TECHNICAL ARCHITECTURE / CTO, CISO & SECURITY ENGINEERING

AS-OS Architecture Deep Dive

Autonomous Security Operating System

A governed runtime for security decisions, authorized action and verifiable outcomes.

5 ORCHESTRATION LAYERS
8 RUNTIME SYSTEMS
1 GOVERNED RUNTIME KERNEL



GOVERNED RUNTIME KERNEL

CLAWOLF
ART OF GOVERNED AUTONOMY

CLAWOLF AS-OS
AUTONOMOUS SECURITY OPERATING SYSTEM

ARCHITECTURE MODEL

Five Layers, Eight Runtime Systems

The layers define the operating lifecycle. Runtime systems provide distinct technical capabilities.

Telemetry & Data Fabric
Source signals and provenance

UNIQUE METHOD **Domain Intelligence Layer**
Sector context and constraints

UNIQUE METHOD **Autonomous Decision Engine**
Analysis and decision packages

UNIQUE METHOD **Decision DNA**
Linked decision lineage

GOVERNED RUNTIME KERNEL

UNIQUE METHOD **Autonomous Trust Engine**
Policy and execution authority

UNIQUE METHOD **Autonomous Validation Fabric**
Independent checks and evidence

Continuous VulnOps
Remediation and revalidation

UNIQUE METHOD **Autonomous Customer Deployment**
Placement and lifecycle readiness

DETECTION DECISION ACTION CONTROL LEARNING

Investigation crosses layers. The kernel is the substrate, not a ninth system or a sixth layer.

PRIMARY ORCHESTRATION

Five-Layer Operating Contract

Control governs action eligibility before dispatch and evaluates the resulting evidence afterward.

Layer	Operational Job	Produced / Retained Evidence
Detection	Normalize source observations and identify suspicious activity.	DetectionSignal and source provenance
Decision	Correlate context, assess risk and investigate the case.	DecisionBundle, graph and evaluator proofs
Action	Prepare bounded intents and dispatch permitted actions.	ExecutionIntent and ActionReceipt
Control	Apply policy, approvals, verification and rollback governance.	ControlVerdict, VerificationResult and audit references
Learning	Record verified outcomes and recurrence feedback.	LearningEvent and revalidation lineage

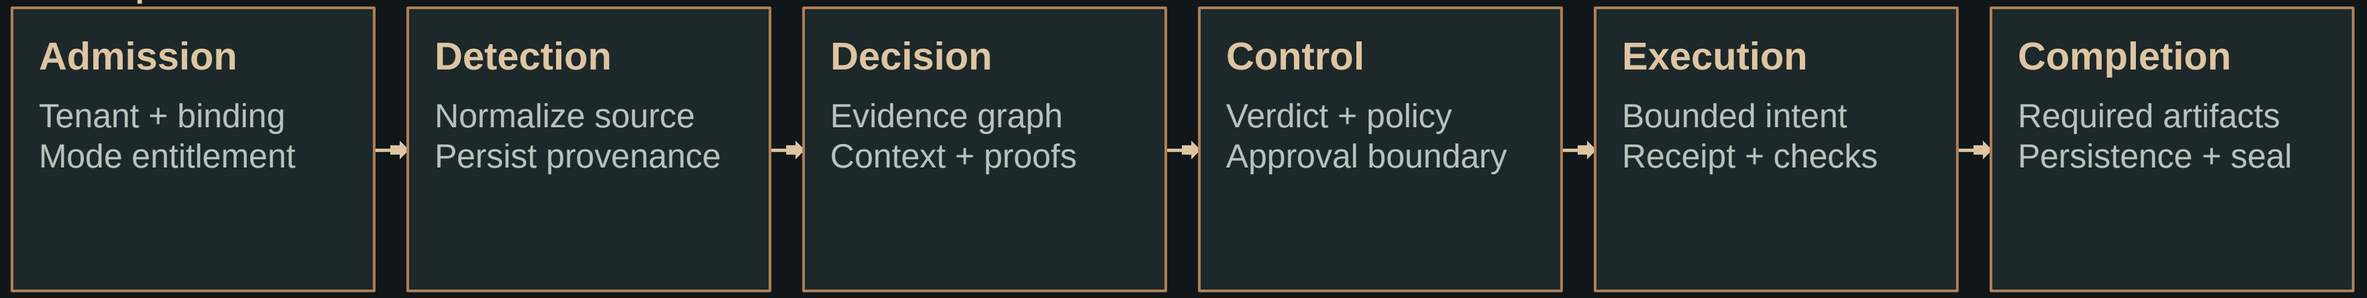
The lifecycle order describes responsibility. It does not permit Action to bypass Control.

GOVERNED RUNTIME KERNEL

Canonical Runtime & Completion Gates

Customer-facing execution enters the same runtime authority and retains a trace of each required stage.

POST /api/asos/runtime/run



awaiting_approval

Pause without dispatch. Resume only after a valid approval.

blocked / failed

Retain the blocker, failing stage and available evidence.

completed

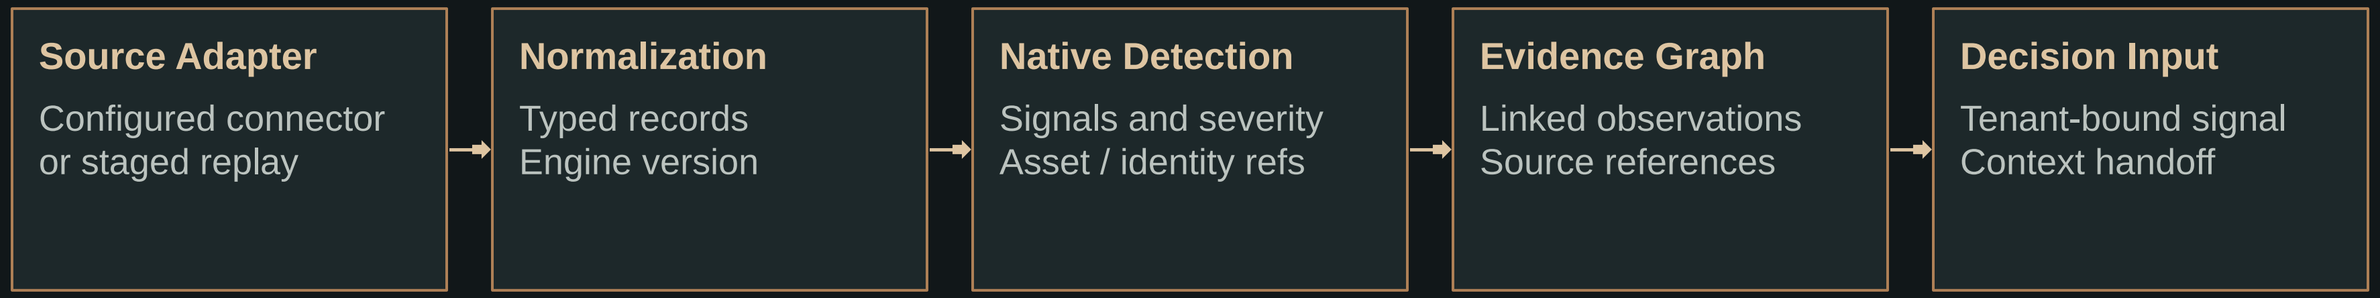
Required artifacts satisfy the selected mode contract.

A completed run does not imply an executed action or verified containment. Read the terminal artifacts.

RUNTIME SYSTEM 01

Telemetry & Data Fabric

A traceable input boundary turns heterogeneous observations into typed detection evidence.



PROVENANCE ENVELOPE

Identity	Reproducibility	Evidence Link
tenantId / runId **source record references**	Source hash + timestamp range Normalization / detector versions	Provenance ID + integrity hash Audit reference

Replay ground truth supports evaluation only. It must not supply the runtime verdict.

RUNTIME SYSTEM 02

UNIQUE METHOD

Domain Intelligence Layer

Business context constrains the decision before AS-OS selects a response.

Knows your business. No blind decision.

01 Organization Facts
Sector, jurisdiction, regulated activity and deployment

02 Applicability Evaluation
Calculated requirements, manual selections and missing facts

03 Runtime Context
Operating model, business semantics and threat context

04 Decision Constraints
Applicable policy and sector response restrictions

98

Canonical regulations & instruments

US / EU / Türkiye
Middle East / GCC

Applicability and content readiness remain separate. Missing facts stay explicit.

Catalog inclusion establishes context, not compliance, certification or 98 fully automated control packs.

RUNTIME SYSTEM 03

UNIQUE METHOD

Autonomous Decision Engine

21 versioned deterministic evaluators contribute domain analysis, reasoning and quorum evidence.

01 DDoS	02 Web	03 Endpoint	04 Network	05 Identity	06 Email	07 Cloud
08 Database	09 API	10 DLP	11 Ransomware	12 Supply Chain	13 Zero-Day	14 OT / ICS Safety
15 Mobile	16 Physical	17 Governance	18 Insider Threat	19 Adversarial AI	20 Kill Chain	21 Brain Latency

Per-Core Proof
ID + version + inputs
Verdict + confidence + refs

→

Aggregation
Applicable contributions
Double-counting controls

→

Decision Package
Risk decomposition
Actions + rejected options

Logic cores are deterministic heuristics. Brain Latency is a meta-evaluator, not a threat domain.

DECISION MECHANICS

Quorum, Uncertainty & Explainability

A passing quorum contributes evidence. Execution still depends on policy, authorization and readiness.

HOSTILE ACTIONABLE VERDICTS

hostileVotes / actionableCores

Illustrative gate calculation

6 / 10 = 0.60

Quorum passes a 0.48 configured threshold.
This example is not a benchmark result.

Applicability

Exclude non-actionable and not-applicable votes from the actionable denominator.

Uncertainty

Account for disagreement and gaps in the evidence. High uncertainty can require review.

Explainability

Keep contributing cores, reasons, rejected alternatives and formula versions.

AI-assisted analysis and deterministic execution authority remain distinct.

RUNTIME SYSTEM 04

UNIQUE METHOD

Autonomous Trust Engine

The authority plane turns policy, risk and human obligations into a bound ControlVerdict.

Identity & License

Tenant, role, binding and execution scope

Policy & Risk

Action taxonomy, uncertainty and protected targets

Human Obligation

Operating mode, severity floor and approval requirement

Dispatch Validation

Required checks and a verdict bound to the intent

HITL SEVERITY THRESHOLD

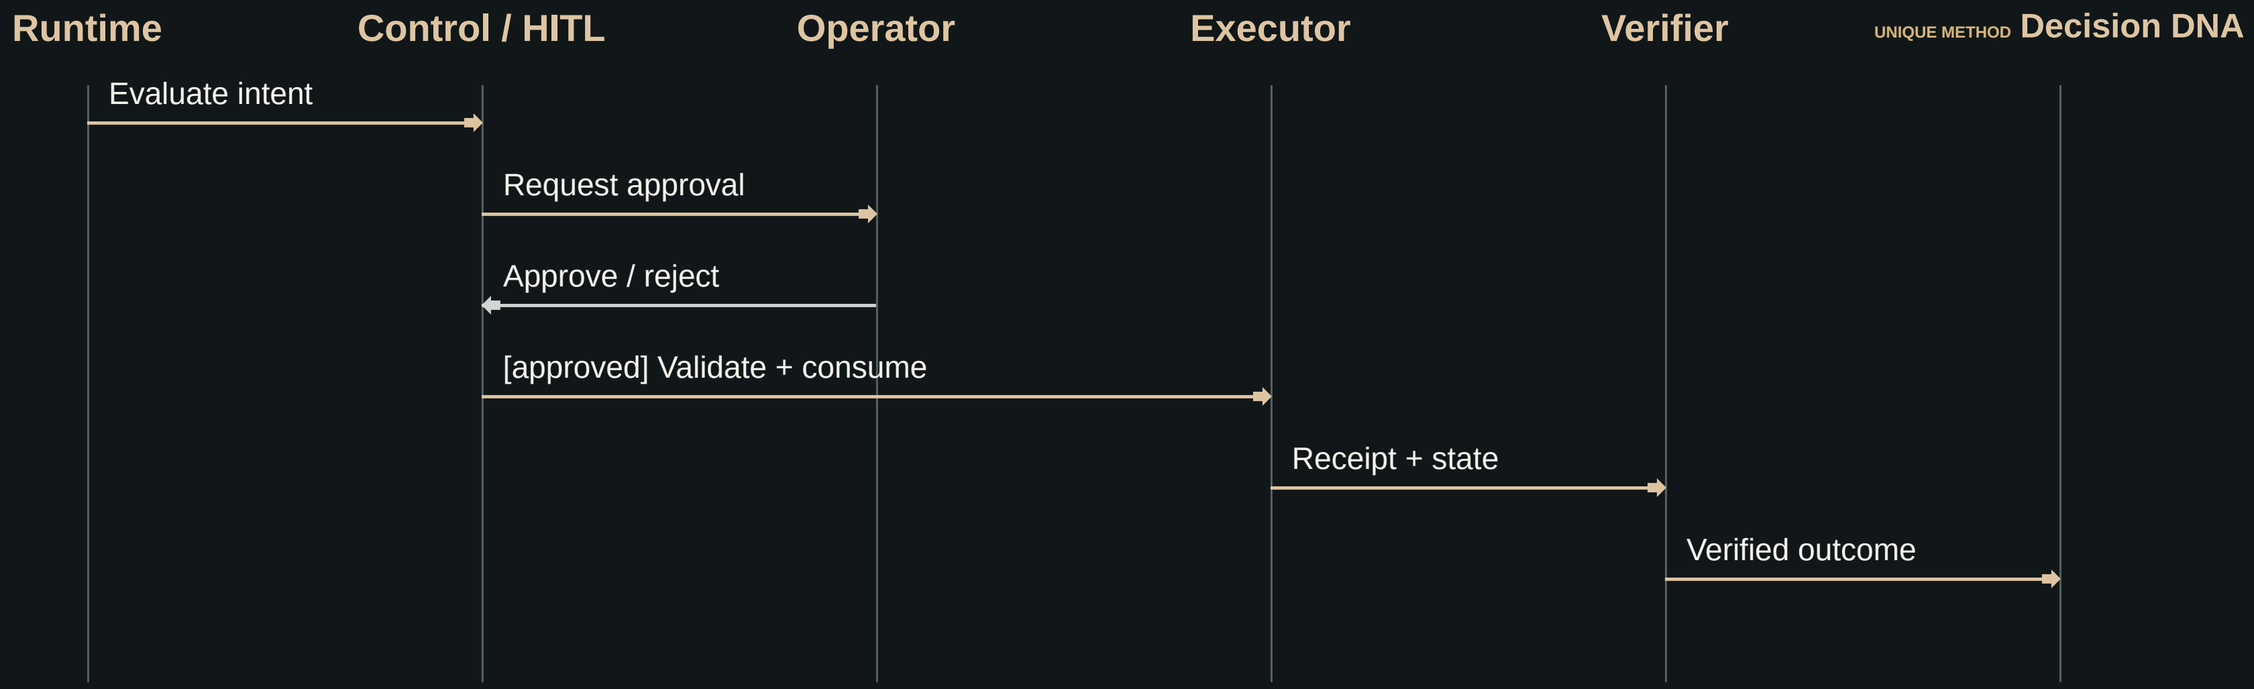
Selected Preset	Human Approval Required
Low	Low, Medium, High, Critical
Medium	Medium, High, Critical
High	High, Critical
Critical	Critical

Low escalates all incident severities. Other mandatory controls can still require approval below the selected threshold.

ControlVerdict: allow / deny / require_human. Full Autonomous never disables mandatory gates.

Approval, Dispatch & Outcome Verification

A human approval satisfies an obligation. It does not create an execution or containment result.



On execution failure: preserve the human acknowledgement, the error and the unfulfilled action.

RUNTIME SYSTEM 05

UNIQUE METHOD

Autonomous Validation Fabric

Independent checks return evidence and an explicit result, including unavailable or failed.

12 VALIDATION DOMAINS

Malware	**Phishing**	**Email**	**Cloud**
Network	**Identity**	**Document**	**Container**
Mobile	**Script**	**Web Application**	**IoT**

SUPPORTED EXECUTION PROFILES

Parse
Artifact and memory / container inspection

Execute
Isolated behavior and malware execution

Network Emulate
Disposable scenario with bounded networking



12 domains describe routing scope. Execution requires a supported profile, enrolled machinery and authority.

RUNTIME SYSTEM 06

UNIQUE METHOD

Decision DNA

A revisioned trust aggregate links context and intent to human decisions, action and observed effects.



*Approval-required paths only.

01 Context Snapshot

02 Evidence Pack

03 Policy Record

04 Reasoning & Logic Cores

05 Confidence & Trust

06 Intent & Outcome

07 Validation & Verification

08 Human Decision Record

09 Execution & Rollback

10 Learning Revision

tenantId + runId + decisionId + revision sequence + previous hash + current hash

DecisionBundle & InvestigationBundle

Runtime outputs preserve the reasoning and recommended next steps an analyst needs to review the case.

DECISIONBUNDLE

- Executive summary
- Decision, confidence and risk score
- Evidence summary and graph reference
- Intent analysis
- 21 logic-core outputs
- Sector intelligence context
- Context intelligence
- Uncertainty assessment
- Recommended next action
- Audit reference

INVESTIGATIONBUNDLE

- Root cause
- Timeline
- Related assets and identities
- Related events
- Attack path
- MITRE mapping
- Evidence expansion
- Investigation findings
- Recommended actions

Runtime Proof Dossier connects package outputs to the trace, action receipts, verification and integrity evidence.

RUNTIME SYSTEM 07

Continuous VulnOps

A finding remains linked to the remediation decision, verification and subsequent recurrence checks.

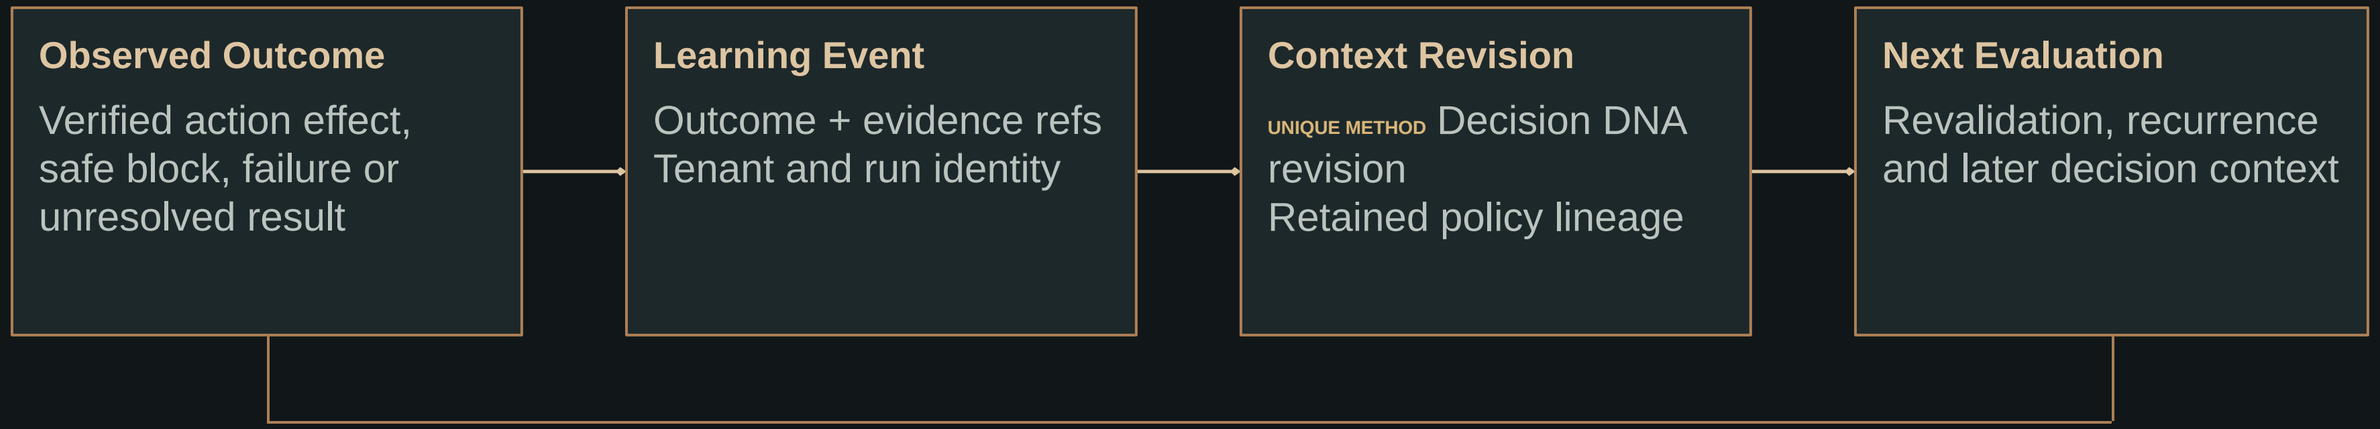


22 capability authorities span this lifecycle, including recovery, report export and audit integrity.

LEARNING LAYER

Outcome Feedback Without Authority Drift

AS-OS retains outcome evidence for later decisions and revalidation while preserving the execution boundary.



No automatic widening of execution scope

A feedback record cannot grant a license, change a human-approval obligation or make an unsupported action executable.

RUNTIME SYSTEM 08

UNIQUE METHOD

Autonomous Customer Deployment

A configuration and lifecycle system that keeps placement, engagement and operating authority separate.

Customer Inputs
Placement + environment
Identity + secrets refs

→

Configuration
Ingestion + runtime contract
Package and governance

→

Readiness
Required dependencies
Validation + enrollment

→

Lifecycle
POC / production context
Release and recovery

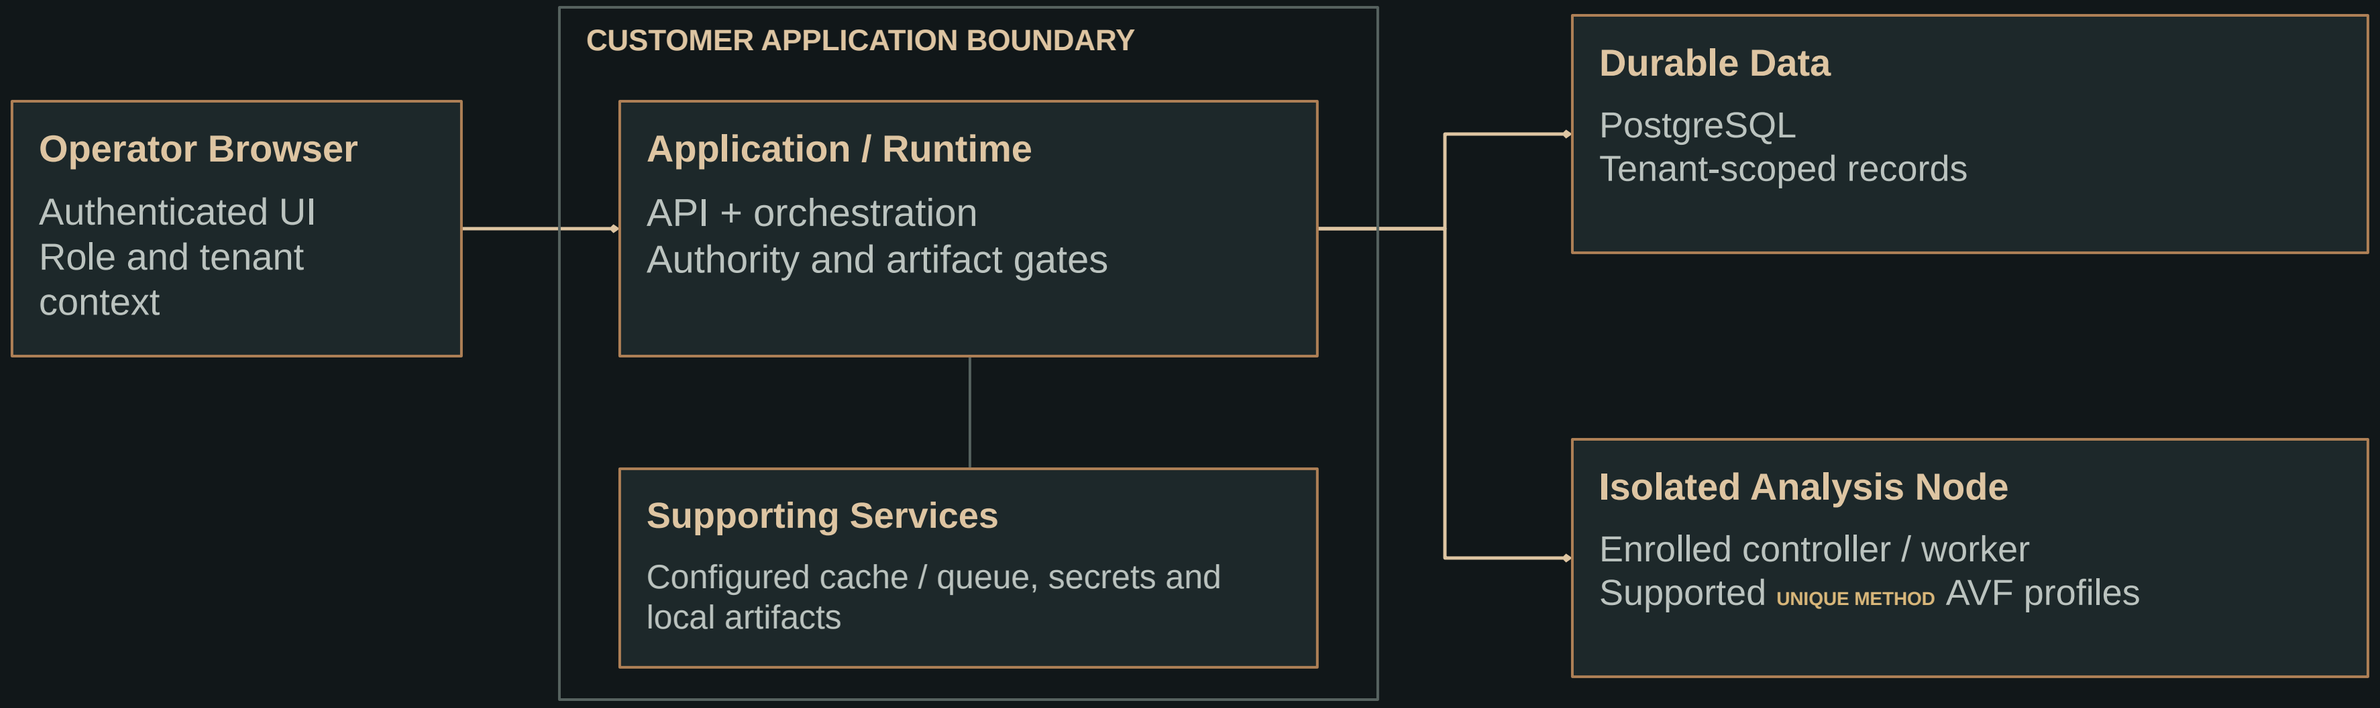
Independent Dimension	Available Context	What It Does Not Grant
Placement	Cloud / on-prem / hybrid / air-gapped	Execution permission
Engagement	POC or production lifecycle	Broader action authority
Operating Control	Observation / Alert / HITL / Full Autonomous	A bypass of policy or human obligations

Update & Release Manager retains verification, preflight, approval, deployment and restoration evidence.

DEPLOYMENT ARCHITECTURE

Application, Data & Validation Boundaries

The topology separates the operator application, durable data and isolated analysis machinery.



Cloud, on-premises and air-gapped placements select different connectivity and package inputs.

WORKFLOW CONTRACTS

Six Runtime Workflows

Each workflow declares its required stages and mandatory artifacts before the runtime can report completion.

Runtime Workflow	Technical Scope	Distinct Completion Evidence
Decision	Evaluate supplied detection signals	Graph, decision, core proofs, quorum, trace and seal
Decision + Investigation	Add investigation to the decision	Decision and investigation report packages
Detection + Decision + Investigation	Run native detection before analysis	Raw provenance plus the investigation output
Decision + Investigation + Action	Add controlled execution	Intent, receipt or safe block, conditional verification
Governed Full Loop	Full governed runtime with learning	Verification, learning and Runtime Proof Dossier
Continuous VulnOps Full Loop	Finding to remediation and revalidation	Exploitability, plan and revalidation evidence

Workflow selection does not override daily operating mode, binding entitlement or the action taxonomy.

TRUST BOUNDARIES

Tenant Identity, Entitlement & Execution Authority

AS-OS resolves the requesting identity and tenant before it admits work or exposes evidence.

Visibility
Can the role discover and open the surface?

≠

Licensing
Does the customer own the capability domain?

≠

Entitlement
Is the binding active for this tenant and workflow?

≠

Execution Authority
May this exact action run under current controls?

Persistent tenant identity

Requests, database transactions, jobs and evidence retain tenant ownership. A browser selection never replaces server-side authorization.

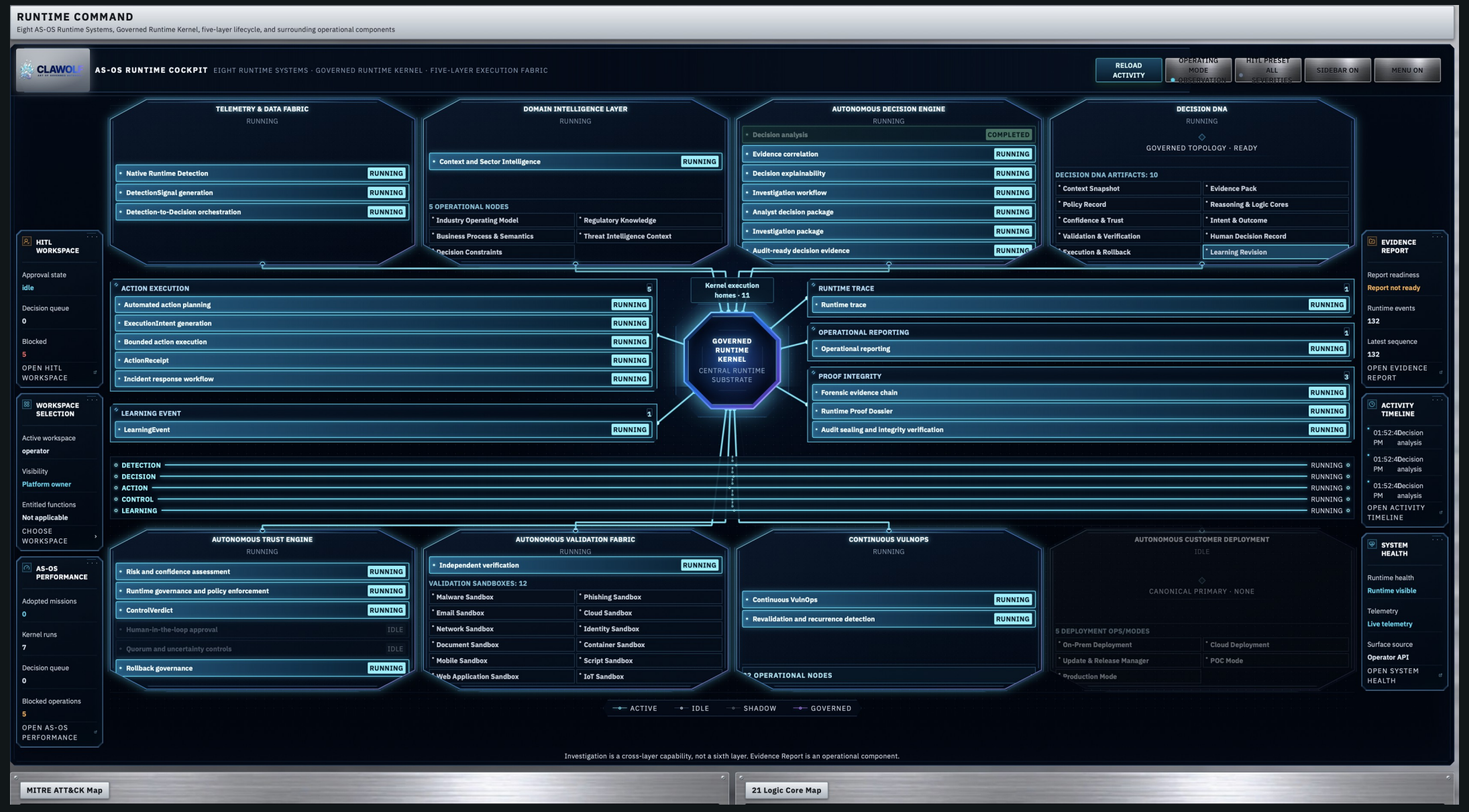
Fail-closed artifact access

Approval, receipt and report reads use the same scoped identities. Foreign-tenant references cannot establish evidence for this run.

OPERATIONAL OBSERVABILITY

Runtime Command & Operator Evidence

The UI projects canonical activity and persisted artifacts instead of inventing execution state.



Run Selection

Latest admitted run and selected tenant remain explicit.

Feature Activity

Canonical function IDs and retained runtime events drive status.

Analyst Surfaces

HITL, Mission Workspace, reporting and trace reference the same artifacts.

CROSS-SYSTEM WALKTHROUGH

Suspicious Identity Activity: Reference Sequence

An architecture example showing how investigation, approval and action share one evidence lineage.

1 / SIGNAL

Identity observation

Telemetry retains actor, target, source reference and canonical severity.

2 / CONTEXT

Business constraint

Domain Intelligence adds sector and protected-asset context.

3 / DECISION

Evidence-backed proposal

ADE evaluates applicable cores and proposes a bounded response.

4 / AUTHORITY

Explicit human boundary

UNIQUE METHOD Trust Engine returns require_human when policy demands approval.

5 / OUTCOME

Verified effect or failure

The configured executor returns a receipt. Verification records the actual state.

6 / PROOF

Linked final record

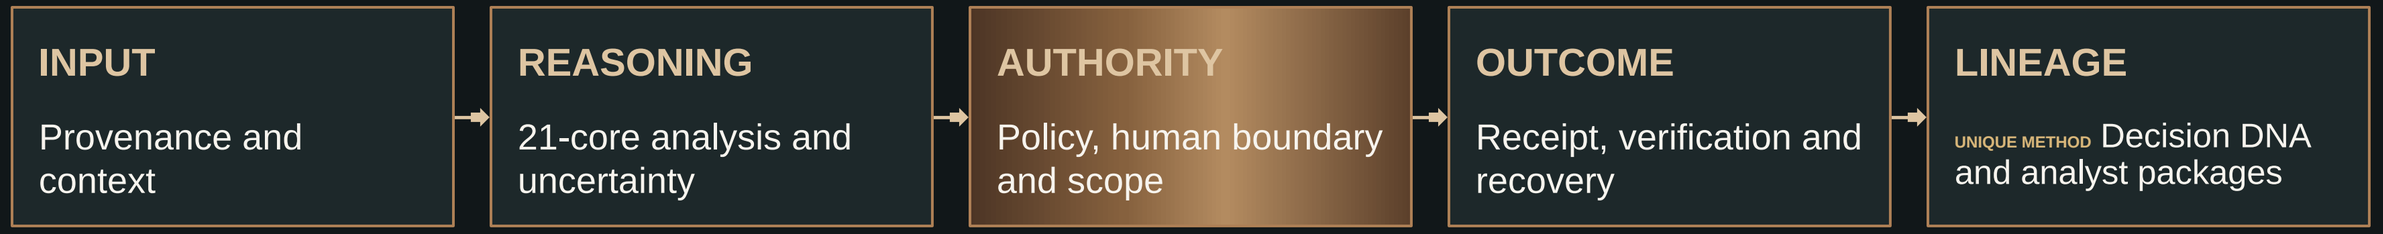
UNIQUE METHOD Decision DNA retains the approval, execution result and remaining gaps.

Illustrative sequence, not a recorded run. Approval, execution, containment and recovery are distinct outcomes.

ARCHITECTURE SUMMARY

A Complete Runtime Evidence Chain

Technical evaluation connects the source observation to the decision, authority, action and measured effect.



Deployment review

Confirm source adapters, executor bindings, provider eligibility, isolated analysis enrollment and restoration paths for the selected environment.

Operational review

Measure workload-specific latency, throughput, recovery and user journeys on the exact delivered version. No architecture diagram substitutes for that evidence.