Proton

A Practical Guide to Security for **Growing Businesses**

By Patricia Egger

Privacy by default

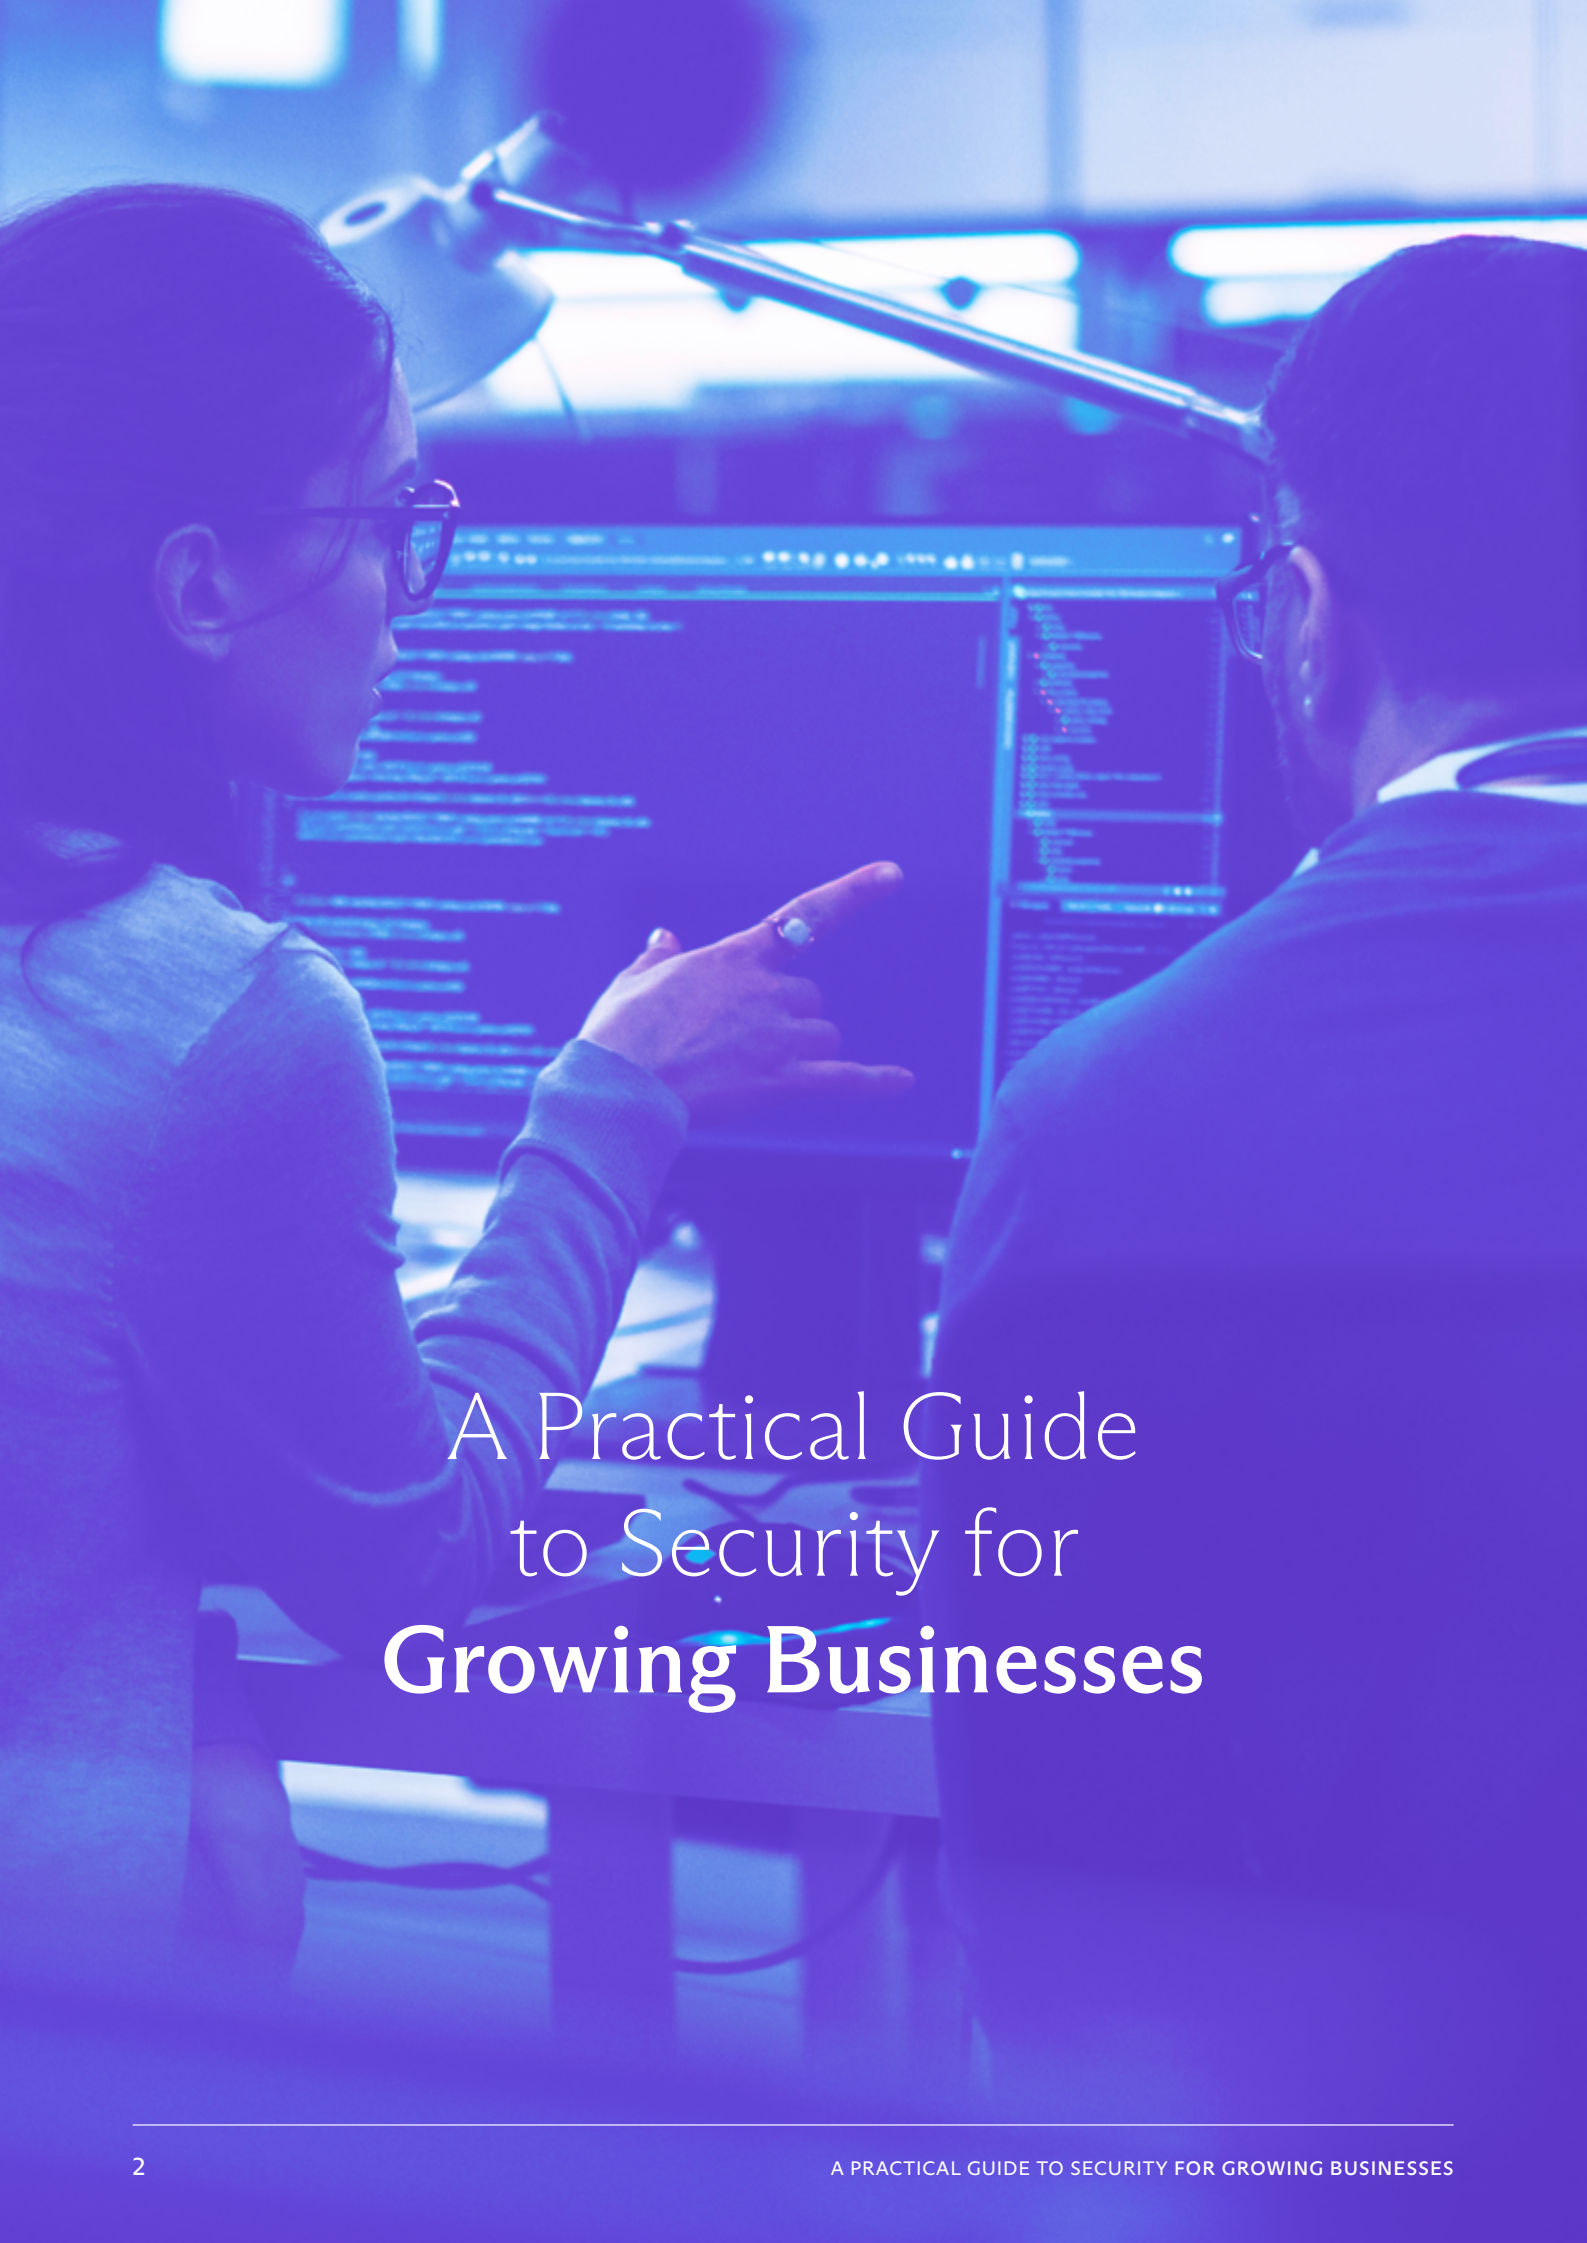

A Practical Guide to Security for **Growing Businesses**

Table of **contents**

Introduction

Chapter sections

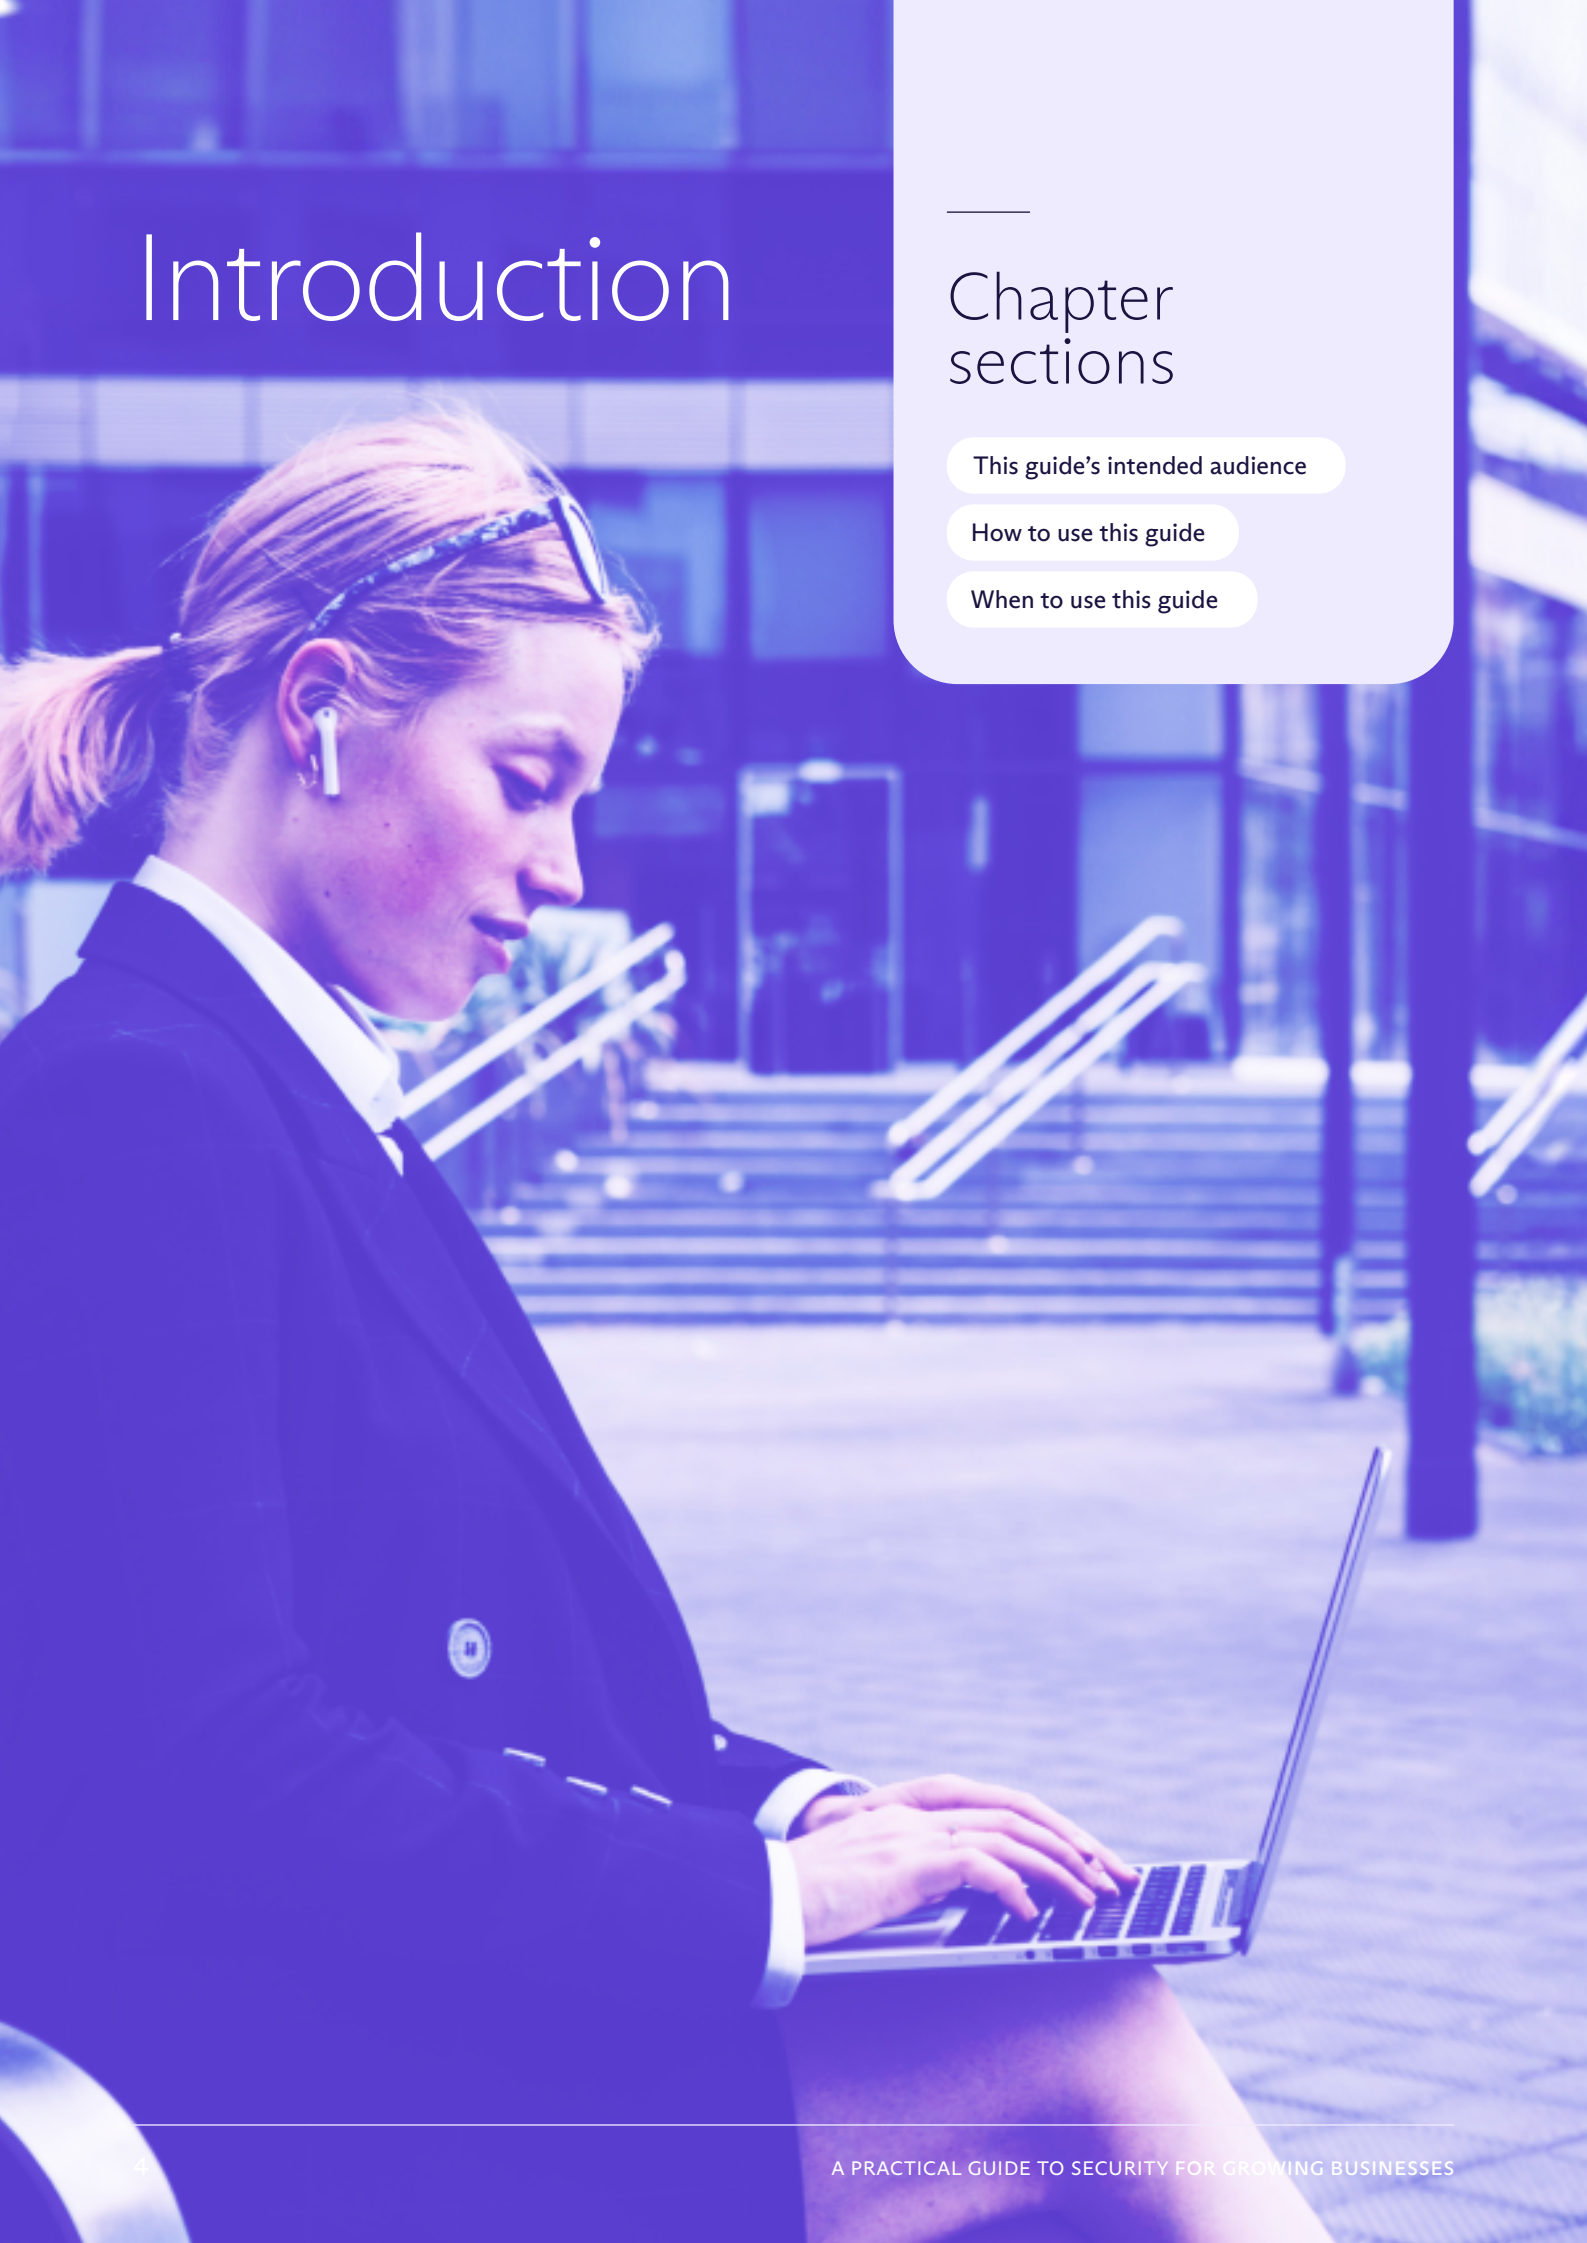

This guide's **intended audience**

While this guide contains a description of security and crucial questions that everyone can use, it will be the most helpful if:

1. Your business uses computers or other internet-connected devices and handles sensitive data.
2. You are the single decision maker or one of the decision makers for a small or medium-sized business.
3. You are not a security expert, but you do have a basic understanding of technology.
4. You want to set up your first security policies or improve your current policies, but you're not sure how to get started.

Regardless of what your business is, security risks are almost inescapable nowadays. Modern security is all about understanding and mitigating those risks to an acceptable level and minimizing the potential damage if anything goes wrong.

If you run a business, you must find ways to operate and serve your customers while facing these risks. You also have data and other assets that are worth something to someone. Even if you think you're too small to be a target, you must consider how to prevent internal breaches. (Unfortunately, no one has yet devised a program that allows you to completely avoid mistakes.)

Knowing what you need to protect and why gives you a chance to come out ahead. Unfortunately, this is easier said than done, and the road to success has no single definitive path. To extend this metaphor, at each fork in the road, the correct answer will depend entirely on what best suits your organization.

Finally, it's also important to remember that your security function doesn't exist in a vacuum. It's an outgrowth of your business. Just as your business depends on your security, your security function is constrained by your organization's resources. A limited security plan you can implement with the resources at hand is infinitely better than an idealized security plan that your organization doesn't have the funds or personnel to put into action.

It is not uncommon to hear that security slows down a business by putting restrictions and rules that staff must navigate and comply with. I don't believe that to be entirely true, and so this guide attempts to teach you how to weigh all of these factors (and indeed many more) so that you can identify how to get the biggest bang for your buck and develop a security plan tailored specifically to your organization.

How to **use this guide**

This guide aims to help you understand the concrete steps you can take to improve your company's security and keep your business safe. Proton has been a leader and innovator in secure communications since we first launched Proton Mail in 2014, now the world's largest encrypted email service. Since then, Proton has built an open-source ecosystem that includes Proton VPN, Proton Calendar, Proton Drive, and Proton Pass, leveraging the same advanced encryption to give individuals and businesses full control over what happens to their data.

We've tackled security both as a provider of secure tools and as a startup that has needed to grow its own security function over the years. Building an end-to-end encrypted ecosystem and going from a small crowdfunded project to an organization with close to 500 employees has given us unique insight into security — both technical and managerial. As a member of the Proton security team, I want to share that insight with you.

Most IT security books are very prescriptive, and while I cover some best practices, **I focus more on the most important questions you should ask yourself about your security.** This is because there is no single correct way to do things — each decision you make will have advantages and disadvantages that have ramifications for your business. The key to security is understanding and navigating these tradeoffs to best achieve your goals.

This might differ slightly from other guides you've read, but it's an introduction to the systems thinking that I believe is best suited to address security at any organization. In this context, **systems thinking means considering the many components that influence each other within your business to get things done. It's not enough to understand only the individual components that make up your organization — you must also know how they interact with each other.**

The fundamental principle to remember is the performance of a system, in this case your security, doesn't depend on how the apps, tools, or individuals perform on their own — it depends on how they perform together. Everything is interconnected. To paraphrase

Russell Ackoff, a pioneer in the development of systems thinking, "A system is never the sum of its parts; it's the product of their interactions."

In practice, security is always more complex than it seems, and many organizations start by doing things that seem evident — either due to the nature of the business or because they're seen as best practices. It may take years for someone to question why your policies contain specific requirements or why you prioritized certain projects or investments. Often, it's because best practices or generic security controls were implemented without much consideration for your organization's specific risks. Typically, it's not until it seems like nothing in your security policies makes sense anymore that teams take a step back and re-evaluate. This is not to say that best practices are bad (they would not be called "best" practices otherwise) — the art lies in not being overly focused on best practices at the expense of your specific risks. The goal should always be to ensure you address all the risks your specific organization faces. In software engineering, there is a term for blindly following best practices: cargo cult programming. As you may imagine, it has a very negative connotation and a similar term should be coined for the security community.

I've tried to make this guide as practical as possible by identifying the most pressing decisions you need to make and an overview of the tradeoffs each decision will present. This is also a quick guide to help beginners get up to speed and develop their own security plan and security function quickly. It is not intended to be an exhaustive guide.

When to **use this guide**

To answer simply, the best time to use this guide is anytime you feel uncertain about the direction of your security efforts. I discuss the questions you must ask yourself and the consequences each potential solution will likely have for your organization. As such, it should be equally helpful to someone revising the direction of their company's security as it is to someone making their first investments in their company's security.

The optimal time to use this guide is when starting your security journey, as this will allow you to establish security policies, processes, and tooling that are set up for your long-term success. Many organizations say that security is a priority but then fail to give it much strategic thought. Instead, security decisions and investments are made and changed organically, typically in response to incidents or new developments. This is understandable, but it's not the best way to create a thorough and effective security function.

> I recommend you make security one of the factors that will affect any decision you make throughout the life of your company. Ideally, this consideration for security will be part of your decision-making process from the very beginning. After all, if you never make a mess, you never need to clean up (although it's never quite that black-and-white).

However, if you're already a medium-sized organization and feel like things are getting out of control, you can always take a step back, revisit these questions, and re-adjust.

Even if it may not seem obvious at the time, the decisions you make regarding your security will have far-reaching implications for your business. No solution will work for every organization. This can make getting started daunting since creating an effective security function is more complicated than simply following a checklist. However, by understanding the consequences of each decision, you can optimize your resources specifically for how your business operates.

Chapter 1

Chapter sections

Developing your **organization's security function**

Throughout this guide, you will see I refer to a "security function" as well as your "security team". Why? Well, your organization's security relies on much more than just your security team. Consider this a concrete example of adopting a systems thinking approach to security. **Your security function encompasses everything that goes into securing what matters to you**. Of course, your designated security team is part of that, but so are the rest of your employees, the tools you use, the policies you put in place, and how they all interact.

To keep your business secure, I suggest you use a systems thinking mindset to identify what you must protect, what threats you must defend against, and what tools or activities you can use to do so.

The questions **you should ask yourself**

What are you protecting? And from whom?

This is probably **the most important question you'll ask yourself** when setting up your security. You should also think about it yourself and have discussions with your security, legal, privacy, finance, and/or engineering teams (if you have them). While I recommend you get input from all of them, I do **not** recommend you let them decide what is important and what is not.

Decisions will need to be made. Resources and projects will need to be prioritized. This must be a business decision made at the very top and must be done explicitly. I also recommend you put these decisions on paper, making it crystal clear what your overall security objectives are. If it helps, you can think of the CIA triad — confidentiality, integrity, and availability — when defining your security objectives, although there can be more to it than that.

Not every company will have the same definition of what is included in "security", but practically every security function will aim to protect both the business and its customers. There's an overlap, of course but there may be parts of your business that aren't related directly to your customers that you still want to protect.

Like in any strategic area, it is important to focus. If you don't separate what's important from what's not, everything will wind up in a pile of indistinguishable value. This is particularly important at the start of a company, when resources may be particularly scarce. That's not to say that you should never review and update this focus. On the contrary, any major change in the business should lead to a review of priorities. For example, launching a new product, entering a new jurisdiction, catering to a new user base, etc., are all occasions that would require a review and perhaps an update. And if you go a year without any major new developments, it's probably a good idea to at least talk about your security objectives anyway.

If I can give you one piece of advice, let it be this: **Be very clear on what data you care about and what you can let go. If everything is important, then nothing is.**

What does security mean for my organization?

Let's start at the ground level. **Security is the umbrella term covering the people, processes, and technology used to protect what you care about. In a company, security should identify risks and put defenses in place to manage those risks to an acceptable level, allowing the business to achieve its desired results.**

In theory, deciding what security activities and investments to choose is straightforward and follows the steps below:

1. Determine security risks. .
   In this context, **a security risk is anything that could happen that has the potential to derail your security objectives.** There are many risk frameworks and standards that you may take inspiration from, such as ISO 27005, COBIT 5 for Risk, or the NIST Risk Management Framework. However, they can be quite complex, especially for small organizations (and even SMEs). **I suggest you keep your risk assessment as simple as possible to keep things manageable.**

   Start with "categories" of things that could go wrong (for instance, business disruption, operations disruptions, data leaks, or non-compliance). For each category, think of "scenarios", or single-sentence descriptions of what could realistically go wrong in your organization. For example, within the "business disruption" category, you could have a scenario stating:

> Let's start at the ground level. Security is the umbrella term covering the people, processes, and technology used to protect what you care about. I suggest you keep your risk assessment as simple as possible to keep things manageable

If you want to go further, you can add one more level — risk sources. Risk sources are ways in which a specific scenario could come to pass. Keeping the previous example (a business disruption caused by a service provider), you could have the following risk sources:

"One of our main service providers is offline — either due to a technical disruption or going out of business — leading to a disruption of our services".

Whichever is the last level of detail you decide to use, those are your "security risks".

"The core of our business (for example, our website) relies entirely on provider XYZ, with which we have no signed Service Level Agreements (SLAs) and which has had outages in the past. If it suffers any disruption, our services would be significantly slowed down or completely halted."

Note that there's a lot of confused nomenclature around information security out there. Risk, threat, and vulnerability are often used interchangeably. These are not the same thing, but if you're just starting out, you probably don't need to get into threats and vulnerabilities (at least not yet).

When determining the security risks you face, consider being the target of a malicious attacker trying to obtain sensitive information (your intellectual property, your customers' data, etc.), having an employee make a big mistake (deleting a database, introducing a vulnerability in a product, sending sensitive data to the wrong person, etc.), or having your operations disrupted by the absence of certain employees with specific knowledge or access rights.

2. **Assign an owner to each risk.** The owner of a risk is someone (a role or title, not a named individual) who is accountable (see "Defining security responsibilities") for the consequences to the organization if the risk were to materialize. Basically, if this risk happens, its owner is who gets the blame. It's extremely important to understand that this person must have the authority to make a decision regarding that risk and the resources to do something about it. If they don't, someone higher up should be the owner. In any case, there must be an owner for each risk to perform the next step.

3. **Measure each one of these risks.** This doesn't need to be scientific, but you should generally create levels that describe the likelihood of each risk happening and its potential impact (you can use a simple "high, medium, or low" scale for both). You then combine those evaluations to create a risk level (also "high, medium, or low" to keep things simple). This can be done by assigning a value to each level (for example, 1 for low, 2 for medium, up to 5 for high) and multiplying the likelihood and impact values to obtain the risk level. Personally, I prefer to work with a risk matrix, as shown below. The values assigned to each combination of likelihood and impact are up to you.

		Impact		
		Low	**Medium**	**High**
Likelihood	**Low**	Low	Low	Medium
	Medium	Low	Medium	High
	High	Medium	High	High

As you can see from the matrix, a risk with a low likelihood could still end up with a medium risk level if the impact is high. Once you have risk levels for each risk, you can (and should) do a comparison check. For example, if risk X comes out as high and risk Y as medium, make sure risk Y is indeed not as bad as risk X. **Risk owners must work with those overseeing the risk measurement process and should agree on its risk level.**

4. **Decide what you want to do with each of the identified risks.**
   This is referred to as "risk treatment" and generally results in one of four options:
   a. **Avoid** - Stop doing the thing that leads to the risk
   b. **Share** - Get insurance or work with some other entity to put some of the risk burden on them
   c. **Mitigate** - Find ways to reduce the likelihood and/or impact of this risk
   d. **Accept** - Indentify the risk as one you can live with.

The decision of what risk treatment to select should belong solely to the risk owner.

5. **Treat your risks Add period.**
For risks you decide to mitigate, do something new or different. This may take many forms, such as creating or changing a process, training staff, implementing a new tool or control, etc.

Note that accepting a risk doesn't mean you don't need to do anything. You want to ensure that the level of risk you accept remains the same (or at least that the situation doesn't get worse), especially if that would make the risk no longer acceptable. So when you accept a risk, be sure to think about what you must do to maintain its current level.

Everything else you do in security should be a logical consequence of these five points.

What is in the scope of **your security function?**

You can break your security function down into three sections:
1. Your policies, which impact everyone. (Rules)
2. Your security team, which is the heart of your security function. (Personnel)
3. Your controls, which are the protections and activities you've put in place. (Defenses)

Policies

These are documents (yes, plural) that describe what is permitted and/or what is not in your organization. You may have an overarching "security policy", but these tend to be so high level they bring no real value. Instead, I recommend breaking these rules down into categories. identity and access management, acceptable use, data and systems handling, etc. Each document should have the shortest and clearest list of requirements and restrictions possible. The shorter they are, the more likely your team is to adhere to them.

While security policies tend to grow into existence organically (typically in response to incidents or questions), I recommend thinking about the categories in advance and even creating blank draft documents until you are ready to tackle them. This is because a set of documents that's created organically tends to have overlaps and gaps, both of which cause confusion.

Trust me, it's worth spending a few hours thinking about your security policies before getting started, as this will ultimately save you countless hours, headaches, and seemingly endless discussions later on. Remember, your security policies are the ultimate rules about what is allowed or not. They should send a clear message to your staff.

Security team

The heart of your security function will be your security team. In practice, security teams will vary significantly from one company to another, both in size and in scope. For instance, some security teams cover physical security, while others will have a dedicated and independent team for that. Similarly, your security team may or may not handle business continuity. How about disaster recovery? Often, this is seen as a responsibility for IT. I recommend you explicitly define what your security function covers and what part(s) are managed by the security team and size the team accordingly.

Regardless of what you define as in or out of scope, your security function will be bigger than your security team (see also "How much do you want to centralize security responsibilities and controls?"). If you plan on assigning the entirety of your security to one team, you are in for a rough ride. This is what

led to the saying, "Security is everybody's responsibility". A small group of people will never be able to cover everything, however good they are. This is just a reality that must be accepted.

Whatever you have defined as being in the scope of the security function but outside the security team needs at least a secondary manager or someone who can provide clear communication, consultancy, and accountability to the security leader. This is because the security leader generally will remain accountable for the entire function, regardless of who does the work (see "Which controls do you want to centralize?").

Controls

Security controls are a subset of what is generally known as "internal controls". To drop the jargon, security controls are anything you do to manage (maintain or reduce) security risks. There are all kinds of controls, such as requiring approval for risky tasks, restricting access to sensitive data, requiring the use of specific tools to protect laptops, etc. You may very well have implemented controls without really knowing it and without having anything documented. This is what's called an "undocumented control." It is difficult, however, to ensure that an undocumented control is well designed and operating as expected.

Generally speaking, every control should have an **owner** (the person who defines the control and is accountable for it) and at least one **operator** (the person or people who do what the control says). These two roles may or may not be filled by the same person.

Security people like to group security controls into categories, such as "access management", which would typically include the following:

- Access provisioning: Providing people with the access they need
- Least privileges: Not providing more access than what is required to get the job done
- Need-to-know: Not providing more information than what is required to get the job done
- Access recertification: Checking that people only have necessary access
- Access de-provisioning: Removing access for people who no longer need it

In small organizations, it's also quite common to have "**God" users** who have excessive privileges — they can create, modify, audit, delete, etc. their organization's data, apps, systems, etc. This level of access rights is dangerous and violates some fundamental principles of security. If an attacker somehow took over your "God" user account (or if that employee becomes disgruntled, corrupted, or coerced), the damage may be catastrophic. Therefore, I don't recommend it.

Security specialists should always advocate for **Should this be "defense-in-depth"?, a basic acknowledgment that things will go wrong — someway and somehow — and so it's best to have layers of security. That way, if one layer breaks down, is ineffective, or doesn't cover everything it should, you have another layer (or several) to fall back on.** We can lose this sentence to get to the advice related to security layers, which is the marrow in this bone. You likely want three layers of protection for your most sensitive assets or highest risk levels.

Checklist

You should determine	
... what security risks you face	✓
... what you want your security function to cover	✓
... who owns each risk	✓
... what threat level each risk carries	✓
... how you'll react to each threat	✓

Proton's **approach**

Our promise to the Proton community has always been that we will protect their data above all else. This is our top security priority. Now that we're a much larger organization, uptime has also gained importance, and we invest a lot of time and resources to ensure our resilience.

Because of the sensitive nature of this information, I can't provide too much detail on Proton's approach to the points in this chapter. However, I can share that we have a defined risk framework with seven categories, including business and operations disruptions and data leaks. These are broken down into 12 risk scenarios and 35 risk sources. We have seven distinct risk owners.

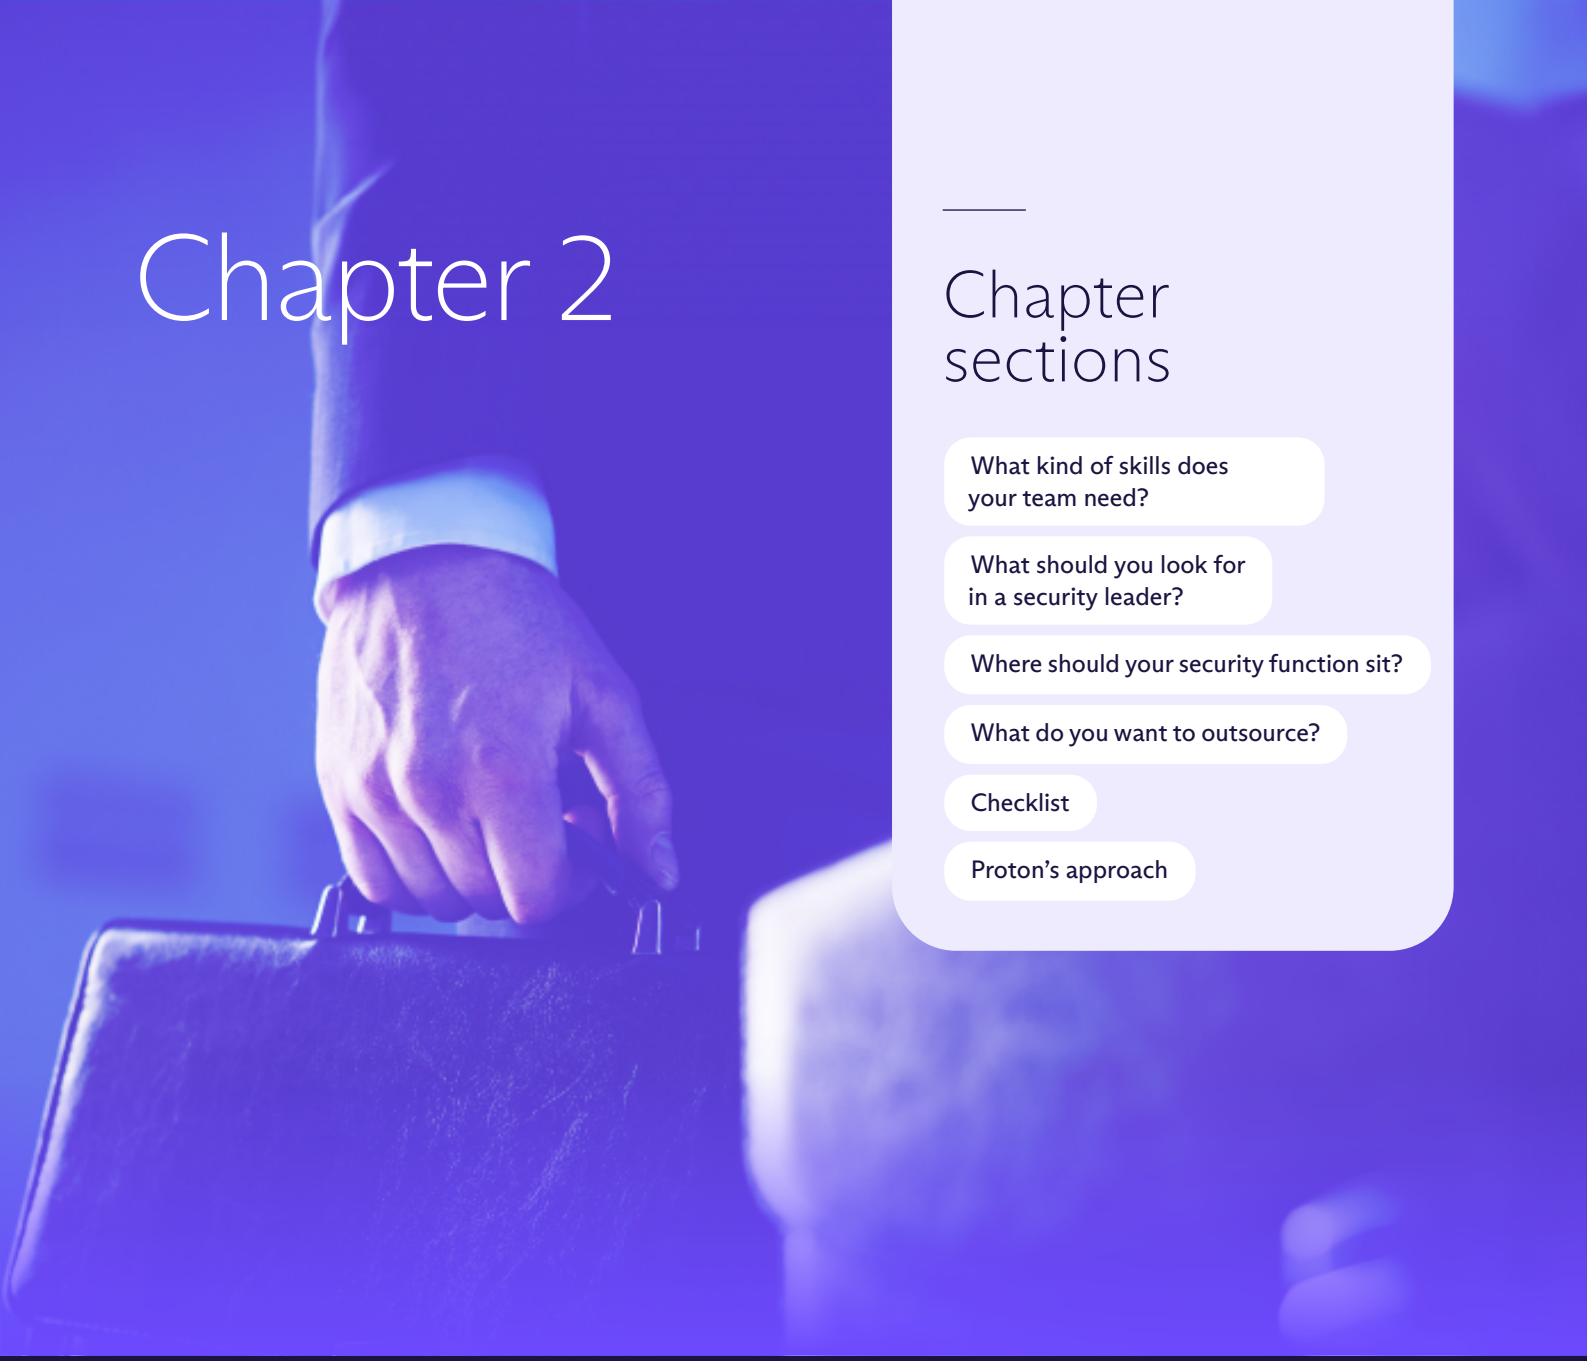

Chapter 2

Chapter sections

Building your **security team**

While your security function is dispersed across your entire organization, your security team is the command center. And even if your team is only a handful of people, you should still designate someone as the security lead. From there, you'll need to grow your team to help address the security risks you are managing from Chapter 1.

The way you structure your security team, both in terms of its personnel and where it sits in the internal organization, can have a profound impact on how it functions. In this chapter, I'll share some simple steps you can take to help ensure your security team is effective and proficient.

What kind of skills does your team need?

Even if you already have a security team, it's a good idea to think about the people who make up that team and the security function more broadly. Just like any other part of your organization, it must make good use of its limited resources, meaning its composition deserves your consideration. This section might help you identify gaps you have or give you insight into how to grow.

Like anything important, your security team should be well rounded. I mean that it shouldn't be made up entirely of the same type of people with the same or similar skills. In security, it's common to overlook the more human and managerial (sometimes referred to as "non-technical") capabilities in favor of the technical ones. I see this as a strategic mistake that can wind up costing you time and money.

So what kind of people do you need?

The main attribute you should look for is someone who has a security mindset. While the average person would see a tool, process, or event as just that, someone with a security mindset should see all the ways in which it could go wrong — whether by malice, accident or neglect. They should also be able to correctly assess which tools, processes, and events matter to security and which don't.

If you don't currently have a security team, you'll need to build one person by person. Be sure to find a good fit for the first member of your security team. They will be the person you consult with on security issues until the team grows, and, like any founder, they will set the tone for the function itself, so make sure they're on the same page as you.

It's common to have a technical person (generally an engineer) as the first member of a security team. This person typically manages (and/or builds) your security tools and is responsible for all security-related activities. Probably the next most common person to receive security responsibilities in the early stages of a company is someone from the IT team (see "Where should the security function sit?").

Technically inclined staff often specialize in a certain topic. This expertise brings great value to any team but can introduce blind spots.

Technical people tend to trust tools to solve problems over training staff or raising awareness in the company. Tools are not the solution to all security problems, and security is a broad topic that no single person can completely master (nor should that be the expectation). This is why making a single person responsible for all your security is risky, especially if they're highly specialized in one topic.

So while starting with a few engineers and/or IT staff is an understandable approach, particularly when budgets are tight, you shouldn't wait too long before you add someone with security management skills. I recommend naming a security officer before your team reaches four to five security/IT engineers. The security officer could very well be a technical person with this additional skillset or relevant experience. You just want to make sure that someone is thinking of security while also looking at the big picture and accounting for your business's goals.

Once you start mixing generalists and specialists, you need to make sure they're able and willing to work together. They should each understand the skills that the others bring to the team and want to collaborate.

You should also consider having security champions throughout the business. These are people from different teams or departments who are interested in and/or knowledgeable about security. They can act as advocates for security efforts and decentralized eyes and ears that help the security team understand the concerns other staff members have and how practices are actually implemented on a day-to-day basis.

What should you look for in a security leader?

You want your security leader to be pragmatic. Security people can be very passionate about their work, which is great, but it sometimes leads to unrealistic expectations about how the rest of the company should do things. You should make sure that, even if some team members have the tendency to set idealistic or unrealistic expectations, you have a leader who can help them focus on what is most important: Understanding how their work fits into your company's overall objectives and plans.

In the early days of a company, you'll likely want the leader to also be an individual contributor. They might be a (senior) security architect or security engineer, perhaps the person who manages your most important security systems. If you can find one, look for someone with experience in risk and project management (and loves what they do).

Keep in mind also that security requires a lot of influence, often without equivalent authority. Your security leader must have the appropriate people skills and come across as approachable, focused, and intentional.

Where should your security function sit?

This is a classic question, and we're not the first ones to discuss it. If you're already familiar with the implications of having the security team sit within one department, you can skip this section.

Let's start with our least favorite approach: Having your security team report to IT. This is just not a good idea for many reasons. In summary, the goal of your IT team is to provide your staff with the tools they need to get their work done efficiently with the best user experience possible. Your IT team needs to make everyone's life (at work) easy. Unfortunately, security practices rarely improve efficiency or your employees' user experience. This clear conflict of interest is why **your security and IT teams should be separate.** Having your security team report to IT is like taking the foremen of a construction site who have been tasked with getting a building done on time with minimal costs and also making them in charge of worker safety and workspace inspections. They'll always be tempted or pressured to bypass safety requirements to maintain their deadline and budget. That's how costly oversights and accidents happen.

Basically, if the IT leader can decide what security controls are applied to your IT systems, you've ensured security will always take a back seat.

Other options include having the security team report to Finance, Legal, or Risk (although, if you're a small business, you probably won't have a risk department yet). These are all fine and may make more or less sense depending on what your business does, how it's structured, who makes up your personnel, etc.

One of your best options is to have your security team report directly to the CEO. This should ensure that security has a direct relationship to the business and facilitate making security a business-wide priority.

Note that if you choose this route, it's essential to have at least one security team member who can make security issues understandable to an executive-level audience.

What do you want to outsource?

A common approach nowadays is to (almost entirely) outsource security. Managed security services abound and cover everything you may want and need. Typical outsourced services include Security Operation Centers (SOC), security awareness and training, security tooling integrations, and even Chief Security Officer-as-a-service. Using these services might be a no-brainer in some cases. In others, you should think hard about whether this will be a good decision in the long run.

When thinking about what to outsource, keep in mind:

1. **Legal/regulatory/brand requirements**
   You should consider not only whether you face any constraints that limit where your data can go or where your systems can be accessed from, but also what your customers would think about you outsourcing potentially sensitive activities and data. Also, many jurisdictions require organizations to inform their customers about all subprocessors with access to personal data. Even if you do everything correctly, having a list of numerous unrecognizable subprocessors can undermine your customers' trust. Just like consumers prefer foods with a small list of ingredients they can pronounce and recognize, they also prefer to use products and services with a small list of subprocessors they know and trust.
2. **Maintaining the necessary competence**
   Although it may seem appealing to have an expert set something up for you, if you rely on external expertise, how will you maintain and adapt these tools and policies as time passes? Consultancies have notoriously high turnaround rates, so having a good consultant today is no guarantee you'll have a good one in the future. When the consultant is gone, their knowledge will be too (unless you make a great effort to prevent this from happening by requiring them to document things and train others).
3. **Access**
   If you want to outsource certain activities, what kind of access will you need to give to this external entity? Will you need to open up your internal network? You should not put yourself in a position of weakening your security to outsource part of it.

Checklist

You should determine	
... what you want your security function to cover	✓
... what skillset you need to effectively protect your business	✓
... if you are ready for a security leader	✓
... who has the skills to lead your team/function	✓
... where the security team should sit within the organization	✓
... what you need to insource and what you can afford to outsource	✓

Proton's **approach**

At Proton, we hired our first full-time security engineer in our fourth year, but we've always had a strong culture of security (any attempt to set up an end-to-end encrypted email service will drill that mindset into your team). After hiring a second engineer, it was time to hire a manager. Since then, the team has steadily grown. At first, this small team reported to the CTO, and a few years later, it shifted to reporting to the CEO.

One of our main goals is to provide resilient services, and because we manage our own infrastructure, business continuity was originally the infrastructure team's responsibility. As the company grew, security has increasingly provided support, particularly on the policies and processes part. Similarly, our main concerns about physical security were related to our data centers. Therefore, while this was also originally within our infrastructure team's purview, it has now become a shared effort with security.

As you might expect, we don't outsource much of our security. The only third parties we use in security are SaaS for security awareness, security researchers (see our bug bounty program), and external security auditors.

Chapter 3

Chapter sections

Defining **security responsibilities**

As mentioned in the previous chapter, your security function encompasses much more than just your security team. This means the people who contribute to your business's security will be scattered throughout your organization (indeed, if you're doing it correctly, everyone at your organization will contribute). The question is: Whom do you want to give security responsibilities to?

You will often hear, "Security is everyone's responsibility". But in practice, this system breaks down as soon as an incident or problem arises and can quickly turn into a blame game where everyone is searching for the initial culprit. That's why I recommend you give everyone clearly defined roles from the very beginning. These roles will explain who is **responsible** and who is **accountable** for the different parts of your security. What is the difference, you ask? Read this chapter to find out.

- **Who is accountable?** This person has ultimate authority over the issue at hand. Accountability cannot be delegated and must be assigned to a single person, not a team or group. This is where the buck stops. The person accountable for your security or any subset of it makes strategic decisions, guides the direction of your security function, and must justify themselves if their decisions or actions (or lack thereof) have negative consequences. The person accountable for your security cannot be just anyone — it must be someone senior or empowered in your organization. They do not need to do all the work, but they do need to oversee the work being done and make sure it is under control.
- **Who is responsible?** This person has been given a specific task to carry out. They're expected to complete that task and to escalate issues if they encounter any. Crucially, responsibility can be delegated. Generally, the person accountable for an issue will delegate the corresponding tasks (responsibilities) to other team members.

Examples. In many organizations, large or small, the head of security is accountable for security awareness (assuming the topic is relevant). That means that she must understand what level of awareness is appropriate for the organization, allocate resources to define an awareness program and oversee its roll-out and outcomes. She will not be providing security awareness sessions herself; she may not even define the program. But she needs to make sure that one is defined and that whatever objectives were set are being met. She will delegate the responsibilities of defining the program, providing the information and reporting results to someone else.

Now, say the organization suffers a massive security incident which, after investigation, is linked to a significant lack of awareness. The head of security is accountable; she will need to explain what went wrong, looking at why she believed her program was appropriate and effective and where it failed.

It's a mistake to think that responsibility and accountability are interchangeable concepts. When discussing how to structure teams and assign tasks, they are not, and everyone on your team should understand the difference.

When it comes to security, it's a matter of when, not if, an issue will arise. If your team doesn't have a clear understanding of who is accountable for issues and who is responsible for carrying out tasks, you make the following outcomes much more likely:

- Nobody takes ownership, meaning nothing gets done, which is not good.
- Someone takes ownership, but they don't have the necessary authority or competencies to see it through. This makes it more likely that efforts will stall or mistakes will be made, which is also not good.

You can avoid this situation entirely by clearly outlining who on your team is accountable for your security and who is responsible for which tasks.

Questions **you should ask yourself**

Who can make decisions related to security?

You might wonder what this question even means. Isn't the security leader in charge of all things security? That may be the case if your company has only a handful of people. However, most organizations will have many people responsible for different bits and pieces of your company's security function (see "What is in the scope of your security function?"). So you need to ask yourself: Who can make decisions about what? Remember, if someone doesn't have the authority to make decisions regarding a risk, they cannot be held accountable for that risk.

A related question is how much do you want to centralize your security authority into one person or team? There's no cut-and-dry correct answer, as each option has strengths and weaknesses. Centralizing accountability will make it significantly easier to administer your security function. However, it has the potential (especially if not done well) to slow down operations. On the other hand, decentralized accountability can lead to more flexibility and agility at the expense of cohesiveness, visibility, and appreciation of the bigger picture.

Regardless of the route you choose, whoever you empower as accountable for your security must have the authority to make decisions and the budget to implement practices regarding your security function.

You can't avoid these decisions. If you choose not to decide, you've simply delegated your decision to someone else. Nature abhors a vacuum.

How much do you want to share about security?

This question goes hand-in-hand with the question of how much you want to centralize security authority. What you want to share about your security practices will depend heavily on your company culture. Some people will accept security controls without fuss, while others may be more reluctant. I generally recommend being as transparent as possible without compromising your security. This applies both to sharing information internally within the company and with external stakeholders.

If you manage a smaller organization, you may not be ready to share much with the outside world. However, you may be able to share quite a lot with your staff, and I recommend you do so for several reasons:

1. **Effective security requires buy-in from all your employees.** Security is often seen as a burden or just a box-ticking exercise. By being open about what you are doing and why, you should alleviate some of the skepticism and reluctance you may face. Remember, security is a defensive position: it's only as strong as your weakest link (and this could wind up being a process, tool, or person).
2. **People are more likely to follow processes and requirements if they understand why they exist.** Your team is likely a group of smart, motivated individuals (otherwise, why would you have recruited them?), so let them ask questions and understand why certain security precautions are necessary. People are much more likely to take their security responsibilities seriously if they know there are real risks you're trying to mitigate.
3. **Receive feedback from the people who actually do the work.** This comes back to the first point. Your security relies upon your staff, and if you never discuss why you made the security decisions you did, you're unlikely to get timely and useful feedback. You want to know immediately if a certain tool or process is so onerous that people will avoid it because that is a security incident waiting to happen.

While the security function's goals will inevitably put restrictions on what staff can do and how they go about their work, this openness shows that you understand the role everyone plays in keeping your organization safe and that without their buy-in, most controls will be ineffective.

Do you want to enforce or check your policies and controls?

Once you have defined your security policies (see Chapter 1), you should ask yourself how you will ensure everyone follows them. As a leader, you remain accountable for this (see "Defining security responsibilities"), and — if your policies were well defined — you should truly care about them.

So you have two choices:

- Ensure your policy requirements are followed by forcing your staff to adopt desired practices and avoid unwanted practices (for example, you force multi-factor authentication on sensitive applications and block USB ports on staff's laptops).
- Trust your team to follow the specified requirements and perform regular checks to ensure they are.

The former is easier in many ways: You design centralized controls that only allow options that follow your policy requirements, and you're done. The cost might be that your employees don't agree with the controls in place and try to work around them or are unsure what to do when they encounter issues not covered by your controls.

In the latter, you give your staff much more freedom to meet your policy requirements as they see fit. The issue is that they can choose to not follow your policy requirements at all, and it requires quite a bit of work to follow up on their behavior, especially as you grow. This also requires you to educate your staff on how to perform certain activities.

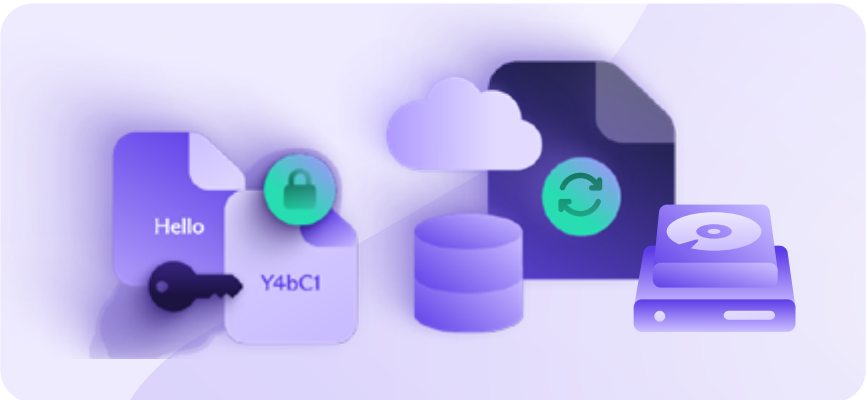

Checklist

You should determine	
... what security roles and responsibilities need to be centralized	✓
... how much you can and want to share about your internal security practices	✓
... whether you will enforce security controls or check that they are independently implemented	✓

Proton's **approach**

At Proton, we've taken a dual approach to security responsibilities. Some security responsibilities are centralized within the security team (such as security monitoring and risk management), while others are distributed across the business (such as vulnerability management and secure development). As we grow, the security team continues to mature and provide more centralized services. However, while we have decentralized some of the security responsibilities, the head of our security team is ultimately accountable for Proton's security function.

We have made a particular effort to ensure our staff understands the changes we've implemented along the way. We encourage everyone, especially those who feel strongly about certain topics, to ask questions and challenge decisions. Communicating about new controls and additional responsibilities taken by the security team has been instrumental in getting people on board.

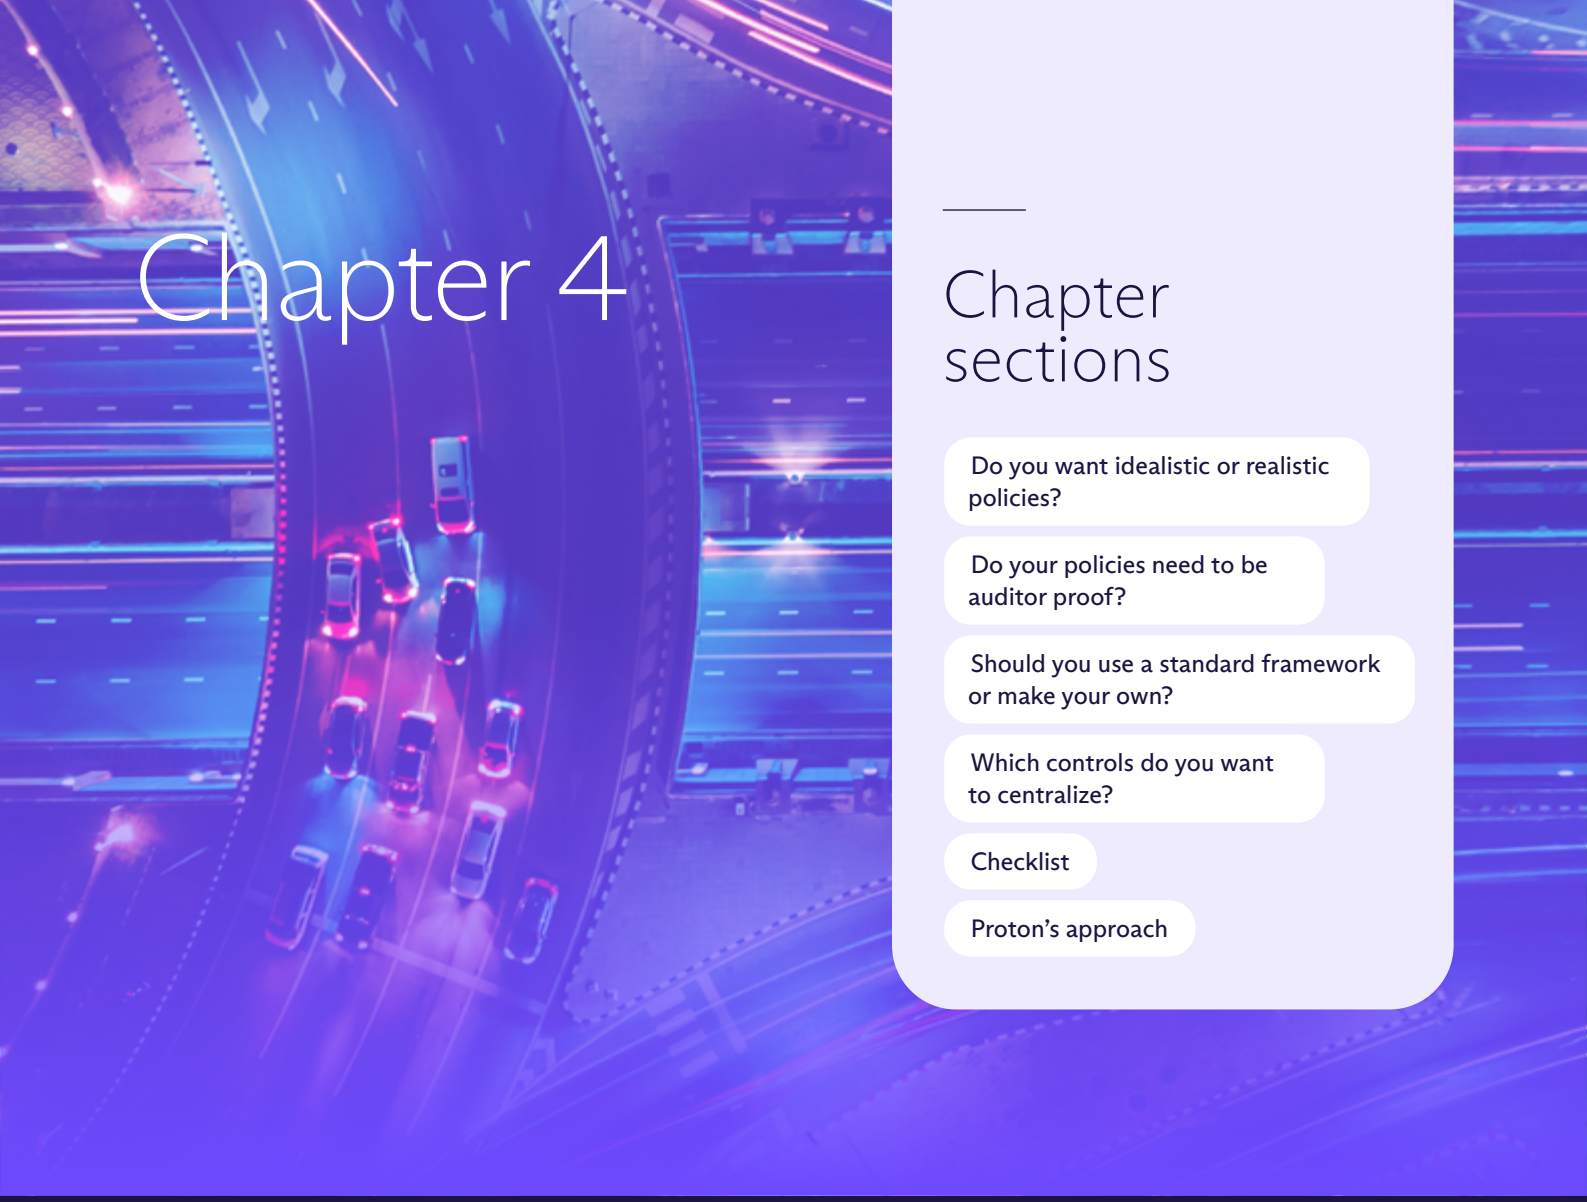

Chapter 4

Chapter sections

Bringing structure **to your security policies and controls**

Organic growth is common in many aspects of most startups, and security often follows the same pattern. This may be appropriate for a while, but chances are it won't be scalable or sustainable.

As much as people love to hate them, frameworks can be really helpful. Even if you don't want to follow them to the letter, they can help you check that you haven't forgotten anything important and can provide you with ideas about how you want to structure things. While you should constantly be re-evaluating your security risks, I encourage you to minimize the significant changes you make to how you structure your organization's security policies and controls. This is why the questions in this chapter are so important. The answers may seem trivial today, but they likely won't be in a year or three.

Questions **you should ask yourself**

Do you want idealistic or realistic policies?

There are two main approaches to defining what goes into a security policy. The first is, "This is how we want to do security". The second is, "This is how we actually do security". And make no mistake — these are completely different schools of thought. The advantage of the former is that it sets a strong tone. It shows a desired state that everyone should strive to achieve. The downside, however, is that it tends to convey the message that what is written is not mandatory right now. And unless you make the way forward clear, your staff may think they will never be.

The opposite is true for the latter approach. If you only write down requirements that are to be followed as of today, it should be easy for everyone to understand and comply. However, by doing so, your policies will tend to be more lenient than you would like them to be and require updates each time you make an improvement.

Mixing achievable requirements with aspirational ones may also be confusing (unless you explicitly mention which ones you're building towards), and the last thing you want is people to think that all requirements in those documents are stretch goals to be attained at a later date.

Again, there is no right or wrong — I simply recommend you think about which one works for your business and consider how to address the drawbacks of the chosen approach.

Do your policies need to be auditor-proof?

You can try to dodge auditors but dealing with them is increasingly a requirement for businesses, be it to obtain a certification, meet regulatory requirements, or fulfill contractual obligations. In any case, if you think you will have an auditor looking into your security any time soon, you should consider preparing your security policies accordingly (which means having them documented by your security team and reviewed and approved by top management).

If you run a very small business, writing down the rules may seem like a poor use of time, but auditors will want to see them. If you can (or want to) wait until your company is bigger (and thus unwritten rules are no longer viable), you can do so. My guidance would be that as long as your company is small enough that everyone knows each other, you can probably wait to make your security policies explicit — unless an auditor is already knocking at your door.

Should you use a standard controls framework or make your own?

You'll need to define security controls at some point. This may be to prepare for an external audit, obtain a security certification, or simply improve your internal visibility and operations. If you're starting from scratch (or looking to significantly alter your current approach), you'll probably be encouraged to look to well-established frameworks for inspiration.

This can be a good idea, although I recommend you regard them strictly as starting points and avoid bending over backward to make your business fit into their model. Instead, you should adapt these frameworks to suit your organization and how it operates.

Some common frameworks are:

- ISO 27002 (AKA ISO 27001 Annex A): This framework is hotly contested by security professionals. Critics (including myself) claim it focuses too much on certain areas while almost completely ignoring others. The 2022 version contains 93 controls grouped into 4 categories.
- NIST CSF (AKA The US National Institute of Standards and Technology's Cybersecurity Framework): This framework isn't as specific and doesn't give as detailed descriptions as ISO 27002. However, it is slightly more comprehensive, with 108 controls grouped into 23 categories and 5 functions. Version 2.0 of the CSF was released in February 2024, its first major update since its creation in 2014.
- SOC 2 (AKA System and Organization Controls, as defined

by the American Institute of Certified Public Accountants — no joke): SOC 2 is itself based on another framework known as COSO. It covers controls for security, availability, processing integrity, confidentiality, and privacy. It was last updated in 2022. There is a lot to be said about SOC 2, but this is not the place for that. Just know that, if you're at the very early stages of your security journey, SOC 2 may be a bit much for you.
- CIS Controls (AKA The Center for Internet Security's Critical Security Controls): This is a list of 18 controls that are seen by the security community as being priorities for organizations to protect themselves (and if you read the beginning of this guide, this should sound wrong to you). Version 8 was released in 2021.

Each of these frameworks has its own set of controls and they are all designed to cover general best practices that could be applied to any organization in any industry. As such, they will all be mediocre for your specific business, and following any one of them (or any other public framework) exactly is never something I would recommend. This dogmatic approach will leave you under-covered in some areas and over-covered in others, which is a bad use of scarce resources no matter how you look at it.

I know the idea of creating your own controls framework seems intimidating — that's why many organizations opt to follow one of the ones listed above. While they contain many good ideas, I recommend that you remember precisely what you identified as your security objectives and your risk management decisions (see Chapter 1). Focusing on controls that contribute to your specific circumstances is a much better use of your time than trying to copy a set of generic controls designed to apply to companies of all sizes and industries.

Another consideration in making this decision comes back to audits. If you're planning on having your security controls audited (either because you have to or because you want to — see "Having something to show your customers"), you'll likely need to justify your framework. If you use one of the ones above, this will be much easier. Unfortunately, if you create a custom framework, this audit may become more difficult (although not at all impossible). So if you go down that road, make sure you have someone at the wheel who knows what they are talking about and can defend your position to an auditor (who will likely prefer strict adherence to one of these best-practice frameworks) and explain why your custom framework is appropriate for you.

Which controls do you want to centralize?

This question is really about the tradeoff between giving your team flexibility to do their work and giving them control over the risks they take — the more flexibility you give them, the more risk they'll be accountable for. Centralized controls will give you the best visibility and control. This path is often chosen by organizations that face high risks or operate in highly regulated industries. It's also useful in organizations where staff aren't expected to spend much time on security or have the knowledge or skills to do so. With more disliked controls (such as restricting access to resources, monitoring for suspicious activities, etc.), centralization is likely your only option.

Decentralized controls have the advantage of giving the operators of the control some flexibility in how, when, and indeed if those controls are actually exercised. Most organizations still in their early stages opt for decentralized controls as they seem easier to set up at the time.

It may be possible to start with a decentralized control and centralize it later. But remember that how you define and implement a security control at the start is likely to stick around for a long time — there is rarely something as permanent as a temporary solution. And keep in mind that what may seem like a reasonable request today while your team is five people strong could turn out to be a nightmare when your organization has grown to a team of 500.

Checklist

You should determine	
... what you are protecting and from whom	✓
... if your security policies are strict rules or targets to strive towards	✓
... if you are writing your policies for your staff only or if an auditor will be looking at them	✓
... the extent to which you want to use an existing controls framework	✓
... which controls you want to centralize	✓

Proton's **approach**

At Proton, we've opted for realistic security policies. It was a topic of discussion within the security team, and ultimately, we decided this would be the best approach for us. This means that all the requirements in our security policies are "doable" and we update them regularly to reflect changes related to improvements we make on a regular basis.

We also feel strongly that using a generic security controls framework is not in our best interest, so we created our own that more accurately reflects our work, our priorities, and our culture. Many controls are currently decentralized, though we're working on centralizing some of them. For example, we have centralized security monitoring, risk management, and incident management, and we're working towards centralizing vulnerability management and access management. As with any organization, we need to find the right balance between assurance and speed, which is something we constantly revisit.

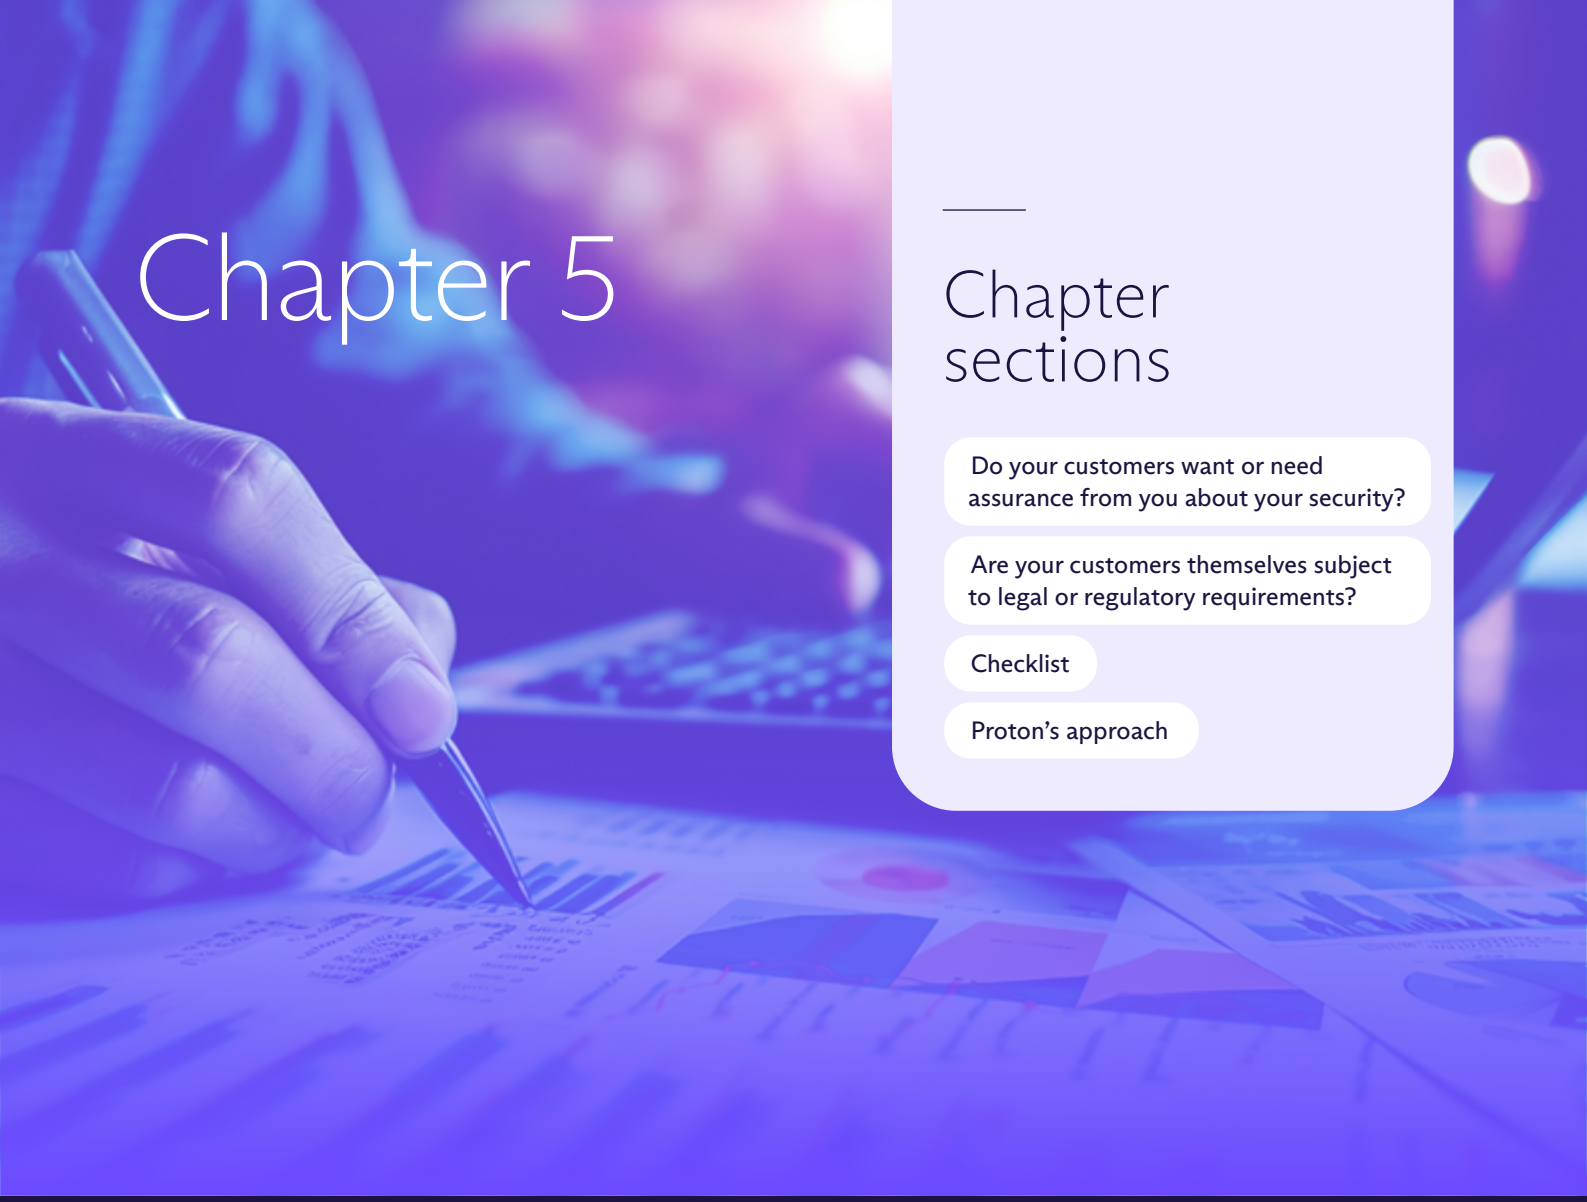

Chapter 5

Chapter sections

Providing your customers **with security assurance**

Depending on what your business is and who you sell to, you may already know that you need to provide your customers or partners with some assurances regarding your security practices. In many cases, service providers themselves are not regulated or subject to legal or regulatory requirements, but they want to cater to organizations that are. If your company processes electronic health data, you'll be subject to certain regulations (for example, there's the US federal regulation Health Insurance Portability and Accountability Act, also known as HIPAA) that you must follow or you'll risk breaking the law. That's fairly straightforward. But keeping the example of HIPAA, what are your obligations if your organization doesn't process any health data but wants to sell services or products to companies that do (and are therefore regulated)? That business relationship may hinge on your ability to demonstrate to your customers that they won't risk breaking the law by entering into a contract with you.

Additionally, and independently of any jurisdiction and industry, if your company processes any sensitive data or has access to critical systems, obtaining some form of assurance as to how you manage your security will still likely be a good investment, even if regulations don't strictly require it.

Questions **you should ask yourself**

Do your customers want or need assurance from you about your security?

If your business handles data of any kind, particularly personal or otherwise sensitive data, chances are your customers will want to know you're taking security seriously. Of course, you can respond to questions ad hoc, but it's generally preferable to be proactive and answer your customers' questions before they ask (or even think about asking).

Companies commonly show their commitment to security by obtaining an ISO 27001 certification. This is a certification of an organization's "information security management system" or ISMS. To be certified, an organization must demonstrate to an external auditor that they manage security according to the principles of the ISO 27001 standard. It's very important to understand that this certification is about "management" (as its name should hint to). It's not a prescriptive standard that tells you how complex your passwords must be. In fact, it doesn't tell you anything about the controls you must have other than that they must be appropriate for your organization.

ISO 27001 is especially popular in Europe, so keep that in mind if you want to obtain customers there. Also, if you are looking at this certification, make sure you go for the latest version (currently 2022).

Another popular option is SOC 2, especially in North America. SOC 2 is very different from ISO 27001 and comes in various flavors: type I, type II, and type III. In short, type I can (and should) pretty much be ignored as it only looks at whether security controls are well designed at a point in time (it's essentially a paper-based exercise). Type III is a public version of type II (it's a document you can make publicly available on your website). Type II reports are the ones that deliver real value to a company's customers. It gives assurance that your security controls are well designed and operate effectively.

It's important to note that SOC 2 is not a certification but an audit report. In short, you describe your security controls and how you implemented them, and an external auditor will come and check that you operate them as described. Therefore you should be aware that by going down this road, you will inform everyone who receives your report about your security controls.

Note that it's never too early to go through the process of obtaining a security certification or audit report (assuming they will help your business). For ISO 27001 specifically, the cost of certification depends on the number of people in scope and the complexity of the relevant processes. If you're still a small organization and don't have complex processes, the cost will be minimal, and you'll learn about the certification process, how to find and work with auditors, and what's required to achieve and maintain the certification (and, just as importantly, what's not required). This is very helpful as it can help you shape your processes and organization so that it's easier (or at least not harder) to maintain certification in the long run.

For SOC 2, you should keep in mind that getting a type II report operates on a rather fixed timeline. That's because you must have a first audit (type I) to assess the design of your controls. Then, you need to operate these controls for several months. Only then can you start your second audit to obtain your type II report.

And remember, whether you choose ISO 27001 or SOC 2, once you're certified, you'll have the pleasure of an auditor visit every year.

Are your customers themselves subject to legal or regulatory requirements?

It may very well be that while your organization itself is not regulated, your customers are. In many cases, their requirements will place limitations (or a strong bias) on the companies and products they can work with. By understanding these requirements and limitations, you can ensure you don't lose potential customers. Typically, you must prove their use of your business won't hurt their compliance and provide them with the information they may need for their own audits. Healthcare records may be the best illustration of this. In the US, there's HIPAA, which originally regulated organizations that process sensitive patient health information. Basically, if you were a healthcare provider, you had to protect your patients' health information. A few years later, these responsibilities were extended to "business associates" (any person or organization using or disclosing individually identifiable health information to perform or provide functions, activities, or services for a covered entity — a basic example of

a "function, activity, or service" would be data analysis).

This means that if you sell online storage, data analysis, or any other data-reliant service and want actors in the health sector as customers, you need to comply with their security requirements. And, just as importantly, your ability to demonstrate your compliance will make it much easier for them to buy your services.

Unfortunately, not all regulations have corresponding certifications. For example, you cannot officially demonstrate compliance with HIPAA or the EU's GDPR (General Data Protection Regulation) with a certification. But even when this is the case, you can document your controls and practices that support your customers' compliance — and having a security certification or audit report can definitely help.

Checklist

You should determine	
... if you will need to provide assurance to external stakeholders	✓
... if your customers are subject to legal or contractual requirements that affect you	✓

Proton's **approach**

Proton started as a B2C company. As we've grown over time, companies have become increasingly interested in our services. And because of our privacy-by-default business model, we receive interest from companies across industries, including highly regulated ones. As such, we are working hard to support our B2B customers in their own security and compliance efforts. This is why we've invested (and continue to invest) in different ways to provide security assurances to all of our users, and especially to our business customers.

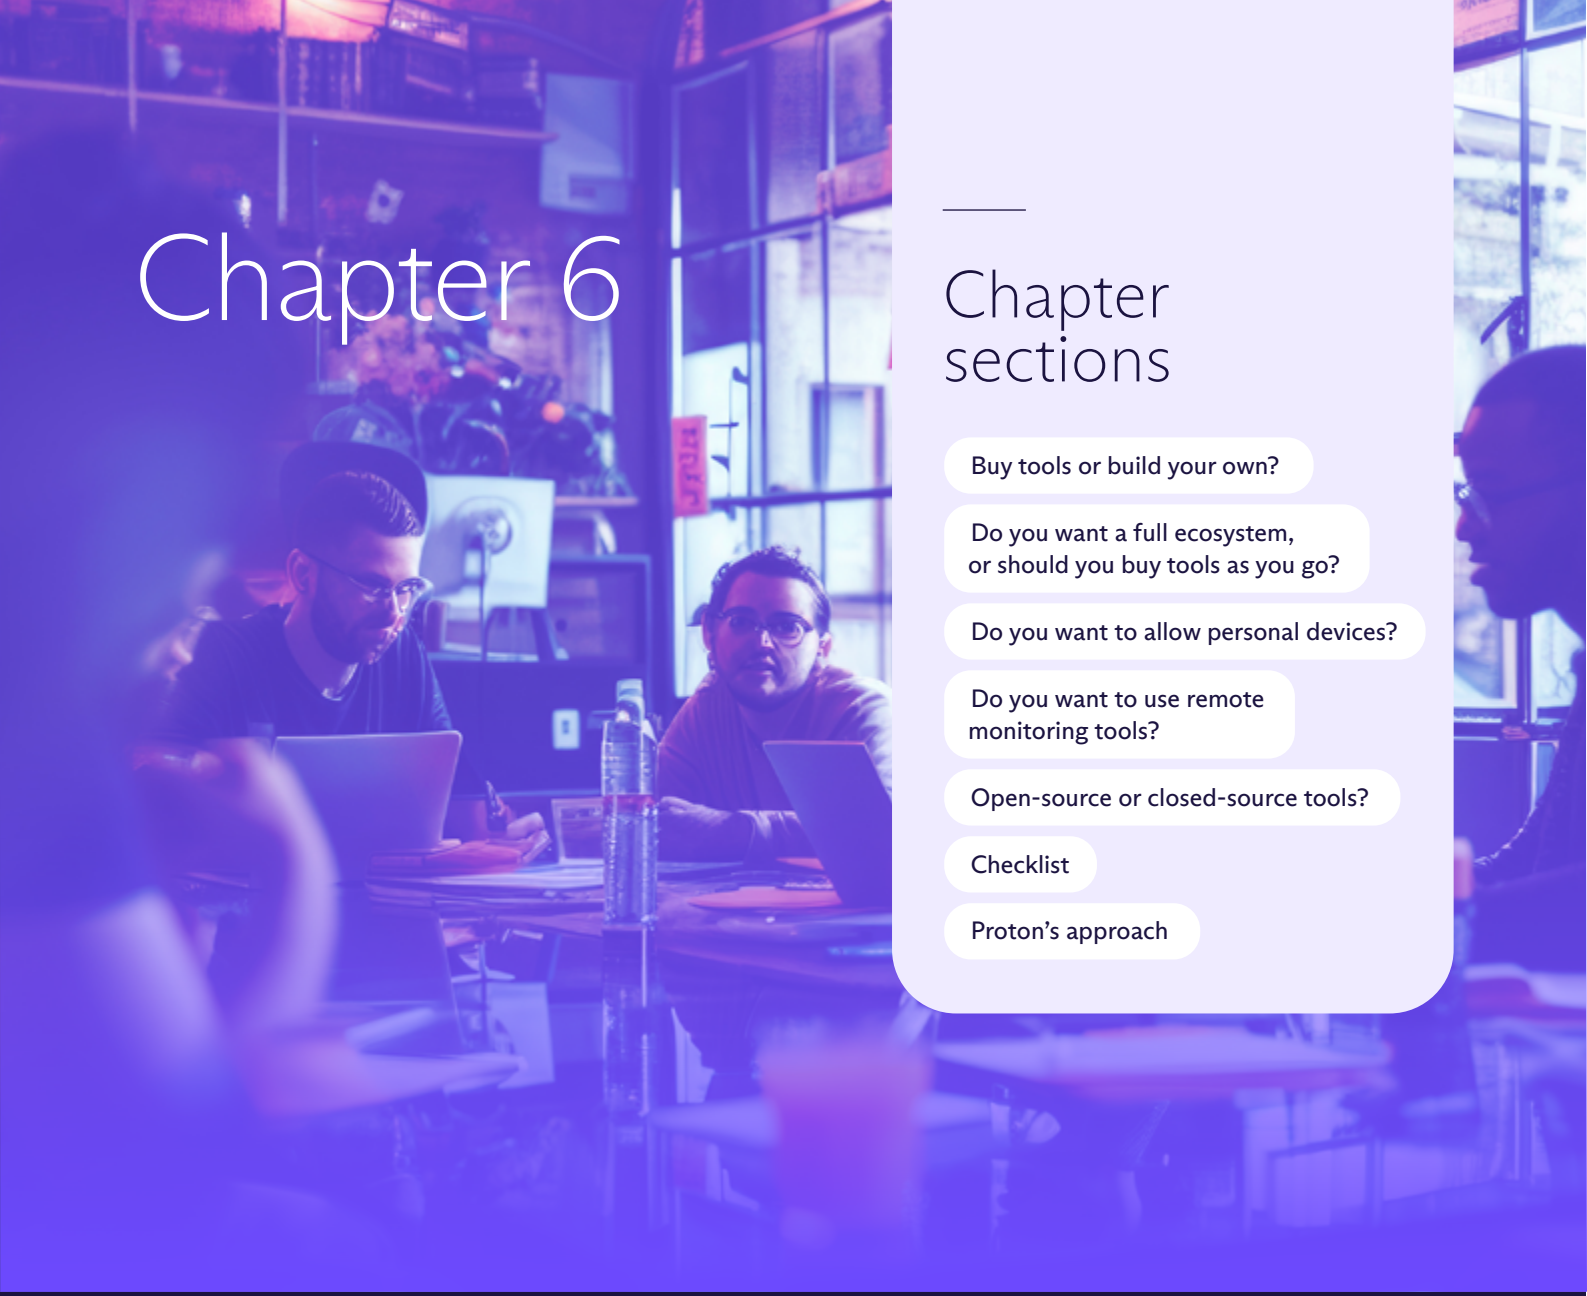

Chapter 6

Chapter sections

How to select **tools that are right for you**

You will eventually need to invest in tools (if you haven't already) both to run and secure your business. To go back to systems thinking, you cannot simply apply your security framework to your security tools. You must consider all your tools and how they interact. Therefore, you should apply the questions below (when relevant) to any tool you use, regardless of whether its main goal is security.

If you work with third parties (see "What do you want to outsource?"), they'll likely be technology providers or partners with technology providers. In any case, make sure you have a good understanding of what you're already paying for. It's not uncommon to double or even triple-pay for tools included in multiple contracts. It's also common for teams to evaluate security tools (like an access control management or an endpoint monitoring system) according to your security priorities but to forget to subject tools whose primary function isn't security (such as accounting, customer relationship management, and HR software) to the same scrutiny. If you have a good security team (see "What kind of skills do you need?"), they should be able to identify the risks associated with any tool relatively easily.

Questions **you should ask yourself**

Buy tools or build your own?

If you are in the tech space or have very specific needs, you may consider building most of your own tools. If not, you will likely turn to off-the-shelf products. And, as always, there are pros and cons to both approaches.

When building tools from scratch (or using open-source components), you can focus on what you're really interested in and avoid paying for extra features you don't want or need. Your in-house tools can be as unique as your business is. This also means that you will have to explicitly and intentionally invest some resources in security functionalities. You might be able to defer this, but you cannot avoid it. Sooner or later, your in-house tools will need security functionalities (in other words, functionalities that are "only" there for security reasons, like a customer support tool that can enforce multi-factor authentication), and implementing them will take time and/or money.

New vulnerabilities will also appear from time to time, and even if you build your own tool, you will surely have external dependencies (for example, on operating systems or libraries) that require updating. These updates will require work, and when the updates break existing features, they'll require even more work. So when is the best time (or least bad time) to make this investment? Remember that change is always more difficult than we'd like, and retrofitting security features into existing technology and processes often doesn't work well. Of course, the need for security functionalities and the timeline will depend on how fast you plan on growing and in which direction(s). In comparison, commercial products should have all of these features baked into them, ready to use (or not).

There's a lot to consider here, but you can boil it all down to this general rule: **If the tool is an integral part of your business, is unique to your business, or can be sold to your customers, build it. Otherwise, buy it.** Of course, it's never quite that simple, and there will always be nuances to consider. If you have another form of return on investment that turns out to be positive (for example, if maintenance costs are prohibitive) and/or a risk rating that is extremely high, the effort to build your own tool may be worth it as well. In general, you should assume that someone with a dedicated product and a large enough customer base will be able to build tools and maintain them much more efficiently than you could (economies of scale, efficient markets, and all that). But an external product has to support a plethora of use cases, while you might just need one very specific function — this is the scenario in which it may actually be more effective, all things considered, to build your own solution.

Also keep in mind that some commercial products can be hosted on-premises (in your own data centers versus in a third-party cloud). This can be a good option for security or otherwise sensitive tooling if you don't want to build it but also don't want it to ever leave your control. Just remember that, in this case, the vendor won't be able to push updates or patches for you. You'll have to do that yourself.

Do you want a full ecosystem, or should you buy tools as you go?

If you're going for the "buy" approach described above, you could consider going for an ecosystem. By this, I mean a family of tools that work together.

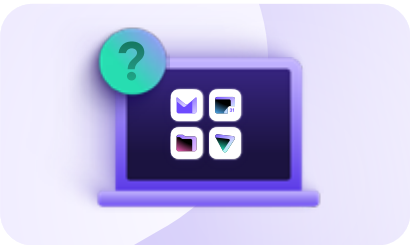

This can be a great approach as you generally get more functionality for less money than if you were to buy each component separately. However, it also means that you may be paying for functionalities you don't want or need.

Going down the ecosystem route also brings a higher risk of vendor lock-in. If you rely on a single provider for all your tools, you'll be forced to think long and hard before switching to another provider, even if it offers a better deal. And make no mistake, ecosystem providers know this. Whether they always act on that knowledge, I cannot say for sure.

If you prefer a pay-as-you-go approach, you must think about interoperability. Can all your tools work (seamlessly) together? Or will configuring them be a pain that your team will find clever ways to avoid? And again, make sure you know what you are paying for and that you're not needlessly paying for duplicates, as discussed above.

Do you want to allow the use of personal devices?

Using personal devices for work (or BYOD or "bring your own device") has become common in many companies. Leaders see it as a way to reduce hardware costs and ensure staff use devices they like (not to mention have with them at all times...).

With the prevalence of BYOD, the concept of zero trust has been gaining momentum. This is the idea that you should be able to use any device to do your work and that whenever a device connects to your network, systems, applications, etc., there will be a check that ensures it's safe to do so (for example, a scan to ensure the device is not compromised or a login to ensure the user has the appropriate authorizations). While most discussions of zero trust focus on the device's security, limiting network access to team members with the correct authorization is just as important. Simply letting your staff access any data with any device without any checks is not only not zero-trust — it's bad security.

Similar to what was discussed in "Which controls do you want to centralize?", you need to think about whether you want to have some "central" control over the devices your staff uses or decentralize checks to individual systems, networks, and applications. With zero trust, you will have some decentralization by definition, as each connection to your system, network, or application will require a check.

If you are considering BYOD, don't forget about mobile phones. Many applications have mobile versions, and your staff may want to install them on their phones (both personal and professional). So, even if you are strictly against BYOD for laptops, you may still have people who want to install work-related applications (like emails, instant messaging, or file sharing) on their phones. You can have amazing security on laptops but still have big problems if you forget that people use their phones to do much of the same work as they do on their laptops. The point is this: **Don't forget about — or ignore — phones.**

Finally, you may find it helpful to give your staff extensive training, documentation, and examples about what you consider "acceptable" and/or "unacceptable" use of their business devices. If you have very tight control over what can be done on devices, this may not be as necessary as it is if you don't. Regardless, consider that with the diversity of your workforce comes an even greater diversity of risk appetites, knowledge, skills, and levels of attention. Remember, people, including your team, can and will do things with their devices that you could never have anticipated. Your security plan and function should be set up to weather, or at least detect and mitigate, some of these outlandish scenarios.

Do you want to use remote monitoring tools?

Finally, think about what visibility you want (or need) of what your staff's devices are doing. Nowadays, plenty of tools with three-letter acronym names (MDM or "mobile device management", EDR or "endpoint detection and response", XDR, which is just fancy detection and response) can be installed on laptops to enforce secure configurations and/or monitor for suspicious activity. These tools can be invaluable from a security perspective and are easy to implement on all devices, company or personal. Whether or not your staff will be eager to have company-controlled software on their personal device is another story...

If you are considering these tools, you should make clear to your team why they're needed and what they'll be used for. Some people hate the idea of these tools. But just like any tool, it can be used or abused. Think of CCTV cameras as an example: They are generally used by companies to monitor who enters and exits their offices and other potentially sensitive areas. This information can be vital in the event of a breach. However, a determined employer could also use those same cameras to track when employees arrive and leave, whether they take smoking breaks, who they walk in and out with, etc. Most people would find nothing objectionable about the first use case and would be (rightfully) creeped out by the second. But the CCTV system, the tool, is the same in both examples — it comes down to how you apply the tools.

The point is this: **Your staff's perception must be a consideration, along with your risks and security controls.** How sensitive are the systems and data you're protecting? What kind of threat actors are you facing? How do you monitor for strange behavior on your network(s), application(s), etc.? You may not need monitoring tools, but if you decide you do, I suggest you pay particular attention to how you will govern the use and configuration of the tooling and the access rights to the data it generates.

Open-source or closed-source tools?

Another classic security question is whether open or closed-source products are "better". It's a question that even appears in some security interviews. The answer is (again): It depends. Note that in this section, I'm referring to tools used within your business, not the tools your business provides to its customers.

Open-source software makes its source code publicly available so that anyone can inspect, modify, or enhance it. While there are many technicalities to layer on top of this definition (like whether the license terms allow you to download, modify, or publish modified versions), this is the basic premise of open source. Closed-source software is any program that does not make its source code publicly available. As a basic example, think Linux (open source) vs. Windows (closed source).

Open source is as much of a philosophy as it is a type of software. Any definitive, universal answer about whether it's better or worse than closed source, in one direction or the other, should be greeted with skepticism. And in reality, this is the wrong question. Instead, you should ask yourself which tools are best suited for your business and specific use case (this does require you to have a somewhat defined use case and a corresponding list of functionalities and requirements, which is easier said than done).

Open-source tools are great if you're building your own tooling as you can customize as much as you need. It's also a good approach if you have more engineers than cash. Some argue that open-source tools are more secure as they can be reviewed and tested by anyone at any time. While this is true in theory, in practice, you should ask yourself if the specific code you're considering is of enough interest to a broad community to take advantage of this characteristic. If nobody else uses that piece of open-source code, chances are nobody will check it for security issues. If not (for example, if you build a limited-scale tool tailored to specific incidents your business faces and doesn't have much applicability elsewhere), I'd argue this justification falls apart. If you want to be able to check it yourself, then of course open source is great. If you're hoping someone else will check it for you, it's a different story.

Another aspect to consider with open source is whether there's any support available. If you are going to rely heavily on an open-source project, you might want to look for ones that offer paid support programs (they're relatively common). Even if you don't need that right now, you might in a year or two.

Checklist

You should determine	
... which security tools you want to build and which you want to buy	✓
... if you want to use a full ecosystem or buy independent tools	✓
... if you want to allow personal devices and what this means for your security	✓
... which tools you want to have as open or closed-source tools	✓

Proton's **approach**

It should come as no surprise that Proton extensively leverages open-source technologies. We've custom-built many of our tools to meet our specific needs, such as our security incident and event monitoring (SIEM) and anti-abuse systems. However, we do use commercial products in areas where open-source options lack the features we need, and the tool isn't central to our business. This is the case for our customer support ticketing and payment processing (see our Privacy Policy). This also means we have selected the buy-as-you-go approach rather than having all our eggs in the same basket.

Planning **for growth**

The priorities of growing your business and securing it can sometimes feel at odds with each other, but this doesn't need to be the case. There will, of course, be tradeoffs that you need to make and risks you need to take. Some security controls may seem cumbersome today but may become crucial in a few years. The real challenge often lies in knowing when to pull the trigger. If I can give you one piece of advice, let it be this: **Consider the cost of change.**

It may not seem like a good use of resources to set up processes to manage access across your business today, but trying to change once you have hundreds of employees across many teams, all of whom have established their own processes independently of one another, will be exponentially more difficult and time consuming. The same goes for any other security control, and the more restrictive, cumbersome, or annoying the control, the more difficult it will be to change.

This is why I say that there is nothing more permanent than a temporary solution. And there is nothing worse than temporary solutions, except for when they become permanent.

Conclusion

It may seem a bit overwhelming to think about all of these questions and how you want to answer each. The point of this book is not to confuse but rather to show how all of these elements are interlinked. This may seem like it makes things more complicated, but it's fairly straightforward if you can name what you need to protect. If you don't know where to start, I recommend you do the following:

1. Think about why (and how much) you care about security.
2. Determine what you want to protect and what you can let slide. Remember, there can be no overlaps and no gaps.
3. Define your security philosophy: How much control, visibility, and assurance do you want?
4. Go through each of the questions in the previous chapters of this book and answer each one, keeping in mind these three points. It should all make sense (if it doesn't, re-think these three points in the conclusion and adjust if needed).

Creating a security function is unavoidable — your business's security will impact its performance even (and perhaps especially) if you choose to ignore it.

Let's take a minute to review through the entire process, as described in Chapter 1, taking the example of a fictitious startup that provides software-as-a-service to its customers in a wide range of industries.

Identify priorities

This startup's first priority is the resilience of their service. This means its security team could have a risk category called "business disruption".

Analyze risks

Now let's assume that a relevant risk scenario for this company is that, when coding a change, an engineer makes a mistake that is pushed to production creating a (self-inflicted) denial of service (DoS). Going one level deeper, the security team identifies two risk sources associated with this scenario. The first is that most engineers haven't been trained on secure coding. The second risk source is that code changes can be pushed to production without a peer review.

Assign risk ownership

The head of the HR department owns the first risk (lack of secure coding training), as she is in charge of the training program for all teams. The likelihood of business being disrupted due to an engineer's lack of training is rated as "Medium", and the impact if it were to materialize is rated as "High". Thus, the overall risk level is "High" (according to the risk matrix this company has defined).

The head of Engineering owns the second risk (lack of peer review), as she sets all the rules related to software development. This time, the likelihood of a self-inflicted DoS due to a lack of peer review is rated as "Low", and the impact of the services being unavailable as well as the consequences if that were to happen is rated as "High". Thus, the resulting risk level is "Medium".

Implement controls

The head of HR decides that she wants to mitigate the risk of outages due to coding mistakes by reducing their likelihood. As a result, she updates the training program to include mandatory secure coding training for all new joiners. This is an administrative and preventative security control. She also works with the head of Engineering to select and integrate security scanning tools into the development pipelines. This is a technical and preventative security control.

The head of Engineering is more concerned with the speed of delivery than she is with the lack of peer reviews. She therefore decides to accept the risk. To ensure the situation does not get worse, she keeps track of any disruptions that could have been caught through a peer review as well as any changes to the complexity of the engineering tasks that may increase the likelihood of a self-inflicted DoS. This is an administrative and detective security control.

See? It wasn't all that bad.

Appendix:

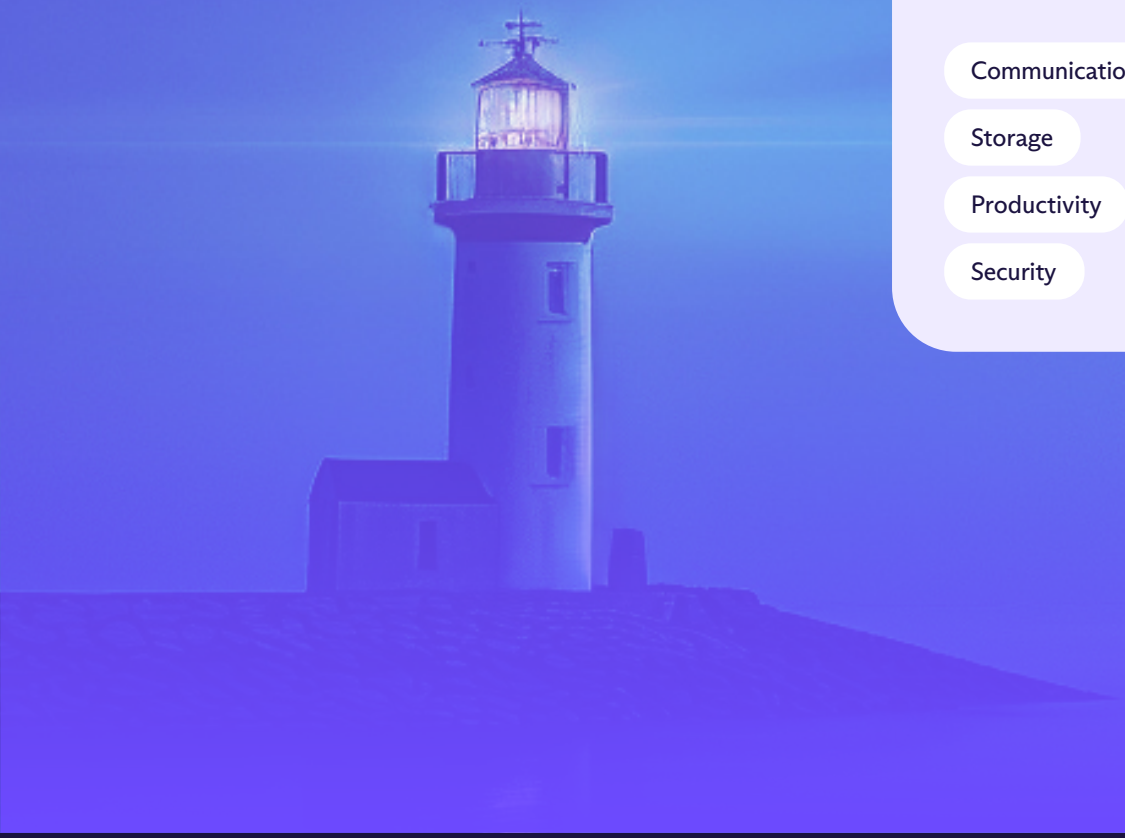

Chapter sections

- Communications
- Storage
- Productivity
- Security

Security tools **for small businesses**

READ THIS APPENDIX to see all the different apps, programs, and services that could help your company increase its information security. This list includes both free and paid services for:

1. Communications
2. Storage
3. Productivity
4. Security

The Proton security team made this an Appendix because information security is primarily about identifying the data you want to protect, the risks your organization faces, and the controls you will put in place to mitigate those risks. Merely switching to encrypted services will not solve your security issues. The previous seven chapters describe how to create a security team, security function, and sound security policies.

However, the following encrypted services will reduce your company's exposure and, when paired with a security-conscious workforce, can go a long way to preventing a data breach or hack. We should note that we don't have any special insider information on these companies. We relied upon their own reporting when coming up with this list.

Communication

Email provider

Most small or medium-sized businesses rely on emails to handle both their internal and external communications. Email security best practices are essential to keeping your business's and employees' data safe, but some email providers can offer your company more security than others.

Proton Mail

Proton Mail offers its users automatic end-to-end encryption. Your emails are encrypted before they leave your device so that only you and your intended recipient can access them. You can even secure your messages to non-Proton Mail users by sending password-protected emails. Proton Mail also includes advanced security features such as Proton Sentinel, an account protection program that provides maximum security for those who need it. It mitigates security threats by combining AI with human analysis. Proton Mail also released Key Transparency for email, a pioneering security feature that enables Proton Mail users to verify the integrity and authenticity of the people users are in contact with.

Platforms: Android, iOS, Linux, macOS, Windows, and web app. Also has Bridge integration with Microsoft Outlook, Mozilla Thunderbird, and Apple Mail.
Open source: Yes
Price: Mail Essentials plan begins at $6.99 per user per month.
Headquarters: Geneva, Switzerland

Team collaboration

Many businesses have employees and contractors working remotely. This can make coordinating a challenge unless you use a team collaboration app. Given the amount of information that can be exchanged and stored on these platforms, using one that is encrypted is a necessity.

Wire

Wire is one of the only end-to-end encrypted services that allows for group calls (up to 25 participants) and video calls (up to 12), which makes it more secure than Slack when trying to manage team communication. Wire has been independently audited, giving you some assurance that Wire's code is doing exactly what they say it is.

Platforms: Android, iOS, Linux, macOS, Windows, and web browser add-ons.
Open source: Yes
Price: Has a free option. Premium plans begin at €6 per user per month.
Headquarters: Zug, Switzerland

Element

Element is an end-to-end encrypted messenger and collaboration tool that runs on the Matrix protocol. It's free for businesses of up to 200 employees and allows for group calls with up to 100 members. It is a more secure version of Slack.

Platforms: Android, iOS, Linux, macOS, Windows, and web browser apps.
Open source: Yes
Price: Has a free option. Premium plans begin at $5 per user per month.
Headquarters: London, England

Messaging

For companies that do not need all the functionality of a collaboration app but still want their communications to be secure, there are end-to-end encrypted messaging apps.

Signal

Signal is widely considered to be the most secure encrypted messaging app, with many messaging apps relying on the Signal Protocol. It supports texts, group texts, voice and video calls, and group calls, but doesn't have a business app.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: Yes
Price: Free
Headquarters: Mountain View, California, USA

Threema

Threema offers anonymous messaging and business-specific app. The company headquarters is in Switzerland, giving its service strong legal privacy protections.

Platforms: Android, iOS, Linux, macOS, Windows, and web app.
Open source: Yes
Price: Starts at 1.50 CHF per user per month
Headquarters: Zurich, Switzerland

Olvid

Olvid is an end-to-end encrypted messaging service that provides a dedicated business app. Olvid offers secure group chats with audio calls, multiple profiles, ephemeral messages, and remote deletion of messages.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: Yes
Price: Business plans begin at €9.90 per user per month.
Headquarters: Paris, France

Storage

Cloud storage has redefined how offices can work. By storing files on the cloud, your business can maintain a backup of all critical documents in case of a catastrophic system failure as well as easily share documents and sync progress between different employees. Protecting these files and the data they contain should be one of your business's top priorities.

Proton Drive

Proton Drive is an end-to-end encrypted and open-source cloud storage solution. In addition to the files itself, Proton Drive also encrypts file metadata such as filenames, folders, and file extensions. It gives your employees control over their file-sharing links even after they've distributed them, allowing them to password protect or disable them with one click.

Platforms: Android, iOS, macOS, web, and Windows.
Open source: Yes
Price: Included in bundle with Mail Essentials plan, which begins at $6.99 per user per month.
Headquarters: Geneva, Switzerland

Tresorit

Tresorit is an end-to-end encrypted cloud storage service. It has optimized its service for businesses, allowing you to create different levels of access for various documents and to revoke users' and devices' access to files.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: No
Price: Starts at 27.50 CHF per user per month.
Headquarters: Zurich, Switzerland

Sync

Sync is another end-to-end encrypted cloud storage service, similar to Tresorit. It gives businesses admin control, allowing supervisors to create different levels of access for different employees. Sync also allows you to preview your files before you open them.

Platforms: Android, iOS, macOS, and Windows.
Open source: No
Price: Starts at $6 per user per month.
Headquarters: Toronto, Canada

Cryptomator

Cryptomator allows you to encrypt your documents before you save them on a separate cloud service, like Dropbox or Google Drive. With Cryptomator, your employees can create a virtual hard drive that is connected to a folder (called a "vault") on their cloud storage service and protect it with a password. Any document they drag and drop into the virtual hard drive is automatically encrypted and backed up in the vault. There is also Cryptomator Hub, for larger businesses looking to add encryption to the files on their company servers.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: Yes
Price: Free (there is a one-time fee of $13.99 to download the Android app and $9.99 to download the iOS app).
Headquarters: Sankt Augustin, Germany

Other cloud services

- **pCloud** is an end-to-end encrypted cloud service. It allows two-factor authentication, and its business subscriptions start at $7.99 per user per month.

Productivity

Also known as a “text editor,” a notepad is a program that allows you to write and edit plain text. A notepad can be used to keep notes, write documents, and alter configuration files, or programming language source code.

Standard Notes

Standard Notes is a simple, end-to-end encrypted note-taking app that can sync your notes across all your devices, and it's now a part of the Proton ecosystem. Its clean interface and numerous extensions mean that you can use Standard Notes for everything from writing yourself reminders to coding.

Platforms: Android, iOS, Linux, macOS, Windows, and web browser add-ons.
Open source: Yes
Price: Has a free option. Premium plans begin at $9.99 per user per month.
Headquarters: Chicago, Illinois, USA

Joplin

Joplin is another end-to-end encrypted note-taking app, but unlike Standard Notes, users must manually activate the end-to-end encryption feature. Joplin relies on the Joplin Cloud or external services, like NextCloud or Dropbox, to synchronize across devices.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: Yes
Price: Joplin Cloud starts at €2.40 per user per month. Free versions are also available
Headquarters: Paris, France

Obsidian

Obsidian is an end-to-end encrypted note-taking app that lets you easily publish notes online. It also allows you to make your own personal wikis and graph the relations between your notes.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: No
Price: Commercial plan is $50 per user per year
Headquarters: Irvine, California, USA

Security

VPN

A virtual private network is an effective way to add a layer of encryption to your online activity, allow your employees to safely work on public WiFi while they are on the road, and securely access your internal network.

Proton VPN

Proton VPN secures your internet connection with AES 256-bit encryption, the industry gold standard, and its use of Perfect Forward Secrecy means that even if your traffic is intercepted and saved, it can never be decrypted at a later date. All Proton VPN apps are open source, and we use the latest VPN protocols, including WireGuard and our own Stealth protocol.

It also offers dedicated IP addresses, meaning you can create secure gateways for your remote employees to access your internal network. You can also segment your employees' access using the IP address you give them, only allowing them to access the data they need for their job.

Platforms: Android, iOS, Linux, macOS, and Windows.
Open source: Yes
Price: Business plans begin at $5.99 per user per month.
Headquarters: Geneva, Switzerland

Password manager

Creating strong, unique passwords or passphrases for your accounts is one of the basics of IT security, but no employee can remember all the passwords necessary to log in to all the platforms they need to use for work. A password manager changes all that. By safely encrypting all your passwords, a password manager allows you to create passwords that are impossible to crack, without having to remember them all. Using a trustworthy password manager to secure your passwords is one of the easiest ways to improve your company's security.

Proton Pass

Proton Pass lets your employees create, store, and autofill their usernames and passwords while protecting this information with end-to-end encryption. It also supports passkeys, which are even easier to use than passwords. Pass has been independently audited and is part of the Proton encrypted suite.

Platforms: Android, iOS, Chrome, Firefox, Brave, and Edge.
Open source: Yes
Price: Business plans begin at $2.99 per user per month.
Headquarters: Geneva, Switzerland

Bitwarden

Bitwarden is an end-to-end encrypted password manager. It helps your employees create randomly generated passwords for all of their accounts, and then syncs those passwords across all their devices.

Platforms: Android, iOS, Linux, macOS, Windows, and web browser add-ons.
Open source: Yes
Price: Starts at $4 per user per month.
Headquarters: Florida, USA

1Password

1Password is another end-to-end encrypted password manager, but it has a few more bells and whistles. While it is only a paid service, it is considered to be one of the most secure password managers. Its Watchtower feature will alert you if any of your passwords have been exposed in recent data breaches.

Platforms: Android, iOS, Linux, macOS, Windows, and web browser add-ons.
Open source: No
Price: Starts at $19.95 for up to 10 users per month.
Headquarters: Toronto, Canada

Dashlane

Dashlane is also a premium end-to-end encrypted password manager. It will scan known security breaches and will send you an alert if it finds any of your passwords among those exposed. Its business plan also comes with an admin console that allows you to set permission levels for all your employees.

Platforms: Android, iOS, macOS, Windows, and web browser add-ons.
Open source: In the process of publishing the code of its apps. Has shared the source code of its Android and iOS apps publicly available.
Price: Starts at $8 per user per month.
Headquarters: New York City, USA

Other password managers

- **KeePass / KeePassXC**: These are both free, open-source password managers, but neither of them offer official mobile apps.

Two-factor authentication

To ensure your critical accounts are secure, you should enable two-factor authentication (2FA) in addition to using a strong, unique password. The site Two Factor Auth will help you identify which services you can use 2FA on. By using 2FA on your accounts, you can prevent intruders from accessing your accounts even if they get a hold of your passwords.

Yubico

The Ybico is a hardware token (a specialized USB stick) that you can plug into your device to confirm your identity. While it is thought to be the most secure form of 2FA, relatively few services support hardware token 2FA. You can sign in to your Proton Account on the web using a hardware key, such as a Yubico, as long as it adheres to the U2F or FIDO2 standard.

Platforms: YubiKey 5 NFC works with macOS, Windows, and NFC-equipped Android and iOS devices.
Open source: YubiKey hardware with its integral firmware has never been open sourced, whereas almost all of the supporting applications are open source
Price: A YubiKey 5 NFC costs $50.
Headquarters: Palo Alto, California, USA

Proton Pass

Proton Pass also includes a built-in 2FA authenticator, enabling it to quickly and easily handle your entire secure log-in process.

Platforms: Android, iOS, Chrome, Firefox, Brave, and Edge.
Open source: Yes
Price: Business plans begin at $2.99 per user per month.
Headquarters: Geneva, Switzerland

Duo

Duo offers several 2FA solutions, including ones that incorporate Yubikey hardware tokens, confirmation requests delivered to the Duo app that foil man-in-the-middle attacks, and time-based one-time passcodes.

Platforms: Duo app is available on Android and iOS.
Open source: Yes
Price: Has a free option. Premium plans begin at $3 per user per month.
Headquarters: Austin, Texas, USA

Other two factor authentication services

Google Authenticator app: Google offers a free authenticator app that creates time-based one-time passcodes for 2FA purposes. It does not have the same functionality as Duo or a YubiKey.

Disk encryption

All your devices should use some form of disk encryption to prevent unauthorized access to your devices' data storage in the event they are stolen or lost. By encrypting your smartphone or computer's hard drive, you turn your sensitive data into illegible code that can only be decrypted by your password. All the options discussed below are examples of disk encryption software.

VeraCrypt

VeraCrypt is a disk encryption service. Using VeraCrypt, your employees can encrypt the hard drive on their device, encrypt their USB flash drive, or even hide how much volume they have on their hard drive.

Platforms: Linux, macOS, and Windows.
Open source: Yes
Price: Free
Headquarters: N/A

Other disk encryption services

- **FileVault:** FileVault is available on macOS X Lion and later. You can use it to fully encrypt your startup disk.
- **BitLocker:** BitLocker is available on most Windows 7 and Windows 10 devices. It is a strong, full disk encryption service.
- **LUKS:** LUKS is a free, open source hard disk encryption service for Linux.

Native encryption for Android and iOS: Any 3G iOS device or later and any Lollipop (5.x) Android devices or later are equipped with their own native disk encryption services. To learn how to encrypt your Android device, click here. To encrypt your iOS device, click here.

Personal antivirus software

Antivirus software (AVS) is a preventative measure meant to keep your devices clean. AVS scans your device for any malware, from ransomware to rootkits. If it detects any, it will attempt to remove them. More modern AVS also provides malware prevention measures.

Bitdefender

Bitdefender has strong antivirus protection, but it is light enough that it won't slow your device down. The AVS receives daily updates so that no malware can take it by surprise. They also have a more advanced option for larger offices. It will protect your servers and all your endpoint workstations without bogging down your network.

Platforms: Android, macOS, Windows are available free. iOS is available with a paid plan.
Price: Has a free option. Small office security starts at $99 for one year for five devices.
Headquarters: Romania

Windows Defender Antivirus

Windows Defender is an outgrowth of the antivirus software offered as part of Windows. It offers real-time protection and access control, and supports programs for Android, Mac, and iPhones.

Platforms: Android, iOS, macOS, and Windows.
Price: Business plans start at $3 a user a month.
Headquarters: USA

F-Secure

F-Secure offers an entire cybersecurity suite, including antivirus solutions, a VPN, password manager, and more. It sells itself as a single-app solution for your online privacy and security needs.

Platforms: Android, iOS, macOS, and Windows.
Price: Business plans start at $69 per device per year.
Headquarters: Finland

About **the author**

Patricia Egger joined Proton in 2021 and is our Security, Risk, and Governance Manager. She received a Master's degree in Applied Mathematics at the École polytechnique fédérale de Lausanne (EPFL). Before joining Proton, Patricia worked as a senior security consultant for Deloitte and the security officer for Kudelski Security. She is also the co-founder and Vice President of Women in Cyber Switzerland.