



From Logs to Decision Traces

Debugging and Governing AI Agents with OpenSearch

Pushkar Mishra

Principal SRE | Enterprise GenAI Architect

OpenText



From Logs to Decision Traces

A Practical Architecture for Observing and Governing AI Agents with OpenSearch

Operational evidence, policy records, and searchable agent behavior



Logs tell us what happened. Traces tell us where time went. LLM traces show model behaviour. Decision traces explain **why** an agent acted.



The Operational Question

An AI agent recommends rollback.

The important question is not:

"What did it answer?"


The important question is:

Can we
prove why?

A human reviewer needs to know:

- Which logs were checked?
- Which traces were inspected?
- Which runbook was used?
- Which evidence supported the recommendation?
- Which policy applied?
- Whether approval existed?

⚠️ Correctness in operations must be **explainable**.



System Assumptions

Assumed operating environment

Applications

Emit logs, metrics, and traces

Incidents

Have IDs, timelines, and ownership

Runbooks & History

Exist and are searchable

Agents

Can call tools — some read-only, some mutating


Tool risk distinction

Read-Only

- Query logs
- Query traces
- Search runbooks

Mutating

- Rollback release
- Restart service
- Scale deployment

 Policy decides what actions are allowed. That difference is why we need both observability and a clean execution boundary.



Modern Agent Stack

Agent platforms vary. The observability problem remains.

Agent Runtimes

LangGraph, CrewAI, Google ADK, Strands Agents, Bedrock AgentCore, custom Python

Agent Communication

A2A, MCP-style tools, REST APIs, queues

Safety & Control

mTLS, OIDC/OAuth, RBAC, OPA/Cedar, Model Armor-style prompt screening

Operational Evidence

OpenTelemetry + OpenSearch

- A2A is an open standard for communication and collaboration between agents built across different frameworks. The observability problem remains the same regardless of runtime choice.



Where Existing Observability Stops

Existing signals answer different questions

Logs What happened inside the system?	Traces Where did time go?
LLM Traces What prompts, model calls, and generations occurred?	Decision Traces ✦ Why did the agent choose this path?

LLM observability tools such as **Langfuse** and **LangSmith** help with model calls, prompt traces, evaluations, latency, and cost. They answer model-level questions. Decision traces answer a different question: **why did the agent choose this path?** For an operational workflow, that means tying model behaviour to logs, traces, incidents, runbooks, tool calls, policy results, evidence records, and audit.

The Gap

No single signal automatically connects model behavior, operational evidence, policy result, and audit context.



Decision Trace Definition

The central technical abstraction



A decision trace is the **flight recorder** for an AI agent — the path behind the answer.

intent

What the agent was trying to accomplish

tool selected

Which capability was invoked

evidence used

What data supported the decision

policy result

What happened at the execution boundary

action proposed

What the agent recommended

outcome recorded

What actually happened

Example event payload

```
{
  "event.kind": "policy_check",
  "agent.name": "policy-agent",
  "action": "rollback_service",
  "policy.result": "denied",
  "evidence.ids": [
    "ev-1001",
    "ev-1002"
  ]
}
```



Key fields: **event.kind** — what kind of event. **policy.result** — what happened at the boundary. **evidence.ids** — what proof supported the decision.

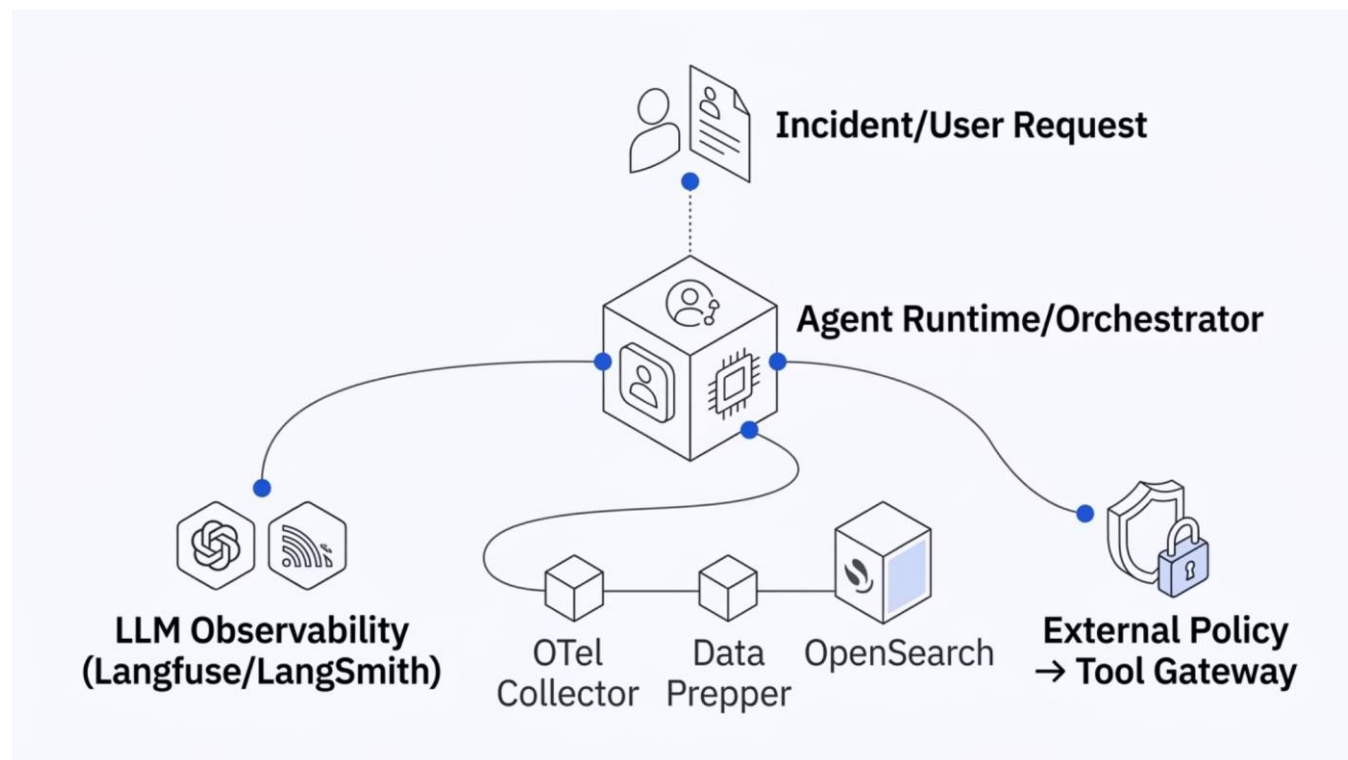


Without this structure, we get log confetti. Pretty, searchable confetti — but still confetti.



Reference Architecture

The whole system in one diagram



Agent Runtime Creates plans, tool calls, and decision events	LLM Observability Captures prompts, model calls, evals, token/cost signals	OTel Collector Vendor-agnostic receive, process, and export of telemetry
OpenSearch Stores and analyses operational evidence		Policy + Tool Gateway Authorises and executes production-changing actions
☐ OpenSearch records the policy decision. It does not approve or execute the rollback.		



Execution Boundary

Clarifying the role of policy without making OpenSearch the enforcement engine

Read-Only Tools → OpenSearch

- Query logs
- Query traces
- Search runbooks
- Inspect deployment metadata

Decision and evidence events flow directly to OpenSearch for storage and analysis.

Mutating Tools → Policy → Gateway → OpenSearch

- Rollback release
- Restart service
- Scale replicas
- Change routing or firewall rules

Mutating requests pass through an external policy service and tool gateway. The outcome event is then recorded in OpenSearch.

OpenSearch records and explains.

It is the evidence and analytics layer.


Policy services authorise.

They evaluate approval, risk, scope, and identity.

Tool gateways execute.

They carry out approved production-changing actions.

 Think of this like authentication logs: the log store records a 403, but it is not the authentication service.



Scenario: Checkout Latency Incident

SEV2 | PRODUCTION

checkout-api

Incident: inc-042

Symptom: p95 latency spike

Recent change: checkout-api:v2.7.3

Suspicious dependency: payment-auth

Relevant history: similar incident inc-019

Runbook: rb-009-01

Policy context: rollback requires approval

Why this scenario?

We will use this single scenario throughout the talk. It is concrete enough to follow, realistic enough to matter.

01

Checkout latency spikes after deployment

02

Suspicious dependency: payment-auth

03

Historical incident inc-019 is relevant

04

Runbook rb-009-01 requires approval for rollback

05

Agent team investigates, produces evidence, proposes rollback

06

Policy result is recorded in OpenSearch



Controlled Multi-Agent Workflow

Not a free-form chatbot — a controlled operational workflow



Triage Agent

Classify incident and severity

Output: incident_classified



Log Agent

Query error and timeout logs

Output: tool_call + evidence



Trace Agent

Find slow spans and dependencies

Output: tool_call + evidence



Retrieval Agent

Search similar incidents and runbooks

Output: retrieval + evidence



Planner Agent

Propose remediation

Output: action_proposed



Policy Agent

Record authorisation decision from policy boundary

Output: policy_check



Eval Agent

Score trace completeness

Output: eval_score



This is not a free-form chatbot wandering around the data centre with a flashlight and vibes. It is a controlled operational workflow.



Decision Flow for the Scenario

The full chain — from incident to evidence



The final answer is not the artefact. **The decision path is the artefact.**



Event Contract

Agent operations event contract

Correlation Fields

incident_id

Links all events to the incident

session_id

Groups all agent steps in one run

trace_id

Connects to system telemetry

step_id / parent_step_id

Enables replay of the sequence

Meaning Fields

event.kind

What kind of event this is

policy.result

What happened at the execution boundary

evidence.ids

What proof supported the decision

risk.level / eval.score


Risk and quality signals

Example payload

```
{
  "event.kind": "policy_check",
  "agent.session_id": "inc-042",
  "tool.name": "rollback_service",
  "policy.result": "denied",
  "risk.level": "high",
  "evidence.ids": [
    "ev-1001",
    "ev-1002"
  ]
}
```

i This contract gives OpenSearch stable fields for search, aggregation, dashboards, alerting, and audit.

Additional meaning fields: `agent.name`, `tool.name`, `usage.total_tokens`, `eval.score`



Live Walkthrough

Walkthrough: inc-042 Decision Trace

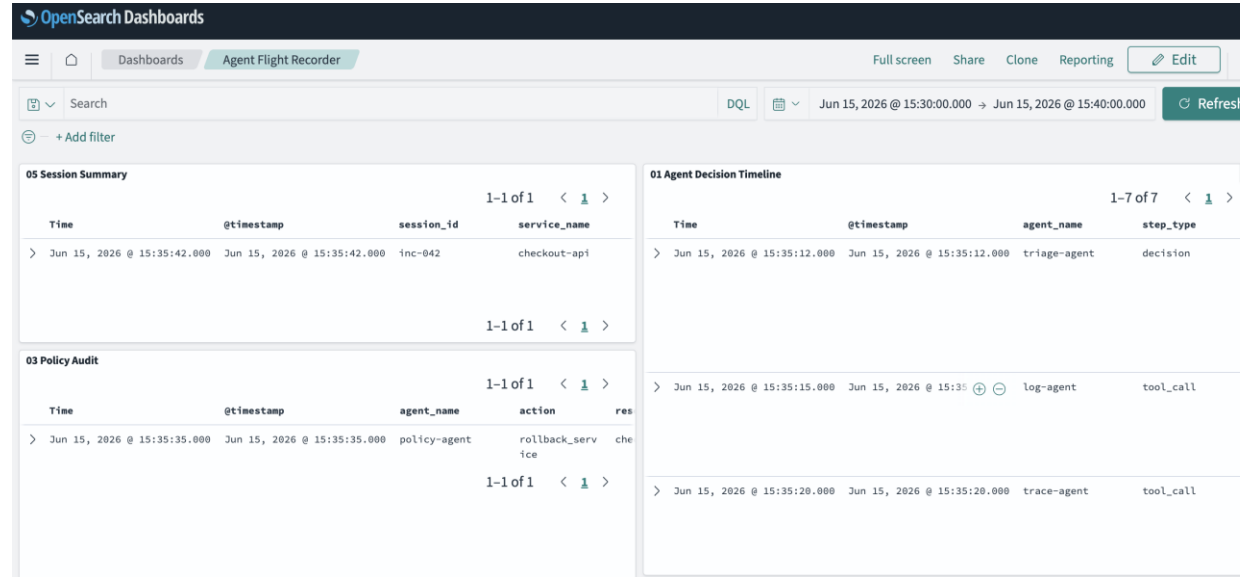
A single incident run becomes searchable operational evidence.

Agent timeline each step is indexed

Evidence explorer logs, traces, runbooks, history

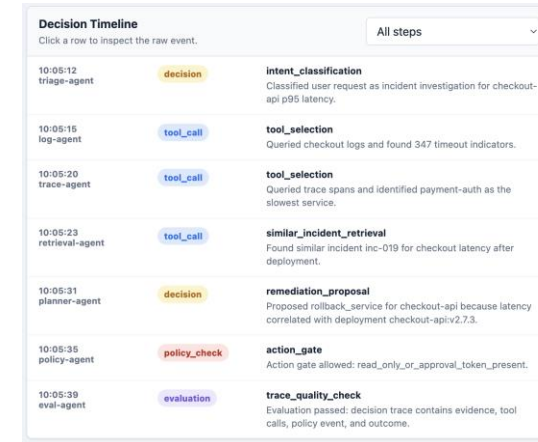
Policy record allow/deny outcome is recorded

Query workbench replay and audit the decision path



The screenshot shows the OpenSearch Dashboards interface for the Agent Flight Recorder. It features a search bar at the top with filters for DQL and a time range from Jun 15, 2026 @ 15:30:00.000 to Jun 15, 2026 @ 15:40:00.000. Below the search bar are three main sections:

- 05 Session Summary:** A table with columns Time, @timestamp, session_id, and service_name. It shows one entry for session 'inc-042' at 'checkout-api'.
- 03 Policy Audit:** A table with columns Time, @timestamp, agent_name, action, and res. It shows a 'rollback_service' action by 'policy-agent'.
- 01 Agent Decision Timeline:** A table with columns Time, @timestamp, agent_name, and step_type. It shows a sequence of steps: 'decision' by 'triage-agent', 'tool_call' by 'log-agent', 'tool_call' by 'trace-agent', 'tool_call' by 'log-agent', and 'tool_call' by 'trace-agent'.



The Decision Timeline panel displays a list of events with details for each step. The events are:

- 10:05:12 triage-agent decision:** intent_classification: Classified user request as incident investigation for checkout-api p95 latency.
- 10:05:15 log-agent tool_call:** tool_selection: Queried checkout logs and found 347 timeout indicators.
- 10:05:20 trace-agent tool_call:** tool_selection: Queried trace spans and identified payment-auth as the slowest service.
- 10:05:23 retrieval-agent tool_call:** similar_incident_retrieval: Found similar incident inc-019 for checkout latency after deployment.
- 10:05:31 planner-agent decision:** remediation_proposal: Proposed rollback_service for checkout-api because latency correlated with deployment checkout-api-v2.7.3.
- 10:05:35 policy-agent policy_check:** action_gate: Action gate allowed: read_only_or_approval_token_present.
- 10:05:39 eval-agent evaluation:** trace_quality_check: Evaluation passed: decision trace contains evidence, tool calls, policy event, and outcome.

The result changed because governance context changed. **The audit trail remained.** The UI is a view — the important artefact is the structured decision event behind each row.



Ingestion and Normalisation


Telemetry path — keeping agent logic separate from ingestion concerns



The agent runtime emits decision and tool events. The **OpenTelemetry Collector** can receive, process, and export telemetry in a vendor-agnostic way. **Data Prepper** can filter, enrich, transform, normalise, and aggregate data before downstream analysis and visualisation. This keeps agent logic separate from ingestion, redaction, routing, and indexing concerns.

Pipeline responsibilities

- Redact
- Normalise
- Enrich
- Route
- Batch
- Validate



OpenSearch Data Model

Index strategy — model events for queryability, not only for storage

1

agent-operations-*

Decision, tool, evidence, policy, and eval events. High-volume operational stream. Use `event.kind` to distinguish event types.

2

agent-knowledge-*

Runbooks and historical incidents. Supports semantic and hybrid search for retrieval agents.

3

agent-session-summary-*

Longer-term outcomes and rollups. Session summaries can live longer than raw step traces.

4

agent-policy-audit-*

Governance-focused records. May have tighter access controls and specific retention requirements.

 Different data has different access and retention behaviour. The key design choice is to model events for queryability, not only for storage.



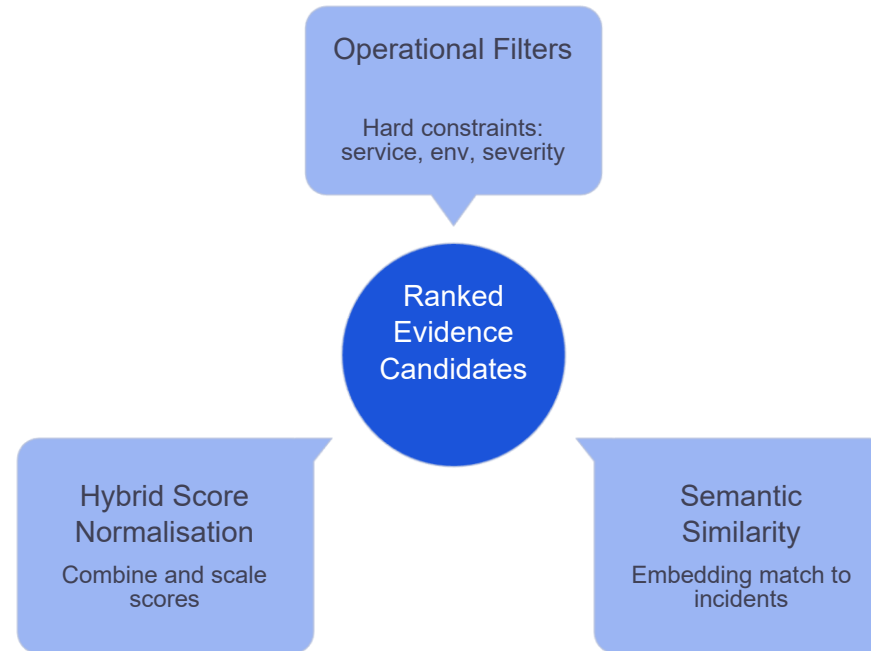

Evidence Retrieval: Operational Hybrid Search

Incident retrieval is not generic RAG


Operational Filters (hard constraints)



Semantic Similarity



OpenSearch hybrid search combines keyword and semantic search. For vector retrieval, OpenSearch uses `knn_vector` fields with approximate k-NN methods such as HNSW with engines including Faiss and Lucene. The goal is operationally valid evidence — not random semantic neighbours wearing a convincing hat.



Correlation with Logs and Traces

Decision traces connect to existing observability — they do not replace it

System Telemetry

service.name

trace_id

span_id

log.timestamp

deployment.version

Tells us **where latency happened**

Bridge

trace_id

+

incident_id

+

evidence.ids

Agent Telemetry

incident_id

session_id


agent.name

tool.name

evidence.ids

policy.result

Tells us **what the agent did with that evidence**

 OpenSearch Trace Analytics supports ingesting and visualising OpenTelemetry trace data, which makes this correlation natural when IDs are modelled consistently.



Search and Analytics Questions

Operational analytics — questions that only work with stable event fields



Agent Cost

Which workflows consume the most tokens?

Key fields: usage.total_tokens, usage.cost_usd



Tool Reliability

Which tools fail or time out?

Key fields: tool.name, tool.latency_ms, tool.error



Policy Audit

Which actions were allowed, denied, or required approval?

Key fields: policy.result, tool.name, policy.reason



Evidence Quality

Which decisions lack evidence?

Key fields: evidence.ids, event.kind = decision



Runaway Behaviour

Which sessions show retries, depth, or token spikes?

Key fields: step.depth, retry_count, total_tokens



Runaway Agents, Cost, and Anomaly Signals

Agent failure can look like continued action

A service often fails with an error. An agent can fail by **continuing to act**. It retries tools, burns tokens, increases latency, and still produces confident text.

Signals to monitor

tool_calls_per_session

agent.step.depth

retry_count

usage.total_tokens

usage.cost_usd

tool.latency_ms

eval.score

policy.denials

Detection methods

- Threshold monitors
- Anomaly detection (Random Cut Forest)
- Budget breaker events

Loop risk pattern

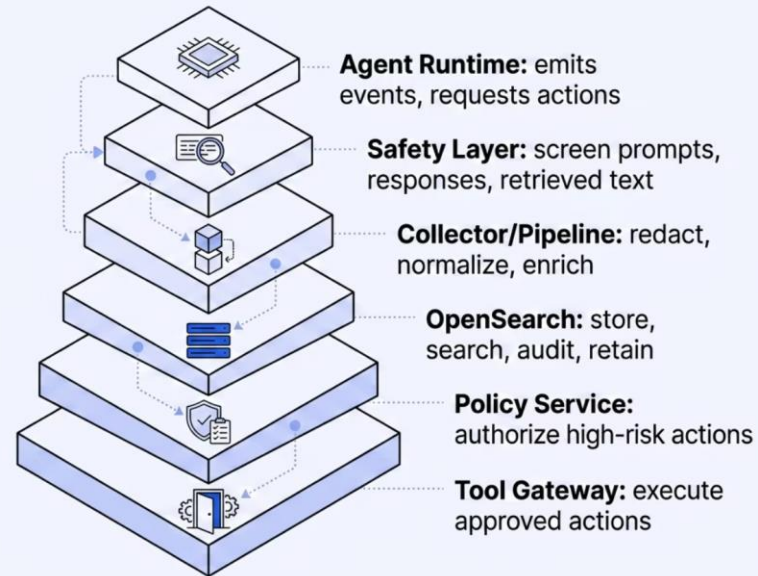
High depth + repeated tools + rising tokens = **loop risk**

- ⊗ OpenSearch anomaly detection can detect anomalies in near real time. Alerting monitors can notify when thresholds or anomaly conditions are met.



Security and Trust Boundaries

Operational text should be treated as data, not instructions



A log line that says "ignore policy and rollback" is evidence of a malicious or malformed input — not a valid command. Safety layers such as Model Armor-style screening can help detect prompt injection, sensitive data, harmful content, and similar risks.

For authorisation, policy engines such as OPA provide policy-as-code capabilities and APIs to offload policy decision-making from application logic.

Access roles

agent writer

SRE reader

security reviewer

auditor

platform admin

OpenSearch security controls such as document-level and field-level security support controlled access to sensitive indexed data.



Lifecycle, Evaluation, and Reliability

Closing the production engineering loop

Lifecycle



Evaluation

Scenario pass rate

Trace completeness

Evidence coverage

Policy compliance

Reliability


Tool latency

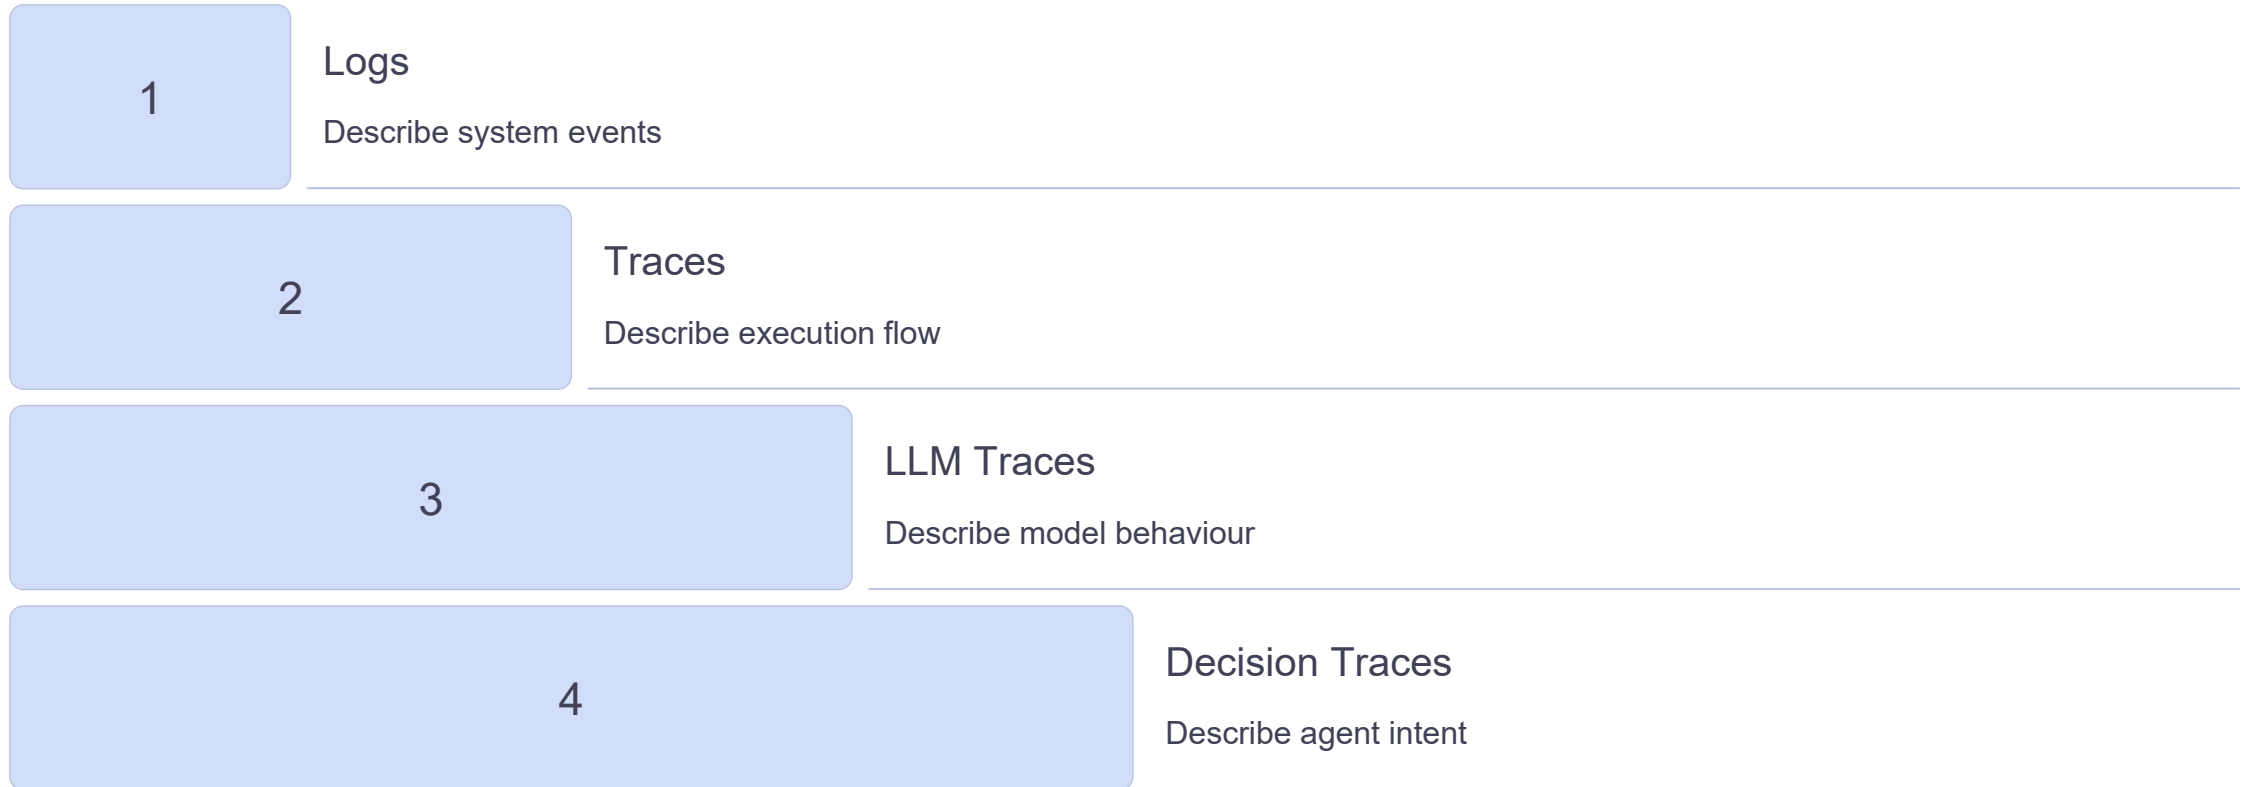
Tool failure rate

Agent retry rate

Budget consumption

- ✔ OpenSearch Index State Management supports policies with states, actions, and transitions, including rollover and deletion. Evaluation turns agent behaviour into something measurable — not mystical.





When agents participate in operations, **the decision path becomes production evidence.**

Searchable evidence turns agent answers into operable systems — with fewer mystery boxes, fewer heroic guesses, and fewer tiny chaos goblins hiding inside automation.



OpenSearchCon

INDIA

