

Technical Information Report

AAMI TIR57: 2016/(R)2023

Principles for medical
device security—Risk
management

Principles for medical device security—Risk management

Approved 5 June 2016 and reaffirmed 3 September 2019 and 31 January 2023 by
AAMI

Abstract: Provides guidance on methods to perform information security risk management for a medical device in the context of the Safety Risk Management process required by ISO 14971. The TIR incorporates the expanded view of risk management from IEC 80001-1 by incorporating the same key properties of Safety, Effectiveness and Data & Systems Security with Annexes that provide process details and illustrative examples.

Keywords: medical device, information security, risk management

AAMI Technical Information Report

A technical information report (TIR) is a publication of the Association for the Advancement of Medical Instrumentation (AAMI) Standards Board that addresses a particular aspect of medical technology.

Although the material presented in a TIR may need further evaluation by experts, releasing the information is valuable because the industry and the professions have an immediate need for it.

A TIR differs markedly from a standard or recommended practice, and readers should understand the differences between these documents.

Standards and recommended practices are subject to a formal process of committee approval, public review, and resolution of all comments. This process of consensus is supervised by the AAMI Standards Board and, in the case of American National Standards, by the American National Standards Institute.

A TIR is not subject to the same formal approval process as a standard. However, a TIR is approved for distribution by a technical committee and the AAMI Standards Board.

Another difference is that, although both standards and TIRs are periodically reviewed, a standard must be acted on—reaffirmed, revised, or withdrawn—and the action formally approved usually every five years but at least every 10 years. For a TIR, AAMI consults with a technical committee about five years after the publication date (and periodically thereafter) for guidance on whether the document is still useful—that is, to check that the information is relevant or of historical value. If the information is not useful, the TIR is removed from circulation.

A TIR may be developed because it is more responsive to underlying safety or performance issues than a standard or recommended practice, or because achieving consensus is extremely difficult or unlikely. Unlike a standard, a TIR permits the inclusion of differing viewpoints on technical issues.

CAUTION NOTICE: This AAMI TIR may be revised or withdrawn at any time. Because it addresses a rapidly evolving field or technology, readers are cautioned to ensure that they have also considered information that may be more recent than this document.

All standards, recommended practices, technical information reports, and other types of technical documents developed by AAMI are *voluntary*, and their application is solely within the discretion and professional judgment of the user of the document. Occasionally, voluntary technical documents are adopted by government regulatory agencies or procurement authorities, in which case the adopting agency is responsible for enforcement of its rules and regulations.

Comments on this technical information report are invited and should be sent to AAMI, Attn: Standards Department, 901 N. Glebe Road, Suite 300, Arlington, VA 22203.

Published by

AAMI
901 N. Glebe Road, Suite 300
Arlington, VA 22203

© 2016 by the Association for the Advancement of Medical Instrumentation

All Rights Reserved

This publication is subject to copyright claims of AAMI. Publication, reproduction, photocopying, storage, or transmission, electronically or otherwise, of all or any part of this document without the prior written permission of the Association for the Advancement of Medical Instrumentation is strictly prohibited by law. It is illegal under federal law (17 U.S.C. § 101, et seq.) to make copies of all or any part of this document (whether internally or externally) without the prior written permission of the Association for the Advancement of Medical Instrumentation. Violators risk legal action, including civil and criminal penalties, and damages of \$100,000 per offense. For permission regarding the use of all or any part of this document, contact the Copyright Clearance Center.

Printed in the United States of America

ISBN 978-1-57020-612-2

Contents

Page

Glossary of equivalent standards	iv
Committee representation	v
Foreword	vii
Introduction	viii
1 Scope.....	1
2 Terms and definitions.....	1
3 General guidance for performing security risk management.....	5
4 Security risk analysis	9
5 Security risk evaluation	12
6 Risk control	12
7 Evaluation of overall residual security risk acceptability.....	13
8 Security risk management report	14
9 Production and post-production information	14
Annex A (informative) Security engineering principles and nomenclature	16
Annex B (informative) Security risk assessment	21
Annex C (informative) Generating cybersecurity requirements	37
Annex D (informative) Questions that can be used to identify medical device security characteristics	39
Annex E (informative) Security risk examples applied to a medical device	49
Annex F (informative) A comparison of terminology between key referenced standards	65
Bibliography	68
Tables	
Table A.1 – Examples of security attributes and comparison between conventional IT and a medical device	17
Table B.1 - Description of Threat Tiers	27
Table E.1 - Security risk evaluation table	56
Table E.2 - Risk estimation analysis example	60
Table E.3 - Residual risk estimation analysis example	60
Table F.1 - Related terms in security standards/technical reports.....	65
Figures	
Figure 1 - Schematic representation of the risk management process (ANSI/AAMI/ISO 14971:2007).....	ix
Figure 2 – A Venn diagram showing the relationship between security and safety risks	x
Figure 3 - Schematic representation of the security risk management process	6
Figure 4 – Relationships between the security risk and safety risk management processes	7
Figure B.1 - A basic high-level risk assessment process	22
Figure B.2 - Security risk is assessed using three primary factors	25
Figure B.3 - Security risk assessment process	25
Figure B.4 - Cyber Threat Taxonomy.....	27
Figure B.5 - An example Threat-oriented Security Risk assessment approach	34
Figure B.6 - An example Asset-oriented Security Risk assessment approach.....	34
Figure B.7 - An example Vulnerability-oriented Security Risk assessment approach	35

Glossary of equivalent standards

International Standards or Technical Reports adopted in the United States may include normative references to other International Standards. AAMI maintains a current list of each International Standard that has been adopted by AAMI (and ANSI). Available on the AAMI website at the address below, this list gives the corresponding U.S. designation and level of equivalency to the International Standard.

www.aami.org/standards/glossary.pdf

Committee representation

Association for the Advancement of Medical Instrumentation

Medical Device Security Working Group

The publication of AAMI TIR57 as a new American Technical Information Report was initiated by the AAMI Medical Device Security Working Group.

At the time this document was published, the **AAMI Medical Device Security Working Group** had the following members:

Cochairs: Ken Hoyme, Adventium Labs
Geoff Pascoe, Deloitte Advisory

Members: Mike Ahmadi, Synopsys Inc
Pat Baird, Baxter Healthcare Corporation
Andrew Dean, Amgen Inc
Harsh Dharwad, Hospira Worldwide Inc
Sherman Eagles, SoftwareCPR
Scott Eaton, Mindray DS USA Inc
Plamena Entcheva-Dimitrov, Preferred Regulatory Consulting
Charles S. Farlow, Medtronic Inc.
Phil Fisk, Baxter Healthcare Corporation
Brian Fitzgerald, FDA/CDRH
Alan Fryer, Micro Systems Engineering Inc
Kevin Fu, The University of Michigan
Ken Fuchs, Center for Medical Interoperability
Bill Hagestad, Smiths Medical
Ed Heierman, Abbott Laboratories
Mike Jaffe, Cardiorespiratory Consulting LLC
Michelle Jump, Stryker Instruments Division
Joshua Kim, Hill-Rom Holdings
Insup Lee
Yimin Li, St Jude Medical Inc
Dan Lyon, Cigital Inc
Melissa Masters, Battelle Medical Products
Jill McCormick, Department of Veteran Affairs
Mary Beth McDonald, Mary Beth McDonald Consulting
Michael McNeil, Philips Electronics North America
Dale Nordenberg
Andrew O'Keeffe, Draeger Medical Systems Inc
Brodie Pedersen, Logic PD
Arnab Ray
Larry Schwartz, Smiths Medical
Michael Seeberger, Boston Scientific Corporation
Lynette Sherrill, Department of Veteran Affairs
Ferry Tamtoro, Amgen Inc
Tom Vaccaro, Becton Dickinson & Company
Fubin Wu, GessNet
Daidi Zhong, Chongqing University

Alternates: Tushar Dharampal, St Jude Medical Inc
Leo Espindle, Amgen Inc
Dawn Flakne, Micro Systems Engineering Inc
Elisabeth George, Philips Electronics North America
Roberta Hansen, Abbott Laboratories
Karen Kazak, Baxter Healthcare Corporation
Tara Larson, Medtronic Inc.

Nick Sikorski, Deloitte Advisory
Nikhil Thakur, FDA/CDRH
J.S. Wiley, Draeger Medical Systems Inc

NOTE—Participation by federal agency representatives in the development of this document does not constitute endorsement by the federal government or any of its agencies.

Foreword

This technical information report (TIR) was developed by the Device Security Working Group.

It is widely recognized that there is little existing guidance for conducting cybersecurity risk assessment of medical devices.

The objective of this TIR is to provide guidance on how medical device manufacturers can manage risks from security threats that could impact the confidentiality, integrity, and/or availability of the device or the information processed by the device. Because medical device manufacturers are already familiar with ANSI/AAMI/ISO 14971:2007, this guidance follows the basic structure of that standard.

Suggestions for improving this recommended practice are invited. Comments and suggested revisions should be sent to Technical Programs, AAMI, 4301 N. Fairfax Drive, Suite 301, Arlington, VA 22203-1633.

NOTE This foreword does not contain provisions of the AAMI TIR57, *Principles for medical device security– Risk management* (AAMI TIR57:2016), but it does provide important information about the development and intended use of the document.

Introduction

Medical device manufacturers are familiar with the requirements of ANSI/AAMI/ISO 14971:2007/(R)2010 *Medical devices — Application of risk management to medical devices*. This standard is an integral part of the safety risk management processes required by many regulatory authorities. ANSI/AAMI/ISO 14971 specifies a process for a manufacturer to identify the hazards associated with medical devices, including *in vitro* diagnostic (IVD) medical devices, to estimate and evaluate the associated risks, to control these risks, and to monitor the effectiveness of the controls (see Clause 1 of ANSI/AAMI/ISO 14971:2007).

NOTE In 2012, the European Committee for Standardization (CEN) adopted EN ISO 14971:2012 as the European harmonized standard, superseding EN ISO 14971:2009. This document does not address content deviations included in Annex ZA of EN ISO 14971:2012. Specifically, the “as far as possible” requirement is not included in the evaluation of security risks. Instead, security risks are to be assessed and controlled to a level that is considered acceptable, taking into account the impact of a threat event and potential vulnerabilities.

Specific clauses of ANSI/AAMI/ISO 14971:2007 define a risk management process consisting of the following elements:

- risk analysis (Clause 4);
- risk evaluation (Clause 5);
- risk control (Clause 6);
- evaluation of overall residual risk acceptability (Clause 7);
- risk management report (Clause 8); and
- production and post-production information (Clause 9).

Figure 1 of ANSI/AAMI/ISO 14971:2007 (see Figure 1) provides a schematic representation of the risk management process. Central to the definition of risk are the concepts of probability of an occurrence of harm and the severity of that harm. Harm is defined in ANSI/AAMI/ISO 14971:2007 as “physical injury or damage to the health of people, or damage to property or the environment”.




Figure 1 - Schematic representation of the risk management process (ANSI/AAMI/ISO 14971:2007)

The increased use of connected medical devices (i.e., devices directly connected to other systems or devices directly or via a computer network) has created a new source of risk for their safe operation. Security risks are associated with reduction of effectiveness and breach of data and systems security as included in the broader definition of harm in this document. While information security has been considered from the patient data privacy perspective for several years, there is no framework for security risk management for medical devices. This document describes a means of applying the risk management principles presented in ANSI/AAMI/ISO 14971 to the management of security risk.

The definition of harm is considered from the perspective of ANSI/AAMI/ISO 14971, as well as from healthcare information technology (IT) standards, such as the ANSI/AAMI/IEC 80001 family. Because a security risk management process that narrowly focuses on the traditional “physical injury or damage” definition may limit the scope of security risk mitigation, this document incorporates the broader considerations that risks include effects outside the traditional scope of patient physical harm and may include “reduction of effectiveness” and “breach of data and systems security” as extended in the ANSI/AAMI/IEC 80001 family of standards. Figure 2 shows the high-level relationship between these two classes of risk.




Figure 2 – A Venn diagram showing the relationship between security and safety risks

Considerations for security risk management are mapped to the risk management process steps of ANSI/AAMI/ISO 14971, as well as the security risk management process of NIST SP 800-30 Revision 1 *Guide for Conducting Risk Assessments* (see Bibliography [53]), with a particular emphasis on risk assessment methods. While this TIR utilizes NIST standards as a basis, manufacturers may use other generally accepted alternatives if they are better suited to their development processes.

Supporting annexes contain the following:

- Annex A: Security engineering principles and nomenclature – A high level overview of security engineering principles focused on their application to the medical device domain.
- Annex B: Security risk assessment – A more detailed description of the process for performing security risk assessment for a medical device based on the principles described in NIST SP 800-30 Revision 1 (see Bibliography [53]).
- Annex C: Generating cybersecurity requirements – A discussion on how to create effective, testable requirements for security properties that avoid complexities associated with “shall not” requirements. This is important since mitigating controls for security risks need to be expressed in the requirements for the device so proper verification of effectiveness can be shown.
- Annex D: Questions that can be used to identify medical device security characteristics – A detailed list of questions for manufacturers to assist in exploring security aspects of their devices, organized along the lines of ANSI/AAMI/ISO 80001-2-2:2010.
- Annex E: Security risk examples applied to medical products – A more detailed example of a fictional medical device and analysis of its security risks.
- Annex F: A comparison of terminology between key referenced standards – A comparison of the basic terms of risk management as defined in ANSI/AAMI/ISO 14971:2007, NIST SP 800-30 Revision 1, IEC 80001-1:2010 and this TIR.

While safety risk involves evaluating the probability and severity of a hazard leading to harm, security risk is based on an assessment of the likelihood that a threat will successfully exploit a device vulnerability, an event that could lead to an adverse impact due to a compromise of system confidentiality, integrity, and/or availability. This loosely parallels the concepts of probability and severity described in ANSI/AAMI/ISO 14971. Organizations should exercise caution when attempting to quantify likelihood of a future adverse impact in traditional probabilistic terms. Instead, they should focus on the skills and motivations of an attacker, and whether the effort required to exploit a vulnerability is less than the perceived gain the attacker will achieve by compromising the system (see B.2.1.1).

Security is an emergent property of a system. As such, the security of a device needs to be considered in the context of the broader system. This document directs manufacturers to understand, evaluate, and document the operating environment of a device so security risks are evaluated in their operational context. It is recognized that the manufacturer cannot ensure aspects of this operating environment that are outside the manufacturer’s control. However, it is incumbent on the manufacturer to understand the methods that Health Delivery Organizations (HDOs)

use to manage the risks of networked medical devices as specified in ANSI/AAMI/ISO 80001-1, as well as the methods used to communicate cybersecurity needs, risks, and controls as documented in ANSI/AAMI/ISO 80001-2-2. Manufacturers should understand the concept of a Responsibility Agreement¹ that may be negotiated with the purchasing HDO, and be prepared to communicate any specific security expectations of the network that the device is connected to through methods such as the HIMSS/NEMA Manufacturer Disclosure Statement for Medical Device Security (MDS2). Manufacturers should also recognize that a poorly secured and updated device could be a source of security vulnerability to other computing systems and other devices to which the device is connected, directly or on a shared network. Such a device can be used as a “pivot” to attack other systems and devices. A reasonable expectation of organizations deploying the device is that such risks have been considered and appropriately mitigated.

As a result, this document uses the broader definition of risk as used in the ANSI/AAMI/IEC 80001 family of standards. Manufacturers that follow the recommendations of this report should consider those risks that come from outside the device, mitigate those that are feasible, and document the expectations of the organization that becomes responsible for integrating the device into a broader network.

The framework described in this document is presented as a companion process to the safety risk management requirements of ANSI/AAMI/ISO 14971. This is similar to the risk management approach documented in ANSI/AAMI/IEC 62366-1:2015 for usability engineering. The key to successful integration of a device security risk management process into the larger risk management process is to bring together personnel with expertise in traditional device development, human factors, and cybersecurity product development. This collaboration will be essential to the ultimate goal of developing secure medical devices.

¹ Responsibility Agreements are described in ANSI/AAMI/ISO 80001-1.



Principles for medical device security – Risk management

1 Scope

This TIR provides guidance for addressing information security within the risk management framework defined by ANSI/AAMI/ISO 14971.

This guidance is intended to assist manufacturers and other users of the standard in the following:

- identifying threats, vulnerabilities, and assets associated with medical devices;
- estimating and evaluating associated security risks;
- controlling security risks; and
- monitoring effectiveness of the risk controls.

This document is based on an application of ANSI/AAMI/ISO 14971 with an expanded consideration of the possible impacts that a security compromise can have on the medical device, people, the environment, the manufacturer, and the information processed and stored by the device. This report also incorporates several principles from NIST SP 800-30 Revision 1 (see Bibliography [53]), a security risk management process developed for traditional IT systems.

The guidance provided by this document is applicable to all stages of the life-cycle of a medical device.

2 Terms and definitions

For the purposes of this document, the following terms and definitions apply:

2.1

asset

person, structure, facility, information, and records, information technology systems and resources, material, process, relationships, or reputation that has value

[SOURCE: NICCS Glossary of Common Cybersecurity Terminology – As accessed on June 15, 2015.]

2.2

authentication

verifying the identity of a user, process, or device, often as a prerequisite to allowing access to resources in an information system

[SOURCE: SP 800-53; SP 800-53A; SP 800-27; FIPS 200; SP 800-30]

2.3

authenticity

property of being genuine and being able to be verified and trusted; confidence in the validity of a transmission, a message, or message originator (see Authentication)

[SOURCE: SP 800-53; SP 800-53A; CNSSI-4009; SP 800-39]

2.4

authorization

access privileges granted to a user, program, or process, or the act of granting those privileges

[SOURCE: CNSSI-4009]

2.5

availability

ensuring timely and reliable access to and use of information

NOTE 1 to entry: The phrase “use of information” encompasses delivery of intended functionality.

[SOURCE: SP 800-53; SP 800-53A; SP 800-27; SP 800-60; SP 800-37; FIPS 200; FIPS 199; 44 U.S.C., Sec. 3542]

2.6

confidentiality

preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information

[SOURCE: SP 800-53; SP 800-53A; SP 800-18; SP 800-27; SP 800-60; SP 800-37; FIPS 200; FIPS 199; 44 U.S.C., Sec. 3542]

2.7

data and systems security

operational state of a medical device in which information assets (data and systems) are reasonably protected from degradation of confidentiality, integrity, and availability

NOTE 1 to entry: Security, when mentioned in this document, should be taken to include data and systems security.

[SOURCE: ANSI/AAMI/IEC 80001-1:2010, 2.5, modified - “MEDICAL IT-NETWORK” has been replaced with “medical device” and Note 2 was redacted.]

2.8

effectiveness

ability to produce the intended result for the patient and the care provider

[SOURCE: ANSI/AAMI/IEC 80001-1:2010, 2.6, modified - “RESPONSIBLE ORGANIZATION” has been replaced with “care provider”.]

2.9

emergency access

process or mechanism by which a device user can quickly and easily access the intended functionality in urgent (emergency) situations, bypassing the device’s established access controls; the ability of the device user to access the intended functionality in case of an emergency situation that requires immediate access to the medical device

NOTE 1 to entry: Other access methods (e.g., “break glass”) fall under this general definition but have varying levels of credentials and audit requirements.

[SOURCE: Adapted from HIMSS/NEMA Standard HN 1-2013 Manufacturer Disclosure Statement for Medical Device Security]

2.10

encryption

conversion of plaintext to ciphertext through the use of a cryptographic algorithm

[SOURCE: FIPS 185]

2.11

harm

physical injury or damage to the health of people, or damage to property or the environment, or reduction in effectiveness, or breach of data and systems security

[SOURCE: IEC 80001-1:2010, definition 2.8]

2.12

hazard

potential source of harm

[SOURCE: ISO/IEC Guide 51:1999, definition 3.5]

2.13

hazardous situation

circumstance in which people, property, or the environment are exposed to one or more hazard(s)

NOTE 1 to entry: See Annex E, ANSI/AAMI/ISO 14971:2007, for an explanation of the relationship between “hazard” and “hazardous situation”.

[SOURCE: ISO/IEC Guide 51:1999, definition 3.6]

2.14

information security

protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide confidentiality, integrity, and availability

[SOURCE: SP 800-37; SP 800-53; SP 800-53A; SP 800-18; SP 800-60; CNSSI-4009; FIPS 200; FIPS 199; 44 U.S.C., Sec. 3542]

2.15

integrity

guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity

NOTE 1 to entry: This definition includes therapy, device, and application integrity.

[SOURCE: SP 800-53; SP 800-53A; SP 800-18; SP 800-27; SP 800-37; SP 800-60; FIPS 200; FIPS 199; 44 U.S.C., Sec. 3542]

2.16

likelihood of occurrence

weighted factor based on a subjective analysis of the probability that a given threat is capable of exploiting a given vulnerability

NOTE 1 to entry: Likelihood of occurrence combines an estimate of the likelihood that the threat event will be initiated with an estimate of the likelihood of impact (i.e., the likelihood that the threat event results in adverse impacts).

[SOURCE: CNSSI-4009, modified - the phrase "In Information Assurance risk analysis," was removed.]

2.17

non-repudiation

assurance that the sender of information is provided with proof of delivery and the recipient is provided with proof of the sender's identity, so neither can later deny having processed the information

[SOURCE: CNSSI-4009; SP 800-60]

2.18

password

protected/private string of letters, numbers, and/or special characters used to authenticate an identity or to authorize access to data

[SOURCE: CNSSI-4009]

2.19

personally identifiable information (PII)

any information about an individual maintained by an agency, including (1) any information that can be used to distinguish or trace an individual's identity, such as name, social security number, date and place of birth, mother's maiden name, or biometric records; and (2) any other information that is linked or linkable to an individual, such as medical, educational, financial, and employment information

NOTE 1 to entry: Personally identifiable information is a superset of Protected Health Information (PHI).

[SOURCE: NIST SP 800-122; GAO Report 08-536]

2.20

predisposing condition

condition that exists within an organization, a mission/business process, enterprise architecture, or information system, including its environment of operation, which contributes to (i.e., increases or decreases) the likelihood that one or more threat events, once initiated, will result in undesirable consequences or adverse impact to organizational operations and assets, individuals, or other organizations

NOTE 1 to entry: Identical to NIST definition (SP 800-30 Revision 1) with the phrase "or the Nation" redacted.

[SOURCE: NIST SP 800-30 Revision 1]

2.21

residual risk

risk remaining after risk control measures have been taken

NOTE 1 to entry: Adapted from ISO/IEC Guide 51:1999, definition 3.9.

NOTE 2 to entry: ISO/IEC Guide 51:1999, definition 3.9 uses the term “protective measures” rather than “risk control measures.” However, in the context of this Technical Information Report, “protective measures” are only one option for controlling risk as described in 6.2.

2.22

risk

combination of the probability of occurrence of harm and the severity of that harm

[SOURCE: ISO/IEC Guide 51:1999, definition 3.2]

2.23

risk analysis

systematic use of available information to identify hazards and to estimate the risk

NOTE 1 to entry: Risk analysis includes examination of different sequences of events that can produce hazardous situations and harm.

[SOURCE: ISO/IEC Guide 51:1999, definition 3.10]

2.24

risk assessment

overall process comprising a risk analysis and a risk evaluation

[SOURCE: ISO/IEC Guide 51:1999, definition 3.12]

2.25

risk control

process in which decisions are made and measures are implemented by which risks are reduced to, or maintained within, specified levels

[SOURCE: ANSI/AAMI/ISO 14971, definition 2.19]

2.26

risk evaluation

process of comparing the estimated risk against given risk criteria to determine the acceptability of the risk

[SOURCE: ANSI/AAMI/ISO 14971, definition 2.21]

2.27

safety

freedom from unacceptable risk

[SOURCE: ISO/IEC Guide 51:1999, definition 3.1]

2.28

threat

any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, or other organizations through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service

NOTE 1 to entry: Identical to NIST definition (SP 800-53) with the phrase “or the Nation” redacted.

[SOURCE: SP 800-53; SP 800-53A; SP 800-27; SP 800-60; SP 800-37; CNSSI-4009]

2.29

threat actor

individual, group, organization, or government that conducts or has the intent to conduct detrimental activities

NOTE 1 to entry: Synonymous with threat agent.

[SOURCE: NICCS Glossary of Common Cybersecurity Terminology – As accessed on June 15, 2015.]

2.30

threat analysis

examination of threat sources against system vulnerabilities to determine the threats for a particular system in a particular operational environment

[SOURCE: SP 800-27]

2.31

threat event

event or situation that has the potential for causing undesirable consequences or impact

[SOURCE: SP 800-30]

2.32

threat source

intent and method targeted at the intentional exploitation of a vulnerability or a situation and method that may accidentally trigger a vulnerability

[SOURCE: FIPS 200; SP 800-53; SP 800-53A; SP 800-37]

2.33

top management

person or group of people who direct(s) and control(s) a manufacturer at the highest level

NOTE 1 to entry: Adapted from ISO 9000:2005, definition 3.2.7.

2.34

vulnerability

weakness in an information system, system security procedures, internal controls, or implementation that could be exploited or triggered by a threat source

[SOURCE: SP 800-53; SP 800-53A; SP 800-37; SP 800-60; SP 800-115; FIPS 200]

3 General guidance for performing security risk management

3.1 Security risk management process

ANSI/AAMI/ISO 14971 requires that device manufacturers establish a risk management process for safety hazards to ensure that hazards are identified, risks are estimated, monitoring is established, and controls are put in place for those risks that require mitigation.

Security risk management should follow a similar path, as illustrated in Figure 3. While organizations may choose to incorporate security risks into their general safety risk management process, this document recommends the creation of a separate risk analysis process focused specifically on impacts that are identified by a security analysis. Because harm, as defined in this document, may include breach of data and system security and reduction in device effectiveness (as defined in ANSI/AAMI/IEC 80001-1), it is more appropriate to create a companion security risk management process to allow the organization to assess the additional risks associated with effectiveness and system/data security. If the processes are integrated, there could be an inclination to drop the evaluation of those risks that do not lead to harm as narrowly defined in ANSI/AAMI/ISO 14971, which can lead to incomplete or inconsistent security controls.

This segregation is further supported by the fact that security risk assessment models typically use assessment factors that are different from the model described by ANSI/AAMI/ISO 14971. In addition, integrating safety and security risk assessment into a single general risk management process may result in major modifications to a well-functioning safety risk management process.

Security risks that impact safety, according to the definition of harm in ANSI/AAMI/ISO 14971 (“physical injury or damage to the health of people, or damage to property or the environment”), should also be captured in the organization’s safety risk management process. It is recognized that treatment of risk in these two risk assessment models, safety risk assessment and security risk assessment, may differ. A specific risk assessed as “must mitigate” in one model might be assessed as “does not need further mitigation” in the other. Risk control measure(s) should be applied to bring the risk into the acceptable range in both assessment models.




Figure 3 - Schematic representation of the security risk management process

3.1.1 Relationship between security and safety risk management

This document recommends that manufacturers establish a companion security risk management process to their existing ANSI/AAMI/ISO 14971-based safety risk management process. The requirements of ANSI/AAMI/ISO 14971 indicate that all hazards need to be managed through the safety risk management process – including those hazards that result from security compromises. However, because the definition of “harm” in this document is broader than the definition used in ANSI/AAMI/ISO 14971, there will be risks managed in the security risk assessment that are not propagated to the safety risk management process. An example would be a risk of compromise of the confidentiality of protected health information that is not considered harm in the context of ANSI/AAMI/ISO 14971, but clearly requires mitigation by the security risk management process. There are also business and reputation risks associated with a security compromise that are not considered harm in the safety sense.

When security risks may lead to safety risks, security and safety staff should work jointly to contribute to the security risk analysis and transfer safety related hazards to the safety analysis. An example of a security risk that is also a safety risk is a malicious attacker gaining access to a medical device's code, altering that code, and causing the device to malfunction. This malfunction may have the potential to cause harm to the patient. However, when a security risk does not involve safety risk, personnel qualified in security are typically sufficient for risk analysis. An example of a security risk that is not a safety risk is a malicious attacker who gains access to patient data. The patient data obtained poses a security risk but does not directly affect safety.

Because the security controls that mitigate safety risks caused by security compromise may simultaneously mitigate non-safety related security risks, it is logical to manage all security risks in a single process. That is, a security compromise that leads to harm (as more narrowly defined in ANSI/AAMI/ISO 14971) should be managed within the security risk management process and propagated for assessment using the organization's safety risk management process.

As illustrated in Figure 4, when a risk control measure is introduced to a design, the design must be reassessed to determine if the control measure introduces a new form of risk. It should be recognized that there is a coupling between safety and security risk assessment processes, so when control measures are introduced for one type of risk (e.g., safety), the manufacturer needs to assess the impact on the other type of risk (e.g., security) and vice versa. For example, the decision to add risk control measures for authentication might introduce risks that the device cannot be accessed in an emergency. The overall risk management process should identify those points of coupling and ensure that assessment of any newly identified source of risk is performed in both domains. This is illustrated in Figure 4, showing the logical connection points between the security and safety risk management processes.




Figure 4 – Relationships between the security risk and safety risk management processes

Typically, these assessments iterate between the security and safety risk management processes. Figure 4 shows an example of one such iteration.

3.2 Management responsibilities

Top management should provide adequate and qualified personnel (see 3.3) to perform the security risk management process. In addition, top management needs to define and document the policy for determining the criteria for security risk acceptability, including criteria for non-safety related security risks, such as business and reputation risks. This policy should ensure that criteria are based upon applicable national or regional regulations and relevant international standards, and take into account available information, such as the generally accepted state of the art for security risk management and known stakeholder security concerns. Stakeholders include medical practitioners, Health Delivery Organizations (HDOs), governments, industry, patients, and members of the public.

Top management should ensure that the organization implements a coordinated vulnerability disclosure policy so security researchers and others have a means to communicate vulnerabilities identified in fielded devices.

Top management should review the suitability of the security risk management process at planned intervals to ensure the continuing effectiveness of the process, and they should document any decisions and actions taken.

3.3 Qualification of personnel

Persons performing security risk management tasks need to have the knowledge and experience appropriate to the tasks assigned to them. These should include knowledge of the particular medical device (or similar medical device) and its intended use, operating environments, security risk control technologies, and the risk management process.

Personnel should be knowledgeable in evaluating security threats and vulnerabilities, and be aware of past and emerging changes in the security risk landscape, both for medical devices, as well as general purpose computing systems. The team performing security risk management tasks should have experience with hardware and software architecture, design, and security test methods.

Appropriate qualification records should be maintained.

3.4 Security risk management plan

Security risk management should be incorporated into the manufacturer's risk management process. As discussed in 3.1, this document recommends that the security risk management process be maintained as a companion process to the safety risk management process. Because the two processes exchange information on the identified risks and candidate risk controls, the risk management plans for the two processes should be coordinated. However, it is acceptable for a manufacturer to incorporate the security risk management plan as part of the overall risk management plan developed in compliance with ISO 14971.

The scope of risk management activities includes the identification of the medical device features and life-cycle phases where security risks need to be managed. To properly determine the scope of security risk management activities, the medical device must be assessed to determine which features expose a potential security risk based on the intended use and operating environment. Because threats change over time and new vulnerabilities in operating systems, middleware and components are discovered on regular basis, security risks are frequently identified after a device is released to the market. Consequently, as threats evolve, it is critical to have a plan for post-market monitoring, reporting, analyzing, and responding to newly identified threats and vulnerabilities.

The criteria for risk evaluation, impact assessment, and residual risk acceptability should be documented in the plan. Because a vulnerability requires exploitation before harm can occur, there is a behavioral aspect to risk estimation that is not amenable to mathematical probability estimation. Because vulnerability requires exploitation before harm can occur, and exploitation involves human behavior, risk estimation is not amenable to mathematical probability estimation. Top management plays a critical role in defining a policy for determining the criteria for security risk acceptability (see 3.2).

The plan should describe the methodology for security risk verification. Some possible alternatives include the following:

- a) verification testing of security controls that mitigate risks, including effectiveness of the mitigation;
- b) robustness testing (e.g., "fuzz testing");
- c) penetration testing (typically performed in a production equivalent or final configured state).

Manufacturers should also document any additional development methodologies that should be applied to reduce the likelihood of unintended functionality being present in the code that could become a vulnerability once deployed. This may include code analysis (e.g., inspections, static code checking) or particular verification testing strategies to increase confidence in the robustness of the implementation.

Consideration should be given to appropriate security testing for the intended use and operating environment (e.g., physical and network).

The security risk management plan should address how components (hardware and software) will be monitored for security performance and how vulnerabilities identified in those components are elevated for risk evaluation in the component's use in the finished medical device. Sources of information on new vulnerabilities for components include threat intelligence, reports from suppliers and the MITRE Common Vulnerabilities and Exposures (CVE) list. The plan should ensure that mechanisms for post-market cybersecurity surveillance of all third-party components will be identified.

The plan should describe how product security performance information will be collected in fielded devices and reported back to the manufacturer for assessment and potential control. The plan should reference the organization's coordinated vulnerability disclosure policy as an additional input source.

The security risk management plan should include the requirements for an appropriate review interval for newly identified threats and vulnerabilities.

NOTE See Clause 9 for additional guidance on recommended post-market surveillance activities and methods.

Manufacturers should plan for the appropriate disclosure of identified vulnerabilities.

Manufacturers should plan for a secure and efficient software update process to manage vulnerabilities in fielded devices and enable prompt response when required.

3.5 Security risk management file

The security risk management file may be integrated with the overall risk management file or be maintained separately.

The security risk management file should contain, at a minimum, the following:

- a) security risk management plan;
- b) security risk analysis components, such as:
 - 1) intended use and reasonably foreseeable misuse;
 - 2) qualitative and quantitative characteristics of the system that could affect the security of the medical device and, where appropriate, their defined limits; and
 - 3) risk controls and their implementation and verification;
- c) security risk controls (typically captured and flowed down as system requirements specifications).

4 Security risk analysis

4.1 Security risk analysis process

Security risk analysis should be performed for the medical device as described in 4.2 to 4.4. The results of the security risk analysis should be recorded in the security risk management file.

Security needs to be assessed in the context of the larger system in which the device operates. The security risk analysis needs to consider the intended use and document the anticipated operating environment for which the device was designed and tested. The security architecture (see A.6) includes the complete operating environment of the device, as well as other important factors. These steps will facilitate communication to end user(s) responsible for device configuration.

If a security risk analysis, or other relevant information, is available for a similar medical device, then that analysis or information can be used as a starting point for the new analysis. The degree of relevance depends on the differences between the devices and whether these introduce new risks or significant vulnerability differences. The extent of reuse should be based on a systematic evaluation of the effects the changes have on exposing the patients, users, and manufacturer to additional risks.

In addition to the records recommended in 4.2 to 4.4, the documentation of the conduct and results of the security risk analysis should include at least the following:

- a) a description and identification of the medical device that was analyzed;
- b) identification of the person(s) and organization carrying out the security risk analysis; and
- c) scope and date(s) of the security risk analysis.

4.2 Intended use and identification of characteristics related to the security of the medical device

For the particular medical device being considered, the manufacturer should document the intended use and reasonably foreseeable misuse. Reasonably foreseeable misuse should include efforts that normal users might make to circumvent security controls when they are perceived as preventing their use of the device.

NOTE 1 In this context, misuse is intended to mean incorrect or improper use of the medical device.

NOTE 2 The exploration of potential malicious abuse by attackers will be accomplished during the threat analysis activity documented in 4.3.

NOTE 3 Annex D contains questions that can serve as a useful guide in identifying medical device characteristics that could have an impact on security.

The manufacturer should document the assumed operating environment and security architecture for which the device is designed to operate, along with any assumptions on external security controls that must be provided by the end user. This documentation should be maintained in the security risk management file.

HDOs may have different and varying cybersecurity needs, risks and controls depending on the budget and sophistication of the organization. The manufacturer should perform a needs assessment with a representative sample of HDOs prior to initiating product design.

The manufacturer should document characteristics of the system that rely on user configuration to ensure the security of the device. The device's intended use and the experience level of the intended users should be understood to more accurately understand the likelihood of proper configuration by the end user.

The risk analysis should address characteristics of the device and its expected operating environment (physical and IT). For example, the risk analysis should consider whether the device is mobile and/or expected to be physically accessible to unauthorized users. Interoperability requirements and constraints of the operating environment should also be considered and documented. For medical devices that use wireless technology with a discovery mode or similar active connection mode, the manufacturer should implement appropriate technical controls to prevent unauthorized users from sensing or connecting to the medical device. In some cases, a manufacturer may document recommended compensating controls for implementation in the user environment.

4.3 Identification of threats, vulnerabilities, assets, and adverse impacts

In order to assess security risk, several factors need to be identified and documented.

The analysis of threats, vulnerabilities, assets, and adverse impacts can be performed in any order and may require participants with different skills and experiences. For example, understanding the adverse impacts may require members with deeper clinical understanding of the device's intended use, while the understanding of threats and their characteristics will require members with deeper security backgrounds.

This is expected to be an iterative process where continued brainstorming and discussion results in new aspects of the medical device being identified that could be affected by additional threats and vulnerabilities. The conclusion of the process is expected to provide a carefully considered and extensive listing of threats, vulnerabilities, and adverse impacts that feeds into the next step of estimating the risks.

NOTE See Annex B for examples of threats, vulnerabilities, assets, and adverse impacts.

4.3.1 Identification of threats

When assessing risk, special consideration should be given to threat actor capability and the benefits they might gain from a potential exploit. The risk analysis should document potential threats and the means a threat actor might use to exploit a vulnerability.

Threats can be identified using information from a variety of sources, including the manufacturer's IT security department, as well as from the security departments of key customers (e.g., large hospitals or clinics where the devices may be deployed). Information should also be gathered from reports of threats that have been experienced by the healthcare system, and from similar adjacent domains, such as industrial control systems. Governmental agencies, such as ICS-CERT, Information Sharing and Analysis Centers (ISACs) and Organizations (ISAOs), exist for the purpose of sharing this information. Information from third-party security information aggregators may also be helpful in gaining a thorough understanding of threat actors, their capabilities, and typical means of creating a threat event.

NOTE See Annex D and Annex E in NIST SP 800-30 Revision 1 (see Bibliography [53]) for additional material regarding threat identification.

4.3.2 Identification of vulnerabilities

The risk analysis should document known and potential vulnerabilities in the device. Vulnerabilities come from three distinct sources. First, vulnerabilities may be introduced by conscious design decisions, such as whether to include a security control or allow its bypass (e.g., in emergency access situations). The second source of vulnerabilities is from errors in the design, implementation, manufacture, or configuration of the device, by the manufacturer, their suppliers, or end users. These design errors are cases where the device does not meet the intent of the designers through a process deficiency. A third source of vulnerabilities is from a design characteristic that at the time of design release was not known to be a vulnerability, but subsequent to design release, a means of exploiting the design characteristic for malicious purposes was discovered. When assessing risk, special consideration should be given to the exploitability of the vulnerability (i.e., the capability necessary for a threat actor to exploit the vulnerability).

Vulnerabilities can be identified by incorporating top-down analysis methodologies (e.g. attack trees, threat modelling) to identify ways that a threat can cause loss of security properties for the device. Standards such as ANSI/AAMI/ISO

80001-2-2 list security capabilities, and this information can also identify potential vulnerabilities in the requirements. The identification process is best started early in the concept phase of the device development lifecycle but should also be applied to products in advanced stages of development or already in service.

Identification of possible points where a design/implementation error may introduce a vulnerability can be initiated with an analysis of the system to determine system nodes (components and users) and communication pathways. Once these are known for a system, they can be assessed for potential vulnerabilities that can affect each of these nodes and/or communication links. Various sources of information may also be used to identify vulnerabilities of the system under review. These sources include knowledge gained from past products, academic publications, industry news sources, independent security firms, independent research, and future projections. In addition, security testing (e.g., vulnerability scanning and penetration testing) may be used in early design phases to assist in identifying potential vulnerabilities. It is critical that the team performing the security risk assessment be knowledgeable on these and other pertinent sources of information so that a comprehensive and realistic identification of threats and vulnerabilities occurs.

NOTE See Annex F in NIST SP 800-30 Revision 1 (see Bibliography [53]) for additional material regarding vulnerabilities.

4.3.3 Identification of assets

Document the assets of the device. Assets can include information (e.g. patient data, diagnostic data, therapy parameters) the device itself, or components of the device including the device software. Assets may also be the physical interfaces to the device, in particular connections to external networks which could present risks to other systems on the network if unprotected at the device level.

Asset identification can be done through a systematic process, looking at each asset on the device and any interface that could be manipulated by an attacker. Asset identification should include those assets needed by the user to manage, diagnose, and treat the patient, as well as on those assets required to keep the device operating safely and securely. The process should consider assets needed to maintain the device or any used as part of post-market surveillance. Particular attention should be afforded to any stored authentication or cryptographic credentials that are used by the device, as their compromise could lead to the subsequent compromise of many other assets.

4.3.4 Identification of adverse impacts

For each identified asset, consider the impact that loss of confidentiality, loss of integrity, or loss of availability might have on safety, effectiveness, or data or system security.

Impact of asset compromise involves examining each asset and answering at least the following questions:

- a) What is the impact if that asset's confidentiality was compromised, and the information it contained was available to an attacker?
- b) What is the impact if that asset's integrity was compromised, and an attacker could change the information or software in a way that wasn't immediately obvious to the user (clinician or patient)?
- c) What is the impact if that asset could be made unavailable?
- d) Could the immediate impact of an asset compromise lead to another type of attack or vulnerability?

NOTE See Annex F in NIST SP 800-30 Revision 1 (see Bibliography [53]) for additional material regarding impact assessment.

4.4 Estimation of the risk(s) for each applicable threat and vulnerability combination

In the NIST SP800-30 Revision 1 model (see Bibliography [53]), security risk is estimated by a combination of the threat, vulnerability, and impact of asset compromise. When compared to the safety risk analysis process, the combination of threat and vulnerability is analogous to the probability of occurrence. The impact of asset compromise is analogous to the severity of harm.

The device manufacturer should establish a documented and repeatable risk estimation process that allows the security risk to be evaluated for acceptability. While the measures for the individual factors may be qualitative, it should be possible for subject matter experts to agree on the valuations of each factor, and the process by which they are combined results in a ranking of risk that can be used to decide which risks require further mitigation.

The outcome of this activity depends on the method employed by the manufacturer, but it is expected to be in the form of a ranking or grouping of the threats and vulnerabilities with possible additional details provided on rationale for the scoring.

NOTE See Annex G, H and I in NIST SP 800-30 Revision 1 (see Bibliography [53]) for additional material regarding risk estimation.

5 Security risk evaluation

Each identified risk needs to be evaluated for whether security risk reduction is required, according to the criteria documented in the security risk management plan. Each risk should also be evaluated to determine if a risk must be promoted to the safety risk file for analysis. As discussed in 3.1, security risks that can lead to hazardous situations need to be assessed in the safety risk analysis as well. The risk level should be mitigated (i.e., risk controlled) to an acceptable level in both analyses.

If reduction is determined to not be required, the manufacturer should document the rationale behind that decision and what changes in the threat environment would trigger reevaluation.

6 Risk control

6.1 Security risk reduction

Paralleling ANSI/AAMI/ISO 14971, when security risk reduction is required, the risk control activities in this clause should be performed.

6.2 Security risk control option analysis

Manufacturers should use a security risk control hierarchy that parallels the safety hierarchy presented in ANSI/AAMI/ISO 14971. Risk control options in order of preference are listed below:

- a) Inherent security by design - Inherent security by design may include options, such as preventing access and/or modification of device settings from other systems or over the networked interface, implementation of secure certificates for authentication and code signing, incorporation of defense-in-depth, and various hardening mechanisms.
- b) Protective measures in the medical device itself or in the manufacturing process - Examples of protective measures include requiring physical proximity, ensuring that malware detection is part of the software manufacturing and deployment process, and the addition of intrusion detection systems to detect attempted attacks.
- c) Information for security - Information for security would include documenting security requirements that are placed on the entity or individual patient that connects the device directly to other systems or integrates the device on a network. Just as labeling for safety is considered a weak mitigation, leveraging excessive requirements on the network configuration to ensure security is least desirable, as one cannot be assured that all networks will be so configured, and it raises the potential that different devices on the same network are subject to incompatible requirements.

Manufacturers should be aware of the ANSI/AAMI/ISO 80001-1 concept of Responsibility Agreements and be prepared to supply the HIMSS/NEMA Manufacturer Disclosure Statement for Medical Device Security (MDS2) to those users of networked devices that request this document.

Manufacturers should consider the ability to respond to future unknown security risks as a key risk control measure, as there may be unknown vulnerabilities present in the designed medical device that become known as threats become more sophisticated. This includes the ability to securely deliver patches after the device is in service.

If, during risk control option analysis, the manufacturer determines the required security risk reduction is not practicable, the manufacturer should conduct a risk/benefit analysis of the residual risk (see 6.5).

6.3 Implementation of risk control measure(s)

The security risk control measures selected need to be implemented and verified with the results recorded in the security risk management file. These control measures will typically be expressed as testable requirements. Annex C provides guidance on documenting security requirements.

6.4 Residual risk evaluation

Residual risk is re-evaluated after the security risk controls have been applied. A risk control measure may mitigate either the threat, vulnerability, or impact. This should be documented to reflect the residual risk estimation. Further risk controls may need to be applied if the residual risk is unacceptable. Manufacturers need to determine what information about acceptable residual risks should be disclosed in the device labeling.