ENGINEERING STANDARDS LIBRARY / LEARNING AND WORKFORCE

Adaptive Learning, Skill Graphs & Cyber Lab Standard

The architecture of technical education and capability tracking has failed.

PERMANENT DOCUMENT ID

MYTHOS-EDU-0001

PUBLICATION

Candidate Standard
Version 1.0.0-candidate

ASSURANCE PROFILE

Standard / Advanced
High Assurance

PERMANENT PUBLICATION IDENTITY

Adaptive Learning, Skill Graphs & Cyber Lab Standard

DOCUMENT ID
MYTHOS-EDU-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-EDU-0001 inside the learning and workforce constellation

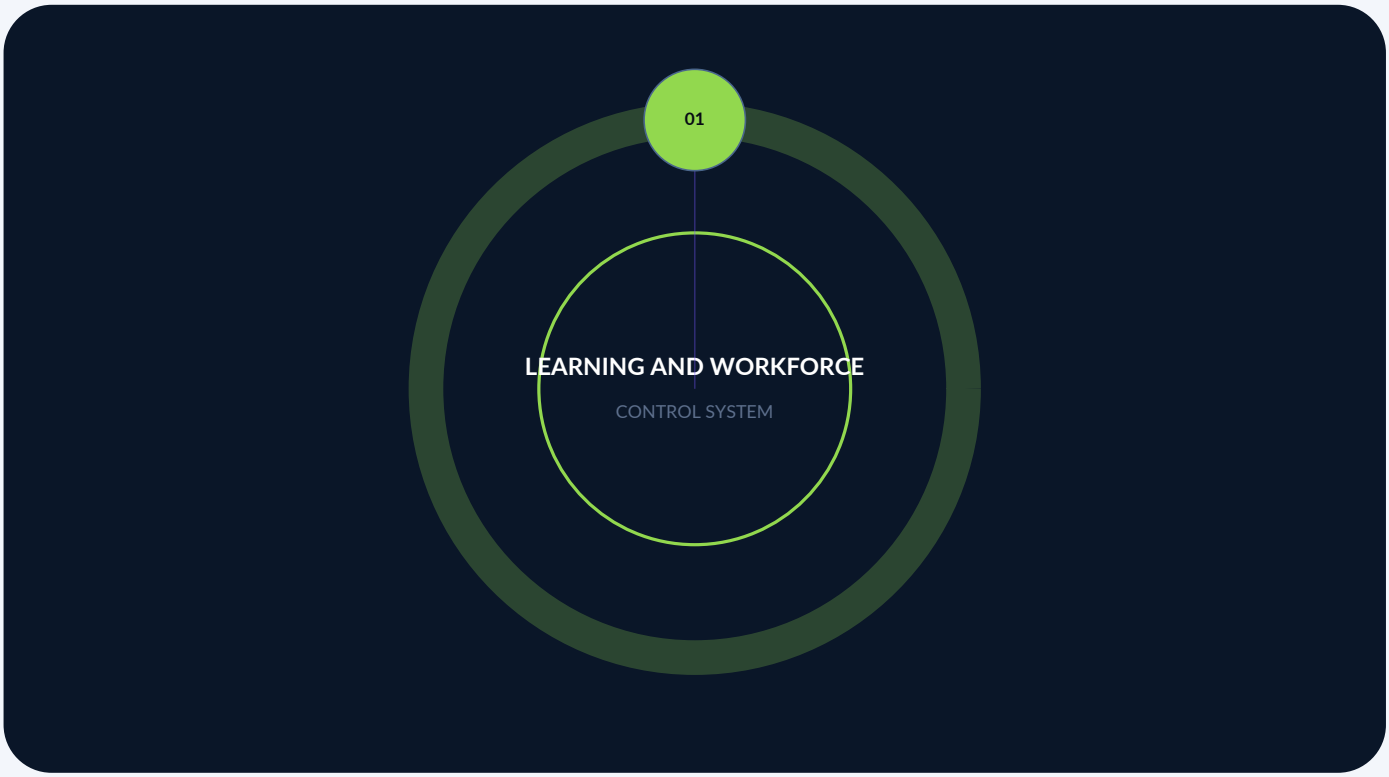

Connected assurance	Architecture, implementation, operations, evidence, and lifecycle are evaluated as one controlled system.
Specialization rule	Companion standards may add precision, but cannot silently weaken this publication.

Executive orientation

Defines adaptive learning, skill graphs, competency evidence, cyber labs, assessment, accessibility, credential portability, safety, and workforce-assurance requirements.

WHY THIS STANDARD EXISTS	The architecture of technical education and capability tracking has failed.
OPERATIONAL SCOPE	This standard covers the evaluation of: - AI-driven adaptive learning engines and dynamic curriculum generation - Semantic Skill Graphs and competency ontologies - Ephemeral cyber range architectures (containerized, microVM) - Evidence-based assessment (automated grading of live infrastructural changes) - Zero-trust isolation for training environments - Integration of operational telemetry (e.g., SIEM, Git) into skill verification
PRIMARY AUDIENCE	- Security engineering leads and capability managers - Platform engineers building cyber ranges and ephemeral lab infrastructure - AI engineers building adaptive tutoring systems - Learning and Development (L&D) directors modernizing technical training - Standards bodies seeking a reference education methodology

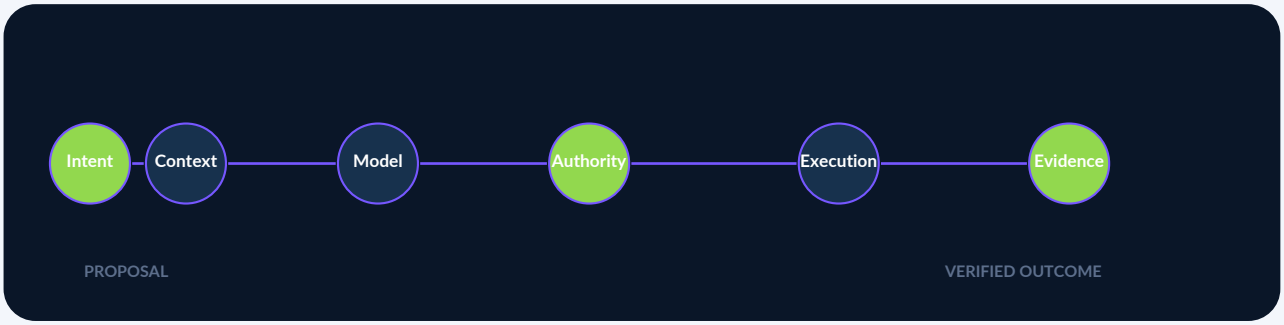

Assurance principle. Model capability, data quality, identity, authority, operational evidence, and human accountability are treated as a connected system. A high aggregate score cannot compensate for a failed mandatory safety or authority boundary.

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	Personalized Learning
4.2.1	Lab Provisioning
4.3.1	Competency Verification
4.4.1	Capability Mapping
4.5.1	Real-World Relevance
4.6.1	Boundary Enforcement
5.1.1	MEST-Lab
5.1.2	MEST-Adaptation
6.1.1	Lab-to-Production Pivot
6.1.2	Assessment Fraud (The Fake Skill)
6.1.3	Stale Curriculum (The Obsolete Defense)
6.1.4	Graph Spoofing (HR System Hack)

MYTHOS-EDU-0001-4.1.1

4.1.1 Personalized Learning

Normative source requirement

Is the curriculum static, or does it adapt to the learner?

Score	Criteria
0	One-size-fits-all video modules and PDFs
1	Branching logic based on quiz scores (e.g., if < 80%, review module)
2	Basic AI recommendations for next courses
3	Real-time adaptive engine: AI analyzes learner's lab performance and dynamically generates new scenarios targeting specific weaknesses
4	Dynamic content generation: RAG over live threat intelligence (CISA KEV) creates zero-day scenarios not pre-built by instructors
5	Perfect adaptation: formally verified cognitive load modeling, AI-driven curriculum generation, mathematical proof of optimal learning paths

---

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. Context manifests, retrieval traces, source citations, freshness records, conflict decisions, and memory provenance.
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-EDU-0001-4.2.1

4.2.1 Lab Provisioning

Normative source requirement

How are hands-on training environments deployed?

Score	Criteria
0	Persistent VMs shared among students; manual resets required
1	Dedicated VMs per student, but take 5+ minutes to boot
2	Containerized labs, but lack complex network topologies (e.g., BGP, firewalls)
3	Ephemeral microVM/container hybrid ranges spun up in < 10 seconds; completely destroyed on session end
4	Infrastructure-as-Code (IaC) based ranges mirroring production environments (e.g., a simulated the governed reference platform agent runtime)
5	Perfect architecture: formally verified ephemeral bounds, zero drift between sessions, mathematical proof of topological fidelity

---

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. Model and dataset cards, lineage, licenses, evaluation reports, release approvals, and rollback artifacts.
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-EDU-0001-4.3.1

4.3.1 Competency Verification

Normative source requirement

How is a learner's skill measured?

Score	Criteria
0	Multiple-choice quizzes at the end of a module
1	Fill-in-the-blank CLI commands
2	Instructor manually reviews the state of the lab VM
3	Automated grading: system checks the live state of the lab (e.g., API verifies RPKI is configured and BGP session is up)
4	Behavioral assessment: system evaluates the process (e.g., checking command history to ensure safe practices were used)
5	Perfect verification: formally verified assessment bounds, zero ability to cheat, cryptographic attestation of competency

---

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-EDU-0001-4.4.1

4.4.1 Capability Mapping

Normative source requirement

Are skills tracked as a flat list or a semantic graph?

Score	Criteria
0	Flat spreadsheet of skills ("Tier 1 Analyst", "Knows Python")
1	Tagging system in an LMS, but unstructured
2	Basic hierarchical taxonomy of skills
3	Semantic Skill Graph: competencies mapped as nodes, linked to specific tools, tasks, and verified evidence edges
4	Automated graph updates: passing a lab automatically updates the graph and recommends new operational tasks
5	Perfect mapping: formally verified ontology, zero gap between graph state and actual operational capability, real-time skill parity with production needs

---

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. Trace schemas, immutable event records, integrity verification, retention policy, and operator queries.
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-EDU-0001-4.5.1

4.5.1 Real-World Relevance

Normative source requirement

Does the training reflect the actual production architecture and threat landscape?

Score	Criteria
0	Generic cyber ranges (e.g., Metasploitable) disconnected from corporate reality
1	Custom scenarios, but updated annually
2	Scenarios mapped to MITRE ATT&CK, but lack production tooling
3	Training environments replicate the enterprise architecture (e.g., using internal Terraform modules and the governed reference platform agents)
4	AI automatically ingests the organization's actual SIEM alerts and generates training scenarios based on real internal incidents
5	Perfect alignment: formally verified threat parity, zero gap between training and operational reality, mathematical proof of readiness

---

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. Model and dataset cards, lineage, licenses, evaluation reports, release approvals, and rollback artifacts.
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-EDU-0001-4.6.1

4.6.1 Boundary Enforcement

Normative source requirement

Can a learner's exploit escape the cyber range?

Score	Criteria
0	Labs on the same VLAN as corporate network
1	Labs on a dedicated VLAN, but rely on software firewalls
2	Labs in an isolated cloud VPC, but allow outbound internet
3	Strict air-gapped lab networks; zero-trust microsegmentation per learner
4	Egress firewalls block all outbound traffic; malware analysis requires a local internet simulator
5	Perfect isolation: formally verified zero-trust boundaries, zero ability for lab exploits to reach corporate networks, cryptographic attestation of session containment

---

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-EDU-0001-5.1.1

5.1.1 MEST-Lab

Normative source requirement

• Provisioning Stress Test: 100+ tests spinning up 1,000 concurrent ephemeral cyber ranges, verifying they boot in < 10 seconds and are strictly isolated from one another.

• Grading Accuracy Test: 50+ tests executing complex configurations (e.g., implementing zero-trust mTLS), verifying the automated grading engine correctly detects both successes and subtle misconfigurations.

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
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-EDU-0001-5.1.2

5.1.2 MEST-Adaptation

Normative source requirement

• Dynamic Generation Test: 100+ tests feeding a novel CVE to the AI engine, verifying it dynamically generates a working lab environment and assessment criteria on the fly.

• Isolation Break Test: 50+ tests attempting to pivot from a compromised learner VM to the host infrastructure or another learner's VM, verifying the microVM/container boundary blocks the attack.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Skill graph, assessment blueprint, learner evidence, lab isolation record, accessibility results, and credential metadata.
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-EDU-0001-6.1.1

6.1.1 Lab-to-Production Pivot

Normative source requirement

Severity: Critical Description: A learner is practicing advanced exploitation techniques in a cyber range. Due to a misconfigured VLAN or lack of microsegmentation, the learner (or a malicious actor in the training network) pivots from the lab into the corporate production network, deploying ransomware.

Required Controls: 1. Labs MUST run in ephemeral microVMs (e.g., Firecracker) with strict network isolation. 2. Lab networks MUST be physically or cryptographically air-gapped from production.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Model and dataset cards, lineage, licenses, evaluation reports, release approvals, and rollback artifacts.
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-EDU-0001-6.1.2

6.1.2 Assessment Fraud (The Fake Skill)

Normative source requirement

Severity: High Description: An employee clicks "Next" through a 50-slide presentation and passes a multiple-choice quiz. Their HR profile is updated to "Certified Kubernetes Administrator." They are later assigned to a production incident and cause an outage due to lack of actual competence.

Required Controls: 1. Abolish multiple-choice quizzes for technical roles. 2. Assessment MUST be evidence-based, verifying live system state via API.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Trace schemas, immutable event records, integrity verification, retention policy, and operator queries. Skill graph, assessment blueprint, learner evidence, lab isolation record, accessibility results, and credential metadata.
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.

MYTHOS-EDU-0001-6.1.3

6.1.3 Stale Curriculum (The Obsolete Defense)

Normative source requirement

Severity: High Description: The organization's training covers defending against Log4j, but a new critical zero-day (e.g., xz-utils backdoor) emerges. The training will not be updated for 6 months, leaving the workforce unprepared.

Required Controls: 1. AI-driven dynamic curriculum generation utilizing RAG over live CISA KEV and internal threat intel.

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. Model and dataset cards, lineage, licenses, evaluation reports, release approvals, and rollback artifacts.
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-EDU-0001-6.1.4

6.1.4 Graph Spoofing (HR System Hack)

Normative source requirement

Severity: Medium Description An attacker compromises the LMS database and manually alters their Skill Graph to grant themselves "Tier 3 SOC Analyst" privileges, leading to a role escalation.

Required Controls: 1. Skill Graph updates MUST be cryptographically signed by the ephemeral lab's grading engine. 2. Manual overrides must require dual-control approval.

CONTROL INTENT	REQUIRED EVIDENCE
Create a measurable, reviewable control that can be verified independently of model or vendor claims.	Skill graph, assessment blueprint, learner evidence, lab isolation record, accessibility results, and credential metadata.
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.

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 - Artificial Intelligence Risk Management Framework (AI RMF 1.0)	NIST AI 100-1; current framework, revision underway Expires 2026-10-24	Voluntary risk-management framework; does not establish product certification or implementation effectiveness.
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.
1EdTech - Open Badges	Open Badges 3.0 final Expires 2027-01-24	Credential interoperability does not establish assessment validity, issuer quality, or workforce acceptance.
1EdTech - Comprehensive Learner Record Standard	CLR 2.0 final Expires 2027-01-24	Record portability does not establish competency, authenticity of evidence, or equitable assessment.

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 technical education and capability tracking has failed.

Organizations attempt to train engineers and security operators using static videos and multiple-choice quizzes hosted on legacy Learning Management Systems (LMS). They track competency via a flat spreadsheet of "skills" that are neither verified nor mapped to actual operational tasks. When an engineer is assigned to an incident, their actual capability rarely matches their "certified" status, leading to catastrophic failure under pressure. Furthermore, hands-on cyber training relies on brittle, persistent virtual machines that take minutes to boot and share a flat network, allowing a single learner's mistake to break the environment for the class.

This standard exists because:

1. Compliance Training is Not Competence: Passing a multiple-choice quiz on "How to Mitigate BGP Hijacking" proves the user can read; it does not prove they can configure RPKI on a live router. Systems MUST utilize evidence-based, hands-on cyber labs where assessment is derived from the successful execution of tasks in a live environment.

2. Static Curriculum is Obsolete: Threat landscapes and infrastructure paradigms change weekly. Training content MUST be dynamically generated and updated by AI (RAG over current threat intelligence) and adapted in real-time to the learner's specific cognitive gaps.

3. Skills Must Be a Formal Graph: A flat list of skills (e.g., "Knows Python") is unstructured data. Systems MUST implement a Semantic Skill Graph, mapping competencies to specific tools, architectural patterns, and verifiable operational evidence (e.g., "Proven ability to implement Pydantic AI tool schemas").

4. Cyber Labs Must Be Ephemeral and Isolated: Training environments that persist for days invite cross-contamination and drift. Labs MUST be ephemeral, spun up in seconds via containers/microVMs, strictly isolated per learner, and destroyed upon completion.

This document establishes the Mythos Education Standard (MEST), a comprehensive, evidence-based methodology for evaluating learning architectures against adaptive AI rigor, semantic skill mapping, and zero-trust cyber range resilience.

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- AI-driven adaptive learning engines and dynamic curriculum generation
- Semantic Skill Graphs and competency ontologies
- Ephemeral cyber range architectures (containerized, microVM)
- Evidence-based assessment (automated grading of live infrastructural changes)
- Zero-trust isolation for training environments
- Integration of operational telemetry (e.g., SIEM, Git) into skill verification

1.2 Out of Scope

- General HR onboarding and compliance tracking (e.g., workplace harassment training)
- General academic university curricula (unless operating sovereign labs)
- General platform engineering (covered in MYTHOS-PLT-0003)

1.3 Audience

- Security engineering leads and capability managers
- Platform engineers building cyber ranges and ephemeral lab infrastructure
- AI engineers building adaptive tutoring systems
- Learning and Development (L&D) directors modernizing technical training
- Standards bodies seeking a reference education methodology

Part II. Definitions and Taxonomy

2.1 Education Dimensions

Dimension	Definition
Adaptive Learning Engine	An AI-driven system that dynamically alters the difficulty, topic, and format of training based on real-time analysis of the learner's performance and cognitive load.
Semantic Skill Graph	A knowledge graph mapping technical competencies (nodes) to operational tasks, tools, and verified evidence (edges), replacing flat spreadsheets.
Ephemeral Cyber Range	A fully isolated, virtualized training environment (utilizing containers or microVMs) that is provisioned instantly for a specific task and destroyed immediately upon completion.
Evidence-Based Assessment	The practice of grading a learner not on quizzes, but by programmatically verifying the state of a live system (e.g., checking if a firewall rule was correctly applied via API).
Dynamic Curriculum Generation	The use of RAG (Retrieval-Augmented Generation) over current threat intelligence and architecture standards to generate training scenarios on the fly.

2.2 The Education Doctrine

> A multiple-choice quiz is a test of reading, not competence. A static PDF is a liability. A persistent training VM is a security risk. If the skill is not a graph mapped to evidence, it is a guess.

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; static LMS, multiple-choice, flat skills, persistent VMs
1	Basic hands-on labs, but manual grading and static curriculum
2	Functional cyber range, but lacks adaptive AI or semantic skill mapping
3	Solid production performance; AI adaptive learning, ephemeral ranges, evidence-based grading, semantic skill graph
4	Strong performance; AI-generated scenarios, automated evidence ingestion from production Git/SIEM
5	Best-in-class; sets the standard for mathematically verified, autonomous capability development

Weighting

Category	Weight	Rationale
Adaptive AI & Dynamic Curriculum	20%	Static content is obsolete the day it is published
Ephemeral Cyber Range Architecture	20%	Persistent labs drift, break, and invite cross-contamination
Evidence-Based Assessment	20%	Quizzes cannot verify infrastructural competency
Semantic Skill Graph Integration	15%	Flat skill matrices cannot map to complex operational needs
Operational Alignment & Threat Intel	15%	Training must reflect the actual production environment
Isolation & Security of Labs	10%	A learner's exploit must never reach the corporate network

Total: 100%

A system MUST score at least 3.0 in Ephemeral Cyber Ranges and Evidence-Based Assessment to be considered for production use.

Part IV. Evaluation Categories

4.1 Adaptive AI and Dynamic Curriculum (Weight: 20%)

4.1.1 Personalized Learning

Is the curriculum static, or does it adapt to the learner?

Score	Criteria
0	One-size-fits-all video modules and PDFs
1	Branching logic based on quiz scores (e.g., if < 80%, review module)
2	Basic AI recommendations for next courses
3	Real-time adaptive engine: AI analyzes learner's lab performance and dynamically generates new scenarios targeting specific weaknesses
4	Dynamic content generation: RAG over live threat intelligence (CISA KEV) creates zero-day scenarios not pre-built by instructors
5	Perfect adaptation: formally verified cognitive load modeling, AI-driven curriculum generation, mathematical proof of optimal learning paths

4.2 Ephemeral Cyber Range Architecture (Weight: 20%)

4.2.1 Lab Provisioning

How are hands-on training environments deployed?

Score	Criteria
0	Persistent VMs shared among students; manual resets required
1	Dedicated VMs per student, but take 5+ minutes to boot
2	Containerized labs, but lack complex network topologies (e.g., BGP, firewalls)
3	Ephemeral microVM/container hybrid ranges spun up in < 10 seconds; completely destroyed on session end
4	Infrastructure-as-Code (IaC) based ranges mirroring production environments (e.g., a simulated the governed reference platform agent runtime)
5	Perfect architecture: formally verified ephemeral bounds, zero drift between sessions, mathematical proof of topological fidelity

4.3 Evidence-Based Assessment (Weight: 20%)

4.3.1 Competency Verification

How is a learner's skill measured?

Score	Criteria
0	Multiple-choice quizzes at the end of a module
1	Fill-in-the-blank CLI commands
2	Instructor manually reviews the state of the lab VM
3	Automated grading: system checks the live state of the lab (e.g., API verifies RPKI is configured and BGP session is up)
4	Behavioral assessment: system evaluates the process (e.g., checking command history to ensure safe practices were used)

Score	Criteria
5	Perfect verification: formally verified assessment bounds, zero ability to cheat, cryptographic attestation of competency

4.4 Semantic Skill Graph Integration (Weight: 15%)

4.4.1 Capability Mapping

Are skills tracked as a flat list or a semantic graph?

Score	Criteria
0	Flat spreadsheet of skills ("Tier 1 Analyst", "Knows Python")
1	Tagging system in an LMS, but unstructured
2	Basic hierarchical taxonomy of skills
3	Semantic Skill Graph: competencies mapped as nodes, linked to specific tools, tasks, and verified evidence edges
4	Automated graph updates: passing a lab automatically updates the graph and recommends new operational tasks
5	Perfect mapping: formally verified ontology, zero gap between graph state and actual operational capability, real-time skill parity with production needs

4.5 Operational Alignment and Threat Intel (Weight: 15%)

4.5.1 Real-World Relevance

Does the training reflect the actual production architecture and threat landscape?

Score	Criteria
0	Generic cyber ranges (e.g., Metasploitable) disconnected from corporate reality
1	Custom scenarios, but updated annually
2	Scenarios mapped to MITRE ATT&CK, but lack production tooling
3	Training environments replicate the enterprise architecture (e.g., using internal Terraform modules and the governed reference platform agents)
4	AI automatically ingests the organization's actual SIEM alerts and generates training scenarios based on real internal incidents
5	Perfect alignment: formally verified threat parity, zero gap between training and operational reality, mathematical proof of readiness

4.6 Isolation and Security of Labs (Weight: 10%)

4.6.1 Boundary Enforcement

Can a learner's exploit escape the cyber range?

Score	Criteria
0	Labs on the same VLAN as corporate network
1	Labs on a dedicated VLAN, but rely on software firewalls
2	Labs in an isolated cloud VPC, but allow outbound internet
3	Strict air-gapped lab networks; zero-trust microsegmentation per learner
4	Egress firewalls block all outbound traffic; malware analysis requires a local internet simulator

Score	Criteria
5	Perfect isolation: formally verified zero-trust boundaries, zero ability for lab exploits to reach corporate networks, cryptographic attestation of session containment

Part V. The Mythos Education Suite (MEST)

Traditional LMS benchmarks measure course completion rates. The MEST evaluates ephemeral lab resilience, automated grading accuracy, and semantic graph fidelity.

5.1 Suite Composition

5.1.1 MEST-Lab

• Provisioning Stress Test: 100+ tests spinning up 1,000 concurrent ephemeral cyber ranges, verifying they boot in < 10 seconds and are strictly isolated from one another.

• Grading Accuracy Test: 50+ tests executing complex configurations (e.g., implementing zero-trust mTLS), verifying the automated grading engine correctly detects both successes and subtle misconfigurations.

5.1.2 MEST-Adaptation

• Dynamic Generation Test: 100+ tests feeding a novel CVE to the AI engine, verifying it dynamically generates a working lab environment and assessment criteria on the fly.

• Isolation Break Test: 50+ tests attempting to pivot from a compromised learner VM to the host infrastructure or another learner's VM, verifying the microVM/container boundary blocks the attack.

5.2 Execution Protocol

```
1. Deploy adaptive learning architecture with ephemeral ranges and skill graph
2. Assign a complex, multi-stage architectural task to a simulated learner
3. Run MEST suite (automated provisioning + dynamic generation + isolation tests)
4. Score each category 0-5
5. Check for multiple-choice assessments or persistent VMs (instant disqualification)
6. If all thresholds met: approve for production
```

Part VI. Security Threat Model for Education Systems

6.1 Threat Catalog

6.1.1 Lab-to-Production Pivot

Severity: Critical Description: A learner is practicing advanced exploitation techniques in a cyber range. Due to a misconfigured VLAN or lack of microsegmentation, the learner (or a malicious actor in the training network) pivots from the lab into the corporate production network, deploying ransomware.

Required Controls:

1. Labs MUST run in ephemeral microVMs (e.g., Firecracker) with strict network isolation.

2. Lab networks MUST be physically or cryptographically air-gapped from production.

6.1.2 Assessment Fraud (The Fake Skill)

Severity: High Description: An employee clicks "Next" through a 50-slide presentation and passes a multiple-choice quiz. Their HR profile is updated to "Certified Kubernetes Administrator." They are later assigned to a production incident and cause an outage due to lack of actual competence.

Required Controls:

1. Abolish multiple-choice quizzes for technical roles.

2. Assessment MUST be evidence-based, verifying live system state via API.

6.1.3 Stale Curriculum (The Obsolete Defense)

Severity: High Description: The organization's training covers defending against Log4j, but a new critical zero-day (e.g., xz-utils backdoor) emerges. The training will not be updated for 6 months, leaving the workforce unprepared.

Required Controls:

1. AI-driven dynamic curriculum generation utilizing RAG over live CISA KEV and internal threat intel.

6.1.4 Graph Spoofing (HR System Hack)

Severity: Medium Description An attacker compromises the LMS database and manually alters their Skill Graph to grant themselves "Tier 3 SOC Analyst" privileges, leading to a role escalation.

Required Controls:

1. Skill Graph updates MUST be cryptographically signed by the ephemeral lab's grading engine.

2. Manual overrides must require dual-control approval.

Part VII. The Mythos Education Standard

7.1 The Mythos Education Doctrine

```
A multiple-choice quiz is a test of reading, not competence.
A static PDF is a liability.

A persistent training VM is a security risk.
If the skill is not a graph mapped to evidence, it is a guess.

Training may be educational.
Operational readiness must be absolute.
```

7.2 Selection Rules

1. No education architecture enters production without passing MEST.

2. No architecture is approved if it relies on multiple-choice assessments for technical roles.

3. No architecture is approved if its cyber ranges are persistent and shared.

4. No architecture is approved without AI-driven dynamic curriculum generation.

5. No architecture is approved if skills are tracked as a flat list rather than a semantic graph.

7.3 The Prime Directive for Education Architecture

> Adapt the curriculum, do not static the PDF. > Destroy the VM, do not persist the state. > Verify the config, do not quiz the theory. > Graph the evidence, do not guess the competence.

```
Academic benchmarks measure course completion rates.
Applied benchmarks measure ephemeral lab isolation.
Security benchmarks measure evidence-based grading rigor.
Operational benchmarks measure semantic skill graph fidelity.

> Learning may be continuous.
> Competence must be absolute.
```

End of Standard MYTHOS-EDU-0001

End of controlled publication

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