MYTHOS SYSTEMS

PERMANENT LIBRARY PUBLICATION

ENGINEERING STANDARDS LIBRARY / CYBERSECURITY AND NETWORK SECURITY

Adversary Tactics and Attack Methodology - Network, Wireless, and Supply Chain Standard

Defines controlled techniques and assurance for networks, wireless systems, and supply chains.

Response, and Assurance Evidence

PERMANENT DOCUMENT ID
MYTHOS-SEC-0014

PUBLICATION
Candidate Standard

ASSURANCE SURFACE
6 Atomic Criteria / 6 Families

PERMANENT PUBLICATION IDENTITY

Adversary Tactics and Attack Methodology - Network, Wireless, and Supply Chain Standard

Defines controlled techniques and assurance for networks, wireless systems, and supply chains.

threat-defense incident-readiness wireless

DOCUMENT ID
MYTHOS-SEC-0014

VERSION
1.0.0-candidate

STATUS
Candidate Standard

PUBLISHER
Mythos Systems

PUBLICATION DATE
2026-07-23

SCHEDULED REVIEW
2027-01-23

CLASSIFICATION
Public

6	6	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-0014 inside the cybersecurity and network security standard constellation

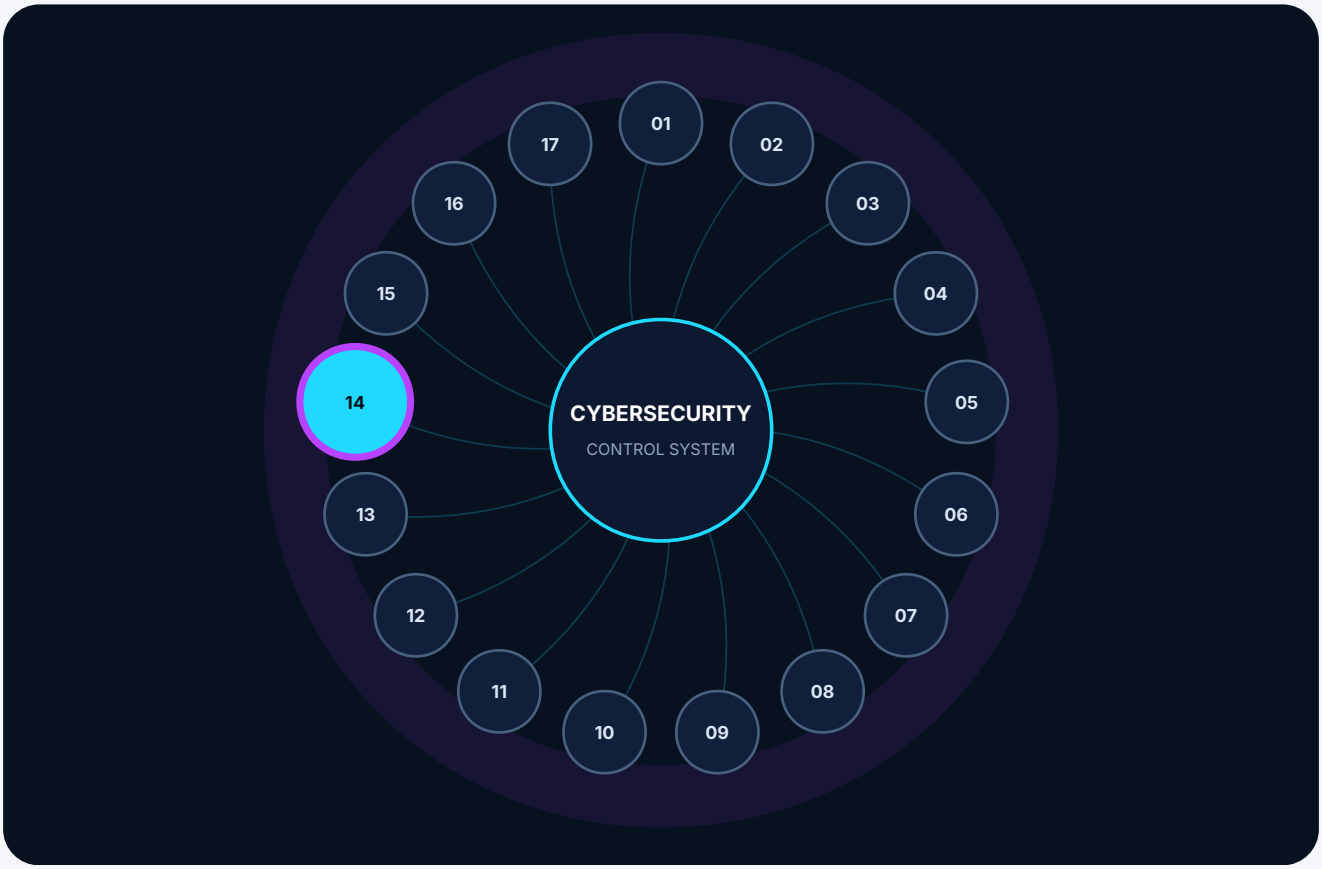

Connected assurance

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

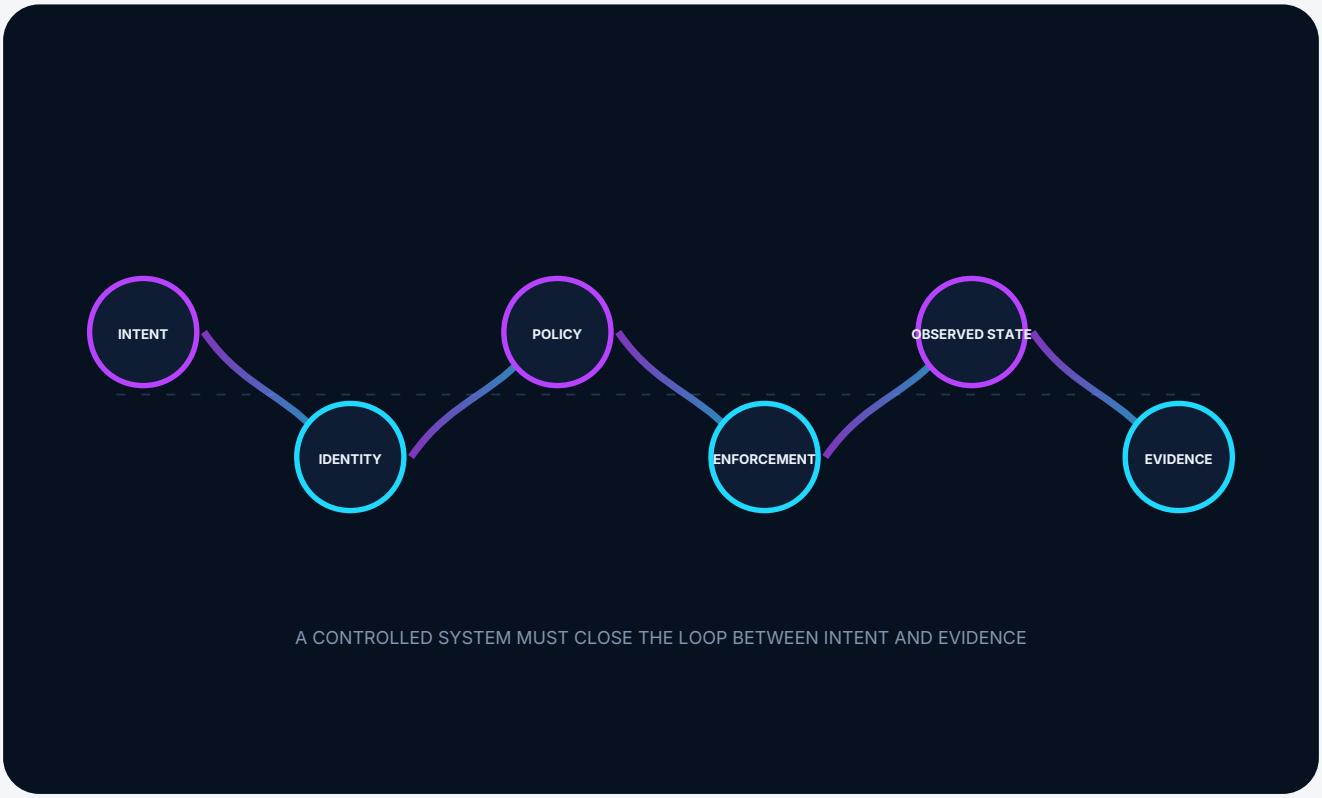

State separation

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-0014

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 - L2/L3 Routing and Spoofing

4.1.1 - Mechanism Mastery

4.2 - Core Infrastructure (DNS & DHCP)

4.2.1 - Name Resolution Integrity

4.3 - Wireless and RF Exploitation

4.3.1 - Airwave Integrity

4.4 - Dependency and Package Poisoning

4.4.1 - Supply Chain Integrity

4.5 - CI/CD and Build Pipeline Compromise

4.5.1 - Build Orchestrator Security

4.6 - Hardware and Firmware Interdiction

4.6.1 - Physical Trust Boundary

Part V. The Mythos Adversary Suite (MATS-Net)

ENGINEERING DOCTRINE

0. Executive Mandate

The practice of network and supply chain defense has failed.

Organizations deploy firewalls at the edge and trust the "magic" of the BGP routing table. They assume that if traffic arrives at their front door, the global routing infrastructure is uncompromised. They consume thousands of open-source NPM, PyPI, and Cargo packages daily, assuming that if a package is public, it is safe. They build CI/CD pipelines with elevated cloud credentials and assume the build server is a trusted, sacrosanct environment.

This standard exists because:

1. **The Internet is a Trust-Based Anarchy:** BGP and DNS were built on the assumption that all participants are honest. An attacker announcing a more specific BGP route can hijack an entire IP prefix, rerouting global traffic through their servers to intercept credentials or drop malware in transit. Defenders MUST understand route hijacking mechanics and enforce RPKI.
2. **Wireless is a Physical Attack Surface:** Wi-Fi relies on RF transmissions that anyone can intercept. WPA3 fixed many WPA2 flaws, but attackers force "downgrade attacks" (Dragonblood) or spin up "Evil Twins" to bypass captive portals. Defenders MUST understand 802.11 frame manipulation and EAP-TLS certificate pinning.
3. **The Build Server is the Crown Jewel:** Code is meaningless; the compiled artifact is what runs in production. If an attacker injects malicious code into a CI/CD runner (e.g., via a poisoned pull request or compromised dependency), the build process itself becomes the malware distribution vector (e.g., SolarWinds, `xz-utils`).
4. **Public Registries are Hostile Territory:** "Dependency Confusion" attacks exploit package managers that look up names in public registries before private ones. Defenders MUST treat all third-party code as untrusted, requiring hermetic builds and SLSA Level 3+ provenance.

This document establishes the **Mythos Adversary Tactics Standard (MATS-Net)**, a comprehensive, transparent methodology for evaluating an organization's defensive posture against network interception, RF manipulation, and supply chain poisoning.

ENGINEERING DOCTRINE

Part I. Scope and Applicability

1.1 In Scope

This standard covers the granular mechanics of:

- Network Attacks: BGP Hijacking, ARP Spoofing, DNS Cache Poisoning, DHCP Starvation
- Wireless Attacks: WPA3 Downgrade (Dragonblood), Evil Twin / Rogue AP, Deauthentication, KRACK
- Supply Chain Attacks: Dependency Confusion, Typosquatting, Compromised CI/CD Runners, Build System Backdoors
- Firmware & Hardware: Malicious USB interdiction, Firmware rootkits

1.2 Out of Scope

- Web and client-side application attacks (covered in MYTHOS-SEC-0013)
- Rules of engagement for penetration testing (covered in MYTHOS-SEC-0015)

1.3 Audience

- Network engineers and architects designing global routing and wireless fleets
- DevSecOps and platform engineers managing CI/CD and artifact registries
- Security analysts and threat hunters monitoring for lateral movement
- Standards bodies seeking a reference infrastructure attack methodology

---

ENGINEERING DOCTRINE

Part II. Definitions and Taxonomy

2.1 Adversary Dimensions

Dimension	Definition
BGP Hijacking	The act of falsely announcing IP prefixes to the global BGP routing table to intercept or blackhole traffic intended for another network.
RPKI (Resource Public Key Infrastructure)	A cryptographic framework that binds IP address prefixes to their rightful ASN, allowing routers to drop invalid BGP routes.
Evil Twin	A fraudulent wireless access point that imitates a legitimate one, often using the same SSID and BSSID, to capture credentials or perform MITM attacks.
Dependency Confusion	An attack where a malicious package with the same name as an internal private package is published to a public registry, exploiting misconfigured package resolvers to inject malicious code into the build.
Hermetic Build	A build process that is completely isolated from the internet and external, untrusted dependencies, fetching only from verified, internal, cached mirrors.

2.2 The Adversary Doctrine

> A BGP route is a claim, not a fact. A public package is a stranger. A build server with internet access is a compromised server. If the routing table cannot be cryptographically verified, the traffic is intercepted.

ENGINEERING DOCTRINE

Part III. Evaluation Methodology

3.1 Scoring Model

This standard evaluates the *defensive architecture* against specific attack mechanics. Each capability is scored from 0 to 5.

Score	Meaning
0	No defense; engineers unaware of mechanics; vulnerable by default
1	Basic awareness, but relies on "trust" rather than cryptographic verification
2	Functional defenses, but vulnerable to advanced bypasses (e.g., RPKI enabled but not Route Origin Validation enforced)
3	Solid production defense; RPKI strict, 802.1X enforced, SLSA L3 provenance
4	Strong defense; hermetic builds, client certificate pinning, zero public registry access
5	Best-in-class; mathematically verified routing boundaries, formally verified build pipelines

Weighting

Category	Weight	Rationale
L2/L3 Routing & Spoofing (BGP/ARP)	20%	The foundation of network trust; hijacks cause global breaches
Core Infrastructure (DNS/DHCP)	15%	DNS poisoning bypasses all application security
Wireless & RF Exploitation	15%	Wi-Fi is a physical boundary; Evil Twins capture credentials
Dependency & Package Poisoning	20%	The #1 vector for initial codebase compromise
CI/CD & Build Pipeline Compromise	20%	The build server is the root of trust; if compromised, all is lost
Hardware & Firmware Interdiction	10%	Physical access bypasses all network controls

Total: 100%

A system MUST score at least 3.0 in L2/L3 Routing and CI/CD Pipeline Security to be considered for production.

EVALUATION AND ASSURANCE SYSTEM

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

6 atomic criteria across 6 capability families define the evaluation surface of this standard.

4.1 / CAPABILITY FAMILY

L2/L3 Routing and Spoofing

SOURCE WEIGHT 20%

4.1.1
Mechanism Mastery

MYTHOS-SEC-0014-EVAL-4.1.1

Mechanism Mastery

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
L2/L3 Routing and Spoofing	4.1.1	20%	Evidence-led

Does the architecture definitively prevent global and local route manipulation?

Attack Granularity Example (BGP IP Prefix Hijack): Victim ASN owns `203.0.113.0/24`. Attacker ASN announces `203.0.113.0/25` to the global table. Because BGP prefers the most specific route, global traffic destined for the victim is routed to the attacker. The attacker intercepts the traffic, harvests cookies/credentials, and forwards the traffic back to the victim to avoid detection. *Defense: RPKI creates Route Origin Authorizations (ROAs) stating only the Victim ASN can announce `/24`. Routers drop the attacker's `/25`.*

0
Unfiltered BGP announcements; no ARP inspection on switches

1
BGP MD5 passwords; basic port security

2
RPKI signed prefixes, but Route Origin Validation (ROV) not enforced on inbound routes

3
Strict RPKI/ROV enforcement dropping invalid routes; Dynamic ARP Inspection (DAI) enabled

4
BGPsec enabled cryptographically validating the AS_PATH; strict L2 isolation

5
Leading isolation: formally verified routing boundaries, zero ability to announce unauthorized prefixes

ENGINEERING INTERPRETATION

This criterion evaluates whether mechanism mastery 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

- routing policy, RIB/FIB snapshots, adjacency state, and path-analysis output
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. simulate route withdrawal, policy conflict, convergence, and rollback
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- route leaks, asymmetric paths, stale advertisements, and unbounded convergence
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- route-policy compliance
- convergence time
- unexpected path count
- 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.

4.2 / CAPABILITY FAMILY

Core Infrastructure (DNS & DHCP)

SOURCE WEIGHT 15%

4.2.1
Name Resolution Integrity

MYTHOS-SEC-0014-EVAL-4.2.1

Name Resolution Integrity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Core Infrastructure (DNS & DHCP)	4.2.1	15%	Evidence-led

Can an attacker redirect traffic via name resolution manipulation?

Attack Granularity Example (DNS Cache Poisoning / Kaminsky): Attacker queries the victim's recursive resolver for a random subdomain (e.g., `random123.victim.com`). Simultaneously, they flood the resolver with spoofed UDP responses claiming to be the authoritative server, providing a malicious IP for `victim.com`. If the resolver accepts the spoofed response, the cache is poisoned. *Defense: DNSSEC cryptographically signs responses, making spoofed responses unverifiable.*

0
Open recursive DNS resolvers; unauthenticated DHCP

1
Basic DNSSEC enabled on zones, but not validated by resolvers

2
DNSSEC validated end-to-end; DHCP snooping enabled on access switches

3
DoH/DoT (DNS over HTTPS/TLS) enforced for all clients; strict DHCP snooping with port-level MAC limits

4
Zero-trust DNS: all internal queries routed through encrypted, authenticated enterprise resolvers

5
Leading integrity: formally verified DNS resolution paths, zero ability to spoof responses

ENGINEERING INTERPRETATION

This criterion evaluates whether name resolution integrity 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

Wireless and RF Exploitation

SOURCE WEIGHT 15%

4.3.1
Airwave Integrity

MYTHOS-SEC-0014-EVAL-4.3.1

Airwave Integrity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Wireless and RF Exploitation	**4.3.1**	**15%**	**Evidence-led**

Are wireless clients protected from RF manipulation and rogue access points?

Attack Granularity Example (WPA3 Downgrade / Evil Twin): Attacker spins up a Rogue AP with the same SSID as the victim, but broadcasting WPA2 only. Client devices configured for "WPA3/WPA2 Transition Mode" automatically downgrade to WPA2, allowing the attacker to capture the WPA2 handshake and perform offline dictionary attacks. *Defense: Enforce WPA3-Only mode on endpoints and use EAP-TLS so the client validates the AP's certificate.*

0
Open Wi-Fi or WPA2-PSK; no management frame protection

1
WPA3 enabled, but supports WPA2 transition mode (vulnerable to downgrade)

2
WPA3-Only enforced; 802.11w (PMF) enabled

3
WPA3-Enterprise with EAP-TLS (client certificates); Evil Twin detection via WIDS

4
Strict certificate pinning on client devices preventing Evil Twin MITM; RF geofencing

5
Leading integrity: formally verified cryptographic handshake, zero ability to downgrade or spoof APs

ENGINEERING INTERPRETATION

This criterion evaluates whether airwave integrity 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

- RF survey, channel and power plan, authentication logs, and client telemetry
- approved architecture or policy record
- current configuration or platform export

ASSESSMENT PROCEDURE

1. test roaming, capacity, interference, authentication, and degraded-band behavior
2. Examine authoritative design, policy, configuration, and lifecycle records.

EXAMINE - INTERVIEW - TEST

COMMON FAILURE PATTERNS

- coverage holes, sticky clients, co-channel interference, and insecure fallback
- The control exists only in diagrams or policy text.

OPERATIONAL MEASURES

- coverage compliance
- authentication success
- roaming interruption
- airtime utilization
- 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

Dependency and Package Poisoning

SOURCE WEIGHT 20%

4.4.1
Supply Chain Integrity

MYTHOS-SEC-0014-EVAL-4.4.1

Supply Chain Integrity

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Dependency and Package Poisoning	4.4.1	20%	Evidence-led

Can an attacker inject malicious code via public package registries?

Attack Granularity Example (Dependency Confusion): Company uses an internal package named `mycompany-auth`. Their package resolver checks the private registry first, then falls back to `npmjs.com`. Attacker publishes a malicious public package named `mycompany-auth` with version `99.0.0`. When the CI/CD pipeline builds the app, the resolver sees `99.0.0` is higher than the internal version, fetches the malicious public package, and injects a backdoor into the build. *Defense: Scope internal namespaces so they NEVER fall back to public registries.*

0
Build pipelines fetch dependencies directly from public internet (npm, PyPI)

1
Basic lockfiles used, but names not verified against internal registries

2
Private registry used, but falls back to public if package not found (vulnerable to Confusion)

3
Strict scoped registries: internal names *only* resolve internally; public fallback disabled for internal namespaces

4
Hermetic builds: all dependencies fetched from internal, SBOM-verified mirrors; no internet access during build

5
Leading integrity: formally verified dependency graphs, zero untrusted external resolution paths

ENGINEERING INTERPRETATION

This criterion evaluates whether supply chain integrity 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

CI/CD and Build Pipeline Compromise

SOURCE WEIGHT 20%

4.5.1
Build Orchestrator Security

MYTHOS-SEC-0014-EVAL-4.5.1

Build Orchestrator Security

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
CI/CD and Build Pipeline Compromise	**4.5.1**	**20%**	**Evidence-led**

Is the build server treated as a highly privileged, untrusted environment?

Attack Granularity Example (The `xz-utils` Backdoor Methodology): An attacker social-engineers their way to becoming a maintainer of a widely used open-source library (`xz`). They inject a malicious build script (`build-to-host.m4`) that only executes *during the build process* on specific Linux distributions. The script modifies the `sshd` binary at compile time, embedding a backdoor. The source code on GitHub looks clean, but the compiled artifact is compromised. *Defense: Hermetic builds, reproducible builds (compiling twice and comparing hashes), and strict isolation of the build environment.*

0
Shared runners used in public cloud CI/CD; long-lived cloud credentials stored in CI variables

1
Private runners, but run as root with internet access

2
Ephemeral runners spun up per job; secrets injected via OIDC federation

3
SLSA Level 3 provenance: builds run in isolated, hermetic environments; artifacts cryptographically signed

4
SLSA Level 4: reproducible builds verified by two independent platforms; zero untrusted PRs can execute scripts on runner

5
Leading security: formally verified build isolation, zero ability to tamper with artifacts post-commit

ENGINEERING INTERPRETATION

This criterion evaluates whether build orchestrator security 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

Hardware and Firmware Interdiction

SOURCE WEIGHT 10%

4.6.1
Physical Trust Boundary

MYTHOS-SEC-0014-EVAL-4.6.1

Physical Trust Boundary

FAMILY	SOURCE REFERENCE	WEIGHT CONTEXT	ASSESSMENT POSTURE
Hardware and Firmware Interdiction	4.6.1	10%	Evidence-led

Are devices protected from physical supply chain tampering?

Attack Granularity Example (Evil Maid / Firmware Rootkit): An attacker with brief physical access to a server plugs in a malicious PCIe card or USB device that flashes a modified SPI chip containing the BIOS/UEFI. The modified firmware boots a hidden hypervisor before the OS loads, making the malware invisible to the OS. *Defense: TPM 2.0 measured boot cryptographically hashes the firmware chain; if the hash changes, the server refuses to boot or decrypts its drives.*

0
Devices shipped directly from vendor to rack without verification

1
Basic asset tagging and inventory

2
Firmware hashes verified against vendor manifests before deployment

3
Secure Boot, measured boot, and TPM 2.0 attestation enforced on all servers

4
Hardware Bill of Materials (HBOM) verified; chips audited for interdiction

5
Leading hardware trust: formally verified silicon roots of trust, zero ability to inject persistent firmware

ENGINEERING INTERPRETATION

This criterion evaluates whether physical trust boundary 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.

LIFECYCLE AND ASSURANCE

Part V. The Mythos Adversary Suite (MATS-Net)

Traditional security benchmarks measure firewall throughput. The MATS-Net suite executes granular protocol and pipeline injections to verify architectural defense.

5.1 Suite Composition

5.1.1 MATS-Routing

- **BGP Injection Simulation:** 100+ tests announcing unauthorized prefixes to edge routers to verify RPKI drops them.
- **ARP Poisoning:** 50+ tests executing `arpspoof` on access VLANs to verify Dynamic ARP Inspection blocks the MAC-IP mismatch.

5.1.2 MATS-SupplyChain

- **Confusion Injection:** 100+ tests attempting to pull public packages with internal names to verify the resolver blocks public fallback.
- **PR Poisoning:** 50+ tests submitting pull requests with `post-install` scripts to verify the CI/CD runner refuses to execute them without sandboxing.

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.