

HOUSTON LABS LLC

ZERO TRUST · CIVILIAN FIELD MANUAL

Zero Trust Phase Two: Integrate and Advance

Zero Trust implementations for high-risk individuals and small organizations

First Edition · 2026 · Adapted from the NSA Zero Trust Implementation Guideline: Phase Two (Jan 2026)

Contents

About This Guide	3
How to Use This Guide	5
Pillar 1 · You: Accounts & Identity	8
Pillar 2 · Your Devices	12
Pillar 3 · Apps & Software	16
Pillar 4 · Your Data	20
Pillar 5 · Network & Environment	24
Pillar 6 · Automation	28
Pillar 7 · Visibility & Awareness	32
Phase Two Completion Checklist	36
Appendix A: NSA Phase Two Crosswalk	39
Appendix B: Sources	41

About This Guide

What this guide is

This is *Houston Labs Civilian Zero Trust: Phase Two, Integrate and Advance*. It is the second hands-on implementation companion in the Houston Labs Civilian Zero Trust series. Where Phase One asked you to build a solid foundation, Phase Two asks you to connect the pieces and raise the bar. You will close the gaps between controls that were deployed separately, add stronger authentication, tighten segmentation, and build the habits that turn a collection of individual settings into a coherent security posture.

This guide assumes you have read the *Houston Labs Civilian Zero Trust Primer* and have completed *Phase One: Build Your Foundation*. If you have not done both, stop here and do them first. Phase Two builds directly on Phase One. Skipping ahead means your integration steps will have no foundation to integrate.

About Houston Labs

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 putting the power of advanced technology into more hands, especially those of ordinary people and small organizations, and helping them use and defend it.

This Civilian Zero Trust series is part of our free advisory and education work. We have no government affiliation and no security product to sell you; our only goal here is to give capable, motivated people the same rigorous thinking professional security teams apply, in a form they can execute without a dedicated staff. The series is vendor-neutral: any tool referenced is for illustration, not commercial endorsement.

How this guide was developed

This guide adapts the *NSA Zero Trust Implementation Guideline: Phase Two* (published January 2026), a U.S. government work in the public domain. The NSA guideline describes 41 activities and 34 capabilities organized across seven pillars for organizations working to integrate distinct Zero Trust solutions. We have translated that framework into concrete, personal-scale steps that a non-enterprise reader can execute, while preserving the pillar structure and the underlying logic of the original.

The NSA guideline was produced for U.S. government agencies and enterprise contractors. This guide is not that document. We have rewritten, reordered, and substantially adapted the material for civilian use. Where the original addresses enterprise IT infrastructure, we address

the tools and contexts a capable individual or small team actually has access to. The intellectual framework is the NSA's. The civilian translation is ours.

Disclaimer

This guide is provided **as is**, without warranty of any kind. It is not legal advice. It is not compliance guidance. It does not constitute a professional security assessment of your situation. The threat landscape changes faster than any printed guide can track. Verify that any tool or service mentioned still exists, is still maintained, and still meets your needs before you rely on it.

Following this guide reduces risk. It does not eliminate it. No guide does.

This is the **First Edition, 2026**.

Independence and affiliation

This is an independent publication by Houston Labs LLC. It is not affiliated with, sponsored by, or endorsed by the National Security Agency, the U.S. government, or any tool vendor, platform, or service mentioned in the text. Any reference to a specific tool is for illustrative purposes only and does not constitute a commercial recommendation or endorsement.

Complete Phase One first

The work in this guide is harder than Phase One. It requires judgment calls, tradeoffs, and in some places a meaningful investment of time. None of it makes sense unless your Phase One foundation is solid. If your password manager is not deployed, your backups are not running, your router firmware is not updated, and your key accounts do not have hardware-backed MFA, Phase Two will not help you. Get Phase One done. Then come back here.

How to Use This Guide

Where Phase Two sits in the series

The Houston Labs Civilian Zero Trust series has three layers. The *Primer* explains what Zero Trust means for civilians, why it matters, and how to think about it. Phase One gets the non-negotiable foundations in place: a password manager, encryption, backups, a hardened router, basic visibility. Phase Two is this guide. It picks up exactly where Phase One left off.

If you have read the Primer, you understand the model. If you have completed Phase One, your foundation exists. Phase Two's job is to connect those foundations into something coherent, then push each pillar to the next level. You are not starting over. You are integrating and advancing.

Do not use this guide as a standalone document. The Primer is required reading. Phase One is a required prerequisite. This guide references both and does not re-explain concepts they already cover.

What "integrate and advance" means

Phase One deploys controls in isolation. You turned on encryption. You set up backups. You installed MFA. Those things are done, but they are separate. Phase Two connects them. Your identity controls start talking to your device controls. Your network segmentation reflects who is actually allowed to reach what. Your visibility tools stop being passive and start being something you actually read.

Advancing means raising the bar in each pillar. SMS two-factor authentication becomes hardware keys. A single backed-up drive becomes a tested, encrypted, offline restore. A router with a changed admin password becomes a network with genuine segmentation. The work in Phase Two is not busywork. Each step closes a specific gap that Phase One left open by design, because Phase One's goal was speed and coverage, not depth.

How this guide is organized

The guide follows the same seven pillars as Phase One and the NSA framework: You, Your Devices, Apps and Software, Your Data, Network and Environment, Automation, and Visibility and Awareness. Each pillar chapter is structured the same way:

- **Why this matters now** ties the pillar to Phase Two specifically and explains what gap you are closing.
- **Do this** gives you the concrete, ordered steps for this phase, followed by a checklist.
- **Scaling to your team or organization** covers what the same work looks like when you are responsible for others.

- **If you are being targeted** addresses higher-risk situations that require more than the standard steps.
- **Common mistakes** describes the errors people make at this stage and how to avoid them.

Each pillar chapter ends with a single line: **You're done with this pillar when:** followed by a plain-language completion criterion. That line is not aspirational. It is the test. If you cannot honestly say you meet it, the pillar is not done.

Personal-first, with an organizational track

The main text of each chapter is written for an individual. The steps, the examples, and the default assumptions are all built around one capable person securing their own setup, possibly with a partner or a small household.

If you also run or secure a small organization, the **Scaling to your team or organization** box in each chapter is written for you. It translates the same pillar into what it means when you have employees, contractors, or volunteers whose devices and accounts you are responsible for. You do not need to read those boxes if you are working only on your personal setup. You should read them if you have any organizational responsibility, even informal.

The three callout boxes

Three types of callout boxes appear throughout the guide. Each has a specific meaning.

Scaling to your team or organization covers what the same work looks like at organizational scale. If you run a team, a nonprofit, a newsroom, or any small organization, these boxes are for you.

If you're being targeted addresses situations where a general threat model is not enough. If you have reason to believe you are under active surveillance or facing a motivated, specific adversary, these boxes tell you what to add or change. Not everyone needs to read them. If you do need to read them, you will know.

⚠ **Common mistakes** describes the errors that are predictable and recurring at each stage. These are not edge cases. They are the things most people get wrong, and knowing them in advance saves real time.

A plain **Note** box appears occasionally for information that does not fit the other categories but is important enough to call out.

Phase Two adds meaningful friction. Pace yourself.

Phase One had a lot of tasks, but most were quick. Install this, turn on that, change this setting. Phase Two has fewer tasks but harder ones. Migrating to hardware security keys, testing a re-

store, auditing your OAuth grants, deciding whether and how to compartmentalize your identities: these take real thought and sometimes real time.

Work through the pillars in order. Each one builds on the last. Do not skip the "You're done with this pillar when" line and move on. If you rush past it, you will find yourself revisiting earlier pillars when a later one depends on something you left unfinished.

This is not a weekend project. Give it the time it deserves.

PILLAR 1 · USER

You: Accounts & Identity

Why this matters now

Phase One gave every account a unique password and turned on MFA. You closed the credential-stuffing door. The next door is still open: SMS codes that a carrier agent can redirect in a phone call, authenticator app codes that a real-time phishing page can relay before they expire, and identities that share enough infrastructure that a breach in one role cascades into another. Phase Two upgrades every interceptable second factor to a phishing-resistant one, and enforces hard boundaries between the identities you carry.

Do this

1. Move critical accounts to passkeys

A passkey is a FIDO2 credential bound to a specific device and to the exact domain that created it. A fake login page receives nothing, because the credential is domain-locked and the cryptographic response only works on the legitimate site. If your email, cloud storage, and financial accounts offer passkeys, enroll now. The setting is typically under Security or Sign-in options. After creating the credential, open a new private browser session and verify the passkey login works before removing any fallback method.

Register at least two hardware security keys as backup second factors on each critical account. A passkey stored on your phone disappears when the phone is lost, reset, or seized. The hardware key is the recovery path. Buy two, register both, and keep them in physically separate locations.

2. Remove SMS 2FA from every account that offers a better option

Sign in to each account where you use an SMS code as a second factor. Check whether that account supports a passkey, hardware key, or authenticator app. Enroll the stronger method first and confirm it works. Then, and only then, remove the SMS option from that account's security settings. Never delete a second factor before its replacement is live and tested.

For accounts where SMS is the only MFA option, flag them. If the service holds sensitive data and has not shipped a stronger MFA method, that is a signal about its security posture. Evaluate whether it belongs in your stack.

3. Harden your hardware key setup

If you have not bought a hardware security key yet, do it now. Look for a key that supports FIDO2 and WebAuthn, works over USB-A, USB-C, and NFC, and has been through independent security review. Register it on your email account, your password manager (if the product sup-

ports hardware key unlock), and every other account where phishing-resistant MFA is available. The backup key goes in a different location from the primary, not in the same bag, not in the same building.

4. Compartmentalize your identities

You need at least three distinct email accounts: one for professional or work use, one for personal accounts and trusted contacts, and one for anything public-facing (forums, mailing lists, event signups). If you do high-sensitivity work (investigative journalism, legal disputes, political organizing, security research), that work belongs on a fourth account hosted by a provider with a strong legal jurisdiction and a documented record of resisting third-party data requests.

Each identity is a separate account, not just an alias on the same inbox. An adversary who compromises one should hit a wall, not a directory of your other personas. Keep credentials for each identity separate in your password manager. Do not sign in to high-sensitivity accounts from the same browser session you use for lower-trust work.

5. Revisit recovery paths

With passkeys and hardware keys now in place, return to the recovery settings on your highest-value accounts. Remove phone-number recovery from any account where you now have a hardware key as the backup path. Confirm that recovery codes are still stored offline and that you know where they are. Update any "trusted phone number" fields to reflect a number you actually control today.

- ☐ Enroll passkeys on your email, cloud storage, and financial accounts.
- ☐ Register two hardware security keys on each critical account; store the backup key at a separate physical location.
- ☐ Audit every account using SMS 2FA; enroll a stronger method and remove SMS once the replacement is confirmed working.
- ☐ Establish separate email accounts for work, personal, public-facing, and (if applicable) high-sensitivity work.
- ☐ Confirm that no single account compromise gives access to credentials for another identity.
- ☐ Remove phone-number recovery from accounts where you now have a hardware key as a fallback.
- ☐ Verify offline recovery-code storage is current, legible, and accessible.

SCALING TO YOUR TEAM OR ORGANIZATION

Phase Two identity work for an organization closes the manual gaps left over from Phase One enforcement.

- **Enforce SSO with phishing-resistant MFA at the policy level.** Configure your identity provider to block any sign-in that does not route through SSO. Set conditional-access policies that require a passkey or hardware key for administrator roles and for access to sensitive systems. Block SMS-only MFA for anyone above the baseline access tier.
- **Build and run a joiner/mover/leaver process every time.** Joiners: provision access scoped to the role, nothing extra, on day one. Movers: update access in the same week a role changes; new access goes in only after old access comes out. Leavers: disable the account and revoke all active sessions on their last day, not the following week. This is a checklist on every HR transition, not an informal memory exercise.
- **Eliminate shared credentials.** Any shared password that more than one person knows cannot be individually revoked. Migrate shared service accounts to per-user provisioning or a secrets manager with per-user audit trails.
- **Run a quarterly access review.** Who holds administrator rights in your identity provider, your cloud services, and your critical tools? Does each person still need them? Remove stale privileges before they become incident reports.

IF YOU'RE BEING TARGETED

- **Hardware keys are not optional.** Passkeys stored on a device can be compromised if that device is seized or forensically imaged. A physical FIDO2 key is harder to replicate and immune to over-the-wire phishing. For any account tied to sensitive work, hardware key enrollment is the baseline.
- **Access your high-sensitivity identity from a dedicated device only.** Never use it on a phone or laptop that also runs personal or public-facing accounts. If that account matters enough to have its own identity, it deserves its own access path.
- **Lock your SIM against carrier-side attacks.** Call your carrier and set a SIM PIN plus an account transfer freeze. For contexts where your real phone number needs to stay private, use a separate data-only eSIM for that work.
- **Treat recovery paths as the primary attack surface.** Sophisticated adversaries often bypass the account and attack the recovery flow, where controls are weaker. Review every recovery option on critical accounts at least twice a year. When in doubt, remove the option and rely on hardware key recovery.

▲ COMMON MISTAKES

- **Enrolling a passkey without a backup hardware key in place.** If your device is lost, the passkey is gone. Register the hardware key before removing SMS as a fallback, not after.
- **Removing SMS 2FA before testing the replacement.** Always confirm the new method works with a test login before deleting the old factor. The time to discover a configuration problem is not when you are locked out.
- **Separate email addresses but one shared password manager vault.** If the vault is compromised, compartmentalization collapses. For a high-sensitivity identity, consider a separate vault instance or a distinct vault passphrase not derived from your main one.
- **Storing both hardware keys in the same bag.** One theft or one loss event eliminates both. Physical separation between primary and backup is the entire point of having two.
- **Treating the joiner/mover/leaver process as a one-time setup.** The process only functions if someone runs it on every transition. If the last three departures were handled informally, the process does not exist in practice regardless of what the documentation says.

You're done with this pillar when: every critical account uses a passkey or hardware key as its second factor, SMS 2FA is removed from every account where a stronger option was available, your identities are separated into distinct accounts with no shared credentials between them, and (for organizations) your identity provider enforces SSO and phishing-resistant MFA with a documented, practiced joiner/mover/leaver checklist.

PILLAR 2 · DEVICE

Your Devices

Why this matters now

Phase One turned on encryption, automatic updates, and a strong lock. Phase Two asks the follow-on question: hardened against what, and against whom? A phone that is encrypted but leaking telemetry to its platform vendor is not the same as one that is network-isolated. A device with three months of update support remaining is not the same as one supported for four more years. This phase forces a deliberate posture choice and gets you to execute it fully.

Do this

1. Choose your posture and commit to it

The two credible paths for high-risk individuals are hardened mainstream and de-Google. Half-measures give you neither the convenience of mainstream nor the isolation of a hardened build. Pick one and configure it completely.

Option A: Hardened mainstream (iOS or stock Android)

Best for: Most people in this audience. Strong security without rebuilding your workflow or losing access to apps your work depends on.

On iPhone: Enable Lockdown Mode when risk is elevated (active targeting, border crossings, high-risk events). Lockdown Mode disables link previews, wired data accessories, complex web rendering, and certain wireless features, eliminating entire classes of remote exploit. Use it for specific threat windows, not necessarily as a permanent state. Set your iCloud backup encryption to Advanced Data Protection so Apple cannot decrypt your backups. Review which apps have background-refresh and location access; disable what you did not consciously grant.

On Android: Choose a device from a manufacturer with a long, committed security update window. Enable full encryption (verify in Settings, Security), restrict installs to the Play Store, and disable features you do not use: Bluetooth off when idle, NFC off if unused, Developer Options locked. Check your Google Account permissions and revoke any connected-app grants you do not recognize.

Option B: De-Googled / GrapheneOS

Best for: People whose threat model includes the platform vendor, or who need per-app network and sensor controls not available on stock Android.

GrapheneOS is an independently audited, hardened Android build that removes Google services, strengthens the permission model, and adds per-app network access controls (you can block any app from reaching the network entirely). It supports sandboxed Google Play, so apps that require Google services can run in an isolated container that cannot access the rest of the system.

To execute this path: Buy a Pixel (the only device GrapheneOS officially supports). Follow the web installer at grapheneos.org, which handles bootloader unlock, image flashing, and re-lock. After installation, install only the apps you verified you need. Set each app's network access, sensors access, and storage scope deliberately. Enable automatic update checks in the GrapheneOS updater app.

What you accept: Installation requires comfort with firmware flashing (roughly two hours) and an ongoing maintenance commitment. Some apps with root-detection require the sandboxed Play layer. If your threat model clearly includes the platform vendor, the effort is worth it; if your primary threats are phishing, malware, and opportunistic theft, Option A hardened properly is excellent.

2. Set up a dedicated device for sensitive work

A device that carries only what it needs for sensitive work bounds what an adversary gets if it is seized or compromised. It does not have to be expensive: a used Pixel running GrapheneOS, or an older iPhone in Lockdown Mode with a minimal app set, both qualify. No personal accounts, no casual apps, no permanently saved sessions. Sign in to accounts as needed and sign out when done.

3. Complete your device inventory and flag end-of-life dates

Open the device inventory you started in Phase One. For each device, record the manufacturer's end-of-support date for the current OS version. Any device within twelve months of end-of-life should be marked for replacement or retirement from sensitive use. Check each device's last OS update date; if any device is more than a few weeks behind the current release, investigate why automatic updates are not applying.

- ☐ Choose Option A or Option B and configure it completely on your primary phone.
- ☐ On iOS: enable Advanced Data Protection for iCloud; set Lockdown Mode for elevated-risk periods.
- ☐ On Android: verify full encryption; disable unused radios; restrict installs to the official store.
- ☐ On GrapheneOS: configure per-app network access for every installed app; enable automatic updates.
- ☐ Set up a dedicated device (phone or laptop) for high-sensitivity work with a minimal app footprint.
- ☐ Update your device inventory with end-of-life dates; flag any device within twelve months of support end.
- ☐ Confirm every device is running the current OS version.

SCALING TO YOUR TEAM OR ORGANIZATION

- **Deploy lightweight MDM.** A mobile device management tool enforces a baseline (screen lock, encryption, minimum OS version) across every device that touches organizational data without requiring a full IT stack. Look for a small-team tier that integrates with your identity provider and enroll all company-issued devices.
- **Gate access on device compliance.** Configure your SSO to check device health before granting access to email, code repositories, or internal tools. A device that is not encrypted, not current, or not enrolled in MDM should not reach company data. Access is conditional on posture, not just credentials.
- **Issue a dedicated device for sensitive roles.** Anyone who handles financial data, legal documents, source code, or investigative work should use a device provisioned and managed by the organization for that work. Personal devices carry personal apps, personal browser history, and personal network exposure. Separation is the control.
- **Offboard devices the same day as the person.** Revoke MDM enrollment and remote-wipe the company profile when someone leaves. Same day means same day, not the following week when someone gets around to it.

IF YOU'RE BEING TARGETED

- **GrapheneOS or iOS Lockdown Mode as the default state, not just for travel.** At nation-state or resourced-adversary threat levels, these are not overcautious configurations. They remove entire classes of remote exploit that affect standard builds.
- **Carry a travel device with minimal data.** A separate phone with no personal accounts, no stored contacts beyond what the trip requires, and accounts signed in only on-site (not permanently saved) limits what is exposed at a border crossing or during a device seizure.
- **Power the device off before reaching a checkpoint.** A locked device may still have decryption keys in RAM. Powered off, it does not. Some devices re-enable biometrics after a reboot; enter your passcode first to reach "before first unlock" state before the device leaves your hands.
- **Treat any device out of your physical control as potentially compromised.** Time in a hotel room, repair shop, or customs inspection is enough. Do not use it for sensitive work until inspected or reset.

▲ COMMON MISTAKES

- **Choosing GrapheneOS but enabling sandboxed Play without per-app network restrictions.** The isolation model breaks down if you grant Play-layer apps unrestricted network access. Configure network permissions deliberately for every app after installation.
- **Enabling Lockdown Mode once and forgetting it can be toggled.** Lockdown Mode is a context-appropriate tool. Leaving it permanently on may break workflows you rely on; understanding when to enable it for specific threat windows is more useful than a blanket setting you override by habit.
- **Treating MDM enrollment as the end of device security work.** MDM enforces a baseline configuration. It does not vet what apps employees install, does not monitor for indicators of compromise, and does not catch misconfigurations that fall within the allowed policy. It is a floor, not a ceiling.
- **Keeping a dedicated sensitive-work device connected to personal accounts.** A "dedicated" device that has your personal email, personal app store account, and personal iCloud signed in is not dedicated. It is a regular device with a different name.
- **Buying a device for GrapheneOS that is not a Pixel.** GrapheneOS supports only Pixel devices. Attempting to install it on other hardware is unsupported and may result in an insecure or broken configuration.

You're done with this pillar when: you have fully executed either the hardened-mainstream or GrapheneOS posture on your primary devices, a dedicated device exists for sensitive work with a minimal app and account footprint, your device inventory shows no device within twelve months of end-of-life without a replacement plan, and (for organizations) MDM is deployed with a compliant-device gate active on access to organizational systems.

PILLAR 3 · APPLICATION & WORKLOAD

Apps & Software

Why this matters now

Phase One removed apps you did not use and tightened permissions on the ones you kept. Phase Two targets access you deliberately granted: every OAuth token you approved, every SaaS tool connected to your Google or Microsoft account, every third-party integration that still holds a live authorization. A password change does not revoke these. They persist until you explicitly remove them, and most have not been reviewed since the day you clicked "Allow." This phase audits those grants, builds a vet-before-you-adopt habit, and gives you a safe path for files that may be weaponized.

Do this

1. Audit and revoke OAuth connected-app grants

OAuth grants are standing authorizations. When you approved a third-party app to read your Gmail, access your Google Drive, or post to your social media, that approval stays active until revoked, regardless of what has happened to that app's security, ownership, or intent since then. Many of these grants date back years and cover more scope than you remember agreeing to.

Go through each account you care about and review its connected apps:

- **Google:** myaccount.google.com, then Security, then Third-party apps with account access.
- **Microsoft:** myaccount.microsoft.com, then Privacy, then Apps and services.
- **Apple:** appleid.apple.com, then Sign in with Apple.
- **GitHub:** github.com/settings/applications, then Authorized OAuth Apps and Authorized GitHub Apps.
- **LinkedIn, X, and other social accounts:** each has a Security or Privacy section listing connected applications.

For each grant: do you still use this app? Does the scope it was granted (read all email, access all files, post on your behalf) match what you would approve today? If the answer to either question is no, revoke it. Revocation is immediate and does not affect your account or the app's functionality for other users. After revoking, watch for any workflow that breaks; if something breaks, re-evaluate whether you actually need that integration and grant only the minimum scope required to restore it.

2. Vet new SaaS before granting access

Before you authorize any new application, ask four questions. Does this tool actually need the access it is requesting (an invoicing app that wants to read all your email is over-reaching)? Who makes it, and have they been independently audited? Where does the data go, and does storing it with this vendor add a copy of your data with their security posture, not yours? What is the deletion process when you stop using it? Know the exit before you sign up.

3. Build a simple app allowlist

Allowlisting does not require enterprise software. It means having a deliberate, short list of applications you have decided are acceptable to install on each device, and treating anything not on that list as a question to answer before installing, not after. Maintain the list in your notes or password manager. When someone recommends a tool or when you encounter a request to install something new, check the list before clicking download. The friction is the point: most compromises from malicious software rely on impulsive installs.

4. Open risky files in a sandbox or throwaway environment

Attachments and documents from unknown or untrusted sources can carry malicious macros, exploits in document rendering engines, or payloads that execute on open. Opening them in your primary working environment is unnecessary risk. Use one of these approaches depending on your setup:

- **A dedicated isolated virtual machine** with no access to your main file system and no persistent network access. Open the file there. If it executes something unexpected, the damage is contained to the VM.
- **A cloud document viewer** (uploading to an online rendering service that opens the file server-side) so the local machine never executes the file's contents directly.
- **An isolated browser profile or container** for web-based files. Firefox containers and Chromium profiles with no access to saved passwords or extensions limit what a compromised page can reach.

- ☐ Review and revoke OAuth connected-app grants on Google, Microsoft, Apple, GitHub, and social accounts.
- ☐ Remove any grant where you no longer use the app or where the scope is broader than the function requires.
- ☐ Apply the four-question vet checklist before authorizing any new SaaS tool or application integration.
- ☐ Write down your app allowlist for each device; commit to asking "is this on my list?" before installing.
- ☐ Set up a sandbox environment for opening untrusted files (VM, cloud viewer, or isolated profile).
- ☐ Disable or remove macro execution in office productivity apps unless your workflow requires it.

SCALING TO YOUR TEAM OR ORGANIZATION

- **Maintain a vetted SaaS inventory.** Keep a list of approved tools that have passed a basic security review. New adoption requests go through a short checklist: who is the vendor, what data will they hold, what is the access scope, and how will the account be offboarded. The inventory also makes departures faster because you know exactly which services to revoke.
- **Audit OAuth grants against the SaaS inventory.** Any connected app not in your approved inventory is shadow IT. Review your identity provider's app-consent logs periodically and require admin approval before new OAuth grants are authorized.
- **Enforce macro and script execution policy.** Disable macros in office productivity tools organization-wide except where a specific workflow explicitly requires them and where those documents come from a verified internal source. Most ransomware delivery through documents relies on macros being enabled by default.
- **Provide sandboxing for the team.** A shared virtual machine or cloud-based file viewer that anyone on the team can use for untrusted attachments removes the individual burden of setting up an isolated environment and creates a consistent, safe practice across the organization.

IF YOU'RE BEING TARGETED

- **Treat every unsolicited file as hostile until proven otherwise.** A document sent by a contact whose account may have been compromised, a PDF from a source you have never interacted with before, a compressed archive from a new collaborator: these are vectors, not courtesies. Sandbox first, every time.
- **Review OAuth grants quarterly, not just once.** Targeted adversaries sometimes use compromised third-party apps as a persistent access path. An app you trusted a year ago may have changed ownership, been acquired, or had its OAuth credentials stolen. Quarterly review catches grants that were legitimate when created but are no longer safe.
- **Minimize your SaaS footprint deliberately.** Every tool that holds your data is a potential breach notification in eighteen months. For sensitive work, prefer tools where data stays on infrastructure you control. When a cloud tool is necessary, verify the vendor cannot read the stored content.
- **Audit browser extensions the same way you audit OAuth grants.** Extensions run with broad access to every page you visit and can read form inputs, including passwords. Remove anything you did not deliberately install or no longer actively use.

▲ COMMON MISTAKES

- **Assuming a password change revokes third-party app access.** It does not. OAuth tokens are independent of your password. Revoking a connected app requires going to the account's security settings and explicitly removing it.
- **Granting broad scope to save a few seconds of setup.** "Allow all" during an OAuth consent screen is fast. Scoping access to only what the app needs takes thirty more seconds and limits the damage if that app is later compromised. Always read the permission list before clicking Authorize.
- **Building an allowlist once and never updating it.** Software changes. A tool that was safe to install last year may have changed ownership, introduced spyware-like features, or gone unmaintained. The allowlist is a living document, not a one-time artifact.
- **Opening suspicious files in a "clean" browser tab instead of a true sandbox.** A browser tab on your main profile is not isolation. The file can still interact with browser extensions, saved credentials, and the local filesystem depending on how the browser handles the content type. Use a real isolated environment.

You're done with this pillar when: every OAuth connected-app grant has been reviewed and anything unused or over-scoped has been revoked, you have a written app allowlist and a habit of checking it before installing, you have a sandbox environment ready for untrusted files, and (for organizations) new SaaS adoption goes through a documented vet process with macro execution disabled by default.

PILLAR 4 · DATA

Your Data

Why this matters now

Phase One started automatic backups and identified your crown jewels. Phase Two asks whether those backups actually work, whether a single event (ransomware, account compromise, platform outage) can wipe all copies at once, and whether you have made a deliberate choice about who can read what you store. A backup you have never restored is an assumption. An encrypted cloud service where you have not verified the encryption model is trust, not control. This phase closes both gaps.

Do this

1. Complete and verify your 3-2-1 backup

The 3-2-1 rule: three copies of your data, on two different media types, with one copy offsite. Phase One may have gotten you to two copies. Complete the third now, and make one of the three offline: a physical encrypted drive stored somewhere separate from your primary device. An offline copy is immune to ransomware propagating through network shares and immune to an account compromise reaching your cloud storage.

After completing the setup, test a restore. Restore a set of critical files to a different folder or device, open them, and confirm the content is intact. A backup that fails the restore test is not a backup. Schedule this test every six months as a hard calendar commitment.

2. Choose your long-term storage strategy

Option A: End-to-end encrypted cloud

Best for: People who need access from multiple devices, easy sharing with collaborators, and managed infrastructure without the overhead of running their own systems.

To execute this path: Migrate primary file storage to a provider independently audited for end-to-end encryption that publishes its key-custody architecture. Your device encrypts files before they leave it; the provider stores ciphertext it cannot read. Confirm the encryption covers everything you store. Generate a recovery key and store it offline: if you lose account access, you are the only one who can recover the data. Evaluate providers by audit history, not marketing. A flaw in their client-side code remains in your threat model.

Option B: Local-first or self-hosted

Best for: People whose threat model includes cloud providers, who have the technical comfort to run and maintain their own systems, or who cannot accept any third-party key custody under any conditions.

To execute this path: Set up a NAS or dedicated home server running file-sync software with full storage encryption. Configure remote access through an encrypted tunnel, not direct internet exposure. Apply OS and application updates to the server on the same cadence as your other devices, store its credentials in your password manager, and include it in your backup rotation (the server is a primary copy, not a backup; it needs its own offsite backup). You are the administrator: patching, hardware failure, and configuration security are your responsibility. Self-hosted systems that go unmaintained become vulnerabilities.

3. Run a data-minimization sweep

Data you do not have cannot be stolen, subpoenaed, or leaked. Work through cloud drives, email archives, note-taking apps, and any shared drives you have access to. For each location: does this data still serve a purpose? Could it cause harm if exposed? Delete what you no longer need. For data you must keep, confirm it is in the right storage tier for its sensitivity. Check which accounts still hold data from services you stopped using; most offer a data-export-and-delete option.

- ☐ Verify you have three backup copies of critical data on two media types, with one offline copy on an encrypted external drive stored offsite.
- ☐ Test a restore from your backup: pick critical files, restore them, open and verify the content.
- ☐ Schedule a restore test every six months on your calendar.
- ☐ Choose Option A or Option B and migrate your primary file storage to that strategy.
- ☐ For Option A: confirm the provider is independently audited, store the recovery key offline.
- ☐ For Option B: enable full encryption on the server, configure remote access through an encrypted tunnel, include the server in backup rotation.
- ☐ Run a data-minimization sweep across cloud drives, email, notes, and old service accounts; delete what you no longer need.

SCALING TO YOUR TEAM OR ORGANIZATION

- **Apply least-privilege access to every shared drive.** No one gets access unless explicitly granted; access is scoped to what the role requires. Review every shared folder and remove anyone who no longer needs it, including former employees, contractors, and past project collaborators.
- **Define a retention and deletion policy and enforce it.** Decide how long different data categories are kept and when they are deleted. Apply those timelines consistently. Data kept longer than necessary can be subpoenaed, breached, or leaked. The policy must cover what happens to data when a team member leaves.
- **Include organizational files in restore tests.** Quarterly testing should cover at least one category of organizational data. The person responsible for recovery should be able to execute a restore within your operational tolerance for downtime.
- **Audit external sharing links.** Pull a report of active shared links from your cloud storage platform. Revoke any "anyone with the link" access on sensitive files. Set expiration dates on new links going forward.

IF YOU'RE BEING TARGETED

- **Your storage provider is a legal target.** A subpoena to a cloud provider can produce your files whether or not you are party to the proceedings. The only complete protection is E2E encryption where the provider holds no key, or local-first storage in a jurisdiction with meaningful legal protection. Consult a lawyer about your specific exposure.
- **Encrypt sensitive files at rest before they reach any cloud storage.** Even with an E2E-encrypted cloud provider, consider encrypting your most sensitive individual files with a separate key before they are uploaded. This creates a second layer that survives a provider key-management failure or a flawed implementation.
- **Treat your offline backup as a sealed archive.** Do not leave an offline backup drive permanently connected to any networked device. Connect it only during scheduled backup windows, then disconnect and store it physically secure and separate from your primary device. A drive that is always connected is reachable by ransomware and by anyone who accesses your device.
- **Data minimization is an active practice, not a one-time sweep.** New sensitive data accumulates continuously. Build a habit of reviewing what you are storing and deleting what is no longer needed at the same interval as your restore tests. The data that does not exist cannot be seized.

▲ COMMON MISTAKES

- **Counting a cloud sync as a backup.** A sync mirrors your current state in real time. If ransomware encrypts a file, the sync replicates the encryption. A backup is a separate, versioned copy you can roll back from. Confirm your solution retains previous versions for a meaningful window.
- **Never testing the restore.** A backup that has never been successfully restored is an assumption that has not been verified. The failure mode (discovering the backup is corrupt or incomplete during an actual crisis) is worse than the fifteen minutes it takes to run a test restore on a calm day.
- **Choosing E2E cloud and then enabling "help me recover my account" options that give the provider a key copy.** Some providers offer account-recovery assistance that requires them to hold a key escrow copy. Enabling that option negates the E2E model. Read the recovery options carefully before enrolling; choose the path that keeps the key solely in your custody.
- **Leaving an offline backup drive permanently connected.** Connected means reachable. A drive that stays plugged in is reachable by ransomware and by anyone who accesses your device. Disconnect it after each backup window and store it separately.

You're done with this pillar when: you have three verified backup copies with one offline on an encrypted drive stored offsite, you have successfully tested a restore and scheduled recurring restore tests, you have migrated primary storage to either an audited E2E-encrypted cloud service or a self-hosted system with full encryption and proper patching, you have completed a data-minimization sweep, and (for organizations) shared drives enforce least-privilege access with a documented retention and deletion policy in effect.

PILLAR 5 · NETWORK & ENVIRONMENT

Network & Environment

Why this matters now

Phase One set a working baseline: you changed your router's admin password, updated its firmware, enabled WPA3 or WPA2 with a strong passphrase, and created a guest network for visitors. That matters. What it does not do is actually isolate the devices on your network from each other. A flat home network places your work laptop, your smart TV, your baby monitor, and every IoT appliance on the same broadcast domain. Any device can reach any other. Compromising a poorly-secured IoT device is often trivial. From there, lateral movement to your primary machine is one hop. Phase Two closes that gap with real segmentation, encrypted DNS with filtering, and a clear-eyed decision about when a VPN actually helps you and when it does not.

Do this

Enforce real network segmentation

A guest-network toggle on a consumer router creates a separate SSID. Most implementations still allow some traffic flow between segments or share an underlying network switch fabric. Real segmentation uses VLANs to hard-isolate traffic at the hardware level. If your router or mesh system supports VLANs, configure three segments: one for work and primary devices, one for IoT, and one for guests. IoT devices should have no path to your work segment. Guests should have internet access only.

- ☐ Check your router's documentation for VLAN support. Prosumer and enthusiast models (Unifi, pfSense-based builds, OpenWRT-capable hardware) support this. Most ISP-supplied equipment and basic consumer routers do not.
- ☐ Create three VLANs: Work (your computers, phones, and primary devices), IoT (smart home devices, cameras, printers, game consoles), and Guest (visitors).
- ☐ Apply a firewall rule blocking IoT devices from reaching Work devices. IoT traffic may exit to the internet. Guest traffic may exit to the internet only. Neither segment reaches your Work VLAN.
- ☐ Move every IoT device to the IoT VLAN. This means re-pairing them to the new SSID. Accept the one-time inconvenience. The security gain is significant and permanent.
- ☐ If your current hardware does not support VLANs, evaluate a router upgrade. As an interim measure, confirm the guest toggle is active and that IoT devices are assigned to it, not to your primary network.

Switch to encrypted DNS with filtering

Your DNS resolver sees every domain name every device on your network queries. By default, those queries travel unencrypted over UDP, visible to your ISP, to anyone on a hostile local net-

work, and to DNS-based surveillance tools. Encrypted DNS (DNS-over-HTTPS or DNS-over-TLS) wraps queries in TLS so they cannot be read in transit. A filtering resolver also blocks queries to known malware distribution and phishing domains before a connection is ever established. This is not a substitute for endpoint protection, but it catches a meaningful class of threats at the network layer with essentially zero ongoing effort.

- ☐ Choose a DNS resolver that supports DoH or DoT and offers malware and phishing category filtering. Look for resolvers with published privacy policies that commit to no query logging or retention.
- ☐ Configure encrypted DNS at the router level rather than per-device. This ensures every device on the network benefits, including IoT devices that cannot be individually configured.
- ☐ On devices you travel with, configure DoH or DoT at the OS or browser level independently of your home router. Protection should follow you to hotel networks and public Wi-Fi, not only apply at home.
- ☐ After configuring, run a DNS leak test from a browser-based tool. Confirm your queries are reaching the resolver you intended and not leaking to an unencrypted fallback.

Decide when a VPN actually helps

VPNs are widely misrepresented. Understand what a VPN does and does not do before relying on one for security.

A VPN encrypts the tunnel between your device and the VPN server, then exits to the internet from the VPN provider's IP address. On an untrusted network (hotel, café, airport, conference), it prevents the local network operator and other users on the same network from reading your traffic. It does not make you anonymous. It does not protect you from malware or phishing. It does not prevent your VPN provider from seeing your traffic. It shifts trust from your ISP to the VPN provider; it does not eliminate trust.

- ☐ Use a VPN when connecting to networks you do not control: public Wi-Fi, hotel ethernet, conference networks, shared coworking space connections. The threat model is a hostile network operator or an attacker on the same local network.
- ☐ Do not use a VPN at home as a security upgrade against remote attackers. You control your home network. A home VPN adds latency and a new trusted third party without improving your security posture against remote threats.
- ☐ If you use a commercial VPN provider, choose one with a publicly audited no-logs policy and a clear jurisdiction. Research the provider before trusting it. Free VPN services frequently monetize user query data, which defeats the stated purpose.
- ☐ Do not install VPN clients from unfamiliar sources. This category includes many apps promoted heavily in ads and social media. A malicious VPN client is a near-total compromise of your network traffic.
- ☐ For high-risk activities involving sources, sensitive reporting, or evading targeted surveillance: a commercial VPN alone is not sufficient. Consult current guidance from EFF. Tor provides stronger anonymity for specific high-sensitivity activities and should be understood as a distinct tool from an everyday VPN.

SCALING TO YOUR TEAM OR ORGANIZATION

The organizational equivalent of network segmentation is Zero Trust Network Access (ZTNA): instead of placing authenticated users inside a perimeter and treating internal traffic as trusted, ZTNA requires every access request to be authenticated and authorized against specific resources regardless of network location. The plain-language goal is: only the right people reach only the specific systems their role requires, and nothing else.

- Retire broad VPN access that places users inside the network with wide lateral movement capability. Replace it with per-application access controls tied to identity and device posture. A contractor at a coffee shop and a staff member in the office should face the same authentication and authorization check.
- Your identity provider, not network location, should gate access to internal systems. Enforce this even for on-premises users. Location is not an identity signal you can trust.
- Segment your internal network so that a compromised workstation cannot reach critical servers, databases, or administrative interfaces without additional authentication. Apply micro-segmentation even on premises.
- Audit which systems are reachable from the internet. Exposed administrative panels, internal dashboards, and legacy applications are frequent initial access points. Remove exposure that is not operationally necessary.

IF YOU'RE BEING TARGETED

- A sophisticated adversary can compromise a router at the firmware level. If you believe you are under active, targeted network-level attack, replacing consumer hardware with a device running open-source firmware (such as OpenWRT) under your direct control reduces the attack surface from manufacturer backdoors and unpatched proprietary code.
- Use a dedicated device for your highest-sensitivity work. That device connects only to your segmented Work VLAN, uses encrypted DNS, and never touches unknown Wi-Fi without a VPN active and verified.
- IMSI catchers impersonate cell towers and are used by some law enforcement agencies and nation-state actors to intercept mobile traffic. In high-risk environments, treat your cellular network as hostile for sensitive communications. Use Wi-Fi over a trusted connection with Signal or an equivalent end-to-end encrypted tool.
- Physical network taps are a real attack vector against high-value targets. In environments where physical access to your network infrastructure is not guaranteed, treat wired connections to shared infrastructure with the same skepticism as public Wi-Fi.

▲ COMMON MISTAKES

- Believing the guest-network toggle fully isolates IoT devices from your primary network. Most consumer router implementations do not enforce this at the hardware level. Verify the implementation or upgrade.
- Running a VPN at home and treating it as a security improvement. On a network you control, a VPN provides no meaningful protection against remote attackers and introduces a new trusted third party into your traffic path.
- Choosing a DNS resolver purely on speed or reputation without reviewing its privacy policy. A fast resolver that logs and sells your queries is worse than your ISP's default in terms of data exposure.
- Treating a filtering DNS resolver as sufficient defense against phishing. DNS filtering blocks known-bad domains. It does not stop phishing on newly registered domains, credential theft delivered over HTTPS, or malware hosted on legitimate cloud infrastructure.
- Leaving administrative interfaces for network hardware reachable from the internet. Your router's admin panel, NAS management page, and any internal dashboard should be accessible only from your local network, and ideally only from a designated management segment.

You're done with this pillar when: your network has three isolated segments (Work, IoT, and Guest), all devices use encrypted DNS with filtering, and you can clearly state when a VPN benefits you and when it does not.

PILLAR 6 · AUTOMATION & ORCHESTRATION

Automation: Make Security Automatic

Why this matters now

Phase One made the right behaviors automatic: updates apply without prompting, backups run on a schedule, your password manager fills credentials without manual intervention. That eliminates a class of failures caused by human inconsistency. Phase Two goes further. The controls in this phase enforce policy, not just convenience. Conditional access means a device in a compromised state is denied entry before a human decides to investigate. Auto-wipe means a stolen device destroys its own contents before an adversary can brute-force the lock. Verified restore tests mean your backup is confirmed insurance rather than an untested assumption. These are active enforcement mechanisms. They act whether or not you are paying attention, and that is precisely the point.

Do this

Configure conditional-access rules

Conditional access gates authentication decisions on real-time signals about the requesting device and context. A login attempt from an unfamiliar country at 3 a.m. on an unmanaged device triggers a step-up challenge or a block. A login from a recognized, compliant device does not. The system enforces the decision. No one needs to be watching a dashboard.

- ☐ Review whether your primary identity providers (Google, Microsoft 365, Apple ID, or an SSO platform such as Okta or Entra ID) support conditional access or risk-based authentication. Enable it at every account that offers it.
- ☐ Configure trusted device lists in your identity providers. Flag your own devices so that logins from them receive reduced friction. Flag unrecognized devices for step-up authentication or an immediate notification to you.
- ☐ Enable new-device and suspicious-activity alerts at every provider that supports them. These are not conditional access, but they are the practical equivalent available to individuals without an enterprise platform.
- ☐ For accounts that provide no conditional-access controls, hardware security keys (covered in Pillar 1, Phase Two) are a strong compensating control. The key must be physically present for authentication to succeed, regardless of where the credential was stolen or from where the login is attempted.

Enable auto-wipe after repeated failed unlock attempts

A stolen phone with a strong passcode takes significant time to crack with dedicated hardware. Auto-wipe after a threshold of consecutive failed attempts is a last-resort control that destroys

local data before indefinite brute-force is possible. It is not for every situation. It is appropriate for anyone whose phone holds sensitive contacts, sources, or communications.

- ☐ Confirm your device backup is current and has been tested before enabling auto-wipe. Once the threshold is triggered, the wipe is immediate and irreversible. You must be able to restore from backup.
- ☐ On iOS: Settings → Face ID (or Touch ID) & Passcode → Erase Data. The trigger is ten consecutive failed unlock attempts.
- ☐ On Android: the equivalent setting depends on manufacturer and OS version. On stock Android, look under Security settings. Some manufacturers surface it through Device Admin apps. Verify the behavior on your specific device before relying on it.
- ☐ Set a passcode of at least six digits. A four-digit PIN is crackable in minutes with dedicated hardware even with auto-wipe enabled if the device permits enough attempts before wiping. A six-digit or longer alphanumeric code substantially increases the time-to-crack window for the attempts that do occur.

Establish recurring automated restore tests

A backup that has never been restored is an assumption. Backup processes fail silently: drives fill up, cloud credentials expire, sync errors go unreported, and software configurations reset after updates. The only way to know your backup works is to restore from it.

- ☐ Schedule a recurring calendar event every ninety days labeled "Restore test." The task: restore at least one file or folder from each backup destination (local drive and cloud). This takes five to ten minutes. A failed restore at this step is a broken backup found before you needed it.
- ☐ Once per year, perform a full restore test to a clean device or a clean OS installation. This confirms you can recover from a total loss, not only retrieve individual files.
- ☐ After any significant change (new cloud provider, new external drive, new backup software version), run an immediate restore test before trusting the new configuration.
- ☐ Log your restore tests. A simple entry, date, source, what was restored, result, gives you a record to consult when an actual recovery event occurs. It also makes it easy to see if testing has lapsed.

Manage account lifecycle deliberately

Organizations automate access provisioning through joiner-mover-leaver workflows. At the personal level, the equivalent is deliberate habit: when you add a new service, provision it intentionally. When you stop using a service, deprovision it fully. Accounts that persist after you have stopped using them are attack surface with no benefit.

- ☐ When signing up for any new service: generate a unique password from your manager, use a dedicated email alias if the service warrants it, enroll MFA immediately, and save recovery codes before closing the registration window.
- ☐ When stopping a service: delete the account, not just the app or the bookmark. Revoke any OAuth grants it holds (covered in Pillar 3, Phase Two). Remove stored payment methods before closing the account. Confirm the deletion confirmation email arrives.

- ☐ Add account closure to your quarterly security review agenda. Flag every service you have not actively used in six months and close it during that session.

SCALING TO YOUR TEAM OR ORGANIZATION

- Conditional access at the organizational level encodes policy as enforcement: "A device that has not received a security patch within 30 days may not access email or internal applications." Your identity provider enforces this without manual review. Configure it once; it runs continuously. Start with your highest-sensitivity applications and expand from there.
- Automate your joiner-mover-leaver process completely. When someone joins, provision exactly the access their role requires from a predefined template, no more, no less. When someone moves to a different role, update their access that day. When someone leaves, revoke all access in a single coordinated action before they leave the building. Manual offboarding that takes days is a documented and recurring vector for post-separation unauthorized access.
- Run scheduled access reviews quarterly. Your identity governance platform, or a spreadsheet reviewed with managers if you lack one, should confirm that every person's current access still matches their current role. Stale permissions accumulate silently. They are found routinely in post-incident reviews.
- Automate restore testing for organizational backups on the same cadence as individual testing, at minimum quarterly. Assign specific ownership to a named individual. An untested organizational backup is a policy document, not a recovery capability.

IF YOU'RE BEING TARGETED

- Auto-wipe is a meaningful control under physical duress. If your device is seized, it limits extractable data before you can act remotely via Find My or a remote wipe command. Know the trigger conditions, have a current verified backup, and understand the wipe is irreversible.
- A dead-man's switch is a mechanism that takes a defined protective action if you fail to check in on a schedule: for example, notifying your lawyer or releasing encrypted materials to a trusted contact if you do not authenticate within a specified window. This is lawful when used to protect journalistic sources or sensitive documents. Services exist for this purpose. Consult a lawyer about jurisdiction-specific implications before deploying one.
- Automation platforms are themselves a high-value target. A sophisticated adversary who gains access to your backup provider, identity provider, or automation rules can disable protections, modify conditional-access policies, or exfiltrate backup archives. Treat credentials for these services with the same rigor as your primary accounts: hardware key, unique strong password, no SMS fallback.

▲ COMMON MISTAKES

- Enabling auto-wipe before verifying a current, working backup. A wipe triggered on a device with a broken or outdated backup is permanent data loss with no recourse.
- Assuming backups are running because they were configured once. Status indicators that show green are not a restore test. The test is a restore. Run one.
- Treating account deletion as app deletion. Deleting the app leaves the account active, the data stored, and all OAuth grants intact. The attack surface remains unchanged.
- Dismissing conditional-access or suspicious-login alerts without investigating. An alert that "a new device signed in from an unknown location" that you dismiss because it is probably you is worth nothing. Investigate every alert that does not match your own behavior, even briefly.
- Automating onboarding without automating offboarding. Permissions accumulate over time. The longer someone holds access they no longer need, the more likely that access is eventually misused or simply forgotten until an incident reveals it.

You're done with this pillar when: conditional access or risk-based authentication is active on your critical accounts, auto-wipe is enabled on your mobile devices with a verified current backup, and you have completed and logged at least one restore test in the past ninety days.

PILLAR 7 · VISIBILITY & ANALYTICS

Visibility & Awareness

Why this matters now

Phase One gave you point-in-time visibility: you checked Have I Been Pwned once, turned on sign-in alerts, reviewed which devices were logged in to your key accounts. That was a snapshot. Phase Two makes visibility continuous. Breaches happen after your last check. Fraudulent accounts get opened in your name without triggering any alert you configured in Phase One. Organizations accumulate audit log data that no one ever reviews. Sophisticated intrusions dwell undetected for months not because signals are absent, but because no one is watching for them. This pillar closes that gap: monitoring that runs continuously, a documented baseline so you can recognize deviations from normal, periodic active checks for signs of compromise, and a simple habit for staying current on threats that are actually relevant to your situation.

Do this

Set up continuous breach and credit monitoring

Have I Been Pwned notifications alert you when your email address appears in a newly catalogued breach dump. That remains important. Layer on top of it: monitoring for new accounts and lines of credit opened in your name. Credential theft and identity fraud are distinct threats that require distinct monitoring.

- ☐ If you have not already done so, place a credit freeze at all three major bureaus (Equifax, Experian, TransUnion) and at NCTUE and Innovis. A freeze prevents new credit from being opened in your name without an active unfreeze. It is free, reversible, and has no negative impact on your existing credit or accounts.
- ☐ Set up credit monitoring at each bureau's own portal or through a service you have researched. The goal is an alert when a new inquiry or new account appears in your name. A freeze blocks new credit; monitoring tells you when someone attempts it.
- ☐ Verify your Have I Been Pwned notification addresses are still active and receiving mail. If you changed email providers since Phase One, re-register your current addresses.
- ☐ Review your password manager's breach-monitoring dashboard during your quarterly security review. Rotate any credentials flagged since the last review.
- ☐ Treat every breach notification as an action item. When you receive one: rotate the compromised credential immediately, check for reuse across other accounts, and watch the affected service for unauthorized activity over the following thirty days. Breach notifications are not informational. They require a response.

Establish your baseline

You cannot reliably detect an anomaly without a documented definition of normal. Establishing a baseline takes one focused session. Every future review becomes faster and more reliable because you have a concrete reference to compare against.

- ☐ Log in to your primary email, cloud storage, and any SSO provider and record: which devices are currently signed in, which geographic regions your logins originate from, which connected apps and OAuth grants are active, and what your typical login times look like.
- ☐ Store this baseline somewhere you can reference: a locked note in your password manager, an encrypted file, or a physical document stored securely. Capture specific device names, approximate geographic locations, and the full list of connected applications.
- ☐ During every subsequent quarterly review, compare current state against your documented baseline. Any device, location, or connected application added since your last review that you did not deliberately add warrants investigation before you move on.

Run periodic signs-of-compromise checks

Active compromise often leaves traces before it causes visible damage. Knowing what to look for and checking on a schedule is the civilian version of threat hunting. This is distinct from waiting for an alert. You are actively looking.

- ☐ Check active sessions and sign-in history on your primary email, cloud storage, and identity providers. Look for unfamiliar device names, unexpected geographic locations, or login times that do not match your activity. Most providers surface this under Security or Privacy in account settings.
- ☐ Check your email account for forwarding rules and filters you did not create. Attackers who compromise email frequently install silent forwarding rules to exfiltrate mail. In Gmail: Settings → See all settings → Filters and Blocked Addresses, then Forwarding and POP/IMAP. In Outlook: Settings → Mail → Rules. Review both locations.
- ☐ If you manage a domain, check your DNS records. Unexpected changes to MX records, SPF, DKIM, or nameservers indicate domain hijacking or a compromised registrar account.
- ☐ On your devices, review recently installed applications, recently added browser extensions, and recently granted accessibility permissions. All three are common persistence mechanisms used by malware and stalkerware.
- ☐ Run this full check at minimum quarterly. Run it immediately after any suspected phishing attempt, after a device is lost or stolen, and after any unexpected breach notification arrives.

Build a personal threat-intelligence habit

Enterprise threat intelligence is a full-time discipline. Personal threat intelligence is a fifteen-minute-per-week habit: stay aware of threats that are actually relevant to your threat model, not every security headline. The goal is targeted awareness, not general anxiety.

- ☐ Identify which platforms and vendors you actually depend on: your OS, your browser, your email and cloud storage providers, your primary communication tools. Follow their security bulletins or advisories. Most publish RSS feeds or mailing lists for this purpose.
- ☐ Follow two or three security-focused sources relevant to your context. For journalists and activists, organizations such as EFF, CPJ, and Access Now publish practical advisories targeted at high-risk individuals. Following the general threat-intelligence industry is noise unless that industry is your field.
- ☐ When a critical vulnerability is announced in software you use, apply the available patch before the end of the day. Do not wait for your next scheduled update window. Critical vulnerabilities are actively exploited within hours of public disclosure.
- ☐ Note recurring threat patterns specific to your sector or community. Activists, journalists, and researchers face spearphishing campaigns designed around their community's context: colleague names, ongoing projects, organizational structures. Awareness of active campaigns lets you apply heightened skepticism to inbound communications during those periods.

SCALING TO YOUR TEAM OR ORGANIZATION

- Turn on audit logging everywhere it is available: your cloud identity provider, your email platform, your cloud storage, your internal applications. Most platforms support logging but have it disabled by default. Log entries collected retroactively after an incident do not exist. Enable logging now, before you need it.
- Designate a specific person to review logs on a schedule, weekly or bi-weekly for a small organization. You do not need a SIEM platform to do this. A filtered export of high-risk events (failed authentication spikes, privilege escalations, bulk downloads, new admin account creation) reviewed in a spreadsheet provides real detection capability at small-org scale.
- Define organizational normal: expected login hours, expected geographic distribution of staff, typical data-access volumes. Document it. Your baseline makes anomaly detection tractable without automated tooling and gives you a reference when something looks wrong.
- Establish a response procedure before you need it. Who gets notified when a suspicious login alert fires? Who decides whether to lock the affected account? Who coordinates the investigation? An alert with no response procedure in place is not a security control.

IF YOU'RE BEING TARGETED

- Breach monitoring and credit monitoring surface broad criminal fraud. They will not surface a targeted intrusion. A targeted adversary does not open credit cards in your name. They access your accounts, read your communications, and monitor your activity quietly. Signs-of-compromise checks (active sessions, email rules, connected apps, device extensions) are the relevant controls for targeted threats.
- If you discover evidence of account compromise, assume every account connected to the same identity is also compromised. Contain first: lock everything. Then remediate: change credentials. Do not remediate one account at a time while the adversary pivots to the others.
- Targeted spearphishing is crafted to match your context with specificity that generic awareness training does not prepare you for. Your threat-intelligence habit should include active monitoring of advisories for your community. EFF, Access Now, and Citizen Lab publish specific campaign warnings targeted at journalists, activists, and researchers. Review them regularly.
- At elevated risk levels, periodic device forensics can detect indicators of compromise including commercial surveillance software. This goes beyond routine self-checks. Access Now's Digital Security Helpline provides direct, confidential technical assistance to journalists, activists, and human rights defenders. Engage them directly if you suspect advanced compromise.

▲ COMMON MISTAKES

- Treating breach monitoring as the whole visibility program. Breach notifications describe historical data exposure. They say nothing about active sessions, account takeover in progress, or unauthorized access happening right now.
- Enabling audit logging and never reviewing it. Logs that are collected but never examined cost storage and produce zero security value. If you cannot commit to regular review, configure automated alerts for high-severity events so that something forces a human response.
- Skipping the baseline step. Anomaly detection without a documented baseline is guesswork. A session review that takes two minutes against a recorded baseline turns into an hour of uncertainty without one.
- Following too many security sources rather than sources specific to your threat model. Headline fatigue is real. Fifteen sources covering enterprise nation-state threat actors are noise for a journalist or nonprofit staffer. Stay specific to what is actually relevant to your situation.
- Applying critical patches on your normal update schedule rather than urgently. A critical severity rating means active exploitation is either underway or imminent. Waiting for a weekly patch window means waiting through live exploitation of a known vulnerability in software you are running.

You're done with this pillar when: continuous breach and credit monitoring are active and verified, you have a documented baseline for your key accounts, and you have completed at least one full signs-of-compromise check against that baseline.

Phase Two Completion Checklist

Use this checklist to confirm you have completed every pillar before considering Phase Two done. Each item maps to concrete steps covered in the pillar chapters. If an item is not checked, go back to that chapter.

Pillar 1: You

- ☐ Enrolled at least one hardware security key on your primary email account and a second high-value account.
- ☐ Removed SMS as a two-factor method on every account that offers a stronger alternative.
- ☐ Audited all remaining TOTP/authenticator app enrollments and confirmed they are backed up securely.
- ☐ Created at least one separate identity (email address and persona) for a distinct role or context you have identified as needing compartmentalization.
- ☐ Confirmed your password manager holds the credentials for your compartmentalized identity and that identity's recovery codes are stored offline.
- ☐ Reviewed and updated account recovery options so none of your primary accounts fall back to SMS or an unprotected secondary address.

Pillar 2: Your Devices

- ☐ Made an explicit posture decision for your primary mobile device: hardened mainstream (iOS or stock Android) or de-Googled/GrapheneOS, using the decision guide in the chapter.
- ☐ Configured your chosen device posture according to the steps in the chapter, including lockdown of biometrics, disable unused radios, and reviewed app permissions post-migration.
- ☐ Identified whether you need a dedicated device for sensitive work and, if so, have either acquired one or have a concrete plan to do so.
- ☐ Confirmed full-disk encryption is verified and active on every device in your inventory from Phase One.
- ☐ Updated your device inventory to reflect posture decisions and any new devices added since Phase One.

Pillar 3: Apps & Software

- ☐ Audited all OAuth "connected app" grants on your primary email, cloud storage, and any other accounts that support OAuth review, and revoked everything you do not actively use.
- ☐ Applied a vetting checklist to any new SaaS tools adopted since Phase One, or adopted during Phase Two, and removed any that failed.
- ☐ Implemented basic application allowlisting or equivalent on at least your highest-risk device.
- ☐ Established a sandbox or throwaway environment (virtual machine, secondary profile, or isolated device) for opening unknown or untrusted files and links.

- ☐ Confirmed that all remaining installed apps still have only the permissions approved during Phase One's permission audit, and revoked any that have crept back.

Pillar 4: Your Data

- ☐ Completed the 3-2-1 backup rule: three copies, two different media, one offsite.
- ☐ Performed and documented a full tested restore from your backup. The restore succeeded and the data was intact.
- ☐ Confirmed at least one backup copy is kept offline and encrypted.
- ☐ Made an explicit decision on cloud storage strategy: end-to-end encrypted cloud, local-first, or self-hosted, using the decision guide in the chapter, and migrated accordingly.
- ☐ Completed a data-minimization sweep: deleted or archived data you no longer need from cloud accounts, devices, and shared drives.
- ☐ Reviewed and cleaned up over-shared links, folders, or documents identified during or after Phase One.

Pillar 5: Network & Environment

- ☐ Enforced genuine network segmentation: work/personal, IoT, and guest networks are on separate VLANs or SSIDs with firewall rules that prevent lateral movement between them.
- ☐ Enabled encrypted DNS with filtering on your router or primary devices and confirmed it is working.
- ☐ Made an explicit decision on VPN use based on the criteria in the chapter: when it helps, when it does not, and what threat it actually addresses for you.
- ☐ Confirmed that all IoT and smart-home devices remain isolated from devices holding sensitive data.
- ☐ Reviewed which devices and services are exposed to the internet from your home network and closed or restricted anything not intentionally public.

Pillar 6: Automation

- ☐ Configured conditional-access or equivalent rules on your primary accounts: lock out unfamiliar locations or devices and require re-authentication after periods of inactivity.
- ☐ Established automated onboarding and offboarding procedures for any organization roles, including immediate access revocation for departing members.
- ☐ Scheduled and conducted at least one access review since completing Phase One: confirmed every account, key, and shared credential is still necessary and held by the right person.
- ☐ Enabled auto-wipe after repeated failed unlock attempts on your primary mobile device.
- ☐ Scheduled recurring automated restore tests at a defined interval (quarterly at minimum) and confirmed the schedule is in your calendar.

Pillar 7: Visibility & Awareness

- ☐ Set up continuous breach monitoring on your primary and compartmentalized email addresses and confirmed alerts are configured and routing correctly.
- ☐ Initiated credit monitoring on all three major bureaus and have a process for reviewing alerts.
- ☐ Turned on and reviewed audit logs for your primary accounts and any organizational systems you manage.
- ☐ Established a documented baseline of normal activity: what devices normally sign in, from where, and at what times.
- ☐ Performed a signs-of-compromise check on your primary devices and accounts using the criteria in the chapter.
- ☐ Built a minimal personal threat-intelligence habit: a defined set of sources you check on a regular schedule for relevant threat information.

Appendix A: NSA Phase Two Crosswalk

This guide adapts the NSA Zero Trust Implementation Guideline: Phase Two (January 2026). The NSA framework describes 41 activities and 34 capabilities organized across seven pillars, with the stated goal of helping organizations "integrate distinct ZT solutions." The table below maps each pillar in this guide to its corresponding NSA pillar and lists representative Phase Two capabilities and activities from the original. This mapping is representative, not exhaustive. The NSA document contains the authoritative and complete list.

Mapping of Houston Labs guide pillars to NSA Phase Two pillars and representative activities

This guide's pillar	NSA pillar	Representative NSA Phase Two activities
You	User	Phishing-resistant MFA integration (hardware keys, passkeys); privileged identity management; identity federation and single sign-on enforcement; joiner/mover/leaver lifecycle automation; continuous identity risk scoring
Your Devices	Device	Endpoint detection and response (EDR) deployment; compliant-device gate (Comply-to-Connect integration); device health attestation feeding access decisions; mobile device management policy enforcement; behavioral analytics on endpoint telemetry
Apps & Software	Application & Workload	Application allowlisting and software supply-chain vetting; OAuth and third-party integration governance; sandboxed execution environments; workload identity and least-privilege service accounts; SaaS security posture management
Your Data	Data	Data loss prevention (DLP) and digital rights management (DRM) integration; data classification automation building on Phase One tagging foundations; encrypted storage with tested recovery; data-minimization and retention enforcement; least-privilege data access controls
Network & Environment	Network & Environment	Micro-segmentation enforcing east-west traffic controls; encrypted DNS with filtering; Zero Trust Network Access (ZTNA) replacing legacy perimeter VPN; software-defined perimeter; continuous network traffic analysis feeding anomaly detection
Automation	Automation & Orchestration	Security orchestration, automation, and response (SOAR) integration; conditional-access policy automation; automated provisioning and deprovisioning; orchestrated access reviews; policy-based auto-remediation on detected violations
Visibility & Awareness	Visibility & Analytics	Cross-pillar analytics correlating identity, device, and network telemetry; security information and event management (SIEM) integration with ZT data sources; behavioral baseline establishment and anomaly alerting; continuous breach and threat-intelligence feeds; extended detection and response (XDR) consolidating endpoint and network signals

NOTE

The NSA Phase Two pillar names used in this table are: User; Device; Application and Workload; Data; Network and Environment; Automation and Orchestration; Visibility and Analytics. These map one-to-one with this guide's pillars. The activities listed above are drawn from the representative capabilities described in the NSA guideline and are intended to help readers locate relevant sections of the source document, not to summarize it completely.

Appendix B: Sources & Acknowledgements

Primary source

NSA Zero Trust Implementation Guideline: Phase Two. National Security Agency, January 2026. This guide's framework, pillar structure, and capability taxonomy are adapted directly from this document. The NSA guideline is a U.S. government work and is in the public domain in the United States. Houston Labs has translated, adapted, and substantially rewritten the material for civilian use. The original document is freely available through the NSA's public cybersecurity resources.

Series predecessors

NSA Zero Trust Implementation Guideline: Phase One. National Security Agency, January 2026. The foundation-level companion to the Phase Two document. Phase One covers 36 activities and 30 capabilities for establishing a secure ZT foundation. Houston Labs adapted this document as the basis for *Houston Labs Civilian Zero Trust: Phase One, Build Your Foundation*, the required prerequisite to this guide.

Houston Labs Civilian Zero Trust: Phase One, Build Your Foundation. Houston Labs LLC, First Edition 2026. The direct prerequisite to this guide. Readers must complete Phase One before working through Phase Two.

Houston Labs Civilian Zero Trust: Primer. Houston Labs LLC. The conceptual introduction to the series. Covers what Zero Trust means for civilians, the threat model, and how to think about the framework before doing any implementation work. Required reading before beginning either Phase One or Phase Two.

Supporting references

NSA "Embracing a Zero Trust Security Model." National Security Agency. The original high-level NSA guidance document introducing the Zero Trust model for U.S. government and contractor environments. A useful reference for the conceptual foundations underlying the implementation guideline series.

NIST Special Publication 800-207: Zero Trust Architecture. National Institute of Standards and Technology, August 2020. The foundational federal definition of Zero Trust architecture, its tenets, deployment models, and use cases. A U.S. government work in the public domain.

CISA Zero Trust Maturity Model, Version 2.0. Cybersecurity and Infrastructure Security Agency, April 2023. A maturity model organized across five pillars (Identity, Devices, Networks, Applications and Workloads, Data) that complements the NSA implementation guideline. A U.S. government work in the public domain.

Public-domain materials

The NSA Zero Trust Implementation Guideline (Phase One and Phase Two), the NSA "Embracing a Zero Trust Security Model" document, NIST SP 800-207, and the CISA Zero Trust Maturity Model are all works of the United States federal government and are in the public domain in the United States under 17 U.S.C. 105. Houston Labs has adapted these works under that public-domain status. Our adaptation, including the civilian translation, organization, prose, checklists, and supplementary guidance, is the independent work of Houston Labs LLC.

Independence disclaimer

This guide is an independent publication by Houston Labs LLC. It is not affiliated with, authorized by, sponsored by, or endorsed by the National Security Agency, the National Institute of Standards and Technology, the Cybersecurity and Infrastructure Security Agency, or any other U.S. government agency. References to government documents are for attribution and reader guidance only.