California Department of Technology

STATE OF CALIFORNIA | EXECUTIVE BRANCH

Defending at the Speed of AI

An Operational Implementation Guide for Public Sector CIOs

Preparing entity environments for offensive uses of AI, aligned to Cal-Secure 2.0

Companion to the Cal-Secure 2.0 Roadmap and Toolkit

Office of Information Security | California Cybersecurity Integration Center

Distribution: Agency and entity Chief Information Officers, Chief Information Security Officers, and cybersecurity program leads across the California Executive Branch

Version 1.0

Contents

1. Purpose and How to Use This Guide

This guide is a practical companion to the Cal-Secure 2.0 Roadmap and Toolkit. The Roadmap sets statewide direction. The Toolkit lets an entity measure where it stands and where it intends to be. This guide sits between them, translating the direction into an ordered set of operational steps that a Chief Information Officer can begin this quarter.

It is written for the CIO who needs to walk into a leadership meeting with a defensible plan for one question: is our environment ready for adversaries who use AI to find and exploit weaknesses faster than we have ever had to respond. The guide does not assume a large security team or a deep budget. Much of the work described here is already part of running a sound program. What has changed is the pace the threat now demands, and the order in which the work pays off.

How the guide is organized

- Section 2 explains what changed and why the patch-and-respond model many programs were built on no longer holds on its own.
- Section 3 maps the twelve implementation areas to the Cal-Secure 2.0 components and to the six functions of the NIST Cybersecurity Framework 2.0, so the work connects directly to how your maturity is measured.
- Section 4 is the core. It covers twelve implementation areas grouped into three themes: strengthen the foundation, shift the mindset, and align people and partners.
- Section 5 sequences the work across thirty, ninety, and three hundred sixty-five days.
- Section 6 shows how to record progress using the Cal-Secure 2.0 Organizational Profile and Maturity Matrix.
- Sections 7 and 8 address your workforce and how to brief agency and executive leadership.

A note on scope

These twelve areas are a foundation, not a complete defense. They raise the floor against the classes of attack the state can measure today. They do not by themselves stop AI-driven social engineering such as synthetic voice and video impersonation of executives, and they will not catch every novel logic flaw. The goal is resilience: slow the attacker, contain the damage, and recover quickly. Assume some attacks will succeed, and test recovery with the same seriousness you test the attack surface. Where an action touches data residency, privacy, or third-party contracts, validate with Legal before enforcement, and pilot aggressive controls before turning them on across production.

2. What Changed, and Why It Matters Now

For two decades, cyber defense rested on three assumptions: that exploit development was slow and limited by human skill, that low-severity vulnerabilities were rarely worth weaponizing, and that detection and response were fast enough to contain an attack. Each of those assumptions has now broken.

AI models can cross-reference a published vulnerability against the software versions running in an environment and produce a working exploit in a fraction of the time a specialist once needed. The skill floor for offensive work has dropped close to a single instruction. The bottleneck for an adversary is no longer talent. It is access to capability, and that access is widening.

The clearest measure is the time between a vulnerability becoming known and a working exploit existing. That window has collapsed: roughly **eight months in 2018**, about **twenty-three days in 2025**, and **under one day in 2026**. The framing in Cal-Secure 2.0 says the same thing in plain terms: the time-to-exploit window has collapsed from months to hours. A manual approval cycle measured in weeks is now itself a security risk.

Mythos-class capabilities make this concrete. A frontier model has demonstrated the ability to find zero-day vulnerabilities in major operating systems and browsers without human guidance, to chain several distinct memory-corruption bugs into a single working exploit path, and to do so from a single prompt rather than an elaborate setup. In one controlled comparison it produced 181 working exploits against a browser where the prior generation produced two under identical conditions. That is a step change, not a trend line.

These capabilities are in restricted release today. A small group of critical infrastructure operators received early access through Project Glasswing so they could harden their own products before broader availability. California is one of the parties working inside that window. The head start is real, and it is temporary. The same capability will reach adversaries. This guide exists to spend that head start well.

What a capability like this will surface in a typical environment

Pointed at an enterprise, a tool of this class behaves less like a scanner and more like a full stress test of posture. It reliably finds the things human-paced assessments have quietly tolerated for years:

- Misconfigurations across cloud and on-premises systems.
- Default, shared, and reused passwords, and unrotated secrets, keys, and tokens.
- Open or exposed storage and data.
- Weak identity and access controls, and management interfaces reachable from the internet.

The honest framing matters here. This class of tool is not all-knowing. In external testing it failed to break a well-configured, modern, fully patched environment, and it struggled against a realistic operational-technology range within its limits. It was highly effective against environments with outdated software, configuration errors, and reused credentials. That is the entire strategic case for this guide. The work below is what turns your environment into the kind a capable attacker cannot easily move through, before the capability becomes widely available.

3. How This Guide Aligns to Cal-Secure 2.0

Every action in this guide maps to a Cal-Secure 2.0 component (People, Process, or Technology) and to one or more functions of the NIST Cybersecurity Framework 2.0 (Govern, Identify, Protect, Detect, Respond, Recover). It also maps to specific capabilities in the Cal-

Secure 2.0 Capability Map, so the work is traceable in your Organizational Profile and Maturity Matrix. The table below is the crosswalk. The detail follows in Section 4.

#	Implementation area	Cal-Secure component	Primary CSF 2.0 function	Capability Map link
1	Aggressively remediate known risk	Technology	Identify, Protect	Vulnerability Management
2	Harden the perimeter	Technology	Protect	Network and Infrastructure Threat Protection
3	Segment to contain a breach	Technology	Protect	Network Threat Protection
4	Realign prioritization and compress SLAs	Process	Govern, Identify	Vulnerability Management, Metrics and Reporting
5	Validate asset inventory and third parties	Process, Technology	Identify	Asset Management, Supply Chain Management
6	Replace end-of-life technology	Technology	Identify, Protect	Asset Management
7	Improve logging for AI-assisted defense	Technology	Detect	Log Management, Continuous Monitoring
8	Improve cyber resilience	Technology	Recover	Business Continuity, Incident and Disaster Recovery
9	Move from vuln management to exploit prevention	Technology, Process	Protect, Respond	Network Threat Protection, Incident Management
10	Use AI for defense	Technology	Detect, Respond	AI SecOps (Detection and Protection)
11	Align accountability and expectations	Process	Govern	Risk Management, Metrics and Reporting
12	Build collaboration and collective defense	Process	Govern	Threat Intelligence, Supply Chain Management

The sequencing matters. Areas one through eight strengthen the foundation and shut down the most common entry points and lateral paths. Areas nine and ten shift the program away from depending on response speed and toward exploit prevention and AI-assisted defense. Areas eleven and twelve align accountability inside the entity and rebalance the attacker advantage through shared defense. Trying to adopt the later areas before the foundation is solid produces dashboards that look modern over a posture that is not.

4. The Implementation Program

Theme 1: Strengthen the Foundation (Areas 1 to 8)

These eight areas are the operational work of running a program that still functions when the time from disclosure to weaponization is measured in minutes. None of them is new. All of them are now urgent.

Area 1. Aggressively remediate known risk

What this means. Your vulnerability backlog is no longer compliance debt. It is a map an adversary can follow. Every item with an available patch is a decision you can make now, with a known fix and a familiar process. What you need is prioritization and speed, not further investigation.

Cal-Secure alignment. Technology component, Accelerate Vulnerability and Patch Management. CSF 2.0 Identify and Protect. Capability Map: Vulnerability Management.

Operational actions

- Patch internet-facing systems first, then internal. External assets are the fastest for an adversary to enumerate and weaponize.
- Close out long-standing exceptions where a patch already exists. For any exception that remains, require a named executive owner and a dated review trigger.
- Treat the backlog as operational risk and report burn-down on the same cadence as service reliability metrics.
- Stop assuming compensating controls buy time. Under a one-day weaponization window, most no longer do.
- Freeze backlog growth. Any new vulnerability with an available patch is remediated inside its SLA or escalated to a sponsored exception.

State shared services you can draw on. Where in-house capacity is limited, consume the managed vulnerability and patch management service offered through CDT and Cal-CSIC rather than building the function from scratch. Smaller entities should treat this as the default path.

Maturity signal. The count of externally exposed, patchable vulnerabilities trends toward zero, with burn-down reported to entity leadership at a regular cadence.

Area 2. Harden the perimeter

What this means. The perimeter was correctly demoted in favor of identity-centric and zero-trust designs, but that was never permission to let it decay. Hardening the perimeter is not about deep prevention. It is about buying time for detection. Every hour added to an attacker's initial-access phase is an hour your detection and containment can use.

Cal-Secure alignment. Technology component. CSF 2.0 Protect. Capability Map: Network Threat Protection, Infrastructure Threat Protection.

Operational actions

- Add distance and friction. Use content delivery, managed hosting, and cloud edge controls to absorb reconnaissance before it reaches your core systems.
- Expand web application firewall coverage and modernize perimeter defenses to include bot management, API gateway protection, and runtime application protection where available.
- Apply a controlled delay before adopting new open-source components or AI models. A thirty-day observation period for new dependencies is an inexpensive guard against supply-chain compromise.

- Deploy internal tripwires: canary accounts, canary files, canary tokens, and honey services. These produce high-fidelity alerts on intrusions that bypass the perimeter.
- Close or proxy every unused management interface, admin console, and debug endpoint reachable from the internet.

Where to start if this has been neglected. Run an external attack surface scan. You almost certainly have exposed assets you do not know about, including forgotten cloud accounts and subsidiary infrastructure. Audit firewall coverage against your inventory of internet-facing applications, and pilot deception on a single high-value segment before scaling.

Maturity signal. A current external attack surface inventory with no unmanaged internet-facing management interfaces, and deception in place on at least one critical segment.

Area 3. Segment to contain a breach

What this means. Once an attacker gains a foothold, a flat internal network lets a single compromised system reach almost anything. Segmentation applies least privilege to network traffic, so systems talk only to what they need. It turns a potential enterprise-wide breach into a contained, single-system incident, and it buys defenders time and visibility.

Cal-Secure alignment. Technology component. CSF 2.0 Protect. Capability Map: Network Threat Protection.

Operational actions

- Set a default-deny policy between network zones. Allow only explicitly justified ports and protocols.
- Deploy micro-segmentation across critical zones with least-privilege controls between segments.
- Isolate the systems an attacker most wants: identity, backup, virtualization management, and secrets platforms. Give them stricter trust boundaries and dedicated administrative paths.
- Block legacy internal protocols that are rarely needed and heavily abused, such as SMBv1, Telnet, and older RDP versions.
- Test lateral-movement containment through red-team and purple-team exercises, not only annual compliance checks.

Where to start if this has been neglected. Map current internal traffic flows before restricting anything; you cannot segment what you cannot see. Start with identity, backup, and virtualization. Run in monitor mode first, log what would be blocked, then move to enforcement. Document every exception with an owner and an expiration date.

Maturity signal. Crown-jewel systems sit behind enforced segmentation, and a recent exercise has confirmed that lateral movement from a compromised endpoint is contained.

Area 4. Realign prioritization and compress SLAs

What this means. Scoring and patch windows built for an era of slow exploit development point attention slowly and at the wrong problems. Assuming exploitation in every prioritization decision is the cheapest mindset change available. It needs no new tooling, only a re-weighting of signals you already have and a compression of timelines.

Cal-Secure alignment. Process component, Standardize and Automate Cybersecurity Metrics and Reporting. CSF 2.0 Govern and Identify. Capability Map: Vulnerability Management, Metrics and Reporting.

Operational actions

- Default to assumed active or imminent exploitation for every vulnerability, especially on externally facing assets.
- Move beyond severity score alone. Weight exposure, reachability, business impact, and current adversary intelligence.
- Compress remediation SLAs to match the threat. A workable starting set: critical external in hours, critical internal in days, high within a week.
- Join vendor early-access and private disclosure programs for the platforms you depend on broadly, so you receive fixes ahead of public release.
- Invest in patch orchestration and automated testing so that faster is not also riskier.

Update statewide-aligned risk models. Cal-Secure 2.0 calls for updating risk models, KPIs, and reporting cadences to reflect AI-accelerated exploit timelines. Audit your current SLAs against recent time-to-exploit data, publish the new tiers with their rationale, and roll out in phases starting with internet-facing assets.

Maturity signal. Published, tiered SLAs that an attacker's pace would not embarrass, with prioritization that accounts for exposure and reachability rather than score alone.

Area 5. Validate critical asset inventory and third parties

What this means. You cannot protect what you cannot see. In most entities the authoritative inventory is stale the day it is published, and the inventory of externally exposed assets, including subsidiary infrastructure and third-party systems in the critical path, is partial. An AI-driven adversary often maps your external surface with higher fidelity than your own records hold.

Cal-Secure alignment. Process and Technology components. CSF 2.0 Identify. Capability Map: Asset Management, Supply Chain Management.

Operational actions

- Maintain a real-time asset inventory with dependencies and connections. A spreadsheet is not a real-time inventory.
- Enable same-day decisioning. A new critical vulnerability should yield an answer about affected assets in hours, not days.
- Know every internet-facing exposure, including third-party systems that form your attack surface from the outside.
- Extend visibility to suppliers and service providers in your critical path. At a minimum confirm their patch cadence, posture, and incident-notification obligations.
- Tie the inventory to your configuration database, endpoint tooling, identity provider, and cloud control planes so the view stays continuous.

Where to start if this has been neglected. Run an external attack surface scan; the gap between its findings and your internal records is your reconnaissance exposure. Deploy

continuous discovery across internal and cloud estates, and commission a supply-chain review of your top twenty suppliers, starting with those that have network or identity reach.

Maturity signal. A continuously updated inventory tied to live control planes, and a documented view of critical-path third parties with their notification obligations on file.

Area 6. Replace end-of-life technology

What this means. Unsupported systems were tolerated when migration was costly and exploitation was rare. That calculation has changed. AI-assisted tools quickly identify which versions you run and cross-reference known vulnerabilities, so outdated systems are effectively pre-labeled targets. Currency is now a leading indicator of risk.

Cal-Secure alignment. Technology component. CSF 2.0 Identify and Protect. Capability Map: Asset Management.

Operational actions

- Update software and hardware to supported versions as a continuous program, not a periodic project.
- Replace unsupported technology with maintained alternatives. Where replacement is not yet possible, isolate aggressively.
- Set a minimum freshness standard, such as no more than two versions behind for internal and third-party software, and measure drift continuously.
- Report currency to technology leadership on the same cadence as availability and performance.
- Price end-of-life risk into procurement and architecture decisions, not only unit cost.

Where to start if this has been neglected. Build an inventory of every asset with a documented end-of-support date, color-coded by time remaining. Publish a replacement roadmap with named owners, dates, and budget. For assets that cannot be replaced in the risk window, place them behind compensating controls and flag them for executive risk acceptance.

Maturity signal. A published end-of-life roadmap with owners and dates, and a measured, falling share of systems beyond the freshness standard.

Area 7. Improve logging to enable AI-assisted defense

What this means. AI-assisted defense is only as good as the data it can read. Inconsistent endpoint coverage, identity logs that are enabled but not retained, control-plane events in one place and application logs in another, and opaque SaaS environments were inefficient under human analysis. At machine speed they are a hard limit on what defense can do. This is the investment that unlocks the other investments.

Cal-Secure alignment. Technology component, supporting Enhance Security Operations with AI and Automation. CSF 2.0 Detect. Capability Map: Log Management, Continuous Monitoring.

Operational actions

- Expand logging across endpoints, identities, networks, and cloud workloads. Close the blind spots before tuning existing alerts.

- Standardize formats and retention so telemetry is machine-consumable. Common schemas make downstream work cheaper.
- Centralize telemetry into a platform built for machine-speed analysis, with the compute to sustain AI-assisted work.
- Close visibility gaps in third-party, SaaS, and legacy systems. Require audit logs in contracts and confirm you can actually ingest them.
- Protect the logs themselves as a critical asset. An attacker who can disable or alter logs can evade behavioral detection.

Where to start if this has been neglected. Standardize retention first, since inconsistent retention is often worse than short retention. Audit SaaS logging; many entities pay for audit logs they never turned on. Wire logs into detection engineering so new detections learn from real telemetry.

Maturity signal. Standardized, centralized telemetry across the major domains, ingestible at machine speed, with log integrity protected.

Area 8. Improve cyber resilience

What this means. Prevention and detection will fail more often and faster under this threat model. Resilience determines how much damage an attack actually causes. Many entities have a backup strategy. Fewer have verified restore. Fewer still have tested restoration under conditions where the backups themselves are assumed compromised. Cal-Secure 2.0 asks entities to treat cyber recovery as distinct from traditional disaster recovery.

Cal-Secure alignment. Technology component, Strengthen Critical Infrastructure, OT, and Cyber Recovery. CSF 2.0 Recover. Capability Map: Business Continuity Integration, Incident and Disaster Recovery.

Operational actions

- Maintain tested, isolated backups and immutable recovery snapshots for identity, configuration, and data.
- Rehearse incident response and recovery through regular tabletop and live drills, including drills where the backup systems are assumed compromised.
- Identify critical business processes and pre-plan alternate operating modes, including manual and degraded modes.
- Invest in redundancy and failover for systems the entity cannot operate without.
- Test break-glass procedures. An unused break-glass procedure is an unreliable one.

Where to start if this has been neglected. Run an end-to-end restore test of a representative service and measure time to operational recovery, not only time to data recovery. Schedule quarterly tabletop exercises with rotating scenarios that include the executive team. Confirm backups are isolated from the production identity and network planes; backups reachable from a compromised endpoint are not backups.

Maturity signal. A recent, successful restore test under adversarial assumptions, with a measured recovery time the business has accepted.

Theme 2: Shift the Mindset (Areas 9 and 10)

The next two areas reduce the program's dependence on response speed, which is exactly what AI-accelerated attacks now exploit. They are architectural and operational changes, not only tooling changes.

Area 9. Move from vulnerability management to exploit prevention

What this means. A program built on detect-and-remediate assumes response can outpace attack. Against AI-assisted attacks that is a losing bet, and the gap widens with each capability release. Shifting emphasis to prevention and pre-planned containment means that when an attack succeeds, its impact is contained before the response team has even assembled, because the containment actions were defined, approved, and automated in advance.

Cal-Secure alignment. Technology and Process components. CSF 2.0 Protect and Respond. Capability Map: Network Threat Protection, Incident Management and Analysis.

Operational actions

- Treat segmentation, access controls, and system isolation as active containment surfaces, not only as hygiene.
- Block exploits in progress with web application firewalls, intrusion prevention, and runtime application protection. Favor controls that intervene over those that only observe.
- Update incident playbooks for rapid containment, and state plainly that some service disruption is an accepted trade-off against breach cost.
- Pre-approve containment actions through change management so they can trigger automatically during an incident without a bottleneck.
- Measure the entity on contained incidents, not only detected ones.

A leadership conversation, not a technical one. Establish an acceptable-disruption policy with the business before an incident, not during one. Then run a red-team exercise that uses time-to-contain as the primary metric, with time-to-detect as a supporting one.

Maturity signal. Pre-authorized, automated containment actions exist for the most likely incident types, and a recent exercise reports a credible time-to-contain.

Area 10. Use AI for defense

What this means. Each generation of offensive model finds more, faster, and more creatively than the last. An entity that relies entirely on human-speed analysis is operating a capability generation behind the adversary. That choice is avoidable. AI-assisted triage compresses alert backlogs, AI-assisted code review catches classes of flaws manual review misses, and AI-assisted red teaming illuminates attack paths human teams do not have time to explore. Cal-Secure 2.0 calls this out directly under Enhance Security Operations with AI and Automation.

Cal-Secure alignment. Technology component. CSF 2.0 Detect and Respond. Capability Map: AI SecOps (Detection and Protection).

Operational actions

- Triage, monitor, and respond to alerts at machine speed using AI-assisted operations tooling, coordinated with the statewide detection and runtime program.

- Turn AI capability inward on your own code and dependencies. Begin by asking an assisted review of any code change, then build toward automated review inside the development pipeline before code merges.
- Equip defenders to use AI in detection engineering, red teaming, and testing. Provide licenses and training, not only policy.
- Trigger automated containment, such as isolating systems, blocking traffic, or revoking credentials, on high-fidelity signals.
- Govern AI tool use with defined guardrails, human oversight on high-impact actions, and explicit accountability. AI for defense is not AI instead of defense.

Defend the agents you deploy. AI agents introduced into security and engineering work are privileged and are not covered by older controls. Before deploying them in or near production, define scope boundaries, blast-radius limits, escalation logic, and human override. Treat the agent harness, including its prompts, tool definitions, and retrieval pipelines, with the same rigor as its permissions. This is also an emerging area of Cal-Secure 2.0 governance, covering runtime security and agentic containment.

Where to start if this has been neglected. Inventory AI security tooling already in use, including anything adopted informally. Resource the three highest-leverage uses first, typically alert triage, threat hunting, and code review. Draft a governance policy covering tool selection, data handling, guardrails, and review expectations.

Maturity signal. At least one AI-assisted defensive use case in production under documented guardrails, with assisted security review integrated into the development pipeline.

Theme 3: Align People and Partners (Areas 11 and 12)

Faster threats require clearer ownership inside the entity and stronger collaboration between entities. Technical controls without aligned accountability decay. Aligned accountability without collective action hits a visibility ceiling.

Area 11. Align accountability and expectations

What this means. In many entities, security outcomes are measured in the security function and funded from the IT function, while the business units that own the systems and data sit outside the accountability model. Under that arrangement, remediation speed is always someone else's priority. AI-speed attacks do not wait for clearer ownership to emerge. Building accountability into performance objectives is what turns policy into outcomes.

Cal-Secure alignment. Process component, Formalize and Evolve Multi-Tiered Cybersecurity Governance. CSF 2.0 Govern. Capability Map: Risk Management, Metrics and Reporting.

Operational actions

- Build security metrics into team objectives across engineering, infrastructure, and the business units that own applications.
- Measure patch velocity and platform currency alongside system performance and reliability metrics.

- Treat remediation speed as a reliability metric, and report it to leadership on the same cadence as uptime and customer-facing service levels.
- Name an accountable executive for every crown-jewel system. Shared accountability is diluted accountability.
- Reset leadership expectations across business, technology, and resilience functions. The new threat model needs a new operating contract.

Where to start. At the next cycle, refresh team objectives to include security metrics rather than waiting for an annual review. Redesign leadership reporting to include a security-operations view that treats remediation speed as a reliability metric, and establish quarterly accountability reviews for the named owners of crown-jewel systems. This aligns with the CalOES Cyber Task Force role in exercising statewide escalation and decision rights.

Maturity signal. Every crown-jewel system has a named executive owner, and remediation speed appears in routine leadership reporting.

Area 12. Build collaboration and collective defense

What this means. No single entity sees the full range of emerging threats or can respond alone. Adversaries already share tools, techniques, and targeting intelligence freely. Defenders too often do not, held back by caution, competition, and legal uncertainty. Under AI-speed attacks that asymmetry is no longer sustainable. A sector that responds as a unit is much harder to target. This is the heart of the statewide model: no single entity defends California alone.

Cal-Secure alignment. Process component, Strengthen Public-Private and Cross-Government Collaboration, and Enhance Statewide Threat Intelligence Fusion with Cal-CSIC. CSF 2.0 Govern. Capability Map: Threat Intelligence, Supply Chain Management.

Operational actions

- Share threat intelligence through the relevant sector information-sharing channels and through Cal-CSIC. Passive membership is not sharing; nominate a named contact and a sharing cadence.
- Engage peer entities and supply-chain partners on emerging risks. Build the relationships before the incident.
- Co-develop and test remediations with vendors and peers ahead of public disclosure where programs allow.
- Participate in joint tabletop and purple-team exercises and shared playbooks for incidents that span entities, coordinated through the statewide governance structure.
- Contribute back. Publish patterns, indicators, and lessons learned where it is legally and operationally appropriate.

Draw on statewide capacity. CDT and Cal-CSIC provide shared services that let entities of any size raise their floor: managed detection and response, managed deception, managed runtime security, managed vulnerability and patch management, and access to AI-assisted analyst capability. Entities below the resourcing line should consume these rather than attempt to build equivalents alone.

Maturity signal. An active sharing relationship with Cal-CSIC and at least two peer entities, with a defined communications policy that says what you can share, with whom, and when.

5. Sequencing the Work: 30, 90, and 365 Days

The plan below adapts the statewide adoption phases (Initiate and Organize, Prioritize and Right-Size, Enable and Equip, Embed and Operate) to an aggressive but realistic timetable. Scale it to your entity. A complex environment may move more slowly on some items, and an entity that consumes statewide shared services may move faster than its headcount alone would suggest. Where two actions appear to conflict, for example a controlled adoption delay against the pressure to patch faster, use judgment and write the nuance into policy rather than following either rule blindly.

First 30 days: clear the decks

The purpose of the first month is to stabilize the foundation and make the decisions no one has been willing to make. It is a focused month, not a busy one.

- Burn down the top ten externally exposed vulnerabilities in the known backlog (Area 1).
- Compress critical-severity external SLAs to hours, and publish the new tiers with the rationale (Area 4).
- Run a one-day tabletop for an AI-driven breach scenario, including Legal, communications, and executive leadership (Area 8).
- Begin the inventory audit across internal assets, external attack surface, and third-party critical path (Area 5).
- Pilot AI-assisted alert triage on a single detection family and measure the change in time-to-triage (Area 10).
- Establish the cross-functional governance mechanism that lets defensive technology be onboarded without approval friction (Areas 9 and 11).

First 90 days: reset operations

The second and third months move segmentation, identity hardening, and logging coverage from intention to implementation.

- Implement micro-segmentation around identity, backup, and virtualization management (Area 3).
- Expand logging across endpoints, identities, cloud, and SaaS to a usable detection baseline (Area 7).
- Enforce multifactor authentication and dedicated admin accounts across all privileged access, and disable legacy authentication (Areas 2 and 3).
- Deploy external attack surface management and wire its findings into your prioritization model (Areas 2 and 5).
- Establish the leadership reporting view that treats remediation speed as a reliability metric (Area 11).

- Engage Cal-CSIC and peer entities with a named contact and a sharing cadence (Area 12).

First 365 days: institutionalize the shift

The first year ends with this work baked into how the entity operates, measured with the same rigor as reliability, with AI for defense as a standard capability rather than a pilot.

- Stand up a dedicated vulnerability operations function with named accountability and measured outcomes, modeled on the discipline of a delivery operations team (Areas 1, 9, and 10).
- Replace the riskiest end-of-life technology identified in the inventory, and publish the closure rate (Area 6).
- Complete continuous AI-assisted red-teaming coverage across the crown-jewel estate (Areas 1 and 10).
- Operationalize AI for defense: assisted operations patterns, automated containment on high-fidelity signals, and assisted review inside the development pipeline (Areas 9 and 10).
- Embed security metrics into team objectives and executive performance expectations (Area 11).
- Re-baseline against an updated threat model. Frontier capability will have moved, and the program needs to move with it (all areas).

6. Tracking Progress in the Cal-Secure 2.0 Toolkit

This guide describes the work. The Cal-Secure 2.0 Toolkit is where you record current and target states and track progress, so OIS can see statewide maturity and act on what the data shows. Two instruments matter most here.

Organizational Profile

For each capability in the Capability Map, the Organizational Profile records whether the capability is relevant to your entity and, if so, its current and target state across maturity, policies and procedures, internal practices, roles and responsibilities, informative references, and supporting evidence. Use the risk-based selection described in NIST SP 800-37 to decide which capabilities to include, weighted to your mission, the systems and data you protect, and the urgency created by known gaps. A useful quality check: each current-maturity statement should be backed by real evidence, and each target statement should describe specific policies, practices, and roles, not an aspiration.

Maturity Matrix

The Maturity Matrix maps each capability to maturity levels using the standard progression of Initial, Managed, Defined, Quantitatively Managed, and Optimized. Record where you are and where you intend to be, then close the gap between current and target. The twelve

implementation areas in this guide map cleanly onto the Capability Map entries in the crosswalk in Section 3, so progress on an area moves a specific, trackable capability.

The capabilities most directly advanced by this guide are Vulnerability Management, Network and Infrastructure Threat Protection, Asset Management, Supply Chain Management, Log Management, Continuous Monitoring, Business Continuity Integration, Incident and Disaster Recovery, AI SecOps for Detection and Protection, Risk Management, Metrics and Reporting, and Artificial Intelligence Governance. Review your Organizational Profile at least annually and after any significant change in mission, systems, or threat environment. More detailed requirements will continue to be issued through updates to the State Administrative Manual and the State Information Management Manual.

7. Your Workforce Through the Transition

The volume and cadence of vulnerability disclosures under this threat model will exceed anything most teams have handled. People feel it. Practitioners across every level are working out what AI means for their roles, and that uncertainty is a normal response to a capability shift, not a sign of crisis. The teams that adapt fastest are the ones that treat AI tooling as something to lean into rather than something to fear. Cal-Secure 2.0 builds workforce readiness into the People component for this reason.

Burnout and attrition are a direct operational risk. The expertise needed to navigate this period is scarce, takes years to develop, and cannot be replaced on short notice. Treat team resilience, meaning sustainable workload, support, and retention, with the same seriousness as the technical work. Plan for reprioritization and added capacity before the first large wave of patches lands, not after.

Practical steps for the CIO

- Require, do not merely permit, the use of AI agents across security functions, with guardrails in place. Optional adoption has not been shown to overcome cultural inertia, and adoption is the limiting factor for most other actions in this guide.
- Give staff structured time and training to become hands-on with assisted tooling. The barrier to entry is lower than most assume.
- Plan capacity ahead of demand. Repurpose staff inside security and engineering, and request reserve headcount or contractor capacity for the anticipated rise in triage, remediation, and incidents.
- Use the People-component initiatives in Cal-Secure 2.0, including role-based training, career pathways aligned to the national workforce framework, and communication-skills development for translating technical risk to non-technical leaders.
- Run statewide tabletop, purple-team, and red-team exercises on a regular cadence, including ransomware and AI-orchestrated intrusion scenarios, and consider structured time for staff to prototype assisted defensive capabilities in a supervised setting.

8. Briefing Agency and Executive Leadership

This shift has reached leadership attention, and that creates an opening. The points below help frame a leadership or oversight update. Align them to your entity's actual situation and programs.

The shift, in plain terms

AI is making the organization faster and more competitive, and the business is already pursuing that value. The same capability in an adversary's hands compresses the time between a vulnerability existing and causing disruption from weeks to hours. This is a permanent acceleration, not a temporary spike. Two things follow. First, several assumptions behind current risk metrics no longer hold and need re-examination, with measurement shifting toward containment and speed of recovery. Second, the same capability that creates the risk also creates a defensive opportunity: the entity can now find its own weaknesses before an adversary does, review code at machine speed, and respond faster than a human-only team can.

Why the existing program matters more, not less

The security program already funded is what makes an AI-era defense viable. The investments in place are what keep a single point of entry from becoming full disruption. In an environment where weaknesses are discovered faster, that containment architecture is more valuable. The changes recommended here are designed to return risk toward where it stood before this capability arrived, and to demonstrate due diligence in response to a documented change in the threat environment.

What you are asking for

This is not an open-ended AI initiative. It is a targeted, time-boxed plan with clear owners and outcomes. A typical ask covers five things: capacity to handle the rising volume of triage and incidents while protecting experienced staff; formal adoption of AI tooling across security functions; hardening of inventory, exposure, segmentation, and authentication across internal systems and key third parties; faster procurement and governance so defensive technology can be onboarded in time; and updated technical and communications playbooks built to execute at the required speed, including pre-authorized containment. Provide regular check-ins through the plan period to capture results and surface roadblocks.

This work has a precedent. A systemic threat with a hard deadline has been met before through coordinated, disciplined effort. This is the same kind of problem, with more capable tools now available to defenders. Being ready is not about reacting to one model or one announcement. It is about permanently closing the gap between how fast weaknesses are found and how fast the entity can respond.

Closing Note

These twelve areas do not lower the standard. They set a survivable one for an environment where adversaries move at machine speed. The time available to prepare is real but finite. The

frontier offensive capability is in restricted release today, and California is inside that window. The most useful thing an entity can do with that window is the unglamorous work above: remediate what is known, harden the fundamentals, contain by design, and adopt AI for defense before the capability is broadly available.

This guide is a starting line. Use it to assess your current maturity in the Cal-Secure 2.0 Toolkit, set target states, and chart a path across People, Process, and Technology. CDT, OIS, and Cal-CSIC will keep refining the underlying roadmap and the shared services that support it, because keeping California safe is a shared commitment.