



CIS RAM Version 1.0

Center for Internet Security®

Risk Assessment Method

**For Reasonable Implementation and
Evaluation of CIS Controls™**

CIS RAM - Center for Internet Security® Risk Assessment Method (Version 1.0)

April 2018

This work is licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License (the link can be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

CIS RAM also incorporates the CIS Controls™ Version 7, which is licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License (the link can be found at <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

To further clarify the Creative Commons license related to the CIS Controls and CIS RAM, you are authorized to copy and redistribute the content as a framework for use by you, within your organization and outside of your organization for non-commercial purposes only, provided that (i) appropriate credit is given to CIS, and (ii) a link to the license is provided. Additionally, if you remix, transform or build upon the CIS Controls or CIS RAM, you may not distribute the modified materials. Commercial use of the CIS Controls or CIS RAM is subject to the prior approval of CIS® (Center for Internet Security, Inc.).

Background and Acknowledgements

The original content of CIS RAM was developed by HALOCK Security Labs. It is based on their extensive experience helping clients and legal authorities deal with cybersecurity and due care issues. Recognizing the universal need for a vendor-neutral, open, industry-wide approach to these issues, HALOCK Security Labs approached CIS to make this work openly available to the entire cybersecurity community. This generous contribution of intellectual property (and the extensive work to generalize and tailor it to the CIS Controls) has been donated to CIS and is now available and maintained as a CIS community-supported best practice.

As with all CIS work, we welcome your feedback, and we also welcome volunteers who wish to participate in the evolution of this and other CIS products.

CIS gratefully acknowledges the contributions provided by HALOCK Security Labs and the DoCRA Council in developing CIS RAM and the CIS RAM Workbook.

Significant contributions to Version 1 of CIS RAM were made by:

Principal Author:

Chris Cronin. Partner, HALOCK Security Labs

Contributing Authors:

Jim Mirochnik, Terry Kurzynski, and David Andrew, Partners, HALOCK Security Labs. Erik Leach and Steve Lawn, HALOCK Security Labs. Paul Otto, Attorney, Hogan Lovells US LLP.

Review and vetting was provided by multiple members of the CIS staff.

Table of Contents

Forward	iv
<i>Who this risk assessment method is for</i>	iv
<i>What this document provides</i>	v
<i>The role of professional judgment</i>	v
Author's Introduction	vi
Structure of the Document	vii
Glossary	viii
Risk Assessment Method Examples	x
Chapter 1: Risk Analysis Primer	2
<i>CIS Risk Assessment Method for Due Care</i>	3
<i>Evolving Risk Analysis Methods</i>	7
<i>Overview of the CIS Risk Assessment Method</i>	9
<i>Selecting A Tier for Your Risk Assessment Instructions</i>	12
Chapter 2: Control-Based Risk Assessment Instructions for Tier 1 Organizations	15
<i>The Risk Assessment Project</i>	15
<i>Defining the Scope & Scheduling Sessions</i>	17
<i>Defining Risk Assessment Criteria</i>	21
<i>Defining Risk Acceptance Criteria</i>	25
<i>A Control-Based Risk Assessment Process</i>	27
<i>Risk Treatment Recommendations</i>	38
Chapter 3: Asset-Based Risk Assessment Instructions for Tier 2 Organizations	48
<i>The Risk Assessment Project</i>	48
<i>Defining the Scope & Scheduling Sessions</i>	49
<i>Defining Risk Assessment Criteria</i>	53
<i>Defining Risk Acceptance Criteria</i>	59
<i>An Asset-Based Risk Assessment Process</i>	61
<i>Risk Treatment Recommendations</i>	74
Chapter 4: Threat-Based Risk Assessment Instructions for Tiers 3 and 4 Organizations ..	84
<i>The Risk Assessment Project</i>	84
<i>Defining Risk Assessment Criteria</i>	85
<i>Defining Risk Acceptance Criteria</i>	92
<i>A Threat-Based Risk Assessment Process</i>	94
<i>Risk Treatment Recommendations</i>	110
Chapter 5: Risk Analysis Techniques	116
<i>Risk Analysis Techniques</i>	116
<i>Introduction</i>	116
<i>Defining Impacts for Tier 1 organizations</i>	116
<i>Defining Impacts for Tier 2, Tier 3, and Tier 4 organizations</i>	121
<i>Estimating Likelihood Through "Defense-Readiness" Analysis</i>	127
<i>Using Probability with Duty of Care Risk Analysis</i>	129

<i>Noting How Realized Risk Might be Detected</i>	<i>133</i>
<i>Leveraging Duty of Care Risk Analysis for Maturity Models.....</i>	<i>135</i>
<i>Interview Techniques</i>	<i>136</i>
<i>Evaluating Inherent Risk</i>	<i>139</i>
<i>Root Cause Analysis.....</i>	<i>140</i>
Helpful Resources.....	142
Contact Information	143

Forward

The objective of the Center for Internet Security® Risk Assessment Method (“CIS RAM”) is to help organizations plan and justify their implementation of CIS Controls™ Version 7, whether those controls are fully or partially operating. Few organizations can apply all controls to all information assets, because – while reducing some risks – security controls also introduce new risks to efficiency, collaboration, utility, productivity, or available funds and resources.

Laws, regulations, and information security standards all consider the need to balance security against an organization’s purpose and its objectives, and require risk assessments to find and document that balance. The risk assessment method described here provides a basis for communicating cybersecurity risk among security professionals, business management, legal authorities, and regulators using a common language that is meaningful to all parties.

The CIS RAM conforms to and supplements established information security risk assessment standards, such as ISO/IEC 27005,¹ NIST Special Publication 800-30,² and RISK IT.³ By conforming to these standards, the CIS RAM helps the reader conduct risk assessments according to established standards. By supplementing these standards, the CIS RAM helps its readers evaluate risks and safeguards using the concept of “due care” and “reasonable safeguards” that the legal community and regulators use to determine whether organizations act as a “reasonable person.”

The CIS® designed and prioritized the CIS Controls so that they would prevent or detect the most common causes of cybersecurity events as determined by a community of information security professionals. As a result, CIS Controls V7 has risk considerations at its core.

But because risks vary from one organization to the next, the risk analysis methods described in this document can assist organizations in applying the CIS Controls so that they reasonably and defensibly address the unique risks and resources at each organization.

Who this risk assessment method is for

Cybersecurity risk assessments are important tools for organizations that help them evaluate and prioritize their risks, but also to determine when their risks are acceptable. This risk assessment method is designed to be practical for a broad population of users, whether they are novices to cybersecurity issues, capable of recognizing cybersecurity concerns, or experts.

Organizations that must demonstrate “reasonable” safeguards and risk management for regulatory, contractual, or security management purposes may benefit from the use of the method. Additionally, the CIS RAM is designed to promote meaningful communications and consensus among technicians, non-technical management, security experts, risk managers, as well as legal and regulatory professionals.

¹ ISO/IEC 27005:2011 provided by the International Organization for Standardization.

² NIST Special Publications 800-30 Rev. 1 provided by the National Institute of Standards and Technology.

³ RISK IT Framework provided by ISACA.

What this document provides

The CIS RAM guides readers to conduct risk assessments in a way that match the expectations stated in laws, regulations, and information security standards. The CIS RAM accomplishes this by providing instructions, templates, examples, and exercises to demonstrate its methods. These substantiate the framework of a risk assessment.

The role of professional judgment

Using CIS RAM, the reader will be able to rapidly develop a risk register that communicates reasonableness to many authorities and experts, but the reader will also need to bring their professional judgment (theirs and the judgment of collaborating experts) to the task.

Professional judgment will help organizations determine the scope and boundaries of the risk assessment, to define the organization's mission, objectives, and obligations, to decide which risks will be evaluated, to identify foreseeable threats, and to recommend risk treatment safeguards.

Author's Introduction

The information security community, regulators, attorneys, and managers all understand that perfect cybersecurity is not possible. Even as organizations implement safeguards that are as practical as CIS Controls V7, there are limitations to the degree that organizations can implement security safeguards. Limited security resources (money, experts, and time), competing business priorities, and the ever-changing threat landscape make it difficult for organizations to completely implement a cybersecurity standard equally to all information assets.

Even without these challenges, organizations must operate in somewhat vulnerable environments to fulfill their mission and achieve their objectives. For example, the security value of encryption is obvious, yet information at some point must be unencrypted to serve its purpose. And sometimes information must be unencrypted to enforce other security safeguards, such as data loss prevention. But how does an organization know whether to accept the risk of those moments and transactions when information is unencrypted? And how does it determine whether other supporting safeguards are appropriately protecting the unencrypted information? There is no single answer to that question or to other “grey area” cybersecurity questions that organizations regularly encounter. To assist organizations in their security efforts, laws, regulations, the courts, and information security professionals tell us to use risk assessments *to answer for ourselves* whether we should accept or reduce risks.

Cybersecurity safeguards must be reasonable and appropriate. They must reduce the risk of harm to organizations and to others, but they also must not create too great a burden on the organizations that use those safeguards. The terms “reasonable” and “appropriate” are loaded with many legal, regulatory, expert, and business meanings. But these meanings can be addressed, documented, and justified using a well-constructed risk assessment.

By using the CIS RAM as part of their cybersecurity program, organizations will be more able to adopt CIS Controls V7 in a way that can be successfully demonstrated as reasonable and appropriate to internal management, authorities, security experts, and legal counsel who have an interest in the organization's success.

The CIS RAM document is designed to guide organizations step-by-step through their risk assessment, regardless of their experience in conducting these assessments. We encourage readers to work through each chapter that is suited for their organization, and to follow along with the exercises, worksheets, and examples until their risk register is complete.

Chris Cronin
Partner, HALOCK Security Labs
Chair, DoCRA Council

Structure of the Document

Center for Internet® Security Risk Assessment Method (CIS RAM) is a documented process for conducting risk assessments that address requirements for security, business, regulations, and duty of care requirements. This document will describe the risk assessment method using the following components:

- Instructions are the major portion of the CIS RAM. Instructions provide step-by-step guidance for conducting a risk assessment as a project. Three sets of instructions are provided that address the risk assessment method for organizations based on their risk management maturity. Instructions may be further customized and adapted by each organization according to their needs. Risk assessment techniques are provided at the end of the document to help organizations further develop their risk assessment capabilities.
- Principles state the necessary and fundamental rules for assessing risks according to this method. The principles are the fundamental characteristics of a risk assessment that translates security concerns to regulatory, legal, and business expectations. As organizations customize instructions and templates for their organization, these principles should remain. Risk assessment processes that are developed and conducted without adherence to these principles cannot be considered as “conforming” to the method.
- Examples demonstrate processes and steps. Examples will be accompanied by explanatory scenarios to show the reader how each step is to be conducted. Examples are provided both in this document, and in a separate document, the *CIS_RAM_Workbook* for ease of use.
- Templates model the risk assessment steps, risk analysis methods, and reporting. Templates will assist in rapid adoption of the method's processes by each organization, and will provide for consistent risk assessment practices between organizations. Templates are provided in a separate document, the *CIS_RAM_Workbook* for easy adoption of CIS RAM.
- Exercises encourage the reader to apply what they've learned in the instructions by using the provided templates to design and conduct their own risk assessment.
- Background notes explain why a risk assessment step is taken, or why a principle is applied. Background commentary enables risk practitioners to describe to interested parties how their risk assessment addresses the needs of interested parties and authorities.
- The Glossary provides definitions for specialized terms used in this document. Because risk management methods vary and audiences have variable experience in risk management, the glossary will ensure consistent term usage and meaning.

--

This guide includes references a selection of controls from CIS Controls V7 as examples of safeguards that are specifically selected to help protect organizations. Since such resources change from time to time, please contact CIS or refer to our website for the most recent information. (www.cisecurity.org)

Glossary

Appropriate: A condition in which risks to information assets will not foreseeably create harm that is greater than what the organization or interested parties can tolerate.

Asset Class: A group of information assets that are evaluated as one set based on their similarity. “Servers,” “end-user computers,” “network devices” are examples, as are “email servers,” “web servers” and “authentication servers.”

Attack Path: A series of activities and information assets within the lifecycle of a security incident.

Attack Path Model: A description of how a specific attack path may occur within an environment.

Burden: The negative impact that a safeguard may pose to the organization, or to others.

Business Owners: Personnel who own business processes, goods, or services that information technologies support. i.e. customer service managers, product managers, sales management.

Constituents: Individuals or organizations that may benefit from effective security over information assets, or may be harmed if security fails.

Control: A documented method for protecting information assets using technical, physical, or procedural safeguards.

Control Objective: The intended outcome of a control.

Due Care: The amount of care that a reasonable person would take to prevent foreseeable harm to others.

Duty of Care: The responsibility to ensure that no harm comes to others while conducting activities, offering goods or services, or performing any acts that could foreseeably harm others.

Impact: The harm that may be suffered when a threat compromises an information asset.

Impact Score: The magnitude of impact that can be suffered. This is stated in plain language and is associated with numeric scales, usually from ‘1’ to ‘3’ or ‘1’ to ‘5’.

Impact Type: A category of impact that estimates the amount of harm that may come to a party or a purpose. The CIS RAM describes three impact types; Mission, Objectives, and Obligations.

Information Asset: Information or the systems, processes, people, and facilities that facilitate information handling.

Inherent Risk: The likelihood of an impact occurring when a threat compromises an unprotected asset.

Key Risk Indicator: Aggregations and trending analysis of measures that management may use to understand their risk status.

Likelihood: The degree to which a threat is expected to create an impact. May be stated in terms of frequency, foreseeability, or probability.

Measure: A repeatable, evidence-based indication that a safeguard achieves its control objective.

Observed Risk: The current risk as it appears to the risk assessor.

Probability: The product of statistical analysis that estimates the likelihood of an event.

Reasonable: A condition in which safeguards will not create a burden to the organization that is greater than the risk it is meant to protect against.

Residual Risk: The risk that remains after a safeguard is applied. This concept is not directly used by CIS RAM, but implies that risk is lowered when a safeguard is applied. Residual risk does not take into account potential negative impacts to the organization when safeguards are applied.

Risk: An estimation of the likelihood that a threat will create an undesirable impact. In terms of this method, risk may be expressed as the product of a likelihood and an impact.

Risk Analysis: The process of estimating the likelihood that an event will create an impact. The foreseeability of a threat, the expected effectiveness of safeguards, and an evaluated result are necessary components of risk analysis. Risk analysis may occur during a comprehensive risk assessment, or as part of other activities such as change management, vulnerability assessments, system development and acquisition, and policies exceptions.

Risk Assessment: A comprehensive project that evaluates the potential for harm to occur within a scope of information assets, controls, and threats.

Risk Evaluation: The mathematical component of risk analysis that estimates the likelihood and impact of a risk, and compares it to acceptable risk.

Risk Management: A process for analyzing, mitigating, overseeing, and reducing risk.

Risk Treatment Option: The selection of a method for addressing risks. Organizations may choose to Accept, Reduce, Transfer, or Avoid risks.

Risk Treatment Plan: A comprehensive project plan for implementing risk treatment recommendations.

Risk Treatment Recommendations: A listing of safeguards or processes that may be implemented and operated to reduce the likelihood and/or impact of a risk.

Safeguard: Technologies, processes, and physical protections that prevent or detect threats against information assets. Safeguards are implementations of controls.

Safeguard Risk: The risk posed by recommended safeguards. An organization's mission or objectives may be negatively impacted by a new security control. These impacts must be evaluated to understand their burden on the organization, and to determine whether the burden is reasonable.

Security: An assurance that characteristics of information assets are protected. Confidentiality, Integrity, and Availability are common security characteristics. Other characteristics of information assets such as velocity, authenticity, and reliability may also be considered if these are valuable to the organization and its constituents.

Standard of Care: A set of practices, controls, or requirements that are known to improve outcomes and reduce failures for practitioners of a specialized field or profession.

Steward: Personnel who are responsible for the security and proper operations of information assets, (e.g. database administrator, records manager, or network engineer).

Threat: A potential or foreseeable event that could compromise the security of information assets.

Threat Model: A description of how a threat could compromise an information asset, given the current safeguards and vulnerabilities around the asset.

Vulnerability: A weakness that could permit a threat to compromise the security of information assets.

Risk Assessment Method Examples

CIS RAM provides three sets of instructions that each describe a full risk assessment project. Each set of instructions is designed for organizations of varying information security management capabilities to increase the method's usefulness.

All three sets of instructions present a fictional organization that is conducting an information security risk assessment, and that improves its risk management capabilities over time. The example organization begins the risk assessment in the first set of instructions as a security novice with little involvement by business management. After a year of improving their security posture and abilities, they assess risk in the second set of instructions using more refined reasoning and methods, and in collaboration with business management. Finally, they mature enough as a capable organization to take on complex risk analysis in the third set of instructions.

The example organization described in this document manufactures and services medical devices ("diary devices") that read biological information from patients that wear the devices. The organization works in clinical environments to support the patients as well as the devices, and as a result carries private health information about the patients. Because they work with military and veterans' organizations, many of their patients are active or former members of the armed forces. As a result, the organization poses heightened risk and requires heightened scrutiny over their cybersecurity controls.

The example organization is hypothetical and is not based on a known organization, technology, or service. But the risks they encounter are commonly seen and managed by many types of organizations. Example materials related to the example organization are provided in the document *CIS_RAM_Workbook* in re-usable templates.

The reader will best develop an understanding of the risk assessment method by following along with the workbook, and by entering their own examples in the spaces provided within each sample worksheet.

Center for Internet Security®

Risk Assessment Method

CIS RAM Version 1.0

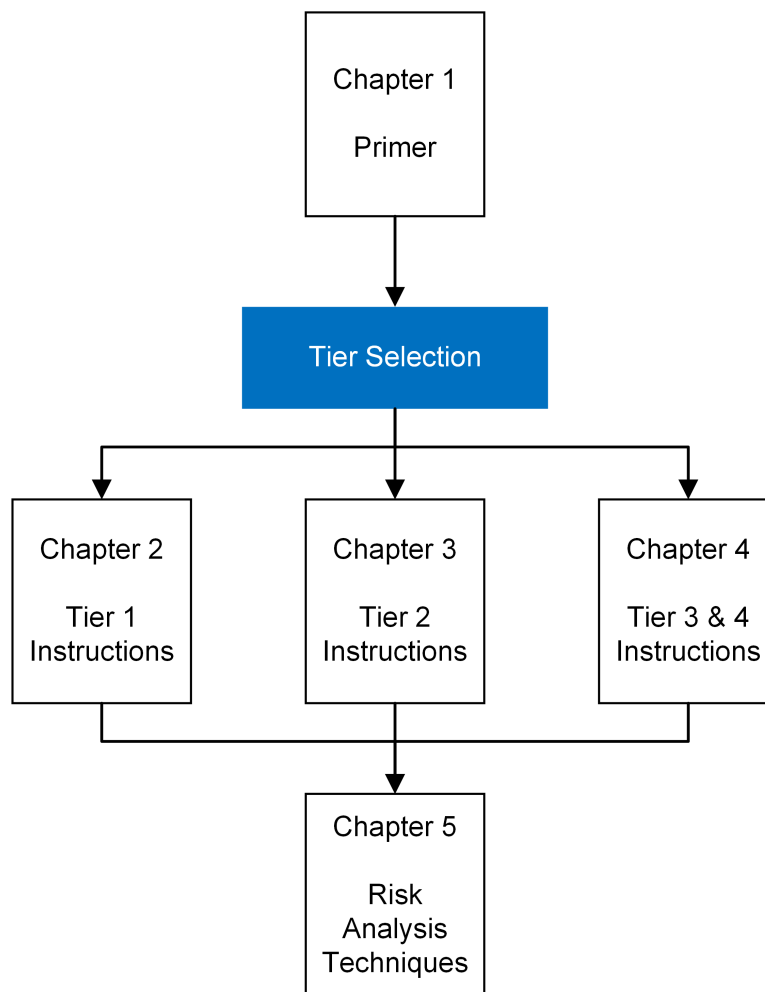
April 2018

Chapter 1: Risk Analysis Primer

CIA RAM describes a method for cybersecurity risk analysis that includes methods that are new to most readers. Chapter 1 will provide an explanation and description of new concepts, language, and processes to provide the reader with a solid foundation for the remaining chapters.

After completing Chapter 1, the reader will be directed to one of three chapters that provide instructions for conducting risk assessments. Chapters 2, 3, and 4 present processes, materials, and examples that are suitable for organizations with varying degrees of capability for conducting risk assessments.

After completing the instructions chapters, all readers will benefit from the guidance, tips, and deep dives presented in Chapter 5.



CIS Risk Assessment Method for Due Care

Introduction

Laws, regulations, and information security standards do not expect that the public can or will prevent all information security incidents. They instead make us responsible for looking ahead to what might go wrong, and to use safeguards that are not overly burdensome to prevent that harm. That is the essence of Duty of Care Risk Analysis⁴ (“DoCRA”) that the CIS RAM is based on.

- Since 1993, all US regulations – whether or not they are related to information security – require risk analysis to achieve a cost-benefit balance while achieving compliance.⁵
- Information security standards have called on the public to use risk analysis when designing security controls that match their environment.⁶
- Judges have used a “duty of care balance test” to determine liability in data breach cases.⁷
- The Federal Trade Commission has consistently required that organizations use risk assessments to determine the reasonableness of their security controls.⁸
- The General Data Protection Regulation (“GDPR”) that requires privacy protections for EU residents, and bases its security requirements on risk analysis.⁹

Experts and authorities consistently require organizations to secure information and systems as much as they can to prevent harm to others, but not to allow safeguards to be overly burdensome to them or the public. And they point to risk assessments as the way to find the balance.

⁴ Also known the DoCRA Standard. <https://www.docra.org>.

⁵ Executive Order 12866” signed in 1993 requires all federal regulation to be enforced using cost-benefit analysis. The Office of Management and Budget enforces the order in part by requiring that regulated organizations use risk assessments to identify effective controls that are “reasonable.”

⁶ See ISO/IEC 27001:2013, NIST SP 800-53 Rev. 4, PCI-DSS v3.2

⁷ See *Dittman v. UPMC*, 154 A.3d 318 (Pa. Super. Ct. 2017), *In re: Target Corporation Customer Data Security Breach Litigation*, Memorandum and Order, MDL No. 14-2522 (D. Minn. 2014)

⁸ Federal Trade Commission. “Commission Statement Marking the FTC’s 50th Data Security Settlement.” www.ftc.gov/system/files/documents/cases/140131gmrstatement.pdf

⁹ General Data Protect Regulation. Directive 95/46/EC.

Basis in Law and Regulation

CIS RAM's risk analysis method was designed to provide common ground for security specialists and business managers, and for legal and regulatory authorities who must evaluate the sufficiency of security safeguards.

Risk analysis became the basis of regulatory law in 1993 when President Bill Clinton signed an executive order, E.O. 12866,¹⁰ that required regulations to be enforced using a cost-benefit analysis. The Office of Management and Budget determined that the best way to achieve cost-benefit analysis was to embed risk analysis in all existing and new regulations.

Starting in 1999 the United States began enforcing regulations with information security and privacy requirements that used risk analysis as the basis for compliance. The Gramm-Leach-Bliley Act Safeguards Rule,¹¹ the HIPAA Security Rule,¹² and the Federal Trade Commission all required that organizations conduct risk assessments to define their own compliance goals. Risk assessments should help organizations determine for themselves the likelihood and impact of threats that could harm the public, and ensure that safeguards would not be overly burdensome. This risk analysis has been commonly described as “Risk = Impact x Likelihood.”

At about this time this same idea emerged independently among a separate set of professionals. Domestic and international information security standards bodies developed risk assessment methods such as NIST Special Publications 800-30, ISO 27005, and RISK IT to help the public assess risks in information technology environments. These information security standards also used the same equation used by U.S. regulators – “Risk = Impact x Likelihood” – to express the foreseeability of harm that might come to information and information systems.

In parallel with the development of these two efforts (and since earlier in the twentieth century) attorneys and judges debated in court rooms and in law journals about how to determine whether someone acted as a “reasonable person” when a plaintiff sued for damages. These debates led to the creation of the “Learned Hand Rule”¹³ (aka “the Calculus of Negligence”). The Hand Rule, as it is now known, states that a **burden** to prevent harm should not be greater than the **probability** of harm times the **liability** after a harmful event; or mathematically stated, $B \leq P \times L$. Courts (variably) have extended this rule to “duty of care balancing tests” that determine whether lack of foresight and less-than-reasonable safeguards led to harm.

But while these disciplines – law, information security, and regulations – all drew on a common definition of risk, each seemed to be unaware of the other's risk analysis methods. Even so, each discipline searched for a universal translator that would allow the entire community of experts and authorities to understand one another.

The CIS RAM provides that universal translator.

¹⁰ *Executive Order 12866 – Regulatory Planning and Review*, 58 FR 51735; October 4, 1993

¹¹ Gramm Leach Bliley Act Safeguards Rule 16 CFR Part 314

¹² HIPAA Security Rule 45 CFR Part 160 and Subparts A and C of Part 164

¹³ *U.S. v. Carroll Towing*, 159 F.2d 169 (2d Cir. 1947)

CIS RAM Principles and Practices

CIS RAM adopts the three principles and ten practices from Duty of Care Risk Analysis. The three principles state the characteristics of risk assessments that align to regulatory and legal expectations. The ten practices describe features of risk assessments that make the three principles achievable.

Principles

1. Risk analysis must consider the interests of all parties that may be harmed by the risk.
2. Risks must be reduced to a level that authorities and potentially affected parties would find appropriate.
3. Safeguards must not be more burdensome than the risks they protect against.

Practices

1. Risk analysis considers the likelihood that certain threats could create magnitudes of impact.
2. Risks and safeguards are evaluated using the same criteria so they can be compared.
3. Impact and likelihood scores have a qualitative component that concisely states the concerns of interested parties, authorities, and the assessing organization.
4. Impact and likelihood scores are derived by a numeric calculation that permits comparability among all evaluated risks, safeguards, and against risk acceptance criteria.
5. Impact definitions ensure that the magnitude of harm to one party is equated with the magnitude of harm to others.
6. Impact definitions should have an explicit boundary between those magnitudes that would be acceptable to all parties and those that would not be.
7. Impact definitions address; the organization's mission or utility to explain why the organization and others engage risk, the organization's self-interested objectives, and the organization's obligations to protect others from harm.
8. Risk analysis relies on a standard of care to analyze current controls and recommended safeguards.
9. Risk is analyzed by subject matter experts who use evidence to evaluate risks and safeguards.
10. Risk assessments cannot evaluate all foreseeable risks. Risk assessments re-occur to identify and address more risks over time.

Table 1 aligns these principles and practices with the three disciplines of law, regulations, and information security standards.

Table 1 - CIS RAM Principles and Practices Alignment to Law, Regulations, and Security Standards

CIS RAM and DoCRA Principles and Practices	Law	Regulations	Security Standards
Risk analysis must consider the interests of all parties that may be harmed by the risk.	●	●	◐
Risks must be reduced to a level that authorities and potentially affected parties would find appropriate.	●	●	◐
Safeguards must not be more burdensome than the risks they protect against.	●	●	◐
Risk analysis considers the likelihood that certain threats could create magnitudes of impact.	●	●	●
Risks and safeguards are evaluated using the same criteria so they can be compared.	○	○	◐
Impact and likelihood scores have a qualitative component that concisely states the concerns of interested parties, authorities, and the assessing organization.	◐	◐	◐
Impact and likelihood scores are derived by a numeric calculation that permits comparability among all evaluated risks, safeguards, and against risk acceptance criteria.	◐	◐	◐
Impact definitions ensure that the magnitude of harm to one party is equated with the magnitude of harm to others.	◐	◐	○
Impact definitions should have an explicit boundary between those magnitudes that would be acceptable to all parties and those that would not be.	◐	○	○
Impact definitions address; the organization's mission or utility to explain why the organization and others engage risk, the organization's self-interested objectives, and the organization's obligations to protect others from harm.	◐	◐	◐
Risk analysis relies on a standard of care to analyze current controls and recommended safeguards.	◐	●	●
Risk is analyzed by subject matter experts who use evidence to evaluate risks and safeguards.	◐	◐	◐
Risk assessments cannot evaluate all foreseeable risks. Risk assessments re-occur to identify and address more risks over time.	○	◐	●

Key: Fully addressed ● Partially addressed ◐ Not addressed ○

Organizations that conduct risk assessments using the CIS RAM will have a plan for implementing CIS Controls V7 that is reasonable, and defensible to authorities and experts alike.

Evolving Risk Analysis Methods

Evolving Classic Risk Concepts

To bridge information security risk analysis with legal and regulatory expectations, CIS RAM builds on and extends a few classic risk analysis concepts. This section will briefly describe how CIS RAM evolves risk evaluation, and definitions for “impact,” risk acceptance, and residual risk.

Calculating Risk Includes Multiple Impacts

CIS RAM uses the classic risk assessment calculation “Risk = Impact x Likelihood” with a few modifications. Most significantly, risk is calculated by multiplying a likelihood value by multiple impact values. These multiple impacts include impacts to the organization’s objectives, its mission, and its obligations to protect others. Organizations should be aware of the many ways that information security risk can create harm.

The risk calculation used by CIS RAM resembles the structure below:

“Risk = Max (Mission Impact, Objectives Impact, Obligations Impact) x Likelihood.”

The instructions provided later in this document clearly describe how this calculation works. Organizations that use this extended calculation will consistently consider the many ways that information security risks can create harm.

Impact Definitions Include Harm to Multiple Parties

To ensure fairness and balance, impact definitions will include potential harm to individuals and organizations that may be impacted by risks. Impacts and impact magnitudes will be stated in qualitative and quantitative form to easily communicate levels of risk to all interested parties, and in a way that matters to each party.

Risk Acceptance Is Clearly Defined

CIS RAM provides organizations with clear guidance for defining acceptable risk that appears fair to authorities and interested parties, and that can be consistently applied to all information security risks.

Acceptable risk will consider whether a observed risk is “appropriate” (all potentially affected parties would agree that the risk is acceptable), and whether a recommended safeguard is “reasonable” (it does not create more of a burden than the risk it protects against).

By expanding the definition of risk acceptance by these two factors, organizations will have an easily communicated rationale for accepting risk, or for prioritizing unacceptable risk.

“Residual Risk” is Known As “Safeguard Risk”

“Residual risk” has traditionally meant the reduced amount of risk that remains after a security control has been implemented. Organizations have generally used “residual” to declare how a planned security safeguard presents acceptable risk. CIS RAM evolves the notion of a “residual risk” to “safeguard risk” to describe the risk that a new safeguard may pose.

The purpose behind evaluating residual risk this way is to address the fact that new controls often have unintended consequences. Recall that impact definitions will be based on multiple factors, such as an organization’s mission, its objectives, and its obligations (described in more detail later in the document). Security controls may reduce the risk to security obligations by controlling access to data, but may increase the risk to the organization’s mission which requires sharing the data. Legal decisions and regulations consider these excessive safeguards as “burdens” because they may harm the organization that is trying to protect the data.

By evaluating safeguard risk using the same criteria that are used to evaluate risks, organizations will be more cognizant of the true cost of controls, and will have a defensible way of stating whether recommended controls are overly burdensome to them or the public.


Evolving Risk Acceptability

Figure 1 illustrates how CIS RAM evaluates “appropriate” risk using a simplified risk statement. In this scenario, an organization is analyzing a risk of a lost device, and estimates the likelihood and expected impact of the loss. Impact definitions estimate potential harm to the organization, and to others.

Using a scale of ‘1’ to ‘3’, the organization multiplies the likelihood score by the higher of the two impact scores to arrive at a risk score of ‘6’. In this example, an acceptable risk would be less than ‘4’, so the score of ‘6’ is not appropriate. “Others” would not accept the possibility of this risk.

Note: The CIS RAM provides extensive guidance on how likelihood and impact scoring and acceptable risk criteria are defined. The values in Figure 1 and Figure 2 are provided simply for illustrative purposes.

Figure 1 - Simplified Risk Model Showing Inappropriate Risk



In Figure 2 the inappropriately high risk is matched with a recommended safeguard to encrypt all devices. Because a safeguard is evaluated using the same criteria as the risk, the organization is evaluating the burden of the safeguard. In this case, they believe there is a small likelihood (‘1’) of a notable cost impact (‘2’). As a result, the safeguard risk calculates as ‘2’ which is lower than the observed risk it is addressing (‘6’). As a result, this safeguard is reasonable.

Figure 2 – Simplified Risk Model Showing a Reasonable Safeguard



Is This Extended Analysis Worthwhile?

Put simply, yes.

Information security controls are very often considered to be a hindrance to business. Users often complain that security controls get in the way of productivity, efficiency, ease of collaboration and communication, and other business-impacting concerns. Organizations should take these complaints seriously. Fortunately, regulators have provided organizations with a means to evaluate these concerns. Moreover, courts consider the burden of safeguards in lawsuits and would understand the reasoning that this risk analysis provides.

By evaluating risks and their recommended safeguards using the same criteria, organizations ensure that risk analysis addresses the concerns of all parties within and outside of their organization, and provides evidence of their conscientious decision to regulators and judges.

Overview of the CIS Risk Assessment Method

Using Risk Assessments to Design and Evaluate CIS Controls V7

CIS Controls V7 was designed to address the most common causes of security incidents in the general public. As a result, the CIS Controls are to a degree risk-prioritized, especially if organizations implement the first five CIS Controls before implementing the remaining 15. However, each organization has special circumstances, including the potential harm they may cause to others, the need to operate somewhat vulnerable systems based on their mission, the needs of their constituents, their available resources, and the foreseeability of threats in their industry.

The risk assessment method described by CIS RAM will help organizations determine whether their implementation of CIS Controls V7 – or their de-prioritization or customization of controls – is reasonable and appropriate given security, legal, and regulatory considerations.

This risk assessment method describes multiple ways that organizations may evaluate, assess, and design safeguards using the CIS Controls.

- In some cases, organizations may start simply and list the Controls to determine whether their information assets are sufficiently resilient against foreseeable threats.
- More capable organizations may list their information assets first, then consider whether associated CIS Controls sufficiently protect those assets against foreseeable threats.
- Organizations with a command of how threats operate may start with a list of known or foreseeable threats against information assets and determine how controls should be implemented to address them.

Each of these approaches relies on the organization's ability to conduct that kind of analysis. And those abilities depend on the involvement of business management in information security, the availability of time and resources to examine information assets and risks, and the expertise of the personnel for conducting the analysis.

Regardless, this risk assessment method will provide a model for organizations to evaluate risk based on the harm they may pose to themselves or their constituency, and to determine whether the burden of each of the CIS Controls – implemented as safeguards – are appropriate.

Risk Assessment Process

A risk assessment is a project that analyzes the risk posed by a set of information assets, and recommends safeguards to address unacceptably high risks.

While the order of events in a risk assessment project will vary from organization to organization, the following activities are generally applied:

Analyze the Observed risk

- Define the Scope: Identify information assets that are being assessed as well as the owners and stewards of the information assets.
- Schedule Sessions: Schedule the interviews and sessions for evidence review.
- Develop the Risk Assessment and Acceptance Criteria: Establish and define the criteria for evaluating and accepting risk.
- Gather Evidence: Interview personnel, review documents, and observe safeguards.
- Model the Risks: Evaluate the current safeguards that would prevent or detect foreseeable threats against the security of information assets.
- Risk Evaluation: Estimate the likelihood and impact of security breaches to calculate the risk score, then determine whether identified risks are acceptable.

Propose Safeguards

- Propose Safeguards: Recommend safeguards from CIS Controls V7 that would reduce unacceptable risks.
- Evaluate Proposed Safeguards: Risk-analyze the recommended safeguards to ensure that they pose acceptably low risks without creating an undue burden.

Risk Assessment Criteria

Risk analysis requires a consistent, repeatable method for estimating and evaluating risk. Risk assessment criteria provide organizations with measures for consistently rating the likelihood and impact of foreseeable threats that may compromise the security of information assets.

Risk assessment criteria are often thought of in terms of a 3 x 3 grid or a 5 x 5 grid, with each dimension representing either “likelihood” values or “impact” values. While scores of ‘1’ through ‘3’ or ‘1’ through ‘5’ are convenient for calculating risk as a product, they are not meaningful by themselves. So criteria must also have a plain-language component that describes levels of impact and likelihood that are meaningful to the organization.

Risk assessment criteria in a simplified format may appear similar to this:

Table 2 - Simplified Impact Criteria

Impact Score	Impact Score Defined
1	No or minimal harm would result.
2	Harm would not be tolerable.
3	Harm may not be recoverable.

Table 3 - Simplified Likelihood Criteria

Likelihood Score	Likelihood Score Defined
1	Not foreseeable.
2	Expected to occur.
3	Regular occurrence.

Risk Acceptance Criteria

Laws and regulations require that organizations apply “reasonable” and “appropriate” safeguards to ensure that the resulting risk is acceptable. The acceptability of risk can be demonstrated using risk analysis that addresses the tolerability of the risk and the burden of safeguards that protect against the risk.

While every organization will define its own risk tolerance, this method provides a process for doing so using plain language and simple math. An example of defining risk acceptability is provided in Table 4 using the simplified impact and likelihood criteria from above. Organizations will develop their risk acceptance criteria by first defining what unacceptable risk is.

In this case the organization has determined its risk acceptance criteria by first deciding that it will not accept a risk that may cause intolerable harm (as indicated by the red box).

Table 4 - Simplified Impact Criteria for Risk Acceptance

Impact Score	Impact Level	Impact Score Defined
1	Acceptable	No or minimal harm would result.
2	Unacceptable	Harm would not be tolerable.
3	High	Harm may not be recoverable.

Then the organization determined that a threat that is expected to occur (and to create harm) must be avoided (as indicated in the red box in Table 5).

Table 5 - Simplified Likelihood Criteria for Risk Acceptance

Likelihood Score	Likelihood Score Defined
1	Not expected to occur
2	Expected to occur
3	Regular occurrence

And finally, the organization combined these limits to express their acceptable risk in both plain language, and in mathematical terms.

Table 6 – Risk Acceptance Criteria

Version	Definitions of Acceptable Risk
Plain language	We must reduce risks that are expected to create intolerable harm.
Mathematical	Acceptable Risk < 2 x 2; or Acceptable Risk < 4

Background – “Reasonableness” and Risk Analysis

The “reasonable person” is used in law as a hypothetical person – or legal fiction – who embodies the sum of our traditions, values, and responsibilities for taking care not to harm others while we engage in public life. The reasonable person has been used in cases to evaluate appropriate behavior for activities such as building and maintaining structures, offering goods and services, or handling assets such as information and information technologies. A reasonable person can engage in activities for their own benefit, but must take care, using appropriate precautions, not to harm others in the process.

In litigation, a judge will often use a “duty of care” or “multi-factor” balancing test to determine the degree to which a defendant was acting reasonably when a plaintiff was harmed. And in regulations organizations must apply “reasonable” safeguards to protect others from harm.

A judge’s duty of care balancing test is very similar in structure to this risk assessment method. An organization will consider foreseeable threats that their business may cause others. They will determine how effectively they prevent that harm by using CIS Controls V7 as a standard for appropriate cybersecurity practices. They will estimate the likelihood and impact of the expected harm of a foreseeable threat, and they will consider alternative safeguards that effectively lower risks without being overly burdensome. In this way, judges and cybersecurity practitioners use the same language to describe reasonable cybersecurity practices.

In similar fashion, a regulator will ask regulated organizations to demonstrate the reasonableness of their safeguards by reviewing the organization’s risk register. Since 1993 US federal regulations require that regulatory rules are not overly burdensome to the public, and that a “cost-benefit” analysis is performed to determine whether regulatory actions are overly burdensome and appropriate to protect the public. Regulatory agencies, including those that govern cybersecurity rules and regulations, require risk assessments as the method for balancing the potential harm to others against the cost of safeguards.

Selecting A Tier for Your Risk Assessment Instructions

This document is designed to be useful for organizations with varying levels of security management capabilities. These capability levels align with Framework Implementation Tiers (“Tiers”) as defined by the NIST Cybersecurity Framework.¹⁴ The Tiers indicate “how an organization views cybersecurity risk and the processes in place to manage that risk.”¹⁵ The Tiers are defined by NIST in the following way (abbreviated).

Tier 1: Partial

- *Risk Management Process* – Informal and ad hoc.
- *Integrated Risk Management Program* – Limited awareness within the organization.
- *External Participation* – Not coordinating with external entities.

¹⁴ *Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0*, National Institute of Standards and Technology. February 12, 2014

¹⁵ *Ibid*, pg. 9.

Tier 2: Risk Informed

- *Risk Management Process* – Informed by organization risk objectives.
- *Integrated Risk Management Program* – Risk-informed, management-approved processes and procedures.
- *External Participation* – Not coordinating with external entities.

Tier 3: Repeatable

- *Risk Management Process* – Enforced through policy, and updated with changes in the environment and threats.
- *Integrated Risk Management Program* – Risk-informed policies and processes are used enterprise-wide. Personnel are skilled and informed to work securely.
- *External Participation* – Receiving information from partners to make internal risk-based decisions.

Tier 4: Adaptive

- *Risk Management Process* – Adaptive through lessons learned and continuous improvement.
- *Integrated Risk Management Program* – Enterprise-wide culture of security awareness and continuous improvement based on lessons learned and external information.
- *External Participation* – Sharing security and threat information with partners.

CIS RAM provides three sets of instructions, templates, exercises, and examples for conducting risk assessments, each with increasing complexity. These three sets of materials are suitable for Tier 1 organizations, Tier 2 organizations, and both Tier 3 and Tier 4 organizations. The reader can determine which of these levels and documentation are best suited to them by reviewing the characteristics provided below.

Tier 1 materials are well-suited to organizations with the following characteristics:

- **NIST Tier:** Tier 1 organizations. Tier 1 materials are best suited for organizations that do not coordinate their information security plans and requirements throughout the organization. Information security is largely driven by technology management.
- **Expertise:** The organization is able to identify generic threats, but not specific methods for hacking systems, devices, and applications.
- **Time:** The organization can absorb the time needed to evaluate information risks at the level of generic systems, devices, and applications.

Tier 2 materials are well-suited to organizations that enjoy more collaboration with business management, and have more resources and capabilities for analyzing security risks and planning programs.

- **NIST Tier:** Tier 2 organizations. Tier 2 materials are best suited for organizations that have at least some collaboration with non-technical business management to define risk criteria.
- **Expertise:** The organization has resources and capabilities to analyze common security threats, and to plan risk-appropriate safeguards. However, they do not have on-hand skills to model how threats would operate within their organization.

- **Time:** The organization is able to invest sufficient time to analyze risks at the level of specific systems, devices, and applications, and sub-components within those assets.

Tier 3 or 4 materials are well-suited to organizations that receive security and threat information from outside sources, that have significant knowledge of information security topics, and time to evaluate threat scenarios that risk assessments are based on.

- **NIST Tier:** Tiers 3 and 4 organizations. Tiers 3 or 4 materials are best suited for organizations that are using risk-based criteria for enterprise-wide policies and processes.
- **Expertise:** The organization has resources and capabilities to analyze security threats, and to plan risk-appropriate safeguards, including the on-hand skills to model how threats would operate within their organization.
- **Time:** The organization is able to invest time to analyze risks at the level of specific systems, devices, and applications within the context of specific threats.

The reader should determine which level materials are best suited to their organization, and should follow the instructions that are provided in that level. They may also decide to learn and use the methods described in other instruction sets, but should stay within their level as much as possible until they are comfortable with their existing risk assessment processes.

Chapter 2: Control-Based Risk Assessment Instructions for Tier 1 Organizations

Tier 1 risk assessment instructions are well-suited to organizations that fit the profile of Tier 1 organizations as described by the NIST Cybersecurity Framework. These organizations can be identified as having the following characteristics:

- **NIST Tier:** Tier 1 organizations. Tier 1 materials are best suited for organizations that do not coordinate their information security plans and requirements throughout the organization. Information security is largely driven by technology management.
- **Expertise:** The organization is able to identify generic threats, but not specific methods for hacking systems, devices, and applications.
- **Time:** The organization can absorb the time needed to evaluate information risks at the level of generic systems, devices, and applications.

This chapter is comprised of sections that each address a specific activity within a risk assessment. Readers should engage this chapter by first reading the text in each section, and then conducting the exercises that are recommended for each section. The material presented in the CIS RAM is substantially different from many other risk assessment standards and models, so the reader should first understand the aim of each section, and then practice what they learn using templates that are provided in the supplementary document *CIS_RAM_Workbook*.

While conducting their first CIS RAM-based risk assessment, organizations should be careful to not try to “boil the ocean.” Regulatory bodies and information security standards alike understand that not all risks can be identified in a single assessment. Organizations should continuously and regularly assess risks to identify, understand, and manage risks over time.

The Risk Assessment Project

Overview

Risk assessments are projects with clear steps for preparing, conducting, and reporting risk analysis. And while risk assessment projects can be modeled with a typical plan, each organization’s project approach will vary depending on factors such as resource availability, and will develop over time as organizations become more capable in their cybersecurity maturity. This section will describe a risk assessment project, its components and variations, and will present guidance for preparing the plan.

The Project Outline

Risk assessments are conducted using a series of steps that include typical actions, and roles as illustrated in Table 7.

Table 7 – Example Risk Assessment Project Outline

Step	Task	Key Roles
1	Defining the Scope & Scheduling Sessions	Executives, Management, Assessor
2	Defining Risk Assessment Criteria	Management, Assessor
3	Defining Risk Acceptance Criteria	Executives, Management, Assessor

Step	Task	Key Roles
4	Risk Assessment (Control-Based)	
4.1	Gather Evidence	Personnel, Management, Assessor
4.2	Model the Threats	Personnel, Management, Assessor
4.3	Risk Evaluation	Assessor
5	Propose Safeguards	
5.1	Evaluate Proposed Safeguards	Assessor, Management

During **Step 1** (Defining the Scope & Scheduling Sessions), the organization will determine which information assets to include in their evaluation. They will also identify business owners and technical stewards who will provide evidence and interviews to assess those assets. The risk assessor will then schedule interview sessions with those owners and stewards.

In **Step 2** (Defining Risk Assessment Criteria) the organization will define the rules by which they assess and score risks. They will define their mission (the value they bring to others), and their obligations (the potential for harm against others) to establish what they are trying to protect. They will then define scoring schemas to be used for impact and likelihood estimation.

In **Step 3** (Defining Risk Acceptance Criteria) the organization will establish their risk tolerance by selecting a combination the likelihood of an impact that would be tolerable to all parties (the organization and parties that may be harmed by realized risks).

In **Step 4** (Risk Assessment – Control-Based) the risk assessor will evaluate the risks of the information assets. For Tier 1 organizations, the analysis includes the following activities:

- “Gather Evidence” involves a review of documents, such as policies, procedures, standards, and benchmarks. It also includes interviews with management and personnel. Evidence gathering also entails observation of configurations, artifacts, facilities, records, and work processes to determine whether they operate in secure or vulnerable ways.

Tier 1 organizations should also consider reviewing the configurations of controls and looking for evidence of their effectiveness. This may be challenging for organizations at this level. Vulnerability scanners and configuration scanners using SCAP policies may provide efficient analysis of technical systems to assist in this analysis.

- “Model the Threats” entails the most variety in approaches that depend on the cybersecurity maturity of the organization. Each organization, however, will model risks with at least these components: considering the CIS Controls that should be in place to protect information assets; determining whether those safeguards are effectively in place to protect information assets; identifying vulnerabilities that may allow breaches of the assets; and identifying threats that could take advantage of those vulnerabilities.
- During “Risk Evaluation” the organization will estimate the likelihood and impact of the risks. The estimates will be based on the scoring and criteria that were established in Step 2. The risk score will be automatically calculated to determine whether the current implementations of CIS Controls are already reasonable.

During **Step 5** (Propose Safeguards) the organization will consider how to address unreasonable risks by selecting CIS Controls that should be implemented to address each risk, and specifically

how the controls will be implemented. These safeguards may include security devices, physical safeguards, training, oversight processes, or other methods. The risk assessor then will test the reasonableness of the safeguards during “Evaluate Proposed Safeguards.” The risk assessor will evaluate the proposed safeguards using the same criteria that were used to evaluate the risks.

A project plan template is available in the supplementary document *CIS_RAM_Workbook*.

Defining the Scope & Scheduling Sessions

Note: To best understand the content of this chapter, the reader should first read each section of the chapter, then go through the recommended exercises at the end of each section to gain practical knowledge of the section’s topics. The reader will be directed to use the templates that are provided in the supplementary document CIS_RAM_Workbook to attempt their exercises.

Defining the Scope

Organizations should conduct risk assessments against a clearly defined scope of information assets. A single theme usually bounds the scope of assets, such as “information assets that contain sensitive information,” “the data center,” “engineering practice areas and technologies that support them,” or a specific business division.

While it is possible to select unrelated information assets for an assessment – or a subset of assets within a larger scope – the organization that receives the assessment and is making investments and prioritization decisions based on its findings will be most comfortable when the information assets are associated with a business entity or a business process. Otherwise, risk assessment findings may seem scattered and unrelated.

Similarly, while assessing the risk of a set of information assets, it makes good sense to consider a set of assets that can directly affect each other’s security. For example, a risk assessment that examines a set of applications should also include the network devices that connect the applications to other assets and other networks, as well as the processes that are used to develop and manage those applications. These systems are directly connected to each other and dependent on each other so their risks are easily associated with one another.

Organizations cannot examine all information assets comprehensively in a single risk assessment, so their scope should consider the time and resources that are available for the assessment. Risk assessors should consult security experts to help them determine which assets to prioritize, and may use inherent risk analysis as described in Chapter 5 to assist in this prioritization.

An example scoping table (Table 8) demonstrates the level of detail that may be appropriate for an initial assessment plan and is provided in the workbook *CIS_RAM_Workbook*

Table 8 - Example Scoping Table

Asset Type	Asset Class	Business Owner	Steward
Information	IP and PII	COO	CIO
Application	Applications	Customer Experience	Prod Mgr, Dev & Dev Ops
Servers	Servers	Dev Ops	DevOps
Network Device	Network Devices	CIO	Network Engineering
Process	Dev, Promotion, Maint.	Dev Ops	Dev, DevOps
Process	Vulnerability Mgt.	CIO	Security Team
Process	Acct Setup, Maint.	Customer Experience	Application Management

Asset Type	Asset Class	Business Owner	Steward
Process	Internal Audit	Compliance	Internal Audit
Process	Device/System set-up	CIO	DevOps
Process	Customer Support	Customer Experience	Application Management

Note that the scoping table includes business owner roles and steward roles. Business owners are the (typically) non-technical managers who are responsible for the information and processes that information assets support. Stewards are (typically) technical managers who are responsible for the functionality and security of the information assets. By identifying information asset ownership up front, the scoping table can be used to help plan interview sessions for the remainder of the risk assessment.

Regardless of how the scope of the risk assessment is established, and how detailed the asset listing is, there are a few helpful practices an organization should keep in mind as they identify their information assets:

- Think of a set of similarly-situated assets as a single asset class. For example, all database servers that use the same technology and the same maintenance and administration methods may be considered one asset class. However, if one set of database servers is different from others (for example, they hold sensitive information in a DMZ while others process less-sensitive information in another zone), these may be considered two assets because their inherent risks will be different, even if they are managed identically.
- Information assets are not only technologies that store and transmit sensitive information. Information assets are any information, technology, process, people, or facilities that may impact the confidentiality, integrity, or availability of information.
- Include in the scope all information assets that are within the same zones (networks, facilities, etc.) as any other in-scope information asset.

Exercise:

The reader should develop their own scoping table using the “Scope - Tier 1” worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. A set of information assets your organization is interested in focusing your security resources on?
 - a. This set can be defined by processes, technologies, a class of information, or a location.
2. What are the boundaries between this set of information assets and other information assets not in this scope?
 - a. What systems, facilities, and network devices link these boundaries together, or separate them?
 - b. Are these “boundary” assets included or excluded from the scope?
3. List information assets or asset classes that are within the scope.
 - a. List information assets or asset classes to a level of detail that the organization has the time and resources to analyze. This may require adjustment during the course of the assessment if the organization realizes it has more time (or less time) than they originally planned to assess the information assets.

Scheduling Interview Sessions

Interview sessions will be topical, and should address one topic or closely related topics for each conversation. Interview sessions can focus on CIS Controls, or information assets and asset classes.

For example, interview sessions that focus on CIS Controls would bring together the personnel and management who know how each control is implemented and operates. One session may be dedicated to CIS Control 1 to understand how devices are inventoried. Another may be scheduled to discuss CIS Control 2 to understand what safeguards are in place to inventory software. Or, if the same personnel are knowledgeable about both safeguards, then perhaps one session could combine both topics.

Similarly, if risk assessors schedule sessions around information assets or asset classes, then it would be appropriate to include business owners and technical owners of those systems to understand how associated CIS Controls are applied to each asset or asset class.

An example session for a web application may include product managers, application developers, application administrators, and business owners. Topics in that session may include CIS Control 14, “Controlled Access Based on the Need to Know,” CIS Control 16, “Account Monitoring and Control,” and CIS Control 18, “Application Software Security.”

How the organization groups and orders these topics is largely up to them, but risk assessors should consider these pointers while scheduling interview sessions:

1. Be respectful of people’s time. While it is important to gather comprehensive information about risks, organizations cannot gather all relevant information in the first one or two risk assessments.
2. Work with managers to determine the most efficient and useful way to schedule interviews, whether by CIS Controls or by information assets and asset classes.
3. Expect that some security safeguards will be applied differently to different information assets and asset classes. For example, CIS Control 5, “Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers,” and CIS Control 16, “Account Monitoring and Control” may be implemented and overseen differently for servers in different environments, or may be centrally controlled for servers, but individually controlled for network devices. Plan on evaluating how these and similar safeguards are applied to different asset classes.
4. Provide an agenda for each interview to help participants prepare any materials or information regarding the information assets or controls that may be discussed.

Scheduling Evidence Reviews

Risk assessors use evidence review sessions to examine information assets and; determine whether they conform to CIS Controls, and evaluate whether they would be effective against foreseeable threats.

The evidence review sessions should be scheduled following the interviews so that risk assessors understand the general landscape of the security environment before trying to understand why certain assets are configured the way they are.

During interviews, the risk assessor will learn about topics that should be more closely understood through a review of configurations, system testing, or records review. Risk assessors should note during or soon after the interview which safeguards they will want to examine further, and inform the participants that they will likely be contacted later in the assessment to participate on those evidence review sessions. Further, the assessor should ask which personnel, processes, or information assets would be appropriate to examine to gather the appropriate evidence. The evidence review sessions can then be scheduled at the end of the interview.

Defining Risk Assessment Criteria

Introduction

Risk assessment criteria are the numerical and plain-language statements that an organization uses to evaluate their cybersecurity risk. The most familiar form of risk calculations, “Risk = Likelihood x Impact,” is the basis for risk analysis in the CIS RAM. But it is just the starting point for risk analysis.

Risk assessment criteria must be meaningful to the organizations that use them, so they must be tied to the potential benefit and harm that the organization may create. The impact of a cybersecurity breach may harm the organization itself, it may harm the organization’s ability to successfully achieve its mission, or it may harm others.


Because cybersecurity failures may harm parties both inside and outside an organization, risk assessment criteria must be universally meaningful and must address the interests of all potentially affected parties. Additionally, risk assessment criteria must demonstrate to authorities that the organization considers the risk of harm to others as much as the risk of harm to themselves, as stated in Principle 1.

While these requirements may seem complex, the method presented in this section will sufficiently address them while using a technique that is simple to develop and use.

Risk Assessment Criteria Foundations

The risk analysis provided in the CIS RAM is at its root a question of balance between the potential of future harm against the certain burden of a safeguard. Regulators and litigators have long considered this balance as key to acting as a “reasonable person.” The core structure of a risk statement is provided below to illustrate the core concept of balance.

Figure 3 - Balance Within Core Risk Analysis



Notice a few things right away with the model risk analysis in Figure 3.

- While organizations typically evaluate the observed risk to determine whether they should address or accept it, this risk statement deliberately compares the observed risk to a proposed safeguard.
- The criteria that evaluates the risk also evaluates the safeguard.
- The impact of the risk estimates the potential of harm to the organization and the potential harm against others.

Risk assessors compare risks to their proposed safeguards to determine whether those safeguards would create a foreseeably lower risk than the current state. To accomplish this, the assessor evaluates the current state risk (or “observed risk”) and the proposed safeguard using the same criteria to ensure comparability.

This comparison prevents organizations from implementing safeguards that are overly burdensome, or that create new, unacceptable risks. For example, an organization that uses software that is no longer supported by the vendor, but relies on that software for critical business

purposes, should find alternative methods for identifying and controlling potential security risks until they replace the software. If management recommends quickly changing out to inferior, but secure software, the organization may suffer a greater impact to their mission than the security risk they are trying to avoid.

While considering CIS Control 18: Application Software Security, a risk statement can be made to estimate the foreseeability of an impactful threat. The risk can be stated as it appears in Table 9 (where the risk score '6' is a product of the likelihood '2' and the highest impact score '3'):

Table 9 - Example Core Risk Statement

Observed risk	Likelihood	Impact to Us	Impact to Others	Risk Score
Hackers may exploit the unsupported, but critical application.	<u>2</u>	2	<u>3</u>	6

A risk assessor should then recommend and evaluate a safeguard to reduce the unacceptably high security risk, as illustrated in Table 10. Here, the organization would realize that the likelihood of a negative impact to their mission is greater than the current state risk. This is an obvious case of the burden being greater than the risk, and a recommended safeguard being unacceptable.

Table 10 - Example Unreasonable Proposed Safeguard

Proposed Safeguard	New Risk	Likelihood	Impact to Us	Impact to Others	Safeguard Risk
Replace application with inferior, secure application.	Application will operate inefficiently.	<u>3</u>	<u>3</u>	1	9

When faced with this analysis, the organization must then find another way to address the risk. This process will be described later in this chapter in the section Risk Treatment Recommendations.

But what should be apparent is that without a definition of risk assessment criteria the likelihood and impact scores are not meaningful. What would impacts or likelihoods of '1', '2', or '3' mean, anyway? The organization will need to create definitions for their likelihood and impact scores so that they are meaningful to all interested parties, and so that they provide a consistent method for risk evaluation.

Impact Definitions

Tier 1 organizations do not have a high degree of attention from management in operating cybersecurity risk. In such organizations, risk assessment criteria can be developed in simple terms that are business appropriate, but that do not use the business justifications that managers often need in order to make decisions.

Risk assessment criteria are composed of impact definitions and likelihood definitions. In its simplest form, an impact definition should consider the organization's mission (the value the organization provides others) and its obligations (the harm that it may cause others without appropriate safeguards). A simple impact model for the example organization described in Chapter 1 can look like Table 11.

Table 11 – Example Impact Definitions

Impact Score	Impact to Our Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Our Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	No harm would come to patients.
2	Some patients cannot access the information they need for good outcomes.	Few patients may be harmed after compromise of information or services.
3	We can no longer provide helpful information to remote patients.	Many patients may be harmed financially, reputationally, or physically, up to and including death.

Note that this example organization, a health technology manufacturer and service provider, has defined the impact to their service as their “mission” and the impact to others as their “obligation.” They are defining their mission in terms of their value to their constituency (patients who use their service) and their obligations to prevent harm to those patients due to an information breach.

Background – Impact Definitions

This document provides instructions for defining impacts and impact scores (magnitudes) in this section with more in-depth instructions and examples in the “Risk Analysis Techniques” chapter. The reader should understand before going further that organizations in most cases should not define impacts exclusively using financial values. While cost is a common and almost necessary consideration while evaluating risks and safeguards, if it is the only criterion, the organization will communicate to their personnel, as well as to interested parties and authorities, that cost is their only concern. The purpose that the organization serves and the harm that may befall others must be part of the evaluation if risk is to be responsibly tied to the potential of harm, and if the evaluation is to be understandable to regulators and legal authorities.

Organizations should also consider having more than three impact types in their impact definitions if they have more than one mission, multiple objectives, and many obligations that they need to consider in their risk analysis. While this expansion may create an increasingly wide risk register, it can help organizations feel comfortable all relevant interests were considered in their risk analysis.

Also note that the impact score of ‘1’ (which is shaded grey to separate it from the higher scores) describes impacts that would be generally understood as acceptable. If a breach led to a situation where patients continued to access helpful information, and there was no foreseeable harm to patients, then of course that would be interpreted as an acceptable impact. The organization should not be satisfied with a breach that had no impact (its incident response procedures should identify a root cause and address it so the breach does not re-occur), but in terms of risk planning, such a risk could be considered “acceptable.”

The impact score of ‘2’ would be used to estimate a risk in which harm would come to the mission of helping remote patients, or if some harm could come to the patients who rely on the confidentiality, integrity, and availability of information. An impact at a level of ‘2’ would be considered ‘not acceptable’ by the organization, its customers, regulatory agencies, or litigators.

So foreseeable risks (a likelihood of ‘2’) that are estimated with an impact of ‘2’ would likely not be considered “acceptable.” But an unforeseeable risk (likelihood score of ‘1’) that would create an impact of ‘2’ would be acceptable since the impact is considered not foreseeable.

The impact score of ‘3’ might be considered ‘catastrophic’ or ‘high’. The mission would fail completely, and the obligations to patient customers could harm many people, up to and including death (presumably because health information was inaccurate, or not available when it was critically needed).

With these impact criteria defined this way, the health information provider could estimate the impact portion of risk consistently. Some amount of knowledge about how risks would create those impacts would be necessary while conducting the risk assessment, but well-informed managers and personnel could provide plausible estimates of risk in a consistent basis using these impact definitions.

An in-depth explanation of how to develop impact definitions with multiple examples is provided in the chapter “Risk Analysis Techniques.”

Likelihood Definitions

This risk assessment method describes likelihood in terms of foreseeability. While risk likelihood is often described in terms of statistical probability, CIS RAM favors foreseeability because it uses simple terminology that aligns with common business practice, as well as the legal and regulatory language used to determine reasonableness of safeguards that reduce risks. Recommendations for aligning this likelihood model to probability methods are provided in the “Risk Analysis Techniques” chapter. By combining probability with foreseeability, organizations may benefit from both data-driven analysis and due care analysis.

The likelihood definition for a Tier 1 organization could be simply constructed, similar in structure and depth to the impact definitions as shown in Table 12.

Table 12 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will occur at some time.

- “Not Foreseeable” implies that a threat is not plausible in the environment that is being assessed. Loss of portable media may not be foreseeable during a risk assessment of a hosted application.
- “Foreseeable” implies something that is plausible, but the organization would be surprised if it occurred. A founding executive taking copies of sensitive data to competitors may be considered foreseeable, even if it is not expected.
- “Expected” implies a threat that is not common, but that would eventually happen. Phishing attacks or other social engineering attacks may be expected in many environments.

When risk assessors estimate the likelihood of a threat, they will select scores ‘1’, ‘2’, or ‘3’ using the foreseeability definition as their guidance. Organizations may add time-based limits to their foreseeability definitions (i.e. “Foreseeable within planning thresholds,” “Expected within the five-year plan” or “Not foreseeable in the next fiscal year”). If organizations do introduce time limits to their likelihood definitions they should prioritize risk treatment investments to meet these

timelines. That may be excessively challenging to many organizations, so they should proceed with caution.

The simplicity of these definitions will assist Tier 1 organizations to quickly and consistently estimate whether they expect impactful risks to occur.

Exercise:

The reader should develop their organization's risk assessment criteria using the "Criteria - Tier 1" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Working with a business management sponsor who can help ensure that the Mission and Obligations definitions are sensible to the organization.
2. Working with legal counsel to help ensure that impact definitions address the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.
3. Referring to guidance for defining and scoring impact types in the "Risk Analysis Techniques" chapter.

The risk assessor will need to use their professional judgment to define impact types, and to describe levels of impact that the organization must manage to. Because risk assessment criteria are a declaration by the organization of what they will manage to in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

Defining Risk Acceptance Criteria

Introduction

Because risk assessments are essentially questions of balance, the criteria for accepting risk should help determine whether balance was achieved. In CIS RAM risk acceptance has two components:

- Appropriate risk: That the likelihood of an impact must be acceptable to all foreseeably affected parties
- Reasonable risk: That the risk posed by a safeguard must be less than or equal to the risk it protects against.

While these components have been demonstrated briefly in Chapter 1, the "appropriate risk" will be described in more detail in this section. "Reasonable risk" will be described later in the Risk Treatment Recommendations section further on.

Recall that impact definitions were worded so that the acceptable impact definitions would seem appropriate to any person who read them. For Tier 1 organizations that use an impact score range of '1' through '3' the range of acceptable impact scores is simply '1'. Definitions for impacts that would score at least a '2' would therefore represent impacts that an organization, and presumably its interested parties, would find unacceptable.

Table 13 - Unacceptable Impacts

Impact Score	Impact to Our Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Our Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information and outcomes are on track.	No harm would come to patients.
2	Some patients cannot access the information they need for good outcomes.	Few patients may be harmed after compromise of information or services.
3	We can no longer provide helpful information to remote patients.	Many patients may be harmed financially, reputationally, or physically, up to and including death.

Similarly, likelihood scores for Tier 1 organizations ranged from '1' through '3' where the score of '2' represented the lowest "foreseeability" score.

Table 14 - Unacceptable Likelihood

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will occur at some time.

Tier 1 organizations should determine that they would safeguard against risks that reached a threshold of unacceptability; for example, risks that could *foreseeably* (likelihood is '2') *prevent patients from getting access to information* (impact is '2') *or cause a breach that may harm patients* (impact is '2'). So if Risk = Impact x Likelihood, then the organization would invest against risks that are scored '4' or greater. *All lower risks can be accepted!*

Table 15 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
2	x	2	=	4
... therefore ...				
Acceptable Risk			<	4

Consider how a *reasonable* risk would be described: If a risk *cannot foreseeably* (likelihood is '1') prevent the organization from *providing helpful information to patients* (impact is '3') then *that is acceptable*. It sounds acceptable to reasonable people, and $1 \times 3 = 3$ which is less than 4.

Also see how the calculation works when an acceptable impact of *no harm to patients* ('1') is *expected to occur* ('3'). $1 \times 3 = 3$, which is again less than '4'. This is an acceptable risk. While risk heat maps are not used in CIS RAM, organizations can now consider that heat maps can represent actual risk acceptability based on organizational requirements, and a duty of care to others.

Figure 4 - Example Risk Heat Map

	Impact		
	1	2	3
Likelihood	3		
	2		
	1		

Exercise:

The reader should define their organization's risk acceptance criteria using the "Criteria - Tier 1" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Working with a business management sponsor who can help ensure that the risk acceptance criteria are sensible to the organization.
2. Working with legal counsel to help ensure that the definition for risk acceptance addresses the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.

The risk assessor will need to use their professional judgment to identify levels of acceptable risk. Because risk acceptance criteria are a declaration by the organization of what they will tolerate in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

A Control-Based Risk Assessment Process

Introduction

Tier 1 organizations that use the CIS Controls, but that have not established a strong capability for managing cybersecurity risk will find that control-based risk assessments are well-suited to their needs.

This CIS Risk Assessment Method is designed to help organizations responsibly use the CIS Controls to the degree that they are able, even if the controls cannot be implemented completely. The risk assessment method helps Tier 1 organizations:

1. Model safeguards based on CIS Controls V7 that addresses their risks, while working within their constraints.
2. Prioritize the safeguards they should implement, based on their organization's risks.

- This section will describe a risk register that is available as a template in the supplementary document *CIS RAM Workbook*.

Up to this point, the organization has identified the information assets or asset classes that they will be risk assessing. They have also developed risk assessment criteria and risk acceptance criteria. Using the risk register template provided in the supplementary document *CIS_RAM_Workbook*, the organization will see a list of CIS Controls and sub-controls that will act as the main index, or driver, for modeling their risks.

Figure 5 – Risk Register Layout Map

[illegible]

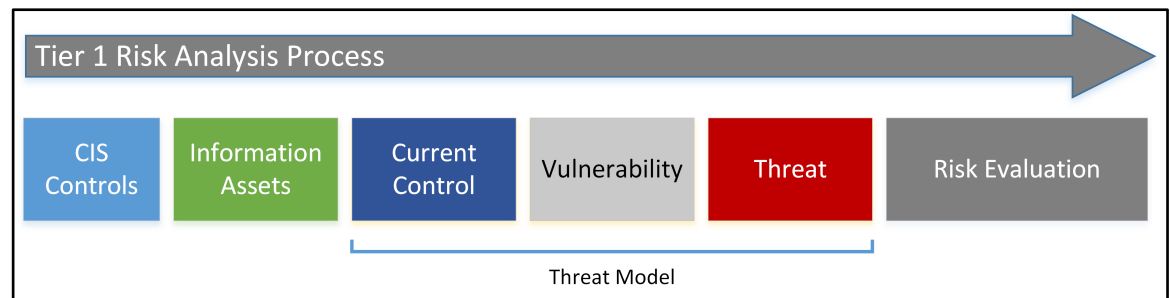
- A. The column headers and guiding text help the reader or risk assessor understand the information that is contained in the column.
- B. The CIS Controls help the risk assessor consider controls that should be in place to protect information assets.
- C. The in-scope information assets or asset classes are identified.
- D. The “threat model” includes the following three items:
 - a. How the organization implements the CIS Control (if they do) to protect the information asset or asset class.
 - b. The vulnerability that may exist if the control is not fully implemented.

- c. The threat that may compromise the asset because of the vulnerability.
- E. The risk evaluation estimates the likelihood that the threat would succeed and the impacts to the mission and obligations if it did. The resulting risk score is then calculated as a product of the likelihood and the higher of the two impact scores.
- F. Risk treatments are recommended for risks that are evaluated as unacceptably high. Safeguards that are based on CIS Controls V7 are described, and they are in turn evaluated for the risk that they may pose to the mission and objectives. A “safeguard risk” score is calculated which should be lower than the risk acceptance criteria, and the risk that it is meant to address.

The Process

The risk assessor will analyze each risk by taking the following steps as diagrammed in Figure 6.

Figure 6 - Tier 1 Risk Analysis Diagram



1. Using the risk register template for Tier 1 organizations that is provided in *CIS_RAM_Workbook*, read and consider the CIS Control that is stated in one of the risk register rows.
2. Select information assets or asset classes that are listed in the asset inventory. Record the selected assets in the “Information Asset” cell in that row.
 - a. If there are multiple information assets or asset classes to consider for each CIS Control, either list them all in one cell in the information asset column, or add multiple rows for the CIS Control so that each information asset is analyzed separately. *The decision should be based on how granular the risk assessor is prepared to be in their analysis and planning, and how useful the distinction between assets would be.*
3. Gather evidence for how well the CIS Control is applied against the selected information assets.
 - a. Evidence may be in the form of interviews, a review of configurations, or a review of evidence, such as records and logs. This step requires knowledge and experience in detecting vulnerabilities and understanding security threats. Methods for gathering evidence are provided in the chapter “Risk Analysis Techniques.”
4. Describe the safeguard that implements the CIS Controls and how the safeguard is applied at the organization in the “Current Control” cell of that row.
5. Consider the difference between the CIS Control and the currently used safeguard and determine whether there is a deficiency in how the control is currently deployed and operating. If the current control is not implemented as described, how would this be described as a vulnerability?

- a. Consider the objective of the CIS Control. For example, CIS Control 2.3 states “Utilize software inventory tools throughout the organization to automate the documentation of all software on business systems.” The sub-control’s objective is to itemize all operating systems and applications that are in use so the organization knows what software it should control. If the current control does not meet the objective, then state the gap as a vulnerability in the vulnerability cell, such as: “We do not have a current, automatically updated list of current operating systems or applications that operate on our systems.”
6. Now consider the threat that could occur because of the vulnerability.
 - a. The above vulnerability could be paired with the threat: “Malware or hackers could take advantage of software vulnerabilities that we did not know about or protect against.”
7. Next, estimate the likelihood that the threat would succeed, and the impact it may create.
 - a. Likelihood estimation can be challenging at first, but the risk assessment criteria was developed to provide some guidance in that estimation process. Further guidance is provided in the “Risk Analysis Techniques” chapter later in this document.
 - b. Impact scores should provide estimates of the impact that such a threat would create. Consider the likelihood and impact scores as a pair. In other words, “What is the likelihood that this impact would result?” Examples below will provide further guidance.
8. The risk score will be automatically computed by multiplying the likelihood score by the higher of the two impact scores.

Tier 1 Risk Assessment Example 1 – Knowing Whether a Current Safeguard is Enough

Let’s examine how this process works using risk analyses for our example Tier 1 organization.

While conducting their risk assessment, the organization starts with CIS Control 1.1 which instructs them to deploy an automated asset inventory discovery tool. They know they don’t have such a tool, and are concerned about the time and potential cost of researching and obtaining one. Additionally, they don’t know whether this should be their top priority, given other items that are on their mind, such as hardening field devices and vulnerability management. While CIS Controls V7 provides clear guidance on the criticality of this important control, there may be other areas of concern that should be addressed earlier, based on the organization’s environment.

Using the risk register template for Tier 1 organizations, the example organization first reviews the control that they are analyzing.

Table 16 - CIS Control Example CIS Control 1.1

CIS Control	Description
1.1	Utilize an active discovery tool to identify devices connected to the organization's network and update the hardware asset inventory.

The risk assessor considers the information asset that the control would protect. In this case, they would apply an asset inventory discovery tool to all network attachable devices. They could focus purely on a class of assets, such as portable workstations, “headless” or “internet-of-things” devices, smart phones, laptops, or devices in a specific network such as the corporate VLAN that is used to connect wireless devices. But to make things simple for their first assessment, they decide that the asset class they will assess will be “all devices.”

Table 17 - Tier 1 Information Asset Example

Information Asset
All devices.

Next, they must think through and record what their current control is, what resulting vulnerabilities may exist, and what threats would be of concern to them. They realize that they don't have an automated tool that provides a regular inventory, but they do have a vulnerability scanning tool that they occasionally use. That may be useful here.

But the control clearly has an objective, which is the automatic detection of all systems that appear on the network. If the organization occasionally uses a vulnerability scanner, then a vulnerability related to this control would be that new systems could join the network and not be detected until the next vulnerability scan occurred. The resulting threat would be obvious: Compromised systems may operate on the network between scans.

So the next three columns would look like this:

Table 18 - Threat Model Example

Control	Vulnerability	Threat
Vulnerability scans occur occasionally and may not identify all systems that have been on the network between scans.	Systems that have joined the network between sporadic scans will not be detected.	Hackers or malware may attack and control systems that have not been detected, controlled, and monitored.

Now that the threat has been modeled, the risk assessor should estimate the likelihood and impact of the risk. The organization must consider the likelihood that an impact would occur if the threat were successful. Risk assessors should think of the likelihood and impact as a dependent pairing. In this example, the organization may believe that multiple systems will join and leave their network without detection. That is a highly likely scenario for them. But they may also believe that the risk is foreseeable but unexpected for an undetected system to cause an impact if the visitors they receive are employees of well-secured partner organizations.

In terms of the impact that people may suffer, the risk assessor considers how harm could be done in this foreseeable but unlikely scenario. If an employee of a secured partner were to bring in a laptop that was infected with malware that could spread to other systems on the network, what harm could that do? If these partner employees join a network that contains a mix of systems, some with highly sensitive information, then is it foreseeable but not expected that the scenario could expose records that could cause harm to few patients, or many patients? Would the mission be reduced to the point that some patients could not get information that could improve health outcomes?

The organization determines that it is foreseeable but not expected that records for many patients may be exposed in the risk scenario they modeled. They do not believe that the risk scenario would affect their mission. So now they will add this information to their risk register to see how the risk evaluates.

Recall the definitions for the impact and likelihood scores that the organization created. Impact scores were defined earlier as shown in Table 19.

Table 19 – Example Impact Definitions

Impact Score	Impact to Our Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Our Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	No harm would come to patients.
2	Some patients cannot access the information they need for good outcomes.	Few patients may be harmed after compromise of information or services.
3	We can no longer provide helpful information to remote patients.	Many patients may be harmed financially, reputationally, or physically, up to and including death.

Impact score '1' is shaded to indicate that it is considered an acceptable impact to all parties. Also recall that likelihood was defined with this next table.

Table 20 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will occur.

A risk that is foreseeable but not expected to occur (likelihood = 2) in a way that creates no impact to the mission (mission impact = 1), but that would create harm to many patients (obligations impact = 3) would appear as below.

Table 21 - Example Risk Estimation

Threat Likelihood	Mission Impact	Obligations Impact	Risk Score
2	1	3	6

The risk score is the product of the likelihood score and the higher of the two impact scores, which in this case is '2 x 3 = 6'.

Also recall that the risk acceptance criteria for the Tier 1 organization looked like this:

Table 22 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
2	x	2	=	4
... therefore ...				
Acceptable Risk			<	4

Because acceptable risk is anything below '4', and the observed risk associated with CIS Control 1.1 is '6', the risk is unacceptably high.

We can bring these elements together to illustrate the point more clearly. (While the risk register in the workbook displays this example in horizontal format, this risk is displayed in vertical format for readability in this document.)

Table 23 - Example Risk for CIS Control 1.1

Risk Analysis	Value
CIS Control	1.1
Description	Utilize an active discovery tool to identify devices connected to the organization's network and update the hardware asset inventory.
Information Asset	All devices.
Control	Vulnerability scans occur occasionally and may not identify all systems that have been on the network between scans.
Vulnerability	Systems that have joined the network between sporadic scans will not be detected.
Threat	Hackers or malware may attack and control systems that have not been detected, controlled, and monitored.
Threat Likelihood	2
Mission Impact	1
Obligations Impact	3
Risk Score	6
Risk Acceptability	Not Acceptable

So the organization realizes, based on its own criteria for scoring and accepting risk, that their occasional use of vulnerability scans is not sufficient for addressing the risk of infected systems joining the network. This risk is not acceptable because it is not "appropriate" (its risk is higher than the acceptable score of "less than 4.")

But they are not sure what to do about this risk, because they don't know whether they will be able to afford the time or budget to implement a more robust solution for CIS Control 1.1 (such as a network access control appliance) and they have many more controls and risks to consider. The method for identifying reasonable ways to implement safeguards will be addressed in the Risk Treatment Recommendations section later in this chapter.

First, however, we will examine a few more of the CIS Controls and see how the Tier 1 organization analyzes them.

Tier 1 Risk Assessment Example 2 – Risk Acceptability in Different Contexts

Further along in the risk assessment, the Tier 1 organization considers CIS Control 3.4 that states, "Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor." This is a challenging control for many organizations. While the objective of automatic patching of vulnerable systems is important, many systems and applications will fail when some security patches interfere with their functionality.

For example, applications that rely on code libraries that are replaced with more secure versions during patching may fail. As a result, many organizations test patches before releasing them to active systems and applications.

The Tier 1 organization believes that they are doing well in this regard. They have two environments; a production environment in which their applications operate on the Internet, and a corporate environment in which they run their business and application development environment. When they run their sporadic vulnerability scans in their production environment, they respond immediately to identified vulnerabilities that can be patched. In fact, they run their vulnerability scans when they receive word of high-risk vulnerabilities from a threat information service they subscribe to (for this example, a fictional service provider named “Threat Info Service”). Their application stack is simple, and based completely on standard implementations of the vendors’ application framework. So when the vendor sends patches, they can be applied rapidly – even if manually – with little risk to the applications.

They believe their risk is low here in terms of security. But they are concerned about their patient customers not being able to use their systems during patching downtimes. So they assess the risk in this way.

Table 24 - Example Risk Analysis for CIS Control 3.4

Risk Analysis	Value
CIS Control	3.4
Description	Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.
Information Asset	All devices in Production Environment.
Control	Vulnerability scans occur when Threat Info Service announces a moderate-to-high vulnerability that needs patching. Team reliably patches systems within 24 hours of announcement.
Vulnerability	A 24-hour window of vulnerability remains with the current process.
Threat	Hackers or malware may attack and control systems that have not been patched within the 24-hour period after the vulnerability was announced.
Threat Likelihood	1
Mission Impact	2
Obligations Impact	1
Risk Score	2
Risk Acceptability	Acceptable

The Tier 1 organization has determined that its risk associated with CIS Control 3.4, at least in its production environment, is acceptable because the risk is “appropriate” (it’s score is less than ‘4’). This allows the organization to keep their current processes in place, and to focus on higher risks first.

But they still have an internal network to consider. In this network they have an important enterprise management application that relies on an older operating system, and that does not work with more advanced patches to the operating system. The application vendor says they will release a more secure version in the next year, and the Tier 1 organization knows that converting

to the new version will be very burdensome. In the meantime, they have operating system vulnerabilities in the enterprise management application environment that need to be addressed. But they are afraid to apply those patches because they may break the application.

They evaluate the risk in their risk register with the values listed below.

Table 25 - Example Risk for CIS Control 3.4

Risk Analysis	Value
CIS Control	3.4
Description	Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.
Information Asset	Enterprise management application in the internal corporate network.
Control	Vulnerability scans occur when Threat Info Service announces a moderate-to-high vulnerability that needs patching. Team patches most systems within 24 hours of announcement.
Vulnerability	Enterprise management application systems are unpatched for more than one year.
Threat	Hackers or malware may attack and control the enterprise management application environment.
Threat Likelihood	2
Mission Impact	2
Obligations Impact	3
Risk Score	6
Risk Acceptability	Unacceptable

They clearly have an unacceptable risk with how this control protects their corporate network (the risk score of '6' is inappropriately high). The enterprise management application is creating a risk that should be addressed in some way. The organization will address this risk in the following section, Risk Treatment Recommendations.

Tier 1 Risk Assessment Example 3 – Prioritized Risks

As the Tier 1 organization later considers its risk related to CIS Control 14.9, they will address a common concern about log management; what events that are captured in logs and repositories should the organization focus on? CIS Control 14.9 states, "Enforce detailed audit logging for access to sensitive data or changes to sensitive data."

The Tier 1 organization suspects that they will need to invest in their log management and SIEM technologies soon, but they are also aware that choosing which log messages to capture, prioritize, and alert on will require careful thought and lots of experience. Their main concern regarding log management is detecting abuse of systems and data access. Morale has been low, and they are in a competitive business that may cause internal employees to steal sensitive data from their enterprise management application and provide it to a competitor.

They are sure that they are not doing well enough with their access log review now, but they will analyze this risk to evaluate and prioritize this risk. They will also use their risk treatment analysis further on to determine what log management configurations are appropriate for their enterprise management application.

Table 26 - Example Risk for CIS Control 14.9

Risk Analysis	Value
CIS Control	14.9
Description	Enforce detailed audit logging for access to sensitive data or changes to sensitive data.
Information Asset	Enterprise Management Application
Control	Access logs are captured and stored locally, and not reviewed.
Vulnerability	The organization is unaware of suspicious or inappropriate use.
Threat	Rogue employees or hackers may use escalated privileges and may access and abuse nonpublic information in the application.
Threat Likelihood	3
Mission Impact	2
Obligations Impact	3
Risk Score	9
Risk Acceptability	Unacceptable

After having entered a few risks in the risk register, this analysis does not surprise them. They were sure that by reviewing this control they would highlight a problem. And now they see that this risk is both inappropriately high, and should be prioritized over the other risks that they analyzed previously. By the time the risk assessment is complete, this risk will rank among the first items that should be addressed (it has the maximum risk score of '9'). They will select and design their risk treatment safeguard in the following step, Risk Treatment Recommendations.

Exercise:

The reader should refer to the template Risk Register – Tier 1 that is provided in the supplementary document *CIS_RAM_Workbook*. They may use the risk register template to enter a set of risks that are associated with the CIS Controls and information assets that are in scope of their assessment.

While doing this exercise, the reader should consider:

1. When one CIS control can be appropriately stated for the whole scope, an asset class, or a stand-alone information asset.
2. Stating at least one risk per CIS control. One CIS Control may appear multiple times if asset classes and information assets use the control differently.
3. Whether a control or information asset requires examination to understand its actual configuration and effectiveness.
4. Whether the organization can tolerate the amount of effort and time that the risk assessment requires.
 - a. Organizations should not try to “boil the ocean.” A risk assessment can only be completed using available resources.
 - b. The organization should use high-level analysis (review of policies and interviews) if they do not have extensive time and resources.
 - c. Information assets should be tested and examined in more detail as time allows.
 - d. The organization should plan recurring risk assessments to identify more risks over time.
5. Collaborating with information security subject matter experts to help model threats that are foreseeable in the environment, and to help evaluate the effectiveness of current safeguards.

The risk assessor will need to use their professional judgment to select the controls and information assets and to model threats that should be analyzed in the risk assessment. Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Tier 1 Risk Assessment Summary

After having analyzed a set of risks, the Tier 1 organization begins to understand a few concepts about risk analysis:

1. The CIS Controls can help Tier 1 organizations model threats by comparing their current practices to known-good practices. The CIS Controls each identify a way to safeguard systems, so they conversely show where and how something might go wrong.
2. Once impacts are defined in terms of acceptable, unacceptable, and high consequences, they support a very rapid, consistent, and simple method for evaluating risks.
3. Similarly, when likelihoods are defined using easy-to-communicate terms, such as foreseeability, risks can be easy to estimate while also being credible.
4. If a risk can be credibly assessed as not foreseeably creating an impact, or foreseeably creating an acceptable impact, then an organization can reasonably accept that risk.
5. If a risk poses a higher risk score than other risks, it will likely be prioritized over other risks during risk treatment design and planning.

Risk Treatment Recommendations

Introduction

Organizations often think of security safeguards as obstacles to business and productivity. Safeguards often cause personnel to take extra steps to get access to systems or information, or to get approval for normal business activities. Safeguards require investments in time and money, which compete with other priorities. And if they become too disruptive to the mission an organization's mission, security safeguards can become disliked and avoided.

In fact, disruptive safeguards often cause personnel to work around them just to get their jobs done, which creates more risk.

But risk treatment recommendations can and should result in safeguards that are demonstrably reasonable. And while obtaining a clear definition for "reasonable safeguards" has been a challenge in the legal, regulatory, and information security communities, the CIS RAM provides a practical solution. Risk assessors evaluate risk treatment recommendations to determine whether a security safeguard is reasonable by; comparing the safeguard to the risk it is meant to reduce, and by comparing the safeguard to the risk acceptance criteria.

Risk treatment recommendations are simple to evaluate once the risk assessment criteria and initial risk analysis have been established. The process occurs over the following steps:

1. While examining an unacceptably high risk, review the CIS Control that corresponds with the risk and recommend a feasible way for the organization to implement or improve that control.
2. If that control is not feasible in the near-term, recommend other CIS Controls related to the risk that can be used to reduce it.
3. Evaluate the risk of the recommended safeguard to understand the burden it would pose to the organization. Then compare that safeguard risk to the risk acceptance criteria to determine whether it is appropriate.
4. Also compare the evaluated risk of the recommended safeguard to the observed risk to determine whether the safeguard is reasonable (safeguards with lower risk scores than the observed risk are reasonable.)
5. Sort the risks by their risk score to prioritize the risks and risk treatments that the organization will invest in.

This section demonstrates these steps in detail by describing the process, then by modeling risk treatments for the unacceptably high risks that were evaluated in previous sections.

The reader should review the definitions of 'reasonable' and 'appropriate' that are provided in the glossary. These terms will be used regularly in this section and have distinct meanings.

1. **Appropriate:** A condition in which risks to information assets will not foreseeably create harm that is greater than the organization or its constituents can tolerate.
2. **Reasonable:** A condition in which safeguards will not create a burden to the organization that is greater than the risk it is meant to protect against.

Risk Treatment Objectives

The objective of well-formed risk treatment recommendations is to create a prioritized list of information security safeguards that would provide appropriate protections while not posing too great a burden on the organization's purpose.

The risk treatment recommendation exercises that are demonstrated in this section examine the unacceptable risks that were illustrated earlier in the document and will select CIS Controls that


would reduce those risks to a degree that is both reasonable (not overly burdensome) and appropriate (not unacceptably harmful).

Recommending Risk Treatment Safeguards from CIS Controls V7

As we examine unacceptably high risks, we will recommend safeguards that are based on CIS Controls V7. But some of the safeguards that an organization is prepared to implement and operate may not be implemented exactly as described in CIS Controls V7. This process takes into account how to select controls that address risks, and how to determine whether they are designed in a way that makes sense in context of both the risk, and the potential burden to the organization.

Recall the relationship between analyzed risks and their recommended risk treatments in Figure 7.

Figure 7 - Balance Within Core Risk Analysis



A risk and its proposed safeguard are both evaluated using the same criteria. If a proposed safeguard has a higher risk (its “safeguard risk”) than the risk acceptance criteria, then it’s not appropriate. If the safeguard has a higher score than the observed risk, then it’s not reasonable.

The exercises in this section will focus on matching completed risk analyses (in blue) with newly recommended safeguards (in green).

Risk Treatment Example 1 – CIS Control 1.1

The first exercise that demonstrated risk analysis in Table 23 showed the Tier 1 organization that its method for identifying and inventorying technical assets in its network was not sufficient. The risk it posed to the organization was too high.

To recommend an appropriate safeguard the organization reviews the description for CIS Control 1.1 and they consider its objective. CIS Control 1.1 intends that organizations should actively and passively inventory all IP systems that join an organization’s networks so they know what technical assets to include in their risk management programs, safeguards, controls, and processes.

They realize that the best solution for them will be to purchase and install an appliance that actively identifies and catalogs IP hosts that join the network, allowing the IT staff to add detailed information about each asset over time. At this time they don’t need to decide what the full feature set should be (i.e. will it operate as a network access control device to enforce safe configurations of systems, will it actively scan for software licenses, will it be integrated into a change management database or service desk application, etc.).

Next they will evaluate what they believe the safeguard risks to be, and will determine whether the safeguard risk is appropriate or not. This step is demonstrated in Table 27.

Note: The risk assessor will record how they will address their risk by stating either “Accept,” “Reduce,” “Transfer,” or “Avoid.” *Accepting* and *reducing* risks will be intuitive to the reader. An

organization may *transfer* a risk by contracting a third party that may handle the risk better, or by acquiring an insurance policy against the risk. The organization may also *avoid* the risk by no longer engaging in the processes, or handling the information assets that cause the risk.

Table 27 - Example Risk for CIS Control 1.1

Risk Analysis	Value
CIS Control	1.1
Description	Utilize an active discovery tool to identify devices connected to the organization's network. This tool shall automatically update the organization's hardware device inventory when devices are discovered.
Information Asset	All devices.
Control	Vulnerability scans occur occasionally and may not identify all systems that have been on the network between scans.
Vulnerability	Systems that have joined the network between sporadic scans will not be detected.
Threat	Hackers or malware may attack and control systems that have not been detected, controlled, and monitored.
Threat Likelihood	2
Mission Impact	1
Obligations Impact	3
Risk Score	6
Risk Acceptability	Not Acceptable
Risk Treatment Option	Reduce
Recommended Safeguard	Purchase and implement an appliance that actively and passively identifies IP hosts in all networks. Implement a process for routinely adding information about assets to the appliance. Appliance should optionally alert on new hosts that join the network.
Safeguard Risk	A moderate cost would have minimal impact on the budget. Installation of the tool is likely not disruptive. Moderate cost in personnel time to add information about IP assets to the appliance database. After a baseline is established, we will be able to distinguish between organization-owned systems, and systems that we do not control. Alerting can be set after baseline is complete.
Safeguard Threat Likelihood	1
Safeguard Mission Impact	1
Safeguard Obligations Impact	3
Safeguard Risk Score	3
Risk Acceptability	Acceptable

Review the risk assessment criteria and risk acceptance criteria for the Tier 1 organization and observe that the organization believes that with the recommended safeguard in place the risk would no longer be foreseeable.

The Tier 1 organization has determined that their recommended safeguard – an appliance that can identify and inventory IP hosts and alert on new hosts – is an acceptable and reasonable safeguard. The estimated safeguard risk is less than their acceptable level of risk, and is lower than the originally evaluated risk that it is addressing.

Background – How Realistic Are Safeguard Risk Estimates?

Critical readers will question how the organization and their risk assessor will know whether their safeguard risk estimations are realistic. After all, how can they know prospectively what their risk would be in such a situation?

There are two important items to keep in mind while gaining comfort with this practice; understanding the legal and regulatory expectations for risk management, and information security standards for evaluating safeguards after they've been implemented.

Law and Regulation: Laws and regulations generally expect risk analysis to evaluate safeguards that are required for achieving compliance, and expect that the risk analysis is performed by appropriately skilled and informed people. These analyses do not guarantee security that is sufficient against any threat, but they do provide a plan toward improved security and compliance that is prioritized by the likelihood of harm, and that has no intolerable harm as its goal.

Information Security Risk Management Standards: Information security risk assessment standards that the CIS RAM is based on operate within fuller risk management programs and cycles. ISO 27005 operates within the ISO 27000 family of standards, and NIST 800-30 works within the NIST Special Publications. Each of these families of standards requires continuous analysis of security safeguards, including analysis of controls after they've been implemented to determine whether they are effective at addressing their security objectives. Recommended safeguards should therefore be risk assessed again after implementation to be sure that they achieve their intended objectives.

Risk Treatment Example 2 – CIS Control 3.4

The Tier 1 organization's second risk analysis involved their method for identifying and addressing system patches in their production environment and their corporate environment. They were comfortable with their patch management process in their production environment. Even though they were not using the control described in CIS Control 3.4 as it was written, they estimated that the likelihood and impact of the risk that their current methods exposed them to met an acceptable level of risk.

Their risks in the corporate environment were higher, though. Their enterprise management application had vulnerabilities that the manufacturer could not provide patches for, but expected to address the known vulnerabilities in a major release soon. The organization believed the release upgrade would be significantly disruptive in the short term, and was hoping to find an alternative safeguard that would be reasonable and appropriate. So the organization will model recommended risk treatments in Table 28 to identify such a safeguard.

Knowing that they could not implement CIS Control 3.4 as written, they consider the objective of the control instead to see if there are alternative means of meeting the objective. They determine that CIS Control 3.3 intends that patches are applied as quickly as is possible, which is not

feasible in this case. Because CIS Control 3.3 is not feasible, they turn to other CIS Controls to see what alternatives are available to them.

The CIS® provides a document titled the *CIS Community Attack Model*¹⁶ that can be helpful to the organization for finding alternative controls. The Community Attack Model (provided in the *CIS_RAM_Workbook* in a tab titled “Attack Path Models,” and at the CIS website) associates the CIS Controls with the way they reduce risks in each stage of an incident based on the functions found in the NIST Cybersecurity Framework.¹⁷ This can be very helpful for identifying related controls.

Figure 8 - Partial Community Attack Model

	Controls	Attack Stages				
		Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege
Functions	Identify	control of HW, SW inventory; Network logs	threat intelligence			control of administrative privilege
	Protect	firewall; mail gateway filtering; web filtering; manage ports, protocols, services; continuous vulnerability assessment	hardened configurations	continuous vulnerability assessment; firewall; mail gateway filtering; web filtering; secure remote access; NIPS	patching; hardened configurations; HIPS; anti-malware; containerization; app whitelisting; Data Execution Protection	control of admin privilege; data security; hardened configuration; continuous vulnerability assessment
	Detect	firewall; honeypot; Network authentication; Network logs	audit logs; threat intelligence	audit logs; Anti-malware; Network Intrusion Detection system	HIPS; anti-malware; containerization; app whitelisting; Data Execution Prevention;	account monitoring; control of admin privilege; audit logs; Configuration Monitoring
	Respond	honeypot			Incident Response - Execution	audit logs; Configuration Management; Account Management
	Recover				Incident Response - Execution; control of HW, SW inventory	

Risk assessors may reference the Community Attack Model to find controls that can be complementary and alternative to the recommended safeguards they are assessing. If an organization struggles to implement a sub-control, they could look for controls that play a similar role in the Community Attack Model to find alternative controls that might help them meet the same security objective. For example, if an organization cannot easily use audit logs to *detect* *delivery* of a kind of threat, they may look to another control in the cell that intersects with the

¹⁶ The *Community Attack Model* may be accessed here: <https://www.cisecurity.org/white-papers/cis-community-attack-model/>

¹⁷ *Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0*, National Institute of Standards and Technology. February 12, 2014

detect row and the *delivery* column to find similar controls – and to eventually see network intrusion detection controls, which may be more useful in their environment.

Given the intent of CIS Control 3 (of which CIS Control 3.4 is a member) to conduct continuous vulnerability assessments, the risk assessor reviews the Community Attack Model and finds “continuous vulnerability assessment” in multiple locations. It appears to be useful for *Protecting* against *Initial Recon*, *Protecting* against *Delivery*, and *protecting* against *Misuse / Escalate Privilege* (as well as other uses that are displayed in the remainder of the model not visible in Figure 8).

The threat they are trying to address could be mitigated by *protecting* the system at *delivery*, and to prevent attackers and malware from knowing what vulnerabilities are present. So the risk assessor considers the possibility of network intrusion prevention systems (NIPS) and refers to CIS Control 12 for “Boundary Defense” to see how that is described. While reviewing CIS Control 12’s sub-controls, they come across CIS Control 12.7 and see a description for an intrusion prevention system (IPS) that might be appropriate for their risk. An IPS could detect and prevent actions that match exploits to known vulnerabilities. Other options, such as putting the enterprise planning application and its users in a separate VLAN could reduce the likelihood of an attack as well, but there would be many systems in that VLAN, so the likelihood reduction might be small.

They recommend and evaluate a safeguard based on CIS Control 12.7.

Table 28 - Example Risk for CIS Control 3.4

Risk Analysis	Value
CIS Control	3.4
Description	Deploy automated software update tools in order to ensure that the operating systems are running the most recent security updates provided by the software vendor.
Information Asset	Enterprise management application in the internal corporate network.
Control	Vulnerability scans occur when Threat Info Service announces a moderate-to-high vulnerability that needs patching. Team patches most systems within 24 hours of announcement.
Vulnerability	Enterprise management application systems are unpatched for more than one year.
Threat	Hackers or malware may attack and control enterprise management application environment.
Threat Likelihood	2
Mission Impact	2
Obligations Impact	3
Risk Score	6
Risk Acceptability	Unacceptable
Risk Treatment Option	Reduce
Recommended Safeguard	(CIS Control 12.7) Purchase and implement an IPS solution to detect, alert on, and prevent attacks on the

Risk Analysis	Value
	enterprise management application, and other vulnerable systems in the environment.
Safeguard Risk	A significant cost would have significant impact on the budget. Installation of the tool in detection mode is likely not disruptive. Installation of the tool in prevention mode is likely disruptive. Moderate cost in personnel time to implement and configure the IPS system.
Safeguard Threat Likelihood	3
Safeguard Mission Impact	2
Safeguard Obligations Impact	1
Safeguard Risk Score	6
Risk Acceptability	Unacceptable

Using a safeguard based on CIS Control 12.7 in this planned deployment still evaluates as unacceptable. The risk assessor is certain that the IPS in terms of budget and potential disruption of service will prevent the organization from properly servicing some of their patient user population (Likelihood = '3'; Mission Impact = '2'; Safeguard Risk Score = '6'). For a Tier 1 organization especially, implementing a tool that may create disruptions to business would be intolerable, and would likely lead to increased frustration by non-technical management with information security.

So the risk assessor considers deploying an open-source IPS in detection and alert mode (Intrusion *Detection* System, or IDS), rather than a commercial IPS in prevention mode. This would allow the team to be aware of suspicious activities and to respond without creating a failure of business functions. After becoming familiar and comfortable with the IPS and seeing what actions it alerts on, the technology team may be able to selectively block access to high-risk systems, such as the enterprise management system.

They find a corresponding CIS Control after reviewing the sub-controls under CIS Control 12. The risk assessor reads CIS Control 12.6 that states, "Deploy network-based Intrusion Detection Systems (IDS) sensors to look for unusual attack mechanisms and detect compromise of these systems at each of the organization's network boundaries." This control can be the initial safeguard, allowing the organization to improve the safeguard to an IPS (CIS Control 12.7) when the organization is ready to block traffic that they understand better.

Table 29 models a variation to this recommended safeguard.

Table 29 - Example Risk Treatment Recommendation CIS Control 12.6 to reduce CIS Control 3.4 risk.

Risk Analysis	Value
Risk Treatment Option	Reduce
Recommended Safeguard	(CIS Control 12.7) Acquire and implement an open-source IPS solution to detect, and alert on attacks on the enterprise management application, and other vulnerable

Risk Analysis	Value
	systems in the environment. After gaining confidence in the types of detected actions and alerts, deploy IPS capability to protect high-risk systems.
Safeguard Risk	Moderate cost in personnel time to implement and configure the IPS system.
Safeguard Threat Likelihood	3
Safeguard Mission Impact	1
Safeguard Obligations Impact	1
Safeguard Risk Score	3
Risk Acceptability	Acceptable

The assessor is certain that the impact to the mission and obligations will not be affected using the IPS as an IDS (in detect mode versus prevention mode). At this point, the risk assessor is confident that they have a plan for implementing a safeguard that aligns with CIS Control 12.6 that would reduce the risk that was evaluated by their review of CIS Control 3.4.

Risk Treatment Example 3 – CIS Control 14.9

Finally, the Tier 1 organization's risk analysis involving CIS Control 14.9 showed an unacceptable risk for how they logged events on the enterprise management application. Given their concern that disgruntled employees are expected to steal and abuse data in the environment, they realize they will need to track access to the application, and be alerted on specific actions, such as large data downloads.

The organization knows that log management systems and SIEMs can be readily implemented as services and management would be friendlier to detecting abuses of access than to invest in the more esoteric aspects of intrusion prevention. So they opt to recommend a commercial SIEM-as-a-Service solution to address this risk.

Table 30 - Example Risk Treatment Recommendation for CIS Control 14.9

Risk Analysis	Value
CIS Control	14.9
Description	Enforce detailed audit logging for access to sensitive data or changes to sensitive data.
Information Asset	Enterprise Management Application
Control	Access logs are captured and stored locally, and not reviewed.
Vulnerability	The organization is unaware of suspicious or inappropriate use.
Threat	Rogue employees or hackers using escalated privileges and may access and abuse nonpublic information in the application.

Risk Analysis	Value
Threat Likelihood	3
Mission Impact	2
Obligations Impact	3
Risk Score	9
Risk Acceptability	Unacceptable
Risk Treatment Option	Reduce
Recommended Safeguard	Implement a SIEM-as-a-Service. To prevent being overwhelmed by log messages and alerts, focus SIEM first on high-risk systems, such as the enterprise management application. Alert on any data manipulation and downloads conducted by administrator accounts.
Safeguard Risk	Initial tuning may be challenging, but will not interfere with our mission or obligations.
Safeguard Threat Likelihood	2
Safeguard Mission Impact	1
Safeguard Obligations Impact	1
Safeguard Risk Score	2
Risk Acceptability	Acceptable

The organization is now comfortable that their risk treatment recommendations are reasonable and appropriate, even though the plans do not include a complete implementation of all CIS Controls to all in-scope assets.

Moreover, the organization knows that they have a basis to explain and defend their plan to interested parties and authorities, and to defend the basis by which they accept certain risks.

Exercise:

The reader should use the template Risk Register – Tier 1 that is provided in the supplementary document *CIS_RAM_Workbook* to enter risk treatment recommendations for each risk that evaluated as unacceptably high.

The reader should consider:

1. Whether an existing safeguard can be improved, and how that would be done.
2. Whether a safeguard based on a different CIS Control would provide reasonable and appropriate risk.
3. Collaborating with information security subject matter experts to help model the potential effectiveness of recommended safeguards.

The risk assessor will need to use their professional judgment to design and recommend information security safeguards, and to evaluate prospectively the risk that they may pose.

Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Risk Treatment Recommendations Summary

Risk treatment recommendations are a critical part of risk assessments to ensure that the organization has developed a plan for addressing risks without creating other risks to the organization or its constituents. Some of the benefits that have been demonstrated about this process are:

1. Organizations can demonstrate to collaborating business managers how recommended security safeguards can be implemented without creating too much of a burden on the business mission.
2. Organizations can demonstrate to regulators and other legal authorities that safeguards are reasonable because the risk of the safeguard (the “burden” to the organization) is not greater than the risk that it is meant to reduce.
3. Organizations can demonstrate that recommended safeguards would be appropriate by showing that they would not foreseeably create an impact that would be intolerable to the organization or its constituents.

The process for evaluating risks and for recommending appropriate risk treatments has been demonstrated at a general level. However, some questions likely remain for the reader about evaluating safeguards, estimating likelihood, and the suitability of probability models in risk analysis. These more detailed topics will be discussed in the forthcoming chapter “Risk Analysis Techniques.”

Chapter 3: Asset-Based Risk Assessment Instructions for Tier 2 Organizations

Tier 2 risk assessment instructions are well-suited to organizations that fit the profile of Tier 2 organizations as described by the NIST Cybersecurity Framework. These organizations can be identified as having the following characteristics:

- **NIST Tier:** Tier 2 organizations. Tier 2 materials are best suited for organizations that have at least some collaboration with non-technical business management to define risk criteria.
- **Expertise:** The organization has resources and capabilities to analyze common security threats, and to plan risk-appropriate safeguards. However, they do not have on-hand skills to model how threats would operate within their organization.
- **Time:** The organization is able to invest sufficient time to analyze risks at the level of specific systems, devices, and applications, and sub-components within those assets.

This chapter is comprised of sections that each address a specific activity within a risk assessment. Readers should engage this chapter by first reading the text in each section, and then conducting the exercises that are recommended for each section. The material presented in the CIS RAM is substantially different from many other risk assessment standards and models, so the reader should first understand the aim of each section, and then practice what they learn using templates that are provided in the supplementary document *CIS_RAM_Workbook*.

While conducting their first CIS RAM-based risk assessment, organizations should be careful to not try to “boil the ocean.” Regulatory bodies and information security standards alike understand that not all risks can be identified in a single assessment. Organizations should continuously and regularly assess risks to identify, understand, and manage risks over time.

The Risk Assessment Project

Overview

Risk assessments are projects with clear steps for preparing, conducting, and reporting risk analysis. And while risk assessment projects can be modeled with a typical plan, each organization’s project approach will vary depending on factors such as resource availability, and will develop over time as organizations become more capable in their cybersecurity maturity. This section will describe a risk assessment project, its components and variations, and will present guidance for preparing the plan.

The Project Outline

Risk assessments are conducted using a series of steps that include typical actions, and roles as shown in Table 31.

Table 31 – Example Risk Assessment Project Outline

Step	Task	Key Roles
1	Defining the Scope & Scheduling Sessions	Executives, Management, Assessor
2	Defining Risk Assessment Criteria	Management, Assessor
3	Defining Risk Acceptance Criteria	Executives, Management, Assessor
4	Risk Assessment (Asset-Based)	
4.1	Gather Evidence	Personnel, Management, Assessor
4.2	Model the Threats	Personnel, Management, Assessor

Step	Task	Key Roles
4.3	Risk Evaluation	Assessor
5	Propose Safeguards	
5.1	Evaluate Proposed Safeguards	Assessor, Management

During **Step 1** (Defining the Scope & Scheduling Sessions), the organization will determine which information assets to include in their evaluation. They will also identify business owners and technical stewards who will provide evidence and interviews to assess those assets. The risk assessor will then schedule interview sessions with those owners and stewards.

In **Step 2** (Defining Risk Assessment Criteria) the organization will define the rules by which they assess and score risks. They will define their mission (the value they bring to others), their objectives (their organizational definitions for success and failure), and their obligations (the potential for harm against others) to establish what they are trying to protect. They will then define scoring schemas to be used for impact and likelihood estimation.

In **Step 3** (Defining Risk Acceptance Criteria) the organization will establish their risk tolerance by selecting a combination of the likelihood of an impact that would be tolerable to all parties (the organization and parties that may be harmed by realized risks).

In **Step 4** (Risk Assessment – Asset-Based) the risk assessor will evaluate the risks of the information assets. For Tier 2 organizations, the analysis includes the following activities:

- “Gather Evidence” involves a review of documents, such as policies, procedures, standards, and benchmarks. It also includes interviews with management and personnel. Evidence gathering also entails observation of configurations, artifacts, facilities, records, and work processes to determine whether they operate in secure or vulnerable ways. The levels of evidence gathering will be directly associated with the maturity of the organization.
- In the “Model the Threats” activity, the organization will model risks with at least these components; CIS Controls that should be in place to protect information assets, safeguards that effectively protect information assets as described by the CIS Controls, and vulnerabilities that result if safeguards are not sufficiently effective. The order in which these components are considered depends on the maturity of the organization, as will be described later in this document.
- During “Risk Evaluation” the organization will estimate the likelihood and impact of the risks. The estimates will be based on the scoring and criteria that were established in Step 2. The risk score will be automatically calculated to determine whether the current implementations of CIS Controls are already reasonable.

During **Step 5** (Propose Safeguards) the organization will consider how to address unreasonable risks by selecting CIS Controls that should be implemented to address each risk, and specifically how the controls will be implemented. These safeguards may include security devices, physical safeguards, training, oversight processes, or other methods. The risk assessor then will test the reasonableness of the safeguards during “Evaluate Proposed Safeguards.” The risk assessor will evaluate the proposed safeguards using the same criteria that were used to evaluate the risks.

A project plan template is available in the supplementary document *CIS_RAM_Workbook*.

Defining the Scope & Scheduling Sessions

Note: The reader should use the worksheets provided in the supplementary document CIS_RAM_Workbook while reading these instructions. The reader will best understand the concepts and their use by practicing the methods described in this chapter.

Defining the Scope

Organizations should conduct risk assessments against a clearly defined scope of information assets. A single theme usually bounds the scope of assets, such as “information assets that contain sensitive information,” “the data center,” “engineering practice areas and technologies that support them,” or a specific business division.

While it is possible to select unrelated information assets for an assessment – or a subset of assets within a larger scope – the organization that receives the assessment and is making investments and prioritization decisions based on its findings will be most comfortable when the information assets are associated with a business entity or a business process. Otherwise, risk assessment findings may seem scattered and unrelated.

Similarly, while assessing the risk of a set of information assets, it makes good sense to consider a set of assets that can directly affect each other’s security. For example, a risk assessment that examines a set of applications should also include the network devices that connect the applications to other assets and other networks, as well as the processes that are used to develop and manage those applications. These systems are directly connected to each other and dependent on each other so their risks are easily associated with one another.

Organizations cannot examine all information assets comprehensively in a single risk assessment, so their scope should consider the time and resources that are available for the assessment. Risk assessors should consult security experts to help them determine which assets, threats, and risks to prioritize.

An example scoping table (Table 32) demonstrates the level of detail that is appropriate for an initial assessment plan and is provided in the workbook *CIS_RAM_Workbook*.

Table 32 - Example Scoping Table

Asset Type	Asset Name	Business Owner	Steward
Information	Application Code	COO	CIO
Information	Patient Information	Customer Experience	CIO
Application	Patient Record (prod)	Customer Experience	Product Manager
Application	Patent Record (dev)	Customer Experience	Software Development
Application	DataMart	Innovations Dept	DevOps
Server	ProductionAppSrvr1	Customer Experience	DevOps
Server	ProuctionDBServer2	Customer Experience	DevOps
Server	DevAppSrvr1	Software Development	DevOps
Server	DevDBServer2	Software Development	DevOps
Server	LDAP1	CIO	DevOps
Server	DNS1	CIO	DevOps
Network Device	Core Router	CIO	Network Engineering
Network Device	DMZ Router	CIO	Network Engineering
Network Device	Firewall 1	CIO	Network Engineering
Network Device	Firewall 2	CIO	Network Engineering
Network Device	Switch	CIO	Network Engineering
Process	AppDev	Customer Experience	Software Development
Process	Code Promotion	Customer Experience	DevOps
Process	Maintenance	Product Manager	DevOps

Asset Type	Asset Name	Business Owner	Steward
Process	Change Management	Product Manager	DevOps
Process	Vulnerability Mgt	CIO	Security Team
Process	Account Setup	Customer Experience	Application Management
Process	Account Maintenance	Customer Experience	Application Management
Process	New Client Onboarding	Customer Experience	Application Management
Process	Internal Audit	Compliance	Internal Audit
Process	Device/System set-up	CIO	DevOps
Process	Customer Support	Customer Experience	Application Management

Note that the scoping table includes the business owner roles and steward roles. Business owners are the (typically) non-technical managers who are responsible for the information and processes that information assets support. Stewards are (typically) technical managers who are responsible for the functionality and security of the information assets. By identifying information asset ownership up front, the scoping table can be used to help plan interview sessions for the remainder of the risk assessment.

Regardless of how the scope of the risk assessment is established, and how detailed the asset listing is, there are a few helpful practices an organization should keep in mind as they identify their information assets:

- Think of a set of similarly-situated assets as a single asset class. For example, all database servers that use the same technology and the same maintenance and administration methods may be considered one asset class. However, if one set of database servers is different from others (for example, they hold sensitive information in a DMZ while others process less-sensitive information in another zone), these may be considered two assets because their risks will be different, even if they are managed identically.
- Information assets are not only technologies that store and transmit sensitive information. Information assets are any information, technology, process, people, or facilities that may impact the confidentiality, integrity, or availability of information.
- Include in the scope all information assets that are within the same zones (networks, facilities, etc.) as any other in-scope information asset.

The reader should develop their own scoping table using the scoping table template in the supplementary document *CIS_RAM_Workbook*.

Exercise:

The reader should develop their own scoping table using the “Scope - Tier 2” worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. A set of information assets your organization is interested in focusing your security resources on?
 - a. This set can be defined by processes, technologies, a class of information, or a location.
2. What are the boundaries between this set of information assets and other information assets not in this scope?
 - a. What systems, facilities, and network devices link these boundaries together, or separate them?
 - b. Are these “boundary” assets included or excluded from the scope?
3. List information assets or asset classes that are within the scope.
 - a. List information assets or asset classes to a level of detail that the organization has the time and resources to analyze. This may require adjustment during the course of the assessment if the organization realizes it has more time (or less time) than they originally planned to assess the information assets.

Scheduling Interview Sessions

Interview sessions will be topical, and should address one topic or closely related topics for each conversation. Interview sessions can focus on CIS Controls, or information assets and asset classes.

For example, interview sessions that focus on CIS Controls would bring together the personnel and management who know how each control is implemented and operates. One session may be dedicated to CIS Control 1 to understand how devices are inventoried. Another may be scheduled to discuss CIS Control 2 to understand what safeguards are in place to inventory software. Or, if the same personnel are knowledgeable about both safeguards, then perhaps one session could combine both topics.

Similarly, if risk assessors schedule sessions around information assets or asset classes, then it would be appropriate to include business owners and technical owners of those systems to understand how associated CIS Controls are applied to each asset or asset class.

An example session for a web application may include product managers, application developers, application administrators, and business owners. Topics in that session may include CIS Control 14, “Controlled Access Based on the Need to Know,” CIS Control 16, “Account Monitoring and Control,” and CIS Control 18, “Application Software Security.”

How the organization groups and orders these topics is largely up to them, but risk assessors should consider these pointers while scheduling interview sessions:

1. Be respectful of people’s time. While it is important to gather comprehensive information about risks, organizations cannot gather all relevant information in the first one or two risk assessments.
2. Work with managers to determine the most efficient and useful way to schedule interviews, whether by CIS Controls or by information assets and asset classes.
3. Expect that some security safeguards will be applied differently to different information assets and asset classes. For example, CIS Control 5, “Secure Configuration for

Hardware and Software on Mobile Devices, Laptops, Workstations and Servers,” and CIS Control 16, “Account Monitoring and Control” may be implemented and overseen differently for servers in different environments, or may be centrally controlled for servers, but individually controlled for network devices. Plan on evaluating how these and similar safeguards are applied to different asset classes.

4. Provide an agenda for each interview to help participants prepare any materials or information regarding the information assets or controls that may be discussed.

Scheduling Evidence Reviews

Risk assessors use evidence review sessions to examine information assets and; determine whether they conform to CIS Controls, and evaluate whether they would be effective against foreseeable threats.

The evidence review sessions should be scheduled following the interviews so that risk assessors understand the general landscape of the security environment before trying to understand why certain assets are configured the way they are.

During interviews, the risk assessor will learn about topics that should be more closely understood through a review of configurations, system testing, or records review. Risk assessors should note during or soon after the interview which safeguards they will want to examine further, and inform the participants that they will likely be contacted later in the assessment to participate on those evidence review sessions. Further, the assessor should ask which personnel, processes, or information assets would be appropriate to examine to gather the appropriate evidence. The evidence review sessions can then be scheduled at the end of the interview.

Techniques for reviewing evidence of the effectiveness of controls are presented later in the chapter “Risk Analysis Techniques.”

Defining Risk Assessment Criteria

Introduction

Risk assessment criteria are the numerical and plain-language statements that an organization uses to evaluate their cybersecurity risk. The most familiar form of risk calculations, “Risk = Likelihood x Impact,” is the basis for risk analysis in the CIS RAM. But it is just the starting point for risk analysis.

Risk assessment criteria must be meaningful to the organizations that use them, so they must be tied to the potential benefit and harm that the organization may create. The impact of a cybersecurity breach may harm the organization itself, it may harm the organization’s ability to successfully achieve its mission, or it may harm others.

Because cybersecurity failures impact parties both inside and outside an organization, risk assessment criteria must be universally meaningful and must address the interests of all potentially affected parties. Additionally, risk assessment criteria must demonstrate to authorities, such as regulators and litigators, that the organization considers the risk of harm to others as much as the risk of harm to themselves.


While these requirements may seem complex, the method presented in this section will sufficiently address them while using a technique that is simple to develop and use.

Risk Assessment Criteria Foundations

The risk analysis provided in the CIS RAM is at its root a question of balance between the potential of future harm against the certain burden of a safeguard. Regulators and litigators have

long considered this balance as key to acting as a “reasonable person.” The core structure of a risk statement is provided below to illustrate the core concept of balance.

Figure 9 - Balance Within Core Risk Analysis



Notice a few things right away with the model risk analysis in Figure 9.

- While organizations typically evaluate the observed risk to determine whether they should address or accept it, this risk statement deliberately compares the observed risk to a proposed safeguard.
- The criteria that evaluates the risk also evaluates the safeguard.
- The impact of the risk estimates the potential of harm to the organization and the potential harm against others.

Risk assessors compare risks to their proposed safeguards to determine whether those safeguards would create a foreseeably lower risk than the current state. To accomplish this, the assessor evaluates the current state risk (or “observed risk”) and the proposed safeguard using the same criteria to ensure comparability.

This comparison prevents organizations from implementing safeguards that are overly burdensome, or that create new, unacceptable risks. For example, an organization that uses software that is no longer supported by the vendor, but relies on that software for critical business purposes, should find alternative methods for identifying and controlling potential security risks until they replace the software. If management recommends quickly changing out to inferior, but secure software, the organization may suffer a greater impact to their mission than the security risk they are trying to avoid.

While considering CIS Control 18: Application Software Security, a risk statement can be made to estimate the foreseeability of an impactful threat. The risk can be stated as it appears in Table 33 (where the risk score ‘12’ is a product of the likelihood ‘3’ and the highest impact score ‘4’):

Table 33 - Example Core Risk Statement

Observed risk	Likelihood	Impact to Us	Impact to Others	Risk Score
Hackers may exploit the unsupported, but critical application.	<u>3</u>	3	<u>4</u>	12

A risk assessor should then recommend and evaluate a safeguard to reduce the high security risk, as illustrated in Table 34. Here, the organization would realize that the likelihood of a negative impact to their mission is greater than the current state risk. This is an obvious case of the burden being greater than the risk, and a recommended safeguard being unreasonable.

Table 34 - Example Unreasonable Proposed Safeguard

Proposed Safeguard	New Risk	Likelihood	Impact to Us	Impact to Others	Safeguard Risk
Replace application with inferior, secure application.	Application will operate inefficiently.	<u>5</u>	<u>3</u>	1	15

When faced with this analysis, the organization must then find another way to address the risk. This process will be described later in this chapter in the section Risk Treatment Recommendations.

But what should be apparent is that without a definition of risk assessment criteria the likelihood and impact scores are not meaningful. What would impacts or likelihoods of '1', '2', '3', '4', or '5' mean, anyway? The organization will need to create definitions for their likelihood and impact scores so that they are meaningful to all interested parties, and so that they provide a consistent method for risk evaluation.

Impact Definitions

Tier 2 organizations generally benefit from more business involvement in managing cybersecurity risk than Tier 1 organizations. Because of that increased involvement, the risk assessment criteria can be – and should be – more explicit and detailed than those used by Tier 1 organizations. Advanced organizations can consider more nuance in terms of business impacts and tolerance, and can employ organizational objectives with more authority.

An impact definition for Tier 2 organizations can have at least three impact types, and five impact scores (magnitudes) like the definition depicted in Table 35.

Table 35 – Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objectives: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

Background – Impact Definitions

This document provides instructions for defining impacts and impact scores (magnitudes) in this section with more in-depth instructions and examples in the “Risk Analysis Techniques” chapter. The reader should understand before going further that organizations in most cases should not define impacts exclusively using financial values. While cost is a common and almost necessary consideration while evaluating risks and safeguards, if it is the only criterion, the organization will communicate to their personnel, as well as to interested parties and authorities, that cost is their only concern. The purpose that the organization serves and the harm that may befall others must be part of the evaluation if risk is to be responsibly tied to the potential of harm, and if the evaluation is to be understandable to regulators and legal authorities.

Organizations should also consider having more than three impact types in their impact definitions if they have more than one mission, multiple objectives, and many obligations that they need to consider in their risk analysis. While this expansion may create an increasingly wide risk register, it can help organizations feel comfortable all relevant interests were considered in their risk analysis.

Tier 2 organizations that previously used Tier 1 risk analysis processes may build on their simpler risk assessment criteria that used three levels of impact scoring. For the purposes of reference to our example organization – a health information provider – they have gone through a year or two of risk management and have gained the attention and confidence of business managers and executives. As a result, their ability to assess cybersecurity risk using business criteria will also improve.

We can see by comparing the risk assessment criteria for Tier 1 organizations in Chapter 2 to Table 35 that the detailed descriptions of impacts have increased in two dimensions; the number of impact scores (magnitudes) increased from three to five, and there is an additional impact type for business objectives.

Tier 2 organizations will find that using a range of five impact scores (magnitudes) increases the utility of risk prioritization at the end of the risk assessment. A three-by-three risk assessment criteria model provides organizations with six possible risk scores; 1, 2, 3, 4, 6, and 9. This leads to a somewhat coarse grouping that may cause risks of somewhat different urgencies to be indistinguishable.

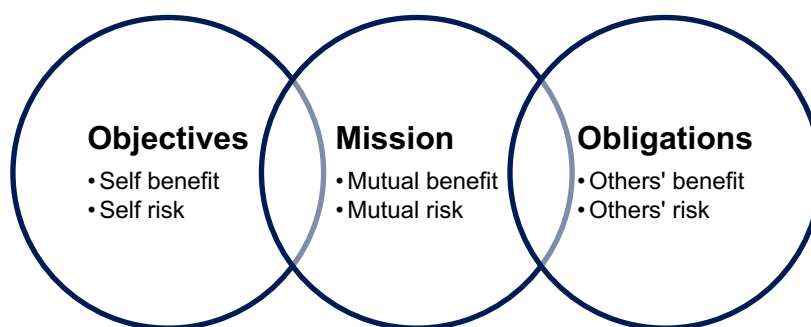
A five-by-five risk assessment criteria model allows for 14 possible risk scores of; 1, 2, 3, 4, 5, 6, 8, 9, 10, 12, 15, 16, 20, and 25. Now risks of somewhat different urgencies will likely be classified into different risk scores and will be more easily distinguished while prioritizing them.

Also note that the impact scores of ‘1’ and ‘2’ in Table 35 are shaded grey to separate them from the higher scores. Scores ‘1’ and ‘2’ describe impacts that would be generally thought of as acceptable. The risk acceptance criteria process will be explained later in the document, but it is useful to consider now that the scores ‘1’ and ‘2’ are consistent in their definition of impact scores across all three impact types, and that the impact scores of ‘3’, ‘4’, or ‘5’ could consistently be thought of as unacceptably high for all three impact types.

The example health information provider also has a new impact to consider in Table 35 to help them include their business objectives in their risk analysis. Business objectives are more self-focused than missions and obligations, and are aligned with success criteria commonly found in business. Some examples include profitability, growth, maintaining accreditations, customer satisfaction, or retaining a position in the marketplace.

Objectives align most directly with what is commonly thought of as the “cost” of a safeguard. But rather than allowing an organization to arbitrarily decide that a safeguard costs too much, this method of including cost in terms of impacts to objectives forces the organization to evaluate why a cost would be excessive. Does the cost of the safeguard impede profitability goals? Does the safeguard limit efficiency or growth? Those are certainly reasonable concerns, as long as the profitability goals are in parity with the mission and obligations impacts. *In other words, an organization should not let profitability be more important than harm to others, or harm to their ability to fulfill their mission. See “Note Regarding Use of Financial Costs as Objectives” in Chapter 5.*

Figure 10 - Objectives, Mission, Obligations



Organizations are well-served with this model because business management, technicians, compliance personnel, and legal counsel all have their interests addressed in risk analysis that uses these criteria.

An in-depth explanation of how to develop impact definitions with multiple examples is provided in the chapter “Risk Analysis Techniques.”

Likelihood Definitions

The likelihood definition for a Tier 2 organization should also increase in nuance from the simpler Tier 1 definition, and can do so by adding two more scores to the table as shown in Table 36.

Table 36 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

- “Not Foreseeable” implies that a threat is not plausible in the environment that is being assessed. Loss of portable media may not be foreseeable during a risk assessment of a hosted application.

- “Foreseeable” implies something that is plausible, but the organization would be surprised if it occurred. A founding executive taking copies of sensitive data to competitors may be considered foreseeable, even if it is not expected.
- “Expected” implies a threat that is not common, but that would eventually happen. Phishing attacks or other social engineering attacks may be expected in many environments.
- “Common” implies something that happens repeatedly, such as mis-addressed emails with sensitive information, malware attacks, or loss of laptops and mobile devices.
- “Current” implies threats that are rarely not present, such as port scanning on perimeter devices, or sharing of information in quasi-public spaces such as pharmacy counters or bank tellers.

When risk assessors estimate the likelihood of a threat, they will select scores ‘1’, ‘2’, ‘3’, ‘4’, or ‘5’ using the foreseeability definition as their guidance. Organizations may add time-based limits in their foreseeability definitions (i.e. “Foreseeable within planning thresholds,” “Expected within the five-year plan” or “Not foreseeable in the next fiscal year”). If organizations do introduce time limits in their likelihood definitions they should prioritize risk treatment investments to meet these timelines. That may be excessively challenging to many organizations, so they should proceed with caution.

Developing the Risk Assessment Criteria

Because risk assessment criteria are meant to describe risk as it applies to the organization that owns the risk, it is appropriate for the most senior management who are responsible for the mission, objectives, and obligations to participate in developing and accepting the criteria.

Table 37 lists roles commonly involved in risk assessment criteria development, and the interested perspective they bring to the definition effort.

Table 37 - Roles Involved in Defining Risk Assessment Criteria

Role	Perspective
Chief Executive Officer Chief Operations Officer	To ensure that the mission, objectives, and obligations of the organization are appropriately defined, and to ensure that a distinction between acceptable and unacceptable impacts are appropriately delineated.
Chief Compliance Officer	To ensure that the interests of regulatory agencies are appropriately included in risk definitions.
Chief Financial Officer	To ensure that objectives are appropriately defined, particularly the distinction between acceptable and unacceptable impacts.
Chief Information Officer Chief Technology Officer	To ensure that technical performance, service, and capabilities are considered, and to include all types of information processes beyond technology.
General Counsel Outside Counsel	To ensure that obligations are appropriately defined and that they compare well with the mission and objectives.
Internal Audit Audit Committee	To ensure that the concerns of interested parties are well represented in all impact definitions and scores.
Key Customers / Clients Key Constituents	To ensure that their interests are included in the obligations definition.

Exercise:

The reader should develop their organization's risk assessment criteria using the "Criteria - Tier 2" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Developing the risk assessment criteria in collaboration with a business managers and legal counsel to ensure that the Mission, Objectives, and Obligations definitions are sensible to the organization.
2. Working with legal counsel to help ensure that impact definitions appropriately address the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.
3. Referring to guidance for defining and scoring impact types in the "Risk Analysis Techniques" chapter.

The risk assessor will need to use their professional judgment to define impact types, and to describe levels of impact that the organization must manage to. Because risk assessment criteria are a declaration by the organization of what they will manage to in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

Defining Risk Acceptance Criteria

Introduction

Because risk assessments are essentially questions of balance, the criteria for accepting risk should help determine whether balance was achieved. In CIS RAM risk acceptance has two components:

- Appropriate risk: That the likelihood of an impact must be acceptable to all foreseeably affected parties
- Reasonable risk: That the risk posed by a safeguard must be less than or equal to the risk it protects against.

While these components have been demonstrated briefly in Chapter 1, the "appropriate risk" will be described in more detail in this section. "Reasonable risk" will be described later in the Risk Treatment Recommendations section further on.

After establishing the impact and likelihood definitions, Tier 2 organizations are now well positioned to state their risk acceptance criteria. Recall that impacts were defined within impact scores that ranged from '1' to '5'. Acceptable impact scores '1' and '2' were defined in a manner that would appear appropriate to interested parties (and are shaded grey to indicate their acceptability), and the impact score '3' was the lowest unacceptable score.

Table 38 – Unacceptable Impacts

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objective: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

And similarly, likelihood scores were within a range of '1' to '5' as below. Once our example organization develops their risk management maturity and are ready to refine their risk distinctions, they may decide to not tolerate unacceptable impacts if they are *foreseeable but not expected* ('2'), or if they are *expected to occur* ('3'). This would be a decision best made by their executives, and especially their compliance team, general counsel, and interested parties. But in this case, the model will assume that they selected an unacceptable likelihood score of '3'.

Table 39 – Unacceptable Likelihood

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

So Tier 2 organizations would define their risk acceptance like this:

Table 40 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
3	x	3	=	9
... therefore ...				
Acceptable Risk				< 9

An example of the heat map for this assessment criteria is shown in Figure 11. While risk heat maps are not used in CIS RAM, organizations can now design heat maps that represent actual risk acceptability based on organizational requirements, and a duty of care to others.

Figure 11 - Example Heat Map

		Impact				
		1	2	3	4	5
Likelihood	5	Green	Yellow	Yellow	Red	Red
	4	Green	Green	Yellow	Yellow	Red
	3	Green	Green	Yellow	Yellow	Yellow
	2	Green	Green	Green	Green	Yellow
	1	Green	Green	Green	Green	Green

Exercise:

The reader should define their organization's risk acceptance criteria using the "Criteria - Tier 2" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Working with a business management sponsor who can help ensure that the risk acceptance criteria are sensible to the organization.
2. Working with legal counsel to help ensure that the definition for risk acceptance addresses the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.

The risk assessor will need to use their professional judgment to identify levels of acceptable risk. Because risk acceptance criteria are a declaration by the organization of what they will tolerate in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

As organizations assess their risk using the CIS Controls V7 (modeled in the following section) they will be able to automatically determine whether the risk evaluates as acceptable or not without needing to consider the question differently in each case. A simple estimation of the likelihood and impact will automatically determine how the organization should prioritize each risk, and whether the organization can safely accept the risk as appropriate.

An Asset-Based Risk Assessment Process

Introduction

Tier 2 organizations enjoy collaborative relationships with non-technical management in managing security risk. They also have the knowledge and experience to plausibly model security threats against information assets. In this section, the example organization has become more capable after a year of risk management and has earned a partnering relationship with non-technical departments. They have learned that information assets should be analyzed based on

their situation and context, rather than as a single class of technical assets (a process used for Tier 1 organizations to simplify their risk assessment efforts).

As a Tier 2 organization, their risk assessments will first consider the information assets that are in scope of their assessment. As they consider those assets, they will think through the CIS Controls that are appropriate for protecting those assets, they will consider vulnerabilities to those controls, and they will model foreseeable threats that can take advantage of the vulnerabilities.

This order of analysis provides a distinct benefit to organizations that have a more mature understanding of security risks and safeguards. It allows organizations to evaluate potential harm that can come to specific assets and the CIS Controls that should protect them. Only then does it evaluate the risks.

Organizations that model risks using CIS Controls V7 may find it challenging to consider how one risk is influenced by other risks related to an asset. For example, if an FTP server allows third-party employees to access it using shared credentials, then those third-party employees may retain access to the FTP server after leaving their employer and raising a risk associated with CIS Control 16.7 which recommends a process for revoking access upon termination. That would likely appear as a high risk to the organization. But if the risk assessor also determines that the FTP service allows only write capabilities (CIS Control 14.7), the risk associated with CIS Control 16.7 is likely reduced. Control-based risk assessments such as those used by Tier 1 organizations may miss that relationship and as a result may exaggerate the risk involved in CIS Control 16.3.

In addition to risk analysis taking this slightly different order in their threat modeling, Tier 2 organizations will also estimate risk with more complex risk criteria. Because of their relationship with business management, Tier 2 organizations will receive some pressure from business managers to refine their prioritization of risks, and to evaluate risk using more explicitly business-based impact criteria.

This section will describe and work through a Tier 2 organization's risk assessment using the CIS Controls, and will show how to evaluate risks using the risk assessment criteria that are appropriate for Tier 2 organizations.

This section of the document will describe a risk register that is available as a template in the supplementary document *CIS_RAM_Workbook*.

The Risk Register

As the Tier 2 organization prepares their risk assessment, they will assemble their list of information assets and align them with the CIS Controls that are appropriate for protecting those information assets. This section will demonstrate this process using the risk register template provided in the supplementary document *CIS_RAM_Workbook* for Tier 2 organizations.

The layout map of a risk register for a Tier 2 organization is depicted in Figure 12.

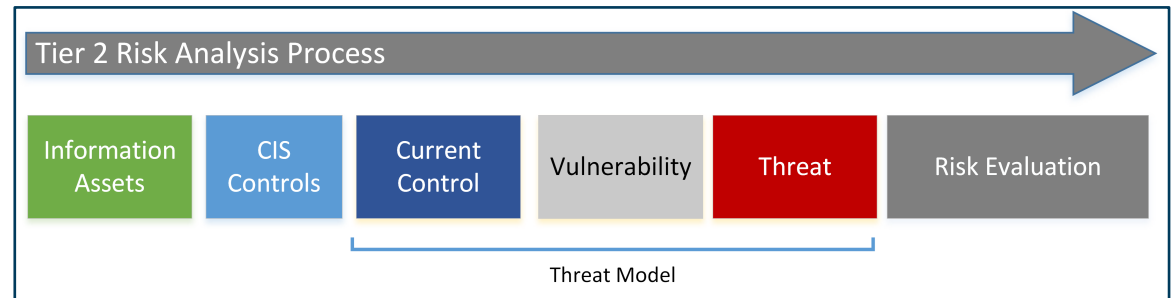
[illegible]

- A. The column headers and guiding text help the reader or risk assessor understand the information that is contained in the column.
- B. The information assets identify the items and processes that are being analyzed.
- C. The CIS Controls help the risk assessor consider controls that should be in place to protect information assets.
- D. The “threat model” includes the following items:
 - a. A description of how the CIS control is implemented in the environment.
 - b. The vulnerability that may exist if the control is not fully implemented.
 - c. The threat that may compromise the asset because of the vulnerability.
- E. The risk evaluation estimates the likelihood that the threat would succeed, and the impacts to the mission, objectives, and obligations if it did. The evaluation includes the resulting risk score, calculated as a product of the likelihood and the highest of the three impact scores.
- F. Risk treatments are recommended for risks that are evaluated as unacceptably high. Safeguards that are based on the CIS Controls are described, and they are in turn evaluated for the risk that they may pose to the mission and objectives. A “safeguard risk” score is calculated which should be lower than the risk acceptance criteria, and the risk that it is meant to address.

The Tier 2 risk assessor will analyze risks using an asset-based approach, starting the assessment by considering the information assets that they intend to protect. Because this process will invariably create more detail in the risk register (in the form of multiple rows on the risk register for each asset and each applicable CIS Control) organizations may wish to start their risk analysis using the Tier 1, control-based method. This method first analyzes how the CIS Controls are generally applied to the environment. The Tier 2 risk assessor can then examine the information assets that are protected differently from the standard practice in a separate row.

The asset-based risk assessment process is illustrated in Figure 13.

Figure 13 - Risk Analysis Process for Tier 2 organizations



The asset-based risk assessment process for Tier 2 organizations is meant to ensure that each information asset is protected by CIS Controls in ways that are appropriate for their particular risk. This asset-based analysis is accomplished by following the steps listed below:

1. Select information assets or asset classes that are listed in the asset inventory. Record the selected assets in the “Information Asset” cell in that row.
 - a. If the Tier 2 organization has already completed a risk register that evaluated all of the CIS Controls as they are applied generally to the organization, then the risk analysis may start by selecting assets that are listed in the risk register.
 - b. Organizations may wish to evaluate asset classes rather than individual assets for efficiency. As an example, there are often a set of similar technologies that are identically configured and managed. In such circumstances, listing these as asset classes is appropriate because each of these items will be similar to the set. Unique items like a core router, a load balancer, a single instance of an application, or a one-off operating system should be analyzed on its own.
2. Model threats using the following approach:
 - a. Consider each information asset or asset class one-by-one.
 - b. Pair each information asset or asset class to the CIS Controls that are appropriate to protect them. This is aided by the “Asset Type” categorization of CIS Controls V7. For example, user desktops, application servers, multi-function printers, and tablet computers are systems that can be paired with CIS Controls that are categorized for “Systems.”
 - c. Add one row to the risk register for each pairing of a CIS Control and the information asset or asset class.
 - d. Note: It will likely be too time-consuming to evaluate all suitable pairings between information assets and applicable CIS Controls. While such analysis would be optimal, it is more fitting for organizations to prioritize the pairing of information assets and asset classes with the top five CIS Controls, and other controls that may be of concern to the organization. Risk management is a continuous cycle that will allow organizations to address more risks over time as their security program matures.
3. Gather evidence for how well each asset is protected by its paired CIS Control.
 - a. Evidence may be in the form of interviews, a review of configurations, a vulnerability test, a penetration test, an evaluation of a system or device using SCAP policies, or a review of evidence such as records and logs.
4. Describe how the control is applied to the information asset or asset class in the “Current Control” cell of that row.

5. Consider the difference between the CIS Control and the current control and determine whether there is a deficiency in how the control is currently deployed and operating. If the current control is not implemented as described, how would this be described as a vulnerability?
 - a. Consider the objective of the CIS Control. For example, CIS Control 10.3 states “Test data integrity on backup media on a regular basis by performing a data restoration process to ensure that the backup is properly working.” Its objective is to ensure that backup data on storage media is retrievable when it is needed. If the current control does not meet the objective, then state the gap as a vulnerability in the vulnerability cell, such as: “We are not confident that our backup data is retrievable from backup media.”
6. Now consider the threat that could occur because of the vulnerability.
 - a. The above vulnerability could be paired with the threat: “System failures may result in a recovery of data as old as one month, or may require manual entry from paper documents.”
7. Next, estimate the likelihood that the threat would succeed, and the impact it may create.
 - a. Likelihood estimation can be challenging at first, but the risk assessment criteria provide some guidance in that estimation process. Guidance is also provided in the “Risk Analysis Techniques” chapter.
 - b. Impact scores should provide estimates of the impact that such a threat would create. Consider the likelihood and impact scores as a pair. In other words, “What is the likelihood that this impact would result?” Examples below will provide further guidance.
8. The risk score will be automatically computed by multiplying the likelihood score by the highest of the three impact scores.

Tier 2 Risk Assessment Example 1 – Risk-Based Business Case

While conducting their risk assessment, the Tier 2 organization comes across a common question that organizations deal with; when does a business reason justify exceptions to security policies? This question arises as they consider CIS Control 15.9 which reads, “Disable wireless peripheral access of devices (such as Bluetooth and NFC), unless such access is required for a business purpose.”

Organizations commonly accept risks that come with policy exceptions, and do so by documenting that risk acceptance. But if they make these decisions without careful, consistent analysis in consideration of foreseeable impacts to themselves and others, then their risk acceptance is itself risky.

In the case of CIS Control 15.9, the organization uses Bluetooth-accessible systems in regional clinics to support their patients. Patients carry electronic “diary devices” that monitor their health status and store health data as diaries. Bluetooth is used in clinics to connect diary device controllers to the diary devices to read them for diagnostics, to receive data from the devices, and to transmit updated firmware to those devices. This is undoubtedly worthy of a documented business need. But is it still unreasonably risky to do so?

Using the risk register template for Tier 2 organizations that is provided in the document *CIS_RAM_Workbook*, the organization first lists the information asset that they are evaluating.

Table 41 - Information Asset Example

Information Asset
Diary device controllers

Then they come across the CIS Control that raises the question about business-acceptable risk.

Table 42 - CIS Control Example

CIS Control	Description
15.9	Disable wireless peripheral access of devices (such as Bluetooth and NFC), unless such access is required for a business purpose.

The risk assessor must then think through and record the threat model for this risk. Recall that the threat model considers their current control, what resulting vulnerabilities may exist, and what threats would be of concern to them. While the diary device controllers must pair with the diary devices, the organization realizes that their Bluetooth-enabled diary device controllers and the files located on them are likely accessible by attackers using readily available attacks methods.

So the next three columns of the risk register would look like this:

Table 43 - Threat Model Example

Control	Vulnerability	Threat
Each diary device is joined to the diary device controller using a one-time, six-digit code that is displayed on the controller and entered at the device. At this point, all file transfers and firmware updates between devices are enabled.	Diary device controllers are using a deprecated version of Bluetooth to support older diary devices. Bluetooth devices can manipulate Bluetooth services on the diary device controllers to gain access to files and commands on the controllers.	Hackers may walk through clinics with Bluetooth devices that are prepared to hack diary device controllers using attacks such as Blueborne, and may access hundreds of patient data files, as well as firmware.

But this may not be the only risk that can be considered in this scenario. Another plausible risk could be that the hacker could put compromised firmware on the diary device controllers to allow them to control the devices after firmware upgrades. Denial of Service attacks may also happen. Man-in-the-Middle attacks are also possible. Risk assessors may be concerned that an endless number of threat models could be created for any pairing between information assets and CIS Controls, making the risk assessment exercise never-ending. Of course, the risk assessor should limit the amount of theorizing they do while modeling threats, focusing primarily on the most plausible threat pairings given the security environment. *As a Tier 2 organization, they should rely on capable personnel and resources that help them focus on the most plausible threats for their environment.*

The threat model is now clearly stated and makes it possible for the organization to estimate the likelihood and impact of the scenario. Recall the definitions for the impact and likelihood scores that the Tier 2 organization created. Impact scores are restated in Table 44.

Table 44 – Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objective: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

Also recall that impact definitions for Tier 2 organizations include criteria for the organization's objectives because those organizations generally benefit from collaboration with business management who are invested in the success of the information security program. These managers often bring to the discussion the organization's strategic and tactical goals for success. But also note that this impact definition contains five magnitudes of impact. Five impact scores help Tier 2 organizations refine their impact estimates in more tangible terms than tables with three scoring levels, and help them refine their risk scoring to better distinguish between risks of varying priority. Acceptable impact scores of '1' and '2' are shaded to set them apart from higher, unacceptable impact scores.

Likelihoods were similarly defined with five potential scores for similar reasons, as shown in Table 45.

Table 45 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

The organization believes that the threat model they documented above – that hackers could hack into diary device controllers using something similar to a Blueborne attack - is foreseeable, and perhaps may be expected to occur. While the scenario would likely not be expected for most organizations, our example organization operates in environments where competitors and

adversarial states have an active interest in compromising their systems, and have proven their capability in the past. So they decide that their likelihood score for this risk will be '3'.

In that scenario, they also expect that their mission would be affected to the point where some (not many) patients would lose their ability to get a functioning diary device, and would therefore not access the information they need to maintain good health outcomes. They select a mission impact of '3'.

Such an occurrence would cause the organization to massively and immediately re-invest in their clinical infrastructure; an investment that they are not prepared to make now and that could take more than a year to recover from. That would create an objectives impact of '4'.

And finally, they believe that each diary device controller would contain no more than 100 patient records at any one time, given their work routines and service schedules. The patient information could be used to harm patients reputationally if hackers knew how to trace the patient IDs in each record to each patient's health conditions, and then acted maliciously with that information. This is not a plausible impact, so they select an obligations impact of '2'.

A risk that is expected to occur (likelihood = 3) in a way that prevents some patients from accessing information (mission impact = 3), that would take more than a year to recover from financially (objectives impact = 4), and that may cause concern but no harm to patients (obligations impact = 2) would appear as such in Table 46.

Table 46 - Example Risk Estimation

Threat Likelihood	Mission Impact	Objectives Impact	Obligations Impact	Risk Score
3	3	4	2	12

The risk score is the product of the likelihood score and the higher of the three impact scores, which in this case is '3 x 4 = 12'.

Also recall that the risk acceptance criteria for the Tier 2 organization looks like this:

Table 47 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
3	x	3	=	9
... therefore ...				
Acceptable Risk				< 9

An acceptable risk would be one that evaluates to anything below '9'. But the risk for how the organization protects their diary device controllers using their implementation of CIS Control 15.9 is '12' and is unacceptably high.

This risk analysis is demonstrated in Table 48 in a single table to bring all of these elements together. For document display purposes, this one risk analysis is shown in vertical format rather than horizontal as it would appear in a risk register. The examples that are described in this section are contained in the workbook *CIS_RAM_Workbook*.

Table 48 - Example Risk Analysis for Devices Protected by CIS Control 15.9

Risk Analysis	Value
Information Asset	Diary device controllers
CIS Control	15.9
Description	Disable wireless peripheral access of devices (such as Bluetooth and NFC), unless such access is required for a business purpose.
Control	Each diary device is joined to the diary device controller using a one-time, six-digit code that is displayed on the controller and entered at the device. At this point, all file transfers and firmware updates between devices are enabled.
Vulnerability	Diary device controllers are using a deprecated version of Bluetooth to support older diary devices. Bluetooth devices can manipulate Bluetooth services on the diary device controllers to gain access to files and commands on the controllers.
Threat	Hackers may walk through clinics with Bluetooth devices that are prepared to hack diary device controllers using attacks such as Blueborne, and may access hundreds of patient data files, as well as firmware.
Threat Likelihood	3
Mission Impact	3
Objectives Impact	4
Obligations Impact	2
Risk Score	12
Risk Acceptability	Not Acceptable

So while there is a documented business need for these devices to operate Bluetooth services, the risk of doing so with this device is still inappropriately high (since risk acceptance criteria is less than '9' and the observed risk score is '12'). The organization will want to address this with a risk treatment safeguard, which will be demonstrated in the Risk Treatment Recommendations section further in this chapter.

Because the risk assessment process for Tier 2 organizations starts at the asset, though, we do have other opportunities to consider how other controls may protect the asset in a manner that may reduce the risk we just observed.

We can examine how this process provides the Tier 2 organization with deeper risk insight with the next two examples.

Tier 2 Risk Assessment Example 2 – Risk Reduction Through Related Controls

The Tier 2 organization is considering risks to its information assets by pairing the asset with CIS Controls that are appropriate for protecting it. CIS Controls V7 assists in this pairing by providing a classification scheme for the controls called "Asset Type." If a diary device controller is a system that has network connectivity, then the Tier 2 organization's risk assessor can look through the CIS Controls that are associated with "System" and "Network" families to identify other controls that would be appropriate to protect the controllers.

They have already evaluated that CIS Control 15.9 does not protect Bluetooth-connected diary device controllers well enough. But they are not without alternatives. The organization thinks through other controls that they use to protect their diary devices and come across CIS Control 16.3 which reads, “Require multi-factor authentication for all user accounts, on all systems, whether managed onsite or by a third-party provider.”

This is interesting to them, because the diary devices use “soft tokens” in the form of certificates that are intended to track the diary devices for inventory purposes. But the certificates are very robust, and are tied to the file access privileges on the device controllers. Could this implementation of CIS Control 16.3 pose an acceptable risk to the diary device controllers, and could they reduce the risk cited for CIS Control 15.9?

The Tier 2 organization’s risk assessor tests this idea in Table 49. Note how the control is now described using more detail than the first attempt at analyzing this risk.

Table 49 - Example Risk Analysis for Devices Protected by CIS Control 16.3

Risk Analysis	Value
Information Asset	Diary device controllers
CIS Control	16.3
Description	Require multi-factor authentication for all user accounts, on all systems, whether managed onsite or by a third-party provider.
Control	While diary devices can connect to diary device controllers over Bluetooth using a one-time, six-digit code, access to existing files with patient information on the controller is granted using the unique soft-cert on each diary device.
Vulnerability	Six-digit codes may be guessed, or soft-certs may be stolen from diary devices and stored on attacker systems.
Threat	Hackers must steal soft-certs from diary devices, then guess one-time six-digit codes to access patient files on diary device controllers.
Threat Likelihood	1
Mission Impact	3
Objectives Impact	4
Obligations Impact	2
Risk Score	4
Risk Acceptability	Acceptable

Given this method for two-factor authentication, the threat model for hackers acquiring patient records from diary device controllers is no longer foreseeable using the threat model that the Tier 2 organization evaluated. And if that’s the case for the two-factor authentication control, then it should also influence the risk associated with CIS Control 15.9 that was evaluated earlier. So the risk assessor re-evaluates that risk again while referring to the controls identified for CIS Control 16.3 to see if they understand the risk differently. Again, note that the risk assessor described the control with more detail than the first attempt, and added a condition to the threat against CIS Control 16.3.

Table 50 - Example Revised Risk Analysis for Devices Protected by CIS Control 15.9

Risk Analysis	Value
Information Asset	Diary device controllers
CIS Control	15.9
Description	Disable wireless peripheral access of devices (such as Bluetooth and NFC), unless such access is required for a business purpose.
Control	[Supplemented by CIS Control 16.3] Each diary device is joined to the diary device controller using a one-time, six-digit code that is displayed on the controller and entered at the device. At this point, all file transfers and firmware updates are enabled. However, files may only be accessed by devices that use soft-certs that are associated with access privileges on diary device controllers.
Vulnerability	Diary device controllers are using a deprecated version of Bluetooth to support older diary devices. Bluetooth devices with seized soft-certs can manipulate Bluetooth services on the diary device controllers to gain access to files and commands on the controllers.
Threat	Hackers may walk through clinics with Bluetooth devices that are prepared with device-specific soft-certs to hack diary device controllers using attacks such as Blueborne. Hackers must steal soft-certs from diary devices, then guess one-time six-digit codes to access patient files on diary device controllers.
Threat Likelihood	1
Mission Impact	3
Objectives Impact	4
Obligations Impact	2
Risk Score	4
Risk Acceptability	Acceptable

As expected, the risk associated with CIS Control 15.9 is reduced because the likelihood of its threat model is reduced when taking into account the implausibility of hackers accessing files on the diary device controllers. In fact, the likelihood of the threat scenario is so low that the risk is acceptable.

The asset-based approach to risk analysis clearly presents an advantage to organizations by providing a more comprehensive and more realistic picture of the actual risk that an information asset is exposed to.

Tier 2 Risk Assessment Example 3 – New Perspectives from Risk

Not all asset-based risk analyses will reduce risk estimations, of course. Some will highlight risks that organizations rarely consider. In this example, the Tier 2 organization paired the CIS Controls for Malware Defenses with the diary device controllers and realized they had a problem. CIS Control 8.1 states, “Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.”

But the diary device controller's vendor does not support or provide an antivirus application for their operating system – a customized Linux distribution. While the diary device controllers are “headless” systems, they have web admin applications that provide administrative functions to controller operators, and can be managed through terminal sessions over console ports. The organization does not want to violate their support agreement with the vendor by compiling and running an antivirus application on the controllers.

By running a risk analysis, the organization will determine whether the malware risk is high enough to require antivirus software. Their finding is illustrated in Table 51.

Table 51 - Example Risk Analysis for Devices Protected by CIS Control 8.1

Risk Analysis	Value
Information Asset	Diary device controllers
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Control	Anti-malware software is not permitted on the diary device controllers.
Vulnerability	<p>Vulnerabilities are limited because common vectors for receiving malware such as email clients and web browsers are not installed on the controllers. Attackers would need to download malware executables from the Internet using scripts or bash commands.</p> <p>Command line, by design, is only accessible over terminal connections to the console port.</p> <p>Bluetooth attacks may still permit malware executables to be uploaded to a file space associated with an anonymous account. The web admin application on each controller has been tested as vulnerable to arbitrary code execution, cross-site scripting, and other attacks.</p>
Threat	Hackers may implant malware on diary device controllers through Bluetooth misuse, and take advantage of web admin application vulnerabilities to execute files or initiate scripts.
Threat Likelihood	3
Mission Impact	3
Objectives Impact	4
Obligations Impact	3
Risk Score	12
Risk Acceptability	Unacceptable

This looks more serious to the organization than they would have originally thought. Still concerned about attacks they have seen in the past, the organization realizes it is as likely for them to be attacked through a web application vulnerability as it would be through a Bluetooth vulnerability. Implanting a usable virus through a Bluetooth exploit or arbitrary file upload would

require a hacker with patience and skill, so the likelihood score of '3' for "Expected" seems more suitable than a '4' for "Common."

The challenges the organization faces are multi-fold in this risk: They cannot install an anti-virus application, nor a more secure web application on the diary device controllers, but the risk of malware is clearly too high. They will evaluate possible safeguards while developing their risk treatment recommendations in the next step of their risk assessment process later in this chapter.

Exercise:

The reader should refer to the template Risk Register – Tier 2 that is provided in the supplementary document *CIS_RAM_Workbook*. They may use the risk register template to enter a set of risks that are associated with the CIS Controls and information assets that are in scope of their assessment.

While doing this exercise, the reader should consider:

1. Ensuring that all in-scope information assets or asset classes are assessed.
2. Not "boiling the ocean." Not all assets can be practically evaluated against all applicable CIS Controls in a single assessment. Prioritize controls that protect information assets that appear vulnerable, or that protect high-value systems and information.
3. Consider assessing all assets against the first five CIS Controls to address the most common causes of cybersecurity incidents.
4. Whether a control or information asset requires examination to understand its actual configuration and effectiveness.
5. Whether the organization can tolerate the amount of effort and time that the risk assessment requires.
 - a. The organization should use high-level analysis (review of policies and interviews) if they do not have extensive time and resources.
 - b. Information assets should be tested and examined in more detail as time allows.
 - c. The organization should plan recurring risk assessments to identify more risks over time.
6. Collaborating with information security subject matter experts to help model threats that are foreseeable in the environment, and to help evaluate the effectiveness of current safeguards.

The risk assessor will need to use their professional judgment to select the controls and information assets and to model threats that should be analyzed in the risk assessment. Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Tier 2 Risk Assessment Summary

After having analyzed a set of risks against a single information asset, the Tier 2 organization realizes the advantages of gaining a more comprehensive view of its risks:

1. When organizations consider impacts to their objectives, risk assessors include business interests in their risk analysis, thus engaging non-technical collaborators in risk decision-making.

2. When organizations add impact scores (magnitudes) and likelihood scores they can make more refined distinctions between risks and can prioritize them more reasonably.
3. Pairing information assets to multiple CIS Controls provides a more comprehensive and perhaps accurate understanding of the actual risk to those assets.
4. Pairing information assets to multiple CIS Controls also prompts risk assessors to consider threats that they may have otherwise neglected.

Risk Treatment Recommendations

Introduction

Organizations often think of security safeguards as obstacles to business and productivity. Safeguards often cause personnel to take extra steps to get access to systems or information, or to get approval for normal business activities. Safeguards require investments in time and money, which compete with other priorities. And if they become too disruptive to an organization's mission and objectives, security safeguards can become disliked and avoided.

In fact, disruptive safeguards often cause personnel to work around them just to get their jobs done, which creates more risk.

But risk treatment recommendations can and should result in safeguards that are demonstrably reasonable. And while obtaining a clear definition for "reasonable safeguards" has been a challenge in the legal, regulatory, and information security communities, the CIS RAM provides a practical solution. Risk assessors evaluate risk treatment recommendations to determine whether a security safeguard is reasonable by; comparing the safeguard to the risk it is meant to reduce, and by comparing the safeguard to the risk acceptance criteria.

Risk treatment recommendations are simple to evaluate once the risk assessment criteria and initial risk analysis have been established. The process occurs over the following steps:

1. While examining an unacceptably high risk, review the CIS Control that corresponds with the risk and recommend a feasible way for the organization to implement or improve that control.
2. If that control is not feasible in the near-term, recommend other CIS Controls related to the risk that can be used to reduce it.
3. Evaluate the risk of the recommended safeguard to understand the burden it would pose to the organization. Then compare that safeguard risk to the risk acceptance criteria to determine whether it is appropriate.
4. Also compare the evaluated risk of the recommended safeguard to the observed risk to determine whether the safeguard is reasonable (safeguards with lower risk scores than the observed risk are reasonable.)
5. Sort the risks by their risk score to prioritize the risks and risk treatments that the organization will invest in.

This section demonstrates these steps in detail by describing the process, then by modeling risk treatments for the unacceptably high risks that were evaluated in previous sections.

The reader should review the definitions of 'reasonable' and 'appropriate' that are provided in the glossary. These terms will be used regularly in this section and have distinct meanings.

1. Appropriate: A condition in which risks to information assets will not foreseeably create harm that is greater than the organization or its constituents can tolerate.
2. Reasonable: A condition in which safeguards will not create a burden to the organization that is greater than the risk it is meant to protect against.

Risk Treatment Objectives

The objective of well-formed risk treatment recommendations is to create a prioritized list of information security safeguards that would provide appropriate protections while not posing too great a burden on the organization's purpose.


The risk treatment recommendation exercises that are demonstrated in this section examine the unacceptable risks that were illustrated earlier in the document and will select CIS Controls that would reduce those risks to a degree that is both reasonable (not overly burdensome) and appropriate (not unacceptably harmful).

Recommending Risk Treatment Safeguards from CIS Controls V7

As we examine unacceptably high risks, we will recommend safeguards that are based on CIS Controls. But some of the safeguards that an organization is prepared to implement and operate may not be implemented exactly as described in CIS Controls V7. This process takes into account how to select controls that address risks, and how to determine whether they are designed in a way that makes sense in context of both the risk, and the potential burden to the organization.

Recall the relationship between analyzed risks and their recommended risk treatments in Figure 14.

Figure 14 - Balance Within Core Risk Analysis



A risk and its proposed safeguard are both evaluated using the same criteria. If a proposed safeguard has a higher risk (its "safeguard risk") than the risk acceptance criteria, then it's not appropriate. If the safeguard has a higher score than the observed risk, then it's not reasonable.

The exercises in this section will focus on matching completed risk analyses (in blue) with newly recommended safeguards (in green).

Risk Treatment Example 1 – CIS Control 8.1

The exercise that demonstrated risk analysis in Table 51 showed the Tier 2 organization that its risk of malware on its diary device controllers (CIS Control 8.1) was too high. The organization was not permitted to add anti-malware software to the device controllers.

To recommend safeguards that are based on alternative controls, the risk assessor reviews CIS Control 8.1 to understand its objective. CIS Control 8.1 intends that all devices should actively search for malware and intrusion activity, block the activity, and report it to security management systems.

The organization faces a challenge when they realize that the diary device controller's vendor does not support malware protection on the controllers, and the organization may violate the service terms of the vendor contract if they install available malware protection.

Until the device controller's manufacturer distributes a more secure distribution of their product, the organization will need to consider some safeguards to protect the vulnerable diary device controllers. IT management recommended that the organization simply install malware protection

on the controllers and take their chances with the vendor. But before committing to this plan, they worked with the risk assessor to analyze the risk of doing that. Table 52 illustrates their analysis.

Note: The risk assessor will record how they will address their risk by stating either “Accept,” “Reduce,” “Transfer,” or “Avoid.” *Accepting* and *reducing* risks will be intuitive to the reader. An organization may *transfer* a risk by contracting a third party that may handle the risk better, or by acquiring an insurance policy against the risk. The organization may also *avoid* the risk by no longer engaging in the processes, or handling the information assets that cause the risk

Table 52 - Example Risk Treatment Recommendation for CIS Control 8.1

Risk Analysis	Value
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Information Asset	Diary device controllers.
Control	Anti-malware software is not permitted on the diary device controllers.
Vulnerability	<p>Vulnerabilities are limited because common vectors for receiving malware such as email clients and web browsers are not installed on the controllers. Attackers would need to download malware executables from the Internet using scripts or bash commands.</p> <p>Command line, by design, is only accessible over terminal connections to the console port.</p> <p>Bluetooth attacks may still permit malware executables to be uploaded to a file space associated with an anonymous account. The web admin application on each controller has been tested as vulnerable to arbitrary code execution, cross-site scripting, and other attacks.</p>
Threat	Hackers may implant malware on diary device controllers through Bluetooth misuse, and take advantage of web admin application vulnerabilities to execute files or initiate scripts.
Threat Likelihood	3
Mission Impact	3
Objectives Impact	4
Obligations Impact	3
Risk Score	12
Risk Acceptability	Unacceptable
Risk Treatment Option	Reduce
Recommended Safeguard	Install anti-malware application and host-based intrusion prevention agents on diary device controllers.

Risk Analysis	Value
Safeguard Risk	Signature-identified malware and common volatile RAM exploits will be detected, prevented, and reported to the central management console. If vendor sees the malware and IPS agents on controllers during their quarterly service sessions, they may cancel those service contracts, delaying service on faulty systems until new images can be installed on the devices.
Safeguard Threat Likelihood	4
Safeguard Mission Impact	3
Safeguard Objectives Impact	3
Safeguard Obligations Impact	2
Safeguard Risk Score	12
Risk Acceptability	Unacceptable

The organization has evaluated that management's recommendation to install the anti-malware agents is just as risky as the observed risk that they are trying to address. The estimated safeguard risk equals that of the observed risk, but based on the organization's concern that the vendor was expected to force re-imaging on the diary device controllers meant that they expected to compromise their mission and objectives.

The risk assessor can now advise management against installing the security software directly, but then should provide alternatives.

Background – How Realistic Are Safeguard Risk Estimates?

Critical readers will question how the organization and their risk assessor will know whether their safeguard risk estimations are realistic. After all, how can they know prospectively what their risk would be in such a situation?

There are two important items to keep in mind while gaining comfort with this practice; understanding the legal and regulatory expectations for risk management, and information security standards for evaluating safeguards after they've been implemented.

Law and Regulation: Laws and regulations generally expect risk analysis to evaluate safeguards that are required for achieving compliance, and expect that the risk analysis is performed by appropriately skilled and informed people. These analyses do not guarantee security that is sufficient against any threat, but they do provide a plan toward improved security and compliance that is prioritized by the likelihood of harm, and that has no intolerable harm as its goal.

Information Security Risk Management Standards: Information security risk assessment standards that the CIS RAM is based on operate within fuller risk management programs and cycles. ISO 27005 operates within the ISO 27000 family of standards, and NIST 800-30 works within the NIST Special Publications. Each of these families of standards requires continuous analysis of security safeguards, including analysis of controls after they've been implemented to determine whether they are effective at addressing their security objectives. Recommended safeguards should therefore be risk assessed again after implementation to be sure that they achieve their intended objectives.

Risk Treatment Example 2 – CIS Control 8.1

As the organization considers an alternative risk treatment, they will turn to other CIS Controls to see what alternatives are available to them.

CIS provides to the public a document titled the *CIS Community Attack Model*.¹⁸ The document lists all of the CIS Controls and associates them with the role they play in each stage of incident preparedness and response; planning, detection, and defense. By using the CIS Community Attack model, an organization can find related and alternative CIS Controls that may substitute for controls that they cannot sufficiently implement.

By reviewing the Community Attack Model (Figure 15), the risk assessor can quickly identify which controls they should consider if CIS Control 8.1 is not feasible.

The Community Attack Model is provided in the *CIS_RAM_Workbook* in a tab titled "Attack Path Models" for the reader's convenience, and can be downloaded at the CIS website.

A partial snapshot of the framework demonstrates what the Tier 2 risk assessor found as they searched for alternative controls.

¹⁸ The *Community Attack Model* may be accessed here: <https://www.cisecurity.org/white-papers/cis-community-attack-model/>

Figure 15 - Partial Community Attack Model

		Attack Stages				
Functions	Controls	Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege
	Identify	control of HW, SW inventory; Network logs	threat intelligence			control of administrative privilege
		Protect	firewall; mail gateway filtering; web filtering; manage ports, protocols, services; continuous vulnerability assessment	hardened configurations	continuous vulnerability assessment; firewall; mail gateway filtering; web filtering; secure remote access; NIPS	patching; hardened configurations; HIPS; anti-malware; containerization; app whitelisting; Data Execution Protection
	Detect	firewall; honeypot; Network authentication; Network logs	audit logs; threat intelligence	audit logs; Anti-malware; Network Intrusion Detection system	HIPS; anti-malware; containerization; app whitelisting; Data Execution Prevention;	account monitoring; control of admin privilege; audit logs; Configuration Monitoring
	Respond	honeypot			Incident Response - Execution	audit logs; Configuration Management; Account Management
	Recover				Incident Response - Execution; control of HW, SW inventory	

Risk assessors may reference the Community Attack Model to find controls that can be complementary and alternative to the recommended safeguards they are assessing. If an organization struggles to implement a sub-control, they could look for controls that play a similar role in the Community Attack Model to find alternative controls that might help them meet the same security objective. For example, if an organization cannot easily use audit logs to *detect delivery* of a kind of threat, they may look to another control in the cell that intersects with the *detect* row and the *delivery* column to find similar controls – and to eventually see network intrusion detection controls, which may be more useful in their environment.

Given the objectives of CIS Control 8.1 to protect systems against malware and identifiable intrusions, the risk assessor reviews the Community Attack Model and finds anti-malware in cells that address *protecting* against *delivery*, and for *protecting* against and *detecting initial compromise*. Detecting delivery of malware seems to be a good place to start their defenses as it is closest to the threat they are addressing in their risk analysis, so they review the CIS Controls that are in the cell at the intersection of *Detect* and *Delivery* to consider options.

Network Intrusion Detection systems and Network Intrusion Prevention Systems (“IDS/IPS”) are interesting to the organization as a network-layer protection. The risk assessor reviews the sub-controls within CIS Control 12 “Boundary Defense” to see how IDS is described. CIS Control 12.6 “Deploy Network-based IDS Sensor” presents a compelling option for them because their current use of the diary device controllers during clinical visits is within mobile wireless LANs that they own and control. A lightweight IDS would be a plausible option in that case.

But would a lightweight IDS/IPS in those portable wireless LANs reduce their malware and intrusion risk on their diary device controllers? They model CIS Control 12.6 as a safeguard against this same risk to see if it would be reasonable.

Table 53 - Example Risk Treatment Recommendation for CIS Control 8.1 using CIS Control 12.6

Risk Analysis	Value
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Information Asset	Diary device controllers.
Control	Anti-malware software is not permitted on the diary device controllers.
Vulnerability	<p>Vulnerabilities are limited because common vectors for receiving malware such as email clients and web browsers are not installed on the controllers. Attackers would need to download malware executables from the Internet using scripts or bash commands.</p> <p>Command line, by design, is only accessible over terminal connections to the console port.</p> <p>Bluetooth attacks may still permit malware executables to be uploaded to a file space associated with an anonymous account. The web admin application on each controller has been tested as vulnerable to arbitrary code execution, cross-site scripting, and other attacks.</p>
Threat	Hackers may implant malware on diary device controllers through Bluetooth misuse, and take advantage of web admin application vulnerabilities to execute files or initiate scripts.
Threat Likelihood	3
Mission Impact	3
Objectives Impact	4
Obligations Impact	3
Risk Score	12
Risk Acceptability	Unacceptable
Risk Treatment Option	Reduce
Recommended Safeguard	[CIS Control 12.6] Add a lightweight IDS/IPS device to portable LANs that operate in clinical visits where diary device controllers are used.
Safeguard Risk	<p>IDS/IPS devices will detect recognizable attacks from other hosts within the LAN that attempt to deliver and deploy malware to controllers.</p> <p>IDS/IPS may not detect malware payload from hosts that connect to controllers over encrypted protocols.</p>
Safeguard Threat Likelihood	2

Risk Analysis	Value
Safeguard Mission Impact	3
Safeguard Objectives Impact	3
Safeguard Obligations Impact	2
Safeguard Risk Score	6
Risk Acceptability	Acceptable

Even though CIS Control 12.6 does not present a direct anti-malware solution this scenario does present an acceptable risk, and a reasonable risk. The recommended risk treatment reduces the likelihood of a successful attack on diary device controllers while not eliminating it in this case.

Management is satisfied that the recommended safeguard is appropriate, but that are also interested in another option presented by CIS Control 12.11 to use two-factor authentication to log into terminal sessions at the diary device controllers. If two-factor authentication provides even lower risk, and the diary device controller vendor supports the option, then this may be a better safeguard than the lightweight IDS/IPS.

Recall that diary devices already store encrypted soft-certs to help authenticate the devices to their accounts on the controllers. Checking with the vendor, the organization sees that the soft-cert option is available for administrator systems that connect to the controllers as well. When the organization's site administrators connect to the diary device controllers while on site, they access terminal sessions using SSH, the only protocol available to them. Multiple soft-certs can be used to both authenticate the SSH sessions, and to execute commands at the controllers.

They model this alternative risk treatment below.

Table 54 - Example Risk Treatment Recommendation for CIS Control 8.1 using CIS Control 12.11

Risk Analysis	Value
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Information Asset	Diary device controllers.
Control	Anti-malware software is not permitted on the diary device controllers.
Vulnerability	<p>Vulnerabilities are limited because common vectors for receiving malware such as email clients and web browsers are not installed on the controllers. Attackers would need to download malware executables from the Internet using scripts or bash commands.</p> <p>Command line, by design, is only accessible over terminal connections to the console port.</p> <p>Bluetooth attacks may still permit malware executables to be uploaded to a file space associated with an anonymous account. The web admin application on each controller has</p>

Risk Analysis	Value
	been tested as vulnerable to arbitrary code execution, cross-site scripting, and other attacks.
Threat	Hackers may implant malware on diary device controllers through web application exploits while they operate in clinical settings.
Threat Likelihood	3
Mission Impact	3
Objectives Impact	4
Obligations Impact	3
Risk Score	12
Risk Acceptability	Unacceptable
Risk Treatment Option	Reduce
Recommended Safeguard	[CIS Control 12.11] Require all usage of SSH and all authentication on diary device controllers to use soft-certs stored on client devices as a second factor of authentication.
Safeguard Risk	All attempts at accessing SSH services in diary device controllers will be blocked unless clients use soft-certs to access SSH sessions. Attackers may seize and re-use soft-certs during 8-hour long clinical visits and may attack controllers as a result.
Safeguard Threat Likelihood	1
Safeguard Mission Impact	3
Safeguard Objectives Impact	3
Safeguard Obligations Impact	2
Safeguard Risk Score	3
Risk Acceptability	Acceptable

It appears that the safeguard risk obtained while using CIS Control 12.11 is much lower than the safeguard risk modeled by the option to use CIS Control 12.6. And because the vendor already supports multi-factor authentication, the solution is almost already configured. The organization chooses to use CIS Control 12.11 as their risk treatment control to protect their diary device controllers against malware until the vendor provides a more robust solution.

Exercise:

The reader should use the template Risk Register – Tier 2 that is provided in the supplementary document *CIS_RAM_Workbook* to enter risk treatment recommendations for each risk that evaluated as unacceptably high.

The reader should consider:

1. Whether an existing safeguard can be improved, and how that would be done.
2. Whether a safeguard based on a different CIS Control would provide reasonable and appropriate risk.
3. Collaborating with information security subject matter experts to help model the potential effectiveness of recommended safeguards.

The risk assessor will need to use their professional judgment to design and recommend information security safeguards, and to evaluate prospectively the risk that they may pose. Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Risk Treatment Recommendations Summary

Risk treatment recommendations are a critical part of risk assessments to be sure that the organization has developed a plan for addressing risks without creating other risks to the organization or its constituents. Some of the benefits that have been demonstrated about this process are:

1. Organizations can demonstrate to collaborating business managers how recommended security safeguards can be implemented without creating too much of a burden on the business mission and objectives.
2. Organizations can demonstrate to regulators and other legal authorities that safeguards are reasonable because the safeguard risk of the safeguard (the “burden” to the organization) is not greater than the risk that it is meant to reduce.
3. Organizations can demonstrate that recommended safeguards would be appropriate by showing that they would not foreseeably create an impact that would be intolerable to the organization or its constituents.
4. Organizations may find it valuable to evaluate multiple safeguards in case one safeguard is more reasonable (creates an even lower risk) than another safeguard.
5. Risk assessors will find that their colleagues will understand and appreciate risks and controls when risk assessors and subject matter experts collaborate on evaluating risk, and planning safeguards.

The process for evaluating risks and for recommending appropriate risk treatments has been demonstrated at a general level. However, some questions likely remain for the reader for evaluating safeguards, estimating likelihood, and the suitability of probability models in risk analysis. These more detailed topics will be discussed in the chapter “Risk Analysis Techniques.”

Chapter 4: Threat-Based Risk Assessment Instructions for Tiers 3 and 4 Organizations

Tiers 3 and Tier 4 risk assessment instructions are well-suited to organizations that fit the profile of Tier 3 and Tier 4 organizations as described by the NIST Cybersecurity Framework. These organizations can be identified as having the following characteristics:

- **NIST Tier:** Tiers 3 and Tier 4 organizations. Tiers 3 and Tier 4 materials are best suited for organizations that are using risk-based criteria for enterprise-wide policies and processes.
- **Expertise:** The organization has resources and capabilities to analyze security threats, and to plan risk-appropriate safeguards, including the on-hand skills to model how threats would operate within their organization.
- **Time:** The organization is able to invest time to analyze risks at the level of specific systems, devices, and applications within the context of specific threats.

This chapter is comprised of sections that each address a specific activity within a risk assessment. Readers should engage this chapter by first reading the text in each section, and then conducting the exercises that are recommended for each section. The material presented in the CIS RAM is substantially different from many other risk assessment standards and models, so the reader should first understand the aim of each section, and then practice what they learn using templates that are provided in the supplementary document *CIS_RAM_Workbook*.

While conducting their first CIS RAM-based risk assessment, organizations should be careful to not try to “boil the ocean.” Regulatory bodies and information security standards alike understand that not all risks can be identified in a single assessment. Organizations should continuously and regularly assess risks to identify, understand, and manage risks over time.

The Risk Assessment Project

Overview

Risk assessments are projects with clear steps for preparing, conducting, and reporting risk analysis. And while risk assessment projects can be modeled with a project plan, each organization’s project approach will vary depending on factors such as resource availability, and will develop over time as organizations become more capable in their cybersecurity maturity. This section will describe a basic risk assessment project, its components and variations, and will present guidance for preparing the project plan.

How and When to Use Threat-Based Risk Analysis

The threat-based risk analysis that is described in this chapter requires considerably more effort and expertise than the analysis methods used by Tier 1 and Tier 2 organizations. This is primarily due an analysis step – attack path modeling – that precedes the risk assessment.

Attack path modeling provides organizations with valuable insight into information security risks by analyzing how their information assets would respond to known attack scenarios.

For example, if an organization wants to understand their susceptibility to a trojan that exfiltrates data from a specific database, they would map out plausible scenarios for how a trojan would enter their environment, would be installed on a workstation, would gain privileges to the database, would access data in the database, and would then send the stolen data to a target system. The attack path model identifies these steps, and lists the information assets that would

be included in that attack. This results in a listing of information assets and threats that would compromise them. Risk assessors would then base their risk analysis on that list of assets and threats.

This chapter will demonstrate how threat-based risk analysis offers useful insight into information security risks, but the reader can understand how a comprehensive risk assessment using this approach could be time-consuming.

Organizations that are beginning to practice threat-based risk analysis may want to answer specific risk questions, rather than to plan an entire risk assessment based on its thorough analysis. Some examples uses may be:

1. After having completed a risk assessment using the approach described for Tier 2 organizations, the risk assessor may want to understand how well-prepared information assets are for preventing specific, current threats.
2. The accuracy of a specific asset's risk score is being questioned because personnel believe that the asset is well-protected by layers of security.
3. The comprehensiveness of a recent risk assessment is in question because interested parties believe some asset vulnerabilities were well not thought out.
4. While planning risk remediation, management wonders whether resolving the risk for one information asset will have a cascading and beneficial effect on a set of other assets.

For these scenarios and others like them, threat-based risk analysis can help organizations evaluate risk within the context of a chain of events without having to apply such deep scrutiny to every asset in the risk assessment scope.

Risk Assessment Project Management Project Outline

Risk assessments are projects that require planning, identification of assets and asset owners, scheduling of sessions, and data gathering. CIS RAM provides detailed instruction for these project management and planning steps in the chapters for Tier 1 and Tier 2 organizations that may benefit from these tactical instructions.

Given the expected maturity of organizations Tier 3 and Tier 4 organizations, these instructions will not be provided in this chapter. However, the reader may benefit from reviewing those materials in those chapters before continuing with Chapter 4.

The supplementary document *CIS_RAM_Workbook* provides templates for project scheduling and scoping Tier 3 and Tier 4 organizations, if needed.

Defining Risk Assessment Criteria

Introduction

Risk assessment criteria are the numerical and plain-language statements that an organization uses to evaluate their cybersecurity risk. The most familiar form of risk calculations, "Risk = Likelihood x Impact," is the basis for risk analysis in the CIS RAM. But it is just the starting point for risk analysis.

Risk assessment criteria must be meaningful to the organizations that use them, so they must be tied to the potential benefit and harm that the organization may create. The impact of a cybersecurity breach may harm the organization itself, it may harm the organization's ability to successfully achieve its mission, or it may harm others.

Because cybersecurity failures impact parties both inside and outside an organization, risk assessment criteria must be universally meaningful and must address the interests of all potentially affected parties. Additionally, risk assessment criteria must demonstrate to authorities,


such as regulators and litigators, that the organization considers the risk of harm to others as much as the risk of harm to themselves.

While these requirements may seem complex, the method presented in this section will sufficiently address them while using a technique that is simple to develop and use.

Risk Assessment Criteria Foundations

The risk analysis provided in the CIS RAM is at its root a question of balance between the potential of future harm against the certain burden of a safeguard. Regulators and litigators have long considered this balance as key to acting as a “reasonable person.” The core structure of a risk statement is provided below to illustrate the core concept of balance.

Figure 16 - Balance Within Core Risk Analysis



Notice a few things right away with the model risk analysis in Figure 16.

- While organizations typically evaluate the observed risk to determine whether they should address or accept it, this risk statement deliberately compares the observed risk to a proposed safeguard.
- The criteria that evaluates the risk also evaluates the safeguard.
- The impact of the risk estimates the potential of harm to the organization and the potential harm against others.

Risk assessors compare risks to their proposed safeguards to determine whether those safeguards would create a foreseeably lower risk than the current state. To accomplish this, the assessor evaluates the current state risk (or “observed risk”) and the proposed safeguard using the same criteria to ensure comparability.

This comparison prevents organizations from implementing safeguards that are overly burdensome, or that create new, unacceptable risks. For example, an organization that uses software that is no longer supported by the vendor, but relies on that software for critical business purposes, should find alternative methods for identifying and controlling potential security risks until they replace the software. If management recommends quickly changing out to inferior, but secure software, the organization may suffer a greater impact to their mission than the security risk they are trying to avoid.

While considering CIS Control 18: Application Software Security, a risk statement can estimate the foreseeability of an impactful threat. The risk can be stated as it appears in Table 55 (where the risk score ‘12’ is a product of the likelihood ‘3’ and the highest impact score ‘4’):

Table 55 - Example Core Risk Statement

Observed risk	Likelihood	Impact to Us	Impact to Others	Risk Score
Hackers may exploit the unsupported, but critical application.	<u>3</u>	3	<u>4</u>	12

A risk assessor should then recommend and evaluate a safeguard to reduce the high security risk, as illustrated in Table 56. Here, the organization would realize that the likelihood of a negative impact to their mission is greater than the current state risk. This is an obvious case of the burden being greater than the risk, and a recommended safeguard being unreasonable.

Table 56 - Example Unreasonable Proposed Safeguard

Proposed Safeguard	New Risk	Likelihood	Impact to Us	Impact to Others	Safeguard Risk
Replace application with inferior, secure application.	Application will operate inefficiently.	<u>5</u>	<u>3</u>	1	15

When faced with this analysis, the organization must then find another way to address the risk. This process will be described later in this chapter in the section Risk Treatment Recommendations.

But what should be apparent is that without a definition of risk assessment criteria the likelihood and impact scores are not meaningful. What would impacts or likelihoods of '1', '2', '3', '4', or '5' mean, anyway? The organization will need to create definitions for their likelihood and impact scores so that they are meaningful to all interested parties, and so that they provide a consistent method for risk evaluation.

Impact Definitions

Tier 3 and Tier 4 organizations generally benefit from more business involvement in managing cybersecurity risk than Tier 1 organizations. Because of that increased involvement, the risk assessment criteria can be – and should be – more explicit and detailed than those used by Tier 1 organizations. Advanced organizations can consider more nuance in terms of business impacts and tolerance, and can employ organizational objectives with more authority.

An impact definition for Tier 3 and Tier 4 organizations can look like the one depicted in Table 57.

Table 57 – Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objectives: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

Background – Impact Definitions

This document provides instructions for defining impacts and impact scores (magnitudes) in this section with more in-depth instructions and examples in the “Risk Analysis Techniques” chapter. The reader should understand before going further that organizations in most cases should not define impacts exclusively using financial values. While cost is a common and almost necessary consideration while evaluating risks and safeguards, if it is the only criterion, the organization will communicate to their personnel, as well as to interested parties and authorities, that cost is their only concern. The purpose that the organization serves and the harm that may befall others must be part of the evaluation if risk is to be responsibly tied to the potential of harm, and if the evaluation is to be understandable to regulators and legal authorities.

Organizations should also consider having more than three impact types in their impact definitions if they have more than one mission, multiple objectives, and many obligations that they need to consider in their risk analysis. While this expansion may create an increasingly wide risk register, it can help organizations feel comfortable all relevant interests were considered in their risk analysis.

Tier 3 and Tier 4 organizations that previously used Tier 1 risk analysis processes may build on their simpler risk assessment criteria that used three levels of impact scoring. For the purposes of reference to our example organization – a health information provider – they have gone through a year or two of risk management, have gained the attention and confidence of business managers and executives. As a result, their ability to assess cybersecurity risk using business criteria will also improve.

We can see by comparing the risk assessment criteria for Tier 1 organizations in Chapter 2 to Table 57 that the detailed descriptions of impacts have increased in two dimensions; the number of impact score options increased from three to five, and there is an additional impact definition for business objectives.

Tier 3 and Tier 4 organizations will find that using a range of five scores increases the utility of risk prioritization at the end of the risk assessment. A three-by-three risk assessment criteria model provides organizations with six possible risk scores; 1, 2, 3, 4, 6, and 9. This leads to a somewhat course grouping that may cause risks of somewhat different urgencies to be indistinguishable.

A five-by-five risk assessment criteria model allows for 14 possible risk scores of; 1, 2, 3, 4, 5, 6, 8, 9, 10, 12, 15, 16, 20, and 25. Now risks of somewhat different urgencies will likely be classified into different risk scores and will be more easily distinguished while prioritizing them.

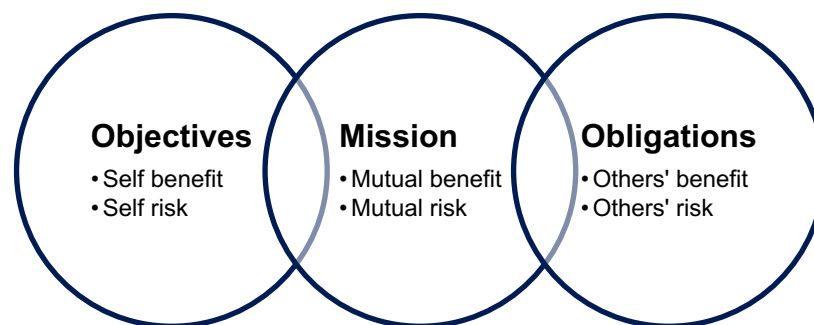
Also note that the impact scores of '1' and '2' in Table 57 are shaded grey to separate them from the higher scores. Scores '1' and '2' describe impact magnitudes that would be generally thought of as acceptable. The risk acceptance criteria process will be explained later in the document, but it is useful to consider now that the scores '1' and '2' are consistent in their definition of impact magnitudes, and that the impact scores of '3', '4', or '5' could consistently be thought of as unacceptably high.

When the example health information provider graduates from using their Tier 1 instructions, that also have a new impact type to consider. Table 57 uses an impact type for business objectives to be considered in the organization's risk analysis. Business objectives are more self-focused than missions and obligations, and are aligned with success criteria commonly found in business. Some examples include profitability, growth, maintaining accreditations, customer satisfaction, or retaining a position in the marketplace.

Objectives align most directly with what is commonly thought of as the "cost" of a safeguard. But rather than allowing an organization to arbitrarily decide that a safeguard costs too much, this method of including cost in terms of impacts to objectives forces the organization to evaluate why a cost would be excessive. Does the cost of the safeguard impede profitability goals? Does the safeguard limit efficiency or growth? Those are certainly reasonable concerns, as long as the profitability goals are in parity with the mission and obligations impacts. *In other words, an organization should not let profitability be more important than harm to others, or harm to their ability to fulfill their mission.*

See "Note Regarding Use of Financial Costs as Objectives" in Chapter 5.

Figure 17 - Objectives, Mission, Obligations



Organizations are well-served with this model because business management, technicians, compliance personnel, and legal counsel all have their interests addressed in risk analysis that uses these criteria.

An in-depth explanation of how to develop impact definitions with multiple examples is provided in the chapter “Risk Analysis Techniques.”

Likelihood Definitions

The likelihood definition for a Tier 3 and Tier 4 organization should also increase in nuance from the simpler Tier 1 definition, and can do so by adding two more scores to the table as shown in Table 58.

Table 58 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

- “Not Foreseeable” implies that a threat is not plausible in the environment that is being assessed. Loss of portable media may not be foreseeable during a risk assessment of a hosted application.
- “Foreseeable” implies something that is plausible, but the organization would be surprised if it occurred. A founding executive taking copies of sensitive data to competitors may be considered foreseeable, even if it is not expected.
- “Expected” implies a threat that is not common, but that would eventually happen. Phishing attacks or other social engineering attacks may be expected in many environments.
- “Common” implies something that happens repeatedly, such as mis-addressed emails with sensitive information, malware attacks, or loss of laptops and mobile devices.
- “Current” implies threats that are rarely not present, such as port scanning on perimeter devices, or sharing of information in quasi-public spaces such as pharmacy counters or bank tellers.

When risk assessors estimate the likelihood of a threat, they will select scores ‘1’, ‘2’, ‘3’, ‘4’, or ‘5’ using the foreseeability definition as their guidance. Organizations may add time-based limits in their foreseeability definitions (i.e. “Foreseeable within planning thresholds,” “Expected within the five-year plan” or “Not foreseeable in the next fiscal year”). If organizations do introduce time limits in their likelihood definitions they should prioritize risk treatment investments to meet these timelines. That may be excessively challenging to many organizations, so they should proceed with caution.

Developing the Risk Assessment Criteria

Because risk assessment criteria are meant to describe risk as it applies to the organization that owns the risk, it is appropriate for the most senior management who are responsible for the mission, objectives, and obligations to participate in developing and accepting the criteria.

Table 59 lists roles commonly involved in risk assessment criteria development, and the interested perspective they bring to the definition effort.

Table 59 - Roles Involved in Defining Risk Assessment Criteria

Role	Perspective
Chief Executive Officer Chief Operations Officer	To ensure that the mission, objectives, and obligations of the organization are appropriately defined, and to ensure that a distinction between acceptable and unacceptable impacts are appropriately delineated.
Chief Compliance Officer	To ensure that the interests of regulatory agencies are appropriately included in risk definitions.
Chief Financial Officer	To ensure that objectives are appropriately defined, particularly the distinction between acceptable and unacceptable impacts.
Chief Information Officer Chief Technology Officer	To ensure that technical performance, service, and capabilities are considered, and to include all types of information processes beyond technology.
General Counsel Outside Counsel	To ensure that obligations are appropriately defined and that they compare well with the mission and objectives.
Internal Audit Audit Committee	To ensure that the concerns of interested parties are well represented in all impact definitions and scores.
Key Customers / Clients Key Constituents	To ensure that their interests are included in the obligations definition.

Exercise:

The reader may develop their organization's risk assessment criteria using the "Criteria - Tier 3 & 4" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Developing the risk assessment criteria in collaboration with a business managers and legal counsel to ensure that the Mission, Objectives, and Obligations definitions are sensible to the organization.
2. Working with legal counsel to help ensure that impact definitions appropriately address the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.
3. Referring to guidance for defining and scoring impact types in the "Risk Analysis Techniques" chapter.

The risk assessor will need to use their professional judgment to define impact types, and to describe levels of impact that the organization must manage to. Because risk assessment criteria are a declaration by the organization of what they will manage to in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

Defining Risk Acceptance Criteria

Introduction

Because risk assessments are essentially questions of balance, the criteria for accepting risk should help determine whether balance was achieved. In CIS RAM risk acceptance has two components:

- Appropriate risk: That the likelihood of an impact must be acceptable to all foreseeably affected parties
- Reasonable risk: That the risk posed by a safeguard must be less than or equal to the risk it protects against.

While these components have been demonstrated briefly above, the first component will be described in more detail in this section. The second component will be described later in the Risk Treatment Recommendations section further on.

After establishing the impact and likelihood definitions, Tier 3 and Tier 4 organizations are now well positioned to state their risk acceptance criteria. Recall that impacts were defined within scores that ranged from '1' to '5'. Acceptable impact scores '1' and '2' were defined in a manner that would appear appropriate to interested parties (and are shaded grey to indicate their acceptability), and the impact score '3' was the lowest unacceptable score.

Table 60 – Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objective: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

And similarly, likelihood scores were within a range of '1' to '5' as below. Once our example organization develops their risk management maturity and are ready to refine their risk distinctions, they may decide to not tolerate unacceptable impacts if they are *foreseeable but not expected* ('2'), or if they are *expected to occur* ('3'). This would be a decision best made by their executives, and especially their compliance team, general counsel, and interested parties. But in this case, the model will assume that they selected a threshold score of '3'.

Table 61 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

So Tier 3 and Tier 4 organizations would define their risk acceptance like this:

Table 62 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
3	x	3	=	9
... therefore ...				
Acceptable Risk			<	9

An example of the heat map for this assessment criteria is shown below. Note that this heat map is defined not just by numbers and colors, but now by a set of criteria that address business issues and a duty of care to protect others.

Figure 18 - Example Heat Map

		Impact				
		1	2	3	4	5
Likelihood	5					
	4					
	3					
	2					
	1					

Exercise:

The reader should define their organization's risk acceptance criteria using the "Criteria – Tier 3 & 4" worksheet that is provided in the supplementary document *CIS_RAM_Workbook*.

The reader should consider:

1. Working with a business management sponsor who can help ensure that the risk acceptance criteria are sensible to the organization.
2. Working with legal counsel to help ensure that the definition for risk acceptance addresses the interests of all potentially affected parties, and to ensure that impact statements appear equitable to all parties.

The risk assessor will need to use their professional judgment to identify levels of acceptable risk. Because risk acceptance criteria are a declaration by the organization of what they will tolerate in terms of harm to themselves and harm to others, organizations should consult with legal counsel before finalizing these criteria and making risk decisions based on them.

As organizations assess their risk using CIS Controls V7 (modeled in the following section) they will be able to automatically determine whether the risk evaluates as acceptable or not without needing to consider the question differently in each case. A simple estimation of the likelihood and impact will automatically determine how the organization should prioritize each risk, and whether the organization can safely accept the risk as a "reasonable" option.

A Threat-Based Risk Assessment Process

Introduction

Tier 3 and Tier 4 organizations benefit from collaboration with business management. They also have refined knowledge of how cyberattacks work. And due to their contact with outside parties have considerably more data about on-the-ground effectiveness of their security safeguards. Because Tier 3 and Tier 4 organizations have these advantages, their ability to analyze and respond to risk should be more refined than their peers with less maturity.

The risk analysis method described in this section is based on an "attack path" model. An attack path, sometimes called a "kill chain," is the route that an attack takes to compromise information assets. For example, ransomware attacks involve many attack stages, starting from the hacker's reconnaissance, through preparation and delivery of exploits, initial compromise, privilege abuse, through to the final control of the targeted storage volume and data.


And while not all attacks are planned (some are automated, drive-by attacks, and some are accidental) they can be modeled in an attack path to understand what chain of safeguards would fail in order for a threat to successfully compromise an asset.

The attack path for the ransomware example above would start with the attacker targeting an organization that would likely pay a ransom to access their critical information. They would research key personnel in the organization to refine their target, and then would develop an exploit for that personnel. They may place the exploit on an Internet server that is accessible to the victim, and would craft an email message for the victim that links to the exploit. When the victim interacts with the email message, they would download the ransomware which then would run on their computer. The attacker could then use the ransomware to encrypt the hard drive, and depending on the variation of ransomware, either lock the information, or copy it to an Internet server that the attacker controls.

This attack path can be drawn to include a set of information assets that the organization could control, and that would be exploited in the attack, such as; information about their employees and their jobs, email servers, email clients, firewalls, content filters, proxy servers, malware protection appliances and software, operating systems, hard drives, and yes, people.

Figure 19 - Example Attack Path and Information Assets

Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege	Internal Recon	Lateral Movement	Establish Persistence	Execute Mission Objectives
Social media	NA	Email server/client	Email client/user computer	User privileges	Local files	NA	Encrypted OS	Deliver ransom instructions



Using the risk analysis steps that were previously described for Tier 1 and Tier 2 organizations, a risk assessor can use risk analysis to determine how well prepared each information asset is for preventing or detecting specific attacks as they are in play.

This “attack path” approach to risk analysis requires a preliminary analysis step before working in the risk register. That analysis step – attack path modeling – documents the lifecycle of an attack, and identifies information assets or asset classes that would be involved in the attack. The information developed in this preliminary analysis provides the risk assessor with a list of information assets that would be involved in a type of attack, and allows the assessor to evaluate the risks they face based on how well their safeguards align with the CIS Controls.

This section will describe and work through a Tier 3 and Tier 4 organization’s attack path analysis and risk register using the CIS Controls to model risk-appropriate safeguards. The risk register and attack path model worksheet described in this section are available as templates in the supplementary document *CIS_RAM_Workbook*.

The Risk Register

The Tier 3 and Tier 4 organization’s risk register is much like the layout of the risk registers shown in Tier 1 and Tier 2 instructions, but establishes attack path models and threats as the basis for analysis. This section will demonstrate the risk assessment processes for Tier 3 and Tier 4 organizations using the risk register template provided in the supplementary document *CIS_RAM_Workbook* for Tier 3 and Tier 4 organizations.

The layout map of a risk register for a Tier 3 and Tier 4 organization is depicted in Figure 20.

Figure 20 - Risk Register Layout Map

Section	Attack Models and Threats	Assets	CIS Controls	Implementation	Evaluation	Recommendation	Safeguard
1	Row 1	Row 1	Row 1	Row 1	Row 1	Row 1	Row 1
2	Row 2	Row 2	Row 2	Row 2	Row 2	Row 2	Row 2
3	Row 3	Row 3	Row 3	Row 3	Row 3	Row 3	Row 3
4	Row 4	Row 4	Row 4	Row 4	Row 4	Row 4	Row 4
5	Row 5	Row 5	Row 5	Row 5	Row 5	Row 5	Row 5
6	Row 6	Row 6	Row 6	Row 6	Row 6	Row 6	Row 6
7	Row 7	Row 7	Row 7	Row 7	Row 7	Row 7	Row 7
8	Row 8	Row 8	Row 8	Row 8	Row 8	Row 8	Row 8
9	Row 9	Row 9	Row 9	Row 9	Row 9	Row 9	Row 9
10	Row 10	Row 10	Row 10	Row 10	Row 10	Row 10	Row 10
11	Row 11	Row 11	Row 11	Row 11	Row 11	Row 11	Row 11
12	Row 12	Row 12	Row 12	Row 12	Row 12	Row 12	Row 12
13	Row 13	Row 13	Row 13	Row 13	Row 13	Row 13	Row 13
14	Row 14	Row 14	Row 14	Row 14	Row 14	Row 14	Row 14
15	Row 15	Row 15	Row 15	Row 15	Row 15	Row 15	Row 15
16	Row 16	Row 16	Row 16	Row 16	Row 16	Row 16	Row 16
17	Row 17	Row 17	Row 17	Row 17	Row 17	Row 17	Row 17
18	Row 18	Row 18	Row 18	Row 18	Row 18	Row 18	Row 18
19	Row 19	Row 19	Row 19	Row 19	Row 19	Row 19	Row 19
20	Row 20	Row 20	Row 20	Row 20	Row 20	Row 20	Row 20
21	Row 21	Row 21	Row 21	Row 21	Row 21	Row 21	Row 21
22	Row 22	Row 22	Row 22	Row 22	Row 22	Row 22	Row 22
23	Row 23	Row 23	Row 23	Row 23	Row 23	Row 23	Row 23
24	Row 24	Row 24	Row 24	Row 24	Row 24	Row 24	Row 24
25	Row 25	Row 25	Row 25	Row 25	Row 25	Row 25	Row 25
26	Row 26	Row 26	Row 26	Row 26	Row 26	Row 26	Row 26
27	Row 27	Row 27	Row 27	Row 27	Row 27	Row 27	Row 27
28	Row 28	Row 28	Row 28	Row 28	Row 28	Row 28	Row 28
29	Row 29	Row 29	Row 29	Row 29	Row 29	Row 29	Row 29
30	Row 30	Row 30	Row 30	Row 30	Row 30	Row 30	Row 30
31	Row 31	Row 31	Row 31	Row 31	Row 31	Row 31	Row 31
32	Row 32	Row 32	Row 32	Row 32	Row 32	Row 32	Row 32
33	Row 33	Row 33	Row 33	Row 33	Row 33	Row 33	Row 33
34	Row 34	Row 34	Row 34	Row 34	Row 34	Row 34	Row 34
35	Row 35	Row 35	Row 35	Row 35	Row 35	Row 35	Row 35
36	Row 36	Row 36	Row 36	Row 36	Row 36	Row 36	Row 36
37	Row 37	Row 37	Row 37	Row 37	Row 37	Row 37	Row 37
38	Row 38	Row 38	Row 38	Row 38	Row 38	Row 38	Row 38
39	Row 39	Row 39	Row 39	Row 39	Row 39	Row 39	Row 39
40	Row 40	Row 40	Row 40	Row 40	Row 40	Row 40	Row 40
41	Row 41	Row 41	Row 41	Row 41	Row 41	Row 41	Row 41
42	Row 42	Row 42	Row 42	Row 42	Row 42	Row 42	Row 42
43	Row 43	Row 43	Row 43	Row 43	Row 43	Row 43	Row 43
44	Row 44	Row 44	Row 44	Row 44	Row 44	Row 44	Row 44
45	Row 45	Row 45	Row 45	Row 45	Row 45	Row 45	Row 45
46	Row 46	Row 46	Row 46	Row 46	Row 46	Row 46	Row 46
47	Row 47	Row 47	Row 47	Row 47	Row 47	Row 47	Row 47
48	Row 48	Row 48	Row 48	Row 48	Row 48	Row 48	Row 48
49	Row 49	Row 49	Row 49	Row 49	Row 49	Row 49	Row 49
50	Row 50	Row 50	Row 50	Row 50	Row 50	Row 50	Row 50
51	Row 51	Row 51	Row 51	Row 51	Row 51	Row 51	Row 51
52	Row 52	Row 52	Row 52	Row 52	Row 52	Row 52	Row 52
53	Row 53	Row 53	Row 53	Row 53	Row 53	Row 53	Row 53
54	Row 54	Row 54	Row 54	Row 54	Row 54	Row 54	Row 54
55	Row 55	Row 55	Row 55	Row 55	Row 55	Row 55	Row 55
56	Row 56	Row 56	Row 56	Row 56	Row 56	Row 56	Row 56
57	Row 57	Row 57	Row 57	Row 57	Row 57	Row 57	Row 57
58	Row 58	Row 58	Row 58	Row 58	Row 58	Row 58	Row 58
59	Row 59	Row 59	Row 59	Row 59	Row 59	Row 59	Row 59
60	Row 60	Row 60	Row 60	Row 60	Row 60	Row 60	Row 60
61	Row 61	Row 61	Row 61	Row 61	Row 61	Row 61	Row 61
62	Row 62	Row 62	Row 62	Row 62	Row 62	Row 62	Row 62
63	Row 63	Row 63	Row 63	Row 63	Row 63	Row 63	Row 63
64	Row 64	Row 64	Row 64	Row 64	Row 64	Row 64	Row 64
65	Row 65	Row 65	Row 65	Row 65	Row 65	Row 65	Row 65
66	Row 66	Row 66	Row 66	Row 66	Row 66	Row 66	Row 66
67	Row 67	Row 67	Row 67	Row 67	Row 67	Row 67	Row 67
68	Row 68	Row 68	Row 68	Row 68	Row 68	Row 68	Row 68
69	Row 69	Row 69	Row 69	Row 69	Row 69	Row 69	Row 69
70	Row 70	Row 70	Row 70	Row 70	Row 70	Row 70	Row 70
71	Row 71	Row 71	Row 71	Row 71	Row 71	Row 71	Row 71
72	Row 72	Row 72	Row 72	Row 72	Row 72	Row 72	Row 72
73	Row 73	Row 73	Row 73	Row 73	Row 73	Row 73	Row 73
74	Row 74	Row 74	Row 74	Row 74	Row 74	Row 74	Row 74
75	Row 75	Row 75	Row 75	Row 75	Row 75	Row 75	Row 75
76	Row 76	Row 76	Row 76	Row 76	Row 76	Row 76	Row 76
77	Row 77	Row 77	Row 77	Row 77	Row 77	Row 77	Row 77
78	Row 78	Row 78	Row 78	Row 78	Row 78	Row 78	Row 78
79	Row 79	Row 79	Row 79	Row 79	Row 79	Row 79	Row 79
80	Row 80	Row 80	Row 80	Row 80	Row 80	Row 80	Row 80
81	Row 81	Row 81	Row 81	Row 81	Row 81	Row 81	Row 81
82	Row 82	Row 82	Row 82	Row 82	Row 82	Row 82	Row 82
83	Row 83	Row 83	Row 83	Row 83	Row 83	Row 83	Row 83
84	Row 84	Row 84	Row 84	Row 84	Row 84	Row 84	Row 84
85	Row 85	Row 85	Row 85	Row 85	Row 85	Row 85	Row 85
86	Row 86	Row 86	Row 86	Row 86	Row 86	Row 86	Row 86
87	Row 87	Row 87	Row 87	Row 87	Row 87	Row 87	Row 87
88	Row 88	Row 88	Row 88	Row 88	Row 88	Row 88	Row 88
89	Row 89	Row 89	Row 89	Row 89	Row 89	Row 89	Row 89
90	Row 90	Row 90	Row 90	Row 90	Row 90	Row 90	Row 90
91	Row 91	Row 91	Row 91	Row 91	Row 91	Row 91	Row 91
92	Row 92	Row 92	Row 92	Row 92	Row 92	Row 92	Row 92
93	Row 93	Row 93	Row 93	Row 93	Row 93	Row 93	Row 93
94	Row 94	Row 94	Row 94	Row 94	Row 94	Row 94	Row 94
95	Row 95	Row 95	Row 95	Row 95	Row 95	Row 95	Row 95
96	Row 96	Row 96	Row 96	Row 96	Row 96	Row 96	Row 96
97	Row 97	Row 97	Row 97	Row 97	Row 97	Row 97	Row 97
98	Row 98	Row 98	Row 98	Row 98	Row 98	Row 98	Row 98
99	Row 99	Row 99	Row 99	Row 99	Row 99	Row 99	Row 99
100	Row 100	Row 100	Row 100	Row 100	Row 100	Row 100	Row 100

The risk register for Tier 3 and Tier 4 organizations is a listing of identified risks and their recommended risk treatments, also known as “safeguards.” Each row represents one risk and its risk treatment recommendation. The parts of the risk register are:

- The column headers and guiding text to help the reader or risk assessor understand the information that is contained in the column.
- Attack models and threats that are actions within an attack path.
- The information assets within the attack path that are being analyzed.
- The CIS Controls text that helps the risk assessor consider controls that should be in place to protect information assets in the context of the attack path.
- How the organization implements the CIS Control (if they do) to protect the information asset, and any vulnerabilities that would allow the threat to compromise the asset.
- The risk evaluation, including the likelihood that the threat would succeed, the impacts to the mission, objectives, and obligations if it did, and the resulting risk score.
- The recommended implementation of CIS Controls that would reduce risks to an acceptable level, and the safeguard risk calculation to estimate the risk of that recommendation.

Attack Path Models

The Tier 3 and Tier 4 organization will develop a set of attack path models to document the detailed steps that foreseeable attacks would follow, and to identify the information assets that would be involved in that attack path. This helps the risk assessor evaluate risks based on how foreseeable threats behave. Attack path modeling allows risk assessors to ask questions such as, “How well positioned are we against this type of attack,” “Am I thinking of threats to assets the way an attacker would?” and “Is my risk evaluation for this asset based on the likelihood and impact of other information assets that would be involved in the attack?”

The attack path models worksheet that risk assessors will use to design attack paths is built upon the *CIS Community Attack Model*, a document provided by CIS® that associates CIS Controls with the stages of cybersecurity planning, detection, and defense. This section will demonstrate the process for modeling attack paths that assessors will use to analyze risks in the Tier 3 and Tier 4 organization's risk register. A template and examples of this process is provided in the supplementary document *CIS_RAM_Workbook*.

The attack path models worksheet is depicted in Figure 21.

Figure 21 - Attack Path Models

Community Attack Model (Top)									
The Community Threat Model (top) aligns the actions within an attack path with CIS Controls that would prevent or detect the actions.									
Attack Path Models (Bottom)									
Attack methods name foreseeable attacks, and describe the threats against assets that would occur in the attack path.									
CIS Controls (V6.0)	Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege	Internal Recon	Lateral Movement	Establish Persistence	Execute Mission Objectives
Identify	control of HW, SW inventory; Network logs	threat intelligence			control of administrative privileges	control of HW, SW inventory			Incident Response - Planning
Protect	firewall; mail gateway filtering; web filtering; manage ports, protocols, services; continuous vulnerability assessment	hardened configurations	continuous vulnerability assessment; firewall; mail gateway; secure remote access	patching; hardened configurations; HIPS; anti-malware; containerization; app whitelisting; Data Execution Prevention	control of admin privilege; data security; hardened configuration; continuous monitoring; control of admin privilege; audit logs; Configuration Monitoring	control of admin privilege; NW segmentation; Manage ports, protocols, services; NW segmentation	control of admin privilege; patching; hardened configurations; anti-malware; NW segmentation	egress filtering; control of HW, SW inventory	egress filtering; NW segmentation; data security
Detect	firewall; honeypot; Network authentication; Network logs	audit logs; threat intelligence	audit logs; Network system	Anti-malware; Intrusion Detection	account monitoring; control of admin privilege; audit logs; Configuration Monitoring	account monitoring; audit logs; Network Monitoring	audit logs; Network Monitoring	NW IDS; Host Intrusion Prevention	Data Execution Prevention; HIPS; Network Monitoring
Respond	honeypot			Incident Response - Execution	Incident Response - Execution			slashhole	Incident Response - Execution
Recover				Incident Response - Execution	Incident Response - Execution				Incident Response - Execution
				control of HW, SW inventory					control of HW, SW inventory
Attack Path Models									
	Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege	Internal Recon	Lateral Movement	Establish Persistence	Execute Mission Objectives
Arbitrary code execution through web application	Our web application is accessible, as is some information about the architecture of the application by reviewing web pages, code objects, and references to linked systems. Asset: Web application and verbose services on the web application stack.	Highly skilled hackers may develop scripts to execute commands through application or database services. Asset: Out of our control.	Attempts at running scripts or direct references to the web server, such as bash. Asset: Application server, database server, and event logs.	Commands executed through application account. Files added, altered, or replaced. Asset: Application server.	Establishment or alteration of existing account. Asset: User accounts.	Directory traversal at the web server. Asset: Application server, event logs.	Commands at the application server. Asset: Application server, event logs.	Installation of executables, establishment of new accounts. Asset: Operating systems, event logs, user accounts, administrative accounts.	Initiation of executables, daemons, services, processes. Asset: Executable processes, daemons, services, event logs.
Ransomware	Hackers determine who in the organization has access to sensitive information. Asset: Public information and social media sites that describe personnel and responsibilities.	Moderately skilled hackers may develop phishing email and ransomware exploits that target selected personnel. Asset: Out of our control.	Hacker sends phishing email to selected personnel. Asset: Email server, SMTP gateway.	Personnel open phishing email and trigger an install of the ransomware payload. Asset: Email client, end-user OS, personnel, proxy server, advanced malware appliance.	Malware encrypts the local storage volume. Asset: End-user OS, storage volume.	Not applicable	Not applicable	See Misuse/Escalate Privilege.	Hackers require payment for release of information back to us. Asset: Out of our control.

The attack path models worksheet is made up of two main parts; the Community Attack Model at the top, and the attack path models at the bottom. The Community Attack Model above depicts the stages of an attack as they relate to CIS Controls that could prevent or detect each stage. The attack path models below list types of information security incidents that may occur, and identify what actions and information assets would be involved in each stage of the incident.

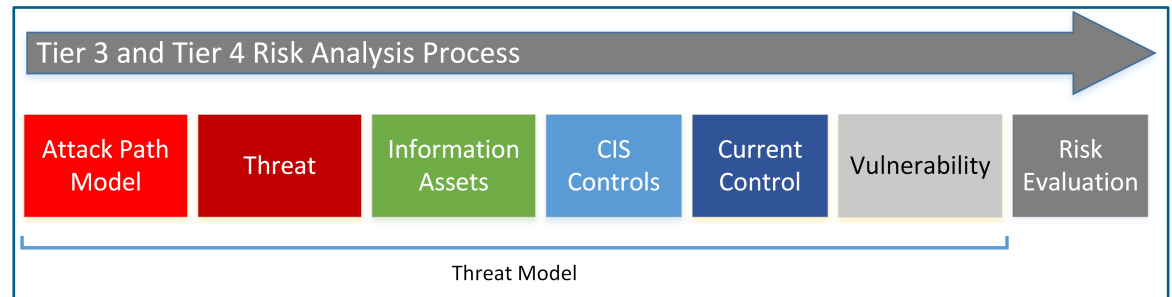
The Process

The risk assessor will start their assessment by modeling attack paths in the attack path model worksheet. The worksheet will result in a set of attack path models (one row per model), and will state the actions and assets that may be involved in each attack.

The risk assessor will then use the risk register to analyze each attack path (one stage in the attack path per row in the risk register). The assessor will evaluate each stage of an attack path just as they would analyze each risk in Tier 1 and Tier 2 assessments.

Tier 3 and Tier 4 organizations may wish to start their risk analysis by first analyzing risks using the processes described for Tier 1 or Tier 2 organizations, then by examining specific threats on an attack path basis. This ensures that all in-scope information assets and all CIS Controls will be addressed in the risk register, along with the more specific and detailed risks that are analyzed using this threat-based process.

Figure 22 - Risk Analysis Process for Tier 3 and Tier 4 organizations



The risk assessment process for Tier 3 and Tier 4 organizations ensures that information assets are analyzed in terms of the risk they pose when multi-phase attacks occur in the environment. This threat-based analysis is accomplished by following the steps listed below:

1. Using the Attack Path Model worksheet, the risk assessor creates a new row in the worksheet to name a type of cybersecurity attack or security incident, such as “Data seizure through web application,” or “Mis-delivery of patient information.”
2. The risk assessor then moves right across the new row to describe each stage of the attack or incident. The description would include an affected information asset, and how the attack would compromise the asset at each stage.
3. The risk assessor can then refer to each cell across an attack path row to populate a risk register. The risk assessor can copy the attack path model name, threats, and information assets that are in each attack path model row to the risk register, with one row per model / threat / asset grouping.
4. While considering the controls that should be addressed in each row of the risk register, the risk assessor should refer to the Community Attack Model grid at the top of the Attack Path Model worksheet to determine which controls should be in place to detect or prevent the threat.
5. The risk assessor will then review the CIS Controls listed in each risk register row, and will gather evidence for how well each asset is protected by safeguards that are associated with the CIS Control.
 - a. Evidence may be in the form of interviews, a review of configurations, or a review of records and logs.
6. The risk assessor will then describe how the control is applied to the information asset or asset class in the “Current Control” cell of that row.
7. Next, the assessor will consider the difference between the CIS Control and the currently applied safeguard to determine whether there is a deficiency in how the control is currently deployed and operating. If the current safeguard is not implemented as described or in a way that is likely deficient against the threat, then the assessor will state this as a vulnerability.
 - a. Risk assessors should consider the objective of the CIS Control as they analyze risks. For example, CIS Control 16.11 states “Automatically lock workstation sessions after a standard period of inactivity.” The control’s objective is to prevent unauthorized people from using unattended user sessions. If the current control does not meet the objective, then the assessor should state the gap as a vulnerability in the vulnerability cell, such as: “Unattended workstations may be used by personnel who are not authorized access to those systems, or to applications that are assigned to the absent user.”
8. Now consider the threat that could occur because of the vulnerability.

- a. The above vulnerability could be paired with this threat: “Malicious personnel may abuse the privileges of authorized users and may execute unauthorized transactions or data downloads.”
9. Next, estimate the likelihood that the threat would succeed, and the impact it may create.
 - a. Likelihood estimation can be difficult, but the risk assessment criteria was developed to provide some guidance in that estimation process. Guidance is provided in the Methods for Evaluating Likelihood section of the “Risk Analysis Techniques” chapter later in the document
 - b. Impact scores should provide estimates of the impact that such a threat would create. Consider the likelihood and impact scores as a pair. In other words, “What is the likelihood that this impact would result?” Examples below will provide further guidance.
10. The risk score will be automatically computed in the risk register by multiplying the likelihood score by the highest of the three impact scores.

Tier 3 and Tier 4 Risk Assessment Example 1 – Ransomware at Email Servers

The Tier 3 and Tier 4 organization has expressed their concern about ransomware and wants to know what their exposure is to it. They want to know what they are doing now to prevent ransomware and what other investments they need to make in order to be appropriately protected. Additionally, the organization is concerned about two web application vulnerabilities, one that would allow data seizure, and the other that will allow arbitrary code execution through vulnerable sites.¹⁹

The risk assessor knows to consider each of these concerns as attack path models, and sets out to document each one.

By using the Attack Path Model worksheet, the risk assessor creates a row titled, “Ransomware” and begins documenting the attack path as shown in Table 63. The table is shown in a vertical format for ease of display in this document, but is in horizontal format in the template provided in *CIS_RAM_Workbook*.

Each stage of an attack path model is described using the Community Attack Model format. For each stage in the model the risk assessor will describe how the attack will compromise an information asset.

¹⁹ This document will only explore how the first attack path model will be defined and risk assessed. However, the other two scenarios are provided in the Attack Path Model worksheet as further examples of the attack path modeling process.

Table 63 - Attack Path Model (Ransomware)

Attack Path Stage	Attack Path Action
Attack Path Model	Ransomware
Initial Recon	<p>Hackers determine who in the organization has access to sensitive information.</p> <p>Asset: Public information and social media sites that describe personnel and responsibilities.</p>
Acquire/Develop Tools	<p>Moderately skilled hackers may develop phishing email and ransomware exploits that target selected personnel.</p> <p>Asset: Out of our control.</p>
Delivery	<p>Hacker sends phishing email to selected personnel.</p> <p>Asset: Email server, SMTP gateway.</p>
Initial Compromise	<p>Personnel open phishing email and trigger an install of the ransomware payload.</p> <p>Asset: Email client, end-user OS, personnel, proxy server.</p>
Misuse/Escalate Privilege	<p>Malware encrypts the local storage volume.</p> <p>Asset: End-user OS, storage volume.</p>
Internal Recon	Not applicable
Lateral Movement	Not applicable
Establish Persistence	See Misuse/Escalate Privilege.
Execute Mission Objectives	<p>Hackers require payment for release of information back to us.</p> <p>Asset: Cash or data.</p>

As a result of this detailed description of the attack path, the risk assessor can now itemize each of these stages in the risk register as the basis of their risk analysis. So they create a set of rows in their risk register that include the information as displayed in the partial risk register shown in Table 64 and provided more fully in *CIS_RAM_Workbook*.

Table 64 - Partial Risk Register with Attack Path Model

Attack Path Model	Threat	Information Asset
Ransomware	Initial Recon: Hackers determine who in the organization has access to sensitive information.	Public information and social media sites that describe personnel and responsibilities.
Ransomware	Delivery: Hacker sends phishing email to selected personnel.	Email server, SMTP gateway.
Ransomware	Initial Compromise: Personnel open phishing email and trigger an install of the ransomware payload.	Email client
Ransomware	Initial Compromise: Personnel open phishing email and trigger an installation of the ransomware payload.	End-user OS
Ransomware	Initial Compromise: Personnel open phishing email and trigger an install of the ransomware payload.	Personnel
Ransomware	Initial Compromise: Personnel open phishing email and trigger an install of the ransomware payload.	Proxy server
Ransomware	Misuse/Escalate Privilege: Malware encrypts the local storage volume.	End-user OS
Ransomware	Misuse/Escalate Privilege: Malware encrypts the local storage volume.	Storage volume
Ransomware	Execute Mission Objectives: Hackers require payment for release of information back to us.	Cash or data

Each row in this risk register establishes a relationship between the attack path model (in this case, “Ransomware”), a threat that could occur at each stage of the ransomware attack, and the information asset or asset class it would occur on. Note that items labeled “Not applicable” and “Not in our control” in Table 63 are not provided rows in the partial risk register in Table 64. This is because there is little the organization can do to address these steps in the attack path for the ransomware scenario.

Also note that many information assets and many kinds of threats can be considered within an attack path. The risk assessor must determine the amount of detail and variety of asset/threat pairings they intend to include in their assessment. Attack path modeling can take significant time. Because of this, risk assessors will need to consider the amount of time and resources they have available to conduct their analysis, and must select a degree of detail based on that

availability. Starting with obvious assets and threats for the first assessment may be enough, knowing that in subsequent assessments more variety can be added to the model.

As a result of this analysis, the organization scours the Internet for mentions of privileged personnel and their roles. They remove as much sensitive information from those sites as they can. They realize that there are other ways to target privileged personnel, but this is considered a prudent step.

The first risk that the assessor analyzes in this attack path is the use of the email server and SMTP gateway that may receive and relay a phishing message to a targeted end-user. The organization believes they have good safeguards to manage this risk, but they check the Community Attack Model to be sure.

Figure 23 - CIS Control Selection from Community Attack Model

		Attack Stages				
Controls	Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege	
Functions	Identify	control of HW, SW inventory; Network logs	threat intelligence		control of administrative privilege	
	Protect	firewall; mail gateway filtering; web filtering; manage ports, protocols, services; continuous vulnerability assessment	hardened configurations	continuous vulnerability assessment; firewall; mail gateway filtering; web filtering; secure remote access; NIPS	patching; hardened configurations; HIPS; anti-malware; containerization; app whitelisting; Data Execution Protection	control of admin privilege; data security; hardened configuration; continuous vulnerability assessment
	Detect	firewall; honeypot; Network authentication; Network logs	audit logs; threat intelligence	audit logs; Anti-malware; Network Intrusion Detection system	HIPS; anti-malware; containerization; app whitelisting; Data Execution Prevention;	account monitoring; control of admin privilege; audit logs; Configuration Monitoring
	Respond	honeypot			Incident Response - Execution	audit logs; Configuration Management; Account Management
	Recover				Incident Response - Execution; control of HW, SW inventory	

This risk was identified in the attack path model in Table 63 in the “delivery” stage, so while assessing the risk the risk assessor will review the email server and SMTP gateway because of their delivery role in the attack, and will name them as assets in the risk register, as shown in Table 64. As the assessor references the Community Attack Model, they will look at the intersection between the *Delivery* column and the *Protect* row to find “Continuous vulnerability assessment,” “firewall,” “mail gateway filtering,” “web filtering,” “secure remote access,” and “NIPS (network intrusion prevention system).”

Considering the threat of email targeting specific users, and the organization’s use of email filtering and sandboxing on their corporate server, the risk assessor reviews sub-controls under CIS Control 7 “Email and Web Browser Protections.” Among those sub-controls the risk assessor considers CIS Control 7.8 “Implement DMARC and Enable Receiver-Side Verification” and CIS Control 7.10 “Sandbox all Email Attachments.” They know that DMARC controls related to CIS Control 7.8 rely on more community cooperation before they can be reliable, and even then could be by-passed by determined attackers. They also know that determined attackers can get past

the organization's email sandboxing technologies. They forego analyzing this risk because they also use anti-malware on their email server and SMTP gateway, which appears one row down the Delivery column in the Detect row.

The risk assessor decides to analyze the risk involved with their anti-malware email module by referring to CIS Control 8.1, as described in Table 65.

Table 65 - Example Risk Analysis for Email Server in a Ransomware Attack Path Model

Attack Path Model	Ransomware
Threat	Hackers may target personnel using their personal email accounts, thus bypassing the corporate email server.
Information Asset	Email server and SMTP gateway
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Control	Advanced malware detection and prevention operates within the SMTP gateway. It detects and quarantines attachments and hyperlinks associated with malicious files and suspicious or blocked URLs.
Vulnerability	End-users may be victims of phishing over personal email services that they can access from offices and on work computers.
Threat Likelihood	
Mission Impact	
Objectives Impact	
Obligations Impact	
Risk Score	
Risk Acceptability	

Note that the assessor has not yet evaluated this risk. Rather, the risk assessor will evaluate each risk in the attack path after considering how that set of risks affect each other. For our example attack path model in Table 63, all nine risks will be written out before each one is evaluated. This is done to ensure that the likelihood and impact of each risk is based on a foreseeable scenario, and not in isolation of each asset which may cause some risks to be estimated arbitrarily high or low.

Working further down the attack path model in Table 63, the organization next considers the risk they may suffer at desktop email clients which are targeted during the *Initial Compromise* in this ransomware attack path. After considering the vulnerability in the previous risk analysis that end-users may still receive phishing messages through their personal email accounts, email client risks are fresh in the risk assessor's mind. The risk assessor reviews the Community Attack Model in Figure 24 to identify CIS Controls that intersect between *Initial Compromise* and *Protect* and they see that for this stage anti-malware and CIS Control 8.1 is again an appropriate control to include in their evaluation. They describe the risk associated with their implementation of CIS Control 8.1 at desktop email clients in Table 66.

Figure 24 - CIS Control Selection from Community Attack Model

		Attack Stages				
	Controls	Initial Recon	Acquire/Develop Tools	Delivery	Initial Compromise	Misuse/Escalate Privilege
Functions	Identify	control of HW, SW inventory; Network logs	threat intelligence			control of administrative privilege
	Protect	firewall; mail gateway filtering; web filtering; manage ports, protocols, services; continuous vulnerability assessment	hardened configurations	continuous vulnerability assessment; firewall; mail gateway filtering; web filtering; secure remote access; NIPS	patching; hardened configurations; HIPS; anti-malware; containerization; app whitelisting; Data Execution Protection	control of admin privilege; data security; hardened configuration; continuous vulnerability assessment
	Detect	firewall; honeypot; Network authentication; Network logs	audit logs; threat intelligence	audit logs; Anti-malware; Network Intrusion Detection system	HIPS; anti-malware; containerization; app whitelisting; Data Execution Prevention;	account monitoring; control of admin privilege; audit logs; Configuration Monitoring
	Respond	honeypot			Incident Response - Execution	audit logs; Configuration Management; Account Management
	Recover				Incident Response - Execution; control of HW, SW inventory	

Table 66 - Example Risk Analysis for Email Client in a Ransomware Attack Path Model

Attack Path Model	Ransomware
Threat	Personnel open phishing email and trigger an installation of ransomware payload. Email may be received through personal email accounts.
Information Asset	Email client
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Control	Signature-based anti-virus on each desktop. Anti-virus filters suspicious web URLs using a dictionary that is updated monthly.
Vulnerability	Advanced malware prevention is not included in endpoint protection applications on end-user workstations (other than URL filtering). End-users may be victims of phishing over personal email services that they can access from offices and on work computers.
Threat Likelihood	
Mission Impact	
Objectives Impact	
Obligations Impact	
Risk Score	
Risk Acceptability	

This risk may evaluate as unacceptably high when the whole attack path is considered. Moreover, this risk appears to be independent of the first risk involving the email server and SMTP gateway. As robust as the corporate SMTP server may be at preventing ransomware phishing, the assessor is concerned that the risk of ransomware phishing may still be unacceptably high because the SMTP gateway only protects corporate email accounts, not personal email accounts hosted by other services.

The risk assessor named a proxy server as an information asset at the “initial compromise” stage in the attack path model in Table 63 because their proxy server blocks outbound requests for known-bad IP addresses and domains, including known malware hosts. So the risk assessor again references the Community Attack model, looking at the *Initial Compromise* column and sees that the “web filtering” function of the proxy server is not in that column. Moreover, the organization is not using the remaining controls in the *Protect* and *Detect* rows of that column, so they cannot refer to those controls in their “Current Controls” column.

The Community Attack Model, while being very helpful to organizations that model attack paths, is a working document that will constantly develop as threat behaviors change, and as controls change to meet new challenges. In the case of this organization, they identify the role of a proxy server (a web filtering tool) for *Protecting* against the *Initial Compromise* of a malware attack. The proxy server prevents the malware from downloading payload from a known-bad web resource. But because the Community Attack Model does not reference web filtering in the intersection of *Protect* and *Initial Compromise*, the organization adds that control to that cell. They have identified a case for using web filtering for disrupting ransomware and should record it for future use.

The risk assessor decides to evaluate the risk associated with the proxy server to see if they may help reduce the risk appropriately. The assessor models the risk in Table 67.

Table 67 - Example Risk Analysis for Proxy Server in a Ransomware Attack Path Model

Attack Path Model	Ransomware
Threat	Personnel open phishing email and trigger an installation of the ransomware payload. Email may be received through personal email accounts.
Information Asset	Proxy server
CIS Control	7.4
Description	Enforce network-based URL filters that limit a system's ability to connect to websites not approved by the organization. This filtering shall be enforced for each of the organization's systems, whether they are physically at an organization's facilities or not.
Control	All Internet traffic for systems within corporate LANs and DMZ have URLs filtered against a subscription service that blocks sessions with known-bad hosts, and blocks URLs not categorized as safe by that service.
Vulnerability	Laptops and mobile devices bypass the proxy server when used outside of the corporate network. Ransomware may attack systems when they are out of the corporate network.
Threat Likelihood	
Mission Impact	

Attack Path Model	Ransomware
Objectives Impact	
Obligations Impact	
Risk Score	
Risk Acceptability	

The risk assessor sees that the proxy server appears more robust than the end-point protection at end-user systems, but it still has shortcomings. The proxy server is not able to enforce URL blocking on systems that are not on the corporate network when a ransomware phishing attack occurs.

While all nine risks in the attack path would be evaluated as a set for the actual risk assessment, this section will evaluate the three example risks to demonstrate the group evaluation process. The remainder of the risks are fully evaluated in the worksheet “Risk Register – Tier 3 & 4” in the *CIS_RAM_Workbook*.

Recall the definitions for the impact and likelihood scores that the Tier 2, and Tier 3 and Tier 4 organizations created. The impact scores are shown in Table 68.

Table 68 – Example Impact Definitions

Impact Score	Impact to Mission <i>Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Operate profitably.</i>	Impact to Obligations <i>Patients may be harmed if information is compromised.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

Likelihoods were similarly defined with five potential scores, as shown in Table 69.

Table 69 – Example Likelihood Definitions

Likelihood Score	Foreseeability
1	Not foreseeable. This is not plausible in the environment.
2	Foreseeable. This is plausible, but not expected.

Likelihood Score	Foreseeability
3	Expected. We are certain this will eventually occur.
4	Common. This happens repeatedly.
5	Current. This may be happening now.

The organization will now go back and review the risks associated with the ransomware attack path to estimate the likelihood and impact of each threat, but in consideration of the other ransomware risks.

The risk assessor will review those risks side-by-side in the abbreviated table below.

Table 70 – Comparative Risks in a Ransomware Attack Path Model

Attack Path Model	Ransomware		
Information Asset	Email Server	Email Client	Proxy Server
Threat	Hackers may target personnel using their personal email accounts, thus bypassing the corporate email server.	Personnel open phishing email and trigger an installation of ransomware payload. Email may be received through personal email accounts.	Personnel open phishing email and trigger an installation of the ransomware payload. Email may be received through personal email accounts.
CIS Control	8.1	8.1	7.4
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.	Enforce network-based URL filters that limit a system's ability to connect to websites not approved by the organization. This filtering shall be enforced for each of the organization's systems, whether they are physically at an organization's facilities or not.
Control	Advanced malware detection and prevention operates within the SMTP gateway. It detects and quarantines attachments and hyperlinks associated with malicious files and suspicious or blocked URLs.	Signature-based anti-virus on each desktop. Filtering of suspicious web URLs using a dictionary that is updated monthly.	All Internet traffic for systems within corporate LANs and DMZ have URLs filtered against a subscription service that blocks sessions with known-bad hosts, and blocks URLs not categorized as safe by that service.
Vulnerability	End-users may be victims of phishing over personal email services that they can access from offices and on work computers.	Advanced malware prevention is not included in endpoint protection applications on end-user workstations (other than URL filtering). End-users may be victims of phishing over personal email services that they can access from offices and on work computers.	Laptops and mobile devices bypass the proxy server when used outside of the LAN. Ransomware may attack systems when they are out of the office LAN.
Threat Likelihood	2	3	3
Mission Impact	3	3	3
Objectives Impact	4	4	4

Attack Path Model	Ransomware		
Obligations Impact	4	4	4
Risk Score	8	12	12
Risk Acceptability	Acceptable	Unacceptable	Unacceptable

After evaluating each of the risks in the attack path model, the risk assessor (and by extension, their organization) is comfortable with the ability of the SMTP gateway to protect users from ransomware phishing attacks that pass through the corporate email server. But they are less comfortable with the risk of those attacks coming in from personal email accounts while end-users use their laptops away from the office. The risk analyses for the email client and proxy server are identical in this case and are shown in Table 71.

Table 71 - Example Risk Estimation

Threat Likelihood	Mission Impact	Objectives Impact	Obligations Impact	Risk Score
3	3	4	4	12

The risk score is the product of the likelihood score and the higher of the three impact scores, which in this case is '3 x 4 = 12'.

Also recall that the risk acceptance criteria for the Tier 3 and Tier 4 organization looks like this:

Table 72 - Risk acceptance criteria

Impact Threshold	x	Likelihood Threshold	=	Risk Threshold
3	x	3	=	9
... therefore ...				
Acceptable Risk			<	9

An acceptable risk would be one that evaluates to anything below '9'. But the risk of ransomware is as high as '12' and is therefore unacceptable.

By analyzing additional risks in the ransomware attack path, such as protections at the end-user's operating system or storage volume, the organization may further mitigate these three risks by other safeguards, such as timely and reliable data backups, or logical controls that prevent sensitive data from being accessed by laptops. But what is known is that in terms of ransomware, there is a continuing risk to the organization that has not been resolved by the SMTP gateway and proxy server.

Exercise:

The reader should refer to the template Risk Register – Tier 2 that is provided in the supplementary document *CIS_RAM_Workbook*. They may use the risk register template to enter a set of risks that are associated with the attack path models, CIS Controls, and information assets that are in scope of their assessment.

While doing this exercise, the reader should consider:

1. That threat-based analysis requires considerable effort and may be implausible as a method for conducting a complete, comprehensive risk assessment.
2. Conducting threat-based analysis within a risk register that was completed using a Tier 1 or Tier 2 approach. Threat-based analysis can be used to answer specific risk questions, such as:
 - a. How realistic is a specific risk score that appears to exaggerate or downplay a stand-alone risk?
 - b. Are there risks we have not already considered in our environment?
 - c. How well positioned are we to protect ourselves against a specific threat?
 - d. What is the most effective investment we can make to reduce a single risk that would only be realized within a threat path?
3. Not “boiling the ocean.” Not all assets can be practically evaluated against all applicable CIS Controls in a single assessment. Prioritize evaluating threats that appear more likely than others, due to past experience or other research.
4. Whether a control or information asset requires examination to understand its actual configuration and effectiveness.
5. Whether the organization can tolerate the amount of effort and time that the risk assessment requires.
 - a. The organization should use high-level analysis (review of policies and interviews) if they do not have extensive time and resources.
 - b. Information assets should be tested and examined in more detail as time allows.
 - c. The organization should plan recurring risk assessments to identify more risks over time.
6. Collaborating with information security experts to help model threats that are foreseeable in the environment, and to help evaluate the effectiveness of current safeguards.

The risk assessor will need to use their professional judgment to select the controls and information assets and to model threats that should be analyzed in the risk assessment. Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Tier 3 Risk Assessment Summary

After having analyzed a set of risks against a single information asset, the Tier 2 organization realizes the advantages of gaining a more comprehensive view of its risks:

1. Modeling threats through attack paths enables Tier 3 and Tier 4 organizations to evaluate cybersecurity risks more comprehensively than by viewing information assets individually.
2. The risk involved at one asset will influence the risk of other assets, which is a more realistic picture of risk within a networked environment.
3. Because attack paths are aligned with the Community Attack Model, risk assessors are assisted by the CIS community in identifying the CIS Controls that are best suited for preventing and detecting attacks at various stages and assets in the attack path.

After evaluating the risks against information assets we have identified many that were unacceptably high and that should be provided with recommended safeguards to reduce their risk. We will walk through and illustrate this process in the Risk Treatment Recommendations section below.

Risk Treatment Recommendations

Introduction

Organizations often think of security safeguards as obstacles to business and productivity. Safeguards often cause personnel to take extra steps to get access to systems or information, or to get approval for normal business activities. Safeguards require investments in time and money, which compete with other priorities. And if they become too disruptive to an organization's mission and objectives, security safeguards can become disliked and avoided.

In fact, disruptive safeguards often cause personnel to work around them just to get their jobs done, which creates more risk.

But risk treatment recommendations can and should result in safeguards that are demonstrably reasonable. And while obtaining a clear definition for "reasonable safeguards" has been a challenge in the legal, regulatory, and information security communities, the CIS RAM provides a practical solution. Risk assessors evaluate risk treatment recommendations to determine whether a security safeguard is reasonable by; comparing the safeguard to the risk it is meant to reduce, and by comparing the safeguard to the risk acceptance criteria.

Risk treatment recommendations are simple to evaluate once the risk assessment criteria and initial risk analysis have been established. The process occurs over the following steps:

1. While examining an unacceptably high risk, review the CIS Control that corresponds with the risk and recommend a feasible way for the organization to implement or improve that control.
2. If that control is not feasible in the near-term, recommend other CIS Controls related to the risk that can be used to reduce it.
3. Evaluate the risk of the recommended safeguard to understand the burden it would pose to the organization. Then compare that safeguard risk to the risk acceptance criteria to determine whether it is appropriate.
4. Also compare the evaluated risk of the recommended safeguard to the observed risk to determine whether the safeguard is reasonable (safeguards with lower risk scores than the observed risk are reasonable.)
5. Sort the risks by their risk score to prioritize the risks and risk treatments that the organization will invest in.

This section demonstrates these steps in detail by describing the process, then by modeling risk treatments for the unacceptably high risks that were evaluated in previous sections.

The reader should review the definitions of ‘reasonable’ and ‘appropriate’ that are provided in the glossary. These terms will be used regularly in this section and have distinct meanings.

1. Appropriate: A condition in which risks to information assets will not foreseeably create harm that is greater than the organization or its constituents can tolerate.
2. Reasonable: A condition in which safeguards will not create a burden to the organization that is greater than the risk it is meant to protect against.

Risk Treatment Objectives

The objective of well-formed risk treatment recommendations is to create a prioritized list of information security safeguards that would provide appropriate protections while not posing too great a burden on the organization’s purpose.


The risk treatment recommendation exercises that are demonstrated in this section examine the unacceptable risks that were illustrated earlier in the document and will select CIS Controls that would reduce those risks to a degree that is both reasonable (not overly burdensome) and appropriate (not unacceptably harmful).

Recommending Risk Treatment Safeguards from CIS Controls V7

As we examine unacceptably high risks, we will recommend safeguards that are based on CIS Controls V7. But some of the safeguards that an organization is prepared to implement and operate may not be implemented exactly as described in CIS Controls V7. This process takes into account how to select controls that address risks, and how to determine whether they are designed in a way that makes sense in context of both the risk, and the potential burden to the organization.

Recall the relationship between analyzed risks and their recommended risk treatments in Figure 25.

Figure 25 - Balance Within Core Risk Analysis



A risk and its proposed safeguard are both evaluated using the same criteria. If a proposed safeguard has a higher risk (its “safeguard risk”) than the risk acceptance criteria, then it’s not appropriate. If the safeguard has a higher score than the observed risk, then it’s not reasonable.

The exercises in this section will focus on matching completed risk analyses (in blue) with newly recommended safeguards (in green).

Background – How Realistic Are Safeguard Risk Estimates?

Critical readers will question how the organization and their risk assessor will know whether their safeguard risk estimations are realistic. After all, how can they know prospectively what their risk would be in such a situation?

There are two important items to keep in mind while gaining comfort with this practice; understanding the legal and regulatory expectations for risk management, and information security standards for evaluating safeguards after they've been implemented.

Law and Regulation: Laws and regulations generally expect risk analysis to evaluate safeguards that are required for achieving compliance, and expect that the risk analysis is performed by appropriately skilled and informed people. These analyses do not guarantee security that is sufficient against any threat, but they do provide a plan toward improved security and compliance that is prioritized by the likelihood of harm, and that has no intolerable harm as its goal.

Information Security Risk Management Standards: Information security risk assessment standards that the CIS RAM is based on operate within fuller risk management programs and cycles. ISO 27005 operates within the ISO 27000 family of standards, and NIST 800-30 works within the NIST Special Publications. Each of these families of standards requires continuous analysis of security safeguards, including analysis of controls after they've been implemented to determine whether they are effective at addressing their security objectives. Recommended safeguards should therefore be risk assessed again after implementation to be sure that they achieve their intended objectives.

The Tier 3 and Tier 4 organization identified two unacceptable risks by modeling a ransomware attack path over several information assets. The first of these unacceptable risks involved email clients that personnel use to access personal email, and how that exposed the organization to ransomware phishing attacks. The risk is shown again in Table 73.

Table 73 - Example Risk Analysis for Email Client in a Ransomware Attack Path Model

Risk Analysis	Value
Threat	Personnel open phishing email and trigger an installation of ransomware payload. Email may be received through personal email accounts.
Information Asset	Email client
CIS Control	8.1
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.
Control	Signature-based anti-virus on each desktop. Filtering of suspicious web URLs using a dictionary that is updated monthly.
Vulnerability	Advanced malware prevention is not included in endpoint protection applications on end-user workstations (other than URL filtering). End-users may be victims of phishing over personal email services that they can access while working from home using work computers.
Threat Likelihood	3

Risk Analysis	Value
Mission Impact	3
Objectives Impact	4
Obligations Impact	2
Risk Score	12
Risk Acceptability	Not Acceptable

But because this risk was identified while evaluating an attack path, the organization considers the recommended safeguards in that same context. They compare this risk with the other unacceptable risks in the attack path in Table 74 to determine whether one safeguard would address both risks in the attack. Both risks have the same risk score, but for different reasons: Laptops are not protected against advanced malware through the end-point protection software, and laptops do not benefit from the proxy server while out of the office.

They can either add advanced malware protection to their endpoints, or they can extend the proxy server to the DMZ and force laptops to resolve network queries through that proxy server while out of the office.

So they model these options below.

Note: The risk assessor will record how they will address their risk by stating either “Accept,” “Reduce,” “Transfer,” or “Avoid.” *Accepting* and *reducing* risks will be intuitive to the reader. An organization may *transfer* a risk by contracting a third party that may handle the risk better, or by acquiring an insurance policy against the risk. The organization may also *avoid* the risk by no longer engaging in the processes, or handling the information assets that cause the risk

Table 74 – Comparative Risks in a Ransomware Attack Path Model

Attack Path Model	Ransomware	
Information Asset	Email Client	Proxy Server
Threat	Personnel open phishing email and trigger an installation of ransomware payload. Email may be received through personal email accounts.	Personnel open phishing email and trigger an installation of the ransomware payload. Email may be received through personal email accounts.
CIS Control	8.1	7.6
Description	Utilize centrally managed anti-malware software to continuously monitor and defend each of the organization's workstations and servers.	Enforce network-based URL filters that limit a system's ability to connect to websites not approved by the organization. This filtering shall be enforced for each of the organization's systems, whether they are physically at an organization's facilities or not.
Control	Signature-based anti-virus on each desktop. Filtering of suspicious web URLs using a dictionary that is updated monthly.	All Internet traffic for systems within corporate LANs and DMZ have URLs filtered against a subscription service that blocks sessions with known-bad hosts, and blocks URLs not categorized as safe by that service.
Vulnerability	Advanced malware prevention is not included in endpoint protection applications on end-user workstations (other than URL filtering). End-users may be victims of phishing over personal email services that they can access from offices and on work computers.	Laptops and mobile devices bypass the proxy server when used outside of the LAN. Ransomware may attack systems when they are out of the office LAN.

Attack Path Model	Ransomware	
Threat Likelihood	3	3
Mission Impact	3	3
Objectives Impact	4	4
Obligations Impact	4	4
Risk Score	12	12
Risk Acceptability	Unacceptable	Unacceptable
Risk Treatment Option	Reduce	Reduce
Recommended Safeguard	Add advanced malware protection module to end-point protection.	Extend proxy server to the DMZ and force laptops to use it as a gateway.
Safeguard Risk	<p>Unexpected cost would be within the budget plan threshold if modules are restricted to laptops this year, and extended to remaining system next year.</p> <p>Threat of malware would no longer be expected. No impact to our mission.</p>	<p>Laptops that use personal VPNs may bypass the proxy service.</p> <p>If the proxy service is unavailable, it may cause users to not use Internet resources while working out of the office.</p> <p>Proxy servers are not able to detect local attacks on systems.</p>
Safeguard Threat Likelihood	2	3
Safeguard Mission Impact	1	2
Safeguard Objectives Impact	2	2
Safeguard Obligations Impact	1	4
Safeguard Risk Score	4	12
Safeguard Risk Acceptability	Acceptable	Unacceptable

The Tier 3 and Tier 4 organization now has the information it needs to determine and document why their recommended safeguard is to add advanced malware prevention at their laptops first, then desktops in the following fiscal year. Desktops will be covered well by the proxy server when operated in their permanent home – the office network.

Exercise:

The reader should use the template “Risk Register – Tiers 3 or 4” that is provided in the supplementary document *CIS_RAM_Workbook* to enter risk treatment recommendations for each risk that evaluated as unacceptably high.

The reader should consider:

1. Whether an existing safeguard can be improved, and how that would be done.
2. Whether a safeguard based on a different CIS Control would provide reasonable and appropriate risk.
3. Collaborating with information security subject matter experts to help model the potential effectiveness of recommended safeguards.

The risk assessor will need to use their professional judgment to design and recommend information security safeguards, and to evaluate prospectively the risk that they may pose. Information security experts may need to be included in the process to ensure that the risk analysis is conducted appropriately.

Risk Treatment Recommendations Summary

Risk treatment recommendations are a critical part of risk assessment to be sure that the organization has developed a plan for addressing risks without creating other risks to the organization or its constituents. Some of the benefits that have been demonstrated about this process are:

1. Organizations can demonstrate to collaborating business managers how recommended security safeguards can be implemented without taxing the business purpose by evaluating the risk of the safeguards against the mission and objectives of the organization.
2. Organizations can demonstrate to regulators and other legal authorities that safeguards are reasonable because the expected risk of the safeguard (the “burden” to the organization) is not greater than the risk that it reduces.
3. Organizations can demonstrate that recommended safeguards would be “appropriate” by showing that they would not foreseeably create an impact that would be intolerable to the organization or its constituents.
4. Recommended risk treatments can be considered in terms of the attack path for Tier 3 and Tier 4 organizations. As the maturity of an organization’s security capabilities grows, so does the sophistication and perhaps efficiency of their risk treatment recommendations.

The process for evaluating risks and for recommending appropriate risk treatments has been demonstrated at a general level. However, some questions likely remain for the reader for evaluating safeguards, estimating likelihood, and the suitability of probability models in risk analysis. These more detailed topics are presented in the next chapter.

Chapter 5: Risk Analysis Techniques

The instructions, examples, and templates described in this section are best understood through experience. The reader will benefit from using the examples that are provided in the supplementary document *CIS_RAM_Workbook* to best understand the instructions in this chapter.

Risk Analysis Techniques

Introduction

The example risk assessment processes described in this document are broadly applicable to many cases and environments. However, there are many reasons why an organization will modify the processes and templates that are provided in the CIS RAM.

Methods for estimating likelihood or probability, assessing safeguards and policies, considering risks of non-technical safeguards, and determining which risks to assess or to ignore all present opportunities for customizing risk assessments to specific environments.

This section describes several customization methods for analyzing risks that organizations may consider as part of their cybersecurity risk assessments.

Defining Impacts for Tier 1 organizations

Perhaps the most important early step in the risk assessment is to develop effective impact definitions. The CIS RAM is based on Duty of Care Risk Analysis principles to enable organizations to make conscientious evaluations of their current and intended risk. A risk assessment that results from the CIS RAM should show whether information security safeguards are appropriate to the public, while being reasonable to the organization. The core of this analysis is the impact definitions, and the balance and consensus that they are meant to establish.

This section will provide guidance for defining impact types effectively.

Table 75 – Summary Evaluation of Impact Definition

Benefits	Limits
<ul style="list-style-type: none"> - Consistent method to evaluate risk impacts. - Satisfies “cost-benefit” analysis that regulators use. - Satisfies “duty of care balance test” that courts rely on. - Balances business interests with public interest. 	<ul style="list-style-type: none"> - Poorly defined impact definitions can frustrate risk assessors. - Poorly “balanced” impact definitions may not reduce legal liabilities.

The primary purpose of impact definitions is to provide risk assessors with a consistent method for scoring risks that is fair to all potentially affected parties. A risk evaluation should demonstrate as much concern for the organization as it does for others.

To make this possible, impact definitions should be designed with the concept of balance firmly in mind.

Recall the impact definitions used for Tier 1 organizations shown in Figure 26. Two columns address the interests of the organization’s purpose and the parties who may be effected by

information security risk. Each of these columns is considered an “impact type.” The “Impact to Our Mission” column addresses the main purpose for the two parties to enjoin in the risk ... the beneficial reason for customers and the organization to share information. The safety of the organization’s patient customers is considered in the “Impact to Obligations” column.

Figure 26 - Example Impact Definitions

Impact Score	Impact to Our Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Our Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	No harm would come to patients.
2	Some patients cannot access the information they need for good outcomes.	Few patients may be harmed after compromise of information or services.
3	We can no longer provide helpful information to remote patients.	Many patients may be harmed financially, reputationally, or physically, up to including death.

Now consider the impact scores, and the red boundary that separates score ‘1’ from scores ‘2’ and ‘3’. This boundary marks the division between impacts that would presumably be acceptable to all parties, and those that would not be.

Consider how impacts for score ‘1’ would appear to the organization, patient customers, and legal authorities. The organization using this impact definition is stating that they would accept impacts from threats if they result in conditions similar to how score ‘1’ is defined. This would indicate that; patients would continually be able to access the information they needed, and patients would not be foreseeably harmed as a result of a threat.

Then consider how impacts for score ‘2’ are defined. Threats that would foreseeably result in an impact score of ‘2’ could mean that some patients who cannot access information may not maintain good health outcomes. That scenario is clearly an unacceptable impact to the organization’s mission, and would not make it worthwhile for patients to entrust their information with that organization. The impact score of ‘2’ for obligations would be unacceptable because some of the patient customers could foreseeably be harmed financially or reputationally as a result of an incident – presumably because of identity theft or a system outage.

All of these features of an impact definition make it effective for estimating risk in a way that is equitable to all potentially affected parties, and even for driving consensus within the organization that uses it. After all, the interests and purpose of the organization are addressed as well as the interests of the public.

Organizations can build effective impact definitions by thoughtfully defining their impact areas, and then by carefully defining their impact scores.

Defining Impact Areas

The CIS RAM describes two impact areas that should be included in a Tier 1 organization’s risk evaluation; Impact to Mission and Impact to Obligations. Each of these impact areas address the interests of people or organizations who may be affected by information security risk. They each play a significant and unique role in analyzing risk, and should be defined within those roles, as described below.

Defining Mission

Definition: An organization's mission is the value it provides others, and that requires that they engage in risk together to achieve that value. A college educates its students, but students must give personal and financial information to those colleges in order to receive the education. Retailers provide products to customers, but in many cases customers turn over their financial information to those organizations to receive those products. Cloud service providers offer Internet-based functionality to business customers, but those customers must turn over business process or operations to those services as a result. So "mission" is a way of asking, "What's in it for the others?" who engage in risk with the organization.

Example Mission Definitions should have the following characteristics:

- Concisely state the benefit the organization provides that encourages others to enjoin them in information security risk.
- Convey a simple fact that can be observed and measured.
- Describe something that the organization already manages to, and that personnel will recognize as important to the organization.

Example 1: A custom manufacturer uses their customers' intellectual property to quickly make components that are perfect upon delivery. Their mission definition may be, "To provide customers with products that meet their unique specifications, without fail." The message is simple, it can be measured, and the organization likely recognizes this as something that they manage to. Moreover, the definition can be useful when a risk assessor tries to determine what core values may be harmed is a risk were to occur, or if a safeguard is too burdensome to the mission. If the risk assessor recommends a control that prevents customers from sending their intellectual capital, or that prevents the manufacturer from storing it or sharing it among drafter and engineers, then the mission would clearly be negatively impacted.

Example 2: A community bank states their mission as, "We promote opportunities to households and small businesses by providing affordable financial products and advisory services." They could say that their mission is to lend and borrow money, but they are thinking ahead to what they do not want to compromise about that mission, and that is to serve their community. But again, they have a concise definition that states why others would enjoin in risk with them, that states a simple, observable, and measurable fact, and that would be familiar to personnel who work at the bank.

Example 3: A telecommunications company is heavily relied upon to provide communication services that are now considered fundamental to a functioning society. Additionally, they carry tremendous amounts of private information about their customers, often at considerable perceived risk by the public. But consumers subscribe to these services for considerable benefit. The telecommunications company defines their mission this way, "To instantaneously and transparently connect our customers with the people, organizations, information, and communication platforms that they care about." This mission definition is less concise, but it may be as concise as they could get, considering the complexity of typical telecommunications services. The definition is measurable, and it would very likely be recognizable by their personnel as an important value.

Defining Obligations

Definition: An organization's obligations, at least in terms of information security, are to prevent foreseeable harm that may come to others as a result of an information security compromise. These types of harm are commonly associated with identify theft, theft of funds, or lost services

and data. But it is critical to keep in mind that obligations should explicitly state the harm that may come to others so that risk assessors, management, and interested parties know that the organization is careful to protect others from harm. So “obligations” are a way of asking, “What foreseeable harm could come to others that we should prevent?”

A Good Obligations Definition should have the following characteristics:

- Concisely state the organization’s intent to prevent harm that information security incidents may cause others.
- Convey a simple fact that can be observed and measured.
- Describe something that the organization already manages to, and that personnel will recognize as important to the organization.

Example 1: The custom manufacturer is concerned about the harm that their customers may suffer if their intellectual capital – their product designs – are leaked to the public and to the customers’ competitors. Their customers often provide specifications for components that would reveal secrets about new products. The manufacturer states their obligations this way, “Our customers’ intellectual property must be kept confidential to preserve their market advantage.” The definition is concise, it states a foreseeable harm to others that must be guarded against, it could be measured, and personnel would know that protection of customers’ intellectual property is important.

Example 2: The community bank knows that their customers are in a particularly vulnerable position. They are often taking on risk while buying homes or starting businesses, and have less room for error. A missed paycheck or even slight misuse of financial information could mean the difference between success and failure for them. So the community bank uses this as their obligations definition, “We must protect our customers’ reputation and financial future against misuse of their financial or personal information.” Again, this definition is concise, it is measurable, and it would be known and recognized by the bank’s personnel.

Example 3: The telecommunications company is a complex organization that can imagine multiple kinds of harm that could result from an information security compromise. An organization with many obligations (or missions or objectives) may state them in their definitions. For example, the telecommunications company realizes that they may foreseeably breach personal communications about their customers, and their communication services may fail when critical applications depend on them. They will state two obligations, “We must protect our customers’ communications records to prevent reputational or financial harm. We must meet our service level agreements with customers to prevent harm that may result from unreliable connectivity.”

As an example, the top row of the impact definitions for the manufacturer would start to take shape as in the table below.

Table 76 - Impact Area Definitions Example (Partial)

Impact to Our Mission	Impact to Obligations
<i>To provide customers with products that meet their unique specifications, without fail.</i>	<i>Our customers’ intellectual property must be kept confidential to preserve their market advantage.</i>

Defining Impact Scores

After defining impact areas, the organization will need to define impact scores for each impact area. Each score ('1' through '3') will have one definition per impact area, as shown in Figure 26. There are a few principles that organizations should consider as they define their risk scores.

Consider the scores as having the following meanings:

Table 77 - Impact Scoring Guidance

Impact Score	Guidance
1	An impact that would be acceptable to the mission, objectives, and obligations. An impact would be noticeable, but is likely unavoidable even after significant investment in controls. It would be considered tolerable by all affected parties.
2	An impact that would be considered unacceptable by any party. While the impact may be recoverable through additional efforts, investment, or time, the organization could have reduced the risk of that impact with security controls.
3	The impact would be catastrophic. The mission, objectives, and/or obligations would no longer be feasible.

Score '1' is shaded to indicate that these impacts should be defined in a way that would be acceptable to all parties.

As impact score definitions are written, the organization should think through how impacts to their mission and obligations would appear at each of these levels.

The manufacturer's impact score definitions are shown below to illustrate the point.

Table 78 - Example Impact Score Definitions

Impact Score	Impact to Our Mission	Impact to Obligations
<i>Defined</i>	<i>To provide customers with products that meet their unique specifications, without fail.</i>	<i>Our customers' intellectual property must be kept confidential to preserve their market advantage.</i>
1	Occasional orders cannot be fulfilled.	Information about jobs may be known, but nothing that can harm customers' market position.
2	Products are delivered outside of spec and customers believe that we cannot produce custom products without fail.	A single customer experiences market repercussions based on a security incident.
3	We can no longer produce reliable, custom products.	Customers can no longer expect confidentiality protection when working with us.

When reading an impact definition horizontally across one score, note that the impact definition in each impact area are equally harmful to all parties. This is a critically important feature of Duty of Care risk analysis. It ensures that risk analysis is equitable. No party's harm is considered more or less tolerable than any other party's harm. What's acceptable to one equates to what is

acceptable to all (the shaded score ‘1’). What is catastrophic to one equates to what would be catastrophic to all.

Example impact definitions for all three example organizations are provided in the supplementary document *CIS_RAM_Workbook* in the tab “Example Impact Definitions” to assist the reader’s comfort and familiarity with this subject.

Defining Impacts for Tier 2, Tier 3, and Tier 4 organizations

Perhaps the most important early step in the risk assessment is to develop effective impact definitions. The CIS RAM is based on Duty of Care Risk Analysis principles to enable organizations to make conscientious evaluations of their current and intended risk. A risk assessment that results from the CIS RAM should show whether information security safeguards are appropriate to the public, while being reasonable to the organization. The core of this analysis is the impact definitions, and the balance and consensus that they are meant to establish.

This section will provide guidance for defining impact types effectively.

Table 79 – Summary Evaluation of Impact Definition

Benefits	Limits
<ul style="list-style-type: none"> - Consistent method to evaluate risk impacts. - Satisfies “cost-benefit” analysis that regulators use. - Satisfies “duty of care balance test” that courts rely on. - Balances business interests with public interest. 	<ul style="list-style-type: none"> - Poorly defined impact definitions can frustrate risk assessors. - Poorly “balanced” impact definitions may not reduce legal liabilities.

The primary purpose of impact definitions is to provide risk assessors with a consistent method for scoring risks that is fair to all potentially affected parties. A risk evaluation should demonstrate as much concern for the organization as it does for others.

To make this possible, impact definitions should be designed with the concept of balance firmly in mind.

Recall the impact definitions used for Tier 2, Tier 3 and Tier 4 organizations shown in Figure 27. Three columns address the interests of the parties who may be effected by information security risk. Each of these columns is considered an “impact type.” The organization itself is considered by evaluating the potential impacts to their ability to succeed in the “Impact to Objectives” column. The safety of the organization’s patient customers is considered in the “Impact to Obligations” column. And the “Impact to Mission” column addresses the main purpose for the two parties to enjoy in the risk ... the beneficial reason for customers and the organization to share information.

Figure 27 - Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objectives: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

Now consider the impact scores, and the red boundary that separates scores ‘1’ and ‘2’ from scores ‘3’, ‘4’, and ‘5’. This boundary marks the division between impacts that would presumably be acceptable to all parties, and those that would not be.

Consider how impacts for scores ‘1’ and ‘2’ would appear to the organization, patient customers, and legal authorities. The organization using this impact definition is stating that they would accept impacts from threats if they result in conditions similar to how scores ‘1’ and ‘2’ are defined. Threats that are scored with an impact as high as ‘2’ would indicate that; not all patients would receive the information they needed, profits would be off-target, but within planned variance, and patients would be concerned about a security incident, but would not suffer harm.

Then consider how impacts for scores of ‘3’ are defined. Threats that would foreseeably result in an impact score of ‘3’ could mean that some patients who cannot access information may not maintain good health outcomes. That scenario is clearly an unacceptable impact to the organization’s mission, and would not make it worthwhile for patients to entrust their information with that organization. In terms of the objectives, the organization’s profitability would be off-plan and would require a fiscal year to recover. Again, this is unacceptable and should be invested against so the scenario is avoided. And finally, the impact score of ‘3’ for obligations would be unacceptable because some of the patient customers could foreseeably be harmed financially or reputationally as a result of an incident – presumably because of identity theft.

All of these features of an impact definition make it effective for estimating risk in a way that is equitable to all potentially affected parties, and even for driving consensus within the organization that uses it. After all, the interests and purpose of the organization are addressed as well as the interests of the public.

Organizations can build effective impact definitions by thoughtfully defining their impact areas, and then by carefully defining their impact scores.

Defining Impact Areas

The CIS RAM describes three impact areas that should be included in risk evaluation; Impact to Mission, Impact to Objectives, and Impact to Obligations. Each of these impact areas address the interests of people or organizations who may be affected by information security risk. They each

play a significant and unique role in analyzing risk, and should be defined within those roles, as described below.

Defining Mission

Definition: An organization's mission is the value it provides others, and that requires that they engage in risk together to achieve that value. A college educates its students, but students must give personal and financial information to those colleges in order to receive the education. Retailers provide products to customers, but in many cases customers turn over their financial information to those organizations to receive those products. Cloud service providers offer Internet-based functionality to business customers, but those customers must turn over business process or operations to those services as a result. So "mission" is a way of asking, "What's in it for the others?" who engage in risk with the organization.

Example Mission Definitions should have the following characteristics:

- Concisely state the benefit the organization provides that encourages others to enjoin them in information security risk.
- Convey a simple fact that can be observed and measured.
- Describe something that the organization already manages to, and that personnel will recognize as important to the organization.

Example 1: A custom manufacturer uses their customers' intellectual property to quickly make components that are perfect upon delivery. Their mission definition may be, "To provide customers with products that meet their unique specifications, without fail." The message is simple, it can be measured, and the organization likely recognizes this as something that they manage to. Moreover, the definition can be useful when a risk assessor tries to determine what core values may be harmed is a risk were to occur, or if a safeguard is too burdensome to the mission. If the risk assessor recommends a control that prevents customers from sending their intellectual capital, or that prevents the manufacturer from storing it or sharing it among drafter and engineers, then the mission would clearly be negatively impacted.

Example 2: A community bank states their mission as, "We promote opportunities to households and small businesses in our community by providing affordable financial products and advisory services." They could say that their mission is to lend and borrow money, but they are thinking ahead to what they do not want to compromise about that mission, and that is to serve their community. But again, they have a concise definition that states why others would enjoin in risk with them, that states a simple, observable, and measurable fact, and that would be familiar to personnel who work at the bank.

Example 3: A telecommunications company is heavily relied upon to provide communication services that are now considered fundamental to a functioning society. Additionally, they carry tremendous amounts of private information about their customers, often at considerable perceived risk by the public. But consumers subscribe to these services for considerable benefit. The telecommunications company defines their mission this way, "To instantaneously and transparently connect our customers with the people, organizations, information, and communication platforms that they care about." This mission definition is less concise, but it may be as concise as they could get, considering the complexity of typical telecommunications services. The definition is measurable, and it would very likely be recognizable by their personnel as an important value.

Defining Objectives

Definition: An organization's objectives are more inwardly focused and more selfish. As people commonly think of the "burden" of a safeguard, they often think of "objectives" and most often of

financial burden, or “cost.” But cost is an overly-narrow and potentially hazardous risk metric. Organizations should want to stay away from analysis that associates financial amounts to levels of harm that others would suffer. “We won’t spend \$200,000 to protect our customers’ privacy,” is a terrible message to send to personnel, to customers, and to the public. Rather, organizations should think about indicators that they are succeeding or failing as an organization, independently of their mission definition. So “objectives” are a way of asking, “How do we know we are a successful organization?”

A Good Objectives Definition should have the following characteristics:

- Concisely state indicators of success that are independent of the mission definition.
- Convey a simple fact that can be observed and measured.
- Describe something that the organization already manages to, and that personnel will recognize as important to the organization.

Example 1: The custom manufacturer has a five-year plan to expand into two new markets and quadruple its production and profits. They know not to associate dollars to potential harm to others, but growth in productivity and profitability can still be properly stated as an objective. Their objectives definition may be, “To quadruple our production and profits in five years through expansion into two new markets.” The message is concise and independent of the mission,²⁰ can be observed and measured, and would be known by personnel, but particularly by management whose goals would be aligned with the five-year plan.

Example 2: The community bank’s mission is a compelling one, and they may want to associate the success of their bank with the success of their community. But they must operate as a viable financial institution if they are to service their community members. So their objectives will be targeted to that goal. They believe that they need to maintain a certain return-on-assets ratio to off-set future financial risk and state it this way, “We must retain a return-on-assets of 1.25% year-over-year.” This definition is a concise indicator of success that can be easily measured, and that personnel, and certainly managers and officers, would already be operating to.

Example 3: The telecommunications company knows one thing very well; that regardless of their earnings, their growth in consumer base, their growth in capital investments, or their reputation ... if they slip below their “top two” status in their competitive market, they will be targeted for acquisition. They define their objectives this way, “To grow our subscriber base, communications capital, and revenue faster than our competition and to remain number one or two in the marketplace.” Again, this is less concise, but the company relies on many moving parts to be successful. The objectives definition can be measured, and personnel are certainly aware of the goals and are responsible for managing to them.

Note Regarding Use of Financial Costs as Objectives: Organizations that wish to state their impacts in terms of financial costs should carefully consider the message they send to colleagues, interested parties, and authorities as they define their objectives impact scores. If an unacceptable impact score (‘3’) for objectives states, for example, “\$100,000,” and the same score (‘3’) for obligations is “Up to 100 customers would have their information stolen and abused” or something similar, then is the organization saying, “We would not spend \$100,000 to

²⁰ The success of the objectives may be dependent on the success of the mission, but the definition of the of the objectives is not relying on the definition of the mission.

protect 100 customers?” If so, are they prepared for how their colleagues, the public, and legal authorities will perceive that message?

By focusing on the magnitude of impacts against objectives (even the noble cause of profitability or financial performance) in terms of harm to the organization (their ability to operate profitably, or recover profitability after an event) they can present their risk balance responsibly.

This qualitative approach to evaluating impacts to financial objectives is still useful by describing the cost of recommended safeguards and evaluating the magnitude that the cost will have on the objectives. Such case-by-case comparisons would still state financial value of the proposed safeguard, but that financial value would be weighed against an operating principal called “profitability” or “financial performance” which regulations and courts already include in their definition of “burden.”

Defining Obligations

Definition: An organization’s obligations, at least in terms of information security, are to prevent foreseeable harm that may come to others as a result of an information security compromise. These types of harm are commonly associated with identify theft, theft of funds, or lost services and data. But it is critical to keep in mind that obligations should explicitly state the harm that may come to others so that risk assessors, management, and interested parties know that the organization is careful to protect others from harm. So “obligations” are a way of asking, “What foreseeable harm could come to others that we should prevent?”

A Good Obligations Definition should have the following characteristics:

- Concisely state the organization’s intent to prevent harm that information security incidents may cause others.
- Convey a simple fact that can be observed and measured.
- Describe something that the organization already manages to, and that personnel will recognize as important to the organization.

Example 1: The custom manufacturer is concerned about the harm that their customers may suffer if their intellectual capital – their product designs – are leaked to the public and to the customers’ competitors. Their customers often provide specifications for components that would reveal secrets about new products. The manufacturer states their obligations this way, “Our customers’ intellectual property must be kept confidential to preserve their market advantage.” The definition is concise, it states a foreseeable harm to others that must be guarded against, it could be measured, and personnel would know that protection of customers’ intellectual property is important.

Example 2: The community bank knows that their customers are in a particularly vulnerable position. They are often taking on risk while buying homes or starting businesses, and have less room for error. A missed paycheck or even slight misuse of financial information could mean the difference between success and failure for them. So the community bank uses this as their obligations definition, “We must protect our customers’ reputation and financial future against misuse of their financial or personal information.” Again, this definition is concise, it is measurable, and it would be known and recognized by the bank’s personnel.

Example 3: The telecommunications company is a complex organization that can imagine multiple kinds of harm that could result from an information security compromise. An organization with many obligations (or missions or objectives) may state them in their definitions. For example, the telecommunications company realizes that they may foreseeably breach personal communications about their customers, and their communication services may fail when critical applications depend on them. They will state two obligations, “We must protect our customers’

communications records to prevent reputational or financial harm. We must meet our service level agreements with customers to prevent harm that may result from unreliable connectivity.”

As an example, the top row of the impact definitions for the manufacturer would start to take shape as in the table below.

Table 80 - Impact Area Definitions Example (Partial)

Impact Score	Impact to Our Mission	Impact to Objectives	Impact to Obligations
	<i>To provide customers with products that meet their unique specifications, without fail.</i>	<i>To quadruple our production and profits in five years through expansion into two new markets.</i>	<i>Our customers' intellectual property must be kept confidential to preserve their market advantage.</i>

Defining Impact Scores

After defining impact areas, the organization will need to define impact scores for each impact area. Each score ('1' through '5') will have one definition per impact area, as shown in Figure 27. There are a few principles that organizations should consider as they define their risk scores.

Consider the scores as having the following meanings:

Table 81 - Impact Scoring Guidance

Impact Score	Guidance
1	An impact that would be negligible for the mission, objectives, and obligations. If any impact were to occur, it would not be noticeable.
2	An impact that would be acceptable to the mission, objectives and obligations. An impact would be noticeable, but is likely unavoidable even after significant investment in controls. It would be considered tolerable by all affected parties.
3	An impact that would be considered unacceptable by any party. While the impact may be recoverable through additional efforts, investment, or time, the organization could have reduced the risk of that impact with security controls.
4	An impact that would be considered large, but recoverable. Significant efforts and investments would need to be made to recover for all parties.
5	The impact would be catastrophic. The mission, objectives, and/or obligations would no longer be feasible.

Scores '1' and '2' are shaded to indicate that these impacts should be defined in a way that would be acceptable to all parties.

As impact score definitions are written, the organization should think through how impacts to their mission, objectives, and obligations would appear at each of these levels.

The manufacturer's impact score definitions are shown below to illustrate the point.

Table 82 - Example Impact Score Definitions

Impact Score	Impact to Our Mission	Impact to Objectives	Impact to Obligations
<i>Defined</i>	<i>To provide customers with products that meet their unique specifications, without fail.</i>	<i>To quadruple our production and profits in five years through expansion into two new markets.</i>	<i>Our customers' intellectual property must be kept confidential to preserve their market advantage.</i>
1	Customers receive excellent products, as needed.	Our growth plan remains on target.	All intellectual property is protected.
2	Occasional orders cannot be fulfilled.	Our annual targets are off year-by-year, but within planned variance.	Information about jobs may be known, but nothing that can harm customers' market position.
3	Contracted work for few customers cannot be completed as planned.	Our growth is too low for one year, but can be recovered to meet the five-year goal.	Information about a job leaks, and a customer needs to investigate whether it created harm. Even if direct harm would not result.
4	Products are delivered outside of spec and customers believe that we cannot produce custom products without fail.	We cannot meet the five-year growth plan.	A single customer experiences market repercussions based on a security incident.
5	We can no longer produce reliable, custom products.	We cannot operate profitably.	Customers can no longer expect confidentiality protection when working with us.

When reading an impact definition horizontally across one score, note that the impact definition in each impact area are equally harmful to all parties. This is a critically important feature of Duty of Care risk analysis. It ensures that risk analysis is equitable. No party's harm is considered more or less tolerable than any other party's harm. What's acceptable to one equates to what is acceptable to all (the shaded scores '1' and '2'). What is catastrophic to one equates to what would be catastrophic to all.

Example impact definitions for all three example organizations are provided in the supplementary document CIS_RAM_Workbook in the tab "Example Impact Definitions" to assist the reader's comfort and familiarity with this subject.

Estimating Likelihood Through "Defense-Readiness" Analysis

The CIS RAM presents a standardized method for estimating the likelihood of an incident by focusing on how a threat could interact with a vulnerability. This concept of foreseeability is easily communicated to broad audiences, and is embedded in legal and regulatory language, so it is a useful construct for likelihood estimation. However, some organizations need more rigor while estimating likelihood and can achieve that by evaluating the characteristics of a successful or failed attack in their environment.

Table 83 – Summary Evaluation of Defense-Readiness Analysis

Benefits	Limits
<ul style="list-style-type: none"> - Consistent method to evaluate likelihood - Supports evidence-based estimation 	<ul style="list-style-type: none"> - Not based on an established standard - Evidence-based criteria are optional

Borrowing from Binary Risk Analysis,²¹ “defense-readiness analysis” asks a series of questions about attacks and safeguards to estimate the ability of a control to detect or prevent foreseeable threats.²² A risk assessor can quickly evaluate defense-readiness by asking a set of questions such as these:

1. Is this threat expected either because it is a common cause of incidents, or because the skills and resources needed to enact the threat are common?
2. Does the control leave the asset exposed to this threat occasionally or partially?
3. Are there no other safeguards or conditions between the asset and the threat?
4. Is the vulnerability present on a frequent or consistent basis?

Organizations that use a 1 through 5 likelihood scale could then derive the likelihood score reducing the maximum score of ‘5’ by ‘1’ for each ‘yes’ response that they provide to the fortitude questions. Organizations that use a 1 through 3 likelihood scale may choose to assign .5 points per response to arrive at scores between 1 and 3.

As an example of this analysis process, an organization is concerned that its software developers have access to the production environment. CIS Control 18.9 states, “Maintain separate environments for production and nonproduction systems. Developers should not have unmonitored access to production environments.” The threat model they are evaluating relates to promotion of unsecure or malfunctioning code to the production environment. The organization currently has a policy that requires code review and approval prior to promoting code to production, but many software developers have access to the production environment in case they need to respond to emergencies.

To estimate the likelihood of the risk, they answer the defense-readiness analysis questions below.

Table 84 – Defense-Readiness Analysis Example

Defense-Readiness Analysis Question	Response (1 = “Yes”, 0 = “No”)
Is this threat expected either because it is a common cause of incidents, or because the skills and resources needed to enact the threat are common?	1

²¹ <http://binary.protect.io> (accessed January 3, 2018)

²² While Binary Risk Analysis provides a rigorous analytic process, it must be modified to address the concepts of “reasonable” and “appropriate” that are central to the CIS RAM and the legal and regulatory sphere.

Defense-Readiness Analysis Question	Response (1 = "Yes", 0 = "No")
Does the control leave the asset exposed to this threat occasionally or partially?	1
Are there no other safeguards or conditions between the asset and the threat?	0
Is the vulnerability present on a frequent or consistent basis?	1
Likelihood score (Sum of responses).	4

The organization arrives at a likelihood score of '4' after their analysis. If a careless or hurried programmer is the threat in this scenario, then the organization responds with these scores for the following reasons.

1. The threat is a common cause for breaches, and requires common skills to access the production environment and promote code, so the organization responds with '1'.
2. The logical access safeguards that protect the production environment always permit programmers to promote code, so the organization responds to this question with a '1'.
3. Because the programmers generally adhere to policies, secure coding practices, and documented workflows, there are other safeguards that would prevent the promotion of harmful code to the production environment. So the organization answers with a '0'.
4. And the vulnerability is consistently present so the organization responds again with a '1'.

Finally, they add the responses to achieve their likelihood score, which is in this case '4'.

Organizations should be aware that defense-readiness analysis is not necessarily evidence-based, and does not comprehensively examine all aspects of control strength. The benefit of defense-readiness analysis is to help organizations systematically, thoughtfully, and consistently estimate their likelihood scores based on internal and external criteria.

For example, an organization can also consistently analyze defense-readiness by asking questions similar to those below:

1. Is the asset in this threat scenario attractive to well-resourced hackers?
2. Is this attack a common cause for breaches in our industry?
3. Has this control been evaluated as effective using a penetration test, or other rigorous test?
4. Is our time to detect-and-prevent this attack's impact shorter than the time the attack needs to create an impact?

Defense-readiness can take evidence into consideration by reviewing information security data about their environment (if they have implemented the tools necessary to gather that information), and by reviewing known causes for information security events, incidents, and breaches.

Using Probability with Duty of Care Risk Analysis

Some organizations have useful sources of data to conduct probability analysis to help them determine the likelihood of risks. Probability analysis requires statistical methods, well-honed estimation skills, and good data to determine the likelihood of an event, or of events with certain impacts.

Table 85 - Summary Evaluation of Using Probability with Duty of Care Risk Analysis

Benefits	Limits
<ul style="list-style-type: none"> - Supports evidence-based estimation - Risk assessors may rely on professional rigor to improve risk estimation over time. 	<ul style="list-style-type: none"> - Statistical estimates must address all impact types to properly evaluate each risk on a due-care basis.

This section will propose an approach to link probability analysis to “duty of care” assessments, but will not present a comprehensive method for doing so, nor will it provide an explanation of the probability analysis that this section refers to. Readers who use probability analysis for information security management are encouraged to explore integration methods like those described here.

Probability analysis is complementary to the risk evaluation methods described in the CIS RAM, but readers should understand their differences.

1. With the right guidance, applying probability models to individual cybersecurity risks is (probably) simpler than the reader may imagine.²³
2. Probability is evidence-based and can help organizations model increasingly realistic threat scenarios, especially as the rigor of their methods increases.
3. Probability is built upon decades of sound methodology, driven by professional statisticians, using universally recognized processes for reducing uncertainty. Organizations that can use probability methods should do so.
4. Probability models often result in ranges or curves, rather than discrete scores such as those produced by the RAM. Ranges of possible outcomes resemble “estimation” better than a discrete value does. However, ranges and curves are more difficult to compare and prioritize than discrete scores.
5. Cybersecurity threat scenarios are varied, and therefore require a variety of probability methods to estimate their likelihood. This makes their raw results hard to compare and rank. A Monte Carlo simulation that provides a single value (for example, a “22%” chance an unencrypted laptop will be stolen) and a Bayes model that provides a curve describing several possibilities of likelihoods of dollar impact (such as “13% chance of \$750,000 loss and 24% chance of 177,000 loss) are not readily comparable. And comparing them to a single risk acceptance criterion will be similarly difficult.
6. Probability models on their own do not naturally align with duty of care questions posed by regulations or litigation. Regulations and litigation insist on evaluating “due care” by balancing different kinds of things (safety of customers versus the value of services provided to them, for example). Statistical models that describe impact in terms of one thing (i.e. cost of an impact) miss other real consequences of cybersecurity breaches, such as harm to others, or the burden of too-stringent safeguards (i.e. reduction in services, risks to personnel safety, difficulty in operating profitably, or the benefits that an

²³ Douglas Hubbard of Hubbard Research has published many books on the subject, the latest of which focuses exclusively on cybersecurity. Hubbard, Douglas and Richard Seiersen. *How to Measure Anything in Cybersecurity*. Hoboken, NJ: John Wiley & Sons, Inc., 2016.

organization provides to others, etc.). When all impacts are evaluated in terms of dollars, juries and judges tend to see the risk assessment unfavorably.²⁴

But organizations that are well-trained in statistics and estimation should use probability modeling and simulations to feed their Duty of Care analysis. How, then, would such an organization accomplish this using the CIS RAM?

Such organizations could “translate” the output of their probability models (such as Bayesian distributions, Monte Carlo simulations, or estimates) into the likelihood scores or risk scores that are used in the risk register.

Consider the case of a Bayesian distribution for the risk of malware causing a breach on an end-user system, given the use of a robust malware prevention system. The risk probability can be calculated as such:

$$P(\text{malware breach} \mid \text{malware prevention}) = P(\text{malware} \mid \text{malware prevention}) P(\text{malware breach} \mid \text{malware}) + (1 - P(\text{malware} \mid \text{malware prevention})) P(\text{malware breach} \mid \sim \text{malware}) = (.05)(.75) + (.95)(.01) = 4.7\%.$$
²⁵

Recall that the impact scoring table for a Tier 2 or Tier 3 and Tier 4 organization has five values that describe levels of foreseeability. An organization can very easily add a column to their likelihood definitions table (as in Table 86) to state how they associate foreseeability with probability.

Note: The probability values provided in Table 86 are illustrative only. Each organization could determine for themselves how to associate foreseeability with probability while defining their risk assessment criteria. Over time, the organization should review and update this table.

Table 86 – Example Likelihood Definitions Aligned with Probability

Likelihood Score	Foreseeability	Probability %
1	Not foreseeable	< .5%
2	Foreseeable but not expected	< 5%
3	Expected to occur	< 10%
4	A common occurrence	< 25%
5	May be happening now	<= 100%

The organization’s probability of malware given their robust malware prevention system is 4.7%, which is below 5% (the organization’s definition for “Foreseeable but not expected” likelihood). This guides them to select the Likelihood value of ‘2’ for the risk of malware infection. They would then select the impact scores for such a scenario to derive their risk score.

²⁴ Viscusi, W. Kip, “Jurors, Judges, and the Mistreatment of Risk by the Courts.” *Journal of Legal Studies*, vol. XXX (January 2001).


²⁵ **Note:** The CIS RAM does not expect that readers understand or use probability analysis. This calculation is shown only to demonstrate how probabilistic analysis can be coordinated with the CIS RAM.

And because some probability analysis results in multiple possibilities within a range (i.e. a 1% chance of a complete system failure, or a 15% chance of a partial system failure) a risk register may simply show the same risk twice, but with two different risk evaluations.

Results such as these – multiple risks within a range – are often caused by a precondition that may or may not be in place when the threat occurs. For example, a low risk of a complete system failure may be associated with normal operating conditions, and a higher risk may be associated with a busy season or other non-typical circumstance. Two risks on the risk register could then describe the two risks differently and call out their dependence on the precondition (i.e. risk one is associated with the busy season, risk 2 is associated with normal operations.).

Many probability models produce a range of both likelihood and impact scores, such as the Bayesian distribution depicted in Figure 28.

Figure 28 - Probability Curve Example



A benefit to such a probability range is that it estimates both the likelihood and impact values of a risk. At its greatest likelihood, there appears to be about a 24.9% probability of a \$1MM impact. At its lowest likelihood, there appears to be below 1% probability of a cost of \$20MM or greater. This allows the client to either state a high likelihood of a lower cost, or a lower likelihood of a higher cost. The organization could model both of these scenarios to determine which creates the greater risk score to address that risk.

Pairing the probability curve to the modified likelihood definitions table (Table 86), the highest probability score (less than 25%) appears to be at the upper edge of a likelihood score of '4'.

If the organization's impact scoring table refers to cost or budget impacts (perhaps in the Objectives column) then \$1MM would be considered within the context of those impact values. For example, the organization may determine that a \$1MM loss would take more than one year to recover from, indicating an impact value of '4' under their Objectives column in Table 87. That

would leave them with a risk score of '16' by multiplying the likelihood score of '4' and the impact score of '4'.

Table 87 – Example Impact Definitions

Impact Score	Impact to Mission <i>Mission: Provide information to help remote patients stay healthy.</i>	Impact to Objectives <i>Objectives: Operate profitably.</i>	Impact to Obligations <i>Obligations: Patients must not be harmed by compromised information.</i>
1	Patients continue to access helpful information, and outcomes are on track.	Profits are on target.	Patients do not experience loss of service or protection.
2	Some patients may not get all the information they need as they request it.	Profits are off target, but are within planned variance.	Patients may be concerned, but not harmed.
3	Some patients cannot access the information they need to maintain good health outcomes.	Profits are off planned variance and may take a fiscal year to recover.	Some patients may be harmed financially or reputationally after compromise of information or services.
4	Many patients consistently cannot access beneficial information.	Profits may take more than a fiscal year to recover.	Many patients may be harmed financially or reputationally
5	We can no longer provide helpful information to remote patients.	The organization cannot operate profitably.	Some patients may be harmed financially, reputationally, or physically, up to and including death.

While organizations may be pleased with the ease by which they can align their probability outputs to their due-care risk assessment criteria, probability should be modeled against all impact criteria. Risk analysis that only considers financial impacts to the organization misses a necessary component of cybersecurity risk analysis; estimating and taking responsibility for the potential harm that others may suffer. As well, the organization must evaluate the risk of safeguards using the same probability models.

For thorough and practical guidance on using probability analysis for cybersecurity decision-making, consult the book, *How to Measure Anything in Cybersecurity* by Douglas Hubbard.²⁶

Noting How Realized Risk Might be Detected

Risk assessments provide organizations with insight into how security incidents and breaches will foreseeably occur. This provides organizations with an opportunity to tie their risk register to their processes for security log management and alerting, for incident response planning, and for security awareness training. All these benefits can be added to a risk register by adding a single column "Realized Risk."

²⁶ Hubbard, Douglas W., Richard Seiersen. *How to Measure Anything in Cybersecurity*. Hoboken, NJ: John Wiley & Sons, Inc., 2016

This column, which could follow the threat model on any of the risk registers that are illustrated in this document, would describe how the organization would know if an attack or error was in progress, or if an incident or event had occurred.

Consider the following risks:

Table 88 - Realized Risks Column Example

Current Control	Vulnerability	Threat	Realized Risk
Vulnerability scans occur occasionally and may not identify all systems that have been on the network between scans.	Systems that have joined the network between sporadic scans will not be detected.	Hackers or malware may attack and control systems that have not been detected, controlled, and monitored.	IP traffic is sent to or from hosts whose MAC addresses are not in the vulnerability scan output.
Vulnerability scans occur when Threat Info Service announces a moderate-to-high vulnerability that needs patching. Team reliably patches systems within 24 hours of announcement.	A 24-hour window of vulnerability remains with the current process.	Hackers or malware may attack and control systems that have not been patched within the 24-hour period after the vulnerability was announced.	Unusual traffic sent to the unpatched systems.
Vulnerability scans occur when Threat Info Service announces a moderate-to-high vulnerability that needs patching. Team patches most systems within 24 hours of announcement.	Enterprise management application systems are unpatched for more than one year.	Hackers or malware may attack and control enterprise management application environment.	SQL command strings sent from client browsers and from form fields and URL requests.

The organization now has a simple way to know what to look out for in terms of log-able events, to add escalation indicators to their incident response plan document, and to use in information security awareness training (especially for those realized risk indicators that are recognizable by general personnel).

Using Realized Risk for Monitoring: In the case of the first “realized risk” notation, the organization may decide to detect risks by comparing MAC addresses in ARP caches to those in their vulnerability scan output. With the right tools or scripting skills, this may be simple to achieve for them.

Using Realized Risk for Incident Response: The organization may note in their incident response plan that they need to take action when a MAC address appears for an unknown device (which would be hyper-vigilant in this case, but to illustrate the point).

Using Realized Risk for Training and Awareness: And the organization may use this as an opportunity to describe to systems engineers and help desk personnel information that may help them detect and investigate other unexpected activities on the network.

Leveraging Duty of Care Risk Analysis for Maturity Models

Many organizations use maturity models to evaluate the security capabilities in their environment. Maturity model evaluations ask questions about specific security controls or control groups so evaluators can determine how formalized an organization's security program is. Maturity models can be aligned to CIS RAM risk assessments by aligning a control in the maturity model with a corresponding CIS control in a risk register to determine the risk acceptability of the maturity score.

Table 89 - Summary Evaluation of Using Duty of Care Risk Analysis for Maturity Models

Benefits	Limits
<ul style="list-style-type: none"> - Organizations may continue to evaluate the formalization of their security programs. - Some moderate maturity scores may be sufficient if aligned with reasonable risks. 	<ul style="list-style-type: none"> - Some authorities who insist on evaluating organizations solely with maturity models may not accept moderate maturity scores that align with reasonable risk.

Maturity model evaluations generally ask organizations a set of control descriptions, and provide multiple-choice values as optional responses to the control description. For example, control descriptions for vulnerability management may state something similar to this:

Vulnerabilities are resolved within three business days.

Optional responses would be similar to this (though responses vary from one maturity model to the next):

0 = Not in place

1 = Ad hoc

2 = Documented

3 = Consistently applied

4 = Tested and corrected

5 = Continuously improved.

If an organization applies this process and they check for success with subsequent weekly tests, but do not improve their test, they would answer with a "4."

If this organization also conducts a CIS RAM risk assessment, they will have examined this control in a risk format in their risk register. In this case, they would have examined CIS Control 3.6 "Compare back-to-back vulnerability scans." They may have determined that the likelihood and impact of foreseeable threats is acceptably low. If their risk analysis tells the organization that their existing review process is acceptable in terms of risk, then they can also note that in their maturity evaluation this way.

Table 90 - Example Maturity Model Mapped to Risk

Control	Existing Control	Maturity Score	Risk Acceptance
<i>Vulnerabilities are resolved</i>	Weekly vulnerability scans review findings to determine whether any vulnerabilities	4	Accept

Control	Existing Control	Maturity Score	Risk Acceptance
<i>within three business days.</i>	were identified in previous scans.		

By aligning a maturity score with a risk score, the organization can determine whether they need to, or wish to, further formalize the control. In this way, the organization can continue to evaluate their security program using maturity scores, but not feel compelled to continuously improve unless there is a risk-related reason for doing so.

Interview Techniques

Risk assessment interviews are critical moments for understanding and modeling risk and should be approached differently than for a compliance assessment or an audit. In a risk assessment, assessors ask about existing safeguards, as in an audit, but in the context of how foreseeable threats may compromise information assets given those safeguards.

Table 91 – Summary Evaluation of Interview Techniques

Benefits	Limits
<ul style="list-style-type: none"> - Increases risk awareness among interviewees. - Observes key personnel knowledge of risk issues. 	<ul style="list-style-type: none"> - Personnel descriptions of safeguards and risks may not be accurate.

In a compliance assessment or audit, the assessor generally reports a “pass” or “fail” or even a “partial pass” by observing whether a required control is present. An auto-closing door with an auditable lock may be “compliant” with a security standard. An operating firewall may qualify as a “pass” for an audit of the network perimeter. Encryption between an application server and a database server may be marked as “green” or “low risk” on a report. But in a risk assessment, the risk assessor brings knowledge to the interview (and later the control configuration review) that stress-tests the described safeguards.

For example, does the firewall’s ruleset contain no more than the minimal rules and policies that are required for business? Who gets to change the firewall’s policies and how is their access tracked? Who has access to the firewall’s CLI and webadmin interfaces, and from what networks? Is the webadmin interface even active? How is firmware upgraded? How are firewall event logs reviewed and analyzed? Is access enforced through multi-factor methods? Does the firewall fail open?

As respondents provide each answer, the risk assessor should then consider a corresponding threat to help them model the risk. It may be appropriate to model risks with the respondent present, or after the interview. This depends on a number of factors, including the risk assessor’s comfort in improvising threat scenarios for each response about a safeguard. But ideally, the risk assessment interview will involve discussions about threats so that full risk discussions inform follow-up questions.

Consider how the interview questions above would play out when an interviewee’s responses are paired with threats.

Table 92 - Interview Threat Pairings

Interview Question	Response	Paired Threat
Does the firewall's ruleset contain no more than the minimal rules and policies that are required for business?	No. Some temporary policies may still be operating.	Hackers like to exploit firewall rules that provide access to internal systems and services. What are the chances they will find permissive policies on the firewall now? (Note: This is an excellent example of a safeguard to observe later).
Who gets to change the firewall's policies and how is their access tracked?	Two administrators have privileges with unique accounts. All firewall changes are logged and alerted to the security team.	Two administrators seems to be a low number. Do they share their passwords with other administrators to help out when needed?
Is the webadmin interface even active?	It is, but only for internal networks.	Attackers or malware may try to log onto the webadmin application using guessed or stolen credentials. How do you prevent that from happening?
Is access enforced through multi-factor methods?	Yes. Each administrator uses an application on their cell phone to receive an out-of-band, six-figure code that lasts one minute.	Could an unauthorized person access their phone while at a system that can log into the webadmin interface with stolen credentials?

By responding to answers with plausible threats, the risk assessor can drive a more interesting interview about whether a control is sufficiently designed, and can give them indicators of what safeguards and configurations they would want to follow up on during subsequent configuration reviews.

In addition to pairing interviewee responses with plausible threats, organizations should plan their risk assessment discussions along the lines of a structure, or planned conversation. This helps ensure that the risk assessor addresses a set of subjects that must be understood during a risk assessment, and allows them to rely on a plan rather than to exhaust themselves with constant improvisation.

Consider using one or more of these rubrics to set the rhythm of risk assessment interviews.

Rubric 1: "Full Lifecycle of Use," or "Process Audit." This method forces risk assessment participants to think of information assets as something that changes over the course of its usage. The process is generally useful for assessors who have experience in security operations, because it requires them to have practical knowledge of how information assets are developed and managed.

An example of a full lifecycle of use interview process will illustrate the point.

1. An information asset's lifecycle starts with defining the business requirements that it is being built for.
2. Then technical specifications are listed to prepare the information asset to satisfy the business requirements.
3. A hardening standard would then be selected to build the asset from

4. The information asset is then built, installed, or deployed with primary access rights.
5. More access rights are established, and the information asset is configured.
6. The information asset is included in common security processes, such as log management, vulnerability management, change management, penetration testing, and patch management schedules and routines.
7. And finally, it ends its lifecycle by being de-commissioned, lost, stolen, or damaged.

The lifecycle of an information asset provides a risk assessor indications of whether and how safeguards are applied to the information asset.

So if early interview questions determine that there is no rigorous process for identifying business requirements or technical specifications for server deployments, then safeguards that will be discussed further on, such as log management and change management, will be assessed with the knowledge that initial requirements for the server are not reliably known. A server's hardening and configurations processes may permit too many services and access privileges because business requirements and technical specifications are the appropriate stages for determining the least capabilities and least privileges for a server.

By the time the risk assessor arrives at questions regarding log management, they would also know to ask, "How do you know that the servers logs are appropriate to the risks and function of the system?" For vulnerability management, the risk assessor could ask, "How do you know the difference between a vulnerability service that is business-appropriate and one that is not?" And "How do you test patches, upgrades, and configuration changes if you don't know which business-critical services you may interrupt?"

Again, this process is appropriate for experienced security professionals, or individuals who have experience in technical operations and understand how previous stages of an information asset's lifecycle can influence the effectiveness of safeguards in later stages.

Rubric 2: "Security Management Lifecycle" The security management lifecycle of an information asset resembles maturity model evaluations used in other security assessment methods. For those not familiar with maturity models, they resemble something similar to the below table.

Table 93 - Example Maturity Model

Maturity Level	Description
1 – Ad hoc	Not consistent in practice.
2 – Documented	Documents state requirements for processes and configurations.
3 – Implemented	Requirements are applied to information assets as safeguards.
4 – Evidenced	Tests demonstrate that safeguards are effective.
5 – Detect & Correct	Management detects when failures occur and improve identified flaws.

Many assessments impose maturity models such as the one depicted above. In these assessments, assessors often rely on the selection of a maturity score as their total description of a safeguard. But without an understanding of an information asset's resilience to foreseeable threats, the acceptability of a control will be difficult to evaluate.

Despite this weakness, a maturity model can be useful in evaluating risk. For instance, while a risk assessor is evaluating a control to determine its resilience against threats, they could use the maturity model to understand the likelihood of current and future risk.

An example risk assessment conversation about server hardening demonstrates this process below.

Assessor (Asking the “1 - Ad hoc” question): Do you harden your servers to a known good standard?

Systems Engineer: Yes, our servers are deployed using images based on SCAP policies that support each operating system version.

Assessor (Asking the “2 - Documented” question): How do you know which SCAP policies to apply to each server?

Systems Engineer: We have standards that list each implemented operating system and version, and that state the SCAP policy version to use for each operating system.

Assessor (Asking the “3 - Implemented” question): How many servers in production now are based on SCAP policies?

Systems Engineer: All servers were built within the last year, so all are configured with its matching SCAP policy as its baseline.

Assessor (Asking the “4 - Evidenced” question): Are they still configured to the hardening standard they started with?

Systems Engineer: I think so. Some settings must have changed for new requirements or troubleshooting. Sometimes we make temporary changes that we don’t always change back.

Assessor (Asking the “5 - Detect & Correct” question): How do you know when a server is no longer configured to match its SCAP policy?

Systems Engineer: Vulnerability assessments show when there are vulnerabilities, so that’s one way. But we are not strictly checking for compliance against the SCAP policy after the servers have been deployed.

Now the organization knows that they started out configuring servers securely, but that the servers begin to diverge from their known-secure configuration and the organization may not detect it to correct it, unless the vulnerability scanning software compares server configurations to their SCAP policies.

Risk assessment interview methods can help drive the discovery and analysis in many ways, and no single method is ideal. Organizations should consider attempting a variety of methods to see which work best in different scenarios. But if there is one rule to apply in developing an interview style it is this: Be flexible to the environment, the capability of the participants and assessors, and the nature of the information you are trying to obtain. Sticking with one method throughout an assessment will miss opportunities to discover risk.

Evaluating Inherent Risk

Some organizations are interested in understanding the “inherent risk” of a scope of information assets. “Inherent risk” is defined in the Glossary as, “The likelihood of an impact when a threat compromises an unprotected asset.” A risk register that evaluates “inherent risk” would be nearly identical to the templates provided in the *CIS_RAM_Workbook* but would not note or consider controls that are in place to protect information assets.

Organizations that look for inherent risk often want to determine things such as:

1. What potential liabilities do we face in Venture A versus Venture B?
2. If we move a business process into a cloud service, what is the risk/reward of option A which uses all of our data, versus Option B with uses some of our data?
3. If we engage this business process / technology / facility, what new regulatory requirements will we bear?
4. What is the current benefit of our existing security program?

And while inherent risk analysis may provide quick answers to these questions, they appear to pose a fictional state for most situations. They appear to imagine environments in which there are truly no security controls. A quick analysis can describe the benefits and limits of inherent risk analysis in Table 93.

Table 94 – Summary Evaluation of Evaluating Inherent Risk

Benefits	Limits
<ul style="list-style-type: none"> - Aligns with some security evaluation standards, such as the FFIEC Cybersecurity Assessment Tool. - Provides a basis for a quick estimation of potential liabilities of certain ventures or technical architecture decisions. - May help raise awareness among non-technical management of the need to continue to invest in information security programs, safeguards and awareness. 	<ul style="list-style-type: none"> - Assumes a fictional condition without controls. - Does not evaluate the potential burdens (whether low or high) of safeguards in business propositions, which be material to the attractiveness of the proposition.

If risk assessors do intend to add inherent risk analysis to their risk registers, they can do so in the templates provided in the *CIS_RAM_Workbook* by adding columns to the left of the observed risk columns. The risk register could then compare the potential of full and immediate compromise of the assets to the current state of controls, and the proposed state of controls. Organizations should first determine what value this analysis provides.

Root Cause Analysis

Because organizations will be identifying weaknesses in safeguards, and will propose risk treatments to address those risks, they will need to understand the actual cause of the vulnerabilities to ensure that the vulnerabilities do not re-occur. This process of identifying the underlying causes for weaknesses or failures is called “root cause analysis.”

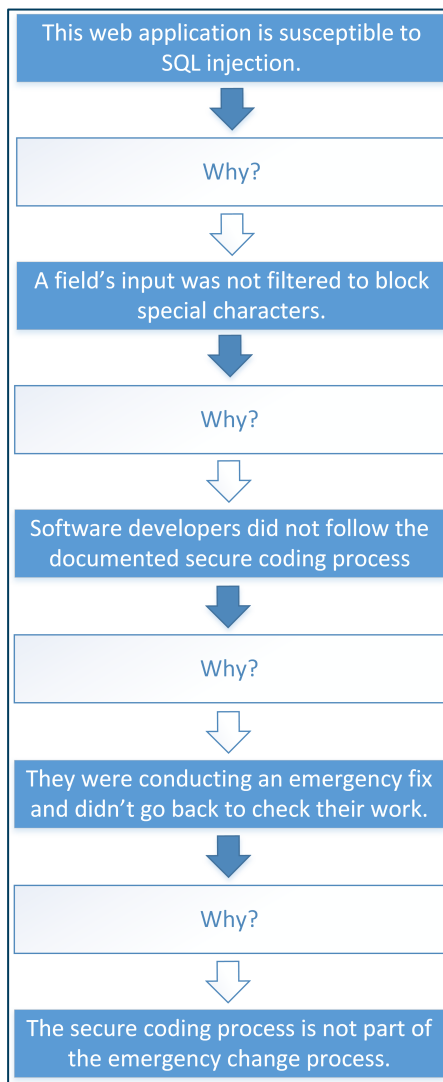
Table 95 – Summary Evaluation of Root Cause Analysis

Benefits	Limits
<ul style="list-style-type: none"> - Helps the organization reduce the likelihood of recurrence of the weakness. - Helps the organization simultaneously address other weaknesses that may result from the root cause. 	<ul style="list-style-type: none"> - May be difficult to identify the actual root cause as the organization starts using the process. - Root cause analysis may lead to one root cause when other root causes may be missed.

A generally accepted practice for identifying root causes is to conduct a “five whys” exercise. This entails the risk assessor to recurrently identify the underlying reason for a problem and the causes of the problem, until a “root cause” of the problem is identified.

Root cause analysis looks similar to the diagram below.

Figure 29 - Root Cause Analysis Model



While root cause analysis is classically conducted by asking, "why" five times to more deeply uncover the root cause of a problem, this example arrived at a plausible root cause after asking "why" four times. This is because an organization may uncover the root cause with more or fewer questions than five.

In this case, the risk assessor will have noted that a web page is vulnerable to a SQL injection attack. If their recommended risk treatment was simply, "filter all form objects on the page to prevent special characters, the application developers would fix the one identified error, and then likely repeat the error the next time they had an emergency application change.

After this root cause analysis, however, the organization knows to both address the identified weakness (the unfiltered input at the application page) and the root cause (the lack of security coding during or immediately after emergency changes).

Root cause analysis should be part of all vulnerability assessments, whether they occur during a risk assessment, an audit, or after a security incident as management determines lessons learned.

Helpful Resources

CIS (Center for Internet Security)

CIS (Center for Internet Security, Inc.) is a forward-thinking, non-profit entity that harnesses the power of a global IT community to safeguard private and public organizations against cyber threats. Our CIS Controls™ and CIS Benchmarks™ are the global standard and recognized best practices for securing IT systems and data against the most pervasive attacks. These proven guidelines are continuously refined and verified by a volunteer, global community of experienced IT professionals. CIS is home to the Multi-State Information Sharing and Analysis Center® (MS-ISAC®), the go-to resource for cyber threat prevention, protection, response, and recovery for U.S. State, Local, Tribal, and Territorial government entities. (www.cisecurity.org)

HALOCK Security Labs

Established in 1996, HALOCK Security Labs is an information security professional services firm based in Schaumburg, IL. For more than 20 years, HALOCK has provided Purpose Driven Security services to help organizations achieve their mission and objectives through sound security practices. HALOCK uses their deep background in the legal and regulatory landscape, security technologies and standards, business governance, and data analytics to provide evidence-based security analysis and guidance to their clients. (www.halock.com)

For guidance in implementing the CIS RAM: (www.halock.com/cisram)

DoCRA Council

The DoCRA Council maintains and educates risk practitioners on the use of the Duty of Care Risk Analysis (DoCRA) Standard that CIS RAM is based on. While DoCRA is applicable to evaluation of information security risk, it is designed to be generally applicable to other areas of business that must manage risk and regulatory compliance. (www.docra.org)

International Organization for Standardization (ISO®)

ISO provides to information security professionals a set of standards and certifications for managing information security through an information security management system ("ISMS"). ISO 27001 is a risk-based method for organizations to secure information assets so that they support the business context, and requirements of interested parties. ISO 27005 is an information security risk assessment process that aligns with CIS RAM. (<https://www.iso.org/isoiec-27001-information-security.html>)

National Institute of Standards and Technology (NIST)

NIST provides a series of standards and recommendations for securing systems and information, known as "Special Publications" in the SP 800 series. NIST SP 800-30 provides guidance for assessing information security risk, NIST SP 800-37 and NIST SP 800-39 each present approaches for managing information security risk within an organization. While these approaches are designed to address federal information systems and reference roles within federal agencies, their principles and practices are generally applicable to many organizations. (<https://csrc.nist.gov/publications/sp>)

NIST also provides the Framework for Improving Critical Infrastructure ("Cybersecurity Framework"). The framework organizes information security controls within a structure that prepares for and responds to cybersecurity incidents. The Cybersecurity Framework aligns its categories and sub-categories of controls with those of other control documents, including the CIS Controls. (<https://www.nist.gov/framework>)

Information Systems Audit and Control Association (ISACA®)

Well known for their IT assurance standards and certifications, ISACA provides an information security risk management framework known as Risk IT. Risk IT bases its risk analysis method on ISO 31000, and adds risk governance and response to the analysis to provide a lifecycle of IT risk management. (<http://www.isaca.org/Knowledge-Center/Risk-IT-IT-Risk-Management/Pages/default.aspx>)

Binary Risk Analysis (BRA)

Binary Risk Analysis is published as version 1.0. The analysis method is presented as a worksheet and an application at the hosting website. The BRA provides risk analysts with a concise and consistent process for evaluating information security risks by breaking down the components of a threat scenario, including the capabilities to defend against variably robust and common threats. (<http://binary.protect.io>)

Fair Institute

Fair Institute maintains and educates risk analysts on the use of Factor Analysis of Information Risk. The FAIR method is similar to BRA in that it provides a consistent method for evaluating information risk based on characteristics of the components of information risks. <https://www.fairinstitute.org/>

All references to tools or other products in this document are provided for informational purposes only, and do not represent the endorsement by CIS of any particular company, product, or technology.

Contact Information

CIS
31 Tech Valley Drive
East Greenbush, NY 12061
518.266.3460
controlsinfo@cisecurity.org

HALOCK Security Labs
1834 Walden Office Sq. Ste 200
Schaumburg, IL 60173
847.221.0200
cisram@halock.com