

October 17, 2007



Information Security Risk Management for Healthcare Systems

This Paper was developed by the
Joint NEMA/COCIR/JIRA Security and Privacy Committee (SPC)

The paper has been approved by:
MITA (Medical Imaging & Technology Alliance)
COCIR (European Coordination Committee of the Radiological and
Electromedical Industry)
JIRA (Japan Industries Association of Radiological Systems)

October 2007

© JOINT NEMA/COCIR/JIRA SECURITY AND PRIVACY COMMITTEE (SPC)

www.nema.org/medical/spc

Secretariat: MITA (Medical Imaging & Technology Alliance) www.medicalimaging.org

1300 North 17th Street, Suite 1752, Rosslyn, VA 22209 USA tel: 703-841-3200 fax: 703-841-5900

Secretary: Richard Eaton, tel: 703-841-3248; fax: 703-841-3348

E-mail: reaton@medicalimaging.org

May be quoted if reference and credit to SPC is properly indicated.

1 Purpose and Scope

This document helps device manufacturers manage IT security risks in healthcare systems by detailing the steps in security risk assessment in the context of security risk management. IT security risks are risks to data and systems. As a best practice, patient and operator safety risk management and IT security risk management processes should be separate but linked. They differ in both the vocabulary and expertise required for proper risk management. If combined as a single assessment process, one or the other is not treated appropriately.

This paper first sets forth typical examples of threats and describes a process that can be used to design secure healthcare systems. A manufacturer's risk analysis starts with the intended use and assumes a hypothetical healthcare facility environment thus it cannot be taken as the healthcare provider risk analysis. The same basic methodology can and should be applied by healthcare facilities to assess and mitigate the risk when combining equipment of different vendors as a healthcare delivery network, when adding new equipment to an existing network, or when significantly changing the configuration of an existing network.

2 Introduction

Risk is inherent in the delivery of healthcare. The security risks associated with healthcare systems have increased as direct (network) and indirect (media) connectivity has increased. With sophisticated equipment, there are always more risks than any organization can afford to fully eliminate. Therefore, the need arises for a systematic, documented method to assign risks so that they can be listed in priority order, mitigated accordingly, and have residual risks documented and accepted.

The process for managing healthcare systems IT security-related risks is very similar to long-standing device safety processes. The medical device industry has been engaged in safety risk analysis for over 30 years. This paper recommends that similar methods be applied to security risks to healthcare systems. These methods support a manufacturer in assisting the healthcare provider and directly support a healthcare provider in maintaining confidentiality, integrity, and availability of protected health information.

The process of IT security risk management as described in this document includes

1. Listing the assets under consideration and understanding their intended use;
2. Collecting security-related requirements for the assets;
3. Elaboration of threats and applying them to systems to determine vulnerabilities including actors, threat paths, and possible outcomes;
4. Scoring of risks;
5. Proposing and implementing mitigations for vulnerabilities appropriate to the healthcare domain;
6. Summarization of residual risks along with the system's role in advancing the healthcare mission, in order to obtain a "go" or "no-go" decision from

manufacturer's executive management to give the authority to proceed with development of the system.

Steps 1 to 5 are an example of Failure Mode Effects Analysis (FMEA) and comprise the first phase of a system security risk assessment. When followed by step 6, summarization and sign off, and built into a sustainable process, a process for healthcare system life cycle security risk management is created.

In detailing security risk management, this paper presents a set of examples which provide insight into healthcare threats and vulnerabilities leading to the relevant steps in how to design for confidentiality, integrity, and availability of systems while maintaining an appropriate level of (1) safety, (2) healthcare effectiveness, (3) privacy for both patient and staff, and (4) interoperability.

3 Some Examples of Emergent Threats

These scenarios provide a brief description of a few healthcare delivery issues and negative impacts which are encountered when dealing with IT-related security threats. Many threats directly impact the protection of privacy. The failure to protect privacy may violate regulations such as the USA HIPAA Privacy and Security rules or the data privacy laws in Europe, Japan and elsewhere. The examples are illustrative and should not be considered as a complete list.

There are some differences in the protection of healthcare systems as compared with many other IT-based businesses. For example, while most banks or businesses may close in the event of a natural disaster or a severe IT exploit, a healthcare facility, in general, needs to remain open. Operation under adverse conditions is essential to treat current patients and to maintain and restore community health in the event of a disaster. Although continuity planning is sometimes seen as purely operational, it is an essential element in security (it should be remembered that security includes protection of confidentiality, integrity, and availability).

3.1 Hospital Service Discontinuity

Scenario: A localized failure within the hospital has disconnected the emergency room (ER) from the network backbone. Thus the typical IT services are unavailable (e.g., network addressing, routing, user authentication) but patients continue to arrive at the ER requiring critical healthcare services. The healthcare systems in the ER need to continue to provide critical healthcare services.

Possible Design Mitigations: To create a robust mitigation to this scenario, the systems could include the ability to create and locally store health information that was created (medical images, records) even in the absence of LAN access. It is advisable for healthcare providers to consider disaster recovery when planning the use of thin-client versus local storage workstations. A careful risk management plan including disaster preparedness will result in a proper balance of these components.

Possible Operational Workarounds: (1) Use additional staffing, (2) Use equipment available at other departments of the same healthcare facility, (3) Route pa-

tients to a nearby healthcare facility or (4) Deploy a secure wireless solution to communicate with the network backbone

At a minimum, patients would experience service degradation via longer waiting times or an off-loading of patients to another healthcare facility.

3.2 Widespread Disaster

Scenario: Provision of healthcare in the aftermath of a widespread disaster. Such a disaster may have been caused by natural (e.g., earthquake, tsunami, hurricane/typhoon, volcano, wildfire) or man-made causes (terror, war, power failure).

During these disasters the general infrastructure (IT networks, roads, electrical power, water) may additionally be disrupted or destroyed. Further, the disaster may have caused damage to the healthcare facility itself and thus may have destroyed parts of the local building or healthcare infrastructure causing a “Healthcare System Failure.” The situation may get worse as the disaster itself increases the number of patients who arrive at the healthcare facility.

Possible Design Mitigations: The system may have an emergency mode that allows for the identification of individuals without authentication to support the local workforce. For procedural measures which may be taken by the healthcare provider to ensure this kind of access, see the SPC White Paper, “*Break-Glass – An Approach to Granting Emergency Access to Healthcare Systems*” December 2004. http://www.nema.org/prod/med/security/upload/Break-Glass-Emergency_Access_to_Healthcare_Systems.pdf. This permits the healthcare provider to use all local trained medical staff as needed.

Possible Operational Workarounds: For this kind of rare disaster, risk mitigation measures are difficult to define, as there may be no other medical equipment available. However, contingency planning is essential in continuing the healthcare mission. Included in any such plans should be the provision that critical healthcare data should be protected against destruction in such a disaster, and, if possible, should be accessible from other facilities that have to serve as backup solutions.

3.3 Indiscriminate Malicious Attack

Scenario: A medical device is being used on a patient (e.g., x-ray, ECG, ventilator, CT, MRI, PET), when a malicious software attack occurs. This may be a side effect of a broad cyber attack where the medical device is not specifically targeted. These broad, sometimes low-skill, technology attack tools are otherwise known as viruses, Trojan horses, or worms, for example. Even under these circumstances, the system should be able to protect patient safety and health. Individual patient and healthcare provider damage may result if the attack leads to the disclosure of personal data.

Possible Design Mitigations: During routine operation, only those services and network protocols that are essential for the proper use of the system should be permitted and remain active. Authentication mechanisms are in place that permit only trusted nodes (e.g., IHE Audit Trail and Node Authentication integration Profile) to communicate, thus blocking attacks. For many systems, when network access fails, the design permits them to fall back into a diminished but still useful

function. For example, a monitoring system without its network will no longer support the central display of patient information, but the bedside monitors will continue their operation and display. Some systems, like PACS workstations, will not function when network access to their storage is disrupted. In this case, systems should be fail-safe to allow healthcare staff to go directly to acquisition devices to view images or print from acquisition systems. This might require carrying media from the acquisition system to workstation.

Possible Operational Workaround: The usual, immediate risk mitigation measure would be to remove the network access from this device. But this may not be appropriate for certain devices that require continuous network access, and it may lead to severe medical consequences for the patient. Although this dependence on network connectivity is rising, these systems are usually in support of efficient workflow rather than directly providing a life-critical function. In general, life-critical devices fail back to proper stand-alone operation even if the network is not functioning. It is important that hospitals have contingency plans for a network malfunction, e.g., rapidly increasing staff during sustained network failure.

3.4 Highly Funded Attack

Scenario: A malicious attacker is highly funded and is highly capable of launching a targeted attack. Typically, the attacker is an outsider and the targets are medical data of VIPs such as athletes or celebrities, stored in a healthcare system. The healthcare systems as target might not be different than any other system that contains information of high value for the highly funded attacker. But compared to the temporary inconvenience of a compromised secret password that may be changed after unauthorized disclosure, the effects of disclosed medical information (e.g., cancer, HIV status) may never be undone and may cause severe social and financial consequences to the victim.

Possible Design Mitigations: The system might include strong user authentication and authorization. It should be able to mark VIP patients and to restrict access to their health information and increase logging activity. Encryption of communication and digital signatures for reports may further increase confidentiality and integrity of stored health information.

Possible Operational Workaround: It is hard for a healthcare facility to protect against this sort of attack. Even a detailed management of access rights may be useless if there are underlying vulnerabilities. Depending on the motivation of the malicious party, financial gain may convince insiders – who usually need and hence have access rights to patient data – to perform such criminal attacks.

3.5 Personal Revenge

Scenario: A threat may originate from angry or vengeful persons (employees, patients, or service staff, for example). The bulk of these attacks come from internal, or formerly internal, people. They have a powerful desire to inflict damage to a specific target inside the healthcare facility or to the healthcare facility as a whole, but are not likely to be sophisticated in terms of knowledge about systems or well funded.

Possible Design Mitigations: The system should include user authentication and access controls, as well as audit logs, to detect and document deviations from internal policies. Changes in access rights need to be effective immediately.

Possible Procedural and Technical Mitigations: Because vengeance is usually manifested in short duration attacks, it is hard to work around. Well administered firewalls, appropriate access rights, separated VLANs, etc. may be adequate. These controls must be more stringent to avoid successful attacks from a knowledgeable person who knows how to obtain access via previous personal contacts or intimate staffing knowledge (i.e., social engineering attack). These scenarios emphasize the importance of personnel training, the proper creation and enforcement of procedures to control access rights, and having processes in place to remove rights immediately upon termination of employment or third-party service contracts.

3.6 New Vulnerability Announced

Scenario: A medical system platform component is discovered to have a vulnerability that is being exploited in the field.

Possible Design Mitigations: Harden the medical system: shutdown unnecessary services, run services at least privilege level, provide minimal access to services unnecessary to the intended use (e.g., email, browsers, etc.). If possible, design to allow remote service access including software update capability and operational validation procedures. For more detail, see SPC White Papers "Patching off-the-shelf Software in Medical Information Systems, October 2004, and Defending Medical Information Systems against Malicious Software," December 2003 at (www.nema.org/medical/spc). Document the normal network behavior (ports, services, etc.) under the intended use and servicing. This may help create alerts that fire when typical network behavior is seen.

Possible Operational Workaround: Isolate from the operational network. If unable to isolate, configure firewalls and, as best practice, intrusion detection systems to provide protection using the documentation of the normal network behavior.

4 Healthcare Security Risk Management Process

The process for managing healthcare system IT security-related risks is very similar to long-standing device safety processes. For example, tools such as Failure Mode and Effect Analysis (FMEA) as applied to safety considerations can be used for security investigations as well. We will not detail safety risk assessment other than to note the relationship where security must be subordinate to patient and operator safety.

To avoid conflict and confusion, we recommend that the security risk assessment process be performed separately from the safety risk assessment, because overall, they have different requirements and involve fundamentally different assets. Whenever a security risk has a credible safety risk, even after proposed mitigation, the safety risk assessment process takes precedence. In general, this means moving the primary discussion of the risk to the safety team accompanied

by a knowledgeable, security team member. In this manner, the two processes generally proceed in parallel throughout the product creation process.

The skills of the risk management team members require specific elaboration. People with general IT knowledge as a background often are not aware of the healthcare-specific issues that may lead to impractical measures at the end. The security risk analysis team should be multi-disciplinary with the following attributes:

- Represents both business and technical aspects of the system (including IT knowledge)
- Understands both clinical processes and manufacturers' development processes
- Understands the healthcare-specific requirements (safety requirements are more important than security requirements)
- Includes a member familiar with the safety risk management process for products.

In addition, the team should be supplemented in an ad hoc manner by visiting experts who can help with network issues, IT security details, vulnerability tool assessments, and other specialized issues as they arise in the risk assessment project.

The IT security risk assessment, as described in the following sections, will answer the following basic questions:

- What are the valuable assets that fall under the intended use of the system?
- What are the security-related requirements for the assets under consideration?
- Who will perform an attack (human and non-human actors)?
- What are the possible threat paths?
- What are the possible impacts of a successful attack?
- What is the score of the initial risk?
- What actions may mitigate the risk?
- What is the score of the residual risk?

4.1 List of Assets

As a basis of the risk management, the assets under consideration that need protection, together with their intended use, must be listed. A typical, but not exhaustive, list of assets include hardware and software used for processing medical information and key data elements, and include different kinds of data:

- Specific components/medical application systems (e.g., image creating modalities, network components) of the IT infrastructure of the hospital
- Unspecific components/medical application systems of the IT infrastructure of the hospital (e.g., denial of service attack may block the whole network traffic)

- Medical application software itself
- Data about configuration of hardware and software
- Personal data of a specific patient
- Personal data of staff and other persons
- Healthcare procedure support information, including history of use and operator/user details.

The list of assets needs to be detailed enough to begin the assignment of direct threats to each of them and to be able to identify and implement appropriate risk mitigation measures. For example, simply to list the hospital information system as one asset would not provide enough specificity to detail a realistic, specific threat. In general, a network diagram (even if it were an agreed upon “typical network”) allows a systematic overview of the IT architecture of the developed equipment or the whole hospital network. It eases the identification of identical systems used at multiple locations in the same installation (or by the same health-care provider) that are exposed to the same risks, and that need the implementation of the same risk mitigation measures. By properly using such an overall network approach, a manufacturer’s single risk-management team can broaden the potential mitigations to account for system use in a wide variety of network implementations.

4.2 Collection of Security-Related Requirements

The assessment team should collect together all materials that detail the system requirements for security, including specifics for all assets on the levels of confidentiality, integrity, availability, accountability (i.e., authentication, and log-file availability/use). This requirements collection can be a specific document collection or be realized as a set of explicit references that are detailed, one-by-one, and documented as part of the risk management process. Input requirements typically would come from:

- Regulatory requirements (HIPAA, Directive 95/46 EC, etc.) directed to the user
- Customer requirements (government agencies, buying groups, etc.)
- Secure platform configuration guides (e.g., NIST, NSA and other guides)
- Internal security/privacy policy documents
- Industry “best practices” white papers
- Requirements from corrective actions based on prior experience.

4.3 Elaboration of Threats and Impacts

With the lists of assets and security-related requirements, the risk assessment team brainstorms, develops, and documents all possible threats for each asset. When a general threat may be exploited in a particular system, it becomes known as system vulnerability. In general, this elaboration of threats for each asset follows the chain of identification access paths, actors, motives, and outcomes and should be documented in a table (see Appendix A)

4.3.1 Possible Actors

There is a wide variety of human and non-human actors. Actors utilizing the threats can be categorized as:

- **Authorized persons:** insiders with valid account who are not authorized to perform a specific task, such as:
 - accidental attacks
 - insiders that are paid by external initiators
 - insiders that are motivated by personal profit or revenge.
- **Persons who are not authorized to access the network infrastructure:** outsiders with no account but some kind of access (physical or logical) to the healthcare provider infrastructure, e.g.,
 - vandals (script kiddies or hackers)
 - paid external people/organized crime
 - journalists seeking stories on VIPs (sports figures, politicians, etc)
 - visitors and patients
 - soldiers and/or terrorists.
- **Non-human events:** these events typically happen on an unpredictable basis without direct human influence
 - local infrastructure failure: Emergency room is disconnected from the network backbone but some emergency help must be provided to patients.
 - major industrial accidents: A large number of injuries must be treated while a power failure caused by that accident hinders provision of health care
 - natural disasters: They may cause injury to the local community as well as to the local infrastructure. A power failure may hinder the operation of the medical equipment, but many injuries flood the emergency room at the same time.

4.3.2 Threat Paths

If people are behind an attack, they may use different ways to access their target in the network. They are different in their ability to be detected: some are viewable by persons, others are not:

- **Direct (physical) access** to the medical system (viewable action)
 - Sitting at a medical system console provides a means to compromise security of the system.
 - Equipment without proper physical security can be stolen.
 - Physical access opens the path to attack from removable media (CD, DVD, USB stick,).
- **Logical access** (non-viewable action) using a network (e.g., IP network or the telephone connection as path)

- From within the enterprise: typically, users known to the system who may have the principal authorization to perform a specific task, such as saving x-ray images on a CD for the purpose of patient treatment, may use this path for illegal activities. However, they may not have the authorization to save x-ray images on a CD for other purposes, such as forwarding the data to a journalist.
- There is an open path that is exploitable from outside the enterprise. For example, an unprotected open port is used by malicious software to take control of the system.

4.3.3 Possible Impacts

In compiling possible system vulnerabilities derived from generic threats, it is important to fully understand the potential impact of a successful attack. The impact can be for a single diagnostic or monitoring event, a single patient, a single diagnostic or monitoring system, or an entire deployed set of systems under a particular software version number. In general, the larger the number of systems impacted, the larger the severity. However, it is important to keep in mind that even a single patient privacy event can be of extreme severity. The irreversible disclosure of damaging private health information, such as certain diseases and conditions, may be financially devastating to individuals, especially when applied to well-known public figures.

The single or multiple systems impacted may be compromised via:

- Compromised safety of the patient or operator
- Unauthorized disclosure, destruction, or modification of personal data
- Monitoring system usage: who is using the equipment, who was treated
- Modification of the system outside of manufacturer specification
- Interruption of system availability
- Theft of services or supplies.

As an aid for the risk management team, and to assure a list as exhaustive as possible, it is sometimes useful to construct a "Threat Matrix" with asset stakeholders across the top and listed target asset as the rows. A good starting point has stakeholders as (a) patients, (b) health service providers, and (c) manufacturers or service organization.

4.4 Scoring of Risk

The scoring for a risk assigned to an identified vulnerability is a combination of the likelihood of a successful attack and the severity of the resulting impact on the assets. Security risk assessment makes use of essential elements of Failure Mode Effects Analysis (FMEA). The goal of this analysis is to arrive at a reasonable categorization of the risk so mitigation activities can be prioritized. This does not require a fine-grained numerical evaluation. In general, three levels for each of likelihood and severity are sufficient. However, based on the local requirements, and the availability and quality of data, the team may want to use 4 or 5 or some other number of levels. Before doing so, the team should appreciate that, in theory, the more levels used, the more detailed the risk scoring. Often, increasing

the number of levels provides a false sense of precision. Risk assessment teams must be careful to avoid speculation about all possible scenarios of attacks-- known as "movie scripting" in the security profession.

4.4.1 Likelihood of a Successful Attack

Realistic likelihood estimations will minimize "movie scripting."*[is this term defined anywhere?]* The following definition is used:

Likelihood is the probability that a vulnerability may be exploited within the construct of the associated threat environment. The following factors should be considered:

- Length of time vulnerability has been exposed
- Threat-source: motivation and capability of the actor.

Even when some of the threats may be amenable to more precise probability estimation, in general it is useful to stick to the three likelihood levels High, Medium, and Low, as described in Table 1.

Table 1: Likelihood levels

	Threat Source	Effect of Controls
High	Highly motivated and sufficiently capable	Are ineffective toward the vulnerability
Medium	Motivated and capable	Are in place and may impede exploitation of the vulnerability
Low	Minor motivation or capability, e.g., incidental attack	Are in place and prevent, or at least significantly impede, the vulnerability from being exploited

When assessing the likelihood of a human attack, the motivation behind the attack is a key factor, because it usually determines the resources and effort that will be expended by the attacker to violate applicable rules or laws. The attacker may actually make a cost benefit assessment to evaluate the cost of mounting the attack against the value of a successful attack. Usually, the greater the motivation of the attacker, the more money and effort will be spent on the attack, and correspondingly, the more sophisticated the countermeasures which will be required to protect the assets. Though fortunately uncommon, there are a few attacks which will appear to be highly motivated, and as a result, may cause severe damage to the healthcare provider.

A few examples will help (situation specific assessments are the user’s job / W. Leetz – agreed by SPC 09.10.2007) in performing the situation-specific assessment of the likelihood of a human attack.

- “Indiscriminate” attacks (mostly without malicious or specific intent) that are typically successful due to negligent attention to security measures:

- lack of clear security policy, objectives or activities
- lack of training or enforcement to adhere to procedures
- excluding the use of necessary security functions, which are short-sidely regarded as a hindrance to the clinical workflow.
- Deliberate attacks (frequently evidenced by malicious intent):
 - terrorism: healthcare may be a specific target
 - financial gain: financial benefit for the attacker may amount to millions of dollars
 - revenge by frustrated employees, patients, relatives of patients, and others.

4.4.2 Severity of a Successful Attack

The **severity** of a security event can be described in terms of loss or degradation of confidentiality, integrity, and availability. The following prioritized list provides a brief description of each security goal and the consequence (or impact) of its not being met:

- **Loss of integrity:** System and data integrity refers to the requirement that information be protected from improper modification. Integrity is lost if unauthorized changes are made to the data or system by either deliberate or accidental acts. If the loss of system or data integrity is not corrected, continued use of the compromised system or corrupted data could result in inaccuracy, fraud, or introduce safety issues. Also, violation of integrity may be the precursor to a successful attack against system availability or confidentiality.
- **Loss of availability:** If a system is unavailable to its end users, the medical facility's mission may be affected. Loss of system functionality and operational effectiveness, for example, may introduce safety concerns or reduce the quantity and/or quality of care.
- **Loss of confidentiality:** System and data confidentiality refers to the protection of information from unauthorized disclosure. Unauthorized, unanticipated, or unintentional disclosure could violate regulatory regional directives for the manufacturer or healthcare provider such as EU Directive 95/46 or, in the United States, cause the manufacturer to violate contractual business associate agreements (and thus lead providers to violate HIPAA regulations).

For all intents and purposes, a rough categorization (High, Medium, Low) will be sufficient to evaluate the potential severity of most adverse security events (Table 2). The number of severity levels used may depend on the specific needs of the system involved.

Table 2: Severity levels.

Severity	Definition
High	Exercise of the vulnerability (1) may result in the highly costly loss of major tangible assets or resources; (2) may significantly violate, compromise, or impede an organization's mission, reputation, or interest; or (3) may result in human death or serious injury.
Medium	Exercise of the vulnerability (1) may result in the costly loss of tangible assets or resources; (2) may violate, compromise, or impede an organization's mission, reputation, or interest; or (3) may result in human injury.
Low	Exercise of the vulnerability (1) may result in the loss of some tangible assets or resources or (2) may noticeably affect an organization's mission, reputation, or interest.

4.4.3 Risk Score

The scoring of a risk is determined by combining likelihood and severity of an attack. It determines the ranking of the risk mitigation measures. There is no simple universal agreement that determines what risk score is acceptable, and what score needs the implementation of risk mitigation measures. The actual priority for risk mitigation depends on the particular healthcare value of the system in its operational context, the system specifics, and other local conditions. Hence, the risk management team and management need to define the acceptable risk level for each system. By combining the likelihood and severity levels into a table such as Table 3, the team can assign risk scores of High, Medium, and Low. (Rem.: after SPC meeting of Oct. 09, 2007)

Table 3: Sample risk scores (in italics) as derived from likelihood and severity levels.

	Severity		
Threat Likelihood	Low	Medium	High
High	Low	High	High
Medium	Low	Medium	High
Low	Low	Low	Medium

The risk management team should agree on the correlation of risk scores that require the definition of follow-up actions before discussing specific risks. In the example above, the risk team was acting conservatively by giving priority to patient safety. Therefore the scores for all high severity cells were rated at a minimum level of “medium.”

A sample description of the risk scores is shown in Table 4 below. It represents the score of a risk to which an asset might be exposed if a given vulnerability were exploited, and the corresponding need for corrective measures.

Table 4: Risk Scores and urgency of mitigation activities

Risk Score	Necessary Actions
High	Strong need for corrective measures. An existing system may continue to operate, but a corrective action plan or other risk mitigation measure must be put in place as soon as possible.
Medium	Corrective actions are needed and a plan must be developed to incorporate these actions within a reasonable period of time.
Low	The system’s owner must determine whether corrective actions are still required or decide to accept the risk.

It is expected that all of these definitions, assumptions, and procedures are captured in organizational process documents. Specific assignments or assumptions developed in the assessment team become part of the product security risk management document.

4.5 Risk Mitigation Measures

After having performed the above described steps, the manufacturer (or the healthcare provider) has the relevant data available to define necessary risk mitigation measures. The goal is to develop measures which will best reduce the risk to an acceptable score for a specific system or for a specific healthcare provider. Once a mitigation plan is developed, the Risk Management Matrix offers a post-

mitigation risk estimation to explicitly determine the final risk. If new risks appear or risk scores are not low enough, then the process should be repeated.

It is important to note that risk mitigation can include system internal technical controls (e.g., network port closure), system external technical controls (e.g., fire-wall appropriately configured), or process description and training for key staff. In general, mitigations span technology, processes, and people. If a risk cannot sufficiently be mitigated in design control, the risk must be properly documented and assigned to the operational environment. External (technical) controls should be applied by the healthcare enterprise (e.g., intrusion detection).

At the conclusion of the process, the risk management team must approve the implemented risk mitigation measures and the summary of the residual risks.

4.6 Residual Risk Documentation and Executive Approval

As a final stage in risk management, a decision-maker with executive approval authority should be presented with a summary of the residual risks and subsequent mitigation plans, if any. The decision maker should take the assessment team's summary, combine it with a knowledge of how the system functions in advancing the mission of the healthcare organization, in order to reach a clear, well supported decision to deploy or not deploy the target system.

After this initial risk assessment, the remaining elements of risk management are integrated into total product life cycle management including:

- Manufacturing
- Sustaining engineering (including new functionality, configuration control, etc.)
- Incident management with Corrective Action/Preventive Action (CAPA)
- Changes to the security landscape
- Servicing
- Documentation
 - internal: risk management file including security assessment, sign-off, etc
 - external: security users guide including a list of residual risks to be managed by user

5 Conclusion

There are regulations and policies put in place in most countries protecting patient and staff safety, healthcare delivery (diagnosis and treatment), and privacy (safeguarding personal data). A careful security risk management process as described in this white paper will help to meet these goals. This process may be used during the development process at the manufacturer, as well as during the network (re)configuration process performed by the healthcare provider, e.g., when adding new networked equipment.

The security risk management process is similar to what is known in industry as a safety risk management process and may use the same tools. Both processes

may run in parallel. And, like safety, security risk assessments are revisited regularly or when a failure occurs. It is important to maintain a good partnership of all stakeholders (safety, security, workflow people) thus ensuring both the effective mitigation of security risks while advancing the healthcare mission

Appendix A – Product Security Risk Management Matrix

This is an example for a product security risk management matrix that fits the process described in this White Paper.

S = Severity of a successful attack, L = Likelihood of a successful attack

No.	Vulnerability (short description)	Cause(s), contributing factor(s)	Initial risk		Risk control type and mitigation	Reference to detailed specification / SW requirement specification	Reference to verification / validation reports (if passed)	Residual risk			
			L	S				L	S	S	
					1. By Design 2. By Protective measures 3. By process control (Manufacturing, Maintenance) 4. By Information for safety						

References and Resources for Risk Management

Alberts, Christopher, and Dorofee, Audrey. OCTAVE Method Implementation Guide v2.0. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2001. see <http://www.cert.org/octave/pubs.html> for more information.

Australian Standard AS4360:2004 Risk Management

Carnegie Mellon Software Engineering Institute, Software Risk Evaluation Method, Version 2.0

<http://www.sei.cmu.edu/pub/documents/99.reports/pdf/99tr029-body.pdf>

IEC 60812 Ed. 1.0: Analysis Techniques for System Reliability - Procedure for Failure Mode and Effects Analysis (FMEA)

IEC 61025 Fault Tree Analysis

IEC 61882 HAZOP Application Guide

IHE, Cookbook for the Security Section of IHE Profiles, Aug 20, 2006,

http://www.ihe.net/Technical_Framework/upload/IHE_ITI_TF_White_Paper_Security_Cookbook_PC_2006_08_30.pdf

ISO 14971:2006: Application of risk management to medical devices

ISO 17799 (2005) Information Technology - Code of practice for information security management

ISO 13335 (2004) Management of Information and Communications Technology Security

MIL-STD-1629A, Procedures for Performing a Failure Mode Effects and Criticality Analysis, November 24, 1980

NEMA/COCIR/JIRA SPC White Papers at www.nema.org/medical/spc

NIST Guidance for securing Microsoft Windows XP systems for IT professionals. http://csrc.nist.gov/itsec/guidance_WinXP.html

NIST SP 800-30: Risk Management Guide for Information Technology Systems <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>

NSA (US government National Security Agency), NSA Security Configuration Guides <http://www.nsa.gov/snac/>

SAE J-1739, Potential Failure Mode and Effects Analysis in Design and Potential Failure Mode and Effects Analysis in Manufacturing and assembly Processes Reference Manual, Aug 1, 2002

=====