

Information Technology

IT disaster

It's not a matter of if— it's a matter of when

By Gil Brodnitz, Gary A. Curtis and Robert Emmel

Complacency, complexity and strained legacy systems are conspiring to raise the risk of an unexpected IT disaster to alarming levels. Standard backup policies and redundant systems are no longer enough. What's needed now is a strategic emphasis on rapid and well-rehearsed recovery.

No one at the global bank panicked in the first few hours after a critical back-office system failed during the regular Monday night batch process. But the mood soon changed when it became clear that the bank's multiterabyte database was corrupted.

Attempts to switch to the hot offsite backup didn't work; the backup had mirrored the corruption. A real crisis was now unfolding. Applications teams and others suspended all other priorities as the four-hour target recovery window came and went with no sign of a root cause or a prompt fix. For more than a day, the

teams brainstormed, well aware that reacting prematurely could only make things worse.

Meanwhile, the bank kept trading. Just keeping up was a huge headache for the IT teams. They first had to find a clean backup. They'd discovered that the corruption had likely occurred nearly two days before the crash, and the only way to be sure that earlier copies were clean was to run a check that would take 36 hours.

Next, the teams had to update the production system, rerunning transaction log files to catch up to the crash point and processing

Today, the only safe course is to assume that in your employment lifetime, the worst will occur.

several days of transactions that had accumulated since then. Senior managers had to meet at all hours to make ad hoc decisions about which processes were required right away and which could be dropped temporarily to make up time.

By Friday, the teams had nearly caught up. But they still weren't sure that the bank could open for business by Monday. While the traders could keep trading, there was concern that it would be too risky to go more than five days without accurate settlement reconciliation. Regulators were alerted.

Although there's a happy ending (catch-up processing was completed over the weekend), the bank came dangerously close to having to suspend trading because of what turned out to be a tiny conflict between a packaged software bug and the server management software—something nobody could have foreseen, let alone tested for in advance. Losses were modest, but had the failure occurred at a different time—at year-end, say, when trading was running at full tilt as investors tidied their portfolios—the incident could have cost the bank millions, triggered regulatory investigations and damaged the institution's reputation.

Long story short: Although the bank had complied fully with internal policies and external regulations, and was ready for the loss of a data site or the failure of a major hardware component, it was ill-prepared to recover from a problem that wasn't mentioned in any standard operations manual.

Spotlight on recovery

An isolated emergency? Hardly. Similar near misses happen all the time in IT operations. Not long ago, point-of-sale transactions at a global retailer froze for 18 hours during the holiday shopping season because of a storage-network software bug that was never identified.

No amount of preparation could have prevented this crisis. And that is precisely the point: Most businesses remain exposed to serious disruption of key IT facilities for reasons that they can never adequately anticipate.

No matter that your IT team routinely analyzes risk and has testing regimes in place. No matter that you have built bulletproof disaster preparation plans. Nor does it matter that you have decentralized key IT facilities and spent millions on remote sites and that you refresh your business continuity plans every year. Today, the only safe course is to assume that in your employment lifetime, the worst will occur.

To be sure, prevention efforts are as vital as ever. But now CEOs and their IT leaders must put far more sweat into smart IT disaster-recovery strategies. The center of gravity of business-continuity efforts must shift quickly toward ensuring reliable recovery from any and all possible events.

A growing gap

In most large corporations, disaster-recovery and business-continuity planning activities have gained visibility and significant funding in recent years. Calamities including Hurricane Katrina and the September 11 attacks in the United States, Europe's killer heat wave and Indonesia's tsunami have sharpened prevention strategies at most businesses of any size. These companies now disperse risk across a range of geographically diverse, resilient networks of data centers, call centers, and operations and manufacturing facilities. (For a related article, see "Risky business," *Outlook*, May 2007.)

These enterprises undertake a broad-based, structured annual review of key business processes to understand how and to what degree the processes are vulnerable to a range of disaster scenarios. They take steps to mitigate the

critical risks and develop plans for continuing or, if necessary, recovering business operations.

More than ever, the conversation about impact assessment, risk mitigation, and business continuity and recoverability begins early in the design of any new process, system or operational location. Most big companies regularly test their plans and their backup centers. And most now have a chief risk officer or other executive responsible for corporate risk management, with staff or consultants who set standards, monitor compliance and report to stakeholders.

As improved as these practices are, they are no longer enough. There is a growing gap between what is and what should be.

Many large businesses are now so dependent on the flawless operation of their systems that they are dangerously vulnerable to substantial, even irreparable, business damage. Moreover, many enterprises have let regulators and other stakeholders direct the thrust of business-continuity efforts toward failure of one or another IT processing site or component. In the rush to comply, business leaders have lost sight of the other risks they face—the myriad smaller incidents that create big losses when servers shut down or decentralized software fails.

When disaster strikes, as it inevitably will, it doesn't matter what caused the problem. What matters—and it matters a lot to a CIO's career—is how quickly and reliably the problem can be resolved and how little business damage is done while the system is down.

The last war

There are two parts to the problem. First, many IT staff tasked with ensuring business continuity tend to react to the most recent interruption—fighting the last war—rather

than actively imagining and planning for rapid recovery from the next software conflict or security breach.

One large financial services firm is just now setting up a dedicated IT risk office—a defensive move that will help the company “check the boxes” in terms of compliance with the business-continuity standards required by regulators, but that will not do as much to help the firm bounce back when a major problem occurs. Senior managers are gathering data so that they can set expectations and prepare for what they expect—but that's all.

Second, when IT staff do plan for and rehearse disaster recovery, they lean toward the immediate and obvious catastrophes—the “loss of a building” scenarios that resonate at a very human (and very understandable) level with senior managers as well as with regulators. However, most real-world problems, as evidenced by our story of the bank and its unseen software bug, are much less spectacular—at least, to begin with (see chart, page 4). And although these problems may not necessarily paralyze entire corporations and will not reach crisis proportions overnight, they can degrade IT performance to the point where there is a substantial impact on the business—not the least of which is considerable management distraction.

The degradation can begin imperceptibly and not necessarily because any one IT component breaks down. Recently, breathtaking expansion at one leading communications service provider had impressed investors but caused significant strains in some areas of the business. Specifically, the company's systems architecture was under severe pressure, with unplanned downtime reaching alarming levels. Some data centers were running at 99 percent of their total capacity; one of every eight backups was failing. IT staff were responding later and later to system

IT staff often cannot deliver the business-continuity levels they have committed to.

alerts. And “red downtime”—when failures were directly affecting customers—was dangerously high.

Fortunately, the problems were solvable. But the point is that they had seemingly and suddenly emerged from nowhere.

Whether the root cause is small and gradual or sudden and significant, the implication for business leaders is the same: IT staff often cannot deliver the business-continuity levels they have committed to. They may be able to keep their promises about recovery from hardware failures, but not so with software problems. In essence, they’re telling top management: “Trust me.”

Escalation

Don’t expect these challenges to lessen anytime soon. In fact, the potential for disruption in IT systems is escalating. Why? Because in many cases, critical IT staff members are


getting closer to retirement; heavily patched legacy systems are still very much in place; and new architectures (such as service-oriented architecture, or SOA), while full of genuine promise for improving the effectiveness of systems development, are inherently complex and involve many layers of new software.

Attendees at a recent conference for IT leaders noted that by 2015, nearly 20 percent of their staff will be 55 or older—up from less than 11 percent during the 1990s. Replacing staff members won’t solve the problem—decades of business and application knowledge will walk out the door as baby boomers retire.

At the same time, many companies are spending some 80 percent of their IT budgets to keep existing applications and their processing components operational, leaving precious little to fund new development efforts. And IT departments are quickly automating many of their

Balancing recovery strategies against potential losses over time

Most business interruptions and extended outages are caused by human interference.



internal support processes, burdening the systems' reliability and exposing companies to the likelihood that there will be nobody with the depth and breadth of experience to troubleshoot effectively.

Things get even more complex when what were once monolithic business applications supporting a single business process are now typically layered composites built from packages and custom applications often developed by many different vendors. Things are really coming to a boil now that architectures such as SOA are gathering momentum, increasing the likelihood that software bugs will surface and that timing-specific software interactions will occur in ways that the best test engineers may never be able to reproduce—and which, in some cases, may not ever reoccur.

The upshot of those converging factors? There's more downtime in your

future—and less likelihood that you'll quickly find what caused it. If CEOs aren't asking tough questions of their CIOs on these issues, there could soon come a time when shareholders and boards of directors are asking the tough questions of *them* (see sidebar, below).

A new mindset

Some companies, trying to get in front of this problem, are pioneering new approaches to recovery. Working with them, Accenture has found seven critical points common to all of these approaches. Each requires a shift in mindset for IT leaders but not a major capital investment. Collectively, they comprise a useful starting point for more detailed business-continuity strategy and action.

1. Start talking business value—and business risk.

We consistently see a general reluctance to talk about business value at

Thirteen questions that CEOs need to ask their CIOs

Sound disaster-recovery planning has to be on the agenda in the boardroom (see story). But the CEO can't give directors conclusive answers without having first talked at length with the CIO. Here are some of the most pressing questions.

1. Tell me about our response simulation and rehearsal plans and activities. When was the last time we had a full-scale rehearsal of an IT disaster recovery?
2. What did we learn from it?
3. How do we learn from others' business-continuity mistakes?
4. How will our recovery plan help the company financially?
5. Have our recovery planning activities made our company more resilient?
6. How can management know how quickly we're responding in a real emergency?
7. What kind of event-monitoring system have we set up to give us some early warning so we don't have to invoke our emergency plans?
8. Who's accountable for IT disaster recovery?
9. How can we be sure our people are trained to respond effectively?
10. What other resources do we have for recovery other than our own staff?
11. I understand we're prepared for hardware failure, but how prepared are we for a large-scale virus or malware attack?
12. What kinds of automated response capabilities do we have to rapidly communicate status and begin response implementation?
13. Do our recovery plans extend to business-support capabilities as well as technology capabilities?

risk—to imagine (if not actually identify) which customers would be affected and how by a system slow-down or shutdown. Much as the population at large avoids detailed discussions about death, IT people in particular are uncomfortable about asking business users what they would lose if specific IT processes were unavailable or severely compromised.

This topic cannot be taboo. Any recovery strategy needs to be considered in light of the resources required—cost included (see sidebar, page 8). The speed at which business operations can be recovered, either in-house or with a third party, is directly related to the willingness to allocate resources to a specific recovery strategy (see chart, page 7).

2. Play more war games.

Once the bank in our story had detected the problem and restored the database, it had to figure out how to quickly process several days of transactions. If its staff had rehearsed a recovery from the storage failure—regardless of cause—they probably would not have had to make ad hoc decisions that involved waking senior managers in the middle of the night and reaching out to vendors half a world away.

In most organizations, scenario planning isn't pushed far or fast enough. At the bank we've described, the technologists knew the databases were at a statistically small risk of corruption and took comfort in the fact that they had built a good backup system. In the midst of the event, however, they discovered that they were not adequately prepared to manage the recovery because they had not rehearsed this scenario. As a result, because complex business decisions that could have been made in advance weren't, the recovery took far longer than imagined.

It costs relatively little to conduct war games—simulations of recovery sce-

narios. They help break the compliance mentality and make visible the value of recovery planning and management. That helps IT leaders build a culture of personal accountability where the buck does indeed stop somewhere. War games also help clarify key roles and responsibilities and chains of command. And they help to remove the social penalties for being the bearer of bad news.

The goal of every simulation should be the fastest and most reliable recovery from the worst possible disruptions—a rapid resumption of normal operations with the least possible impact on customers, revenue, cost and time. Note the emphasis on reliability: It's better to recover reliably in one day from any conceivable failure than to plan to recover in four hours but only from a subset of possible causes.

3. Stay in constant “debrief” mode.

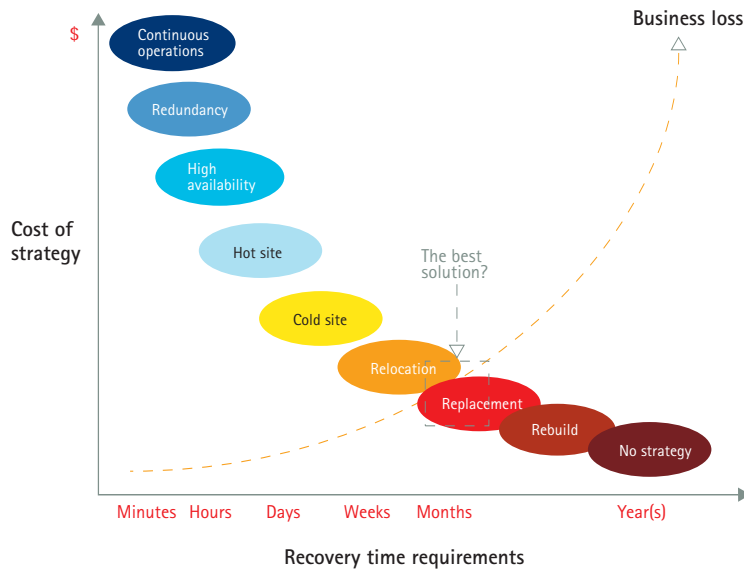
Groups of IT professionals should have processes in place for comprehensively debriefing every time there is a near miss in any relevant company. Only by dissecting problems at others' facilities as well as their own can they capture and catalog the key learnings. IT groups should have designated leadership roles for such knowledge capture. And ideally, they should also have access to knowledge-brokering mechanisms for integrating outside knowledge and third-party perspectives.

4. Appoint an IT risk ombudsman.

While it is important to have a chief risk officer or someone with executive responsibility for risk, some organizations also find it worth considering the role of IT risk ombudsman—a respected senior manager to whom IT staff can raise concerns without fear of personal exposure. Not only will the ombudsman be a problem spotter without agendas or affiliations, he or she will also be a veteran technologist with a deep understanding of the whole IT architecture.

Business-recovery strategies versus business loss over time

Selecting the most appropriate recovery strategy for mission-critical operations involves the trade-off between time and money. (This graph is for discussion purposes only, and is not to scale.)



Recovery strategy	Definition
Continuous operations	Real-time, redundant processing infrastructure (with some geographic separation from the primary location) that is an exact copy of the primary processing infrastructure's capabilities.
Redundancy	Duplication (not necessarily in real time) of information processing and data storage resources.
High availability	Provides data redundancy on a real-time basis at a geographically remote location.
Hot site	A vendor or business-owned remote facility that is equipped to support information processing capabilities in the event a business data center is unavailable.
Cold site	A physical location, other than the primary information processing facility, that contains pre-installed networking, telecommunication and electrical capabilities to allow the installation and connection of alternate processing resources.
Relocation	The movement and installation of all current information processing capabilities to another geographical location.
Replacement	Replace equipment and processing capabilities.
Rebuild	Complete rebuild of a business information processing capability, facility and environment.
No strategy	Recovery capabilities will be determined after a disaster has occurred.

Source: Accenture analysis

The ombudsman has to be an owner of the recovery philosophy, someone whose focus is on retention of business value, and not simply a business-continuity manager. This individual can also be the sponsor of a comprehensive diagnostic that

paves the way for a broad recovery-management initiative.

5. Rethink robustness

Robustness in terms of IT-recovery performance is not simply a product of how many levels of backup there

are. It must also be thought of in terms of labor availability and partner capabilities—almost in “bio-diversity” terms. At the headquarters of a top credit-card processor, call centers were shut down after a hurricane because staff had no access to clean water. Still, a shutdown was averted because the company was able to scramble to shift call volumes to its outsourced centers.

6. “De-average” the data

It’s not prudent to work with average statistics for response times or downtime probabilities. The customer doesn’t care about the average when systems are compromised by downtime that shows up at the worst possible moment. By using averages, you’re effectively saying that if it’s not likely to happen, it’s okay to be poorly prepared. The lesson: IT leaders must deal in bands of risk with clear indications of highs and lows.

7. Fix the whole thing, not just elements in isolation

We have often observed that IT’s typical approach after a disruption is

to try to rebuild processing capabilities application by application and then to wonder why this process continually fails to recover the data center in a timely manner. (The issue is perpetuated in disaster-recovery testing scenarios.) But every application maps to at least one technology platform and to related, usually integrated, applications. They need to be identified and fixed in concert with one another.

Expanding the focus to smart recovery is a never-ending challenge. As long as the business processes, organizations, applications and infrastructure keep growing and evolving, new failure and recovery scenarios will emerge. An effective recovery strategy calls for a permanent shift in mindset away from compliance and complacency and toward a heightened sense of readiness. Although it doesn’t mean behaving as if every moment could be your last, it does call for the kind of alertness to imminent danger and preparation to recover from it that is shared by few who aren’t emergency services personnel.

Relating risk tolerance to recovery speed

Development of a recovery plan calls for an accurate accounting of risk types as well as an understanding of their level of acceptance and potential impact on the business. Four practical factors deserve a mention.

- The speed with which business operations can be recovered, either in-house or with a third party, is directly related to the willingness to allocate resources to a specific recovery strategy.
- A particular recovery strategy should be selected based upon business needs, not solely on technical and/or equipment manufacturers’ capabilities or third-party hot-site vendors’ recommendations.
- When mapping business losses against the costs of recovery, the point at which the lines intersect may not necessarily represent the most prudent overall recovery strategy. In other words, the mathematical result is not always the best answer.
- Supporting the chosen recovery strategy must come with an understanding of which resource will be traded off: time or money.

Recovery done right also presupposes that the right coordinated operating models and business processes are in place—or are being put in place. Unless there are designated roles, and unless staff are being trained in and are practicing comprehensive recovery scenarios, downtime minutes will turn into hours and hours into days, with the growing likelihood that customers will be affected—and will care.

There can be no better reason for CEOs to act.

About the authors

Gil Brodnitz is a senior executive in the Accenture Strategic IT Effectiveness practice. For more than 15 years, Mr. Brodnitz has served senior clients at major financial institutions, with a focus on creating closer integration and alignment between business strategy and the use of information technology. He is the former chief e-commerce officer and CIO of a leading pharmaceuticals marketing services provider. Mr. Brodnitz is based in Washington, D.C.

gil.brodnitz@accenture.com

Gary A. Curtis is the global co-lead for Accenture Technology Consulting. Mr. Curtis has served the top managements of several global investment banks, major media companies and high-technology providers for more than 25 years. He specializes in evaluating the business value of large-scale IT applications and infrastructure portfolios, as well as in creating programs to improve that value over time. Based in San Francisco, Mr. Curtis serves on the advisory boards of several companies that are developing new technologies.

gary.a.curtis@accenture.com

Robert Emmel is the Chicago-based global lead for the Accenture Business Continuity Planning practice. With more than 20 years of experience supporting and leading a broad range of business, risk management and information technology service and operational processes, he has participated in or headed more than 80 recovery planning engagements, 30 data center assessment and enhancement projects, and numerous data center consolidation/migration efforts. Mr. Emmel's work on recovery planning, risk management and data center operational support issues has been published in many business and IT periodicals.

robert.emmel@accenture.com

Outlook is published by Accenture.
© 2008 Accenture.
All rights reserved.

The views and opinions in this article should not be viewed as professional advice with respect to your business.

Accenture, its logo, and High Performance Delivered are trademarks of Accenture.

The use herein of trademarks that may be owned by others is not an assertion of ownership of such trademarks by Accenture nor intended to imply an association between Accenture and the lawful owners of such trademarks.

For more information about Accenture, please visit www.accenture.com