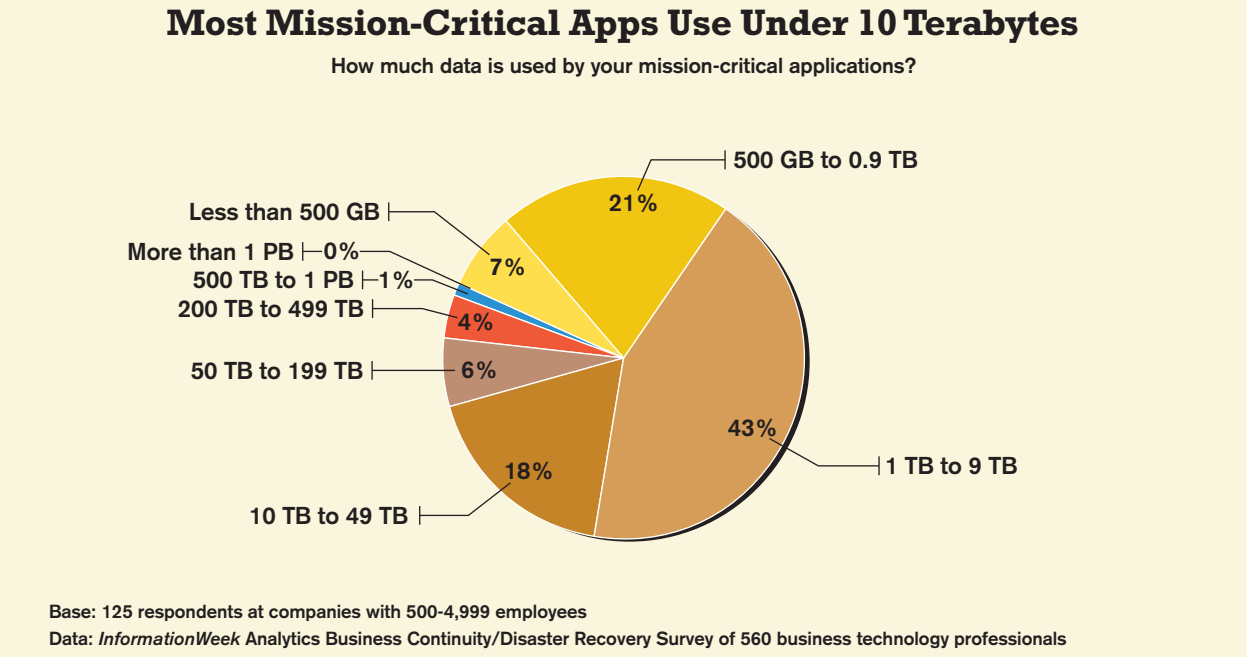

ing that widgets are sold out, the system could show there are still some in stock. Because synchronous replication systems post each write request before the next is processed, they inherently enforce write-order integrity.

Asynchronous replication systems may also include features like data compression, scheduling,

Figure 12

Recovery Time Objective

What is your recovery time objective (RTO) for your mission-critical applications?

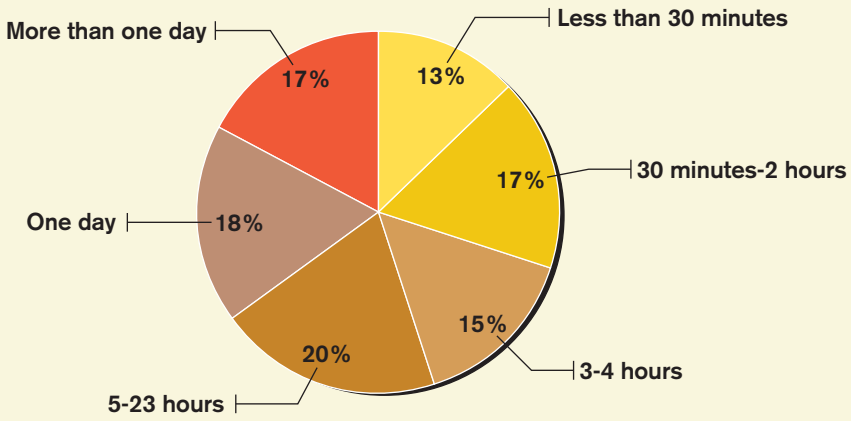

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 13

Recovery Point Objective

What is your recovery point objective (the point in time your data will be rolled back to when you recover)?

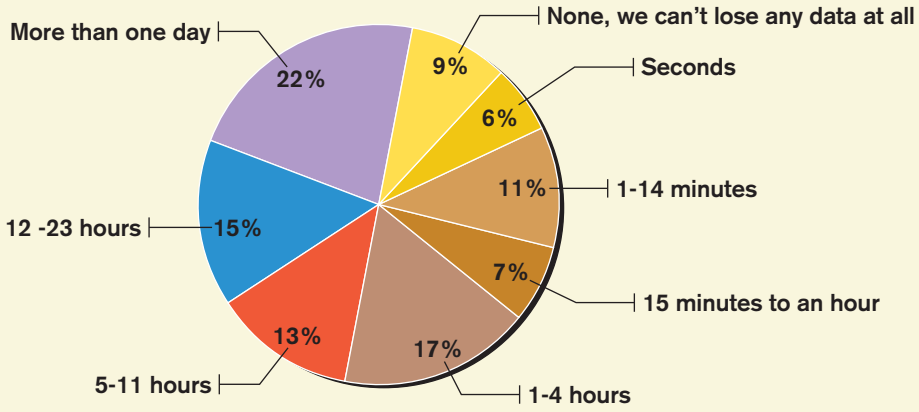

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

and throttling, all of which allow replication traffic to share a WAN link with other applications. Limiting the maximum bandwidth used for replication at any given time enables IT to reserve bandwidth for user traffic as needed, and keep the entire pipe available for replication traffic when user traffic is minimal. If the data being written exceeds the bandwidth allocated, the excess will be logged in a journal and sent later.

Asynchronous replication systems can satisfy RPOs as short as a few seconds if traffic doesn't exceed the available bandwidth. Typical asynchronous replication schemes do bog down at peak periods, making this approach best for applications with RPOs of a few minutes.

Snapshot, or "point-in-time" replication, takes advantage of a disk array or the copy-on-write snapshot provided within the operating system. The copy-on-write snapshot, taken periodically, contains the blocks of data that have changed since the last snapshot was taken. Changes are replicated to the target system and applied to the data there.

When a block is written to several times between snapshots, this type of replication will copy only the last contents of the block. Synchronous or asynchronous replication would capture each change in turn, so snapshot replication systems can be more bandwidth efficient. They can also be integrated with Windows VSS (Volume Shadowcopy Service) to ensure that databases are in a consistent state for each replication. Snapshot replication systems will always lag behind the source system by the snapshot frequency plus the time needed to transfer the snap-

Figure 14

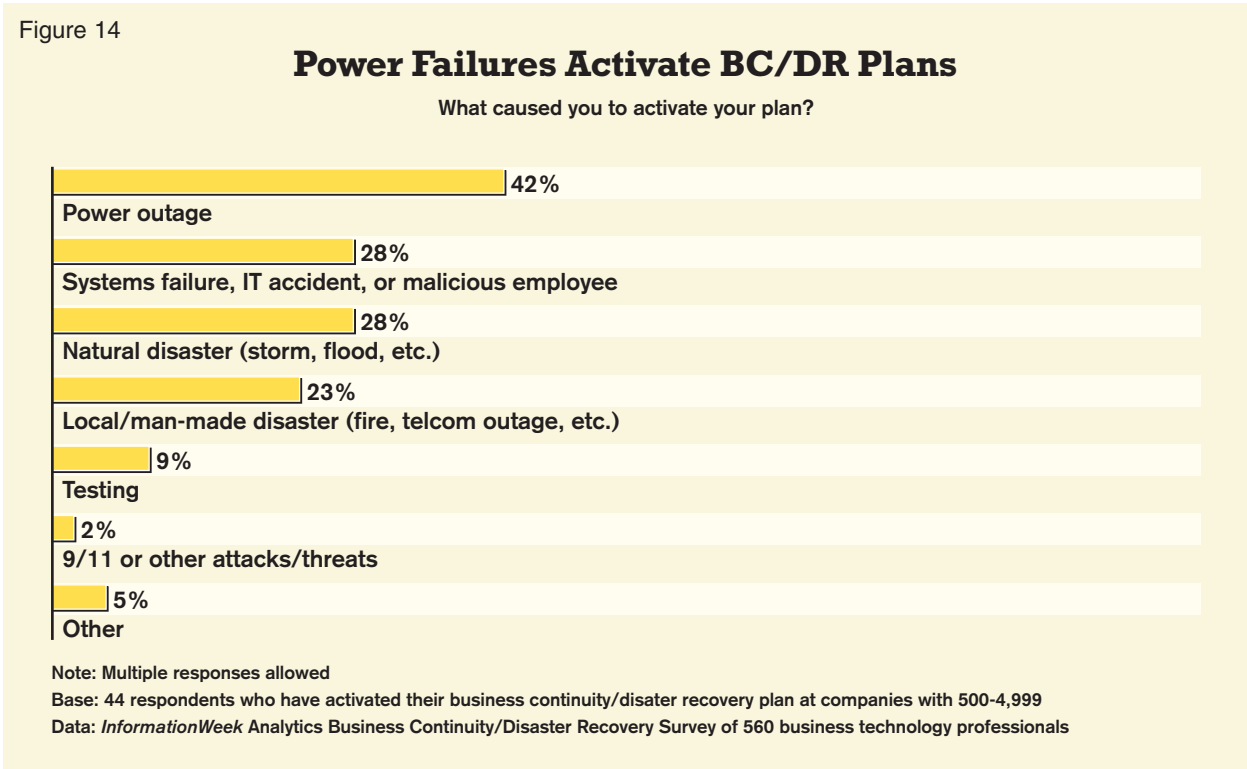

Figure 15

Mission-Critical Apps Primarily Protected Off-Site

What is the data protection method you use on your most mission-critical applications?

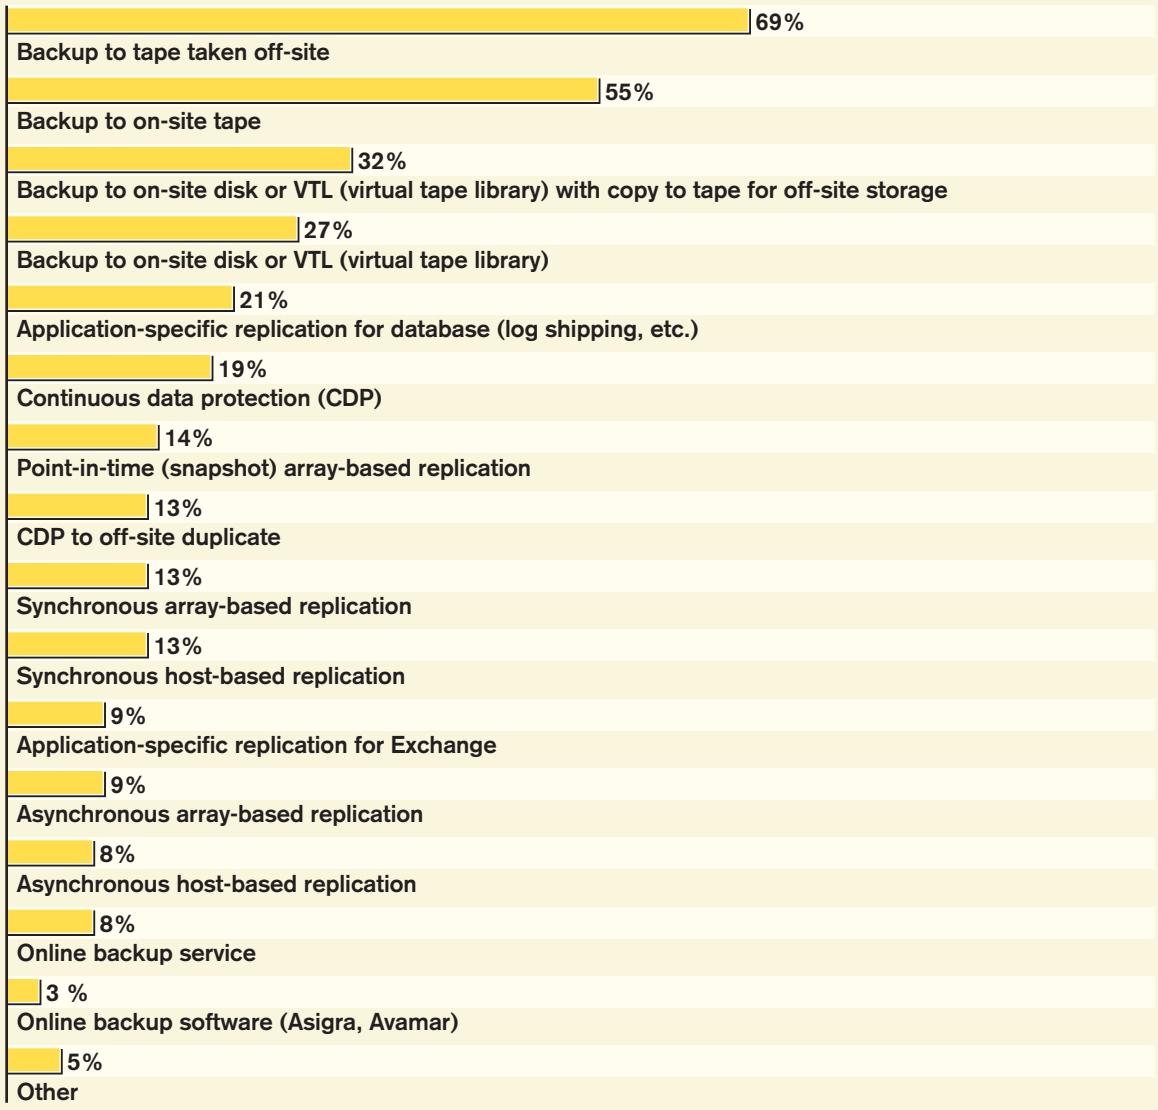

Note: Multiple responses allowed

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

shot data and apply it at the target. Because snapshot replication systems transfer snapshot data asynchronously, vendors will often call them asynchronous replication.

Replication can be performed by disk arrays, software on host servers, or dedicated appliances in a storage area network. Array-based replication lets a system administrator replicate logical drives (LUNs) for multiple host systems in a single replication set while placing no load on the servers. This is the method most often used for large, OLTP applications.

Unfortunately, while most SAN disk arrays support both synchronous and asynchronous replication, there is no standard replication protocol, so array-based replication is limited to arrays from the same family. Most vendors don't even support replicating from their high-end monolithic arrays to their own midrange modular arrays. Connecting Fiber Channel arrays over a wide area network can also require expensive Fiber Channel-to-IP routers or dark fiber connections.

SAN replication appliances allow organizations to use different storage hardware in their primary and standby data centers without the overhead of host replication software. Replication appliances can trap write requests through lightweight server agents that duplicate requests to the appliance and the disk array, or at an intelligent Fibre Channel switch. They then replicate the data across an IP network and write it to a disk array at the target location. Most can also replicate to multiple targets and support continuous data protection; more on that later.

Figure 16

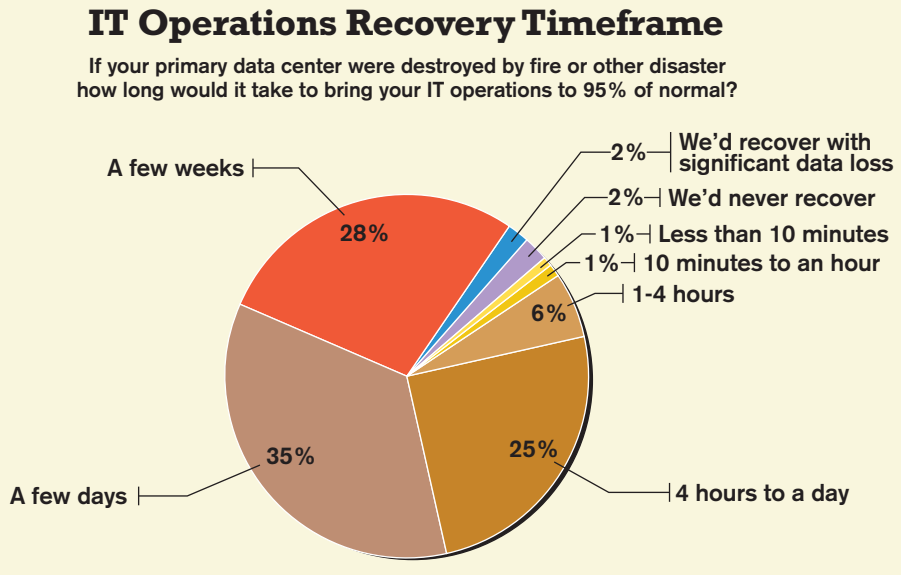

Many smaller organizations can't justify the cost of SAN-based replication but still have a need to protect their data. Using software in the host server to replicate data can provide the same level of protection as SAN-based replication, at a significantly lower cost. Host-based replication also frequently includes server and application failover features, which can speed up recovery.

While host-based replication can be cost-effective, system administrators have to manage each server's replication to its standby doppelganger separately. As organizations grow to need more

Figure 17

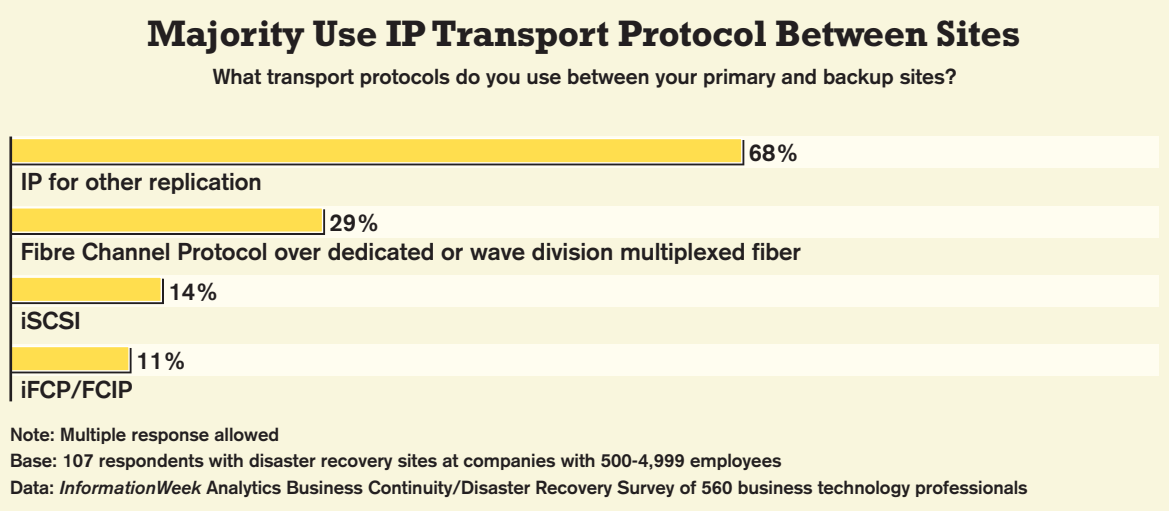

Figure 18

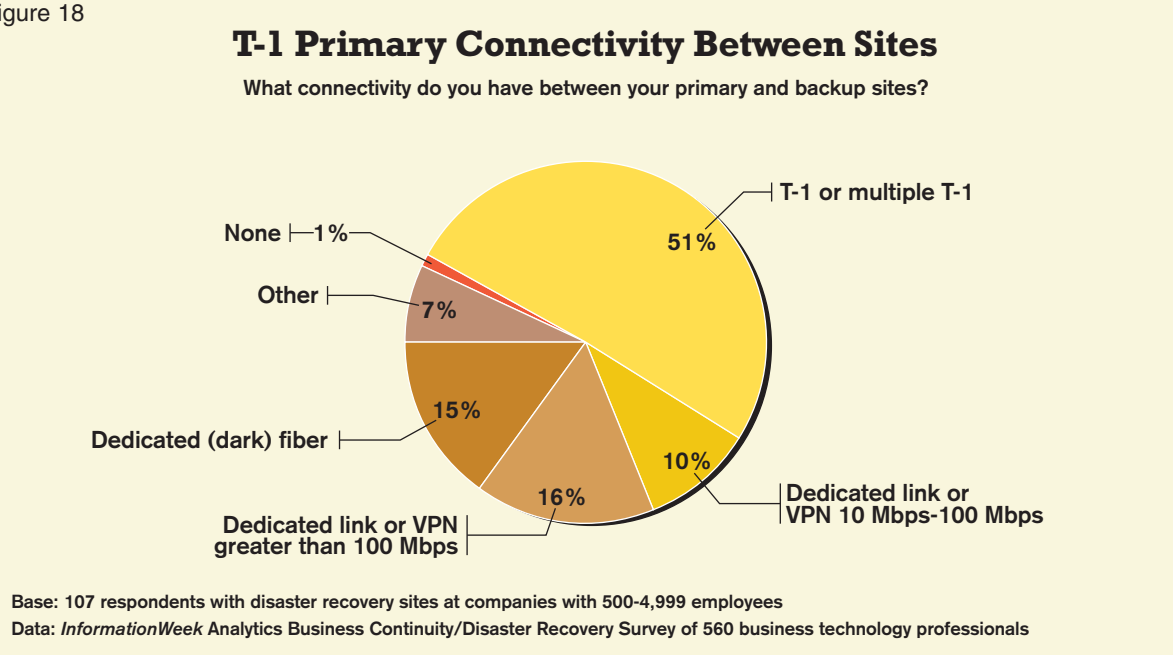

than a handful of critical servers, the administrative overhead of tuning the bandwidth allocations and monitoring replication pairs starts to outweigh the lower acquisition cost of using software to solve the problem.

The DR planner must also decide at what level to acquire the data to replicate. While array and SAN-based systems must replicate each write request to disk block(s), several other options are available for host-based systems. Most commonly, host replication uses a file system filter to trap and replicate file write requests. This allows them to exclude writes to temporary and swap files, which contain many write requests that are not related to critical data.

Another alternative is to replicate application and/or database transaction data rather than disk or file I/O. Application replication must be closely matched to the application it's protecting, but it can also replicate data to a running database or application server. This speeds up recovery and reduces the possibility of a corrupted database at the DR site, since each transaction is verified at the target database server. This replication can be built in to the application server, as with Exchange 2007, SQL Server, and Oracle, or be a third-party application, like Cemaphore's MailShadow or a Teneros appliance for Exchange.

Transaction-level replication can also reduce bandwidth requirements. Modern databases write data to transaction logs as well as the database itself. Systems that replicate at the file or block level must replicate these write requests separately, essentially sending each transaction across the wire twice. A typical database replication will copy the transaction logs from the source to the target server and use the database engine to roll the transaction logs forward into the target

Figure 19

Back-Up Plans For Application Servers Vary

What is your plan for providing application servers in the event of a disaster?

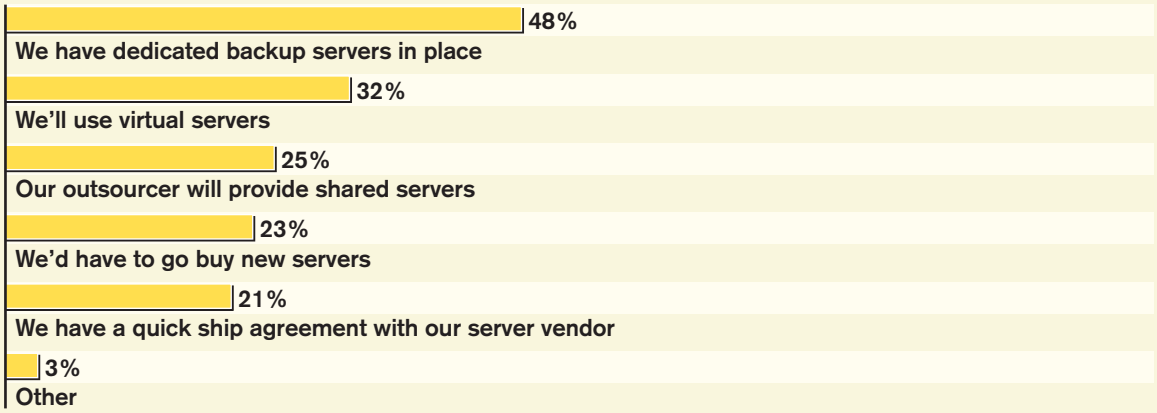

Note: Multiple responses allowed
Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

copy of the database. Transaction replication relies on either transaction log files being filled and copied to the target server, or that they acquire their data from the source database server after it's been written. This means that these systems can't achieve the near-zero RPO that block or file replication products can, often trailing in their recovery time by a few seconds to a few minutes.

CONTINUOUS DATA PROTECTION

Real-time data replication can protect against total server failure, allowing an organization to recover their applications to the moment where the source server stopped sending data. However, one caution is necessary: If the data is corrupted on the primary server, a block- or file-level replication will replicate the data to the backup site, corrupting the data there as well.

Where replication immediately applies that write to a duplicate data store, creating a mirror image of the primary data, continuous data protection (CDP) captures each disk write but keeps a journal of all changes. This allows a system administrator to recover the duplicate data store for any point in time. Vendors including Double-Take, EMC, and NetApp are now adding data journaling to replication products and bandwidth management to CDP products, so that one set of software can serve both purposes.

Our respondents use the full range of products from tape backups to synchronous replication. We allowed multiple answers to these questions, and it would appear our respondents are fans of protection in depth, using multiple methods for their mission-critical applications.

Figure 20

Business Continuity Investment—IT Budget

How much of your overall IT budget does/will your organization devote to business continuity?

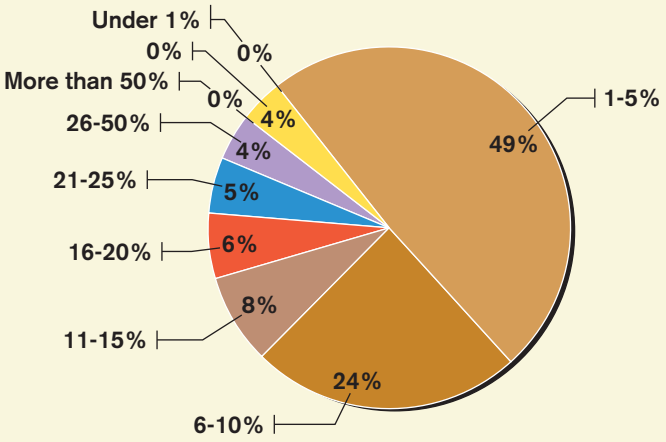

Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

APPLICATION FAILOVER

In many IT organizations, the storage group is given primary responsibility for the DRP. As a result, they spend the majority of their disaster planning time and resources on preserving the organization's data, and not enough on ensuring that applications can be brought back into operation quickly and made available to users.

Even if the system drive of a server has been replicated, bringing that server back into operation is a time-consuming task. If a similar server is at the DR site, it will have to be connected to the replicated system drive and booted. An administrator will then have to log in and reset the network configuration, as the DR site is a different subnet and the recovery server has a different MAC address than the original server.

For applications like Microsoft Exchange that are tightly tied to other network functions, including Active Directory, the process can be significantly more challenging. This is especially true with host-based replication, because the two servers can't share the same boot data.

Application failover or clustering products automate the process of bringing an application and its server(s) back online. They can automatically detect when the primary server goes offline and change the recovery server's name and IP address, update DNS servers and Active Directory, and make whatever other changes are needed to bring the application back online. Applications like VMware's Site Recovery Manager, Neverfail Group's Neverfail, and Marathon Technologies' Everrun all provide these services.

Figure 21

Business Continuity Investment—Storage Budget

What percent of your storage budget does/will your organization devote to business continuity?

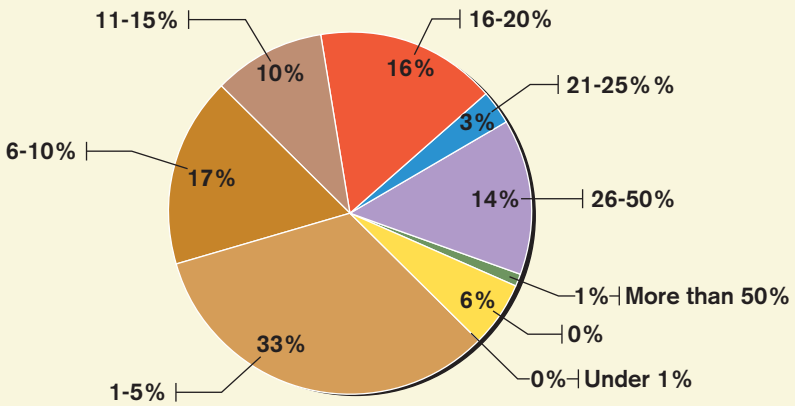

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 22

E-mail Backup Taken Off-Site

What is the data protection method you use on your e-mail/messaging system?

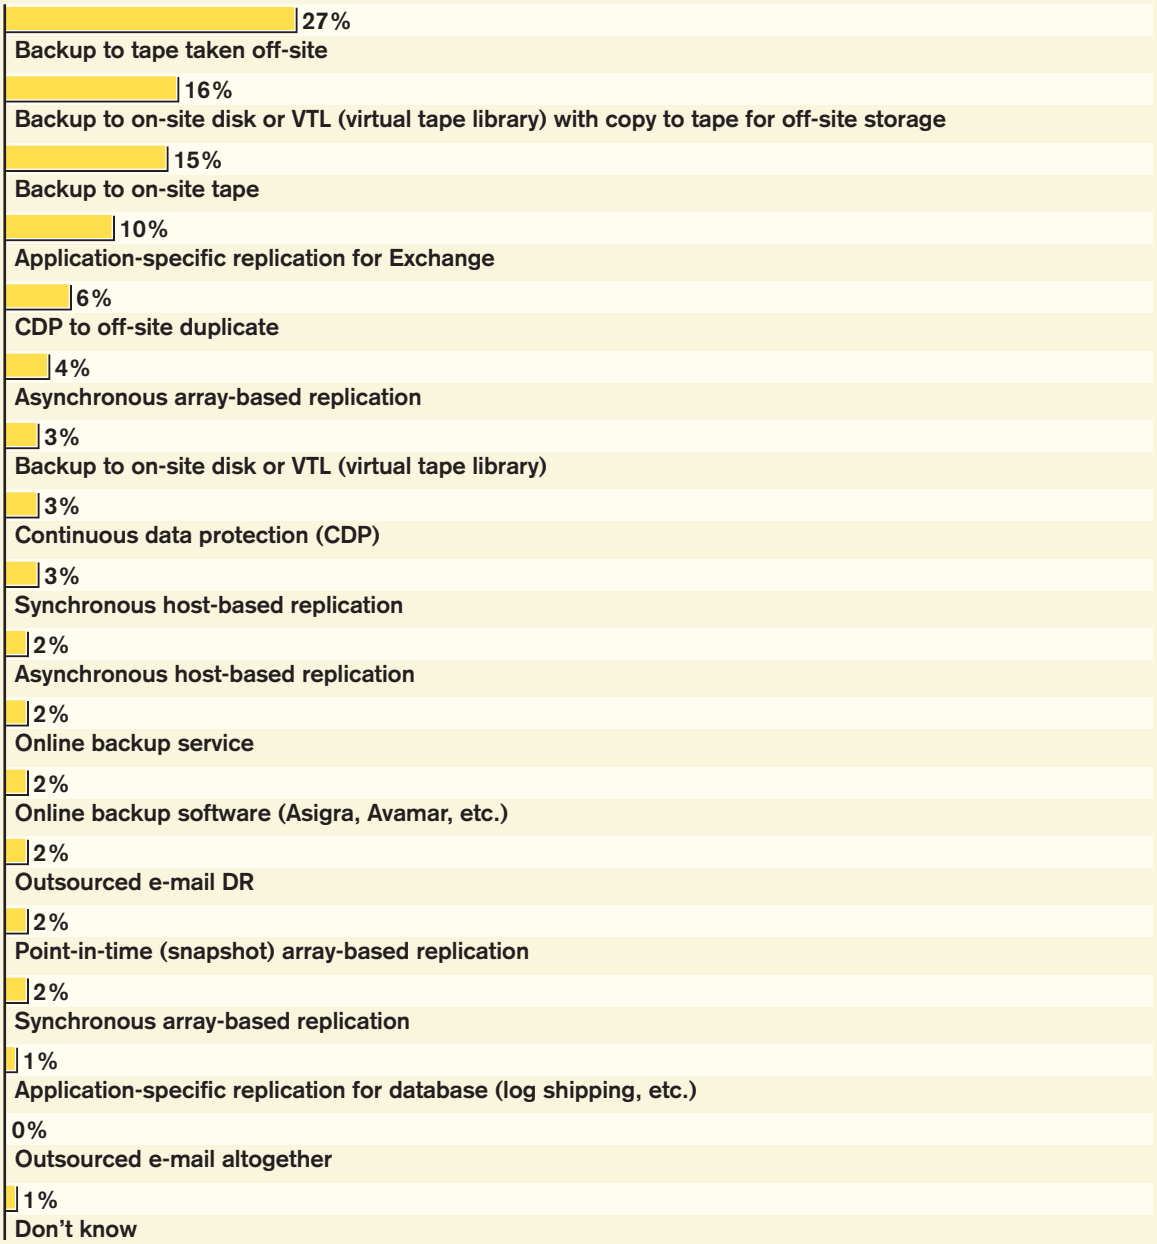

Note: Multiple responses allowed
Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

THE IMPACT OF SERVER VIRTUALIZATION ON DRP

Widespread use of x86 server vitalization has had a huge effect on the disaster recovery process. The obvious impact is a reduction in the number of servers that have to be provisioned, powered, and maintained at a DR site. Five years ago, even the smallest DR site would have had dedicated servers for each application that needed to be recovered quickly, but now, a single virtual server host can handle multiple applications. Not only can organizations save money on server hardware, but the reduced size, power, and cooling footprint of a small blade chassis running several virtual server hosts opens up branch offices and collocation centers as potential DR sites.

Organizations that are virtualizing their production servers can use host-based products, creating a replication pair for each set of hosts rather than for each application. Administrators will have no problem restoring servers to dissimilar hardware. Since the hypervisor hides the underlying hardware, restoring, for example, a Dell PowerEdge 1950's application to an HP DL380 is no longer a struggle.

Virtual server and data migration tools like vMotion, VMware DRS, and XenMotion allow organizations to bring their applications back up on a limited amount of hardware at a DR site. They also enable IT to quickly add additional servers to the cluster and rebalance the load so organi-

Figure 23

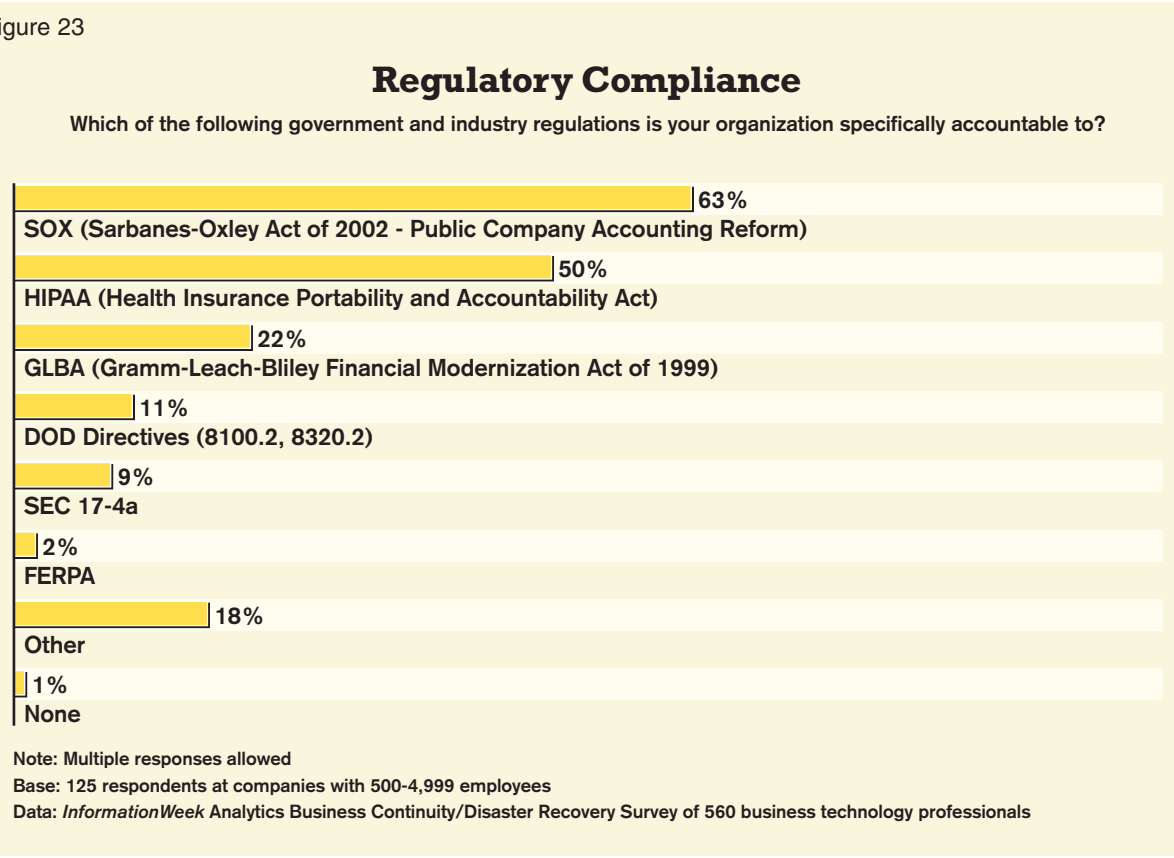

zations can bring systems on line faster and more gracefully, if they take advantage of quick-ship programs and buy more hardware with insurance proceeds.

TESTING 1, 2

Testing a DRP is just like testing a new application and working out the bugs. Disaster recovery tests perform the same function, identifying dependencies and invalid assumptions in the initial draft of the plan. Organizations should fake failures to achieve a full restoration of the service or application to work out the bugs in their plans. However, only 40% of our respondents said that they restored 100% of the applications in their last test within their recovery time objective.

A disaster recovery plan must be tested and revised based on the results of the last test and then retested to verify that the plan reflects the current state of the application. Three-quarters of our respondents test annually or not at all, a frequency that we feel is wholly insufficient. To keep the team sharp and identify changes to systems that haven't been reflected in the DR plan, we believe some testing should be performed at least quarterly, and every application be tested at least once a year.

Creating a disaster recovery plan takes a significant effort, not only from the IT department but also from senior management, who must allocate funds for both the planning project and the systems the planning process identifies as appropriate, as well as setting recovery goals and priorities.

Figure 24

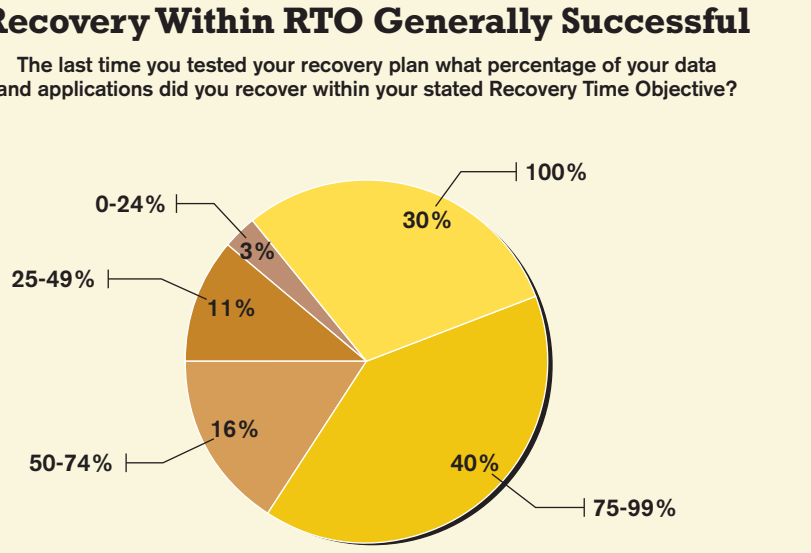

Base: 82 respondents who test their plans at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

High-speed metro Ethernet services, collocation data centers, and server virtualization can combine with data replication and application recovery applications to allow even smaller organizations to build disaster recovery plans that can help them get back online fast, without breaking the bank.

Figure 25

On-Site, Off-Site Methods To Back Up Data

What data protection methods are in use in your organization?

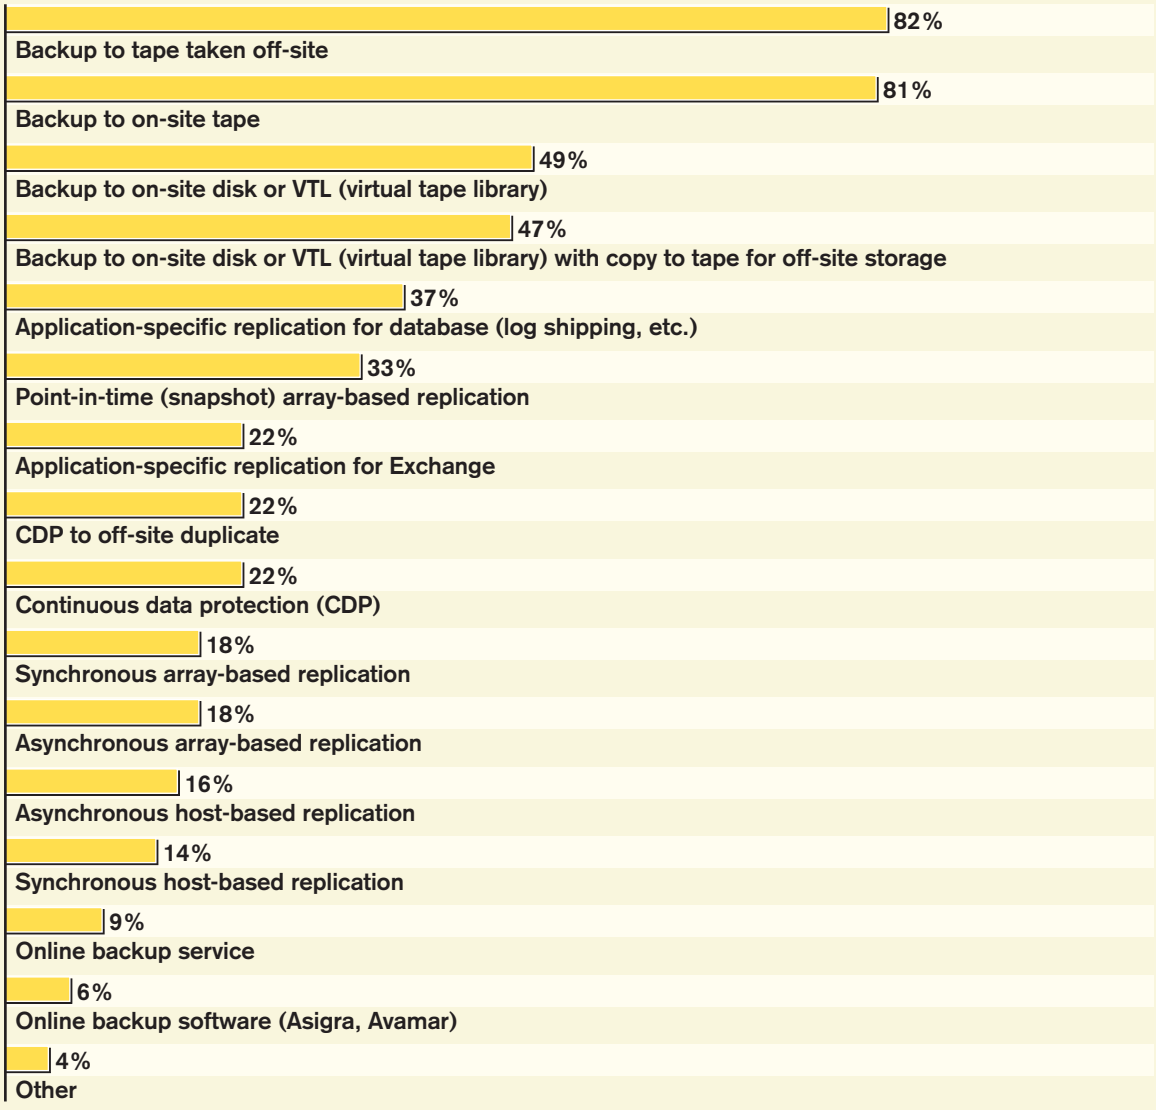

Note: Multiple responses allowed
Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Appendix

Figure 26

Plans Largely Unexecuted

Have you ever had to execute your business continuity/disaster recovery plan in whole or in part?

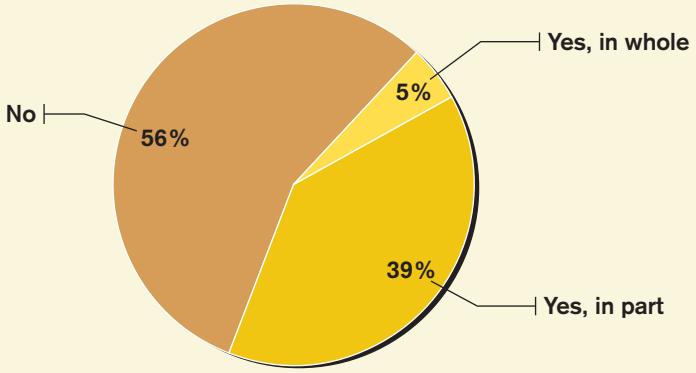

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 27

Recovery Plans Tested Infrequently

How often do you test your recovery plan?

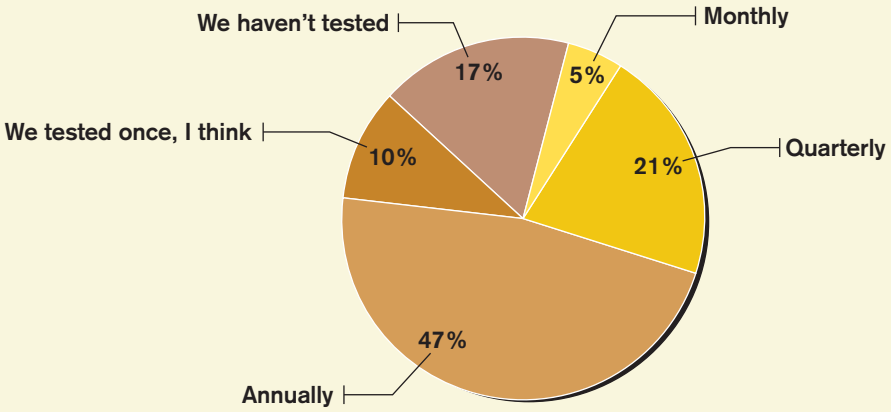

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 28

Time, Staff Limitations Hinder More Frequent Plan Testing

Why not test more often?

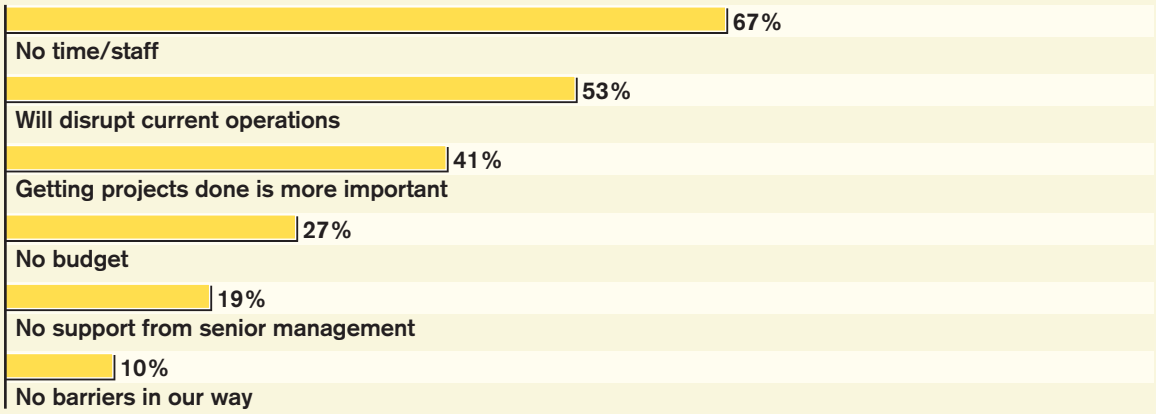

Note: Multiple responses allowed
Base: 94 respondents testing their plans quarterly or less frequently at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 29

Most Plans Updated On An Annual Basis

How often do you update your plan?

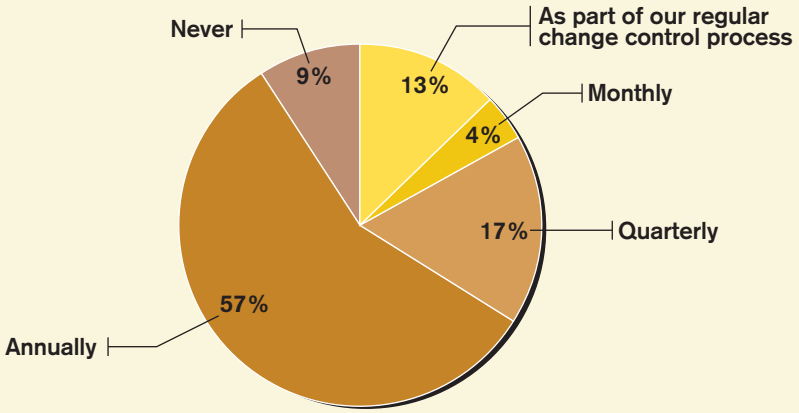

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 30

End-User Systems Less Likely To Be Included In BC/DR Plan

What parts of your IT infrastructure are covered by your BC/DR plan?

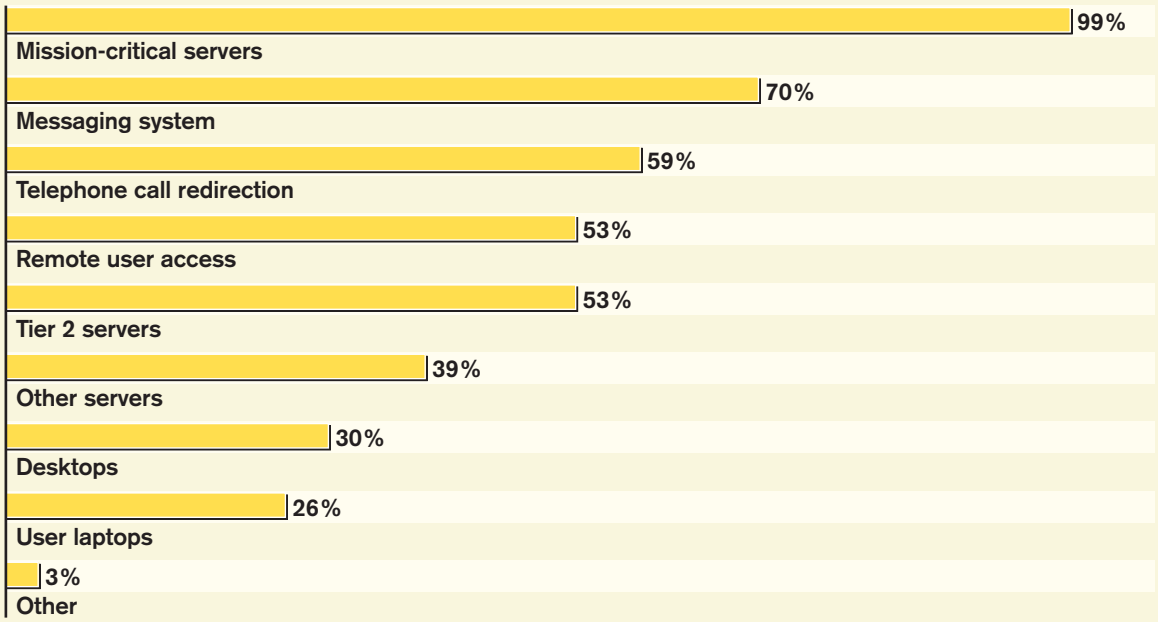

Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 31

Majority Manage Over A Terabyte of Data

How much data does your site manage?

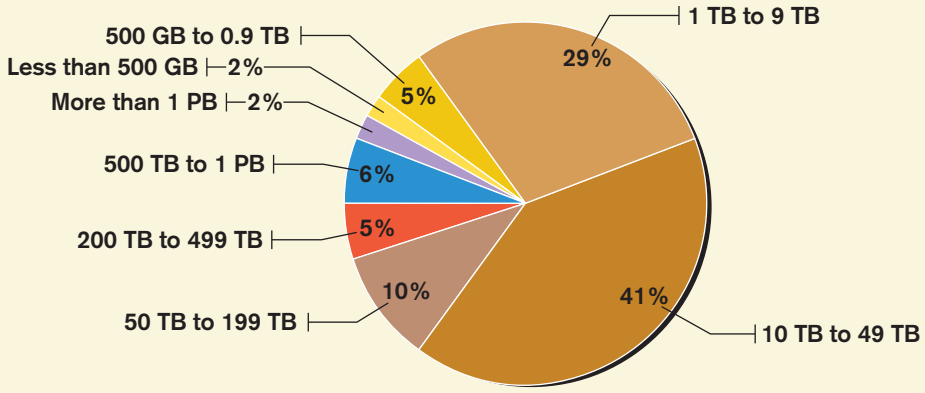

Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 32

Number Of Data Centers

How many data centers does your organization run?

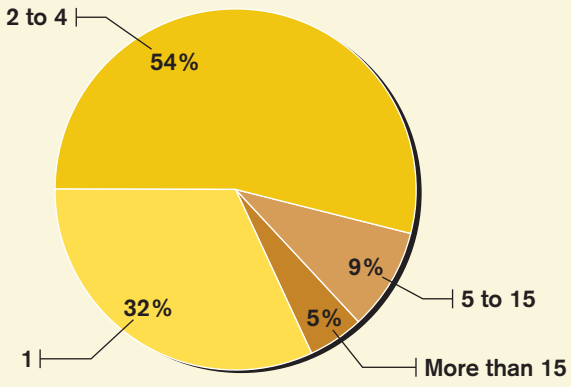

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 33

Organizations Mostly Centralized

Which of the following best describes the basic geographic structure of your organization?

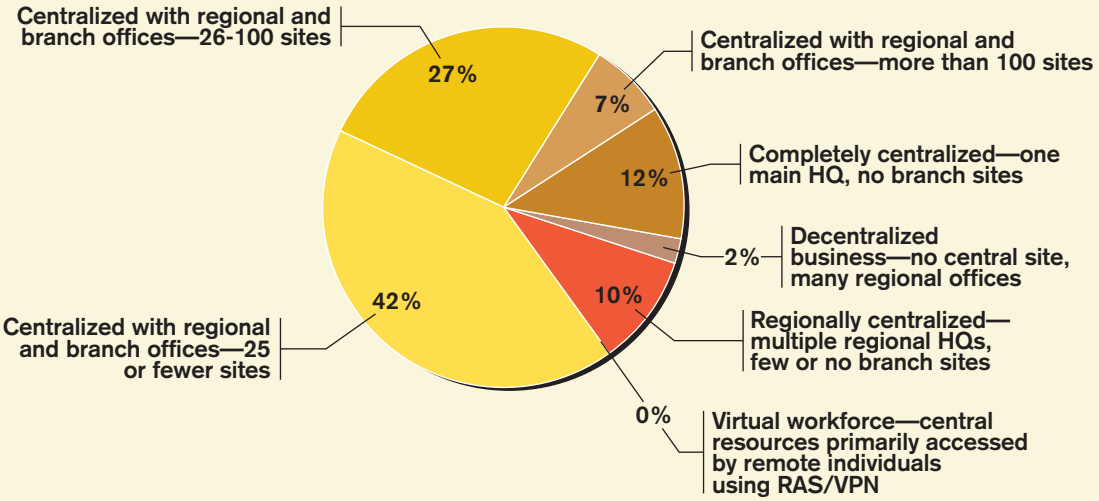

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 34

IBM Primary BC/DR Vendor Currently In Use

Which of the following vendors is your organization currently working with for business continuity/disaster recovery?

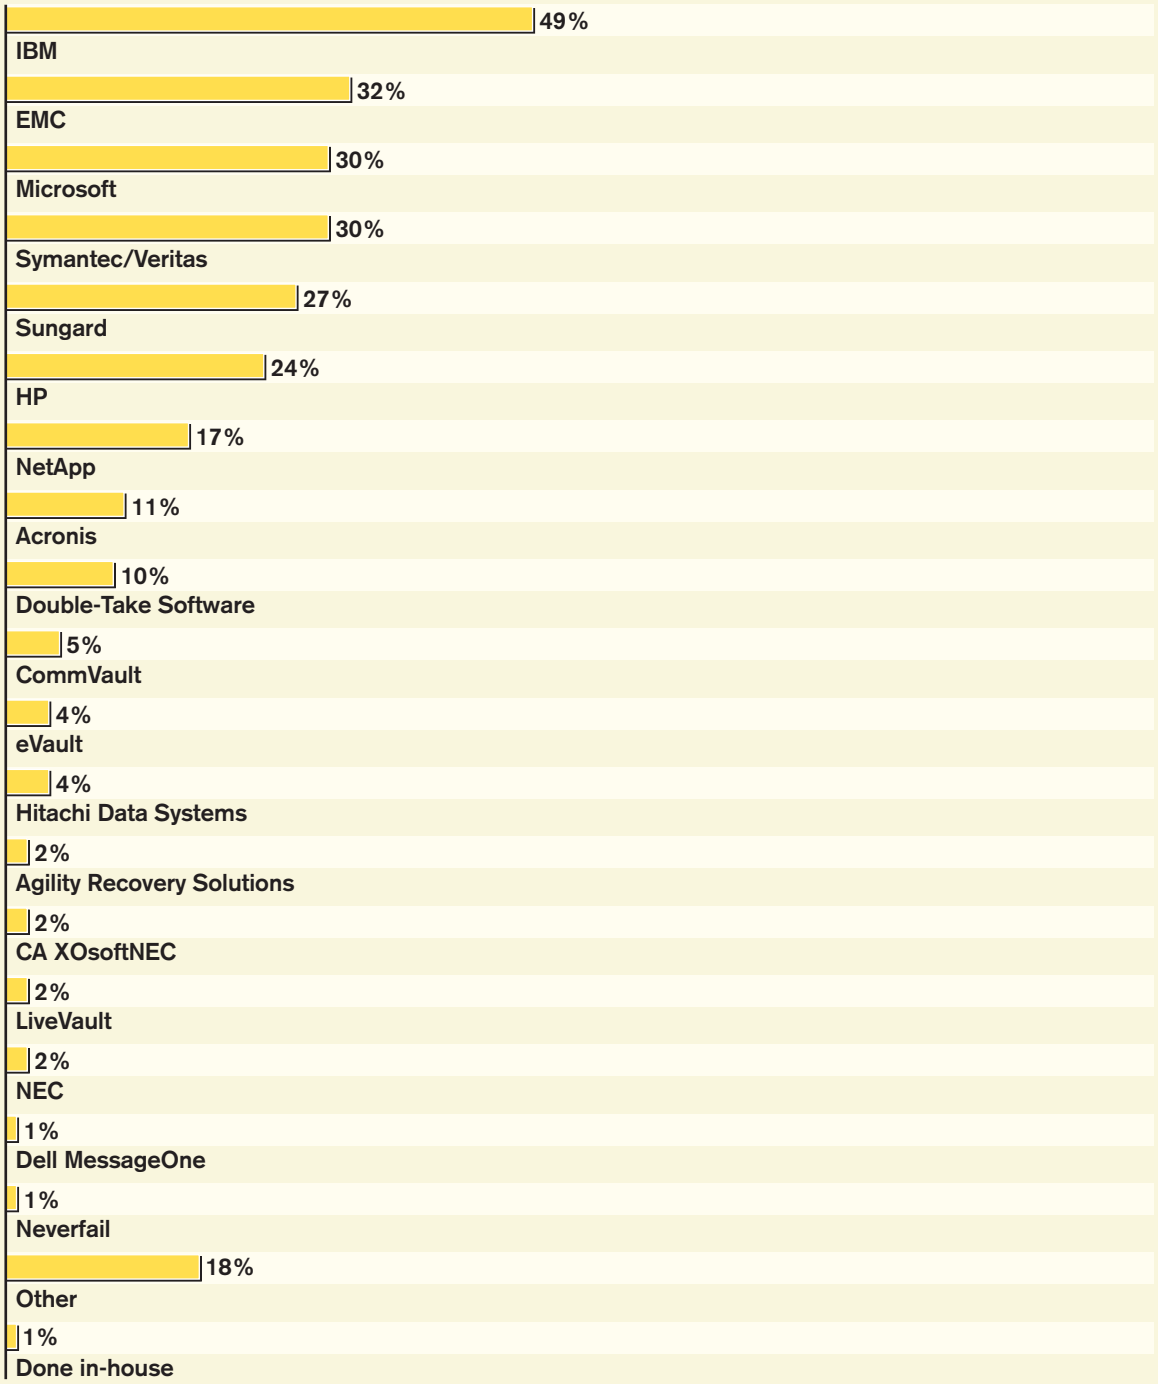

Note: Multiple responses allowed
Base: 99 respondents with a business continuity/disaster recovery plan in place at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 35

IBM Top Vendor Under Consideration For Future BC/DR Deployment

Which of the following would you consider as your top vendors for future business continuity/disaster recovery deployments (new projects or upgrades)?

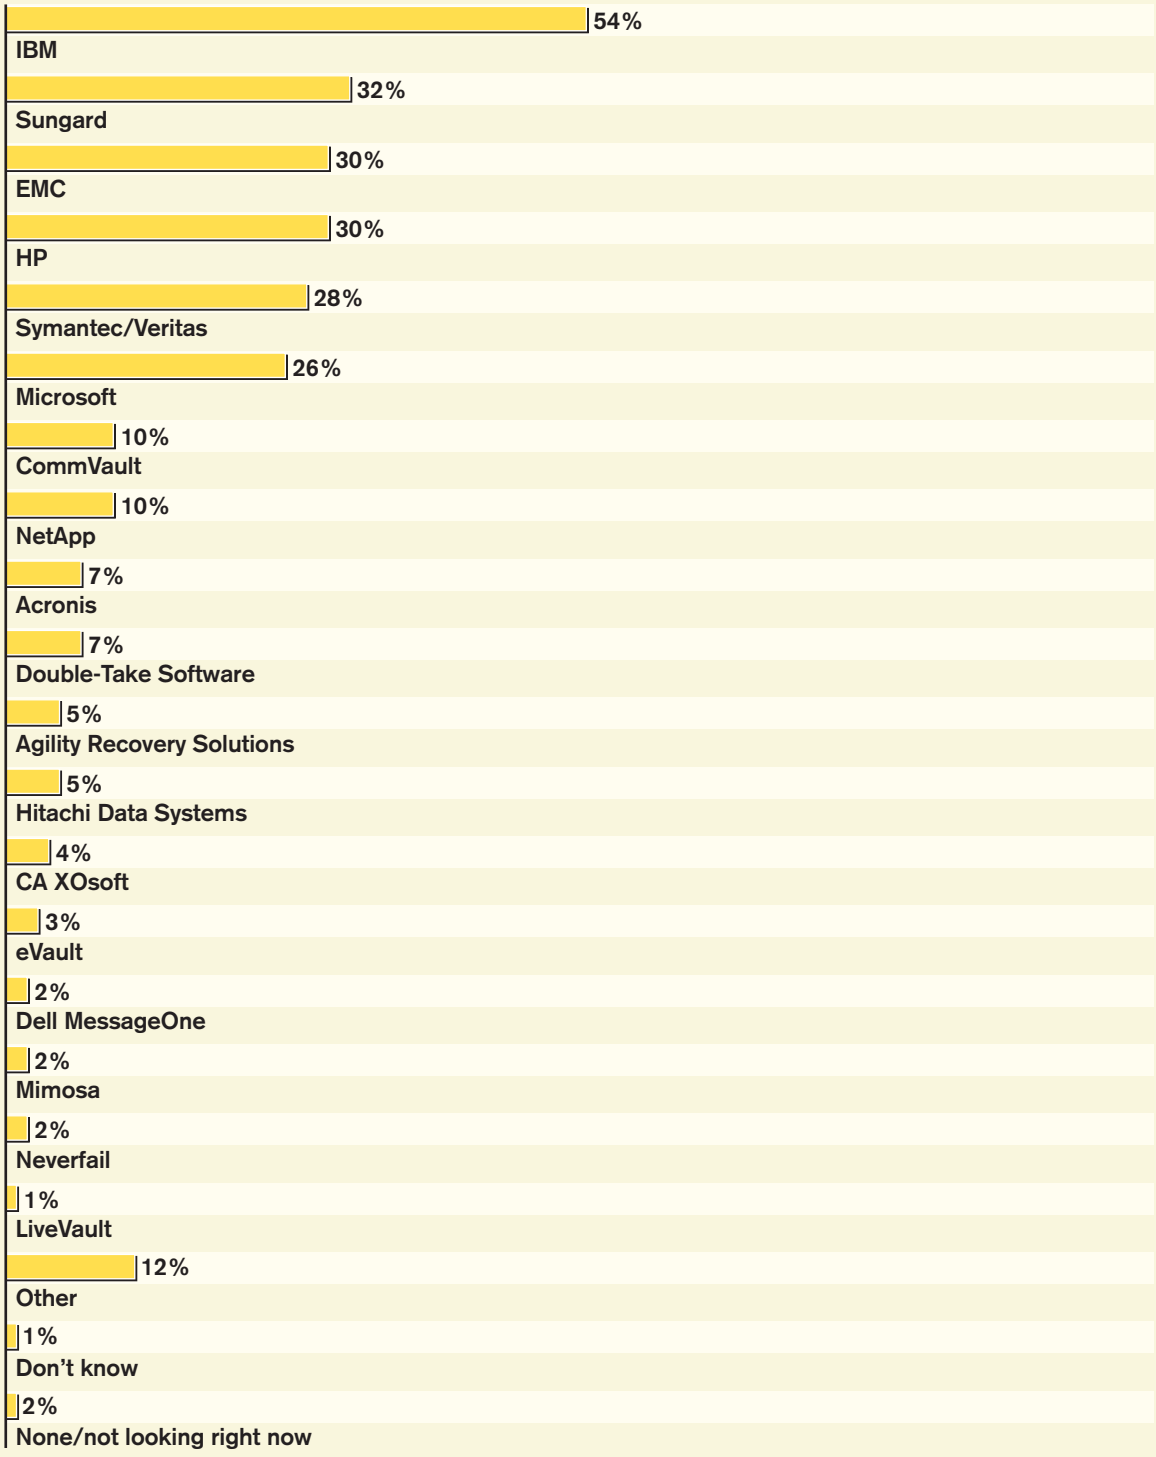

Note: Three responses allowed

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 36

Company Revenue

Which of the following dollar ranges includes the annual revenue of your entire organization?

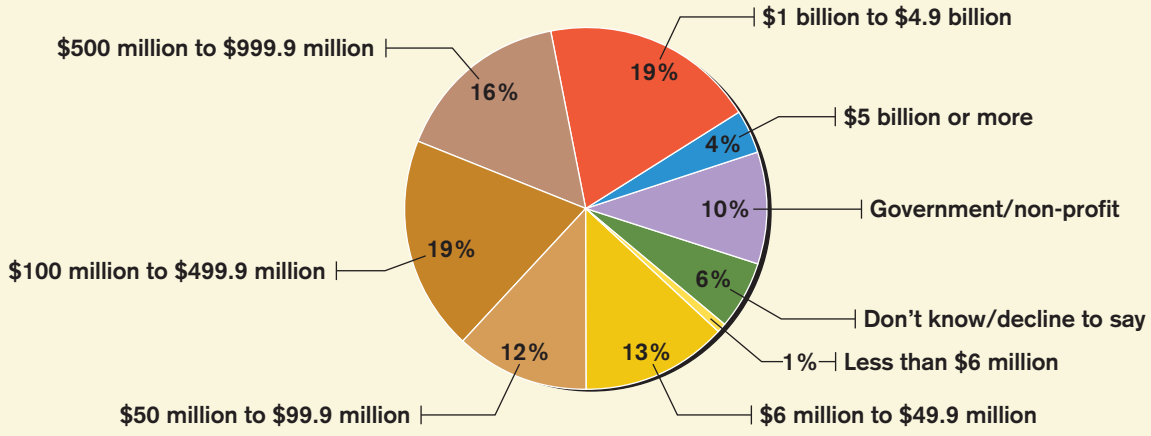

Base: 125 respondents at companies with 500-4,999 employees
Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 37

Company Size

Approximately how many employees are in your organization?

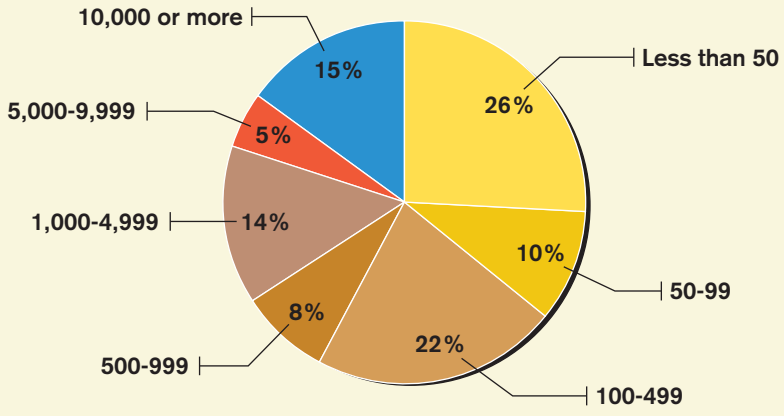

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 38

Job Title

Which of the following best describes your job title?

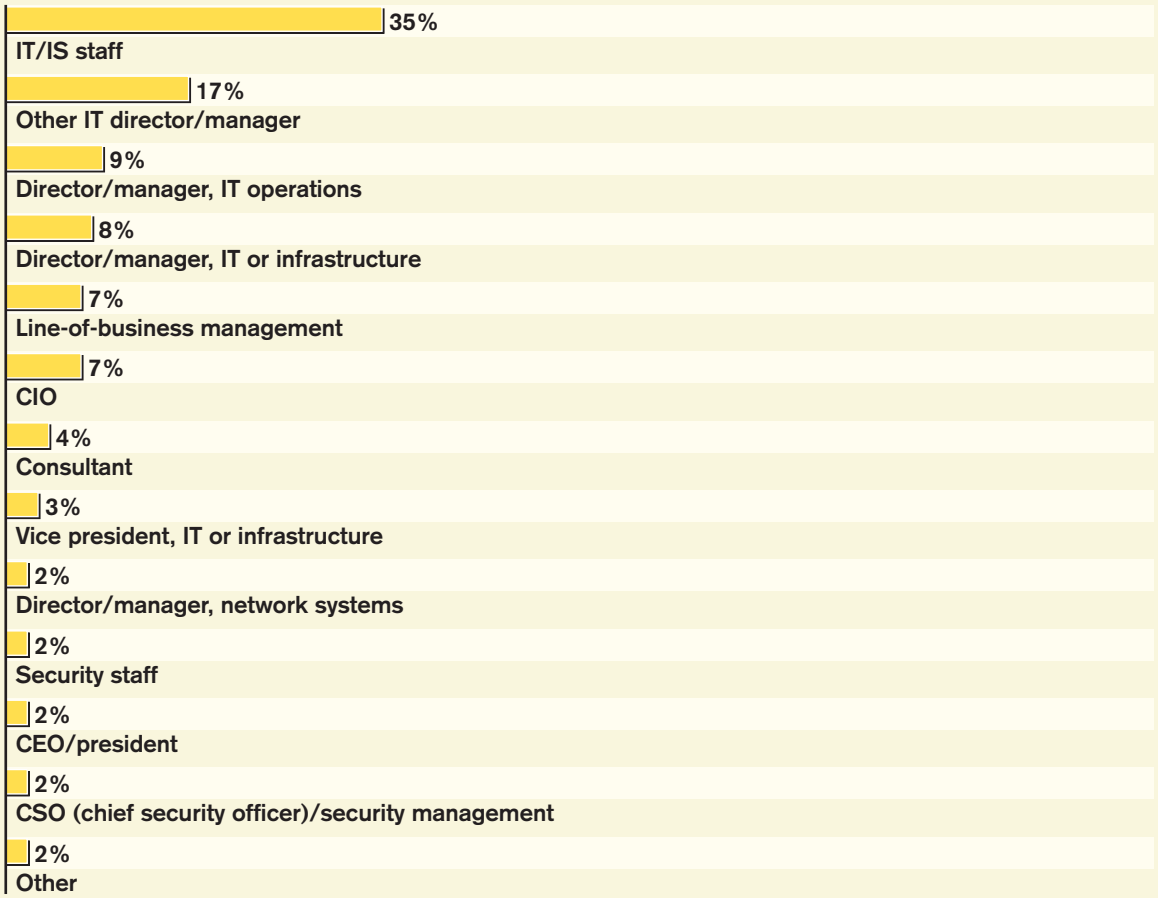

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals

Figure 39

Primary Industry

What is your organization's primary industry?

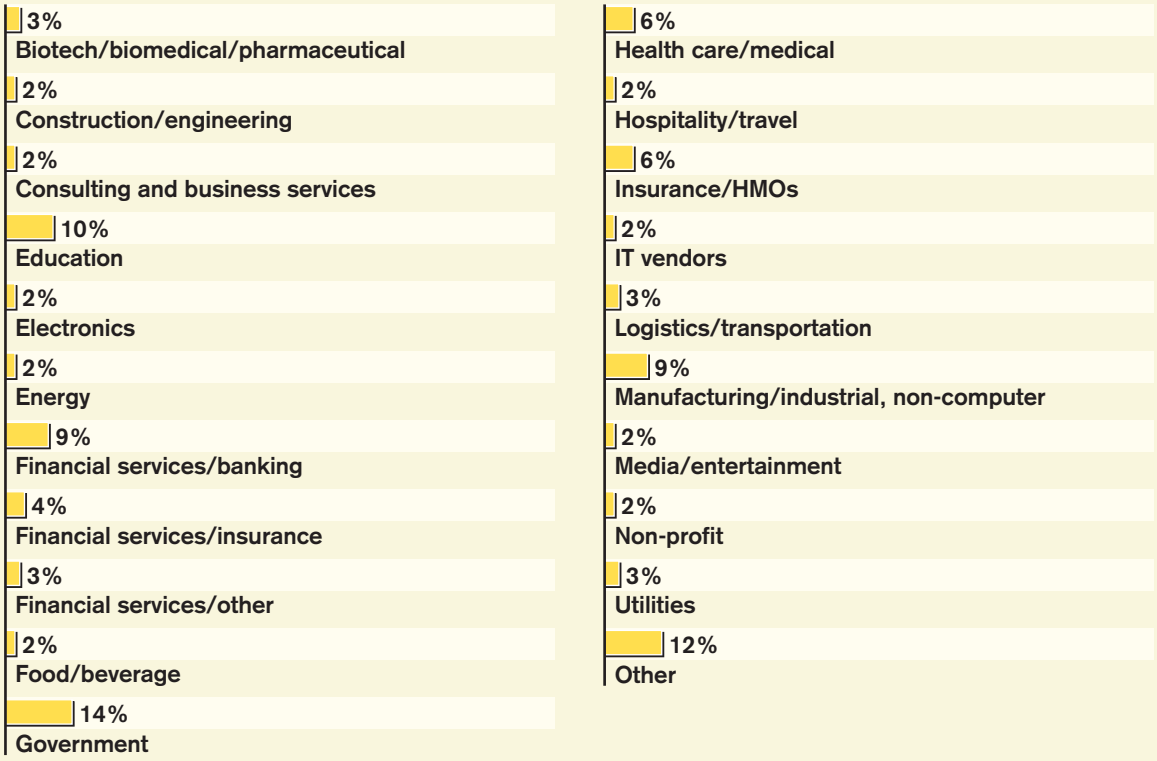

Base: 125 respondents at companies with 500-4,999 employees

Data: *InformationWeek* Analytics Business Continuity/Disaster Recovery Survey of 560 business technology professionals