GOVERNANCE / RESEARCH BRIEF

# HARNEXA Harness Architecture Overview

A technical architecture brief for controlled agent execution in commerce workflows.


| | |
|---|---|
| 01 | 6 enforcement layers |
| 02 | 3 risk classes |
| 03 | 1 audit chain |

EXECUTION ARCHITECTURE

# The model is not the control plane.

Agent architectures often expose too much tool access too early. A commerce workflow needs an execution layer that can block, pause, and explain actions before systems of record are changed. The harness is the product boundary. It turns model output into controlled operations by enforcing identity, permissions, budgets, approval, sensors, and audit records around every agent run.

**01**
The model should never be the only enforcement point.

**02**
Every tool needs risk class, scope, budget, and audit semantics.

**03**
The buyer should see what is blocked as clearly as what succeeds.

## Execution control plane

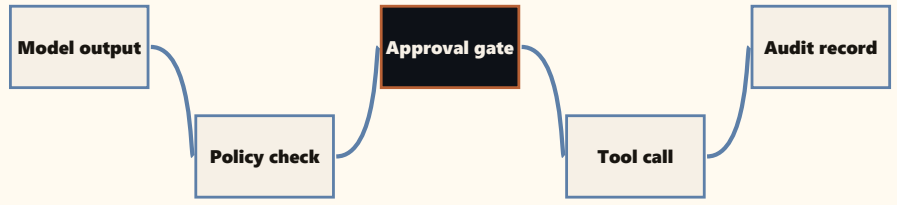


**Risk class, budget, terminal status, and denial are logged as first-class events.**

### MCKINSEY TECHNOLOGY / QUANTUMBLACK / BUILDING THE FOUNDATIONS FOR AGENTIC AI AT SCALE

Agentic AI at scale depends on stable interfaces, shared data meaning, and a controlled execution layer.

### PARLOA / GLOBAL AI AGENT MANAGEMENT & COMPLIANCE UPDATES

Centralized audit logs are becoming table stakes: security and operations teams need to see who changed what and when.

### OECD / THE AGENTIC AI LANDSCAPE AND ITS CONCEPTUAL FOUNDATIONS

The OECD distinguishes agentic AI by coordination, task decomposition, delegation, and sustained operation in less predictable environments.

TOOL BOUNDARY AND APPROVAL SWIMLANE

# Six enforcement layers.

Request to audit path: HARNEXA separates model reasoning from governed execution and review evidence. Read the artifact as an execution blueprint: every request passes through policy, budget, approval, sensors, and audit before the buyer trusts tool access.

## Tool boundary and approval swimlane

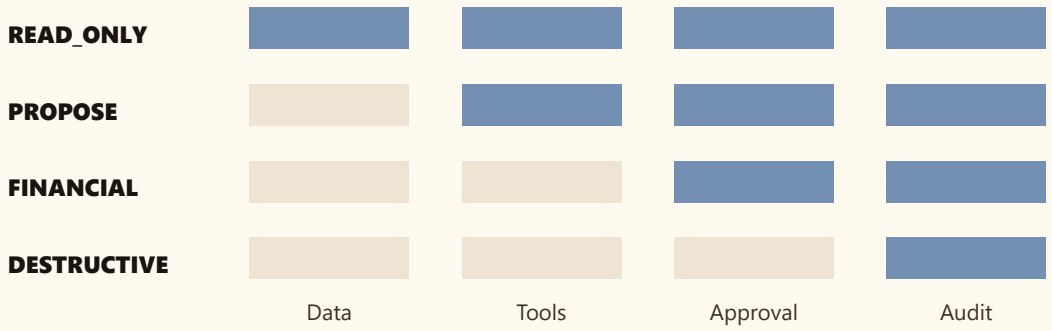


**REQUEST**
Normalize the workflow request and attach profile, purpose, buyer context, and maximum risk class.

**POLICY**
Evaluate permissions, memory namespace, tool scope, and budget before the tool layer is reached.

**GATE**
Pause any financial or destructive operation and record the review requirement before execution.

**AUDIT**
Store events, sensor results, approval state, and terminal status in an append-only path.

ARCHITECTURE REVIEW

# Prove what the system blocks before proving what it does.

Use the review to trace one recommendation from request to terminal state, including denied financial actions and audit evidence.

## INSPECTION QUESTIONS

- Can a reviewer trace one recommendation from request to final status?
- What happens when a tool call exceeds budget or risk class?
- How is approval represented in code rather than copy?
- Which parts of the workflow are read-only, proposal-only, or blocked?

## DO NOT SHIP UNTIL

- Do not expose broad API credentials to agent code.
- Do not confuse a chat transcript with an audit trail.
- Do not ship live-write tools until review gates are tested.

| | | |
|---|---|---|
| **ARCHITECTURE** | Map the workflow into request, model, tool, approval, and audit boundaries. | Harness architecture diagram and tool/action matrix. |
| **IMPLEMENTATION** | Add computational sensors for schema, permission, budget, and terminal status. | Sensor outputs visible in local tests and audit fixtures. |
| **REVIEW** | Run a denied financial action and prove the system pauses instead of executing. | PENDING_REP_REVIEW or equivalent human approval state. |

SOURCE LEDGER

# Source ledger and architecture boundary

The brief combines external research signals with HARNEXA's implementation boundary: identity, permission, approval, audit, CLEAR evaluation, and no autonomous public execution.

MCKINSEY TECHNOLOGY / QUANTUMBLACK

### Building the foundations for agentic AI at scale

McKinsey frames data quality, shared meaning, stable interfaces, observability, and controlled execution as scale requirements for agentic AI.

https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/building-the-foundations-for-agentic-ai-at-scale

PARLOA

### Global AI Agent Management & Compliance Updates

Parloa highlights centralized audit logs that show who changed what and when as agent programs scale across teams and markets.

https://www.parloa.com/blog/parloa_product_release_2026/

OECD

### The agentic AI landscape and its conceptual foundations

The OECD distinguishes AI agents from agentic AI by coordination, task decomposition, delegation, sustained operation, and operation in less predictable environments.

https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/the-agentic-ai-landscape-and-its-conceptual-foundations_a9d4b451/396cf758-en.pdf

RESEARCH STANCE

External research is cited as market signal; HARNEXA interpretation is kept separate from source claims.

BUYER ARTIFACT

Every report includes a concrete visual artifact that can be inspected in a sales, risk, or architecture review.

OPERATING BOUNDARY

Narrow tools, denied-action tests, approval gates, audit semantics. The asset is for governed review, not autonomous production execution.

HARNEXA BOUNDARY

Technical governance aid only. Not legal advice or compliance certification. HARNEXA AI builds audit-ready foundations for qualified review. No PDF in this library claims EU AI Act certification, legal advice, payment readiness, or autonomous production execution.