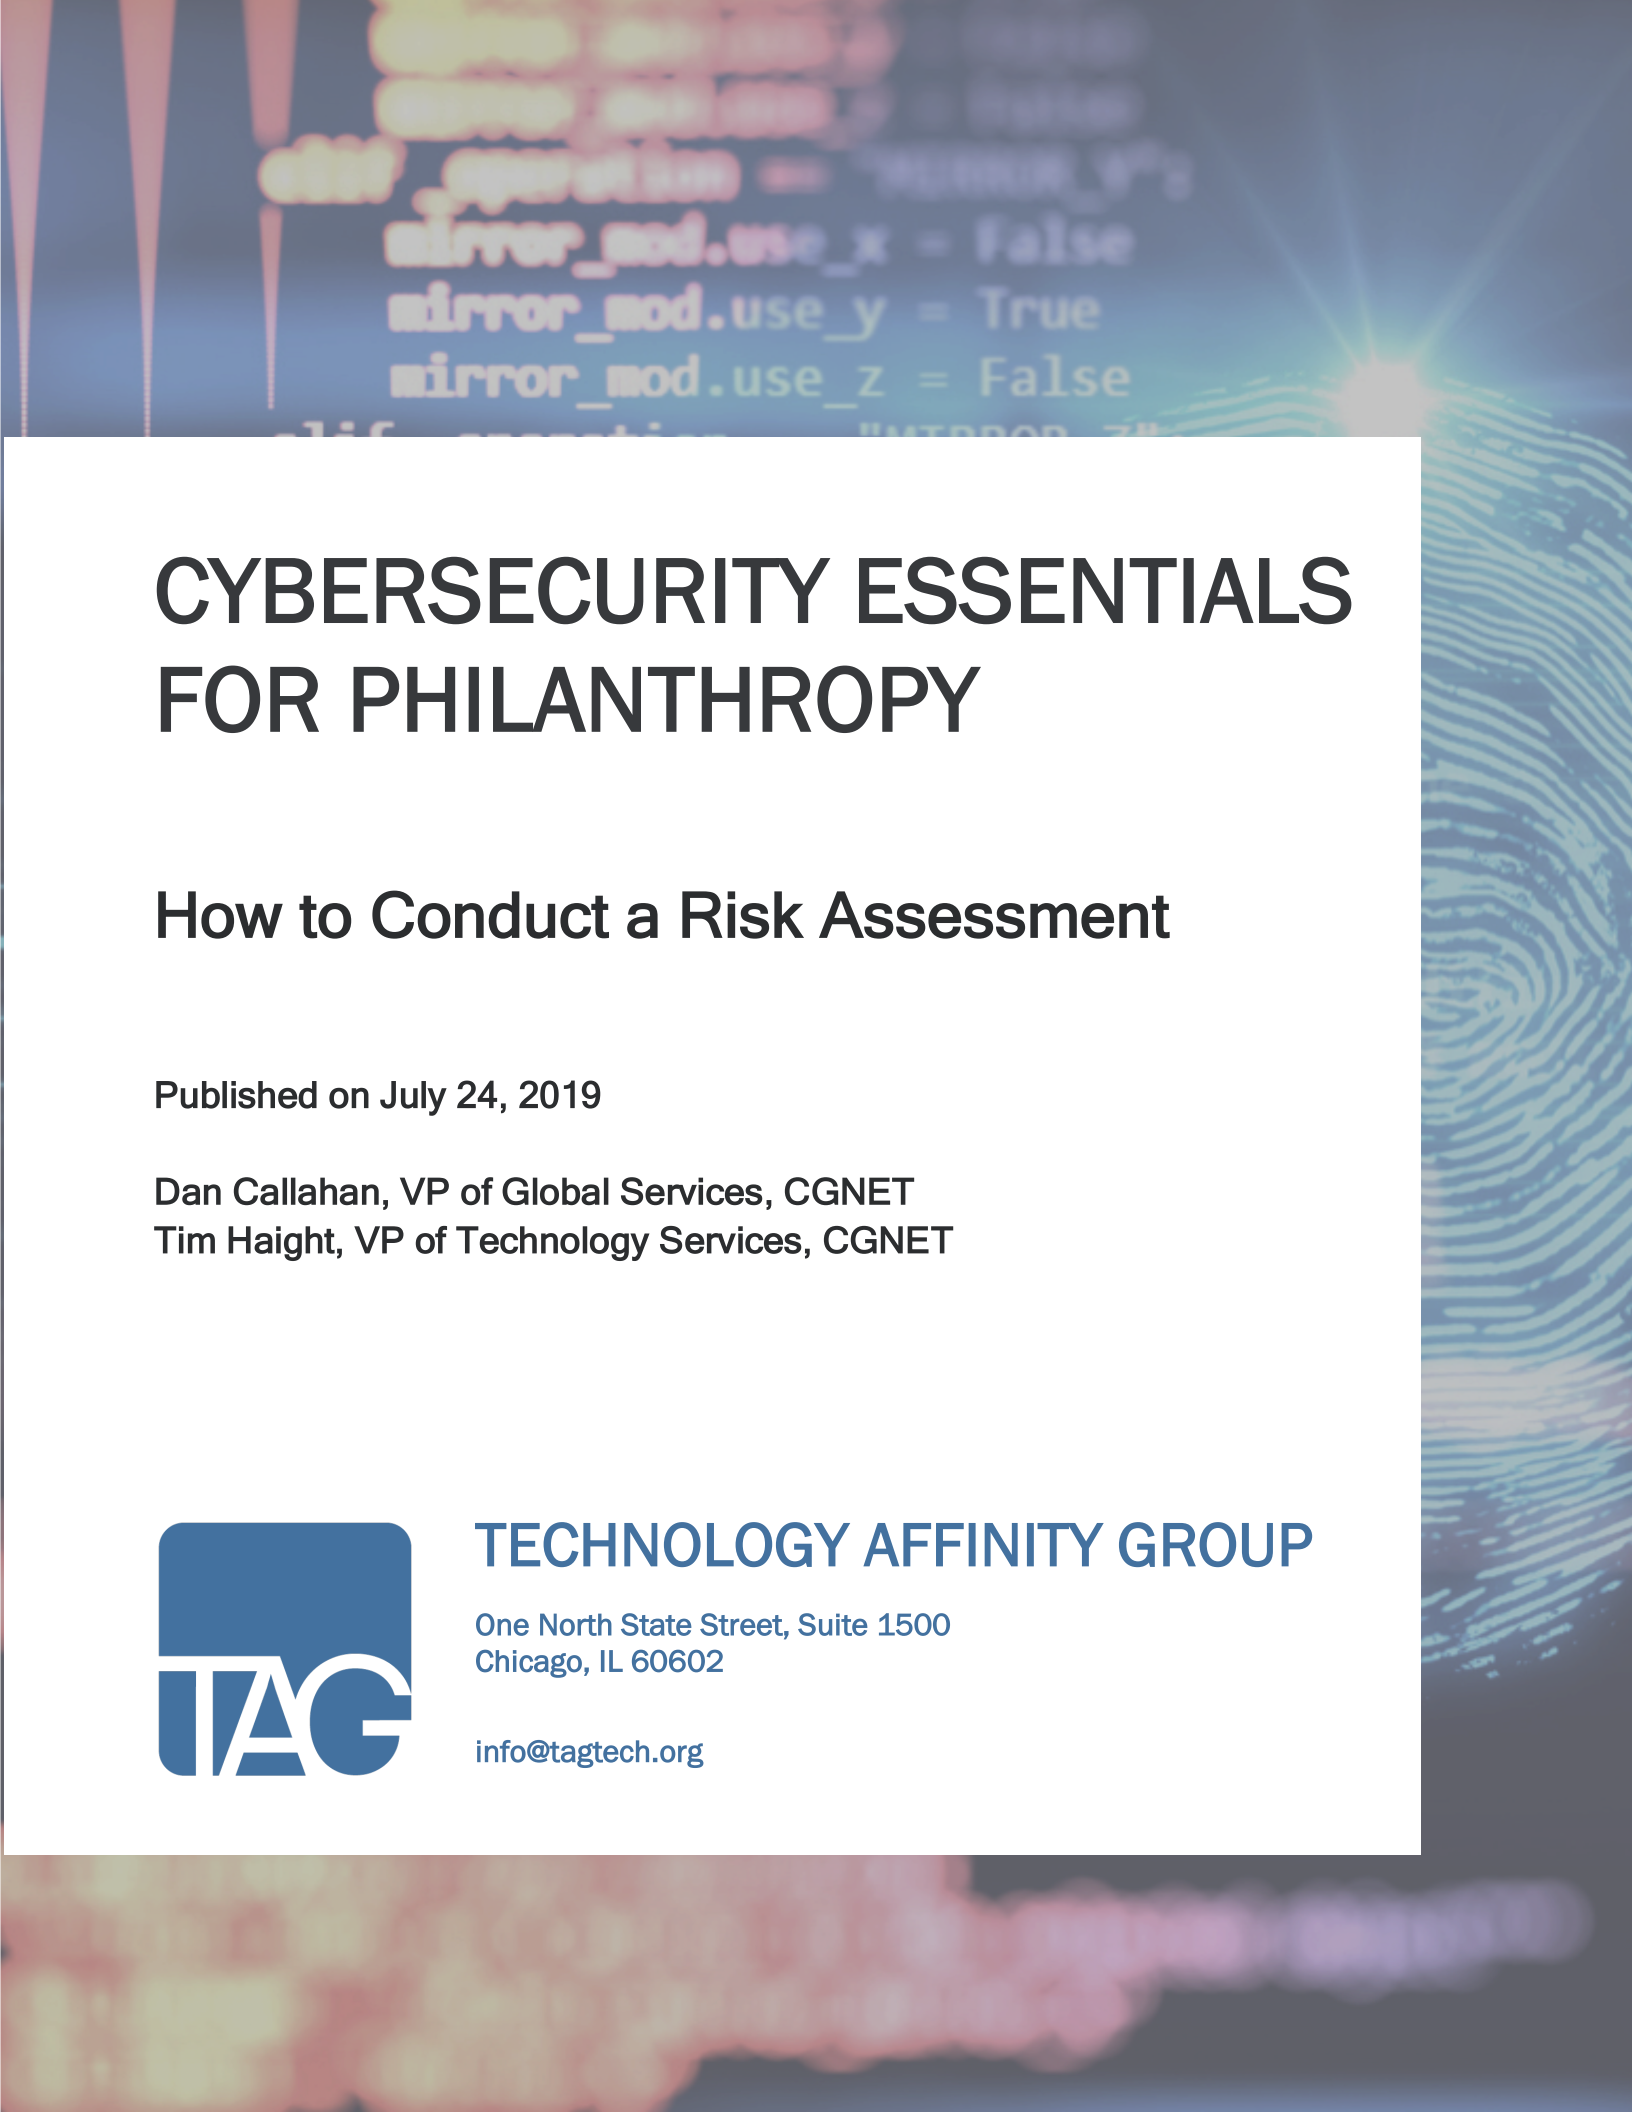

CYBERSECURITY ESSENTIALS FOR PHILANTHROPY

How to Conduct a Risk Assessment

Published on July 24, 2019

Dan Callahan, VP of Global Services, CGNET
Tim Haight, VP of Technology Services, CGNET

TECHNOLOGY AFFINITY GROUP

One North State Street, Suite 1500
Chicago, IL 60602

info@tagtech.org

OVERVIEW

PRAGMATIC INSIGHT FROM IT LEADERS IN PHILANTHROPY

Executives in philanthropy are increasingly concerned about cybersecurity. Phishing attacks are weekly, if not daily, and the stakes of a breach are high. In spite of our best attempts as a sector to develop robust practices, 21% of respondents to TAG's 2018 State of Philanthropy Tech survey reported experiencing a security breach in the past two years. For private independent foundations, the breach rate was even higher at 24%. No wonder there's growing concern.

> **Through the *CyberSecurity Essentials for Philanthropy* series, we aim to reduce your organization's risk and establish best practices throughout the sector.**

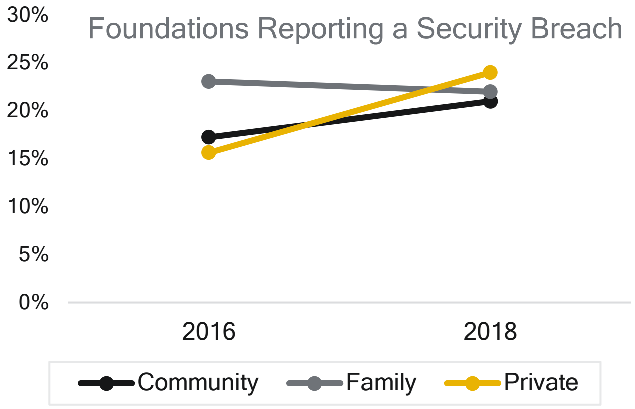

Source: 2018 State of Philanthropy Tech Survey, available at http://www.tagtech.org/philanthropytech2018.

This publication offers best practices and suggestions based on the collective on-the-ground knowledge and experience of your peers at philanthropic organizations across North America. On behalf of the members and directors of the Technology Affinity Group, we're grateful for the authors' generosity and expertise.

JAMES R. RUTT
Chief Information Officer, Dana Foundation
President, Board of Directors, Technology Affinity Group

CHANTAL E. FORSTER
Executive Director, Technology Affinity Group

HERE'S A RECIPE FOR CREATING A CYBERSECURITY RISK ASSESSMENT

This document answers the question: *How do I figure out what information assets are at risk and what would happen if something was compromised?*

> **"We never knew we had so much organizational information in so many places."**
> **- Anonymous Foundation IT Manager**

We've written this document to provide pragmatic strategies and real-world tactics based on our everyday experience working with IT leaders. In this document you'll find:

1. How to discover your information assets.
2. How to assess the threats to these assets.
3. How to estimate the impact if an asset is hacked.
4. How to summarize your information.
5. How to use your summary to act.

The task of putting together a cybersecurity risk assessment can seem overwhelming. You might be asking, "Where do I start? How do I predict what could happen? How do I estimate the (negative) impacts?" Such questions might lead you to put a risk assessment off.

We're here to tell you that you can do this. You're going to have to get into some cybersecurity "weeds" but we won't let you get lost. You already have the practiced judgement to make reasonable estimates of what could happen and how much it could hurt your organization. This checklist is your guide to conducting an actionable risk assessment for your organization.

Let's get started.

INTRODUCTION

WHY CONDUCT A RISK ASSESSMENT?

Why conduct a risk assessment? There are three reasons:

1. Costs.
2. Priorities.
3. The Unknown.

Let's examine them in reverse order.

- **The Unknown**: Cybercrime has become big business. We are constantly told that getting hacked is inevitable. So, are all your bases covered? Where is your list of bases, exactly? Less metaphorically, have you found all the risks? An early part of risk assessment is risk identification, where you break down the unknown into manageable chunks.
- **Priorities**: Once you've found all the risks, how do you react to them? Which do you tolerate? Which do you try to reduce? Risk assessment allows you to make a defensible choice.
- **Costs**: You shouldn't pay more to address a risk than what the consequences of the risk would cost. How do you figure that out?

Ultimately, risk assessment is about setting an agenda. You decide what your risks are, which risks you will address and in what order, and why it's worth doing. It's also about accountability. Your security plan isn't just what you decided to do; you have evidence.

THE OBJECT OF THE GAME

The purpose of a risk assessment is to prioritize risks so you can decide which to spend time and money addressing. Say you have two potential risk events: a user's computer being infected with malware, and a user's email credentials being stolen. In order to assess the relative risk and priority of these two events, you must establish some basis for comparison.

In risk analysis, events are compared along two dimensions: the *impact* of the event on the organization and the *likelihood* that the event will occur. It's typical to assign a number to the impact and likelihood for each risk event. Let's say we are using a one to three range, where 1 equals "low," 2 equals "moderate," and 3 equals "high."

The risk rating is the product of the values on each dimension. An event with an impact of 3 and a likelihood of 3 would have an overall risk value of 9. An event with an impact of 2 and a likelihood of 2 would have a risk value of 4.

Once every event has been assigned a value, they can be put into a matrix to sort them out. The matrix would look something like what you see below.

		Impact of Incident		
		Low (1)	Moderate (2)	High (3)
Likelihood of Incident	High (3)			
	Moderate (2)			
	Low (1)			

Table 1: Risk Assessment Matrix

It is then possible to decide how to treat different risks based on which cell in the chart they occupy. For example, you might decide that all risks in the high-high cell should be addressed as soon as possible. Risks in cells with both values being moderate or above might be set to be addressed after the highest risks, and all cells with at least one "low" might be tolerated and not recommended for efforts to reduce the risk. Below is an example.

		Impact of Incident		
		Low (1)	Moderate (2)	High (3)
Likelihood of Incident	High (3)	🟩	🟨	🟥
	Moderate (2)	🟩	🟨	🟨
	Low (1)	🟩	🟩	🟩

Table 2: Risk Treatment Categories

Getting to this matrix involves collecting the information about each event that is relevant to its impact and likelihood. We will show how to do this in our discussion of the risk assessment checklist in the next section.

The risk assessment matrix is a useful way to summarize your risk situation.

THE RISK ASSESSMENT CHECKLIST

Here is what you must do to carry out an effective risk assessment:

1. Identify the assessment's scope.
2. Identify assets: anything in your organization whose confidentiality, integrity or availability must be protected.
3. Identify the assets' owners.
4. Determine what could threaten each of those assets, what vulnerabilities each threat could exploit, and how your current security system is protecting them.
5. Assign a value to each asset, based on what it would cost your organization if the asset's confidentiality, integrity or availability were lost.
6. Build and populate the risk assessment matrix to determine acceptable and unacceptable risks.
7. Propose ways to mitigate the unacceptable risks.
8. Evaluate the ways to mitigate the unacceptable risks.
9. Prioritize what you must protect.

IDENTIFY YOUR ASSETS

We will consider the organizational scope of the assessment to be your entire organization.

An asset is anything that affects the goals or operations of the organization, or its obligations to others. Your assets to be protected should be within a defined scope. Since we are talking about an information technology assessment, we will consider only assets related to information. **The Center for Internet Security defines an information asset as, "Information or the systems, processes, people, and facilities that facilitate information handling."**

Remember that the definition of information technology security includes the confidentiality, integrity and availability of its assets.

- Confidentiality means limited access to content and data.
- Integrity means that content and data are not subject to unauthorized changes.
- Availability means that content or a service is available for the organization's use when needed.

These definitions mean that things like hot weather, power outages, disaster recovery and business continuity are security concerns. Assets do not have to be tangible, either. They can be procedures such as monitoring access to the organization's building or server room.

Assets can be treated individually, like a firewall. However, they can also be treated as an asset class, such as all your organization's firewalls, if members of the class do not differ significantly. This is important because it allows less granular definitions, such as "confidential documents," which allow you to identify a type of document rather than having to examine each one.

ISO 27005:2018, the international standard for information security risk management, divides assets into two classes:

1. **The primary class** contain the organization's business processes and activities, as well as its information.
2. **The second class** includes the supporting assets on which the primary assets rely, such as hardware, software, networking, personnel, the site, and the organization's structure.

Sources of information about which assets to include can start with inventories of the organization's hardware and software. The next best source of information is the asset's owner.

An owner is the person who is responsible for each asset. The IT Manager might be responsible for the switches, and the Controller or CFO might be responsible for financial data. The idea is that the owner knows the processes and data of the asset, in addition of the importance of the asset to the organization. Depending on the organization, owners may have to be interviewed, unless the organization is small enough that all the owners can gather at a working meeting.

It's good to list all the assets and owners on a spreadsheet. Another good column is the location of an asset. For, example, it is important to know where documents or data are stored. These columns will soon be joined by others. In general, looking deeper into these categories may bring up other assets for consideration.

> **Visualize your assets! Use Visio, LucidChart or another drawing application and create a storage container for each place that's storing information. You can tag the drawing with information and this can help you sort your containers according to risk.**

MODEL POTENTIAL HACKING THREATS

A threat is an event that could compromise the security of an information asset, such as a hacker sending a phishing email to staff. A threat model is a description of how a threat could compromise an information asset, given the current protections and vulnerabilities around the asset. In the case of the phishing email, a current protection could be anti-phishing training for staff. A vulnerability could be the limited display on a smartphone, which makes it harder to identify phishing emails.

The final threat model for each asset allows you to assess the likelihood of the event. Negative impacts are more likely if the asset has vulnerabilities—that is, aspects that make it more susceptible to attack. Negative impacts are less likely if good safeguards are currently protecting the asset.

The identity of the person or organization that is expected to carry out the threat can also make a difference. For example, when a foundation has names of members of groups working to change a country whose government opposes these groups, they are more likely to attract nation-states as threats. Each risk, then, will apply a threat model to a particular asset which allows you to assess the risk's likelihood. To assess the impact of each risk, we need to know more about each asset.

So, how can you check if you've identified enough threats? How can you understand which bad actors are associated with which kinds of attacks?

The standards organizations are a great resource; ISO 27005:2018 has an Annex C that presents examples of typical threats. Also, each year, various organizations present the results of what events have occurred in systems they measure. A good example of this type of document, which has all kinds of information about threats, is Verizon's "Data Breach Investigations Report."

On your spreadsheet, you should assign columns to threats, vulnerabilities and safeguards for each asset or asset class, as well as the asset's location, which may or may not represent a vulnerability.

EVALUATE POTENTIAL HACKING IMPACTS

For each asset, you must ask, "How severe would the impact be to the organization if this asset were compromised?"

Clearly, different assets can be compromised in different ways. A confidential document is compromised if control over access to it is lost. A server is compromised if an unauthorized user gains access to it, or if it no longer functions. Thinking about this, you'll see that a single asset can experience several different negative events. Each one of these event-asset combinations ultimately would become a different entry on the risk assessment matrix, since the events could have different levels of severity.

Another aspect of impact is how the compromise of an asset affects others outside the organization. For example, loss of a confidential document from a partner affects that partner. It also affects the organization if a non-disclosure agreement is in effect and it affects the relationship simply because of damaged trust.

Some risk assessment methodologies rate the effects on your organization and the effect on others separately, then combine the separate ratings into one measure of impact. This is an interesting approach but one we will not use here for the sake of simplicity. Consult some of our references, below, for more discussion about this.

How severe would the impact be to the organization if this asset were compromised? One way of estimating risk, understanding that a "container" is the method/location by which your asset is stored, is to consider:

- Is the container Internet-facing? (higher risk; less if service provider has security controls in place)
- Is the container a "consumer" version of tools like Dropbox? (higher risk; less if business version of tool is used)
- Is someone administering the container? (lower risk if "yes")

CONSTRUCT THE RISK ASSESSMENT MATRIX

By now, you're wondering about how much judgment is involved in these processes of modeling threats and evaluating impacts. Some people may consider damage to the organization's reputation from an information breach the most serious impact. Some organizations would suffer more from their website being down.

It is important to understand that we are discussing qualitative risk assessments here. You are not trying to exactly quantify the impact, for example by assigning each an exact price. You are not assigning exact probabilities to each likelihood. The effort to assign precise impacts and risk probabilities falsely implies a greater precision to the answers than is warranted.

The best we can do is to try to find quantitative measures of impact and the likelihoods that are reasonable, if not exact. Standards organizations that examine and define the categories of impact and the likelihood and then numbers are assigned to each. The Center for Internet Security has presented the following as plain-language definitions to make the impact and likelihood scores more meaningful.

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: Impact Score Definitions

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

Table 4: Likelihood Score Definitions

In the references below, we present the risk assessment discussions for three of the major standards bodies: ISO/IEC, NIST and CIS. You can adopt definitions from these documents, or you can make up your own. The object is to have definitions so that a reasonable person, examining the risk, would rate the risk the same way others would, given their different opinions.

One way or another, you come up with the matrix we showed above, but you supply your own definitions.

POPULATE THE MATRIX

You are now ready to populate the matrix. Rate each event/asset combination based on its impact and likelihood of happening, using whatever definitions you've decided upon for each rating number. If you have been keeping that spreadsheet we've mentioned, add two more columns; one for impact and one for likelihood, and put the appropriate rating numbers in each row.

If you want an example of such a spreadsheet, by the way, CIS has put a nice one in its *CIS RAM Workbook*, listed in the references. The workbook, in fact, is an Excel workbook, full of spreadsheets.

DETERMINE ACCEPTABLE AND UNACCEPTABLE RISKS

Now that all your asset/threat combinations are in the matrix, you can assign each a risk value based on the formula:

$$\mathit{Likelihood\ of\ Event} * \mathit{Impact\ of\ Event}$$

If you're using a high/medium/low categorization and assigning a 1, 2 or 3 to each, then your matrix will have values as shown in the matrix below.

		Impact of Incident		
		Low (1)	Moderate (2)	High (3)
Likelihood of Incident	High (3)	3	6	9
	Moderate (2)	2	4	6
	Low (1)	1	2	3

Table 5: Risk Assessment Matrix with Calculated Values

The next step is to decide which categories of risks are acceptable and which are not. This is extremely practical, because it puts constraints on what you must do.

Most people can differ about where to draw the acceptable/unacceptable line. Some, for example might argue that every risk with a low likelihood can be tolerated. Others may see some high-impact risks are so important that they need to be addressed, regardless of their likelihood. The decision is up to you and your organization, which could include your Board of Directors.

At this point, your risk assessment could be complete. The question now is what you *do* with the risk assessment and *how* does it become a guide to action? One way is to read the assessment and decide on how to address the most important risks based on your experience. If you want a more systematic approach, you can continue reading for more guidance.

EVALUATE WAYS TO MITIGATE THE UNACCEPTABLE RISKS

Now is the time for us to discuss controls. **A control, in security standards circles, is a procedure for mitigating (reducing) a risk, either by reducing its impact or its likelihood of occurrence.**

Three of the best lists of controls are:

1. ISO/IEC 27002.
2. NIST Special Publication 800-53 r4 (or r5 for the most recent draft version.)
3. CIS Controls version 7.1.

A control recommended by one of these standards bodies is also expected to be a best practice. Ideally, adopting the control would be the best thing you could do to reduce the risk of a threat to a particular asset. This is not always the case, however, as the control may be too expensive. Or, a new technology may have appeared since the standard was published that better addresses the risk. Your organization's specific situation can affect whether a control is applicable to your assets or not.

Nevertheless, going through the exercise of applying the one or more controls that could protect one or more of your assets is of benefit. In some cases, the control will turn out to be the best practice. In others, it may cause you to think of the best mitigation of the risk in your organization's situation.

If you don't use controls, you should assign a solution to each unacceptable risk. It's possible a solution against some risks could be getting cyber insurance. In this case, you will enter a discussion with your insurance vendor about the right scope of the insurance. However, cyber insurance providers want to know what you're doing to reduce your risk of attack. Don't think that you can just "buy" your way out of this problem! A general principle of risk assessment is that **the solution should not be more expensive than the cost of the risk, if it occurs.**

Solutions can bring their own risks. For example, one control on access to the network is having strong passwords for logging on. Over time, however, it has been found that users do not like strong passwords, because they are hard to remember. If they are asked to create a password too often, they may simply change one character in the new version. This situation has led some standards bodies to re-evaluate their criteria for strong passwords. In any case, it is a good example of what is called safeguard risk, which is the risk applying a safeguard can add to the overall risk.

Ideally, you can estimate the likelihood and impact of safeguard risks with the same procedure you used for assigning your original risks. For each asset, you examine the impact on the organization and the likelihood of its occurrence. Then you multiply the scores and see whether they fall on the same acceptable/unacceptable division you applied to the original risks.

SUCCEEDING WITH YOUR RISK ASSESSMENT

You have now gotten through this checklist, which can also be an assessment. What remains, however, is some advice on how you carry out your checklist.

This is a good place to get some buy-in by having some members of the staff participate in the assessment. In addition to giving these folks some skin in the game, it provides you with some measure of how you're doing. Do the various committee members seem to understand and similarly use the ratings, for example?

This is also a good place to talk to staff. It is good to talk with the owners of an asset about that asset's risks, but it is also important to talk to people who are not owners of any particular asset or class of assets. You might talk, for example, with the President or CEO, and their opinions on certain subjects should not be ignored. It is likely that the CEO may have something to say about the effect of a breach on the foundation's reputation since they are the ones held responsible for the organization's security.

One of the benefits of holding these conversations is that makes staff think about cyber security. So, try to structure your staff conversations in an open-ended fashion, so that people can be expansive about their thoughts and ideas.

YOU CAN DO THIS!

Somehow, your organization is going to have to decide where to spend money on information security. We have become so dependent on technology, and hacking has become much more prevalent, that an attack is inevitable. As the old auto mechanic's advertisement said, "Pay me now, or pay me later." If you think defending against cyberattacks is expensive, try recovering from one.

The proof of this can be seen elsewhere in the *Cybersecurity Essentials for Philanthropy* series, where foundation IT managers plotted their major activities on a timeline before and after a serious information breach. Before the incident, the IT manager's activities varied. But after, most of their activities addressed security issues.

Having a plan such as this checklist can take a lot of the uncertainty out of deciding what to protect and when to do it. Something that seemed imponderable can become manageable with the right mindset.

RESOURCES

Below are links to the tools and resources referenced in this document:

Calder, Alan and Steve Watkins*, IT Governance: An international guide to data security and ISO27001/ISO27002 (5th ed.)*. Philadelphia: Kogan Page, 2012.

Center for Internet Security, *CIS RAM Version 1.0: Center for Internet Security Risk Assessment Method*. Center for Internet Security, April 2018.

Center for Internet Security, *CIS RAM Workbook for CIS RAM Version 1.0*. Center for Internet Security, April 2018.

ISO/IEC*, International Standard ISO/IEC 27005, (Third edition): Information technology – Security techniques – Information security risk management.* Geneva: ISO/IEC, July 2018.

National Institute of Standards and Technology, *NIST Special Publication 800-30 Revision 1: Guide for Conducting Risk Assessments.* Washington, D.C.: U.S. Department of Commerce, National Institute of Standards and Technology, September 2012.

Verizon, "2019 Data Breach Investigations Report," (accessed at https://enterprise.verizon.com/resources/reports/dbir/ , 7/5/2019)

Idealware: https://www.idealware.org/ There are many papers and resources here!

ABOUT THE AUTHORS

DAN CALLAHAN

Vice President, Global Services
CGNET

in https://www.linkedin.com/in/danielcallahan/

Dan is responsible for development of CGNET's cloud and cyber security services. He oversees all aspects of CGNET's Office 365, Teams/Skype for Business, Azure, Enterprise Mobility + Security and Dynamics CRM Online cloud services. He also oversees all aspects of CGNET's vulnerability testing, GDPR compliance, risk assessment and security consulting services. As a consultant, Dan has conducted many technology planning, security, change management and tool selection projects. Dan served as Director of Marketing and Business Operations at CGNET from 1999 to 2003. Prior to rejoining CGNET in 2011, Dan held Director- and VP-level positions in Product Management and Marketing at iPass (acquired by Parateum), SOMA Networks, Daintree Networks (acquired by GE) and YouSendIt (acquired by OpenText).

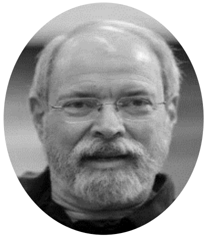

TIM HAIGHT

Vice President, Technology Services
CGNET

in https://www.linkedin.com/in/tim-haight-a1ba2/

Tim has been studying how nonprofit organizations can optimize their use of information technology for more than 30 years, since he was the first evaluator for Apple Computer's Community Affairs Program. He has been at CGNET since 2002 and conducted a great number of organizational analyses.

Tim has done assessments and strategic plans for the Marin Community Foundation, Metta Fund, the Houston Endowment, ClimateWorks Foundation, the California Wellness Foundation, the Carnegie Corporation of New York, Bush Foundation, the Robert Wood Johnson Foundation and many other organizations. Prior to CGNET, Tim was Editor-in-Chief of FTPOnline, Vice President of Communications at OneChannel, Executive Editor of Network Computing, and West Coast Bureau Chief of *CommunicationsWeek*.

ABOUT THIS SERIES

The *CyberSecurity Essentials for Philanthropy* series launched in 2019 is provided by the Technology Affinity Group (TAG) in partnership with member organizations and private sector advisors.

View the full curriculum available for the series at: tagtech.org/cybersecurity

This is an educational publication and is not intended as legal advice. You should contact your attorney for legal advice. The opinions expressed here are the opinions of the individual authors and may not represent the opinions of their employers or of TAG.

TAG CYBERSECURITY WORKING GROUP

This work is led on a volunteer basis by the TAG Cybersecurity Working Group whose members include the following:

Jim Rutt (Chair), Dana Foundation
John Mohr, The MacArthur Foundation
Oleg Bell, Open Society Foundations
Karen Graham, Idealware
Darlene Ott, The Winnipeg Foundation
Dan Callahan, CGNET
Calvin Lewis, Cleveland Foundation
Christopher Jean-Pierre, Wellspring Philanthropic Fund
Steve Jarboe, Accenture
Anthony Putignano, Wizehive
Charles Boname, Vancouver Foundation

FUNDING PROVIDED BY

This series is funded in part through an award from the Robert Wood Johnson Foundation President's Grant Fund at the Princeton Area Community Foundation.