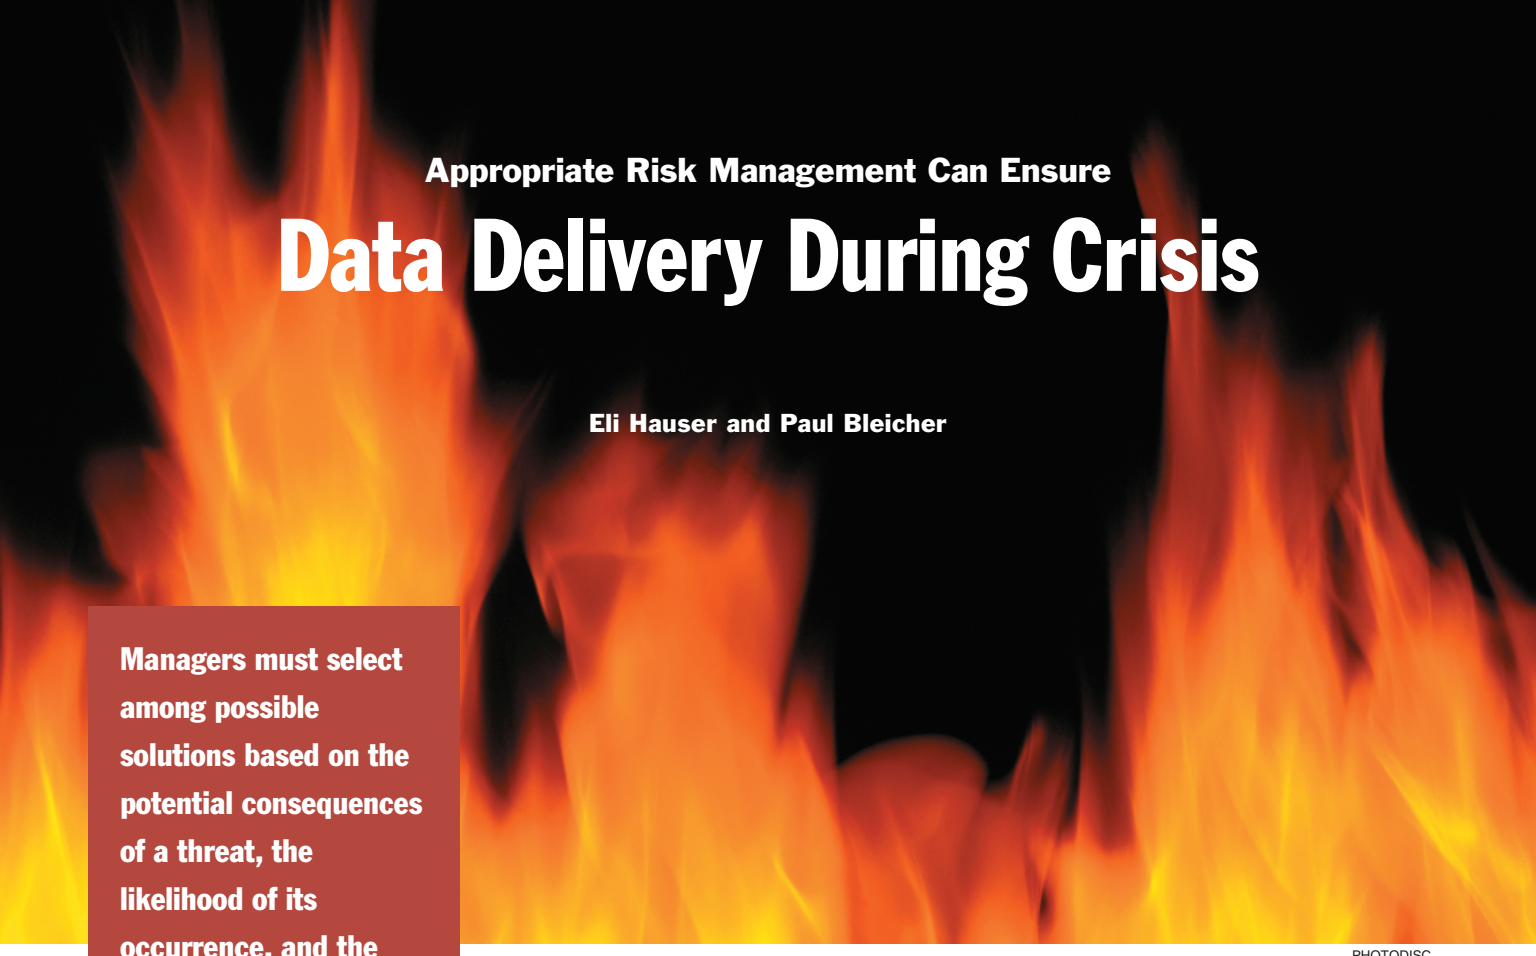

PHOTODISC

Appropriate Risk Management Can Ensure

Data Delivery During Crisis

Eli Hauser and Paul Bleicher

Managers must select among possible solutions based on the potential consequences of a threat, the likelihood of its occurrence, and the risk tolerance of the business component.

Much of the financial infrastructure of the United States, if not the world, operated in or near the two towers of New York's World Trade Center. Despite the near paralysis that resulted from the events of 11 September 2001, many businesses that had operated there were up and running the next day. They were able to do so only because many of them, such as Morgan Stanley, had contingency plans that allowed them to replicate or "mirror" their databases and continue operations in new locations across the Hudson River. The Gartner Group states that two of five companies that experience a disaster of that magnitude are out of business within five years. When disaster strikes, proper planning dictates who succeeds and who fails.

During the ensuing weeks, many clinical researchers asked themselves, "Could our business have resumed operations in 24 hours? What if our clinical trial database was destroyed? Would we have a backup? How long would such a loss delay our next blockbuster? How much would it cost to make sure that could never happen to us?" The answers to those questions lie in each company's analysis of risk, benefit, and cost and how they plan for and implement good risk management practices for business and data continuity.

Critical decisions

Two caveats apply to risk management. First, there is no limit to the amount a company can spend on it. Second, no matter how much a company spends, it can never mitigate all risk. Therefore, corporate executives must establish priorities to make such critical investment decisions.

Given the wide variety of competing alternatives, the best way to analyze those issues is through a framework structured to organize, stratify, and systematically identify threats, assess consequences, and guide choices. Such a framework can also help companies develop and use an exhaustive checklist of considerations, ensuring that they identify and review a broad spectrum of issues. In working through alternatives, the business continuity team can evaluate the company's business, organizational, and technological variables to find the best solutions.

Although risk-management issues are similar in most industries, pharmaceutical companies must comply with unique regulations that affect business continuity. Fortunately, the devices, techniques, and disciplines needed for data center continuity complement those regulations and guidances, including 21 CFR 11 and the International Conference on Harmonisation. For example, 21 CFR 11 and FDA's Guidance on Computerized Systems Used in Clinical Trials seek to ensure system security, data integrity, and dependability in systems used to comply with FDA regulations. A key component of regulations covering electronic records is the documentation of written procedures and management controls; both are also essential elements of a reliable data center. When designing risk-management programs, pharmaceutical company IT managers must satisfy both the company's business continuity needs and the relevant regulatory requirements.

Threat factor

A simple way to understand risk is through the equation: risk = threat × vulnerability × cost. Threat is simply the frequency with which an event occurs, which

can be quantitatively estimated. Experience, of one's own or from similar businesses, can help determine that frequency. Depending on the threat, companies must consider several factors, especially the business location. Vulnerability is measured by the likelihood that a threat to a specific data center or business will occur. That assessment may be highly quantitative or very subjective, depending on the circumstances. Cost estimates, too, can be both quantitative and qualitative. Some threats, such as the cost of a data center shutdown, are highly quantifiable. Others, such as the cost of not following best practices, can be highly subjective. The assessment team needs to be comfortable with that wide variation in approach and perceived accuracy.

Obviously, a serious threat with both a high cost and a high likelihood of occurrence, such as an earthquake in San Francisco, is a significant risk for any company. Another threat, such as a typhoon, can be costly, but because it is not likely to occur in the suburbs of Chicago, it would be a low risk for a site there. However, a fairly unlikely event with a huge cost may be a greater risk than imagined because both cost and the likelihood of the occurrence are factors in determining threat.

Therefore, companies embarking on a risk management program must first develop an itemized list of risk elements by category and rank them by the likelihood of threat success, indicating the cost of the damage for each (see Risk Assessment box).

Identifying threats is a significant exercise and must be well planned. The simplest scenarios—such as a catastrophic fire that destroys a single facility—are easy to document. In the aftermath of September 11, companies must also consider the possibility of several crises occurring at the same time, such as the destruction of a facility along with simultaneous loss of voice and data communications, physical plants, and an off-site backup storage facility.

In search of solutions

The next step in building a data/business continuity program is to develop a plan of action for each identified risk, thereby modifying the environment to the desired risk profile. Responses and selections should be based on cost, consequence, and likelihood of occurrence. Often, a single remedy will address multiple risk factors, so a matrix showing risks vis-a-vis remedies is appropriate.

Armed with a range of possible responses, companies must choose those that are most appropriate for each situation. As the likelihood of an adverse event increases or as the magnitude of its consequence increases, three response models come into play: monitor, manage, and eliminate. Not all risks are created equal. The management team must select among a continuum of possible solutions, matched against the consequence of a threat, the likelihood of its occurrence, and the risk tolerance of the business component.

Multiple power supplies may eliminate the possibility of catastrophic power failure in a data center. On the other hand, a company may choose to manage the possibility of a damaged server hard disk through rapid response and data restoration from a backup. That is a less costly alternative to maintaining a continuously replicated hard disk, which would eliminate the possibility of down time from hard disk loss. And some threats may simply require careful monitoring. A good example of that is the percentage of network bandwidth that the data center is using. If the percentage gets too high, the company can add more bandwidth.

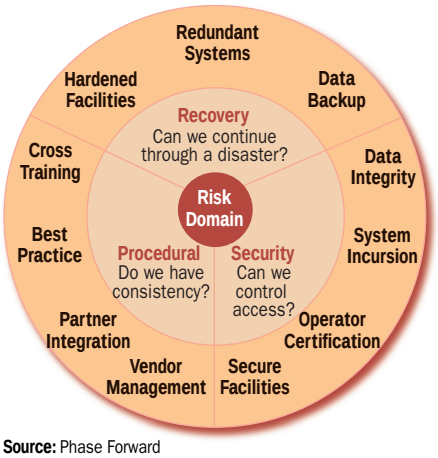

Source: Phase Forward

Risk domain. A large number of disciplines are needed to create a cohesive risk management program. Different risk categories involve different skills, disciplines, and remedies.

Risks and responses

For each risk category (recovery, security, and procedures), the data/business continuity team must develop a plan for computer application support, staff, and computer operations. (See "Risk Domain.")

A remedy in one category may also address an issue in another. Operating within a bonded, formally designed data center, for example, reduces security risks and, at the same time, provides much of the infrastructure required for a highly reliable facility. The choice of a particular remedy to address a risk in one category may influence and possibly increase requirements in another. Thus, the decision to establish multiple facilities to manage the risk of losing a primary facility demands not only a set of procedures to support both the original facility and its backup, but also complementary security measures to guard the overall larger infrastructure.

Recovery risks. On a macro scale, recovery risks deal with physical disasters such as earthquakes, hurricanes, blackouts, and explosions. On a micro level, they involve much smaller events like hard disk crashes. Recovery risks are implied in the question: "Can we continue operations through and after a disaster?"

The core recovery element is well

Risk Assessment

A stratified risk management approach requires a staged sequence of analyses to identify, categorize, and rank each risk. Each stage of the assessment and selection process calls for different skills and personnel.

Recongnize type of threat

- Recovery (loss of data, machines, services, or facilities)
- Security (hackers, fraud, or inadvertent deletion)
- Procedures (staff qualification, training, and adherence to procedures)

Assess Consequences

- Likelihood of occurrence
- Cost of damage

TABLE 1 Recovery risks and solutions[a]

System	Failure Mode	Solution Set
Availability	CPU failure	Dual CPU
	Disk failure	Disk mirroring
	Network failure	Network diversity
	Dataset corruption	Data integrity monitoring
Recovery	Server destruction	Hot switch backup
	Database destruction	Offsite data backup
Facility		
Availability	Building access failure	Short- and long-term building contingency plan
	Internet access failure	Backup internet access
	Phone failure	Phone systems
	Power failure	UPS, diesel engines
Recovery	Earthquake	Reinforced facilities
	Fire or water damage	Environmental monitoring
		Short- and long-term building contingency plan

[a]*Recovery risks are among the most tangible and easily understood.*

TABLE 2 Incursion/Integrity checklist[a]

Data	Failure Mode	Solution Set
Incursion	Authentication failure	Password management
	Transmission failure	Encryption
	Database access failure	Monitoring
Integrity	Inappropriate data access	User registry, rights, log-in
	Database corruption or destruction	Database integrity monitoring
System		
Incursion	Internet access failure	Firewall management, dial-in VPN
	Physical access failure	
	Internal employee access failure	Passkey systems
Integrity	Database corruption or destruction	Hot switch backup
	Operating environment compromise	Certified operators
		Minimal operator rights

[a]*Security risks typically represent the most abstract, amorphous and disconcerting set of risks to most people.*

understood: the creation of an environment designed to recover from a loss of data, machines, services, or facilities within a given timeframe and with a minimal amount of data or connectivity loss. A wide range of alternative strategies, with a primary focus on technology, is available at each level of the system architecture.

Within the recovery risk arena, the planning team can generate a matrix that maps different types of conditions and failure modes. They can separately identify the problems associated with a system failure and a facility failure, and map those against two types of failure-mode availability, representing an access failure, and recovery, representing a physical failure. They then can map each of those into sets of potential solutions (Table 1).

Security risk. The second category is represented by the question: "Can we control access to data?" In examining both the physical and electronic types of control for both data and system access, it is worthwhile to differentiate between incursion (constraining access) and integrity (data or system manipulation from within) failure modes (Table 2).

Security risks include hackers, fraud, and inadvertent deletion, and each requires a different detection mechanism and response. Often discussions about security risks focus on one type of risk, treat all risks the same, or mistakenly address all with one solution. It is important to clearly distinguish between alternative failure modes and to carefully match each solution to its intended use. A great firewall will not prevent an authorized person from changing data. Thus, to develop a complete solution set, planning teams must identify each threat and discuss it within its specific context.

Procedural risk. The third component of the risk-management matrix involves "softer" elements such as staff qualification, training, and adherence to procedures. Those personnel management issues play a critical role in creating or closing security holes and are part of procedural risk. The question companies must ask is: "Does our staff comply with consistent business practices internally and with outside vendors and partners?"

To address procedural risks, companies should design a resource and management strategy to

- minimize the individual roles of key staff
- maximize full procedural training and cross training
- eliminate undocumented, informal, or non-repeatable processes.

Writing and maintaining up-to-date standard operating procedures, adopting best practices, and creating a culture of process adherence, documentation, and auditing are key to managing that risk area.

Because personnel are involved, adherence to procedures as well as training and sustaining the appropriate level of control can be difficult. The risk plan should include the funding and attention to the quality infrastructure required to develop, monitor, and maintain compliance. To accomplish that, companies should consider increased management oversight and control in the form of process and documentation audits and reviews (Table 3).

Solution planning

To prepare a response profile, companies must map each risk or group of risks against the identified available solutions and decide how robust or recoverable each application or system needs to be.

In most cases, the primary determinant in this decision is the required or desired recovery speed and determine the amount of acceptable data loss. Once the team decides how quickly a system or application needs to be recovered, the security and procedural requirements will follow. But, as the capability and speed of recovery increase, so does the cost—not just for hardware or facilities but also for training, procedures, and management. That said, it is important to remember that

TABLE 3 Process and documentation review[a]

Server Operations	Failure Mode	Solution Set
Recovery	Undocumented procedures	Policies and procedures
	Single employee knowledge	Cross training
	Unique implementation	Best practice
	Unknown passwords	System access
Security	Unknown access	Employee certification
	Insecure practices	Security training
	Unintentional actions	Operator qualifications
Business Operations		
Recovery	Inability to make decisions	Delegation of authority
	Knowing what to do	Disaster recovery plans
	Missing knowledge	Key employee
Security	No access procedures	Access rules and regulations
	Access guidelines	Security review

[a]*Procedural risks refer directly to standardizing procedures, managing vendors and partners, and ensuring appropriate staff knowledge, expertise, and cross-training.*

inexpensive solutions might offer significant protection.

Planning teams should consider three general levels of recovery.

Low priority areas. These include non–business critical systems with recovery needs in excess of one half hour. An example is using a backup application to copy files periodically to a second disk, relying on the end user to detect a failure and recover the file.

Fault-tolerant systems. The recovery environment for this highest level of data availability is designed to permit almost instantaneous recovery with no loss of individual transactions. In this configuration, end users may experience a slight delay in system response to a disaster, but otherwise, the user is able to continue working on the system as if nothing occurred. Fault-tolerant systems are very expensive.

High-availability systems have been developed to allow nearly continuous system and application operation at a significantly lower cost. This option employs a set of back-up systems at each level and includes various levels of automatic fault detection. Here, one trades off the fault-tolerant system's capability for near-instantaneous recovery, no loss of transaction, and no loss of user status for a lower investment and nearly analogous performance.

A high-availability system

One base configuration for building a high-availability system might include a server with a dual central processing unit (CPU) and full disk mirroring, allowing a system component failure with no loss in service or data. The company could also add a co-located, online, fully synchronized secondary server with a router that can redirect the server address of both if the first server becomes unavailable.

That configuration can manage a full or partial server failure. In that example, if the primary server's disk or CPU fails, the user should be able to continue online with no interruption and no loss of transactional integrity. Should the primary server fail, the user will probably need to log onto the server again.

A reasonable assumption is that the data center has full backups for power; internal and external network; heating, ventilating and air-conditioning; and telecom. If additional precautions are necessary, the company might require a second facility. In that case, it is appropriate to locate the secondary transactional server at the second data center and move the network rerouting into the network. As in the previous instance, users will most likely experience an interruption of service and have to log back onto their servers.

In considering the many risk issues associated with a data center and business continuity, a risk-management approach must be comprehensive in both scope and scale. Planning teams must carefully categorize and assess possible risks to ensure appropriate and balanced responses.

An awareness of life-cycle issues requires that companies examine operational issues for their data centers at least once or twice a year. A comprehensive list relating to hardware systems, software applications, and work procedures—as well as a sensitivity to assembling, confirming, and sustaining the required knowledge for ensuring successful recovery and security implementations—is imperative. The presumption that things will take care of themselves or that the design staff will build systems that are secure and easy to maintain leads companies to ignore or underinvest in operations that, by their nature, change constantly.

As companies introduce new systems, people, and procedures into their business, they must assess and integrate them into the overall risk management and investment plan. Similarly, as they retire or change procedures and systems, they should reassess the modified environment as well.

Of the three general risk categories, procedural risk includes some of the more difficult items to maintain, because it is highly dependent on ensuring that people adhere to processes. Security and recovery risks are more visible and tangible and therefore easier to understand, justify investing in, and implement.

To redress that, management attention, education, investment, and communication should focus on balancing all three risk areas to achieve the desired risk profile. With an explicit and systematic framework, companies can prioritize issues, direct management attention, and highlight resource, capital, and training investment needs to identify, rate, and rank potential threats to business success and operational quality.

Eli Hauser, *MBA, is executive director of data warehousing solutions and* **Paul Bleicher,*** *MD, PhD, is chairman and chief scientific officer of Phase Forward Incorporated, 1440 Main Street, Waltham, MA 02451, (888) 703-1122, email: paul.bleicher@phaseforward.com, www.phaseforward.com.* *Their article was first published (with the title "When Disaster Strikes") in* Digital Pharma, *a December 2001 supplement to* Pharmaceutical Executive, *a sister publication of* Applied Clinical Trials.

**To whom correspondence should be addressed.*