EDP ANALYZER

© 1980 by Canning Publications, Inc.

APRIL, 1980

VOL. 18, NO. 4

# RISK ASSESSMENT FOR DISTRIBUTED SYSTEMS

As computers are moved out of 'secure' computer centers into 'unsecure' offices, the vulnerability to computer crime will probably increase. How should data processing management deal with this new vulnerability? Can the responsibility for computer security be passed on to user department management, along with the machines? And if so, how should the responsibilities be split between data processing and user management? In this report we will try to answer these questions by discussing the initial steps in the computer security process—risk assessment—and how users can participate in meaningful ways.

Blake Greenlee, of Citibank in New York City, spoke at the 1979 National Computer Conference (Reference 1a) about managing business risks, including data processing risks. He has been in charge of developing and implementing corporate risk assessment standards at Citibank since 1973. He also directs the internal EDP audit consulting staff that supports the bank's audit programs worldwide. What Greenlee had to say is very applicable, we think, in the new distributed processing environment. He believes that not only the processing but also the risk responsibility should be distributed.

As Greenlee explained, Citibank has taken the pragmatic position that their equipment, files, and facilities are important *only* in so far as they affect the continuity of operations. Equipment costs and reconstruction costs are not their prime concern, because these are insurable. Citibank's real concern is *whether the operation can be reconstructed* after a crime, a fire, or whatever.

Citibank, like many companies, does business in a number of countries. That geographical distribution has led them toward distributed processing, in some very challenging and possibly threatening places. Their day-to-day banking operations depend upon the continuity of these distributed data processing operations.

So, Greenlee asks, what kinds of problems should corporations be considering in this distributed environment? "What do you do if your data center is still in operation, your communication facilities are intact, but, for some reason, you can not operate—say, there has been a fire in the building, or there is a local transportation strike? How do you serve your customers

ISSN 0012-7523. Multiple copy prices listed on last page. Photocopying this report for personal use is permitted, providing payment of $2.00 fee *per copy of report* is made to Copyright Clearance Center, P.O. Box 765, Schenectady, N.Y. 12301; please include identifying fee code 0012-7523/80/040001-13$02.00.

under such conditions? It makes no sense at all to have a data center that is running and has backed-up files if you can not do your business," he pointed out.

To answer such questions, Citibank has developed a standardized set of procedures for doing risk analysis that it uses worldwide. These procedures, which we will discuss later, were developed by the corporate audit group. They deal with the entire spectrum of risks of doing business.

As Greenlee explained, many of the company's risks are internal, people-related matters. These include the accidental modification, loss, or omission of information by an employee, damage to equipment (malicious or accidental), loss due to staff shortages, deliberate acts that are well-intentioned but counter to company standards or policies, fraud, and defalcation. The impact of such acts may be errors or processing interruptions, either of which can be difficult to recover from. In banking, Greenlee pointed out, errors or interruptions could result in difficulty with regulatory agencies, lawsuits, or unhappy customers.

People-related threats can also come from outside the company. In many countries, including the United States, strikes, terrorism, civil disorder, and crime may endanger business continuity. This can result in loss of access to company buildings, loss of utilities, loss of transportation, loss of police, fire and sanitation services, damage to equipment, and even kidnapping, extortion, or assassination.

A company must also look to its support facilities, including supplies, spare parts, and so on. Greenlee particularly stressed the problem of utilities, both public and internally-controlled. What happens, for instance, when personnel who are anxious to conserve energy shut down the cooling system on weekends, forgetting that someone else will be using the computer? It takes most mini-computers months to recover from eight hours of over-temperature operation, he noted. Or what if your operation relies on moving paper from one place to another, and the elevators in your high-rise building go out? Or what if the turbines that distribute water in the high-rise go out? This is an especially serious problem in some countries where importing replacement parts may be restricted or takes a very long time. Equipment does wear out and get damaged with startling regularity, he pointed out.

And finally, how does a company deal with natural disasters? What happens to operations—and to company personnel and their families—during such calamities as earthquakes or hurricanes? And what about courting disaster by locating next to an explosive chemical plant, on an earthquake fault line, under an airport landing or take-off path, or in a high crime area? At Citibank they try to consider all of these possibilities in their risk assessment process. It is really quite a spectrum, we think.

## The security problem

Citibank's risks are typical of what most large companies face, we suspect. They can be true even for medium-sized organizations with offices in only a few cities. In terms of data processing, Citibank's geographical distribution is just the kind of environment in which distributed data processing is proliferating. As we have pointed out in the past, distributed systems are going to be widely used in the future by many companies. More and more, companies with varied locations and highly differentiated internal divisions will find themselves putting processing capability out into the hands of the end users.

Certainly the problems of computer security, broadly speaking, are similar everywhere. But when one begins to look at actual specific sites, the problems do vary—because personnel, functions, and computer system configurations differ from one situation to another. Local combinations of assets, vulnerabilities, personnel, and security strategies make risks a local matter.

What happens—or could happen—at an office in one country is bound to be somewhat different from what happens in another.

### A new role for data processing

In this report we will be focusing on the role that a central data processing department might play in analyzing and controlling data processing risks in a distributed system environment. This is a particular instance of the general risk management problem that Greenlee addressed. Greenlee proposes a risk management based on the distribution of risk responsibility. We believe that the same is true in distributed systems—computer security risk responsibility will need to be distributed to user management. But how does one separate the data processing department's role from user management's role?

One role that we consider viable for a data processing department is to serve as technical guide and counselor for users of a distributed system. In this role, the department's function would primarily be to create and distribute technical guidelines for such things as equipment selection (pointing out system security features), risk assessment, contingency planning, departmental computer operations, and so on.

With distributed systems, the data processing department can no longer 'do everything'. The users must take over some of the responsibilities. In risk assessment we see user managers identifying local vulnerabilities, identifying and quantifying potential local exposures, and estimating their operating risks. All of these estimates would be based on guidelines and procedures prepared centrally by the data processing and internal audit departments. User management would also be responsible for putting the most necessary and cost-effective counter-measures into place, again based on centrally-prepared guidelines.

So what we are saying is: In a distributed system environment, risk assessment responsibilities should be distributed too.

## Traditional risk assessment

In the past, risk assessment for data processing has been performed by the central data processing and internal audit departments. The following is a brief description of the traditional risk assessment process. The points of concern apply to both distributed systems and to centralized systems. The key difference is that for centralized systems, the assessment itself has been performed centrally. Special teams from these two departments typically have performed all of the following assessment steps, for centralized systems.

First, the team has asked, "What are our *assets*? What do we have that is valuable, where loss or damage could hamper or destroy our ability to do business?" The team attempts to set a value on each asset identified. Of course, in a centralized system, the assets are right at hand. In a distributed system, however, many of the assets are 'out there somewhere'.

Second, the team has asked, "What are the potential *threats* to these assets? What could happen to compromise these assets?" Assets can be imperiled by many things, including accidents and well-intentioned (but misguided) deliberate acts. When considering computer crime, the questions could include, "Could the assets be stolen, damaged, or harmfully altered? How could these events happen?" We will look at some such threats in a moment.

Third, in looking at their operation and facilities, the team has asked, "Where is our organization *vulnerable* to those threats? For instance, is our computer site physically secure? Are sensitive files controlled and backed up? Are personnel able to gain access to areas outside their responsibilities?"

The fourth step has been for the team to identify and quantify *exposures* to threats. "What harm or loss could result—what consequences—if a threat did come to pass against an area of vulnerability? Would the organization be slightly inconvenienced, highly inconvenienced, or even put out of business?" Exposures involve possible financial losses. More importantly, though, as Greenlee mentioned, a company's main concern may well be recoverability, since financial losses can be covered by insurance.

Fifth, the team has asked, "What are the *risks*? That is, what is the probability of a given threat coming to pass? Low? High? Once in a day, a month or a year? Once in a century?"

Once an analysis of threats, vulnerabilities, and risks has been performed by the team, the result has been a set of figures. These are very approximate figures, since many of the input numbers are, at best, only guesswork. Data processing managers then have before them a list of specific dangers to their computer center and its operations, a set of estimates of how likely those dangers are to occur, and a set of figures for the harm or loss that could result. The dangers are ranked to indicate the relative levels of seriousness they pose for the organization. Based on such an assessment, priorities have been established for instituting security *controls*, or countermeasures.

## The new distributed environment

In the new distributed systems environment, the process of risk assessment will still involve the same elements listed above. But these must be broken up among user, data processing, and internal audit personnel. Central data processing and auditing may be too far removed from the user environment, functionally and perhaps geographically, to be able to assess users' individual risks. Later in this report we will suggest one approach for splitting up this risk assessment task. But first, we look at the distributed system environment and the new crime vulnerabilities it presents.

### New vulnerabilities

The key word in distributed systems, of course, is 'distributed'. Among the features of distributed systems that will lead to new vulnerabilities are the following:

- Distributed equipment
- Distributed data
- Distributed programs
- Distributed knowledge of computers
- Distributed program documentation
- Distributed printed forms

This distributed environment is much different from a centralized system, even one with lots of on-line terminals. Distributed systems move the actual processing capability 'out there'. In the central system, the data processing department controls the terminal network and must, therefore, also control security. But in the distributed system, many of the data processing components are no longer under the physical control of the central data processing department. Not only that, but there may not be much conformity among the numerous departmental systems. This is why we see the data processing department's role likely to shift from *control* to *guidance*—helping user departments create plans for equipment selection, software development, work scheduling, and security.

### New threats

Given the new environment, and the new vulnerabilities it brings, what threats are *most likely* to endanger a company with distributed computing facilities? We single out people-type threats, particularly from employees, to keep the discussion within bounds. The discussion could just as well cover threats of loss of critical services (such as electric power) or natural disasters.

*Threats involving theft.* One prime asset that can be stolen, often with devastating results, is a *master file.* This file can contain trade secrets, customer lists, sensitive financial figures, proprietary computer programs, or the like. Employees can also steal computer time—the *service* of the computer itself. Printouts and other *output* can also be stolen, especially if stored in unsecured areas. And, of course, computer *equipment* can be stolen. Small computers, terminals, modems, etc. can easily be carried out of a building.

*Threats involving destruction.* In this case, the primary targets are likely to be computer *equipment* and *data files.* Attacks on equipment by disgruntled employees, vandals, and even terrorists have been reported. As for data, its destruction might involve destroying floppy disks or writing over existing files.

*Threats involving manipulation.* Input data must be considered vulnerable to manipulation. One threat is the insertion of *false input.* Another is the *suppresion of output*, so that normal output does not appear. An embezzler might suppress output to cover his tracks, for example. Or he might simply *alter output*, for instance, by substituting a false output statement for the real one. Someone pilfering inventory could *alter master files* so that losses do not show on inventory listings. Or an employee with programming knowledge could *modify a program*, perhaps to change programmed decisions, such as account write-offs, inventory scrapping, stock reordering, and so on. Finally, *equipment* itself can be modified, for instance by substituting a spurious logic board to change computer operations. This might be done, for example, to by-pass or eliminate built-in audit checks.

Under the central batch system, it has been easiest to: (1) insert false input via the regular data entry mode (by-passing controls, of course), (2) steal master file data, for instance by walking out with a tape, copying it elsewhere, and returning it, (3) steal services, such as operating the machine at night for private purposes, and (4) damage equipment, for instance by arson.

These are the 'big four' *reported* computer crimes. They are the big four because they may well be the easiest to perform in a centralized batch processing environment—which is still the dominant environment.

When you add remote terminals to this centralized system, additional 'bad things' can happen. Using a terminal in a remote location—to which physical access may be easier—an intruder can: (1) change programs, (2) change master files, and (3) alter input. Most of the existing literature on computer crime focuses on a central batch operation augmented by on-line terminals. Almost none of it has dealt with the special aspects of distributed systems. Yet we think these aspects may cause significant shifts in the nature of computer crime. So let's look at a few of the threats that appear to be much more likely in a distributed systems environment.

*Theft or alteration of output.* With smaller, less expensive hardware involved, the distributed computers are not likely to be located in 'tight fortresses'. That kind of hard security for a small operation just may not appear economical. And when the computer is a mini- or a micro-computer incorporated in a work-station, then physical access to the computer may be even easier.

One thing that will be increasingly distributed along with all the hardware is printed forms. These are likely to be kept in cabinets right in the office. With offices easy to enter and keys to cabinets easy to obtain (as we reported in our July 1979 issue) payroll check forms, invoices, and such may be 'easy pickings'—and so will the equipment and programs with which to run them. An employee or a knowledgeable intruder can, for instance, run an accounts payable program after regular working hours to produce one or more fraudulent checks, which he then carries away and cashes.

*Theft of master file data.* In the centralized system, the data files are located in one place. But in distributed systems, we expect the files to be distributed also. In addition to that, as we reported in July, 1979, managers may use the computer to store information about their own activities: travel plans, appointment calendars, tickler files, and the like. Most such data will be stored on removable media, such as floppy disks. Compared with the standard magnetic tape, or even with many print-outs, floppies are easy to steal; an intruder or an employee can walk out with a floppy disk concealed under his shirt or tucked into a sheaf of papers. The disk could be copied elsewhere and returned; personal computers are making *that* kind of stealing easy to do.

Controlling the information stored on floppy disks will be a real problem, with serious repercussions possible if sensitive ones are stolen. For example, suppose a company has a set of new product specifications. They may use a central word processing system to prepare the specifications. And different departments, such as marketing and engineering, may copy all or part onto their own floppy disks. Without uniform and extensive safeguards, such information might be available in a variety of departments, and accessible to many, many employees.

*Modification of equipment.* In his now-classic book, *Crime By Computer* (Reference 2), Donn B. Parker talked about the Trojan Horse. This is a program, insidiously hidden within a larger, innocuous program, that is designed by an intruder to take over the computer. In distributed systems, with less sophisticated users, the Trojan Horse is still a threat—for theft, damage, or manipulation. The intruder just develops the program on his or her personal computer and then finds a way to get it onto a regular program floppy disk.

Another speaker at the 1979 National Computer Conference, Robert Jacobson, president of International Security Technology (Reference 1b), suggested the following criminal gambit. Assume that the disk controller chip in a small computer system is replaced by a field engineer with a new one. Jacobson asks, "How do we know what that new chip does? The same would be true of a new logic board inserted in a computer. Who knows what alterations could be made to the computer's operating system? Who knows what audit checks could be circumvented?"

*Theft of equipment.* Two factors make this a greater problem in the distributed environment. First, miniaturization has made computer hardware smaller, while at the same time increasing processing power. Thus, it is possible for user departments to have their own computers dedicated to their own functions.

Second, as we have already mentioned, computers are being used in locations that were never designed for the kind of physical security that is often found in centralized shops. Breaking and entering is easier. Hauling off a small computer can be an almost trivial task, only exceeded in ease by the job of selling it. There is a huge market for micro-computers, particularly when 'the price is right'.

So with the new distributed environment presenting such increased risks, we turn our attention to several risk assessment methods and their suitability in the new environment.

## Risk assessment methods in use

Earlier in this report we discussed the elements in the risk assessment process. We look now at how those general elements are combined in risk analysis methods. Several authors relate techniques that are applicable to either centralized or distributed systems. But they are not equally effective in both environments, we think. We have chosen two traditional approaches to illustrate commonly used assessment methods in use today. Both put the analysis task in the hands of a central team, which then makes key decisions for the entire organization.

*Centrally generated information.* FitzGerald (Reference 3a) recommends a 'threat scenario' technique for risk assessment. His approach typifies those that are based on a brainstorming effort by a central team. The work is done by a team composed of user, data processing, and audit department representatives.

Once oriented by written materials prepared by the central staff, the members meet in a number of two-hour sessions to brainstorm potential threats. All threats, no matter how improbable or bizarre, are written on large charts, so they are visible to all.

Using the charts, one team member creates story-form scenarios following the brainstorming meetings. Then all of the team members rank these scenarios by likelihood of occurrence. They use a method that does not involve the usual estimation of probabilities. Instead, threats are paired and ranked against each other, with the rankings then consolidated into one list. Following the ranking, each threat is weighted according to its relative seriousness to the organization. The team then develops possible safeguards or controls, and analyzes their cost effectiveness, in order to prepare an action plan for the organization.

Similar brainstorming techniques can be used to generate estimates of asset values and exposure costs. Such costs, like the probability figures used in most risk assessments, are at best only guesswork.

*Inward-flowing information.* Gerberick (Reference 3b) suggests a risk assessment method that involves a central decision-making committee consisting mainly of data processing people. Most of the input for this committee's work comes from a set of questionnaires filled out by user department personnel.

The user department representatives identify the important data that is maintained by each application system, and assign order-of-magnitude values to the possible loss or modification of that data. The committee gives the user representatives a list of *threat classes* (without trying to identify each specific threat) and the users estimate the likelihood of the occurrence of each threat class. The committee then calculates the *exposure*, in dollars per year, for each type of important data, to obtain company-wide figures and develop controls.

Certainly one of the difficulties with risk assessment, particularly with centralized assessments, is the time involved. Assessment teams draw numerous key people away from their regular work for relatively long periods. Gerberick and others have mentioned the difficulty of getting a commitment from senior management to perform risk analyses, for this very reason. And commitment, especially without senior level support, is equally hard to get from the potential risk assessors themselves. All of this is compounded by the fact that security deals with hypothetical events. Unlike the real events that produce company revenues, the tangible results of security and risk assessment are often hard to measure, especially when none of the threats materialize. Gerberick's answer to this difficulty is to make risk assessment only a part-time duty of the several data processing people on the committee.

### A distributed assessment approach

Centralized risk assessment may be viable in smaller organizations, or in a specific branch of a large one. But for a far-flung organization, we think a more viable solution is to distribute the risk assessment responsibility itself, under central guidance. Greenlee proposed such an approach in his speech. But before recounting the steps in his method, we will briefly describe Citibank's distributed environment.

*Citibank's distributed environment.* Kirschner (Reference 3c) reports on the evolution of Citibank's commitment to distributed processing. After having centralized processing in the 1960s, Citibank began to decentralize in the 1970s, through a controlled distribution of processing. The problem of trying to manage a broken-up transaction flow was well-known to the bank, so they created a system of 'channeling' in which one service group became responsible for all of the needs of one group of customers. Data processing support was dispersed out to the customer-servicing units, and in many cases, minicomputers were dedicated to specific applications.

That trend toward distributed processing continues today at Citibank. They are now moving toward the use of work-stations—dedicating a small computer and, say, one or two employees to service a highly specific customer group. The goal is to provide customized services for specific customers or groups rather than trying to fit customers into a generalized system. Since 1974, Citibank has moved from 25 data processing facilities, mostly using IBM hardware, to some 200 facilities using a variety of systems from several suppliers

*Citibank's earlier risk assessment experience.* Citibank's experience with risk assessment as a matter of corporate policy began in 1973. Their first approach was a centralized one, developed by the corporate audit department. Essentially, a team talked to the users and found out what the problems would be if specific threats materialized. Then they added up the risks, and arrived at a corporate world-wide risk posture.

But, according to Greenlee, there were some problems with this approach. First, the method caused too much paperwork. Second, the method generated wild errors in expected-loss figures. The risk assessment inherently involved many small probabilities, many potentially large losses, and errors in estimating both. Says Greenlee, "No one would believe us, and they were right." Lack of accuracy and credibility were major stumbling blocks in getting high-level support for subsequent security planning.

Their third problem was how to follow up the risk assessment with appropriate action. How were they to embed the risk assessment in a risk containment process for the bank? They found this to be a particularly aggravating problem in a widely distributed system, because the needs and concerns of managers varied greatly. Citibank learned that it is bad practice to try to interpose anyone between the line manager and the person who is supposed to assess the risk.

Due to those difficulties, Citibank moved to its present approach. It has been in place since 1977. Rather than distributing security *requirements* to its outlying departments, Citibank has opted to distribute risk assessment responsibility and procedures.

### How Citibank now assesses its risks

Citibank's risk assessment and contingency planning procedures follow a key principle of distributed processing: Distribution does *not* mean simply 'tossing the ball' to the user departments. Rather, user departments are given the responsibility, *but they are also given guidance.* Risk analysis is part of every department manager's job, and it is evaluated along with his other responsibilities at promotion and raise time. Following are Citibank's risk assessment procedures that Greenlee outlined (but we have used our terminology, rather than Greenlee's, for consistency of discussion).

*Training managers to assess risks.* Citibank's standardized procedures, which apply world-wide, are structured to be educational for the line managers who must comply. A tutorial slide presentation is given to all managers. It explains that in risk assessment they are to look at three categories. The first is *threats* (which Greenlee calls 'exposure types'). We discussed the kinds of threats Citibank considers in its international operations early in this report: threats from people (both employees and outsiders), threats from the failure of supporting services, and threats from natural disasters.

Second, the managers are asked to consider the *consequences* of these threats occurring. The main types of consequences are errors and processing interruptions—which in turn result in intermittent or catastrophic failures—and thus lead to interruptions in business continuity. That is the main concern for Citibank, as Greenlee stresses. The *costs* of recovery are insurable but recovery itself is not insurable.

The third category the managers are asked to look at is *containment measures.* Citibank has developed internal standards for data centers: for contingency planning, back-up (either off-site or redundant hardware on-site), maintenance, and training. Corporate policy requires these measures to be in place, and woe betide the manager who does not have them there.

To give the managers a better feel for the appropriate amount of detail expected by the corporate office, Citibank recommends that the managers try out the assessment procedures at home, for their families. It is suggested that they set an alarm clock for two hours. In the first half hour they should try to define their families' threats. In the second half hour they are to fill out a chart on the consequences if the threats occur. And in the second hour they are to put together a family contingency plan. If they can not do this family risk assessment in that time, then they are trying to tackle the problem at too great a level of detail. This exercise is meant to give the managers a sense of the scope and amount of effort to then expend on their business risk assessment. The managers also find, from this home assessment, that only a handful of counter-measures will be needed to protect their families. And the same is true for their business units.

***The goals of assessment.*** This assessment is not the kind of security procedure calling for elaborate studies and detailed analyses, or for elaborate knowledge of security technology. Those considerations can be part of the guidance available from central data processing and auditing.

Citibank's goal in this kind of assessment, again, is mainly to insure business continuity. So they require the same generic type of contingency plan and risk assessment from a small branch manager in, say, Jericho, New York, as from a manager of a major bank in, say, Sao Paulo, Brazil. The amount of detail will be different, but every manager who has line responsibility must be able to operate through the contingencies which he or she faces.

***The assessment procedures.*** Following their training, the managers are asked to perform a risk assessment of their operation, based on its current security measures. It is to include the likelihood of occurrence of each exposure type (threat) and the consequences of these. The assessment is to be both quantitative and qualitative, where possible. Precise cost figures are not necessary.

To help the managers perform this analysis, Citibank has developed two forms, which the managers fill in. These are to be done in pencil, not typed, since that much effort is not warranted by the kind of accuracy involved. Citibank stresses the ratio of effort to pay-off.

In the first table, the manager lists and ranks the appropriate threats. The ranking is in terms of expected frequency of occurrence. Low frequency is once in ten to fifty years. Medium frequency is once in one to five years (for instance, wear-out of heating, ventilating, and air conditioning equipment). High frequency is occurence more than once a year.

Certainly such estimates are easier to make for equipment failure than for crimes. But the relative difference between the theft of a mini-computer and the theft of a floppy disk (and the data stored on it) needs to be evaluated, as an example. Citibank tells its managers to forget the most infrequent threats, but to document their decisions to do so. They should concentrate on the medium- and high-frequency threats.

In the second chart, the likely consequences, if the threats do occur, are rated in terms of their relative importance to the company. Will the consequence be (1) no loss, (2) minor impact, (3) clear and measurable impact, or (4) put you out of business?

Next, the managers try to quantify the consequences likely to occur per day, and per occurrence, for certain types of threats. Citibank stresses that this analysis be done locally, because of local conditions—for instance, it may take six months in some countries to get replacement parts for repairing equipment. So the 'per occurrence' estimate is very important.

So far, the assessment process has been relatively simple. Given the level of detail and the quality of estimates requested, the manager's task is neither very complex nor very time-consuming. Greenlee emphasizes that the risk assessment should not become more elaborate than necessary, even if the results are not as rigorous as other methods might produce. In some cases, however, local managers might decide to use approaches such as FitzGerald and Gerberick suggest. In that way they can draw upon the expertise of their staffs in preparing the assessments.

The manager's next step is to prepare a brief report, about a page and a half long; Greenlee stresses its brevity. It should include: one paragraph describing the operation, a second one listing the major types of operational threats, and a third telling what actions are planned to reduce the risks. The two informal tables are attached as working papers.

***Signing off for responsibility.*** The risk assessment is signed by the line manager who has the responsibility, typically the person who prepared it. Then it is *reviewed* by the local internal auditor (in that country). Finally, it goes to senior management for *approval.* They must sign off for any residual risk.

### Contingency planning at Citibank

Citibank's risk assessment approach appears to be a good example of the role central data processing and internal audit departments may move toward as distributed systems spread. Following the risk assessment process, data processing should continue to help users create their contingency plans.

At Citibank, following corporate guidelines, each manager produces a chart showing what he plans to do if his operation is shut down for a day, a week, or a longer term. Despite all preventive measures, threats do materialize, and they must be planned for as much as possible. Citibank tries to plan for all appropriate categories of such exposure. Greenlee even recommends including emergency personnel measures, to help the organization's employees and their families recover from disasters.

Contingency planning should not ignore the threat of theft. There can be many types of impact from theft, not just monetary loss. For instance, the theft of some kinds of private or proprietary data could be damaging to an organization's reputation or customer relations. The loss of trade secret data could threaten a competitive position. Altering or destroying both the master files and their backup copies might even bankrupt a company; it can be much more difficult to replace data than to replace the computing equipment.

Also, just replacing hardware and software can be difficult, especially in foreign operations. This is a problem that is aggravated in distributed systems, since distributed configurations may become highly individualized. Replacing equipment can involve several suppliers and, in foreign countries, long delivery times.

Such losses are precisely the kinds that most concern Greenlee. Loss of time while you wait for replacements or seek alternate ways to keep operations going means inability to serve customers. Damage to reputation or entanglement of managers in legal proceedings can also impair the organization's ability to function. The solution to all such losses, of course, is the kind of contingency planning that leaves a manager a way out when disaster strikes. Not only is recoverability crucial, but it must also be as smooth, fast, and efficient as possible. Therefore, Citibank requires managers at all levels to have contingency plans. The plans are based on the signed-off risk assessment. And they are reviewed regularly for completeness, suitability, and accuracy. The auditors rate the plans—and the managers' salary increases and promotions depend on them. The central audit support group also guides the preparation of contingency plans and their maintenance.

## Possible counter-measures

In thinking about risk assessment in distributed systems, what kinds of counter-measures appear appropriate? And what types of guidance should the data processing department give to user management in this area? We will suggest two subjects that we think users will need help with. These are, of course, only two of many possible areas for which data processing can develop standards, technical evaluations, and 'good practices' guidelines.

First, we suggest that data processing provide users with technical evaluations. We will discuss one shortly: evaluations of security features of new systems being considered by users. User management really is not technically equipped to evaluate such offerings. So we suggest that data processing perform this function.

Second, we suggest that the data processing department draw up some 'good practices' guidelines for computers operating in office settings. The one we will discuss is controlling physical access.

***Evaluations of security features of new systems.*** It would be desirable for data processing departments to study hardware offerings and explain the security features (or lack thereof) to users. To find out just what types of security features will typically be available in hardware for distributed systems, we looked at the IBM 8100. This is what we found.

IBM announced its 8100 system in October 1978. It is IBM's entry into the distributed system market. The 8100 can be used as a stand-alone departmental computer or as a node in a distributed network, with a System 370 as the central host processor.

The 8100 system can be coupled with three security products announced in December 1977: the 3845 and 3846 data encryption devices, and the cryptographic sub-system. The 3845 is a table-top data encryption device designed to be connected between a modem and a computer terminal. The 3846 is a similar device for use in the data center. This pair of devices comes with a hand-held key pad with which users can enter and change stored encryption keys. If someone tries to unlock the unit's protection switch, the stored key is automatically erased.

The cryptographic sub-system is for use in a distributed system that uses IBM's system network architecture (SNA). It enciphers and deciphers data and produces, manages, enciphers, and deciphers encryption keys. It can be used to encrypt information moving over communication links or stored on tapes or disks. The sub-system provides data encryption capabilities at the terminals as well as at the central processor.

All three products use the Data Encryption Standard algorithm that has been adopted as a U.S. federal standard; see our December 1978 issue.

These encryption products, that can be installed in an 8100 system, are its most touted security features. In addition, the 8100 has key locks on both its processors and terminals. These are physical locks, to which authorized operators are furnished keys. The equipment must be unlocked to be operational.

Internally, the system contains a control program that allows only specified personnel to make changes to the operating system. And it has logical address isolation, which prevents one user from gaining access to another user's storage and work areas. These are the basic security features of the 8100. When it is connected to a network controlled by a System 370, other protection features are made available to it.

Data processing can *evaluate* such security features for users, in terms of the threats they are designed to guard against. Such evaluations can be particularly helpful if they point out where security features are either weak or lacking.

***Guidelines for controlling physical access.*** After users examine the built-in hardware and software security features as one part of their risk analysis, they will find that most security problems remain human problems—for which technical counter-measures are only part of the answer. So user department managers will also need to take into account the physical security of their computer installations.

Following are some suggested guidelines (taken from one of our other publications) for improving physical security, in the language and depth that might be appropriate for user department managers.

One counter-measure is preventing unauthorized physical access to the computer both during and after normal working hours. Unauthorized access can result in computer equipment being stolen or damaged, data being examined, changed, or stolen, etc.

*During working hours*, access control is relatively simple in small organizations, but more difficult in larger ones. In a very small organization, every employee knows the other employees, so it is a simple task to know who is authorized to use the computer. Unauthorized access in this time period is rarely a problem. However, in larger organizations, it is much harder to tell just who is authorized to use the computer. Such organizations may need some form of physical access control.

Some methods which restrict access include: (1) locking the computer room and having a receptionist at the entrance, (2) using badges and color codes on badges to indicate authorization, or (3) using signed passes for entering the room and particularly for taking things in and out of the room.

*After working hours*, the normal precautions apply, of course—using dead bolt (rather than spring loaded) locks, making sure that all doors and windows are locked, having adequate night lights, and hiring a security service that checks to see that the facility is secure at night and on weekends and holidays.

Even the smaller equipment prevalent in today's distributed processing systems is valuable. A small computer can be easily and quickly removed by a burglar—so the manager might consider using a security consultant to check his safeguards. This consultant might recommend the use of a burglar alarm, plus bolting equipment to the desk, table, or floor, to make removal more difficult.

The procedures used by the janitorial service should also be checked. When janitors enter a building at night to clean it, they often unlock numerous doors, making it easy for an intruder to enter. It may take less than a minute for the intruder to do mischief.

It is also very good practice to 'lower the visibility' of the computer equipment. Putting it near windows or publicizing its existence, for instance, are poor practices.

Key control is a serious matter as well. Too many keys may mean that too many people have access. And what happens to keys at night? Does a secretary simply put them into a desk?

Of course, we have only scratched the surface on the subject of physical security. Physical security involves many factors other than access control, such as protection from fire and water damage. This brief discussion of controlling physical access is the type of guideline we think data processing departments should create for their users. It is not technically confusing, and it concentrates on the major types of threats. But it can get managers thinking about the problems, and the counter-measures that need to be considered.

## Conclusion

In the distributed systems environment, *put yourself in the place of user department management.* All of a sudden you have inherited a major security problem when you installed your own computer. But you probably are not very familiar with evaluating those new threats and risks. You are willing to do your share, but you look to the data processing and internal audit departments for help and guidance.

The responsibility for risk assessment needs to be distributed. That approach is what we have discussed in this report. It looks promising.

REFERENCES

1. The 1979 National Computer Conference, June 4-7, 1979, New York City. The following presentations are not included in the proceedings, but they were recorded on tape. The cassette tapes may be purchased for $5.60 each (plus $1.50 per order for billing and shipping) from On the Spot Duplicators (7309 Fort Hunt Road, Alexandria, Virginia 22307):
   a. "Risk assessment techniques," Session No. 104, contains talk by Blake Greenlee.
   b. "Managing the computer security problem," Session No. 70, contains talk by Robert Jacobson.
2. Parker, Donn B., *Crime by Computer,* Charles Scribner 597 Fifth Avenue, New York, N.Y. 10017; 1976.
3. *EDPACS* (11250 Roger Bacon Drive, Suite 17, Reston, Virginia 22090); price $5.00 for each issue:
   a. FitzGerald, Jerry, "Developing and ranking threat scenarios," September 1978, and "EDP risk analysis for contingency planning," August 1978.
   b. Gerberick, Dahl, "Security risk analysis," April 1979,
   c. Kirschner, Leslie S., "Auditing in a minicomputer environment," July 1978.
4. Browne, Peter, *Security: Checklist for Computer Self-Audits,* AFIPS Press (1815 North Lynn Street, Suite 800, Arlington, Virginia 22209); 1980; price $35.

*Next month, we return to the subject of electronic funds transfer* (EFT), *particularly as it is being used by corporations in the conduct of their business (as opposed to retail* EFT*). Some interesting new services have been introduced since our previous report—such as bank services to support corporate cash management. We will give some user experiences with these services.*

EDP ANALYZER is published monthly and copyright© 1980 by Canning Publications, Inc., 925 Anza Avenue, Vista, Calif. 92083. All rights reserved. While the contents of each report are based on the best information available to us, we cannot guarantee them. Photocopying this report for personal use is permitted under the conditions stated at the bottom of the first page. Prices of subscriptions and back issues listed on last page. Missing issues : please report non-receipt of an issue within one month of normal receiving date; missing issues requested after this time will be supplied at regular rate.

## SUBJECTS COVERED BY EDP ANALYZER IN PRIOR YEARS

**1977 (Volume 15)**

*Number*

1. The Arrival of Common Systems
2. Word Processing: Part 1
3. Word Processing: Part 2
4. Computer Message Systems
5. Computer Services for Small Sites
6. The Importance of EDP Audit and Control
7. Getting the Requirements Right
8. Managing Staff Retention and Turnover
9. Making Use of Remote Computing Services
10. The Impact of Corporate EFT
11. Using Some New Programming Techniques
12. Progress in Project Management

**1978 (Volume 16)**

*Number*

1. Installing a Data Dictionary
2. Progress in Software Engineering: Part 1
3. Progress in Software Engineering: Part 2
4. The Debate on Trans-border Data Flows
5. Planning for DBMS Conversions
6. "Personal" Computers in Business
7. Planning to Use Public Packet Networks
8. The Challenges of Distributed Systems
9. The Automated Office: Part 1
10. The Automated Office: Part 2
11. Get Ready for Major Changes
12. Data Encryption: Is It for You?

**1979 (Volume 17)**

*Number*

1. The Analysis of User Needs
2. The Production of Better Software
3. Program Design Techniques
4. How to Prepare for the Coming Changes
5. Computer Support for Managers
6. What Information Do Managers Need?
7. The Security of Managers' Information
8. Tools for Building an EIS
9. How to Use Advanced Technology
10. Programming Work-Stations
11. Stand-alone Programming Work-Stations
12. Progress Toward System Integrity

**1980 (Volume 18)**

*Number*

1. Managing the Computer Workload
2. How Companies are Preparing for Change
3. Introducing Advanced Technology
4. Risk Assessment for Distributed Systems

*(List of subjects prior to 1977 sent upon request)*

## PRICE SCHEDULE (all prices in U.S. dollars)

| | U.S., Canada, Mexico (surface delivery) | Other countries (via air mail) |
|---|---|---|
| Subscriptions (see notes 1,2,4,5) | | |
| 1 year | $48 | $60 |
| 2 years | 88 | 112 |
| 3 years | 120 | 156 |
| Back issues (see notes 1,2,3,5,) | | |
| First copy | $6 | $7 |
| Additional copies | 5 | 6 |
| Binders, each (see notes 2,5,6) | $6.25 | $9.75 |
| (in California | 6.63, including tax) | |

NOTES

1. Reduced prices are in effect for multiple copy subscriptions and for larger quantities of a back issue. Write for details.
2. Subscription agency orders are limited to single copy subscriptions for one-, two-, and three-years only.
3. Because of the continuing demand for back issues, all previous reports are available. All back issues, at above prices, are sent air mail.
4. Optional air mail delivery is available for Canada and Mexico.
5. We strongly recommend AIR MAIL delivery to "other countries" of the world, and have included the added cost in these prices.
6. The attractive binders, for holding 12 issues of EDP ANALYZER, require no punching or special equipment.

Send your order and check to:
EDP ANALYZER
Subscription Office
925 Anza Avenue
Vista, California 92083
Phone: (714) 724-3233

Send editorial correspondence to:
EDP ANALYZER
Editorial Office
925 Anza Avenue
Vista, California 92083
Phone: (714) 724-5900

Name______________________________

Company ______________________________

Address ______________________________

City, State, ZIP Code______________________________