The Lean AI-Era Cyber Defense

A Phased Blueprint for Cyber Resilience in Growing Organizations

For IT and Business Leaders of Growing Organizations

How to read this paper. This blueprint is architecture-first and vendor-neutral. Each phase opens with the capability you need and what to look for in any solution, then a clearly marked Cisco recommendation for teams who want a single integrated stack. If you run another vendor's infrastructure or security tools today, the phases still apply: adopt the operating model, keep what works, and treat the Cisco sections as one fully integrated way to deliver it. The paper is also deliberately sequenced. You cannot defend what you cannot see, so it starts with assessing what you run. You cannot patch your way out of unsupported gear, so modernization comes next. And identity, not the firewall, now decides how far an attacker moves, so the security phases follow from there.

This blueprint is a distilled, lean-team version of Cisco's fuller practitioner guidance for the AI-exploit era, focused on what a growing organization can realistically operate.

A word on mixed estates. The resilience this blueprint delivers is fully portable. Every phase reduces risk on whatever infrastructure you run today, Cisco or not. Consolidating onto one operating model is an efficiency gain that a lean team will feel, but it is not a prerequisite for getting secure. Nothing here requires a rip-and-replace. Start with what you own, close the highest-risk gaps first, and consolidate only where it reduces operational burden.

Executive Summary

The clock just changed, and with it the math of cyber defense. In 2025, a U.S. government competition run by DARPA proved that autonomous AI systems could find and fix software flaws at machine speed: across 54 million lines of real code, competing systems discovered most of the planted vulnerabilities and patched them in an average of about 45 minutes each, and turned up 18 previously unknown real-world zero-days along the way. What defenders built to strengthen software, attackers can turn around. The same capability that patches a flaw in minutes can weaponize one in minutes.

The effect is already visible in breach data. Verizon's 2025 Data Breach Investigations Report found that exploitation of software vulnerabilities rose 34% year over year to one in five breaches, nearly overtaking stolen credentials as the top way in. The surge concentrated on internet-facing edge devices such as VPNs and firewalls, where exploitation grew almost eightfold, from 3% to 22% of exploitation cases. Most telling for a lean team: for critical new edge-device vulnerabilities, the median time from public disclosure to mass exploitation was effectively zero days, while organizations took a median of 32 days to patch and fully remediated only about half within the year. The race to patch is often lost before it starts.

For a growing organization the exposure is sharpest. The same report found that a large share of attacks now fall on organizations under 1,000 employees, and that smaller companies defend with a fraction of the resources: single-digit percentages of IT budget going to security, and by some measures nearly half of smaller firms with no dedicated security staff at all. You are expected to defend at the pace of a large enterprise, on the leanest team of any segment. The instinct is to buy another security tool. That deepens the problem. What a lean team needs is not more tools, it is a resilient operating model: know what you run, keep it current, and make sure a single compromise goes nowhere.

This paper presents that model in five phases:

1. **Assess.** Build one current picture of every asset and identity, ranked by real exposure, without a manual audit.
2. **Modernize.** Turn patching into a published rhythm, retire what can no longer be defended, and let the platform do the work.
3. **Verify.** Stop identity from being the way in, with phishing-resistant access and least privilege.
4. **Contain.** Shrink the attack surface and segment the network so a foothold has nowhere to go.
5. **Outpace.** Detect and respond at machine speed, without building a security operations center.

The architectural idea in one line: **you cannot patch faster than AI can exploit, so resilience, not reaction, is the strategy.** Disciplined patching still matters, but on its own it cannot win the race. Resilience comes from combining it with visibility, modernization, identity, segmentation, and managed response, run as one operating model with one security policy rather than a shelf of disconnected tools. Reachability decides priority, identity decides blast radius, and a lean team wins by operating a resilient architecture rather than chasing every alert.

[VISUAL 1 - The five-phase journey. A horizontal spine: ASSESS -> MODERNIZE -> VERIFY -> CONTAIN -> OUTPACE, each with its one-line promise beneath. Show Phases 1-2 as "foundation" (assess and reduce exposure), Phases 3-4 as "contain" (identity and segmentation), Phase 5 as "operate" (machine-speed response). Note that each phase stands alone and is independently valuable.]

1. The Problem: Attackers Now Move Faster Than You Can Patch

Security teams have managed long lists of vulnerabilities for decades. What changed is the clock. AI can now read code, find exploitable flaws, and write working exploits at machine speed, and the evidence is no longer speculative. In DARPA's 2025 AI Cyber Challenge, autonomous systems analyzed 54 million lines of real-world code and found the majority of planted vulnerabilities, patching them in an average of roughly 45 minutes each, and surfaced 18 genuinely unknown zero-days in the process. Independent research teams have shown the same capability pointed outward: one AI agent discovered a real, previously unknown flaw in widely used open-source software before attackers could weaponize it, and another has reported finding valid vulnerabilities at roughly eighty times human speed. The tooling that helps defenders find and fix flaws is the same tooling that lets attackers find and exploit them.

The effect already shows up in breach data. Verizon's 2025 Data Breach Investigations Report recorded a 34% year-over-year rise in vulnerability exploitation as an entry point, reaching one in five breaches, tracking a broader surge in disclosed vulnerabilities that NIST logged climbing from roughly 28,000 in 2023 to 40,000 in 2024. The pressure is concentrated where a lean team is most exposed. Exploitation of internet-facing edge devices such as VPNs and firewalls grew almost eightfold, from 3% to 22% of exploitation cases. For critical new edge vulnerabilities, the report found the median gap between disclosure and mass exploitation was effectively zero days, while organizations took a median of 32 days to remediate and closed only about 54% within the year.

Two consequences fall hardest on growing organizations. First, internet-facing services are the front line, and the patch window there has collapsed to nothing. Second, end-of-life and unsupported gear becomes the sharpest exposure of all, because a flaw found there has no patch path at all, and smaller organizations tend to keep hardware in service well past its prime.

The DBIR is blunt that attackers have turned toward smaller businesses to offset diminishing returns elsewhere: a large share of attacks now land on organizations under 1,000 employees, 88% of small-business breaches involved ransomware, and these are the same organizations that fund security with single-digit percentages of their IT budget and, by some measures, often with no dedicated security staff at all. Cisco, which has been running frontier models against its own portfolio, reached the same conclusion from the vendor side and has moved to a predictable, scheduled disclosure model in response; the point for this paper is that the trend is corroborated by independent, non-vendor evidence first.

Why the obvious fixes fail

Two reflexes both misfire. The first is to try to out-patch the problem, treating every new advisory as a fire drill. When the disclosure-to-exploitation window is effectively zero, human-paced, one-off patching cannot win on its own, and CVE-by-CVE triage that worked in a slower era no longer scales. The second reflex is to buy another point security tool. For a lean team that deepens the problem it was meant to solve, adding another console, another integration, and another place the picture breaks. Neither reflex addresses the real task, which is to build an operating model where the infrastructure is known, current, and contained by design, so that the speed of the attacker stops being decisive.

2. The Core Insight: Resilience Is an Architecture, Not a Reaction

Patching is essential hygiene. It is not, by itself, a strategy for an era when exploits are written at machine speed. When you cannot reliably win the race between disclosure and exploitation, the goal shifts from reacting faster to being harder to hurt in the first place. That is resilience, and it is a property of how the estate is built and operated, not a product you buy.

Resilience rests on three ideas a tool-first framing misses. First, you cannot defend what you cannot see, so a current, honest inventory of assets and identities is the precondition for everything else. Second, reachability, not raw vulnerability count, decides priority: a fully patched system on an open path to sensitive data can matter more than a flawed one no attacker can reach. Third, identity now decides blast radius. A vulnerability may give an attacker a foothold, but credentials, tokens, and over-permissioned accounts decide how far that foothold travels. Get those three right and the network becomes something a lean team can actually defend: known, current, and compartmentalized, so a single compromise stays a single compromise.

There is a fourth idea that a security-only framing structurally cannot deliver, and it is the one that matters most for a lean team. The network and the security controls are the same estate. The place a vulnerability lives, the identity that can reach it, the segment that contains it, and the policy that governs all three are not separate problems to be stitched together across

separate tools. When infrastructure and security share one operating model, a single security policy can be written once and enforced everywhere, across data center, cloud, campus, branch, and the devices in between, and the same inventory that tells you a device is unsupported is the inventory that drives the segmentation policy protecting it. A security-only platform can watch the network; it cannot see or retire your unsupported hardware, segment the fabric itself, or tie a refresh to a security requirement. That fusion is the difference between a lean team operating a system and a lean team becoming a systems integrator.

This reframes the whole category. The question is not "which security product," it is "can we operate a resilient model with the team we have." That favors one operating model over point tools, one policy over many rule sets, and a sequence that builds durable assets at every step.

3. The Blueprint

Sequence is the principle. You establish visibility first because it is the precondition for prioritizing anything. You modernize next because unsupported gear has no defense. Then you close identity, contain movement, and finally operate at machine speed. Three properties make this a blueprint rather than a wish list.

1. **Each phase stands alone.** Stop after Phase 1 and you still gained a complete, current inventory and a real risk picture. Stop after Phase 2 and your estate is current and defensible.
2. **Each phase has exit criteria.** You know when the next phase is earned, and when it is not.
3. **Nothing is thrown away.** Inventory, identity, segmentation, and detection investments compound rather than reset.

It is also a growth path. The same operating model that keeps a 300-person company resilient carries it into mid-market and, if the business grows, toward enterprise scale, without a change of model.

The spine: **Assess → Modernize → Verify → Contain → Outpace.**

What to look for: one operating model and one security policy

Whatever products deliver the phases, the operational win comes from running them from as few consoles as possible, with as little specialist effort as possible, and from expressing security policy once rather than re-implementing it per tool. A lean team cannot absorb five security dashboards and the integration work between them. Look for a cross-domain operating plane that unifies the network, its security, and identity under one interface, lets you define intent-based policy centrally and enforce it across every control point, and adds AI-assisted operations that reduce alert load rather than add to it. A delivery model that lets a partner or managed service carry the parts you should not staff for matters just as much.

The Cisco recommendation. This is where Cisco is genuinely differentiated, and it is the through-line of this blueprint: Cisco Cloud Control is the cross-domain operating plane that brings networking, security, identity, and observability under one interface, with unified security policy defined centrally and enforced across data center, cloud, campus, branch, and IoT. No security-only platform spans that range, because none of them own the network fabric, the lifecycle data, and the identity and enforcement layers as one estate. That is the fusion the core insight describes, made operational. Cloud Control is in Controlled Availability today, and each phase below is executable now through the domain platforms it unifies (the Meraki dashboard, Nexus Dashboard, Cisco IQ, Duo, and the rest); see the availability column in Appendix C for what ships today versus what is emerging. The recommendation for each capability follows in the phase that introduces it.

You do not need the whole stack to start

A lean team should read the phases as a menu with a clear minimum, not a mandatory bill of materials. Each phase has a minimum-viable move that delivers most of the risk reduction, and a fuller build you grow into only if scale or complexity warrants it. The table below is the honest starting line.

Phase	Minimum-viable move	Fuller build (grow into)
Assess	Turn on inventory from the platform you already run; get one risk-ranked exposure list	Cisco IQ intelligence across the estate + multivendor exposure analytics
Modernize	Standing change windows on a published cadence; refresh or isolate unsupported gear	Platform-driven updates + compensating controls + financed refresh waves + PQC readiness
Verify	Phishing-resistant MFA for admins, finance, and help desk on your existing identity provider	Risk-based access, device trust, session protection, non-human and agent identity
Contain	Isolate management interfaces and unsupported gear; firewall/IPS on exposed paths	Identity-based segmentation and unified policy across all control points
Outpace	Partner-delivered MDR on the telemetry you have	MDR across network, firewall, and identity + AI visibility + agent security

4. Phase 1 — Assess

The promise: a complete, current picture of every asset and identity you run, and where the real risk is, without a manual audit.

Business outcome: you can answer "what do we run, what is exposed, and what do we fix first" in an afternoon, not a quarter.

Everything downstream depends on this. You cannot patch, segment, or govern what you do not know exists, and a spreadsheet compiled last quarter is already wrong. The goal of this

phase is one current, trusted view assembled by the platforms you already run, not by hand. It has two halves that come together into a single priority list.

The first half is the infrastructure inventory: every device and software version, its lifecycle status, whether it is internet-facing, and whether it is approaching or past its last day of support. The second half is the identity inventory: human accounts, and the service accounts and API keys that quietly outnumber them, with a focus on weak or missing multifactor authentication and accounts no one owns. The output that matters is not a raw list, it is a ranked one. Reachability and business impact decide the order, not severity score alone, because a reachable, exposed asset outranks a theoretically worse one an attacker cannot touch.

What to look for (any vendor)

- **Platform-driven inventory** that assembles itself from what your network already reports, rather than a manual audit that is stale on arrival.
- **Lifecycle and advisory intelligence**: which assets are past or nearing end of support, and which advisories actually affect your running versions.
- **Internet-facing and reachability context**, so exposure, not raw CVE count, drives the priority order.
- **An identity inventory** spanning human and non-human accounts, surfacing weak or missing MFA, dormant accounts, and ownerless credentials.
- **Multivendor coverage**, because attack paths do not respect a single vendor's boundary and partial coverage hides the path that matters.
- **Shared context across domains** — the strongest inventories are not standalone. Look for lifecycle and end-of-support intelligence tied to the same inventory that later drives your segmentation and identity policy, so what you learn here is not re-gathered in every later phase.

The Cisco recommendation. Start with Cisco Cloud Control as the operational plane. It provides inventory, topology, and consolidated event information by aggregating the underlying domain-specific platforms into one view, rather than asking a lean team to stitch them together. Those domain platforms are the ones you may already run: the Meraki dashboard for the cloud-managed estate, Nexus Dashboard for the data center, and Catalyst Center for on-prem campus and branch. On top of that inventory, Cisco IQ adds the intelligence layer, surfacing lifecycle status, last-day-of-support dates, and advisory exposure across the Cisco estate, including devices not attached to any controller, and it is included in the Standard and Signature support tiers you may already pay for, so for most organizations this is an activation decision rather than a purchase. Its Resilient Infrastructure Services framework digitizes the exposure assessment as a Resilient Infrastructure Playbook.

For the identity half, Cisco Identity Intelligence (within the platform, drawing from Duo and your existing directories) surfaces weak MFA, dormant accounts, and over-scoped access. For visibility across non-Cisco assets, Splunk Enterprise Security with Exposure Analytics widens the aperture to

the multivendor estate. The exposure-assessment approach and the hardening guidance this blueprint draws on are published on Cisco's Resilient Infrastructure page (https://www.cisco.com/c/en/us/about/trust-center/resilient-infrastructure.html).

For a mixed estate, do not wait for perfect consolidation. Normalize inventory from the network, endpoint, cloud, and identity platforms you already run. The outcome that matters is one reachability-ranked exposure list, whatever its sources.

What you deploy

- **Nothing new to purchase in most cases.** Turn on inventory in Cisco Cloud Control across the domain platforms you already run, and activate Cisco IQ on your existing support.
- **Existing Catalyst 9000 gear not attached to any platform** can still participate, using existing DNA licenses to enable cloud-managed capabilities through Cisco Cloud Entitlement for DNA (https://documentation.meraki.com/Platform_Management/Product_Information/Licensing/Meraki_Licensing/Cisco_Cloud_Entitlement_for_DNA), so unmanaged campus switches join the inventory without new licensing.

Included at no additional cost

- **Cisco IQ** asset, lifecycle, and advisory intelligence, and the **Resilient Infrastructure Playbook**.
- **Cisco Identity Intelligence** identity visibility, drawing from Duo and your existing directories.

Optional

- **Splunk Enterprise Security with Exposure Analytics** where you need a single exposure view across non-Cisco assets.

Exit criteria. You have a current inventory of assets and running versions, a last-day-of-support list with dates, internet-facing assets tagged, an identity inventory with MFA gaps surfaced, and a single risk-ranked priority list ordered by reachability and business impact.

5. Phase 2 — Modernize

The promise: patching becomes routine maintenance on a published rhythm, and gear that can no longer be defended enters a funded retirement plan.

Business outcome: patch management stops being a fire drill and becomes a planned, budgeted activity your change process already absorbs.

This is the phase a security-only vendor cannot deliver, and that makes it the most defensible content in the blueprint. Retiring undefendable risk is not a refresh line item; it is a security control, and only a vendor that owns the lifecycle data, the platform economics, and the financing can turn it into one. With

a ranked picture in hand, the work is to get current and stay current, on a cadence a lean team can actually run. Two things make that possible now. The first is that the disclosure process itself has become predictable. Cisco moved to a scheduled, twice-monthly security disclosure model on the first and third Wednesdays of each month, with seven days of advance notice of exactly which technologies each release will cover. That turns patching from a series of surprises into a standing calendar item you can pre-stage change windows and approvals against. Set remediation targets by exposure tier, not by CVE count: the internet-facing tier within days of an advisory, the next tier in the next window.

The second is that the platform does the heavy lifting, and the economics are decisive for a lean team. Patching a hundred standalone devices can require more than a hundred decisions; patching a platform-managed fleet requires far fewer, with staged rollout, golden images, and rollback built in. This is why the modernization question is less "which box" and more "is it platform-managed," because a platform-managed estate is one a small team can actually keep current. Where a permanent fix cannot be applied immediately, a validated compensating control buys time on the specific vulnerable path without a reboot.

Some gear, though, will never receive another patch. Past last day of support, there is no software path, and no compensating control is a substitute for retirement. This is the argument a budget owner understands: an internet-facing, unsupported device is undefendable risk, not an IT wish-list item. Retirement is phased in waves tied to budget cycles and change windows, and financing can make the spend track the phased rollout rather than land as one capital shock, which is a lever no security-only vendor can offer. This is also the moment to fold in a forward-looking criterion so you refresh once rather than twice: whether a device can support post-quantum cryptography, so a single modernization wave satisfies both the security and the crypto-agility requirement.

What to look for (any vendor)

- **A predictable disclosure and patch cadence** you can align to standing change windows, rather than ad-hoc advisories that force fire drills.
- **Platform-driven updates**: golden images, staged and canary rollout, and rehearsed rollback, so patching is fewer decisions and less manual effort.
- **A compensating-control option** that reduces exposure on a specific vulnerable path while a permanent fix is scheduled, ideally without downtime.
- **Lifecycle intelligence that drives the refresh**, tied to the same inventory that runs your network, so end-of-support status is a live signal rather than a spreadsheet, and refresh is a security decision, not just a hardware one. This is the criterion a security-only tool cannot meet: it can see traffic, but it cannot see or retire your unsupported gear.
- **Financing that tracks a phased roadmap**, so a needed refresh is not blocked by budget timing.
- **Post-quantum readiness as a refresh criterion**, so one modernization wave carries the estate forward.

The Cisco recommendation. Cisco's scheduled disclosure cadence (first and third Wednesdays, with seven days of advance notice) is designed to make patching a planned activity; align standing change windows to it. For platform-driven updates, operate from Cisco Cloud Control as the single plane and let it drive the domain platforms beneath it: staged firmware in the Meraki dashboard for the cloud-managed estate, and image management in Nexus Dashboard for the data center, both with golden-image and rollback patterns. Where an immediate patch is not possible, Cisco Live Protect inserts Cisco-validated runtime shields as a temporary compensating control with no reboot; it is generally available on Nexus 9000 today, with expansion across the broader portfolio planned. For retirement, Cisco IQ separates legacy-but-supported from unsupported and prioritizes what to replace first; and Quantum Ready Assessments evaluate devices for post-quantum readiness so refresh addresses crypto-agility in the same wave.

What you deploy

- **A modernization plan**, not a product: standing change windows on the disclosure cadence, and a refresh wave schedule for unsupported gear tied to budget cycles.
- **The platform-economics case for consolidation** is developed in depth in the companion Cisco blueprint on operational efficiency [LINK TBD], which shows how a single cloud-managed operating model turns modernization from a per-device project into a continuous, low-effort rhythm.

Included / Activate only

- **Staged firmware and rollback** driven from Cisco Cloud Control across the Meraki dashboard and Nexus Dashboard you already run.
- **Cisco Live Protect** shields where supported, as the interim control on high-exposure paths.
- **Cisco IQ** lifecycle prioritization and **Quantum Ready Assessments**.

Exit criteria. Two consecutive disclosure cycles have been absorbed as routine maintenance, the internet-facing tier is remediated within days of an advisory, remediation targets are set by exposure tier rather than CVE count, and every unsupported internet-facing asset has either a scheduled refresh or a compensating control.

6. Phase 3 — Verify

The promise: identity stops being the way in, with phishing-resistant access and least privilege as the default.

Business outcome: a stolen password or a phished token no longer opens the door, and no single account carries more access than its job requires.

A patched estate still falls to a stolen credential. Identity is both the most common way in and the thing that decides how far an attacker travels once inside, and AI has made credential phishing and social engineering cheaper, faster, and more convincing. This phase makes identity the control point: strongly authenticate every human, narrowly authorize every account, and continuously verify both. The most

effective single move is phishing-resistant multifactor authentication, and it is achievable for a lean team because you do not have to do it everywhere at once.

Stage it by blast radius. Make MFA available everywhere first, including the command-line, remote-desktop, and legacy authentication paths attackers probe, because any second factor defeats commodity credential stuffing and password spray. Then move the small, high-value population, administrators, finance, and help-desk staff, to phishing-resistant methods first. That group is small enough to migrate quickly and is exactly who attackers target. Add single sign-on so the secure path is also the easy path, device-trust so corporate access opens only from healthy devices, and protection against session-token theft at both the MFA and SSO layers.

One light but important touch: the accounts that are not people. Service accounts, API keys, and tokens quietly outnumber human identities and are frequently long-lived, over-permissioned, and ownerless. You do not need an enterprise secrets program to start; you need every non-human identity mapped to an owner and a purpose, with the orphaned and over-scoped ones retired. This is also the on-ramp to the agent era, because AI agents are simply the newest and fastest-moving class of non-human identity, and putting identity at the center now means the agent wave arrives partly pre-handled.

What to look for (any vendor)

- **Phishing-resistant MFA** (FIDO2/WebAuthn passkeys, device-bound for high-privilege roles) with a staged path from MFA-everywhere to phishing-resistant for the admin tier.
- **Single sign-on and automated provisioning/deprovisioning**, so access is live on day one and gone on exit.
- **Device trust** tied to your identity provider, achievable without a full device-management rollout.
- **Session-theft protection** at both the MFA and SSO layers, so a stolen token cannot be replayed from an untrusted device.
- **Non-human identity hygiene**: every service account and key mapped to an owner and purpose, with a runway to governed agent identity.
- **Identity signal that feeds enforcement**, not just authentication. The strongest setups let identity and device posture drive network segmentation directly, so a compromised account is contained by the network, not just flagged by the login. That link between identity and enforcement is hard for a standalone identity tool to close.

The Cisco recommendation. Keep your identity provider. Most growing organizations already standardize on one, and this phase augments it rather than replacing it: add phishing-resistant MFA, device trust, session protection, and identity-risk controls where your current stack has gaps. Cisco Duo makes the corporate account the path of least resistance and delivers phishing resistance across the human access cycle. Duo Essentials covers the core of this phase: phishing-resistant authentication, SSO, MFA, Duo Directory, and complete passwordless. Duo Advantage adds Cisco

Identity Intelligence with identity threat detection and risk-based, device-aware access, and Duo Premier adds VPN-less zero-trust access to private applications and agentic identity, which discovers, governs, and secures AI agents as first-class non-human identities. On the non-human and agent side, Cisco's announced intent to acquire Astrix Security (May 2026) deepens non-human identity security, with that capability folding into Cisco Identity Intelligence, Secure Access, and Duo. Where remote access to private applications is in scope, Cisco Secure Access provides zero-trust access without backhauling traffic. The differentiator to note: because this identity context lives in the same operating model as the network, it can drive the segmentation policy in the next phase, not just gate the login.

What you deploy

- **Cisco Duo Essentials** for phishing-resistant MFA, SSO, and passwordless across the workforce, staged by blast radius.
- **Cisco Secure Access** where zero-trust access to private applications for remote and hybrid users is in scope.

Recommended / Optional

- **Cisco Duo Advantage** for identity threat detection and risk-based access as the program matures.
- **Cisco Duo Premier** where VPN-less ZTNA and agentic identity governance are required.

Exit criteria. Phishing-resistant MFA is enforced for the administrator, finance, and help-desk tier, MFA is available everywhere including legacy paths, SSO and device trust are in place for corporate access, and every non-human identity is mapped to an owner or scheduled for retirement.

7. Phase 4 — Contain

The promise: the attack surface is smaller by default, and a foothold has nowhere to go.

Business outcome: a compromised device or account is boxed in, so one incident stays one incident instead of spreading to critical systems.

Not every compromise can be prevented, so this phase assumes a breach and shrinks its blast radius. It has two complementary halves: reduce the surface, then compartmentalize what remains. The first half is proactive hardening, the attack-surface reduction that lowers risk regardless of which vulnerability is discovered next. Every unused service disabled, every management interface closed, every legacy protocol deprecated, and least-privilege configuration applied removes an attack vector before an advisory is even written. This is now a first-class discipline, with infrastructure moving toward secure-by-default settings and stronger device access as a baseline.

The second half is segmentation. A flat network lets a single foothold move toward sensitive data, privileged identities, and critical applications. Segmentation, enforced by identity rather than by IP address or location, limits who and what can reach the high-value targets you identified in Phase 1 and interrupts the routes an attacker would take between them. For a lean team the key is that this is cloud-delivered and identity-based, not a fleet of dedicated appliances and a specialist to tune them. Segmentation can feel daunting, so start with a practical scope rather than boiling the ocean: isolate management interfaces, restrict what can reach unsupported gear, separate guest and IoT from corporate systems, limit admin access to privileged systems, and put intrusion prevention on exposed application paths. Shielding controls on the allowed paths, a firewall and intrusion prevention at the points that matter, block known exploit activity before it reaches vulnerable systems.

The advantage a lean team should insist on here is one security policy, not many. When hardening, segmentation, and shielding are expressed as one intent-based policy and enforced across every control point, the estate stops being a collection of rule sets that drift apart. That is the fusion the core insight described, and it is where an integrated network-and-security operating model does something a stack of point tools cannot.

Be clear about scope. This paper covers network and identity-based segmentation, which is the right level for lean organizations. Workload-level microsegmentation, enforced inside the application tier, remains largely an enterprise and regulated-industry pattern, and a growing organization does not need it to be resilient. If you grow into that requirement, it is a clean extension of the same model rather than a rebuild.

What to look for (any vendor)

- **Secure-by-default hardening**: unused services and management-plane exposure disabled, legacy protocols deprecated, least-privilege configuration applied as a baseline.
- **Identity-based segmentation** that separates users, devices, and IoT by identity and posture rather than by IP or location, delivered without an appliance fleet.
- **Shielding on the allowed paths**: a firewall and intrusion prevention at the perimeter and key internal points, informed by current threat intelligence.
- **One security policy across control points** — intent defined once and enforced consistently across data center, cloud, campus, branch, and IoT, rather than disconnected rule sets per device. This is the integration-only criterion: it requires a plane that owns both the policy and the network enforcement, which a security-overlay tool does not.
- **A perimeter entry point sized to the segment**, from a cloud-managed edge to a dedicated firewall, depending on scale and complexity.

The Cisco recommendation. For hardening, Cisco's Resilient Infrastructure guidance and product-specific hardening guides define the secure-by-default baseline, with the portfolio moving toward disabled-by-default services and stronger device access (for example, deprecating SNMPv1/v2c, Telnet, and FTP, and adding FIDO2 over SSH). For identity-based segmentation, operate from Cisco Cloud Control and use Cisco Access Manager , which delivers cloud-native network access control natively in the Meraki dashboard, enforcing identity-based segmentation via security group tags with no on-premises RADIUS servers or NAC specialists; for mixed or more complex estates, Cisco Identity Services Engine (ISE) remains the deeper option. Network segmentation has traditionally demanded a fabric that only specialists could build; Meraki cloud-managed fabric now stands up a BGP-EVPN VXLAN fabric from the dashboard (https://documentation.meraki.com/Switching/Cloud_Management_with_IOS_XE/Design_and_Configure/Cloud_Fabric_(BGP-EVPN_VXLAN_Fabric)), putting macro-segmentation within reach of a lean team. For shielding, Cisco Secure Firewall (as a dedicated appliance) or Meraki MX (as the cloud-managed edge for smaller sites) provides the perimeter, with Snort-based intrusion prevention on the paths that matter. The through-line is unified policy: intent-based security policy is defined once in Cisco Cloud Control and enforced across these control points, so a lean team maintains one policy rather than reconciling many.

Scope note stated plainly: this blueprint covers network and identity-based segmentation. Workload microsegmentation is an enterprise and regulated-industry pattern and is intentionally out of scope for this segment; it extends the same model if you ever need it.

What you deploy

- **Cisco Access Manager** (Advantage license) for identity-based segmentation, or **Cisco ISE** for mixed and more complex estates.
- **Cisco Secure Firewall** or **Meraki MX** as the perimeter entry point, sized to the site.

Included / Activate only

- **Meraki cloud-managed fabric** (BGP-EVPN VXLAN) for dashboard-driven segmentation where a fabric is warranted, without specialist fabric engineering.
- **Secure-by-default hardening** per Cisco's Resilient Infrastructure and product hardening guides, an operational activity rather than a purchase.
- **Snort-based intrusion prevention** with the Secure Firewall subscription.

Exit criteria. Secure-by-default hardening is applied to management planes and unused services, identity-based segmentation is enforced on the high-consequence paths from Phase 1, a perimeter firewall with intrusion prevention is in place, and no unsupported asset is reachable without a compensating control.

8. Phase 5 — Outpace

The promise: machine-speed detection and response, without building a security operations center.

Business outcome: threats are caught and contained around the clock by a service and by AI-assisted workflows, not by a night shift you cannot staff.

Attackers operate at machine speed while most lean teams operate at human speed, and that gap is where dwell time and damage accumulate. The honest translation of "operate an agentic security operations center" for a growing organization is: do not try to build one. Consume it. Detection and response delivered as a managed service, correlating signals across your network, firewall, and identity telemetry, gives a lean team around-the-clock coverage and expert triage without a 24x7 rotation it cannot hire for. This is the same handoff a resilient architecture makes everywhere else: keep what compounds in-house, and let a partner or managed service carry what you should not staff.

Two additions round out the phase. First, prioritize by risk, not alert volume: the value is in high-confidence findings on externally reachable, high-consequence paths, fed from the exposure picture you built in Phase 1, not in processing every alert. Second, plan for agents. As your organization adopts AI agents, whether ones you build or local coding agents your developers run, they become a new class of endpoint that reasons and acts. Securing them means discovery on the endpoint, control over the tools and connections they reach, and runtime containment, extending the identity work from Phase 3 into the agent era. Much of this rides on the same client you may already deploy for secure access, so it is an extension rather than a new stack.

If you already have an MDR provider, this phase is not a rip-and-replace. The requirement is telemetry-driven: your MDR should see network, firewall, identity, endpoint, and cloud signals and prioritize by exposure and business impact. If your current provider covers endpoint and cloud, extend it to the network, firewall, and identity telemetry the earlier phases produce rather than swapping it out on day one.

What to look for (any vendor)

- **Managed detection and response (MDR)** correlating network, firewall, and identity telemetry, delivered by a partner so around-the-clock coverage does not require in-house staffing.
- **Risk-based prioritization**: high-confidence findings on reachable, high-consequence paths, fed from your exposure picture, rather than raw alert volume.
- **AI-assisted triage and investigation** that amplifies a small team rather than replacing it, with humans in the loop for containment and policy change.

- **Detection fed by the same estate you already run** — the strongest signal comes when network, identity, and firewall telemetry flow from the operating model itself, so the MDR sees native context rather than a thin overlay bolted on top. That native telemetry is the integration-only advantage.
- **Agent security** for AI agents you build or adopt: endpoint discovery, control over the tools and connections they reach, and runtime containment.

The Cisco recommendation. For lean teams, Cisco XDR is most often consumed as Cisco managed detection and response (MDR) through a Cisco partner, correlating network, firewall, endpoint, and cloud signals with humans retained for high-consequence decisions. Splunk provides the underlying data and risk-based alerting where a deeper security-operations capability is warranted. For AI visibility, AI Access within Cisco Secure Access discovers which AI applications and agents are in use across the organization, including AI embedded in other software, giving you the same north-south vantage the companion AI-acceleration blueprint recommends. For local coding agents specifically, Cisco DefenseClaw, integrated into Cisco Secure Client, secures them on the endpoint with discovery, tool and connection control, and runtime containment. Because Secure Access and Secure Client may already be deployed for the identity and access work in Phase 3, this is an extension of what you run, not a new footprint.

What you deploy

- **Cisco MDR through a Cisco partner** on your network, firewall, and identity telemetry.

Recommended / Optional

6. **AI Access in Cisco Secure Access** for AI usage and agent visibility, and **DefenseClaw in Cisco Secure Client** for local coding-agent security, as AI adoption arrives.
7. **Splunk** where a deeper, data-driven security-operations capability is warranted.

Exit criteria. Managed detection and response is active on network, firewall, and identity telemetry with agreed response handling, prioritization is driven by reachability and business impact, and AI usage is visible with runtime controls wherever agents are in use.

9. Where to Start

The blueprint meets you where you are, on whatever infrastructure and tools you run today.

- **On Cisco support already?** Start Phase 1 by activating Cisco IQ on your existing support contract. The inventory and exposure picture is largely an activation, not a purchase, and it is the fastest first win in the paper.

- **Running an identity provider or productivity suite?** Keep it. Begin the Phase 3 motion on the identity you have, staging phishing-resistant MFA for the admin tier first, and add the network-side controls the suite does not include.
- **On the Meraki dashboard?** Phases 2 and 4 are close at hand: staged firmware for the modernization rhythm, and Access Manager for identity-based segmentation, both from the console you already run.
- **Running a mixed or competitor estate?** The phases still apply. Lead with the vendor-neutral inventory and exposure view, use MDR that spans multivendor telemetry, and consolidate the operating model over time rather than all at once. Be clear-eyed that a single-console outcome is easier the more of the estate shares one platform; where it does not, you gain resilience without the consolidation bonus.
- **Running lean?** The blueprint assumes it. Lead with one operating model and AI-assisted operations so a small team is not buried in consoles, and use a partner or managed service for detection, response, and delivery. Every phase boundary doubles as a natural review point.

The Cisco recommendation. If you would rather not assemble this from parts, Cisco delivers all five phases as one operating model under Cisco Cloud Control, with Cisco IQ and Resilient Infrastructure Services as the assessment and remediation path, available directly or through a Cisco partner as a managed service. The advantage is the fewest consoles, the least specialist effort, and AI-assisted operations that offset the headcount a resilient security posture would otherwise require.

10. Getting Started

Start by seeing, not buying. A short exposure assessment, run against your current estate through Cisco IQ or with a Cisco partner, should return three things: a current inventory of what you run and what is past or nearing end of support, a risk-ranked exposure baseline ordered by reachability and business impact, and a phase-readiness review against the exit criteria in Appendix B. For most organizations on Cisco support, the assessment is an activation rather than a new purchase, and it disrupts nothing.

From there, the sequence is one this paper has already written. To scope the assessment or see the operating model in action, contact your Cisco account team or Cisco partner.

[VISUAL 3 - Cisco reference architecture. Show one Cisco Cloud Control plane across the top, aggregating the domain platforms (Meraki dashboard, Nexus Dashboard, Catalyst Center) beneath it. Left to right by phase: ASSESS (Cisco IQ + Identity Intelligence); MODERNIZE (scheduled cadence + staged firmware + Live Protect on Nexus + Cisco Capital); VERIFY (Duo + Secure Access); CONTAIN (Access Manager / ISE + Secure Firewall or Meraki MX + cloud fabric + hardening); OUTPACE (MDR via partner + AI Access in Secure Access + DefenseClaw).

Mark Cloud Control as Controlled Availability, Live Protect expansion beyond Nexus as roadmap, and workload microsegmentation as out of scope.]

Appendix A — Capability-to-Cisco Mapping

The capability you need first; the Cisco way to deliver it on the right.

Phase	Capability you need (any vendor)	Cisco recommendation
1 — Assess	Platform-driven asset + identity inventory, lifecycle/advisory intelligence, reachability-ranked exposure	Cisco Cloud Control inventory (aggregating Meraki dashboard / Nexus Dashboard / Catalyst Center) · Cisco IQ + Resilient Infrastructure Playbook · Cisco Identity Intelligence · Cloud Entitlement for DNA (unattached Catalyst 9K) · Splunk ES with Exposure Analytics (multivendor)
2 — Modernize	Predictable patch cadence, platform-driven updates, compensating control, lifecycle refresh, PQC readiness, financing	Scheduled disclosure cadence (1st/3rd Wed) · staged firmware via Cloud Control (Meraki / Nexus) · Cisco Live Protect · Cisco IQ lifecycle + Quantum Ready Assessments
3 — Verify	Phishing-resistant MFA (staged), SSO, device trust, session-theft protection, non-human identity hygiene	Cisco Duo Essentials (core) · Duo Advantage (ITDR) · Duo Premier (ZTNA + agentic identity) · Cisco Secure Access · Identity Intelligence (+ Astrix, announced)
4 — Contain	Secure-by-default hardening, identity-based segmentation, perimeter shielding, central policy	Resilient Infrastructure hardening guides · Cisco Access Manager (or ISE for mixed estates) · Meraki cloud fabric (BGP-EVPN VXLAN) · Cisco Secure Firewall or Meraki MX · Snort IPS · Cloud Control policy
5 — Outpace	Managed detection and response, risk-based prioritization, AI visibility and agent security	Cisco XDR consumed as MDR via partner · Splunk (deeper SecOps) · AI Access in Cisco Secure Access · DefenseClaw in Cisco Secure Client

Appendix B — Phase Exit-Criteria Checklist

Phase 1 — Assess complete when: a current asset inventory exists / last-day-of-support list with dates / internet-facing assets tagged / identity inventory with MFA gaps surfaced / one risk-ranked priority list ordered by reachability and business impact.

Phase 2 — Modernize complete when: two consecutive disclosure cycles absorbed as routine maintenance / internet-facing tier remediated within days of advisory / remediation targets set by exposure tier, not CVE count / every unsupported asset has a scheduled refresh or a compensating control.

Phase 3 — Verify complete when: phishing-resistant MFA enforced for the admin/finance/help-desk tier / MFA available everywhere including legacy paths / SSO and device trust in place for corporate access / every non-human identity mapped to an owner or scheduled for retirement.

Phase 4 — Contain complete when: secure-by-default hardening applied to management planes and unused services / identity-based segmentation enforced on high-consequence paths / perimeter firewall with intrusion prevention in place / no unsupported asset reachable without a compensating control.

Phase 5 — Outpace complete when: MDR active on network, firewall, and identity telemetry with agreed response handling / prioritization driven by reachability and business impact / agent inventory with runtime controls wherever AI agents are in use.

Appendix C — Status Legend and Sources

Status legend

GA: Generally Available. CA: Controlled Availability. Roadmap: as labeled. Product availability, license tiers, and feature timing should be confirmed against current Cisco documentation before external publication.

Availability: what ships today vs. near-term

Consolidated here so the narrative stays clean. Confirm all of the following against current Cisco documentation and your region before external publication.

Capability	Available today	Near-term / confirm
Scheduled disclosure cadence	Public; twice-monthly (1st/3rd Wed), 7-day advance notice	Cadence began July 2026; scope by product family
Cisco IQ / Resilient Infra Services	GA (SaaS); included at eligible Standard/Signature support	Not activated at Basic; Quantum Ready Assessments rolling out 2026
Cisco Live Protect	GA on Nexus 9000 (NX-OS), incl. N9300	Campus/branch, Secure Firewall, SD-WAN expansion is roadmap
Cisco Cloud Control	Controlled Availability (US, 2026)	Global to follow; execute via domain platforms meanwhile
Cisco Access Manager	GA (Nov 2025); cloud NAC in Meraki dashboard	Depends on MS/MR Advantage licensing; IdP list expanding
Cisco Duo	Essentials / Advantage / Premier GA	Agentic identity at Premier; packaging evolving, confirm tier

Capability	Available today	Near-term / confirm
Astrix Security	Not yet integrated	Announced intent to acquire (May 2026); treat as future
DefenseClaw / AI Access	In Secure Client / Secure Access	Confirm current feature scope

Sources

- DARPA / ARPA-H, AI Cyber Challenge (AIxCC) Final Results, DEF CON 33, August 2025 (autonomous systems found the majority of vulnerabilities across 54 million lines of code and patched most in ~45 minutes; 18 real zero-days discovered). Primary problem-section source; U.S. government.
- Verizon, 2025 Data Breach Investigations Report (vulnerability exploitation up 34% to 20% of breaches; edge/VPN exploitation grew ~8x from 3% to 22%; median disclosure-to-mass-exploitation effectively zero days for critical edge vulns; 32-day median remediation, ~54% fully remediated; SMB attack share and 88% ransomware; security-budget and staffing gap). Primary problem-section source; independent.
- NIST National Vulnerability Database (CVE volume rising from ~28,000 in 2023 to ~40,000 in 2024). U.S. government.
- Supporting independent evidence on autonomous vulnerability discovery: Google Project Zero "Big Sleep" (AI agent found a real, previously unknown flaw in widely used open-source software); XBow (AI-driven testing reporting valid vulnerabilities at roughly 80x human speed); Veracode 2025 GenAI Code Security Report (45% of AI-generated code samples introduced OWASP Top 10 issues). Named-industry / research sources.
- Cisco Trust Center, Resilient Infrastructure (https://www.cisco.com/c/en/us/about/trust-center/resilient-infrastructure.html): secure-by-default hardening, deprecation of SNMPv1/v2c, Telnet, FTP/TFTP, secure NTP/NTS, guidance on end-of-support exposure. Cited for Cisco capability, not problem framing.
- Cisco Blogs, "Strengthening the Foundation: A Predictable, Customer-Focused Response to AI-Accelerated Vulnerability Discovery" (twice-monthly disclosure cadence, 7-day advance notice, bundled CVEs).

- Cisco Blogs / Cisco IQ FAQ and Cisco Live 2026 coverage (Resilient Infrastructure Services three-phase framework, Standard/Signature inclusion, Quantum Ready Assessments); Cisco Live Protect product pages and Nexus 9000 NX-OS Security Configuration Guide (Live Protect mechanics and GA scope).
- Cisco / Duo, "Cisco Announces Intent to Acquire Astrix Security" and "Introducing Duo Agentic Identity"; Duo Editions and Pricing (Essentials/Advantage/Premier scope).

- Cisco Meraki Documentation: Access Manager licensing/ordering (GA date, cloud NAC, identity-based segmentation), Cloud Fabric (BGP-EVPN VXLAN), and Cisco Cloud Entitlement for DNA (Catalyst 9000 onboarding).

Note on sources: the problem section is anchored on independent government and named-industry research (DARPA AIxCC, Verizon DBIR, NIST) ahead of any vendor source, to establish the trend neutrally. Cisco sources are used only to describe Cisco capability and availability, and are labeled where they appear. This blueprint was written entirely from public sources; internal, not-for-distribution material was not used.