

WHITE PAPER



PRESENTED BY:



INCIDENT RESPONSE AND DISASTER RECOVERY:

**HOW INCIDENT RESPONSE & DISASTER RECOVERY, AS APPLIED TO
INFOSEC, REDUCE VULNERABILITIES THAT RESULT FROM INCIDENT-
CAUSING EVENTS**

John R. Perkins, CISSP

James D. Heimberg, ABC, Ph.D., ISSO

March 27, 2006

INTRODUCTION

Federal Network Services, Inc., Information Security (INFOSEC) Services Division is pleased to present this white paper that describes how incident response and disaster recovery planning in support of INFOSEC significantly reduce the vulnerabilities that result from incident-causing events. There are three levels of documentation and planning that will be discussed in this white paper: incident response, contingency planning, and continuity of operations planning (which is often referred to as disaster recovery planning).

The approach methodology must be holistic to be effective. The entire enterprise and its attendant policies and procedures must evolve with all three levels of incident response and disaster recovery planning in mind. If one area were to fail or become sidetracked, the plans for all three may not work. Leeway and flexibility can be built into the plan.

With proliferation of Information Technology (IT), the Government is faced with the increasingly difficult challenge of protecting its automated information resources. Government-owned computers and networks process, store, and transmit information that is critical to day-to-day operations. Therefore, Federal IT systems are an inviting target for those committed to disrupting the U.S. way of life. Likewise, laws have been passed to guard the infrastructure of local and state governments, Federal contractors, and specific areas of the U.S. economy. These systems, which are under Federal mandate, need to have similar types of protection. Federal IT resources and those under mandate often support critical computing activities and many other operations on which the U.S. population relies. The likelihood of unauthorized access to these IT systems via the Internet and other avenues is high. The security of IT resources requires use of reasonable caution to secure the systems and networks, and the ability to respond quickly and efficiently if system and network security defenses are breached.

Not only is data at risk, but authority to use, add to, manipulate, or remove data that could be a hazard to IT operations, the security of U.S. citizens, or the economic well being of the nation is also at risk. For this reason, any attack or force majeure that may impact the processes under which the information is used should be included in the enterprise's policies and procedures as risk factors. The inclusion of incident response and disaster recovery planning into the policies and procedures used throughout an enterprise is paramount to success. Things happen, they always will, and every enterprise must be prepared to deal with events that slow down productivity. This document will focus on Information Technology (IT) associated organizations, either those of the Federal government or those that are under Federal mandate.

An effective solution must be created that will be used to respond to any event that might pose a hazard to conducting business. Although the phrase "Incident Response and Disaster Recovery Planning" can be misleading to some extent, it is not always a major disaster that strikes and leaves an enterprise on its knees. It could be something as simple as a pinched network cable in an equipment rack that halts Internet access. But with a good plan in place, it can be dealt with accordingly.

Of additional concern in any incident is returning to normal operations. Depending on the length of the incident and its impact on IT systems, the configuration of the general support systems may be changed as a result of normal operations such as patch management, capital equipment replenishment, or any other systems put in place to keep dynamic IT systems working properly throughout their lifecycle.

Thus, it is up to management to determine the scope on which the incident response and disaster recovery planning effort will be focused, but the broader the range, the more likely the plan will be to succeed.

Incident Response & Disaster Recovery Planning Methodology

Overview

Three levels of incident response and disaster recovery planning were referred to in the introduction. Each is defined as follows.

Incident response describes the actions taken at the point where an incident is first detected and mostly has to do with how the root cause is determined, how the effects of the incident are limited, how the vulnerability is remediated, how the problem is overcome (if the remediation is short term), and the reporting system. In other words, incident response is focused on the actions of the user and the first responders to an incident. Typically, incident response has to do with a brief period of time measured in hours and typically less than one business day, even when eradication and reporting take longer.

Contingency planning describes the actions taken from the point of discovery of an incident after the incident is reported during the intervening period during which the incident is overcome. Typically, contingency planning has to do with the time from completion of incident response until the full scope of the incident can be defined and the appropriate organizations are notified of the incident ensuring that any required services are shut down and perhaps even restarted once the incident is corrected. The time covered by the contingency plan should not exceed the initial 72-hour period from the start of or report of the incident.

Continuity of Operations refers to the activities required to keep the enterprise running during a period of displacement or interruption of normal operations at the primary facility according to Stephen Fried, author of *Information Security: The Big Picture – Part IV*, as published in Information Security KickStart Highlights published by SANS for GIAC training in 2001. Under this guideline, the Continuity of Operations Plan is defined as the identification of crucial business processes, definition of how to restore them at a temporary location, and the recognition of resources needed for disaster recovery. Continuity of operations planning further includes the all-important steps to plan for productivity recovery. This is the process of rebuilding an enterprise's operation or infrastructure after the disaster has passed. It is conducting damage assessment, decision-making, operations recovery, data recovery, and restoration of productivity. Continuity of operations planning generally starts approximately 72 hours after interruption of service and continues until the operation is permanently capable of performing its mission and role.

The typical incident response process is made up of six steps as defined by the SANS Institute. The six steps are:

- Preparation
- Identification
- Containment
- Eradication
- Recovery
- Follow Up

Each of the steps to handle the various issues, which might be construed as incidents, must occur for the proper handling of an incident. A plan for each step must be completed before there is a need to actually use the steps. What this means is that management must understand the concept behind each procedure and know how each step coincides with the others.

A holistic view is the only way to prepare for an incident of any type. The following steps are cornerstones in all disaster recovery planning.

- **Preparation**, everything leads to preparation. The preparation phase consists of the planning, written processes and procedures, policies, standard practices, and training that occur before and after an incident. Lessons learned from all steps will be used to validate the preparation process, testing of the processes will reveal how well conceived they are, and adjustments learned from the testing will ensure the ability to smoothly implement the plans. If the preparation phase of incident response were skipped, the remaining steps would be random decisions without focus. The incident may be handled this time, but what about next time? Without planning, decisions may not be effective. By establishing policies, procedures, and training in advance, the organization is not only enabled with regard to the response being tailored to the issue, further incidents may be prevented. The best way to handle an incident, of course, is to prevent it.
- **Identification** is perhaps the simplest of phases. On a basic front, the identification of a problem is often obvious. It may be a virus alert, a fire alarm, or a power failure. All of these create a unique set of issues that must be dealt with.
- **Containment** may be overlooked in some cases. If a virus strikes one workstation out of a hundred, the system administrator may just fix the workstation, either through reloading from the baseline or by replacing the hard disk. However, at the same time, the system administrator might fail to realize the network is at risk or even the root cause of the problem. It is this containment that plays a very important role or else fixing the problem only results in recurrence and further cost associated with containment.
- **Eradication** is probably the most unique of all the procedures in incident handling. It is completely dependant on the cause and will vary greatly. It differs greatly depending on events from malicious codes types to natural disasters.
- **Recovery** is the process of returning to normal operations. It may be as simple as returning control of a workstation to a user for normal use or declaring the event as completed and allowing use of the system for normal operations.
- **Follow up** is the process of reporting on the incident, which may be as simple as filling out a form or completing an event log or if the incident was more complex, it may call for documenting lessons learned so that there is a knowledge base from which future similar incidents can be more easily identified and handled.

Types of Events that can be Categorized as Incidents

Three types of events are typically addressed under Federal guidelines. They sometimes are referred to as Type 1, Type 2, and Type 3 and are described as follows.

Type 1 incidents are primarily administrative or non-disruptive in nature and can be resolved at the enterprise or agency level. Examples include the inadvertent deletion of non-critical data by a user or a pinging incident on a system connected to the Internet.

Type 2 incidents include any intentional disruption of service, any successful penetration, any event that involves a system reported under other rules, or a series of related events (trend analysis). An important criterion for this type of incident is the intent of the person who originated the incident. If it appears the intent is not to cause harm to or to degrade Federal IT systems or activities (or those under Federal mandate), the incident should be reported as a Type 2 incident. Examples would be attempts (either failed or successful) to gain unauthorized access to a system or its data, such as password guessing; unwanted disruption or denial of service; the unauthorized use of a system for the processing or storage of data; and changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent.

Type 3 incidents could be construed as Type 2, but with the intent to harm or degrade IT systems or business workflow or that are of such a serious nature they should be reported to an outside investigative agency (e.g., local law enforcement, Federal Protective Service, FBI, Federal Computer Information Response Center [FedCIRC], or the National Infrastructure Protection Center [NIPC]). Generally speaking, a higher level authority will coordinate the efforts outside of the enterprise for all Type 3 incidents and will determine the proper reporting criteria, format, and jurisdictions. No contact outside enterprise should be made regarding a Type 3 incident without specific instructions from an appropriate higher level authority, which will confirm if an incident reported as a Type 3 is truly defined as a Type 3.

Maintaining efficient productivity is the key goal of incident response and disaster recovery planning. When are these plans warranted? What constitutes an incident that activates an incident response plan, a contingency plan, and a continuity of operations plan? There are various factors to be evaluated.

However, there are some examples of issues that may invoke the use of the plans. A brief, non-all-inclusive list of incident-causing events follows.

- Malicious Code – Virus, Worm, or the like
- Intentional Attacks – Hackers, Terrorists
- Human Error
- Misplaced Data, which may include lost, stolen, or misrouted data
- Natural Disasters
- Hardware Failure
- Software Failure or corruption
- Electrical Power or Other Utility Failure
- Voice or Data Communications Failure

Definition of Some of the Incident Causing Events

Malicious code is the most common event on a network. This type of event most likely occurs on a daily basis and is not revealed. Firewalls stop some hackers at the border of the network. Antivirus applications stop virus's before they reach an inbox in an e-mail system. But, on occasion, a virus or worm will make it past the automated detection tools and infect a computer. System Administrators must recognize the type of code involved, and understand what level of threat it conveys. After ensuring an understanding of the cause, the code can be removed safely and the system(s) restored.

Intentional attacks are usually the work of hackers, although burglars and terrorists also fall into this category. Intruders will use many methods to gain entrance into an enterprise, it just depends what they are after. If they are looking for notoriety, it may be defacing the enterprise's Website. If the intruder(s) is after money, the intruder might look for sensitive or classified information (e.g., an intruder may steal data and hold it for ransom). These are events that can be planned for in advance. Like the human factor, most intrusions can be avoided and repaired, if the resulting events are treated first.

Human error is the least controllable category. Training and competence can only go so far. Preparing for this type of event is difficult, but not impossible. Once the resulting event is treated, things can be restored much quicker.

Misplaced data can have severe repercussions in today's online world. With the recent theft of a laptop from Fidelity Investments (as documented in *Stolen Fidelity Computer Raises Privacy Fears* published on March 23rd 2006 by MSNBC online at Universal Resource Locator [URL] <http://www.msnbc.msn.com/id/11974062/>), data loss can result in a massive disaster recovery effort. This particular laptop contained personal information on more than 196,000 people.

Natural disasters of recent history have caused a devastating effect on governments and businesses. Those without a back-up plan had more issues rebuilding than those who had a solid plan. An enterprise must decide what is important for the business to survive and how to mitigate vulnerabilities when recovering.

Hardware failure is an expansive category including everything from personal computers to fax machines to network communication devices. The response to each will be a bit different. A computer sitting at the desk of an administrative assistant will be treated differently than one sitting at the desk of the contract manager. The requirements are different for everyone in the enterprise and any incident response or disaster recovery plan must take priorities into account.

Software failure or corruption is another area that must be accounted for in all levels of planning. Is all of the enterprise's software properly licensed? Is any of the software subscription based and have a short expiry period? Legal copies of all software need to be reviewed to ensure compliance and that no work stoppage occurs. It also is a good plan to have an offsite storage facility with copies of software, albeit not just for software failures.

Electrical power failures disrupt the lifeblood for advanced electronics. Without it, an enterprise might be forced to send employees home and time is lost from productivity. Similarly, any utility provided to a facility could result in lost productivity and chaotic operations. This includes heat, telephone service, or any other supporting service provided as a utility.

Voice or data communications failures are grouped together because they are so close in usage. Enterprises speak and share data over the Internet or via any direct wired or wireless connections. The telephone lines are not only voice, but carry data signals as well. Whether the problem is with satellites, a provider, or a connection at the enterprise's facility matters little when the communications are down.

When to Implement a Contingency Plan of any Type

Once an issue that requires activating a plan is identified, the procedure must include the recovery process and the point at which an incident is escalated. The timing is not a cut-and-dried circumstance; it depends on the problem, which is one of the reasons that incident response is divided into the three types of planning and actions.

By identifying the impact on the enterprise, it is easier to discover which type of plan to implement. The following categories provide a basic interpretation of how to define the issue at hand. The times are cited in accordance with hours in a business day so that those organizations that are not operational 24 hours per day, seven days a week can have some idea of how to deal with incidents.

- Category 1 – Slight degradation of resources (less than 4-hour recovery)
- Category 2 – Resources degraded (less than 8-hour recovery)
- Category 3 – Facility or resources denied (less than 24-hour recovery)
- Category 4 – Major catastrophe; natural disaster (more than 24-hour recovery)

Who is Involved

All employees direct or contract, are responsible for incident response and disaster recovery. "Commitment" is the major factor to creating and maintaining a good disaster recovery planning. Management at all levels must demonstrate their commitment to incident response and disaster recovery at all times.

Middle management is responsible for managing the efforts involved in incident response and disaster recovery planning. There are some immediate response mechanisms that management will not have direct authority over as they happen rather quickly. For all non-

immediate activities, middle managers are responsible for ensuring the implementation and operational activities included with those efforts. The executive management team should provide support to managers by ensuring they have the resources needed to continue supporting external customers.

Employees should understand what their tasks are during incident response and disaster recovery operations. If an employee's only task is to stay at his or her desk, the employee should know and understand why. This is the reason for sharing plans and running active tests of the plans. Employees also are the key players in discovering a condition that may lead to activating incident response actions and must realize that they have a duty to report any suspicious occurrences. Knowing the signs of a potential problem are very important as are reporting procedures. This is one of the key reasons for a separate incident response guide.

A computer security incident is an event that harms data, systems, or networks; or that otherwise hinders the ability of a system and network to continue operations as intended. However, the scope of a computer security incident does not include end-of-life cycle hardware failures, software misconfigurations, or user accidents. The following are examples of events that are not part of a computer incident.

- A router power supply fails and causes subnet outage.
- A raid controller card fails causing server storage problems.
- An exchange server is not able to replicate information stored due to an incorrectly configured subnet mask, which causes e-mail outage for some users.
- A user accidentally unplugs a power cable for a switch causing network outage.

The following are some examples of the events considered computer security incidents.

- Discovery of a strange process running that consume extensive processor cycles.
- A mysterious system crash.
- Discovery of an intruder logged into the system, new user accounts, or increased activity on previously unused accounts.
- Discovery of a virus infecting the system.
- Discovery of suspicious browsing or network probes.
- Denial of system resources, data modifications, deletions, etc.

There are seven major categories of incidents identified for which handling procedures are explicitly defined; however for all security breaches, the same reporting and handling procedures shall be followed.

The seven major incident categories are:

1. Insider attack
2. Denial of Service
3. Virus
4. Worm
5. Trojan horse
6. Intruder attack
7. Physical asset

It is imperative to accurately report an incident in order to understand the extent and source of the security threat. Without an indication of the scope and impact of the event, it is difficult to determine a correct response. The following are some of the issues to report.

- Is this a multi-site incident?
- How many systems are affected at each site?
- Is sensitive information involved?
- What is the entry point of the incident (e.g., network, phone line, local terminal, etc.)?
- What is the potential damage?

- What is the estimated time for recovery?
- What resources are required to handle the incident?

Process Components

General

Each enterprise that is highly involved with IT will have experience performing risk analysis and from the view of this white paper, that is a necessary assumption. Risk analysis should be used to determine the types of incidents that are likely to occur and the amount of risk attributed to each. For information about risk analysis, refer to the Federal Network Services, INFOSEC Services white paper *Risk Analysis: How Risk Analysis can be Conducted to Establish Mitigation Requirements and Improve System Operations*.

In any situation that invokes an incident response, contingency, or disaster recovery effort, there are priorities that must be considered. The priorities are as follows:

- Priority 1 – Protect Human Life & Safety (always has precedence over other considerations)
- Priority 2 – Protect Sensitive Data
- Priority 3 – Protect Other Data
- Priority 4 – Prevent Damage to Systems
- Priority 5 – Minimize Disruption of Resources

The rest of this section describes how to implement an incident response and disaster recovery planning effort in very general terms.

Documenting Procedures

If the enterprise has an active policy and procedure system, the initial effort should be to review the policies and procedures to determine how comprehensive the procedures are at establishing steps to take during an incident. If the system is found to be lacking, selected items should be modified to allow for the three documents to guide incident response and disaster recovery efforts.

Incident Response Guide Development

The intent of the Incident Response Guide is to act as a guide for first responders and personnel who identify a suspected incident in determining what actions to take in the critical initial period. Users will ask questions like:

- Should I turn off the computer?
- Should the computer be unplugged from the network?
- Is there anything I can do to keep this from spreading?
- What can I do to help?

Those who take the initial calls may not know what the best answers to the above questions are, and that is to be expected. There are various types of incidents that need to be diagnosed and the diagnosis will often guide what actions should be taken, because the answers to the questions above will differ based on the initial diagnosis.

Thus, the Incident Response Guide describes the types of incidents that can be expected and provides assistance in making diagnoses that are accurate. Additionally, the Incident Response Guide should provide documented steps to reduce the effects of the problem until it can be fully resolved and eradicated.

Finally, the Incident Response Guide should provide for a Computer Emergency Response Team (CERT) of some type depending on the size of the IT portion of the enterprise and provide guidance for how to report and record information about the incident. The CERT will be responsible for resolving the problem and returning to normal operations, when further steps are not required as the result of a contingency-invoking event. Whether the instructions are included in the Incident Response Guide or in developed, published policies and procedures, instruction must be prepared that defines when an event is moved from a security incident to a contingency-invoking event. This is the initial escalation to be discussed in the incident response and disaster recovery planning process.

Contingency Plan Development

The Contingency Plan covers the initial one to three days of response to an event that shuts down operations. Typically, this type of event can result in a return to normal operations without long-term negative events. An example of a short-term contingency event is the loss of power as a result of a downed power grid that supplies power to a facility and remains down for a number of hours, even extending into a brief period of days.

The purpose of a contingency plan is to ensure that the proper authorities are contacted, that the facility is protected as needed, and that the return to operations is conducted in a manner that ensures the safety and security of the employees.

At a minimum, a contingency plan should contain the following types of information.

- System Description
- Points of Contact
 - For systems
 - For associated government agents
 - For emergency response agencies
 - For reporting contingency events
- Instructions on system shutdowns
- Instructions on facility shutdown
- Instructions on system restarts
- Instructions on determining that return to the facility is safe for employees
- Instructions for re-opening the facility
- Emergency cut-over procedures for short-term durations

Whether the instructions are included in the Contingency Plan or in developed, published policies and procedures, instruction must be prepared that defines when an event is moved from a contingency-invoking event to disaster status. This is a critical escalation to be discussed in the incident response and disaster recovery planning process.

Continuity of Operations Plan Development

The Continuity of Operations Plan covers the period of time from determination that a contingency event can be classified as a disaster until normal operations can be returned to the original facility or housed in a new facility if the former facility cannot be used again. An example of a disaster is the residue of a natural disaster such as a hurricane that forces operations from its normal residence as a result of damage to the facility or its contents.

The purpose of a continuity of operations plan is to define how business operations will be continued despite the loss of facility or any other disruption that forces operations to be stopped or moved for an extended period of time (usually more than three days and often a longer period measured in weeks or months).

At a minimum, a continuity of operations plan should contain the following.

- Organizational and IT System Descriptions
- Criteria for determining how the Continuity of Operations Plan is invoked
- Primary management responsibilities during the emergency operations period
- Points of Contact
 - Coordinators for all major functions during off-site operations
 - For systems, both primary and back up
 - For associated government agents
 - For emergency response agencies
 - For personnel permanent assigned to the back up facility
 - For all sites and off-site storage facilities
- Instructions on system shutdowns
- Instructions on facility shutdown
- Instructions on system restarts
- Instructions on determining that return to the facility is safe for employees
- Instructions for re-opening the facility
- Policies and procedures for operating from the back-up facility as the primary location
- Plans for moving staff to other locations in support of continued operations
- Management plans for the back-up facility and creation of a tertiary back-up facility
- Logistic plans for supporting out locations
- Service plans for ensuring traveling service personnel are supported during the change
- Crisis communications plan
- Public affairs plan for operation from the back-up facility
- Recovery plans for repair or rebuilding of the primary facility
- Hiring plan for any additional resources needed at the back-up facility
- Procurement plan for assets needed at the primary and back-up facility during moved operations and move back to the primary site
- Management communications plan for the differences caused by disruption and the move

Testing the Plans

All of the planning in the world will not help in the least if the plan cannot be implemented! The only way to know for sure that the plan can be implemented is to test it. That can be accomplished through a series of drills that test each portion of the plan. It is not necessary to stop all workflow in order to test parts, but testing the plan is a necessary part of planning, just to ensure that the processes can, in fact, be followed.

Summary

Management has a responsibility to the various publics it serves to ensure that operations can continue in the case of any type of incident. The responsibility can be fulfilled by creating an environment in which plans for various types of incidents are documented and can be followed. This white paper brings together a number of issues about incident response and disaster recovery that amount to required actions and those that follow from best practice throughout Government and industry. A critical facet of incident response and disaster recovery planning is to ensure that the systems that were impacted by any type of incident are returned to normal operations in the condition they were in when the incident started. Only then should vulnerabilities be patched and managed as a part of the overall INFOSEC process of the enterprise. Thus, one of the key goals is to ensure that there is an adequate portrayal of the system at all times, so that the status of each node of the system (a benchmark that preceded

the incident) can be recreated. The reason for doing this is that any change to the status that existed is cause to infer the potential for additional vulnerabilities. This was one of the reasons that so many systems had problems recovering their nodes that were impacted by Hurricane Katrina in the Gulf Coast states and many other natural disasters that have occurred.

The key to a successful IT security program lies in how well systems are documented and evaluated with regard to security issues. Federal Network Systems' INFOSEC Services Division provides a holistic approach to INFOSEC in general and specifically to C&A efforts for Government and Federally mandated programs to ensure IT systems can be operated safely and securely. In addition, INFOSEC Services provides services that offer Return-On-Investment (ROI) features to increase safe, secure operations of IT systems and to provide for better system protection throughout the life cycle of the system.

For more information, refer to the contact information provided on our Website: www.fnsnet.com and a representative will be happy to provide more information.