ENGINEERING STANDARDS LIBRARY / IDENTITY AND ACCESS

Workload, Agent & Human Identity Standard

The architecture of identity has failed.

PERMANENT DOCUMENT ID

MYTHOS-ID-0001

PUBLICATION

Candidate Standard
Version 1.0.0-candidate

ASSURANCE PROFILE

Standard / Advanced
High Assurance

PERMANENT PUBLICATION IDENTITY

Workload, Agent & Human Identity Standard

DOCUMENT ID
MYTHOS-ID-0001

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-24

SCHEDULED REVIEW
2027-07-24

12	6	4	1
ATOMIC CRITERIA	ASSURANCE VIEWS	IMPLEMENTATION PROFILES	PERMANENT IDENTITY

Publication doctrine. This standard is a permanent library identity. Interpretation must preserve its recorded evidence, controlled exceptions, formal revision history, and explicit relationship to companion standards.

DOMAIN CONTROL SYSTEM

MYTHOS-ID-0001 inside the identity and access constellation

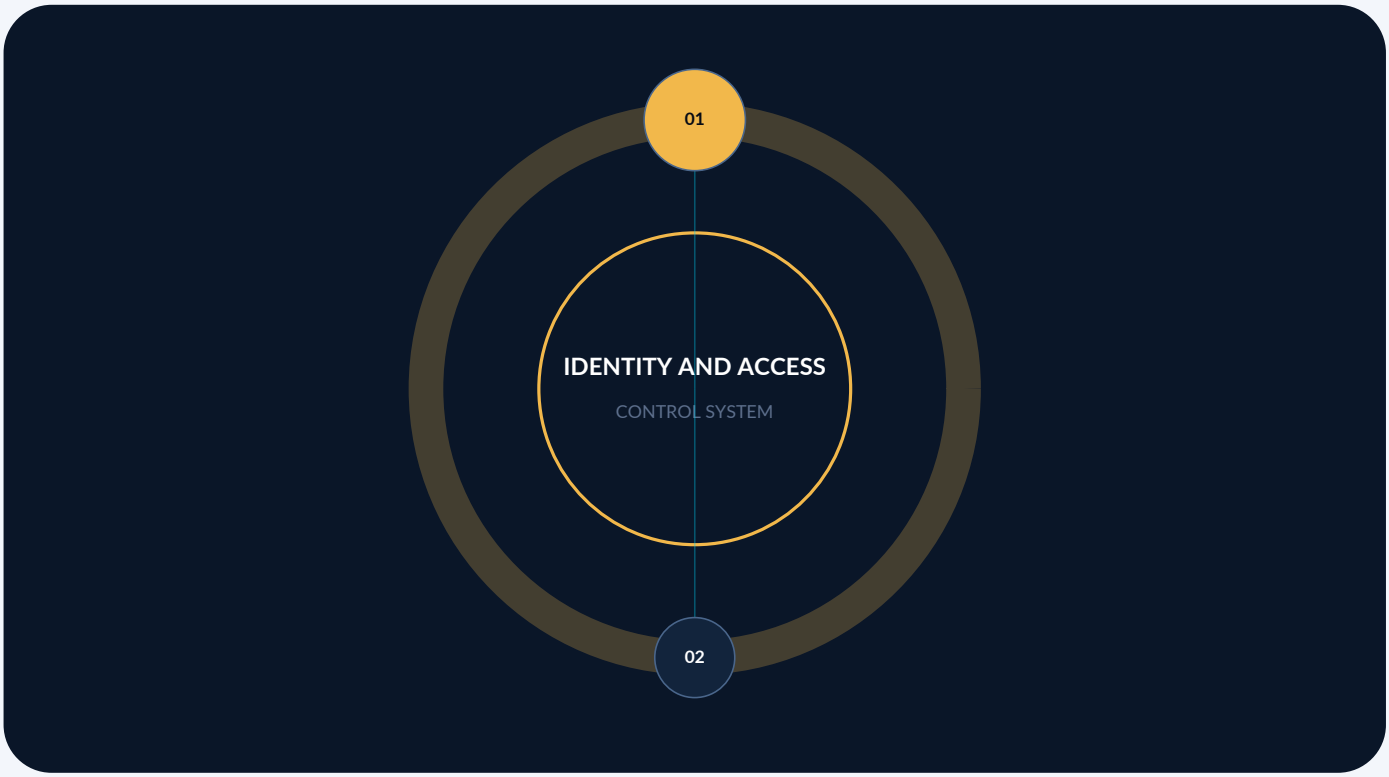

Connected assurance	Architecture, implementation, operations, evidence, and lifecycle are evaluated as one controlled system.
Specialization rule	Companion standards may add precision, but cannot silently weaken this publication.

Executive orientation

Defines identity architecture for humans, workloads, services, agents, tools, and devices with attestation, federation, lifecycle, delegation, authentication, and audit requirements.

WHY THIS STANDARD EXISTS	The architecture of identity has failed.
OPERATIONAL SCOPE	This standard covers the evaluation of: - Human identity: Passkeys (FIDO2, WebAuthn, Apple/Google/Microsoft sync) - Workload identity: SPIFFE/SPIRE, mTLS, cloud IAM instance profiles - Agent identity: Ephemeral tokens, Macaroons, capability-based delegation - Authorization protocols: OAuth 2.1, OIDC, token exchange - Privileged Access Management (PAM): Just-in-Time (JIT) access, step-up auth - Cryptographic token lifecycle: Issuance, validation, rotation, revocation
PRIMARY AUDIENCE	- Security architects and IAM engineers - Platform engineers building zero-trust infrastructure - AI engineers building agent authentication and delegation - Compliance and audit teams managing access governance - Standards bodies seeking a reference identity methodology

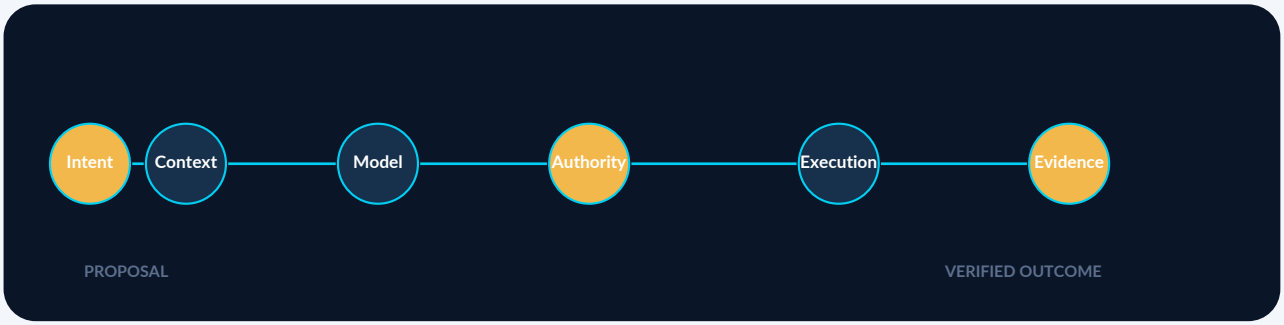

Assurance principle. Model capability, data quality, identity, authority, operational evidence, and human accountability are treated as a connected system. A high aggregate score cannot compensate for a failed mandatory safety or authority boundary.

Contents

Evaluation criterion index

This standard contains 12 atomic criteria. Each criterion is independently testable and requires retained evidence.

ID	EVALUATION CRITERION
4.1.1	Phishing Resistance
4.2.1	Machine Secret Elimination
4.3.1	Authority Scoping
4.4.1	Standing Privilege Elimination
4.5.1	Token Validity
4.6.1	Identity Observability
5.1.1	M-ID-S-Human
5.1.2	M-ID-S-Agent
6.1.1	Adversary-in-the-Middle (AiTM) Phishing
6.1.2	Agent Token Hijacking
6.1.3	Standing Privilege Abuse
6.1.4	Confused Deputy (Scope Creep)

MYTHOS-ID-0001-4.1.1

4.1.1 Phishing Resistance

Normative source requirement

Are human users authenticated via cryptographic possession rather than shared secrets?

Score	Criteria
0	Passwords only
1	Passwords + SMS/TOTP MFA (phishable)
2	Push-based MFA (phishable via MFA fatigue)
3	FIDO2/WebAuthn passkeys (phishing-resistant)
4	Device-bound passkeys requiring local biometric user verification for every auth
5	Perfect resistance: formally verified key bounds, hardware-bound nonces, zero shared secrets

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-4.2.1

4.2.1 Machine Secret Elimination

Normative source requirement

How do microservices and workloads authenticate to each other?

Score	Criteria
0	Static API keys, passwords, or cloud IAM keys in environment variables
1	Cloud-managed instance profiles (AWS IAM Roles for Service Accounts)
2	Internal PKI with long-lived certificates
3	SPIFFE/SPIRE issuing short-lived SVIDs (SPIFFE Verifiable Identity Documents)
4	Automated SVID rotation (< 5-minute TTLs); zero static secrets in the environment
5	Perfect identity: formally verified workload attestation, zero-trust mTLS mesh, cryptographic proof of workload binary signature before SVID issuance

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records. Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-4.3.1

4.3.1 Authority Scoping

Normative source requirement

How is an autonomous agent's authority constrained when acting on behalf of a user?

Score	Criteria
0	Agent is passed the user's full OAuth token or password
1	Agent uses a generic service account with broad API access
2	Agent uses OAuth token exchange (RFC 8693) to reduce scopes, but token is long-lived
3	Ephemeral agent identity: short-lived token bound to a specific task and user
4	Cryptographic attenuation: agent cannot delegate more authority than it possesses (Macaroons/Biscuits)
5	Perfect delegation: formally verified capability chains, zero scope expansion possible, cryptographic proof of user intent bound to agent action

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records. Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-4.4.1

4.4.1 Standing Privilege Elimination

Normative source requirement

Are administrative rights granted continuously or only when needed?

Score	Criteria
0	Users have persistent "admin" or "root" roles
1	Break-glass accounts exist, but manual checkout process
2	PAM solution used to rotate passwords for standing admin accounts
3	Just-in-Time (JIT) access: zero standing privilege; rights granted dynamically via approval
4	JIT access with automatic expiration and real-time anomaly detection on session behavior
5	Perfect privilege: formally verified zero-standing access, sub-minute provisioning, cryptographic attestation of PAM session boundaries

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records. Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-4.5.1

4.5.1 Token Validity

Normative source requirement

How are authorization tokens issued, validated, and revoked?

Score	Criteria
0	Long-lived JWTs with no expiration or revocation mechanism
1	Short-lived JWTs, but no revocation (must wait for expiry)
2	Refresh tokens with immediate revocation capability
3	OAuth 2.1 DPoP (Demonstrating Proof-of-Possession) binding tokens to client keys
4	Transaction-bound tokens: token contains a hash of the specific action it authorizes
5	Perfect lifecycle: formally verified token bounds, zero replay capability, cryptographic non-repudiation of token use

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-4.6.1

4.6.1 Identity Observability

Normative source requirement

Can security teams reconstruct the complete chain of identity and delegation?

Score	Criteria
0	Logs only record "user X logged in"
1	Logs record API calls, but no token context
2	Logs link actions to specific OAuth tokens
3	Full audit trail: user -> OIDC -> token -> delegation -> agent -> workload
4	Real-time anomaly detection on identity events (impossible travel, token reuse)
5	Perfect observability: formally verified identity provenance graphs, cryptographic signing of all identity events, zero blind spots in delegation chains

---

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Reproducible benchmark results, workload definitions, raw outputs, and variance records. Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Run a version-pinned workload with representative inputs, record distributions rather than a single score, repeat under failure and load, and compare reproduced results with the stated acceptance threshold.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-5.1.1

5.1.1 M-ID-S-Human

Normative source requirement

- Phishing Simulation: 100+ tests attempting to intercept credentials via reverse proxies, verifying Passkeys/FIDO2 resist the relay.
- MFA Fatigue: 50+ tests spamming push notifications, verifying the system requires transaction-bound user presence.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Architecture records, implemented configuration, test evidence, operational telemetry, exception decisions, and accountable approval.
ASSESSMENT METHOD	MATURITY SIGNAL
Inspect the architecture and configured control, execute normal and adverse scenarios, verify measurable outcomes, sample retained evidence, and confirm that exceptions are time-bound and approved.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-5.1.2

5.1.2 M-ID-S-Agent

Normative source requirement

• Scope Escalation: 100+ tests attempting to use an agent's delegated token to access unauthorized APIs, verifying the capability scoping blocks the action.

• Token Replay: 50+ tests attempting to reuse an expired or already-consumed agent token.

• Attenuation Break: 50+ tests attempting to forge a Macaroon/capability token to gain higher privileges, verifying the cryptographic signature fails.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Trace the decision path from identity and policy through execution, attempt bypass and privilege-escalation scenarios, verify revocation behavior, and inspect the resulting evidence record.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-6.1.1

6.1.1 Adversary-in-the-Middle (AiTM) Phishing

Normative source requirement

Severity: Critical Description: An attacker uses a reverse proxy (e.g., Evilginx) to steal session cookies and bypass traditional MFA.

Required Controls: 1. FIDO2/WebAuthn Passkeys. The cryptographic challenge is origin-bound, meaning a proxy cannot relay the response to the legitimate backend.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Architecture records, implemented configuration, test evidence, operational telemetry, exception decisions, and accountable approval.
ASSESSMENT METHOD	MATURITY SIGNAL
Inspect the architecture and configured control, execute normal and adverse scenarios, verify measurable outcomes, sample retained evidence, and confirm that exceptions are time-bound and approved.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-6.1.2

6.1.2 Agent Token Hijacking

Normative source requirement

Severity: Critical Description: A prompt injection or memory poisoning attack causes an agent to exfiltrate its delegated OAuth token to an attacker-controlled server.

Required Controls: 1. Tokens MUST be short-lived (sub-15 minutes). 2. Tokens MUST be bound to the specific task (transaction-bound), rendering them useless to the attacker. 3. Where possible, use DPoP so the token is bound to the agent's private key, which the attacker does not possess.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Context manifests, retrieval traces, source citations, freshness records, conflict decisions, and memory provenance.
ASSESSMENT METHOD	MATURITY SIGNAL
Trace the decision path from identity and policy through execution, attempt bypass and privilege-escalation scenarios, verify revocation behavior, and inspect the resulting evidence record.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-6.1.3

6.1.3 Standing Privilege Abuse

Normative source requirement

Severity: High Description: A compromised admin account is used to create backdoors or exfiltrate data because the account had continuous root access.

Required Controls: 1. Just-in-Time (JIT) PAM. The account has zero standing privilege and must request elevation, triggering an approval workflow and audit trail.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Policy configuration, identity or trust records, authorization decisions, revocation evidence, and adversarial test results.
ASSESSMENT METHOD	MATURITY SIGNAL
Trace the decision path from identity and policy through execution, attempt bypass and privilege-escalation scenarios, verify revocation behavior, and inspect the resulting evidence record.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

MYTHOS-ID-0001-6.1.4

6.1.4 Confused Deputy (Scope Creep)

Normative source requirement

Severity: High Description: An agent is granted an OAuth scope to "read files," but the API implicitly allows it to "delete files" because the scope was too broad.

Required Controls: 1. Granular capability-based scoping (Macaroons). 2. The token must explicitly state the exact action (`action:read`) and target (`file:123`).

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Architecture records, implemented configuration, test evidence, operational telemetry, exception decisions, and accountable approval.
ASSESSMENT METHOD	MATURITY SIGNAL
Exercise crash, retry, duplicate, reorder, partition, and restore scenarios; compare authoritative state before and after replay; confirm idempotency and recovery evidence.	0 absent | 1 ad hoc | 2 repeatable | 3 controlled | 4 measured | 5 continuously verified

Decision rule. A criterion is not satisfied by documentation alone when the claim concerns runtime behavior. Reproduced evidence, responsible ownership, known limitations, and controlled exceptions are required.

Current authority and freshness register

These sources inform interpretation but do not convert this candidate publication into certification, legal compliance, or implementation evidence. Rolling sources must be version-pinned at adoption time.

AUTHORITY	VERSION / FRESHNESS	APPLICABILITY LIMIT
NIST - Digital Identity Guidelines	SP 800-63-4 final, July 2025 Expires 2027-01-24	Federal guidance must be tailored for non-federal risk, jurisdiction, usability, privacy, and operational context.
W3C - Verifiable Credentials Data Model v2.0	W3C Recommendation, May 15 2025 Expires 2027-01-24	Data model conformance does not establish issuer trust, credential truth, or ecosystem governance.
W3C - Decentralized Identifiers (DIDs) v1.0	W3C Recommendation; DID 1.1 remains a Candidate Recommendation Expires 2027-01-24	DID method security, governance, resolution, and privacy properties vary materially.
SPIFFE - SPIFFE Specifications and Workload API	Rolling specification snapshot; SPIRE 1.15.2 documentation Expires 2026-10-24	Implementation and deployment assurance depend on attestation plugins, trust-domain design, and operational controls.

Source lineage appendix

The following text preserves the substantive source draft in a normalized public form. Where it conflicts with the permanent publication identity, candidate status, current authority register, or evidence requirements above, the controlled publication layer governs.

0. Executive Mandate

The architecture of identity has failed.

Organizations secure multi-cloud architectures and autonomous agent fleets using static API keys, shared passwords, and long-lived service accounts. When a secret is compromised, it is manually rotated. When an agent acts on behalf of a user, it inherits the user's entire scope of authority with no mechanism for attenuation. This is not security; it is accountability theater.

This standard exists because:

1. Passwords and Static Keys are Architectural Liabilities: A shared secret is a single point of failure. If a secret is stolen, the attacker becomes indistinguishable from the legitimate user. Human identity MUST be based on phishing-resistant, asymmetric cryptography (FIDO2/WebAuthn Passkeys). Machine identity MUST be based on ambient, short-lived cryptographic identity (SPIFFE).

2. Agents are Not Humans, Nor are They Service Accounts: An autonomous agent is a new class of entity. It requires an ephemeral identity that is bound to a specific task, a specific user delegation, and a strict time-to-live (TTL). Agents MUST NOT inherit blanket user credentials; they MUST operate via cryptographically attenuated delegation chains.

3. Standing Privilege is an Attack Vector: A user or agent with continuous "admin" access is a compromised asset waiting to happen. Privileged Access Management (PAM) MUST enforce Just-in-Time (JIT) access, elevating privileges only for the duration of a specific, approved task.

4. OAuth 2.0 is Insufficient Without Strict Scoping: Traditional OAuth often results in "scope creep," where a token grants access to an entire API rather than a specific action. Token issuance MUST be bound to granular capabilities, and high-risk actions MUST require transaction-bound step-up authentication.

This document establishes the Mythos Identity Standard (M-ID-S), a comprehensive, evidence-based methodology for evaluating identity architectures against zero-trust principles, cryptographic non-repudiation, and ephemeral authority.

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Human identity: Passkeys (FIDO2, WebAuthn, Apple/Google/Microsoft sync)
- Workload identity: SPIFFE/SPIRE, mTLS, cloud IAM instance profiles
- Agent identity: Ephemeral tokens, Macaroons, capability-based delegation
- Authorization protocols: OAuth 2.1, OIDC, token exchange
- Privileged Access Management (PAM): Just-in-Time (JIT) access, step-up auth
- Cryptographic token lifecycle: Issuance, validation, rotation, revocation

1.2 Out of Scope

- General zero-trust network architecture (covered in MYTHOS-NET-0004)
- General HITL workflow design (covered in MYTHOS-UX-0002)
- Post-Quantum Cryptography algorithms (covered in MYTHOS-SEC-0001)

1.3 Audience

- Security architects and IAM engineers
- Platform engineers building zero-trust infrastructure
- AI engineers building agent authentication and delegation
- Compliance and audit teams managing access governance
- Standards bodies seeking a reference identity methodology

Part II. Definitions and Taxonomy

2.1 Identity Dimensions

Dimension	Definition
Passkey	A phishing-resistant credential based on FIDO2/WebAuthn, utilizing public-key cryptography where the private key is bound to a hardware secure enclave or synced OS keychain.
SPIFFE	Secure Production Identity Framework for Everyone. A framework for issuing cryptographically verifiable identity documents (SVIDs) to workloads, eliminating the need for static secrets.
Ephemeral Agent Identity	A short-lived, cryptographically signed identity token issued to an autonomous agent, bound to a specific user delegation and task scope, expiring automatically upon completion.
Attenuation	The process of reducing the scope or authority of a token when it is delegated. A child agent cannot have more authority than the parent agent or delegating user.
Just-in-Time (JIT) PAM	An access model where users/agents have zero standing privileges and must request and be granted elevated access for a specific, time-boxed window.

2.2 The Identity Doctrine

> A static secret is a liability. A password is a shared key waiting to be leaked. An agent that holds a user's blanket token is an unconstrained delegation. If the privilege is standing, the breach is already underway.

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; passwords, static API keys, standing privileges
1	Basic MFA and service accounts, but no workload identity or agent scoping
2	Functional OAuth/OIDC, but relies on long-lived tokens or manual PAM
3	Solid production performance; Passkeys, SPIFFE, JIT PAM, token attenuation
4	Strong performance; transaction-bound step-up auth, zero-standing privilege
5	Best-in-class; sets the standard for mathematically verified, zero-trust identity

Weighting

Category	Weight	Rationale
Human Authentication (Passkeys)	15%	Passwords are the #1 initial access vector
Workload Identity (SPIFFE/mTLS)	20%	Static secrets in CI/CD and microservices cause cascading breaches
Agent Delegation & Attenuation	25%	Agents are the new perimeter; unscoped agents are catastrophic
Privileged Access (PAM/JIT)	15%	Standing admin privileges guarantee total system compromise
Token Lifecycle & Cryptography	15%	Long-lived tokens are stolen tokens
Operational Lifecycle & Auditing	10%	Identity must be observable and verifiable

Total: 100%

A system MUST score at least 3.0 in Workload Identity and Agent Delegation to be considered for production use.

Part IV. Evaluation Categories

4.1 Human Authentication (Passkeys) (Weight: 15%)

4.1.1 Phishing Resistance

Are human users authenticated via cryptographic possession rather than shared secrets?

Score	Criteria
0	Passwords only
1	Passwords + SMS/TOTP MFA (phishable)
2	Push-based MFA (phishable via MFA fatigue)
3	FIDO2/WebAuthn passkeys (phishing-resistant)
4	Device-bound passkeys requiring local biometric user verification for every auth
5	Perfect resistance: formally verified key bounds, hardware-bound nonces, zero shared secrets

4.2 Workload Identity (SPIFFE/mTLS) (Weight: 20%)

4.2.1 Machine Secret Elimination

How do microservices and workloads authenticate to each other?

Score	Criteria
0	Static API keys, passwords, or cloud IAM keys in environment variables
1	Cloud-managed instance profiles (AWS IAM Roles for Service Accounts)
2	Internal PKI with long-lived certificates
3	SPIFFE/SPIRE issuing short-lived SVIDs (SPIFFE Verifiable Identity Documents)
4	Automated SVID rotation (< 5-minute TTLs); zero static secrets in the environment
5	Perfect identity: formally verified workload attestation, zero-trust mTLS mesh, cryptographic proof of workload binary signature before SVID issuance

4.3 Agent Delegation and Attenuation (Weight: 25%)

4.3.1 Authority Scoping

How is an autonomous agent's authority constrained when acting on behalf of a user?

Score	Criteria
0	Agent is passed the user's full OAuth token or password
1	Agent uses a generic service account with broad API access
2	Agent uses OAuth token exchange (RFC 8693) to reduce scopes, but token is long-lived
3	Ephemeral agent identity: short-lived token bound to a specific task and user
4	Cryptographic attenuation: agent cannot delegate more authority than it possesses (Macaroons/Biscuits)

Score	Criteria
5	Perfect delegation: formally verified capability chains, zero scope expansion possible, cryptographic proof of user intent bound to agent action

4.4 Privileged Access (PAM/JIT) (Weight: 15%)

4.4.1 Standing Privilege Elimination

Are administrative rights granted continuously or only when needed?

Score	Criteria
0	Users have persistent "admin" or "root" roles
1	Break-glass accounts exist, but manual checkout process
2	PAM solution used to rotate passwords for standing admin accounts
3	Just-in-Time (JIT) access: zero standing privilege; rights granted dynamically via approval
4	JIT access with automatic expiration and real-time anomaly detection on session behavior
5	Perfect privilege: formally verified zero-standing access, sub-minute provisioning, cryptographic attestation of PAM session boundaries

4.5 Token Lifecycle and Cryptography (Weight: 15%)

4.5.1 Token Validity

How are authorization tokens issued, validated, and revoked?

Score	Criteria
0	Long-lived JWTs with no expiration or revocation mechanism
1	Short-lived JWTs, but no revocation (must wait for expiry)
2	Refresh tokens with immediate revocation capability
3	OAuth 2.1 DPoP (Demonstrating Proof-of-Possession) binding tokens to client keys
4	Transaction-bound tokens: token contains a hash of the specific action it authorizes
5	Perfect lifecycle: formally verified token bounds, zero replay capability, cryptographic non-repudiation of token use

4.6 Operational Lifecycle and Auditing (Weight: 10%)

4.6.1 Identity Observability

Can security teams reconstruct the complete chain of identity and delegation?

Score	Criteria
0	Logs only record "user X logged in"
1	Logs record API calls, but no token context
2	Logs link actions to specific OAuth tokens
3	Full audit trail: user -> OIDC -> token -> delegation -> agent -> workload
4	Real-time anomaly detection on identity events (impossible travel, token reuse)
5	Perfect observability: formally verified identity provenance graphs, cryptographic signing of all identity events, zero blind spots in delegation chains

Part V. The Mythos Identity Suite (M-ID-S)

Traditional IAM benchmarks measure authentication latency. The M-ID-S evaluates phishing resistance, workload secret elimination, and agent delegation bounds.

5.1 Suite Composition

5.1.1 M-ID-S-Human

- Phishing Simulation: 100+ tests attempting to intercept credentials via reverse proxies, verifying Passkeys/FIDO2 resist the relay.
- MFA Fatigue: 50+ tests spamming push notifications, verifying the system requires transaction-bound user presence.

5.1.2 M-ID-S-Agent

- Scope Escalation: 100+ tests attempting to use an agent's delegated token to access unauthorized APIs, verifying the capability scoping blocks the action.
- Token Replay: 50+ tests attempting to reuse an expired or already-consumed agent token.
- Attenuation Break: 50+ tests attempting to forge a Macaroon/capability token to gain higher privileges, verifying the cryptographic signature fails.

5.2 Execution Protocol

```
1. Deploy IAM architecture with Passkeys, SPIFFE, and Agent delegation
2. Execute complex, multi-step agentic workflows requiring JIT access
3. Run M-ID-S suite (automated phishing + token replay + escalation)
4. Score each category 0-5
5. Check for static API keys in code or long-lived admin roles (instant disqualification)
6. If all thresholds met: approve for production
```

Part VI. Security Threat Model for Identity

6.1 Threat Catalog

6.1.1 Adversary-in-the-Middle (AiTM) Phishing

Severity: Critical Description: An attacker uses a reverse proxy (e.g., Evilginx) to steal session cookies and bypass traditional MFA.

Required Controls:

1. FIDO2/WebAuthn Passkeys. The cryptographic challenge is origin-bound, meaning a proxy cannot relay the response to the legitimate backend.

6.1.2 Agent Token Hijacking

Severity: Critical Description: A prompt injection or memory poisoning attack causes an agent to exfiltrate its delegated OAuth token to an attacker-controlled server.

Required Controls:

1. Tokens MUST be short-lived (sub-15 minutes).

2. Tokens MUST be bound to the specific task (transaction-bound), rendering them useless to the attacker.

3. Where possible, use DPoP so the token is bound to the agent's private key, which the attacker does not possess.

6.1.3 Standing Privilege Abuse

Severity: High Description: A compromised admin account is used to create backdoors or exfiltrate data because the account had continuous root access.

Required Controls:

1. Just-in-Time (JIT) PAM. The account has zero standing privilege and must request elevation, triggering an approval workflow and audit trail.

6.1.4 Confused Deputy (Scope Creep)

Severity: High Description: An agent is granted an OAuth scope to "read files," but the API implicitly allows it to "delete files" because the scope was too broad.

Required Controls:

1. Granular capability-based scoping (Macaroons).

2. The token must explicitly state the exact action (`action:read`) and target (`file:123`).

Part VII. The Mythos Identity Standard

7.1 The Mythos Identity Doctrine

```
A static secret is a liability.
A password is a shared key waiting to be leaked.

An agent that holds a user's blanket token is an unconstrained delegation.
If the privilege is standing, the breach is already underway.

Humans may be unpredictable.
Identity must be absolute.
```

7.2 Selection Rules

1. No identity architecture enters production without passing M-ID-S.
2. No architecture is approved if it relies on passwords for human authentication.
3. No architecture is approved if workloads use static API keys.
4. No agent is approved if it holds a user's persistent OAuth token.
5. No architecture is approved with standing administrative privileges.

7.3 The Prime Directive for Identity Architecture

> Bind the key to the hardware, do not share the secret. > Issue the SVID to the workload, do not embed the credential. > Attenuate the agent token, do not delegate the blanket scope. > Expire the privilege, do not grant the standing role.

```
Academic benchmarks measure authentication latency.
Applied benchmarks measure workload secret elimination.
Security benchmarks measure phishing resistance.
Operational benchmarks measure delegation attenuation.

> Identity may be federated.
> Authority must be absolute.
```

End of Standard MYTHOS-ID-0001

End of controlled publication

MYTHOS-ID-0001 remains a Candidate Standard until technical review, security and privacy review, accessibility review, current-source validation, and formal approval are recorded.