Credence Solutions Group, LLC

Agile | Intelligent | Adaptive

Mythos-Ready Security Checklist for CISOs

Branded PDF edition for Credence Solutions Group, LLC

A practical, board-ready framework for security leaders operating in an era of AI-accelerated vulnerability discovery, exploitation, and remediation.

Prepared for download and client distribution

June 2026

This checklist is designed for a security operating environment where AI systems can rapidly discover, validate, and help weaponize vulnerabilities. Mythos-ready means the organization can identify exposure, prioritize exploitable risk, patch or isolate quickly, detect compromise, and recover without relying on heroic manual effort.

Context: Anthropic describes Project Glasswing as a defensive initiative using Claude Mythos Preview to help secure critical software, and its May 2026 update said partners had found more than 10,000 high- or critical-severity vulnerabilities using the model. Source: Anthropic Project Glasswing.

Program shift: The CSA / SANS / OWASP GenAI Security Project briefing frames this as a security-program shift: CISOs should adjust risk calculations for higher patch volume, shorter remediation windows, more persistent complex attacks, and should double down on segmentation, egress filtering, MFA, dependency management, automated assessments, and AI-enabled defensive operations. Source: CSA Mythos-ready briefing.

Readiness scoring

Green: Implemented, measured, and tested.

Amber: Partially implemented or not consistently measured.

Red: Missing, manual, unowned, or dependent on heroic effort.

A program is Mythos-ready only when all P0 items are green, all P1 items are at least amber with funded remediation, and no critical business service has unresolved internet-exposed high-risk vulnerabilities without compensating controls.

1. Board governance and risk posture

☐ P0 - Board briefing completed. Board and executive team understand that AI-assisted vulnerability discovery compresses the time between discovery, exploit development, and exploitation.

☐ P0 - Risk appetite updated. The company has explicitly accepted that emergency patching, temporary downtime, degraded service, or feature freezes may be necessary to reduce exposure.

☐ P0 - Mythos-era risk register created. Include AI-accelerated exploitation, unpatched OSS dependencies, third-party software exposure, unmanaged internet-facing assets, AI-agent misuse, and remediation capacity limits.

☐ P0 - Named executive owners. CIO, CTO, CISO, CPO, GC, engineering leadership, and business-unit owners have clear decision rights for emergency remediation.

☐ P1 - Security budget reframed. Security funding is tied to business resilience and AI strategy enablement, not just audit compliance.

☐ P1 - Crisis decision model approved. Pre-approved conditions exist for taking systems offline, disabling features, blocking vendors, rotating secrets, or isolating environments.

Evidence: Board deck, approved risk appetite statement, board minutes, funded remediation plan, emergency-change policy, updated cyber risk register.

2. Asset, exposure, and software inventory

☐ P0 - Complete external attack surface inventory. All internet-facing apps, APIs, cloud services, VPNs, identity endpoints, SaaS admin portals, exposed storage, and third-party-hosted services are continuously discovered.

☐ P0 - Critical asset inventory. Crown-jewel systems are tagged by business criticality, data sensitivity, revenue impact, operational dependency, and recovery priority.

☐ P0 - Software bill of materials coverage. SBOMs or equivalent dependency inventories exist for critical first-party applications, commercial software, containers, cloud workloads, and AI/ML systems.

☐ P0 - Ownership mapped. Every critical app, repo, package, cloud account, data store, and runtime environment has an accountable technical owner.

☐ P1 - Runtime context linked to code. Security teams can connect a vulnerability to the running workload, exposed service, data classification, owner, exploitability, and business impact.

☐ P1 - Shadow IT discovery. Unknown SaaS, cloud accounts, rogue domains, unmanaged APIs, abandoned apps, and orphaned workloads are continuously detected.

☐ P1 - End-of-life technology register. Unsupported OS, libraries, appliances, databases, network devices, and application frameworks are tracked with retirement dates.

Evidence: CMDB, EASM reports, cloud asset inventory, SBOM repository, dependency graph, ownership map, criticality tagging, exposed-services report.

3. Vulnerability management rebuilt for AI-speed exploitation

CISA says its Known Exploited Vulnerabilities catalog is the authoritative source for vulnerabilities exploited in the wild and should feed vulnerability-prioritization programs. In a Mythos-class environment, use KEV status, exploitability, exposure, asset criticality, and compensating controls - not CVSS alone.

☐ P0 - Exploitability-based prioritization. Prioritize by active exploitation, KEV, public exploit availability, internet exposure, reachable attack path, privilege required, data sensitivity, and business criticality.

☐ P0 - Emergency patch SLAs approved. Suggested starting targets are actively exploited or KEV on internet-facing critical assets within 24 hours; critical internet-facing vulnerabilities within 48 to 72 hours; critical internal vulnerabilities with lateral-movement potential within 7 days; and high-risk vulnerabilities on critical systems within 14 days.

☐ P0 - Exception governance. Every missed SLA requires a named risk owner, expiry date, compensating controls, and CISO-visible escalation.

☐ P0 - Continuous patch pipeline. Patching is treated as an always-on production process with testing, canary deployment, rollback, validation, and post-deploy verification.

☐ P0 - Emergency change path. Security patches are not blocked by normal change-board cadence when exploitation risk exceeds agreed thresholds.

☐ P1 - Patch-to-production measured. Track time from vendor advisory to production remediation, not just scan finding to ticket closure.

☐ P1 - Virtual patching and compensating controls. WAF rules, EDR policies, network segmentation, service isolation, feature disabling, and egress restrictions are pre-planned when patching is delayed.

☐ P1 - Vulnerability deduplication. Security tooling normalizes findings across SAST, DAST, SCA, container, cloud, IaC, EDR, and pen-test sources.

☐ P1 - Remediation capacity modeled. Engineering teams have capacity allocation for urgent security work, including explicit tradeoffs against roadmap delivery.

Evidence: SLA dashboard, KEV correlation, patch-latency metrics, risk exceptions, emergency-change records, compensating-control playbooks, remediation burn-down.

4. Secure SDLC and product security

NIST's SSDF defines secure software development practices for preparing the organization, protecting software, producing well-secured software, and responding to vulnerabilities. NIST also notes SSDF can support procurement and supplier communication.

☐ P0 - Security gates in CI/CD. Critical SAST, SCA, secrets, IaC, container, and malware checks block release unless formally risk-accepted.

☐ P0 - Secrets never enter code. Pre-commit scanning, CI scanning, repo history scanning, and automated secret revocation are active.

☐ P0 - Secure design reviews for critical systems. Threat modeling is mandatory for internet-facing systems, authentication flows, payment flows, admin functions, AI features, and high-sensitivity data paths.

☐ P0 - Dependency update automation. Critical library, container base image, package-manager, and framework updates can be tested and merged quickly.

☐ P1 - AI-assisted security testing. Defensive AI tools are used to review code, generate tests, identify exploit paths, triage findings, and propose patches under human review.

☐ P1 - Fuzzing for exposed parsers and protocols. Fuzz high-risk attack surfaces such as file parsers, APIs, auth flows, network protocols, deserialization, media processing, and cryptographic integrations.

☐ P1 - Abuse-case testing. Test authorization bypass, business logic abuse, tenant isolation failure, payment fraud, prompt injection, SSRF, data exfiltration, and privilege escalation.

☐ P1 - Secure release evidence. Each release has traceable evidence for security checks, dependencies, build provenance, known exceptions, and approval.

Evidence: SDLC policy, CI/CD gate logs, threat models, code-scanning dashboards, dependency update metrics, release attestations, secure design review records.

5. Software supply chain and third-party risk

☐ P0 - Critical supplier inventory. Identify vendors whose compromise could affect authentication, endpoint management, cloud control planes, CI/CD, source code, customer data, financial systems, or production availability.

☐ P0 - Supplier patch commitments. Critical suppliers contractually commit to vulnerability disclosure, emergency patch timelines, SBOM/VEX where appropriate, breach notification, and security contact escalation.

☐ P0 - Build pipeline hardened. CI/CD systems use MFA, least privilege, protected branches, signed commits or equivalent controls, isolated runners, dependency pinning, and restricted release credentials.

☐ P0 - Artifact integrity. Production artifacts are signed, provenance is captured, and deployments verify artifact integrity.

☐ P1 - Open-source intake process. New OSS components require license, maintenance, vulnerability, popularity, maintainer-risk, and transitive-dependency review.

☐ P1 - Dependency drift monitored. Track outdated packages, abandoned libraries, vulnerable transitive dependencies, typosquatting risk, and unpinned versions.

☐ P1 - Vendor concentration risk reviewed. Identify single points of failure across identity, cloud, endpoint, network, observability, and incident-response providers.

Evidence: Supplier criticality register, security addenda, SBOMs, VEX records, signed build logs, provenance attestations, OSS approval workflow, vendor risk reviews.

6. Identity, privilege, and secrets

CSA's Mythos-ready briefing explicitly calls out MFA and defense-in-depth as near-term controls. Anthropic's Project Glasswing materials emphasize that defenders must harden before Mythos-class capabilities proliferate.

☐ P0 - Phishing-resistant MFA for privileged access. Require FIDO2/passkeys or equivalent phishing-resistant MFA for admins, developers, production access, cloud consoles, CI/CD, source control, VPN, and security tooling.

☐ P0 - Privileged access management. Admin access is just-in-time, time-bound, approved, logged, and reviewed.

☐ P0 - Service-account inventory. All service accounts, API keys, OAuth apps, tokens, certificates, and machine identities have owners, scopes, rotation dates, and usage monitoring.

☐ P0 - Secrets rotation playbook. The organization can rapidly rotate credentials after code exposure, vendor compromise, endpoint compromise, or suspicious access.

☐ P1 - Least privilege enforced. Standing privileges are minimized across cloud IAM, SaaS admin roles, Kubernetes, databases, production consoles, and developer tooling.

☐ P1 - Identity threat detection. Detect impossible travel, token replay, suspicious consent grants, MFA fatigue, privilege escalation, anomalous API use, and dormant-account activation.

☐ P1 - Break-glass accounts secured. Emergency accounts are monitored, vaulted, tested, and excluded from normal federation dependencies only where necessary.

Evidence: MFA coverage report, PAM logs, IAM review records, service-account inventory, token rotation metrics, privileged-access exceptions, identity detection rules.

7. Network, cloud, and zero-trust containment

Segmentation, egress filtering, and defense-in-depth are explicitly recommended in the CSA Mythos-ready briefing. Network-focused guidance also emphasizes environmental separation, blocking lateral-movement paths, ringfencing critical applications, and rapid isolation.

☐ P0 - Production, development, and test separated. No broad network trust, shared credentials, shared databases, or unmanaged administrative paths across environments.

☐ P0 - Critical asset ringfencing. Crown-jewel applications, identity systems, payment systems, backups, source code, CI/CD, and security tooling have restrictive allowlist access.

☐ P0 - Egress filtering. Workloads cannot freely communicate outbound to the internet; high-risk destinations, protocols, and data-transfer paths are controlled.

☐ P0 - Lateral movement controls. High-risk ports and unnecessary east-west traffic are blocked or tightly controlled.

☐ P1 - Cloud control-plane hardening. Cloud org policies, SCPs, guardrails, logging, encryption, public exposure controls, and admin boundaries are enforced.

☐ P1 - Kubernetes and container isolation. Cluster admin access, pod privileges, network policies, image provenance, secrets, and admission controls are locked down.

☐ P1 - Rapid isolation capability. Security operations can quarantine a host, workload, subnet, account, application, or identity within minutes.

☐ P1 - Attack-path management. Continuously identify paths from low-value footholds to high-value assets.

Evidence: Segmentation diagrams, firewall and policy rules, cloud guardrail reports, attack-path graphs, egress policy, isolation runbooks, network-flow baselines.

8. Detection, response, and security operations

☐ P0 - Logging coverage for critical systems. Identity, endpoint, cloud, SaaS, network, DNS, email, application, database, CI/CD, source control, and security tools feed a central detection platform.

☐ P0 - Detection mapped to current attack paths. Detections cover exploitation of internet-facing systems, credential theft, token abuse, privilege escalation, lateral movement, data staging, exfiltration, destructive actions, and supply-chain compromise.

☐ P0 - 24/7 triage path. Alerts on critical assets, exploited vulnerabilities, identity compromise, and production anomalies have an always-available response path.

☐ P0 - Incident playbooks updated. Include AI-accelerated exploitation, mass vulnerability disclosure, third-party zero-day, code-signing compromise, CI/CD compromise, credential leakage, and SaaS breach.

☐ P1 - AI-assisted SOC workflows. Use defensive AI for alert enrichment, timeline construction, detection engineering, malware triage, log summarization, and response drafting - with human approval for high-impact actions.

☐ P1 - Threat hunting cadence. Hunt for exploitation of newly disclosed vulnerabilities, suspicious scanning, anomalous service behavior, identity abuse, and low-noise persistence.

☐ P1 - Detection engineering SLA. New high-risk vulnerabilities and attack techniques trigger detection updates within defined timeframes.

☐ P1 - Purple-team validation. Tabletop and technical exercises validate whether controls detect and contain realistic exploit chains.

Evidence: SIEM/EDR coverage map, detection library, MITRE ATT&CK; mapping, playbooks, on-call rosters, incident tickets, tabletop reports, purple-team findings.

9. Resilience, recovery, and business continuity

☐ P0 - Immutable backups. Critical systems and data have immutable, offline, or logically isolated backups.

☐ P0 - Recovery tested. Recovery is tested for identity systems, core business apps, customer data stores, production cloud environments, endpoint management, and security tools.

☐ P0 - RTO/RPO validated. Actual recovery times and data-loss tolerances are measured, not assumed.

☐ P0 - Clean-room recovery plan. The organization can rebuild core infrastructure from trusted sources if AD, cloud IAM, CI/CD, or endpoint tooling is compromised.

☐ P1 - Crisis communications ready. Legal, privacy, regulatory, customer, employee, insurer, law-enforcement, and board communications are templated and assigned.

☐ P1 - Business service degradation plan. Business owners know which services can be disabled, rate-limited, isolated, or run manually during emergency remediation.

☐ P1 - Supplier outage scenarios. Plans account for compromise or outage of critical SaaS, cloud, identity, EDR, logging, and communications providers.

Evidence: Backup reports, restore-test records, BCP/DR plans, RTO/RPO test results, crisis comms templates, clean-room rebuild procedures, supplier contingency plans.

10. AI-agent and GenAI security governance

☐ P0 - Approved AI tooling inventory. All coding agents, SOC copilots, LLM APIs, browser agents, MCP servers, plugins, and SaaS AI features are inventoried and risk-rated.

☐ P0 - Agent permissions scoped. AI agents have least-privilege access to repos, tickets, terminals, cloud consoles, documents, email, chat, and production systems.

☐ P0 - No unlogged autonomous production actions. Agents cannot modify production, rotate credentials, delete data, approve releases, alter security controls, or contact customers without policy-based approval.

☐ P0 - Sensitive data controls. Prompts, tool outputs, files, logs, and embeddings are classified, retained, and protected according to data policy.

☐ P1 - Prompt-injection defenses. Agents that read emails, tickets, documents, web pages, code, or third-party content are treated as exposed to hostile instructions.

☐ P1 - Agent supply-chain review. MCP servers, plugins, tools, packages, prompts, skills, and model integrations go through vendor and security review.

☐ P1 - Human-in-the-loop thresholds. Define which AI-suggested actions require security, engineering, legal, or business approval.

☐ P1 - AI activity monitoring. Monitor agent actions, data access, tool calls, code changes, failed attempts, privilege use, and anomalous behavior.

Evidence: AI tool register, agent access reviews, model/data-flow diagrams, prompt-injection test results, AI audit logs, approval policies, acceptable-use policy.

11. Metrics CISOs should report monthly

Use these as board and operating metrics.

Metric	Target posture
Internet-facing asset inventory coverage	95-100% known and owned
Critical asset ownership	100% owner assigned
KEV / actively exploited vulnerability SLA	95%+ within SLA
Critical internet-facing patch latency	Median under 72 hours
Emergency patch success rate	Tracked with rollback rate

Metric	Target posture
Vulnerability exception age	No open-ended exceptions
SBOM / dependency inventory coverage	90%+ for critical apps
MFA coverage for privileged users	100% phishing-resistant where feasible
Service-account ownership	100% owner and rotation policy
Segmentation coverage for crown jewels	100% enforced or exception-approved
Logging coverage for critical systems	95-100% feeding detection stack
Tested recovery of critical services	At least quarterly for top-tier systems
AI-agent inventory coverage	100% approved or blocked
Supplier emergency patch commitments	100% for critical suppliers

12. 30 / 60 / 90-day execution plan

First 30 days - reduce existential exposure

☐ Brief board and executive team on Mythos-class risk.

☐ Create a Mythos-era risk register and emergency remediation governance.

☐ Freeze or restrict unmanaged internet-facing assets.

☐ Identify top 25 crown-jewel systems and assign business/technical owners.

☐ Correlate internet exposure with KEV, critical vulnerabilities, identity paths, and sensitive data.

☐ Enforce phishing-resistant MFA for privileged access.

☐ Establish emergency patch SLAs and exception process.

☐ Build rapid isolation playbooks for critical workloads and compromised identities.

☐ Inventory AI agents, coding assistants, MCP servers, and AI-enabled SaaS tools.

Days 31-60 - increase remediation velocity

☐ Automate dependency, container, and OS patch workflows for critical systems.

☐ Add exploitability-based prioritization to vulnerability management.

☐ Implement release-blocking checks for critical code, dependency, secret, IaC, and container findings.

☐ Segment production from development and testing.

☐ Ringfence identity, CI/CD, source code, backups, payment systems, and customer-data stores.

☐ Harden cloud control planes and remove excessive standing privileges.

☐ Update incident playbooks for AI-accelerated exploitation and mass zero-day disclosure.

☐ Run a tabletop exercise on a critical third-party zero-day with no immediate patch.

Days 61-90 - operationalize Mythos readiness

☐ Deploy continuous attack-surface management and ownership reconciliation.

☐ Expand SBOM/dependency inventory to all critical applications.

☐ Implement AI-assisted defensive code review and vulnerability triage under human supervision.

☐ Conduct purple-team validation against top attack paths.

☐ Test recovery of at least one tier-1 business service from clean backups.

☐ Measure and report patch-to-production latency.

☐ Establish supplier emergency vulnerability-response commitments.

☐ Present board-level Mythos-readiness dashboard with red/amber/green status, funding gaps, and accepted risks.

CISO definition of done

You are Mythos-ready when your organization can answer yes to these five questions:

1. Do we know what we run? Assets, software, dependencies, owners, exposure, and criticality are continuously known.

2. Can we fix faster than attackers can exploit? Emergency patching, compensating controls, and isolation are operationally proven.

3. Can we contain failure? Segmentation, least privilege, egress control, and recovery prevent one exploit from becoming enterprise compromise.

4. Can we detect and respond at machine speed? Telemetry, detection engineering, AI-assisted SOC workflows, and 24/7 response paths are in place.

5. Can we prove it to the board? Metrics, evidence, risk exceptions, and funding gaps are visible in business terms.

Sources

Anthropic - Project Glasswing - https://www.anthropic.com/project/glasswing

Cloud Security Alliance - AI Vulnerability Storm: Mythos-Ready Security Program - https://labs.cloudsecurityalliance.org/research/ai-vulnerability-storm-mythos-ready-security-program/

CISA - Known Exploited Vulnerabilities Catalog - https://www.cisa.gov/known-exploited-vulnerabilities-catalog

NIST - Secure Software Development Framework (SSDF), SP 800-218 - https://csrc.nist.gov/pubs/sp/800/218/final