



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at <http://www.giac.org/registration/gsec>

Disaster Recovery in Healthcare Organizations: The Impact of HIPAA Security

James C. Murphy

November 24, 2003

GSEC Practical Assignment, Version 1.4b

Option 1

© SANS Institute 2004, Author retains full rights.

Disaster Recovery in Healthcare Organizations: The Impact of HIPAA Security

Table of Contents

Abstract	1
Introduction.....	1
The Ecology of Disaster Recovery	2
Focus on information	3
Security disciplines	3
Disaster as an event	5
HIPAA and Disaster Recovery	6
Requirements of the Rule	7
Limitations of the rule	9
Designing the Disaster Recovery Plan	10
Distributed Disaster Definition	10
First Things	11
Disaster Mitigation Preparation.....	11
Pre-Disaster Planning	13
Declaration	18
Recovery	18
Testing	19
Conclusions	20
Resources	20
References cited	21

Disaster Recovery in Healthcare Organizations: The Impact of HIPAA Security

James C. Murphy – GSEC Practical Assignment, v. 1.4b – November 24, 2003

Abstract

Healthcare organizations are facing increasing regulatory burdens, and the latest to demand response is HIPAA Security. One major aspect of HIPAA Security is the disaster recovery plan, which seeks to restore appropriate access to information after a major calamity. Disaster recovery has a place among other organizational security processes, including information security in general, physical security, and business continuity. Each of these security disciplines also is differentiated based on focus and organizational responsibility. Disaster recovery focuses on information, and within healthcare organizations, the focus of HIPAA disaster recovery is the electronic protected health information. This does mean that a strict HIPAA disaster recovery plan will be inadequate, since non-electronic information needs protection as well. The disaster recovery plan begins with modifying management practices to mitigate the effects of disaster, then documenting all elements of a distributed computing environment, including policies, procedures, infrastructure technology, and applications. The plan will address the requirements for a recovery location and a recovery environment and steps to take to set up the recovery environment and implement the applications. Finally, the plan will also include steps to return to the original location after reconstruction.

Introduction

Picture this scenario: You are the administrator of a small healthcare facility on the Eastern United States seaboard. The latest and greatest named Atlantic Storm rears its ugly head and crashes down on you, your facility, and your community. Power is lost, flooding disrupts and corrupts the interior of your facility, and all your desktop computers are damaged beyond repair. Because the patients are being cared for and relocated to new facilities by emergency medical staff and volunteers, thankfully no one suffers long-term effects of the storm. Before you completely assess the damage, a family member of one of your patients makes contact with a simple request, based on authorized permission, that you produce a copy of the patient's medical records that were stored on your desktop workstations. Your answer is easy – the computers are damaged, we will have to see if the data can be recovered off the hard drives. The data may be un-recoverable. This leaves the family member unsettled and concerned about ultimate proper treatment. Before the storm, your facility was in full compliance with HIPAA Privacy, all of your Protected Health Information was safe, policies governing use and disclosure were documented and in force, and you were able to produce records for authorized family members within the authorized notice period. With the destruction of the facility and the computer records, you expect that this "act of nature" and your logs of the disaster's destruction will provide your escape from liability in this circumstance. On the contrary, if you do not have an adequate backup of your information and a properly conducted Disaster Recovery Plan, you stand to be in violation of HIPAA Security!

This scenario may not be far-fetched, especially within the areas of the Southeastern US where hurricanes have wreaked havoc over the last 8-10 years, including this past September 2003, with the invasion of Hurricane Isabel. The Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule has defined what kind of information has to be protected and under what circumstances protected information may be used or disclosed. HIPAA Security details *how* the information, initially only the electronic versions, is to be protected. In fact, strict interpretation of HIPAA actually means that HIPAA Privacy compliance is not complete without compliance to HIPAA Security. The portion of the Privacy Rule that contains instructions about security, often called the “mini” security rule, are only a foretaste of the actual [Security Rule](#), which has now been finalized.

Disaster recovery as a process has had a long history and is not far from our experience in the current Information Age, while we struggle through the natural and man-made disasters and heightened tensions that have plagued our society in the last few years. Recent disastrous events have heightened the awareness and the need for business continuity and disaster recovery (Tissot, 2002), which should begin to address the problems identified by O'Brien (1999), prior to the tragedies of September 11, 2001, who found in an IBM survey that “92 per cent of Internet businesses are not prepared for a computer system disaster.” And Kelly (2001) reported a TechRepublic survey of IT professionals of whom “87% say their firms' IT systems lack the redundancies and/or protection in case of emergencies.” Brown (2003) has found that the tragedies of the last few years, including weather disasters as well as the 9/11 tragedies has begun to take action, but with some degree of reluctance. Webster (2002) also found that some IT executives were treating disaster recovery as an insurance policy that was too expensive to afford. Surprisingly, his main source was the manager of a hospital who received such feedback from the hospital's administration! Part of the problem may be that there is still ambiguity about the nature of the disaster recovery process and how it relates to other organizational challenges such as security and business continuity. Freeman (2002) actually described Business continuity as a subset of disaster recovery.

In this document I will describe the relationship of disaster recovery to business continuity and other organizational security disciplines. I will also identify the unique challenges that healthcare organizations face with disaster recovery, not the least of which is the regulatory burden of HIPAA Security. Finally, I will describe a methodology for developing and implementing a disaster recovery plan within healthcare environments.


The Ecology of Disaster Recovery

The realities of the current environmental and political climates have brought the business and industrial world to the point that information disasters are not the exception, but unfortunately the rule! It is not a matter of *if* a disaster will strike, but *when!* And, since more and more data resides on organizational networks for access and processing, the risk of loss has been increasing (Tissot, 2002). The types of

disasters are not difficult to discover - naturally occurring circumstances such as winds, water, snow, earthquakes and lightning; and human-caused circumstances, either accidental or deliberate. Brown (2003) lists power failures and surges the most frequent source of disasters (more than 33%), followed by storm damage (20%), floods (16%), and fires and bombs (9%). The information-centered reality of doing business in the 21st century means that companies lacking appropriate business continuity and disaster recovery plans will eventually find it selves out of business when disaster strikes (Brown, 2003).

The definition of ecology – the study of the home, or place, and the relationships therein – is well understood by biologists. Disaster recovery also has a “place” within an organizational environment, which is often not completely appreciated or understood or appropriately identified. Disaster recovery also has relationships among other security related business disciplines and a disaster itself is part of a continuum of events that disrupt information flow

Figure 1.



Focus on information

First and foremost, disaster recovery is targeted at the organizational information resource (Figure 1). The primary function of a disaster recovery plan is to rebuild the information technology resource to provide access to the necessary information immediately after a major disaster or other business interruption, and to set up a replacement resource in a temporary location for an indefinite time if in fact the original location was damaged. As it is developed, a thorough disaster recovery plan takes all information systems and collections into account. However, a healthcare organization is required by HIPAA Privacy to pay strict attention to protected health information (PHI), which means paper and electronic information. HIPAA Security narrows the focus even further, requiring attention only to electronic protected health information (EPHI). HIPAA Security requires that the organizational disaster recovery plan address EPHI, but the plan itself would be incomplete without addressing all of the information resource. I will return to HIPAA’s disaster recovery expectations later.

Security disciplines

Disaster recovery is one of a number of interrelated and overlapping business disciplines involved in the protection of corporate assets (Figure 2). Information security defines the structure and practices for protecting the corporate information resource. This includes the hardware and network infrastructure requirements, the software configuration, the processes of administration and management, the practices of the technology users and managers, and the documentation for all the above. Information

security defines the terms and concepts for which a disaster recovery plan is developed to address. (Table 1) Disaster recovery is an obvious subset of this discipline, and a disaster recovery plan is built from the organizational security practices and documentation.

Physical security defines the protection of organizational assets including but also beyond the information resource infrastructure. These include property, personnel, facilities, and the documented practices employed while actively working within the facilities. Disaster recovery is also a subset of this discipline as well, since the plan takes into account structural resources for protection of information before disasters and those necessary for recovering the information after a disaster. Disaster recovery itself is often confused with the larger effort of business continuity – also referred to as *business resumption planning* or *contingency planning* – that documents the process for continuing critical business functionality during and after a disaster. Bahan (2003) refers to the business continuity plan as “...an umbrella plan whose major subcomponents include the Disaster Recovery Plan.” In addition to managing the information resource, business continuity planning takes into consideration the workplace temporary relocation; accounting functions – receivable, payable, and payroll; and the interaction with the public – publicity, clients and customers, debtors and creditors – during and after the disaster. The business continuity plan also includes the facility emergency response to the disaster involving personnel safety and evacuation (Meyer, 2002), which is distinct from the information disaster recovery plan.

Figure 2.



TABLE 1. INFORMATION SECURITY DEFINITIONS	
Confidentiality	Information protected from unintended disclosure - secrecy, privacy.
Integrity	Maintained in an unimpaired condition; kept to specifications - consistency, authenticity.
Availability	Usable when needed - enhanced by redundancy, prevention of denial of service.
Accessibility	Information resources available appropriately - possession.
Timeliness	Resources available when needed - real-time systems.
Utility	Usable for intended purpose.
Vulnerability	Weakness; potential for problems in any of the above areas.
Threat	Danger that vulnerability can lead to undesirable consequences.
Risk	Harm from threats, or measure of extent of the harm or loss of value.

In actual disaster situations, business continuity and disaster recovery plans are incomplete without each other. A disaster recovery plan that restores information access is not enough to bring the overall organization back into operation. A business continuity plan that relocates employees and organizational functions and does not bring back information is incomplete.

In a healthcare environment, business continuity planning has additional challenges based on the unique nature of the industry. Besides the common

business functions and information requirements listed above, healthcare business

continuity plans have to address resident patients and any critical life support technology and structures associated with the care of the patients. This patient aspect of healthcare organizational continuity is obviously the first priority in bringing back business functionality.

Each of these security-related disciplines differs in the focus of responsibility for building and implementing plans, even though all levels of the organization will participate and provide input (Table 2). Business continuity is the responsibility of the executive level,

Business Continuity	Executive level
Physical Security	Corporate security
Disaster Recovery	Technology staff
Information Security	All personnel

because of the decisions and expenses that will need to be made. In larger organizations, physical security is the responsibility of corporate security personnel, either as employees or from a contracted agency. Since the disaster recovery plan is concerned with information and the technology, the organizational technical staff has the primary responsibility for

developing and implementing the plan.

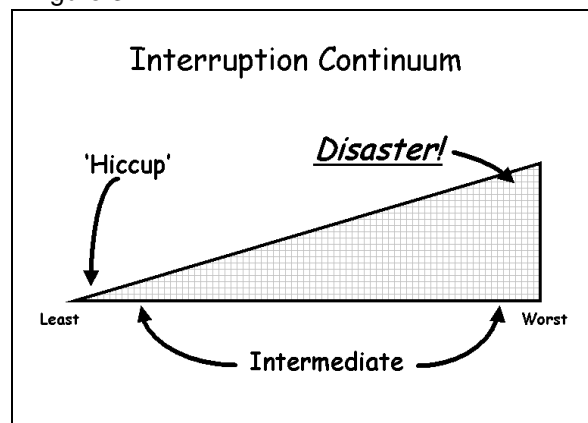
Information security in general is a responsibility of all levels of the organization - all employees and contractors who use the network resources, internal and external to the physical location. Users of the network are the source for the majority of vulnerabilities to organizational information security.

Disaster as an event

In Figure 3, we can see that an information disaster is within a continuum of interruptions that affect access to information. The vertical axis indicates severity of the interruption and the horizontal axis is a ranking of interruptions from least severe to most, or worst depending on the nature and overall cost of the interruption in terms of time, access to information, organizational productivity, and ultimately money. At the low end of the continuum are 'hiccups', such incidents as network or electrical cable breaks, plugs pulled, power fluctuations, or accidental file deletions. These interruptions are local to individual users, but to the naive user, they temporarily interfere with the flow of or access to necessary information. Usually, these incidents can be resolved in minutes with no actual loss of time or productivity.

In the middle range of interruptions, we can include failed components of the desktop workstation, network devices or the servers themselves. These components include disk drives, network interface cards, power supplies, processor boards, or any other device that results in a relatively local interruption. These incidents can usually be repaired or replaced in a matter of hours or days, and the effects are relatively local, depending on the number of users involved.

Figure 3.



None of these interruptions constitute an actual disaster, which is at the high end of the continuum and affects relatively large numbers of users, if not the whole organization, for a much longer period of time. A disaster also requires a formalized plan of action to restore services, as opposed to a simple repair or replacement of temporarily disabled devices. Later in this paper I will provide a working definition of a disaster, which will help clarify the development of a disaster recovery plan. In the distributed network world in which we now operate, data and information are no longer restricted within a centralized computer room on a large mainframe, therefore the terms of reference for developing a disaster recovery plan are different.

HIPAA and Disaster Recovery

Any exercise in information security is expensive and inconvenient. When access is restricted and networks are filtered, legitimate work is slowed, which can often be translated to monetary cost. However, organizational information and the technology surrounding it has become the foundation for organizational persistence through time. Protecting the information resource is expensive, but much less expensive than the wholesale loss of the organization after a major disaster. The value and importance of organizational data and information is often overlooked when considering the need for disaster recovery protection. Webster (2002) equates the importance of money *and* data to the survival of an enterprise after a crisis.

Also, as Bogen (2002) describes, an inadequate plan can harm patients from loss of information, organizational reputation, and ultimately public mistrust (see also Widup, 2003). As the introductory scenario illustrated, HIPAA brings even more of a monetary risk to the protection of information. If adequate protection according to HIPAA is not initiated and maintained, then the potential is set for compounding monetary loss with monetary penalties, either from regulatory fines or potentially from civil proceedings by individuals or their families. Major healthcare accreditation organizations, such as the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) and the Accreditation Association for Ambulatory Health Care (AAAHC) require protection of health information afforded by a disaster recovery plan (Spath, 2002; Burlington-Brown and Hughes, 2003). And, the American Health Information Management Association recommends a disaster recovery plan for all health records and information for Long Term healthcare facilities (Dougherty, 2001).

The HIPAA Privacy rule has already raised emotions and some resistance from healthcare organizations and individuals seeking medical care (Box 1). For many who have begun the implementation of HIPAA Privacy, and are attending to the other

Box 1. HIPAA in the News

“Hipaa [sic] is what happens when the federal government gets its tentacles into a problem cited by a few people. The impetus for all this cost and futile nonsense was one incident in which someone's medical records were accidentally revealed. The incident ends up repeated with outrage twelve times in the Congressional record. Now we all stand behind yellow lines in a flurry of Hipaa [sic] paperwork.”

Marianne M. Jennings

[HIPAA: The federal government strikes health care again.](#)

Jewish World Review 7/25/03

regulations involving data transactions, the HIPAA Security rule is not yet on the active horizon. The rule was finalized in February 2003, and allowing for two months notice and two-year implementation period, the rule will take effect in April 2005. To many, this time frame may allow the development of inaccurate expectations, since the plans for responding to HIPAA Security may take more time and expense than anticipated.

Requirements of the Rule

HIPAA Security has built in a measure of flexibility to allow for variability in responses to some standards and implementation specifications of the rule within smaller healthcare organizations. Many of the standards are required, allowing no flexibility; all organizations must implement the required standards and implementation specifications. There are a number of standards that are addressable, which allows each organization some flexibility in determining how to respond. An organization can choose to:

- Implement one or more of the specifications
- Implement one or more of the alternatives
- Implement a combination of specifications or alternatives
- Implement none of the specifications

The rule is clear; the selection of the response for any addressable standard or specification must be based on the initial Security Risk Analysis. This allows for justifying the responses based on the cost, or impact of the risk involved with the addressable standards (Box 2). The HIPAA Security rule has targeted *electronic* PHI, as indicated in the text for the Risk Analysis. Within the introductory text, the rule explains that though the Privacy Rule gives attention to paper as well as electronic PHI, the creators of the rule have essentially tabled requirements for paper records. This means that all the standards and implementation specifications direct a narrower view of security for an organization, since in reality paper information requires protection as well.

Box 2. Risk Analysis

§164.308 Administrative safeguards

(a)(1)(i) **Standard: Security management.**

(ii) **Implementation specifications.**

(A) **Risk Analysis (Required).**

Conduct an accurate and thorough assessment of the potential **risks and vulnerabilities** to the **confidentiality, integrity, and availability** of **electronic** protected health information held by the covered entity.

All of HIPAA Security compliance begins with the Risk Analysis. Within the Administrative Safeguards portion, it is the first implementation specification in the first standard, and it is required. The ability to select alternatives to the Addressable

specifications depends on the justification based on the Risk Analysis. There are many sources, from government and private organizations, for risk analysis methodology, all of which describe gathering organizational vulnerabilities to damage or loss of EPHI and assessing the specific threats that would exploit the vulnerabilities. The appropriate combining of vulnerabilities and threats will provide a measure of risk. Ultimately, the impact of each risk can be evaluated to create a sequence of risks ranked by severity.

Box 3. Contingency Plan

§164.308 Administrative safeguards

(a)(7)(i) **Standard: Contingency Plan.**

(ii) **Implementation specifications:**

(A) **Data backup plan (Required).** Establish and implement procedures to create and maintain retrievable exact copies of electronic protected health information.

(B) **Disaster recovery plan (Required).** Establish (and implement as needed) procedures to restore any loss of data.

(C) **Emergency mode operations plan (Required).** Establish (and implement as needed) procedures to enable continuation of critical business processes for protection of the security of electronic protected health information while operating in emergency mode.

(D) **Testing and revision procedures (Addressable).** Implement procedures for periodic testing and revision of contingency plans.

The portions of the HIPAA Security Rule pertaining to disaster recovery are brief and to the point (Box 3). There are four specifications, quite simply and tersely stated, within Administrative Safeguards under the Standard called *Contingency Plan*. Specification 'A' refers to the backups of electronic data, 'B' refers to recovering that backed up data when necessary, 'C' refers to a plan for providing access to data during an emergency situation, and 'D' refers to the testing of all of the contingency specifications. Note that 'A', 'B', and 'C' are Required, while 'D' is Addressable.

These are the only portions of the Security Rule that refer to disaster recovery. In the section called *Physical Safeguards*, following the Administrative Safeguards, under a Standard called *Facility access controls*; an implementation specification addresses access to a facility during or shortly after a disaster to attempt to retrieve resident data. This specification only refers to controlling appropriate access to the facility and is not actually part of the disaster recovery plan itself (Box 4).

As mentioned above, the specific specifications applying to disaster recovery in the Security Rule are brief and to the point. This allows a great deal of flexibility in how the specific plans are designed and carried out, and removes a great deal of the ambiguity in the Proposed rule, which had disaster recovery specifications in two different sections. However, it would perhaps be more appropriate if the overall Standard were called '*Disaster Recovery Plan*' instead of '*Contingency Plan*', and specification 'B' were called '*Data recovery plan*' instead of '*Disaster recovery plan*'. This is because the four implementation specifications appropriately constitute elements of the actual disaster recovery plan, not the more inclusive organizational contingency, or business continuity plan, which is not specifically addressed within HIPAA.

Box 4. Contingency Operations

§164.310 – Physical safeguards

(a)(1) **Standard: Facility access controls.**

(2) **Implementation specifications:**

(i) **Contingency operations (Addressable).**

Establish (and implement as needed) procedures that allow facility access in support of restoration of lost data under the disaster recovery plan and emergency mode operations plan in the event of an emergency.

Limitations of the rule

It must also be pointed out that if strictly followed, the HIPAA Security disaster recovery specifications will result in an inadequate and incomplete disaster recovery plan. This is because:

- The Security Rule does not formally address business continuity, apart from implication based on the overall context of the rule. The use of the words 'Contingency Plan' within the Standard is perhaps misleading.
- The Security Rule only addresses EPHI; an appropriate disaster recovery plan needs to address paper records and other non-electronic information.
- HIPAA itself only addresses PHI – there is certainly more organizational information.
- The specification 'D', Testing and revision procedures, is only Addressable, and a disaster recovery plan without required testing is inadequate.

In order to address the organizational responsibilities delegated to them, the Information Technology management will be compelled to develop disaster recovery plans and tests beyond the specific requirements of HIPAA Security. This also includes participating in the overall business continuity and physical security planning as well. Adhering strictly to the requirements within HIPAA Security will not address all of the organizational needs. Implementing this broader set of plans will most assuredly address the HIPAA Security requirements as well.

HIPAA also requires careful documentation of the processes involved in information security management. Virtually every standard and implementation specification begins with "establish and implement procedures...". The processes that correspond to the standards and implementation procedures are defined by policy statements and carried out according to procedure statements (Table 3).

Process	Business Event
Policy	Reason for the Business Event
Procedure	Methods for carrying out the Business Event.

These policies and procedures are also important to the development of a disaster recovery plan and vital to the recovery, in the event that key personnel become lost. Most larger organizations will have existing documented policies and procedures, which is part of "best practices" for information technology departments. Smaller organizations will likely have repeatable processes for technology management, but may not have them documented. Lack of documented policies and procedures constitute a major vulnerability in the overall security protection of EPHI. A careful Risk Analysis will begin with an examination and review of the documentation or policies and procedures governing information security as well as the documentation of policies and procedures addressing employee awareness and training of security responsibilities.

Designing the Disaster Recovery Plan

Distributed Disaster Definition

Most modern networks are based on the historic client-server model of distributed servers and desktop workstations located throughout the organization. In a past article (Murphy, 1995), I identified specific distinctions in disaster recovery planning in client-server environments from main frame environments. Though concepts are similar in this paper, the information technology world has grown since then.

Some locations have multiple workstations within a single office or laboratory. Servers are not necessarily restricted to centralized computer rooms, and even the descendants of the main frame have been reduced in size. This reflects the reality of the aphorism “the network is the computer.” Because of high-speed local and wide area networks, large organizations can afford to be spread across campuses or cities, and users from any organizational unit can easily access data from any other unit. The reality of widely distributed information does certainly complicate the overall disaster recover planning. Also, accessing data by wired or wireless devices creates more overall exposure to the EPHI targeted by HIPAA. Box 5 lists the elements of such a distributed environment that are subject to a disaster and need to be addressed in the analysis and planning.

Box 5. Elements of a Distributed Information Environment.

1. **The server(s)**, e.g., production, development, database, etc.
2. **The client environment**, PCs, other workstations, wireless devices
3. **The software environment**, databases or other applications
4. **The backup environment**, acquiring and maintaining file and database backups
5. **The print environment**, a print server and one or more printers, scanners, etc.
6. **The network environment**, cables, switches, routers, etc.
7. **The non-electronic record environment**, medical records, notes, film
8. **The people environment**, users responsible for collecting and processing information.

It is not as easy to define the term 'disaster' in a distributed environment as it once was for centralized systems. Martin (2002) describes a disaster as “any event that can cause a significant disruption in operational and/or computer processing capabilities for a period of time...” That is sufficient for a small organization with a small number of servers in relatively close possibility, but not for a large campus environment with distributed technology across a wide area network. None of the lesser interruptions described above in the continuum (Figure 3) constitute a disaster by themselves, even though the service may be disrupted for a few days. A

total failure of an individual client, printer or server will not constitute a disaster, even if all components of the main server are totally lost, since each component can be replaced within a day or two. Therefore, in most cases, a disaster in a distributed environment can be defined as:

The calamitous loss of more than one (or all) of the elements at the same time,

This could mean the loss of the computer room housing the servers and the network hub, the loss of most or all the client workstations and the printer environment, which

may be in a different building, or worse, the loss of both locations and all the individual elements (Murphy, 1995).

It must also be clarified that in a distributed environment, true information disasters may be *localized* – one organizational unit may be totally wiped out by fire, while the rest of the organization may be untouched. This incident must still be considered a disaster for the purpose of setting off the recovery plans, even though the rest of the organization is still able to function. The obvious deduction to make is that in truly distributed environments, *multiple disaster recovery plans must be developed*.

First Things

Disaster recovery plan preparation begins with and relies upon executive level approval. Spending decisions about extra equipment, external disaster locations, and staff time spent in preparation, will have to be made by those with authority. Tragically, it often requires a major disaster at a separate organization to trigger executive involvement. Even then, challenges persist (Webster, 2002). With executive approval, a disaster recovery team can be assembled, and the logical choice for the team leader and an alternative are individuals from the technical support group. All organizational units should be represented by active users of each system that will be recovered after a disaster. Additional technical staff representing different aspects of technology support should be added as well. For the sake of the plan itself, all team members must include home and work contact information, which must be periodically updated.

The initial responsibility of the team will be to interview users responsible for an impact assessment of major systems that will require recovery after a disaster (Spath, 2002).. Members of the committee representing the unit systems may be able to supply the needed information. Ultimately, decisions will have to be made about the priority sequence of restoration of each of the organizational unit systems. The technical support staff will translate the application systems requirements into appropriate network, server and client requirements for the recovery mode environment.

The disaster recovery team will also conduct the initial information risk analysis, which for healthcare organizations, will follow the HIPAA Security Rule outline. Some of this information may exist within a larger Corporate Risk Analysis, but for HIPAA compliance, this risk analysis must be focused on the EPHI. This risk analysis will provide the basis for determining the requirements of the disaster recovery plan. The first time an analysis is performed will identify the most glaring vulnerabilities and potential risk. Subsequent periodic risk analysis updates will provide a review of the improvements.

Disaster Mitigation Preparation

The disaster recovery team can be the impetus for stimulating operational changes within the technical support function. One basic underlying perspective for anticipating loss from a disaster is the overall reduction of vulnerabilities leading to the disruption of

information flow as identified by the risk analysis (Martin, 2002). Of course, the most severe interruptions are those classified as disasters, but others, as illustrated by the interruption continuum (Figure 3, above), if not addressed may have the potential for leading to serious, near-disaster interruptions of data access. Addressing some of the vulnerabilities may also reduce the impact of a disaster when it strikes. These issues constitute disaster mitigation preparation, since they address problems before they accumulate in severity and since they constitute appropriate systems and network management responsibilities.

The interruptions at the low end of the continuum, short-duration localized breaks can be addressed by uninterruptible power supplies to correct power fluctuations; adequate help desk and other user support staff to address user errors; and user training for proper use of equipment and software, which is also part of the HIPAA Security requirements.

For intermediate interruptions, file backup and recovery practices can assist users in accidental loss or deletion of critical data files. Also, adequate vendor maintenance contracts can insure the replacement of failed components in a reasonable time frame. In fact, appropriate server and network design can include component redundancy, such as RAID disk arrays allowing data mirroring (Hagland, 2003), and even server redundancy with fail-over software. This is a technology whose time has come, based on the low cost of powerful rack-mounted servers and software tools (Tissot, 2002). Rack-mounted servers and Storage Area Networks allow for physical separation of processing power and stored data, allowing for independent service and replacement. Another variation of this idea is to set up redundant servers in offsite repositories with daily data duplication (Meyer, 2002).

In large organizations with several organizational locations connecting multiple local area networks through a wide area network, one of the vulnerabilities to information access is often the single location of network access points. While creating a totally redundant network can be too expensive (Tissot, 2002), creating a duplicate network hub within the campus that can share the load during normal hours but take over the load if one hub is lost can be a way of eliminating that vulnerability. Finally, these services require a well-trained systems support staff to perform the necessary support and to document the overall management planning. The cost of all these services and practices should be weighed against the overall cost of replacing information and of potential penalties if regulations are violated.

Another underlying perspective for healthcare organizations is the determination and documentation of the resident EPHI. If EPHI is distributed across multiple desktop workstations in multiple locations, re-creating the desktop environment and reconstituting the EPHI may be difficult if not impossible. The same can be said for the use of wireless devices to access and download EPHI. Strong consideration should be given to restricting the storage of EPHI to protected servers and disk arrays, allowing only authorized remote access to the information, rather than allowing directories or

Table 4. INFORMATION CLASSIFICATION: BUSINESS REQUIREMENTS				
	Access	Direction	Risk	Access
Registered	Sr. Execs	Strategic	Critical	Not on Line
Restricted	VPs	Tactical long	High Risk	Standalone
Confidential	Directors	Tactical short	Sensitive	Restricted
Internal	Employees	Minimal	Delegated	Internal

databases to be resident on desktops. These decisions require documented policy statements, which can then be used in employee awareness training.

For the sake of efficiency, information can be classified based on importance to organization. HIPAA will of course require the protection of all EPHI, but other information, such as employee data, non-patient accounting records and external contract records need to be protected as well. Information can be classified according to business requirements as well as control requirements, seen in Tables 4 and 5. Some information at the high ends of each table may never be placed on a network accessible computer, and some of the information at the low ends may not need

Table 5. INFORMATION CLASSIFICATION: CONTROL REQUIREMENTS					
	Copy	Disclosure	Disposal	Physical	Transmit
Registered	No	No	Return	Security	No
Restricted	Approval	Executive	Supervision	Lock	Encryption
Confidential	Approval	Need	Supervision	Lock	Optional
Internal	Yes	Mgmt	Optional	Optional	Text

disaster protection at all. The two sets of intermediate categories in each table will likely be the location of EPHI and other information that will require disaster recovery protection.

Pre-Disaster Planning

In proper priority, the organizational business continuity plan will be initiated along with, or prior to the disaster recovery plan, since they are interdependent. Also, at this point in the effort, the risk analysis will be completed in order to be used as the basis of the plan.

Several sources have outlined disaster recovery plans to assist project planners (Burrington-Brown and Hughes, 2003; Bahan, 2003; Martin, 2002; Freeman, 2002; Syong, 2001). Most of them have the same basic steps, similar to the steps outlined in Box 6 and described below. Since the purpose of a disaster recovery plan is to restore access to information that was lost during the disaster, the plan begins with the application development environments, hardware and software, that hold and display the information. Representatives of the technical support groups and every application system in active use will be involved in documenting the environments. This list can be provided in the form of a checklist for simplification. Burrington-Brown and Hughes (2003) offer an alternative set of sample checklists for gathering data and determining process details.

1. **Alternative Disasters.** The recovery planning must include discussions of alternative disasters and disaster scenarios based on the threats identified within the risk analysis. Most efficiently, the plan itself should be based on the worst case, and the alternative strategies can be built on or derived from the main plan.

2. **Recovery Site.** The technical staff will take the initiative on this task. Historically, large mainframe sites were restricted to reserving space at a vendor hot site. With smaller servers, more opportunities are available. Ideally, the recovery site will be another company-owned location within a few miles; with pre-configured communications lines back to the main site. Rented temporary buildings or trailers could be used at the original site, if conditions are favorable.

3. **Technology Recovery Steps.** This involves the setup of the temporary location, including network connections, servers, and potentially workstations necessary to recover all systems lost during the disaster. The technical systems staff must have the installation and configuration of the hardware and the recovery from the backup environment clearly documented in the form of a step-wise script in the event that recovery takes place without certain key individuals. Any competent systems person should be able to perform the steps as written.

4. **Application Recovery Team.** This is a subset of the disaster recovery team (with perhaps additional user representatives) based on specific application systems or sets of systems. The recovery team needs to include technical representatives, developers and users of the system. All must be familiar enough with the whole process to drive the events, and to train additional ad hoc team members if they are needed.

5. **Pre-Disaster Working Environment.** The technology support staff will be responsible for documenting the server, storage and network environment, including any client workstations and peripheral devices (printers, scanners) necessary for the implementation of the application. The application recovery team and technical staff must be aware of the specific, detailed elements of the application system environment as it exists independent of any disaster. This must include network numbers and names, user identifier names, application vendor contact information, and application licensing information. This could be

Box 6. Pre-Disaster Planning

1. The **response to alternative disasters** must be addressed.
2. A **recovery site** must be located.
3. The **technology recovery configuration steps** must be spelled out.
4. An **application recovery team** must be identified.
5. The **pre-disaster working environment** must be defined.
6. A **minimum recovery configuration** must be defined.
7. The **source of the minimum configuration** must be identified.
8. An **alternative work process**, if needed, must be developed.
9. A **recovery window** must be defined.
10. The **impact of periodic business** cycles must be evaluated.
11. For multiple application systems, a **recovery priority** must be determined.
12. The **application recovery plan** must be documented.
13. The **restoration to normal operations plan** must be documented.

the impetus for developing inventories of resources as part of pre-disaster planning.

6. **Minimum Recovery Configuration.** For the duration of the recovery, until complete services can be replaced, it is not likely that an exact duplicate of all hardware will be needed. The users can identify the most important functionality needed to begin a minimal operation of the application and the technical staff can recommend a hardware configuration to meet those needs. The developers and users also have to determine the need for additional software licenses for the minimum configuration.
7. **Source Of Minimum Configuration.** There are three main options for providing the minimum recovery configuration, each could be considered separately, or the solution could be a combination of the three:
 - A. **Acquire the minimum recovery equipment** and place it in the alternative location, where it will remain idle until a disaster recovery test or an actual disaster occurs. This can be the easiest and most flexible solution, but is also likely to be cost prohibitive. Such an alternative has seldom been a consideration of mainframe disaster recovery plans, due to the exorbitant expense of maintaining a separate, idle mainframe. Though the cost would be significantly less for a distributed environment, if an organization has invested heavily in client-server technology, the cost for duplicating each environment could be comparable to the cost of a second mainframe.
 - B. **Use existing hardware wherever possible.** This solution is certainly less expensive, but it has the complication of 'commandeering' equipment in use and possible de-configuring the existing environment and reconfiguring the systems for the recovered environment. If an available recovery server were designated, an inexpensive suggestion would be to add additional pre-configured empty disk drives, enough to hold the minimum configuration, and leave them idle on the active server. At recovery time, the drives could be activated and made ready to receive the restored backup files.
 - C. **Contract with a disaster recovery vendor** for a service to replace the minimum recovery configuration (or a subset) within a reasonable time frame, such as one to two days. The cost for such a service will be a monthly fee, which could be budgeted as a maintenance service contract. The service could include a 'hot site' in a distant city where the recovery configuration is set up, a mobile computer room with the systems set up driven to the company site, or the delivery of the replacement configuration to a company location. This effort requires little or no capital expenses, but may require a multi-year lease. As with any contract, let the buyer beware.

- 8. *Alternative Work Process.*** The application recovery team is primarily responsible for this and the following four steps. If it is necessary, each team will need to develop a non-electronic process for continuing the work required by the application system in the event that the technology is not restored in a timely fashion.
- 9. *Recovery Window.*** The recovery window is the minimum time the particular application can be performed by hand, or be done without completely. It is not the length of time it takes to bring up the minimum configuration. This obviously will be different for different applications - financial systems and active patient treatment files will likely have shorter, more critical recovery windows than most other applications.
- 10. *Periodic Business Cycles.*** If a disaster occurs at a particular time of month, the impact on a given application system, e.g., accounts payable/receivable, ledger balancing, billing activities, etc., may be different. The time of the disaster may therefore change the priority of recovery for the group of applications systems
- 11. *Recovery Priority.*** The priority for recovery among multiple application systems is important to determine in this planning stage, lest post-disaster debate slow down the process. This will be based on the need for alternative processes, and will require periodic review, and will differ with periodic business cycles.
- 12. *Application Recovery Plan.*** This includes recovering of the database system itself and the applications around the database, if appropriate. This must be documented plainly in the likelihood that replacement staff are involved in the recovery. This will be the basis for the applications testing plan, which is ultimately necessary to validate the recovery process, and to test periodically the recovery procedures.
- 13. *Restoration Plan.*** After a disaster recovery process occurs, adequate plans must be in place to migrate the environment back to the normal production location. This is often overlooked until it is late in the game, but such a move can present problems, especially if the recovery site is a great distance away. This does not necessarily mean reversing the steps or the recovery plan, since new physical company locations may be involved. It does, however require the same attention to detail in choosing the timing and priority of the application systems.

All of these elements must be spelled out in detail as part of the specific plan for each disaster recovery effort, based on location within the organization and/or different application systems. It is not unusual for the initial plan in a large organization to take several months to complete, since many details have to be included. Subsequent

application system project teams can follow the basic outline of the first plan, saving a good bit of time. The documentation of this overall plan can optionally be built using a disaster recovery vendor tool, but a carefully managed word processing document will suffice. Ideally, one main system plan should be maintained on-line and designated as the foundational plan from which all copies will be made, and a duplicate could be placed on a small laptop and located in the off site location. This means that all printed plans will be considered unofficial copies. The plan will be a dynamic document, requiring frequent changes and modifications. Therefore, the plan can be maintained on-line and a duplicate.

All documentation and plans must be kept in an off site locations, e. g., at the recovery location. The creation of the plan will undoubtedly stimulate the need for appropriate information technology policies and procedures, such as network designs and conventions, user registration and security. Copies of each of them should accompany the plan. Table 6. is a suggested list of policy and procedure categories. Phoenix Health Systems also offers a set of information management policies and procedures for customers (see Resources below). These documents will feed directly into HIPAA Security requirements as discussed earlier. From this plan, periodic tests can be performed, and ultimately, the recovery process itself will be accomplished.

TABLE 6. INFORMATION MANAGEMENT POLICIES AND PROCEDURES	
SUBJECT	EXPLANATION
Network Structure	How is the network built and what is its current inventory?
Data Ownership, Responsibility	Who is in charge of specific data stores?
Data Classification/Encryption	Which data is most important and how is it sent outside of the network?
Intranet/Internet	What information is shared internally and what is kept from external web viewers?
Systems Administration	What are the responsibilities of the Technical Support staff ?
Service Level Agreements	How does Technical Support commit to system availability and services provided?
User Registration, Passwords	How are users placed onto the network and how are passwords managed?
Network Access Authorization	Who can and cannot use the network?
Data Backup And Recovery	How and when is data from servers backed up, stored, and recovered?
Vendor Maintenance	What vendor contracts are in place and what are the specific services provided?
Technology Change Management	How is the network changed/expanded? How are systems modified/upgraded?
System Failure	How are system failures defined and what are the corrective steps?
Incident Response/Notification	What happens when a security incident occurs?
Audit, Disaster Recovery	How are audits and current disaster recovery plans carried out?

This plan design has been based on relatively large, complex organizations and for the worst-case disaster scenario when all information systems and infrastructure have been lost. Smaller sized organizations with less complicated configurations may not require as complex a plan. The detail of the plan on the basis of application system recovery still would apply, but without the complexities of large networks, multiple applications systems and large numbers of users, small facilities would be able to streamline the plan appropriately. This would also include selecting a subset of the list

of policy/procedure topics. Once the HIPAA Security standards for a risk analysis has been performed, a small facility may be able to build a relatively simple disaster recovery plan that still accomplishes the requirements.

Declaration

Even though such a calamity may be obvious to any observer, there still is a need to declare the disaster in order to begin the recovery procedures. A central phone number, such as the corporate hot line or corporate security needs to be the starting point for reporting potential disasters. Backup numbers should also be in place in case the primary numbers are themselves part of the disaster. After receiving the initial call, the hot line will call someone with authority, e.g., the head of the disaster recovery team, or his or her delegate who must assess the situation, declare the disaster, contact the recovery team, and initiate and assume ownership of the recovery process. Prior to the declaration, it is important to consider the incident a potential disaster for the sake order and It is important that ***the declaration of the disaster and recovery begins from the decision of a single person*** and not from several potentially competing individuals or groups who may have different levels of understanding of the situation. The crisis at the worst nuclear disasters in American history, Three Mile Island, PA, was nearly exacerbated beyond control because five different groups were attempting to steer the effort: the builders of the plant, the managers of the plant, the local State government, the Nuclear Regulatory Commission, and the local and national media (Burns, 1991).

A declaration sequence beyond the initial contact also must be spelled out in the event of the absence of the disaster recovery team manager. For example, the process may start with the manager, and in his or her absence, the pre-identified assistant to the manager. Beyond the absence of the assistant, the decision should move to the Manager of the Technical Support, then the Director of the Information Technology Department, and ultimately the CIO of the company – or other organizational-specific chain of authority. Beyond the declaration, a notification escalation chain of contacts within the organization itself should be set up to spread the news. The documentation in the Disaster Recovery Plan should include phone and pager numbers and addresses of all key individuals in the sequence chains.

Recovery

Having previously defined and acquired the minimum recovery environment, the recovery plan after a disaster basically means assembling the recovery environment at the selected remote site, configuring the system, loading the backups - all within the recovery window and following the priority list. The recovery team and application recovery teams will obviously divide the responsibilities of the recovery based on their existing roles. Once a disaster has been declared, Box 7 suggests the main recovery steps to follow.

Box 7. Recovery Steps

1. Contact the Disaster Recovery, Business Continuity, and Application Recovery Teams and initiate the processes.
2. Contact the disaster recovery vendor for the recovery system (if selected), or commandeer the alternative equipment from its current usage.
3. Contact the vendors of the original equipment to begin building the recovery environment.
4. Configure user IDs and network names and numbers within the recovery environment
5. Acquire the backup tapes for the server software and files.
6. Reconfigure the alternative server for the needed environment and load backup files.
7. Reconfigure the workstations to be used in the recovery test.
8. Establish the alternative print environment.
9. Connect all the elements to the network for client access, either locally or back to the main site.
10. Attempt to perform routine tasks on the recovery environment.

Assuming no personnel absences, the bottlenecks for the recovery will likely be the configuring the network appropriately and acquiring of the backup tapes from the backup server or off-site storage. Next will come the configuring of the backup server for the recovery environment, and the reloading of the tapes onto the recovery server. This will be followed by the application support groups who will initiate the rebuilding of the application system. The time to completion can be several days, more depending on the delivery of the replacement equipment. If the plan has been carefully constructed and tested, the actual recovery after a disaster could be almost anti-climactic, though the significance of the disaster and the potential losses of personnel and business processes are not to be ignored.

Testing

Once the agreement for an alternative location is in place, including the space allocation and wiring, a disaster recovery effort can be tested. Tests are invaluable to the process; many loose ends will be exposed and a realistic recovery time period can be established. The tests should be made frequently enough to keep the procedures and sequences current, since application systems are dynamic, involving changes in hardware, software, and the size of the user population, and changes in the makeup of the disaster recovery team. This is why strong consideration should be made to keeping the official copy of the recovery plan on-line. In such a case, printed copies would be considered unofficial by nature.

Rothstein (1993) pointed out that the truest measure of a test is not successful recovery, but the identifying of flaws and weaknesses. This is important to keep in mind, and to forewarn users and executives about, because the first test will most assuredly fail, likely in many ways. Disaster recovery vendors usually include one or two tests per calendar year in their contracts, subject to negotiations. The test itself will begin at the declaration of the designated person as stated above, and will end when the particular services can be provided on the minimal configuration. The steps for the test are virtually identical to those of the recovery process, save the need to replace the original equipment. The challenge for environments with multiple application systems is to balance time and frequency; perhaps all projects should not be tested at the same time.

Conclusions

A disaster that takes out significant portions of a central computing resource is hopefully rare. However, several disasters over the past few years have heightened the sensitivities of all IT professionals. The onset of HIPAA Security will affect all healthcare organizations regardless of size, but the overall effort for disaster recovery will be less complex for smaller organizations. The effort is worth the wise expense of time and money, even if the disaster will not affect anyone outside the company. The amount of expense is directly related to the value a company places on its information resources, and in the case of HIPAA, the amount of penalty for failure to comply. A disaster recovery effort is enhanced with careful planning and testing, which can only take place after the events specific to distributed computing environments have been defined, organized and placed in proper sequence.

Resources

There are an abundance of resources on the internet for designing and planning disaster recovery plans. For healthcare organizations, some of the References Cited (following page) will lead to outlines and suggestions similar to this paper. Several healthcare vendors and consortia now offer disaster recovery planning services and tool kits, some of which are listed below. The Centers for Medicare and Medicaid Services (CMS) held a recent conference in which two presentations offer additional perspectives on disaster recovery and contingency planning, also listed below.

[A⁴ Health Systems Releases HealthMatics[®] Assure](#). November 10, 2003.

Centers for Medicare and Medicaid Services. [2003 National Medicaid HIPAA & MMIS Conference](#), February 9 - 13, 2003

[Contingency Planning](#): Addressing Critical Business Processes that Support Implementation of HIPAA Transactions (PDF 1.4MB)
Marie Margiottiello, CMS; Henry Chao, CMS

[Security: Disaster Recovery and Business Applications](#) (PDF 97K)
Brenda Rose, Maryland Medicaid

[HIMSS Disaster Preparation for Healthcare IT](#).

[HIMSS Emergency Preparedness and Disaster Recovery CD-ROM](#)

[Kodak Launches New Family of Healthcare IT Services](#). October 30, 2003.

[Phoenix Health Systems' HIPAA Security Policies Templates Suite](#). 2003

References cited

Bahan, Chad. "[The Disaster Recovery Plan.](#)" SANS Reading Room, June, 2003.

Bogen, Jon. "Implications of HIPAA on business continuity and disaster recovery practices in healthcare organizations." *Healthcare Review*, May 28, 2002.

Brown, Jennifer. "Survey warns IT managers to make disaster recovery top priority." *Computing Canada*, Feb 28, 2003.

Burns, Christopher. Three Mile Island: The Information Meltdown. *In: Great Information Disasters*, Forrest W. Horton, Dennis Lewis, eds. ASLIB. London. 1991.

Burrington-Brown, Jill, and Gwen Hughes. "[AHIMA Practice Brief: Disaster Planning for Health Information](#)" (Updated June 2003). Copyright © 2003 American Health Information Management Association. 2003.

Dougherty, Michelle. "Long Term Care Health Information Practice And Documentation Guidelines." AHIMA/FORE Long Term Care Taskforce. 2001.

Freeman, Wayne. "[Business Resumption Planning: A Progressive Approach.](#)" SANS Reading Room, February, 2002.

Hagland, Mark. "Disaster Recovery And Storage: Hospitals make progress on a crucial need." *Healthcare Informatics*, March 2003.

"[Health Insurance Reform: Security Standards; Final Rule.](#)" 45 CFR parts 160, 162, and 164. *Federal Register* 68, no. 34, page 8377. February 20, 2003.

Kelly, Sean. "A disaster waiting to happen." *Communications News*, August 01 2001.

Martin, Bryan C. [Disaster Recovery Plan Strategies and Processes.](#) SANS Reading Room, February, 2002.

Meyer, Matthew. "Healthcare facility planning for disaster recovery." *Healthcare Review* Jan 8, 2002.

Murphy, James C. "Client-Server Disaster Recovery: Process Organization." [SANS IV Proceedings and Workbook.](#) USENIX Assoc. 1995.

O'Brien, Jennifer M. "Business warned to plan: Quantum heads initiative to educate companies about need for disaster recovery planning." *Computer Dealer News*, March 5, 1999.

Rothstein, Philip. J. "Hard evidence: Contingency plans as tools for recovery." *Enterprise Systems Journal*, July, 1993.

Syong, Gan Chee. "[Introduction to Business Continuity Planning.](#)" SANS Reading Room, October 1, 2001

Spath, Patrice. "Health information disaster planning 101. (The Quality-Cost Connection)." Copyright 2002, A Thomson Healthcare Company.

[**Note:** Paper listed in [LookSmart's FindArticles](#), but link no longer active.]

Tissot, Marlene. "Be prepared. (Special Report: Disaster Recovery)." *Communications Week International*, March 4, 2002

Webster, John. "[Disaster recovery — insurance policy or survival plan?](#)" *Storage Network World, Online*. September 30th, 2002.

Widup, Suzanne. "[Business Continuity Planning In Difficult Economic Times.](#)" Sans Reading Room, 2003

© SANS Institute 2004, Author retains full rights.