ENKIDU

—— MIOVSKI GROUP ——

CONCEPTUAL ARCHITECTURE

# Lamassu

## A Two-Tier Urban-Sensing Platform

A tiered-intelligence platform for ambient urban sensing — sensor-tuned micro-models and local adjudication hosted at regional Site IFCs, deep all-source correlation and fine-tuning hosted at one Central IFC, ingesting only from already-existing and already-deployed sensor technologies, structured by Tangram, and governed by a strong policy framework and Origi, the Trust Assurance Framework.

VERSION 4.4 · FOURTH REVISION
4 JULY 2026
PREPARED BY MIOVSKI GROUP
CLASSIFICATION · CONFIDENTIAL · CONCEPTUAL ARCHITECTURE

# Contents

# Glossary

Terms below are used throughout this document with the specific meanings defined here. The Site IFC / Central IFC tiering follows the equivalent tier concepts of a sibling defence architecture (Northwind's FOB-IFC and Central IFC), recast for civilian, non-kinetic use. The trust harness is a deployment of Enkidu's Origi Trust Assurance Framework, also used in that sibling architecture.

### Site IFC (Site Intelligence Fusion Center)

A Lamassu-operated regional fusion site serving one geographic or organizational service area. It ingests directly from the already-existing, already-deployed sensor technologies within its area, runs sensor-specialist models and a local coordinator, and holds its own local deterministic adjudication engine — so it commits a regional picture on its own authority. It reaches back to the Central IFC for enterprise-wide correlation, fine-tuned model updates, and direction, but is not dependent on that link: if reachback is cut, the Site IFC continues ingesting, inferring, and committing on its own, and reconciles with the wider picture when the link returns. No Lamassu-owned infrastructure exists below the Site IFC — sensors and the systems that operate them are already-existing and independently supplied.

### Central IFC (Central Intelligence Fusion Center)

The single rear, enterprise-wide tier that every Site IFC reaches back to. It receives regional pictures and observations from all Site IFCs (and may itself directly ingest already-existing enterprise-wide or cross-jurisdiction sources), runs the largest models for deep, cross-site all-source correlation, fine-tunes the models deployed to every Site IFC, holds corporate memory, and issues direction — retasking, priority changes, updated models — back down to the Site IFCs. Its own local adjudication engine commits indicators that span or exceed the scope of any single Site IFC.

### Already-existing sensor technology

Sensor hardware, sensor platforms, or sensor-data systems — cameras, IoT devices, OT/facility instrumentation, and the platforms that operate them — that are already installed and operated by a city, facility, or third party, independent of and prior to Lamassu. A Site IFC connects to their outputs, either directly over the sensor's native protocol or through the existing system that already manages them, but Lamassu never specifies, procures, or deploys sensor or processing infrastructure of its own below the Site IFC.

### Sensor-specialist micro-model

A compact language or multimodal model (≈1–8B parameters), fine-tuned for the telemetry vocabulary of a single sensor modality, served at the Site IFC that ingests that modality's already-existing sources. One specialist owns one modality. Specialization is preserved as a modelling pattern, and its execution now sits at the regional tier closest to the sources it reasons over.

### Coordinator

A reasoning model that plans a run, dispatches work to sensor specialists, composes their findings across modalities, and frames a candidate indicator. A **site coordinator** composes within one Site IFC's regional scope; the **cross-site coordinator** at the Central IFC composes across Site IFCs. Neither ever assigns the final verdict.

### Tangram

The platform layer. A zoning architecture that bounds operations into zones with explicit entry/exit conditions and supplies the observability pattern across every ingested source and every tier. Zones are logical and data-domain constructs, and a zone's membership can span both tiers — a Site IFC and the Central IFC — with Tangram tracking which tier(s) a zone's activity is currently resolved at. Each zone is also a three-layer governed execution context (policy / conops / reporting).

### Policy framework

The set of machine-readable guardrails Tangram attaches to every agent and zone, and that the trust harness evaluates on every action, at either tier. It states what an agent may do — never what it should do — and denies by default.

### Trust Assurance Framework / trust harness (Origi)

Enkidu's policy-based, evidence-backed framework for asserting and measuring trust in AI-augmented systems, branded **Origi**, and the runtime governance layer through which every agent action passes, at both the Site IFC and the Central IFC. Origi assigns trust-operation zones to agents, controls sensor and actuator access, secures every ingestion link, every reachback link between a Site IFC and the Central IFC, and every internal service-to-service link at each tier, and captures the chain of custody used to score confidence. It applies a zero-trust posture — nothing an agent produces is trusted by default — enforced through checkpoints at every trust boundary, and a strict separation of duties between the components that observe, judge, and act.

### Assurance case (Claim → Argument → Evidence)

Origi's structure for making a trust position checkable rather than merely asserted: a **claim** bound to a policy with an explicit acceptance threshold, decomposed by an **argument**, and supported by **evidence** carrying a signed, tamper-evident provenance record. A claim is satisfied only when both the measured outcome meets its policy threshold and the supporting evidence chain validates. Every committed indicator in Lamassu is itself an assurance case (Section 5.3).

### Adjudication engine

Deterministic, code-enforced logic — not any model — that converts proposed findings into a committed, operator-facing indicator under bound policy, applying Origi's assurance-case structure: the candidate indicator is the claim, the specialists' and coordinator's composition is the argument, and the chain of custody is the evidence. A commit occurs only when both the claim meets its policy threshold and the evidence chain validates. Each Site IFC holds its own adjudication engine for its regional picture; the Central IFC holds one for cross-site indicators. The model proposes; the engine commits. The architectural keystone at both tiers.

### Adjudication algorithm

The specific confidence-aggregation function, per-mission acceptance thresholds, conflict-resolution rule, and partition-reconciliation rule that run inside an adjudication engine's assurance-case gate. Distinct from the engine itself: the engine's gate structure (Section 5.3) is a fixed architectural invariant; the algorithm's parameters are determined empirically through the simulation-driven methodology of Section 14.3 before detailed design proceeds.

### Typed handoff

A strictly-schematized object passed between agents, models, tools, or Tangram zones — including the handoff between a Site IFC and the Central IFC over reachback. No free-form data crosses a boundary; the receiver validates.

### Chain of custody / chain of evidence

Origi's provenance record applied to Lamassu: the signed, hash-linked record of every observation and inference — what was produced, how, by which system, service, or agent, at which tier, from what inputs, and when — used to compute confidence in a detection rather than trusting a number a model emits. Evidence carries an expiry, so stale evidence cannot satisfy a claim. Reconciliation after a reachback partition is resolved through this chain: divergence is surfaced, never silently overwritten.

### Data-sensitivity & privacy framework

The civilian classification scheme — Public, Internal, Sensitive, Personal Data — with privacy and handling caveats (data-residency, retention, consent basis, redaction, sharing limits) that the trust harness enforces at the point of distribution, at either tier.

### Trust dossier

The running record, per committed indicator, of its claim, bound policy, composing argument, and validated evidence chain — Origi's trace tree presented as a position an operator, auditor, or regulator can inspect on demand (Section 10.3).

### Civil actuator

A permitted, non-kinetic effect an agent may cue at the highest trust zone: steer or zoom a camera, retask or retune a sensor, raise an alert to the operator console, or push a notification to an integrated system. Any effect that reaches the outside world requires explicit human approval.

### High-availability (HA) redundancy

Redundant, load-balanced service instances with automated failover, operating independently within each Site IFC and within the Central IFC. HA redundancy protects each tier's own infrastructure from a server or service-instance failure; it is distinct from, and complementary to, the two-tier partition tolerance described below.

### Reachback

The link from a Site IFC up to the Central IFC. Observations and regional assessments flow up; direction, priority changes, and updated model weights flow down. A Site IFC is not dependent on reachback for its own regional commit.

### Partition tolerance / reconciliation

The property that a Site IFC continues ingesting, inferring, and committing its regional picture on its own authority if reachback to the Central IFC is cut, and reconciles its chain of custody with the wider picture when the link returns — surfacing any divergence rather than overwriting it.

# Abstract

Lamassu turns a city's already-existing sensors into operator-reviewable intelligence. It runs on two physical tiers, adapted from the equivalent forward/rear

tier concepts of a sibling defence architecture (Northwind's FOB-IFC and Central IFC), recast for civilian use: the **Site IFC** and the **Central IFC**.

A **Site IFC** is a regional fusion site serving one service area. It ingests directly from the already-existing sensor technologies already deployed within that area, runs sensor-specialist models and a local coordinator, and holds its own local adjudication engine — so it commits a regional picture under its own authority, and keeps doing so if its link to the Central IFC is cut, reconciling when the link returns. Many Site IFCs feed one **Central IFC** — the rear, enterprise-wide tier, which correlates across every Site IFC's regional picture with larger models, fine-tunes the models each Site IFC deploys, holds corporate memory, and issues direction back down. Lamassu deploys no sensor, no edge node, and no field-processing infrastructure of its own below the Site IFC: every already-existing sensor technology is ingested at the Site IFC serving its area, never at anything Lamassu itself installs in the field.

The architecture follows one premise: **determinism where it matters, reasoning where it helps, and inference tiered between site and centre.** All sensing data enters through the Site IFC serving its area; sensor-specialist micro-models and a site coordinator compose a regional candidate; the site's own deterministic adjudication engine commits the regional indicator. Regional pictures reach back to the Central IFC, where a cross-site coordinator and larger models correlate across every Site IFC and a central adjudication engine commits enterprise-wide indicators. Tangram structures both tiers; the policy framework bounds what every agent may do at either; **Origi**, Enkidu's Trust Assurance Framework, governs every agent and secures every ingestion link, every reachback link, and every internal link. Every commit an adjudication engine makes is itself an Origi assurance case — a claim bound to policy, argued by composition, and satisfied only when its supporting evidence chain validates — so the result is explainable, because every indicator at either tier resolves to signed provenance through a fully observable trace tree; resilient, because a Site IFC does not depend on the Central IFC to keep producing a committed local picture; and scalable, because the platform grows by adding Site IFCs to cover new service areas and by adding compute at either tier, never by deploying anything below the Site IFC.

**Lamassu is a civilian platform.** It informs and enables operators — utility and facility security teams, environmental and emergency-management agencies, transport and event operators — and it never commands a kinetic effect. The most consequential thing an agent may do unaided is raise an alert; any effect that reaches the outside world is gated behind a human decision.

# 1. Purpose and Scope

## 1.1 What this document is

Lamassu is an agentic urban-sensing platform: it decomposes sensing-to-insight analysis into cooperating agents that ingest, infer, compose, and adjudicate, rather than a single monolithic model or a fixed rules pipeline. This document describes that platform — its two physical tiers, its planes, its policy framework, and its trust model — as one architecture with mission-specific bindings, rather than separate builds per use case.

Lamassu runs on two physical tiers: a **Site IFC** serving each service area, and a single **Central IFC** that every Site IFC reaches back to. The tiers are adapted from the equivalent forward/rear tier concepts of a sibling defence architecture — Northwind's FOB-IFC and Central IFC — recast into two civilian tiers.

## 1.2 The three capabilities this architecture must deliver

Beyond its analytic core, Lamassu is required to deliver three operational capabilities. Each is named here, designed for in a dedicated section, and summarised against the architecture in Section 13.

| CAPABILITY | WHAT IT REQUIRES OF THE ARCHITECTURE |
|---|---|
| **Adaptive sensor ingestion** | The ability to connect to diverse, already-existing and already-deployed sensor technologies using standard protocols, at the Site IFC serving each area, and to keep producing a useful picture when individual sources degrade or go offline (Section 6). |
| **Explainable & actionable outputs** | Every detection, at either tier, carries enough context for an operator to understand what was detected, which sources contributed, and the confidence of the assessment (Section 10). |
| **Scalable architecture** | A design that grows by adding Site IFCs to cover new service areas, by ingesting a broader inventory of already-existing sources per area, and by scaling compute at either tier — and is architecturally compatible with international standards (Section 11). |

TABLE 1 · THE THREE OPERATIONAL CAPABILITIES LAMASSU MUST DELIVER, AND WHERE EACH IS DESIGNED FOR.

# 2. Design Premise and Principles

## 2.1 The core premise

Lamassu rests on one sentence: **determinism where it matters, reasoning where it helps, and inference tiered between site and centre.** Determinism owns anything that must be auditable — above all the committed indicator, which only an adjudication engine may write, at either tier. Reasoning is reserved for the genuinely hard part — composing many modalities into a candidate picture — and given to a coordinator: a site coordinator for a Site IFC's own regional scope, a cross-site coordinator for the Central IFC's enterprise-wide scope. Tiering by locality, alongside specialization by modality, is how the platform stays modular and explainable: each sensor family is owned by its own specialist model, hosted at the Site IFC closest to the already-existing sources it reasons over, while the Central IFC is reserved for correlation and reasoning that only makes sense across many sites at once.

## 2.2 The load-bearing invariants

Six invariants hold everywhere in the architecture, at both tiers.

| # | INVARIANT | WHAT IT MEANS IN LAMASSU |
|---|---|---|
| **1** | Determinism owns the verdict | The committed indicator — regional at a Site IFC, enterprise-wide at the Central IFC — is produced only by that tier's deterministic adjudication engine. No model — specialist or coordinator, at either tier — can write a committed verdict or its gating confidence. |
| **2** | Every handoff is typed | No free-form data crosses an agent, model, tool, or zone boundary — including the reachback handoff between a Site IFC and the Central IFC — only schematized objects. The receiver validates. |
| **3** | Provenance is mandatory and validated | Every finding carries the identifiers of the source telemetry it rests on, traceable through whichever tier produced it. References that do not resolve to actually-ingested, consented inputs are filtered before they can |

| # | INVARIANT | WHAT IT MEANS IN LAMASSU |
|---|---|---|
| | | reach either adjudication engine. |
| **4** | The topology is observable | Every agent boundary, model call, tool call, and engine decision emits a span, at either tier. A run is a trace tree, not a flat log, and a cross-tier run's trace resolves down through the Central IFC into the originating Site IFC's evidence. |
| **5** | A specialist sees only its modality | A sensor-specialist agent and its micro-model can read only their own modality's telemetry within their Site IFC and cannot write fused state. Scope is bounded structurally, not by convention. |
| **6** | Authority is explicit and human-owned | What every agent may do, at either tier, is set by a machine-readable policy it cannot widen; any effect that reaches the outside world is owned by a human. The policy layer of every zone is human (Section 8). |

TABLE 2 · THE SIX LOAD-BEARING INVARIANTS.

### 2.3 A civilian platform

Lamassu is, by design and without exception, a **civilian** urban-sensing platform. It serves operators such as utility and facility security teams, environmental and emergency-management agencies, transport authorities, and event and venue operators. It does not target, and it commands no kinetic effect. Every concept in this document is recast for civil use — including the Site IFC / Central IFC tiering itself, adapted from a defence fabric's forward/rear tiers into two civilian regional/enterprise tiers: sensor modalities replace intelligence disciplines; a data-sensitivity and privacy framework replaces security classification; and the only effects an agent may cue are non-kinetic civil actuators — steering a camera, retasking a sensor, raising an alert — each gated, for any external action, behind a human decision.

## 3. The Two Physical Tiers — Site IFC and Central IFC

Lamassu runs on exactly two physical tiers: a **Site IFC** for each service area, and a single **Central IFC** that every Site IFC reaches back to. Below the Site IFC, Lamassu deploys nothing of its own — it ingests only from sensor technologies that are already deployed and already operated by the city, facility, or third party it serves. The tiers themselves are adapted from the equivalent forward/rear tier concepts of a sibling defence architecture (Northwind's FOB-IFC and Central IFC), recast here for civilian, non-kinetic use.

## 3.1 A two-tier architecture, with a firm floor

A **Site IFC** is a regional fusion site: it ingests directly from the already-existing sensor technologies within one service area (a district, campus, facility cluster, or comparable footprint), runs the sensor-specialist models and coordinator for that area, and holds its own local adjudication engine. Many Site IFCs — one per service area — reach back to a single **Central IFC**, the rear, enterprise-wide tier that correlates across all of them, fine-tunes the models every Site IFC deploys, and issues direction back down.

The floor of the architecture is firm: no sensor, camera, or field device is ever wired into hardware that Lamassu deploys below the Site IFC, because Lamassu deploys none there. A Site IFC is the forward-most compute Lamassu operates — a proper compute site scoped to one service area — not a piece of field-adjacent edge hardware.

## 3.2 Two properties that make the Site IFC the right forward tier

- **Locality.** A Site IFC sits close to the already-existing sensor technologies of the area it serves, so specialist inference and the first, regional commit happen near the sources they reason over, without waiting on a round trip to a distant centre.
- **Independence from reachback.** A Site IFC reaches back to the Central IFC for enterprise-wide correlation and updated models, but it is not dependent on that link: if reachback is cut, the Site IFC sustains its own ingestion, inference, and commitment, exactly as the wider picture is reconciled once the link returns (Section 6.4, Section 9).

## 3.3 Representative compute classes

The figures below anchor the platform design as representative classes — not a bill of materials — and should be confirmed against current datasheets and workload sizing in detailed design.

| TIER | COMPUTE CLASS | ROLE IN THE PLATFORM |
|---|---|---|
| **Site IFC** | Mid-scale GPU server(s) (single or clustered, sized to the service area's sensor volume and model roster), sited at or near the service area | Ingestion from the area's already-existing sensor technologies; that area's sensor-specialist model serving; the site coordinator; the site's own deterministic adjudication engine; local short-term stores and the operator console for that area. |
| **Central IFC** | Data-centre GPU cluster (server-class, multi-PFLOP class, e.g. L40S / H100 / GB200-class), scaled horizontally as more Site IFCs are onboarded | Cross-site coordinator; the largest correlation models; the central adjudication engine for cross-site indicators; fine-tuning and corporate memory; long-term stores; the enterprise-wide operator console. |

TABLE 3 · THE TWO COMPUTE TIERS IN LAMASSU. COMPUTE CLASSES ARE ILLUSTRATIVE; CONFIRM AGAINST DATASHEETS AND WORKLOAD SIZING IN DETAILED DESIGN.

## 3.4 High-availability redundancy, and partition tolerance between tiers

Two distinct resilience mechanisms apply, and they protect against two distinct failures. **High-availability redundancy** operates independently within each tier: services at a Site IFC, and services at the Central IFC, each run as redundant, load-balanced instances behind automated failover, so the loss of a server or a service instance at either tier does not interrupt that tier's own ingestion, inference, or commitment.

**Partition tolerance** protects against loss of the reachback link itself, rather than loss of any tier's own infrastructure. If the link between a Site IFC and the Central IFC drops, the Site IFC keeps ingesting from its already-existing sources, keeps running its specialists and coordinator, and keeps committing regional indicators through its own adjudication engine — it does not wait on the Central IFC to function. When the link returns, the Site IFC's chain of custody reconciles with the wider, Central-IFC-held picture, and any divergence between what the Site IFC concluded alone and what the wider picture would otherwise show is surfaced to the operator rather than silently overwritten.

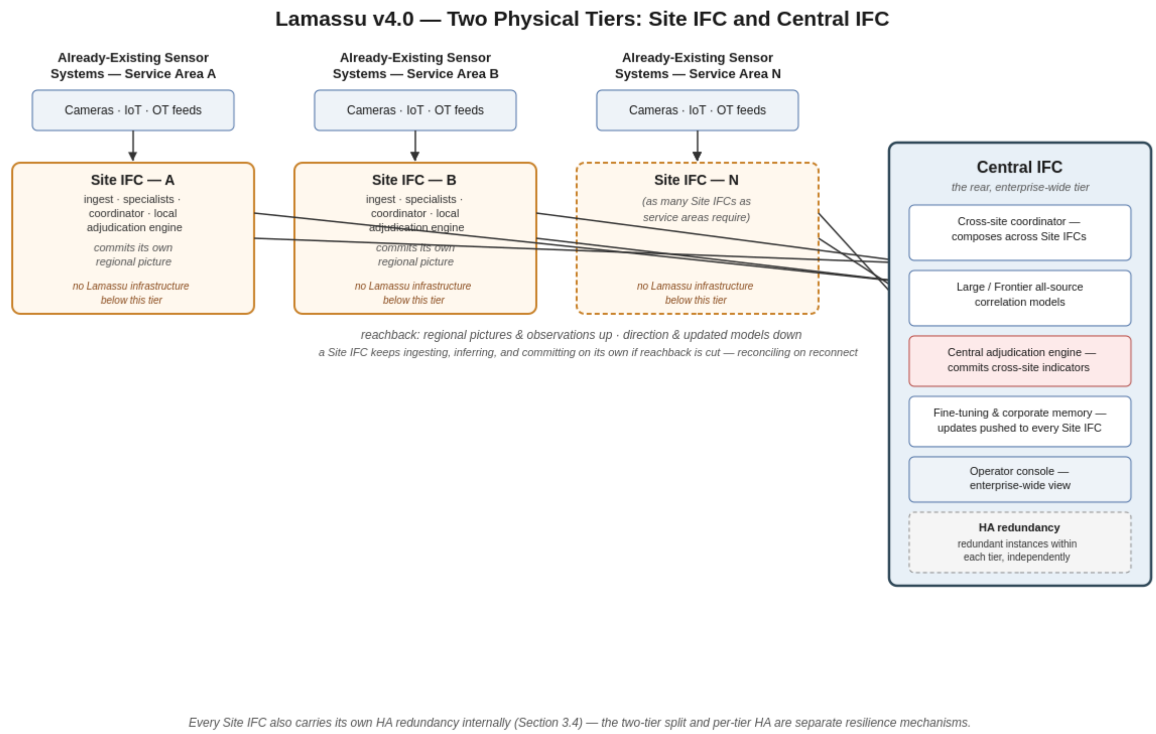


FIGURE 1 · THE LAMASSU ARCHITECTURE. ALREADY-EXISTING, ALREADY-DEPLOYED SENSOR SYSTEMS FEED EACH SITE IFC DIRECTLY; MANY SITE IFCS REACH BACK TO ONE CENTRAL IFC. EACH TIER CARRIES ITS OWN HA REDUNDANCY; THE TWO-TIER SPLIT ADDS PARTITION TOLERANCE ACROSS THE REACHBACK LINK.

## 3.5 A rapidly deployable configuration

A Site IFC may be realized as a fixed installation or as a **rapidly deployable configuration** — the same ingestion pipeline, sensor specialists, site coordinator, local adjudication engine, policy framework, and Origi, packaged for fast stand-up and relocation rather than for a permanent data-centre footprint. The two are functionally identical; only the physical packaging, power, and connectivity envelope differ. A deployable Site IFC still ingests only from already-existing sensor technologies already present at or near the site it serves — no sensor, camera, or field device is added below it, exactly as for a fixed Site IFC (Section 3.1).

The deployable configuration reuses the compute, ruggedization, and power envelope of the equivalent deployable tier in a sibling defence architecture (Northwind's FOB-IFC), so the same transit cases, mounts, and power adapters can serve either. It comes in three options, matching the compute demands of the service area it will serve:

| OPTION | CONFIGURATION | ROLE / FIT |
| --- | --- | --- |
| **Single deskside unit** | One deskside GPU superchip system (≈750 GB coherent memory class, ≈20 PFLOPS FP4, ≈1.6 kW, air-cooled); transit-case, shelter-, or vehicle-mounted | Baseline deployable Site IFC: runs the full sensor-specialist roster and site coordinator for one service area. |
| **Dual linked units** | Two deskside units linked over an 800 Gb/s interconnect (≈1.5 TB coherent memory class, ≈40 PFLOPS FP4, ≈3.2 kW) | Higher-capacity or redundant deployable Site IFC: more concurrent specialist models and larger context, for a larger or busier service area. |
| **Rugged shelter rack** | 2–4 GPU modules in a rugged, rack-mounted shelter enclosure | Where a shelter rack and higher throughput are preferred over a deskside chassis, for the same role. |

TABLE 3A · DEPLOYABLE SITE IFC COMPUTE OPTIONS. ALL THREE ARE AIR-COOLED, TRANSIT-CASE/SHELTER/VEHICLE-MOUNTABLE, AND GENERATOR- OR FACILITY-POWER COMPATIBLE — THE SAME HARDWARE CLASS AS THE EQUIVALENT DEPLOYABLE TIER THEY ARE COMPATIBLE WITH.

The physical envelope follows the same ruggedization baseline as the reused hardware, expressed here against international and civil standards:

| PROPERTY | BASELINE | STANDARD |
| --- | --- | --- |
| **Environmental sealing** | Dust-tight, water-resistant enclosure; sealed connectors | IEC 60529 (IP67-class) |
| **Shock, vibration, temperature** | Transit and operating shock/vibration; wide operating temperature range | IEC 60068-2 (shock/vibration/temperature testing); IEC 60721-3 (environmental classification) |
| **Power quality & EMC** | Runs on facility or generator power without dedicated conditioning | IEC 61000-6 (EMC generic standards); ISO 8528 (generator-set compatibility) |
| **Connectors** | Sealed, screw-coupling circular | EN 3645 / IEC 61076-2-101 |

| PROPERTY | BASELINE | STANDARD |
|---|---|---|
| | connectors for power, data, and video | (M12-class) |

TABLE 3B · RUGGEDIZATION BASELINE FOR THE DEPLOYABLE CONFIGURATION, CITED AGAINST INTERNATIONAL AND CIVIL STANDARDS.

Estimated power draw and weight follow the compute option chosen: roughly 1.6 kW for a single unit and 3.2 kW for a dual-linked pair, drawn from facility or generator power; roughly 20–40 kg for the compute chassis alone, before transit case, cabling, and mounting hardware. Neither figure assumes a permanent utility connection or a built data-centre space — the deployable configuration is designed to stand up wherever facility or generator power and a network bearer reach.

## 3.6 Connectivity for the deployable configuration

Sensor ingestion at a deployable Site IFC is unchanged from Section 6: the same connector model, the same standard protocols, and the same already-existing-sources-only floor apply, regardless of which compute option or which reachback bearer is in use. What the deployable configuration adds is a choice of standards-based bearers for reachback to the Central IFC, selected by what the site actually has available:

| BEARER | STANDARD | WHEN USED |
|---|---|---|
| **Wired** | IEEE 802.3 Ethernet / fibre | Preferred wherever a wired connection to an existing network is available at the site. |
| **Commercial cellular** | 3GPP LTE / 5G | Where commercial cellular coverage reaches the site. |
| **Commercial / civil satellite** | DVB-S2X-class VSAT or equivalent civil broadband satellite service | Beyond wired or cellular coverage. |
| **Civil mesh radio (fallback)** | Commercial mesh-radio product conforming to IEEE 802.11s-class mesh networking | Where none of the above reaches the site directly; relays reachback traffic to the nearest point where a wired, cellular, or satellite bearer is reachable. |

TABLE 3C · REACHBACK BEARER OPTIONS FOR THE DEPLOYABLE CONFIGURATION. EVERY OPTION CARRIES THE SAME STANDARD IP TRANSPORT, SECURED BY ORIGI (SECTION 9), REGARDLESS OF WHICH BEARER IS IN USE.

The bearer choice is a transport decision only: it changes nothing about how a deployable Site IFC ingests, infers, or commits, and nothing about the policy framework or Origi governing it. Partition tolerance (Section 3.4) applies identically — a deployable Site IFC continues committing its own regional picture on any bearer, or on none, and reconciles with the Central IFC once a bearer is reachable again.

# 4. The Four Planes Across Two Tiers

Lamassu's logical architecture is four cooperating planes observed end to end by a single tracing layer. Planes 1–3 run once per Site IFC for that area's own scope, and run again at the Central IFC for the enterprise-wide scope, while plane 4 wraps both tiers and the reachback link between them.

## 4.1 The four cooperating planes

| PLANE | WHAT IT DOES | WHERE IT RUNS |
|---|---|---|
| **1 · Deterministic substrate** | Ingest, normalize, tag, and store telemetry with provenance. No model in the loop. | Each Site IFC, for its area's already-existing sources. |
| **2 · Sensor-specialist agents** | One specialist per sensor family, each owning one tuned micro-model, emitting cited, typed findings. | Each Site IFC, close to the sources it reasons over. |
| **3 · Reasoning & orchestration** | A coordinator composes findings into a candidate; an adjudication engine commits. | Site coordinator + local adjudication engine at each Site IFC (regional scope); cross-site coordinator + central adjudication engine at the |

| PLANE | WHAT IT DOES | WHERE IT RUNS |
|---|---|---|
| | | Central IFC (enterprise-wide scope). |
| **4 · Trust assurance & observability** | The policy framework and trust harness wrap and observe everything; provenance and the trace tree. | Both tiers, and the reachback link between every Site IFC and the Central IFC. |

TABLE 4 · THE FOUR PLANES ACROSS TWO TIERS. PLANE 3 RUNS TWICE — ONCE PER SITE IFC AND ONCE, AT WIDER SCOPE, AT THE CENTRAL IFC. PLANE 4 WRAPS EVERYTHING, INCLUDING REACHBACK.

## 4.2 How the planes interact across tiers

Within a Site IFC, planes 1–3 interact directly: ingestion feeds specialists, specialists feed the site coordinator, the site coordinator proposes to the site's own adjudication engine, and the engine alone commits the regional indicator. The committed regional indicator, and the observations behind it, then flow up over reachback to the Central IFC, where they become one input among many to the cross-site coordinator. The Central IFC's own adjudication engine commits a separate class of indicator: one that spans or exceeds what any single Site IFC could conclude on its own — for example, a pattern that only becomes visible when several service areas' pictures are compared side by side.

Two failure modes are now distinguished and handled by two different mechanisms. A **service failure** within a tier — a specialist instance, a coordinator, an adjudication engine process — is routed around by that tier's own high-availability redundancy (Section 3.4), and is invisible outside the tier. A **reachback failure** between a Site IFC and the Central IFC is handled by partition tolerance (Section 3.4, Section 9): the Site IFC's own commitments continue uninterrupted, and only the Central IFC's cross-site view is affected, recovering by reconciliation once the link returns.

# 5. The Tiered Intelligence Model

The defining choice is to replace one large model with many small ones, each an expert in a single sensor modality, coordinated by a larger reasoner — a bet that specialization plus composition beats generalization for urban sensing. Specialization runs along two axes: modality and locality. Sensor specialists and the first coordinator sit at the Site IFC nearest the sources they reason over; the largest models and the deepest correlation sit at the one Central IFC, reserved for what only makes sense to reason about across many areas at once.

## 5.1 Sensor-specialist micro-models, at the Site IFC

Each sensor family is analysed by a dedicated agent running a micro-model of roughly one to eight billion parameters, tuned for that modality's telemetry and served at the Site IFC that ingests it, with predictable latency. The roster spans the civilian urban modalities: vision (EO/IR video), acoustic, RF spectrum, spatial (LiDAR / point-cloud), and environmental (air-quality, gas, weather). A specialist reads only its own modality within its own Site IFC, cannot write fused state, and emits cited, typed findings (Invariant 5).

## 5.2 Two coordinators, two scopes

A **site coordinator** runs at each Site IFC. It plans a run, dispatches typed requests to that area's sensor specialists, composes their cited findings across modalities into a candidate regional event with its provenance, and frames a proposal for the site's own adjudication engine. A **cross-site coordinator** runs once, at the Central IFC. It composes the committed regional indicators — and, where relevant, the underlying observations — arriving from every Site IFC into a candidate enterprise-wide event, and frames a proposal for the central adjudication engine. Neither coordinator ever re-reads raw telemetry to second-guess a specialist's modality call, and neither ever commits a verdict itself.

## 5.3 Propose, compose, commit — an Origi assurance case at each tier

The keystone three-step transition holds at both tiers, and no model may short-circuit it at either. At a Site IFC: **specialists propose** modality findings; **the site coordinator composes** them into a regional candidate; **the site's adjudication engine commits**, defers, or rejects under bound policy. At the Central IFC: **Site IFCs propose** (their committed regional indicators and supporting observations); **the cross-site coordinator composes** them into an enterprise-wide candidate; **the central adjudication engine commits**, defers, or rejects under bound policy. In both cases the committed indicator is the only output its operator sees, and a model never writes it — an engine does. Because both engines are deterministic and sit outside every model, every committed result, at either tier, is reproducible and auditable.

Each commit decision is itself an assurance case in the structure of Origi, Enkidu's Trust Assurance Framework (Section 9): the candidate indicator is the **claim**, bound to the mission's policy and its acceptance threshold; the specialists' cited findings and the coordinator's composition of them are the **argument** that decomposes the claim into checkable parts; and the chain of custody behind every cited observation is the **evidence**. An adjudication engine commits a claim only when both conditions Origi requires are met — the composed indicator satisfies its policy threshold, **and** the supporting evidence chain validates: every signature verifies, no hash link is broken, and no evidence has expired. A claim is never committed on a passing threshold alone if the evidence behind it cannot be trusted; it is instead deferred or rejected, and the deferral or rejection is itself recorded to the chain of custody with the reason. This is what makes a committed indicator a defensible position rather than an asserted one — the same distinction Origi draws between a number someone typed and a number a named, verifiable process produced.

This subsection fixes the adjudication engine's **structure** — the assurance-case gate itself — as an architectural invariant. It does not fix the adjudication **algorithm**: the specific confidence-aggregation function, per-mission thresholds, conflict-resolution rule, and partition-reconciliation rule that run inside that gate. Those are determined empirically, through the simulation-driven methodology of Section 14.3, before detailed design proceeds.

# 6. Adaptive Sensor Ingestion

> ***CAPABILITY ADDRESSED · Adaptive sensor ingestion*** *— One pipeline, replicated per Site IFC, connects to diverse, already-existing sensor technologies over standard protocols, and keeps producing a useful picture when individual sources — or the reachback link itself — degrade or go offline.*

All telemetry enters Lamassu at the Site IFC serving the area it comes from. Every source is already deployed and already operated by the city, facility, or third party Lamassu serves — reached either directly over the sensor's own protocol or through the existing system (a VMS, an IoT platform, an OT historian) that already manages it. Lamassu never installs a sensor and never hosts ingestion or inference on infrastructure of its own below the Site IFC; the pipeline normalizes whatever arrives into a common, tagged representation the platform can reason over, under Tangram's structure and Origi from the first touch. The Central IFC does not, as a rule, ingest raw sensor telemetry directly — its primary inputs are the committed regional pictures and supporting observations that Site IFCs send it over reachback, though it may also directly ingest already-existing enterprise-wide or cross-jurisdiction sources (e.g. a national or regional data platform) where such a source is not naturally scoped to any one service area.

## 6.1 The common ingestion pipeline

| STAGE | WHAT HAPPENS |
|---|---|
| **Acquire** | A reading or feed arrives at the Site IFC serving its area, over a standard protocol connector, from an already-existing sensor or the already-existing system that operates it. |
| **Authenticate & admit** | Origi verifies the source — device identity, system credential, or feed credential — and admits the datum only if it passes. This is the first link in the chain of custody. |
| **Normalize & tag** | The datum is converted to the common representation and tagged with modality, time, location, originating source, reliability, and a data-sensitivity / privacy caveat (Section 8.5). |
| **Zone (Tangram)** | Tangram assigns the datum to the relevant zone(s) — the area it pertains to, the track it belongs to, the processing state it has reached, and which tier(s) that zone currently resolves at — so its movement is observable and the right services are notified. |
| **Enrich & correlate** | The datum is linked to existing tracks and entities and run through the relevant sensor-specialist model, at that Site IFC, for detection / classification — the point at which raw data becomes a candidate observation. |
| **Record (chain of custody)** | Origi writes a signed, hash-linked provenance entry for the datum and every transformation, so the evidence behind any later assessment — at either tier — is reconstructable and tamper-evident. |
| **Distribute under** | The tagged observation is shared with the site coordinator |

| STAGE | WHAT HAPPENS |
|---|---|
| **policy** | and, over reachback, with the Central IFC's cross-site coordinator — each share brokered against policy and privacy caveats. |

TABLE 5 · THE COMMON INGESTION PIPELINE, RUN ONCE PER SITE IFC. BECAUSE EVERY SOURCE IS EXTERNAL AND ALREADY-EXISTING, THE SAME STAGES APPLY UNIFORMLY REGARDLESS OF WHICH SITE IFC OR WHICH SYSTEM A DATUM ORIGINATED FROM.

## 6.2 Standard protocols and the connector model

Lamassu does not invent a private bus and does not require a sensor owner to adopt one. Every Site IFC speaks the protocols urban sensors and their existing platforms already use, behind a connector model: each connector adapts one protocol to the common ingestion representation, and connectors are versioned, signed artifacts that can be added at a Site IFC without changing the pipeline itself.

| PROTOCOL / STANDARD | WHAT IT BRINGS IN |
|---|---|
| **ONVIF / RTSP** | Already-existing IP cameras and video management systems — the most common urban video estate. |
| **MQTT / OGC SensorThings API** | Already-existing IoT and civic sensor networks — air-quality, acoustic, environmental, occupancy. |
| **NGSI-LD (ETSI / FIWARE context broker)** | Already-existing smart-city context data and cross-domain entities, with a standard data model. |
| **OPC-UA / Modbus** | Already-existing building-management and industrial / OT sensors at facilities and utilities. |
| **DDS** | Already-existing real-time publish/subscribe telemetry where low latency matters. |
| **OpenAPI / REST + gRPC** | Request/response and streaming feeds from already-existing external services and analytics. |

TABLE 6 · REPRESENTATIVE STANDARD PROTOCOLS EACH SITE IFC INGESTS THROUGH ITS CONNECTOR MODEL. EVERY ONE OF THESE PATHS REACHES AN ALREADY-EXISTING SENSOR OR AN ALREADY-EXISTING SYSTEM THAT OPERATES ONE — NONE REACHES HARDWARE LAMASSU DEPLOYED.

## 6.3 Onboarding already-existing feeds; dynamic discovery

On onboarding and continuously thereafter, a Site IFC's ingestion layer connects to already-existing sensors and systems nominated by the operator within its service area and proposes attaching the feeds it recognises. New feed types are added by registering a connector at the relevant Site IFC, not by installing anything in the field, so the set of sources can evolve over the life of a deployment without any change to Lamassu's own footprint. Every newly attached feed enters in a low-trust zone and earns standing as its data proves consistent (Section 9.2); any feed carrying personal data is held for operator confirmation and a privacy caveat before it is admitted (Section 8.5).

## 6.4 Graceful degradation, safe-abstain, and partition tolerance

Sources fail — a camera drops, a feed goes stale, a sensor drifts — and now, so can the reachback link between a Site IFC and the Central IFC. Lamassu is built to keep producing a useful, honestly-qualified picture in every case, rather than failing silently or asserting confidence it has not earned:

- **Reweight and continue.** When a feed degrades, a site coordinator widens its fusion window and reweights the remaining specialists at that Site IFC; coverage degrades smoothly and measurably rather than collapsing.
- **Safe-abstain.** A specialist on out-of-distribution or missing input abstains rather than guessing, and the chain of custody records the abstention, so a gap is visible instead of hidden.
- **Redundant coverage, where the existing estate provides it.** If two already-existing sensors within a Site IFC's area happen to cover overlapping ground, the ingestion pipeline can draw on both. This is a property of the sensor estate the operator already owns, not a Lamassu-owned mesh.
- **Partition tolerance on reachback loss.** If the link from a Site IFC to the Central IFC drops, the Site IFC does not pause or degrade its own regional ingestion, inference, or commitment — it continues exactly as if connected, and its chain of custody reconciles with the Central IFC's wider picture once the link returns, with any divergence surfaced rather than silently overwritten (Section 3.4, Section 9).
- **Confidence reflects coverage.** Because confidence is computed from the chain of evidence (Section 10.2), a thinner evidence base — whether from a degraded source or a partitioned Site IFC's necessarily narrower view — yields a lower, explicit confidence at whichever tier is assessing it.

# 7. Platform Layer — Tangram

Tangram is the platform layer. It does two jobs — it zones the operational space, and it provides the observability pattern that tracks events through those zones and notifies the services that need to know — and it carries the policy framework that turns a zone from a boundary into a governed execution context (Section 8).

## 7.1 Zoning across two tiers

A Tangram zone is a bounded operational space with one goal and explicit entry and exit conditions, defined by its boundary conditions rather than its physical form. A zone is defined by the data it governs — a sector of a campus as represented in ingested feeds, a track's lifecycle state, a processing stage, a mission phase — never by which piece of Lamassu-owned hardware sits where. A zone's membership can span both tiers: a zone can be scoped to one Site IFC's service area, or it can be an enterprise-wide zone whose members are drawn from several Site IFCs and resolved at the Central IFC. Tangram tracks which tier a zone's activity currently resolves at, so a regional zone escalating into a cross-site concern is itself an observable, explicit transition — the same mechanism that governs a track crossing from one source's coverage area into another's.

## 7.2 Tangram as the observability pattern

Every zone transition is an observable event — a track entering a sector, a contact changing classification, a specialist abstaining, a regional indicator escalating to the Central IFC. Tangram publishes these events across both tiers so the services with a stake in a zone are notified the moment something material changes, providing threat/event tracking with continuous custody across ingested sources and across the reachback link, relevance-scoped notification rather than flooding, and a common operating picture — regional at a Site IFC, enterprise-wide at the Central IFC — assembled from shared zone vocabulary. Zoning and observability are the same mechanism seen from two sides: defining a zone's boundaries is what makes an event's movement across it, or across tiers, observable, and what makes notifying exactly the right services possible.

# 8. The Policy Framework

A Tangram zone is not only a boundary; it is a governed execution context. This is the strong policy framework Lamassu adopts so that what every agent may do —

with which data, under whose authority — is explicit and machine-checkable, never merely asserted. The framework is the operational expression of Invariant 6.

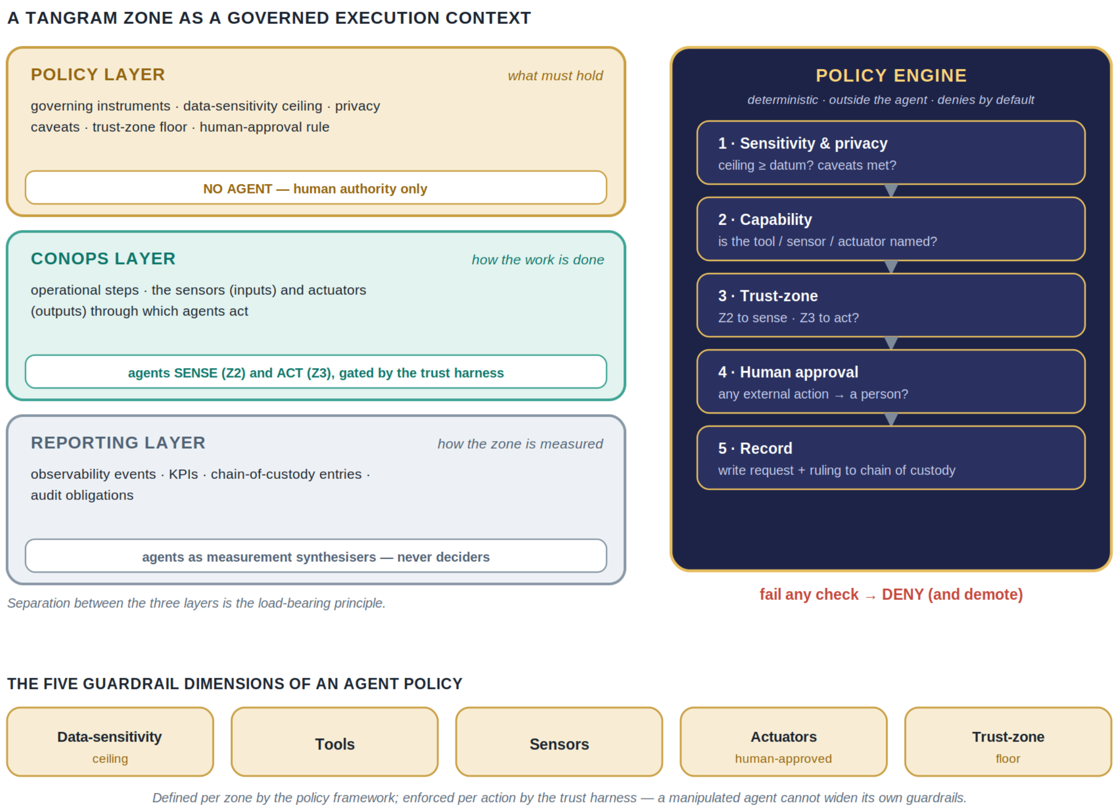


FIGURE 2 · THE POLICY FRAMEWORK. EVERY ZONE IS A THREE-LAYER EXECUTION CONTEXT (LEFT); EVERY AGENT ACTION IS RESOLVED BY A DETERMINISTIC, DENY-BY-DEFAULT POLICY ENGINE (RIGHT) AGAINST THE FIVE GUARDRAIL DIMENSIONS OF ITS POLICY (BOTTOM). THIS MECHANISM IS UNCHANGED BY THE MOVE TO TWO TIERS — IT GOVERNS AGENTS AND DATA, NOT PHYSICAL PLACEMENT, AND APPLIES IDENTICALLY AT EVERY SITE IFC AND AT THE CENTRAL IFC.

## 8.1 The three-layer zone: policy, conops, reporting

Every zone is structured into three layers, and the separation between them is the load-bearing principle of the architecture. The policy layer states what must hold; the conops layer is how the work is done; the reporting layer is how the zone is measured.

| ZONE LAYER | WHAT IT HOLDS IN LAMASSU | WHERE AGENTS MAY OPERATE |
|---|---|---|
| **Policy** | The governing instruments and constraints for the zone: the data- | No agent. Policy determinations — and any |

| ZONE LAYER | WHAT IT HOLDS IN LAMASSU | WHERE AGENTS MAY OPERATE |
|---|---|---|
| | sensitivity ceiling, the privacy caveats, the permitted tools, sensors, and actuators, the trust-zone floor, and the human-approval rule. | decision for action — are human. Agents read policy as a constraint; they never set it. |
| **Conops** | The operational steps of the zone, and the sensors (inputs) and actuators (outputs) through which the work is sensed and acted on. | Agents act here as sensors and actuators, gated by trust zones (Z2 to sense, Z3 to act). |
| **Reporting** | The observability events, KPIs, chain-of-custody entries, and audit obligations the zone emits. | Agents act here as measurement synthesisers — correlating, scoring, summarising; never as deciders. |

TABLE 7 · THE THREE LAYERS INSIDE EVERY TANGRAM ZONE. AGENTS INHABIT THE CONOPS AND REPORTING LAYERS ONLY; THE POLICY LAYER IS HUMAN.

## 8.2 The policy layer is human

No agent — generative, predictive, or deterministic — occupies the policy layer of any zone. An agent may sense, correlate, score, and recommend, producing a structured, logged input to a determination; the determination itself, and any decision for action, is made by a human who is accountable for it. Systems do not make decisions for action because they cannot be responsible for them. This is the architecture's overriding principle, and the policy framework is how it is enforced rather than merely asserted.

## 8.3 The five guardrail dimensions of an agent policy

Within the policy layer, the guardrails an agent runs under are expressed as a policy — a machine-readable object that Tangram attaches to the agent and that Origi evaluates on every action. A policy does not describe what an agent should do; it bounds what it may do. Five dimensions are specified.

| GUARDRAIL DIMENSION | WHAT THE POLICY SPECIFIES | ENFORCED BY |
|---|---|---|
| **Data-sensitivity** | The highest sensitivity, and the privacy caveats, of data the agent may read in | Trust harness, against the datum's sensitivity |

| GUARDRAIL DIMENSION | WHAT THE POLICY SPECIFIES | ENFORCED BY |
|---|---|---|
| **ceiling** | this zone — for example Sensitive, no-export, redaction-required. | tag (Section 8.5). |
| **Tools** | The named inference tools, models, and connectors the agent may invoke. | Guardrails (execution rails) + Origi. |
| **Sensors** | The sensor classes the agent may read — video, acoustic, RF, LiDAR, environmental — within this zone. | Trust-zone Z2 gate (Section 9.2). |
| **Actuators** | The civil effectors the agent may cue — steer a camera, retask a sensor, raise an alert — always with the human-approval rule attached for external action. | Trust-zone Z3 gate + human approval (Section 9.2). |
| **Trust-zone floor** | The minimum trust zone the agent must hold for the policy to grant anything; a hard-gate event drops it to Z0 and voids the grant. | Trust harness (Section 9.2). |

TABLE 8 · THE FIVE GUARDRAIL DIMENSIONS OF AN AGENT POLICY. THE FRAMEWORK DEFINES THEM PER ZONE; ORIGI ENFORCES THEM PER ACTION.

## 8.4 The deny-by-default policy engine

When an agent requests an action, the policy engine resolves the request against the policy of the zone the agent is operating in, in a fixed order, and denies by default: an action is permitted only if every guardrail dimension explicitly grants it, and any ambiguity or missing grant resolves to denial, never to a permissive fallback.

## 8.5 The data-sensitivity and privacy framework

| LEVEL | WHAT IT COVERS | HANDLING |
|---|---|---|
| **Public** | Data that carries no risk if disclosed. | No special handling. |

| LEVEL | WHAT IT COVERS | HANDLING |
|---|---|---|
| **Internal** | Operational data not intended for public release, but with low individual risk. | Standard access control; no export restrictions. |
| **Sensitive** | Data whose disclosure could aid a bad actor — e.g. exact sensor placements, vulnerability detail. | Restricted to authorised operators and Z2+ agents; sharing brokered. |
| **Personal Data** | Any datum identifying a person — faces, plates, device identifiers, location traces. | Privacy-regulated: lawful basis, minimization, retention limit, and redaction enforced; export blocked without authority. |

TABLE 9 · THE CIVILIAN DATA-SENSITIVITY SCHEME. PERSONAL DATA IS GOVERNED BY PRIVACY REGIMES AND CARRIES THE STRICTEST HANDLING.

Several caveats ride on top of a level rather than forming one: data-residency (the datum may not leave a jurisdiction), retention limit (it auto-expires), consent / lawful basis (the ground on which personal data is held), redaction-required (faces or plates blurred before sharing), and no-export / restricted-sharing. Origi carries each as a machine-readable tag and enforces it at the point of distribution.

## 8.6 Agent authorization levels

Just as operators hold access rights, agents are designated with their own authorization levels, so that sharing inside the automated system can be reasoned about deterministically. An agent's authorization is the ceiling its policy may grant; an agent may never read, correlate, or carry data above its level, and Origi will not pass higher-sensitivity data to a lower-authorized agent.

| AGENT AUTHORIZATION | MAY ACCESS | TYPICAL PLACEMENT |
|---|---|---|
| **Public-data agent** | Public and Internal data under need-to-know. | Support and housekeeping agents at either tier. |
| **Operational agent** | Up to Sensitive data and fused tracks; no personal data. | Most sensor-specialist and coordinator agents at a Site IFC; most cross-site inference |

| AGENT AUTHORIZATION | MAY ACCESS | TYPICAL PLACEMENT |
|---|---|---|
| | | agents at the Central IFC. |
| **Privileged agent** | Personal data, only under an approved privacy assessment and explicit need-to-know. | Agents in dedicated, audited zones, at either tier. |

TABLE 10 · AGENT AUTHORIZATION LEVELS. AN AGENT'S LEVEL BOUNDS WHAT ITS POLICY CAN GRANT; ORIGI NEVER SHARES ABOVE IT, AT EITHER TIER.

### 8.7 Need-to-know, need-to-use, need-to-act

Lamassu distinguishes three relationships a party may have to a datum, and they are not the same authority. An agent may have need-to-know (to read) or need-to-use (to correlate) without ever having need-to-act. Need-to-act is reserved for humans: when a fused observation raises a requirement for action in the policy layer, Origi deconflicts the picture and presents the situation to an authorised operator — entitled to see every datum bearing on it — who then owns the decision and the action. No agent crosses from knowing or using into acting.

## 9. Trust Harness — Origi

Lamassu's trust harness is a deployment of **Origi**, Enkidu's Trust Assurance Framework: a policy-based, evidence-backed approach to asserting and measuring trust in AI-augmented systems, applied at inference time. Origi is the harness through which every agent action passes, at every Site IFC and at the Central IFC. Its remit is four-fold — govern what each agent may do through trust-operation zones; control sensor and actuator access; secure every ingestion link, every reachback link, and every internal link at either tier; and capture the chain of custody that determines confidence in a detection and reconciles it across a partition. Origi does not assume any agent is trustworthy: it applies a zero-trust posture — never trust, always verify — continuously checking that each step satisfies its bound policy before allowing it to proceed.

### 9.1 Two layers of defense

The harness is the outer of two layers. An inner, content-level guardrail layer provides programmable input, dialog, retrieval, execution, and output rails around each agent — rejecting or transforming inputs and outputs and gating tool calls.

Origi adds an outer layer concerned with authority and accountability rather than content: whether this agent, at its current trust, may take this action at all, and proving it later. Keeping deterministic authority checks outside the agent and above the probabilistic rails bounds the blast radius when an inner check is wrong — a guard model is never the sole gate.

## 9.2 Three enforcement phases

Origi places a checkpoint at every trust boundary Lamassu crosses — where a request enters, where one agent hands off to another, where an agent calls a tool or cues an actuator, and where an indicator leaves the system — and brokers every crossing against the request's policy. These checkpoints resolve into three enforcement phases that map directly onto Lamassu's own pipeline:

| PHASE | WHAT IT CHECKS IN LAMASSU | VERDICT |
|---|---|---|
| **Pre-flight (ingress)** | The ingestion pipeline's Authenticate & admit stage (Section 6.1): classifies the incoming datum, assigns its policy profile, and checks that its source holds the authority to contribute. | Allow, sanitize, or refuse. |
| **In-flight (per step)** | Every specialist-to-coordinator handoff, every coordinator-to-adjudication-engine proposal, and every actuator cue: validated against policy, bounded by least privilege and the agent's trust zone, with irreversible or externally-visible actions routed to a human-in-the-loop hold. | Allow, warn, block, or hold for human review. |
| **Post-flight (egress)** | The operator handoff (Section 10.3): checks the committed indicator is faithful to what actually happened, attaches its calibrated confidence, and attaches the provenance manifest behind it. | Release, hold, or block. |

TABLE 10A · ORIGI'S THREE ENFORCEMENT PHASES, MAPPED ONTO LAMASSU'S INGESTION, INFERENCE, AND OPERATOR-HANDOFF STAGES. SECURITY AND HIGH-STAKES CHECKS FAIL CLOSED: IF A CHECK CANNOT COMPLETE, THE ACTION IS BLOCKED RATHER THAN ALLOWED.

## 9.3 Trust-operation zones

Every agent operates inside a trust-operation zone — a capability envelope earned through demonstrated, evidence-backed good behaviour. The zone determines what classes of resource the agent may touch, and the highest zone is the only one from which an agent may cue a civil actuator.

| TRUST ZONE | WHAT AN AGENT MAY DO | ADMISSION BAR |
|---|---|---|
| **Z0 · Quarantined** | Nothing; isolated, evidence preserved. | Demotion target after any hard-gate event (e.g. detected injection). |
| **Z1 · Observe** | Read its own context; produce assessments. No sensor or actuator access. | Clean admission; identity and scope validated. |
| **Z2 · Sense** | Read from sensors; contribute to shared state. | Sustained Z1 good standing; clean security scans. |
| **Z3 · Act** | Cue a civil actuator — steer a camera, retask a sensor, raise an alert. | Sustained Z2 standing + explicit human approval for any external action. |

TABLE 11 · TRUST-OPERATION ZONES. SENSOR ACCESS IS GATED AT Z2; CUEING A CIVIL ACTUATOR IS RESERVED FOR Z3, WITH A HUMAN IN THE LOOP FOR ANY EXTERNAL EFFECT.

Two safeguards carry directly from the Trust Assurance Framework. Separation of duties: the component that assesses an agent's trust, the one that rules on its zone, and the one that grants capability are distinct, so no agent — and no single compromised service — can promote itself. And trust is earned slowly and lost instantly: a single security or alignment violation forces immediate demotion, revoking sensor or actuator rights before the next action.

## 9.4 Separation of duties

Origi's separation of duties — Observer, Judge, Controller — maps onto the platform, each isolated from the agents and from one another so that no single component can both decide an agent deserves more capability and grant it. Specialists and the coordinator are Observers that propose. The policy logic acts as Judge, ruling against zone policy on whether a step may proceed and whether an agent may be promoted. The adjudication engine and the operator handoff act as Controller, committing or releasing what the Judge allowed. No single component both proposes and commits — which, because inference is decomposed into specialist and coordinator roles, is largely true by construction already. Every assessment, verdict, and zone transition is itself recorded as signed provenance, so who granted an agent the capability it used is always answerable.

## 9.5 Secured communication, reachback, and chain of custody

Origi secures every link that touches either tier: the ingestion link from every already-existing external sensor or system into its Site IFC, the **reachback link between every Site IFC and the Central IFC**, and every internal service-to-service link within each tier — all with encryption and mutual authentication, so observations cannot be read or forged in transit, and a compromised source, Site IFC, or service can be cryptographically excluded. As observations and inferences flow, the harness records each as a signed, hash-linked entry in a chain of custody: what was observed or inferred, by which system, service, or agent, at which tier, from what inputs, and when. Because entries are hash-linked, tampering upstream invalidates everything downstream; because they are signed, every contributor is attributable; and because evidence expires, stale evidence cannot satisfy a claim. Privacy and handling caveats ride as machine-readable tags on every entry and are enforced at the point of distribution, at either tier.

When a reachback link is cut and later restored, Origi is also what performs reconciliation: the Site IFC's chain of custody, built up while partitioned, is merged against the Central IFC's wider record, and any conflict — a regional conclusion that the enterprise-wide picture would have revised — is surfaced to the operator as an explicit divergence rather than silently overwritten. Origi's remit spans every Site IFC's own internal service mesh, the Central IFC's internal service mesh, and the reachback link connecting each Site IFC to the Central IFC.

## 9.6 What Origi assures

Origi assures nine components of trustworthiness, each carrying its own policy, claim, and evidence chain. In Lamassu's civilian sensing context, the nine resolve as follows:

| COMPONENT | WHAT IT ASSURES IN LAMASSU |
|---|---|
| **Competence & Reliability** | Specialists and coordinators perform accurately across approved sensor conditions, and abstain gracefully at the edges (Section 6.4). |
| **Transparency & Explainability** | Every committed indicator resolves to the trace tree behind it, at a level the operator can act on (Section 10.3). |
| **Calibrated Confidence** | A stated confidence reflects the actual evidence behind a detection, computed from the chain of custody, never |

| COMPONENT | WHAT IT ASSURES IN LAMASSU |
|---|---|
| | asserted by a model (Section 10.2). |
| **Predictability & Consistency** | Behaviour is stable under benign variation; model or policy changes are versioned and detected, never silent (Section 8.4). |
| **Accountability & Recourse** | Every decision traces end-to-end; high-stakes and externally-visible actions require human approval (Section 8.2, 8.7). |
| **Fairness & Bias Management** | Detections are evaluated for equitable treatment across the areas and populations a Site IFC serves, not only in aggregate. |
| **Security & Robustness** | Every link is encrypted and mutually authenticated; the harness resists manipulation within a stated threat model (Section 9.5). |
| **Alignment with Intent** | Agents pursue the mission's actual objectives under bound policy, never a harmful proxy for it (Section 8.3, 8.4). |
| **Human Agency & Reliance** | Humans retain need-to-act; agents inform and recommend, never decide for a human (Section 8.7). |

TABLE 10B · THE NINE TRUST COMPONENTS ORIGI ASSURES, MAPPED TO THEIR TREATMENT IN LAMASSU. A RECURRING DISCIPLINE: A CONFIDENT WRONG DETECTION IS TRACKED AND PENALIZED MORE HEAVILY THAN AN HONEST ABSTENTION, BECAUSE IT IS THE FAILURE MODE THAT MOST ERODES CALIBRATED TRUST.

# 10. Explainable and Actionable Outputs

> ***CAPABILITY ADDRESSED · Explainable & actionable outputs*** *— Every detection carries what was detected, which sources contributed, and a confidence the operator can trust — computed from evidence, not emitted by a model.*

## 10.1 What every indicator carries

The committed indicator the operator receives is not a bare alert. Because it is produced by the deterministic adjudication engine from typed, cited findings, it carries, by construction, the context an operator needs to act:

- **What was detected.** The event type, its location and time, and its lifecycle state in plain terms.
- **Which sources, and which tier, contributed.** The specific already-existing sensors, feeds, and systems whose observations support it, and whether the indicator is a regional commit from one Site IFC or an enterprise-wide commit from the Central IFC — each resolvable to its provenance, so the operator can see the basis, not just the conclusion.
- **How confident, and why.** A confidence derived from the chain of evidence (below), with the corroboration and any abstentions or gaps that produced it.
- **What is recommended.** The proposed next step — and, where it would reach the outside world, an explicit prompt for the human decision the policy framework requires.

## 10.2 Confidence from the chain of evidence

Confidence is a property of the evidence, computed from the chain of custody — not a number a model emits unaccountably. A track corroborated by several independent, well-attested already-existing sources earns high confidence in both its existence and its severity; a single-source, weakly-attested track is flagged as tentative. Before any detection drives a recommendation, the harness deconflicts sources — reconciling and de-duplicating reporting so one source is not mistaken for independent corroboration.

## 10.3 The operator handoff and the trace tree

Every run is a trace tree, not a flat log: every agent boundary, model call, tool call, and engine decision is a span, at whichever tier produced it, and the committed indicator is the root that resolves down through composition to the individual cited observations. For an enterprise-wide indicator committed at the Central IFC, the trace resolves further still — down through the cross-site coordinator into the specific Site IFCs' regional commits, and from there into each area's original observations. An operator can expand any indicator to the exact evidence and the exact policy decisions behind it, at either tier; an auditor can reconstruct any committed result after the fact; and a regression points at the span that produced it. Explainability is therefore not a report bolted on at the end — it is the same provenance both engines required to commit, surfaced to the human.

Because every committed indicator is an Origi assurance case (Section 5.3), the trace tree behind it is also a **trust dossier**: the claim, its bound policy, the argument that composed it, and the validated evidence chain behind it, assembled into a running record an operator, auditor, or regulator can inspect on demand rather than a summary reconstructed after the fact.

# 11. Scalable Architecture and Open Integration

> ***CAPABILITY ADDRESSED · Scalable architecture*** *— The central site scales by ingesting a broader inventory of already-existing sensor technologies and by adding compute, and every interface is a published, standards-conformant contract — so the platform is compatible with international standards by design.*

## 11.1 Scaling by coverage, breadth, and depth — never below the Site IFC

Lamassu scales in three directions, and none of them involve deploying anything below the Site IFC. **Coverage:** adding a new Site IFC to bring a new service area into the platform, reaching back to the same Central IFC as every other Site IFC — the primary way the platform's footprint grows. **Breadth:** within any one Site IFC's area, onboarding more already-existing sensor technologies and systems — additional camera networks, additional IoT platforms, additional facility or OT systems — each added by registering a connector against the standard protocols in Section 6.2, with no field visit required. **Depth:** adding compute at either tier — more specialist model-serving capacity or coordinator throughput at a Site IFC, or more cross-site correlation and adjudication capacity at the Central IFC — as the volume and variety of ingested sources grows. There is still no fleet of field hardware to grow, size, or maintain.

Because ingestion, zoning, and custody are uniform across every source and every Site IFC (Section 6), growing the number of service areas treats the whole, wider body of evidence as one coherent, governed picture at the Central IFC — the same way a single Site IFC treats its first few feeds.

## 11.2 Open at the interface, governed at the boundary

Lamassu is an open platform, not a sealed product. Its value grows as new capability — a vendor's already-existing sensor system, an analytic model, a third-party agent — can be plugged in and made to contribute. The governing idea is simple: open at the interface, governed at the boundary. Any conforming capability can connect, but everything that connects enters through one place that Tangram zones, the policy framework bounds, Origi governs, and secured comms carries. A published interface lets a capability connect, not command.

## 11.3 Conforming to international and open standards

Lamassu meets vendors and authorities on the standards their already-existing sensor systems already use, and wraps each so it conforms to the platform's zones and to Origi. This is also what makes the platform architecturally compatible with international deployment.

| STANDARD / FAMILY | WHERE IT LANDS IN LAMASSU |
|---|---|
| **Modular open-systems approach (published key interfaces)** | The overall design contract — every integration surface is a published, verifiable interface. |
| **OGC SensorThings, NGSI-LD, MQTT, oneM2M** | The sensor- and context-ingestion surface for already-existing IoT and smart-city feeds (Section 6.2). |
| **ONVIF / RTSP** | The ingestion surface for the already-existing urban camera estate. |
| **OGC (WMS/WFS/WMTS), STAC, CityGML** | Geospatial foundation data and the shared map / common operating picture. |
| **OPC-UA / Modbus / DDS** | Facility, industrial, and real-time data integration with already-existing OT systems. |
| **OpenAPI / REST, gRPC; MCP / A2A** | The two software integration surfaces — published APIs for systems and models, agent wrappers for agents. |
| **OASIS Common Alerting Protocol (CAP)** | Emitting operator alerts into external emergency-management and notification systems. |

| STANDARD / FAMILY | WHERE IT LANDS IN LAMASSU |
|---|---|
| **IEEE 802.3; 3GPP LTE / 5G; DVB-S2X-class satellite; IEEE 802.11s-class mesh** | Standards-based reachback bearers for the deployable Site IFC configuration (Section 3.6) — wired, cellular, satellite, and civil mesh-radio fallback, all carrying the same secured IP transport. |
| **IEC 60529, IEC 60068-2, IEC 60721-3, IEC 61000-6, ISO 8528, EN 3645 / IEC 61076-2-101** | Environmental sealing, shock/vibration/temperature, EMC, generator compatibility, and connector standards for the deployable Site IFC's physical envelope (Section 3.5). |
| **ISO/IEC 27001 / 27701, GDPR; ISO/IEC 42001 / 23894, NIST AI RMF; ISO/IEC TS 5723** | Information-security, privacy, and AI-management conformance the platform is designed to evidence, and the trustworthiness vocabulary Origi's nine components are mapped to (Section 9.6). |
| **NIST SP 800-207 (Zero Trust Architecture); NIST SP 800-53; W3C PROV; OWASP LLM/Agentic Top 10** | The zero-trust posture, separation of duties, provenance model, and threat taxonomy Origi is built against (Section 9). |

TABLE 12 · INTERNATIONAL AND OPEN STANDARDS LAMASSU CONFORMS TO, AND WHERE EACH LANDS. EVERY ROW CONNECTS TO AN ALREADY-EXISTING SENSOR TECHNOLOGY OR SYSTEM — MEETING FEEDS AND AUTHORITIES ON STANDARDS THEY ALREADY SPEAK IS WHAT MAKES THE PLATFORM SCALABLE WITHOUT REQUIRING LAMASSU TO DEPLOY ANYTHING OF ITS OWN INTO THE FIELD.

## 11.4 Two integration surfaces and the onboarding gate

Everything external enters through one of two surfaces. Published APIs carry systems — already-existing sensors and models, the latter packaged as signed inference microservices discovered through a gateway. Agent wrappers carry agents — a third-party agent is wrapped so it never talks to the platform directly; it talks to its wrapper, which sits behind the content guardrails and under Origi's Observer-Judge-Controller loop, with its tools, inputs, and outputs mediated and its entire trace signed. Both surfaces share one onboarding gate: a capability is published to a contract, run against a conformance harness, issued a cryptographic identity with a recorded bill of materials, admitted to a low-trust zone with least privilege, observed and promoted as field evidence accrues, and revoked instantly if it fails. The platform is open without being porous.

# 12. One Platform, Many Missions

The same central site hosts many missions at once, each a virtual Tangram execution zone running under its own policies, sensitivity ceiling, and guardrails, isolated from the others by the machinery of Sections 8 and 9. A Lamassu mission is a sensing-and-fusion mission: it sets what is to be concluded, with which data, and under whose authority; the action that may follow belongs to the operators it serves, never to Lamassu.

## 12.1 Mission archetypes

| MISSION ARCHETYPE | WHAT IT FUSES, AND WHICH TIER RESOLVES IT |
|---|---|
| **Critical-infrastructure protection** | A standing mission around a fixed site — a utility, substation, water plant, airport perimeter, or port. Already-installed sensors (video, acoustic, RF, ground sensors) already operated at the site feed that area's Site IFC directly; the Site IFC classifies approaches and raises them before they reach the boundary, and holds the standing regional picture on its own authority. |
| **Environmental early warning** | Air-quality, gas, flood, and wildfire-smoke sensing from already-existing environmental sensor networks, fused at each area's Site IFC into local early warnings, with the Central IFC correlating across areas to detect a wider pattern (e.g. a smoke plume or contamination plume crossing several service areas) that no single Site IFC would see alone. |
| **Public-safety event awareness** | Crowd density, flow, and anomaly sensing from already-existing cameras and sensors at a large event or transport hub, fused at that event's Site IFC into operator-facing situational awareness, with privacy caveats (redaction, retention limits) applied by default to any personal data. |
| **Airspace safety** | Detection of unauthorized drones over sensitive civil sites (airports, stadiums) by fusing already-existing RF, acoustic, and visual sensor feeds at the site's own Site IFC — a civil air-safety mission — with the Central IFC correlating a track that moves between the airspace of two adjacent Site IFCs' areas. |

TABLE 13 · CIVILIAN MISSION ARCHETYPES. EACH IS PRIMARILY RESOLVED AT THE SITE IFC SERVING ITS AREA, WITH THE CENTRAL IFC RESERVED FOR PATTERNS THAT ONLY EMERGE ACROSS SEVERAL SERVICE AREAS.

### 12.2 Adding a mission

A new mission is defined, not built: it binds a set of objectives to the policies, sensitivity ceiling, and guardrails needed to pursue them, and is instantiated as a virtual execution zone at the relevant Site IFC, the Central IFC, or both, depending on whether the mission is naturally regional or naturally cross-site. Because the planes, the ingestion pipeline, the policy framework, and Origi are mission-agnostic and identical in kind at every Site IFC, adding a mission — or adding a new Site IFC to extend an existing mission to a new area — is a configuration act under governance, not a re-architecture, and never requires deploying anything below the Site IFC.

## 13. Capability Coverage Summary

The three required capabilities are not appendices to the architecture; each is delivered by named, load-bearing mechanisms, distributed across the Site IFC and Central IFC tiers. This section maps the requirement to the mechanism and the section that specifies it.

| CAPABILITY | HOW THE ARCHITECTURE DELIVERS IT | SECTIONS |
|---|---|---|
| **Adaptive sensor ingestion** | One common pipeline replicated at every Site IFC; a signed connector model speaking ONVIF, MQTT, SensorThings, NGSI-LD, OPC-UA, DDS against already-existing sensor technologies only; dynamic onboarding; reweighting, safe-abstain, and partition tolerance under source or reachback loss. | 6.1–6.4 |
| **Explainable & actionable outputs** | Deterministic adjudication over typed, cited findings at both tiers, each commit an Origi assurance case (claim, argument, evidence); confidence computed from the chain of evidence; an operator handoff that resolves every indicator — regional or enterprise-wide — to its trust dossier through the trace tree. | 5.3, 9.6, 10.1–10.3 |
| **Scalable** | Horizontal scaling by adding Site IFCs for new | 11.1– |

| CAPABILITY | HOW THE ARCHITECTURE DELIVERS IT | SECTIONS |
|---|---|---|
| **architecture** | service areas; breadth scaling by onboarding more already-existing sources per area; vertical scaling of compute at either tier; published, standards-conformant interfaces; international information-security, privacy, and AI-management standards by design. | 11.4 |

TABLE 14 · HOW LAMASSU V4.4 DELIVERS EACH REQUIRED CAPABILITY, AND WHERE EACH IS SPECIFIED.

# 14. Maturity and Forward Path

## 14.1 Phased roadmap

| PHASE | FOCUS | WHAT IT ESTABLISHES |
|---|---|---|
| **Phase 1** | One Site IFC | A single Site IFC's ingestion pipeline stood up against a small number of already-existing feeds in one service area, two or three sensor specialists, a site coordinator, and a local adjudication engine — proving a regional commit end to end, with no infrastructure below the Site IFC and no dependency yet on a Central IFC. The simulation harness of Section 14.3 is built alongside it, and the site adjudication algorithm's initial parameters are selected from it before Phase 1 exits. |
| **Phase 2** | Reachback and the Central IFC | The Central IFC stood up alongside the first Site IFC: the reachback link, the cross-site coordinator, the central adjudication engine, fine-tuning and corporate-memory functions, and partition tolerance proven by deliberately cutting and restoring the link. The simulation harness is extended to cross-site scenarios and used to select the central adjudication algorithm's parameters (Section 14.3). |
| **Phase 3** | Many Site IFCs, full scale, open | Additional Site IFCs onboarded to cover further service areas, each reaching back to the same |

| PHASE | FOCUS | WHAT IT ESTABLISHES |
|---|---|---|
| | integration | Central IFC; the full specialist roster and policy framework at every tier; high-availability redundancy within each tier; the open-integration surfaces and onboarding gate; standards-conformance evidence; operator consoles at both regional and enterprise scale. |

TABLE 15 · AN INDICATIVE THREE-PHASE PATH FROM A SINGLE SITE IFC TO A FULLY-ONBOARDED, MULTI-SITE PLATFORM REACHING BACK TO ONE CENTRAL IFC. THERE IS STILL NO FIELD-DEPLOYMENT PHASE BELOW THE SITE IFC.

## 14.2 What this architecture commits to

This document fixes the load-bearing decisions: two physical tiers only — a Site IFC per service area and one Central IFC that every Site IFC reaches back to — with no Lamassu-owned sensor or processing infrastructure below the Site IFC; ingestion restricted to sensor technologies that are already deployed and already operated by the city, facility, or third party Lamassu serves, reached directly or through the existing system that manages them, at the Site IFC serving each area; the four planes, split between per-site regional scope and Central-IFC enterprise-wide scope; tiered sensor-specialist inference with a site coordinator and a local adjudication engine at every Site IFC, and a cross-site coordinator and a central adjudication engine at the Central IFC, each commit resolved as an Origi assurance case; a common, adaptive ingestion pipeline over standard protocols, replicated per Site IFC; a strong, deny-by-default policy framework with a human policy layer, at both tiers; Origi, the Trust Assurance Framework, at runtime, with per-tier high-availability redundancy and cross-tier partition tolerance as two independent resilience mechanisms; explainability grounded in the chain of evidence and the trust dossier behind every indicator, at both tiers; and an open, standards-conformant integration posture. It deliberately leaves model checkpoints, exact schemas, connector inventories, the hardware bill of materials for either tier, and detailed capacity planning to detailed design.

## 14.3 Specifying the adjudication algorithm: a simulation-driven methodology

Section 5.3 fixes the adjudication engine's structure as an architectural invariant: a commit occurs only when a composed indicator meets its bound policy threshold and its supporting evidence chain validates. What that section deliberately leaves open is the adjudication algorithm — the specific confidence-aggregation function, per-mission thresholds, conflict-resolution rule, and partition-reconciliation rule that run inside that structure. These are empirically-determined parameters, not properties derivable analytically in advance: the right balance between detection accuracy, false-commit rate, latency, and fairness is a set of trade-offs to be explored, not calculated. Heading into detailed design, Lamassu specifies the adjudication algorithm through simulation rather than through up-front derivation.

## Why simulation, not analysis

An adjudication algorithm's parameters interact in ways that resist closed-form reasoning: how a confidence-aggregation function should weight three corroborating specialists against one dissenting one, where a per-mission threshold should sit between missed detections and false alarms, and how a reconciliation rule should resolve a Site IFC's partitioned conclusion against the Central IFC's wider picture are all questions best answered by measuring candidate algorithms against realistic scenarios, including adversarial ones, rather than by derivation from first principles.

## The methodology

| STEP | WHAT IT ESTABLISHES |
|---|---|
| **1 · Parameterize, don't hardcode** | The confidence-aggregation function, thresholds, conflict-resolution rule, and reconciliation rule are expressed as a swappable module family inside the fixed assurance-case gate — never as one hardcoded design — so simulation is selecting parameters and modules, not re-deciding the architecture. |
| **2 · Build one fixed simulation harness** | A synthetic and replayed multi-modal, multi-site telemetry environment with labeled ground truth — corroborating, conflicting, degraded, and adversarial cases, plus partition and reconnection scenarios — built once, so every candidate algorithm is scored against the same test bed. |
| **3 · Score on Origi's trust components, not accuracy alone** | Detection accuracy, calibration error, confident-wrong-answer rate, false-commit and false-defer rate, time-to-commit latency, fairness across service areas and subgroups, and reconciliation correctness after partition — a multi-objective scorecard, not a single number. |
| **4 · Search the parameter space** | Structured design-of-experiments or multi-objective search (e.g. Bayesian or evolutionary methods) over the algorithm family, using the simulation harness as the fitness function, producing |

| STEP | WHAT IT ESTABLISHES |
| --- | --- |
| **systematicall y** | a Pareto frontier of candidate algorithms rather than one scalar-optimal answer. |
| **5 · Adversarially stress-test the leading candidates** | The strongest candidates from the frontier are red-teamed before selection — attempts to force a false commit or a missed detection, hardened in response — reusing Origi's own independent adversarial validation practice (Section 9.6). |
| **6 · Select and document as an Origi assurance case** | The chosen algorithm's claim ("this algorithm meets its bound policy's performance, safety, and fairness thresholds"), argument (the simulation study's design), and evidence (the simulation run's signed, reproducible results) are recorded exactly as any other committed indicator would be — the choice of algorithm is itself defensible, not asserted. |
| **7 · Retain the harness as a standing regression suite** | Re-run whenever a specialist model, mission, or policy threshold changes after fielding, consistent with Origi's evidence-expiry principle — the algorithm is periodically re-justified, not justified once and assumed forever. |

TABLE 15A · THE SIMULATION-DRIVEN METHODOLOGY FOR SPECIFYING THE ADJUDICATION ALGORITHM, APPLIED INDEPENDENTLY AT THE SITE IFC AND CENTRAL IFC SCOPE BEFORE DETAILED DESIGN PROCEEDS.

This document deliberately stops at the methodology. The specific confidence-aggregation function, threshold values, and conflict-resolution rule are simulation outputs, not something a conceptual architecture can specify in advance; the simulation harness's tooling, synthetic-data generation approach, and search algorithm are themselves detailed-design decisions, to be made by the team that builds the harness.

# 15. Conclusion

Lamassu rests on a trustworthy analytic core, a strong governance layer, and a firm floor — no Lamassu-owned sensor, edge node, or field-processing infrastructure, ever — and places that core on two physical tiers where doing so earns its keep: regional inference next to the already-existing sensors it reasons over, at every Site IFC, and enterprise-wide correlation, fine-tuning, and corporate memory at one rear

Central IFC. Every input Lamassu reasons over comes from a sensor technology that is already deployed and already operated by the city, facility, or third party it serves, ingested at the Site IFC serving that area — never at anything Lamassu itself installs in the field.

The platform is explainable because every indicator, at either tier, is committed as an Origi assurance case and resolves to signed provenance through an observable trace tree; adaptive because every Site IFC ingests diverse, already-existing feeds over the standards they already speak; resilient because a Site IFC does not depend on the Central IFC to keep producing a committed local picture, and because high-availability redundancy protects each tier's own infrastructure independently of that; and scalable because it grows by adding Site IFCs to cover new service areas and by adding compute at either tier — without ever relaxing the trust boundary Origi holds around either, and without ever requiring Lamassu to deploy a sensor, an edge node, or any processing infrastructure of its own below the Site IFC.