



Autonomous Decision, Planning, and Human Oversight Standard



Architecture & Assurance Map

A compact reasoning model for implementation, assessment, and evidence review.



OPERATING POSITION

Physical AI must remain bounded by safety, authority, uncertainty, and recoverability.

EVIDENCE PATH

sense -> interpret -> plan -> authorize -> act -> observe -> recover

READINESS SIGNALS

hazard controls, human override, perception confidence, geofence, fail-safe state

PERMANENT PUBLICATION IDENTITY

Autonomous Decision, Planning, and Human Oversight Standard			DOCUMENT ID MYTHOS-ROB-0002 VERSION 1.0.0-candidate STATUS Candidate Standard PUBLISHER Mythos Systems PUBLICATION DATE 2026-07-24 SCHEDULED REVIEW 2027-01-24 CLASSIFICATION Public
18 ATOMIC CRITERIA	6 ASSURANCE VIEWS	4 IMPLEMENTATION PROFILES	1 PERMANENT IDENTITY

Publication doctrine. This publication establishes durable requirements for its defined scope. Interpretation must preserve the permanent document identity, required evidence, controlled exceptions, source currency, and formal revision history.

EMERGING TECHNOLOGY AND ASSURANCE TOPOLOGY

Robotics and Autonomous Systems

A trustworthy autonomous decision and planning program is a governed path from research claim through baseline, experiment, safety and security challenge, transition decision, bounded operation, and preserved evidence.

CLASSIFY

Separate established practice, candidate standard, reproduced result, vendor claim, implementation hypothesis, and open research.

EXPERIMENT

Declare baseline, configuration, uncertainty, scenario coverage, stop conditions, and independent reproduction before conclusions.

CONSTRAIN

Identity, policy, safety envelopes, physical boundaries, cryptographic authority, and human intervention limit consequence.

TRANSITION

Adoption, pilot, production, rollback, supplier exit, lifecycle, and retirement decisions remain evidence-bearing and reversible.

EVIDENCE - BOUND TRANSITION FLOW



No novelty claim becomes an operational claim without reproduced evidence and an explicit authority decision

INTERPRETATION AND ASSURANCE

How to apply this publication

Normative intent

Requirements define outcomes and constraints. Implementations may vary, but evidence must demonstrate equivalent control.

Evidence over assertion

Vendor documentation and design intent are not equivalent to reproduced behavior. Record evidence grade and uncertainty.

Risk-proportional depth

Higher consequence, autonomy, sensitivity, and irreversibility require stronger isolation, review, verification, and recovery.

Controlled exception

Exceptions require an owner, rationale, compensating control, residual-risk acceptance, expiration, and reevaluation trigger.

Normalized assessment scale



A numeric score must never hide a failed mandatory gate, missing evidence, untested workload, or unacceptable residual risk.

INTERACTIVE NAVIGATION

Contents

Executive mandate	6
Scope and applicability	7
Definitions and taxonomy	8
Architecture and trust model	9
Governance and accountability	10
Normative control system	11
Evidence and assessment	16
Assurance views and metrics	17
Failure modes and adversarial scenarios	19
Implementation profiles and lifecycle	20
Adoption roadmap and conformance	22
Source authority and revision control	24

Navigation doctrine

Bookmarks and internal links are conveniences. Permanent document identity, criterion identifiers, evidence references, and revision history remain authoritative.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

0. Executive Mandate

Research value, transition discipline, and non-negotiable evidence boundaries

MANDATE

Organizations evaluating autonomous decision and planning **MUST** separate research from operational authority, establish a reproducible baseline, constrain safety and security consequence, disclose uncertainty, and require a documented transition or rejection decision.

Required outcomes

- Establish a durable, product-neutral evidence system for autonomous decision and planning.
- Prevent novelty, investment, vendor prestige, roadmap language, or laboratory success from substituting for operational proof.
- Require representative performance, adversarial, failure, recovery, safety, privacy, and lifecycle evidence.
- Keep pilots bounded by identity, authority, environment, duration, budget, stop conditions, and rollback.
- Preserve rejected alternatives, limitations, residual risk, source currency, and reassessment triggers.

Decision boundary: conformance supports a bounded engineering assurance position. It does not prove scientific consensus, regulatory approval, universal safety, quantum advantage, cryptographic perfection, autonomous competence, or commercial viability.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

1. Scope and Applicability

Research, pilot, transition, operational, and retirement boundaries

IN SCOPE

- Research, architecture, acquisition, experiment, pilot, transition, operation, and retirement of autonomous decision and planning.
- Hardware, software, models, data, protocols, facilities, people, suppliers, providers, physical processes, and evidence.
- Laboratory, simulation, testbed, shadow, pilot, canary, production, and field deployment states.
- Claims, benchmarks, uncertainty, safety, security, privacy, economics, interoperability, support, and exit.

OUT OF SCOPE

- Scientific truth inferred from a single benchmark, simulation, provider, hardware generation, or laboratory configuration.
- Safety, regulatory, export, spectrum, airspace, clinical, cryptographic, or standards approval not granted by the responsible authority.
- Roadmap forecasts, token price, patent counts, press releases, or investment treated as engineering evidence.
- Hidden production use of research components outside the declared transition and assurance boundary.

Applicability triggers

- The capability is novel, rapidly changing, weakly standardized, difficult to reproduce, or dependent on immature suppliers.
- A claim concerns safety, security, privacy, physical action, critical infrastructure, cryptographic trust, or consequential autonomy.
- The system requires specialized facilities, scarce hardware, spectrum, foundry, cryogenic, optical, robotics, or field-test dependencies.
- Failure or lock-in may create long-lived data, capital, infrastructure, workforce, or regulatory consequences.
- A research component is proposed for pilot, customer, production, public, or externally relied-upon use.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

2. Definitions and Taxonomy

Controlled language for architecture, governance, and assessment

AUTHORITY

The right to approve, deny, constrain, or execute an action. Model output is not authority.

EVIDENCE

A reproducible source, trace, measurement, test, record, signed artifact, or documented observation supporting a claim.

ASSURANCE VIEW

A defined perspective such as governance, architecture, security, operations, privacy, or human oversight.

ATOMIC CRITERION

A uniquely identified requirement that can be assessed, evidenced, excepted, and revised independently.

IMPLEMENTATION PROFILE

A risk-proportional set of required depth, evidence, and gates for a class of use.

CONTROLLED EXCEPTION

A time-bounded deviation with owner, rationale, compensating control, residual risk, expiration, and review trigger.

DETERMINISTIC MECHANISM

A component whose inputs, policy, state transition, and side effects can be constrained and verified.

SOURCE AUTHORITY

The documented basis for treating a source as reliable for a defined claim, context, jurisdiction, and time.


REASSESSMENT TRIGGER

A change, incident, drift signal, dependency revision, or time boundary requiring renewed evaluation.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

3. Architecture and Trust Model

Separation of research, authority, physical systems, runtime, and evidence



Mandatory trust boundaries

- Research systems **MUST NOT** inherit production identity, secrets, authority, data, networks, or physical control by convenience.
- Claims **MUST** retain the exact hardware, software, model, protocol, calibration, environment, workload, and observation boundary.
- Safety, cryptographic, physical, spectrum, autonomy, and human-intervention constraints **MUST** remain independently enforceable.
- Transition **MUST** have a named decision authority, acceptance criteria, monitored pilot, rollback, supplier exit, and retirement path.
- Evidence **MUST** remain reviewable when a vendor, cloud, token network, laboratory, model, robot, or quantum backend is unavailable.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

4. Governance and Accountability

Decision rights, separation of duties, and lifecycle ownership

ACCOUNTABLE OWNER

Accepts scope, risk tolerance, funding, and residual risk.

SYSTEM OWNER

Maintains architecture, inventory, operating boundaries, and change control.

SECURITY / PRIVACY

Defines threat, identity, data, isolation, monitoring, and incident requirements.

EVALUATION AUTHORITY

Designs representative tests, gates releases, and preserves reproducibility.

OPERATIONS

Maintains availability, rollback, recovery, evidence, and incident handling.

HUMAN OVERSIGHT

Reviews consequential decisions, exceptions, disputed outcomes, and harm signals.

DECISION PRINCIPLE

No single actor should be able to define the objective, grant authority, execute the consequential action, suppress evidence, and approve the result. Separation must be technical where consequence requires it.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

5. Normative Control System

Atomic criteria 01-04 of 18

MYTHOS-ROB-0002-C01

Autonomy Scope, Authority, and Decision Boundary

The organization **MUST** define, implement, verify, and maintain autonomy scope, authority, and decision boundary for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE robot and application boundary; ODD; hazard analysis; safety and cybersecurity ownership

MYTHOS-ROB-0002-C02

Operational Design Domain and Exclusion Conditions

The organization **MUST** define, implement, verify, and maintain operational design domain and exclusion conditions for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE safety function validation; stop, intervention, safe-state, and stored-energy tests

MYTHOS-ROB-0002-C03

Goal, Utility, Constraint, and Safety Specification

The organization **MUST** define, implement, verify, and maintain goal, utility, constraint, and safety specification for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE sensor calibration; dataset provenance; simulation validity; scenario and field-test coverage

MYTHOS-ROB-0002-C04

World Model, State Estimation, and Belief Uncertainty

The organization **MUST** define, implement, verify, and maintain world model, state estimation, and belief uncertainty for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE identity, command, network, remote-access, software, firmware, model, and update evidence

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

5. Normative Control System

Atomic criteria 05-08 of 18

MYTHOS-ROB-0002-C05

Planner, Policy, Search, and Optimization Governance

The organization **MUST** define, implement, verify, and maintain planner, policy, search, and optimization governance for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE normal, degraded, adversarial, lost-link, lost-localization, and recovery test results

MYTHOS-ROB-0002-C06

Action Feasibility and Safety-Envelope Enforcement

The organization **MUST** define, implement, verify, and maintain action feasibility and safety-envelope enforcement for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE operator training; human factors; takeover, rescue, maintenance, and incident procedure

MYTHOS-ROB-0002-C07

Human Role, Supervision, and Responsibility

The organization **MUST** define, implement, verify, and maintain human role, supervision, and responsibility for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE fleet telemetry, mission record, decision trace, near-miss, maintenance, and change history

MYTHOS-ROB-0002-C08

Intervention, Takeover, Stop, and Safe-State Behavior

The organization **MUST** define, implement, verify, and maintain intervention, takeover, stop, and safe-state behavior for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE acceptance, revalidation, regulatory review, residual risk, retirement, and evidence archive

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

5. Normative Control System

Atomic criteria 09-12 of 18

MYTHOS-ROB-0002-C09

Confidence, Abstention, and Out-of-Distribution Detection

The organization **MUST** define, implement, verify, and maintain confidence, abstention, and out-of-distribution detection for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE robot and application boundary; ODD; hazard analysis; safety and cybersecurity ownership

MYTHOS-ROB-0002-C10

Explainability, Decision Trace, and Contestability

The organization **MUST** define, implement, verify, and maintain explainability, decision trace, and contestability for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE safety function validation; stop, intervention, safe-state, and stored-energy tests

MYTHOS-ROB-0002-C11

Reward, Objective, Proxy, and Goodhart Risk

The organization **MUST** define, implement, verify, and maintain reward, objective, proxy, and goodhart risk for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE sensor calibration; dataset provenance; simulation validity; scenario and field-test coverage

MYTHOS-ROB-0002-C12

Adversarial, Deceptive, and Conflicting Instruction Handling

The organization **MUST** define, implement, verify, and maintain adversarial, deceptive, and conflicting instruction handling for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE identity, command, network, remote-access, software, firmware, model, and update evidence

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

5. Normative Control System

Atomic criteria 13-15 of 18

MYTHOS-ROB-0002-C13 Simulation, Formal Analysis, and Scenario Coverage

The organization **MUST** define, implement, verify, and maintain simulation, formal analysis, and scenario coverage for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE normal, degraded, adversarial, lost-link, lost-localization, and recovery test results

MYTHOS-ROB-0002-C14 Field Trial, Shadow Mode, and Progressive Release

The organization **MUST** define, implement, verify, and maintain field trial, shadow mode, and progressive release for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE operator training; human factors; takeover, rescue, maintenance, and incident procedure

MYTHOS-ROB-0002-C15 Monitoring, Drift, Learning, and Update Control

The organization **MUST** define, implement, verify, and maintain monitoring, drift, learning, and update control for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE fleet telemetry, mission record, decision trace, near-miss, maintenance, and change history

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

5. Normative Control System

Atomic criteria 16-18 of 18

MYTHOS-ROB-0002-C16**Failure Recovery, Mission Abort, and Post-Incident Review**

The organization **MUST** define, implement, verify, and maintain failure recovery, mission abort, and post-incident review for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE acceptance, revalidation, regulatory review, residual risk, retirement, and evidence archive

MYTHOS-ROB-0002-C17**Legal, Ethical, Rights, and Public-Impact Review**

The organization **MUST** define, implement, verify, and maintain legal, ethical, rights, and public-impact review for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE robot and application boundary; ODD; hazard analysis; safety and cybersecurity ownership

MYTHOS-ROB-0002-C18**Decision Evidence, Approval Gate, and Reassessment**

The organization **MUST** define, implement, verify, and maintain decision evidence, approval gate, and reassessment for the declared autonomous decision and planning boundary. The record **MUST** identify accountable ownership, evidence state, exact configuration, assumptions, affected people and systems, failure and recovery behavior, uncertainty, transition authority, and reassessment triggers. The operational design domain, hazard controls, human intervention, safe state, physical consequence, field validation, and incident evidence **MUST** be testable.

MINIMUM EVIDENCE safety function validation; stop, intervention, safe-state, and stored-energy tests

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

6. Evidence and Assessment

Reproduce, compare, challenge, observe, and trace

EXAMINE

Research papers, standards, designs, bills of material, protocols, source, models, data, calibration, configurations, contracts, and claims.

REPRODUCE

Repeat the experiment or assessment with captured versions, seeds, parameters, environment, uncertainty, and acceptance criteria.

COMPARE

Use an honest incumbent, classical, manual, or alternative baseline with equivalent outcome, data, scale, and resource accounting.

CHALLENGE

Apply adversarial input, environmental variation, dependency loss, fault injection, safety cases, security tests, and boundary conditions.

OBSERVE

Measure field or production-like quality, latency, energy, drift, physical behavior, interventions, failures, maintenance, and user impact.

TRACE

Link claim, authority, experiment, configuration, evidence, decision, pilot, change, incident, rollback, and reassessment.

Evidence grades

A

Independently reproduced, current, representative, uncertainty-bounded, integrity-protected, and operationally relevant.

B

Reproduced by the implementing organization with strong provenance, comparator, challenge, and declared limitations.

C

Partial, laboratory-only, simulation-heavy, inherited, or indirect evidence with material uncertainty and dependencies.

D

Roadmap, assertion, demonstration, benchmark excerpt, patent, token metric, press release, or vendor statement without...

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

7. Assurance Views

Six perspectives required for a complete assurance claim

GOVERNANCE

Ownership, purpose, risk tolerance, policy, exception, and decision rights.

ARCHITECTURE

Boundaries, dependencies, failure modes, state, data flow, and portability.

SECURITY

Identity, authorization, isolation, secrets, abuse resistance, and incident response.

OPERATIONS

Reliability, performance, cost, observability, change, recovery, and support.

PRIVACY / RIGHTS

Purpose limitation, minimization, consent, access, correction, deletion, and egress.

HUMAN / SOCIETAL

Usability, accessibility, disclosure, contestability, impact, and human control.

Conformance requires a defensible position across every applicable view. A strong aggregate score cannot compensate for an unowned system, a failed mandatory gate, untested recovery, or evidence below the accepted grade.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

7.2 Evaluation Metrics and Decision Gates

Measures expose utility, uncertainty, safety, physical limits, transition cost, and lifecycle burden

UTILITY

Task outcome, comparator delta, coverage, quality, accuracy, mission success, and stakeholder value.

UNCERTAINTY

Confidence, calibration, variance, sample coverage, reproducibility, sensitivity, and unresolved assumptions.

PERFORMANCE

Latency, throughput, fidelity, range, precision, bandwidth, scale, density, and availability.

SAFETY / SECURITY

Hazard exposure, unsafe action, exploitability, spoofing, physical consequence, intervention, and containment.

RESOURCE

Energy, power, cooling, spectrum, shots, qubits, memory, materials, facility, operator, and maintenance demand.

MATURITY

Standards state, supply chain, interoperability, tooling, skills, support, qualification, and lifecycle readiness.

ECONOMICS

Acquisition, operation, provider, token, facility, migration, insurance, support, and opportunity cost.

TRANSITION

Pilot escape, rollback time, incumbent fallback, portability, lock-in, retirement, and evidence completeness.

Decision gate: thresholds **MUST** be declared before assessment, include an incumbent baseline and uncertainty statement, and identify any safety, security, cryptographic, physical, reproducibility, or rollback failure that blocks transition.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

8. Failure Modes and Adversarial Scenarios

Challenge assumptions before the system is trusted with consequential work

UNSAFE AUTONOMY

The robot acts outside its operational design domain or safety envelope.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

PERCEPTION COLLAPSE

Occlusion, weather, spoofing, drift, calibration, or domain shift produces hazardous state estimates.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

HUMAN INTERVENTION FAILURE

Operators cannot understand, stop, take over, or recover the system in time.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

CYBER-PHYSICAL COMMAND ABUSE

Compromised identity, network, software, model, or tool issues unsafe motion or process actions.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

FLEET-CORRELATED FAILURE

A map, model, update, scheduler, battery, cloud, or communications dependency affects many robots at once.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

LEARNING REGRESSION

A model or skill update degrades safety, task quality, fairness, explainability, or recovery behavior.

REQUIRED: PREVENT | DETECT | CONTAIN | RECOVER | EVIDENCE

Adversarial testing MUST protect people and production systems. Use isolated environments, synthetic or minimized data, bounded authority, kill switches, explicit legal and ethical review, and documented stop conditions.

9. Implementation Profiles

Risk-proportional implementation without weakening mandatory boundaries

FOUNDATIONAL	ADVANCED	HIGH ASSURANCE
TYPICAL SCOPE Low consequence, limited autonomy, non-sensitive data, reversible outputs.	TYPICAL SCOPE Business-critical workflow, sensitive data, multiple tools or providers, moderate autonomy.	TYPICAL SCOPE Safety, security, regulated, financial, identity, or broadly consequential use.
MINIMUM DEPTH Named owner; inventory; basic evaluation; access control; logging; rollback; annual review.	MINIMUM DEPTH Formal threat model; representative evaluation; approval gates; isolated execution; tested recovery; quarterly review.	MINIMUM DEPTH Independent challenge; strongest identity; deterministic enforcement; comprehensive evidence; canary release; continuous monitoring.


Profile selection rule

Choose the highest profile triggered by consequence, autonomy, sensitivity, scale, irreversibility, external dependency, or inability to rapidly detect and recover. Tailoring may strengthen controls; it may not erase a mandatory gate without a controlled exception.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

9.2 Operational Lifecycle

Evidence-bearing operation from concept through retirement



LIFECYCLE INVARIANT

Every stage **MUST** leave sufficient evidence for the next authority to understand what was intended, what changed, what was tested, what failed, what was accepted, what remains uncertain, and how to recover or exit.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

10. Adoption Roadmap

A sequenced path from uncontrolled capability to evidence-bearing operation

0-30 DAYS**Establish ownership and truth**

Inventory systems and dependencies; classify risk; freeze unsupported claims; identify immediate privilege, data, and evidence gaps.

31-90 DAYS**Create enforceable boundaries**

Define policy; isolate execution; implement identity and approval gates; establish representative evaluation and baseline telemetry.

3-6 MONTHS**Prove recovery and operating control**

Exercise failure, rollback, revocation, incident, provider exit, and state-recovery scenarios; mature evidence packaging.

6-12 MONTHS**Scale assurance**

Automate control checks; expand workload coverage; independent challenge; portfolio metrics; controlled exception expiry; continuous reassessment.

Adoption is complete only when operating evidence demonstrates the selected profile in the actual deployment context. Documentation alone is not conformance.

MYTHOS-ROB-0002 / ROBOTICS AND AUTONOMY

10.2 Conformance and Exception Statement

What an assurance claim must contain

SYSTEM / SERVICE

Named, versioned capability and deployment boundary.

PROFILE

Foundational, Advanced, or High Assurance with trigger rationale.

SCOPE

Workloads, users, data, tools, regions, providers, and exclusions.

CRITERIA

Pass, partial, fail, not applicable, and exception status for every criterion.

EVIDENCE

Artifact references, evidence grade, date, environment, and reproducibility.

LIMITATIONS

Known failure modes, unsupported workloads, uncertainty, and residual risk.

EXCEPTIONS

Owner, rationale, compensating control, expiration, and approval.

VALIDITY

Assessment date, change boundary, review date, and reassessment triggers.

Prohibited claim: “compliant” without the named scope, profile, evidence date, limitations, exceptions, and authority accepting residual risk.

11. Source Authority and Reference Register

Primary sources informing interpretation; verify current status at assessment time

NIST CSF 2.0

Cybersecurity Framework 2.0

<https://doi.org/10.6028/NIST.CSWP.29>

Enterprise governance and cybersecurity outcome framework.

ISO 13482

Robots and Robotic Devices - Safety Requirements for Personal Care Robots

<https://www.iso.org/standard/53820.html>

Safety framework for service and personal-care robot classes.

NIST SP 800-53 Rev. 5

Security and Privacy Controls for Information Systems and Organizations

<https://doi.org/10.6028/NIST.SP.800-53r5>

Broad organizational control baseline for technology systems and services.

NIST AI RMF 1.0

Artificial Intelligence Risk Management Framework

<https://doi.org/10.6028/NIST.AI.100-1>

Risk management structure for trustworthy AI capabilities used in autonomous systems.

ISO 10218-1:2025

Robotics - Safety Requirements - Industrial Robots

<https://www.iso.org/standard/73933.html>

Current industrial robot design and inherent-safety requirements.

NIST Response Robot Test Methods

Standard Test Methods for Response Robots

<https://www.nist.gov/el/intelligent-systems-division-73500/standard-test-methods-response-robots>

Repeatable mobility, sensing, manipulation, endurance, and operator-proficiency tests.

ISO 10218-2:2025

Robotics - Safety Requirements - Industrial Robot Applications and Cells

<https://www.iso.org/standard/73934.html>

Current integration, commissioning, operation, maintenance, and disposal requirements.

Source currency rule: authorities, specifications, laws, provider services, hardware, protocols, standards status, and research evidence MUST be checked at assessment time. This publication records a 2026-07-24 source snapshot.

11.2 Revision History and Decision Register

Permanent identity, controlled change, and publication integrity

VERSION	DATE	STATE	SUMMARY
1.0.0-candidate	2026-07-24	Candidate	Permanent-publication edition; source refresh; expanded criteria, evidence, assurance, and lifecycle guidance.
Next review	2027-01-24	Scheduled	Review source currency, threat evolution, operating evidence, adoption feedback, and proposed revisions.

Publication decisions

- PERMANENT IDENTITY** The document ID remains stable across revisions; batch or production sequence metadata is not public identity.
- CHANGE CONTROL** Normative changes require rationale, impact analysis, affected-criterion mapping, and approval.
- SUPERSESION** A superseded edition remains traceable; current routes and catalogs identify the active edition.
- CORRECTION** Material errors require a recorded correction, affected-evidence analysis, and notification appropriate to impact.
- ARCHIVAL INTEGRITY** Published artifacts receive hashes, metadata, and a release manifest sufficient for independent verification.