ivanti

Evaluating Cyber Risk Objectively:

A Guide to Data-Driven Risk Assessments

Table of Contents

About this e-book

Exposure management puts security risk in the context of the organization’s overall risk appetite, allowing security leaders to make informed recommendations that will actually win their fellow executives’ commitment.

This e-book provides the tools for a critical component of exposure management: **data-driven risk assessment**.

We’ll present a straightforward method for conducting risk assessments and exhibiting their results, showing the direct impact on business-critical assets of choosing to reduce risks versus choosing to accept them.

We’ll offer guides to apply the risk assessment process to various scenarios, provide numerous examples and tips for presenting your assessment and mitigation proposals and help you find a common language with other business leaders about risk and exposure.

Introduction

You know the cyberattacks your organization is exposed to, the “cracks” that cause issues time and time again and the initiatives that can strengthen your company’s security posture.

Your colleagues across the business know their objectives and the calculated risks they want to take to achieve them.

You both ultimately want the same thing: a healthy, profitable business. But if you don’t share a common understanding of risk, you’re talking past one another.

Getting on the same page requires a detailed, objective and quantifiable cyber risk assessment.

A risk assessment and its outcomes will present conclusions in language that other business leaders understand, which will help you rally support for critical security initiatives.

And it may also challenge your assumptions about what risks to mitigate and what risks to accept.

So how do you put numbers on a security event that hasn’t happened? That’s what we’re going to dig into.

Risk assessment

Key terms and general risk modeling

A general risk model goes something like this:

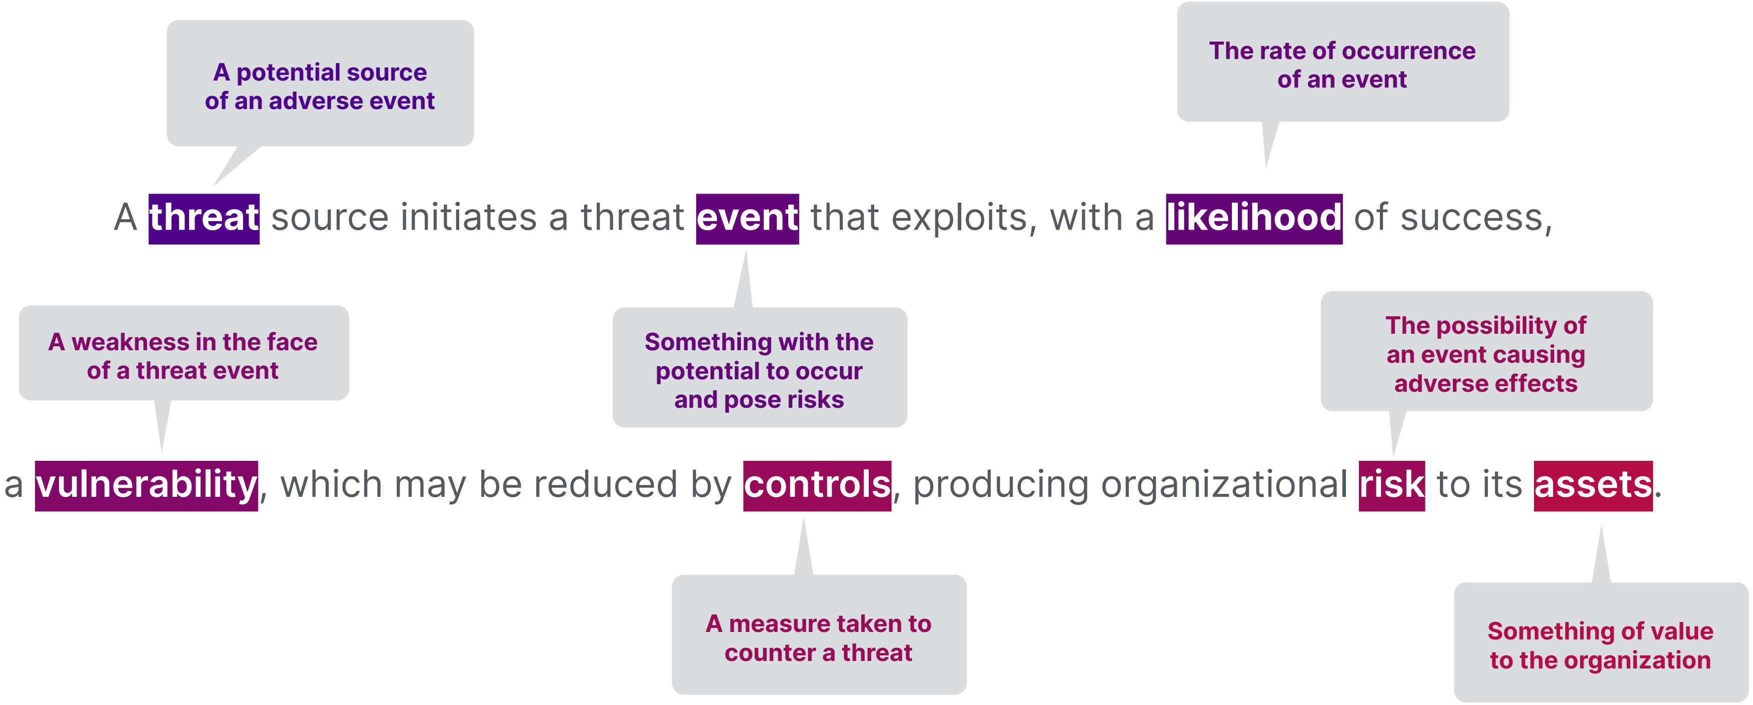

A potential source of an adverse event

The rate of occurrence of an event

A **threat** source initiates a threat **event** that exploits, with a **likelihood** of success,

A weakness in the face of a threat event

Something with the potential to occur and pose risks

The possibility of an event causing adverse effects

a **vulnerability**, which may be reduced by **controls**, producing organizational **risk** to its **assets**.

A measure taken to counter a threat

Something of value to the organization

Real-life events follow this model, for example:

The **ALPHV group** was behind several **ransomware attacks**, which were **successful** on several companies across the US, Europe and APAC, exploiting **weak access controls and user training**, causing **operational downtime** for days, in some cases leading to huge backlogs due to lack of frequent **backups**, producing **financial, reputational, and legal risks** for the company and its **information**, which may have been breached and released to the public.

Risk assessment steps

There are five straightforward steps to your risk assessment, which we'll cover in this section.

1. Identify the assets in scope for your risk assessment.
2. Assign value to these assets.
3. Identify vulnerabilities and threats.
4. Calculate risks (likelihood multiplied by impact).
5. Perform a cost-benefit analysis.

Then, you can present recommendations for the risk response, backed up by data from your risk assessment. There are four possible risk responses, which we will address later in more detail.

Avoid the risk by stopping the risky activity or shutting down the vulnerable system altogether.

Accept the risk and decide to live with it, taking no action.

Transfer the risk, usually by contracting an insurance company to cover the financial exposure.

Mitigate the risk by applying additional safeguards.

Step 1 Identifying assets

What are we looking for?

Assets can be categorized as **tangible** (building, people, equipment) or **intangible** (information, patents, brand, licensing, customer list, R&D).

Assets can also be categorized by ownership. On a high level, these are **business assets** (cash, land, inventory, factory, employee expertise) or **IT assets** (software, servers, firewalls, laptops, monitoring tools).

But IT assets do not merely exist for themselves: they ultimately exist to protect business assets. Your company may have a trade secret (an intangible business asset) residing on a server (a tangible IT asset). Security controls applied on this single server may make or break the company.

Let's say you're concerned with data loss prevention. You need to assess what is in your scope (work laptops, BYOD, phones, printers) and determine their number, which you ideally have in an up-to-date inventory. Then you can explain why securing x number of endpoints is actually equal to securing business assets involving confidential information, such as customer data, credit card information or employee salaries.

You want to reduce the risk of this information leaving the company — but if you turn that risk into numbers, you'll be better prepared to come up with (and win support for) an appropriate mitigation strategy.

Crown jewels

Crown jewels are the most valuable company assets, and out of all assets, their compromise would cause the greatest business impact.

You probably think of **information** as a crown jewel, and you're right. But other leaders may define the crown jewels a bit differently. They may point to certain business processes and systems, without always making the connection that without information, their value would be null. Information lives many times longer than systems and is, in fact, what needs protection.

You are ultimately referring to the same assets, because business processes and information systems go together. But to be impactful, your risk analysis needs to be very specific.

Let's assume one of the company's crown jewels is its CRM system and the customer information it processes. See the difference between these two statements:

*"We will lose **information** due to bad actors capturing it by spoofing our DNS."*

*"As a result of a data breach, which affected 45% of the companies in our industry last year, **customer PII** and transactions processed by Salesforce could be publicly disclosed, which would cause estimated fines of $4.5B, never-ending overhead costs caused by lawsuits and a reputational risk that would take 5-10 years to recover from."*

Clearly, the second one holds more power, and your risk assessment will help you get to that level of specificity.

Step 2
Determining asset value

Whether tangible or intangible, an asset's value needs to be clear. While it's not necessarily up to the security team to assign asset value, it is an important — if not the most important — parameter when making business decisions about implementing safeguards for that respective asset.

Determining the value of an asset should consider:

- Cost to acquire or develop the asset.
- Cost to maintain and protect the asset.
- Value of the asset to owners and users.
- Value of the asset to adversaries.
- Price others are willing to pay for the asset.
- Cost to replace the asset if lost.
- Operational and production activities that would be affected if the asset were unavailable.
- Liability issues if the asset were compromised.
- Usefulness and role of the asset in the organization.

Each company asset should have an owner, and the asset value should be determined by that owner, possibly under the guidance of the security team.

Then, a very important question is: How much could it cost the company to NOT protect the asset? Answering that question begins with a business impact analysis, or BIA.

An ideal business impact analysis

A BIA is usually a step in the business continuity planning lifecycle. It's a key activity where technology and business come together to discuss the criticality of activities and their supporting assets.

Asking *"What exactly would happen if this process stopped working?"* makes process owners and function heads think about and indicate possible impacts, including financial impacts. This will get you near-real numbers that you can use in your risk assessment.

An ideal BIA covers:

- Process purpose, owners, inputs and outputs.
- The impact of an interruption, measured in financial, operational, legal/regulatory and reputational impact.
- The worst possible scenarios and times of an interruption (such as peak periods).
- Recovery time objectives (RTOs), recovery point objectives (RPOs) and maximum tolerable downtime (MTD) for activities.
- Resources needed to conduct/support the process: other units, information, people, infrastructure, vendor locations and IT assets. These will inherit the process RTOs and RPOs and will be prioritized for recovery in a disruption.

Tips

There are numerous templates out there, but it's not one-size-fits-all. We offer a simplified template below (which you can download as editable slides) that you can use as a starting point.

Business continuity and information security may sit in different units. In this case, information should be shared, to ensure a consistent approach to risk management.

Assigning criticality of assets with BIA

The BIA identifies which of the company's critical systems are needed for survival and estimates the outage time that can be tolerated by the company. The outage time that can be endured by a company is referred to as the maximum tolerable downtime (MTD).

Some examples of MTD:

- **Nonessential:** 30 days
- **Normal:** 7 days
- **Important:** 72 hours
- **Urgent:** 24 hours
- **Critical:** minutes to hours

Each business function and asset should be placed in one of those categories, depending upon how long the company can survive without it. These estimates will help the company determine what controls are necessary to ensure the availability of these resources.

For example, if being without a web server for four hours would cost the company $120,000, the web server could be considered **critical** and thus the company should consider having a redundant web server.

But if a reporting tool for building occupancy becomes unavailable for three weeks, and it would not cost the company more than $500, it's safe to assume that the tool is **non-essential**, and you can either count on the vendor's SLA or best effort to fix it.

Tip

Process and system RTOs, correlated with financial loss values, can be used as a parameter in quantitative risk analysis. Chances are you already have this information captured in existing BIAs and in your disaster recovery plan. If these are not available, you can simply conduct the analysis narrowed to your scope.

A simplified BIA may look like this.

Process Name	Process Interruption	Financial Impact	Legal Impact	Reputational Impact
Same Day Payments	<4h	$-$$$	n/a	n/a
	4h-8h	$$$-$$$	n/a	Low
	2-5 workdays	$$$-$$$$	Medium	Medium
	1-2 weeks	$$$-$$$$	Medium	High
	2-4 weeks	$$$$-$$$$$$	High	Critical
Justification	■ Loss in reputation and public confidence: ... ■ Loss of competitive advantages: ... ■ Increase in operational expenses: ... ■ Violations of contract agreements: ... ■ Violations of legal and regulatory requirements: ... ■ Delayed-income costs: ... ■ Loss in revenue: ... ■ Loss in productivity: ... ■ Other: ...			
Resources Needed	Staff: min. four staff first week; min. six staff after one week Payment system: must be available, no manual workarounds Payment tokens: min. two tokens IT systems: payment system (servers + database), internet, connection to the clearing house, email Location: can perform from home the first day, then need an office space Office supplies: printer			

Tip

You can download this BIA template as editable slides.

How BIA relates to risk assessment

BIA is part of a business continuity-focused risk assessment. It zooms in on the impact of unlikely but major scenarios on business activities, and it clearly shows business priorities. These results help you determine your organization's crown jewels, including their supporting information systems.

Risk assessment usually takes the form of the equation: ***risk = threat x impact x probability.*** However, the BIA adds the dimension of time to this equation. This means it is geared toward those threats that might most rapidly disrupt critical business processes. For this reason, BIA is a critical input for determining controls that protect the **availability** component of the **CIA triad: confidentiality, integrity and availability**.

A **ransomware attack** fits perfectly into one of the continuity scenarios because it makes the data unavailable. However, not all security risks affect the **availability** component.

A **data breach** (exfiltrated without leveraging ransomware) may leave the systems and processes functioning as usual and keep the data available — or rather, too available, breaching the **confidentiality** component of the CIA triad. Processes that manipulate the data that has been breached will not suddenly stop in this scenario.

If you have a trade secret, it is most likely one of your crown jewels. It being maliciously modified (the **integrity** component) or disclosed to the public is another example that may not get captured in the BIA.

That is why you'll need to approach your risk assessment from additional angles, but a BIA is your best starting point.

CIA rating assessment

With the BIA we have determined which assets are time sensitive. A **CIA rating assessment** can add a valuable dimension to this picture.

ISACA proposes a methodology of asset valuation and categorization that allows you to weight an asset's value according to its sensitivity.

For each of the three components of the CIA triad, you assign a value of 1 (low), 2 (medium) or 3 (high). The CIA rating assessment is the sum of those three values (C + I + A), for a score somewhere between 3 and 9.

Each company asset would have an assigned CIA rating, which will be directly linked to a risk-based control implementation. A rating of C3-I2-A2 (7) will require different controls than an asset rated C1-I1-A2 (4).

Both methods are done with business owners, who are accountable for the information flowing through their assets, which makes this is an opportunity to solidify a common awareness of asset values and reduce the disconnect between risk awareness and business orientation.

Tips

A CIA rating assessment isn't necessary for each configuration item or VDI. Assets should be grouped in a way that makes sense for the company, for instance, according to a common business goal they have.

If your organization has a published risk appetite statement, use it as a point of reference to align other business leaders on the values you assign in your CIA rating assessment. (Don't have one? Use this editable template as a starting point.)

Step 3

Threat modeling and vulnerability assessment

The next step in risk analysis is to determine high-priority threats and vulnerabilities. Although you may already have a good idea of what these are, validating those assumptions will add further credibility to your risk assessment.

Threat modeling

Threat modeling must be grounded by threat intelligence: looking at the global threat landscape, the types of threats leveled against other organizations in your field and the particularities of your business. This will focus your efforts on threats that are more likely to occur, instead of dissipating efforts on less likely threats.

There are many methodologies for conducting threat modeling.
Two common methodologies are:

Attack trees:
A threat source has multiple avenues to reach an objective. An attack tree creates branches and leaf nodes, showing these different paths and indicating conditions that have to be met for threat actors to achieve their objective. Each node will be assessed for vulnerabilities.

Reduction analysis:
This approach builds upon attack trees, finding commonalities among the leaf nodes. This then allows you to identify potential controls that could mitigate more than one vulnerability each.

Tip

The examples of attacks that you surface as part of your threat modeling exercise can add color to your risk assessment and mitigation recommendations, helping make threats tangible to leaders outside security.

Vulnerability assessment

Whereas threat modeling works outside in, starting with potential threats and identifying the ways in which they could penetrate your organization, vulnerability assessment works inside out, finding weaknesses in your attack surface that could be exploited by threat actors.

A central principle of exposure management is a broad definition of the attack surface. When assessing vulnerabilities in your attack surface, remember that they are broader than information and system vulnerabilities. These can also include processes (e.g., late or improper patching) and people (e.g., susceptibility to social engineering, using weak passwords).

Software such as network vulnerability scanners, application security testing and external attack surface management will surface valuable information on potential vulnerabilities. You may also have sources of information such as audit reports, risk registers, control testing, pen test results, incident reports, policy studies and simple observation of the company's culture.

When it comes to system vulnerabilities, vendor vulnerability ratings and scoring systems such as the Common Vulnerability Scoring System (CVSS) are a first step in identifying high-priority vulnerabilities. But to further refine your list (and potentially surface those lower-scored vulnerabilities that would be more severe for your organization), you need to layer on two additional filters: **threat context** and **risk context**, both of which you've already laid the groundwork for.

Threat context looks at the general threat landscape:
Is this vulnerability being actively exploited?
A vulnerability may be rated Critical, but if it is not being actively exploited, it diminishes in importance. Your threat modeling exercise has already begun this research for you.

Risk context looks at your organization:
If this vulnerability were exploited, what would the impact be on our organization?
A vulnerability may have a low rating, but if it would take a critical system offline if it were exploited, it merits more attention. Thanks to your impact assessment, you already know what assets, systems and processes would cause the greatest organizational impact were they to be compromised or disrupted.

Tip

Download this checklist to ensure you are considering all facets of your attack surface.

Step 4 Determining risk

By now we have the asset values (AV), the threats and the vulnerabilities. We can finally start looking at likelihood and impact to calculate risks. There are two approaches for measuring risk: **quantitative** and **qualitative**.

A quantitative risk analysis uses equations to assign monetary values to components within the analysis.

A qualitative risk analysis assigns values and ratings such as low/medium/high, 1-10 or a stoplight.

You can determine which methodology works best for your organization. For this e-book, we will focus primarily on quantitative risk analysis.

Tip

Automated risk analysis tools can be used to reduce the risk assessment time associated with manual work and to calculate benefits of different security controls.

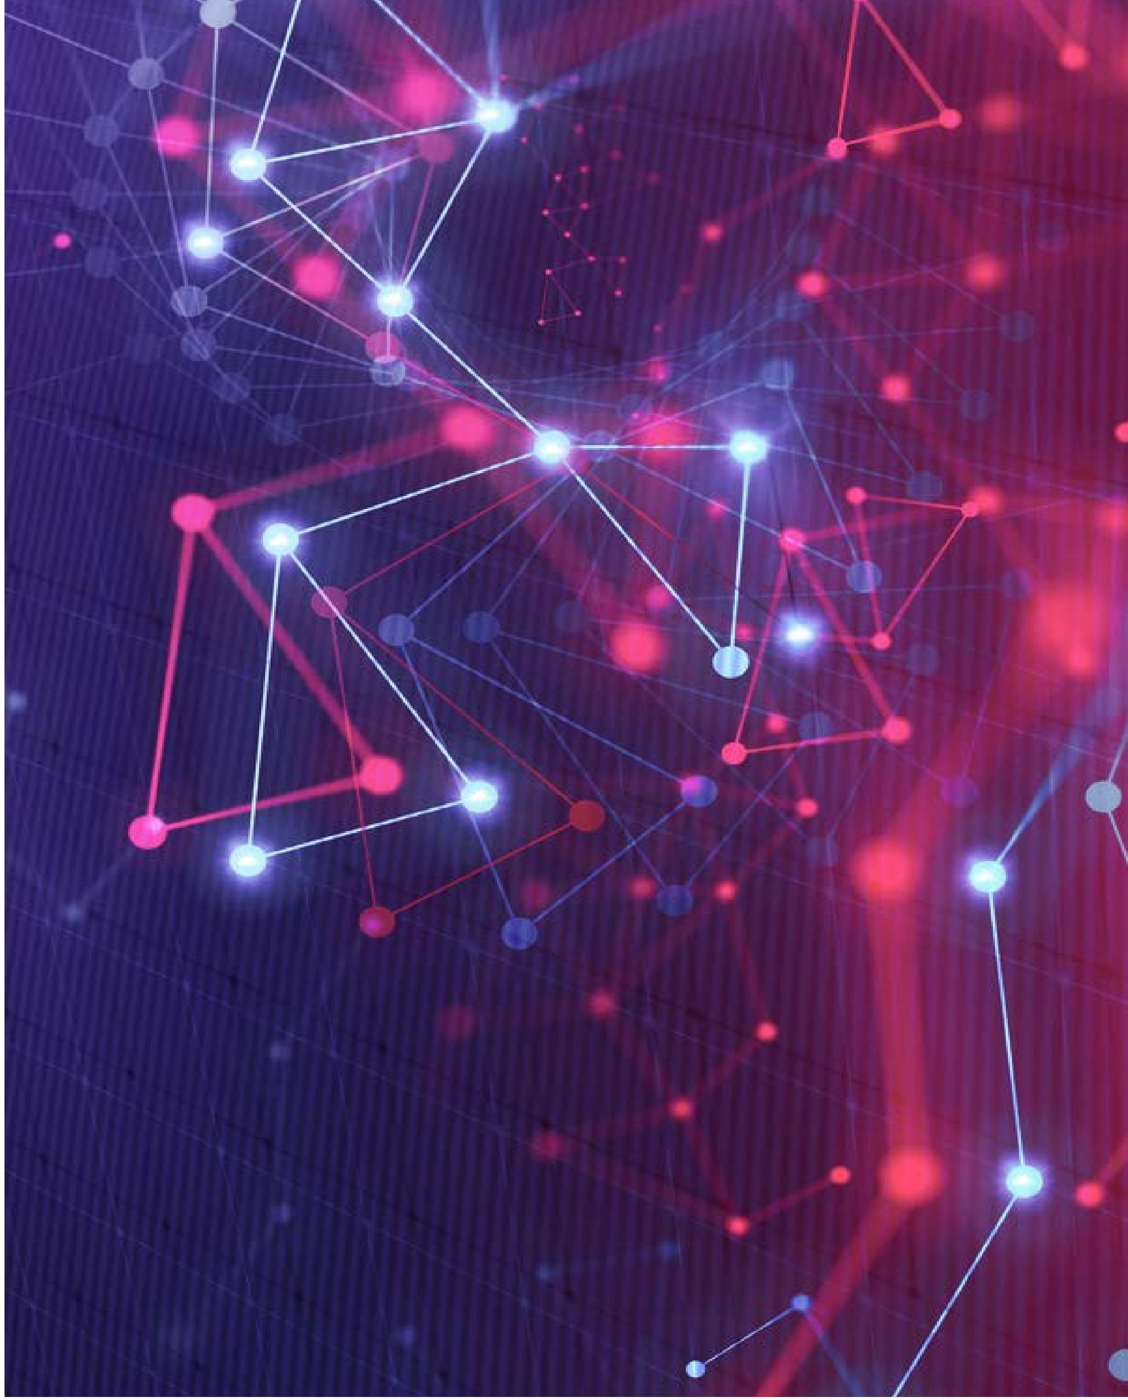

Quantitative risk analysis

Your CEO and board will want your analysis stated in monetary, quantitative terms whenever possible. It is one thing to know that the risk of a cyber attack is high. It's another to know that, in the case of a successful ransomware attack that encrypts 15% of the data on a critical system, the company would be at risk of not only losing competitive advantage but interrupting critical business commitments for up to four days, costing the company $225,000 each day.

The widely accepted equations to quantify risk are actually quite simple.

Annual loss expectancy (ALE) = **single loss expectancy (SLE)** x **annualized rate of occurrence (ARO)**
where **single loss expectancy (SLE)** = **asset value (AV)** x **exposure factor (EF)**

The **exposure factor (EF)** represents the loss percentage that a threat could have on a certain asset. For example, if you have a trade secret that may be leaked, the loss of competitive advantage may be worth 10% of the asset value — in other words, the EF would be 10%. Assuming the **asset value (AV)** of the trade secret is $4M, you can use a simple equation to calculate the **single loss expectancy (SLE):**

AV ($4M) x EF (10%) = SLE ($400K)

The **annualized rate of occurrence (ARO)** is the estimated frequency of the threat taking place within one year. Using the same example, you may look at comparable events at companies like yours and estimate that a trade secret leak may occur every 20 years. In other words, its ARO would be 0.05 (1 year/20 years). Now we can calculate the **annual loss expectancy (ALE)**:

SLE ($400K) x ARO (0.05) = ALE ($20K)

The **ALE** result helps you decide whether the company should implement any controls to mitigate the risk or simply accept the risk. In our example, the company could reasonably justify controls that cost less than $20K per year.

The outcome of a quantitative risk analysis is a key tool in your arsenal to build a business case for a risk mitigation strategy — or to disqualify mitigation strategies whose costs exceed the ALE.

Examples of Quantitative Risk Analysis Results

Asset	Threat	SLE	ARO	ALE
Trade secret	Leaked	$4M	0.05	$20K
File server	Corrupted	$13.5K	0.1	$1.35K
Plant	Fire	$250K	0.1	$25K
Data	Malware	$7.5K	1.0	$7.5K
Customer credit card no.	Disclosed	$300K	4.0	$1.2M

Let’s look at another example.

You’ve recently been hired to lead IT and security at a rapidly growing AI start-up. You are well aware of the imminent risks of a successful phishing attack. In a worst-case scenario, a threat actor could exfiltrate and leak the proprietary algorithms you are developing, hurting your competitive advantage and therefore your valuation. You started running phishing campaigns, and results consistently show that 10% of employees click on links or open attachments.

You’ve got to explain to your CEO and investors — who are laser-focused on showing engineering progress — that it’s worthwhile to divert funds and employee attention to measures like an anti-phishing tool and an information security training campaign in order to mitigate that risk. It’s time for a quantitative RA.

Determine SLE:

According to the 2024 IBM Cost of a Data Breach Report, phishing attacks were the second-costliest initial threat vector (tied with business email compromise, and second only to malicious insiders), reaching an average of **$4.88 million** in 2024. This is an average **SLE**, which you’ll want to adjust according to your own **EF** and **AV** variables.

For the purpose of this example, let’s recalibrate the SLE based on the fact that you work for a much smaller business. Your **EF**, or the percentage of your data that would be compromised, would likely remain constant, but the **AV** is likely much smaller — say 10% the size. This makes your SLE **$488K**.

Determine ARO:

In the course of your research, you’ve learned that similarly sized companies in your industry are hit with phishing attacks on a weekly basis, or 52 annually. If these are successful 10% of the time — consistent with the phishing campaigns you’ve been running — the **ARO** is **5.2**.

Determine ALE:

SLE ($488K) x ARO (5.2) = ALE ($2.5M)

(Numbers are rounded.)

Upfront, we can tell that accepting the risk of this ALE does not seem to be a wise option, and you can make a strong business case for an aggressive risk mitigation strategy.

Qualitative risk analysis

A qualitative RA describes the risk within a scale of ranges, such as from very low to very high and/or from 1 to 10. You may often see a qualitative risk matrix in this format, or similar:

		Likelihood				
		1 Rare	2 Unlikely	3 Possible	4 Likely	5 Almost Certain
Impact	5 Very High	5	10	15	20	25
	4 High	4	8	12	16	20
	3 Medium	3	6	9	12	15
	2 Low	2	2	6	8	10
	1 Very Low	1	2	3	4	5

Risk Legend	Low	Medium	High	Extreme

Let's return to the phishing example above. In your qualitative RA, you deem the scenario **likely**: as a smaller company, you aren't targeted with the frequency of larger, more well-known companies, but you know phishing is among the most common threat vectors. You deem the impact **medium**: while the worst-case scenario could be extremely damaging to your reputation and therefore your valuation, a more realistic scenario would not involve such high exposure. Your result is a **high** risk.

You can bring more precision and repeatability to this exercise by determining the high and low end of your scope for **likelihood** based on historical events — what are the most and least extreme scenarios that are still within the realm of plausibility? And you can give a financial value to each category under **magnitude of impact**.

Both quantitative and qualitative risk assessments have their advantages and disadvantages. A mix of both is possible, and it offers an opportunity to consider risk that can't be as easily quantified as financial or operational risk, but that are still very serious, such as regulatory or reputational risks.

A table such as this one, which assigns estimated indirect costs to non-financial risks, can be a starting point:

Impact	Monetary	Regulatory	Reputational
Very High	> $10M	Criminal charges/cases; potential withdrawal of license; personal liability of management	Trust completely lost; continuous negative press coverage on local/global scale; brand image smeared
High	$1M - $10M	Major fines, lawsuits and settlements with public prosecutor; license to operate could be restricted	Trust severely damaged and never fully recoverable; negative press coverage increasing in frequency and spreading to an international level
Medium	$100k - $1M	Public warnings and fines; immediate actions required for remediation	Trust diminished and recoverable at considerable cost; increased negative press coverage
Low	$10k - $100k	Non-public warnings from regulators/supervisors	Trust broken but recoverable with time; press coverage national with some neutral coverage
Very Low	< $10k	No or limited attention from regulators or supervisors	Trust questioned but recoverable in short time; press coverage only local one-off basis

Whichever approach you decide to use, the risk assessment has one more challenge up its sleeve: uncertainty.

Both the quantitative RA and the qualitative RA methods are imperfect. Both — even the numerical one — have a level of subjectivity tied closely to uncertainty.

Uncertainty refers to the degree to which you lack confidence in an estimate. Capturing the degree of uncertainty when carrying out a risk assessment is important because it indicates the level of confidence you have in the resulting figures.

A purely objective analysis is difficult (or even impossible) to attain. Uncertainty can arise when there aren't enough historical events or research to fully confirm your data, when you don't have time to collect more data or when you simply can't predict a black swan event.

Listing the specific areas of uncertainty can help you define your confidence level — and also illuminate areas where you can potentially reduce that uncertainty.

Let's take implementing a bring-your-own-device (BYOD) policy as an example of uncertainty.

If you work in an industry with extremely low security risk tolerance — the defense industry, for example — you are likely disinclined to allow BYOD. But let's say employees are increasingly requesting BYOD, and leadership is opening up to the idea in order to boost employee job satisfaction. You need to figure out if the incremental risk is worthwhile and what controls you'd need to put in place to reduce that risk.

So, what might you not know?

- If you're a government agency or contractor, will future regulatory changes impact your company's policy?
- How many employees will take advantage of this policy? What types of roles?
- How many hours would it take IT staff to register and provide support for devices?

You may not be able to answer these questions definitively, but there are actions you can take to reduce your uncertainty — for instance, surveying employees to better predict the uptake should you loosen your BYOD policy.

Step 5

Cost/benefit analysis

Now that you've determined how much risk you're facing, you can start to weigh your mitigation options. While we're focused here on security controls, the same principles apply to any sort of mitigation, including patching software vulnerabilities and correcting misconfigurations.

Control costs

To make an informed recommendation about your risk response, you have to answer a critical question: what is the real cost — including hidden costs — of mitigating this risk?

A security control may be technical (such as DLP implementation), administrative (such as an infosec awareness program) or physical. You place controls between assets and threats to reduce their vulnerability, reducing the potential impact of a threat occurrence.

It's important to remember that every mitigation has a cost, even if that control doesn't require you to spend money. For example, taking a server offline while you update role-based access controls will cost the business money, just like implementing a new security tool will. The only difference between the two is that one will be a line item on your budget, and the other will be a hidden cost — unless you remember to account for it.

As a general rule, the cost of securing an asset should be lower than that asset value. Otherwise, mitigation does not make sense.

To calculate a control's cost (and also explicitly consider hidden costs), you should consider:

- Product costs.
- Design and planning costs.
- Implementation costs, including cost of downtime.
- Environment modifications.
- Compatibility with other countermeasures.
- Maintenance requirements.
- Testing requirements.
- Repair, replacement or update costs.
- Operating and support costs (including staff training).
- Effects on productivity.
- Subscription costs.
- Extra person hours for monitoring and responding to alerts.

Control value

The conceptual formula for control value:

(ALE before implementing control) - (ALE after implementing control) - (annual cost of control) = value of control

Let's take the earlier example of combating data exfiltration via a phishing attack by implementing an **anti-phishing tool** coupled with an **infosec training campaign.**

To determine the value of the controls, we have to first calculate costs. Let's say after you've considered all direct and indirect costs, you come up with $90K as your total cost.

Anti-phishing tool	Infosec training campaign
Deployment (one-off)	Deployment (one-off)
Integration	Integration
Licensing/subscription fees	Licensing/subscription fees
User training	Materials for campaigns
Maintenance and support	Maintenance and support
Productivity loss	Productivity loss (employee time to participate in the training)
Est.: $70K/year	**Est.: $20K/year**

The initial ALE was:

SLE ($488K) x ARO (5.2) = ALE ($2.5M)

(Numbers are rounded.)

To determine the new ALE, we must account for a reduced ARO once the controls are in place.

Your initial ARO assumed weekly phishing attempts, with a 10% success rate based on the regular campaigns you had run. The frequency of attempts won't change, but let's say after doing your research into both vendors, you estimate that the combined solution will reduce the success rate to 2%. This reduces your ARO to one-fifth the original calculation, or 1.04.

Now you have your new ALE:

SLE ($488K) x new ARO (1.04) = ALE ($507.5K)

(Numbers are rounded.)

You can then calculate the control value:

Initial ALE ($2.5M) - new ALE ($507.5K) -
annual cost of control ($90K) = control value ($1.9M)

(Numbers are rounded.)

Total risk vs. residual risk

The reason we are implementing security controls is to reduce the company's overall (total) risk to an acceptable level. But since no system or environment is 100% secure, there is always some risk left to deal with: the residual risk.

The ALE after implementing controls is also called the residual ALE, and it can express the residual risk.

Quantitative residual ALE = residual ARO x residual SLE

In our example, SLE has stayed the same after implementing controls, but ARO has been reduced in likelihood from 5.2 to 1.04, leading to a residual ALE much lower than the initial (inherent) ALE. This makes an excellent case for a positive risk response, i.e. mitigation.

Tip

You may also see
or use these formulas:

threats x
vulnerability x
asset value = **total risk**

threats x
vulnerability x
asset value x
control gap = **residual risk**

Total risk -
controls = **residual risk**

Control benefits

Many times, a control's benefit goes beyond mitigating a threat, i.e. the control value, as discussed earlier. Controls can also save money on **operational costs**, and calculating these savings may influence your analysis.

Example: **Automated patching**

Let's say you identified the opportunity to implement an automated patching process. This will ensure patching is done much faster, which will reduce the window of open vulnerabilities and the likelihood of human error.

But that's not all. Let's assume that your company manages 75 applications that need two bi-weekly patches each. Each requires an average of four hours to pack and deploy, leading to 600 hours per year spent on this activity. Assuming a fully burdened tech staff hourly rate of $100, we can calculate the following operational cost savings:

Number of applications	75
Number of bi-weekly updates per application	2
Number of hours to package an application and deploy	4
Number of hours spent on manual packaging	600
Fully burdened tech staff hourly rate	$100
Annual cost of manual packaging	**$60K**

Especially with automation, time savings run high, allowing your team to focus on more meaningful work and ultimately saving the company money.

Here are some other examples of benefits that can be quantified:

Benefit	Quantification
Threat detection: early detection of anomalies	Reduction in detection times of potential threats
Incident response: swift response to alerts and incidents	Reduction in incident response times
Real-time monitoring: SIEM doesn’t wait for an administrator to check for recent activity	Increase in employee’s time efficiency
Integrated view: all in-scope devices send data into the same dashboard	Reduction in time to toggle between applications; saving time to correlate events
Accuracy: less false positives	Time saved to assess whether an event is true positive or not
Compliance: enhanced reporting	Time saved to produce on-demand reports during an audit
Forensics: quicker incident analysis; e-discovery requests can be honored more easily	Time saved to reconstruct activity logs and trace the sources of threats
Behavior analytics: learns inside activity patterns	Time saved to determine whether normal user activity is conducted

Risk response

Now that you have collected the relevant information and have conducted the risk assessment, you can make a data-backed recommendation for your response.

Response options

As mentioned, there are four types of responses to risk:

Avoid
risk by closing the risk-bearing business/process/system/location.

Accept
the risk; note that this will require you to periodically revisit this decision.

Transfer
risk to insurance; note that this option is only a financial risk transfer, and the company is still accountable for a materialized risk and will have to deal with other types of losses, such as reputational loss.

Mitigate
the risk, acknowledging that the cost of risk materialized is higher than the cost of controls.

Let's return one more time to our example of combating data exfiltration via a phishing attack, where we've calculated the total risk, the cost of implementing controls and the residual risk after applying controls.

Avoiding the risk would require stopping any external email exchange — clearly not a feasible response.

Accepting the risk makes sense in cases where the asset value is not high enough to justify the cost of protecting it. In our example, we calculated the cost of risk materialized, i.e. ALE, as $2.5M, which could be reduced by approximately 80% by spending $90K/year on controls — a strong argument against accepting the risk.

Transferring the risk to an insurance company might have a low upfront cost, but it isn't a great option in this case. Cyber liability coverage would likely cover the financial cost, but it doesn't account for other types of impact. In our example, we are worried about the impact on our next funding round — a reputational risk that isn't addressed by transferring the risk to a third party.

Finally, we can **mitigate** the risk by applying security controls — in this case a combination of an anti-phishing tool and an infosec training campaign — a conclusion our risk assessment clearly supports.

Incorporating risk appetite

Not every risk assessment will offer a clear-cut choice of response, even when the cost of controls is clearly less than the cost of risk materialized. This is because resources are always limited, and mitigation strategies are always going to come with an **opportunity cost**. Your choice will not always be between taking action or accepting a risk — frequently, it will be between making one investment or another.

Understanding your organization's **risk appetite** will help to sort out the edge cases, where the risk assessment supports a risk mitigation strategy, but the opportunity cost is heavy.

Risk appetite is the level of risk an organization is willing to accept in pursuit of its objectives. A **high risk appetite** means being open to accepting greater risks for possibly higher rewards, while a **low risk appetite** means the organization prefers reducing risk as much as possible.

Risk appetite will vary greatly by industry, company size, growth objectives, etc. It also exists on multiple dimensions: an organization may have a high appetite for operational risk, but a low appetite for compliance risk.

At a high level, a risk appetite statement typically takes a form like the one below:

General Risk Appetite **[Company XYZ] adopts a balanced approach to risk, recognizing that not all risks are equal and that some level of risk is necessary to achieve our strategic goals.**	
Innovation Risk	We have a high risk appetite for investing in advanced technologies and innovative solutions that differentiate our products in the competitive landscape. We understand this requires accepting a degree of uncertainty in R&D and product development.
Operational Risk	We maintain a low to moderate risk appetite. While striving for operational excellence, we prioritize initiatives that improve efficiency and service quality without compromising our delivery standards.
Security Risk	We have an extremely low risk appetite for security threats and breaches. Our commitment to network security and data protection is paramount, and we invest substantially in safeguarding our systems and our clients' data.
Compliance Risk	We have a low risk appetite for non-compliance with legal and regulatory requirements. Ensuring adherence to relevant laws, standards and best practices in all operational areas is critical.

Within each of these dimensions, there are several key factors to consider:

- **Risk capacity** is the maximum amount of risk that an organization can bear, typically decided by financial resources, operational capabilities and regulatory restraints.
- **Risk tolerance** is an acceptable deviation from its target.
- **Risk thresholds** are "red lines" that indicate the need for a change of strategy.

The threshold between tolerance and capacity, or even between degrees of tolerance, can help you sort through the "edge cases" where it is unclear if acceptance or mitigation is the appropriate response.

To return to our phishing example, there is a potential scenario in which your high-growth startup has such a high risk appetite and such a lean cybersecurity budget that the risk of a successful phishing attack is tolerable, despite the fact that controls are comparatively inexpensive.

Tips

If your organization has a dedicated enterprise risk function, ask them to provide direction, since they will usually have risk appetite documented as part of their risk management framework.

If your organization doesn't have a risk appetite statement, use this editable template as a starting point.

Examples

So far, we've been able to assign values to assets, assess threats and weaknesses, determine likelihood and impact using qualitative methods and determine risk using ALE and estimate costs of controls (and the costs of not implementing them).

Now that we have the data, let's play with some end-to-end examples.

Tip

Download an editable template to present your own assessment in this format.

Example 1: BYOD data leakage prevention

Problem statement:

Employees are increasingly taking advantage of our bring-your-own-device (BYOD) policy. We do not currently apply the same security controls on BYOD as we do on corporate-owned devices, and confidential data egress is a concern.

Employees may make use of this lack of controls to process company information outside of our environment, such as transferring data to personal emails or clouds, data transfer websites, AI tools, etc.

Recommendation:

Extending endpoint DLP, which we currently have on corporate-owned devices, to BYOD would minimize data leakage incidents (mistaken or intentional) and provide enhanced compliance.

Risk Assessment

AV: Through our BIA, we calculate the asset value of the company's sensitive data to be $1M.
EF: We can assume that in the case of a data breach, we could lose 60% of the value of our sensitive data.
ARO: By analyzing our risk register and help desk tickets, we conclude that a dangerous incident is likely to occur twice a year.

AV x EF = SLE ⇒ $1M x 0.6 = $600K
SLE x ARO = ALE ⇒ $600K x 2 = $1.2M

Uncertainty: Confidence level in the data is 85%.
There are factors that can affect this level, such as:

- The nature and extent of the leaked sensitive data.
- Potential disgruntled employees.
- Potential new employees skipping awareness training (a compensating measure we have in place now).

<table>
<tr><th colspan="5">Risk Response</th></tr>
<tr><th>Option</th><th colspan="3">Feasibility</th><th>Residual Risk</th></tr>
<tr><td>Avoid</td><td colspan="3"><b>Possible:</b> Avoidance would reduce the risk to 0, but it would require us to end our BYOD policy, which would lead to disgruntled employees.</td><td>0</td></tr>
<tr><td>Accept</td><td colspan="3"><b>Possible:</b> Acceptance is possible, but it may be outside of our risk appetite. Risk acceptance means accepting direct financial costs and indirect costs, such as reputational loss, lawsuits, fines, regulatory scrutiny, negative social media, etc.</td><td>ALE = $1.2M</td></tr>
<tr><td>Transfer</td><td colspan="3"><b>Possible:</b> Transfer to insurance would cover costs up to a certain amount. Cyber liability coverage limits typically range between <b>$500,000 and $5 million per occurrence.</b> However:
■ Costs of damaged reputation, regulatory fines and customer lawsuits will be extras unaccounted for in the insurance.
■ Insurance might stop coverage for reoccurrences if mitigating controls are not implemented.</td><td>Uncertain. Depends on insurance premium and number of events. Only financial risk is reduced.</td></tr>
<tr><td rowspan="7">Mitigate
(recommended)</td><td colspan="3">Possible – see below</td><td rowspan="7">An 80% reduction in ARO gives us a new ARO of 0.4. Residual SLE remains the same.
<b>New (residual)</b>
<b>ALE = $600K x 0.4 = $240K</b></td></tr>
<tr><td colspan="2">Costs</td><td>Benefits</td></tr>
<tr><td>Additional licenses</td><td>$30K</td><td rowspan="5">Est. 80% reduction in likelihood of insider threat incidents</td></tr>
<tr><td>Person hours required for deployment</td><td>$10K</td></tr>
<tr><td>User training</td><td>$10K</td></tr>
<tr><td><b>Total first year</b></td><td><b>$50K</b></td></tr>
<tr><td><b>Total subsequent years</b></td><td><b>$30K</b></td></tr>
</table>

Example 2: Role-based access controls overhaul

Problem statement:

The role-based access controls on our CRM system have not been updated in five years. In that time, our organizational structure has changed significantly, and the role types and their respective privileges no longer align with the actual roles of the individuals accessing the system, in many cases providing unnecessary privileges that could be exploited were a threat actor to gain credentials.

Recommendation:

While remapping role-based access will require a significant cross-functional effort, making these changes will not require system downtime, and any short-term productivity loss while users adjust to their new roles will be resolved quickly. These costs are well outweighed by the annualized loss expectancy. We therefore recommend that we move forward with this effort and are asking for support from our colleagues in sales, marketing and operations.

Risk Assessment

AV: Through our BIA, we calculate the value of the data in our CRM to be $1.5M.
EF: Considering other controls that we have in place, such as logging and monitoring and multi-factor authentication, a successful attack would be able to access and exfiltrate around 10% of the sensitive data.
ARO: By looking at pen test findings from the past five years (two tests per year with different scopes), we conclude that we can expect one incident/year involving the use of privileged access.

AV x EF = SLE ⇒ $1.5M x 0.1 = $150K
SLE x ARO = ALE ⇒ $150K x 1 = $150K

Uncertainty: Confidence level in the data is 80%. There are factors that can affect this level, such as:

- The nature and extent of the leaked customer data.
- Potential disgruntled administrators.
- Leaked credentials due to another breach.
- Threat actors evolving attack techniques to circumvent privileged accesses.

<table>
<tr><th colspan="5">Risk Response</th></tr>
<tr><td>Option</td><td colspan="3">Feasibility</td><td>Residual Risk</td></tr>
<tr><td>**Avoid**</td><td colspan="3">**None:** Avoidance would reduce the risk to 0, but we would not be able to perform any privileged actions in our CRM system.</td><td>0</td></tr>
<tr><td>**Accept**</td><td colspan="3">**Possible:** Acceptance is possible, and it may be within our risk appetite; however we must also account for indirect costs, such as reputational loss and regulatory scrutiny.</td><td>ALE = $150K</td></tr>
<tr><td>**Transfer**</td><td colspan="3">**Possible:** Transfer to insurance would cover the direct costs of an occurrence (SLE = $150K). Cyber liability coverage limits typically range between **$500K and $5M per occurrence**. However, insurance might stop coverage for reoccurrences if mitigating controls are not implemented.</td><td>**Uncertain. Depends on insurance premium and number of events.**</td></tr>
<tr><td rowspan="5">**Mitigate
(recommended)**</td><td colspan="3">**Possible – see below**</td><td rowspan="5">ARO remains the same.
A 50% reduction in EF gives us a residual SLE $75K.
New (residual) ALE = $75K x 1 = $75K</td></tr>
<tr><td colspan="2">**Costs**</td><td>**Benefits**</td></tr>
<tr><td>Staff time across IT, operations, marketing and sales to map new roles</td><td>$10K</td><td rowspan="3">Est. 50% reduction in exposure in the event of a successful attack.</td></tr>
<tr><td>Productivity loss while CRM users adapt to new roles</td><td>$5K</td></tr>
<tr><td>**Total**</td><td>**$15K**</td></tr>
</table>

Conclusion

You're now armed with the tools you need to perform a data-driven risk assessment and present your findings. You're a major step closer to reconciling business-level objectives with security objectives and building a shared understanding of risk.

And when you speak the same language as the CEO and the board, you'll achieve more than your immediate goal: you'll earn attention, awareness and sponsorship.

Throughout this e-book we've linked to additional tools and templates that can help you in your risk assessment, which you can **click here to download**:

.pdf
Business impact assessment template

.doc
Risk appetite statement template

.pdf
Attack surface checklist

.ppt
Risk assessment summary template

Good luck!

About Ivanti

Ivanti breaks down barriers between IT and security so that Everywhere Work can thrive. Ivanti has created the first purpose-built technology platform for CIOs and CISOs — giving IT and security teams comprehensive software solutions that scale with their organizations' needs to enable, secure and elevate employees' experiences. The Ivanti platform is powered by Ivanti Neurons - a cloud-scale, intelligent hyper automation layer that enables proactive healing, user-friendly security across the organization, and provides an employee experience that delights users. Over 40,000 customers, including 85 of the Fortune 100, have chosen Ivanti to meet challenges head-on with its end-to-end solutions. At Ivanti, we strive to create an environment where all perspectives are heard, respected and valued and are committed to a more sustainable future for our customers, partners, employees and the planet. For more information, visit ivanti.com and follow @GoIvanti.

For more information,
or to contact Ivanti,
please visit ivanti.com.