MA-S2

Mission Assurance Security Standard for Software

v1.0

May 2026

“

Cyber-tuned LLMs are increasingly capable of finding novel vulnerabilities in complex environments, and our adversaries are increasingly capable of wielding them. This has increased the surface area of attacks for mission-critical institutions, who are facing a new bottleneck of vulnerability triage, remediation, and orchestration. Palantir has taken on the responsibility of proposing a new standard for the security community as the bar that we believe any serious institution should hold.

Shyam Sankar — **Chief Technology Officer, Palantir Technologies**

0.0 Abstract

Palantir Technologies operates software and AI operations across the most sensitive environments in the world: classified government networks, active combat theaters, air-gapped intelligence systems, and critical commercial infrastructure. Over 20+ years, we have designed our deployment infrastructure and security posture to support these environments on an ongoing basis, with readiness for all new developments and technological advancements, such as the most recent advent of cyber-tuned LLMs.

This document proposes a candidate Mission Assurance Security Standard (MA-S2) framework for software, one to which we are actively working towards. Beyond meeting established frameworks (e.g, SOC 2, FedRAMP, IL5 / IL6, NIST SP 800-53, ISO 27001, etc.), the proliferation of AI-assisted vulnerability discovery has underscored the need for a heightened standard across the software landscape. MA-S2 is a proposed framework that outlines the security controls, identification and remediation capabilities, orchestration capabilities, and other attestation requirements for software.

We are publishing this proposed standard publicly and for comment because we view it as our responsibility to any enterprise that wishes to secure its own digital footprint and ecosystem, and we view it as our contribution to the broader security community. For any CISO, CFO, procurement officer, and security architect that believes its mission is as critical as those that are responsible for the United States' national security, and wish to protect their environments to this same standard, we encourage them to urgently prioritize meeting these standards.

For the broader security community, we actively invite reviews and further contributions to ensure a collectively high and growing standard.

1.0 Background

1.1 From Vulnerability Identification to Automated Remediation & Orchestration

The Office of the CISO and the cybersecurity industry has operated for decades under a stable assumption: that finding vulnerabilities in complex software systems was the single most difficult and important challenge. Individual software solutions with large teams and programs (e.g., penetration testing cycles, bug bounties, etc.) have been faced with addressing the expensive, slow, and challenging process of vulnerability discovery.

The discovery of zero-day vulnerabilities in major operating systems and browsers, which previously required elite research teams, has become an automated, high-volume capability available to any sufficiently resourced adversary. The actors who are already exploiting this shift at a national level are already deploying these capabilities.

AI models are now capable of rapidly surfacing novel vulnerabilities in complex, multi-stage cyber attack paths. Organizations can identify more vulnerabilities and at a faster pace, but lack domain awareness and automated remediation infrastructure. It is no longer acceptable to be unable to answer the most basic questions of what software is running, where, and on what version. Industry must be able to respond at velocity and scale.

1.2 Critical Standard for Software & AI Security

The Mission Assurance Security Standard (MA-S2) is organized around four key controls that represent the minimum viable posture for any software vendor operating in a modern, secure environment:

[0] Continuous, AI-Augmented Vulnerability Scanning (CVS)

[1] Attack Path Modeling and AI-Assisted Adversarial Simulation (APM)

[2] Real-Time Software Inventory and Domain Awareness (INV)

[3] Autonomous Remediation Orchestration (ARO)

2.0 Applicability and Attestation

2.1 Applicability

This standard should apply to mission-critical software regardless of deployment model, including public cloud, private cloud, on-premises, air-gapped, and other environments.

2.2 Complementary Frameworks

MA-S2 is designed to complement existing attestation frameworks. Vendors holding SOC 2 Type 2, FedRAMP Moderate or High, or DISA IL5 / IL6 authorization are expected to have foundational controls in place. MA-S2 addresses the next frontier of security requirements, brought on by the proliferation and speed of AI-native attack vectors.

2.3 Purpose and Use

MA-S2 is intended for use by third parties, including customers, procurement organizations, regulators, and independent assessors, to evaluate the security posture of software vendors operating in their environments. A vendor's adherence to MA-S2 is established through an independent assessment of evidence against each control domain, producing an MA-S2 attestation report that the vendor may share with customers and prospects under appropriate confidentiality terms.

Organizations may also use MA-S2 internally as a self-assessment framework to evaluate their own software posture and identify gaps.

2.4 Attestation and Evidence

Vendors must produce attestation evidence for each control domain in this standard. Evidence may take the form of: third-party audit reports, platform-generated telemetry, architecture documentation reviewed by a qualified independent assessor, or contractual SLA commitments with defined measurement methodologies. Self-attestation without supporting evidence is not acceptable.

3.0 Control Domain 0: Continuous, AI-Augmented Vulnerability Scanning (CVS)

3.1 Control Objective

Vendors must demonstrate the capability to continuously scan all software components they ship or operate for known and novel vulnerabilities, regardless of packaging format or deployment model, using tooling that integrates AI-assisted analysis alongside deterministic scanning methods. Point-in-time scanning is insufficient.

Control Domain 0 requires vendors to maintain a continuous, living record of their vulnerability landscape across the entirety of the software footprint.

3.2 Required Controls

CVS–0.1 Automated Container and Dependency Scanning

Every software release artifact must undergo automated vulnerability scanning prior to deployment, regardless of packaging format (e.g., container images, VM images, binaries, serverless functions, firmware, etc.). Scanning must cover the artifact itself, all declared dependencies, and all transitive dependencies.

CVS–0.2 Severity Classification Using Current Standards

All identified vulnerabilities must be classified using CVSS v3.x or later standards, with explicit assignment of severity tiers: Critical, High, Medium, Low, and Informational. Severity classification must be supplemented by EPSS scoring and by cross-reference against the CISA Known Exploited Vulnerabilities (KEV) catalog. Vendors who classify vulnerabilities solely on CVSS score without EPSS or KEV enrichment are not meeting this control.

CVS–0.3 AI-Augmented Vulnerability Analysis

Vendors must demonstrate integration of AI-assisted tooling in their vulnerability identification pipeline. AI tooling may take the form of fine-tuned security models, general-purpose frontier models with security-specific context harnesses, or validated third-party AI scanning services. The purpose of this requirement is not to prescribe a specific model but to establish that vendor scanning posture is not limited to deterministic pattern matching against known CVE databases. Novel, zero-day-class vulnerabilities that do not yet have CVE assignments must be discoverable within the vendor's pipeline.

CVS–0.4 Automatic Detection, Mitigation, and Recall of Critical Findings

When a Critical or High severity vulnerability with a known exploit is identified in a deployed release, the vendor's platform must satisfy three requirements:

- [•] **Automated Detection and Escalation**
  Identification of the finding and escalation to the appropriate response workflow must occur without manual initiation.

- [•] **Automated Interim Mitigation**
  The platform must automatically apply compensating controls that reduce exposure (e.g., network isolation, feature disablement, or configuration changes) while the full remediation is in progress.

- [•] **One-Action Recall or Quarantine Across the Fleet**
  The platform must be capable of recalling or quarantining the affected release across all affected environments as a single orchestrated action. Human approval may gate execution of the recall, but the orchestration mechanism itself must not depend on manual, environment-by-environment intervention.

CVS–0.5 Contextual Vulnerability Prioritization and Remediation SLAs

Vendors must establish, publish, and demonstrate adherence to remediation SLAs for identified vulnerabilities. SLAs must satisfy two requirements:

First, prioritization must incorporate exploitability context, not raw severity alone.

[•] CVSS scores, particularly when sourced from CNAs without independent enrichment, are insufficient as a sole prioritization input. Vendors must demonstrate that their SLA assignment incorporates exploitability context, including at minimum:

- **EPSS and KEV enrichment (per CVS-0.2)**
- **Reachability analysis at both the source code and dependency level (is the vulnerable function actually called by first-party or transitively included code?) and the runtime level (is the vulnerable code path reachable from an attacker-controllable surface in the deployed environment?)**
- **Asset and data context (criticality of the affected component, sensitivity of data processed, blast radius of compromise)**
- **Attack path context as defined in APM-1.3**

[•] Vulnerabilities determined to be unreachable or otherwise non-exploitable in the deployed context may be deprioritized or risk-accepted through the suppression mechanism defined in ARO-3.4, with documented justification. Vendors who apply uniform SLAs based on CVSS score alone, without contextual enrichment, are not meeting this control.

Second, SLAs must be defined by severity tier and adhered to in practice.

[•] Vendors must publish remediation SLAs covering, at minimum, Critical, High, and Medium severity tiers as determined after contextual enrichment. SLA durations must be defensible relative to the exploitability of the finding — vulnerabilities with active known exploits (KEV-listed or with high EPSS scores) must carry materially shorter SLAs than vulnerabilities of equivalent CVSS score without active exploitation evidence.

[•] SLA adherence must be demonstrated through platform telemetry per ARO-3.6, not manual reporting. Vendors must report SLA compliance rates by severity tier as part of their MA-S2 attestation evidence.

3.3 Disqualifying Deficiencies

A vendor fails this control domain if they cannot demonstrate:

- [•] **Automated per-release scanning, or reliance purely on periodic manual scanning**
- [•] **EPSS or KEV enrichment of vulnerability classifications**
- [•] **AI-assisted component(s) in the discovery pipeline**
- [•] **Automatic escalation workflows for Critical or High severity findings**
- [•] **Documented and measurable remediation SLAs**

4.0 Control Domain 1: Attack Path Modeling and AI-Assisted Adversarial Simulation (APM)

4.1 Control Objective

Vendors must demonstrate that their security posture accounts for the use of AI by adversaries to traverse multi-stage attack paths through software systems. This includes the ability to model how AI-assisted adversaries would sequence vulnerabilities across their software stack and to prioritize remediation based on exploitability in those respective contexts.

Control Domain 1 requires vendors to model and handle their exposure at the attack path level across the entire landscape of components, micro-services, and other elements of their architecture, such that chained findings that together constitute a certain severity are appropriately handled despite how individual findings may be handled in isolation.

4.2 Required Controls

APM–1.1 Attack Path Modeling Capability

Vendors must demonstrate the capability to model multi-stage attack paths across their architecture. This may be implemented through dedicated threat modeling platforms, AI-assisted attack graph generation, or documented manual threat modeling methodologies reviewed on a regular cadence. The output of attack path modeling must be technically integrated into vulnerability prioritization tooling and decisions.

APM–1.2 Adversarial AI Simulation

Vendors must demonstrate that they have subjected their software to testing by AI-assisted tooling that simulates adversary behavior, including chaining of vulnerabilities across attack path stages. Evidence may include outputs from internal red team exercises using AI tooling, third-party AI-assisted penetration testing reports, or participation in AI-native bug bounty programs. While this standard does not prescribe specific tooling, it requires going beyond traditional penetration testing and providing evidence of adversarial AI simulation.

APM–1.3 Contextual Triage Integration

Findings from Control Domain 1 must be enriched with attack path context in parallel with the prioritization of remediations. A Critical CVE in a component that is not reachable from any external attack surface may be appropriately deprioritized. A Medium CVE in a component that is the first link in a traversable attack path to a privileged credential store must be treated as urgent. Vendors must demonstrate that their triage workflows reflect this contextual logic, rather than relying solely on a raw CVE severity score.

APM–1.4 Threat Intelligence Integration

Vendors must demonstrate integration with current threat intelligence feeds, including nation-state actor TTPs aligned to frameworks such as MITRE ATT&CK, as inputs to attack path modeling. Static threat models that are not updated with current intelligence are not acceptable. Vendors must demonstrate that their adversarial modeling reflects the capabilities and targeting priorities of the primary threat actors relevant to their customers' industries.

4.3 Disqualifying Deficiencies

A vendor fails this control domain if they cannot demonstrate:

- [•] **Any attack path modeling capability beyond individual CVE management**
- [•] **Ongoing AI-assisted adversarial simulation**
- [•] **A triage workflow that incorporates context, rather than relying solely on raw CVE scores**
- [•] **Integration of external threat intelligence into prioritization decisions**

5.0 Control Domain 2: Real-Time Software Inventory and Domain Awareness (INV)

5.1 Control Objective

Vendors must maintain a complete, real-time, machine-readable inventory of all software components across every environment in which their software operates. This inventory must include version, runtime configuration, and deployment location. It must be continuously updated, not generated on demand. Finally, it must be comprehensive across heterogeneous and all environments.

Control Domain 2 requires vendors to a maintain real-time and fully comprehensive inventory of their entire software supply chain, including first-party and inherited open-source code and all upstream systems / tooling upon which additional code is written.

5.2 Required Controls

INV–2.1 Software Bill of Materials (SBOM) at Release

Every release artifact must be accompanied by a machine-readable SBOM in a current standard format (SPDX or CycloneDX). The SBOM must include all direct and transitive dependencies with explicit version pinning. Generic or range-versioned dependencies are not acceptable.

INV–2.2 Runtime Inventory Continuous Reconciliation

Vendors must demonstrate a continuous reconciliation capability that compares the SBOM of each deployed release against the actual runtime state of each deployment environment. Discrepancies (e.g., components present at runtime that are not in the declared SBOM, components that are in the declared SBOM but are not present at runtime, etc.) must generate automatic alerts. This reconciliation must function across heterogeneous environments, including those with limited or intermittent connectivity.

INV–2.3 Environment-Level Deployment Visibility

Vendors must provide visibility into the precise release version deployed in each environment, the timestamp of the most recent successful deployment, the timestamp of the most recent successful scan against that deployed version, and any open findings against that environment. This visibility must be available through a machine-readable API, in addition to human-readable / human-usable operational tooling.

INV–2.4 Supply Chain Visibility

Vendors must demonstrate visibility into the security posture of their upstream software supply chain, including all vendors from whom they source components and all vendors upon which they build whitelabeled tooling. Evidence of upstream SBOM collection, upstream vulnerability monitoring, and supplier security attestation requirements must be produced.

INV–2.5 Air-Gapped and Disconnected Environment Coverage

Inventory and monitoring capabilities must extend to disconnected, air-gapped, and limited-connectivity environments. Vendors whose inventory systems require continuous outbound connectivity are not meeting this control. Inventory updates must be synchronized when connectivity is available, with defined maximum allowable drift between last-known and current-state inventory.

5.3 Disqualifying Deficiencies

A vendor fails this control domain if they cannot demonstrate:

- [•] **Machine-readable SBOMs for current and prior releases**
- [•] **Ongoing runtime reconciliation capability**
- [•] **Environment-level deployment state visibility via API**
- [•] **Comprehensive supply chain visibility beyond their own first-party code**
- [•] **Ongoing inventory coverage for air-gapped or disconnected environments**

6.0 Control Domain 3: Autonomous Remediation Orchestration (ARO)

6.1 Control Objective

Vendors must demonstrate the capability to autonomously orchestrate remediation actions across their full deployment fleet. This includes the deployment of patched releases, rollback of vulnerable releases, environment-specific suppression of findings, and more. Human authorization and / or involvement may be prompted or even required for specific classes of actions, but it must not be the bottleneck.

Control Domain 3 requires vendors to close the vulnerable code chain across every environment where it is deployed on an ongoing basis with autonomous remediation orchestration.

6.2 Required Controls

ARO–3.1 Automated Patch Deployment

Vendors must demonstrate the capability to automatically deploy patched releases to affected environments upon successful scan validation and approval workflow completion. Deployments must be zero-downtime where operationally required. Rollback capability must be tested and documented, with a demonstrated restoration time objective (RTO) and consistent and expedient expected rollback RTO.

ARO–3.2 Fleet-Wide Remediation Orchestration

Patch deployment must be orchestratable across the vendor's full deployment fleet from a single control plane. Specifically, vendors must have the ability to orchestrate patch deployments across all customer environments simultaneously or in a controlled, sequenced rollout. This includes fleet-wide orchestration capability across disconnected and air-gapped environments.

ARO–3.3 Compliance-Aware Change Management

Remediation orchestration must be compliance-aware, such that environment-specific requirements can be modeled and respected. This includes the modeling and automation of approval requirements, change freeze windows, accreditation constraints (e.g., FedRAMP, IL5 / IL6), and operator-defined deployment policies. Where there may exist discrepancies between environment-specific constraints and proposed security remediation actions, the orchestration layer must resolve these constraints automatically.

ARO–3.4 Vulnerability Suppression with Audit Trail

Vendors must provide a formal mechanism for time-bounded, risk-accepted suppression of specific vulnerabilities, including a full audit trail. Suppressions must be subject to automated expiration, and be made available upon request.

ARO–3.5 Developer-Operator Workflow Integration

Where developer teams responsible for producing patched code and operator teams responsible for deployment authorization are organizationally distinct, the vendor must demonstrate a workflow that produces measurable, timestamped records for three milestones, outlined below. These milestones must be captured in a system of record (such as a ticketing system, workflow platform, or release management tool) and must be reportable on demand. Informal communications (chat, email, code review discussion) are not in scope of this control.

- [•] **Vulnerability Discovery**
  The point at which a vulnerability is identified and assigned to a responsible developer for resolution.

- [•] **Patch Availability**
  The point at which a resolved version of the affected software is available for deployment.

- [•] **Deployment Authorization(s)**
  The point(s) at which the operator(s) responsible for each affected environment authorize deployment of the patched release.

ARO–3.6 Remediation Metrics and Reporting

Vendors must produce on a mutually agreed cadence a remediation report that includes critical metrics such as mean time to remediate (MTTR) by severity tier, SLA compliance rate, open finding count by severity and age, suppression inventory, and fleet-wide deployment coverage. These metrics must be generated from platform telemetry.

6.3 Disqualifying Deficiencies

A vendor fails this control domain if they cannot demonstrate:

- [•] **Automated patch deployment**
- [•] **Autonomous fleet-wide orchestration capability**
- [•] **Compliance-aware change management**
- [•] **A formal, auditable suppression mechanism**
- [•] **Measurable MTTR and SLA compliance reporting from platform telemetry**

7.0 Final Note

The current landscape of adversarial threats is extreme, and the advent of cyber-tuned LLMs is skyrocketing the rate of vulnerability identification. With hundreds of vendors rushing in to be the harness for specific models, CISOs must grapple with new bottlenecks: the need to create real-time visibility into exactly what versions of software are running across every type of environment, the ability to achieve rapid remediation with exact precision, and having the capacity to reliably close the remediation chain through autonomous orchestration.

We invite any organization that shares this assessment of the threat environment to contribute to MA-S2 as a standard for the software landscape, and as a vendor requirement for organizations' software supply chains.

Appendix

MA-S2 Control Summary

Control Domain	Control ID	Control Overview
Continuous, AI-Augmented Vulnerability Scanning	CVS–0.1	Automated Container and Dependency Scanning
	CVS–0.2	Severity Classification Using Current Standards
	CVS–0.3	AI-Augmented Vulnerability Analysis
	CVS–0.4	Automatic Detection, Mitigation, and Recall of Critical Findings
	CVS–0.5	Contextual Vulnerability Prioritization and Remediation SLAs
Attack Path Modeling and AI-Assisted Adversarial Simulation	APM–1.1	Attack Path Modeling Capability
	APM–1.2	Adversarial AI Simulation
	APM–1.3	Contextual Triage Integration
	APM–1.4	Threat Intelligence Integration
Real-Time Software Inventory and Domain Awareness	INV–2.1	Software Bill of Materials (SBOM) at Release
	INV–2.2	Runtime Inventory Continuous Reconciliation
	INV–2.3	Environment-Level Deployment Visibility
	INV–2.4	Supply Chain Visibility
	INV–2.5	Air-Gapped and Disconnected Environment Coverage
Autonomous Remediation Orchestration	ARO–3.1	Automated Patch Deployment
	ARO–3.2	Fleet-Wide Remediation Orchestration
	ARO–3.3	Compliance-Aware Change Management
	ARO–3.4	Vulnerability Suppression with Audit Trail
	ARO–3.5	Developer-Operator Workflow Integration
	ARO–3.6	Remediation Metrics and Reporting

Appendix

Proposed Minimum Questions for Software & AI Procurement Processes

Any organization incorporating MA-S2 requirements into internal assessments of their own IT exposure and current software landscape, and any organization looking to evaluate their software vendors against this standard, should at minimum require sufficient evidence of answers to the following questions.

[0] Describe your automated vulnerability scanning process. At what point in the release pipeline does scanning occur? What tooling is used, and does it include AI-assisted discovery of novel vulnerabilities beyond known CVE pattern matching?

[1] How do you classify and prioritize vulnerabilities? Do you incorporate EPSS scores and the CISA KEV catalog in addition to CVSS scores?

[2] Describe your software inventory capability. Can you produce a real-time, machine-readable inventory of what version of what software is running in each permutation of an environment where your software is deployed, including air-gapped environments? How is this inventory kept current?

[3] Have you subjected your software to AI-assisted adversarial simulation whereby vulnerabilities are chained across multi-stage attack paths? Describe the methodology and provide summary results or third-party attestation.

[4] How do you orchestrate patch deployment across your fleet? Can patched releases be deployed to all affected customer environments from a single control plane? What is your demonstrated rollback time?

[5] Who is responsible for remediating a vulnerability in your product, and how long does it take from identification to deployment of a fix in a customer environment?

[6] How does your remediation process account for customers with air-gapped, classified, or compliance-constrained environments? Can you demonstrate automated, compliance-aware change management across those environments?