



SANS Institute

Information Security Reading Room

Business Continuity Planning In Difficult Economic Times

Suzanne Widup

Copyright SANS Institute 2020. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Business Continuity Planning In Difficult Economic Times

Suzanne Widup

GSEC Practical

Abstract

Business continuity planning on an enterprise-wide level is an involved and costly process. Certainly, each company would do well to have some kind of comprehensive plan to handle crises. However, in these challenging economic times, many companies are choosing to postpone or eliminate business continuity planning from their budgets, even though they may have no plan at all in place. This being the case, security professionals may want to try a more modular approach to safeguarding the company's assets. A Critical System Recovery Plan documents each step required to recover an application deemed vital to the well being of the organization. Since it focuses on only one application environment at a time, the scope—and thus the cost—is much reduced. Compiling these plans on a gradual, system by system basis will move the corporation slowly towards the goal of being able to document their recovery processes, while still having a lesser impact to the reduced revenues that businesses are seeing today.

Introduction

“The primary objective of a Business Resumption Plan is to enable an organization to survive a disaster and to reestablish normal business operations. In order to survive, the organization must assure that critical operations can resume normal processing within a reasonable time frame.” (C.N.S.) This being the case, a business continuity effort would seem to be an essential function at any company. However, in many businesses, managers are reluctant to release funding for this activity due, at least in part, to current economic difficulties.

According to Cutter Consortium, 20.4% of organizations have not developed a disaster recovery plan. Of those business that had a plan, however, 26.1% indicated that their plans had not been tested. The reluctance of management to fund the planning for recovery of critical systems is sometimes based on the idea that it's only applicable to large disasters. It's convenient to assume the odds of an earthquake or terrorist attack taking out the data center are very small, but that doesn't account for the most common occurrences that cause interruption of critical services and systems. “Note that 80% of **all** downtime is caused by either hardware failures or operational errors caused by people. The impact to the business may be no less damaging whether caused by human error, an unanticipated power outage or a terrorist attack.”. (StorageTek)

Given the above statistics, the need for some kind of plan is evident. However, with the pressures on managers to control costs, it can be difficult for a security professional to convince management that this kind of activity is prudent at this time, especially on an enterprise-wide scale.

Consequences of Doing Nothing

It is uncomfortable to be a security professional unable to effect change in the arena of disaster recovery, knowing that if something does happen, the company may be at risk of failure. However, this is the situation many find themselves faced with due to budget limitations. Information assurance personnel may advocate a full enterprise-wide plan to mitigate the effects of a failure, but they increasingly encounter resistance, with cost listed as a major factor.

Many businesses believe crises only happen to others and that their size or some other feature makes them immune. They genuinely believe ‘It will not happen to us’. Others firmly believe that insurance will cover the cost but insurance does not win back lost market share. The case studies at business schools are littered with such examples. Whilst bombs, fires and flood capture the headlines almost 90% of crises are ‘quiet catastrophes’. It is these ‘quiet catastrophes’ that also have the potential to damage an organisation's most valuable assets; its brand and reputation. Recent research indicates that where an organisation has successfully dealt with

a crisis their shareholder value price has increased in the long-term in contrast to those who did not or were perceived not to have managed the crisis well.

IN ESSENCE AN ESTABLISHED AND SUCCESSFUL BRAND OR PUBLIC IMAGE, REPUTATION AND TRUST OF EITHER A PRIVATE OR PUBLIC SECTOR ORGANISATION CAN BE DESTROYED IN MINUTES UNLESS VIGOROUSLY DEFENDED AT A TIME WHEN THE SPEED AND SCALE OF EVENTS CAN OVERWHELM THE NORMAL OPERATIONAL AND MANAGEMENT SYSTEMS (The Business Continuity Institute)

Many managers feel that tape backup is sufficient planning to avert disaster, but if there is no documentation covering the steps to recover a particular system--or access to the personnel who know how to recover the systems by heart (if that person even exists)--tapes may be less useful than hoped. Basing the recovery plan on one person or even a small group of people puts too much emphasis on their presence.

In fact, the failure need not destroy the data center to have serious impact to the company's reputation and future. As outlined above, even the short-term loss of service can cause long term damage to market share and consumer trust. An event that limits the personnel's access to the building for a time may be sufficient to be categorized as a disaster if systems are down and people are unable to reach them to fix the problem. Certainly, this was the case for buildings after the 1989 earthquake in northern California. People were not allowed back into buildings until they had been cleared as safe to enter by officials. However, if the systems that are vital to a business are housed in a building that is not yet considered safe (whether it is or not is academic if access is denied), and there are critical systems down in that facility, then the effect is the same as if the building were destroyed in the short term.


If the failure is catastrophic, then the figures are even more disturbing. "After only 2 days of a total data center failure, the average business experiences a 30% drop in its essential business activities. By the 5th day, 70% of its capabilities are lost, and by the 10th day, the typical business is functioning at only 10% capacity." (Fulmer) A large percentage of the companies who suffer a catastrophic failure never recover at all, or are out of business in a year.


Finally, there are legal repercussions to not having a plan in place. One law in particular--the Foreign Corrupt Practices Act of 1977 "imposes penalties on publicly traded companies and their directors for not appropriately accounting for an safeguarding corporate assets". (Fulmer) There are other laws that touch on accounting requirements from the IRS—what happens if your accounting data cannot be provided upon request and you're a publicly traded company?

The Impact of the Current Economic Environment

In this difficult economy, businesses must scrutinize their spending carefully. While having a company-wide business continuity plan may be desirable, budgetary constraints may make it less than feasible at this time. Managers are reluctant to commit resources to plan for an event they hope will not occur during the best economic times—the current environment makes them even more reluctant.

Of the 300 business-technology executives interviewed in December about their 2003 budget plans and IT strategies, 65% say that business-continuity planning and improved disaster preparedness are priorities. A year ago, 77% named these areas as top priorities.





Business-continuity planning has fallen in priority for 2003. Some 61% of executives at small companies surveyed report continued commitment, down 14% compared with a year ago. Responses among managers at midsize companies show a similar decline, and larger-sized sites are down just 7% year over year. (D'Antoni)

Certainly, an argument can be made that many companies were making business continuity planning a higher priority in the wake of the September 11th attacks. This event, above all, has brought home to many companies that their systems are vulnerable not only to acts of nature, but also acts of man. Some of the decline can be attributed to either a waning of interest as the immediacy of the attacks fades over time. Some of the decline could be attributed to company's completing the efforts they began last year. However, a substantial amount of the decrease in spending must be attributed to the decline in the revenues of the companies due to the economic difficulties currently facing the country.

John Ervin, a systems administrator at Tessy Plastics LLC in Lynchburg, Va., said a lack of funding has forced him to buy used equipment to back up his systems. "We've implemented a used tape drive on our main server and do good backups," he said. "If I had to purchase the stuff new, I couldn't have done it. ... Right now, money is tight."

And Ervin isn't alone, according to a study released last week by Dataquest Inc. in San Jose. The study, "Investment Decisions: Preparing for Organizational Disasters," found that IT managers from 205 companies representing eight vertical industries in the U.S. aren't investing appropriately in disaster plans because of inadequate budgets.

"Budget constraints are forcing an average of 40% of respondents to rely on a best guess to determine potential risk rather than obtaining formal

assessments, which would be too costly," said Tony Adams, principal analyst at Dataquest's IT Services group.

"Preparation is key, and without adequate investment for protection of critical systems, the repercussions of disasters will be lengthier and more costly," he said. (Verton)

Indeed, many managers are choosing to omit business continuity efforts altogether in the interests of meeting their budgets. Certainly, traditional business continuity efforts that are company-wide--while clearly valuable--are costly both in resources and capital. The formal approach has many steps and must involve many people. Many companies do not have the expertise to do this in-house and must pay consultants to handle it for them. This drives the costs up very quickly, and is sufficient to discourage some companies from attempting the effort at all.

The Traditional Approach to Business Continuity Planning

According to Andrew Hiles, in "Business Continuity: Best Practices", these are the steps in a traditional Business Continuity Project:

- Understanding and "selling" BC issues – raising awareness and gaining commitment
- Planning the BC project
- Risk Analysis and Risk Reduction
- Business Impact Analysis
- Defining Continuity Strategy
- Developing the BC Plan(s)
- Developing and Implementing BC Plan Maintenance Procedures.
- Developing and Implementing BC team training and plan testing programs.

The traditional approach to business continuity planning would include many phases, such as getting management buy-in (a daunting and time consuming task in itself), and assembling a team of qualified people who can do a risk analysis and business impact analysis. Events are classified by their impact to the business in the amount of damage that they would cause if they occurred, coupled with the likelihood of their occurrence. The plan is then drawn up and (ideally) tested. Revisions are made based on the lessons gleaned from the tests. The traditional approach, when the costs are assigned to each phase, may be more than many companies are willing to pay in this environment. This is understandable in the short term, but should be strongly scrutinized as a long-term policy, as the consequences can be devastating for not having an effective plan in place.

Certainly, the formal methodology will help the company understand the return on investment for various activities associated with the disaster recovery planning process. However, while this approach has considerable value in a formal business continuity effort, when a company is pressed for time and resources, a more direct approach may be increase the likelihood that management will approve this as a task that should be undertaken during times of lower revenue.

An Alternative Approach

In the absence of the ability to fund a company-wide business continuity/disaster recovery program, a modular approach may be better than nothing. A Critical System Recovery Procedure (CSRP) can be written for each critical application. Diagrams can be included to show interrelationships between hardware that form application architectures. This documentation can serve as a step-by-step procedure for recovering a particular application that was deemed critical. In this manner, limited progress can be made towards business continuity in an organization, while still keeping focus on costs.

These plans should not take into consideration the time it may take the company to assemble the necessary components required to enact the plans, but assume that this has been accomplished, and that the business is now ready to start the recovery of that particular application. The overall decisions of where to locate the hardware, hot site vs. cold site, how to handle getting the hardware in place, etc. belong in the overall company-wide plan, when that is completed.

In the meantime, these procedures can be provided with the assumption that if the data center suffers a catastrophic event that renders it unusable, at some point, these applications will need to be recovered. In the future, when a business continuity program has been endorsed, these kinds of documents will be invaluable, and can then be easily incorporated as tested and verified components of an overall plan. So while the larger program may not be feasible at this time of economic downturn, it does not mean that progress cannot be made towards an eventual goal of being able to recover critical business functions a larger scale. Even small steps towards the goal will leave the business in a better position than if no action was taken.

Basically, this involves narrowing the scope of a business continuity plan to just one application at a time. Approached on a modular level, each CSRP is written with some initial assumptions defined at the beginning of the document.

Here is an example of the assumptions section of a CSRP:

Assumptions

- The Data Center has experienced an event that has rendered it unusable.
- The System Administration staff are not the people following this plan.
- The Business Continuity Plan/Disaster Recovery plan for the company has been activated and followed to the point that the following necessary prerequisites are in place prior to the use of this plan:
 - Facilities
 - HVAC
 - Power
 - Network Infrastructure
 - Server Hardware
 - Operating System media
 - Tape reader
 - Tapes from backup vendor

Given these assumptions, the focus of the document should be stated next—this will serve to limit the expectation of the audience for the document, as well as keep the focus firmly on the application to be recovered, without the need for covering those items that are in the assumptions list. The document should state the scope of this particular recovery plan, and not concern itself with items to recover that are outside of its reach. This is very important, as if “scope creep” enters into the picture; the likelihood of being able to successfully finish a document that will recover the critical system intended is lessened. Remember, one of the goals is to make certain that this is a limited effort, and thus get approval of the budget conscious management to undertake.

A sample scope statement would be:

Scope

Given the assumptions listed above, the scope of this plan is limited to recovering the Clear Case software development environment at ABC Company in the event that the systems involved are no longer usable.

Determining Which Systems to Document

Traditional business continuity planning involves risk analysis and identification of critical systems. This can be time-consuming process and until costs are associated with the system to be restored; there can be a tendency to try and include every system. When there is no price tag attached, every manager thinks his or her system is a high priority to be recovered. When the costs of the

planning process and the necessary maintenance are communicated, however, managers may make the decision that the cost required isn't something they can afford just now, and perhaps those systems can wait longer to be recovered.

A good approach would be to assign a dollar figure to each application for the cost of recovery planning and maintenance. This would be compared to the dollar figure for downtime. The return on investment is then easier to determine, and decisions can be made as to which applications should have the first CSRPs.

An alternative approach would be at the system administration level. The people that configure, install and manage the servers that these applications run on should be in a unique position to document their restoration in the event of a catastrophic failure. Even though this will not be a substitute for an actual business continuity plan, with the personnel identified and its locations for restoration in the event data center is unavailable designated, it is a starting point.

Typically, the system support personnel are aware of which systems are the highest profile and can tolerate the least downtime. Usually this is because if they do go down, these systems generate the highest number of people screaming for their access to be restored. While this isn't always an indicator of criticality to the business, having the system administration personnel involved in identifying which systems absolutely must be restored in the event of a catastrophic failure is something of a shortcut in the risk analysis process.

An additional shortcut would be to talk to the people who manage the company help desk. There may already be service level agreements established for each of the applications in production. If this is the case, those with the shortest amount of downtime tolerable are typically classified with a higher priority when they go down. Finding out which of the applications can stand only small amounts of downtime may provide a starting point for which applications should be documented and in what order.

One final suggestion if you're still at a loss as to which systems are potentially high priority to the business and should be candidates for a CSRPs—ask the questions “Is this revenue impacting?” “Is this customer facing?” and “Are there legal ramifications to losing access to this data?”

For the systems that are revenue impacting, obviously downtime will result in revenue losses to the business and the return on investment for documenting the CSRPs would be a matter of doing the math.

For the systems that are customer facing, your company's public perception is affected, which may have lasting impact on the reputation and future market share of the company.

For the systems that your Legal department has determined will have create legal ramifications for the business if the data is either inaccessible for a period of time, or if the data is destroyed and not recoverable, the business can make the decision as to which of the applications/systems should have CSRP's developed.

The Contribution of the Support Organizations

One of the most valuable sources of expertise in compiling the CSRP is the support organization(s) charged with keeping vital applications up and running. Not only do these people have enormous expertise in the installation and support of applications and operating systems, they may already have an appreciation of which systems are the most critical to the business.

The system administration staff should be of tremendous help in determining which systems need immediate disaster recovery procedures written for them, and providing information on how to recover those applications. These people are your subject matter experts on the operation systems, and possibly also the applications. The application support personnel are another excellent resource for information on the configuration and relationships the application depends on. This is especially true if the application is a homegrown one. Getting the management approval to work with this group of people is essential for successfully creating these kinds of plans.

According to the (Dataquest) survey, another flaw in organizations' approach to disaster planning is that too few actually involve IT managers in the planning process. Though executives often weigh in on which business disruptions would most hurt the bottom line, IT managers actually know which people and systems run the business, and how. Dataquest urges companies to get their IT managers involved in the disaster planning, as well as helping evaluate every new corporate initiative. Much like information security in general, the analyst firm cautions that organizations must begin factoring in business continuity and disaster planning whenever evaluating a new project, instead of just treating recovery as an afterthought or add-on. (Schwartz)

In fact, the ideal time to validate a CSRP is when the application is first being rolled out into production. At this time, the details of configuration have passed through testing and QA, and are ready to be finalized. All application configuration information can then be documented while it's fresh in everyone's minds. Just remember that subsequent changes should be recorded in the document so it doesn't become stale.

While the time to verify the CSRP is when the system is ready for prime time, the time to start planning for it is in the initial stages of the project, not as it's rolled out onto the raised floor. If the continuity planning were to be incorporated into

the project lifecycle of your company, the future systems would be documented as they go into production. Certainly this is better than having to go in after the system is in place and try and document how to recover in the event of a failure. After all, it's not as though once even an enterprise-wide continuity plan is compiled, no further systems are put into production that require planning for continuity. Business continuity planning is a process, not an activity that gets done once and put on the shelf (unless you want that plan to fail when it is most needed.)

Types of Information to Gather

The types of documents that would be valuable include:

- Design documents
- Data flow diagrams
- Dependency documents
- Architectural documents showing relationships between systems
- Installation specifications for the operating systems and all applications
- Vendor contact information
- User process flow diagrams

If possible, obtain installation and configuration documentation from the vendors on the operating systems and applications, and incorporate them as an appendix. They will be priceless as reference material for the person doing the recovery. Original media is also valuable—the original operating systems media, for example, will be necessary for your recovery process. This is especially important for those companies that currently automate their system builds through image servers. You aren't going to have those resources when recovering a system at a remote location in the event of a stoppage. Your hardware may not be exactly what the system is currently running either, so a full installation is most likely going to be necessary. Also, if a particular version of the operating system or application is required, be sure not to just put in the most current version of the media.


Consequently, some kind of provision for storage of the media and the plan should be made. Off-site is recommended, preferably far enough away to not be affected by the event, yet able to be accessed quickly. Many companies use branch offices to store these kinds of things, and for documentation, hosting companies may be an option if you use web sites for your documentation.

Also important is communication of the plan and where it resides to the people in your company. The time to ask "Where's the plan?" is before the disaster strikes. If people have to lose time finding the recovery documents and materials, you have delayed critical recoveries that much longer.

Handling communications during the crises is something you may also want to address. One company setup a central phone number to be used in case of a crises. Emergency cards were issued to all employees to keep in their wallets with the phone number. (Incidentally, this phone number was also for the families of employees to call in and get the status of their loved ones in the event of a disaster—of considerable comfort to the people involved, as well as being good organization.) On cards such as these could be listed the location of the plans and how to obtain them. This is most beneficial for an enterprise-wide effort, but can have good consequences for the CSRPs as well.

What Belongs in a Critical System Recovery Procedure Document

So let's assume you're charged with writing the CSRP document for the following application:



Based on the diagram above, you will need to write a procedure for each of the following components that the application architecture uses:

- Solaris operating system, plus configuration to make the system an NIS master and restore existing maps for authentication
- Solaris operating system plus installation of clear case
- Network appliance installation and configuration, plus restore of data
- Windows 2000 operating system for client, with clear case client configuration

As you can see, when you're assigned the development of the CSRP for a particular application, it can be much more complex than just one system. The first step is to document the existing dependencies with a diagram that shows what the components of the system are. Once that is done, breaking down the individual components into recovery documents is less overwhelming.

In the above example, this would involve procedures for two Solaris systems, one Network Appliance, and one Windows system. You will need to provide detailed and technically simplified procedures for installing the operating system for each host. You should document all relevant configuration of each system, and the steps required to affect that configuration (i.e., what files need to be edited, what parameters set, etc.). Screen shots of an operating system build would be particularly valuable to the person who is tasked with recovering the system.

Document the application installation procedures for each application that should be installed. Include information on your company's particular configuration and how to set up the application to reflect that configuration, keeping in mind the person who is performing the steps has never seen this before.

Include a section to document the steps required to get the data back into the appropriate places. This should include a description of the current backup process, and instructions on how to restore the data in a manner that the application will be able to access it. For instance, if the data is exported from a database, you will need to provide specific commands on how to have the database import that data back into the database—just restoring files may not be sufficient to get the system into a usable state.

Finally, if there are any other steps required for final processing or configuration before the system is considered restored, be sure to document those as well.

One important section in the recovery documentation is the contact information for all relevant vendors. This should include the vendor web sites where appropriate (but don't count on them being accessible for your plan to work—especially if you're within the same geographic location where a disaster could affect both companies), and their technical support contact information. License

information for both operating systems and applications is also important, as this is sometimes required for installation. Information on support contracts and service level agreements should be documented here as well. Expectation should be set on turnaround from the vendor where appropriate.

One easy to overlook item to include would be the contact information for the vendor that handles the offsite tapes, if your company uses such a service. Also, since tape drives are not always present on the systems when you order them, contact information for obtaining a tape drive that will handle the necessary recovery tapes, as well as instructions on configuration and use would be vital to any recovery plan. If this information is out of scope, or planning for this eventuality would exceed budget constraints, include it in your assumptions section that this will be present already before the plan is enacted, as shown in the example above.

Who Should Test Each Procedure

One common problem with many company's continuity plans is that the procedures for recovering systems assume familiarity with the operating systems, frequently of a system administrator level. The uncomfortable fact is that it may be someone with little or no computer experience who is left to follow those plans. Note that in the case of the World Trade Center attacks, technical people from all over the country were pulled in because the pool of local resources was insufficient (or too traumatized) to meet the demand caused by the disaster. For this reason, a company cannot assume that a person with a particular expertise will be the one who is actually restoring the systems. Consequently, the recovery plans should be very detailed and easy to follow. Assume your audience is a non-technical manager who is trying to recover a particular system.

Define terms so that the person following the plan doesn't get hung up on the jargon. A glossary included as an appendix would be of help for the person following the plan to have as a reference for terminology that may be unfamiliar to them. A list of common operating system commands may also be of help for the more cryptic operating systems. Navigation commands for the command line, as well as editor commands if applicable for editing configuration files is also helpful. Remember, the audience is a non-technical person who is not familiar with the operating system at all.

When you have the recovery plan tested, make certain the person is as non-technical as possible, so that prior knowledge assumptions (and there will be some—it's almost guaranteed) can be identified and clarified in the plan. This implies that the plan should be tested more than once, and by more than one person. That way, the clarifications can be tested by someone who hasn't already followed the plan once, and thus won't be familiar with it.

While testing an enterprise-wide recovery plan can be costly and have an impact on many resources, testing a CSRP involves a much smaller scale, and thus lower cost. Equipment may be rented or borrowed from vendors for the duration of a single test to further reduce the cost. Fewer personnel need be committed to a CSRP test than to an enterprise-wide disaster recovery test, making it (hopefully) more likely that the plan will actually be tested—and possibly allowing the lessons learned in one test to provide immediate feedback to another test if time allows.

Integration Into Eventual Enterprise Wide Plan

The goal of these CSRPs is to have procedures that will eventually be able to roll up into an overall enterprise-wide plan. They should be modular in nature, and able to stand on their own, given their initial assumptions. In this manner, systems can be documented and tested one application at a time, at a substantially reduced cost to the organization.

When the enterprise endorses a full-scale business continuity effort, these plans can be incorporated easily, and much progress has already been made towards the ability of a company to recover in the event of a disaster or failure.

Chances are, if you've identified the critical systems well, the eventual business impact analysis will show that the systems you have documented are the ones that the company can stand the least downtime from. This being the case, the work that has already been done will shorten the time to complete the enterprise-wide effort as well.

The other benefit is that the assumptions are spelled out in each plan document, pointing the way to the planners of the large-scale endeavor to the items that are still in need of clarification. The company-wide effort can focus on those things that would be overarching, such as the need for the facilities to receive the replacement hardware, getting that hardware, and all the other items that are necessary before the CSRPs can be enacted.

Some may say that this is going about business continuity planning from the wrong direction, and certainly the formal process is preferable for a company that can afford it. However, barring that, progress can be made towards safeguarding the assets of the company in the event of a failure.

With a more modular approach to the process, there may be a tendency to assume that the entire plan need not be tested as a whole. This would be a mistake—the assumptions that various pieces of infrastructure are present should also be tested as part of a full test of the plan. Do the assumptions cover all that was really assumed, or were some things missed. The best time to find

this out is during a full test of the enterprise plan, and not once the plan has been activated.

Conclusion

While even a collection of CSRPs are no substitute for an eventual plan of larger scope to cover such assumptions as where the systems are to go, how will they be acquired, powered, networked, etc., these documents are better than not having any information on how to recover critical systems in the event of a failure. They provide a less costly approach to making a company better able to recover from an event, while remaining modular in nature and thus easier to incorporate into a company-wide plan when developed.

© SANS Institute 2003, Author retains full rights.

Works Cited

The Business Continuity Institute. Business Continuity Management – Good Practice Guidelines. 1/11/02. <http://www.thebci.org/BCI%20GPG%20-%20Introduction.pdf>

C. N. S. Computing and Networking Services. Disaster Recovery Planning. University of Toronto. 4/2003. <http://www.utoronto.ca/security/drp.htm>

Cutter Consortium. More Than 20% of Companies Have No Disaster Recovery Plan. October 12, 1999. <http://www.cutter.com/research/1999/crb991012.html>

D'Antoni, Helen. Business Continuity Slides Down The Priority Scale. Jan 13, 2003 <http://www.informationweek.com/story/IWK20030109S0002>

Fulmer, Kenneth L. Business Continuity Planning: A Step-by-Step Guide. Brookfield, Connecticut: Rothstein Associates, Inc. 2000. pp. 1-2

Hiles, Andrew. Business Continuity: Best Practices. Brookfield, Connecticut: Rothstein Associates, Inc. 2000. pp. 2

Storagetek. Business Continuity Solutions. 05/09/03. http://www.storagetek.com/solutions/bus_continuity/

Schwartz, Mathew. Dataquest to Business: You're Not Ready. March 19, 2003. <http://www.esj.com/News/article.asp?EditorialsID=461>

Verton, Dan. Tight IT Budgets Impair Planning as War Looms. March 10, 2003. <http://www.computerworld.com/securitytopics/security/recovery/story/0,10801,79176,00.html>



Upcoming SANS Training

[Click here to view a list of all SANS Courses](#)

SANS Sydney 2020	Sydney, AU	Nov 02, 2020 - Nov 14, 2020	Live Event
SANS Secure Thailand	Bangkok, TH	Nov 09, 2020 - Nov 14, 2020	Live Event
APAC ICS Summit & Training 2020	Singapore, SG	Nov 13, 2020 - Nov 28, 2020	Live Event
SANS Community CTF	,	Nov 19, 2020 - Nov 20, 2020	Self Paced
SANS Local: Oslo November 2020	Oslo, NO	Nov 23, 2020 - Nov 28, 2020	Live Event
SANS OnDemand	OnlineUS	Anytime	Self Paced
SANS SelfStudy	Books & MP3s OnlyUS	Anytime	Self Paced