



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

Check out the list of upcoming events offering
"Implementing and Auditing CIS Controls (Security 566)"
at <http://www.giac.org/registration/gccc>

Realistic Risk Management Using the CIS 20 Security Controls

GIAC (GCCC) Gold Certification

Author: Andrew Baze, abaze@hotmail.com
Advisor: Manuel Humberto Santander Peláez
Accepted: July 18, 2016

Template Version September 2014

Abstract

If a forest fire was closing in on your town, putting your family and friends at risk, how many of your firefighting resources would you send to protect a distant town just in case they might have a fire someday? Does your organization spend an inordinate amount of time “managing” risk, when the current state of security is known to be poor, with far too few resources available to deal with the top issues? Risk management deals with uncertainty. It involves managing potential negative outcomes in the context of what is known. In other words, something negative could happen at some point, and the measure of "negative" is relative to a baseline. And considering the lack of adequate security in many organizations, even those which may meet their compliance requirements, one could assume that those baselines often haven't been sufficiently established. Without a clearly-defined known state, there is no basis for comparison, making the magnitude of the risk difficult to establish. This paper will describe how to create and execute a security control effectiveness evaluation program in the context of assessing the CIS 20 Critical Security Controls, with a goal of creating a foundation for realistic risk management.

1. Introduction

While risk management is often a compliance-focused exercise with debatable outcome beyond support for achieving certifications or passing audits, realistic cybersecurity capability assessment and subsequent prioritization and management of risks is possible if applied in the context of the CIS 20 Critical Security Controls.

This document will explore risk management definitions and program components, examples of ineffective risk management approaches, risk management benefits, how to effectively change risk the typical risk assessment approach.

These topics will provide the context of general risk management programs and why some of them do not work effectively. Later, they will describe how cybersecurity risk management programs can integrate with broader programs, and how a cybersecurity practitioner or risk manager would benefit from evaluating the current state of control effectiveness within an organization, closing gaps identified via that process, and simply representing risks as extensions of the existing gaps. With this approach, prioritization will be more realistic, enabling more effective use of resources to tackle the most important security problems.

2. Risk Management Definitions

In order to explore how realistic cybersecurity risk management should take place, some frequently misused and misunderstood definitions will be reviewed.

A very simple definition of risk is the “possibility of loss or injury” (Merriam-Webster, 2016). However, this definition misses some important nuance relative to the definition used in a cybersecurity context, as proposed by the NICCS (National Institute for Cybersecurity Careers and Studies), which is “the potential for an unwanted or adverse outcome resulting from an incident, event, or occurrence, as determined by the likelihood that a particular threat will exploit a particular vulnerability, with the associated consequences” (Explore Terms: A Glossary of Common Cybersecurity Terminology), 2016). As described later in this document, enterprise-level risk management programs commonly focus on impact and likelihood, and some cybersecurity risk management programs only focus on threat and

vulnerability. The NICCS definition accurately describes the relationship between these variables.

The following sentence illustrates this mapping with added italics for emphasis. Risk is the potential for an unwanted or adverse outcome resulting from an incident, event, or occurrence, as determined by the likelihood that a particular threat will exploit a particular vulnerability, with the associated consequences (or impact).”

Risk management should also be defined simply. While it may appear complex at first, this definition is concise: “Risk management is the identification, assessment, and prioritization of risks followed by coordinated and economical application of resources to minimize, monitor, and control the probability and/or impact of unfortunate events” (Hubbard, 2009). While this document will not provide a roadmap for the “coordinated and economical application of resources,” it will describe how to effectively identify, assess and prioritize risks.

Unfortunately, many risk management programs are more focused on meeting compliance requirements than effectively managing risks. The following quote accurately illustrates the situation: “Since the late nineties, thousands of corporations have been poring over their financial documents, consulting legal experts, overhauling their IT infrastructure, hiring compliance chiefs, and doing everything else humanly possible to comply with regulations like Sarbanes Oxley (SOX), the Gramm Leach Bliley Act (GLBA) and HIPAA (the Healthcare Insurance Portability and Accountability Act). Add to these regulations SEC 17, Europe's Basel II and complex regional laws such as California's Security Breach Information Act, and it's easy to see how the explosion in regulatory compliance requirements has bred its own cottage industry, replete with corporate consultants, IT solutions and revenues in the billions” (Carpenter, 2016).

In the same article, Carpenter goes on to illustrate the relationship that should exist between security and compliance: “For an organisation to truly achieve its information security goals, a different, more global perspective is needed: compliance should be one of many byproducts of a global policy management initiative whose aim is to safeguard the entirety of the organisation's intellectual property assets.” In other words, it is the focus on security, not compliance, that should drive risk management activities.

Andrew Baze, abaze@hotmail.com

Considering this common risk management focus on compliance, the definition used by Nassim Nicholas Taleb may be more accurate: “Risk management as practiced is the study of an event taking place in the future, and only some economists and other lunatics can claim – against experience – to “measure” the future incidence of these rare events, with suckers listening to them – against experience and the track record of such claims” (Taleb, 2014).

Taleb goes on to advocate for a discussion of fragility instead of risk, with the justification that risk is inherently predictive, intended to represent a potential future state, while fragility requires no prediction, and represents a system’s current state. Using this approach, focusing on the current state of an organization’s cybersecurity controls implementation (its fragility), a cybersecurity risk management practitioner can significantly increase the impact of his or her program.

But prior to demonstrating how to evaluate controls, understanding where the evaluation fits in a broader risk management program must be reviewed.

3. Typical Risk Management Program Components

To understand how capability evaluation can enable more effective risk prioritization, the traditional components of a risk management program are described below. They are often represented by a deceptively simple diagram, as shown below in Figure 1, adapted from ISACA’s “Continuous Risk Management Steps” (CISM Review Manual 2015, 2014).




Figure 1. Risk management program basic components

Another view from ISO 31000 starts with establishing the context, then moves into risk identification, analysis, evaluation, and risk treatment, all while monitoring and reviewing as part of ongoing assessment, and communicating risk information (ISO 31000:2009, 2009). None of these components are unique to the International Organization for Standardization, although the order and specific wording may change from standard to standard.

If an organization's risk management program has support, a clear charter, a defined scope, and other prerequisites (outside the scope of this document), its operation can be condensed to the following four steps (as also indicated in the definition for a risk management program earlier):

- 1) Identify risks and determine to what extent they could affect the organization.
- 2) Create prioritized plans to eliminate, reduce, or shift the risks.
- 3) Implement those plans.
- 4) Monitor the plan implementation and expected gap closure, continuously updating stakeholders and updating or improving the plan as needed.

The view in Figure 1 is simple and may seem intuitive. The area discussed at length in this document, risk assessment, is not. More context will help explain some of the reasons for this complexity.

4. Examples of Ineffective Risk Management Approaches

Douglas Hubbard, in his book "The Failure of Risk Management," describes five levels of risk management, a spectrum of program relevance. Those levels are represented below in Table 1.

Table 1. The five levels of risk management, from "The Failure of Risk Management" (Hubbard, 2009)

Level	Type	Description
1	Best	In this level, quantitative models are used, simulations are run, empirical measurements are used, and efforts are made to identify all possible risks
2	Better	This is the first level in which quantitative models are built.

3	Baseline	Intuition (explicitly, not hiding behind unproven models that often dilute the intuitive assessment or otherwise decrease its accuracy) is used, and no formal risk management is attempted.
4	Worse	Worse, aka “the merely useless.” In this level, detailed (though not quantitative) scoring methods are used, but management does not rely on the results, so the main negative consequence is the waste of time and money.
5	Worst	“Worse than useless.” In this level, ineffective methods are used that result in increased error, higher confidence, and poor decision-making that is no improvement over using pure intuition.

Many risk management programs fall into the Level 5, “worse than useless” bucket, although they still meet various certification and audit requirements because they follow a process. First, information is collected, risks are assessed and analyzed (regardless of how accurately). Then management reviews and incorporates this information into its decision-making processes (the mitigation component), and the board of directors gets their report. The result is that the organization is imbued with a sense of confidence. Ironically, this approach could be used to pass an audit if the process is well-documented, simply because it can demonstrate that management actively incorporates the output, which is viewed as proof of its “effectiveness.” However, compliance does not equal information security (Carpenter, 2016).

As with Level 5, the “merely useless” Level 4 program also focuses on attempting to take poorly-defined data (detailed, but not quantitative, open to interpretation) and score it. But as long as it isn’t used to make decisions, it isn’t as harmful.

This document is intended to improve on the baseline program, Level 3.

But there is one more common approach to cybersecurity risk management that will be reviewed in the next section, and it is important because it uses nomenclature that is critical to the effective alignment of the cybersecurity risk manager’s approach and the business’ broader risk management program.

4.1. Another Common Approach to Security “Risk Management”

An unfortunately common approach to security risk management is described by Andrew Jaquith in “Security Metrics,” represented in the diagram below, which he refers to as the “Hamster Wheel of Pain”:




Figure 2. Andrew Jaquith’s Hamster Wheel of Pain

As Jaquith describes, “The fundamental problem with the Hamster Wheel model is simple: it captures the easy part of the risk management message (identification and fixing things) but misses the important parts (quantification and triage based on value)... To put this simply, for most vendors, ‘risk management’ really means ‘risk identification,’ followed by a frenzied process of fixing and re-identification” (Jaquith, 2007).

In this model, risk management is not really taking place, since no assessment of impact and likelihood are included. Instead, vulnerability assessments drive the work efforts. As described later in this document, it is the likelihood variable in isolation (as opposed to cross-referencing with impact and control effectiveness) that determines the priority. This is an urgency-based (versus importance-based) approach. In most cases, an intuitive understanding of existing threats is also factored in, which informs the limited prioritization.

Andrew Baze, abaze@hotmail.com

This approach usually falls below the “baseline” (intuitive, no formal risk management attempted) level as described in the previous section, because the intuition of subject matter experts is not what drives prioritization. Instead, it is often either the “vulnerability of the day” or “threat of the day” that determines the highest priority. However, if an organization relies heavily on experts to determine where their resources are spent, then the risk management “baseline” level may be achieved. However, this level is equivalent to not doing any formal risk management.

5. General Risk Management Benefits

While the risk management picture painted above looks bleak, especially in the cybersecurity domain, risk management is usually scoped more broadly, across various domains such as accounting, strategic planning, operations management, and compliance. This broader approach, usually in the form of an enterprise-wide program, can still have benefits, regardless of program maturity level.

For example, causing leadership to consider their risk appetite or causing risk-focused conversations to take place (whether during an initial evaluation or discussing a final report), a risk management program can help remove blinders that would otherwise remain in place.

In addition to causing key stakeholders to think differently about risk, meeting compliance requirements can be “a foot in the door,” the first exposure to systematic controls for many organizations. While it should not be viewed as the desired outcome, since standards compliance is not equal to security, it may be viewed as the de facto beginning of a program.

And in some cases, when risks are so obvious to all involved and put the organization in enough danger, the documentation process, if followed up with reporting and a drive for accountability, can improve the resilience of the organization. Some of the benefits may be coincidental, and some may provide a false sense of security, but real risk mitigations may still take place.

In the terms of the COSO Enterprise Risk Management Integrated Framework, stakeholder value is obtained by helping align risk appetite and strategy, enhancing risk response decisions, reducing operational surprises and losses, identifying and managing multiple and cross-enterprise risks, seizing opportunities, and improving deployment of

capital (Committee of Sponsoring Organizations for the Treadway Commission, 2004). Each of these areas is a potential area of opportunity, and should not be ignored when building out a robust cybersecurity assessment program.

While there are potential benefits to a broad risk management program, this high-level, usually more intuitive approach is not sufficient in the cybersecurity risk domain, because of the catastrophic consequences of not understanding or managing existing security control gaps. The variables reviewed earlier (impact, likelihood, threat, vulnerability, and control effectiveness) still need to be evaluated. But they do not need to be evaluated in the commonly prescribed order. The following sections describe how the process can be improved.

6. Risk Identification and Assessment

As described earlier as the first of the basic risk management program components, the area of Risk Identification and Assessment is the foundation for a relevant program. The old adage “proper planning prevents poor performance” certainly applies in this case. Proper identification and assessment allows practitioners to plan gap closure work. It is this effective, data-driven prioritization that will enable appropriate education and the best application of resources.

6.1. Risk Identification

Risk identification typically happens two ways. The first is that a risk manager consults existing risk documentation, industry organizations, peer groups, or simply performs an online search for “top risks for [industry area].” Note that this doesn’t provide prioritization, only identification. The second way is to systematically apply experience and intuition, whether the risk manager brainstorms (alone or with peers) or interviews relevant stakeholders. Well-informed, expert opinions, whether of experienced risk managers or business subject matter experts, can be an effective way to identify high-impact and high-probability risks.

Also note that identification can happen at any time, just as ongoing reassessment and clarification of impact or likelihood may take place.

6.2. Risk Assessment

Risk assessment is a different story. Many models exist, but most measure risks using the following three measures:

- 1) Impact: What is the negative outcome if this risk were to happen?
- 2) Likelihood: What is the probability of this risk happening?
- 3) Control effectiveness: How well do the existing controls (e.g., plans that have been implemented, rules the business currently follows) reduce the potential impact, likelihood, or both?

Assessment often involves in-person interviews, written surveys, or a combination of the two. While it may start with a list of known threats, common vulnerabilities, or a pre-populated, curated risk list that the interviewee can prioritize, the result is usually a set of opinions.

Impact, likelihood, and control effectiveness terms must be clearly understood in a cybersecurity context in order to show the connection between control effectiveness and realistic prioritization.

6.2.1. Variable #1: Impact

Impact is one of the more intuitive components of a risk assessment. It represents consequences. It is the effect on the business of a risk becoming a reality. Impacts are usually considered in the context of categories such as cost, schedule, strategy, operations, reputation, legal and compliance.

In terms of scoring, MITRE, a not-for-profit organization sponsored by the United States federal government, provides a detailed view of how one could implement an impact scoring system at <https://www.mitre.org/publications/systems-engineering-guide/acquisition-systems-engineering/risk-management/risk-impact-assessment-and-prioritization>.

However, any impact scoring system that allows one to repeatedly, consistently differentiate impacts relevant to a business' industry should suffice.

Andrew Baze, abaze@hotmail.com

6.2.2. Variable #2: Likelihood

Likelihood describes the probability of a risk happening.

According to OWASP (Open Web Application Security Project), two factors are the main determiners of likelihood: threats and vulnerabilities (OWASP Risk Rating Methodology, 2006).

Concerning the threat component of the OWASP equation, threat identification and prioritization does not need to be a complex, time-consuming operation. For many organizations, the Open Threat Taxonomy (OTT) may be more than sufficient. In addition to providing a low-cost set of choices for assessment purposes, this clear categorization and prioritization of threats may also be useful in justifying or communicating the seriousness of a security gap that requires funding. The OTT breaks down threats into physical, resource, personnel, and technical groupings. Since its ratings are based on industry data from over 150 organizations internationally, it is quite comprehensive in nature. In the words of the project's coordinator, James Tarala, "Why should every organization have to identify threats on their own? We all face the same threats, possibly to differing degrees. If we can agree on a common set of threats, we are free to focus on defending ourselves against them" (Tarala, Open Source Threat Taxonomy, 2016).

Using the OTT list of threats as a baseline, combined with an organization's known vulnerability state (or combined with the lack of data, and the assumption of vulnerability), can enable a well-grounded likelihood measurement.

See a sample of OTT threats with names and ratings (with 5.0 rated as the worst) below in Table 2 below:

Table 2 Technical threats example from Open Threat Taxonomy (Tarala, Open Source Threat Taxonomy, 2016)

Threat ID	Threat Action Name	Threat Rating
TEC-007	Credential Discovery via Sniffing	4.0
TEC-008	Credential Discovery via Brute Force	4.0
TEC-009	Credential Discovery via Cracking	4.0
TEC-010	Credential Discovery via Guessing	2.0
TEC-011	Credential Discovery via Pre-Computational Attacks	3.0
TEC-012	Misuse of System Credentials	3.0
TEC-013	Escalation of Privilege	5.0
TEC-014	Abuse of System Privileges	4.0
TEC-015	Memory Manipulation	4.0
TEC-016	Cache Poisoning	3.0
TEC-017	Physical Manipulation of Technical Device	2.0
TEC-018	Manipulation of Trusted System	4.0
TEC-019	Cryptanalysis	1.0
TEC-020	Data Leakage / Theft	3.0
TEC-021	Denial of Service	2.0
TEC-022	Maintaining System Persistence	5.0
TEC-023	Manipulation of Data in Transit / Use	2.0

As described earlier in the Hamster Wheel of Pain cycle, where vulnerability is the motivator, the likelihood variable should not consume all of the risk evaluation focus (without sufficiently considering impact and control effectiveness), or the result will be reactive, ineffective prioritization.

More importantly, however, than impact and likelihood, from a day-to-day cybersecurity operations perspective, is the measurement of control effectiveness.

6.2.3. Variable #3: Control Effectiveness

The control effectiveness evaluation allows an assessor to overlay reality on the otherwise fuzzy and necessarily unreal (risk has to be expressed in a future state) risk

assessment process. It is via this measure that one can bridge the gap between what could happen and what is happening currently.

According to Stephen Northcutt of SANS, “Security controls are technical or administrative safeguards or counter measures to avoid, counteract or minimize loss or unavailability due to threats acting on their matching vulnerability...” (Northcutt, 2016).

Whether these safeguards are implemented in the way they were intended is the focus of the control effectiveness evaluation. The recommended set of technical security controls is the Center for Internet Security 20 Critical Security Controls (CIS 20 CSC). More context and a detailed evaluation description are described in later sections.

6.2.4. Last Step: Prioritization

The conclusion of the identification and assessment portion of a risk management program is the prioritization of the risks. This prioritized output is the input to the next portion of the program, in which risk treatment plans are developed (see “Develop Risk Treatment Plans,” Figure 1).

Effective risk management cannot happen without prioritization. In a less mature organization, this may happen intuitively. And in many cases, when few serious problems exist, prioritization is simple. For example, it should be easy to determine that removing a persistent threat that has already compromised sensitive servers and is actively exfiltrating data is more important than rewriting the security incident response procedures.

However, when many serious problems exist, prioritization is critical, because there are always limited resources, whether time, money, or people, and the prioritization will help define where the “cut line” is, determining which work will be done and not done in the next spending or work cycle.

As described earlier, a typical method of prioritization in risk management is to take the expected impact to the organization and cross-reference it with the likelihood of that risk happening, which results in an inherent risk value.

Note that the term “inherent risk” is the combination of impact and likelihood without considering any existing controls. In other words, it represents the risk in an environment

where no controls exist whatsoever. “Residual risk” describes how much risk remains after controls (e.g., mitigation plans) have been implemented.

In the cybersecurity risk management realm, inherent risk and the view it provides has relatively little value. It could be very useful for an audit organization that needs to determine where to spend its resources, relative to other areas (Working with Inherent and Residual Risk, 2011), but when looking at security risks in isolation, this view is seldom useful. Without the overlay of control effectiveness, security risks will almost always be reflected with high impact and high probability. When all of the potential risks are closely clustered (see Figure 3 below), it is difficult to effectively differentiate. This is where control effectiveness and residual risk play a critical role.

Once controls have been applied, residual risk can be determined. The evaluation of the effectiveness or capability of those controls is the determining factor for how important a risk is relative to others.

7. How to Prioritize Better by Considering Control Effectiveness

A risk matrix is a commonly-used diagram often used to communicate risk importance. In most cases, only impact and likelihood are represented. This is insufficient. If the effectiveness of current controls is not taken into consideration (and the idea of measuring cybersecurity risk as if there were no firewalls, scanning, antivirus, and other commonplace controls in place may seem ludicrous), effective prioritization will not be possible. The resulting inaccurate prioritization could result in valuable resources being spent on risk areas that may not be very important, and which may not deserve any resources (i.e., the risk would be accepted with no mitigation).

For the purposes of this section, the following sample data set will be used to illustrate the value of control effectiveness in risk mitigation prioritization.

Table 3. Risk Assessment Values (hypothetical)

Risk ID	Impact (1-10)	Likelihood (1-10)	Impact-Likelihood Ratio (I*L*0.1)	Security Control Implementation (1-10, 1 is best)
1	10	10	10	2
2	9	10	9	8
3	8	9	7.2	5
4	5	9	4.5	1
5	6	9	5.4	10

A quick review will help clarify the following charts. In Table 3 above:

- “Impact” describes the amount of expected harm, with 1 being little and 10 being worst (if no controls were in place).
- “Likelihood” describes how likely the risk will become real (if no controls were in place).
- “Impact-Likelihood Ratio” is $\text{Impact} * \text{Likelihood} * 0.1$ (used in the risk matrix in Figure 4 below).
- “Control Implementation” describes how effectively security controls have been implemented, with 1 indicating that best-in-class security controls are in place, and 10 indicating that no effective security controls are in place (relative to the individual risk).

Risks 1-5 in Table 3 above have been plotted on a 5x5 risk matrix, organized to show the most severe risks in the upper right section. This grid format is commonly used in risk management, and could just as easily be represented by a 4x4 or 6x6 grid. The X or Y axis could go from best to worst, or vice versa. As long as it is used consistently within an organization, the exact layout is not important. Additional examples with similar layouts can be seen at MITRE (Sample - Safety Management System Risk Matrix, 2016), Washington State Department of Transportation (Risk Management Plan, 2016), OWASP (OWASP Risk Rating Methodology, 2016), and The Australian Centre for Healthcare Governance (Risk Matrix, Consequence And Likelihood Tables, 2016).

The Inherent Risk Matrix below in Figure 3 shows only impact and likelihood data from Table 3 factored in.

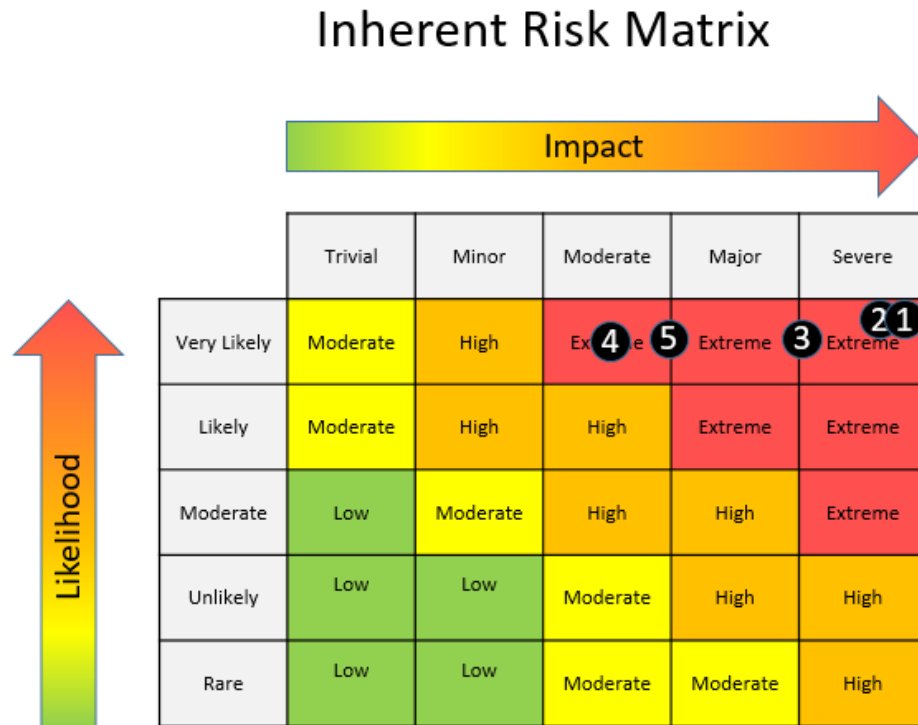


Figure 3. Impact and likelihood charted, an undifferentiated view of inherent cybersecurity risk

In this view above, all risks appear to be serious, and in need of urgent attention. Obviously, there is only limited value in showing that all of the organization's cybersecurity risks are plotted in the upper right corner of the matrix. It could be significant, if the goal is to show inherent cybersecurity risk relative to other risk types (which may be inherently much less impactful or likely) as part of a much broader risk management program. However, for the purposes of cybersecurity risk management, there is little value in placing all risks in one very small area, because it makes prioritization difficult.

This model needs improvement, in order to enable realistic prioritization. This can be done by representing the impact and likelihood ratio on the Y-axis, and control implementation effectiveness on the X-axis. As illustrated below, incorporating control effectiveness results in the ability to show residual risk, and will enable much better mitigation prioritization.

Please note that the control implementation as shown in the X-axis of Figure 4 below worsens from left to right.

With the Security Controls data integrated, a different picture emerges.




Figure 4. Overlaying control effectiveness enables differentiation between risks and illustrates residual risk

In Figure 4, now that a realistic evaluation of security control effectiveness has been incorporated, more accurate prioritization can take place. The most extreme risks will obviously take precedence over the low, moderate, and high-risk areas.

When only impact and likelihood were considered, risk prioritization order by Risk ID would have been 1, 2, 3, 5, and 4. All were categorized as “Extreme.”

Now that security control implementation has been factored in, the order is quite different. Only risks 2 and 5 are “Extreme”, risk 3 is “High,” and risks 1 and 4 trail at Moderate/High and Low, respectively. Depending on how the organization treats risks, Risk

4 may not receive any resources. According to Figure 3, it was in need of urgent attention. According to Figure 4, it is last on the list.

This is a simple illustration, with few risks and some broad differentiation introduced for dramatic effect. However, it illustrates the value of paying special consideration to the effectiveness of control implementation as part of a risk assessment.

Given the reactive nature of many organizations (see Hamster Wheel of Pain example earlier, with a simple focus on vulnerability identification), the control evaluation and subsequent mapping are key to effective prioritization.

8. The Crux: A Detailed Control Evaluation

A security practitioner does not need to understand risk management in order to assess security controls' effectiveness. Based on an evaluation, existing controls may be very effective on one scale, ineffective on another, may meet NIST 800-53 standards, may not meet ISO 27002, and may still pass an internal audit. In any case, all kinds of evaluations can be done without understanding what risks exist to the organization. But focusing on controls and how well they work can change that.

Control effectiveness is something a security operator deals with on a daily basis. Take the CIS Critical Security Control #4 (more on this framework later) for example: "Continuous Vulnerability Assessment and Remediation." The following questions allow one to quickly determine whether the control is implemented, and a review of the evidence (e.g. log files) determines whether the organization is capable in this context:

- 1) Are scans run weekly?
- 2) Is automated patch-management software running?
- 3) Are scan results compared back-to-back, and to logs?

Regardless of what risk management programs may exist, formal or not, determining the current level of control effectiveness should absolutely be within the cybersecurity team's charter, and built into its ongoing operation. There is no shortage of control standards, checklists, and internal policy requirements that could be compared against, and more likely than not, they are already being used. A comparison of any organization's existing

Andrew Baze, abaze@hotmail.com

cybersecurity program (assuming it is at least moderately competent) and the CIS 20 CSC will undoubtedly show many commonalities in controls and processes, regardless of whether the prioritization is aligned.

With this in mind, the organization's CISO or less formal equivalent (depending on the size of the business) should have a clear understanding of the policy, standards that must be met, and the organization's ability to meet those standards. When this can be evaluated, the most important variable in the risk assessment equation is covered.

While management may have a strong grasp of desired security outcomes, if the current state cannot be clearly described, then the risk manager has no basis by which to measure the size of any gaps identified. In other words, the adage "You can't manage what you can't measure" applies here. Since the program should be able to clearly identify gaps, determine the impact and likelihood of those gaps being exploited, and determine the cost of shrinking, closing, or transferring those gaps, measurement is required.

9. Measuring Control Effectiveness – the CIS 20 CSC

When assessing an organization's technical security controls, the assessor has a very thorough set of industry-standard controls to compare against: the CIS (Center for Internet Security) 20 Critical Security Controls (CSC). The 20 CSC are designed with five tenets in mind (Tarala, Welcome to the CIS Controls, 2016):

- 1) **Offense informs defense**
 - a. This means that the controls are all realistic. Offensive actions are taking place right now, in the wild, and they could be mitigated by these controls. If those actions were ever to cease, then the corresponding control would eventually be removed.
- 2) **Prioritization**
 - a. The controls are listed in priority order. Note: this means that initial prioritization of risks, when mapped one-to-one with controls, has already been done at a high level.
- 3) **Metrics**

- a. Each one of the controls is designed to be readily measured, which obviously aids in the controls' evaluation.
- 4) Continuous Diagnostics and Mitigation
- a. Each of the controls is designed to be monitored on an ongoing basis, which should then result in a clearer mitigation approach.
- 5) Automation
- a. Controls should be automatable, enabling an organization's ability to scale to managing multiple controls.

Considering that real-world, impactful and likely (applicable risk terms) offensive actions have already been considered during the prioritization of these controls, a rough risk prioritization has also been completed.

Of course, a tailored impact and likelihood evaluation for a unique organization must still take place, but it can more realistically take place after a controls assessment has been completed.

The entire list of 20 controls are represented in Table 1 below:

Table 4. Center for Internet Security 20 Critical Security Controls

1	Inventory of Authorized and Unauthorized Devices
2	Inventory of Authorized and Unauthorized Software
3	Secure Configurations for Hardware and Software
4	Continuous Vulnerability Assessment and Remediation
5	Controlled Use of Administrative Privileges
6	Maintenance, Monitoring, and Analysis of Audit Logs
7	Email and Web Browser Protections
8	Malware defenses
9	Limitation and control of network ports
10	Data Recovery Capability
11	Secure Configurations for Network Devices
12	Boundary Defense
13	Data Protection
14	Controlled Access Based on the Need to Know
15	Wireless Access Control
16	Account Monitoring and Control

17	Security Skills Assessment and Appropriate Training to Fill Gaps
18	Application Software Security
19	Incident Response and Management
20	Penetration Tests and Red Team Exercises

Another important factor to be aware of, in addition to the built-in prioritization of the controls by CIS, is that the first five are considered foundational.

As one could easily intuit, determining what hardware exists on a network is foundational to securing that hardware (CSC #1). Secondly, determining what software runs on the network will determine how that hardware can be configured (CSC #2). The next logical step is to configure the hardware and software securely (CSC #3). After that, looking for new problems (vulnerabilities) on the hardware and software makes sense, which is followed by ensuring that only certain people have access to the hardware, software, and their configurations (CSC #4 and #5 respectively).

The Center for Internet Security is not the only organization that has invested in such research. The Australian Department of Defense (DSD) has determined that “implementing the Top 4 [of their list of 35 controls] will mitigate at least 85% of the intrusion techniques that the Australian Cyber Security Center responds to” (Top 4 Strategies to Mitigate Targeted Cyber Intrusions: Mandatory Requirement Explained, 2013).

These Top 4 DSD controls (and their loose mappings to CIS Critical Security Controls) are:

- 1) Application whitelisting (CSC #2, Inventory of Authorized and Unauthorized Software).
- 2) Patch applications (CSC #4, Continuous Vulnerability Assessment and Remediation [patching]).
- 3) Patch the operating system (CSC #4).
- 4) Minimize administrative privileges (CSC #5 Controlled Use of Administrative Privileges).

CIS control #1 (Inventory of Authorized and Unauthorized Devices) is a prerequisite for implementing DSD controls #1-4. CIS control #3 (Secure Configurations for Hardware and Software) should significantly reduce the effort of implementing DSD controls #2 and #3.

In fact, the language of DSD control #3 (Patch the operating system) includes references to discovering (inventorying) all devices in the environment as well as ensuring that configurations and deployments already contain relevant patches (Top 4 Strategies to Mitigate Targeted Cyber Intrusions: Mandatory Requirement Explained, 2013).

If the list of 20 controls seems daunting, consider an initial assessment that focuses just on the first five. Not only will they significantly reduce a disproportionate amount of risk (in the 85% range), they will be challenging enough to implement on their own. This doesn't mean that mitigation activities which tie to the other controls need to cease, but that they might be represented as being lower priority relative to those top five.

10. Detailed Review: Evaluating CSC #1

Evaluating a technical security control in more detail will provide a useful example. The first of the CIS 20 CSC is titled "Inventory of Authorized and Unauthorized Devices." Its evaluation answers the question "Does the organization effectively inventory its devices?" The well-known security problem in this context is that when an organization does not know what devices it has in its inventory, there is no way to ensure the devices are secure.

The following actions of CSC #1 have associated metrics and evaluation questions, and are designed to be measurable (see tenets 1-5 mentioned earlier).

Table 5. CSC #1 Evaluation Actions

1.1	Deploy an automated asset inventory discovery tool and use it to build a preliminary inventory of systems connected to an organization's public and private network(s). Both active tools that scan through IPv4 or IPv6 network address ranges and passive tools that identify hosts based on analyzing their traffic should be employed.
1.2	If the organization is dynamically assigning addresses using DHCP, then deploy dynamic host configuration protocol (DHCP) server logging, and use this information to improve the asset inventory and help detect unknown systems.
1.3	Ensure that all equipment acquisitions automatically update the inventory system as new, approved devices are connected to the network.

1.4	Maintain an asset inventory of all systems connected to the network and the network devices themselves, recording at least the network addresses, machine name(s), purpose of each system, an asset owner responsible for each device, and the department associated with each device. The inventory should include every system that has an Internet protocol (IP) address on the network, including but not limited to desktops, laptops, servers, network equipment (routers, switches, firewalls, etc.), printers, storage area networks, Voice-Over-IP telephones, multi-homed addresses, virtual addresses, etc. The asset inventory created must also include data on whether the device is a portable and/or personal device. Devices such as mobile phones, tablets, laptops, and other portable electronic devices that store or process data must be identified, regardless of whether they are attached to the organization's network.
1.5	Deploy network level authentication via 802.1x to limit and control which devices can be connected to the network. The 802.1x must be tied into the inventory data to determine authorized versus unauthorized systems.
1.6	Use client certificates to validate and authenticate systems prior to connecting to the private network.

Core evaluation tests and measures (metrics) to evaluate this control (A Measurement Companion to the CIS Critical Security Controls (Version 6), 2015) are as follows:

Evaluation tests:

- Place ten unauthorized devices on various portions of the organization's network unannounced to see how long it takes for them to be detected
 - They should be placed on multiple subnets
 - Two should be in the asset inventory database
 - Devices should be detected within 24 hours
 - Devices should be isolated within 1 hour of detection
 - Details regarding location, department should be recorded

Measures:

- Number of unauthorized devices presently on network
- Average time to remove unauthorized devices from network
- Percentage of systems on network not using NLA (Network Level Authentication)

- Number of hardware devices recently blocked from connecting to the network by NLA system
- Time to detect new devices added to organization's network (minutes)
- Time to isolate/remove unauthorized devices from organization's network (minutes)

Walking through each of the actions (1.1-1.6), testing how well those controls have been implemented with the tests above, and applying the described measures on an ongoing basis will provide a realistic assessment of the effectiveness of the current implementation of CSC #1.

Using the free resources available at CISecurity.org (which includes all controls and associated activities, evaluation tests, measures, and much more), the remaining 19 controls can be just as thoroughly evaluated.

11. Mapping Controls to Risks

Even though an organization will benefit most, immediately, from the detailed control evaluation, understanding the risk of allowing the current state to continue is still important. And constructing risk statements need not be time-consuming or complicated.

Following the "If – Then" guideline from MITRE can easily be used to create a risk statement. It is constructed in this manner:

"IF - THEN Risk Statement

- Example:
 - Requirement reads: "Use Common Operational Picture (COP) in DII COE Release 1.5"
 - Identified risk: availability of DII COE version 1.5 when needed
- Risk statement:
 - IF DII COE version 1.5 is more than 1 month late, THEN Program xyz release 1 will experience a day for day schedule slip" (Risk Management Toolkit, 2016)

Andrew Baze, abaze@hotmail.com

To summarize, if a particular condition exists, then the following consequences will occur.

A simplified risk statement for CSC #1 is “Inaccurate or improper inventory of authorized and unauthorized devices could result in a data breach.”

In other words, “If devices aren’t inventoried, the organization is open to being breached.” Of course, there will be many consequence-related steps between an un-inventoried device connecting to the network and the exfiltration of sensitive data, but this statement is intended to be as simple as possible, and it does represent both the condition and the consequences. With a control this important, and a risk this impactful and likely, a simplification can be entirely appropriate.

A more detailed risk statement for an organization could be reflected as “The lack of implementation of an automated asset discovery tool that connects to the organization’s public and private networks combined with the lack of an asset inventory system could result in missing or outdated inventory data used by the security team to scan, patch, deploy new services, and other key functionality. This, in turn, could result in insecure systems on the network that provide ingress to attackers, resulting in a data breach, fines, and loss of customer trust.”

Similar risk statements (simple, exhaustive, or somewhere in between, depending on the organization’s preference) can be constructed for Critical Security Controls 2-20, and will comprise a comprehensive technical security risk register. And when they have been evaluated and prioritized, next steps for mitigation should be easier than ever to clearly identify.

12. Rearranging the Risk Assessment Processes

What if an organization’s entire cybersecurity risk management program consisted of nothing other than a control effectiveness assessment? Could the other key components be realistically evaluated? Remember that inherent risk is often intended to be evaluated in absence of existing controls, which usually has limited value from a cybersecurity perspective. A business with an Internet presence and no cybersecurity controls would probably not last more than a few hours before being seriously compromised. This illustrates

Andrew Baze, abaze@hotmail.com

how a security-focused view of risk and a more broadly-scoped (e.g. enterprise-level or audit-focused) risk management program can have fundamental differences.

Instead of a more traditional approach to risk management, where risk assessment focuses initially on impact and likelihood, then overlays control effectiveness, the approach in the cybersecurity domain should be reversed, initially focusing on the control effectiveness assessment, which can often better inform both impact and likelihood. A more specific example of this approach is that evaluating CSC #4 (Continuous Vulnerability Assessment and Remediation) should result in a better understanding of existing vulnerabilities. This better understanding can then be used to more accurately define the likelihood (which incorporates threat and vulnerability) of a risk occurring, and could also provide more context for the expected impact.

Below is a common risk assessment approach (see ISACA and ISO references earlier in Section 2), starting with risk identification. This traditional method focuses on starting with a list of risks, working with partners to determine the right impact and likelihood levels (see Inherent Risk Matrix, Figure 3), then doing a separate overlay of control effectiveness in order to prioritize risk treatment (see Residual Risk Matrix, Figures 4).




Figure 5. A common approach to risk assessment

However, this approach can be improved upon in a cybersecurity context. An initial assumption that most risks in this domain will have a high impact and high likelihood reduces the value of that portion of the assessment process, at least at first. Instead, the focus should be on a thorough evaluation of security controls. The process of evaluating these controls will much more effectively inform an accurate evaluation of impact and likelihood simultaneously.

In addition, a subsequent mapping of security controls to corresponding risks will clarify the risk prioritization process, since data on a per-control basis can be directly applied on a per-risk basis. This approach, illustrated below, changes the focus of the discussion, starting with a control effectiveness evaluation. Risk prioritization is a secondary outcome. While secondary, however, the risks in this approach will be more accurately evaluated, since they will be grounded in real, current state of security control implementation.

Figure 6 below shows the reordering of the evaluation process, beginning with the control effectiveness assessment, resulting in improvements in both accuracy and prioritization.




Figure 6. A more effective approach to cybersecurity risk assessment

Changing the approach and starting a risk assessment by focusing on control effectiveness, then using that output to better inform impact, likelihood, and remediation prioritization will enable security operators and risk practitioners to produce more impactful

and relevant results for the organization. They will be able to circumvent the Hamster Wheel of Pain and mature beyond the “merely useless” and “no formal risk management” approaches. And they will be able to realistically improve their organization’s security by focusing on industry-vetted controls, prioritizing based on the current, known gaps versus theoretical ones.

13. Conclusion

Evaluating an organization using the 20 Critical Security Controls, or even just the top five, will provide concrete guidance on a topic that many security risk management programs struggle with: how to spend money where it matters most, based on what is wrong currently versus what may be a problem later.

Just as inventorying devices is the first priority of the 20 CSC, a control assessment should be the first and most important step in a security risk assessment. Not only is it valuable in its own right, but it will also provide valuable context for the other components of the assessment.

As described in the phrase often attributed to Einstein, “make everything as simple as possible, but not simpler.” In this case, grounding a risk management program by starting with reality management involves simply using a realistic control set and initially focusing on the known over the unknown. This approach will produce more tangible, actionable, and defensible results.

14. References

- (2015). A Measurement Companion to the CIS Critical Security Controls (Version 6). Center for Internet Security.
- Carpenter, B. C. (2016, June 24). Why Compliance is NOT the Answer. Retrieved from ITSecurity: <http://www.itsecurity.com/security.htm?s=14930>
- CISM Review Manual 2015. (2014). Rolling Meadows: ISACA.
- Committee of Sponsoring Organizations for the Treadway Commission. (2004, September). Enterprise Risk Management - Integrated Framework. Retrieved from COSO -

- Committee of Sponsoring Organizations of the Treadway Commission:
http://www.coso.org/documents/COSO_ERM_ExecutiveSummary.pdf
 Explore Terms: A Glossary of Common Cybersecurity Terminology). (2016, June 24).
 Retrieved from NICCS National Initiative for Cybersecurity Careers and Studies:
<https://niccs.us-cert.gov/glossary>
- Hubbard, D. W. (2009). *The Failure of Risk Management*. New Jersey: John Wiley & Sons.
- ISO 31000:2009. (2009). Retrieved from ISO:
<https://www.iso.org/obp/ui/#iso:std:iso:31000:ed-1:v1:en>
- Jaquith, A. (2007). *Security Metrics: Replacing Fear, Uncertainty, and Doubt*. Upper Saddle River: Pearson Education, Inc.
- Kaplan, M. (2016, June 4). *Managing Risks: A New Framework*. Retrieved from www.HBR.org: <https://hbr.org/2012/06/managing-risks-a-new-framework/ar/1>
- Merriam-Webster. (2016, June 24). Retrieved from Merriam-Webster: <http://www.merriam-webster.com/dictionary/risk>
- Northcutt, S. (2016, June 4). *Security Laboratory*. Retrieved from www.SANS.edu:
<http://www.sans.edu/research/security-laboratory/article/security-controls>
- OWASP Risk Rating Methodology. (2016, June 4). Retrieved from OWASP.org:
https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology
- Risk Management Plan. (2016, July 5). Retrieved from Washington State Department of Transportation:
<http://www.wsdot.wa.gov/publications/fulltext/ProjectMgmt/PMOG/RiskManagementPlan.xls>
- Risk Management Toolkit. (2016, June 6). Retrieved from MITRE.org:
<http://www2.mitre.org/work/sepo/toolkits/risk/procedures/RiskStatements.html>
- Risk Matrix, Consequence And Likelihood Tables. (2016, July 5). Retrieved from The Australian Centre for Healthcare Governance:
<http://www.healthcaregovernance.org.au/docs/risk-matrix.doc>
- Sample - Safety Management System Risk Matrix. (2016, July 7). Retrieved from [MITRE](http://MITRE.org):
https://www.mitrecaasd.org/SMS/doc/Sample_Risk_Matrix.pdf
- Taleb, N. N. (2014). *Antifragile*. New York: Random House LLC.
- Tarala, J. (2016). *Audit/Security 566.1*. The SANS Institute.

Tarala, J. (2016, Jun 4). Open Source Threat Taxonomy. Retrieved from
www.AuditScripts.com:

http://www.auditscripts.com/resources/open_threat_taxonomy_v1.1a.pdf

Tarala, J. (2016, June 6). Welcome to the CIS Controls. Retrieved from Center for Internet
Security: <https://www.cisecurity.org/critical-controls.cfm>

Top 4 Strategies to Mitigate Targeted Cyber Intrusions: Mandatory Requirement Explained.
(2013, July). Retrieved from Australian Government Department of Defense:
<http://www.asd.gov.au/infosec/top-mitigations/top-4-strategies-explained.htm>

Working with Inherent and Residual Risk. (2011, November 5). Retrieved from Exploring the
Black Box: <http://exploringtheblackbox.net/exploringtheblackbox/2011/11/5/working-with-inherent-and-residual-risk.html>

Appendix A: Control Effectiveness / Risk Assessment Toolkit

To summarize the overall process, the following checklist describes how to conduct a thorough cybersecurity risk assessment. The tools referenced below (and earlier in this document) can save immense amounts of time and research, although it will still take significant time and effort to effectively evaluate an organization.

- ✓ Assess Control Effectiveness
 - Use Center for Internet Security Top 20 Critical Security Controls, which includes detailed descriptions and measurement tools
 - <https://www.cisecurity.org/critical-controls.cfm>
- ✓ Identify Risks
 - Map risks 1:1 to each of the CSC
 - Additional risks could be mapped to CSC sub-controls as needed
 - Consider the simple If-Then model, which could be described using the current control state:
 - “If [the current state of this control] continues, the organization will continue to be [as broken as the control evaluation indicated], potentially resulting in [more of what is probably happening now, or whatever might be imminent].”
 - <http://www2.mitre.org/work/sepo/toolkits/risk/procedures/RiskStatements.html>
- ✓ Review Impact
 - Use MITRE or any other applicable impact measurement scale
 - <http://www2.mitre.org/work/sepo/toolkits/risk/procedures/RiskStatements.html>
- ✓ Review Likelihood
 - Consider Threat and Vulnerability variables

- Use Open Threat Taxonomy and its prioritized threat list
 - http://www.auditscripts.com/resources/open_threat_taxonomy_v1.1a.pdf
- ✓ Review prioritized list and plan accordingly
 - Execute on most critical risk areas, closing existing security gaps in a defensible order
 - See Sections 5 and 6 in the OWASP Risk Rating Methodology: decide what to fix and customize the model as needed
 - https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology

This list uses free materials and countless hours of research by many organizations whose sole goal is to improve security. Whether using these sources or something comparable, focusing on an evaluation of current control effectiveness will enable a cybersecurity risk management program to focus on solving real problems in a reasonable order, better applying its scarce resources to secure the organization.