MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Fully Homomorphic Encryption and Zero-Knowledge Proof Standard

Defines FHE and ZKP applicability, design, implementation, performance, threat models, and validation.

Response, and Assurance Evidence

PERMANENT DOCUMENT ID
MYTHOS-SEC-0010

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
11 Atomic Criteria / 7 Families

PERMANENT PUBLICATION IDENTITY

Fully Homomorphic Encryption and Zero-Knowledge Proof Standard

Defines FHE and ZKP applicability, design, implementation, performance, threat models, and validation.

DOCUMENT ID
MYTHOS-SEC-0010

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

11	7	4	1
ATOMIC CRITERIA	CAPABILITY FAMILIES	ASSESSMENT PROFILES	PERMANENT IDENTITY

Publication doctrine. This standard establishes durable requirements for its defined scope. Interpretation must preserve the permanent document identity, recorded evidence, and formal revision history.

DOMAIN CONTROL SYSTEM

MYTHOS-SEC-0010 inside the cybersecurity and network security standard constellation

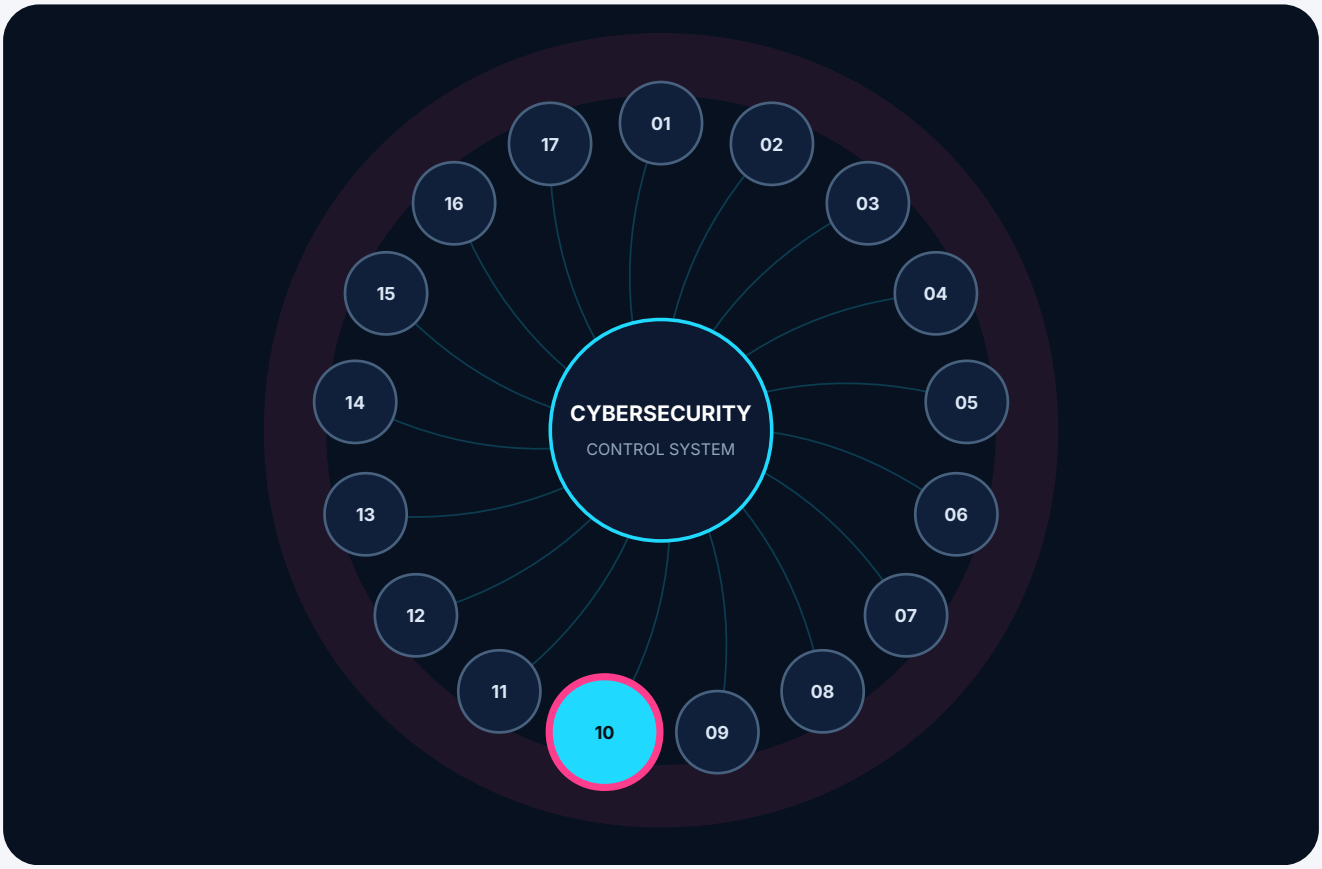

Connected assurance

Architecture, policy, implementation, operations, and evidence are evaluated as one controlled system.

Specialization rule

Companion standards may add precision, but they cannot silently weaken this publication's requirements.

CLOSED-LOOP ARCHITECTURE

Intent becomes trustworthy only when observed state returns as evidence

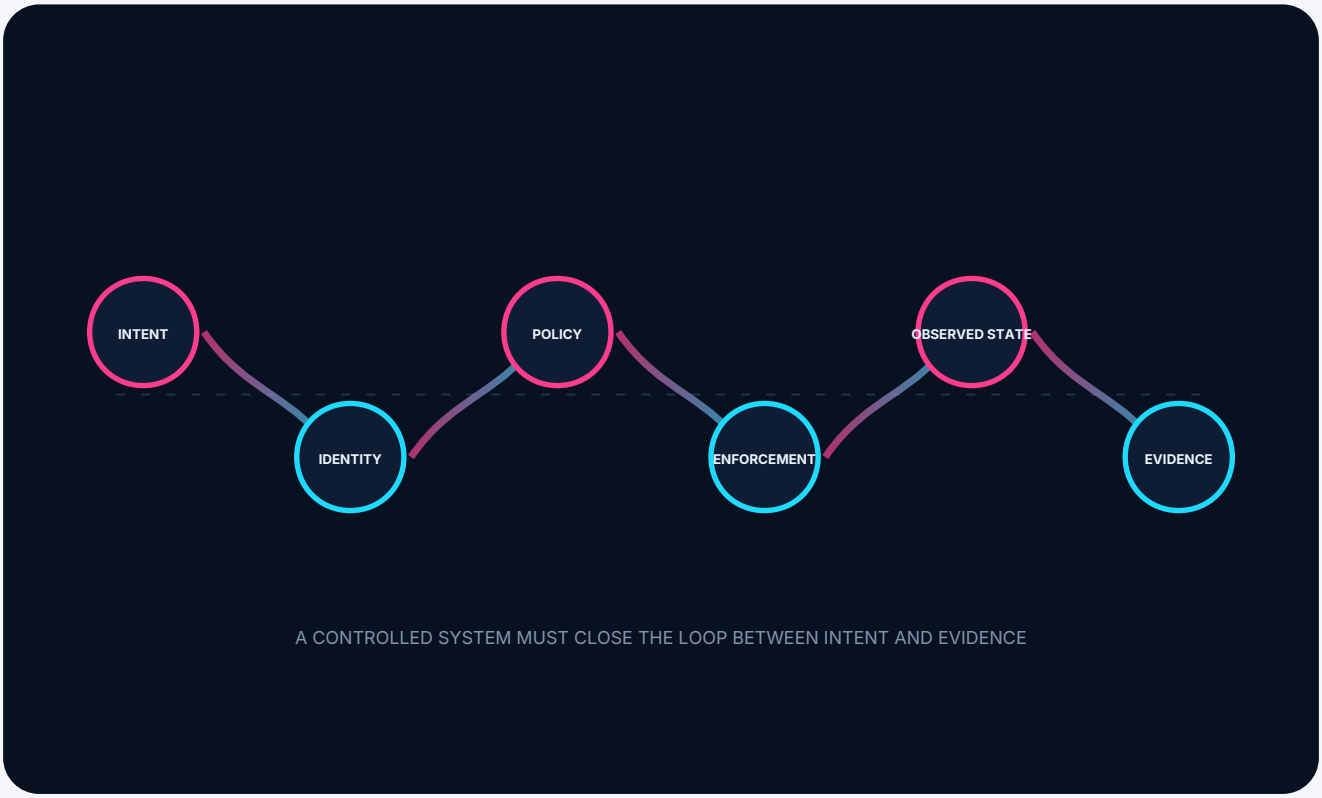

State separation

Desired, approved, deployed, observed, historical, and simulated states must remain distinguishable.

Recoverable change

Material change requires validation, authorization, monitoring, containment, and a credible recovery path.

EVALUATION TOPOLOGY

The capability families and atomic criteria of MYTHOS-SEC-0010

Capability families

Families expose the major engineering surfaces and preserve the source weighting model.

Atomic criteria

Each criterion carries a question, maturity spectrum, evidence expectations, assessment method, and operational signals.

NAVIGATION

Contents

0. Executive Mandate

Part I. Scope and Applicability

Part II. Definitions and Taxonomy

Part III. Evaluation Methodology

Evaluation system

4.1 - FHE Inference and RAG Capabilities

- 4.1.1 - Encrypted Inference
- 4.1.2 - Encrypted RAG

4.2 - ZKP Verifiable Computation and zkML

- 4.2.1 - Model Attestation (zkML)
- 4.2.2 - Agentic Auditability

4.3 - Performance and Hardware Acceleration

- 4.3.1 - FHE/ZKP Latency Overhead
- 4.3.2 - Hardware Acceleration

4.4 - Encrypted Data Architecture and Vector Search

- 4.4.1 - Encrypted Vector Database

4.5 - Secure Multi-Party Computation and Federated Learning

- 4.5.1 - Privacy-Preserving Training

4.6 - Key Management and Noise Mitigation

- 4.6.1 - FHE Key Lifecycle
- 4.6.2 - Noise Management (Bootstrapping)

4.7 - Standardization and Cryptographic Agility

- 4.7.1 - Algorithm Agility

Part V. The Mythos FHE & ZKP Suite (MFZS)

ENGINEERING DOCTRINE

0. Executive Mandate

The pursuit of AI sovereignty has reached its cryptographic limits.

Confidential Computing (TEEs) protects data in use by isolating it in hardware enclaves. But TEEs require trusting the hardware manufacturer (Intel, AMD, AWS). For truly sovereign AI - where an organization can process highly classified data on untrusted infrastructure without ever decrypting it - Fully Homomorphic Encryption (FHE) is the only architectural solution. FHE allows mathematical operations to be performed directly on ciphertext, generating an encrypted result that, when decrypted, matches the result of the same operations performed on the plaintext.

Simultaneously, as AI agents take autonomous actions in financial, legal, and infrastructure domains, their decision-making processes are opaque. Stakeholders cannot verify how an agent arrived at a conclusion without exposing the proprietary prompts, model weights, or private data used. Zero-Knowledge Proofs (ZKPs) solve this by allowing one party (the agent) to prove to another (the auditor) that a computation was executed correctly, without revealing the inputs or the model.

This standard exists because:

1. **Data in Use Must Be Encrypted:** RAG over proprietary documents currently requires decrypting them into vector embeddings. FHE allows vector similarity search to occur while the documents and the query remain fully encrypted, enabling sovereign AI deployment on adversarial infrastructure.
2. **AI Execution Must Be Verifiable:** When an AI agent proposes a critical configuration change, stakeholders need cryptographic proof that the proposal was generated by the approved model and adheres to safety constraints, without revealing the prompt. ZKPs provide this mathematical assurance.
3. **Trust Must Be Mathematical, Not Hardware-Based:** TEEs are a stopgap. Hardware fails, side-channels leak, and hypervisors can be compromised. FHE removes the hardware trust boundary entirely, replacing it with mathematical guarantees.
4. **Performance is No Longer a Total Blocker:** FHE and ZKP tooling (e.g., Concrete-ML, zkML, Zama, Ingonyama) have matured to the point where specific, high-value AI workloads (inference, verification, RAG) can be executed cryptographically with acceptable latency using hardware acceleration (GPUs, FPGAs, ASICs).

This document establishes the **Mythos FHE & ZKP Standard (MFZS)**, a comprehensive, evidence-based methodology for evaluating encrypted computation and verifiable execution in production AI environments.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the evaluation of:

- Fully Homomorphic Encryption (FHE) for AI inference and RAG
- Zero-Knowledge Proofs (ZKPs) for verifiable ML inference (zkML)
- Encrypted vector databases and similarity search
- Privacy-preserving model training and fine-tuning
- ZK attestation for agentic decision-making and auditability
- Hardware acceleration for FHE/ZKP (FPGAs, ASICs, GPU kernels)
- Secure Multi-Party Computation (MPC) for federated learning
- Cryptographic agility for post-quantum FHE/ZKP migration

1.2 Out of Scope

- Trusted Execution Environments and hardware enclaves (covered in MYTHOS-SEC-0009)
- General data encryption and Post-Quantum Cryptography (covered in MYTHOS-SEC-0001)
- General data sovereignty policies (covered in MYTHOS-DAT-0003)

1.3 Audience

- Cryptographers and security engineers designing FHE/ZKP pipelines
- Principal engineers selecting privacy-enhancing technologies for AI
- AI governance, risk, and compliance teams requiring verifiable audit trails
- Hardware architects evaluating FHE/ZKP acceleration
- Standards bodies seeking a reference FHE/ZKP methodology

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Cryptographic Dimensions

Dimension	Definition
FHE	Fully Homomorphic Encryption. The ability to compute arbitrary functions on encrypted data (ciphertext) without decrypting it, producing an encrypted result that matches the plaintext computation.
ZKP	Zero-Knowledge Proof. A cryptographic method by which one party (the prover) can prove to another (the verifier) that a statement is true, without conveying any information apart from the fact that the statement is true.
zkML	Zero-Knowledge Machine Learning. The application of ZKPs to verify that a specific AI model produced a specific output, without revealing the model weights or the input data.
Ciphertext Overflow / Noise	The accumulation of mathematical noise during FHE operations. If noise exceeds a threshold, the ciphertext becomes undecryptable. Noise management (bootstrapping) is the primary performance bottleneck in FHE.
Verifiable Computation	The ability to mathematically verify that a computation was executed correctly by an untrusted party (e.g., an AI agent or cloud provider).

2.2 The FHE/ZKP Doctrine

> Hardware fails. Math endures. A computation that cannot be verified is a hallucination. A query that exposes the plaintext is a breach. Sovereignty is not a cloud region; it is a mathematical guarantee.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

Each capability is scored from 0 to 5.

Score	Meaning
0	Fails basic task; theoretical only, no implementation
1	Prototype exists but performance is unusable for production
2	Functional but lacks hardware acceleration or RAG integration
3	Solid production performance; acceptable latency for targeted workloads
4	Strong performance; hardware-accelerated, integrated with vector DBs/zkML
5	Best-in-class; sets the standard for sovereign, verifiable computation

Weighting

Category	Weight	Rationale
FHE Inference & RAG Capabilities	20%	Encrypted RAG is the primary sovereign AI use case
ZKP Verifiable Computation & zkML	20%	Verifiable agent execution is mandatory for regulated industries
Performance & Hardware Acceleration	15%	FHE/ZKP are computationally prohibitive without acceleration
Encrypted Data Architecture & Vector Search	15%	Encrypted similarity search enables sovereign RAG
Secure Multi-Party Computation & Federated Learning	10%	Privacy-preserving training unlocks distributed data
Key Management & Noise Mitigation	10%	FHE key management and bootstrapping are operationally complex
Standardization & Cryptographic Agility	10%	FHE/ZKP algorithms are evolving; agility is mandatory

Total: 100%

A system MUST score at least 3.0 in FHE Inference and ZKP Verifiable Computation to be considered for production use in sovereign AI environments.

EVALUATION AND ASSURANCE SYSTEM

Capability claims are bounded by evidence, context, and reproducible assessment.

11 atomic criteria across 7 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

FHE Inference and RAG Capabilities

SOURCE WEIGHT 20%

4.1.1
Encrypted Inference

4.1.2
Encrypted RAG

MYTHOS-SEC-0010-EVAL-4.1.1

Encrypted Inference

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
FHE Inference and RAG Capabilities	**4.1.1**	**20%**	**Evidence-led**

Can the system execute AI model inference on fully encrypted data?

0
No FHE inference capability

1
Theoretical support but no working implementation

2
Basic FHE inference on simple models (logistic regression)

3
FHE inference on standard transformer models with high latency (>10s per token)

4
Optimized FHE inference with acceptable latency (<2s per token) for targeted workloads

5
Leading FHE inference: production-latency computation on encrypted data, verified by FHE evaluation suite

ENGINEERING INTERPRETATION

This criterion evaluates whether encrypted inference is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0010-EVAL-4.1.2

Encrypted RAG

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
FHE Inference and RAG Capabilities	4.1.2	20%	Evidence-led

Can the system perform vector similarity search over encrypted embeddings?

0
Decryption required for vector search

1
Theoretical encrypted search but no implementation

2
Encrypted exact match (not semantic similarity)

3
Encrypted approximate nearest neighbor (ANN) search with latency overhead

4
Optimized encrypted RAG pipeline with FHE-compatible embedding models

5
Leading encrypted RAG: FHE vector search, zero plaintext exposure, production latency

ENGINEERING INTERPRETATION

This criterion evaluates whether encrypted rag is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.2 / CAPABILITY FAMILY

ZKP Verifiable Computation and zkML

SOURCE WEIGHT 20%

4.2.1
Model Attestation (zkML)

4.2.2
Agentic Auditability

MYTHOS-SEC-0010-EVAL-4.2.1

Model Attestation (zkML)

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
ZKP Verifiable Computation and zkML	4.2.1	20%	Evidence-led

Can the system cryptographically prove which model generated a specific output?

0
No verification capability; trust the API provider

1
Basic hash check of model weights

2
ZKP proves model architecture but not specific weights

3
ZKP proves specific model weights produced the output (high latency)

4
Optimized zkML proof generation with acceptable latency for audit

5
Leading zkML: real-time proof generation, instant verification, zero information leakage about weights

ENGINEERING INTERPRETATION

This criterion evaluates whether model attestation (zkml) is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- cryptographic inventory, algorithm registry, key records, certificate data, and compatibility tests
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test rotation, hybrid negotiation, rollback, revocation, and algorithm-policy enforcement
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- hardcoded algorithms, unknown dependencies, missing rollback, and unverifiable keys
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- inventory coverage
- deprecated algorithm count
- rotation success rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0010-EVAL-4.2.2

Agentic Auditability

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
ZKP Verifiable Computation and zkML	4.2.2	20%	Evidence-led

Can an AI agent prove it adhered to policy constraints without revealing its private context?

0
No auditability; agent actions are opaque

1
Agent outputs text logs (unverifiable)

2
Agent outputs signed logs

3
ZKP proves agent stayed within a defined action boundary

4
ZKP proves policy compliance and safety constraint adherence

5
Leading auditability: ZKP proves compliance, intent, and model identity without revealing user data or proprietary prompts

ENGINEERING INTERPRETATION

This criterion evaluates whether agentic auditability is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.3 / CAPABILITY FAMILY

Performance and Hardware Acceleration

SOURCE WEIGHT 15%

4.3.1
FHE/ZKP Latency Overhead

4.3.2
Hardware Acceleration

MYTHOS-SEC-0010-EVAL-4.3.1

FHE/ZKP Latency Overhead

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Performance and Hardware Acceleration	4.3.1	15%	Evidence-led

What is the computational overhead compared to plaintext execution?

0
>100,000x overhead (unusable)

1
10,000-100,000x overhead

2
1,000-10,000x overhead

3
100-1,000x overhead (usable for offline/ high-value tasks)

4
10-100x overhead with GPU/FPGA acceleration

5
<10x overhead with dedicated ASIC acceleration; acceptable for interactive workloads

ENGINEERING INTERPRETATION

This criterion evaluates whether fhe/zkp latency overhead is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0010-EVAL-4.3.2

Hardware Acceleration

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Performance and Hardware Acceleration	4.3.2	15%	Evidence-led

Does the system support hardware acceleration for cryptographic operations?

0
CPU only

1
Basic GPU support

2
Optimized GPU kernels (CUDA/OpenCL)

3
FPGA support for FHE bootstrapping or ZKP proof generation

4
Dedicated ASIC emulators

5
Native ASIC integration with zero-copy memory architectures

ENGINEERING INTERPRETATION

This criterion evaluates whether hardware acceleration is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- cryptographic inventory, algorithm registry, key records, certificate data, and compatibility tests
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test rotation, hybrid negotiation, rollback, revocation, and algorithm-policy enforcement
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- hardcoded algorithms, unknown dependencies, missing rollback, and unverifiable keys
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- inventory coverage
- deprecated algorithm count
- rotation success rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.4 / CAPABILITY FAMILY

Encrypted Data Architecture and Vector Search

SOURCE WEIGHT 15%

4.4.1
Encrypted Vector Database

MYTHOS-SEC-0010-EVAL-4.4.1

Encrypted Vector Database

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Encrypted Data Architecture and Vector Search	4.4.1	15%	Evidence-led

Does the system support storing and querying encrypted embeddings?

0
Plaintext vector database only

1
Encrypted at rest, decrypted for search

2
Property-preserving encryption (deterministic, leaks order)

3
Order-preserving encryption (limited security)

4
Fully Homomorphic vector database with approximate search

5
Leading encrypted DB: FHE-native, zero leakage, production-scale indexing

ENGINEERING INTERPRETATION

This criterion evaluates whether encrypted vector database is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.5 / CAPABILITY FAMILY

Secure Multi-Party Computation and Federated Learning

SOURCE WEIGHT 10%

4.5.1
Privacy-Preserving Training

MYTHOS-SEC-0010-EVAL-4.5.1

Privacy-Preserving Training

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Secure Multi-Party Computation and Federated Learning	4.5.1	10%	Evidence-led

Can models be trained on distributed, private datasets without data movement?

0
Centralized plaintext training only

1
Basic federated learning but model updates visible

2
Secure aggregation of model updates

3
MPC-based training with moderate overhead

4
FHE-based training on encrypted gradients

5
Leading federated training: FHE/MPC, zero data leakage, mathematically proven

ENGINEERING INTERPRETATION

This criterion evaluates whether privacy-preserving training is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.6 / CAPABILITY FAMILY

Key Management and Noise Mitigation

SOURCE WEIGHT 10%

4.6.1
FHE Key Lifecycle

4.6.2
Noise Management (Bootstrapping)

MYTHOS-SEC-0010-EVAL-4.6.1

FHE Key Lifecycle

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Key Management and Noise Mitigation	**4.6.1**	**10%**	**Evidence-led**

Are FHE keys (Public, Secret, Evaluation) managed securely?

0
Keys stored in plaintext files

1
Basic KMS integration

2
HSM-backed key storage

3
TEE-protected key generation and storage

4
Distributed key generation (no single point of failure)

5
Leading lifecycle: HSM/TEE-backed, distributed, PQC-agile, and zeroization

ENGINEERING INTERPRETATION

This criterion evaluates whether fhe key lifecycle is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

MYTHOS-SEC-0010-EVAL-4.6.2

Noise Management (Bootstrapping)

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Key Management and Noise Mitigation	**4.6.2**	**10%**	**Evidence-led**

How efficiently does the system manage ciphertext noise / bootstrapping?

0
No bootstrapping; circuit depth strictly limited

1
Bootstrapping exists but takes minutes

2
Bootstrapping takes seconds

3
Optimized bootstrapping with dynamic noise estimation

4
Hardware-accelerated bootstrapping (<100ms)

5
Leading noise management: continuous, hardware-accelerated bootstrapping with zero plaintext exposure

ENGINEERING INTERPRETATION

This criterion evaluates whether noise management (bootstrapping) is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- approved architecture or policy record
- current configuration or platform export
- change and exception records
- monitoring, test, or assessment evidence

ASSESSMENT PROCEDURE

1. Examine authoritative design, policy, configuration, and lifecycle records.
2. Interview the accountable owner and operators responsible for the control.
3. Test a representative positive, negative, failure, and recovery scenario.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- The control exists only in diagrams or policy text.
- Observed state differs from approved intent without a governed exception.
- Evidence cannot establish scope, time, owner, or operating effectiveness.

OPERATIONAL MEASURES

- coverage of in-scope assets or flows
- age and closure rate of unresolved findings
- percentage of changes passing pre-deployment validation
- time to detect and contain control failure

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

4.7 / CAPABILITY FAMILY

Standardization and Cryptographic Agility

SOURCE WEIGHT 10%

4.7.1
Algorithm Agility

MYTHOS-SEC-0010-EVAL-4.7.1

Algorithm Agility

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Standardization and Cryptographic Agility	4.7.1	10%	Evidence-led

Is the system architected to swap FHE/ZKP schemes without rewriting application logic?

0
Hardcoded FHE/ZKP scheme (e.g., BGV, CKKS, Groth16)

1
Configurable but requires code changes

2
Pluggable cryptography layer

3
Dynamic scheme selection based on workload (e.g., CKKS for RAG, BGV for integer math)

4
PQC-compatible FHE/ZKP schemes

5
Leading agility: formally verified crypto boundaries, automated scheme selection, and PQC migration paths

ENGINEERING INTERPRETATION

This criterion evaluates whether algorithm agility is governed as an operating capability rather than a one-time configuration. Assessment must distinguish intended design, approved implementation, observed behavior, historical evidence, and unresolved risk.

REQUIRED EVIDENCE

- cryptographic inventory, algorithm registry, key records, certificate data, and compatibility tests
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test rotation, hybrid negotiation, rollback, revocation, and algorithm-policy enforcement
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- hardcoded algorithms, unknown dependencies, missing rollback, and unverifiable keys
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- inventory coverage
- deprecated algorithm count
- rotation success rate
- coverage of in-scope assets or flows

DECISION BOUNDARY

A maturity score is valid only for the assessed scope and evidence period. Documentation alone does not establish operating effectiveness.

LIFECYCLE AND ASSURANCE

Part V. The Mythos FHE & ZKP Suite (MFZS)

Traditional security benchmarks evaluate software configuration. The MFZS evaluates mathematical guarantees and sovereign computation performance.

5.1 Suite Composition

5.1.1 MFZS-FHE

- **Encrypted RAG:** 100+ tests querying an FHE-encrypted vector database, verifying zero plaintext leakage.
- **Inference Accuracy:** 50+ tests verifying FHE inference outputs match plaintext inference outputs (noise tolerance).

5.1.2 MFZS-ZKP

- **zkML Verification:** 50+ tests verifying ZKP proofs of model inference against known model weights.
- **Agentic Audit:** 50+ tests verifying ZKP proofs of policy compliance for autonomous actions.

5.2 Execution Protocol

PUBLICATION BOUNDARY

Candidate standard. Permanent identity. Evidence before authority.

Approval requires designated technical, security, legal, accessibility, and governance review with recorded findings and dispositions.