


CMP-004

Agentic Frameworks

Agent runtimes, orchestration and governed infrastructure automation



CURRENT EVIDENCE SNAPSHOT

Official product sources refreshed through 2026-09-17

Controlled Mythos laboratory rerun: NOT ASSERTED

VERSION 1.2.0-candidate

CANDIDATE COMPARATIVE EVALUATION / PUBLIC

BENCHMARK DOCTRINE

Gates before scores. Evidence before certainty. Profile fit before universal ranking.

Agentic Frameworks

Agent runtimes, orchestration and governed infrastructure automation

DOCUMENT ID

CMP-004

VERSION

1.2.0-candidate

STATUS

Candidate Comparative Evaluation

PUBLICATION DATE

2026-09-17

SCHEDULED REVIEW

2026-12-16

CANONICAL ROUTE

/library/evaluations/cmp-004/

3

CANDIDATES

18

ATOMIC CRITERIA

9

MANDATORY GATES

3

WORKLOAD PROFILES

E2

CURRENT MAX EVIDENCE

CURRENT DECISION BOUNDARY

Reopen the comparison. AgentCore and Itential changed materially after the July baseline while LangGraph continued to evolve. No inherited ranking should be promoted to a current recommendation until durable execution, authority, observability, replay, migration, and failure tests are rerun.

NO CURRENT UNIVERSAL WINNER IS PUBLISHED FROM THIS REFRESH.

July 24 baseline scores are preserved as lineage and sensitivity evidence, not silently reissued as September laboratory results.

LangGraph

Python 1.2.11 / JS 1.4.15

REFRESH TRIGGER

Durable graph runtime and checkpoint ecosystem continue to evolve across Python and JavaScript lines.

Amazon Bedrock AgentCore

Managed service / Sept 2026 updates

REFRESH TRIGGER

Batch Evaluations and A/B Testing are now GA, materially strengthening the observe-evaluate-improve proof loop.


Itential Automation Platform

Platform 6.5 / Cloud Jul-Aug 2026

REFRESH TRIGGER

FlowAI, Work Center, asynchronous workflow support, and auditable agent sessions expand governed infrastructure-agent operation.

FINAL BENCHMARK SYSTEM / DECISION ARCHITECTURE



What changed since the July benchmark snapshot

Changes below are sourced from current official release documentation. They identify revalidation triggers; they are not translated into new numeric scores without controlled proof.

LangGraph

Python 1.2.11 / JS 1.4.15

MATERIAL DELTA

Durable graph runtime and checkpoint ecosystem continue to evolve across Python and JavaScript lines.

CRITERIA TOUCHED

C01 / C04 / C07 / C11 / C12 / C14

MODERATE REVALIDATION

Amazon Bedrock AgentCore

Managed service / Sept 2026 updates

MATERIAL DELTA

Batch Evaluations and A/B Testing are now GA, materially strengthening the observe-evaluate-improve proof loop.

CRITERIA TOUCHED

C01 / C03 / C07 / C08 / C11 / C18

MATERIAL REVALIDATION

Itential Automation Platform

Platform 6.5 / Cloud Jul-Aug 2026

MATERIAL DELTA

FlowAI, Work Center, asynchronous workflow support, and auditable agent sessions expand governed infrastructure-agent operation.

CRITERIA TOUCHED

C01 / C03 / C05 / C06 / C07 / C11

MATERIAL REVALIDATION

MANDATORY GATES / CURRENT REVALIDATION POSTURE

Mandatory gates remain non-compensable

No current release is marked PASS solely from documentation. E2 means the official source supports the capability path; LAB and RETEST identify proof still required. HOLD blocks a current decision.

MANDATORY GATE	LANGGRAPH	AMAZON BEDROCK	ITENTIAL AUTOMA
C01 Functional coverage	RETEST	RETEST	RETEST
C02 Security / trust	E2	E2	E2
C03 Governance / approval	E2	RETEST	RETEST
C04 State integrity	RETEST	E2	E2
C07 Observability / evidence	RETEST	RETEST	RETEST
C09 Resilience / continuity	LAB	LAB	LAB
C11 Change safety / rollback	RETEST	RETEST	RETEST
C16 Portability / exit	LAB	LAB	LAB
C18 Assurance confidence	HOLD	HOLD	HOLD

LEGEND

E2

official-source support only

LAB

controlled proof required

RETEST



release delta touches this gate

HOLD

decision gate remains closed

Historical profile sensitivity - preserved, not reissued

Values below are the 2026-07-24 evidence-adjusted baseline. They show sensitivity to operating-model weights. The September refresh does not recompute rank because E3 lab proof has not been reproduced.



CURRENT RANK STATUS: WITHHELD PENDING CONTROLLED PROOF.

Evidence confidence is separate from capability

E5

Independent repeatable validation

confidence modifier 1.00

E4

Production evidence in adopting environment

confidence modifier 0.91

E3

Controlled Mythos laboratory result

confidence modifier 0.78

E2

Official documentation / release / source

confidence modifier 0.64

E1

Vendor assertion or marketing

confidence modifier 0.38

E0

Unsupported or unverified

confidence modifier 0.00

CURRENT EVIDENCE STATE

LangGraph

E2 REFRESHED

Historical adjusted confidence: 59.7%

E3 LAB: NOT REPRODUCED

Amazon Bedrock AgentCore

E2 REFRESHED

Historical adjusted confidence: 58.2%

E3 LAB: NOT REPRODUCED

Itential Automation Platform

E2 REFRESHED

Historical adjusted confidence: 58.2%

E3 LAB: NOT REPRODUCED

ASSURANCE RULE

A current release note can change capability context, but it cannot raise assurance beyond E2. Current recommendation strength is therefore bounded until the test plan produces reproducible artifacts.

MIGRATION, LIFECYCLE AND IRREVERSIBLE-COMMITMENT RISK

Lifecycle risk is evaluated independently of features

The benchmark models migration, upgrade, compatibility, support, staffing, data/state export, downtime and exit. Current release changes below are explicit proof triggers.

LangGraph

MODERATE REVALIDATION

Durable graph runtime and checkpoint ecosystem continue to evolve across Python and JavaScript lines.

RETEST: C01 / C04 / C07 / C11 / C12 / C14

Amazon Bedrock AgentCore

MATERIAL REVALIDATION

Batch Evaluations and A/B Testing are now GA, materially strengthening the observe-evaluate-improve proof loop.

RETEST: C01 / C03 / C07 / C08 / C11 / C18

Itential Automation Platform

MATERIAL REVALIDATION

FlowAI, Work Center, asynchronous workflow support, and auditable agent sessions expand governed infrastructure-agent operation.

RETEST: C01 / C03 / C05 / C06 / C07 / C11

IRREVERSIBLE COMMITMENT GATE

Before migration or procurement lock-in, prove rollback, state portability, operational reconstruction, and supplier-exit assumptions.

Controlled proof is the next decision-producing step

REPRODUCTION STATE / NOT EXECUTED IN THIS EVIDENCE-REFRESH RELEASE

01

CONTROLLED SCENARIO

Run a long-lived mission with branching, retry, interruption, approval, restart, replay, and compensation.

02

CONTROLLED SCENARIO

Integrate local and frontier models, MCP tools, secrets, identity, and policy.

03

CONTROLLED SCENARIO

Inject model timeout, malformed tool output, duplicate event, lost checkpoint, and operator cancellation.

04

CONTROLLED SCENARIO

Measure task throughput, state growth, token use, tool latency, and evidence completeness.

05

CONTROLLED SCENARIO

Export mission state, traces, memory, policies, and tool definitions.

06

CONTROLLED SCENARIO

Red-team prompt injection, confused-deputy behavior, excess authority, and unsafe recovery.

EVIDENCE PACKAGE

Preserve exact versions, architecture, configuration, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

Current decision posture

Reopen the comparison. AgentCore and Itential changed materially after the July baseline while LangGraph continued to evolve. No inherited ranking should be promoted to a current recommendation until durable execution, authority, observability, replay, migration, and failure tests are rerun.

CURRENT RANKING: NOT REISSUED / CONTROLLED LAB REQUIRED

CURRENT OFFICIAL-SOURCE REGISTER / CHECKED 2026-09-17

LangGraph	E2 / OFFICIAL	
LangGraph GitHub releases		github.com

Amazon Bedrock AgentCore	E2 / OFFICIAL	
Amazon Bedrock AgentCore release notes		docs.aws.amazon.com

Itential Automation Platform	E2 / OFFICIAL	
Itential Platform 6.5 release notes		docs.itential.com

REVISION LINEAGE

1.2.0-candidate / 2026-09-17 - current-source evidence refresh; no lab rescore issued.

1.1.0-candidate / 2026-07-24 - baseline benchmark using the final comparative evaluation system.

Trigger earlier review after major release, acquisition, licensing change, critical vulnerability, material outage, failed PoC, or workload change.

BASELINE ANALYTIC RECORD CONTINUES ON PAGE 11 FOR TRACEABILITY.

ATOMIC CRITERIA MODEL

WEIGHTED CAPABILITY WITH EXPLICIT MINIMUM EVIDENCE

ID	CRITERION	WEIGHT	GATE	MINIMUM EVIDENCE
C01	Agent orchestration, execution, memory, tool, and human-control coverage	5	YES	Feature documentation, APIs, configuration exports, and scenario demonstration.
C02	Agent identity, tool authority, isolation, and policy architecture	5	YES	Threat model, identity flows, RBAC tests, audit records, and security configuration.
C03	Governance and approval controls	5	YES	Approval workflow, policy controls, separation-of-duties tests, and decision records.
C04	Durable mission, state, checkpoint, and memory integrity	5	YES	Backup, restore, rollback, drift, and corruption-recovery evidence.
C05	Automation and API quality	4	NO	API specifications, authentication tests, rate limits, event samples, and automation runs.
C06	Integration ecosystem	4	NO	Supported integrations, connector tests, failure behavior, and lifecycle ownership.
C07	Observability and evidence	4	YES	Logs, traces, metrics, reports, export tests, and evidence-retention controls.
C08	Agent, task, model, tool, and concurrency scale	4	NO	Controlled load tests, topology scale, saturation behavior, and capacity model.
C09	Resilience and continuity	5	YES	Failure injection, HA tests, recovery timing, degraded-mode operation, and runbooks.
C10	Usability and operator cognition	3	NO	Task observation, error-rate measures, navigation tests, and operator interviews.
C11	Interrupt, approval, replay, compensation, and rollback safety	5	YES	Plan or preview output, canary tests, rollback execution, and post-change verification.
C12	Extensibility and programmability	3	NO	Plugin, module, provider, workflow, or code-extension tests and upgrade impact analysis.
C13	Deployment and sovereignty options	3	NO	Deployment architecture, data-flow mapping, residency evidence, and offline tests.
C14	Lifecycle and upgrade discipline	4	NO	Release notes, upgrade test, compatibility matrix, support policy, and migration plan.
C15	Economics and cost predictability	3	NO	Bill-of-materials, licensing model, staffing estimates, metering, and sensitivity analysis.
C16	Portability and exit	4	YES	Export tests, format analysis, replacement workflow, and decommission evidence.
C17	Vendor and ecosystem risk	3	NO	License analysis, roadmap evidence, bus-factor indicators, and dependency inventory.
C18	Assurance confidence	5	YES	Source provenance, lab artifacts, production evidence, independent replication, and review date.

SCORE SCALE

0 absent; 1 materially deficient; 2 partial; 3 adequate with limits; 4 strong; 5 leading for the defined profile. Scores remain provisional until the required evidence grade is achieved.

RAW CAPABILITY SCORECARD

DOCUMENTED CAPABILITY BEFORE EVIDENCE DISCOUNT

CRITERION	LANGGRAPH	AGENTCORE	ITENTIAL
C01 Agent orchestration, execution, memory, tool, and human-control coverage	4.5	4.7	4.0
C02 Agent identity, tool authority, isolation, and policy architecture	4.0	4.6	4.4
C03 Governance and approval controls	4.1	4.5	4.7
C04 Durable mission, state, checkpoint, and memory integrity	4.6	4.4	4.2
C05 Automation and API quality	4.5	4.6	4.5
C06 Integration ecosystem	4.3	4.6	4.8
C07 Observability and evidence	4.4	4.7	4.4
C08 Agent, task, model, tool, and concurrency scale	4.2	4.7	4.2
C09 Resilience and continuity	4.1	4.5	4.3
C10 Usability and operator cognition	4.0	4.4	4.3
C11 Interrupt, approval, replay, compensation, and rollback safety	4.8	4.4	4.7
C12 Extensibility and programmability	4.8	4.2	4.2
C13 Deployment and sovereignty options	4.5	3.2	3.8
C14 Lifecycle and upgrade discipline	4.1	4.2	4.1
C15 Economics and cost predictability	4.4	3.5	3.4
C16 Portability and exit	4.7	3.3	3.5
C17 Vendor and ecosystem risk	4.2	3.4	3.6
C18 Assurance confidence	3.5	3.4	3.4

WEIGHTED BASELINE / 100

LangGraph

86.1

AgentCore

84.6

Itential

83.4

EVIDENCE-ADJUSTED SCORECARD

CAPABILITY DISCOUNTED BY CURRENT EVIDENCE CONFIDENCE

CANDIDATE	RAW	EVIDENCE CONFIDENCE	ADJUSTED	CURRENT STATE
LangGraph	86.1	59.7%	51.5	CONTROLLED LAB PENDING
AgentCore	84.6	58.2%	49.8	CONTROLLED LAB PENDING
Itential	83.4	58.2%	49.2	CONTROLLED LAB PENDING

EVIDENCE SCALE

E0 / 0.00

Unsupported or unverified

E1 / 0.38

Vendor assertion or marketing statement

E2 / 0.64

Official documentation, release notes, or source repository

E3 / 0.78

Controlled Mythos laboratory result

E4 / 0.91

Production evidence from the adopting environment

E5 / 1.00

Independent repeatable validation

CURRENT CONFIDENCE BOUNDARY

Most candidate criteria are supported at E2: official documentation or source evidence. Environment-specific laboratory, production, and independent evidence remain future gates.

PROFILE-SPECIFIC RANKINGS

EVIDENCE-ADJUSTED SENSITIVITY ANALYSIS

CUSTOM GOVERNED AGENT RUNTIME

Prioritizes deterministic graphs, durable execution, interrupts, model independence, portability, and deep engineering control.

1. LangGraph	52.2
2. Amazon Bedrock AgentCore	50.4
3. Itential Automation Platform	49.8

MANAGED CLOUD AGENT OPERATIONS

Prioritizes managed runtime, identity, gateway, memory, observability, scale, and reduced platform burden.

1. LangGraph	50.8
2. Amazon Bedrock AgentCore	50.3
3. Itential Automation Platform	49.3

INFRASTRUCTURE AUTOMATION CONTROL PLANE

Prioritizes approvals, integration, deterministic execution, infrastructure workflows, audit, and enterprise operations.

1. LangGraph	51.9
2. Amazon Bedrock AgentCore	50.9
3. Itential Automation Platform	50.3

RANK SENSITIVITY

Small score differences are not decision certainty. Treat near-ties as a requirement for deeper PoC, commercial, migration, and operating-model evidence.

CANDIDATE DEEP DIVE / 01

LANGGRAPH

OPERATING MODEL

Low-level open-source orchestration framework and runtime for long-running, stateful agent graphs.

DOCUMENTED STRENGTHS

- Fine-grained control over graph, state, branching, interrupts, and persistence.
- Open-source and model-provider flexibility.
- Strong fit for custom deterministic agent runtimes.
- Developer-centric composability and testability.

CAUTIONS AND PROOF OBLIGATIONS

- Production platform services must be assembled or adopted separately.
- Security, tenancy, identity, policy, and operations remain architectural responsibilities.
- Low-level flexibility can become complexity without strong engineering standards.

BASELINE SCORE PROFILE



GATE DISPOSITION

PASS - custom governed agent runtime | CONDITIONAL - production platform services must be engineered

CANDIDATE DEEP DIVE / 02

AMAZON BEDROCK AGENTCORE

OPERATING MODEL

Managed framework-agnostic agent platform providing runtime, gateway, memory, identity, policy, observability, and operational services.

DOCUMENTED STRENGTHS

- Managed runtime and scaling.
- Integrated gateway, memory, identity, policy, and observability services.
- Framework and model flexibility within the AWS operating boundary.
- Reduced infrastructure burden for production agent operations.

CAUTIONS AND PROOF OBLIGATIONS

- AWS dependency and service maturity must be accepted.
- Portability, data boundaries, and cost require controlled validation.
- Managed abstractions can constrain low-level control and debugging.

BASELINE SCORE PROFILE



GATE DISPOSITION

PASS - managed cloud agent operations | CONDITIONAL - AWS boundary, portability, and cost

CANDIDATE DEEP DIVE / 03

INTENTIAL AUTOMATION PLATFORM

OPERATING MODEL

Enterprise network and infrastructure orchestration platform with integration, workflow, governance, and operational automation capabilities.


DOCUMENTED STRENGTHS

- Strong network and infrastructure orchestration orientation.
- Enterprise workflow, integration, approval, and operational controls.
- Useful bridge between intent, existing automation, controllers, and ITSM.
- Fits environments where deterministic infrastructure execution is primary.

CAUTIONS AND PROOF OBLIGATIONS

- Not a general-purpose low-level agent runtime.
- Model, memory, reasoning, and multi-agent features require integration or custom architecture.
- Commercial platform coupling and connector lifecycle require governance.

BASELINE SCORE PROFILE



GATE DISPOSITION

PASS - infrastructure automation control-plane profile | NOT EQUIVALENT - general-purpose agent-runtime comparison

CROSS-CANDIDATE CONTRASTS

WHERE ARCHITECTURAL DIFFERENCES MATTER MOST

CONTROL-LOOP MODEL

How authority moves from intent to plan, approval, execution, observation, reconciliation, and recovery.

STATE OWNERSHIP

Which system owns desired state, operational state, evidence, history, and conflict resolution.

MANAGED VERSUS SOVEREIGN BOUNDARY

Which runtime, identity, data, policy, relay, or service dependencies remain outside organizational control.

EXTENSION MODEL

Plugins, providers, modules, applications, workflows, SDKs, APIs, and custom code introduce different lifecycle risks.

FAILURE TRANSPARENCY

Candidates differ in how clearly they expose partial execution, degraded mode, retries, inconsistency, and recovery state.

EXIT MECHANICS

Export formats, state semantics, policy portability, knowledge transfer, and replacement effort are materially different.

CANDIDATE	CURRENT ADVANCEMENT POSITION
LangGraph	PASS - custom governed agent runtime
AgentCore	PASS - managed cloud agent operations
Itential	PASS - infrastructure automation control-plane profile

ARCHITECTURE AND INTEGRATION FIT

THE PRODUCT IS ONLY ONE PART OF THE OPERATING SYSTEM

REFERENCE OPERATING LAYERS

01 Intent, task, and mission model

02 Supervisor, graph, workflow, and durable state

03 Model, memory, context, and knowledge services

04 Tool gateway, identity, policy, approvals, and execution

05 Observation, evaluation, evidence, replay, and recovery

INTEGRATION BOUNDARIES

- Model providers and local runtimes
- MCP, APIs, CLIs, infrastructure controllers, and automation tools
- Identity, secrets, policy, and approval systems
- Vector, graph, object, event, and relational data stores
- Observability, evaluation, incident, and evidence platforms

ARCHITECTURE GATE

Every external dependency requires an owner, identity boundary, failure mode, evidence path, version policy, recovery procedure, and exit strategy.

SECURITY, OPERATIONS, LIFECYCLE, ECONOMICS, AND EXIT

CROSS-DOMAIN DECISION RECORD

SECURITY AND AUTHORITY

Identity, credentials, secrets, roles, privileged actions, signing, approvals, isolation, audit, and incident response must be tested as an integrated system.

RELIABILITY AND RECOVERY

Control-plane, data-plane, dependency, region, database, queue, network, and operator failures require defined degraded modes and recovery objectives.

CHANGE AND LIFECYCLE

Version, compatibility, migration, canary, rollback, content, plugin, provider, firmware, and decommission procedures are part of product fitness.

ECONOMICS AND CAPACITY

Licenses, metering, data, compute, hardware, storage, support, staffing, training, integration, migration, and opportunity cost require sensitivity analysis.

SOVEREIGNTY AND PRIVACY

Data location, support access, telemetry, subprocessors, encryption, key control, disconnected operation, and legal obligations must align with policy.

PORTABILITY AND EXIT

Configuration, policy, data, state, evidence, workflows, identities, and operational knowledge must be exportable and reconstructable.

DECISION OWNERSHIP

Architecture, security, operations, finance, procurement, legal, data, and service owners jointly accept the final profile, evidence, residual risk, and exit obligations.

RISK AND FAILURE REGISTER

SCENARIOS THAT CAN INVALIDATE A FEATURE-BASED SELECTION

ID	SEVERITY	FAILURE SCENARIO	REQUIRED TREATMENT
R-01	CRITICAL	A flexible framework is mistaken for a complete production platform.	PREVENT / DETECT / RECOVER / EVIDENCE
R-02	HIGH	A managed platform is accepted without testing portability, cost, data boundaries, and failure transparency.	PREVENT / DETECT / RECOVER / EVIDENCE
R-03	HIGH	Agent authority exceeds tool, data, or environment permissions.	PREVENT / DETECT / RECOVER / EVIDENCE
R-04	HIGH	State replay duplicates irreversible actions.	PREVENT / DETECT / RECOVER / EVIDENCE
R-05	MEDIUM	Human approval is cosmetic because plans, diffs, and consequences are not understandable.	PREVENT / DETECT / RECOVER / EVIDENCE
R-06	MEDIUM	Model quality is used to compensate for weak deterministic execution and verification.	PREVENT / DETECT / RECOVER / EVIDENCE

RISK RULE

A candidate with an unresolved critical failure path cannot advance because optional capability, market position, or commercial discount cannot compensate for missing safety or recovery evidence.

CONTROLLED PROOF-OF-CAPABILITY PLAN

REPEATABLE SCENARIOS, INSTRUMENTATION, AND ACCEPTANCE

TEST 01

Run a long-lived mission with branching, retry, interruption, approval, restart, replay, and compensation.

TEST 02

Integrate local and frontier models, MCP tools, secrets, identity, and policy.

TEST 03

Inject model timeout, malformed tool output, duplicate event, lost checkpoint, and operator cancellation.

TEST 04

Measure task throughput, state growth, token use, tool latency, and evidence completeness.

TEST 05

Export mission state, traces, memory, policies, and tool definitions.

TEST 06

Red-team prompt injection, confused-deputy behavior, excess authority, and unsafe recovery.

EVIDENCE PACKAGE

Preserve versions, architecture, configuration, data set, scenario, expected result, instrumentation, raw output, timing, failures, operator actions, deviations, acceptance state, owner, and artifact hashes.

CONDITIONAL RECOMMENDATION AND REVALIDATION

DECISION RECORD, ACCEPTANCE, AND NEXT EVIDENCE

RECOMMENDATION POSITION

- Select the architecture layer before selecting the product.
- Require durable-state, authority, replay, approval, and failure-injection tests.
- Preserve model and tool portability through typed contracts and evidence.
- Revalidate monthly because managed agent platforms and framework capabilities are changing rapidly.

CANDIDATE	ADJ. SCORE	GATE POSITION	LAB STATE
LangGraph	51.5	PASS - custom governed agent runtime; CONDITIONAL - production platform services must be engineered	PENDING
AgentCore	49.8	PASS - managed cloud agent operations; CONDITIONAL - AWS boundary, portability, and cost	PENDING
Itential	49.2	PASS - infrastructure automation control-plane profile; NOT EQUIVALENT - general-purpose agent-runtime comparison	PENDING

REVALIDATION SCHEDULE

Review no later than 2026-10-24. Review earlier after a major release, commercial change, critical vulnerability, material outage, architecture transition, failed acceptance test, or change to the target workload profile.

SOURCE AUTHORITY AND REVISION

CURRENT EVIDENCE SNAPSHOT AND PUBLICATION HISTORY

SOURCE ID	PUBLISHER	TITLE	CHECKED
NIST-AI-RMF	NIST	AI Risk Management Framework	2026-07-24
NIST-SP800-53	NIST	Security and Privacy Controls for Information Systems and Organizations	2026-07-24
LANGGRAPH	LangChain	LangGraph Overview	2026-07-24
LANGGRAPH-GH	LangChain	LangGraph Source Repository	2026-07-24
AGENTCORE	Amazon Web Services	Amazon Bedrock AgentCore Overview	2026-07-24
AGENTCORE-GATEWAY	Amazon Web Services	Amazon Bedrock AgentCore Gateway	2026-07-24
ITENTIAL	Itential	Itential Automation Platform Documentation	2026-07-24

SOURCE POSITION

Official documentation and source repositories support the current capability model. Source records preserve publisher, title, URL, purpose, and review date in the machine-readable authority register. Commercial terms, performance, and environment-specific behavior remain unverified until controlled proof.

REVISION HISTORY

1.1.0-candidate / 2026-07-24 - permanent-publication rebuild using the final benchmark system, profile-specific rankings, evidence adjustment, mandatory gates, and controlled PoC requirements.