MYTHOS-GRC-0001

Federal, DoD, and Sovereign Compliance Standard

Establishes an evidence-driven operating model for systems that process federal, defense, sovereign, or Controlled Unclassified Information, emphasizing current authorization, technical enforcement, data boundaries, zero trust, supply-chain assurance, and cryptographic transition.

6	8	GRC
EVALUATION CRITERIA	CURRENT AUTHORITY RECORDS	CANONICAL SERIES

CANDIDATE STANDARD

Mythos Systems | Engineering Standards Program | 2026-07-22

This candidate publication is complete for technical review but does not by itself establish certification, legal compliance, implementation conformance, or final approval.

Contents

Use the interactive entries below or the PDF bookmark outline to navigate the publication.

Document Control

Field	Value
Document ID	MYTHOS-GRC-0001
Version	1.1.0 - Enterprise Candidate Edition
Status	Candidate Standard
Classification	Public
Publisher	Mythos Systems
Owner	Aaron Stovall
Domain	Governance, Risk, and Compliance
Initial source date	2026-07-16
Enterprise reconstruction	2026-07-22
Review cadence	Quarterly for dynamic authorities; at least annually for stable controls
Supersedes	Source draft retained in Appendix A
Publication lineage	Permanent canonical series

CANDIDATE STATUS

The source document identified itself as a definitive published standard. This enterprise edition corrects that status. It remains a candidate until current-source validation, qualified domain review, assessment engineering, accessibility review, and formal approval are recorded.

Part I - Executive Direction

1. Executive Overview

Establishes an evidence-driven operating model for systems that process federal, defense, sovereign, or Controlled Unclassified Information, emphasizing current authorization, technical enforcement, data boundaries, zero trust, supply-chain assurance, and cryptographic transition.

2. Principal Findings

- Replace audit-snapshot compliance with continuously observed and evidence-backed control operation.
- Treat CUI and sovereign-data boundaries as enforceable architecture, not document labels.
- Automate hardening, drift detection, remediation, evidence collection, and exception handling.
- Maintain exact program-state awareness because defense and federal requirements change on short cycles.

3. Required Leadership Decisions

- Name an accountable authorizing and compliance engineering function.
- Fund evidence automation, policy-as-code, boundary validation, and continuous control monitoring.
- Adopt a controlled authority register with version, validation date, expiration, and invalidation triggers.
- Require documented applicability review by qualified federal, legal, and contracting specialists before compliance claims.

4. Enterprise Target State

TARGET OPERATING MODEL

Organizations applying MYTHOS-GRC-0001 operate the subject as a governed lifecycle: inventory and classify; establish authoritative policy and architecture; enforce controls technically; observe actual state; verify outcomes independently; retain evidence; manage exceptions; and revalidate when authorities, platforms, risks, or operating conditions change.

5. Business, Safety, and Operational Impact

- Reduces reliance on manual evidence and undocumented expert knowledge.
- Makes risk acceptance, exceptions, and unresolved uncertainty visible to accountable leaders.
- Creates measurable implementation and assurance outcomes rather than a one-time score.
- Supports procurement, architecture review, operational readiness, and independent assessment from one source-controlled publication.

Part II - Standard Foundation

6. Purpose and Scope

Establishes an evidence-driven operating model for systems that process federal, defense, sovereign, or Controlled Unclassified Information, emphasizing current authorization, technical enforcement, data boundaries, zero trust, supply-chain assurance, and cryptographic transition.

6.1 Applicability

Apply this standard according to system purpose, deployment model, data classification, user population, safety impact, regulatory obligations, and organizational risk tolerance. Tailoring must be documented. A lower implementation profile cannot silently waive a mandatory legal, safety, accessibility, contractual, or security obligation.

6.2 Out of Scope

This publication does not replace legal advice, contracting interpretation, regulator guidance, platform-specific engineering, human-factors testing, safety certification, or implementation evidence. Companion Mythos standards remain applicable where their domains intersect.

7. Requirements Language and Conformance

The terms MUST, MUST NOT, SHOULD, SHOULD NOT, and MAY are normative only when written in uppercase. Conformance requires applicable mandatory criteria to be implemented and supported by current evidence. A weighted score is informative and cannot compensate for a failed mandatory gate.

State	Meaning
Conformant	All applicable mandatory requirements are satisfied with sufficient current evidence.
Partially Conformant	Some applicable requirements are implemented, but material gaps remain.
Not Conformant	One or more mandatory requirements are not satisfied.
Compensating Control Accepted	A formally approved alternative achieves the required outcome.
Exception Active	A time-bounded, accountable risk acceptance exists.
Not Assessed	Evidence is insufficient to reach a conclusion.

8. Assurance and Maturity

Level	Description
M0 Unmanaged	The outcome is not intentionally controlled.
M1 Documented	Responsibilities and procedures exist but are inconsistent or manual.
M2 Implemented	The control operates in the applicable environment.
M3 Measured	Performance, exceptions, and effectiveness are measured.
M4 Assured	Independent testing and continuous evidence support the control.
M5 Adaptive	The control responds safely to changes in risk, environment, and evidence.

Part III - Risk and Architecture

9. Risk Model

- Authority drift: external requirements, specifications, programs, and platform behavior change faster than the publication is reviewed.
- Control illusion: documentation or visual state indicates compliance while runtime behavior differs.
- Automation overreach: a generated plan or score is treated as authorization or proof.
- Evidence gaps: conclusions cannot be reconstructed from source, configuration, event, test, and approval records.
- Human factors: users cannot understand, interrupt, challenge, or safely recover from consequential behavior.
- Dependency concentration: a provider, runtime, device, framework, or external authority becomes a hidden single point of failure.

10. Reference Control Architecture

ARCHITECTURE SEPARATION

Source authority defines applicable obligations and constraints. Human and organizational governance establishes accountability. Typed policy and architecture encode the target state. Runtime enforcement controls behavior. Independent testing verifies outcomes. Evidence systems preserve what was proposed, approved, executed, observed, and verified.

11. Roles and Accountability

Role	Accountability
Executive sponsor	Approves risk posture, funding, and unresolved material exceptions.
Domain authority	Owns interpretation, architecture, and control applicability.
Engineering owner	Implements controls and operational lifecycle.
Security/privacy/accessibility/safety reviewer	Challenges design and verifies domain-specific obligations.
Independent assessor	Evaluates evidence and tests without relying only on implementer assertions.
Publication owner	Maintains version, sources, traceability, expiration, and change record.

Part IV - Normative Evaluation and Control System

12. Evaluation Doctrine

The preserved source uses a weighted 0-5 evaluation model. This edition retains that material but corrects its interpretation. Score 5 means optimized, strongly evidenced, and independently assured within the assessed scope. It never means perfect, mathematically infallible, universally compliant, or permanently secure.

Score	Enterprise interpretation
0	Absent, unsafe, or unable to perform the required outcome.
1	Initial or ad hoc capability with substantial manual dependence.
2	Repeatable partial implementation with material gaps.
3	Production baseline: implemented, governed, and evidenced.
4	Advanced: automated, measured, resilient, and independently tested.
5	Assured: optimized for the defined profile, continuously evidenced, and subject to recurring independent validation.

MANDATORY GATE

An organization cannot average away a failed mandatory criterion. Applicability, safety, legal requirements, accessibility, security boundaries, authorization, and evidence sufficiency must be evaluated before weighted scoring is used.

13. Criterion Index

ID	Criterion	Weight	Primary outcome
MYTHOS-GRC-0001-C01	NIST RMF and cATO Enforcement	20%	Maintain a current control implementation model that links system components, control owners, evidence, exceptions, changes, and authorization state.
MYTHOS-GRC-0001-C02	CUI Isolation and DFARS 7012	20%	Define and technically enforce CUI boundaries, identity domains, storage, transit, administration, incident reporting, and approved external-service use.
MYTHOS-GRC-0001-C03	DISA STIG Automation (Policy-as-Code)	20%	Convert applicable technical baselines into versioned, reviewable, testable policy and retain approved exceptions when automation is not possible.
MYTHOS-GRC-0001-C04	DoD Zero Trust Architecture	15%	Implement identity, device, workload, data, network, automation, and analytics controls as an integrated access and trust system.
MYTHOS-GRC-0001-C05	Supply Chain (CISA BODs and SBOMs)	15%	Maintain software, firmware, model, infrastructure, and service provenance with vulnerability and directive response tied to deployment decisions.
MYTHOS-GRC-0001-C06	Cryptographic Sovereignty (NSA CNSA 2.0)	10%	Maintain a cryptographic inventory, algorithm policy, key ownership model, migration roadmap, and evidence for approved national-security profiles.

MYTHOS-GRC-0001-C01 - NIST RMF and cATO Enforcement

NORMATIVE REQUIREMENT

The organization MUST maintain a current control implementation model that links system components, control owners, evidence, exceptions, changes, and authorization state.

Objective and Risk Addressed

Maintain a current control implementation model that links system components, control owners, evidence, exceptions, changes, and authorization state.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.
- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- Automated control telemetry
- approved SSP and architecture
- POA&M and exception records
- authorization decision
- evidence freshness report.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.
- Sample control evidence, reproduce drift detection, validate evidence lineage, and confirm changes invalidate stale authorization assumptions.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
control evidence freshness	Track trend, threshold, exception, and accountable response.
unresolved high-risk findings	Track trend, threshold, exception, and accountable response.
drift detection latency	Track trend, threshold, exception, and accountable response.
automated evidence coverage	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	System Security Plan (SSP) is a static Word document; controls verified manually annually
1	Controls mapped in a GRC tool (e.g., eMASS), but evidence collection is manual
2	Automated scripts pull config data monthly for audit prep

Score	Preserved source criterion
3	cATO achieved: real-time control telemetry streamed to AO dashboard; automated drift detection
4	Automated remediation: non-compliant drift is instantly reverted to meet NIST SP 800-53 bounds
5	Perfect compliance: formally verified control execution, zero gap between policy and runtime, mathematical proof of continuous authorization

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

MYTHOS-GRC-0001-C02 - CUI Isolation and DFARS 7012

NORMATIVE REQUIREMENT

The organization MUST define and technically enforce CUI boundaries, identity domains, storage, transit, administration, incident reporting, and approved external-service use.

Objective and Risk Addressed

Define and technically enforce CUI boundaries, identity domains, storage, transit, administration, incident reporting, and approved external-service use.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.
- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- data-flow diagrams
- boundary rules
- encryption and key records
- egress logs
- provider eligibility
- incident exercise results.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.

- Trace representative CUI through storage, processing, backup, support, and egress; test denied paths and privileged access.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
unapproved egress events	Track trend, threshold, exception, and accountable response.
boundary-policy coverage	Track trend, threshold, exception, and accountable response.
CUI asset inventory accuracy	Track trend, threshold, exception, and accountable response.
incident notification readiness	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	CUI stored in standard multi-tenant SaaS without boundary isolation
1	Logical separation via VPCs, but shared identity/IAM planes
2	Dedicated CUI enclaves, but manual egress filtering
3	Cryptographic CUI tagging; strict egress firewalls (zero-trust proxy) blocking all non-DFARS 7012 compliant endpoints
4	Hardware-backed CUI isolation (Confidential VMs/TEEs) ensuring cloud admins cannot access data
5	Perfect isolation: formally verified CUI data lineage, zero ability for CUI to traverse non-approved boundaries, cryptographic attestation of tenant isolation

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

MYTHOS-GRC-0001-C03 - DISA STIG Automation (Policy-as-Code)

NORMATIVE REQUIREMENT

The organization MUST convert applicable technical baselines into versioned, reviewable, testable policy and retain approved exceptions when automation is not possible.

Objective and Risk Addressed

Convert applicable technical baselines into versioned, reviewable, testable policy and retain approved exceptions when automation is not possible.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.
- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- STIG/SRG applicability register
- policy repository
- scanner results
- remediation evidence
- exception approvals
- release gates.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.
- Select high-impact STIG checks, compare policy output to authoritative content, introduce drift, and verify detection and remediation.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
automated check coverage	Track trend, threshold, exception, and accountable response.
mean remediation time	Track trend, threshold, exception, and accountable response.
exception age	Track trend, threshold, exception, and accountable response.
false-positive rate	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	STIGs reviewed manually via SCCM/ACAS scans; remediation is ad-hoc
1	Some STIGs automated via Ansible/Puppet, but runbooks are manual
2	Baseline infrastructure coded in Terraform, but STIG compliance not gated in CI/CD
3	Policy-as-Code (OPA/Sentinel) in CI/CD blocks the deployment of any resource failing STIG compliance
4	Continuous runtime STIG validation (e.g., eBPF/Falco) detecting and killing non-compliant processes
5	Perfect enforcement: formally verified STIG mapping, zero ability to deploy non-compliant infrastructure, cryptographic proof of configuration integrity

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

MYTHOS-GRC-0001-C04 - DoD Zero Trust Architecture

NORMATIVE REQUIREMENT

The organization MUST implement identity, device, workload, data, network, automation, and analytics controls as an integrated access and trust system.

Objective and Risk Addressed

Implement identity, device, workload, data, network, automation, and analytics controls as an integrated access and trust system.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.
- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- zero-trust roadmap
- policy decisions
- identity/device posture
- segmentation tests
- data controls
- telemetry.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.
- Test access under changing identity, device, location, workload, and data conditions; verify no implicit trust path bypasses policy.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
policy-evaluated access percentage	Track trend, threshold, exception, and accountable response.
stale trust duration	Track trend, threshold, exception, and accountable response.
privileged standing access	Track trend, threshold, exception, and accountable response.
segmentation escape findings	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	Perimeter-based security (VPN + flat network)
1	Basic MFA and network segmentation
2	Microsegmentation and identity-based access (IL4 level)
3	DoD Zero Trust Target Level: Continuous user/device authentication, data-centric security, automated policy enforcement
4	Advanced Zero Trust: AI-driven behavioral analytics, dynamic capability grants (Macaroons), cross-domain solutions
5	Perfect alignment: formally verified zero-trust pillars, zero implicit trust, mathematical proof of least-privilege execution

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

MYTHOS-GRC-0001-C05 - Supply Chain (CISA BODs and SBOMs)

NORMATIVE REQUIREMENT

The organization MUST maintain software, firmware, model, infrastructure, and service provenance with vulnerability and directive response tied to deployment decisions.

Objective and Risk Addressed

Maintain software, firmware, model, infrastructure, and service provenance with vulnerability and directive response tied to deployment decisions.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.

- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- SBOMs/AIBOMs
- signed provenance
- KEV response records
- supplier attestations
- deployment gates
- exception record.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.
- Inject a prohibited or known-exploited component into a test build and verify prevention, detection, notification, and evidence.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
inventory coverage	Track trend, threshold, exception, and accountable response.
KEV remediation time	Track trend, threshold, exception, and accountable response.
unsigned artifact rate	Track trend, threshold, exception, and accountable response.
supplier evidence freshness	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	No SBOM; manual patching cycles
1	Basic vulnerability scanning (Nessus)
2	SBOM generated for all applications; CISA KEV catalog monitored manually
3	Automated compliance with CISA BODs (e.g., BOD 22-01) within mandated timelines; SBOMs cryptographically signed
4	Continuous dependency analysis blocking deployments if a CISA KEV vulnerability is detected
5	Perfect resilience: formally verified supply chain provenance, zero unpatched KEVs, real-time threat intelligence integration

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

MYTHOS-GRC-0001-C06 - Cryptographic Sovereignty (NSA CNSA 2.0)

NORMATIVE REQUIREMENT

The organization MUST maintain a cryptographic inventory, algorithm policy, key ownership model, migration roadmap, and evidence for approved national-security profiles.

Objective and Risk Addressed

Maintain a cryptographic inventory, algorithm policy, key ownership model, migration roadmap, and evidence for approved national-security profiles.

Applicability

Apply this criterion wherever the evaluated capability, data, workflow, interface, user, environment, provider, or authority is material. Document exclusions and any compensating control.

Minimum Acceptable Implementation

A production baseline corresponds to source score 3: the capability is deliberately implemented, technically enforced where practical, owned, monitored, and supported by current evidence. Source scores 4 and 5 represent increasing automation, resilience, continuous assurance, and independent validation.

Implementation Direction

- Define authoritative state, owner, interfaces, dependencies, lifecycle, and failure behavior.
- Encode enforceable policy and typed contracts instead of relying only on prose or visual convention.
- Provide safe fallback, interruption, exception, recovery, and decommissioning behavior.
- Collect evidence at the point of action and verify the observed result independently.
- Revalidate after material source, platform, model, device, provider, architecture, or risk change.

Required Evidence

- crypto inventory
- key and HSM records
- algorithm policy
- migration milestones
- interoperability tests
- exception approvals.

Assessment and Test Procedure

- Confirm applicability, ownership, current architecture, and approved policy.
- Identify prohibited or transition algorithms, validate negotiated protocols and artifact signatures, and test rollback or interoperability paths.
- Inspect failures, exceptions, and negative tests rather than reviewing only successful examples.
- Rate evidence confidence and record untested conditions, unsupported platforms, and unresolved uncertainty.

Metrics and Assurance Indicators

Metric	Expected use
approved-algorithm coverage	Track trend, threshold, exception, and accountable response.
unknown crypto instances	Track trend, threshold, exception, and accountable response.
key rotation compliance	Track trend, threshold, exception, and accountable response.
migration completion	Track trend, threshold, exception, and accountable response.

Preserved Source Scoring Anchors

The following anchors are retained from the source draft. Absolute words such as “perfect” are historical wording and are normalized by the enterprise interpretation above.

Score	Preserved source criterion
0	Commercial cryptography (RSA/SHA-256); no PQC readiness
1	CNSA 1.0 Suite (ECC/P-384) implemented
2	CNSA 2.0 Suite algorithms (ML-KEM/ML-DSA) implemented
3	Hardware-backed (HSM/TPM) CNSA 2.0 keys; PQC signatures on all artifacts
4	Hybrid classical+PQC TLS enforced for all CUI data in transit
5	Perfect sovereignty: formally verified PQC bounds, zero deprecated algorithms, mathematical proof of quantum resistance

Exceptions and Compensating Controls

Any exception must identify the unmet outcome, affected scope, owner, risk, compensating control, evidence, expiration, and remediation plan. Exceptions do not change the underlying requirement.

Part V - Implementation and Operations

14. Implementation Profiles

Profile	Required posture
Baseline	Documented ownership, minimum production controls, evidence, and exception process.
Enterprise	Central policy, automation, integration, continuous monitoring, and independent review.
High Assurance	Strong isolation, high-assurance identity, formal or adversarial verification, short evidence windows, and two-person governance where material.
Sovereign or Air-Gapped	Local control planes, approved dependencies, offline update and evidence workflows, restricted egress, and explicit supply-chain custody.

15. Operational Lifecycle

Stage	Required activities
Discover and classify	Inventory scope, authority, data, users, dependencies, risks, and current evidence.
Design	Define target state, architecture, policies, tests, metrics, and ownership.
Implement	Deploy controls through reviewed changes with rollback and evidence.
Operate	Monitor state, exceptions, incidents, drift, performance, and provider changes.
Verify	Perform recurring independent assessment and negative testing.
Change	Reassess applicability and invalidate stale evidence after material changes.
Retire	Revoke authority, remove data and credentials, preserve required evidence, and update the registry.

16. Anti-Patterns and Prohibited Practices

- Declaring conformance from a document review without implementation evidence.
- Using one weighted score to override a failed mandatory requirement.
- Treating model output, interface state, scanner output, or tool success as independent verification.
- Using stale or unversioned external authorities for current decisions.
- Hiding exceptions, unsupported platforms, incomplete testing, or low-confidence findings.
- Hard-coding a single provider, device, framework, or vendor as a universal requirement unless the publication is explicitly vendor-scoped.

Part VI - Assurance, Evidence, and Traceability

17. Evidence Model

Evidence must identify subject, scope, source, version, time, actor, method, result, integrity, retention, and relationship to the assessed criterion. Screenshots and generated summaries are supporting artifacts, not sufficient evidence by themselves.

18. Assessment Confidence

Confidence	Basis
Low	Self-assertion, incomplete samples, stale evidence, or untested material conditions.
Moderate	Current documentation and representative evidence with limited independent testing.
High	Current direct evidence and repeatable independent testing across material conditions.
Very High	Sustained evidence, broad negative testing, independent review, and verified operation over time.

19. Current Authority Register

Authority	Status and scope	Valid until
NIST SP 800-171 Rev. 3	Final - CUI security requirements for nonfederal systems	2027-01-18
NIST SP 800-171A Rev. 3	Final - Assessment procedures for SP 800-171 Rev. 3	2027-01-18
NIST SP 800-53 Rev. 5, Release 5.2.0	Final - Federal security and privacy control catalog	2027-01-18
NIST SP 800-172 Rev. 3	Final - Enhanced CUI requirements for advanced threats	2027-01-18
DoD CMMC Program	Dynamic program - DIB assessment and contracting requirements; current status must be revalidated	2026-08-21
DISA STIG and SCAP Content	Rolling catalog - Technical hardening and assessment content	2026-08-21
DoD Zero Trust Strategy and Roadmap	Current strategy - Defense zero-trust outcomes and execution direction	2027-01-18
NSA CNSA 2.0	Current advisory suite - National-security cryptographic transition direction	2027-01-18

AUTHORITY LIMITATION

The register identifies current technical and governance authorities relevant to this candidate standard. It does not establish legal applicability, contractual compliance, certification, or clause-level equivalence. Dynamic sources must be refreshed before material decisions.

20. Publication Gate

- Qualified domain technical review completed.
- Current authorities and versions independently confirmed.
- Criterion requirements, evidence, and tests reviewed for measurability.
- Legal, contracting, regulatory, safety, privacy, and accessibility applicability reviewed where relevant.
- Machine-readable criterion catalog validated.
- PDF accessibility and visual quality reviewed.
- Formal approval, effective date, and review triggers recorded.

Appendix A - Complete Preserved Source Draft

The following source draft is preserved in full for completeness and traceability. Conversational framing, “definitive” status, universal claims, obsolete program assumptions, and “perfect” score language are historical source content and are not controlling statements in this enterprise candidate edition.

Document ID: MYTHOS-GRC-0001 Version: 1.0.0 (Definitive Registry Edition) Status: Published Standard Classification: Public Organization: Mythos Systems (MYTHOS AI) Owner: Aaron Stovall Effective Date: 2026-07-16 Review Cadence: Quarterly, and immediately after a material shift in NIST SP 800-171/172, CISA BODs, DISA STIGs, or DoD Zero Trust execution roadmaps Applies To: All organizations evaluating, selecting, or operating information systems processing, storing, or transmitting Controlled Unclassified Information (CUI), Federal data, or DoD workloads, including cloud service providers (FedRAMP) and defense industrial base (DIB) contractors Companion Standards: MYTHOS-SEC-0001 (PQC), MYTHOS-SEC-0003 (Zero-Trust & Capability Grants), MYTHOS-DAT-0003 (Data Sovereignty & Egress), MYTHOS-ID-0002 (PKI & ABAC), MYTHOS-PLT-0001 (DevSecOps & Supply Chain)

0. Executive Mandate

The architecture of federal and defense compliance has failed.

Organizations attempt to secure Controlled Unclassified Information (CUI) and pass audits by maintaining massive, static System Security Plans (SSPs) and executing manual DISA STIG checklists via spreadsheets. They achieve an Authority to Operate (ATO) via a grueling 18-month review process, only to immediately drift out of compliance the day after the audit because the controls were never technically enforced.

This standard exists because:

1. A Checklist is Not a Control: A human verifying that "audit logging is enabled" via a manual screenshot is theater. Controls MUST be enforced mathematically via Policy-as-Code (OPA/Sentinel) in the CI/CD pipeline and runtime, rendering non-compliant infrastructure impossible to deploy.
2. Point-in-Time ATOs are Obsolete: An ATO that expires after three years provides zero assurance of current security. Systems MUST achieve continuous ATO (cATO) status by streaming real-time control telemetry (evidence bundles) directly to the Authorizing Official (AO), eliminating the audit cycle.
3. CUI Data Sovereignty is Non-Negotiable: CUI cannot be treated as standard cloud data. It MUST be cryptographically tagged, isolated in zero-trust enclaves, and subjected to strict egress firewalls. If CUI can be copied to an unmanaged laptop or a third-party SaaS API, the architecture has failed.
4. Zero Trust is a Military Mandate: The DoD Zero Trust Strategy requires architecture to abandon perimeter-based security entirely. Systems MUST implement the 7 pillars of Zero Trust (User, Device, Application/Workload, Data, Network/Environment, Automation/Orchestration, Visibility/Analytics) as foundational, non-bypassable requirements.

This document establishes the Mythos Federal Compliance Standard (MFCS), a comprehensive, evidence-based methodology for evaluating sovereign and defense architectures against continuous compliance, automated STIG enforcement, and zero-trust CUI isolation.

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- NIST SP 800-171 Rev. 3 / SP 800-171A Rev. 3 / SP 800-172 (Enhanced CUI)
- NIST SP 800-53 / Risk Management Framework (RMF)
- DoD Cybersecurity Maturity Model Certification (CMMC) 2.0 (Dynamic Status)

- DISA Security Technical Implementation Guides (STIGs) and Security Requirements Guides (SRGs)
- FedRAMP and cATO (Continuous Authority to Operate) criteria
- DoD Zero Trust Strategy and Reference Architecture (Target Level)
- DFARS 252.204-7012 (Cyber Incident Reporting & Cloud requirements)
- CISA Binding Operational Directives (BODs)

1.2 Out of Scope

- General AI Governance and Risk (covered in MYTHOS-GRC-0002)
- General international privacy laws like GDPR (covered in MYTHOS-GRC-0002)
- Post-Quantum Cryptography algorithms (covered in MYTHOS-SEC-0001)

1.3 Audience

- System Security Engineers and Information System Security Officers (ISSOs)
- Cloud architects designing FedRAMP High / IL4/IL5/IL6 environments
- DevSecOps teams automating STIG and RMF controls
- Authorizing Officials (AOs) and CISOs managing federal risk
- Standards bodies seeking a reference DoD compliance methodology

Part II. Definitions and Taxonomy

2.1 Federal & DoD Dimensions

Dimension	Definition
cATO (Continuous ATO)	An authorization paradigm where an system maintains its Authority to Operate by providing real-time, automated evidence of continuous control compliance, eliminating the 3-year re-authorization cycle.
CUI (Controlled Unclassified Information)	Information that requires safeguarding or dissemination controls pursuant to and consistent with applicable law, regulations, and government-wide policies but is not classified under Executive Order 13526.
STIG (Security Technical Implementation Guide)	Cybersecurity methodologies for the standardization of security protocols within networks, servers, computers, and logical designs to enhance overall security.
Policy-as-Code (PaC)	The practice of writing security and compliance rules (e.g., STIGs) in a domain-specific language (Rego/Cedar), enforced automatically in the CI/CD pipeline before infrastructure is provisioned.
DoD Zero Trust Target Level	A security model requiring continuous verification of all users, devices, and data flows, eliminating implicit trust based on network location.

2.2 The Federal Compliance Doctrine

SOURCE STATEMENT

A screenshot is not evidence. An audit is a lagging indicator of failure. A CUI environment without automated egress control is breached. If the STIG is not enforced in code, it is a suggestion.

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; manual SSPs, no CUI isolation, perimeter-based security
1	Basic RMF documentation, but manual STIG checks and no cATO
2	Functional CUI boundary, but relies on periodic audits and manual remediation
3	Solid production performance; Policy-as-Code STIGs, cATO telemetry, DoD Zero Trust Target Level
4	Strong performance; cryptographic CUI lineage, real-time drift remediation, FedRAMP High/IL5+
5	Best-in-class; sets the standard for mathematically verified, sovereign defense infrastructure

Weighting

Category	Weight	Rationale
NIST RMF & cATO Enforcement	20%	Point-in-time audits guarantee drift; compliance must be continuous
CUI Isolation & DFARS 7012	20%	CUI leakage to non-CMMC/IL compliant systems is a critical breach
DISA STIG Automation (PaC)	20%	Manual STIG checklists are unenforceable and obsolete
DoD Zero Trust Architecture	15%	Perimeter-based defense is explicitly banned by DoD strategy
Supply Chain (CISA BODs/SBOMs)	15%	Federal systems are prime targets for supply chain attacks
Cryptographic Sovereignty (NSA CNSA)	10%	Federal data requires highest-tier cryptographic resilience

Total: 100%

A system MUST score at least 3.0 in RMF/cATO and DISA STIG Automation to be considered for CUI production use.

Part IV. Evaluation Categories

4.1 NIST RMF and cATO Enforcement (Weight: 20%)

4.1.1 Continuous Compliance

Is compliance a continuous, automated state, or a point-in-time paperwork exercise?

Score	Criteria
0	System Security Plan (SSP) is a static Word document; controls verified manually annually
1	Controls mapped in a GRC tool (e.g., eMASS), but evidence collection is manual
2	Automated scripts pull config data monthly for audit prep
3	cATO achieved: real-time control telemetry streamed to AO dashboard; automated drift detection
4	Automated remediation: non-compliant drift is instantly reverted to meet NIST SP 800-53 bounds
5	Perfect compliance: formally verified control execution, zero gap between policy and runtime, mathematical proof of continuous authorization

4.2 CUI Isolation and DFARS 7012 (Weight: 20%)

4.2.1 Data Boundary Enforcement

is CUI protected from mixing with non-CUI or commercial data?

Score	Criteria
0	CUI stored in standard multi-tenant SaaS without boundary isolation
1	Logical separation via VPCs, but shared identity/IAM planes
2	Dedicated CUI enclaves, but manual egress filtering
3	Cryptographic CUI tagging; strict egress firewalls (zero-trust proxy) blocking all non-DFARS 7012 compliant endpoints
4	Hardware-backed CUI isolation (Confidential VMs/TEEs) ensuring cloud admins cannot access data
5	Perfect isolation: formally verified CUI data lineage, zero ability for CUI to traverse non-approved boundaries, cryptographic attestation of tenant isolation

4.3 DISA STIG Automation (Policy-as-Code) (Weight: 20%)

4.3.1 Technical Implementation Guides

Are DISA STIGs enforced architecturally, or hoped for via checklists?

Score	Criteria
0	STIGs reviewed manually via SCCM/ACAS scans; remediation is ad-hoc
1	Some STIGs automated via Ansible/Puppet, but runbooks are manual
2	Baseline infrastructure coded in Terraform, but STIG compliance not gated in CI/CD
3	Policy-as-Code (OPA/Sentinel) in CI/CD blocks the deployment of any resource failing STIG compliance
4	Continuous runtime STIG validation (e.g., eBPF/Falco) detecting and killing non-compliant processes
5	Perfect enforcement: formally verified STIG mapping, zero ability to deploy non-compliant infrastructure, cryptographic proof of configuration integrity

4.4 DoD Zero Trust Architecture (Weight: 15%)

4.4.1 Target Level Alignment

Does the architecture meet the DoD Zero Trust Reference Architecture Target Level?

Score	Criteria
0	Perimeter-based security (VPN + flat network)
1	Basic MFA and network segmentation
2	Microsegmentation and identity-based access (IL4 level)
3	DoD Zero Trust Target Level: Continuous user/device authentication, data-centric security, automated policy enforcement
4	Advanced Zero Trust: AI-driven behavioral analytics, dynamic capability grants (Macaroons), cross-domain solutions
5	Perfect alignment: formally verified zero-trust pillars, zero implicit trust, mathematical proof of least-privilege execution

4.5 Supply Chain (CISA BODs and SBOMs) (Weight: 15%)

4.5.1 Federal Supply Chain Resilience

How does the system comply with CISA BODs and federal SBOM mandates?

Score	Criteria
0	No SBOM; manual patching cycles
1	Basic vulnerability scanning (Nessus)
2	SBOM generated for all applications; CISA KEV catalog monitored manually
3	Automated compliance with CISA BODs (e.g., BOD 22-01) within mandated timelines; SBOMs cryptographically signed
4	Continuous dependency analysis blocking deployments if a CISA KEV vulnerability is detected
5	Perfect resilience: formally verified supply chain provenance, zero unpatched KEVs, real-time threat intelligence integration

4.6 Cryptographic Sovereignty (NSA CNSA 2.0) (Weight: 10%)

4.6.1 National Security Systems Alignment

Does the cryptography meet NSA CNSA 2.0 and Post-Quantum mandates?

Score	Criteria
0	Commercial cryptography (RSA/SHA-256); no PQC readiness
1	CNSA 1.0 Suite (ECC/P-384) implemented
2	CNSA 2.0 Suite algorithms (ML-KEM/ML-DSA) implemented
3	Hardware-backed (HSM/TPM) CNSA 2.0 keys; PQC signatures on all artifacts
4	Hybrid classical+PQC TLS enforced for all CUI data in transit
5	Perfect sovereignty: formally verified PQC bounds, zero deprecated algorithms, mathematical proof of quantum resistance

Part V. The Mythos Federal Compliance Suite (MFCS)

Traditional compliance audits measure documentation completeness. The MFCS evaluates architectural enforcement, CUI isolation, and automated STIG rigidity.

5.1 Suite Composition

5.1.1 MFCS-STIG

- Non-Compliant Deploy Test: 100+ tests attempting to deploy an S3 bucket lacking server-side encryption or a port open to 0.0.0.0/0, verifying the Policy-as-Code engine blocks the deployment.
- Runtime Drift Test: 50+ tests attempting to manually disable a firewall rule on a production server, verifying the system detects the STIG violation and auto-remediates it within 5 minutes.

5.1.2 MFCS-CUI

- Egress Exfiltration Test: 100+ tests attempting to copy a file tagged as CUI to a non-IL4 compliant cloud storage endpoint, verifying the zero-trust proxy blocks the transfer.
- cATO Telemetry Test: 50+ tests verifying that real-time control evidence (e.g., successful MFA logs, encryption status) is successfully streaming to the Authorizing Official's dashboard.

5.2 Execution Protocol

text

1. Deploy architecture in a FedRAMP/DoD compliant environment
2. Run MFCS suite (automated STIG injection + CUI egress attempts)
3. Score each category 0-5
4. Check for manual SSPs, open perimeters, or lack of SBOMs (instant disqualification)
5. If all thresholds met: approve for CUI production

Part VI. Security Threat Model for Federal/DoD Systems

6.1 Threat Catalog

6.1.1 The "Audit Snapshot" Rot

Severity: High Description: An organization passes its NIST SP 800-171 assessment. The next day, a developer misconfigures a CI/CD pipeline, disabling audit logging to save space. The system remains "certified" for 3 years while actively non-compliant.

Required Controls:

1. cATO (Continuous ATO) with real-time telemetry.
2. Automated remediation of control drift.

6.1.2 CUI Boundary Crossing (The SaaS Leak)

Severity: Critical Description: A defense contractor integrates a commercial AI tool (e.g., ChatGPT) to summarize CUI documents. The CUI is egressed to the third-party SaaS API, violating DFARS 7012 and CMMC requirements.

Required Controls:

1. Strict egress proxies blocking all non-approved endpoints.
2. Context-aware DLP identifying and redacting CUI before it leaves the boundary.

6.1.3 Advanced Persistent Threat (APT) via Supply Chain

Severity: Critical Description: A foreign intelligence agency compromises an open-source library used by a defense contractor's CI/CD pipeline, establishing a backdoor into the CUI enclave.

Required Controls:

1. Mandatory SBOM generation and Sigstore signing.
2. Hermetic builds and continuous CISA KEV monitoring.

6.1.4 CMMC Status Fluidity

Severity: Medium Description: A contractor assumes CMMC Phase II requirements are static. The requirements are suspended or altered by the DoD (as occurred on July 13, 2026), and the contractor's compliance architecture is misaligned with the new Phase I/II reality.

Required Controls:

1. Dynamic tracking of CMMC rule sets; architecture must adapt to DoD policy shifts without requiring complete re-architecture.
2. Adherence to the highest baseline (SP 800-172) to ensure resilience against downgrades/upgrades.

Part VII. The Mythos Federal Compliance Standard

7.1 The Mythos Federal Doctrine

text A screenshot is not evidence. An audit is a lagging indicator of failure.

A CUI environment without automated egress control is breached. If the STIG is not enforced in code, it is a suggestion.

Compliance may be mandated. Architectural enforcement must be absolute.

7.2 Selection Rules

1. No Federal/DoD architecture enters CUI production without passing MFCS.
2. No architecture is approved if it relies on manual STIG checklists.
3. No architecture is approved without continuous ATO (cATO) telemetry.
4. No architecture is approved if CUI can egress to non-approved boundaries.
5. No architecture is approved without meeting the DoD Zero Trust Target Level.

7.3 The Prime Directive for Federal & Sovereign Compliance

SOURCE STATEMENT

Stream the evidence, do not write the SSP. Enforce the STIG, do not check the box. Isolate the CUI, do not trust the VPC. Prove the zero trust, do not build the perimeter.

text Academic benchmarks measure documentation completeness. Applied benchmarks measure STIG automation rigor. Security benchmarks measure CUI boundary isolation. Operational benchmarks measure cATO telemetry latency.

SOURCE STATEMENT

Mandates may be bureaucratic. Sovereign enforcement must be absolute.

End of Standard MYTHOS-GRC-0001

Appendix B - Source Disposition

Source	Disposition
MYTHOS-GRC-0001_Federal_DoS_Sovereign_Compliance_Standard.md	Preserved in full; normative status corrected; evaluation anchors retained; enterprise architecture, implementation, evidence, assessment, authority, and publication controls added.
Original “Published Standard” label	Superseded by Candidate Standard status pending formal review and approval.
Original score-5 “perfect” wording	Preserved as source lineage but normalized to optimized and independently assured within defined scope.

Current authority validation

Validated 2026-09-17 / source-bound freshness record

This appendix records authority state verified for the permanent canonical edition. It does not replace the preserved normative body or imply certification, legal compliance, or implementation effectiveness.

NIST - SP 800-171 Rev. 3 - Protecting Controlled Unclassified Information in Nonfederal S...

Revision 3 / Final

Current NIST CUI security-requirement baseline; supersedes SP 800-171 Rev. 2.

https://csrc.nist.gov/pubs/sp/800/171/r3/final

DoD CIO / Acquisition.gov - Cybersecurity Maturity Model Certification (CMMC) program and...

Current implementation state / Volatile - implementation sequencing under review

Phase II requirements were suspended on 2026-07-13 while Phase I self-assessment requirements remained in place. Revalidate immediately before contractual reliance.

https://dodcio.defense.gov/CMMC/