



CANONICAL LIVING OVERVIEW GUIDE

WebAssembly, WASI, Edge Sandboxing, and Portable Compute

A modern WebAssembly guide covering core modules, validation, linear memory, the Component Model, WIT, canonical ABI, WASI 0.2 and 0.3, native async, capability security, Wasmtime, language tooling, edge deployment, performance, and supply-chain assurance.

PERMANENT PUBLICATION IDENTITY

WebAssembly, WASI, Edge Sandboxing, and Portable Compute

Document ID
MYTHOS-FG-FRONTIER-0007

Version
1.0.0-candidate

Status
Candidate Focused Field Guide

Review date
2026-10-26

Publication position

This is a living field guide. It is designed for frequent revision while retaining a stable permanent identity. Candidate status means the content is ready for structured use and review but is not represented as an external certification or authority publication.

LEVEL L1-L4

Foundational through frontier

CLASS FIELD GUIDE

Practical and educational

CADENCE QUARTERLY

Interim updates as needed

EVIDENCE SOURCE-BOUND

Limitations remain visible

Reader orientation

Dimension	Engineering position
Primary domain	MYTHOS-EMT - Emerging Technology & Sovereign Infrastructure
Secondary domain	MYTHOS-SWE; MYTHOS-CLD; MYTHOS-SEC
Audience	Software, platform, edge, plugin, serverless, security, language-runtime, and agent-capability engineers.
Prerequisites	Systems programming, compilers, operating systems, APIs, distributed systems, security, and Rust or another supported language.
Publisher / owner	Mythos Systems / Aaron Stovall
Public classification	Public technical guidance

Expected outcomes

- Explain the WebAssembly execution and validation model precisely.
- Build typed components using WIT and compose them across language boundaries.
- Use WASI and host capabilities without recreating ambient operating-system authority.
- Deploy sandboxed workloads to browsers, servers, edge sites, plugins, and agent runtimes.
- Evaluate portability, performance, isolation, toolchain maturity, and lifecycle risk.

How to use this guide

Read the guide as a progressive system rather than a disconnected encyclopedia. Foundational material establishes the mental model; practitioner sections convert it into repeatable implementation and operations; advanced sections expose architecture, scale, security, and failure behavior; frontier sections describe technologies whose evidence or maturity is still evolving.

Suggested reading path

New practitioners should follow the chapters in order and complete the laboratories. Experienced engineers may begin with the architecture, failure, validation, or frontier sections, then use the related-document map to deepen a specific topic.

Learning and assurance model

Level	Reader outcome
L1 - Foundational	Terminology, first principles, component roles, packet or state flow, and simple working examples.
L2 - Practitioner	Implementation, configuration, automation, testing, troubleshooting, and operational evidence.
L3 - Advanced	Architecture, scale, concurrency, resilience, threat models, performance, and complex migrations.
L4 - Frontier	Emerging systems, research directions, technical uncertainty, readiness gates, and adoption signals.

Contents

How to use this guide	3
Contents	3
Chapter 01 - WebAssembly core execution model	5
Chapter 01 laboratory and review	10
Chapter 02 - Component Model, WIT, and canonical ABI	11
Chapter 02 laboratory and review	16
Chapter 03 - WASI 0.2, WASI 0.3, and native async	17
Chapter 03 laboratory and review	22
Chapter 04 - Capability security and sandbox design	23
Chapter 04 laboratory and review	28
Chapter 05 - Language toolchains and developer workflow	29
Chapter 05 laboratory and review	34

Chapter 06 - Edge, serverless, plugin, and agent architectures	35
Chapter 06 laboratory and review	40
Chapter 07 - Performance, networking, and observability	41
Chapter 07 laboratory and review	46
Chapter 08 - Supply chain, lifecycle, and adoption	47
Chapter 08 laboratory and review	52
Integrated learning roadmap	53
Source authority register	54
Limitations, freshness, and revision control	55

CHAPTER 01

WebAssembly core execution model

Understand what a validated module can and cannot do by itself.

Chapter operating position

Understand what a validated module can and cannot do by itself.



1.1 Modules, types, and validation

Modules, types, and validation must be engineered as an observable mechanism rather than a frontier label or speculative promise. Define functions, tables, memories, globals, imports, exports, instructions, and static validation. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for modules, types, and validation.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating modules, types, and validation as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for modules, types, and validation.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

1.2 Linear memory and references

Linear memory and references must be engineered as an observable mechanism rather than a frontier label or speculative promise. Manage pages, loads, stores, bounds, memory growth, tables, references, and host interaction. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for linear memory and references.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating linear memory and references as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for linear memory and references.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.



1.3 Execution and traps

Execution and traps must be engineered as an observable mechanism rather than a frontier label or speculative promise. Handle stacks, control flow, calls, numeric semantics, traps, determinism limits, and embedding. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for execution and traps.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating execution and traps as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for execution and traps.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

1.4 Binary and text formats

Binary and text formats must be engineered as an observable mechanism rather than a frontier label or speculative promise. Use wasm binaries, WAT, custom sections, names, debugging metadata, and reproducible builds. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for binary and text formats.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating binary and text formats as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for binary and text formats.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 01 laboratory and review

Practical laboratory

Create and inspect a small module, then test validation failure, out-of-bounds memory, and missing imports.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore garbage collection, typed references, multiple memories, threads, and future core proposals.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 02

Component Model, WIT, and canonical ABI

Move from low-level modules to language-neutral composable components.

Chapter operating position

Move from low-level modules to language-neutral composable components.

2.1 Components and worlds

Components and worlds must be engineered as an observable mechanism rather than a frontier label or speculative promise. Define imported and exported interfaces, worlds, instances, resources, and composition. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for components and worlds.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating components and worlds as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for components and worlds.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

2.2 WebAssembly Interface Types

WebAssembly Interface Types must be engineered as an observable mechanism rather than a frontier label or speculative promise. Model records, variants, enums, flags, options, results, lists, strings, resources, futures, and streams. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for webassembly interface types.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating webassembly interface types as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for webassembly interface types.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.



2.3 Canonical ABI

Canonical ABI must be engineered as an observable mechanism rather than a frontier label or speculative promise. Lower and lift high-level values across module boundaries with memory, realloc, encoding, and ownership rules. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for canonical abi.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating canonical abi as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for canonical abi.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

2.4 Composition and virtualization

Composition and virtualization must be engineered as an observable mechanism rather than a frontier label or speculative promise. Link components, adapt versions, wrap legacy modules, virtualize dependencies, and expose host services. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for composition and virtualization.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating composition and virtualization as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for composition and virtualization.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 02 laboratory and review

Practical laboratory

Define two WIT components in different conceptual languages and compose them through a versioned interface.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore component registries, typed capability marketplaces, and verified composition.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 03

WASI 0.2, WASI 0.3, and native async

Use portable system interfaces while tracking rapidly evolving maturity.

Chapter operating position

Use portable system interfaces while tracking rapidly evolving maturity.

3.1 WASI interface families

WASI interface families must be engineered as an observable mechanism rather than a frontier label or speculative promise. Understand CLI, filesystem, clocks, random, sockets, HTTP, configuration, key-value, and command worlds. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for wasi interface families.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating wasi interface families as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for wasi interface families.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

3.2 Preview 2 and WASI 0.2

Preview 2 and WASI 0.2 must be engineered as an observable mechanism rather than a frontier label or speculative promise. Use component-based interfaces, resources, polling, streams, HTTP proxy, and current toolchain support. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for preview 2 and wasi 0.2.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating preview 2 and wasi 0.2 as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for preview 2 and wasi 0.2.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.



3.3 WASI 0.3 native async

WASI 0.3 native async must be engineered as an observable mechanism rather than a frontier label or speculative promise. Use async functions, future and stream primitives, runtime scheduling, and simplified HTTP service interfaces. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for wasi 0.3 native async.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating wasi 0.3 native async as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for wasi 0.3 native async.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

3.4 Version pinning and migration

Version pinning and migration must be engineered as an observable mechanism rather than a frontier label or speculative promise. Align WIT packages, wit-bindgen, runtime, adapters, language targets, and compatibility tests. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for version pinning and migration.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating version pinning and migration as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for version pinning and migration.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 03 laboratory and review

Practical laboratory

Create a migration matrix from WASI Preview 1 and 0.2 to 0.3 and identify toolchain pins and fallback paths.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore portable structured concurrency and distributed components.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 04

Capability security and sandbox design

Grant only the resources a component needs.

Chapter operating position

Grant only the resources a component needs.



4.1 Default-deny host resources

Default-deny host resources must be engineered as an observable mechanism rather than a frontier label or speculative promise. Deny filesystem, environment, network, clocks, random, processes, devices, and secrets until granted. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for default-deny host resources.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating default-deny host resources as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for default-deny host resources.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

4.2 Preopened and virtual resources

Preopened and virtual resources must be engineered as an observable mechanism rather than a frontier label or speculative promise. Expose bounded directories, sockets, key-value stores, HTTP clients, clocks, and synthetic capabilities. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for preopened and virtual resources.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating preopened and virtual resources as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for preopened and virtual resources.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.



4.3 Identity and policy

Identity and policy must be engineered as an observable mechanism rather than a frontier label or speculative promise. Bind component digest, publisher, tenant, mission, requested world, target, duration, and resource limits. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for identity and policy.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating identity and policy as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for identity and policy.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

4.4 Untrusted code and side channels

Untrusted code and side channels must be engineered as an observable mechanism rather than a frontier label or speculative promise. Address sandbox escapes, runtime bugs, denial of service, timing, cache, covert channels, and supply chain. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for untrusted code and side channels.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating untrusted code and side channels as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for untrusted code and side channels.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 04 laboratory and review

Practical laboratory

Run a component under progressively broader grants and verify prohibited filesystem and network access.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore capability-secure component operating systems and confidential WASM enclaves.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 05

Language toolchains and developer workflow

Build portable components without hiding language-specific limitations.

Chapter operating position

Build portable components without hiding language-specific limitations.



5.1 Rust and wit-bindgen

Rust and wit-bindgen must be engineered as an observable mechanism rather than a frontier label or speculative promise. Generate bindings, implement guest interfaces, select targets, package components, and manage async support. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for rust and wit-bindgen.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating rust and wit-bindgen as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for rust and wit-bindgen.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

5.2 Python, JavaScript, Go, and other languages

Python, JavaScript, Go, and other languages must be engineered as an observable mechanism rather than a frontier label or speculative promise. Compare componentization, runtime size, garbage collection, async, exceptions, and ecosystem. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for python, javascript, go, and other languages.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating python, javascript, go, and other languages as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for python, javascript, go, and other languages.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

5.3 Testing and debugging

Testing and debugging must be engineered as an observable mechanism rather than a frontier label or speculative promise. Use unit tests, host mocks, component tests, WAT inspection, traces, stack information, and conformance suites. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for testing and debugging.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating testing and debugging as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for testing and debugging.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

5.4 Packaging and registries

Packaging and registries must be engineered as an observable mechanism rather than a frontier label or speculative promise. Publish WIT packages, components, metadata, signatures, provenance, dependencies, and compatibility. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for packaging and registries.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating packaging and registries as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for packaging and registries.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 05 laboratory and review

Practical laboratory

Build a language-neutral component test contract and validate one implementation against it.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore universal package management for signed components and interfaces.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 06

Edge, serverless, plugin, and agent architectures

Use portable sandboxing where process or container weight is undesirable.

Chapter operating position

Use portable sandboxing where process or container weight is undesirable.



6.1 Edge functions

Edge functions must be engineered as an observable mechanism rather than a frontier label or speculative promise. Deploy small startup-efficient workloads with local resources, intermittent links, and controlled updates. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for edge functions.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating edge functions as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for edge functions.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

6.2 Serverless components

Serverless components must be engineered as an observable mechanism rather than a frontier label or speculative promise. Serve HTTP components, compose middleware, scale, cache, isolate tenants, and manage state. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for serverless components.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating serverless components as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for serverless components.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

6.3 Plugins and extensions

Plugins and extensions must be engineered as an observable mechanism rather than a frontier label or speculative promise. Load third-party logic into applications through typed interfaces, quotas, versioning, and revocation. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for plugins and extensions.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating plugins and extensions as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for plugins and extensions.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

6.4 Agent capabilities

Agent capabilities must be engineered as an observable mechanism rather than a frontier label or speculative promise. Package tools and skills as components with explicit effects, schemas, evidence, and runtime isolation. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for agent capabilities.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating agent capabilities as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for agent capabilities.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 06 laboratory and review

Practical laboratory

Design a WASM-based plugin and agent-tool system with separate read, write, network, and secret capabilities.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore distributed component meshes and portable sovereign application stacks.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 07

Performance, networking, and observability

Measure actual portability cost and runtime behavior.

Chapter operating position

Measure actual portability cost and runtime behavior.

7.1 Compilation and startup

Compilation and startup must be engineered as an observable mechanism rather than a frontier label or speculative promise. Compare interpretation, baseline, optimizing compilation, caching, AOT, cold start, and code size. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for compilation and startup.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating compilation and startup as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for compilation and startup.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

7.2 Memory and calls

Memory and calls must be engineered as an observable mechanism rather than a frontier label or speculative promise. Measure linear-memory access, copies, canonical ABI crossing, resource handles, allocation, and garbage collection. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for memory and calls.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating memory and calls as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for memory and calls.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

7.3 Networking and async

Networking and async must be engineered as an observable mechanism rather than a frontier label or speculative promise. Evaluate HTTP, sockets, streams, futures, runtime scheduler, backpressure, cancellation, and timeout. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for networking and async.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating networking and async as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for networking and async.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

7.4 Telemetry and profiling

Telemetry and profiling must be engineered as an observable mechanism rather than a frontier label or speculative promise. Trace host and guest calls, component identity, resource use, traps, errors, effects, and user outcome. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for telemetry and profiling.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating telemetry and profiling as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for telemetry and profiling.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 07 laboratory and review

Practical laboratory

Benchmark native, containerized, and WASM implementations of a small service with equivalent security boundaries.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Explore hardware-assisted component isolation and zero-copy canonical interfaces.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

CHAPTER 08

Supply chain, lifecycle, and adoption

Treat portable components as governed software artifacts.

Chapter operating position

Treat portable components as governed software artifacts.



8.1 Artifact integrity

Artifact integrity must be engineered as an observable mechanism rather than a frontier label or speculative promise. Record source, toolchain, WIT, dependencies, component digest, signature, SBOM, provenance, and tests. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for artifact integrity.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating artifact integrity as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for artifact integrity.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

8.2 Runtime qualification

Runtime qualification must be engineered as an observable mechanism rather than a frontier label or speculative promise. Validate conformance, sandbox, CVEs, resource controls, performance, platform support, and upgrade behavior. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for runtime qualification.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Implementation sketch

```
component:
  id: mythos:capability/analyzer@1.0.0
  world: analyzer
  imports:
    - wasi:filesystem/preopens@0.2
    - mythos:evidence/store@1.0
  exports:
    - analyze: func(input: list<u8>) -> result<report, error>
  grants:
    filesystem: [/workspace/read-only]
    network: deny
    environment: deny
  limits:
    memory_mib: 128
    cpu_ms: 5000
```

Failure patterns

Treating runtime qualification as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for runtime qualification.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

8.3 Compatibility and migration

Compatibility and migration must be engineered as an observable mechanism rather than a frontier label or speculative promise. Version worlds, interfaces, resources, adapters, hosts, registries, and consumers. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for compatibility and migration.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating compatibility and migration as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for compatibility and migration.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

8.4 Adoption decision

Adoption decision must be engineered as an observable mechanism rather than a frontier label or speculative promise. Compare WASM with native processes, containers, microVMs, language VMs, eBPF, and browser execution. The design should identify physical and logical assumptions, authoritative state, interfaces, identities, resource limits, uncertainty, telemetry, failure behavior, recovery, and acceptance evidence.

This guide traces the subject through the source, toolchain, module or component, WIT contract, runtime grant, execution, host effect, telemetry, and portable outcome chain. A laboratory demonstration, simulation result, standards reference, valid packet, successful kernel, or fluent AI interaction is not sufficient proof. Intended behavior must be reconciled with measured state, negative tests, degraded conditions, lifecycle constraints, complete-system cost, and user or mission outcomes.

Implementation begins with a small known-good baseline and explicit invariants. Introduce one device, algorithm, optical engine, robot skill, transfer path, component interface, GPU kernel, approval mode, or network feature at a time; preserve before-state evidence; test prohibited and failed behavior; and retain a reversible fallback.

Troubleshooting locates the first divergence from the expected state machine or physical model. Account for calibration, loss, noise, timing, version skew, hidden dependencies, resource exhaustion, stale state, synchronization, permissions, model uncertainty, supply-chain changes, and missing evidence. Closure requires causal explanation, reproducibility, validated recovery, and disclosed limitations.

Engineering practices

Define the objective, authority, physical assumptions, interfaces, resources, risks, and acceptance evidence for adoption decision.

Separate stable standards and deployable practice from drafts, experiments, prototypes, and research hypotheses.

Version hardware, firmware, software, models, schemas, calibration, data, policy, and evaluation artifacts.

Test nominal, prohibited, noisy, degraded, overloaded, disconnected, failed, recovery, and rollback conditions.

Preserve source, configuration, measurement, trace, artifact, approval, experiment, and user-outcome evidence.

Adopt through staged pilots, independent baselines, explicit exit criteria, and reversible architecture.

Failure patterns

Treating adoption decision as mature because a component or laboratory demonstration exists.

Comparing a specialized accelerator or protocol against a weak or incomplete classical baseline.

Ignoring physical loss, calibration, timing, packaging, power, thermal, network, and operational constraints.

Allowing AI, autonomy, local compute, or disconnected operation to bypass identity, policy, safety, and evidence controls.

Using average performance while hiding tails, failed trials, environmental sensitivity, or complete-system overhead.

Declaring adoption success without interoperability, maintenance, recovery, workforce, supplier, and exit evidence.

Validation and evidence

Method	Expected evidence
Examine	Architecture, standards maturity, physical assumptions, identity, interfaces, versions, resources, and risks for adoption decision.
Observe	Signals, states, measurements, packets, events, traces, resource use, errors, artifacts, approvals, and user-visible outcomes.
Test	Expected, prohibited, noisy, adversarial, overloaded, disconnected, partially failed, recovery, and rollback cases.
Measure	Correctness, fidelity, accuracy, latency, throughput, energy, availability, safety, cost, scalability, and operator effort.
Reconcile	Claimed capability against measured physical, software, network, security, operational, and user evidence.

Chapter 08 laboratory and review

Practical laboratory

Create an adoption scorecard for an edge plugin platform and define exit and rollback conditions.

Laboratory evidence package

Architecture or topology showing boundaries, identities, dependencies, and authoritative inputs.

Versioned configuration or code with explanatory comments and reproducible setup instructions.

Nominal-path validation plus at least two injected failure or boundary conditions.

Structured logs, traces, metrics, packet captures, test output, or other appropriate evidence.

A concise decision record covering findings, limitations, residual risks, and next actions.

Frontier extension

Develop policy-verifiable portable compute fabrics spanning browser, device, edge, cloud, and air-gapped systems.

Review gate

Advance only when another engineer can reproduce the result, explain the architecture, identify the failure boundary, and distinguish measured evidence from assumptions or projections.

Integrated learning roadmap

The guide is intended to be revisited. First establish fluency, then build a small implementation, operate it under failure, automate repeatable work, and finally evaluate emerging techniques against a stable baseline. Progress is demonstrated through evidence rather than reading completion alone.

Stage	Demonstrated capability
Foundation	Explain the system from first principles and trace its critical path without relying on product terminology.
Build	Implement the smallest representative system with explicit contracts, identities, telemetry, and tests.
Operate	Diagnose injected failures, recover safely, and preserve evidence of the causal chain.
Automate	Convert a proven manual workflow into bounded, idempotent, observable automation.
Architect	Design for scale, resilience, security, interoperability, and lifecycle change.
Explore	Evaluate frontier techniques using stated hypotheses, limitations, and adoption gates.

Source authority register

Source policy

Official standards, specifications, and project documentation remain with their issuing organizations. Mythos Systems provides original engineering interpretation, synthesis, implementation guidance, and relationship mapping. Source verification dates do not guarantee that an authority has not changed since review.

01. W3C - WebAssembly Core Specification
Role: WebAssembly execution, validation, modules, memory, tables, instructions, and embedding semantics.
Verified: 2026-07-26
Official source: <https://www.w3.org/TR/wasm-core-2/>

02. Bytecode Alliance - The WebAssembly Component Model
Role: Components, WIT interfaces, composition, canonical ABI, and language tooling.
Verified: 2026-07-26
Official source: <https://component-model.bytecodealliance.org/>

03. Bytecode Alliance - WASI and Component Model FAQ
Role: Current WASI 0.3 status, native async, streams, futures, and migration context.
Verified: 2026-07-26
Official source: <https://component-model.bytecodealliance.org/reference/faq.html>

04. Bytecode Alliance - Wasmtime Component Model
Role: Reference runtime behavior, default-deny resource access, and WASI 0.2/0.3 execution.
Verified: 2026-07-26
Official source: <https://component-model.bytecodealliance.org/running-components/wasmtime.html>

Authority application and refresh triggers

Authority class	Required library action
Protocols and specifications	Recheck interoperability, security considerations, defaults, and migration guidance when a referenced specification changes.
Security and assurance frameworks	Update control mappings, evidence expectations, risk language, and assessment procedures after material framework revision.
Languages, runtimes, and projects	Re-run examples, tests, deprecation checks, and operational recommendations against supported releases.
Benchmarks and evaluations	Re-execute claims when workloads, methods, models, compilers, drivers, hardware, or scoring rules change.
Operational evidence	Incorporate validated incidents, failure tests, reader corrections, and measured implementation outcomes into the next revision.

Limitations, freshness, and revision control

Boundary	Statement
Publication limitation	This guide synthesizes broad engineering practice. It does not replace product documentation, legal advice, safety certification, or environment-specific design review.
Evidence limitation	Not every pattern is benchmarked in every environment. Examples illustrate engineering methods and must be validated against actual versions, workloads, hardware, policy, and failure conditions.
Freshness limitation	Vendor behavior, software libraries, AI models, security threats, and emerging standards can change quickly. Volatile claims require source revalidation before operational adoption.
Adoption limitation	Frontier content is not an implicit adoption recommendation. Adoption requires defined value, qualification gates, controlled proof, rollback, ownership, and residual-risk acceptance.
Status boundary	Candidate Focused Field Guide describes the publication, not certification of a product, implementation, engineer, or organization.

Revision record

Version	Revision statement
1.0.0-candidate	Initial candidate living-guide edition. Published 2026-07-26; scheduled review 2026-10-26.

Search and relationship metadata

Dimension	Engineering position
Primary domain	MYTHOS-EMT - Emerging Technology & Sovereign Infrastructure
Secondary domain	MYTHOS-SWE; MYTHOS-CLD; MYTHOS-SEC
Publication class	Field Guide / Canonical Living Overview Guide
Skill levels	L1 Foundational; L2 Practitioner; L3 Advanced; L4 Frontier
Tags	WebAssembly, WASI 0.3, component model, WIT, Wasmtime, sandbox, portable compute, edge functions, plugins, capability security
Canonical route	/library/field-guides/mythos-fg-frontier-0007/

Living publication

Corrections, expanded laboratories, improved diagrams, authority changes, and new engineering evidence should revise this permanent publication rather than create disconnected replacements.