WithSecure™ Consulting Guide

A **risk-based formula** for security testing

Identify critical assets and test according to real-world threats and business risk

Authors:
Tinus Green and Will Jardine

Why do we need a new formula for testing?

The tools and processes used in modern system design and development are exponentially increasing the number of assets organizations must identify, manage, and secure. This is an opportunity to rethink the formula for security testing, so it addresses the real risk and impact of an attack, while being regulation compliant.

In this paper, we'll break down the process involved, explaining how to harden your operations and improve resilience by aligning security spend with real risk. The approach gives rationale to your testing program, both in defining which assets to prioritize for testing and how to test them.

Common approaches to scoping and testing are led by volume: how much can be tested within the scoped budget and how many vulnerabilities can be found? Though efficient in theory, they fail to address how attackers view an organization's estate and assets, what they will target, and how they would pivot between systems to achieve their motives. These testing approaches are also often broad, unfocused, and slow to deploy. Testing without context leads to an inaccurate understanding of posture and comes at the cost of long-term resilience. Testing without efficiency is unsustainable. Both lead to you under-securing critical assets by focusing on the wrong threats, or over-securing assets that don't need it.

Risk-based prioritization is an alternative means to decide if and how to test assets. The approach contemplates which assets require scrutiny, based on specific, real-world threats, especially those that would threaten business continuity. As a result, security testing becomes more goal- oriented and shows tangible results.

This paper explains how to perform risk-based prioritization and the testing that follows. Delivered by your security team, the exercise can increase accuracy of security spend, expedite remediation, streamline resource, and provide proof of impact for key stakeholders.

Risk-based prioritization in practice

The exercise described is used by our team on client engagements and can be applied in any organization. Organizations that are conscious of shortcomings in their system and vulnerability management are likely to see the most tangible benefits. This is especially true if you have a large estate and/or environments where pentesting has been neglected over time.

The 3 goals of the exercise are, in order:

1. **Gather context on the assets in your estate** by answering specific questions about them, via stakeholder workshops or through hands-on exploration
2. **Apply a risk rating** to each asset to help you prioritize them for testing
3. **Identify the best approach for testing** based on that prioritization

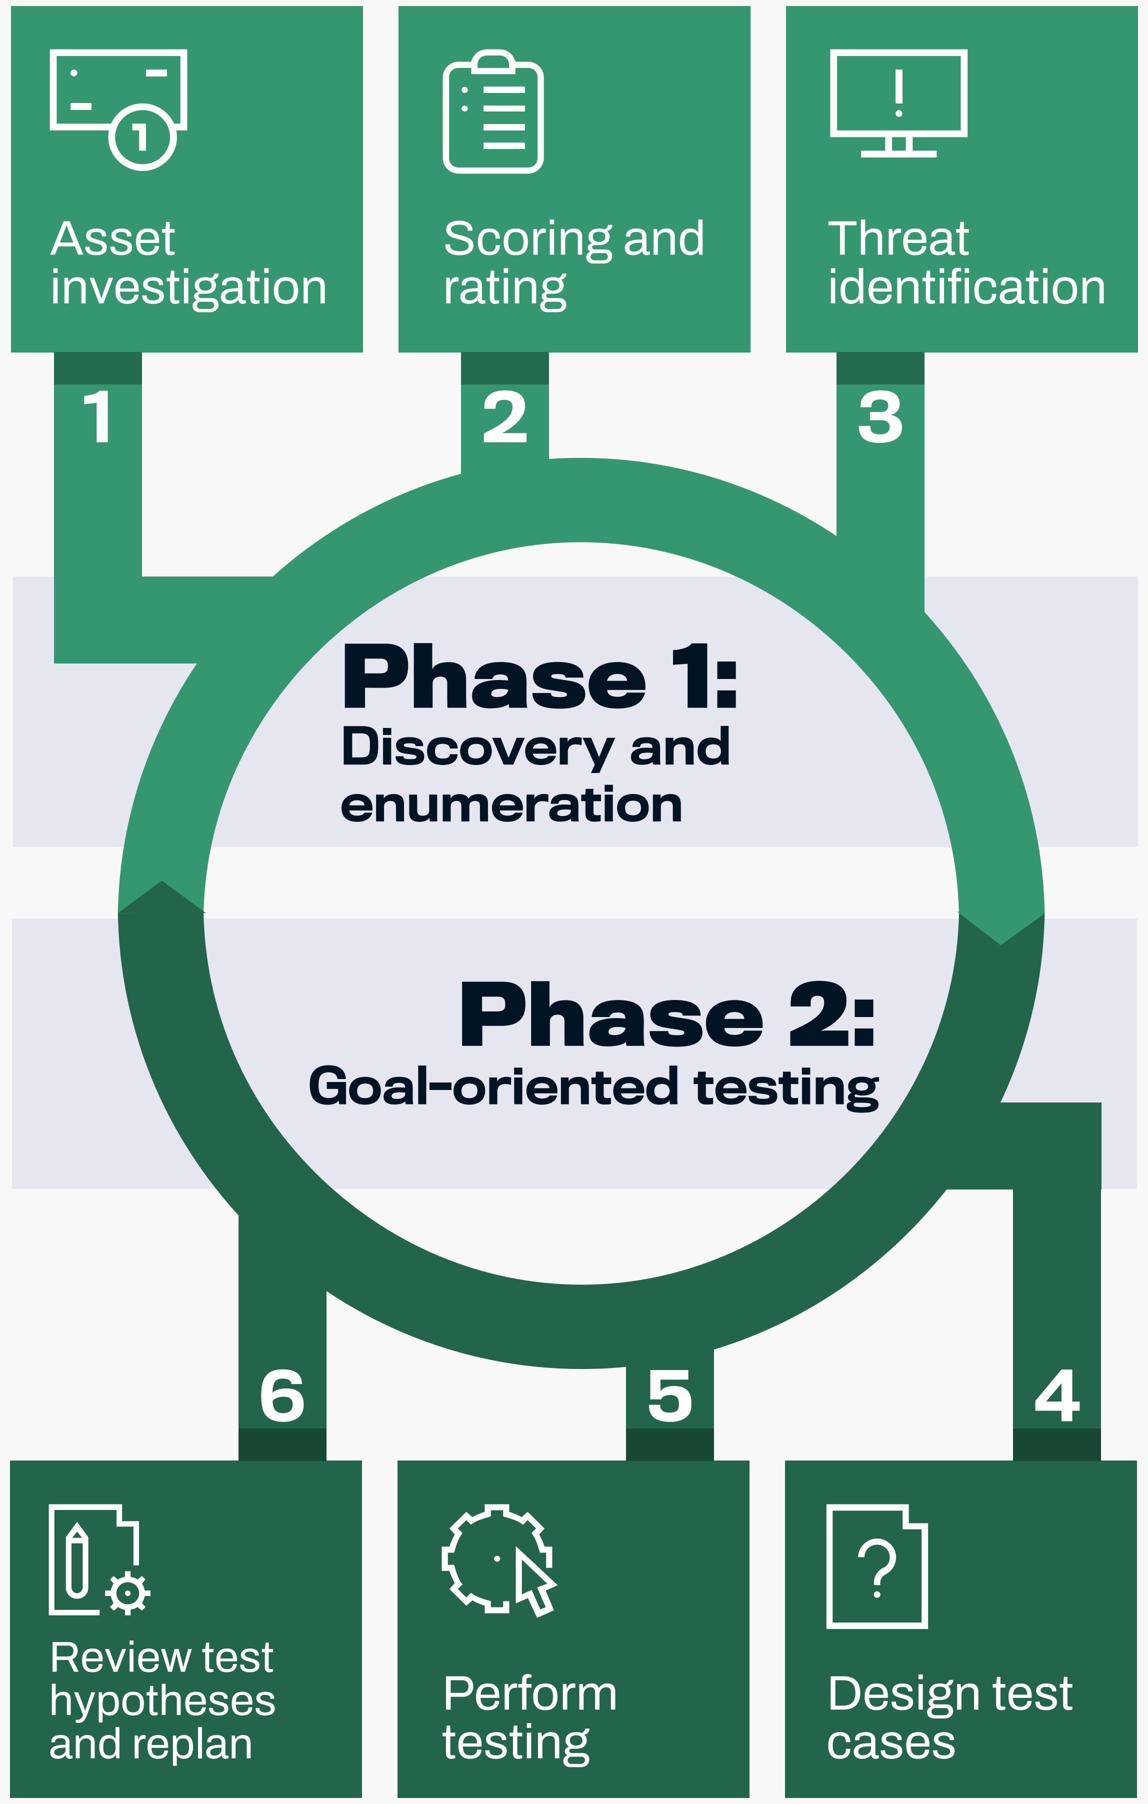

Fig. 1. The risk prioritization cycle

Risk-based prioritization can be delivered as a point-in-time exercise, or with a continuous and agile approach. Where the former is used in more selective scenarios (e.g. "We need to test all the apps related to this process"), the latter involves working with development teams over time to understand how assets are affected by their changing environment. This may seek to improve your testing approach by applying the correct context, or may introduce more assets as development continues.

From hands-on pentesters to risk/vulnerability managers, your security team is suited to perform the exercise collaboratively. The senior-most member of that team should be assigned as team lead, due to the requirement for risk awareness, understanding of business logic and architectures, and so on. Experience with threat modelling or similar approaches often proves a useful indicator of participants' ability to think in a more riskfocused and pragmatic way.

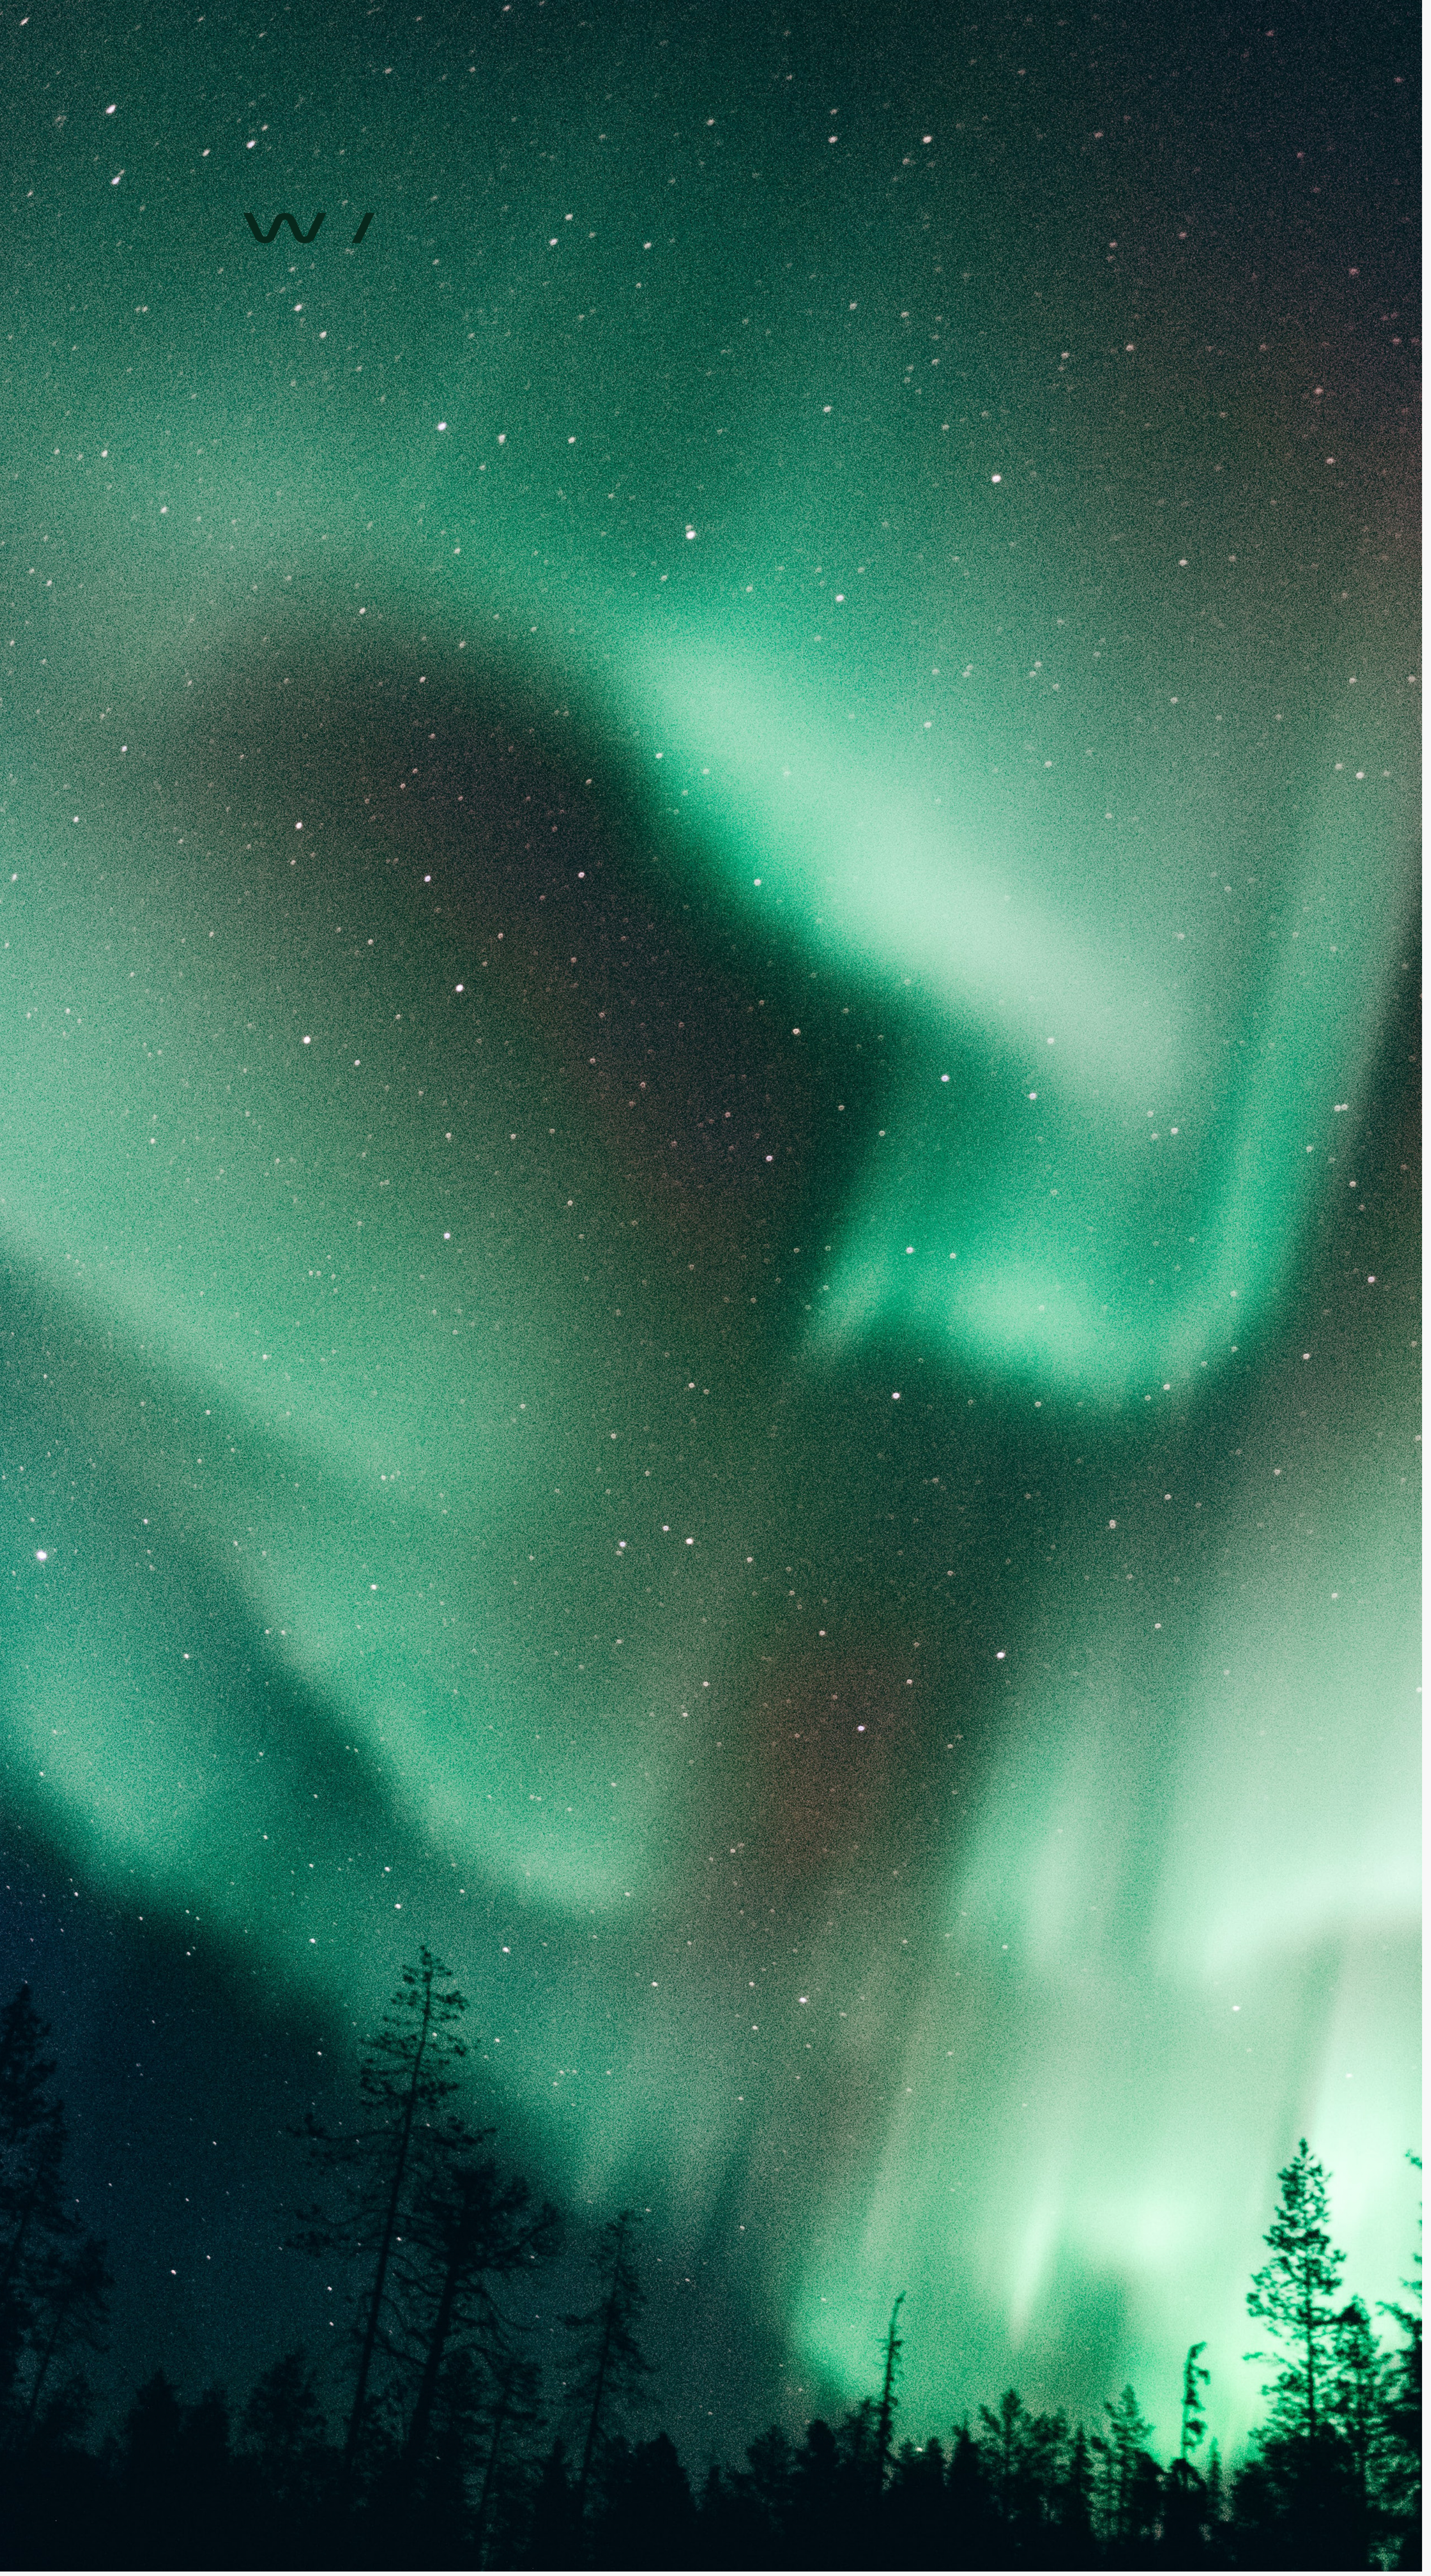

Getting started: Naming conventions

To avoid confusion, the team carrying out the exercise will benefit from aligning their understanding of the terms used throughout. The definitions on the right here are a useful baseline, but others may be needed.

- **Asset**: a system or collection of services
- **Connected asset**: a secondary asset that interfaces with an asset
- **Threat**: something that can cause harm to an asset or assets
- **Threat actor**: an attacker who uses a threat/ threats as a means to perform an attack
- **Vulnerability**: a flaw in the design, implementation, or configuration of a system that could be exploited by a threat actor
- **Goal**: the malicious actions taken by a threat actor to achieve a specific motive
- **Risk**: a combination of a) the impact a threat has on the business or its assets and b) the likelihood of a successful threat
- **Control**: any prevention or detection measure that lowers either the impact or likelihood of a threat, reducing the risk as a result

Phase 1: Discovery and enumeration

Asset investigation

The exercise begins by identifying a number of critically important assets in your estate, as well as the employees who act as their stakeholders. A discovery phase might precede this investigation if critical assets are not known. A critically important asset can be understood as *any asset that, if impacted in a cyber attack, would significantly disrupt business continuity.*

There are 2 routes to assess each asset or assets—via a workshop or hands-on exploration (explained on the next page). The process will continually expand from asset to asset until the most critical areas of your estate have been mapped out, or a decision is made to stop the exercise. Discretion should be applied to avoid assessing everything.

Key questions

For the asset

- Who interacts with the asset?
- How do users get access to the asset?
- If the asset was compromised, what would be the most likely goal of the attacker?

  a) Service disruption
  b) Financial gain
  c) Exfiltration of client information
  d) Exfiltration of employee information

For each connected asset

- What services consume resources from this asset?
- What services are consumed by this asset?
- What protocols are used for these connections?
- How is these connections orchestrated? For example, are credentials stored in a configuration file or hard-coded?
- Who are the relevant stakeholders for this asset?

The questions here are based on those we use during workshops. For the most part, yours are likely to be similar. The exception is “If the asset was compromised, what would be the most likely goal of the attacker?”, which will be unique to every organization and its risk profile.

Your answers and findings should be logged, ready for scoring and rating.

Route I. Workshop

To avoid confusion, the team carrying out the exercise will benefit from aligning their understanding of the terms used throughout. The aim of a workshop is to understand each asset, the context it operates in, and how an attacker may target it. This is done through discussion and iteration. It's less practical than hands-on exploration, but benefits from being accessible to non-technical stakeholders. Observations are documented so they can be used to assign an importance rating and a plan for testing each asset.

The following information is documented for each asset:

1. A description of the asset (its role, functionality, user interactions, and any other significant insights, such as common security concerns or uncommon connection methods)
2. A description of the asset's integrations or connected assets
3. Relevant stakeholders and asset owners
4. The goal that would lead a threat actor to attempt to compromise it
5. An importance rating (calculated using the process on the next page)
6. Recommendations for the type of testing and for how long it should be performed

Route II. Hands-on exploration

A hands-on approach can be taken as an alternative. This requires available resource in your technical team to experiment with apps or infrastructure, scope the same information covered in the workshops, and ask additional questions if necessary. It's an effective solution for teams unable to collaborate in person, i.e. if they are operating remotely, or where asset stakeholders are unavailable or uncooperative.

Hands-on exploration involves direct interaction with an asset, assessing its functionality by testing its use cases and connections to other assets. It won't be possible to answer everything here. A better objective is to formulate an initial, strategic plan of action with the flexibility to change.

An example of hands-on exploration would be logging in to an app to explore its functionality (e.g. uploads, payments, direct object references, and interactivity with another app). This gives the team an attackers' perspective of the asset from the inside out. While less cooperative, we've often found it to be a fast and efficient means of gathering the same information as a workshop.

Phase 1: Discovery and enumeration

Scoring and rating

Once all assets have been documented, they are each scored and rated according to the risk and potential impact they pose.

Risk: what is the likelihood of the asset being targeted by an attacker?

- Is the asset internal or external?
- How many users have access?
- How sensitive is the asset? (e.g., How high-profile is it? Is it financially valuable to an attacker?)

Impact: what are the potential impacts of a compromise? (e.g., financial loss, all customer information leaked, full control of the estate)

The scoring system is defined by the risk and impact factors specific to your organization. These vary business to business, but the two key metrics that should always be incorporated are "effect if compromised" and "likelihood of being targeted". Our own example is shown here.

Risk matrix		Likelihood of being targeted				
		Highly unlikely	Unlikely	Possible	Likely	Highly likely
Impact if compromised	Catastrophic	Medium	High	Critical	Critical	Critical
	Significant	Medium	Medium	High	High	Critical
	Moderate	Low	Low	Medium	Medium	High
	Insignificant	Low	Low	Low	Low	Low

Fig. 2. WithSecureTM risk matrix

Phase 1: Discovery and enumeration

More metrics can be added as per the needs and interests of your organization. Banks, for example, would likely emphasize loss of financial or client data. For telecommunications providers, it may be service disruption and control of their operational technology (OT) environment. Essentially, the risk matrix needs to sync with your internal risk register. For instance, even if a catastrophic effect is highly unlikely, it might still be rated as a critical asset in some organizations.

This risk assignment is most effective if it can respond to the business's needs dynamically. Even if stakeholders' opinions of an asset's importance contradict the levels proposed by the risk matrix, these should still be fed into the process. You may not be able to fully quantify everything, and stakeholders are likely to have personal opinions about assets' importance based on a wealth of organizational experience.

Based on the results from the scoring exercise, an importance rating from critical to low can be assigned to each asset, as in the table on the right.

With the assets listed in order of risk, security tests may now be tactically scoped. The list can be mapped to strategic security and business goals for the buy-in of other stakeholders, as well as to show the increasing maturity/resilience of assets over time.

Critical	The possibility of the asset being targeted by a threat actor is likely or almost certain and the impact of a compromise is either critical or catastrophic.
High	The possibility of the asset being targeted by a threat actor is likely and the impact of a compromise is either major or critical.
Medium	The possibility of the asset being targeted by a threat actor is possible or likely and the impact of a compromise is major.
Low	The possibility of the asset being targeted by a threat actor is unlikely, although possible, and the impact of a compromise is either insignificant or moderate. This rating can also be assigned to assets that have high importance but don't belong to your organization.

Fig. 3. Table of risk ratings

Phase 1: Discovery and enumeration

Threat identification

The information gathered from workshops/ hands-on exploration and rating exercises can be used to create models of your organization's environment. These provide the security team with a visual reference showing how assets are connected, i.e., how the risk of one asset affects the others around it. There are examples of these models on this page and the next.

The first shows the connections between 4 assets in an estate. It also indicates the exposure of the assets. A threat actor's ability to modify data in Application A would be a medium risk vulnerability should it be tested in isolation. However, the model shows how this vulnerability would go on to affect Application B's users, contextualizing the vulnerability risk as a significant business risk too. Based on the sensitivity of Application B, the organization can decide whether remediation of the medium risk vulnerability should be prioritized, perhaps above that of other, high-risk issues.

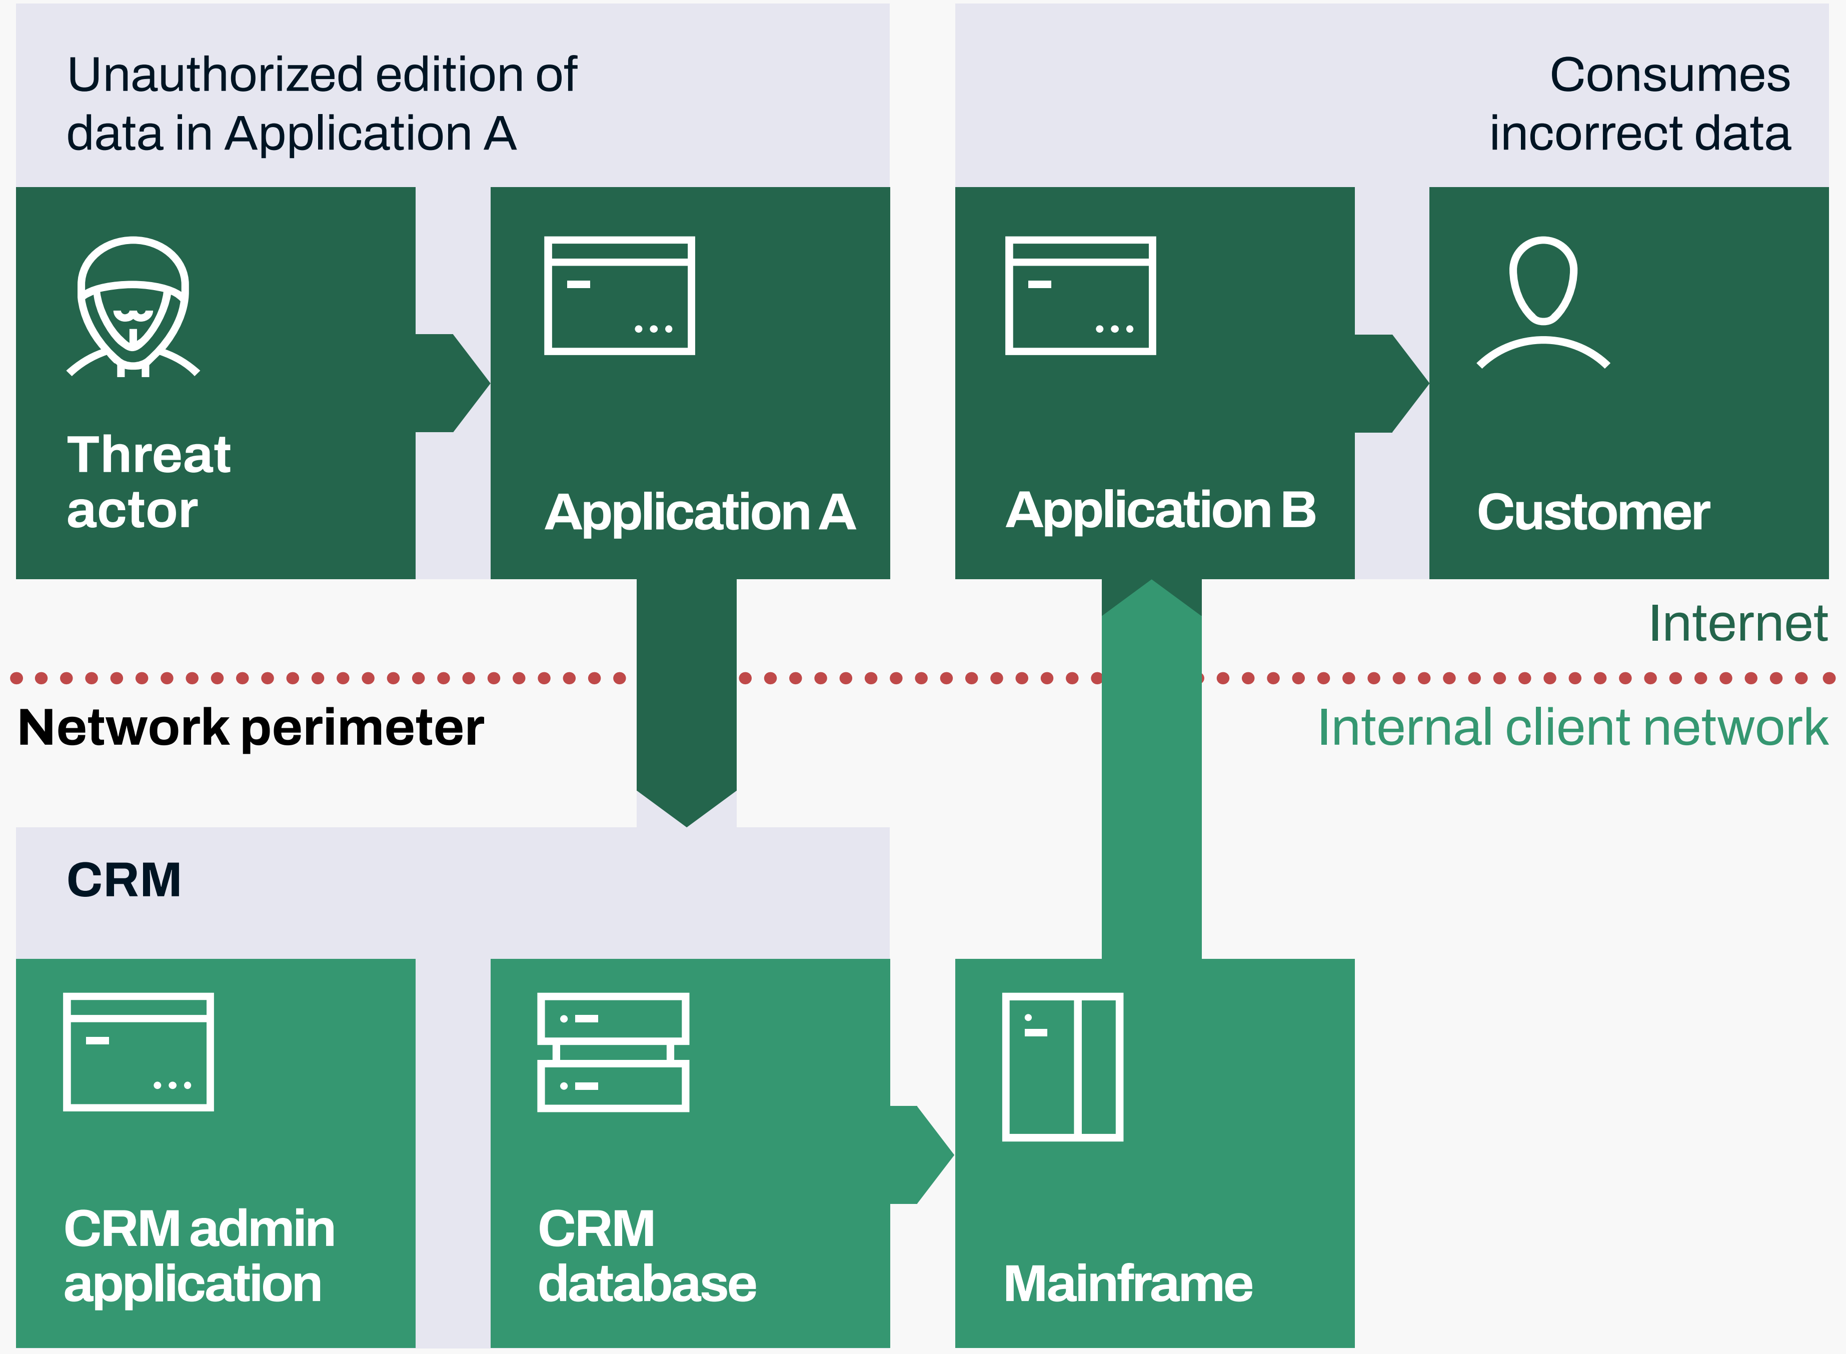

Fig. 4. Model of asset connections and attacker goals

Phase 1: Discovery and enumeration

Fig. 5. indicates another 4 assets and their interconnection. It also shows the potential attack injection points, including the compromise of a middleware admin and access to one of the communication channels between assets.

The likelihood of these attacks can be determined according to the exposure of said channels and the number of middleware admins. For example, the likelihood of attack on the middleware asset is low, because it is internal-facing, with access restricted to few users. Yet, the impact may be considered significant due to the unrestricted access it provides. Overall, this would constitute a medium importance rating.

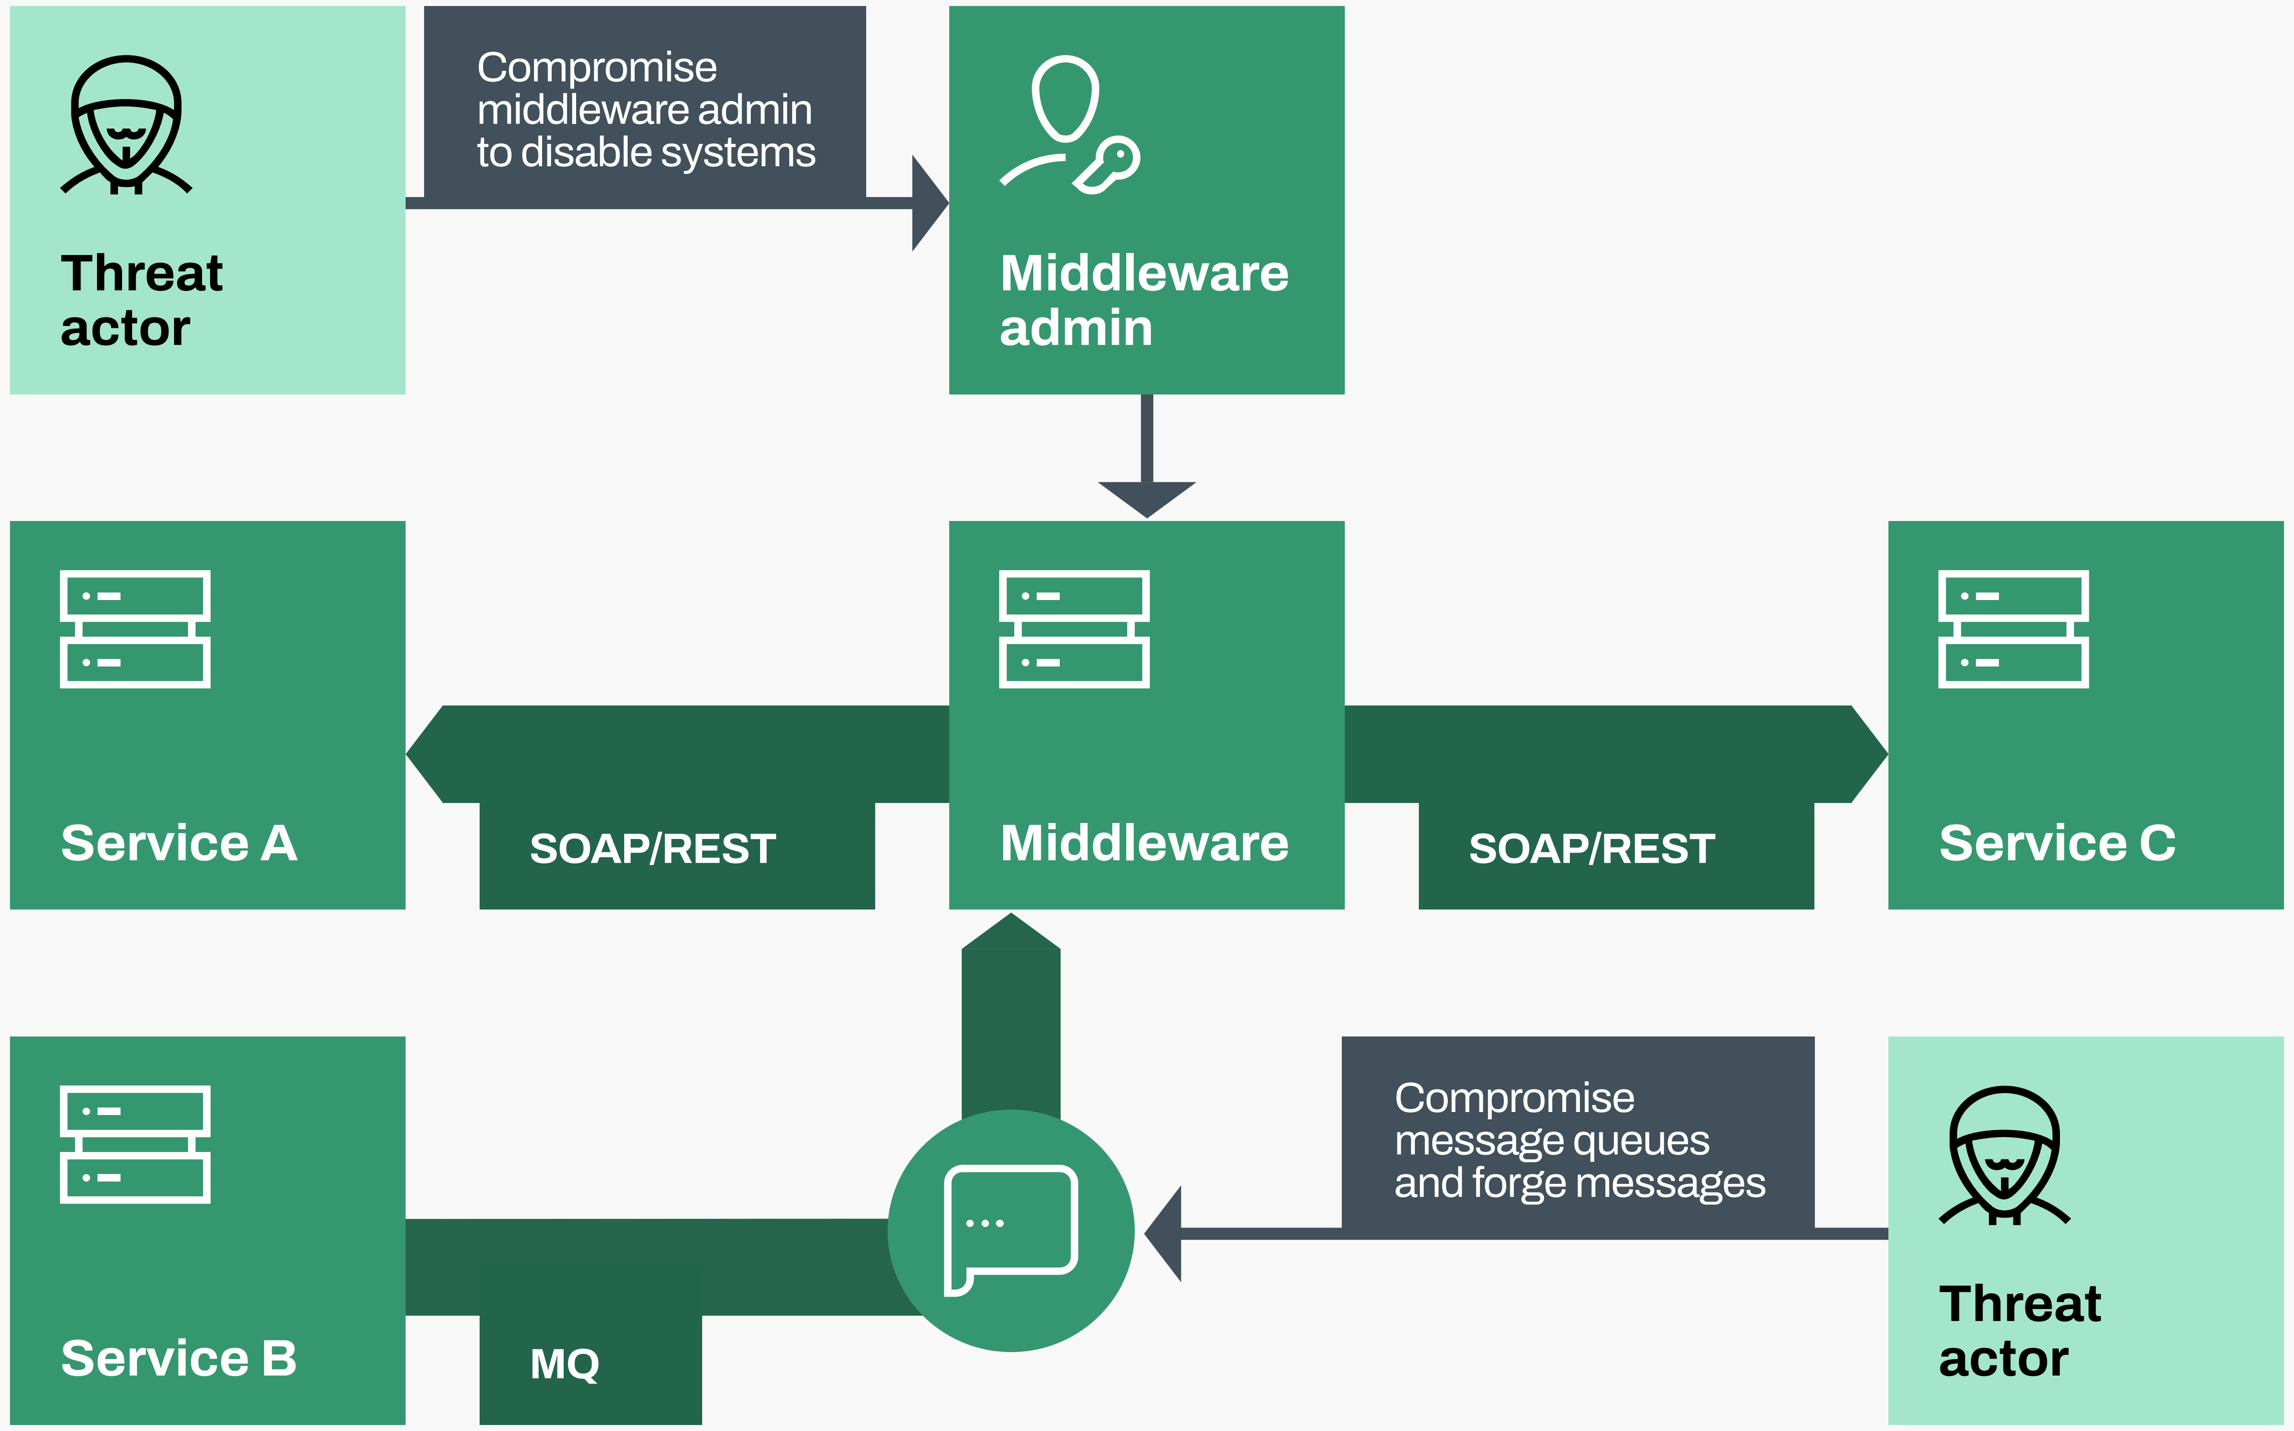

Fig. 5. Model of asset connections and input points that could be compromised by threat actors

The analysis exercise on this page and page 8 is the culmination of understanding:

- What assets you own
- How they are connected
- The likelihood and business impact of assets being compromised
- Their risk rating and the consequent importance of testing them

This context can be used to hypothesize which assets will likely be targeted, and how, based on their attractiveness to threat actors. This is exemplary of the offensive “attacker mindset” organizations should be striving for in their security work.

Phase 2: Goal-oriented testing

Test case design and testing

Discovery and enumeration complete, you can now start to design tests around how your assets may be targeted by threat actors harboring specific malicious goals.

Step 1 of the exercise is to build a test case for each asset, starting with the most critical. Each case forms the basis of a test to uncover potential threats/goals (why an attacker would target an asset), such as “We can bypass MFA to steal PII from customers”. The ultimate objective is to define the most contextually-relevant way to test assets against each goal. In doing so, you can prioritize your effort, resource, and testing methodologies towards the most impactful areas - where successful execution of a goal would meaningfully impact the organization - thus removing any focus on less relevant areas.

Threat/goal	Testing approach
Theft/bypass of intended e-commerce/purchase logic	• Looking for logic flaws in the order process • Insecure direct object references (IDORs) in basket functionality • Access control issues • Considering third-party technologies or APIs used for payments
Accessing other customers’ baskets or personally identifiable information (PII)	• Access control issues • Client-side issues (e.g. Cross-Site Scripting (XSS), Cross-Site Request Forgery (CSRF)) allowing theft of data • Compromise of database via SQL injection or similar
Abusing complex access control/identifier logic to escalate privileges	• Privilege escalation testing, horizontally (user to user) and vertically (user to admin) best accomplished through collaborative testing • Identifying single sign-on (SSO) solutions or similar that have control over authentication and authorization for the asset • Multi-factor authentication (MFA) bypasses
Reviewing the use of third-party and potentially insecure JavaScript dependencies as a means of compromise	• Mapping out any JavaScript use • Reviewing use of Subresource Integrity (SRI) or similar
Compromising back-end development via outdated versions, injection vulnerabilities, or similar	• Reviewing all disclosed software and versions • Exploring relevant CVEs and exploits • Exploring injection points for issues such as SQL injection or command injection • Focusing on sensitive functionality, such as file uploads

Fig. 6. Table showing examples of threats and goals alongside corresponding contextual testing approaches

Phase 2: Goal-oriented testing

Review test hypotheses and replan

A test case proven incorrect is evidence of controls working effectively, which demonstrates successful work performed by development and security teams. Contrasting the often negative, defect-focused approach of standard pentesting, it recognizes where things are working, so the same measures may be applied elsewhere. It also quantifies prevention budget spent wisely.

These incorrect test cases are one output of risk-prioritized testing and can be used as the starting point for other testing. What this testing is and where it's focused will be unique to your organization. Fundamentally though, risk-prioritization is the foundation needed to deliver a testing program that is efficient and effective overall by instilling the risk/threat mindset.

Testing must be delivered cyclically to reliably increase your security posture. This will keep it current, in relation to changing assets and evolving threats, and account for discoveries made along the way, which demand continued investigation.

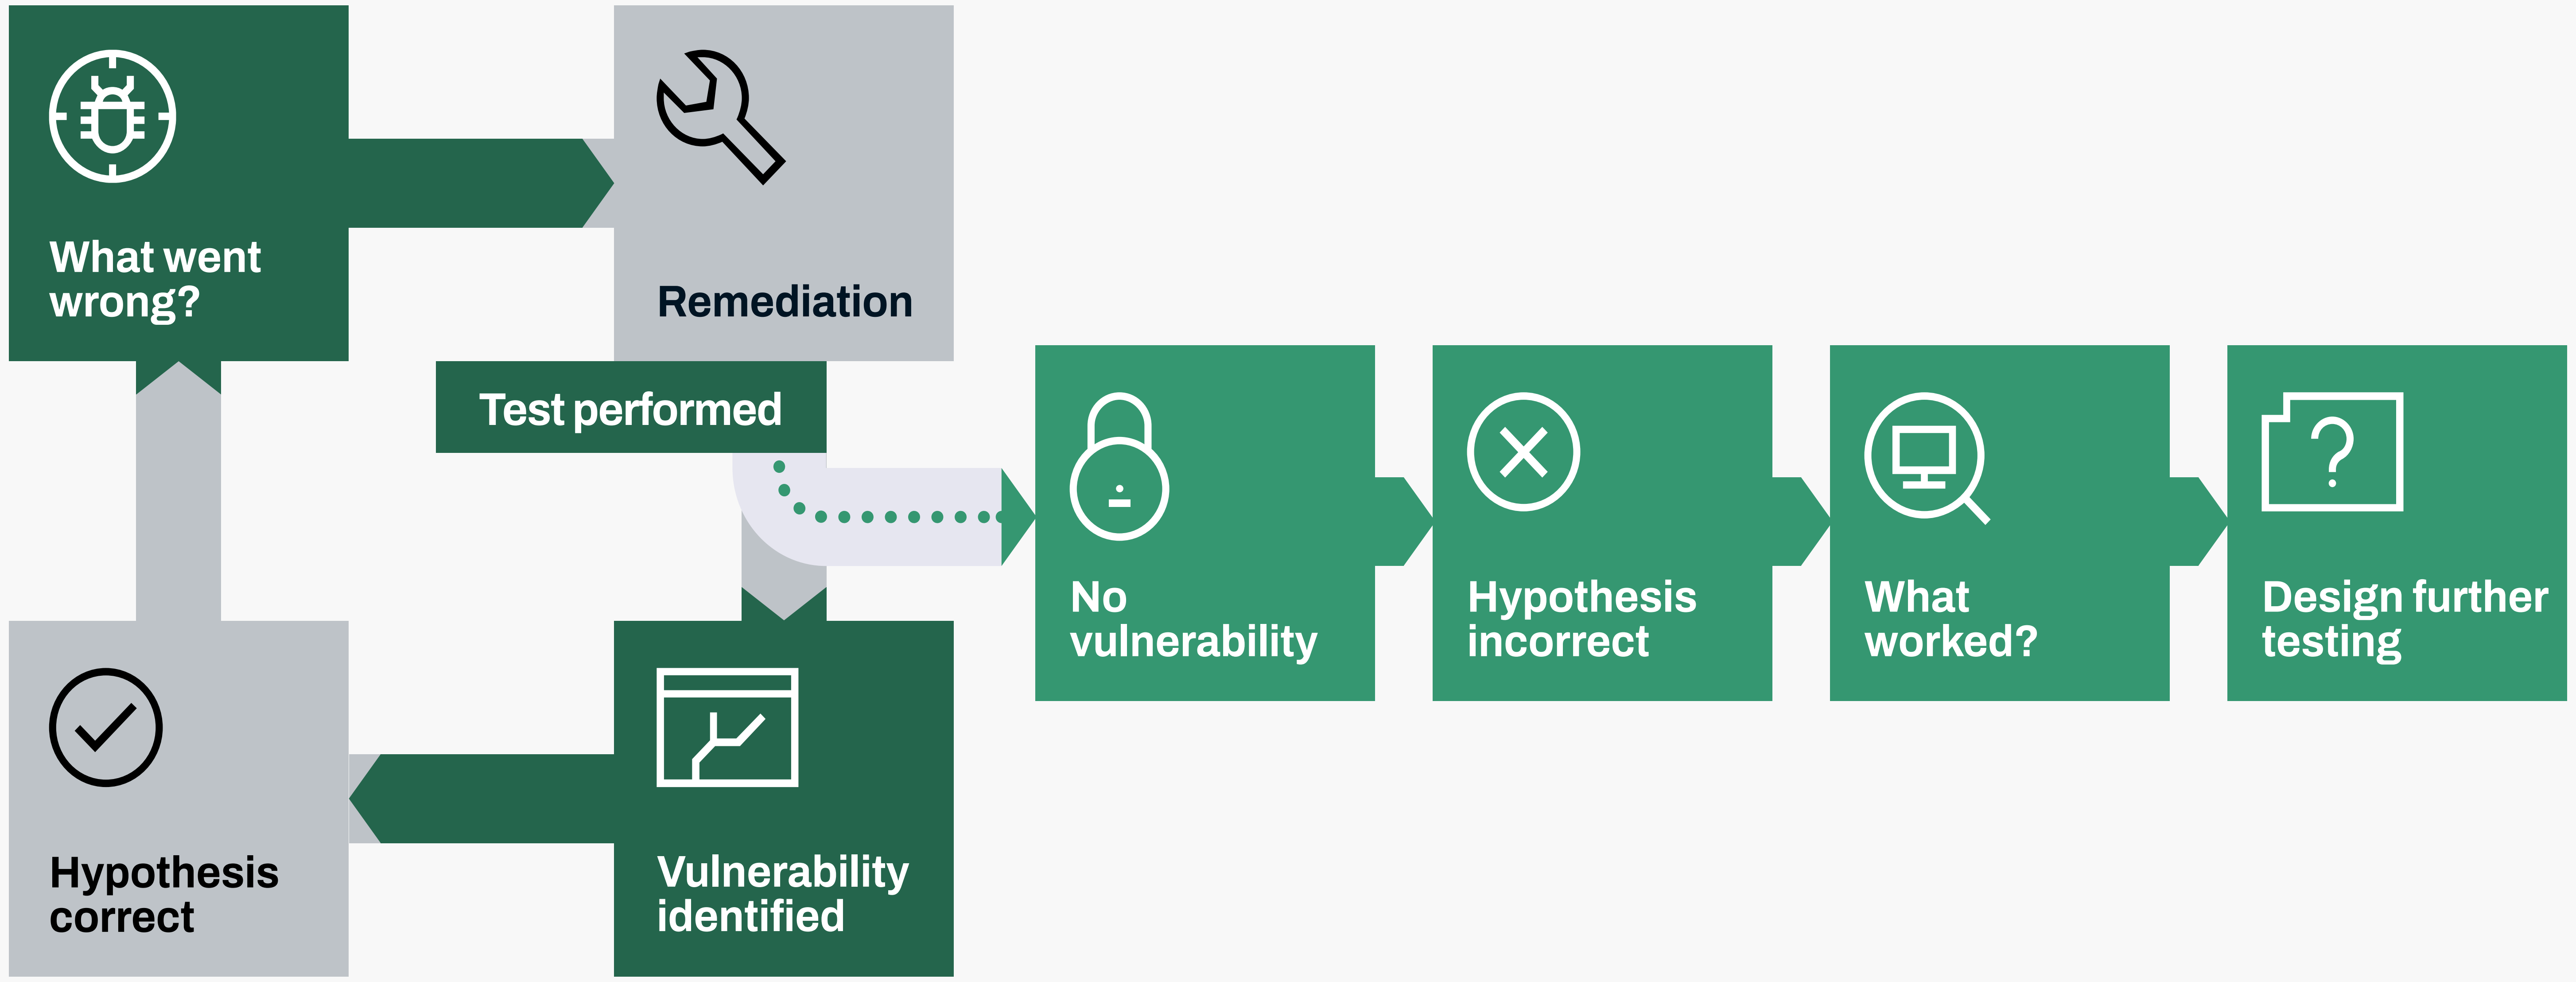

Fig. 7. The review and replan cycle for incorrect and correct hypotheses

Long-term gains

First and foremost, a risk-prioritized approach to security testing enables organizations to focus on the assets most in need of attention. Over a longer period, it does even more to make your security testing work harder.

These long term gains include:

Focus objectively on the threats that matter most

Each asset is reviewed from the perspective of an attacker, acknowledging that assets supporting critical business processes are likely to pose a significant threat. This means the risk profile of the wider business is accounted for. High-risk/impact assets are prioritized and subjected to appropriate testing and more of it than assets that pose little to no material risk. Testing can then focus on key threat actor goals for each of those assets—goals which would meaningfully affect them—rather than just performing generic pentesting and achieving generic results. Resource and time can be scoped for maximum investment. Critical threats can be highlighted and remediated, contributing to reduced risk exposure and overall resilience.

Overcome the static nature of point-in- time testing

Delivered iteratively, prioritization takes care of assets that would otherwise go untouched for months between engagements and be exposed to threats in that time. By exploring your assets iteratively, reviewing findings, and investigating further, there's a greater chance emerging vulnerabilities will be flagged and assets get tested before they can be exploited.

Put past testing to use

You can optimize your risk-prioritized testing over time by consuming information from previous security assessments and threat analysis exercises. This data may be used to tailor and adjust future tests so you don't duplicate the effort, thus adding another layer of efficiency and improving resilience.

Summary

Risk-based prioritization reveals your most critical assets and the top risks they face. This enables you to focus on prevention in these areas specifically. You will know which assets matter most, which threats are of greatest concern, and have evidence of your ability to prevent them. You can then start to consider the detection of attempts to utilize those key threats, increasing your overall confidence in those assets.

This is a solid and accepted logic within infosec, yet the knowledge hasn't been available through standard scoping and testing processes before. It is generally applied to infrastructure assets, but less so to applications, perhaps due to decreased maturity of the focus on those key goals/threats for many app testing programs.

The overall objective of this prioritized way of working is operational resilience. Organizations can move away from traditional approaches that no longer suit the scale and complexity of their environments, nor the evolution of the threats they face. Build a resilient approach to testing—the bread and butter of your security—and you've got solid foundations for making further intelligent, targeted investments elsewhere.

Who We Are

WithSecure™ is cyber security's reliable partner. IT service providers, MSSPs and businesses along with the largest financial institutions, manufacturers, and thousands of the world's most advanced communications and technology providers trust us for outcome-based cyber security that protects and enables their operations. Our AI-driven protection secures endpoints and cloud collaboration, and our intelligent detection & response is powered by experts who identify business risks by proactively hunting for threats and confronting live attacks. Our consultants partner with enterprises and tech challengers to build resilience through evidence-based security advice. With more than 30 years of experience in building technology that meets business objectives, we've built our portfolio to grow with our partners through flexible commercial models.

WithSecure™ is part of F-Secure Corporation, founded in 1988, and listed on the NASDAQ OMX Helsinki Ltd.

W/TH
secure