

HOUSTON LABS LLC

ZERO TRUST · CIVILIAN FIELD MANUAL

Zero Trust Implementations for High-Risk Individuals and Small Organizations

A practical security field manual for real targets

First Edition · 2026 · Adapted from the NSA Zero Trust Implementation Guideline Primer (Jan 2026)

Contents

About This Guide	3
How to Use This Guide	6
Start Here: Do These First	9
The Zero Trust Mindset	11
Your Threat Landscape	15
Pillar 1 · You: Accounts & Identity	19
Pillar 2 · Your Devices	25
Pillar 3 · Apps & Software	31
Pillar 4 · Your Data	36
Pillar 5 · Network & Environment	40
Pillar 6 · Automation (Make It Automatic)	45
Pillar 7 · Visibility & Awareness	50
When You're Targeted Personally	56
Physical & Travel Security	61
Family & Dependents	65
Financial Self-Defense	69
Incident Response & Recovery	73
Keep Going	78
Master Checklist	80
Glossary	85
Appendix A: NSA Pillar Crosswalk	90
Appendix B: Resources	92
Appendix C: Sources	94

About This Guide

What This Guide Is

Security guidance has a gap problem. On one side: enterprise playbooks written for organizations with dedicated IT departments, six-figure tooling budgets, and the authority to mandate software across a workforce. On the other: beginner advice so generic it rarely applies to people who face genuine threats.

This guide lives in neither of those places. It is written for capable, security-minded people (journalists, activists, researchers, founders, organizers, attorneys, clinicians, and others) who hold sensitive information, face real adversaries, and have no security department to call. You are not a beginner. You already understand how your tools work. What this guide gives you is a coherent threat-defense framework, translated out of government doctrine and into the decisions you actually make.

That framework is called Zero Trust. Not a product, not a vendor checklist, a way of designing your security so that a single breach does not cascade into a catastrophe. The seven-pillar structure in this guide maps directly to the seven domains government security architects use to build classified-system defenses. The steps are adapted for a person working alone, without an IT department, at civilian scale.

Who Houston Labs Is

Houston Labs LLC is a technological research and development company, based in New York and established in 2025. We build the tools, systems, and ideas at the frontier where advanced technology meets human creativity, across applied research, creative technology, products and ventures, and advisory and education.

Our mission is to further humanity by closing the adoption gap: the widening distance between what advanced technology makes possible and what ordinary people and small organizations can actually use. Most of that power stays with those who already hold the advantage. We exist to put it in everyone's hands, through products built to be genuinely usable and through free guidance like this.

This guide is part of our free advisory and education work. The individuals and small organizations we build for, the ones we want to keep competitive in a fast-changing age, are also the ones being targeted, and the ones with no security department to call. So we took the doctrine governments use to defend their most sensitive systems and translated it into something a capable person can execute alone. We publish it independently and vendor-neutrally: we have no financial relationship with any security tool or service named in these pages, and our guidance is capability-based so you can choose what fits your situation.

The Promise

Government-grade Zero Trust, translated for the people who get the real work done.

Every chapter in this guide ends with a prioritized checklist. Every concept comes with a concrete civilian example. Every pillar chapter is honest about the difference between baseline steps that everyone should take and elevated measures that only matter if you are actively targeted.

You do not have to do everything at once. One chapter, one checklist item, makes you harder to compromise than you were yesterday. Start where you are and build from there.

A Note on Tools

This edition is capability-based. It describes what a good tool does (what criteria it should meet, what questions to ask before you trust it) rather than naming and ranking specific products. That approach ages better; the threat landscape shifts faster than any product recommendation can keep up with. Verify any tool against current community sources before you adopt it for sensitive use.

A product-specific edition (naming and evaluating specific tools by category against the criteria laid out here) is planned as a follow-up to this guide. It will cover the same seven pillars with concrete software and service recommendations.

Disclaimer

This guide is provided "as is," without warranty of any kind, express or implied. Houston Labs LLC makes no representations about the completeness, accuracy, or fitness for any particular purpose of the information contained herein. Security is a fast-moving field. Threats evolve, tools change, legal contexts vary by jurisdiction, and recommendations that are sound at publication may be outdated by the time you read this. Verify recommendations against current sources before relying on them for high-stakes decisions. This guide is not legal advice, not professional security consulting, and is not a substitute for either. Consult qualified professionals for your specific situation.

First Edition, 2026.

This guide is adapted from the *NSA Zero Trust Implementation Guideline Primer* (January 2026), a work of the United States federal government produced in the course of official duties and therefore in the public domain in the United States. The civilian adaptation, original text, examples, framing, and editorial voice are the original work of Houston Labs LLC.

This is an independent publication. It is not affiliated with, authorized by, sponsored by, or endorsed by the National Security Agency, the United States Department of Defense, or any other agency or instrumentality of the United States government. References to NSA

doctrine, government frameworks, or federal security guidance are for informational and educational purposes only and do not imply any relationship with or approval by those entities.

How to Use This Guide

The Structure: Seven Pillars Plus Special Chapters

The guide is built around seven security pillars, the seven domains that together define a complete personal security posture. Each pillar gets its own chapter, written to stand alone. You can read them in any order and return to any one as a reference when a specific concern comes up.

The seven pillars, in civilian terms:

1. **You (User)**, your accounts, identity, authentication, and account recovery
2. **Your Devices**, phones, laptops, and the hardware your data lives on
3. **Apps & Software**, what you install, what permissions it holds, and how you keep it current
4. **Your Data**, where it lives, who can reach it, and how it is encrypted and backed up
5. **Network & Environment**, the connections you use and the physical environments you work in
6. **Automation**, scripts, APIs, tokens, and background processes that act on your behalf
7. **Visibility & Awareness**, knowing what is happening on your accounts and devices so you catch trouble early

Surrounding the pillars are several special chapters: this one; a *Mindset* chapter explaining Zero Trust doctrine in civilian terms; a *Threat Landscape* overview; the *Start Here* chapter with the highest-leverage quick wins; a chapter on *Operating Under Elevated Risk*; and a resources appendix. The special chapters do not follow the pillar structure (they use whatever organization fits the material) but they use the standard callouts where appropriate.

The Personal-First Spine

Every pillar chapter is written for a single person first. The core text assumes you are securing your own accounts, devices, and data. You have no authority over anyone else's setup and no IT department to call.

If you also manage a small team or organization, look for the **Scaling to your team or organization** boxes. These appear in every pillar chapter and translate the personal-first steps into group settings: shared infrastructure, onboarding processes, access policy baselines, and the minimum viable security program for a team of up to roughly fifty people. If you are only securing yourself, skip those boxes freely without losing continuity.

The Three Callout Boxes

Each chapter uses three standard callout boxes. Knowing what each one signals helps you read faster and skip what does not apply to you right now.

Scaling to your team or organization

Group-level guidance for small-org operators. Covers shared accounts, team access policies, onboarding baselines, and low-cost tooling appropriate for organizations without a dedicated security team. Skip freely if you are only securing yourself.

If you're being targeted

Elevated measures for readers who face active, sophisticated adversaries, targeted surveillance, stalking, harassment campaigns, or nation-state-level threats. These steps go beyond what most people need and add meaningful friction to daily workflows. If you work in high-risk journalism, political opposition, human rights advocacy, sensitive research, or any role that has drawn serious adversarial attention, read these carefully. If none of those describe you, note where the box is and come back if your situation changes.

⚠ Common mistakes

The errors that appear most consistently in real security incidents, phrased as concrete warnings. Worth reading even if you skip the rest of a chapter; these describe the specific gaps that attackers actually exploit.

Reading by Role

If you are an everyday individual

Start with *Start Here*, then work through the *You (User)* and *Your Devices* pillar chapters. Those three chapters together cover the vast majority of attack paths that affect people who are not specifically targeted. The other pillars are available when a specific concern arises: unusual network behavior, a suspicious app, a question about where your data actually lives.

If you are a small-org operator

Read all seven pillar chapters once through, paying attention to the boxes. Then work back through the checklists, starting with the pillar that matches your highest-risk current gap. Pay particular attention to the *Automation* and *Visibility* chapters, small teams consistently underestimate how many tokens and automated access grants they have issued and how little visibility they have when something goes wrong.

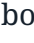


If you are a high-risk target

Read the full guide, including every box. Read *Mindset* and *Threat Landscape* before the pillar chapters, the doctrine and adversary model make everything else more legible. Then work through *Operating Under Elevated Risk* alongside the pillar checklists. If possible, work through

the implementation steps with a trusted technical peer who can help with the more complex configurations.

How Each Pillar Chapter Is Structured

Every pillar chapter follows the same six-part pattern. Once you have read one, you know how to navigate all of them:

1. **Why it matters**, the specific threat, in civilian terms
2. **What good looks like**, the target state you are working toward
3. **Do this**, prioritized steps with a checklist
4. **Scaling to your team or organization**, inside the  box
5. **If you're being targeted**, inside the  box
6. **Common mistakes**, inside the  box

You do not have to do everything at once. Security is a practice, not a project with a finish line. Every step you complete narrows the attack surface available to an adversary. Start with the next chapter, *Start Here*, and build from there.

Start Here: Do These First

You do not have to read this entire manual before you become meaningfully harder to attack. Security has an 80/20 rule: a small set of high-leverage moves blocks the most common attack paths. The seven actions below are those moves. They are ordered by the risk they close, not by how convenient they are. Start at the top.

- ☐ **Set up a password manager and create a unique, random password for every account.**
Credential stuffing (taking a leaked username and password from one breach and trying it everywhere) is the single most common account-takeover method. A password manager (an encrypted vault that generates and fills credentials for you) eliminates this attack class entirely. If you reuse even one password, every site that holds it is a single point of failure for every other account you own.
- ☐ **Turn on phishing-resistant MFA (multi-factor authentication, a second proof beyond your password) or passkeys everywhere you can; remove SMS 2FA where alternatives exist.**
A stolen or guessed password alone is not enough to break in when strong MFA is active. Hardware security keys and passkeys are immune to phishing; SMS codes are not, SIM-swapping (convincing a carrier to redirect your number) is a well-documented attack used against high-value targets. Upgrade to app-based TOTP (time-based one-time passwords) at minimum; use hardware keys or passkeys for your most critical accounts.
- ☐ **Turn on automatic updates for your phone, computer, apps, and router.**
The majority of successful malware infections exploit known vulnerabilities, bugs that already have patches available. Attackers specifically target people who have not yet applied updates. Automatic updates remove the human delay. This is true at every threat level, from opportunistic scammers to sophisticated actors.
- ☐ **Encrypt your devices and set a strong screen lock (at minimum a long PIN or passphrase).**
Full-disk encryption (FDE) means a lost or seized device is a brick, not a data breach. Modern iPhones and most Android phones encrypt by default when a screen lock is set, verify yours is actually on. Laptops often do not encrypt by default; enable it now. A weak PIN (4 digits, all zeros, your birth year) defeats the purpose.
- ☐ **Back up your important data using the 3-2-1 rule, and test that you can actually restore from it.**
Ransomware, device seizure, hardware failure, and accidental deletion are all solved by a good backup. The 3-2-1 rule: three copies of the data, on two different media types, with one copy off-site or offline. An untested backup is not a backup, run a restore drill. If you have never practiced recovering a file from scratch, do it before you need to.

☐ **Lock down your account recovery, treat your primary email address and phone number as master keys.**

Every "forgot your password?" flow leads back to email or SMS. Whoever controls your inbox controls every account linked to it. Secure your primary email account first (strong unique password, the best MFA you have). Then audit every other account's recovery contact and remove phone numbers from recovery wherever possible. Your email account deserves more protection than your bank account, it is the skeleton key to everything else.

☐ **Check your exposure on Have I Been Pwned (haveibeenpwned.com) and turn on login alerts for critical accounts.**

You cannot defend what you cannot see. Have I Been Pwned (a free, widely-trusted breach-notification service) tells you which of your email addresses appear in known data breaches, and which passwords were exposed. Login alerts from your email and financial providers give you a real-time signal when someone else tries to access your account. Check both now; set up alerts; revisit the breach check whenever you hear about a major incident.

These seven actions harden the highest-probability attack paths against you. None of them are advanced. All of them are reversible if you change your mind about a tool. Doing them imperfectly is still far better than not doing them.

Then work through the pillars at your own pace.

The Zero Trust Mindset

The Old Model: The Castle and the Wall

Traditional security is built around a single idea: keep attackers out. Build a strong perimeter (a firewall, a corporate VPN, a walled network) and trust everything inside it. Once something or someone is inside the castle, it has earned the right to move freely. The gate is the defense. If the gate holds, you are safe.

This model made sense when data lived on servers in one building, employees worked from one office, and the threat was mostly outsiders trying to breach the front wall. It does not describe the world you actually work in.

Your data lives across three cloud services, two email accounts, a project management tool, a shared drive, and your phone. You work from home, coffee shops, airports, shared offices, and hotel Wi-Fi. Your devices connect to dozens of services through APIs and background sync processes you have probably never audited. The people inside your "perimeter" (collaborators, contractors, anyone who has ever had access) may have moved on, changed loyalties, or been compromised themselves. And the threat is not just outsiders at the gate. One of your passwords is likely in a breach database right now. A phishing message will land in your inbox this month. The attackers are already in the courtyard, and the castle wall has a dozen unmapped doors in it.

The castle model fails not because it was wrong for its time, but because the castle no longer exists.

The New Model: Zero Trust

Zero Trust starts from a different premise: *there is no safe perimeter*. The network you are on, the device you are using, the app making a request (none of these earn automatic trust just because they are familiar. Every access request is verified at the moment it is made. Every connection is treated as potentially hostile until it can be authenticated. And the architecture is designed so that when something does get through) and something will, the damage is contained rather than total.

Zero Trust is not paranoia. It is engineering. You are not trying to build a system that can never be breached; you are building a system that survives being breached. Those are very different design goals, and the second one is actually achievable.

Five principles make the doctrine concrete.

Principle 1: Never Trust, Always Verify

Do not grant access based on location or familiarity. Being inside your home network does not make a device trustworthy. Using your usual laptop in your usual coffee shop does not make that connection safe. A colleague's request coming from a known email address does not confirm that your colleague sent it. Trust is not a property of where something is, it is earned by verifying identity and state at the moment of access, every time.

In practice: every sensitive account requires strong authentication every time, not just on first login. Your cloud storage is not wide-open to anything on your home Wi-Fi. An unusual request from a known contact still gets scrutiny. The question is never "does this look familiar?", it is "can this specific request be verified, right now, by evidence I trust?"

Principle 2: Assume Breach

Operate as though an attacker already has a foothold somewhere in your setup. Not because you are being alarmist, but because it is almost certainly true. Credential databases leak routinely. Apps are compromised. Devices get lost, seized, or infected. One of your passwords, from one service you used years ago, is almost certainly sitting in a breach database. The question is not whether a breach will happen, it is whether you have designed your setup to limit what an attacker can do once it does.

Assume breach changes how you architect everything. Your primary email credentials alone should not give an attacker access to your backups, your financial accounts, and your encrypted messaging history. Your phone being seized should not expose three years of private conversations. A single compromised cloud account should not unravel your entire digital life. Assume breach means you design for the worst reasonable case and then live normally inside that design. It is not about fear; it is about blast radius.

Principle 3: Verify Explicitly

When something requests access, use all the evidence available, not just a password. Strong verification uses multiple signals simultaneously: something you know (a password or PIN), something you have (a hardware security key, an authenticator app, a passkey stored on a device), and signals about the state of the thing making the request (is the device encrypted? is it running current software? is this request geographically plausible?).

In civilian terms, a password alone is not enough for any account that matters. MFA (multi-factor authentication (a second proof beyond your password, such as an authenticator app code or a hardware key) is the minimum. Phishing-resistant MFA) using a hardware security key or a passkey, is better still, because it cannot be intercepted by a fake login page the way a one-time code can. Verify explicitly means you never stake account security on a single signal, no matter how strong it seems.

Principle 4: Least Privilege

Give every account, app, and process only the access it actually needs to do its job, nothing more, and revoke it when the job is done. This is one of the oldest principles in computer security and one of the most consistently ignored, because granting broad access is easier and over-access rarely has obvious immediate consequences.

That photo-editing app does not need access to your contacts, location, and microphone. That shared folder you set up for a collaboration last year does not need to stay open to all your collaborators indefinitely. Your personal email account and your work accounts should not be linked in a way that lets a compromise of one cascade into the other. The less access any single account, app, or credential holds, the smaller the blast radius when it is compromised. Shrink access. Audit it periodically. Revoke what is no longer needed.

Principle 5: Architect Inside-Out

Protect what matters most, first. Start at the center (your most sensitive accounts, your most critical data, your highest-risk communications) and build outward. Do not try to harden everything at once. You will run out of energy before you finish, and your most important assets will receive the same attention as your least important ones.

Identify your crown jewels: the accounts, files, and communications that would cause the most harm if an adversary obtained them. Your primary email account is usually first on the list, because it controls password recovery for almost everything else. Build the strongest protections around those first, then expand outward as capacity allows. This inside-out approach is how governments prioritize security for classified systems, and it scales down to the individual precisely because it forces explicit choices about what matters.

Why the Mindset Comes First

Tools are implementations of this mindset. A password manager enforces unique credentials so that a single breach does not cascade. An authenticator app enforces verify explicitly. Device encryption enforces assume breach. End-to-end encrypted messaging enforces least privilege in transit. But if you do not understand why you are using a tool, you will misconfigure it, underuse it, or abandon it the moment it adds friction.

The mindset also gives you a framework for decisions that no checklist covers. A new phishing technique appears. A service you use announces a breach. A collaborator asks for access that seems slightly broader than necessary. The five principles tell you how to reason about each of these, even before specific guidance exists.

Everything in the pillar chapters is an application of these five principles. Keep them in mind as you work through the checklists. They are the "why" behind every item.

NOTE

Zero Trust is not a product you can buy. Any vendor that markets "a Zero Trust solution" is selling a tool that supports Zero Trust principles in one specific domain. No single product implements the full framework. The framework emerges from how you combine and configure the pieces, and from the decisions you make every day about access, verification, and privilege.

Your Threat Landscape

Who Is Actually Coming for You

Security training often skips the adversary question. Most guidance assumes a vague "bad actor" and jumps to countermeasures. That approach fails because different adversaries have different capabilities, different goals, and different weak points. A scammer running bulk phishing campaigns is stopped by different defenses than an ex-partner who already knows your security questions and your email address. Both are real threats. Both require a specific response.

The table below maps the adversary types that governments and enterprises defend against to their civilian equivalents, the people and organizations that actually target individuals like you.

Government and enterprise adversary types mapped to civilian equivalents

Government / enterprise adversary	Civilian equivalent	What they are after
Nation-state actor	Targeted surveillance of journalists, activists, human rights workers, researchers, lawyers, and political figures	Sources and contacts, communications content, location history, unpublished documents, operational plans, identifying information about others in your network
Insider threat	Stalkers, abusive partners, doxxers, hostile former employees or collaborators	Your location, your private communications, compromising personal information, account control, your contact list as leverage
Credential theft	Phishing, account takeover, SIM swapping (hijacking your phone number to intercept SMS codes)	Access to your email, financial accounts, social media, and every account that uses those for password recovery
Data exfiltration	Data brokers, breach databases, aggregators, and scrapers	Your personal information at scale (home address, employer, family relationships, biometrics, behavioral patterns) packaged and sold
Supply chain attack	Malicious apps, browser extensions, and compromised software updates	A persistent foothold on your device, often used to harvest credentials, log keystrokes, or exfiltrate data over time, sometimes months after initial infection
Ransomware	Ransomware (same at every scale)	Your files, encrypted and held until you pay, or sold to other parties if you do not

Most people face some version of every row in that table, simultaneously. These adversary types are not mutually exclusive. A sophisticated attacker may use multiple techniques in sequence: a phishing message to steal credentials, then those credentials to harvest personal data, then that data to run a targeted impersonation attack on someone in your network. The rows describe attack categories, not distinct attackers.

Size Your Own Threat Model

A threat model is not a paranoid exercise. It is four honest questions about your specific situation. Answer them before you work through the pillar chapters and the checklists will make more sense. Return to these questions when your work, your relationships, or your risk profile changes.

Question 1: What Are My Crown Jewels?

Crown jewels are the accounts, files, and communications that would cause the most harm if an adversary obtained them. Be specific. Examples: your primary email account, because it controls password recovery for almost everything else; source communications if you are a journalist; client files if you are an attorney; years of contact history that would expose other people; financial accounts and credentials; unpublished research; a private encryption key or authentication credential.

Write the list down. Keep it short, five to ten items at most. If everything feels equally critical, you do not have a threat model yet; you have undifferentiated anxiety. Force yourself to rank. The items at the top of the list get the most protection first.

Question 2: Who Wants Them?

Name the adversary type as specifically as you can. Opportunistic scammers running automated phishing campaigns? A specific person with a personal grievance who already knows things about you? A commercial interest (a data broker, a former employer, a litigant) with resources and patience? A government or law enforcement interest?

The specificity matters because it determines capability and method. Bulk phishing campaigns are automated and untargeted, you stop them with a password manager and MFA. A determined adversary who has already researched you and knows your recovery email is a different problem that requires different defenses. Most people face more than one adversary type. List all of them.

Question 3: What Can They Do?

Be honest about the capabilities of each adversary you named. A bulk-phishing operation has a list of email addresses and a convincing fake login page. A stalker may know your home address, your workplace, your phone number, and some of your passwords from a previous relationship. A data broker has purchased information about you from dozens of sources and can

build a detailed profile. A well-resourced attacker may have access to commercial spyware, social engineering against people in your network, and in some cases the legal or political leverage to compel service providers to disclose your data.

Honest adversary capability assessment tells you where to focus effort. If your realistic threat is primarily automated bulk attacks, the *Start Here* chapter covers you well. If your threat includes someone with existing personal knowledge of you, or institutional leverage, the elevated-risk measures in the boxes become relevant.

Question 4: What Happens If They Get In?

Walk through the realistic consequences of a successful attack, starting from your crown jewels. If your email is compromised: which other accounts does an attacker now control via password recovery? Who in your contact list gets exposed? What documents are in your inbox? If your devices are seized: what data is accessible without your cooperation? Which apps can be opened without authentication? If your location history leaks: who else in your network is identifiable from it?

This question reveals blast radius, the total damage a single breach can cause. Everything in this guide is ultimately about shrinking that blast radius. The smaller it is, the more resilient you are to the breaches that will inevitably happen.

IF YOU'RE BEING TARGETED

If your adversary is a government, a well-resourced private actor, or a sophisticated stalker (not a possibility you are speculating about but a realistic assessment based on your work, your public profile, or prior contact) your threat landscape has additional dimensions that consumer security guidance rarely addresses.

Sophisticated adversaries use techniques that are rarely discussed in mainstream security advice: IMSI catchers (devices that impersonate cell towers to intercept calls and SMS messages), commercial spyware deployed through zero-click exploits that require no action from you to install, legal compulsion of cloud service providers to disclose account data without notifying you, and social engineering attacks aimed at people in your network rather than at you directly.

This does not mean you are helpless. It means the countermeasures go further than the baseline, and that implementing them correctly matters more. Every box in the pillar chapters covers the elevated measures relevant to that domain. The *Operating Under Elevated Risk* chapter addresses this threat tier directly. If you are in this category, also consider reaching out to organizations that provide direct support to journalists, activists, and high-risk individuals: [Access Now's Digital Security Helpline](https://www.accessnow.org/help/) (<https://www.accessnow.org/help/>) offers free, confidential technical assistance.

The Threat Landscape Is Not Static

Revisit your threat model when things change. A new role, a new publication, a new political position, a relationship ending badly, a lawsuit, a prominent interview, or a breach notification from a service you use, all of these can shift your adversary set or your crown jewels. The four questions above take fifteen minutes. They are worth revisiting once or twice a year even if nothing dramatic has happened.

The pillar chapters that follow are written against the full threat range in the table above. The core steps address the threats that affect most people. The boxes address the elevated tier. You do not have to defend against every adversary simultaneously. Start with the most likely and the most damaging. Build from there.

▲ COMMON MISTAKES

- Assuming you are not a target because you are not famous or politically prominent. Opportunistic attacks do not require fame, they require a reachable account and a weak credential.
- Treating the threat model as a one-time exercise. Your adversary set changes as your work and relationships change.
- Skipping the crown-jewels question and trying to harden everything at once. You will run out of energy and attention before you finish, and your most critical assets will not receive the most protection.
- Assuming the elevated measures in the boxes are only for people in immediate physical danger. Stalking, doxxing, targeted harassment campaigns, and account takeover affect ordinary people in ordinary circumstances.
- Confusing threat modeling with a security to-do list. The model tells you who and what, the pillar chapters tell you how to respond.

PILLAR 1 · USER

You: Accounts & Identity

Why it matters

Every account you own is a door. Passwords are the lock. The problem is that most people use the same lock on every door, and billions of those keys are already loose.

Here is what happens after a data breach: attackers sell or publish databases of leaked email addresses and passwords. Automated tools called **credential stuffers** take those lists and try every combination against every major service (banking, email, social media, work tools) within hours. If you reused a password from that breached site anywhere else, those accounts are compromised without any targeted effort at all. The attacker never needed to single you out; you were caught in a net cast at millions of people simultaneously.

Check your exposure now at [Have I Been Pwned](https://haveibeenpwned.com) (<https://haveibeenpwned.com>). Enter your email address. You will almost certainly find it in at least one breach. This is your current state, not a future threat.

Beyond reused passwords, your **identity layer** (your password, your second factor, and your account-recovery options) is the primary target of attackers at every level of sophistication. Phishing campaigns, SIM swaps, and account-recovery social engineering all funnel toward this layer because controlling it means controlling everything downstream. Your email account deserves special attention: whoever controls it can reset every other password you own. It is the master key to your digital life, and adversaries know it.

Compartmentalization (keeping separate identities for separate contexts) is the civilian version of "deny by default." When a breach hits one compartment, it does not cascade into the others.

What good looks like

A hardened identity posture has four properties:

- **No reuse.** Every account has a unique, randomly generated password. You could not recite most of them from memory, and that is exactly right, the password manager remembers them so you do not have to.
- **A strong second factor.** MFA (multi-factor authentication, a second proof of identity beyond your password) is enabled on every account that supports it. Critical accounts use phishing-resistant methods: passkeys or a hardware security key.
- **Hardened recovery paths.** Your email account and phone number are protected with your strongest MFA method. Recovery codes are stored offline. There are no weak side doors, no easily guessed security questions, no SMS-based recovery on high-value accounts.

- **Separated contexts.** Work, personal, and public-facing activities each use distinct email addresses and, where the stakes warrant it, distinct personas. A breach or leak in one context does not expose the others.

Do this

1. Get a password manager and use it for everything

Look for an end-to-end-encrypted password manager that has been independently audited, offers apps for every device you use, and syncs over an encrypted vault. Set a long, memorable master passphrase, a random sequence of four or five words works well. Then, every time you log in to a site, let the manager generate a new unique password and save it. Within a few weeks you will have replaced most reused credentials without any single heroic effort.

2. Enable MFA on every account, starting with email

Your email provider's MFA setting is the single highest-leverage action in this chapter. Do it first. Then work through your other accounts in order of sensitivity: financial accounts, work accounts, cloud storage, social media.

Not all MFA methods are equal. Here is the hierarchy from strongest to weakest:

MFA methods compared by phishing resistance

Method	Phishing-resistant?	Notes
Passkey (FIDO2/WebAuthn)	Yes	Strongest widely available option; built into modern phones, computers, and hardware keys
Hardware security key (FIDO2/U2F)	Yes	Physical token; immune to remote phishing; use as a backup alongside passkeys
Authenticator app (TOTP codes)	No, but far better than SMS	Time-based codes can be phished in real time; still significantly stronger than no MFA
SMS text message code	No	Vulnerable to SIM swap; move off this wherever a stronger option exists

SIM swap is an attack where someone convinces your mobile carrier (usually through a fraudulent call or forged ID) to transfer your phone number to a SIM card they control. From that moment, every SMS verification code you would normally receive goes to the attacker instead. Carriers' identity-verification processes are often weak enough for a determined adversary to beat. It is a well-documented vector in targeted account takeovers, including against journalists, executives, and activists.

The goal is phishing-resistant MFA on critical accounts. Use passkeys where supported, major platforms are rolling them out as the default sign-in option. If you use a hardware security key, buy two: register both, store the backup somewhere secure and physically separate from the primary.

3. Store recovery codes offline

When you enroll MFA, most services give you one-time recovery codes for use if you lose your second factor. Print them or write them down and store them somewhere physically secure, a locked drawer, a fireproof box. Do not save them in a cloud document or note-taking app. An attacker who compromises your cloud storage should not find the bypass codes to your MFA sitting right there.

4. Lock down account-recovery paths

Account recovery (the "I forgot my password" flow) is frequently the weakest link. Audit it for each critical account. Remove phone-number recovery on high-value accounts if you can replace it with a recovery code or secondary email. Replace knowledge-based security questions (mother's maiden name, first car, childhood street) with randomly generated nonsense answers stored in your password manager. The real answers to those questions are often findable through public records, old social media posts, or a quick phone call to a family member posing as you.

5. Separate your contexts

Use a distinct email address for work, for personal accounts, and for anything public-facing (newsletters, forums, social media signups. This is the civilian version of least-privilege access: a breach at one newsletter you subscribed to years ago does not expose your work inbox, and your work inbox does not expose your personal accounts. Many email services and password managers support alias addresses) one-click throwaway addresses that forward to your real inbox, for signing up to lower-trust services.

- ☐ Choose and set up a password manager; generate unique passwords for all existing accounts over the next two weeks.
- ☐ Enable phishing-resistant MFA (passkey or hardware security key) on your primary email account today.
- ☐ Enable MFA on all financial, work, and cloud accounts; prefer an authenticator app over SMS where passkeys are not yet available.
- ☐ Move off SMS 2FA on every account where a stronger option exists.
- ☐ Print or write down MFA recovery codes for critical accounts and store them offline in a physically secure location.
- ☐ Review and harden account-recovery settings on email and financial accounts; replace or randomize security-question answers.

- ☐ Search your email addresses on Have I Been Pwned; change the password for any account flagged in a breach.
- ☐ Set up separate email addresses (or aliases) for work, personal, and public-facing use.
- ☐ Register a backup hardware security key if you use one; store it in a different physical location from your primary key.

SCALING TO YOUR TEAM OR ORGANIZATION

For a team of any size, individual good behavior is not enough, the weakest account is the entry point. These controls apply whether you run a five-person nonprofit or a fifty-person company.

- **Mandate MFA organization-wide.** Google Workspace and Microsoft 365 both allow you to enforce MFA for all accounts and block sign-in without it. Do this before anything else. Use conditional access policies to require phishing-resistant methods for anyone with administrative privileges.
- **Centralize authentication through SSO.** Single sign-on (SSO, one identity provider handles sign-in for all your tools) gives you one place to enforce policy, one place to audit logins, and one place to cut off a departing employee. It also reduces the number of separate passwords your team manages.
- **Run a tight identity lifecycle.** Joiners get access scoped to what their role actually needs, nothing more. Movers (people who change roles) get old access removed and new access granted promptly. Leavers get all accounts disabled on their last day, before they hand in their badge. Deprovisioning must be a checklist on every offboarding, not an afterthought discovered weeks later.
- **Eliminate shared logins.** Every person who has ever known a shared password for a service account is a standing residual-access risk who cannot be individually deprovisioned. Use a shared-credential manager with per-user audit trails, or provision individual accounts for each person.
- **Audit privileged access quarterly.** Who has admin rights in your identity provider, cloud services, and critical tools? Does each person still need them? Remove stale privileges as roles change.

IF YOU'RE BEING TARGETED

If your threat model includes targeted adversaries (hostile governments, law enforcement in high-risk jurisdictions, resourced stalkers, or corporate espionage) standard MFA hygiene is the floor, not the ceiling.

- **Hardware security keys become non-negotiable.** Passkeys stored on a device you carry can be coerced out of you or compromised if the device is seized. A separate FIDO2 hardware key is harder to replicate remotely and immune to phishing. Register two and store the spare in a trusted, separate physical location.
- **Dedicate a hardened email account for your most sensitive work.** Use a provider with a strong legal jurisdiction and a proven track record for resisting third-party data requests. Keep that address strictly compartmentalized, do not use it to sign up for unrelated services. Never access it from a device or network you do not trust.
- **Use email aliases aggressively.** Never give your real address to any service that does not genuinely need it. Aliases contain breaches to single services and prevent adversaries from correlating your activity across contexts.
- **Lock your SIM against carrier-side attacks.** Call your carrier and set a SIM PIN and an account transfer freeze if offered. This raises the bar for SIM-swap attacks considerably. Consider a prepaid or data-only eSIM for contexts where minimizing your real phone number's exposure matters.
- **Audit recovery attack surfaces routinely.** Adversaries attempting account takeover often target recovery flows rather than the account directly because recovery flows frequently have weaker controls. Periodically review and tighten every recovery path on your critical accounts. Remove any phone-number or SMS recovery option where you have a stronger alternative.
- **Consider a clean separate identity for high-risk work.** A distinct email address, a separate dedicated device, and a different persona for investigative work, activism, or legal disputes limits the blast radius of a single compromise and makes it significantly harder for an adversary to link your contexts together.

▲ COMMON MISTAKES

- **Using a "strong" password on multiple accounts.** Complexity does not protect you when the site you used it on was breached and the hash was cracked. Uniqueness is the only property that limits credential-stuffing damage. Strength without uniqueness is theater.
- **Enabling MFA but leaving SMS as the backup recovery method.** If an attacker can SIM-swap your number, the SMS fallback undoes your stronger primary factor. Audit all fallback and recovery options, not just the primary MFA method.
- **Hardening social media before securing email.** People spend time locking down Instagram while the inbox that can reset their Instagram password sits wide open. Email is the master key. Harden it before anything else.
- **Storing recovery codes in cloud notes or cloud drives.** A recovery code saved in a cloud note is only as secure as that cloud account. Print the codes and store them physically, separate from the devices they protect.
- **Using personal email for work or high-risk communications.** Personal accounts typically carry more attack surface, years of connected apps, old recovery options, a lower security bar. Keep contexts separated from the start; retrofitting separation is painful.
- **Leaving dormant accounts alive.** Old accounts with reused passwords and no MFA are soft targets that grant access to whatever data was stored in them and sometimes to connected services. Periodically search for accounts you no longer use and delete them, or at minimum rotate their passwords and revoke connected app access.

PILLAR 2 · DEVICE

Your Devices

Why it matters

Your phone knows where you sleep, who you talk to, and what you believe. Your laptop holds contracts, source code, and years of correspondence. When a device is lost, stolen, or seized (and eventually one will be) the question is not "did something bad happen?" but "how much did the adversary get?"

Most device compromises do not involve exotic exploits. They exploit one of three failures: software that is months out of date and carries a known vulnerability, a device with no encryption so anyone with a USB cable can image the storage, or a screen lock so weak it is bypassed in under a minute. Nation-state actors use these exact same gaps, they reach for the cheap, reliable method first.

The threat is also physical. Devices are seized at borders, grabbed during a raid, or copied overnight in a hotel room while you sleep ("evil maid" attack, an adversary with brief physical access installs persistent malware or reads unencrypted storage). A device left in a car in a bad neighborhood is a realistic vector for a journalist or activist, not just a cautionary tale.

What good looks like

A well-secured device posture has these properties:

- Every device installs security patches automatically, within days of release.
- Full-disk or full-device encryption is active. The storage is unreadable without the unlock credential.
- A strong screen lock (ideally a six-digit minimum PIN or a passphrase) activates after a short idle timeout.
- Biometrics are set up for convenience but can be bypassed by passcode when required.
- Remote find-and-wipe is enrolled and you know how to trigger it.
- Apps come from the platform's official store or a known, audited source. Sideloaded apps are the exception, not the default.
- Devices that no longer receive security updates are retired or air-gapped from sensitive use.
- Smart-home and IoT devices (cameras, smart speakers, door locks) live on a separate network segment from your working devices.

Do this

Automatic updates first

This is the single highest-return action in this entire chapter. Most successful device compromises exploit vulnerabilities that already have patches, the device just was not updated. Enable automatic updates everywhere: the operating system, the app store, the browser, and your router's firmware. Do not wait for a convenient moment; attackers do not wait either.

Verify encryption

Modern iPhones and most Android phones encrypt storage by default when you have a screen lock set. Verify this is actually on (check Security settings. For laptops: macOS provides FileVault (full-disk encryption tied to your login password); Windows provides BitLocker on Pro/Enterprise editions and Device Encryption on compatible Home editions. Enable whichever applies. Store the recovery key somewhere safe and offline) losing it means losing access to your own drive. A printed copy in a locked drawer beats a file named "recovery-key.txt" on the same machine.

Screen lock and biometrics

Set a PIN or passphrase with at least six characters, longer is meaningfully better. Enable biometrics (Face ID, fingerprint) for daily convenience, but be aware of two specific scenarios where you should prefer the passcode:

- **Border crossings and checkpoints.** In many jurisdictions, authorities can compel you to unlock a device with biometrics more easily than they can compel you to disclose a passcode. Consult the EFF's border-crossing guides and speak with a lawyer before travel to understand your rights. Before crossing, power off or enable a "lockdown" mode that disables biometric unlock.
- **Protests and public confrontations.** Biometric unlock can be applied by force. If you face a situation where someone could physically direct your face or finger at the device, know how to trigger emergency/lockdown mode quickly on your specific phone model.

Set auto-lock to activate after two minutes of idle time or less.

Find My and remote wipe

Enroll every device in your platform's find-and-locate service (Apple's Find My, Google's Find My Device). Know how to initiate a remote wipe from a browser on someone else's device. Remote wipe is the last resort (it destroys the data rather than recovering it) but it is the right move when a device is definitely gone and contains sensitive information. Practice the steps once so you are not learning them in a crisis.

App hygiene

Install apps only from your platform's official store unless you have a specific, justified reason to do otherwise. Review app permissions when you install: a flashlight app does not need contacts access. Periodically audit installed apps and delete what you no longer use. Fewer apps mean fewer attack surfaces and fewer apps silently phoning home.

Retire end-of-life devices

An end-of-life device (one whose manufacturer no longer issues security updates) is a liability. It cannot be patched against newly discovered vulnerabilities. For phones: most manufacturers support a device for three to five years. When support ends, treat it as insecure. For laptops: operating system support timelines vary; when your OS stops receiving updates, either upgrade the OS (if supported on the hardware) or replace the machine. Do not use an end-of-life device for anything sensitive.

Isolate IoT and smart-home devices

Smart TVs, thermostats, security cameras, and voice assistants typically receive poor security support and share your home network. Put them on a separate Wi-Fi network (most consumer routers support a "guest" network or VLAN for this purpose). This way, a compromised smart speaker cannot reach your laptop or your NAS (network-attached storage). Treat your phone and laptop as a separate tier from your IoT devices.

- ☐ Enable automatic OS and app updates on every device.
- ☐ Confirm full-disk or full-device encryption is active (FileVault, BitLocker, or Android/iOS built-in).
- ☐ Store the disk-encryption recovery key somewhere secure and offline.
- ☐ Set a strong screen lock PIN or passphrase (six or more characters); set auto-lock to two minutes.
- ☐ Enroll in Find My / Find My Device and test remote-wipe steps.
- ☐ Audit installed apps; remove what you do not use; check permissions.
- ☐ Note the security-update end-of-life date for each device you rely on.
- ☐ Move smart-home and IoT devices to an isolated network segment.
- ☐ Know how to disable biometric unlock on your phone quickly.

Two paths: choose your device posture

Option A, Hardened mainstream

Best for: Most people in this audience. You want strong security without rebuilding your workflow from scratch.

iOS: Enable all updates, strong passcode, Lockdown Mode for elevated-risk periods (it restricts attack surface by disabling features like link previews, wired accessories, and complex web rendering, enable it before travel to risky environments or during active targeting).

Android: Prefer a device from a vendor with a strong update track record and a long committed support window. Enable full encryption, strong lock, and restrict app installs to the Play Store. Disable features you do not use (Bluetooth when not in use, NFC if unused).

Trade-off: Apple and Google still have access to metadata, backup data, and some telemetry. You accept that in exchange for convenience and a well-maintained security model.

Option B, De-Googled / GrapheneOS

Best for: People who want to minimize the platform vendor's data access, and who are willing to spend real time on setup and maintenance.

GrapheneOS is a hardened Android distribution that removes Google services, adds a stronger permission model, and provides per-app network and sensor controls. It supports sandboxed Google Play for apps that require it, so app compatibility is better than older de-Googled approaches.

Trade-off: Installation requires comfort with firmware flashing. Some mainstream apps behave differently or require configuration. Push notifications and some banking apps need the sandboxed Play layer. Expect a few hours of setup and an ongoing maintenance commitment.

Which is for you: If the question "who is my threat model, and does it include the platform vendor?" has a clear yes, Option B is worth the effort. If the answer is "primarily phishing, malware, and opportunistic theft," Option A hardened properly is excellent and leaves more time for the other six pillars.

SCALING TO YOUR TEAM OR ORGANIZATION

Even a five-person organization benefits from lightweight device management. Key moves:

- **MDM (Mobile Device Management):** A lightweight MDM lets you enforce a baseline config (screen lock, encryption, OS version minimums) on company-issued phones and laptops without turning your org into an enterprise IT shop. Many MDM tools offer small-team tiers.
- **Device inventory:** Maintain a simple spreadsheet (or a tool with an asset-tracking feature) listing every device that touches company data: device type, owner, OS version, encryption status, last update date, end-of-life date. Review it quarterly.
- **Compliant-device gate:** Configure your identity provider so that access to company email, code repositories, or internal tools requires a device that meets minimum standards (up-to-date OS, encryption confirmed). Most modern SSO providers support device compliance policies.
- **Separation of personal and work:** Require separate accounts or profiles for work use. "Use your personal phone for Signal and your work profile for Slack" is a reasonable minimum. Full separation (a dedicated work device) is better for sensitive roles.
- **Offboarding:** When someone leaves, revoke their device from MDM and remote-wipe any company profile. Do this the same day, not the next week.

IF YOU'RE BEING TARGETED

At nation-state or sophisticated-adversary threat levels, your device posture escalates:

- **GrapheneOS or iOS Lockdown Mode as the default baseline**, not just for travel. These are not overcautious, they eliminate entire classes of remote compromise that affect standard configurations.
- **Dedicated travel device.** Carry a separate phone with minimal data, accounts, and apps when crossing borders, attending high-risk events, or traveling to adversarial jurisdictions. Leave your daily driver at home.
- **Carry as little as possible.** The travel device should contain only what you genuinely need for the trip. Log in to accounts on-site rather than storing credentials on the device.
- **Power off at checkpoints.** Power the device off before reaching a border checkpoint, not just lock it. Full power-off means decryption keys are not in RAM, and certain attacks that work on a locked device do not work on a powered-off one. Know your device's behavior, some phones re-enable biometrics after a reboot; enter your passcode first to set the "before first unlock" state.
- **Physical tamper awareness ("evil maid" attack).** If a device is out of your control for any period (hotel room, repair shop, customs inspection) treat it as potentially compromised before using it for sensitive work. At a minimum, power off before surrendering it. If you have reason to suspect physical access occurred, do not use the device for sensitive communications until it can be inspected or replaced. Tamper-evident stickers or nail polish on port covers can reveal whether a device was physically accessed, though they do not prevent access.
- **Consult counsel before border travel.** The legal landscape for device searches varies significantly by country and changes over time. Organizations like the EFF publish current guidance on US border rights. In other jurisdictions, local legal advice is essential before travel.

▲ COMMON MISTAKES

- Assuming encryption is on without verifying it in settings, especially on older Android devices where it may require manual activation.
- Storing the disk-encryption recovery key in cloud storage on the same device, so it is compromised along with the machine.
- Using a four-digit PIN and treating that as "a strong screen lock." Six digits is the floor; a passphrase is better.
- Keeping devices that stopped receiving security updates in active use for sensitive work because they "still work fine."
- Leaving smart-home devices on the same network as working devices, creating a flat network where a compromised camera can reach a laptop.
- Confusing "cloud backup is on" with "the device is protected." Backups protect your data from loss; encryption protects it from an adversary who has your hardware.
- Forgetting to revoke remote-wipe enrollment when you sell or recycle a device, the new owner could trigger a wipe, or your account remains linked to hardware you no longer control.

PILLAR 3 · APPLICATION & WORKLOAD

Apps & Software

Why it matters

Every app on your device is code running with some level of trust. That trust is often wider than the app's actual job. A flashlight app that reads your contacts. A notes app that silently requests microphone access. A browser extension that can read every page you visit, including your bank, your email, your password manager. These are not hypothetical edge cases; they are the normal state of consumer software.

The attack surface here is three-layered. First, the app itself might be malicious (a fake or lookalike app engineered to impersonate a legitimate one). Second, a legitimate app might be compromised after you install it, either via a poisoned update pushed through the developer's own supply chain or through a malicious dependency bundled in by the developer's build tools. Third, an app might be entirely benign but ask for) and receive, permissions far beyond what its job requires, meaning any future vulnerability becomes a much bigger breach.

There is a fourth vector that most people miss entirely: OAuth grants. When you click "Sign in with Google" or "Allow this app to access your calendar," you are handing a third-party application a live credential that works even after you change your password. Many users have dozens of forgotten OAuth grants dangling off their primary accounts, any one of which could be abused if the third-party app is ever compromised or sold to a malicious buyer.

What good looks like

A well-managed app posture looks like this: your device runs only software you consciously chose to install, every app has the minimum permissions needed for the feature you actually use, your browser profile is lean (few extensions, each reviewed), and you can name every third-party application that has standing access to your Google, Microsoft, or Apple account. Updates apply automatically. Apps you no longer need are gone, not forgotten-about, gone.

Supply-chain hygiene means you treat an install prompt with the same mild suspicion you apply to an unexpected email attachment. You install from the canonical source: the official app store, the developer's own website with a verified checksum, or your organization's managed software catalog. You do not install from a third-party mirror, a link in a chat message, or a "helpful tool" offered by a stranger during a support call.

Do this

Install discipline

- ☐ Install apps only from official app stores or the developer's own verified site, never from third-party mirrors, links in messages, or files emailed to you.
- ☐ Before installing, check the publisher name carefully. Fake apps impersonate real ones with near-identical names and icons; the publisher field is harder to spoof.
- ☐ On Android, keep "Install from unknown sources" disabled unless you have a specific, understood reason to enable it, and re-disable it immediately after.
- ☐ Treat any "support tool," remote-access app, or screen-sharing software offered to you by an unsolicited caller as malicious until proven otherwise, legitimate support teams do not cold-call you and ask you to install things.

Permissions: minimize and review

- ☐ On your phone, go to Settings → Privacy → Permission Manager (iOS: Settings → Privacy & Security) and review location, microphone, camera, and contacts access. Revoke anything that isn't obviously necessary for a feature you actually use.
- ☐ Set location to "While Using" rather than "Always" by default. Persistent background location is almost never required.
- ☐ Deny permissions on first install; grant them only if the app stops working without them, and only at that moment.
- ☐ Revisit permissions quarterly, apps update and sometimes quietly request new ones.

Remove what you don't use

- ☐ Audit installed apps now: delete anything you haven't opened in 90 days.
- ☐ Fewer apps means a smaller attack surface, an uninstalled app cannot be compromised.
- ☐ On desktop, uninstall rather than just removing shortcuts. Check your package manager or app list for forgotten installs.

Keep apps updated

- ☐ Enable automatic updates for apps on both phone and desktop. Updates frequently patch actively exploited vulnerabilities.
- ☐ When a major version update is available and automatic update hasn't triggered, install it promptly rather than deferring.

Browser hardening and extensions

Your browser is one of the highest-privilege apps on your device, it runs code from every site you visit and, if extensions are present, those extensions can read and modify every page you load, including your webmail and online banking.

- ☐ Keep your browser count low: ideally one primary browser, kept on the latest stable release.
- ☐ Audit extensions. Open your browser's extension manager and remove everything you didn't explicitly install or don't actively use. This includes toolbars, "helpers," and anything you don't recognize.
- ☐ Before installing an extension, check the permissions it requests. An extension asking to "Read and change all your data on the websites you visit" has effectively full read access to your online life.
- ☐ Prefer extensions from large, audited publishers or open-source projects with active maintenance. Avoid one-person or abandoned extensions, these are frequently acquired and repurposed for data harvesting.
- ☐ Enable your browser's built-in phishing and malware protection.

OAuth grants: review and revoke connected apps

OAuth (Open Authorization) is the protocol behind every "Sign in with..." or "Allow access to..." prompt. When you approve one, the third-party app receives a token (a credential) that lets it act on your account, sometimes indefinitely. Review these periodically and revoke anything you don't recognize or no longer use.

- ☐ **Google:** Go to myaccount.google.com/permissions (<https://myaccount.google.com/permissions>) and remove apps you don't recognize or no longer use.
- ☐ **Microsoft:** Go to myapps.microsoft.com (<https://myapps.microsoft.com>) or the My Account portal → Apps & services.
- ☐ **Apple:** Settings → your name → Password & Security → Apps Using Apple ID.
- ☐ When in doubt, revoke the grant. Legitimate apps will simply ask again the next time you use them.
- ☐ Prefer narrow-scope grants. When an app offers "read-only" access as an option, take it. Don't grant write access unless the app explicitly needs it.
- ☐ Treat OAuth grants from apps you no longer use as stale credentials, revoke them the same way you would rotate a forgotten password.

SCALING TO YOUR TEAM OR ORGANIZATION

At the organizational level, your biggest risks are shadow IT (staff connecting personal apps to company accounts without review) and unchecked OAuth sprawl (a long tail of third-party integrations no one remembers approving).

- **Vet SaaS before adopting it.** Build a lightweight intake process: who's requesting it, what data does it touch, is there a BAA or DPA if needed? Even a shared spreadsheet beats no process.
- **Control OAuth app consent.** In Google Workspace Admin and Microsoft Entra ID (Azure AD), you can require admin approval before any user grants a third-party app access to organizational data. Turn this on. It is one of the highest-value admin controls available for free.
- **Maintain a software inventory.** Know what is installed on company devices and what has access to company accounts. An endpoint management tool (MDM) helps; so does a quarterly manual review for smaller teams.
- **Basic app allowlisting** on managed devices (only pre-approved software can run) is the gold standard. It is operationally demanding for small teams, but even a documented "approved software list" combined with admin rights removed from standard user accounts closes the most common malware installation vector.
- **Offboard completely.** When someone leaves, revoke their OAuth tokens and review which third-party apps they personally authorized that still connect to shared accounts.

IF YOU'RE BEING TARGETED

A targeted adversary may attempt to install monitoring tools on your device through a malicious app, a poisoned update, or a social-engineering scenario ("here's a tool you need to verify your identity / join this call / view this protected document"). The goal is persistent access, a foothold that survives reboots and survives you changing your passwords.

- **Minimize your app footprint aggressively.** The fewer apps installed, the smaller the surface. On a high-risk device, challenge every single app: is this truly necessary on this device?
- **Scrutinize every permission and every OAuth grant.** An adversary-controlled app requesting contacts and location is a surveillance tool. Treat any app you didn't deliberately choose with suspicion.
- **Never install software offered during an unexpected interaction,** an unsolicited call, a chat message from an unverified contact, an email from someone you don't know. Legitimate organizations do not operate this way.
- **Open risky documents in a sandbox or throwaway environment.** If you receive a file from an untrusted or unknown source (a PDF, a Word document, a spreadsheet) open it in an isolated environment rather than your primary device. Options include a separate virtual machine, a Chromebook in Guest mode, or an online document viewer that doesn't execute macros.
- **Review installed apps after any suspicious incident.** If a device behaved oddly, look for apps you didn't install. On iOS, look for MDM (Mobile Device Management) profiles under Settings → General → VPN & Device Management, legitimate MDM profiles come only from your employer; anything else is alarming.
- **Consider a separate device** for your highest-risk communications or work. Compartmentalization limits blast radius: if the primary device is compromised, the work on the secondary one is not exposed.

▲ COMMON MISTAKES

- Granting an app every permission it asks for at install time, without reading the list.
- Forgetting about OAuth grants after approving them, "I signed in with Google once two years ago" leaves a live token sitting there.
- Installing the same extension in every browser profile "for convenience," multiplying the exposure.
- Treating app store presence as a security guarantee. Malicious apps do appear in major stores; publisher name and review date matter more than mere listing.
- Leaving apps on a device after you stop using them, "I might need it someday" is how forgotten software becomes an unpatched, persistent risk.
- Running a browser extension from an abandoned or unknown publisher because it "seems useful." Extensions are acquired and repurposed regularly; check when the last update was and who currently maintains it.

PILLAR 4 · DATA

Your Data

Why it matters

Most security advice tells you to defend the perimeter, lock the door, patch the software, don't click the link. Data-layer security asks a harder question: if all of that fails and an adversary is already inside, what do they find? This is the inside-out architecture principle: design your defenses around what matters most first, not around the edges of a network.

The threats are ordinary and common. Ransomware encrypts your files and demands payment (a tested backup makes the ransom pointless. A stolen laptop exposes everything on it) encryption at rest means the disk contents are unreadable without your credentials. An overshared Google Drive link set to "anyone with the link" sits accessible long after the project that created it ended. A journalist's source is identified because the document contained metadata (author name, GPS coordinates in a photo) that nobody thought to strip.

Data minimization is the civilian equivalent of enterprise data-loss prevention: the simplest way to protect data you don't have is to not keep it. Don't retain what you no longer need. Deleted data cannot be subpoenaed, stolen, or leaked.

What good looks like

You have identified your crown jewels. Those assets are encrypted at rest, backed up in at least two independent locations, and shared only with people who have a genuine need. You have tested a restore recently enough to know the backups work. Devices that leave your hands have been wiped properly. You can name which cloud folders are broadly shared and which are locked down, and you know how to check.

In transit, sensitive files travel over encrypted connections. For sensitive conversations you use a messaging tool that provides end-to-end encryption, E2EE (meaning only the sender and recipient can read the message; the provider cannot, and cannot be compelled to hand over content it can't read).

Do this

Identify your crown jewels

- ☐ Name the three to five assets that would cause the most harm if lost, exposed, or destroyed: source materials, financial records, private correspondence, legal documents, credentials backups.
- ☐ Locate where each currently lives, device, cloud folder, email attachment, physical only. You cannot protect what you haven't located.
- ☐ Apply stricter controls to these assets first, before spending effort on everything else.

Back up with the 3-2-1 rule, and test a restore

The 3-2-1 backup rule: keep **3** copies of important data, on **2** different media types, with **1** copy offsite (or in a separate cloud account). A backup you've never tested is an assumption, not a safety net.

- ☐ Set up automated backups, manual backups don't happen consistently under stress.
- ☐ Verify the offsite copy is current: check the last backup date, not just that a service is running.
- ☐ Test a restore. Open a backed-up file from your backup destination to confirm recovery actually works.
- ☐ Keep at least one backup disconnected from your main device and account (an air-gapped external drive or a separate cloud account) so ransomware or account compromise cannot reach all copies simultaneously.

Choose your data storage path

Option A, E2EE cloud storage + 3-2-1 backups

What it is: A cloud provider that offers genuine end-to-end encryption, your files are encrypted on your device before upload, and the provider holds no key. You complement this with automated backups to a second cloud account and a local or external drive.

Best for: Users who work across multiple devices, travel frequently, or need remote access to files. Lower operational burden once set up. Good resilience against device loss or failure.

Trade-offs: Dependent on the provider's continued operation and business model. Monthly cost. Losing your encryption key or recovery code means losing your data, no provider reset can help. Requires trusting the provider's audit history and cryptographic implementation.

Verify E2EE is real: The provider should publish independently audited cryptographic documentation. "Encrypted" without "end-to-end" typically means the provider holds the key and can be compelled to hand it over.

Option B, Local-first / self-hosted + encrypted external drives

What it is: Your primary data lives on devices you control. Backups go to encrypted external drives (using full-disk encryption with a strong passphrase) and optionally to a self-hosted server or NAS (network-attached storage) on your own hardware. No third-party cloud provider holds your files.

Best for: Users with high sensitivity to provider-disclosure risks (journalists, activists, researchers with confidential sources) or those who want zero dependency on external services.

Trade-offs: Higher operational burden, you are responsible for physical security of drives, hardware failure, and disaster recovery. No automatic remote access. Encrypted drives that fail without a second copy are gone permanently. Keeping the offsite copy current is human-driven, not automatic.

Physical security matters: An encrypted drive left on a desk defeats the purpose. Offsite means a genuinely separate location, not a second drawer.

Encryption, sharing hygiene, and data minimization

- ☐ For sensitive archives, add a second layer beyond full-disk encryption: an encrypted container with a long, unique passphrase stored in your password manager.
- ☐ Audit shared links in your cloud storage. Search for items shared "with anyone who has the link" and restrict or remove sharing where the need has passed.
- ☐ Default to sharing with specific people, not open links. Set expiration dates on shared links where your platform supports it.
- ☐ Delete data you no longer need. Finished-project drafts, scans of shredded documents, old correspondence, delete rather than archive indefinitely.
- ☐ Secure deletion on SSDs is technically complex. Use full-disk encryption (so deleted data is already encrypted garbage) combined with your platform's erase features before discarding a device.

SCALING TO YOUR TEAM OR ORGANIZATION

For small organizations, data security failures typically come from three places: overly permissive shared drive folders, data that walks out the door when someone leaves, and no policy on what to keep or delete.

- **Least-privilege shared drives.** Audit access in Google Workspace and Microsoft 365. Every folder the whole organization can access is a folder an attacker (or a disgruntled former member) can reach. Apply permissions by role and project, not by default-open.
- **Light data-loss prevention (DLP).** Both Workspace and M365 include basic DLP controls, rules that flag or block external sharing of content matching patterns like credit card numbers. Even basic rules are worth enabling.
- **Retention and deletion policy.** Decide how long different data types are kept and document it. A written policy also demonstrates good-faith compliance if you face a legal request. Then actually enforce it.
- **When staff leave:** immediately revoke access, transfer ownership of shared files to a manager, and archive data from accounts being closed. Do this on day one of departure. Data in a closed account is often irretrievable.

IF YOU'RE BEING TARGETED

A targeted adversary is after specific data: your sources, communications, unpublished work, legal strategy. They may subpoena cloud providers, seize physical devices, or compromise accounts. Defense means making that data hard to find, hard to read, and hard to attribute.

- **Encrypt sensitive archives separately.** A separately encrypted archive, protected by a passphrase stored only in your head or on a hardware key you control, is inaccessible even if your cloud account is compromised or a provider complies with a legal order.
- **Compartmentalize data.** Keep active sensitive projects on a separate device, or in an encrypted container that is closed and locked when not in use. Mix sensitive and routine data as little as possible.
- **Strip metadata before sharing.** Photos embed GPS coordinates and timestamps in EXIF metadata. Office documents record author names and revision history. Verify your tool actually removes metadata rather than just hiding it before sending anything that could reveal identity.
- **Securely wipe devices before disposal.** A factory reset is not a secure erase for flash storage. Use your platform's built-in encrypted erase. For physical drives without built-in tooling, physical destruction of the storage medium is definitive.
- **Protect sources.** If you work with confidential sources, don't retain identifying information you don't need. Discuss operational security expectations before they share anything sensitive. The EFF's Surveillance Self-Defense guide (<https://ssd.eff.org>) covers journalist and activist data practices in depth.

▲ COMMON MISTAKES

- Treating a sync service as a backup, if you delete a file from a synced folder, the deletion propagates everywhere. Sync and backup are different things.
- Never testing a restore. The backup is not real until you've confirmed recovery from it.
- Confusing "encrypted" with "end-to-end encrypted." Standard cloud encryption protects against external attackers and other customers, not against the provider, not against legal compulsion directed at the provider.
- Leaving overshared links active indefinitely. That draft-budget link from two years ago is still live unless you explicitly removed it.
- Sharing photos that contain GPS location metadata with people who don't need to know your location, or your source's.
- Storing the only backup copy in the same location as the original. A house fire or ransomware attack reaches both simultaneously.

PILLAR 5 · NETWORK & ENVIRONMENT

Network & Environment

Why it matters

Your network is the pipe everything flows through. Every device in your home or office (laptops, phones, thermostats, cameras, speakers) communicates over it. An attacker who establishes a foothold on your network can intercept unencrypted traffic, probe every device attached to it, and pivot from a cheap smart plug to your laptop. Most consumer routers ship from the factory with default administrator credentials (often something like `admin / admin`) that are publicly documented and identical across thousands of units. Anyone who can reach that admin panel (whether they're on your network, in your building, or exploiting a remote-management feature you left enabled) can reconfigure everything silently.

Public Wi-Fi is a different category of threat: you have no knowledge of or control over the equipment. A malicious hotspot with a plausible name is trivially easy to stand up. Even a legitimate network may have a negligent or compromised operator. Treat any network you didn't personally configure as potentially hostile, not because every coffee shop is run by spies, but because operating that way costs you almost nothing and protects you against the real cases.

The good news: most network attacks exploit default configurations. Ten minutes of router housekeeping eliminates the most common exposure.

What good looks like

- Your router has a unique, strong admin password, not the factory default. WPA3 (preferred) or WPA2-AES is enabled with a strong passphrase. WEP and WPA-TKIP are broken; never use them.
- Firmware is current. WPS (Wi-Fi Protected Setup, the push-button pairing feature) is disabled. Remote administration from outside your network is off.
- IoT devices (smart TVs, cameras, thermostats, speakers, game consoles) and guests are on a separate network segment, isolated from your main devices.
- Your DNS queries are encrypted so your ISP cannot see every hostname you visit.
- You use end-to-end-encrypted (E2EE) messaging for sensitive conversations, meaning only you and the recipient can decrypt the content, not the platform or provider.
- You have a clear, accurate model of what a VPN does and does not do, and use one only where it genuinely helps.

Do this

Lock down your router

Log into your router's admin panel, usually found at `192.168.1.1` or `192.168.0.1` in your browser; the address and default credentials are printed on the device label. Do all of the following in a single session:

- ☐ Change the admin username and password to something unique and strong. Save it in your password manager, you will not remember it otherwise.
- ☐ Set Wi-Fi encryption to WPA3 if available, or WPA2-AES. Disable WPS, it has well-documented PIN-cracking vulnerabilities and is not needed for normal device setup.
- ☐ Change the Wi-Fi passphrase from the factory default to a strong one. Don't use the printed key.
- ☐ Disable remote management and remote administration. You almost certainly do not need to access your router's admin panel from outside your home.
- ☐ Check for a firmware update and apply it. Enable automatic firmware updates if your router supports it. Bookmark a quarterly firmware-check reminder as a calendar task (see Pillar 6).
- ☐ Change the default network name (SSID) to something that doesn't include your name, apartment number, or anything identifying.

Create a guest network for IoT and visitors

A guest network is the civilian version of network segmentation, a foundational enterprise security concept. The principle: devices that have no reason to communicate with your laptop should not be able to reach it. Smart devices are notoriously under-maintained by their manufacturers; a compromised camera or speaker on your main network is a beachhead. Move it to the guest network, where it can still reach the internet but cannot reach your other devices.

- ☐ Enable the guest network in your router settings and give it its own strong, separate passphrase.
- ☐ Move all IoT devices to the guest network: smart TVs, speakers, cameras, thermostats, robot vacuums, game consoles, anything that isn't a computer or phone you actively use for sensitive work.
- ☐ Hand guests the guest-network passphrase, not your main one.
- ☐ Enable "client isolation" on the guest network if your router offers it. This prevents devices on the guest segment from seeing each other as well as your main devices.

Treat public Wi-Fi as hostile

This doesn't mean you can never use a coffee-shop network. It means you operate with the assumption that someone on the same network may be watching traffic, and the hotspot itself may be spoofed or compromised.

- ☐ Use HTTPS everywhere. The padlock (or "https://" prefix) in your browser means the connection to the site is encrypted. Never enter credentials on an HTTP site; modern browsers warn you, pay attention to those warnings.

- ☐ Avoid accessing sensitive accounts (banking, primary email, work systems) on public Wi-Fi unless you have a trusted VPN active (see below).
- ☐ Consider using your phone's cellular connection as a personal hotspot instead of joining an unknown Wi-Fi network entirely.
- ☐ Enable your device's built-in firewall when on untrusted networks. On macOS: System Settings → Network → Firewall. On Windows: Windows Security → Firewall.

Encrypt your DNS

DNS (Domain Name System) translates domain names like "example.com" into IP addresses. By default, these queries travel unencrypted, giving your ISP (and any passive observer on the path) a log of every hostname you look up, even when the site uses HTTPS.

Encrypted DNS (DoH (DNS over HTTPS) or DoT (DNS over TLS)) wraps those queries in encryption. Most modern operating systems and browsers support it natively. Enable it in your OS network settings and your browser's privacy or security settings, pointing to a resolver you trust. Some routers support it network-wide, protecting every device at once. A product-specific edition covering resolver selection criteria is planned as a follow-up.

- ☐ Enable DNS over HTTPS in your browser's privacy or security settings.
- ☐ Enable encrypted DNS at the OS level: macOS Ventura and later, Windows 11, iOS 14+, and Android 9+ all support it natively.

Understand VPNs, and use one only where it helps

A VPN (Virtual Private Network) creates an encrypted tunnel between your device and a VPN server. All traffic from your device to that server is encrypted and appears to originate from the server's IP address. That is the complete picture of what it does.

What it does not do: it does not make you anonymous. Sites you visit still see a connection (from the VPN's IP address). The VPN provider can see all your traffic. You have simply moved your trust from your ISP to the VPN company. A VPN does not protect you from malware, phishing, account takeover, or data the destination site collects. It does not hide metadata from your ISP; they see that you're connected to a VPN, for how long, and how much data you transfer.

Where a VPN genuinely helps: on hostile public Wi-Fi, it encrypts traffic from local eavesdroppers. In jurisdictions with invasive ISP data-retention laws, it reduces what your ISP can log. For hiding your home IP address from a specific destination site.

If you use one, choose a reputable provider with a published, independently audited no-logs policy. A product-specific edition covering selection criteria is planned.

- ☐ Don't install a VPN on the assumption it makes you "private" or "secure", understand the specific threat it addresses.

- ☐ If you use a VPN, pick one with an audited no-logs policy from a provider in a jurisdiction with strong privacy laws.
- ☐ Connect before joining any public or untrusted Wi-Fi network.

Use end-to-end-encrypted messaging

E2EE (end-to-end encryption) means only you and the intended recipient hold the keys to decrypt a message. The platform, the provider, and anyone intercepting the transmission cannot read the content. Standard SMS, unencrypted carrier calls, and many popular messaging apps are not E2EE. Signal is the widely-accepted benchmark: open source, independently audited, and E2EE by default for messages, calls, and video.

- ☐ Use Signal (or an equivalent app that is open source and independently audited for E2EE) as your default for sensitive conversations.
- ☐ Enable disappearing messages on sensitive threads so a seized or lost device doesn't expose your message history.
- ☐ Verify safety numbers with high-value contacts to confirm you're talking to who you think you are, and not an impersonator or man-in-the-middle.

SCALING TO YOUR TEAM OR ORGANIZATION

Network segmentation is foundational at the organizational level:

- Maintain at minimum three separate network segments: staff or corporate devices, guest or visitor access, and IoT or infrastructure devices. A device on one segment should not be able to initiate connections to another segment without an explicit firewall rule permitting it.
- Apply ZTNA (Zero Trust Network Access) principles in plain language: "only the right people reach the right systems." A contractor does not need access to your internal file server. Your conference-room display does not need to reach your accounting database. Start by mapping what talks to what, then remove any connection that has no business justification.
- Keep your router or firewall firmware on a patching schedule with a named owner. A managed firewall appliance with active support is appropriate for organizational use, consumer routers are not designed for the threat model a small organization faces.
- Consider logging outbound DNS queries and flagging unusual or newly-registered domains. This is a low-cost early-warning indicator for compromised devices calling home.

IF YOU'RE BEING TARGETED

- Treat every network (including your home network) as potentially hostile. A sophisticated adversary can compromise a router remotely, via firmware exploits, or physically if they have access to your premises.
- Conduct all sensitive conversations exclusively over E2EE channels. Assume any non-E2EE channel is readable by an adversary with ISP-level or platform-level access.
- A VPN does not hide metadata from your ISP or from a national-level observer, they see that you're connected to a VPN, the timing, and the data volume. Timing analysis can reveal patterns even without content.
- When genuine network anonymity is required (where even your VPN provider knowing your identity is unacceptable) Tor (the Tor network) exists. It routes traffic through multiple volunteer-operated relays so no single node knows both your identity and your destination. Tor has real trade-offs: it is significantly slower, and some sites block Tor exit nodes. It is not appropriate for everyday browsing; use it when your threat model specifically warrants it.
- On sensitive travel, consider cellular data over hotel or conference Wi-Fi. Do not plug in unknown USB devices, cables, or charging adapters, "juice jacking" via modified cables is a documented attack vector.

▲ COMMON MISTAKES

- Never changing the router's admin password, factory defaults are publicly documented and often identical across thousands of units of the same model.
- Leaving WPS enabled because "it makes connecting easier." WPS has known PIN-brute-force vulnerabilities; disable it.
- Believing a VPN provides privacy or security. It relocates trust from your ISP to the VPN provider, it does not eliminate the trust requirement.
- Putting IoT devices on the main network because setup was faster. A compromised smart device on your main network can reach every other device on it.
- Using SMS as an MFA channel and treating your phone number as a secure credential. Phone numbers can be hijacked via SIM-swapping. Prefer authenticator apps or hardware security keys. (See Pillar 1 for detail.)
- Skipping router firmware updates for years. Routers run continuously and accumulate unpatched CVEs just like any other software.

PILLAR 6 · AUTOMATION & ORCHESTRATION

Automation: Make Security Automatic

Why it matters

Security measures that depend on willpower fail. They fail during busy stretches, during travel, during stress, exactly when adversaries are most likely to probe. The strongest intentions don't survive "I'll do that later." The most important insight in this pillar: the goal is not discipline. It is architecture. Design your systems so the secure path is the automatic path, and the insecure path requires deliberate effort to choose.

Enterprise security teams call this SOAR (Security Orchestration, Automation, and Response), automated playbooks that detect and respond to threats without a human in the loop. The civilian version is simpler: turn on every available automatic protection and replace manual habits with scheduled events. Remove decisions that shouldn't need to be decisions.

Automation also enforces consistency. A human following a checklist is probabilistic. A configured system is deterministic. You want your security controls to be deterministic.

What good looks like

- Devices, operating systems, and apps update automatically, patches apply before you think to schedule them.
- Encrypted backups run on a schedule with no manual trigger required.
- Your password manager generates a unique, random password for every account and fills it automatically, no memorization, no reuse.
- Breach alerts are active, so you hear about compromised credentials before an attacker exploits them.
- Devices lock after a short idle period; sessions expire on their own.
- Quarterly security reviews are calendar events, not intentions.

Do this

Turn on automatic updates, everywhere

This is the single highest-return security action available to most people. The vast majority of successful malware and ransomware attacks exploit known, already-patched vulnerabilities on systems that haven't applied the update. "Keep firmware updated" appeared in Pillar 5 for your router; this is the same principle applied universally and without exception.

- ☐ Enable automatic OS updates on your phone. On iOS: Settings → General → Software Update → Automatic Updates. On Android: Settings → System → System Update, then enable auto-download and install.

- ☐ Enable automatic updates on your computer, the OS, the browser, and all installed applications.
- ☐ Enable automatic app updates through your app store. Stagnant apps accumulate unpatched CVEs.
- ☐ Confirm your router is on automatic firmware update if the model supports it (see Pillar 5). If it doesn't, add a quarterly firmware-check reminder to your calendar.
- ☐ Stop deferring update prompts. Each week you delay an available patch is a week of unnecessary exposure to a documented, exploitable vulnerability.

Set up automatic encrypted backups

A backup that isn't automatic is not reliable, you will skip it eventually, and the skip will happen at the worst time. The 3-2-1 backup rule (three copies, two different media types, one offsite) was covered in Pillar 4. The automation piece: your primary backup should be continuous or nightly, triggering without any input from you.

- ☐ Enable your OS's built-in scheduled backup to an encrypted external drive. On macOS, this is Time Machine. On Windows, this is File History or Windows Backup.
- ☐ Enable automatic cloud backup for your phone, photos, contacts, and critical documents. On iOS, iCloud Backup runs automatically when plugged in. Enable iCloud Advanced Data Protection if you want end-to-end encryption of your iCloud backup. On Android, Google One backup runs automatically; review what it covers in Settings → System → Backup.
- ☐ Put a monthly five-minute task on your calendar: confirm the backup completed successfully and do a spot-restore of one file. An untested backup is an assumption, not insurance.

Let your password manager do the work

A password manager (covered in depth in Pillar 1) earns most of its security value through automation. It generates a unique, random, strong password for every account, stores it encrypted, and fills it for you. This eliminates password reuse (the root cause of the majority of account takeovers) without requiring your memory or your attention. The manager makes good password hygiene the path of least resistance.

- ☐ Install your password manager's browser extension and enable mobile autofill. It should fill credentials automatically when you reach a login page, you shouldn't need to open the app.
- ☐ Always let the manager generate new passwords. Never type your own. Generated passwords are long, random, and unguessable; your improvised ones (even good ones) are not.
- ☐ Enable breach monitoring in your password manager if it offers it. Most will flag saved credentials when the associated site appears in a known data dump.
- ☐ Store your password manager's recovery kit or emergency sheet separately from the manager itself, in a printed copy in a physical safe, or on a secondary secure device. You cannot use the manager to find the key that unlocks the manager.

Subscribe to breach alerts

Credentials from past breaches circulate on criminal forums for years after the original incident. Proactive monitoring means you hear about a compromised credential before an attacker uses it.

- ☐ Register every email address you actively use at [Have I Been Pwned](https://haveibeenpwned.com) (<https://haveibeenpwned.com>) and turn on notifications. You will receive an email when any of your addresses appears in a newly catalogued breach dump.
- ☐ When you receive a breach notification: change the compromised password immediately (your manager makes this a single-click action), then check whether you reused that password anywhere else and change those too.
- ☐ If your password manager includes a built-in breach dashboard, review it during your quarterly check-in and rotate any flagged credentials.
- ☐ Credit monitoring services often include personal-data breach monitoring beyond credentials. See the Financial chapter for credit-monitoring and credit-freeze guidance.

Auto-lock and auto-logout

A device you walk away from is a device anyone nearby can access. An idle browser session left open is a free entry point for anyone with physical access or a stolen session token. Both are solved the same way: configure systems to lock and log out automatically.

- ☐ Set your computer to lock automatically after five minutes of inactivity, less if your environment warrants it. Also lock manually whenever you step away: Windows Key + L on Windows; Control + Command + Q on macOS.
- ☐ Set your phone's auto-lock to thirty to sixty seconds.
- ☐ Use hardware security keys for your most critical accounts. A security key is a physical device (it plugs into USB or taps over NFC) that performs cryptographic authentication in a single touch. It authenticates faster than typing a six-digit TOTP code and is phishing-resistant by design. It reduces MFA friction while raising your security ceiling. See Pillar 1 for MFA context.
- ☐ Enable automatic session expiry for browser sessions where it's available. On shared or untrusted computers, explicitly sign out of every account before leaving.

Schedule recurring security reviews

Automation handles the repeatable controls. Some things require human judgment: reviewing which apps are connected to your accounts, auditing which devices are trusted, rotating a handful of key credentials on a schedule. The trick is to make these calendar events, not intentions.

- ☐ Create a recurring calendar event ("Quarterly Security Review," 30 minutes) every three months. Agenda: review active sessions and connected apps on email, cloud storage, and other critical accounts; check for unrecognized devices; verify backups completed; review breach-monitoring dashboard; check for overdue software updates. Pillar 7 includes a detailed checklist for this review.

- ☐ Review your password manager's health report at the same time: weak passwords, reused passwords, flagged breaches. Rotate anything that shows up.

SCALING TO YOUR TEAM OR ORGANIZATION

- Automated patching is non-negotiable at the organizational level. An MDM (Mobile Device Management) platform enforces OS and application update compliance across every enrolled device and can report which devices are out of policy. Every unpatched device on your network is a liability you own.
- Conditional access rules enforce security posture automatically, without manual review: "a device that has not been patched within 30 days may not access company email." Configure these rules once; they run continuously without intervention.
- Automate onboarding and offboarding. When someone joins, they should receive exactly the access their role requires, no more, provisioned from a defined template. When someone leaves, all their access should be revoked in a single coordinated action, immediately. Manual offboarding processes that take days or require remembering which systems to update are a documented source of post-separation unauthorized access.
- Schedule automated access reviews on a quarterly basis. Send managers a list of who holds what permissions and ask them to confirm it's still correct. Most identity-governance platforms support this natively. Even a spreadsheet review is better than none, stale permissions accumulate silently and are routinely exploited.
- Automated alerting on log events (covered in Pillar 7) catches anomalies without anyone watching a dashboard all day. Configure it once, then handle what it surfaces.

IF YOU'RE BEING TARGETED

- Automate everything you can so nothing depends on you making the right call while under stress, coercion, or exhaustion. Cognitive load degrades security decisions. Automation doesn't get tired or pressured.
- Automation can itself become a single point of failure. If your automated cloud backup flows into a single account that is then compromised, your backups are compromised with it. Maintain a second backup path that is offline or under a separate identity entirely, an encrypted drive stored offsite, or a backup account that is not linked to your primary identity.
- Enable auto-wipe after a set number of failed unlock attempts on your phone, typically ten consecutive failures. This is a last-resort protection that destroys local data before an adversary can brute-force it. Ensure your cloud backup is current and verified before enabling it, and understand the trigger is irreversible. iOS has this in Settings → Face ID (or Touch ID) & Passcode → Erase Data. Android equivalents depend on device and OS version.
- A dead-man's switch is a mechanism that automatically takes a defined protective action if you fail to check in on a schedule, for example, releasing an encrypted set of documents to a trusted contact, or notifying your lawyer, if you don't authenticate within a specified window. This approach is entirely lawful when used to protect journalistic sources or sensitive materials. Services exist for this purpose. Research options that fit your threat model, and consult a lawyer about the legal implications in your specific jurisdiction before deploying one.

⚠ COMMON MISTAKES

- Dismissing update prompts because they're inconvenient. Every deferred update is a known, documented vulnerability that remains open on your device.
- Assuming a backup exists because you set it up once. Backup processes fail silently over time, drives fill up, cloud credentials expire, settings reset. Test a restore at least quarterly.
- Storing your password manager's master password or emergency recovery kit inside the password manager itself. If you lose access, that becomes a circular dependency with no exit.
- Setting auto-lock intervals too long, fifteen minutes because locking feels disruptive. Five minutes of inconvenience is the trade for an unlocked device not sitting accessible on a café table.
- Treating breach notifications as informational. A notification means act now. Change the password, check for reuse, and monitor for unauthorized activity, today, not next week.
- Keeping only cloud backups. If your cloud account is compromised, so are your backups. Maintain at least one local encrypted copy that an attacker cannot reach through your credentials.

PILLAR 7 · VISIBILITY & ANALYTICS

Visibility & Awareness

Why it matters

In enterprise security, this pillar covers SIEM (Security Information and Event Management (centralized log collection and automated alerting), SOC (Security Operations Center) a team watching dashboards around the clock), and threat intelligence feeds. Most individuals and small organizations have none of that. What you do have is the ability to configure alerts, re-view activity logs, and notice when something doesn't look right, if you set up the signals in advance and know what normal looks like.

Visibility is where the other pillars close the loop. Strong identity controls and encryption reduce attack surface. Visibility is what tells you when something got through anyway. The average time between a breach and its discovery is measured in weeks or months for organizations. For individuals it is often years, or never. They find out when a credential appears in a public dump or a fraudulent charge shows up on a statement.

The goal is not to turn yourself into a security analyst. It is to configure the signals that will alert you if something goes wrong, then act decisively and quickly when they fire.

What good looks like

- You receive immediate alerts on new login activity for email, cloud storage, and any other account that matters.
- Breach monitoring is active on your key email addresses, you'll hear about exposed credentials before an attacker exploits them.
- You review active sessions, connected apps, and authorized devices on a quarterly schedule and revoke anything you don't recognize or use.
- You can recognize the most common tells of phishing and social engineering in real time, because you are the first sensor in every attack chain.
- You know what normal activity looks like for your accounts, so anomalies stand out instead of blending into background noise.
- You know the behavioral signs of a compromised account and can move quickly when you see them.

Do this

Turn on login alerts and account-activity notifications

Every major platform (Google, Apple, Microsoft, Meta, Dropbox, and others) offers notifications when your account is accessed from a new device or an unusual location. Many of these are

not enabled by default. They are your earliest warning of account takeover: in many cases, the alert fires before the attacker has had time to change your password or enable their own MFA.

- ☐ In your email provider's security settings, enable notifications for new-device logins and suspicious sign-in activity.
- ☐ Do the same for your primary cloud-storage account, financial accounts that offer it, and any social media account connected to your real identity or professional work.
- ☐ Right now, open the "Recent Activity" or "Security Events" page for your email and primary accounts. Most providers log recent logins with device type, approximate location, and timestamp. Look for anything unfamiliar.
- ☐ When an alert fires and you didn't just log in: treat it as a potential breach, not a glitch. Change your password immediately from a clean device, review active sessions, and revoke any you don't recognize. Then verify your MFA methods haven't been changed.

Monitor for data breaches

Credentials from breached services circulate on criminal forums for years. An attacker may hold valid credentials for accounts you haven't used or thought about recently. Proactive monitoring is how you find out before the attacker acts on it.

- ☐ Register every email address you actively use at Have I Been Pwned (<https://haveibeenpwned.com>) and enable notifications. You will receive an email the next time any of your addresses appears in a newly catalogued breach.
- ☐ Check your credit reports on a regular schedule. In the United States, the three major bureaus (Equifax, Experian, and TransUnion) are required to provide free annual reports at AnnualCreditReport.com. Look for accounts, hard inquiries, or addresses you don't recognize. See the Financial chapter for guidance on credit freezes, which are the strongest proactive control available.
- ☐ During your quarterly review (see Pillar 6), check your password manager's breach-monitoring dashboard and rotate any flagged credentials.

Periodically review sessions, connected apps, and trusted devices

Over time, your accounts accumulate authorizations: third-party apps you granted access once and forgot, devices you stopped using, active sessions that never expired. Each is a potential entry point. A quarterly review across your five most critical accounts takes under thirty minutes and removes a substantial amount of stale attack surface.

- ☐ For your email account: go to Settings → Security → "Third-party apps with account access" or equivalent. Revoke any app you no longer actively use or don't recognize.
- ☐ Review "Trusted devices" or "Remembered computers", remove any device you no longer own or can't account for.
- ☐ Check "Active sessions", sign out of any session showing an unexpected location, an old timestamp, or a device name you don't recognize.

- ☐ Repeat for your primary cloud, work accounts, and any social media platform connected to your real identity. Focus on the accounts whose compromise would cause the most harm.

Recognize phishing and social engineering

You are the first sensor in every attack. Technical controls block a significant fraction of threats; social engineering is specifically designed to route around them by targeting the human. Recognizing an attempt in real time is the skill that closes that gap. The following are the most reliable tells:

- **Engineered urgency.** "Your account will be permanently suspended in 24 hours." "You must act immediately." Urgency is a psychological technique to bypass analytical thinking. Slow down, legitimate services almost never require instantaneous action.
 - **Display name / sending address mismatch.** "Apple Support <support@randomdomain.net>." Always check the actual sending address, not just the display name that appears in your inbox. Display names are trivially forged.
 - **Hover before you click.** In email and web pages, hover over a link to see the real destination URL before clicking. If the domain doesn't match the expected service, do not click it. On mobile, press and hold to preview the link.
 - **Requests that bypass established process.** IT staff will not ask for your password over email. Banks do not verify your account by sending you a link in a text message and asking you to call the number in the message. A manager (or someone claiming to be your manager) will not urgently request gift cards over chat.
 - **Unexpected attachments or credential prompts.** Treat any unexpected attachment as suspicious, especially if it asks you to enable macros, install a plugin, or grant elevated permissions. An unexpected login screen in the middle of a workflow is a red flag, stop and verify independently.
 - **Vishing, voice phishing.** Callers who identify as your bank, the IRS, your carrier, or tech support. Hang up. Find the organization's official number independently and call back. Do not use a number provided by the caller.
- ☐ When in doubt, do not click or respond. Open a new browser tab and navigate directly to the service in question.
 - ☐ Report suspected phishing in your email client, most providers have a "Report phishing" option. This improves filtering for everyone.
 - ☐ Use your password manager as a passive phishing detector: if it doesn't auto-fill credentials on a login page, that page may not be the site you think it is. Investigate before typing anything.

Know your "normal", and the signs of compromise

Visibility requires a baseline. Spend five minutes now forming a clear picture of what normal looks like for your accounts: which devices you use, roughly from where, at what times, which

apps have authorized access. That baseline is what makes anomalies legible rather than confusing.

Review this checklist if you suspect something is wrong:

- ☐ A login alert fired for a device or location you don't recognize.
- ☐ Sent messages or emails you didn't write.
- ☐ Contacts telling you they received strange messages from you.
- ☐ Password-reset or verification emails you didn't request.
- ☐ Account recovery settings (backup email, phone number, MFA methods) changed without your action.
- ☐ Unrecognized charges, new subscriptions, or transactions on financial accounts.
- ☐ Your password no longer works and you didn't change it.
- ☐ Security software disabled, or apps installed that you didn't install.
- ☐ Device behaving unexpectedly: slowness, battery drain, settings changes, new apps you don't recognize.

If you see multiple of these signs, move quickly. Change passwords from a different, clean device first. Revoke all active sessions. Contact the platform's account-recovery process. Then work systematically through the other pillars, this event is the signal to tighten everything.

SCALING TO YOUR TEAM OR ORGANIZATION

- Enable audit logging in your collaboration platform. Google Workspace, Microsoft 365, Slack, and similar tools maintain event logs (logins, file access, sharing changes, admin actions) that are often disabled or inaccessible by default. Turn them on and confirm their retention period. You cannot investigate an incident with logs you don't have.
- Configure alerts for high-risk events: new administrator privilege grants, bulk file downloads or deletions, logins from new countries or impossible travel (a login from one city and then another city an hour later), mass email forwarding rules. Most platforms support this natively through their security dashboards or admin consoles.
- Decide who reviews what, and put it on a schedule. In a small organization this might be one person spending thirty minutes per month reviewing the flagged events. The critical thing is that it happens regularly and that someone is accountable for acting on what surfaces.
- Conduct a simple periodic log review: sort by highest-severity event type, look for patterns (repeated failed logins on a single account, access from unexpected locations, bulk data movements) and investigate anything you cannot explain with a specific business reason. Formal SIEM tooling is overkill for most small organizations; the built-in dashboards are sufficient when used consistently and with someone paying attention.

IF YOU'RE BEING TARGETED

- Heighten your monitoring attention during known elevated-risk periods: when a sensitive story is being published, during litigation or legal proceedings, around travel to high-risk jurisdictions, or following receipt of threats or hostile contact.
- Be aware that modern commercial spyware (of the kind deployed against journalists, activists, and civil society figures by nation-states and sophisticated private actors) is engineered to be invisible. It does not produce obvious battery drain, does not appear in your app list, and does not change device behavior in ways you'd notice. If you suspect device-level compromise by a capable adversary, do not rely on your device's own security tools to detect it. The spyware may be operating at a layer below what those tools can inspect.
- If you have specific reason to believe you are targeted by sophisticated spyware (not just phishing, but implant-level malware) seek expert help before doing anything else. The Access Now Digital Security Helpline (accessnow.org/help) provides free, confidential technical assistance to journalists, activists, and civil-society organizations at elevated risk. Organizations such as Citizen Lab (citizenlab.ca) investigate and document targeted digital threats. Do not attempt to investigate or clean a suspected advanced compromise on your own; you risk alerting the attacker or destroying forensic evidence.
- Watch for physical surveillance indicators alongside digital ones: unfamiliar vehicles appearing regularly near your home or workplace, individuals who appear across multiple unconnected locations, any attempt to get you to leave your device unattended. Sophisticated operations typically combine physical and digital access, a device left alone for minutes is long enough for a hardware implant or a "evil maid" attack.
- If you find an unknown active session, a device on your network you can't account for, or physical evidence of unauthorized access to your devices, treat it as a confirmed breach and respond accordingly. Isolate, document what you see without disturbing it, and get expert help before taking further action. Speed matters, but so does not tipping off an adversary who may still have access.

▲ COMMON MISTAKES

- Ignoring login alert emails by treating them as routine noise. Train yourself to read each one. Each alert deserves five seconds: "Did I just do this?" If the answer is no, treat it as an incident.
- Forgetting to revoke third-party app access after you stop using a service. Connected apps with stale permissions are a persistent, invisible attack surface that grows over time.
- Conflating "I don't see any problem" with "there is no problem." The absence of a visible anomaly means your monitoring didn't catch anything, not that nothing happened. This is why proactive alert configuration matters.
- Dismissing compromise indicators as coincidental technical glitches. Unexpected login failures, strange items in sent mail, and unfamiliar active sessions are not coincidences by default. Investigate first.
- Relying solely on antivirus as your visibility layer. Antivirus detects commodity malware reliably. It will not catch a targeted spear-phishing attack, a credential phish, a session-hijack, or an authorized third-party application that has been repurposed against you.
- Waiting too long to act after receiving a breach notification or security alert. Speed matters: every hour of delay is additional time for lateral movement, data collection, or account lockout. Respond to breach notifications the same day.

When You're Targeted Personally

Lead with Safety

This chapter covers situations where someone is actively targeting *you* specifically, your location, your identity, your family connections. The advice below moves from prevention through response, but if you are in physical danger or leaving an abusive relationship, stop and call for help before touching any settings on your device.

NOTE

If you suspect a device is being monitored by someone who controls your safety or living situation: **do not remove the monitoring software yet**. Doing so can alert the installer and escalate danger. Instead, call the National Domestic Violence Hotline (1-800-799-7233 / text START to 88788) from a device you trust, a friend's phone, a library computer. The NNEDV Safety Net project specializes in technology-facilitated abuse. Access Now's Digital Security Helpline can connect you with a digital-safety advocate who understands both the technology and the personal risk calculus.

Preventing Doxxing Before It Happens

Doxxing is the deliberate publication of your private information (home address, employer, phone number, family members' names) to expose, harass, or endanger you. The damage arrives seconds after publication. Prevention is a sustained practice.

- ☐ Audit what a stranger can find. Open a browser window you are not logged in to and search your full name, username, phone number, and home address. Note every surface.
- ☐ Use a P.O. box or CMRA (Commercial Mail Receiving Agency) address for any public registration: business filings, domain WHOIS records, voter registration where state law permits.
- ☐ Scrub public social profiles: remove your city, employer, phone number, birth year, and tagged location photos from anything visible outside your confirmed contacts.
- ☐ Enable WHOIS privacy protection on domain registrations. If a domain was registered before you added privacy, historical WHOIS records may already be archived, check and act accordingly.
- ☐ Use a dedicated email alias for public accounts, forums, and sign-ups you don't fully trust. Most major email providers and standalone alias services support this.
- ☐ Reverse-image-search your own photos periodically to find unexpected placements.

Data-Broker Opt-Outs

Data brokers compile profiles from public records, app tracking data, and commercial purchases, then sell them. Your home address, phone number, and relatives' names are already in multiple broker databases.

- ☐ Manually opt out of the highest-traffic brokers: Spokeo, Whitepages, BeenVerified, Intelius, MyLife, Radaris, and PeopleFinder. Each has an opt-out page; most require you to provide an email address to complete the request.
- ☐ Repeat the process every few months. Brokers re-list removed profiles as new data sources feed their systems.
- ☐ Paid data-removal services automate opt-outs and monitor for re-listing. They reduce the labor cost significantly but do not guarantee zero footprint. Spot-check results yourself regardless of what the service reports.
- ☐ Opt out of the credit bureau marketing databases (Equifax, Experian, TransUnion, Innovis) separately from your credit freezes, these are distinct programs.

Responding to an Active Doxx

If your information has been published, time matters and so does sequence. Act in this order.

1. **Do not engage or retaliate publicly.** Any response amplifies reach.
2. **Document before removing.** Screenshot with the full URL, page title, and your device timestamp visible. Use a web archive service to create a timestamped, independently hosted copy, harder to dispute than screenshots you control.
3. Report the content to the platform under their targeted-harassment or doxxing policy. Many platforms now prioritize these reports.
4. Alert family members or colleagues named in the publication so they are not surprised by contact attempts.
5. Notify your local law enforcement non-emergency line and file a report, even if they cannot act immediately. The report creates a dated record.
6. Tighten account security immediately: rotate passwords, audit recovery options, confirm MFA (multi-factor authentication, a second proof of identity beyond your password) is active on all major accounts.

Stalkerware and Spyware: Detection and Safe Removal

IF YOU'RE BEING TARGETED

Stalkerware (software installed secretly on your device to track your location, messages, calls, or camera) is most often planted by a partner, ex-partner, or family member who had brief physical access to your unlocked device. Removing it can immediately alert the installer. If you are in an abusive situation, contact a digital-safety advocate before making any device changes.

Indicators that stalkerware may be present: battery draining faster than expected, unexplained data-usage spikes, the device warming while idle, unfamiliar apps in the app list, or an abuser referencing conversations you had privately.

- ☐ On Android: check Settings → Apps for apps you don't recognize, especially any holding Device Administrator, Accessibility Service, or "Display over other apps" permissions. Legitimate apps rarely need all three.
- ☐ On iOS: Apple's sandboxed architecture limits most stalkerware to iCloud-level access. Check Settings → Privacy & Security → Location Services for unrecognized apps, and verify no one else knows your Apple ID password.
- ☐ When safe to act: a factory reset is the most thorough remediation. Before resetting, change all account passwords from a separate trusted device first so the restored phone starts clean.
- ☐ Restore only data you can inspect (contacts and documents) not a full device backup, which may restore the malware along with everything else.
- ☐ If you have reason to believe firmware-level compromise is possible (nation-state targeting), replace the device entirely rather than resetting it.

Hidden Bluetooth Trackers

Small Bluetooth location tags can be concealed in a bag, coat pocket, car, or piece of luggage to track your physical movements without your knowledge. Tracker hardware is inexpensive and widely available.

- ☐ iPhone users receive automatic alerts when an unknown tag has been traveling with them. Enable notifications and investigate any alert promptly.
- ☐ Android users: install a dedicated Bluetooth scanner app designed to detect unknown tracking devices. The major tag manufacturers have published detection apps; independent options also exist.
- ☐ Physically inspect wheel wells, bumpers, and undercarriage of your vehicle if you suspect physical surveillance. Trackers fit in tight spaces and are often held with magnets.
- ☐ If you find an unknown tracker, do not destroy it, it may be evidence. Photograph it in place, note the location, and consult law enforcement before removing it.

Abusive and Shared-Access Scenarios

Shared accounts, family mobile plans, smart-home devices, and location-sharing apps create digital transparency that someone with controlling intent can exploit. Access granted during a relationship does not disappear when the relationship changes.

- ☐ Audit every service where another person holds credentials, is listed as a co-owner, or can see your activity: email recovery contacts, family-sharing plans, cloud family groups, carrier account portals, smart-lock and doorbell apps.
- ☐ Review live location sharing in every messaging and map app. These are often set once and forgotten. Each can be revoked individually without alerting the other person that you checked.
- ☐ Smart-home devices set up by someone else likely still have their account as the primary owner, even if you use the device daily. To fully remove their access, factory-reset the device and set it up under your own account.
- ☐ If you need to leave a shared phone plan, contact the carrier directly. Most carriers have domestic-violence provisions that allow account separation without the account holder's consent. Ask specifically about those provisions.

Lawful Temporary "Going Dark"

There are legitimate reasons to temporarily reduce your digital visibility: an imminent threat, a relocation you haven't disclosed, or a period of elevated risk. This is different from destroying evidence, it means limiting new exposure going forward.

- ☐ Pause or tighten social media visibility. Disable check-ins, location tags, and stories for the duration.
- ☐ Ask trusted contacts not to tag you in photos or mention your location in public posts.
- ☐ Use end-to-end encrypted messaging with disappearing messages set to a short interval. This limits future exposure; it does not alter existing records.
- ☐ Disable "last seen" indicators and read receipts in messaging apps to reduce your observable activity patterns.
- ☐ Do not announce publicly that you are going quiet, that signals the exact moment your visibility dropped and may invite more attention.

Preserving Evidence

Evidence degrades quickly. Screenshots disappear when platforms remove content. Capture everything before you report or request removal.

- ☐ Screenshot with the full URL and your device's date-time visible in frame.
- ☐ Archive the URL with a web archiving service, this creates a timestamped, independently held copy that is much harder to dispute.
- ☐ Keep a running dated log: what happened, when, what you observed, what you did. Plain text in a secure notes app works fine.

- ☐ Do not alter collected evidence or add commentary after the fact.

Who to Contact

Key resources for individuals under targeted threat

Situation	Resource	Contact
Intimate-partner or domestic abuse	National Domestic Violence Hotline / NNEDV Safety Net	1-800-799-7233 · text START to 88788 · thehotline.org
Digital security under active threat	Access Now Digital Security Helpline	accessnow.org/help
Legal rights and policy guidance	Electronic Frontier Foundation (EFF)	eff.org/issues/security
Identity theft and fraud	FTC Identity Theft resource	IdentityTheft.gov
Criminal harassment or stalking	Local law enforcement	Non-emergency line for reports; 911 for immediate danger

Physical & Travel Security

Physical Device Security

A device in your hands is defended by software, biometrics, and your own attention. A device on a café table for ninety seconds is defended by nothing. Physical access is the fastest bypass to every layer of digital security you have built.

The **evil maid attack** (named for the untrusted hotel worker who has brief unsupervised access to your room) describes any scenario where an adversary has short physical access to your device. In ninety seconds, a prepared adversary can install a hardware keylogger, copy unencrypted storage, or implant firmware-level malware. The countermeasures are boring and reliable: never leave devices unattended, power off rather than sleep, and use full-disk encryption so an offline copy of your drive is worthless.

- ☐ Lock your screen every time you step away, even for thirty seconds. Configure the auto-lock timer to one minute or less.
- ☐ Never leave a device unattended in a public space or an untrusted room. If you must leave it, power it off.
- ☐ Use a privacy screen (a physical filter that narrows the viewing angle so only the person directly in front of the screen can read it) on laptops and phones when working in public. Shoulder-surfing (someone reading your screen from an angle) is a real and simple attack.
- ☐ Do not plug into unknown USB ports: airport kiosks, hotel-room USB chargers, shared charging cables. Use a USB data blocker (a pass-through adapter that cuts the data lines, leaving only power), or carry your own charger and cable.
- ☐ Keep firmware and OS updates current. Physical access attacks often exploit known firmware vulnerabilities that patches already address.

Travel Posture: Carry Minimal Data

The data on your device is the prize. Minimize what you carry, and you minimize what can be lost, seized, or copied.

- ☐ Before any significant trip, audit what is on your devices. If you don't need a file during the trip, remove it or leave the device at home.
- ☐ A dedicated travel device (a laptop or phone provisioned fresh for the trip) is the strongest practical solution for high-risk travel. Provision it with only the accounts and data you need. Upon return, wipe and reprovision before connecting it back to your main environment.
- ☐ Use remote-desktop or on-demand cloud access to reach sensitive files from the travel device rather than copying them onto it. If the device is seized, it contains no local data worth taking.

- ☐ Confirm full-disk encryption is enabled and active. Full-disk encryption protects data at rest only when the device is powered off. A device in sleep mode has already unlocked its encryption keys in RAM. Power off before any situation where the device might leave your hands.

Border Crossings

Device search at international borders (including U.S. ports of entry) is a situation where your rights and the practical reality diverge. Understand both before you travel.

U.S. Customs and Border Protection officers have broad legal authority to inspect electronic devices at the border. Courts have reached different conclusions about how deep that authority extends (whether it covers a basic manual search of the screen vs. a full forensic extraction) and the law in this area is still unsettled. The practical distinction between U.S. citizens and visa-holders is significant: a citizen cannot be denied entry for refusing to unlock a device, but a non-citizen may face denial of admission, and both face the possibility of device seizure and delayed entry.

NOTE

This section describes the situation neutrally. It does not advise you to comply or to refuse, that is a judgment call based on your citizenship status, what is on your device, your destination, and your risk tolerance. Consult EFF's "Digital Privacy at the U.S. Border" guide (eff.org) and, if you have specific concerns, an immigration or civil-liberties attorney before you travel. Never provide false statements to CBP officers; that creates legal exposure far worse than a device search.

- ☐ Power devices fully off before reaching the checkpoint. Full-disk encryption is only fully engaged when the device is off.
- ☐ Know what is on your device. If you carry source material, confidential communications, or data belonging to clients, sources, or your organization, understand the exposure before you board.
- ☐ A clean travel device eliminates the most sensitive content from the equation entirely.
- ☐ If a device is searched or seized, note the officer's name and badge number, the date and location, and what devices were taken. Report to your organization's security contact and consult legal counsel promptly.

Protests and Demonstrations

Protests involve dense crowds, law enforcement presence, and potential device seizure, a combination that rewards preparation over improvisation.

- ☐ Write your lawyer's or legal observer's phone number on your arm in pen before you go. If your phone is seized or runs out of power, that number is still accessible.

- ☐ Disable biometric unlock before you arrive. A passcode cannot be compelled from you in most U.S. contexts in the same way a fingerprint or face can be compelled passively. On iPhone: pressing the side button and either volume button simultaneously disables Face ID until the passcode is entered. Set this up before you are in a crowd.
- ☐ Enable full-disk encryption and use a strong passcode (six digits minimum; alphanumeric is stronger).
- ☐ Carry minimal data. Leave at home any device containing sensitive source materials, client information, or data you cannot afford to have seized.
- ☐ Use end-to-end encrypted messaging for coordination. Know the app's behavior under a locked screen, some show message previews even when locked.
- ☐ Keep all actions lawful. This guide does not advise on anything beyond lawful participation.

Limiting Location Exposure

Your phone's cellular radio continuously connects to towers. Wi-Fi probing broadcasts your device's identifier to nearby access points. Both generate location records held by carriers, app providers, and potentially third parties. You cannot eliminate this entirely while using a phone; you can narrow the window.

- ☐ Airplane mode cuts both cellular and Wi-Fi simultaneously. It is the fastest way to stop active broadcasting.
- ☐ Leaving the device at home eliminates the location trail entirely for that period. This is the strongest option for high-sensitivity situations.
- ☐ Audit location permissions: review which apps have always-on location access vs. only-while-using. Revoke always-on for anything that doesn't clearly require it.
- ☐ Download offline maps before traveling to a sensitive area so navigation does not require an active connection.
- ☐ Disable Wi-Fi when not actively using it in public. The Wi-Fi radio probes for familiar networks even when you are not connected, creating a detectable pattern.

Hotel and Short-Rental Security

Hotel rooms and short-term rentals are shared infrastructure with unknown prior occupants, building management with physical access, and Wi-Fi operated by a third party. Treat them accordingly.

- ☐ Treat hotel Wi-Fi as an untrusted public network. All traffic on it is potentially visible to the network operator or other guests. Use a VPN (Virtual Private Network, a service that encrypts your outbound traffic and routes it through a server you trust) for anything sensitive. Choose a VPN provider with an independently audited no-logs policy.
- ☐ Do not leave devices unattended in the room. Room safes are not secure: hotel staff and safe manufacturers typically hold override codes. Carry your device or lock it in luggage with a TSA-approved cable lock as a deterrent, not a guarantee.

- ☐ Power devices off before leaving the room for any extended period.
- ☐ Smart TVs, in-room speakers, and smart thermostats are internet-connected devices you do not control. Treat them as ambient microphones for any conversation you would not want recorded. Unplug smart speakers when not using them if the conversation warrants it.
- ☐ In high-risk contexts, request a room change if you notice signs of tampering: door-lock irregularities, unfamiliar small devices attached to power strips, or unusual hardware near the desk.

▲ COMMON MISTAKES

- Sleeping the device instead of powering it off at a border checkpoint, leaving encryption keys in RAM and accessible.
- Using biometrics as the sole unlock method while traveling, when a passcode provides stronger legal protection in many scenarios.
- Connecting to hotel Wi-Fi for sensitive work without a VPN, assuming the network is private because it has a password.
- Leaving a laptop in a hotel room overnight without powering it off, assuming the lock on the door is sufficient protection against the evil maid threat.
- Packing full production credentials and sensitive data "just in case" instead of provisioning a travel device with only what the trip requires.

Protecting Family & Dependents

Your Security Chain Includes the People You Love

Adversaries know that targeting you directly is harder than targeting someone adjacent to you. A phished partner, a kid with malware on a shared home network, or a scammed parent who gives up your phone number, these are real, common attack vectors. Securing yourself while ignoring your household leaves the back door open.

This is not about blame. People without your security background are not failing; they simply have different knowledge. Your job is to lower the effort required to be secure for the people around you, not to lecture them into compliance.

Threats That Arrive Through Family

The most common paths an adversary takes through family members:

- **Phishing a partner** who shares a home network, accounts, or knows your schedule. A compromised partner device on the same Wi-Fi can expose your traffic, your shared services, and your location.
- **Compromising a child's device.** Children are frequently targeted through games, social platforms, and peer-forwarded links. A child's device on your home network has the same network access as yours.
- **Scamming a parent** out of credentials or funds. Older adults are the primary target of phone and email fraud. A scammer who gets your parent's phone unlocked may find your number, your address, and enough relationship context to impersonate you.
- **Social engineering through relationships.** An adversary who knows your family's names, schedules, and patterns (available on social media) can construct a convincing pretext for calling you or impersonating you to others.

Securing Children

Children's digital security is genuinely different from adult security: the threat model includes peer-sourced risks, age-inappropriate content, commercial data harvesting, and predatory contact, alongside the household network risk you share.

- ☐ Set up a separate user account or profile on every shared device. A child's account should have the minimum permissions needed, no administrative rights, no unrestricted app installation.
- ☐ Use parental controls not as surveillance but as guardrails: content filtering, screen-time limits, and app approval. Build in age-appropriate trust increases as kids demonstrate judgment.

- ☐ Talk explicitly about sextortion: the technique where someone, often posing as a peer, pressures a minor into sending an explicit image and then threatens to share it. FBI data shows this targeting children and teenagers at scale. The message to give your child: *if someone pressures you for images or threatens to expose you, come to me first, you are not in trouble.*
- ☐ Treat a child's device on your home network as an untrusted device. Use network segmentation (a guest Wi-Fi network, or a VLAN if your router supports it) to separate kids' devices from your work machines.
- ☐ Review app permissions together as a teaching moment, not an inspection. Show what location-always-on vs. location-while-using means. Kids who understand the mechanics make better decisions independently.

Partners and Shared Digital Lives

Shared calendars, shared locations, shared streaming accounts, and family password managers are all legitimate features of modern life. They also create access and visibility that should be intentional, consensual, and revocable.

- ☐ Any shared location access should be mutual and explicitly agreed upon, with both parties knowing it exists. Walk through the settings together so neither person discovers access they didn't know about.
- ☐ Shared accounts should use the household's password manager, not a single shared password both people have memorized. This keeps credentials strong and makes it possible to revoke access cleanly if circumstances change.
- ☐ Encourage your partner to have their own accounts wherever possible rather than sharing a single login. Two separate accounts with strong independent credentials are more resilient than one shared credential that, if compromised, takes down both of you.
- ☐ If your partner's device gets compromised, act on the assumption that anything on your shared home network needs inspection. Change shared passwords, review shared account activity, and check for unusual access.
- ☐ Have a plan for account recovery if one of you is incapacitated. A sealed, printed recovery-codes sheet stored in a safe is a legitimate household contingency, not a security hole.

Aging Parents

Phone and email fraud disproportionately targets older adults. The mechanics are consistent: urgency, fear, authority, and an unusual payment method (gift cards, wire transfer, cryptocurrency). Your job is to make the right thing easier than the wrong thing, not to take over their digital lives.

- ☐ Set up caller ID and spam-blocking at the carrier or device level. Many modern smartphones do this automatically; confirm it is on.

- ☐ Create a shared "call me first" rule: if someone calls claiming to be from a bank, the IRS, Medicare, or a grandchild in trouble and asks for money or personal information, they call you before doing anything. Make this easy, keep your number on their home screen.
- ☐ Set up automatic updates on their devices. Unpatched devices are the most common entry point for malware that reaches older adults through forwarded links.
- ☐ Help them set up a simple password manager with a handful of critical accounts, email, bank, medical portal. You do not need to manage it for them; you need to reduce their reliance on reused, simple passwords.
- ☐ Enable login alerts on their email and bank accounts so unusual access triggers a notification. Make sure they know what the alerts look like so they can distinguish a real alert from a phishing alert.
- ☐ Do not take over their accounts or demand their credentials. That removes their autonomy and creates a single point of failure if something happens to you. Help them be capable, not dependent.

Shared Households and Devices

Roommates, household staff, and guests create access that is easy to underestimate. Physical presence is access; network presence is access.

- ☐ Create separate user profiles on every shared computer. Each profile should have its own password and only the permissions that user needs. Never use an administrative account for daily work on a shared machine.
- ☐ Use a guest Wi-Fi network for household visitors and IoT devices (smart speakers, doorbells, cameras, thermostats). Guest traffic should not be able to reach your work devices on the primary network.
- ☐ Log out of financial and sensitive accounts when you step away from a shared device. Browser password autofill does not distinguish between authorized users.
- ☐ Audit which apps and services have been authorized to access your accounts from shared devices. Remove authorizations that are no longer needed.
- ☐ For home offices: treat the room itself as a physical security zone when working on sensitive material. A closed door is not paranoid; it is professional.

Talking to Non-Technical Loved Ones Without Lecturing

Security conversations that feel like lectures do not produce behavior change. They produce resentment and avoidance.

A few principles that work better:

- **One thing at a time.** "Can I help you set up two-step login on your email today?" beats "here are the twelve security things you should do."
- **Frame it as care, not criticism.** "I'd feel better knowing your phone would lock itself if you left it somewhere" lands differently than "your phone isn't locked and that's a problem."

- **Make it concrete and specific.** Show them what a scam call script sounds like. Walk through a fake phishing email together. Abstract threats do not motivate; specific examples do.
- **Respect their autonomy.** They will make some decisions you disagree with. Your role is to inform, not to control. Provide the information, do the easy setup together, then let them own their choices.
- **Celebrate the wins.** When a parent spots a scam call and hangs up, that is a real success. Acknowledge it.

SCALING TO YOUR TEAM OR ORGANIZATION

Everything in this chapter applies to small organizations with a close team. Treat your staff like trusted household members: give them their own accounts and credentials, never share administrative logins, set up a guest network for contractors and visitors, and make "call me first" a team policy for any unusual financial request, especially urgent wire transfers or gift-card payments, which are the signature of business email compromise (BEC) fraud.

Financial Self-Defense

Credit Freezes: The Single Best Move

A credit freeze (also called a security freeze) instructs each credit bureau to block any new credit inquiry without your explicit approval. When your credit is frozen, an identity thief who has your Social Security number, address, and date of birth still cannot open a new credit card, take out a loan, or finance a purchase in your name. The creditor cannot pull your file.

It is free by federal law. It does not affect your credit score. You can unfreeze temporarily (for a specific creditor or time window) and refreeze at no cost. The primary cost is inconvenience: you must remember to unfreeze when you legitimately apply for new credit, then refreeze promptly.

- ☐ Freeze at all three major bureaus separately: Equifax (equifax.com), Experian (experian.com), and TransUnion (transunion.com). Each has an online freeze portal. Create an account at each bureau and retain the PIN or account credentials in your password manager.
- ☐ Also freeze at Innovis (innovis.com) (a smaller bureau used by some lenders) and ChexSystems (chexsystems.com), which is used for bank account opening.
- ☐ Freeze credit for your children, who have no credit history and are therefore invisible to monitoring, making child identity theft hard to detect for years. Parents can request a freeze on behalf of a minor.
- ☐ After any freeze, set a calendar reminder to verify the freeze is still active quarterly. Breaches at the bureaus themselves have occurred; verify rather than assume.

Hardening Financial Logins

Your bank, brokerage, and payment accounts are high-value targets. The attacker who controls your email already controls most of your password resets; the attacker who controls your phone number can intercept SMS verification codes.

- ☐ Use a unique, strong password for every financial account. Store these in an end-to-end encrypted, independently audited password manager, never in a spreadsheet, a note app, or your browser's built-in save feature unless you have audited its security model.
- ☐ Enable MFA (multi-factor authentication) on every financial account. Prefer app-based authenticators or hardware security keys over SMS. SMS verification codes can be intercepted via SIM-swapping, a social-engineering attack where a criminal convinces your carrier to port your phone number to a SIM they control.
- ☐ Set up transaction alerts. Most banks allow you to receive an immediate notification for any transaction above a threshold you set. Set that threshold to one dollar so any charge triggers an alert.
- ☐ Call your mobile carrier and ask about a SIM-lock or port-freeze: a PIN requirement that must be provided before your number can be ported to a new carrier or device. Not all carriers offer this, but many do.

- ☐ Review authorized apps and third-party connections in your financial accounts annually. Revoke anything you no longer use.

Recognizing Fraud and Social-Engineering Money Scams

Financial social engineering works because it exploits emotion faster than critical thinking can engage. Every effective money scam uses at least one of these four elements: urgency, fear, authority, and an unusual payment method.

Common patterns to recognize and stop:

- **"Your account has been compromised."** A caller or pop-up claims your bank or brokerage account has been hacked and you must act immediately, move your funds to a "safe" account they provide, or allow remote access to your computer to "fix" the problem. Legitimate banks do not ask you to move your own funds and never ask for remote access.
- **Gift-card payments.** Any instruction to pay a legitimate debt, tax obligation, or fee using gift cards is a scam. No government agency, utility, or legitimate business accepts gift cards as payment.
- **Romance and "pig-butcher" scams.** An online relationship builds trust over weeks or months, then pivots to an investment opportunity, usually in cryptocurrency, with fabricated returns visible in a fake app or platform. The investment is unrecoverable. Treat any online relationship that steers toward investment advice as a fraud signal.
- **Fake invoices and business email compromise (BEC).** An email that appears to be from a vendor, colleague, or executive requests a payment change, a wire transfer, or urgent action on an invoice. Verify any new payment instructions by calling the requestor at a known, pre-established phone number, not one provided in the email.
- **Urgency plus secrecy.** Any scenario where you are told not to tell your family, your bank, or your lawyer ("they will complicate this") is a scam. Legitimate situations do not require secrecy from your own advisors.

Check, ACH, and Wire Risks

Paper checks contain your full account number and routing number, printed in plain sight. Anyone who handles a check you write has the information needed to create a counterfeit check or initiate an ACH (Automated Clearing House) debit from your account.

- ☐ Use electronic bill pay from within your bank's own portal rather than mailing checks wherever possible. This exposes your account numbers to fewer parties.
- ☐ Monitor your accounts for unauthorized ACH debits. ACH fraud is common and many banks allow you to block or whitelist ACH originators.
- ☐ Wire transfers are irreversible within minutes. Before initiating any wire, verify the receiving account by phone using a number you looked up independently, not one in the email requesting the wire.

- ☐ Consider a separate low-balance account for bill payment or regular spending. If this account is compromised, the attacker's access is limited to what you keep there, not your primary savings.

Crypto Self-Custody Basics

If you hold cryptocurrency in self-custody (meaning you control the private keys rather than leaving assets on an exchange) the seed phrase is the money. It is not a password that can be reset. Anyone who has your seed phrase can drain your wallet completely and irreversibly.

- ☐ Use a hardware wallet (a dedicated offline signing device) for any amount of cryptocurrency you cannot afford to lose entirely. Software wallets on internet-connected devices are meaningfully less secure.
- ☐ Write your seed phrase on paper or metal. Never photograph it, type it into any computer or phone, store it in a cloud service, or share it with anyone, including customer support representatives claiming to be from the wallet manufacturer. No legitimate service ever needs your seed phrase.
- ☐ Store the written seed phrase in a physically secure location separate from the hardware wallet itself. A fireproof safe is a reasonable minimum; consider a second copy in a different physical location.
- ☐ Treat any message, pop-up, or email telling you to "validate your wallet," "sync your recovery phrase," or "connect your hardware wallet to this site" as a phishing attempt. These are the dominant crypto scam vectors.
- ☐ For smaller amounts you access frequently, a reputable exchange account with strong MFA is a reasonable trade-off between security and convenience. Understand that exchange-held funds are subject to exchange insolvency, regulatory seizure, and the exchange's own security posture, risks you do not control.

Recovering from Financial Fraud

Discovery is not the end of the problem; it is the start of the recovery process. Move methodically.

- ☐ Report identity theft at [IdentityTheft.gov](https://www.identitytheft.gov) (<https://www.identitytheft.gov>) (FTC). The site generates a personalized recovery plan and pre-filled dispute letters for your specific situation.
- ☐ File a police report with your local department. Some creditors and bureaus require a police report number to process disputes. Bring your FTC report to the filing, it documents the fraud pattern.
- ☐ Contact each affected financial institution's fraud department directly (use the number on the back of your card or on the institution's official website, not a number from a suspicious email). Dispute unauthorized transactions in writing and retain copies.
- ☐ Dispute fraudulent accounts with the credit bureau(s) where they appear. Under the Fair Credit Reporting Act (FCRA), bureaus must investigate and respond within 30 days. Include your police report and FTC report in every dispute.
- ☐ Freeze your credit at all three bureaus immediately if you have not already.
- ☐ Document your recovery timeline: dates, names, case numbers, what was said. This record is essential if disputes escalate or legal action becomes necessary.

▲ COMMON MISTAKES

- Freezing credit at only one or two bureaus, leaving the others open to fraudulent inquiries.
- Using SMS-based two-factor codes on bank accounts while the phone number is vulnerable to SIM-swapping.
- Photographing a cryptocurrency seed phrase "for safekeeping" and storing it in a cloud photo library.
- Verifying a wire transfer recipient using contact information from the same email that requested the wire.
- Waiting days before reporting suspicious account activity, reducing the window for a chargeback or recovery.

When Something Goes Wrong: Incident Response

The Framework: Contain → Assess → Recover → Learn

Incidents feel like emergencies because they are. The natural response is to move fast, which often makes things worse. A calm framework costs you a few minutes and saves you hours of compounded damage.

- **Contain.** Stop the bleeding. Disconnect the affected device from the network. Change compromised credentials from a clean device before the attacker can use them further. Freeze credit if financial data was exposed.
- **Assess.** Understand what actually happened before you start recovering. What was accessed? What was exfiltrated? What accounts or devices are affected? Who else might be impacted?
- **Recover.** Restore from known-good backups. Rebuild compromised accounts. Re-enable services once you have confirmed they are clean.
- **Learn.** Document what happened, how it happened, and what you will change. The incident is only useful if it produces a more resilient posture afterward.

One rule applies to every scenario below: **work from a known-clean device**. Using a compromised device to recover a compromised account is not recovery, it is feeding the attacker your new credentials in real time.

Per-Scenario Playbooks

Account Takeover

An account takeover occurs when an attacker gains access to one of your accounts (email, social media, financial, cloud storage) and uses it before you detect it.

First Hour

- ☐ On a clean device: attempt to log into the account. If you are locked out, use the account's official recovery flow, not a link from a suspicious email.
- ☐ If you regain access: immediately change the password to a new, unique one. Check and update MFA settings, remove any authenticator apps, phone numbers, or recovery emails you did not add. Review active sessions and terminate all others.
- ☐ If you cannot regain access: use the account provider's identity-verification recovery process. This takes time; start it now.
- ☐ Change the password on any other account that used the same or a similar password.

First Day

- ☐ Audit what the attacker could access using this account: password reset emails sent to it, linked accounts, stored files, payment methods.
- ☐ Alert anyone the attacker may have contacted impersonating you. Review sent-mail folders and any connected messaging apps for outbound messages you did not send.
- ☐ Check your email's filter and forwarding rules, attackers frequently set up silent forwarding rules that persist even after you recover the account.
- ☐ File a report with the platform's abuse or trust-and-safety team.

Follow-Up

- ☐ Audit every account that used this email address as its login or recovery contact. Verify each is still secured.
- ☐ Enroll in phishing-resistant MFA (hardware security key or passkey) on your email, removing SMS verification.
- ☐ Check Have I Been Pwned (haveibeenpwned.com) to determine whether the credential appeared in a breach.

Lost, Stolen, or Seized Device**First Hour**

- ☐ Use your platform's remote-wipe feature immediately if the device contains sensitive data: Find My on Apple devices, Find My Device on Android. Do this before the attacker has time to exfiltrate or disable the device's network access.
- ☐ Change the password on the primary account linked to the device (Apple ID, Google account) from a clean device.
- ☐ Revoke active sessions for any accounts you were logged into on the device.

First Day

- ☐ File a police report if the device was stolen. You will need this for insurance and potentially for carrier account actions.
- ☐ Notify your mobile carrier to suspend the SIM card, preventing calls, texts, and data from flowing to the device.
- ☐ Review and revoke any OAuth tokens or app authorizations granted by the device.

Follow-Up

- ☐ If a device was seized by law enforcement or at a border, consult a lawyer before attempting remote wipe, destruction of evidence laws and the specific context matter.
- ☐ Confirm your full-disk encryption was active and the device was in a locked state when it left your control. This determines the practical data exposure risk.

Malware or Ransomware

First Hour

- ☐ Disconnect the affected device from all networks immediately: unplug ethernet, disable Wi-Fi, disable Bluetooth. Ransomware spreads laterally across networks; isolation limits the blast radius.
- ☐ Do not pay a ransom without consulting a professional. Payment does not guarantee decryption and funds criminal operations. Law enforcement and cybersecurity firms sometimes have decryption tools for known ransomware families.
- ☐ Do not attempt to use the compromised device for anything else until it has been professionally assessed or wiped.

First Day

- ☐ From a clean device: change passwords for accounts you were logged into on the compromised machine.
- ☐ Identify your last known-good backup. Confirm it predates the infection and has not been connected to the infected machine since.
- ☐ Report ransomware to the FBI's Internet Crime Complaint Center (IC3.gov) and, if you are a business, to CISA (cisa.gov/report). These reports inform decryptor development.

Follow-Up

- ☐ Wipe and reinstall the OS from a trusted source, or restore from your verified clean backup.
- ☐ Identify the infection vector: a phishing email, a malicious download, an unpatched vulnerability. Close it before reconnecting to the network.
- ☐ Review your backup strategy. The 3-2-1 rule (three copies, on two media types, one off-site or offline) ensures ransomware cannot reach all copies simultaneously.

Doxxing or Data Exposure

First Hour

- ☐ Document the exposure before taking anything down: screenshots, archived URLs, dated log. See the "When You're Targeted Personally" chapter for full evidence-preservation guidance.
- ☐ Alert people named in the exposure so they can prepare for potential contact.
- ☐ Report to the platform. File a law enforcement report even if immediate action is unlikely.

First Day

- ☐ Tighten account security across all major platforms: unique passwords, MFA enabled, recovery contacts reviewed.
- ☐ Contact the National Domestic Violence Hotline (1-800-799-7233) or Access Now Digital Security Helpline (accessnow.org/help) if the exposure is part of an abusive or stalking situation.

Follow-Up

- ☐ Begin data-broker opt-out process to reduce ongoing exposure surface.

- ☐ Consider a CMRA or P.O. box address for future public-facing registrations.

Clicked a Phishing Link

First Hour

- ☐ Did you enter credentials on the page? Change those credentials immediately from a clean device, assume they are compromised.
- ☐ Did you download a file or allow an install prompt? Treat the device as potentially compromised. Disconnect from the network and assess.
- ☐ Did you only click and see the page without entering anything or downloading? The risk is lower, but check the account the link arrived in for follow-up phishing activity.

First Day

- ☐ Run a reputable anti-malware scan on the device if a file was opened.
- ☐ Review recent account activity on any service whose credentials may have been entered.
- ☐ Report the phishing attempt to your email provider and to the Anti-Phishing Working Group (reportphishing@apwg.org).

Follow-Up

- ☐ Enable phishing-resistant MFA on the compromised account to prevent credential stuffing even if the password was captured.
- ☐ Brief the people you communicate with most: phishing campaigns often pivot quickly to impersonate a compromised account to its contacts.

Keep an Offline Recovery Kit

When your digital environment is compromised, you cannot rely on it to recover itself. An off-line recovery kit (physical, stored somewhere safe) is the foundation of a real recovery plan.

- ☐ Print the backup recovery codes for your most critical accounts: your primary email, your password manager, and your MFA authenticator app. Store them sealed in a waterproof envelope in a physically secure location.
- ☐ Write down emergency contact numbers: your lawyer, a trusted technical advisor, your most important financial institution's fraud hotline, and family contacts. Do not rely on a compromised phone to hold these.
- ☐ Keep your password manager's emergency-access instructions (how to invoke account recovery without your device) printed with the recovery codes.
- ☐ Include the serial numbers and IMEI numbers of your devices. You will need these for police reports and carrier actions.
- ☐ Test your recovery kit once a year: can you recover your email account from scratch using only the printed materials and a clean device?

When to Call a Professional

Some incidents exceed what any individual should handle alone. Call a professional when:

- You have reason to believe a nation-state or sophisticated threat actor is involved. Signs include persistence after clean reinstalls, firmware-level anomalies, and coordinated access across multiple unrelated accounts.
- Ransomware has encrypted critical business or professional data and you have no clean backup.
- A device seizure has legal implications (law enforcement, border crossing, litigation) and you need to understand your rights and obligations before acting.
- Financial fraud has resulted in wire transfer or ACH losses that require institutional escalation beyond what you can initiate yourself.

Resources: Access Now's Digital Security Helpline serves journalists, activists, and civil society without charge. The EFF's Surveillance Self-Defense resources (eff.org) include referrals. For financial fraud, engage a consumer law attorney alongside the FTC and law enforcement processes. For business incidents, a retained incident response firm is worth the cost before an incident happens, not after.

▲ COMMON MISTAKES

- Using the compromised device to change passwords, feeding new credentials to the attacker in real time.
- Paying ransomware demands without consulting a professional, which funds the operation and does not guarantee recovery.
- Remote-wiping a seized device before consulting a lawyer about the legal context.
- Skipping the forwarding-rule audit after an email account takeover, leaving a silent exfiltration channel running.
- Not maintaining an offline recovery kit and discovering during an incident that all recovery paths go through the compromised device.

Keep Going

You made it through the whole guide. That means something. You now understand your threat model, your attack surface, and the levers that actually move the needle. You have a working mental model that most people (including most professionals) never develop.

Here is the most important thing to carry forward: you do not have to be perfect. You never could be. Neither can enterprises with million-dollar security budgets and dedicated teams. The goal is not perfection. The goal is to be a harder target than you were last month, and harder still next quarter.

Adversaries optimize too. They go where the effort-to-reward ratio is best. Most attacks are not bespoke operations against a specific individual; they are automated sweeps looking for the easiest door. Lock your doors. That alone eliminates the vast majority of risk.

The 80/20 that actually sticks

If you did nothing else after reading this guide except the seven "Start Here" actions (password manager, phishing-resistant MFA, automatic updates, full-disk encryption, tested backups, locked-down account recovery, and a breach-exposure check) you would be dramatically safer than you were before. Those seven steps address the attacks that succeed against capable, security-aware people every day.

Everything else in this guide is about going deeper on the specific pillars where your threat model demands it. Not every chapter applies to everyone at the same urgency. Use your judgment.

Build a cadence, not a marathon

The most durable security practice is also the simplest: a short, regular review. Four times a year, block 30 minutes on your calendar. Call it a security check-in. Run through these questions:

- ☐ Any new accounts since last time? Get them into the password manager with a unique password and MFA turned on.
- ☐ Any accounts you no longer use? Delete them or at minimum revoke app permissions.
- ☐ Pending software or firmware updates you have been deferring? Do them now.
- ☐ Any new devices on your home or office network? Confirm they belong there.
- ☐ Have I Been Pwned, any new breaches involving your email addresses? Change the affected passwords.
- ☐ Recovery codes for your most critical accounts, still stored somewhere safe and accessible?

- ☐ Any life change (new job, travel, relationship change, organizational role shift) that changes your threat model? Adjust accordingly.

That is it. Thirty minutes, four times a year. Two hours of deliberate attention annually produces compounding returns in resilience.

Resilience, not paranoia

The goal of all of this is not to make you anxious. It is to give you options. When an account gets compromised (and statistically, one will) you want the blast radius contained: strong unique passwords mean one breach does not cascade into ten; MFA means a stolen password alone is not enough; backups mean ransomware is an inconvenience rather than a catastrophe.

Resilience means you can take the hit and keep going. That is the real measure of a strong security posture: not that bad things never happen, but that when they do, you recover faster and with less damage than you would have otherwise.

You have the tools. You have the model. The rest is practice.

Keep going.

Master Checklist

One consolidated list covering all seven pillars and every special-topic chapter. Work through it in order the first time; revisit individual sections as your situation changes. Items are phrased as actions, check them off when done, not when you've merely read them.

Do These First

- ☐ Set up a password manager and create a unique, random password for every account.
- ☐ Turn on phishing-resistant MFA or passkeys everywhere; remove SMS 2FA where alternatives exist.
- ☐ Turn on automatic updates for your phone, computer, apps, and router.
- ☐ Encrypt your devices and set a strong screen lock.
- ☐ Back up important data following the 3-2-1 rule and test a restore.
- ☐ Lock down account recovery, secure your primary email first, then audit every other account's recovery contact.
- ☐ Check your exposure on Have I Been Pwned and turn on login alerts for critical accounts.

Pillar 1 · You & Your Accounts

- ☐ Audit every account in your password manager; delete or deactivate accounts you no longer use.
- ☐ Remove your phone number from recovery on accounts where email recovery is available.
- ☐ Enable hardware security key or passkey authentication on your email, cloud storage, and financial accounts.
- ☐ Review third-party app authorizations (OAuth grants) on your email and social accounts; revoke anything unused.
- ☐ Set up a separate alias or catch-all email address for sign-ups and newsletters to isolate your primary inbox.
- ☐ Confirm your password manager's own master password is long, unique, and backed by MFA.

Pillar 2 · Your Devices

- ☐ Verify full-disk encryption is active on every laptop and desktop you own.
- ☐ Set your lock screen to engage after no more than two minutes of inactivity.
- ☐ Enable remote wipe on every phone and laptop.
- ☐ Confirm your router firmware is up to date and change the admin password from the default.
- ☐ Disable features you do not use: Bluetooth when not in use, AirDrop set to contacts-only or off.
- ☐ Review which apps have access to location, microphone, camera, and contacts; revoke what you do not actively need.

- ☐ Write down (on paper, stored securely) what you would do if your primary phone were seized or destroyed today.

Pillar 3 · Apps & Software

- ☐ Uninstall apps you have not opened in six months.
- ☐ Use a browser with strong tracker-blocking defaults or install a reputable content-blocking extension.
- ☐ For sensitive communications, use an end-to-end-encrypted messaging app (Signal is a widely-trusted example) instead of standard SMS.
- ☐ Verify your email provider offers end-to-end encryption options for sensitive messages, or adopt an E2EE email provider for high-risk correspondence.
- ☐ Check browser extensions, remove any you did not intentionally install or that have broad "read all data" permissions.
- ☐ Prefer apps from official stores or verified developer sources; avoid sideloading from untrusted sites.

Pillar 4 · Your Data

- ☐ Identify your most sensitive files (financial records, legal documents, source material, contacts list) and confirm they are encrypted at rest.
- ☐ Test your backup restore: actually recover a file end-to-end, not just confirm the backup exists.
- ☐ Review cloud storage sharing permissions; revoke links you shared but no longer intend to be active.
- ☐ Purge data you no longer need, old tax returns beyond retention period, copies of sensitive docs on devices that don't need them.
- ☐ Ensure any encrypted archive or backup has its key or passphrase documented somewhere you (or a trusted person) can access if you are incapacitated.
- ☐ Check data broker removal: search your name on a few major people-search sites and submit opt-out requests.

Pillar 5 · Network & Environment

- ☐ Use a VPN (virtual private network, encrypts your traffic from your device to the VPN server) on untrusted networks such as hotels and coffee shops.
- ☐ Change your home router's admin username and password from the factory defaults.
- ☐ Separate IoT devices (smart TVs, cameras, thermostats) onto a guest or dedicated network segment so they cannot reach your computers.
- ☐ Use an encrypted DNS provider (DNS-over-HTTPS or DNS-over-TLS) to reduce passive surveillance of which sites you visit.
- ☐ Know what to do before connecting to a network you do not control, have a VPN ready, not just installed.

Pillar 6 · Automation

- ☐ Audit any automation workflows (scripts, cloud functions, webhook integrations, IFTTT/Zapier-style connectors) that have access to sensitive accounts or data.
- ☐ Rotate or revoke API keys and service tokens you no longer actively use.
- ☐ Ensure automated backup jobs are actually running, check the last-success timestamp, not just the job definition.
- ☐ Apply least-privilege to service accounts: each automation should only have access to what it strictly needs.
- ☐ Store secrets (API keys, tokens, passwords) in a secrets manager or encrypted store, never in plaintext config files or code repositories.

Pillar 7 · Visibility & Awareness

- ☐ Turn on login notifications for email, cloud services, and financial accounts.
- ☐ Set a calendar reminder to check Have I Been Pwned every three to six months.
- ☐ Review account activity logs (login history) for your email and primary cloud accounts at least quarterly.
- ☐ Know what a phishing attempt looks like, verify sender domains, do not click login links in email, go directly to the site instead.
- ☐ Follow one reliable security news source so major incidents reach you within days, not weeks.
- ☐ After any major public breach affecting a service you use, change your password and audit connected sessions immediately.

When You're Targeted

- ☐ Move all sensitive communications to end-to-end-encrypted channels with disappearing messages enabled.
- ☐ Enable Lockdown Mode (iOS) or equivalent high-security mode on devices that support it.
- ☐ Audit and minimize your public digital footprint: social profiles, old forum posts, publicly listed phone numbers.
- ☐ Use a hardware security key for all critical accounts, do not rely solely on authenticator apps.
- ☐ Assume your adversary has access to your phone metadata (call logs, location) even if not your content, compartmentalize accordingly.
- ☐ Consult a lawyer and a digital security trainer (organizations like EFF's Surveillance Self-Defense program provide referrals) before taking drastic defensive steps.
- ☐ Document incidents: date, what happened, what you observed, screenshots, this record matters if you need legal recourse.

Physical & Travel

- ☐ Know your rights at your country's border regarding device searches; consult EFF or a lawyer before international travel with sensitive material.
- ☐ Consider traveling with a minimal or wiped device and restoring from a clean backup after crossing into a high-risk jurisdiction.
- ☐ Enable a strong device passphrase (not biometrics alone) before any situation where you may be compelled to unlock a device.
- ☐ Use a privacy screen on laptops and phones in airports, transit, or shared workspaces.
- ☐ Be aware of your physical surroundings when discussing sensitive matters, assume you can be overheard or filmed.

Family & Small Organization

- ☐ Get everyone in your household or small team onto a password manager, shared vaults for shared accounts.
- ☐ Establish a shared understanding of what to do if someone receives a suspicious message or call: do not click, do not reply, report internally first.
- ☐ Set up family or team login alerts so a breach affecting one person is caught by the group.
- ☐ Review which family members or colleagues have access to shared accounts and remove former members promptly.
- ☐ Create a written "if something happens to me" document covering how to access critical accounts, backed up and held by a trusted person.

Financial

- ☐ Place a security freeze on your credit file with all major bureaus, this is free and blocks most new-account fraud.
- ☐ Turn on transaction alerts for every bank and credit card account.
- ☐ Use virtual card numbers for online purchases where your provider supports them.
- ☐ Review your credit report for unfamiliar accounts at least once a year.
- ☐ Treat financial account credentials as your highest-priority MFA upgrade target.
- ☐ Know your financial institution's fraud reporting number before you need it.

Incident-Response Readiness

- ☐ Write down the steps you would take if your email account were taken over right now, before it happens.
- ☐ Keep printed (not only digital) recovery codes for your most critical accounts in a physically secure location.

- ☐ Know the account-recovery process for your email provider, password manager, and phone carrier before you are locked out.
- ☐ Identify one trusted person who knows how to reach you if your primary phone and email are both unavailable.
- ☐ Test your incident plan once: simulate losing your phone and walk through recovery from scratch.
- ☐ After any incident, change credentials for the affected account and every account sharing a password with it; review logs to understand scope.

Glossary

Terms are defined as they are used in this guide. Where a term has a precise technical meaning that differs slightly from casual usage, the technical meaning is given.

Attack surface

The total set of points through which an adversary could try to enter or extract data from a system, every account, device, app, network connection, and person. Reducing your attack surface means eliminating or hardening those points before an attacker can exploit them.

Breach

An incident in which unauthorized parties gain access to data or systems they are not permitted to access. A credential breach specifically means a database of usernames and passwords has been stolen and is typically sold or published online.

Credential stuffing

An automated attack in which stolen username/password pairs from one breach are tried against other services, exploiting the fact that many people reuse passwords. A password manager with unique passwords per site eliminates this risk entirely.

Data broker

A company that collects personal information (from public records, purchase history, social media, location data, and other sources) and sells or licenses it to third parties. Data brokers are a primary source of the personal details used in targeted social engineering and doxxing.

Dead-man's switch

An automated mechanism that triggers a pre-set action (such as sending a message, publishing a document, or notifying a contact) if the owner fails to check in within a defined period. Used by journalists, activists, and others as a contingency against detention or incapacitation.

Doxxing

The act of researching and publicly publishing an individual's private information (home address, workplace, family members, phone number) typically to expose them to harassment, threats, or physical harm.

Encryption at rest

The encryption of data while it is stored on a device or server, so that the raw files are unreadable without the correct key. Full-disk encryption is the most common implementation for personal devices.

Encryption in transit

The encryption of data as it moves across a network, so that anyone intercepting the traffic cannot read it. HTTPS and TLS are the most common implementations; they protect the connection between your browser and a web server.

End-to-end encryption (E2EE)

A form of encryption in which only the sender and intended recipient hold the keys, not the service provider in the middle. If the messaging service's servers are breached, or if the provider is compelled by a court, E2EE-protected messages remain unreadable. Signal is a widely cited example.

EXIF / metadata

Data embedded in a file that describes when, where, and how it was created. Photos taken on a smartphone typically include GPS coordinates, device model, and timestamp in their EXIF metadata. This information can reveal your location and identity even when the photo itself does not.

Full-disk encryption (FDE)

A security measure that encrypts the entire storage drive of a device, making all data unreadable without the correct credentials or key. On a lost or seized device, FDE prevents an adversary from reading your files simply by plugging in the drive. BitLocker (Windows), FileVault (macOS), and the default encryption on modern iPhones and Pixels are common implementations.

GrapheneOS

A hardened, open-source Android operating system focused on privacy and security, available for Google Pixel devices. It strengthens Android's security model with additional sandboxing, permission controls, and network isolation, and is recommended for high-threat-model users who need an Android device.

Hardware security key

A physical device (typically a USB or NFC token) that generates a cryptographic proof tied to a specific website and your physical presence. Because the proof is bound to the legitimate site's domain, hardware keys are immune to phishing. FIDO2/WebAuthn-compliant keys (such as YubiKey or Google Titan) are the strongest widely available MFA method.

IoT (Internet of Things)

The broad category of internet-connected devices beyond computers and phones: smart TVs, cameras, doorbells, thermostats, routers, appliances, and similar hardware. IoT devices often have poor security update track records and weak default configurations, making them a common entry point on home and office networks.

Least privilege

The security principle that any user, device, or application should have access only to the specific resources and permissions it needs for its current task, nothing more. Least privilege limits the damage an attacker can do if any single account or component is compromised.

Lockdown Mode

An optional, high-restriction operating mode available on Apple devices (iPhone, iPad, Mac) that disables or limits features known to be exploited by sophisticated spyware, in-

cluding certain message attachment types, FaceTime call features, and web browser technologies. Designed for people at risk of mercenary spyware attacks.

MDM (Mobile Device Management)

Software that allows an organization to remotely configure, monitor, and manage devices, enforcing encryption, requiring screen locks, pushing updates, and wiping devices that are lost or no longer authorized. In an enterprise, IT runs MDM on company devices; in a small organization context, MDM lets you apply consistent policies across a team's phones and laptops.

MFA (Multi-factor authentication)

A login method that requires two or more distinct proofs of identity: typically something you know (a password), something you have (a code from an app or a hardware key), or something you are (biometric). MFA stops most account takeover attempts even when your password is already known to an attacker.

OAuth

An open authorization standard that allows you to grant a third-party application limited access to your account on another service, without giving the third party your password. "Sign in with Google" and "Connect your Twitter account" flows use OAuth. Revoking OAuth access cuts off an app's access immediately.

Passkey

A FIDO2-based login credential that replaces the password entirely. A passkey is a cryptographic key pair: the private key lives on your device and never leaves it; the public key is registered with the service. Authentication requires both possession of the device and your biometric or PIN to unlock it, making passkeys both phishing-resistant and passwordless.

Password manager

Software that generates, stores, and fills strong unique passwords for every account, protected by one strong master password (and ideally MFA). A password manager is the single highest-leverage security tool available to individuals; it eliminates password reuse and makes credential-stuffing attacks ineffective against you.

Phishing

A social engineering attack delivered via email, SMS (smishing), voice call (vishing), or fake website that tricks you into revealing credentials, installing malware, or authorizing a fraudulent transaction. Phishing is the most common initial access technique in both mass and targeted attacks.

Ransomware

Malware that encrypts your files and demands payment (typically in cryptocurrency) in exchange for the decryption key. Tested offline backups are the primary defense; they make ransomware an inconvenience rather than a catastrophe, because you can restore from a clean copy without paying.

Recovery codes

One-time backup codes provided when you set up MFA, intended for use if you lose access to your primary second factor (e.g., your phone is lost or broken). Recovery codes must be stored securely offline (in a password manager, printed in a safe, or memorized) because anyone who has them can bypass MFA.

SIM swap

An attack in which a criminal convinces your mobile carrier (through social engineering or bribery) to transfer your phone number to a SIM card they control. Once they have your number, they receive any SMS-based MFA codes sent to you, enabling account takeover. Mitigated by using authenticator apps or hardware keys instead of SMS for MFA, and by placing a PIN or port-freeze on your carrier account.

Single sign-on (SSO)

An authentication arrangement that lets you log in once (to an identity provider such as Google, Apple, or Microsoft) and then access multiple connected services without re-authenticating. SSO centralizes your security: excellent MFA on the SSO account protects all dependent services, but a compromised SSO account becomes a master key.

Social engineering

Any attack that manipulates people rather than exploiting technical vulnerabilities, impersonating a trusted person or institution to extract credentials, money, information, or access. Phishing is one form; others include pretexting (fabricating a scenario), vishing (voice-based), and in-person impersonation.

Spyware / stalkerware

Malicious software that covertly monitors a device (recording location, messages, calls, keystrokes, or camera and microphone activity) and transmits the data to a remote party. Commercial-grade spyware (such as Pegasus) is sold to governments; stalkerware is typically used by abusive partners. Both operate without the device owner's knowledge or consent.

Threat model

A structured analysis of who might want to harm you or your data, what they want, what capabilities they have, and what the consequences of compromise would be. A clear threat model lets you prioritize security investments rationally (hardening what matters most against realistic adversaries) rather than trying to defend equally against everything.

Three-two-one (3-2-1) backup

A backup strategy: keep at least three copies of your data, on at least two different types of media, with at least one copy stored off-site (or in a separate cloud account). The rule ensures that no single failure (hardware death, ransomware, fire, theft) can destroy all copies simultaneously.

Two-factor authentication (2FA)

A common term for MFA when exactly two factors are required: typically your password plus a time-based code, push notification, or hardware key confirmation. The terms 2FA

and MFA are often used interchangeably, though MFA is the more precise term when more than two factors may be involved.

VPN (Virtual Private Network)

A service that encrypts your device's traffic and routes it through a server operated by the VPN provider, masking your IP address from sites you visit and from your local network or ISP. A VPN shifts trust from your local network to the VPN provider, it is most useful on untrusted networks (hotels, airports) and for obscuring your location, not as a comprehensive privacy solution.

WPA3

The current Wi-Fi security standard, replacing WPA2. WPA3 uses stronger encryption and protects against offline dictionary attacks on captured handshakes, meaning an attacker who records your Wi-Fi traffic cannot later crack your passphrase by brute force. Set your router to WPA3 or WPA2/WPA3 transition mode if your devices support it.

Zero Trust

A security architecture and philosophy that replaces the traditional "trusted inside, untrusted outside" perimeter model with continuous verification: no user, device, or network location is inherently trusted. Every access request must be authenticated and authorized, access is limited to what is needed, and all activity is logged for anomaly detection.

ZTNA (Zero Trust Network Access)

A technology that provides application-level access to specific resources rather than broad network access, replacing traditional VPNs in enterprise Zero Trust deployments. Instead of connecting a device to the whole internal network, ZTNA grants access only to the specific application being requested, after verifying device health and user identity.

Appendix A: NSA Zero Trust Pillar Crosswalk

This guide maps one-to-one to the NSA's seven Zero Trust pillars, translated into civilian terms and civilian-scale implementation. In an enterprise, many of the Target-level activities in the table below are handled for you by an IT or security team, MDM enrollment, SIEM deployment, SOAR playbooks. As an individual or small organization, you handle equivalent functions yourself with the tools and practices described in each chapter. The table below shows the correspondence explicitly.

This guide's chapters mapped to NSA Zero Trust pillars and representative Target-level activities

This guide's chapter	NSA pillar	Representative NSA Target-level activities
Chapter 1, You (User)	User	MFA & Identity Provider (IdP) integration; Privileged Access Management (PAM); User Behavior Analytics (UBA); phishing-resistant credential enforcement
Chapter 2, Your Devices	Device	Device Health & Comply-to-Connect; MDM/UEM enrollment; Endpoint Detection & Response (EDR); firmware integrity verification; device posture assessment
Chapter 3, Apps & Software	Application & Workload	Application allowlisting & control; container and workload integrity; API gateway & mutual TLS (mTLS); software supply chain verification; app-level access policies
Chapter 4, Your Data	Data	Data tagging & classification; Data Loss Prevention (DLP); encryption at rest and in transit; rights management; data access logging and auditing
Chapter 5, Network & Environment	Network & Environment	Macro- and micro-segmentation; DNS filtering & sinkholes; Software-Defined Perimeter (SDP); encrypted DNS (DoH/DoT); network traffic inspection
Chapter 6, Automation	Automation & Orchestration	SOAR (Security Orchestration, Automation & Response); policy enforcement automation; automated patch deployment; dynamic access policy adjustment; playbook-driven incident containment
Chapter 7, Visibility & Awareness	Visibility & Analytics	Log aggregation & SIEM; User and Entity Behavior Analytics (UEBA); continuous monitoring & alerting; threat intelligence correlation; anomaly detection

Special chapters and their pillar mappings

This guide includes four chapters beyond the seven pillars that address civilian-specific concerns with no direct enterprise equivalent. Each maps primarily to pillar functions as follows:

Special chapters and their closest NSA pillar functions

Special chapter	Closest NSA pillar function(s)
Incident Response	Assume-breach posture (all pillars); response automation (Automation & Orchestration); forensic log review (Visibility & Analytics)
Family & Household	Identity lifecycle management (User); insider risk and shared-account controls (User + Application & Workload)
Physical Security & Travel	Facility and physical access controls (Device + Network & Environment); operational security (OPSEC) supporting all pillars
Financial Security	Fraud prevention and asset protection (Data + User); access credential hygiene for financial accounts (User)

NOTE

The NSA pillar framework was designed for enterprise and government environments with dedicated security teams, centralized tooling, and formal change-management processes. This guide deliberately keeps the pillar structure intact so that readers moving into or working alongside enterprise security programs can map their personal practices to the organizational framework they will encounter. The civilian implementations are functionally analogous, not identical.

Appendix B: Resources & Where to Go Next

This is a curated short list, not exhaustive. Every resource here is free, maintained by credible organizations, and directly actionable for the audience this guide addresses. For product-specific tool recommendations (the "which password manager, which VPN, which hardware key" questions), see the tool directories below and watch for the planned opinionated-tools edition of this guide.

Tool directories

Two community-maintained directories evaluate and list privacy- and security-respecting software across categories: password managers, browsers, messaging apps, VPNs, email providers, and more. They apply consistent criteria (open source where possible, independent audits, transparent funding) and update their recommendations as the landscape changes.

- [Privacy Guides \(privacyguides.org\)](https://www.privacyguides.org) (<https://www.privacyguides.org>), the actively maintained successor, run by an independent nonprofit. This is the current authoritative version of the directory.
- [PrivacyTools \(privacytools.io\)](https://www.privacytools.io) (<https://www.privacytools.io>), the original site; check that recommendations are current before acting on them.

NOTE

A product-specific edition of this guide (with opinionated, tested picks in each category and direct setup instructions) is planned as a follow-up. The current edition is deliberately capability- and criteria-based so it does not go stale between product-landscape changes.

Learn & defend

- [EFF Surveillance Self-Defense \(ssd.eff.org\)](https://ssd.eff.org) (<https://ssd.eff.org>), practical, scenario-based security guides from the Electronic Frontier Foundation. Covers threat modeling, secure communications, device security, and special topics for journalists, activists, and others in high-risk roles. Free, regularly updated.
- [CISA Secure Our World \(cisa.gov/secure-our-world\)](https://www.cisa.gov/secure-our-world) (<https://www.cisa.gov/secure-our-world>), the U.S. Cybersecurity and Infrastructure Security Agency's public-facing guidance on the highest-impact individual security actions: MFA, strong passwords, software updates, and phishing recognition.

If you're being targeted or abused

If you are facing a targeted threat (stalking, domestic abuse, harassment, or state surveillance) these organizations provide direct, confidential help from people trained in both safety planning and digital security.

- [Access Now Digital Security Helpline \(accessnow.org/help\)](https://www.accessnow.org/help) (<https://www.accessnow.org/help>), free, confidential digital security support for activists, journalists, human rights defenders, and civil society organizations facing targeted digital threats. Available in multiple languages.
- [NNEDV Safety Net \(techsafety.org\)](https://www.techsafety.org) (<https://www.techsafety.org>), the National Network to End Domestic Violence's technology safety project. Provides resources specifically for survivors of intimate partner violence dealing with stalkerware, account abuse, and location tracking.
- [National Domestic Violence Hotline \(thehotline.org\)](https://www.thehotline.org) (<https://www.thehotline.org>), 24/7 crisis support with safety planning that includes digital and device safety. Call 1-800-799-7233 or text START to 88788.

Identity & fraud

- [FTC IdentityTheft.gov \(identitytheft.gov\)](https://www.identitytheft.gov) (<https://www.identitytheft.gov>), the U.S. Federal Trade Commission's official resource for identity theft recovery. Generates a personalized recovery plan, pre-filled letters to creditors and agencies, and step-by-step guidance.
- [AnnualCreditReport.com](https://www.annualcreditreport.com) (<https://www.annualcreditreport.com>), the federally mandated free credit report portal. Request your reports from all three major bureaus (Equifax, Experian, TransUnion). You can place a free credit freeze directly with each bureau from here.

Check your exposure

- [Have I Been Pwned \(haveibeenpwned.com\)](https://haveibeenpwned.com) (<https://haveibeenpwned.com>), Troy Hunt's free breach notification service. Enter any email address to see which known data breaches include it, and subscribe for future alerts. Referenced in the "Start Here" chapter as one of the seven highest-leverage actions you can take immediately.

Appendix C: Sources & Acknowledgements

This guide is an independent civilian adaptation of a body of U.S. government Zero Trust publications. The source documents are listed below. United States government works produced by employees of the federal government in the course of their official duties are not subject to copyright in the United States and are in the public domain. Houston Labs LLC has adapted, re-organized, and rewritten this material for a civilian individual and small-organization audience; the adaptation itself is the work of Houston Labs LLC.

This guide is not affiliated with, sponsored by, or endorsed by the National Security Agency, the Cybersecurity and Infrastructure Security Agency, the National Institute of Standards and Technology, or any agency or instrumentality of the United States government.

Primary source documents

- National Security Agency. *Zero Trust Implementation Guideline Primer*. January 2026. The primary organizing framework for this guide's seven-pillar structure and implementation sequencing.
- National Security Agency. *Embracing a Zero Trust Security Model*. Cybersecurity Information Sheet. February 2021. The foundational NSA statement of Zero Trust principles and the never-trust/always-verify doctrine.

NSA Cybersecurity Information Sheets, Zero Trust pillar series

Each of the following NSA information sheets addresses one of the seven Zero Trust pillars in detail. This guide adapts the Target-level activities and implementation priorities from these documents into civilian-scale guidance.

- National Security Agency. *Advancing Zero Trust Maturity Throughout the User Pillar*. Cybersecurity Information Sheet.
- National Security Agency. *Advancing Zero Trust Maturity Throughout the Device Pillar*. Cybersecurity Information Sheet.
- National Security Agency. *Advancing Zero Trust Maturity Throughout the Network and Environment Pillar*. Cybersecurity Information Sheet.
- National Security Agency. *Advancing Zero Trust Maturity Throughout the Data Pillar*. Cybersecurity Information Sheet.
- National Security Agency. *Advancing Zero Trust Maturity Throughout the Visibility and Analytics Pillar*. Cybersecurity Information Sheet.
- National Security Agency. *Advancing Zero Trust Maturity Throughout the Application and Workload Pillar*. Cybersecurity Information Sheet.

- National Security Agency. *Advancing Zero Trust Maturity Throughout the Automation and Orchestration Pillar*. Cybersecurity Information Sheet.

Supporting standards and frameworks

- National Institute of Standards and Technology. *SP 800-207: Zero Trust Architecture*. August 2020. The foundational NIST definition of Zero Trust architecture, tenets, and deployment models. Available at csrc.nist.gov (<https://csrc.nist.gov/publications/detail/sp/800-207/final>).
- Cybersecurity and Infrastructure Security Agency. *Zero Trust Maturity Model, Version 2.0*. April 2023. CISA's maturity model mapping Traditional, Initial, Advanced, and Optimal levels across the five CISA pillars (Identity, Devices, Networks, Applications and Workloads, Data). Available at [cisa.gov/zero-trust-maturity-model](https://www.cisa.gov/zero-trust-maturity-model) (<https://www.cisa.gov/zero-trust-maturity-model>).

About this adaptation

Houston Labs LLC produced this guide as an independent work. The adaptation involved substantial rewriting, reordering, and original authorship to address the civilian individual and small-organization context, which differs significantly from the enterprise and government environments the source documents address. Errors of translation, omission, or interpretation are the responsibility of Houston Labs LLC, not of the source agencies.

Feedback and corrections are welcome. Contact information is available at the Houston Labs LLC website.