



E N K I D U

€ 2 4 M I O V S K I G R O U P € 2 4

C O N C E P T U A L A R C H I T E C T U R E

Enlil

Conceptual Architecture for a distributed Fusion-in-a-Box edge platform mesh fabric,
Anshar zoning, and the Kittu trust harness.

English · Complete edition

Version 0.14 · Fourteenth Draft

3 July 2026

Prepared by MIOVSKI GROUP

Draft for Internal Review

Glossary

Terms used throughout this conceptual architecture. Code names are shown in bold. A note on language: in this document inference denotes the analytical work performed by AI agents – correlating observations, assessing threats, and producing assessments – not human intelligence tradecraft.

Enlil

The overall solution: a distributed Fusion-in-a-Box platform built on a blended mesh / hub-and-spoke fabric, the Anshar platform layer, and the Kittu trust harness.

Fusion-in-a-Box (FIB)

A self-contained edge compute unit – a FIB node or a FIB hub – that ingests local sensor data, performs on-board inference, and shares observations with neighbours. (The fabric has four tiers: FIB nodes and FIB hubs at the edge, the FOB-IFC forward fusion centre above them, and the rear Central IFC.) A FIB is the onboard intelligence core for its host: for a crewed unit, its virtual combat assistant – an intelligence and decision-support assistant to the combatant, not a participant in the fight; for an autonomous combat platform, its perception-and-decision core – the perception and intelligence-decision core, never the combat decision-maker.

FIB node

A FIB carried by a dismounted team, a crewed vehicle, or an autonomous combat platform (UAS, UGV, or USV), built on a smaller member of the NVIDIA Jetson Orin family (Orin Nano, Orin NX, or AGX Orin), capable of limited local inference.

FIB hub (local hub)

A higher-capacity FIB built on the NVIDIA Jetson Thor T5000, acting as a higher-inference fusion center for a cluster of local nodes.

Central Intelligence Fusion Center (Central IFC)

A rear, data-centre-class facility (NVIDIA GB200/GB300 NVL72-class) that receives observations and assessments from the forward fusion centres, runs Large and Frontier models for deep all-source correlation, fine-tunes the models the edge deploys, and issues direction back down the fabric. Stood up in Phase 3 (Section 1.3).

FOB-IFC (Forward Operating Base Intelligence Fusion Center)

A forward fusion centre at a Forward Operating Base, built on NVIDIA Grace-Blackwell hardware (DGX Station GB300, single or dual). It connects several dozen FIB hubs, runs models up to ~80B parameters to correlate intelligence across them, and reaches back to the Central IFC. The middle tier of the fabric, between the FIB hub and the Central IFC (Section 3.3).

Autonomous combat vehicle

An uncrewed air, land, or maritime platform (UAS, UGV, USV) that hosts a Enlil node or hub as its perception-and-decision core, operating within Kittu's trust-operation zones. Enlil may task it autonomously for intelligence collection; the authority to engage remains with a human operator or the platform's own combat function.

Mesh fabric

A peer-to-peer layer in which neighboring nodes exchange observations directly, with no single point of failure, so the picture survives the loss of any one node or link.

Distributed hub-and-spoke

A layered topology overlaid on the mesh: nodes (spokes) cluster around a local hub, and hubs connect upward to Central IFCs.

Deployment profile

The bearer arrangement a fielding chooses for the fabric: LCSS-integrated (riding existing CAF bearers) or infrastructure-denied (an organic, self-forming MANET). The logical overlays are identical in both (§11.4).

Anshar

The Enlil platform layer, running on every tier of the fabric from the smallest node to the Central IFC. It has two layers: a physical layer of zones where agents access tools, skills, and devices, and where security controls run and a logical layer of continuums that run on top of them, where policy is applied. Anshar also provides the observability pattern across the fabric, tracking threats and notifying the relevant units.

Anshar zone

Anshar's physical layer: a physical construct of deployed fusion compute nodes, hubs, and IFCs, at minimum a set of nodes and a hub where agents reach tools, skills, devices, sensors, and actuators, and where access and security controls are enforced. Geography is an attribute of a zone, not an operating constraint (§4.1).

Anshar continuum

Anshar's logical layer: a logical construct that runs on top of the zones and carries a mission's objectives, clearances, and continuum policy. A continuum can span one zone or many, and many continuums can share a zone; continuum policy is applied here (§4.2).

Kittu

The Enlil trust harness: the runtime instantiation of the Trust Assurance Framework across every tier of the fabric. It adds a trust and security layer on top of NVIDIA NeMo Guardrails, governs agent trust-operation zones, secures communication across the fabric, and captures the end-to-end chain of custody used to score threat confidence.

Trust-operation zone

A graduated capability envelope an agent earns through demonstrated good behavior. It governs whether an agent may read a sensor, write to shared state, or at the highest zone command a collection actuator.

Actuator authority

The authority to command a collection actuator to task, reposition, or steer a collection asset to gather data. An agent that has earned a high trust zone may exercise it autonomously within mission policy. It never extends to a combat or kinetic effect: an engagement is always cued to, and decided by, a cleared human operator or the platform's own combat function.

NeMo Guardrails

NVIDIA's open-source toolkit providing programmable input, dialog, retrieval, execution, and output rails around an LLM-based agent. Kittu sits above it as a second, trust-based control layer. [11][12]

Chain of custody / chain of evidence

The signed, tamper-evident record of every observation and inference who or what produced it, from what inputs used to determine confidence in the existence and severity of a threat.

Enlil Observation

The ingestion pipeline's unit product: a normalized, tagged object carrying identity, origin, payload descriptor, classification, reliability, custody head, and zone/continuum bindings (§6.4, Table 7A). Correlated

observations form tracks; inference over tracks produces assessments each with parent pointers, so lineage is structural.

HMI

Human-Machine Interface. Enlil targets an AI head-up display (HUD), audio command and feedback, and ATAK (or equivalent) integration.

ATAK / CoT

Android Team Awareness Kit and its Cursor-on-Target messaging standard the geospatial situational-awareness fabric Enlil integrates with for operator-facing display and tactical data exchange. [13][14]

CF Int

Canadian Forces Intelligence Branch the military-intelligence authority within the CAF. [CAF-5]

CJOC

Canadian Joint Operations Command the command authority for joint and deployed Canadian operations; its J-2 staff own intelligence tasking and all-source fusion.

CSE

Communications Security Establishment Canada's national SIGINT and cyber-security agency.

DI

Director of Intelligence the intelligence lead within a given command.

PIR

Priority Intelligence Requirement the tasking authority's top intelligence need, per CFAO 20-204. [CAF-6]

RPAS

Remotely Piloted Aircraft System the Canadian term per CFAO 11-025, equivalent to the NATO term UAS. [CAF-7]

Five Eyes / FVEY

The US, UK, Canada, Australia, and New Zealand intelligence alliance; FVEY denotes shared product released among the five partners.

DND Instruction 2004/4

Classification, Handling, and Declassification of Defence Information defines Protected A/B/C, Secret, and Top Secret and their handling procedures. [CAF-3]

SAR

Special Access Required a compartmented-information marking applied above the base classification.

DND Intelligence Policy

CFAO 20-204 governs tasking, collection management, all-source fusion, and Priority Intelligence Requirements. [CAF-6]

Abstract

Enlil is a distributed Fusion-in-a-Box (FIB) platform that places sensing, inference, and decision support at the tactical edge and binds the whole into one coherent, survivable picture. It is an intelligence-fusion platform, not a combat platform: it fuses intelligence from many sources and, through inference, turns it into the picture and the decision support that the forces it serves—dismounted and crewed troops, or autonomous combat platforms—need in order to act. Enlil informs and enables the fight; it does not itself prosecute it. This document is a first-draft conceptual architecture: it establishes the layered model, the node, hub, and fusion-centre tiers, the fabric, and the two Enkidu frameworks at the core, at a level intended to anchor subsequent detailed design and a needs-assessment.

The infrastructure and fabric layers blend a mesh architecture with a distributed hub-and-spoke architecture. Localized hubs—built on the NVIDIA Jetson Thor T5000—act as higher-inference fusion centers for clusters of local nodes, while the nodes themselves, built on the Jetson Orin family (Orin Nano, Orin NX, AGX Orin), each perform limited inference and exchange observations peer-to-peer across the mesh. Above the hubs, a forward intelligence fusion centre (FOB-IFC) at the Forward Operating Base correlates across dozens of hubs, and behind it a rear Central Intelligence Fusion Center (Central IFC) supplies the deepest all-source reasoning—one continuous fabric from the smallest node to the rear data centre.

The platform layer is Anshar: a zoning architecture that spans every tier of the fabric—nodes, hubs, the FOB-IFC, and the Central IFC—and also supplies the observability pattern to the underlying FIBs, tracking the movement of threats and notifying the nodes, hubs, and fusion centres that need to know. Zoning applies to the physical world (sectors of terrain in active combat) and equally to abstract spaces, provided their boundaries and entry/exit conditions are well defined.

The trust harness is Kittu, the Trust Assurance Framework that runs at every tier of the fabric—node, hub, FOB-IFC, and Central IFC. Kittu adds a trust and security layer on top of NVIDIA NeMo Guardrails. It assigns trust-operation zones to agents wherever they run; controls whether an agent may use a sensor or an actuator; ensures every link across the fabric—node-to-node, node-to-hub, hub-to-FOB-IFC, and FOB-IFC-to-Central-IFC—is encrypted and secured; and captures the end-to-end chain of custody of every inference, which it uses to determine confidence in the existence and severity of a threat. At the highest trust zone, an agent may autonomously command collection actuators—repositioning a collection platform or steering a sensor to gather data—but it never commands a combat effect.

Enlil is distributed inference. Each node infers on what it can see; that view is shared with neighbors and hubs and weighed in their inference. Hubs pass observations and assessments up to a forward intelligence fusion centre (FOB-IFC)—a Grace-Blackwell system at the Forward Operating Base that correlates across dozens of hubs with models up to ~80B parameters—which in turn reaches back to the Central IFC; direction flows back down the same path—all under Anshar's zoning and Kittu's trust harness.

The desired operational effect is that each FIB acts as a virtual combat assistant to its operator: it informs, it observes, it records, and it distributes local observations among neighbors seamlessly. For a crewed unit, operators interact through an AI HUD, audio commands and spoken feedback, and ATAK integration; an autonomous combat platform (air, land, or maritime) consumes the same fused picture and tasking machine-to-machine, with the FIB as its perception-and-decision core.

Enlil is, by design, an open platform. It exposes published interfaces so that two communities can integrate their technologies behind the same trust harness: intelligence-asset developers, who bring new sensors, collection sources, and inference models into the fused picture; and combat-platform developers, whose crewed and autonomous systems consume that picture and decision support. Integration is governed rather than open-ended—a published interface lets a capability connect, not command (Section 10)—so the platform grows by accretion without ever relaxing the trust boundary around it.

1. Context and Scope

Enlil exists to give a Canadian forward unit the fused picture and decision support that, today, only a rear Joint Intelligence Centre can provide and to do it on hardware a platoon can carry, surviving the loss of the network and of individual units. Enlil is an intelligence-fusion platform, not a combat platform: it produces the fused picture and the decision support, while the decision to act and the action itself remains with the human commander or the combat platform it serves. It is the conceptual successor to the centralized fusion model and aligns with DND/CAF modernization priorities under the Digital Campaign Plan and the DND/CAF Artificial Intelligence Strategy (2023). Instead of moving all data rearward to be fused, Enlil fuses forward and reconciles upward.

Although one prominent use case is a virtual combat assistant for dismounted and crewed troops integrated with ATAK and visual and audio interfaces (Section 9) it is not the only one, nor the primary one. The Department's requirement centres on autonomous combat platforms across the air, land, and maritime domains: a Enlil FIB is designed to serve equally as the onboard perception-and-decision core of an uncrewed combat platform (UAS, UGV, USV) and as an assistant to a human crew. Throughout this document, a reference to an operator should be read as the host a crew or an autonomous platform except where a human is explicitly required notably a combat effect, which Enlil never authorises (Section 5.2). Because both kinds of host, and the intelligence assets that feed them, evolve independently of Enlil, the platform is built to be open: it offers published interfaces for intelligence-asset developers and for combat-platform developers alike to integrate their technologies behind a single governed trust boundary (Section 10).

Enlil is designed to support Canadian Armed Forces (CAF) operations across multiple domains:

Joint operations under Canadian Joint Operations Command (CJOC).

NORAD air-defence and North American aerospace defence.

Royal Canadian Navy (RCN) maritime domain awareness, including Arctic sovereignty operations.

CAF cyber operations and space-domain awareness.

Expeditionary operations with distributed, forward-deployed platoons and sections in denied or austere environments (Arctic, remote deployment).

Coalition interoperability with Five Eyes partners (US, UK, Australia, New Zealand) and NATO members in joint task forces.

It directly supports the Canadian Defence Policy (2022) technology-modernization pillars, the DND/CAF Artificial Intelligence Strategy (2023), the Digital Campaign Plan, NORAD modernization priorities (Arctic and early-warning), and the Cyber Forces mandate under the Canadian Armed Forces Cyber Operations Centre (CAFCOC). [CAF-1][CAF-2][CAF-13][CAF-14]

Illustrative context Arctic and remote deployment

A Canadian Ranger patrol in the High Arctic operates with limited connectivity and pre-positioned sensors (UGS, RPAS). The FIB node fuses local SIGINT, EO/IR, and acoustic detections, maintains situational awareness over the mesh without satellite reachback, and shares observations with neighbouring nodes and a local hub. When reachback is available, observations are sent to CJOC J-2 for all-source fusion and direction; when connectivity is lost, the local picture survives and reconciles when the link returns.

How to read this draft. An executive pass: the Abstract, this section, the compliance map at §1.5, and the worked scenarios of §9. An operational pass adds §3 (the fabric), §6 (what is collected and how it enters), and §4 (zones, continuums, and policy). The engineering read is linear. Technical depth arrives progressively operational concept first, mechanism second, hardware last.

1.1 What this draft covers

This document is organized around Enlil's layers and the two frameworks at its core:

The architectural principles that govern every layer (Section 2).

The infrastructure and fabric layer – the blended mesh / distributed hub-and-spoke topology and the FIB node and hub tiers (Section 3).

The platform layer (Anshar) – zoning across nodes and hubs, and the observability pattern that tracks threats and notifies the right units (Section 4).

The trust harness (Kittu) – trust-operation zones over NeMo Guardrails, sensor/actuator control, secure comms, and the chain of custody behind threat confidence (Section 5).

Intelligence data and ingestion – the eight source disciplines, what each produces, where it is collected (field FIB vs. Central IFC), the sensors required, and the ingestion pipeline (Section 6).

The inference layer – the model tiers Enlil fine-tunes (tiny to medium) for raw-data inference and for the Anshar and Kittu agents, and how the larger tiers are phased (Sections 1.3 and 7).

Distributed inference – how nodes, hubs, and Central IFCs share observations and assessments and exchange direction (Section 8).

The human-machine interface – AI HUD, audio, and ATAK integration (Section 9).

The conceptual hardware specification – ruggedized node and hub configurations, their interfaces and inter-FIB communications, and estimated power, operating time, and weight (Section 11).

It closes with the open questions and the research a detailed design must resolve (Section 13). As a first draft, it deliberately states assumptions and flags where hardware-capability research is still required. This architecture aligns with the Canadian Defence Policy (2022), the DND/CAF Artificial Intelligence Strategy (2023), and the CFSP classification and handling procedures (DND Instruction 2004/4).
[CAF-1][CAF-2][CAF-3]

This draft is written for three readerships at once: DND/IDEaS evaluators scoring the architecture against the challenge outcomes (the map at §1.5 is theirs), CAF operational sponsors placing the system in their formations (§1.4 and §9 are theirs), and the engineering team of record that will carry the concept into detailed design.

1.2 What this draft does not yet cover

Exact model selections and quantization budgets, network/RF waveform down-selection, accreditation mapping, and the formal interface specifications between layers are out of scope for this draft and are listed as follow-on work in Section 13. Section 11 gives a conceptual hardware specification – representative configurations, interfaces, communications, power, and weight – but not a final bill of materials. The HMI options in particular are pending detailed research into the specific capabilities of each Jetson tier.

1.3 Delivery in three phases

Enlil is delivered in three phases. Nothing in this document is out of scope; rather, the architecture is sequenced so the forward capability lands first and the deep rear capability follows.

Phase 1 – Forward fusion. The Fusion-in-a-Box hubs and the FOB Intelligence Fusion Centres (FOB-IFC): cluster fusion at the hub and forward all-source correlation at the FOB-IFC with models up to ~80B parameters. This is the focus of the current draft.

Phase 2 – Edge proliferation. The nodes pushed onto dismounted teams, crewed vehicles, and autonomous combat platforms (UAS, UGV, USV) across the air, land, maritime and space domains.

Phase 3 Central depth. The Central Intelligence Fusion Centres and the fine-tuning of the Large and Frontier models that the edge consumes the data-centre tier (GB200/GB300 NVL72-class).

Phase 1 is specified here in detail; the nodes (Phase 2) and the Central IFC with its Large- and Frontier-model fine-tuning (Phase 3) are described at the conceptual level and built out in later increments. Where earlier drafts said a capability was out of scope, it should now be read as a later phase.

1.4 Where Enlil sits in the CAF ecosystem

Enlil is not a new network, a new radio, or a new common operating picture. It is a software fusion-and-compliance layer that rides the C5ISRT infrastructure the CAF already operates on land, the Land Command Support System (LCSS) and its bearers; above battalion, the joint data links; behind it all, the reachback enterprise. Concretely, the fabric's four tiers map onto the LCSS domain model a CAF signaller already carries: [CAF-18]

- " Enlil tier CAF / LCSS placement What rides where
- " Micro / standard node Soldier domain ISSP end-user devices running ATAK; small platforms and emplaced sensors The FIB is the operator's and the platform's on-board intelligence core; it emits tracks and alerts into the operator's existing ATAK COP over CoT (§9).
- " Heavy node / FIB hub Mobile domain battlegroup and combat-team command posts; LAV 6-class platform compute and comms The hub aggregates its node cluster and hosts the compliance seam where local fusion meets the unit's C2.
- " FOB-IFC HQ domain Brigade / Task Force HQ; adjacency to the All-Source Sensors Integration Cell (ASIC) Theatre-level fusion beside the ASIC, and the natural joint data-link gateway point Link 16 is integrated in the LCSS at the TF HQ level (§10.3).
- " Central IFC Strategic enterprise CJOC J-2 / NDOIC reachback National all-source direction and feeds down; assessments, provenance, and exportable audit up (§12.6).

Table 1A · Enlil tiers on the CAF C5ISRT / LCSS map. The fabric occupies the ecosystem's existing seams; it does not create new ones.

Figure 1A · Enlil tiers overlaid on the LCSS domain model, from the Soldier domain to the strategic enterprise. Observations propagate upward; direction and national feeds flow down; the operating claim rides the footer.

Read against the CAF's four planning levels government policy, strategic, operational, tactical the same mapping tells each echelon what it consumes. The tactical level consumes the live local picture on the node and hub; the operational level consumes the FOB-IFC's fused theatre picture; the strategic level consumes Central IFC assessments and the exportable audit and provenance record. Government policy consumes none of it directly but the audit chain is what lets the department answer for what the system did.

The operating claim, stated once: Enlil adds fusion, compliance, and provenance to the CAF's existing infrastructure. It does not replace networks, radios, or the COP it rides existing bearers (§3.1, §11.4) and reports into the existing picture (§9, §10.3).

1.5 Compliance map against the IDEaS challenge outcomes

The challenge states its essential outcomes (4.3.1 4.3.9) and desired outcomes (4.4.1 4.4.4). The map below ties each to the mechanism that satisfies it and the section that carries the evidence the navigation aid for an evaluator scoring this document against the requirement list. [1][CAF-18]

- " Outcome Requirement (abridged) Where in this architecture mechanism

- " 4.3.1 Modular AI-enabled component enforcing classification rules and policy during multi-sensor fusion §4.4, §4.5, §5, §6.4 policy bound per continuum, enforced per action at the zone by Kittu; classification tags applied at ingest and checked at every share.
- " 4.3.2 Enforcement from machine-readable policy definitions §4.4 (Table 3B) five guardrail dimensions as machine-checkable policy; formal schema in detailed design (§13).
- " 4.3.3 At least two sensor modalities §6.1 §6.2 eight disciplines with sensor classes per tier; both §9 scenarios fuse two or more.
- " 4.3.4 Security domains (at least network security) §5.3, §6.4 stages 3 and 9 authenticated admission, transport security over any bearer, CDS-brokered domain crossings; the §9.3 scenario crosses two domains.
- " 4.3.5 Classification levels (at least Protected B) §6.6, §4.5 Protected B tagging, handling, and zone/continuum binding end to end.
- " 4.3.6 Programmatic checks and enforcement without human approval for predefined conditions §4.4, §5.2, §6.4 stage 8 deterministic per-action resolution; autonomous within policy at high trust; humans enter only where policy names them.
- " 4.3.7 Provenance records: source sensor, classification markings, timestamps, domain of origin §6.4 stages 4 and 7, Table 7A the Observation carries all four; signed, hash-linked custody from acquisition.
- " 4.3.8 Audit logs: rules applied, enforcement actions (permit, restrict, downgrade, segregate), dispositions §12.6, §6.4 stage 8 the enforcement vocabulary is the pipeline's own; every disposition logged with policy_id and policy_version.
- " 4.3.9 Lineage traceability from ingestion through fusion output; exportable §12.6, §6.4 the derivation ladder makes lineage structural; JSON/XML/CSV export for accreditation, forensic, and ATIA review.
- " 4.4.1 Real-time enforcement suitable for tactical decision-making §6.5, §7 §8, §11 stated per-stage latency budgets; inference forward on the edge tiers.
- " 4.4.2 Policy updates without system restart; rapid context transition §4.4 zone and continuum policies are live-versioned; a context transition is a policy binding, in minutes.
- " 4.4.3 SWaP and compute limits incorporated; compliance maintained under them §11, §6.5 reference profiles per tier; the compliance stages are load-shed-exempt by design.
- " 4.4.4 Explainability, human-readable decisions, controlled override with accountability §9.2, §12.6 a one-line rationale per enforcement action; cleared-operator override as a first-class audit event.

Table 1B · Compliance map: every essential and desired outcome of the challenge, the mechanism that answers it, and where the evidence lives.

Two of the desired outcomes are demonstrated rather than merely claimed: the §9.3 and §9.4 scenarios walk outcome 4.3.4 and 4.3.8 end to end, audit line included, and §6.5 turns outcome 4.4.1 into testable numbers.

2. Architectural Principles

Eight principles govern every layer of Enlil and should be used to adjudicate design trade-offs in the detailed phase.

Infer forward. The inference that decides the next minute runs on the node, not in the rear. Reachback enriches; it is never a dependency for tactical decision-making.

Survive partition. The mesh has no single point of failure. Any node, hub, or link may be lost; the surviving fabric continues to observe, infer, and inform.

Share by default, within policy. Observations propagate to neighbors and hubs seamlessly, and locally-obtained observations always propagate upward through the fabric; every downward or cross-domain share, by contrast, is brokered by Kittu against trust and classification policy.

Never trust an agent or a peer implicitly. Capability is earned and continuously re-verified. An agent's zone, and a peer's standing, can be revoked the instant evidence turns.

Bound everything in a zone. Every operation runs on a physical zone and within a logical continuum that has explicit entry and exit conditions, so behavior is legible, measurable, and observable.

Evidence before confidence. No threat is asserted without a chain of custody. Confidence in a threat's existence and severity is a function of the evidence behind it, not the fluency of the assessment.

The FIB serves its host. Each unit is the onboard intelligence core of its host—a crewed unit's virtual combat assistant or an autonomous combat platform's perception-and-decision core. It informs, observes, and records; it keeps a human in control of every combat effect, while intelligence collection it may task autonomously once it has earned a high trust zone.

Fuse intelligence; never authorise combat. Enlil fuses intelligence data; it does not fight. Its actuator authority reaches only to intelligence collection—tasking, repositioning, or steering collection assets to gather data—which an agent that has earned a high trust zone may issue autonomously within mission policy. It never authorises a kinetic or combat effect: an engagement is cued to, and decided by, a cleared human operator or the combat platform's own function. And locally-obtained observations always propagate upward through the fabric; classification and releasability controls govern downward and cross-domain sharing, never that upward flow.

2.1 Canadian security and compliance context

Enlil operates within the Canadian defence security ecosystem and is designed to adhere to its governing instruments:

DND Instruction 2004/4—Classification, Handling, and Declassification of Defence Information (defines Protected A/B/C and Secret/Top Secret categories). [CAF-3]

CFSO.04.035—Communications Security Orders: encryption, key management, and secure-comms protocols for the CAF. [CAF-4]

DITSR (DND IT Security Requirements)—system security controls, vulnerability management, and access control. [CAF-8]

ITSP.40.111—Government of Canada cloud security profile, for any central-fusion-centre cloud components. [CAF-10]

PIPEDA and the Access to Information Act—personal-information handling and ATIA-disclosure protocols. [CAF-11]

DND accreditation framework the formal accreditation pathway for classified systems (DND Trusted IT Product Evaluation Program TIPTOP or equivalent).

The Kittu trust harness (Section 5) instantiates these policies as machine-readable constraints, and the chain of custody (Section 5.3) supports accreditation audits and forensic review under ATIA and Privacy Act requests. This Canadian control baseline is the standard against which the detailed design is accredited.

3. Infrastructure and Fabric Layer

Enlil's lowest layers – the physical compute units and the network that joins them – blend two topologies deliberately. A mesh gives peer-to-peer resilience; a distributed hub-and-spoke overlay gives the inference hierarchy. The two coexist: nodes are simultaneously mesh peers to their neighbors and spokes to their local hub.

3.1 The blended topology

At the lowest level every FIB is a mesh peer. Neighboring nodes exchange observations directly over a self-healing, low-power mesh, so the loss of any node degrades coverage gracefully rather than catastrophically, and there is no central relay whose failure is fatal. Overlaid on that mesh is a hub-and-spoke hierarchy: a cluster of nodes is served by a local hub that performs heavier, higher-confidence inference on the cluster's pooled observations; hubs connect upward to a forward intelligence fusion centre (FOB-IFC) at the Forward Operating Base, which correlates across dozens of hubs and reaches back to the Central IFC. This is the localized hubs as higher-inference fusion centers for local nodes' pattern, with the mesh underneath guaranteeing the picture survives even when the hierarchy is severed.

"	Topology element	Role in Enlil	Failure behavior
"	Mesh (peer-to-peer)	Direct node-to-node observation sharing among neighbors; the substrate of distributed inference.	Self-healing; loss of a node or link reroutes around it. No single point of failure.
"	Spoke (node -> hub)	A node passes pooled observations and local assessments to its local hub for higher inference.	If the hub is lost, nodes continue on the mesh and can re-affiliate to another hub.
"	Hub (local fusion)	Higher-confidence inference over a cluster; relays up to, and direction down from, central fusion.	Degrades to node-only inference; a neighboring hub can absorb orphaned nodes.
"	Reachback (hub -> FOB-IFC -> Central IFC)	Raw observations and assessments up; direction and corporate memory down.	On loss, hubs and nodes operate autonomously and reconcile when the link returns.

Table 1 · The blended mesh / distributed hub-and-spoke fabric and how each element behaves under loss.

Figure 1 · The blended fabric. Nodes are mesh peers to their neighbors and spokes to a local Thor-based hub; hubs reach back to Central IFCs. The mesh underlay keeps the picture alive when the hierarchy is cut.

One property of the topology deserves an explicit sentence: the mesh underlay and the hub-and-spoke overlay are logical constructs. Any IP bearer that meets the fabric's latency, throughput, and partition envelope can carry them – the LCSS data bearers, SATCOM, LTE/5G, or an organic tactical mesh where no infrastructure exists (§11.4). The fabric does not assume, and does not require, a particular radio.

3.2 FIB node and hub tiers

Enlil defines several FIB tiers, each mapped to a member of the NVIDIA Jetson family so that size, weight, power, and inference capacity match the role. The smaller Orin modules share a common form factor and software stack, so a node can be specified at the right cost/power point and scaled without re-architecting; the Thor T5000 hub is a generational step up in compute for local fusion. [6][7][8][9][10] The figures below are drawn from NVIDIA and ecosystem documentation and should be re-confirmed against current datasheets in detailed design.

"	FIB tier	Jetson module	Approx. AI compute	Power	Role in the fabric
"	Micro node	Orin Nano (4.8 GB)	~20.67 TOPS	7.25 W	Dismounted / wearable sensing; single-stream perception; minimal local inference; mesh peer.

- " Standard node Orin NX (8 16 GB) ~70 157 TOPS 10 40 W Vehicle, squad, or autonomous-platform (UAS/UGV/USV) node; concurrent multi-sensor perception; small on-board models; mesh peer + spoke.
- " Heavy node AGX Orin (32 64 GB) ~200 275 TOPS 15 60 W Section/platoon node or autonomous-vehicle core; multi-camera fusion; larger local models; candidate sub-hub.
- " Local hub Thor T5000 ~1000+ TOPS (up to ~2070 FP4 TFLOPS) 40 130 W Higher-inference fusion center for a node cluster; runs the heavier Kittu and Anshar services.
- " FOB-IFC DGX Station GB300 ~20 PFLOPS FP4; 748 GB coherent ~1.6 kW Forward all-source correlation across dozens of FIB hubs; runs up to ~80B models; reaches back to Central IFC (Section 3.3).

Table 2 · FIB tiers Jetson node and hub tiers plus the Grace-Blackwell FOB-IFC. Compute/power figures from NVIDIA and ecosystem sources; confirm against datasheets. [6][7][8][9][10][98]

Two consequences matter for the architecture. First, inference is tiered: a micro node may run only a quantized detector and abstain to its hub for anything harder, while the Thor hub can run multiple concurrent models and reason over the cluster s pooled picture. Second, the form-factor commonality across the Orin tiers means a single carrier/enclosure design can host Nano/NX and an AGX variant can be a near-drop-in heavy node, easing logistics and sparing.

Figure 2 · FIB tiers from micro node to Thor-based local hub, with the inference each tier owns and where it abstains upward.

3.3 The FOB Intelligence Fusion Center (FOB-IFC)

Above the FIB hubs sits a third compute tier: the FOB Intelligence Fusion Center (FOB-IFC), a forward fusion centre deployed at a Forward Operating Base. A FOB-IFC connects several dozen FIB hubs over the tactical network and runs models up to ~80B parameters to correlate their assessments into a forward all-source picture the job that, until now, only a rear Central IFC could do. It is the middle tier of the fabric: nodes feed hubs, hubs feed the FOB-IFC, and the FOB-IFC reaches back to the Central IFC.

The FOB-IFC is built on NVIDIA Grace-Blackwell hardware. The baseline is the DGX Station built on the GB300 Grace Blackwell Ultra Desktop Superchip a deskside system with data-centre-class compute that is air-cooled and transit-case deployable, unlike the liquid-cooled rack systems of the Central IFC.

- | Option | Configuration | Role / fit |
|--------|--------------------------------|--|
| " | DGX Station GB300 (single) | 1× GB300 Grace Blackwell Ultra superchip; 748 GB coherent memory (252 GB HBM3e + 496 GB LPDDR5X); ~20 PFLOPS FP4; ConnectX-8 800 Gb/s; ~1.6 kW, air-cooled. [98] Baseline FOB-IFC: runs up to ~80B models; correlates across dozens of FIB hubs. |
| " | DGX Station GB300 (dual) | 2× DGX Station linked over ConnectX-8 (800 Gb/s); ~1.5 TB coherent; ~40 PFLOPS FP4; ~3.2 kW. [98] High-capacity or redundant FOB-IFC: more concurrent correlation models and larger context. |
| " | GB200-class rack (alternative) | 2 4 Blackwell GPUs (GB200 NVL2/NVL4, or HGX B200/B300) in a rugged 2 4U shelter rack. [100][101] Where a shelter rack and higher aggregate throughput are preferred over a deskside box. |
| " | GB200/GB300 NVL72 (boundary) | Rack-scale, 72 GPUs, ~13.4 TB, ~120 kW, liquid-cooled. [99] Not FOB-deployable this is the Central IFC (Phase 3), shown for context. |

Table 2A · FOB-IFC hardware options on NVIDIA Grace-Blackwell. The DGX Station GB300 (single or dual) is the baseline; the NVL72 rack is the Central-IFC boundary, not a FOB system. [98][99][100][101]

Two properties make the FOB-IFC the right forward tier. First, memory: a single DGX Station s 748 GB of coherent memory holds an 80B-class model with room for long context and several concurrent correlation models, so a FOB can fuse the output of dozens of hubs in one place. Second, deployability: it is shelter-, vehicle-, or transit-case-mounted and runs on Forward Operating Base or generator power, air-cooled unlike the liquid-cooled, 120 kW NVL72 rack of the Central IFC. The FOB-IFC reaches back to the Central IFC for the deepest all-source reasoning and for updated model weights, but it is not dependent on that link: if reachback is cut, the FOB-IFC sustains forward fusion on its own, exactly as the hubs and nodes do below it.

4. Platform Layer Anshar

Anshar is the Enlil platform layer, and it runs on every tier of the fabric from the smallest node up to the Central IFC. Anshar has two layers. The physical layer is a set of zones (§4.1): each zone is a grouping of the fabric's own compute nodes, hubs, and IFCs and it is where agents actually reach tools, skills, devices, sensors, and actuators, so the access-and-permissions controls live here. The logical layer is a set of continuums (§4.2): a logical fabric that runs on top of the zones, where each mission's policy is applied. A continuum can span one zone or many, and many continuums can share the same zone. Alongside the two layers, Anshar provides the observability pattern (§4.3) that tracks threats across the fabric and notifies the units that need to know. Two frameworks govern what runs: a policy framework (§4.4) that sets the operational guardrails applied at each continuum and enforced at the zone and a security-classification framework (§4.5) that binds intelligence classifications to continuums, zones, and agents. Together, Anshar and the Kittu trust harness (§5) make explicit, and machine-checkable, what each agent may do, with which data, under whose authority.

4.1 The physical layer: Anshar zones

A Anshar zone is the platform's physical unit: a grouping of the fabric's own compute nodes, hubs, and IFCs with, at a minimum, a set of nodes and a hub. A zone is where agents run and where they reach the tools, skills, devices, sensors, and actuators available to them, so it is the layer at which access and permissions are enforced. Kittu's trust-operation zones (Z0-Z3, §5.2) are the per-agent capability envelopes applied within a zone; they are not the zone itself.

Geography is an attribute of a zone and of its members, not an operating constraint. A zone records where its nodes, hubs, and IFCs are, but it is not defined by a patch of ground and it is not confined to one; its membership can move as units move. What a zone is for the objective it serves, the policy it runs under, the entry and exit conditions of the work belongs to the logical layer that runs on top of it (§4.2), not to the physical grouping itself.

- " zone attribute Meaning in Enlil
- " Members The nodes, hubs, and IFCs that make up the zone at a minimum, a set of nodes and a hub.
- " Location Where the members are a recorded attribute of the zone and its components, not a limit on what the zone may do.
- " Access surface The tools, skills, devices, sensors, and actuators the zone exposes to the agents that run on it.
- " Security control The trust-operation zones (Z0-Z3) and least-privilege grants Kittu enforces for every agent in the zone.
- " Owner The node or hub accountable for the zone (can move as units move).

Table 3 · A zone and its attributes. A zone is a physical grouping of fabric compute; its purpose is carried by the continuums that run on it (§4.2).

4.2 The logical layer: Anshar continuums

The zones are Anshar's physical layer; the continuums are the logical layer that runs on top of them. A continuum is the unit in which Enlil bundles a set of objectives with the policies, clearances, and guardrails needed to pursue them architecturally, a logical agent-execution context that runs over one or more zones. A continuum can span a single zone or many, and the same zones can host many continuums at once, each running under its own parameters, policies, and classification ceiling, isolated from the others by the very machinery of §4.4 and §4.5. This is where a mission's policy is applied; the zone underneath is where that policy is enforced as concrete access to tools, skills, devices, and sensors.

A Enlil continuum is an intelligence-fusion mission, not a combat mission. It is the context in which

intelligence drawn from many sources and disciplines – the observations of nodes, the cluster assessments of hubs, and the all-source correlation of a FOB-IFC – is brought together and turned, through inference, into a conclusion a commander can act on: a threat identified, a track confirmed, a position resolved. The mission sets what is to be concluded, with which data, and under whose authority; the action that may follow belongs to the troops or the autonomous platforms the mission supports, never to Enlil itself. The examples below are therefore fusion problems, distinguished by the sources they draw on and the tiers of the fabric they exercise.

Four mission archetypes recur across the Canadian use cases:

Perimeter defence. A standing mission around a fixed site – a Forward Operating Base, a logistics node, an airfield. Pre-positioned and fixed sensors (EO/IR, acoustic, ground radar, RF) feed a FIB that fuses their detections into a continuous read of who and what is approaching, classifies potential threats, and raises them to the operator before they reach the wire. Because the sensors and the picture are local and persistent, perimeter defence sits naturally at the Forward Operating Base and its FOB-IFC, where the standing picture can be held and correlated over time – a clear use case for the FOB tier.

Search and rescue. A time-critical mission to locate friendly personnel or assets – a downed aircrew, a separated patrol, a vessel in distress. It fuses every source that might fix a position: last-known tracks, beacon and SIGINT hits, EO/IR sweeps, and the observations of nodes already in the area. The work is inherently collaborative: the FIB hubs nearest the search area pool what their clusters can see, while the FOB-IFC correlates across those hubs and against rear holdings to narrow the area and direct the next look. It is a strong example of FOB-IFC and FIB hubs working the same problem from different vantage points.

Hunter-seeker. The adversarial counterpart of search and rescue: a mission to detect and localise hostile combatants or assets that are actively trying not to be found. It fuses the same breadth of sources but weights movement, deception, and pattern-of-life more heavily, and it holds custody of contacts as they move between sectors. Like search and rescue, it leans on FOB-IFC-and-hub collaboration – the hubs maintain local track custody and hand contacts across zone boundaries, while the FOB-IFC fuses their reporting into a coherent assessment of where the adversary is and where it is going. The conclusion is presented to a cleared operator; whether and how to act on it remains the operator's decision.

Combat assistant. A mission that fuses intelligence in direct support of an action in progress. The FOB-IFC and the FIB hubs analyse the live picture and deliver it where it is needed – to a crewed unit as a virtual combat assistant, pushing the fused picture, cueing, and recommendations through the HUD, audio, and ATAK; or to an autonomous combat platform machine-to-machine, as the perception-and-decision input to that platform's own systems. Enlil supplies the intelligence and the decision support; the actioning the engagement itself – remains with the human crew or with the platform's own combat function, and need-to-action stays with a cleared human (§4.5).

- " **Mission attribute** – What it sets
- " **Objectives** – What the mission is to achieve – the Priority Intelligence Requirements, named areas, tracks, or effects it owns.
- " **Duration** – The mission's authorised lifetime; expiry triggers teardown and capability revocation.
- " **Resources** – The nodes, hubs, sensors, actuators, models, and agents allotted to the mission.
- " **Trust level** – The Kittu trust ceiling agents in the mission may earn (Z0–Z3), and the human-approval rules that apply.
- " **Clearance ceiling** – The maximum classification, and caveats, the mission may handle – the ceiling for every agent enrolled in it.
- " **Policy set** – The policies (§4.4) in force for the mission – the guardrails its agents execute under.

Table 3E · A continuum descriptor binds objectives and resources to a trust level, a clearance ceiling, and a policy set. Instantiating it creates a continuum – a logical agent-execution context over one or more zones.

Figure 3A · Two continuums running concurrently over shared zones. Each binds its own clearance ceiling, trust level, and policy set over the same physical fabric; Kittu enforces the boundary between them, and a

cross-mission read is denied and logged.

Because a continuum is a logical construct over the zones, its lifecycle is the lifecycle of applying policy:

Authorise. A mission is assigned its objectives, duration, resources, trust level, clearance ceiling, and policy set by a cleared authority – a human act, recorded in the chain of custody.

Instantiate. Anshar instantiates the continuum and binds the allotted resources to it; the mission's policy set and clearance ceiling become the constraints in force inside it.

Enrol. Agents, sensors, and actuators are admitted to the mission; each agent inherits the mission's clearance ceiling as its policy ceiling and starts at the trust floor, earning capability through good standing (§5.2).

Execute. Agents act inside the continuum under continuous Kittu enforcement; every action is checked against the mission's policy and classification, and every observation is custody-recorded against the mission.

Terminate. On objective completion or expiry, the continuum is torn down; capabilities are revoked instantly, agents return to baseline, and the mission's chain of custody is sealed and retained.

Isolation between concurrent missions is not a convenience; it is enforced exactly as classification and trust are enforced everywhere else. An agent enrolled in a Secret mission cannot read a Top Secret mission's tracks even on the same hub, because its clearance ceiling forbids it and Kittu will not pass the data. Data crosses a mission boundary only through Kittu releasability brokering, never implicitly. And because trust is earned slowly and lost instantly (§5.2), mission teardown – or any hard-gate event inside a mission – revokes capability before the next action, so a compromise is contained to the mission in which it occurred.

Two missions, one hub. A Thor hub simultaneously hosts MISSION ARTEMIS – a Secret ISR task holding custody of vehicle tracks in a named sector – and MISSION BOREALIS – a Top Secret, CAN-EYES-ONLY SAR-compartmented task. Both draw on the same EO/IR and SDR feeds, but each runs as its own continuum. ARTEMIS's agents are cleared to Secret and may correlate and recommend on their tracks; BOREALIS's agents, cleared to Top Secret and inside the compartment, see reporting ARTEMIS cannot. An ARTEMIS agent that requests a BOREALIS track is denied at the classification check (§4.4) and the request is logged. Neither mission's human decision-maker is removed from the loop – need-to-action remains theirs. When ARTEMIS expires, its continuum is torn down and its agents' capabilities revoked, while BOREALIS continues untouched.

Taken together, a continuum is the place where Enlil's policy engine, its classifications, its zones, and its trust harness meet: the continuum descriptor names the policies and clearances; Anshar instantiates it as a logical construct over the zones; the agent execution layer runs inside it; and Kittu enforces every classification, capability, and trust boundary at the point of action and records it in the chain of custody. The detailed wiring of these frameworks into the trust-assurance harness is developed in Section 5 and refined in the trust-assurance work that follows.

4.3 Anshar as the observability pattern

Anshar is also how Enlil watches itself and the threat picture. Every zone transition is an observable event: a track entering a sector, a contact changing classification, a node abstaining to its hub. Anshar publishes these events across the fabric so that the units with a stake in a zone are notified the moment something material changes within it. Concretely, the observability pattern provides:

Threat tracking. As a threat moves, it crosses zone boundaries; Anshar records each crossing and maintains continuous custody of the track across the nodes that observe it.

Relevance-scoped notification. When a zone's state changes, Anshar notifies the nodes and hubs for which that zone is relevant – the neighbor about to inherit the track, the hub that owns the sector – rather than flooding the whole fabric.

A common operating picture. Because every unit shares the same zone vocabulary, the local pictures compose into one coherent whole without a central server having to assemble it.

Zoning and observability are the same mechanism seen from two sides. Defining a zone's boundaries is what makes a threat's movement across it observable and what makes notifying exactly the right units possible.

Figure 3 · The Enlil layered model. Infrastructure and fabric at the base; the Anshar platform layer split into its physical zones where access and security controls run and the logical continuums that run on top of them, where policy is applied; the Kittu trust harness wrapping the agents; and the HMI on top.

4.4 The policy framework

Enlil carries two policy frameworks, one at each Anshar layer. A zone policy is the baseline: it is set when a zone is configured and deployed, and it governs what any agent on that zone may do by default which tools, sensors, and devices are reachable, and the security controls that apply. A continuum policy is mission-specific: it is bound to a continuum when the mission is instantiated, it runs on top of the zones the continuum spans, and it may tighten or override the zone policy for the life of the mission. Where the two speak to the same question, the continuum policy wins; where the continuum is silent, the zone policy holds. The result is a single effective policy, resolved per agent and per action at the point of use.

Both frameworks are live. A zone-policy amendment propagates to the zone's members and takes effect per action, with no process restart; a new or amended continuum policy binds at the next instantiation or is pushed to a running continuum as a version bump. Every policy change is itself a chain-of-custody event versioned, signed, and auditable and every enforcement record names the policy version it was resolved against (§12.6). The operational meaning: transitioning between operational contexts a new coalition caveat set, revised collection constraints is a policy binding, performed in minutes, not a maintenance window (outcome 4.4.2).

Whichever framework a rule comes from, it governs what happens inside a continuum (§4.2) and is enforced as concrete access and permissions at the zone (§4.1). A continuum is a three-layer execution context, structured into a policy layer, a conops layer, and a reporting layer, and the separation between them is the load-bearing principle of the whole architecture. The policy layer states what must hold the governing instruments, classifications, and constraints that bound the mission. The conops layer is how the work is done the operational steps, and the sensors and actuators through which agents act. The reporting layer is how the mission is measured the observability events, KPIs, and audit obligations Anshar already publishes (§4.3). The structure is adapted from a zone model proven in a regulated medical deployment and recast here for the intelligence-and-effects domain. [2]

- " Layer What it holds in Enlil Where agents may operate
- " Policy The governing instruments and constraints for the zone: the intelligence-data classification ceiling, the releasability caveats, the tools, sensors, and actuators permitted, the trust-zone floor, and the human-approval rule. No agent. Policy determinations and any decision for a combat effect are human. Agents read policy as a constraint; they never set it.
- " Conops The operational steps of the mission, and the sensors (inputs) and actuators (outputs) through which the work is sensed and acted on. Agents act here as conops-layer sensors and actuators, gated by Kittu trust zones (Z2 to sense, Z3 to act).
- " Reporting The observability events, KPIs, chain-of-custody entries, and audit obligations the mission emits (§4.3, §5.3). Agents act here as measurement synthesisers correlating, scoring, summarising; never as deciders.

Table 3A · The three layers inside every continuum. Agents inhabit the conops and reporting layers only; the policy layer is human.

The policy layer is human. No agent edge, generative, or predictive inference alike 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 a combat effect, is made by a cleared

human who is accountable for it. Systems do not make combat decisions, because they cannot be responsible for them; deciding to task a collection asset is not such a decision, and an agent that has earned a high trust zone may take it autonomously. This is the architecture's overriding principle, and the policy framework is how it is enforced rather than merely asserted.

Within the policy layer, the guardrails an agent runs under are expressed as a policy – a machine-readable object that the Anshar agent execution layer attaches to the agent and that Kittu 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
- " Intelligence-data classification ceiling The highest classification, and the caveats, of intelligence the agent may read under this mission – for example Secret, CAN-EYES-ONLY. Kittu, against the datum's classification tag (§6.4).
- " Tools The named inference tools, models, and connectors the agent may invoke. NeMo Guardrails execution rails + Kittu (§5.1).
- " Sensors The sensor classes the agent may read – EO/IR, SDR/DF, MASINT, UGS – under this mission. Kittu trust-zone Z2 gate (§5.2).
- " Actuators The collection actuators the agent may command – task an RPAS to collect, steer a sensor. A combat effect is never Enlil's to command; it is handed to a human or the platform's combat function. Kittu trust-zone Z3 gate; autonomous for collection within mission policy (§5.2).
- " Trust-zone floor The minimum Kittu trust zone the agent must hold for the policy to grant anything; a hard-gate event drops it to Z0 and voids the grant. Kittu (§5.2).

Table 3B · The five guardrail dimensions of an agent policy. A zone policy sets their baseline and a continuum policy may override them; Kittu enforces the result per action at the zone.

When an agent requests an action, the policy engine resolves the request against the policy of the mission the agent is operating in, in a fixed order, and denies by default:

Classification check. Is the agent cleared, and is its policy ceiling at or above the datum's classification and caveats? If not, deny – and never substitute lower-classification context to make the action fit.

Capability check. Is the requested tool, sensor, or actuator named in the policy? If not, deny.

Trust-zone check. Does the agent currently hold the trust-zone floor the policy requires – Z2 to sense, Z3 to act? If not, deny.

Combat-effect check. Would the request produce a combat or kinetic effect? If so, it is not Enlil's to authorise – cue a cleared human or the platform's own combat function, and do not act. Collection-tasking, by contrast, an authorised Z3 agent may issue autonomously within policy.

Record. Whatever the outcome, write the request, the policy evaluated, and the determination to the chain of custody (§5.3).

Because the policy engine is deterministic and sits outside the agent, a manipulated or mistaken agent cannot widen its own guardrails; at worst it is denied and demoted. The policy framework is, in this sense, the operational expression of the Canadian Forces Security Policy least-privilege principle across the whole fabric. [CAF-3]

4.5 The security-classification framework

The policy framework bounds what an agent may do; the security-classification framework defines the currency those bounds are written in. Every datum, every zone and continuum, every agent, and every human carries a classification, and Anshar and Kittu treat that classification as a hard constraint, not a label. Enlil operates as a Top Secret system overall: intelligence flows inward to be fused and exploited. Locally-obtained data always propagates upward through the fabric; the default direction of

release is otherwise one-way data is admitted and centralised, but is not pushed back down or out except under explicit, brokered releasability. [CAF-3]

- " Canadian level NATO / Five Eyes equivalent Handling in Enlil
- " Top Secret (incl. SAR compartments) COSMIC Top Secret (CTS) The system ceiling. Compartmented (SAR) data is bound to dedicated missions; access requires the compartment, not merely the level.
- " Secret NATO Secret (NS) The working level for most fused tracks and assessments at hub and FOB-IFC.
- " Confidential-equivalent €24 NATO Confidential (NC) €24 Sensitive but below Secret; handled with need-to-know segregation by mission.
- " B €24 NATO Restricted (NR) €24 Personal and administrative data (Protected B, §6.6). iOS / iPadOS 26 devices are default-authorized for NATO Restricted material.
- " NATO Unclassified (NU) €24 NU remains NATO property and is not for public release without authority.

Table 3C · Canadian classifications mapped to NATO / Five Eyes equivalents, which share the same levels (allies occasionally add a Restricted tier at the bottom). [CAF-3][CAF-4][CAF-17]

Several markings ride on top of a classification rather than forming a level of their own. Kittu carries each as a machine-readable caveat and enforces it at the point of distribution (§5.3):

CAN-EYES-ONLY. A national rider on any level; it is not itself a classification, and restricts a datum to Canadian access regardless of level.

RELEASABLE TO / NOFORN / NOCONTRACTOR. Releasability caveats that name, or exclude, the partners and contractors a datum may reach.

ATOMAL. The NATO rider for atomic information (for example CTS-A), added when applicable.

ECL Export Control List. Export-controlled technical data under the Export and Import Permits Act, replacing the US ITAR construct. [CAF-4]

SOI Special Operational Information. Defined under the Security of Information Act, not a classification per se; it covers sources, methods, covert operations, and cryptographic means, and unauthorised release is a heightened breach of trust. Units and persons who routinely handle SOI may be designated Persons Permanently Bound to Secrecy (PPBS). [CAF-16]

Just as humans hold clearances, agents are designated with their own clearance levels so that sharing inside the automated system can be reasoned about deterministically. An agent's clearance is the ceiling its policy may grant (§4.4); an agent may never read, correlate, or carry intelligence above its clearance, and Kittu will not pass higher-classification data to a lower-cleared agent.

- " Agent clearance May access Typical placement
- " PB and B data; personal and administrative records under need-to-know.
Node-level administrative and support agents.
- " Secret Up to Secret intelligence and fused tracks; no compartmented or Top Secret material. Most node and hub inference agents.
- " TS · Top Secret Up to Top Secret; compartmented (SAR) data only within its compartment. FOB-IFC / Central-IFC correlation agents in dedicated missions.

Table 3D · Agent clearance levels (PB, Secret, TS). An agent's clearance bounds what its policy can grant; Kittu never shares above an agent's clearance.

Three rules govern the movement of classified intelligence across the fabric, and Kittu enforces them on every share:

No down-classification leak. Higher-classification data is never shared to a lower-classification agent, mission, or level; releasability is brokered, never assumed (§5.3).

Compartment containment. Sub-nodes and compartments may read from the central store, but the reverse

does not hold automatically a compartment's data does not flow outward without explicit authority.

Source deconfliction. Before fused intelligence drives any need for action, the harness deconflicts sources reconciling and de-duplicating reporting so that one source is not mistaken for independent corroboration.

Need-to-know, need-to-use, need-to-action. Enlil 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-action. Here need-to-action means the authority to commit a combat effect, and it is reserved for humans: when fused collection raises a requirement for combat action in the policy layer, Kittu deconflicts the picture and presents the situation to a human operator cleared to the appropriate level entitled to see every datum bearing on the scenario who then owns the decision and the engagement. No agent crosses from knowing or using into combat action; collection-tasking, which is not a combat action, an authorised agent may perform autonomously.

5. Trust Harness Kittu

Kittu is the Trust Assurance Framework realized at the FIB layer. It is the harness through which every agent action passes. Its remit is four-fold: govern what each agent is allowed to do through trust-operation zones; control sensor and actuator access; secure all inter-FIB communication; and capture the chain of custody that determines confidence in a threat. It does this as an additional layer on top of NVIDIA NeMo Guardrails.

5.1 Kittu on top of NeMo Guardrails

NeMo Guardrails provides programmable rails around an LLM-based agent input, dialog, retrieval, execution, and output rails that can reject or transform an agent's inputs and outputs and gate its tool calls. [11][12] Enlil uses those rails as the inner, content-level safety layer. Kittu adds a second, outer layer concerned with trust and authority, in the context of DND Instruction 2004/4 classification handling, DITSR access controls, and the Canadian Forces Security Policy (CFSP) least-privilege principle, rather than content:

"	Layer	Concern	Examples
"	NeMo Guardrails (inner)	Content and behavior of a single agent.	Jailbreak/prompt-injection detection on ingested reporting; output moderation; tool-call input/output validation. [12]
"	Kittu (outer)	Whether this agent, at its current trust, may take this action at all and proving it later.	Trust-zone check before a sensor read or actuator command; encrypted peer comms; signed chain of custody; revocation on a hard-gate event.

Table 4 · Defense in depth: NeMo Guardrails governs agent content; Kittu governs agent authority and accountability. A guard model is never the sole gate.

Classification tags flow with every datum and bound observation. Representative Canadian tags, per DND Instruction 2004/4, include Unclassified, Protected B (PII), Secret (CAF personnel records), and Top Secret (SAR CSIS); Kittu treats the tag as a hard constraint on who and what may read, share, or act on the data. [CAF-3]

The division matters because a guard model can itself be manipulated. By keeping deterministic authority checks (trust-zone membership, least privilege, signing) in Kittu, outside the agent and above the probabilistic rails, Enlil bounds the blast radius when an inner check is wrong.

5.2 Trust-operation zones for agents

Every agent in a node or hub operates inside an Kittu trust-operation zone a capability envelope earned through demonstrated, evidence-backed good behavior. The zone determines what classes of resource the agent may touch, and the highest zone is the only one from which an agent may command collection actuators. No zone permits an agent to command a combat effect: that authority is never Enlil's.

"	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 (cameras, RF, acoustic); contribute to shared state.	Sustained Z1 good standing; clean security scans.
"	Z3 · Act	Command collection actuators task or reposition a collection platform, steer a sensor to gather data. Never a combat effect.	Sustained Z2 standing; autonomous for collection-tasking within mission policy.

Table 5 · Kittu trust-operation zones. Sensor access is gated at Z2; collection actuation is reserved for Z3,

exercised autonomously within mission policy. A combat effect is never Enlil's to command.

Two safeguards from the Trust Assurance Framework carry directly into Kittu. First, 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. Second, 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.

Figure 4 · Kittu trust-operation zones gating sensor and actuator access, sitting above the five NeMo Guardrails rails.

5.3 Secured communication and chain of custody

Kittu secures every link in the fabric—node-to-node, node-to-hub, and hub-to-node—with encryption and mutual authentication, so that observations shared across the mesh cannot be read or forged in transit, and a captured unit can be cryptographically excluded. Where the fabric rides ATAK/CoT or a mesh bearer, Kittu layers its own transport security over the bearer's. [13][14]

As observations and inferences flow, Kittu records each as a signed, hash-linked entry in a chain of custody: what was observed or inferred, by which node or agent, from what inputs, and when. Because entries are hash-linked, tampering upstream invalidates everything downstream; because they are signed, every contributor is attributable. This chain is then used directly to score confidence in a threat: a track corroborated by several independent, well-attested sources across multiple nodes earns high confidence in both its existence and its severity, while a single-source, weakly-attested track is flagged as tentative. Confidence is therefore a property of the evidence, computed from the chain of custody—not a number the model emits unaccountably.

Kittu carries Canadian classification and handling caveats as machine-readable tags on every link and every chain-of-custody entry. Handling caveats such as CANADIAN EYES ONLY, NOFORN (no foreign-national access), NOCONTRACTOR, and RELEASABLE TO a named Five Eyes partner are enforced at the point of distribution, exactly as DND Instruction 2004/4 and CFSO.04.035 require. Technical data subject to export control is additionally flagged as an Export Control List (ECL) item under the Export and Import Permits Act, replacing the US ITAR construct. [CAF-3][CAF-4]

6. Intelligence Data and Ingestion

Enlil fuses eight intelligence disciplines. Each produces a different kind of data, with a different latency, reliability, and sensitivity, and critically for this architecture each is collectable in a different place. Some intelligence can be sensed forward, by a node or hub inside the operating arena; some can only come from the Central IFC, which has the reachback, the platforms, and the authorities the edge does not. This section establishes what each discipline produces, where Enlil collects it, and what sensor or technology is required the requirements on which the sensor and data-integration design rests and then defines the ingestion pipeline that brings it all into the fabric under Anshar and Kittu.

A note on terminology, consistent with the rest of this document: the disciplines below are the sources of data; the inference performed on that data is the work of Enlil's AI agents. This section is about getting the data in; Sections 7-8 are about reasoning over it.

6.1 The eight disciplines: what they produce and where they are collected

The table below summarizes each discipline. The Collection locus column is the key architectural judgment: Field means a node or local hub can collect it organically inside the operating arena; Central means it arrives from the Central IFC over reachback; Both means a tactical slice is collectable forward while the deeper or wider product comes from the centre.

"	Discipline	Representative data produced	Collection locus	Why
"	Military Int.	(CF Int / CJOC J-2) Order of battle, the common operating picture, Priority Intelligence Requirements (PIRs) per CFAO 20-204, tasking and direction, all-source assessments. [15][16][CAF-6]	Central	This is the staff product and direction itself produced by the DND/CAF all-source intelligence enterprise (CF Int / CJOC J-2) and pushed down to the edge, not sensed in the field.
"	SIGINT	COMINT (communications intercept), ELINT (radar/emitter parameters), direction-finding / geolocation, traffic analysis, electronic order of battle. [17][18][19]	Both	Tactical COMINT/ELINT and DF are collectable forward with man-portable receivers; national/strategic SIGINT and deep exploitation come from the Communications Security Establishment (CSE) at the centre.
"	HUMINT	Spot/SALUTE reports, source reporting on intent and disposition, tactical questioning, document/media exploitation, source reliability ratings. [20][21][22]	Both	Field teams generate reporting forward (entered digitally); the wider vetted source corpus, deconfliction, and analysis reside centrally.
"	IMINT	EO/IR full-motion video, still imagery, synthetic-aperture radar (SAR) imagery, change detection, geolocated targets. [23][24][25]	Both	Organic RPAS EO/IR per CFAO 11-025 (and SAR on heavier RPAS) is collected forward; satellite and high-altitude IMINT come from national or allied partners at the centre.
"	GEOINT	Foundation maps, terrain and elevation models, line-of-sight and mobility analysis, geospatial feature data. [25][26]	Central (field-updated)	Foundation geospatial data is produced centrally and pushed down; nodes refine it locally with organic survey and imagery-derived updates.
"	MASINT	Acoustic, seismic, magnetic, RF, IR/spectral signatures; CBRN/radiation detection; unattended ground-sensor (UGS) detections and classification. [27][28][29][30]	Both	UGS and tactical MASINT sensors are inherently forward-emplaced; specialized and national MASINT (e.g. strategic nuclear) come from the centre.
"	OSINT	Social-media posts, commercial satellite imagery, AIS (vessels) / ADS-B (aircraft) tracks, news and publicly available information (PAI); per CFAO 20-204, OSINT must clear operational-security and legal review before fusion with classified material. [31][32][33][CAF-6]	Central (limited field)	OSINT needs connectivity and scale best provided centrally; a node can do limited local collection (e.g. scanning local civil RF / broadcast).
"	Allied / Partner	Five Eyes (US, UK, AUS, NZ) and NATO intelligence feeds under bilateral and multilateral sharing agreements and NSA/GCHQ/CSE protocols. [16][CAF-4]	Central (field liaison)	Partner feeds are brokered centrally for releasability by the DND International Intelligence Policy Office; forward liaison elements may inject partner reporting locally where authorized.

Table 6 · The eight disciplines: data produced, where Enlil collects it, and why. Field = organic to nodes/hubs in the arena; Central = from the Central IFC over reachback. [15] [33]

The pattern that emerges is the backbone of the ingestion design. The disciplines that are sensed SIGINT, IMINT, MASINT, and the forward slice of HUMINT are collectable by FIB nodes and hubs in the arena and drive the on-board sensor requirement. The disciplines that are produced or aggregated Military Intelligence direction, GEOINT foundation data, OSINT at scale, and Allied/Partner feeds arrive primarily from the Central IFC and drive the reachback and integration requirement. Most disciplines are mixed, which is exactly why a blended, distributed fabric is the right structure.

6.2 Field-collected intelligence: the on-board sensor requirement

For the disciplines collectable forward, the following sensor and technology classes are the requirements for the sensor-integration section of the detailed design. They are sized to the FIB tiers of Section 3.2 a micro node carries one or two; a hub or heavy node hosts several or aggregates from many.

- " Discipline (field) Sensor / technology required Typical FIB host
- " SIGINT COMINT/ELINT/DF Software-defined radio (SDR) receivers; wideband scanning receivers; multi-antenna direction-finding arrays; emitter-signature libraries for classification. [18][19] Standard / heavy node; hub for multi-node DF geolocation.
- " IMINT EO/IR FMV Gimballed EO/IR camera payloads (day camera + LWIR/MWIR thermal), optical zoom, laser rangefinder/illuminator; on organic RPAS (per CFAO 11-025) or mast/vehicle mounts. [24][25] Any node; richer gimbals on heavy node / UAS.
- " IMINT SAR / wide-area Synthetic-aperture radar and ground-moving-target-indicator (GMTI) payloads; wide-area motion imagery higher-SWaP, larger RPAS only (per CFAO 11-025 clearance). [24][25] Heavy-node-class UAS; product fused at hub.
- " MASINT unattended ground sensors Air- or hand-emplaced UGS combining seismic, acoustic, magnetic, and passive-IR detectors for personnel/vehicle detection and classification (REMBASS-class). [29][30] Emplaced sensors reporting to nearest node/hub.
- " MASINT other signatures Acoustic gunshot/UAS detectors; RF spectrum sensors; CBRN/radiation detectors; multi/hyperspectral imagers for material signatures. [27][28] Node-mounted or vehicle-mounted; hub aggregation.
- " HUMINT field reporting Ruggedized digital reporting device / app for structured SALUTE and spot reports, with source-reliability tagging and media attachment at point of collection. [20][22] Operator device feeding the local node.
- " Position / navigation / timing Anti-jam/anti-spoof GNSS plus inertial navigation (GNSS-aided INS) the geo-reference every observation above depends on. [25] All tiers (foundational).

Table 7 · Field-collection sensor requirements by discipline, mapped to FIB tiers. These are the inputs to the sensor-integration design. [18][19][24][25][27][28][29][30]

The on-board sensor suite is modular by tier: a micro node may carry only EO/IR and a UGS link; a heavy node or hub adds SDR/DF, SAR feeds, MASINT, and the compute to fuse them. Every sensor is integrated through the same ingestion pipeline and governed by the same Kittu trust zones.

These classes are not hypothetical: each maps to sensors the CAF already fields or is procuring WESCAM MX-series EO/IR turrets across RCAF platforms on the IMINT rows; the ELM-2084 Multi-Mission Radar and Falcon Shield RF sensors on the radar and counter-UAS MASINT rows; sonobuoys, geobuoys, and Arctic seafloor arrays for acoustic MASINT; and the ISSP smartphone running ATAK as the HUMINT reporting device. Integration with the in-service sensor estate, not the procurement of new sensors, is the design goal. [CAF-18]

6.3 Centrally-sourced intelligence: the reachback requirement

The disciplines that arrive from the Central IFC are not sensed by Enlil hardware; they are ingested as data feeds over reachback. The requirement here is integration, not sensing: secure connectors to the central enterprise that deliver Military Intelligence direction and the COP, GEOINT foundation data, OSINT

products at scale, national SIGINT/IMINT/MASINT, and releasability-brokered Allied/Partner feeds. Because these cross the central-to-edge boundary, they are precisely the feeds whose classification and releasability Kittu must enforce on the way down, and whose volume the hub must triage before pushing only what is relevant to its nodes. Foundation GEOINT in particular is pre-positioned (cached forward) so that nodes retain terrain and mapping during a communications blackout, then refresh it on reconnection.

6.4 The ingestion pipeline

All intelligence field-sensed or centrally-sourced enters Enlil through one common ingestion pipeline that runs on every node and hub. The pipeline normalizes wildly different inputs (a thermal video frame, an intercept, a SALUTE report, a foundation map tile) into a common, tagged representation that the fabric can reason over, and it does so under Anshar's structure and Kittu's harness from the very first touch. The first stage runs once per source; the eight that follow are traversed by every datum. The stages:

Onboard (once per source). A sensor type or central feed enters service through a typed ingest adapter: the adapter normalizes the native output to the common representation, declares the source's default classification and reliability profile, and registers its identity with Kittu. Adapters are conformance-tested on the §10.4 harness before fielding, and every observation is time-stamped against fabric time (GNSS-disciplined, with specified hold-over) and geo-registered with declared uncertainty—an admission requirement, not a hope, because correlation quality downstream depends on it. [CAF-18]

Acquire. A sensor reading or a reachback feed arrives at the nearest FIB. Field sensors connect through tier-appropriate interfaces; central feeds arrive over the secured reachback link.

Authenticate, inspect, and admit (Kittu). Kittu verifies the source sensor identity, peer signature, or central-feed credential and the payload: schema and format validation and content sanitization run on every input, so an authenticated source cannot deliver a malformed or hostile payload unchecked. Data that fails inspection is quarantined to trust-operation zone Z0 with its custody record preserved (§5.2). Where a flow crosses an accredited security-domain boundary—a coalition ingest, a Secret-to-Protected-B exchange—Enlil enforces policy up to the boundary and brokers the transfer through an accredited cross-domain solution (CDS), with Kittu's custody records spanning both sides. This is the first link in the chain of custody: every datum is bound to who or what produced it before anything else happens. [3]

Normalize and tag. The datum is converted to the common representation—the Enlil Observation, whose concept-level model Table 7A gives below—and tagged with discipline, time, location, originating sensor/node, reliability, and classification/releasability caveat. Under DND Instruction 2004/4 the classification tag carries the level (Unclassified, Protected A/B/C, Secret, Top Secret), the handling caveat (CAN-EYES-ONLY, NOCONTRACTOR, and the like), the mark-for-release authority (for example CJOC J-2 or the DND/CAF classification authority), and a PIPEDA/ATIA compliance flag where the datum is personal or access-restricted. Classification tagging here is what lets Kittu enforce sharing rules downstream. [CAF-3][CAF-11]

Classify the zone (Anshar). Anshar assigns the datum to the relevant zone(s)—the sector it was collected in, the track it pertains to, the processing state it has reached—so its movement through the fabric is observable and the right nodes and hubs can be notified. [2]

Enrich and correlate. The datum is linked to existing tracks and entities and, where the tier allows, run through on-board models for detection/classification—the point at which raw data becomes a candidate observation for inference (Sections 7-8). Observations of the same event arriving by different paths are recognized—content hash plus a spatiotemporal window—and merged with both custody chains preserved.

Record (chain of custody). Kittu writes a signed, hash-linked provenance entry for the datum and every transformation applied to it, so the evidence behind any later assessment is reconstructable and tamper-evident. [3]

Distribute under policy. The tagged observation is shared with neighbors and the hub over the mesh, and

(from the hub) to central fusion each share brokered by Kittu against trust-zone and releasability policy. The broker's vocabulary is the audit vocabulary of §12.6 permit, restrict, downgrade, or segregate and every disposition is recorded with the policy version it was resolved against (policy_id, policy_version), so a compromised node leaks the least possible and every decision is answerable. [3]

Transform and release. Where a share is authorized but the recipient's domain, caveat set, or message format differs from the source's, the pipeline produces a releasable derivative rather than releasing the original: metadata stripping, imagery masking or resolution reduction, text extraction into the recipient's format CoT for the ATAK COP, J-series via a Link 16 gateway, AdatP-3/APP-11 for NATO text traffic (§10.3) and classification re-marking. The derivative is recorded as a child object in the custody chain: parent unreleased, child released, lineage intact. Transform rules are policy, not code they live in the continuum policy, so a new partner profile is a policy change (§4.4), not a software release. This is downgrade as the challenge means it: not block-or-release, but the production of exactly what the recipient may have. [CAF-18]

Figure 5 · The common ingestion pipeline. Field-sensed and centrally-sourced intelligence enter the same nine-stage flow; Kittu authenticates, records custody, and brokers distribution and release, while Anshar zones and tracks every datum. Inference (Sections 7-8) consumes the tagged observations this pipeline produces.

The pipeline's product deserves a name and a shape. Every datum that clears stage 4 exists thereafter as a Enlil Observation the common, tagged object the rest of the fabric reasons over. Its concept-level model is fixed here; the formal, machine-readable schema is detailed-design work (§13).

- " Block Fields (concept level) Purpose
- " Identity UUID; content hash Deduplication and integrity anchoring across paths and partitions.
- " Origin Sensor/feed ID; node ID; collection time (fabric time); location with declared uncertainty Who, where, when the provenance base the challenge names (outcome 4.3.7).
- " Payload descriptor Discipline; modality; format; size What the datum is, without touching what it says.
- " Classification Level; handling caveats; releasability; PIPEDA/ATIA flag The block every sharing decision reads (§4.5, §6.6).
- " Reliability Source rating; sensor confidence The evidence weight fusion assigns (§5.3).
- " Custody head Pointer to the signed, hash-linked chain Tamper-evident history from acquisition (§5.3).
- " Bindings Zone(s); continuum(s) Where it lives and which mission may see it (§4.1 §4.2).

Table 7A · The Enlil Observation concept-level object model. The formal schema is follow-on work; the fields and their purposes are fixed here.

Above the Observation sits a short derivation ladder: correlated observations of one entity form a track; inference over tracks produces an assessment. Each derived object carries pointers to its parents, so data lineage (outcome 4.3.9) is structural reconstructable from the objects themselves, not reconstructed from logs after the fact. Residency follows the ladder: nodes hold observations and local tracks; hubs hold cluster tracks; the FOB-IFC holds theatre tracks and assessments. Raw payloads stay as far forward as possible; derived objects are what move which is also the fabric's first bandwidth answer.

Two properties of this pipeline matter for the rest of the architecture. First, ingestion is where provenance begins: because Kittu authenticates and records custody at acquisition, the confidence Enlil later places in a threat is grounded in an unbroken chain that starts at the sensor or the feed, not at the inference. Second, ingestion is uniform across the fabric: the same pipeline runs on a micro node and on a Thor hub, so a datum collected by the smallest node is tagged, zoned, and custody-tracked identically to one from the Central IFC which is what lets the distributed inference of Section 8 treat the whole fabric's data as one coherent, governed body of evidence.

6.5 Pipeline behaviour under stress

The pipeline of §6.4 describes what happens to a datum; an edge system is judged by what happens when there are too many of them distance, mobility, contested RF, and the bandwidth seam between forward collectors and rear analysis are the operating condition, not the exception. Enlil therefore defines its overload behaviour as architecture, stated here so it can be tested, rather than discovered in the field. [CAF-18]

Priority and drop policy. Everything in flight carries one of three priority classes: threat-bearing observations and alerts first; routine tracks and telemetry second; bulk raw payloads last. Every stage runs bounded queues with a declared per-class drop policy. Perishable streams full-motion video frames drop newest-wins under pressure; HUMINT reports and custody records are never dropped: if they cannot move, they persist and wait.

The degradation ladder. When compute or power saturates, the fabric sheds work in a fixed order: enrichment and on-board model inference first, raw-payload retention second, distribution frequency third. What is never shed: authentication, tagging, and custody stages 3, 4, and 7 are load-shed-exempt by design. Under any load, nothing moves untagged, unauthenticated, or unrecorded. That sentence is itself a compliance claim (outcome 4.4.3): SWaP and compute limits degrade the picture's richness, never its governance.

Partition. Partition is a normal state (§2, principle 2). During partition, tagged observations queue durably at the node with custody intact; on reconnection they replay upward in priority order, and divergent post-partition pictures reconcile under the semantics of §13. Store-and-forward is a property of the pipeline, not an application bolted beside it.

The latency budget. The working budget on the reference tiers, stated as targets to be validated in detailed design (§13): acquisition to tagged-and-recorded in tens of milliseconds; acquisition to distributed (stage 8) in low hundreds of milliseconds for a priority observation; enrichment asynchronous and never blocking the compliance path. These are estimates, flagged as such stated as budgets so they can be tested against the challenge's real-time requirement (outcome 4.4.1) rather than left as an adjective.

6.6 Protected B compliance in the ingestion pipeline

Protected B (PB) defined in DND Instruction 2004/4 and engaged by the Access to Information Act and the Privacy Act is information whose unauthorized disclosure could cause serious injury to an individual or organisation. It is central to the DND/CAF challenge, and Enlil enforces PB handling through the same Anshar-and-Kittu machinery that governs everything else: [CAF-3][CAF-11]

1. Classification tagging at ingestion. Every datum from personal records (personnel files, medical records, performance reviews) is tagged Protected B (ATIA-restricted) or Protected B (Privacy Act) at the Normalize-and-tag stage of Section 6.4.
2. Segregation by zone. PB data is bound to zones and continuums where access is restricted to cleared personnel with an explicit need-to-know.
3. Trust-zone enforcement. Access to PB data requires at minimum the Z2 (Sense) trust zone; any collection actuator command taken in a PB context requires Z3 and an audit entry, and may be issued autonomously within mission policy; a combat effect is never Enlil's to command (Section 5.2).
4. Export restriction. PB data cannot be exported or shared outside the Canadian system without explicit releasability authority; Kittu enforces this at the Distribute under policy stage.
5. Audit trail for disclosure. The chain of custody records every access, transformation, and export of PB data, supporting Access-to-Information requests and Privacy Impact Assessments (PIAs).

7. Inference Layer

Sections 4-6 establish what runs on the fabric, how it is governed, and what data flows in. This section is about the inference itself – the models that turn raw intelligence data into observations and that drive the agents behind Anshar and Kittu. The defining commitment of Enlil’s inference layer is small, specialized, and fine-tuned: rather than one general model, Enlil fields a family of models from tiny to medium, each fine-tuned for a specific job and sized to the FIB tier it runs on. The largest models – Large and Frontier – live at the Central IFC and are delivered in Phase 3 (Sections 1.3 and 7.5).

7.1 The model-size ladder and Enlil’s scope

The industry has converged on a rough ladder of model sizes. Enlil’s deployed scope is the bottom five rungs – the models that run on the Jetson node, Jetson hub, and Grace-Blackwell FOB-IFC tiers of Sections 3.2 and 3.3, up to ~80B at the FOB-IFC – because by 2026 sub-10B models routinely match 2024-era frontier models on the narrow tasks they are tuned for, at a fraction of the compute – supporting CAF operational timelines and the Digital Campaign Plan modernization schedule. [34][35][36][CAF-1]

Tier	Parameter range	In Enlil scope?	Where it runs
Tiny models	10M – 500M	Yes	Micro / standard nodes; the hot perception path (detectors, signal/acoustic classifiers).
Micro language models	~100M – 2B	Yes	Standard / heavy nodes; compact text and small vision-language tasks.
Small language models (SLMs)	1B – 10B	Yes	Heavy nodes and hubs; report understanding, tool-calling agents, edge VLMs.
Medium language models	10B – 30B	Yes	hub only – Thor T5000 hub; cluster-level fusion reasoning and the heavier guard models.
Large models	30B – 80B	Yes	FOB-IFC – Forward IFC (DGX Station GB300); cross-hub correlation (Section 3.3).
Frontier & Large	80B+ / 80B+ / MoE	Phase 3	Central IFC – Central IFC data centre; see Sections 1.3 and 7.5.

Table 8 · The model-size ladder. Tiny through Medium run on the node and hub tiers and Large (to ~80B) on the FOB-IFC; Large 80B+ and Frontier are Central-IFC (Phase 3). [34][35][36][98]

Figure 6 · The model-size ladder. Enlil’s scope is Tiny through Medium, fine-tuned and deployed across the FIB tiers; the Large and Frontier rungs sit at the Central IFC and are delivered in Phase 3 (Section 1.3).

7.2 Fine-tuning, not training from scratch

Enlil does not pre-train foundation models. It fine-tunes existing open-weight bases for its specific deployment scenarios – a small base plus parameter-efficient fine-tuning (LoRA/QLoRA) can outperform a general large model on a vertical task with only a few thousand labeled examples and hours of single-GPU training. [37][39] The discipline is consistent across every model Enlil fields:

Start from a strong open base. Qwen3-class models are among the strongest fine-tuning bases at small sizes; Phi-4 family for reasoning-heavy tasks; Gemma 3/4 for multimodal; SmoLLM/SmoLVLM at the smallest viable sizes. [34][35][36][38]

Adapt with LoRA / QLoRA. Parameter-efficient fine-tuning keeps the adaptation small, auditable, and cheap to re-issue as the mission or the threat changes. [37][39]

Quantize to the tier. AWQ for GPU/NPU deployment (best accuracy retention), GGUF for CPU, INT4 as the

mobile/edge starting point chosen per FIB tier's memory and power envelope. [36]

Distill where the hot path demands it. For the lowest-latency perception tasks, distill a tiny student from a larger teacher so the node runs a model it can actually sustain. [34]

Every model in Enlil is a fine-tuned, quantized specialist anchored to a FIB tier. This is what makes on-board inference feasible within a node's power budget and what keeps each model small enough to be governed, signed, and updated under Kittu.

7.3 Models for raw intelligence data types

Each raw data type from Section 6 has a model class suited to it. The table maps the field-collectable disciplines to representative model families and the tier they target. These are candidate bases for fine-tuning, to be confirmed in detailed design against the latest open releases.

"	Raw data type	Model class & representative bases	Tier
"	IMINT EO/IR FMV, stills	Tiny object detectors (YOLO-nano / -tiny family, e.g. YOLOv11n/YOLO26n, LMW-YOLO ~2.6M for small-object remote sensing) for the hot path; a small edge VLM (SmoI/VLM 256M 2B, PaliGemma 3B, Phi-4-Multimodal) for scene reasoning and captioning. [40][41][44][45][46]	Tiny + Small
"	SIGINT COMINT/ELINT/IQ	Tiny/Micro task-specific CNNs on raw IQ for automatic modulation classification (ResNet-18/MobileNetv2, CLDNN CNN-LSTM, complex-valued CVNN/CDSCNN); optionally a small VLM fine-tuned on RF spectrograms for explainable recognition. [47][48][49][50]	Tiny + Micro
"	MASINT acoustic / seismic / UGS	Tiny CNNs on raw waveforms / spectrograms for signature classification (TinyChirp-class acoustic CNNs; tiny matched-filter CNNs for 1-D sensor streams). [42][43]	Tiny
"	HUMINT / OSINT text	Micro / Small language models for parsing SALUTE reports, named-entity extraction, and summarization (Phi-4-mini 3.8B, Qwen3-4B, Gemma 3 4B). [34][35][38]	Micro + Small
"	GEOINT imagery / geospatial	Small vision and geospatial models for feature extraction and change detection, paired with the IMINT detectors above. [40][45]	Tiny + Small
"	Cross-source fusion (hub)	Medium language model for correlating multi-source observations into assessments (Gemma-3-27B-class, Qwen 14 30B). Hub-only. [34][35]	Medium

Table 9 · Representative model classes per raw intelligence data type, with target FIB tier. Candidate bases for fine-tuning; confirm in detailed design. [34][35][38][40] [50]

7.4 Models for the foundational agents (Anshar and Kittu)

Enlil's foundational components are themselves agentic, and they need models too. The Anshar zone agents (classifying a datum into a zone, deciding which units a transition is relevant to) and the Kittu harness agents (trust-zone classification, tool/actuator-call validation, content guarding over NeMo Guardrails) are best served by small fine-tuned function-calling and classification models. Qwen3-class bases fine-tuned for tool use are well suited to the Anshar and Kittu tool-calling agents; for Kittu's content-guard role specifically, dedicated guard models in the 1B-27B range (Llama-Guard / ShieldGemma-class) provide the moderation layer that sits beneath the trust-zone checks. [38][51][52] All are fine-tuned with LoRA/QLoRA and quantized to their host tier, exactly as in Section 7.2, so that the agents which enforce Enlil's governance are themselves small, inspectable, and signed.

The same fine-tuning discipline serves both halves of the system: the models that read the world (raw-data inference) and the models that govern the system (Anshar and Kittu agents) are all small, specialized, and quantized to the tier they run on.

7.5 Phase 3: Large and Frontier models at the Central IFC

Large language models above ~80B and Frontier models (200B+ or Mixture-of-Experts) are delivered in Phase 3 (Section 1.3), at the Central IFC. They cannot run within the power, memory, and thermal envelope

of a forward FIB; Large models up to ~80B run forward on the Grace-Blackwell FOB-IFC (Section 3.3), but larger Large models and all Frontier models cannot, and Enlil's design principle is to infer forward on models the edge can actually sustain. These models reside at the Central IFC, where the deep all-source reasoning, the fine-tuning of the models the edge deploys, and the corporate-memory functions of Section 8 take place. Enlil consumes the products of central Large/Frontier inference—direction, enrichment, updated model weights—over reachback, but it does not host or run them. Any requirement that depends on a Large or Frontier model is, by definition, a Phase 3 / Central-IFC capability, sequenced beyond this draft's boundary.

8. Distributed Inference

Enlil is distributed inference: there is no single brain. Each node infers on the information it can access; that view is shared with neighboring nodes and with its hub and is weighed in their inference; hubs pass both the raw observations they collect and their own assessments up to Central IFCs, take direction back, and relay it to their nodes. The picture emerges from the fabric rather than being assembled in one place.

8.1 The inference flow

Node-local inference. A node runs the models its tier can support on its own sensors – detection, classification, tracking – and forms a local assessment with a confidence drawn from its chain of custody.

Neighbor sharing (mesh). The node shares its observations and assessment with neighbors over the mesh. Each neighbor folds the incoming view into its own inference, so a contact seen weakly by one node and strongly by another is reconciled into a better joint estimate.

Hub fusion (higher inference). The local Thor hub pools the cluster's observations and runs heavier models to produce a higher-confidence assessment than any single node could – resolving conflicts, corroborating across sources, and updating zone custody.

Reachback and direction. The hub passes its observations and assessments to the FOB-IFC, which correlates across the Forward Operating Base's dozens of hubs with models up to ~80B and, in turn, reaches back to the Central IFC (CJOC J-2 or the theatre fusion centre) for the deepest all-source and historical context and for direction. Direction – re-tasking, priorities, releasability – flows back down through the FOB-IFC and hub to the nodes.

Reconciliation after partition. If a node or hub was cut off, it operated autonomously on what it had; when the link returns, its chain of custody reconciles with the wider picture, and any divergence is surfaced rather than silently overwritten.

Figure 7 · Distributed inference. Nodes infer locally and share peer-to-peer; hubs fuse the cluster and exchange observations and direction with Central IFCs.

8.2 Why distribute

Distributing inference buys three things a centralized fusion model cannot. It survives partition, because every node keeps inferring when isolated; it is fast, because the decision happens where the data is, with no round trip; and it degrades gracefully, because losing a node loses that node's view, not the whole picture. The cost – reconciling many partial views into a coherent whole – is exactly what Anshar's shared zoning and Kittu's chain of custody are there to pay.

9. Human-Machine Interface

For a crewed unit, a FIB is, above all, a virtual combat assistant to its operator. It informs, it observes, it records, and it distributes local observations among neighbors seamlessly – all under Anshar’s zoning and Kittu’s harness. The interface is therefore central, not peripheral. Enlil targets three complementary modalities. The specific capabilities each Jetson tier can support – display pipelines, audio I/O, concurrent model headroom for on-device speech – are pending detailed research and will shape what each tier offers.

- " **Modality** | What it provides | Notes for detailed design
- " **AI HUD** | A heads-up display overlaying the fused picture, tracks, zone alerts, and assistant prompts on the operator’s view. Depends on display interfaces and concurrent vision-model headroom per tier; heavy nodes/hubs drive richer overlays.
- " **Audio command & feedback** | Spoken queries to the assistant and spoken alerts/answers back – hands-free, eyes-up operation. On-device speech recognition/synthesis sizing per Jetson tier is an open research item; smaller nodes may rely on the hub.
- " **ATAK integration** | Sharing tracks, zones, and alerts into the operator’s existing ATAK common operating picture via Cursor-on-Target. CoT is an open XML standard with a plugin ecosystem and mesh bearers; the natural integration surface. [13][14]

Table 9A · Enlil HMI modalities. Final capability per modality is pending Jetson-tier research.

These modalities serve the crewed case – a dismounted soldier or a vehicle crew. For an autonomous combat vehicle (UAS, UGV, USV), the FIB is the platform’s perception-and-decision core and the interface is machine-to-machine: the same fused picture and tasking flow over the mesh as structured data rather than to a HUD; collection-tasking the platform may act on autonomously, while any combat effect is referred to a remote operator or supervising crew, who owns that decision (Section 5.2). The virtual-combat-assistant HMI and the autonomous-platform interface are two faces of the same FIB.

ATAK deserves emphasis as the integration anchor. It is a mature, widely fielded geospatial situational-awareness platform with an extensible plugin architecture and a common messaging schema, Cursor-on-Target (CoT), for exchanging position, status, and sensor data between systems – and it runs over mesh and off-grid bearers. [13][14] By emitting Enlil tracks, zone transitions, and threat alerts as CoT, a FIB can drop its fused picture into the operator’s existing ATAK view without bespoke integration, and consume CoT from other TAK-connected units in return. Where ATAK is not the chosen client, an equivalent CoT-speaking surface serves the same role.

The operator should never feel they are running a fusion system. They should feel they have an assistant that sees more, remembers everything, and quietly tells them – and their neighbors – what matters, when it matters.

Figure 8 · The FIB as virtual combat assistant: AI HUD, audio command and feedback, and ATAK/CoT integration, all driven by the same fused picture.

9.1 Worked vignette

A standard node on Orin NX, deployed with a Canadian forward platoon, detects an unfamiliar contact at the edge of its sector. Its local model classifies it tentatively; the chain of custody marks the assessment single-source, low confidence. Over the mesh, two neighbouring nodes – section and squadron elements – confirm the contact from different angles; their observations fold into a joint estimate and confidence rises. As the contact crosses into the next sector, Anshar records the zone exit/entry and notifies the inheriting node and the Thor hub. The hub fuses the cluster’s views, raises confidence to high, and relays a corroborated assessment to CJOC J-2. The operator’s HUD shows the track with severity; ATAK displays it in the common operating picture for the wider unit. An agent proposes cueing a local RPAS to maintain custody – an actuator action – so Kittu requires Z3 and explicit operator approval before the command is

issued. Throughout, every step is signed into the chain of custody per DND Instruction 2004/4.

9.2 Explainable enforcement and controlled override

Every enforcement action the pipeline takes permit, restrict, downgrade, segregate renders on the operator surface as a one-line, human-readable rationale: the rule that fired, the policy version, what was withheld or transformed, and why. The line is a rendering of the same §12.6 audit entry, not a separate path what the operator reads is what the accreditor will read.

A cleared operator may override a restrict or downgrade disposition. Kittu checks identity, clearance, and role before offering the control; the override is a first-class audit event capturing who, what, the stated justification, and the policy it overrode; and it is scoped this datum, this continuum, or time-boxed never standing. One disposition is not overridable: the check on any combat effect (§2, principle 8) offers no override, because that authority is never Enlil s to delegate.

Mechanically, an override is a continuum-policy exception bound through the same two-framework resolution of §4.4 no new subsystem, and therefore no new accreditation surface. Together, the rationale line and the override loop are the architecture s answer to outcome 4.4.4: the operator can always see why, and a cleared human can always take responsibility on the record.

9.3 Worked scenario: coalition release under caveat

Adapted from the SME advisory review of this architecture [CAF-18]. A Canadian battle group is deployed within a NATO multinational brigade. Its organic RPAS carries an EO/IR turret; the resulting track runs at NATO Secret on the battle group s own network. A partner infantry company two kilometres away on previous-generation radios, on a lower security domain is directly threatened by what the RPAS is watching: an armoured column approaching its position.

Released raw, the feed violates information-security rules and exposes the platform s capabilities. Withheld, it costs the partner the warning. Manual review cannot arbitrate at the speed of the approach. This is precisely the seam the challenge describes: compliance enforcement inside the fusion path, at machine speed.

The walk through the pipeline (§6.4). The track is ingested and tagged at stages 2 5: NATO Secret, national caveats, the platform signature carried in the payload descriptor, an acoustic-MASINT corroboration merged at stage 6. The coalition continuum s policy names the partner s caveat profile, so at stage 8 Kittu resolves the share as downgrade not permit, not restrict. Stage 9 produces the releasable derivative: a sanitized structured text alert position, count, heading in the partner s message format; the video stream masked; the platform s identifying signature stripped; the classification re-marked to the releasable level. The accredited CDS brokers the domain crossing. The partner receives a warning it can decode, in time to act.

The §12.6 record carries the whole event: parent object unreleased, child derivative released, the rule that fired, policy_id and policy_version, disposition downgrade, and the custody linkage between parent and child. The battle group commander s HMI shows the one-line rationale (§9.2); the accreditor later reads the same entry.

Outcomes evidenced end to end: 4.3.3 (two modalities fused), 4.3.4 (two security domains crossed under enforcement), 4.3.6 (no human approval required the condition was predefined in continuum policy), 4.3.8 (the downgrade disposition and its rule in the log), and 4.4.4 (the operator-readable rationale).

9.4 Worked scenario: the Arctic seam

A second scenario exercises the strategic seam, again adapted from the advisory [CAF-18]. In the Arctic air-defence identification zone, a low-observable collector holds a passive track on intruding aircraft position, speed, electronic signature but its emissions posture forbids broadcasting on the common

tactical data link: transmitting would give away exactly what its passivity protects.

The collector's product arrives at a northern hub and the FOB-IFC over its own directional link, is fused with NORAD-context feeds arriving from the Central IFC, and becomes a validated track. Stage 9 emits the releasable form a J-series track message through the certified Link 16 gateway at the TF HQ integration point (§10.3) with the collector's signature and sensing parameters masked. Consumers on the common link receive the track without receiving the source.

Protected B handling rides the same pass: any incidental Canadian-person data in the fused picture carries the PIPEDA/ATIA flag from stage 4 and is segregated from the coalition release automatically (§6.6). The full audit lineage from passive collection through fusion to release exports for air-defence sector review (§12.6).

Outcomes evidenced: 4.3.2 (the masking and release rules are machine-readable policy), 4.3.5 (Protected B enforced inside a Secret-and-above flow), 4.3.7 (provenance across three collection sources), 4.3.9 (exportable lineage), and 4.4.1 (the §6.5 latency budget is what makes the track tactically relevant).

10. Open Platform Integration

Enlil is built as an open platform, not a sealed product. Its value to a coalition comes from how readily new capability can be plugged into it – a vendor's sensor, a counter-drone effector, an uncrewed platform, a specialised analytic model, or a third-party agent – and made to contribute to the common picture and the field fight. The threat evolves faster than any one vendor, and faster than any one procurement cycle, so the platform is designed to absorb outside capability quickly while never relaxing the trust boundary around it. This section sets out the two ways external systems integrate – through published APIs and through agent wrappers – and how every one of them is brought under Enlil's zones (Anshar, Section 4), its trust harness (Kittu, Section 5), and its secured communications (Section 5.3) before it can influence anything an operator sees or does.

A note on terms, consistent with the rest of this document: intelligence means the INT disciplines and their data (Section 6); inference means the analytical work the models perform. A third-party system is a sensor, effector, or platform; a third-party agent is an external AI agent. Offensive and defensive effectors – observation and strike drones, counter-UAS, jammers – may be integrated as capabilities that Enlil feeds with intelligence and, at most, tasks for collection; Enlil never authorises their engagement. A command that would produce a combat effect is cued to a cleared human or the platform's own combat function, which owns that decision (Section 5.2).

10.1 Why an open platform

The case for openness is both operational and doctrinal. Operationally, the value of integration is speed: the ability to bring in the capability you need at the moment you need it – rather than waiting months for a bespoke build, by which time the threat has changed – is itself an advantage. Allied trials of open sensor standards have shown heterogeneous systems connecting in effectively zero integration time once they conform to a common interface. [91] Doctrinally, allied and NATO defence acquisition has converged on the Modular Open Systems Approach (MOSA): loosely-coupled, severable modules behind published, consensus-based key interfaces, with conformance that can be independently verified, adopted to widen competition, avoid vendor lock-in, and accelerate capability insertion. [89] Enlil takes that posture by design.

Concretely, an open platform buys Enlil four things:

Coalition reach. Any allied vendor that can meet a published interface can contribute, without bespoke, per-vendor integration work.

Competition without lock-in. Modules from different suppliers compete on capability and can be swapped or upgraded independently, because the interface – not the vendor – is the fixed point.

Rapid capability insertion. New sensors, effectors, and models click in and out as the mission and the threat change, rather than being welded into a monolith.

Sovereignty of provenance. Enlil owns the trust boundary even around parts it did not build; an external capability is hosted on Enlil's terms, not the other way round.

The governing idea

Open at the interface, governed at the boundary: any conforming vendor can connect, but everything that connects enters through one place that Anshar zones, Kittu governs, and secured comms carries.

10.2 What integrates: the capability classes

Enlil is built to host four broad classes of vendor capability. Each enters at a defined surface and is admitted with a defined trust posture; the more consequential the capability, the tighter the gate.

- " Capability class Representative technologies How it plugs in Zone / trust posture
- " ISR sensing & monitoring (intelligence-gathering) EO/IR, SAR, RF / SIGINT, radar, acoustic, lidar, unattended ground sensors As sensor modules emitting observations and declarations, not raw streams SAPIENT-style, carried as Cursor-on-Target events Sensing zone; read-mostly, low actuator privilege
- " Defensive effectors Counter-UAS detect-track-defeat, RF jammers, directed energy, defence drones As effectors that take tasking and return status, behind an actuator interface High-trust zone (Kittu Z3) for collection-tasking; any engagement is decided by a human or the platform's own combat function
- " Offensive effectors Observation and strike UAS, loitering munitions Intelligence and collection-tasking, plus custody hand-off; the engagement decision is never Enlil's Highest-trust for collection-tasking; engagement authorised and executed by a human or the platform's combat function
- " Uncrewed platforms UAS, UGV, USV carriers and their ground-control stations Command-and-control and telemetry via STANAG 4586 (DLI / CCI) or MAVLink Platform zone; control authority graduated by Level of Interoperability
- " Analytic & model services Third-party detection, classification, and fusion models Packaged as signed inference microservices (NVIDIA NIM / OpenAPI), discovered through a gateway Inference zone; outputs treated as observations pending corroboration
- " Third-party agents Vendor autonomy, planners, and decision aids Wrapped into Kittu through the agent toolkit (MCP / A2A) Run inside the trust harness; full Observer-Judge-Controller mediation

Table 19 - Vendor capability classes Enlil is built to host, with the integration surface and trust posture for each. Enlil provides effectors with intelligence and, at most, collection-tasking; it never authorises their engagement that decision stays with a human or the platform's own combat function.

The effector rows deserve emphasis, because they are the most consequential. An integrated drone or counter-UAS system gets the same rule as a Enlil-native capability: Enlil may task it for collection, at the high-trust zone and within mission policy, but it never issues the command that produces a combat effect. Weapons release is not Enlil's to grant not even with a human present. A command that would engage is cued to a cleared human operator or the platform's own combat function, which authorises and executes it. Enlil hosts offensive and defensive systems as intelligence consumers and collection assets; it does not command their fires.

10.3 Conforming to open standards

Enlil does not invent a private bus and force vendors onto it. It speaks the interoperability standards the coalition already uses, and wraps each so it conforms to Enlil's zones and trust harness. Table 20 maps the principal standards to where they land in the architecture.

- " Standard What it governs Where it lands in Enlil
- " MOSA (NDAA §805 / allied open architecture) [89] Modular design, published key interfaces, certified conformance The platform's overall design contract every integration surface is a published, verifiable key interface
- " SAPIENT (NATO STANREC / AEDP-4869; BSI Flex 335) [91] Autonomous sensor modules feeding a fusion / decision module; concise declarations, not raw streams; a conformance test harness Maps directly onto FIB node (edge sensor autonomy) to hub (fusion); the observations-not-raw-feeds model is already Enlil's
- " STANAG 4586 [90] UAS command and control: Data Link Interface and Command-and-Control Interface; Levels of Interoperability (LOI 1-5) The platform-control surface; LOI tiers map onto Kittu trust zones, from monitor-only to full C2

- " Cursor-on-Target / TAK (MIL-STD-2525) [92] A shared event format so any sensor can cue any display or effector; the de-facto tactical common operating picture Enlil's shared-picture wire format already the basis of the HMI's ATAK/WinTAK link (Section 9)
- " MAVLink / Open Drone ID (Remote ID) [93] Drone telemetry and control; broadcast drone identification Drone integration, plus Remote-ID ingest for counter-UAS friend/foe discrimination
- " DDS (Data Distribution Service) [94] Data-centric publish/subscribe transport for real-time systems An option for the real-time data fabric between modules, alongside the mesh
- " NIM / OpenAPI; MCP / A2A [95][96] Model-serving endpoints; agent tool and agent-to-agent interfaces The two software integration surfaces detailed in Section 10.4
- " Link 16 / JREAP-C · AdatP-3 / APP-11 [CAF-18] Joint and coalition messaging: the J-series tactical data-link picture above battalion level, and NATO structured text formats Egress and ingress at the hub and FOB-IFC tiers through existing certified gateways at the TF HQ integration point Enlil produces and consumes messages (stage 9); it does not host a terminal

Table 20 - Established open standards Enlil conforms to, and where each lands in the architecture. The platform meets vendors on standards they already use.

In the Canadian context these map onto domestic equivalents: NATO STANAGs are mirrored by Canadian Standards and Recommended Practices (CSRPs) published by DND; STANAG 4586 RPAS control is governed by CFAO 11-025 (Remotely Piloted Aircraft Systems Operations); and MOSA aligns with the DND Modular Open Systems Approach initiative in Canadian defence procurement. Canadian vendors on the DND Approved Vendor List (AVL) integrate sensors, effectors, and platforms against these published interfaces, while alliance partners (Five Eyes and NATO) contribute validated intelligence feeds under sharing agreements brokered by the DND International Intelligence Policy Office. [CAF-7]

The SAPIENT alignment is the deepest, and worth drawing out. SAPIENT's architecture is a network of autonomous sensor modules that push concise threat declarations to a high-level decision-making module, defined by a freely-available interface control document and checked by a conformance test harness. In allied trials it linked many sensor modules from different vendors to several decision modules, some achieving genuine plug-and-play integration. [91] That is, in miniature, Enlil's own design: edge nodes that infer locally and share observations rather than raw data, fused at a hub. A SAPIENT-compliant sensor is therefore close to a drop-in Enlil node the gap is the trust wrapper, not the data model.

Cursor-on-Target plays the complementary role: originated by a US Air Force Research Laboratory and MITRE effort and deliberately licence-free and extensible, it became the tactical edge's lingua franca on the principle that any sensor should be able to cue any shooter or display through one common event structure. [92] Because Enlil both consumes and emits CoT, a third-party sensor's tracks appear on the very map the operator already uses.

10.4 Two integration surfaces: published APIs and agent wrappers

Everything external enters through one of two surfaces. The first carries systems sensors, effectors, platforms, and models. The second carries agents. They differ in mechanism but share one rule: the vendor brings the capability; Enlil keeps the trust boundary.

Published APIs for systems. Enlil publishes versioned interface contracts: OpenAPI / REST and gRPC for request-response and streaming, plus the domain protocols above (CoT, the SAPIENT ICD, STANAG 4586, MAVLink) for systems that already speak them. Each contract is a stable, documented key interface in the MOSA sense, with a schema a vendor can build and test against. For vendor-supplied models, the package is a signed inference microservice: an NVIDIA NIM or any OpenAPI-compatible container that bundles the model and exposes a standard endpoint, discovered through a gateway or proxy. NIM containers are cryptographically signed, ship with a vulnerability (VEX) record, and run in an air-gapped enclave, which is exactly what a fielded, classified deployment needs. [95] API governance is uniform across every contract: versioned schemas with validation, mutual-TLS identity, least-privilege scopes, rate limiting, and signed artifacts. Nothing reaches the mesh except through a contract Enlil has published and a credential Enlil has issued.

Agent wrappers for agents. Enlil's agent runtime is built on the NVIDIA NeMo Agent Toolkit, which is deliberately framework-agnostic: a vendor agent written in LangChain, LlamaIndex, CrewAI, Semantic Kernel, Google ADK, or plain Python is wrapped as a first-class component without a rewrite. [96] Interoperability rides on two open protocols the Model Context Protocol (MCP) for exposing and consuming tools, and the Agent2Agent (A2A) protocol for authenticated agent-to-agent collaboration. [96] The wrapper is the control point. Every wrapped agent runs inside Kittu: it never talks to the mesh directly, it talks to its wrapper, which sits behind NeMo Guardrails and under Kittu's Observer to Judge to Controller loop, with its tools, inputs, and outputs mediated and its entire trace signed. A third-party agent that misbehaves is contained at its wrapper, in an error-zone, before it can affect anything downstream.

One rule, two surfaces

A third party brings the capability; Enlil keeps the trust boundary. Meeting the interface is theirs to do. The governance zone, harness, secured comms is non-negotiable and identical whether the contributor is a sensor, a drone, a model, or an agent.

10.5 Governing the edge: zones, trust harness, secured comms

Integration is not admission. A published interface lets a vendor connect; it does not make them trusted. Trust is earned along a governed onboarding path, and it is graduated and revocable. Three layers govern every external capability, and an onboarding gate ties them together.

Zones (Anshar). Every external system or agent is bound to a zone with an explicit policy and least-privilege observability. A new vendor capability starts in a low-trust zone and is promoted only as it proves itself, on the graduated-trust model of Section 4; error-zones contain misbehaviour without taking down the picture.

Trust harness (Kittu). The wrapper runs the third party under Observer-Judge-Controller, with guard models screening its inputs and outputs and every action signed and attributable the same chain of custody (Section 5.3) applied to an outside actor. The integration analogue of Kittu's chain of evidence is exact: claim (the vendor's stated capability), argument (the conformance tests it passed), evidence (its signed traces in the field).

Secured comms (Section 5.3). Mutual-TLS identity for every node and agent, signed messages, and encrypted transport over the MANET fabric. An unidentified or unsigned participant is simply not on the network.

Those three layers are enforced through a single onboarding and conformance gate the integration-time twin of the training-data ingestion gate (Section 12.5). The stages, and how each maps onto Kittu's chain of evidence, are below.

"	Stage	What happens	Chain-of-evidence tie
"	1 - Publish	Vendor builds to a published Enlil interface contract (OpenAPI, CoT, SAPIENT ICD, or STANAG 4586). Interface entity (the claim)	
"	2 - Conform	Run against the conformance test harness functional plus adversarial on the SAPIENT / STANAG test-harness pattern. Argument (claim verified)	
"	3 - Sign & register	Issue a cryptographic identity; record container signature, software bill of materials, and vulnerability (VEX) record. Signed provenance	
"	4 - Admit	Bind to a low-trust zone with least-privilege scopes; effectors receive collection-tasking only, never an engagement command. Least privilege (AC-5 / AC-6)	
"	5 - Observe & promote	Run under Kittu; monitor behaviour against the contract; graduate trust as field evidence accrues. Graduated trust / freshness	
"	6 - Revoke	Credential and zone membership are revocable instantly; a failing or compromised module is cut off and quarantined. Kill switch / containment	

Table 21 - The third-party onboarding and conformance gate. It is the integration-time twin of the

training-data ingestion gate (Section 12.5): the same chain-of-evidence discipline, applied to capabilities entering the platform rather than data entering the corpus.

The result is a platform that is open without being porous. A coalition vendor a sensor maker, a counter-drone company, a drone manufacturer, an autonomy start-up can integrate against documented interfaces and standard protocols and be operational quickly. But everything it contributes is zoned by Anshar, governed by Kittu, carried over secured comms, and admitted only through a gate that holds it to the same standard of evidence as everything else Enlil trusts. The open items this raises which conformance suites to mandate, how allied effectors are certified for release authority, and the formal interface specifications themselves are carried into Section 13.

11. Reference Hardware Profiles

A posture statement first, because it determines how to read everything below. Enlil is a software-defined capability: Anshar, Kittu, the ingestion pipeline, and the inference stack run on commercial off-the-shelf (COTS) edge compute, and the configurations in this section are reference profiles designs that prove the size, weight, and power envelope per tier not a bespoke fleet. The platform targets a compute envelope (a CUDA-class accelerator, memory, sustained TOPS, and a power class per tier); any conformant module hosts it, and the Jetson family named below is the current best fit, not a dependency. Software moves on agile release cycles, independent of hardware refresh the deliberate answer to the procurement stagnation that bespoke green fleet acquisition inherits [CAF-18]. Ruggedization, likewise, is a property of the enclosure and carrier procured COTS to MIL-STD-810 and IP baselines not a militarized redesign of the compute.

This section gives a conceptual hardware specification for the FIB node and hub tiers introduced in Section 3.2: representative ruggedized configurations, the interfaces each needs to take in data, the communications that link nodes and hubs, and estimates of power, operating time, and weight. The figures are planning estimates drawn from current ruggedized-Jetson, tactical-radio, and military-power products [53] [63]; they are intended to bound the design space and feed the needs-assessment, not to specify a final bill of materials.

Every node and hub is ruggedized and field-ready: sealed against dust and water, tolerant of shock, vibration, and temperature extremes, and built around screw-locking circular connectors that survive ballistic shock and constant vibration where commercial RJ45/USB would fail. [53][57][61]

11.1 Ruggedization baseline

All Enlil hardware shares a common ruggedization baseline so that a node or hub can be carried, mounted, or emplaced in the field without modification. Representative targets, taken from fielded ruggedized-Jetson systems:

Environmental sealing. IP67 (dust-tight, water-immersion) enclosures with sealed connectors; passive, fanless cooling so the unit is silent and fully sealed against sand, salt fog, and debris. [53][55][57]

Shock, vibration, temperature. MIL-STD-810 shock and vibration; wide operating temperature (roughly -40 °C to +70 °C, module-dependent). [53][54][57]

Power and EMI compliance. MIL-STD-1275 / MIL-STD-704 vehicle/aircraft power input and MIL-STD-461 EMI where vehicle-integrated; wide DC input (commonly 9-36 V). [54][55]

Connectors. Circular MIL-DTL-38999 (or M12-class) screw-coupling connectors with IP67 sealing for power, data, and video, preventing accidental disconnection under maneuver. [53][57]

11.2 Node and hub configurations

Five representative configurations map to the FIB tiers. Each is sized so its interfaces and radios match its role; smaller tiers favor wireless and a minimal sealed connector set, larger tiers add wired sensor and video I/O.

- " Configuration Compute module Form factor & use Ruggedization
- " Micro node Jetson Orin Nano Wearable / dismounted; pocket- or vest-mounted; minimal sealed I/O. IP67, fanless, MIL-STD-810; sealed M12 connectors.
- " Standard node Jetson Orin NX Squad / vehicle node; compact sealed box; drone-payload-capable (low SWaP). IP67, fanless, MIL-STD-810; 38999 connectors.

- " Heavy node Jetson AGX Orin Section / platoon node; vehicle- or mast-mounted; rich sensor and video I/O. IP67, MIL-STD-810H; 38999; MIL-STD-1275/704 vehicle power.
- " Local hub Jetson Thor T5000 Cluster fusion center; vehicle-, shelter-, or transit-case-mounted. IP67, MIL-STD-810H, MIL-STD-461; full vehicle-power compliance.
- " FOB-IFC DGX Station GB300 (1 2×) FOB intelligence fusion centre; shelter-, vehicle-, or transit-case-mounted; air-cooled; runs up to ~80B correlation models. Transit-case rugged; MIL-STD-810 transit; FOB / generator power.

Table 10 · Representative ruggedized FIB configurations mapped to the Jetson tiers. [53][55][56][57]

All ruggedized Jetson carriers and associated hardware are procured and accredited to Canadian standards: the DND Approved Products List (APL) for computing modules, the DND Approved Vendor List (AVL) for procurement, DITSR hardware-security requirements, CFSO.04.035 for radios and encrypted links, and ISED spectrum-allocation approval for MANET frequencies. [CAF-8][CAF-4][CAF-15]

11.3 Data interfaces

Each configuration must take in data from organic sensors (Section 6.2) and from the network. Ruggedized Jetson carriers expose a wide interface set; Enlil selects per tier, preferring wireless and sealed connectors at the small end and adding wired sensor/video I/O at the large end. Representative interfaces available across the family:

- " Interface class Examples Primary use in Enlil
- " Camera / video in MIPI CSI, GMSL2 (long-cable), 3G/HD-SDI, HDMI in EO/IR and FMV capture; mast and vehicle cameras; UAS video. [55][56][58]
- " High-speed data USB 3.2, M.2 (PCIe) for SDR/NVMe, 1/2.5/10 GbE SDR/DF receivers, MASINT sensors, local storage, hub uplinks. [53][55][58]
- " Serial / bus RS-232, RS-422, CAN bus, GPIO, MIL-STD-1553 Legacy sensors, UGS, effectors, vehicle integration. [53][55]
- " Wireless Wi-Fi 6, 4G LTE / 5G, GNSS/RTK (anti-jam) via M.2 Reachback where available; PNT; tethered-device links. [55][58]
- " Tactical radio MANET radio via Ethernet / OEM module The primary node/hub link see Section 11.4. [64][65]

Table 11 · Representative data interfaces across the FIB family; selected per tier. Smaller tiers favor wireless + minimal sealed I/O, larger tiers add wired sensor and video. [53][55][56][58]

11.4 Communications between nodes and hubs

The fabric of Section 3 is bearer-agnostic: it requires an IP transport with a stated latency, throughput, and partition envelope, and it runs over whichever bearer a deployment provides (§3.1). Two deployment profiles cover the range. In the LCSS-integrated profile the expected norm on CAF operations nodes and hubs ride the bearers the force already fields: LCSS data bearers, SATCOM, and LTE/5G where present, with Kittu s transport security layered over each. In the infrastructure-denied profile, every node and hub carries a tactical MANET radio, and the radios self-form a resilient mesh with no fixed infrastructure what keeps the fabric alive where nothing else exists. Modern tactical MANET radios using MIMO waveforms link hundreds of nodes ground, vehicle, UAS with high-bandwidth video, voice, and IP data, at the same time providing electronic-warfare resilience and built-in encryption that Kittu layers its own transport security over. [64][65][66][67]

- " Link Bearer Purpose & character
- " Node <-> node MANET mesh (primary); short-range RF backup Peer observation sharing; self-healing; every node is a relay. [64][66]

- " Node <-> hub MANET mesh (same network, higher-throughput radio at hub) Pooled observations up, direction down; hub aggregates the cluster. [64][67]
- " Hub <-> hub MANET mesh + directional/long-range link where needed Inter-cluster bridging; mesh continuity across the area of operations. [64][66]
- " Hub <-> Central IFC Beyond-line-of-sight: SATCOM / LTE-5G / long-range MANET Reachback for raw data, assessments, and direction. [65][68]

Table 12 · Communications across the fabric. In the infrastructure-denied profile a single MANET mesh carries node, hub, and inter-hub traffic; in the LCSS-integrated profile the same logical links ride existing CAF bearers; beyond-line-of-sight reachback uses SATCOM or cellular. [64][65][66][67][68]

Representative MANET radios (e.g. the Silvus StreamCaster family and Persistent Systems Wave Relay) deliver on the order of tens to ~100 Mbps of mesh throughput, come in handheld, mounted, and OEM-module form factors, are built to IP67/IP68 and MIL-STD-810, and interoperate natively with ATAK/WinTAK aligning directly with the HMI of Section 9. [64][65][66] In Canadian service, radios are drawn from the DND Approved Vendor List and frequency-licensed under ISED, with encryption and key management meeting DITSR and CFSO.04.035 (FIPS 140-2 Level 2+ for Secret data). [CAF-4][CAF-8][CAF-15] Final bearer-profile validation, and radio down-selection for the infrastructure-denied profile, are follow-on work (Section 13).

11.5 Power, batteries, and solar

Power is the binding constraint at the edge. The table below gives estimated draw per configuration a nominal figure for steady sensing and inference and a stressed figure for peak concurrent models including a carrier/I/O/radio overhead on top of the bare module. Battery sizing is referenced to the BB-2590 (the Bren-Tronics BB-2590/U fielded by the CAF and procured through the DND supply chain), the most widely fielded military battery (~298 Wh, ~1.4 kg, rechargeable, smart-bus), with the Conformal Wearable Battery (CWB) for dismounted micro nodes and the six-cell BB-2590 6-Pack APU (~1.2 1.75 kWh) for hubs. [69][70][71][72][73]

Configuration	Nominal draw	Stressed draw	Battery (typical)	Est. runtime (nom / stress)
Micro node	~18 W	~33 W	1× CWB or 1× BB-2590	BB-2590: ~16 h / ~9 h
Standard node	~32 W	~52 W	1× BB-2590	~9 h / ~6 h (×2 ~ 18 h nom)
Heavy node	~58 W	~93 W	1 2× BB-2590	1×: ~5 h / ~3 h; 2×: ~10 h nom
Local hub	~120 W	~160 W	6-Pack APU (~1.5 kWh)	~12 h / ~9 h
FOB-IFC	~1.6 kW	~3.2 kW (dual)	FOB / generator power	Mains-class not battery-portable

Table 13 · Estimated power draw and battery runtime per configuration. Draw includes carrier/I/O/radio overhead; runtimes are planning estimates at the stated battery capacity. [69][70][71][72]

Three power conclusions shape the design. First, a micro or standard node runs the better part of a day on one BB-2590, so a dismounted operator carries spares rather than a generator. Second, the Thor hub is power-hungry at ~120 160 W it needs the 6-Pack APU or vehicle power for a sustained watch, which is why the hub is vehicle-, shelter-, or transit-case-mounted rather than worn. Third, solar recharge is viable for the lighter tiers: foldable military thin-film panels (~30 W each, packable and durable) and solar briefcases/charge-controllers can recharge BB-2590s in the field, and a small array (~100 120 W in good sun) can offset much of a micro/standard node s average draw extending endurance or sustaining an emplaced node indefinitely. Solar cannot, on its own, sustain a Thor hub under load. [72][73] The FOB-IFC is a mains/generator-class system (~1.6 kW single, ~3.2 kW dual) and is not battery- or solar-portable; it is sited at the Forward Operating Base on facility or generator power.

Power sizing reinforces the tiering: the small nodes are battery- and solar-sustainable for dismounted use; the Thor hub trades endurance for inference power and is sited where vehicle or APU power is available.

11.6 Estimated weight

Weight follows the same tier logic and drives how each FIB is carried. The estimates below are for a ruggedized unit plus one battery; add ~1.4 kg per additional BB-2590 and the radio where not integrated. [54][57]

- " Configuration Device (rugged) + 1 battery Carry mode
- " Micro node ~0.5 1.0 kg + CWB (~0.9 kg) or BB-2590 (~1.4 kg) Worn / pocket / vest.
- " Standard node ~1.5 2.5 kg + ~1.4 kg Backpack / vehicle mount / drone payload.
- " Heavy node ~3 5 kg + ~1.4 2.8 kg Vehicle / mast / section kit.
- " Local hub ~5 9 kg (device) + APU (~9 12 kg) Vehicle / shelter / transit case.
- " FOB-IFC ~20 40 kg (1 2× DGX Station) n/a (mains powered) Shelter / vehicle / transit case; wheeled or two-person.

Table 14 · Estimated weights per configuration (ruggedized device plus one battery). Planning estimates; confirm against selected enclosures and batteries in detailed design. [54][57][69]

The takeaway for force design: a micro or standard node is a one-person carry including power, a heavy node is a vehicle or two-person load, and a hub is a vehicle/shelter installation consistent with the roles assigned to each tier in Section 3.2. Exact weights depend on the final enclosure, connector set, and battery count, which are follow-on work (Section 13).

12. Training and Fine-Tuning Data Sources

Section 7 commits Enlil to a family of small, specialized models that are fine-tuned rather than trained from scratch (Section 7.2). That commitment only holds if the data behind each fine-tune is itself trustworthy. This section names credible, openly-licensed candidate sources for the two distinct jobs the inference layer must learn: the per-discipline models that turn raw intelligence data into observations (Section 7.3), and the foundational agents that run Anshar and Kittu (Section 7.4). It then sets out the integrity discipline that keeps adversarial or low-provenance data out of the corpus. The sources below are a starting point for a fine-tuning data strategy – not a final, validated training manifest. Every candidate must clear the ingestion gate of Section 12.5 before it trains a fielded model.

A note on terms, consistent with the rest of this document: intelligence here means the INT disciplines and their data (Section 6); inference means the analytical work the models perform. The datasets in Section 12.2 are intelligence data used to train models that perform inference.

12.1 Sourcing principles and the poison-pill problem

Data poisoning is the deliberate corruption of training data to bend a model's learned behaviour. It comes in three shapes: availability attacks that degrade accuracy broadly; targeted attacks that corrupt a narrow slice – a label, a rule, a phrase; and backdoor attacks that implant a hidden trigger which flips behaviour only when the trigger appears, leaving ordinary evaluation clean. [74] The danger for Enlil is specific: a poisoned sensor or guard model can pass every benchmark and still misclassify on cue in the field. Recent work shows that very small poisoned fractions can implant durable backdoors, that the harm stays dormant through normal validation, and that after-the-fact machine unlearning does not reliably remove it. [74] The defence is therefore preventive and front-loaded – get the provenance right going in, because it is far harder to clean a model afterwards.

Five rules govern what Enlil admits as training data, consistent with DND security-classification policy and the National Security Policy on Responsible Artificial Intelligence (2024): [CAF-12]

Authoritative provenance. Prefer DND, Canadian and allied academic, and recognized-benchmark curators, pulled from the canonical source – not scraped copies or anonymously-redistributed mirrors of uncertain integrity.

Licence fit. Confirm each dataset's licence permits Enlil's intended use. Several strong datasets are non-commercial or research-only; those are usable to pre-train and benchmark, but a fielded commercial model needs either a clean-licensed source or a separately-licensed one.

Signed, hashed ingestion. Every artifact is content-hashed and signed on entry, exactly as Kittu treats evidence – the same chain of custody described in Section 5.3.

Independent validation. Held-out, trusted benchmarks plus anomaly, duplicate, label-flip, and backdoor-trigger scanning run before any artifact is allowed to train a model.

Freshness and expiry. Each dataset carries a review date; drift-prone modalities (RF, OSINT text) are re-validated on a cycle, mirroring the Trust Assurance Framework's freshness rule.

The core discipline

Source from authoritative curators, sign and hash every artifact on ingest, and hold the training corpus to the same chain of custody as everything else Kittu governs.

12.2 Field-collected intelligence: sensor-data sources

These sources feed the Tiny-to-Small models that run on the nodes themselves (Section 7.3), matched to the on-board sensor requirement of Section 6.2 and Table 7. The aim at this tier is not a single large model but a set of compact, discipline-specific ones – IMINT, SIGINT, and MASINT – each fine-tuned on data that resembles what its sensor will actually see.

"	Discipline / modality	Candidate open datasets (curator)	Licence & caveat	Model tier
"	IMINT EO / overhead	xView (DIU + NGA; 60 classes, WorldView-3 0.3 m) [75]; DOTA (academic; 15-18 classes, oriented boxes) [77]	xView non-commercial; DOTA embeds Google-Earth imagery	pretrain / benchmark only Tiny-Small
"	IMINT SAR	MSTAR (AFRL / DARPA public release; 10 ground-vehicle classes) [76]; xView3-SAR (DIU + Global Fishing Watch; Sentinel-1, ~243k objects) [75]; SSDD / OpenSARShip (academic)	MSTAR & xView3 open / credentialed; ship sets research-use	Tiny-Small
"	IMINT EO/IR thermal	Teledyne FLIR ADAS (26k aligned RGB-thermal frames, 15 classes) [78]; LLVIP (15k low-light RGB-IR pairs) [78]	FLIR free under TOS; LLVIP research-use; both suit EO/IR fusion	Tiny-Small
"	SIGINT RF modulation	RadioML 2018.01A (DeepSig; 24 modulations, 2.56 M I/Q frames) [79]; TorchSig / Sig53 (open generator) [79]	RadioML CC BY-NC-SA (non-commercial); TorchSig permissive	generate own corpus for production Tiny-Small
"	MASINT acoustic	Military Audio Dataset (7 classes incl. gunshot, vehicle, helicopter, shelling) [80]; ESC-50, UrbanSound8K, FSD50K, AudioSet [80]	MAD / ESC-50 / FSD50K open CC; AudioSet labels CC, audio via source links	Tiny
"	MASINT seismic / UGS	Sparse open data; lean on self-collection and synthesis; reference seismic-ML corpora for method transfer	Largely build-your-own	flagged open (Section 13) Tiny

Table 16 - Candidate open datasets for the field-collected intelligence models, by discipline, with licence caveats and target model tier.

Two caveats shape how these are used. First, licensing: the strongest RF and overhead sets (RadioML, xView, DOTA) carry non-commercial or imagery-redistribution constraints, so Enlil treats them as pre-training and benchmarking material and, for production, regenerates equivalent data under controlled pipelines it owns – TorchSig for RF, and synthetic / measured-paired SAR (for example the SAMPLE approach) for radar – so the provenance of the fielded model is Enkidu's own. Second, coverage: MSTAR is the long-standing SAR-ATR reference for ground-vehicle targets and is an AFRL/DARPA public release, while FLIR and LLVIP give the aligned visible-infrared pairs that map directly onto the EO/IR fusion of Section 6. Seismic and unattended-ground-sensor data is the sparsest modality – little credible open data exists – so it is largely a self-collection and synthesis problem, carried forward as an open question in Section 13.

12.3 Centrally-sourced intelligence: text and OSINT corpora

The text models in Enlil's scope – Small-to-Medium language models for HUMINT and OSINT (Section 7.3), and the reachback-fed material of Section 6.3 – need entity, event, and relation extraction more than open-ended generation. High-quality labelled INT text is genuinely scarce and frequently sensitive, so the practical path is a cold-start on open, well-provenanced corpora followed by careful, access-controlled specialisation at the Central IFC:

Named-entity and relation extraction from open, established corpora (multilingual NER such as WikiANN-style sets; event-extraction schemas where licensing permits, typically via recognized data consortia).

Entity grounding against open knowledge bases (Wikidata / Wikipedia) for places, organisations, and equipment.

Open multilingual news for register and domain coverage, used for adaptation rather than as ground truth.

The discipline of Section 12.1 applies with extra force here: open-web and social text is the highest poisoning-surface source there is, so anything scraped is treated as untrusted and is provenance-checked, deduplicated, and held out of the trusted training set unless it clears the gate. The function-calling text that lets these models act the reasoning behind Anshar and Kittu agents is covered separately in Section 12.4.

12.4 Training data for the foundational agents: Anshar and Kittu

The foundational agents have two distinct data needs. The first is function-calling and tool-use reasoning the Qwen3 function-calling bases of Section 7.4 that let Anshar s zone and message agents, and Kittu s orchestration, choose and invoke tools correctly. The second is guard and detector training for Kittu s synchronous checks at the gateway: prompt-injection and jailbreak detection, PII and secret screening, faithfulness scoring, and policy classification the purpose-trained small and medium models described in the Trust Assurance Framework.

There is an important structural point here. Per the Trust Assurance Framework, the harness manufactures most of its own training data: the signed run-trace and provenance logs it already records are labelled examples of agent behaviour, mediated tool calls, injection attempts, and adjudicated zone transitions. That self-generated corpus, augmented with synthetic adversarial data and above all the University of Waterloo red-team corpus from the framework s red-team / blue-team programme, is the long-run source. The public datasets below are the cold-start blend that bootstraps the detectors before that self-generated corpus exists; they are not the steady-state source.

Agent / role	Task	Candidate datasets (curator)	Provenance & licence
Anshar	zoning / messaging agents; Kittu orchestration	Function-calling, tool-use reasoning	Salesforce xLAM / APIGen-60k (execution-verified) [81]; Berkeley Function-Calling Leaderboard [81]; ToolACE, Glaive xLAM research-use; APIGen pipeline reproducible regenerate under own provenance
Kittu	ingress / handoff detector	Prompt-injection, jailbreak, instruction-vs-data	NVIDIA Shedu / Nemotron Content-Safety v1-v2 (built for NeMo Guardrails) [82]; WildGuardMix / WildJailbreak [83]; deepset-prompt-injections, HarmBench, JailbreakBench [84] Shedu CC-BY-4.0 (commercial-ready); others mixed open / research
Kittu	tool-call / policy classifier	Least-privilege, parameter & policy screening	AgentDojo, InjecAgent (agentic indirect injection at tool / handoff boundaries) [85]; xLAM relevance-detection split [81] Open research benchmarks also red-team seed data
Kittu	PII / secret detector	Egress and context scanning	ai4privacy open-pii-masking (~1.4 M samples, 23 languages, GDPR / EU-AI-Act tagged) [86] Open (per-release licence); multilingual, synthetic
Kittu	faithfulness / grounding scorer	Output-vs-trajectory grounding	RAGTruth (human-annotated hallucination spans) [87]; HaluEval [87] Open research-use
Kittu	Observer / Judge	Trust-vector & zone adjudication	Primarily self-generated signed run-traces + UWaterloo red-team corpus; guard sets above for cold-start Enkidu-owned provenance (preferred)

Table 17 - Candidate datasets for the Anshar and Kittu agents, by harness role. The Observer and Judge are trained chiefly on the harness s own signed evidence; public sets bootstrap the synchronous detectors.

The NVIDIA Shedu / Nemotron Content-Safety set deserves note: it is CC-BY-4.0, was built specifically to be commercially usable, and was curated by the NeMo Guardrails team the same substrate Kittu sits on (Section 5.1) which makes it an unusually clean fit for the ingress detector. The agentic benchmarks AgentDojo and InjecAgent matter because they exercise indirect injection at exactly the boundaries Kittu is built to broker: agent-to-agent handoffs, agent-to-tool calls, and shared state. Their attack cases double as red-team seed data.

Where the cleanest data comes from

Kittu s cleanest training corpus is its own signed run-trace provenance Enkidu controls end to end. Public guard datasets bootstrap the detectors; the red-team loop sustains them.

One caution carries over from the Trust Assurance Framework: a guard model is itself an attack surface, and public guard datasets can themselves be poisoned. So no single detector is ever a sole gate – multiple models vote, deterministic controls (allow-lists, schema and least-privilege checks) bound the blast radius, and every public set above passes the same ingestion gate as any other source before it trains a detector.

12.5 Integrity controls: the anti-poisoning ingestion gate

Every external dataset, before it trains anything, passes a gate that mirrors Kittu's own evidence pipeline so the training corpus is held to the same standard as live evidence. Because tiny poisoned fractions can implant dormant backdoors that survive ordinary validation and resist unlearning [74][88], the controls are preventive and front-loaded rather than reactive.

[89] U.S. Department of Defense, Modular Open Systems Approach (MOSA); 10 U.S.C. 2446a (NDAA FY2017 §805); OUSD(R&E) MOSA Reference Frameworks. Loosely-coupled severable modules, published key interfaces, certified conformance.

[90] NATO STANAG 4586, Standard Interfaces of UAS Control System (UCS) for NATO UAV Interoperability Data Link Interface (DLI), Command-and-Control Interface (CCI), and Levels of Interoperability (LOI 1-5).

[91] UK Dstl, SAPIENT (Sensing for Asset Protection with Integrated Electronic Networked Technology); Interface Control Document published as BSI Flex 335; NATO STANREC / AEDP-4869. Autonomous sensor modules to decision modules; plug-and-play trials.

[92] Cursor-on-Target (CoT) event message format, US AFRL / MITRE; Team Awareness Kit ecosystem (ATAK / WinTAK / TAK Server); MIL-STD-2525 symbology.

[93] MAVLink micro air vehicle communication protocol; ASTM F3411 / Open Drone ID (Remote ID) broadcast identification standard.

[94] Object Management Group, Data Distribution Service (DDS) for real-time publish/subscribe connectivity; used as a MOSA data-centric fabric (e.g. RTI Connex).

[95] NVIDIA NIM (NVIDIA Inference Microservices): containerized, OpenAPI-compatible inference endpoints; model signing, VEX vulnerability records, least-privilege hardening, air-gapped self-hosting. developer.nvidia.com/nim.

[96] NVIDIA NeMo Agent Toolkit: framework-agnostic agent library (LangChain, LlamaIndex, CrewAI, Semantic Kernel, Google ADK, custom/Python); Model Context Protocol (MCP) and Agent2Agent (A2A) protocol support; public plugin/wrapper API. github.com/NVIDIA/NeMo-Agent-Toolkit.

[97] Model Context Protocol (MCP) and Agent2Agent (A2A) protocol – open standards for agent-to-tool and agent-to-agent interoperability with authentication.

[98] NVIDIA. DGX Station (GB300 Grace Blackwell Ultra Desktop Superchip): 72-core Grace CPU + Blackwell Ultra GPU via NVLink-C2C; 748 GB coherent memory (252 GB HBM3e + 496 GB LPDDR5X); ~20 PFLOPS FP4; ConnectX-8 800 Gb/s (multi-Station linking); ~1.6 kW, air-cooled. [nvidia.com](https://www.nvidia.com), 2026.

[99] NVIDIA. DGX GB300 and GB200/GB300 NVL72: rack-scale, liquid-cooled Grace Blackwell (Ultra) systems – 72 GPUs, ~13.4 TB unified memory, ~120 kW per rack; data-centre / Central-IFC class. [nvidia.com](https://www.nvidia.com), 2025 2026.

[100] NVIDIA. GB200 Grace Blackwell Superchip and GB200 NVL2 / NVL4 nodes (2 4 Blackwell GPUs with Grace CPUs over NVLink-C2C). [nvidia.com](https://www.nvidia.com), 2025 2026.

[101] NVIDIA. B200 / B300 (Blackwell / Blackwell Ultra) GPUs – 192 GB / 288 GB HBM3e; HGX B200 / B300 eight-GPU servers. [nvidia.com](https://www.nvidia.com), 2025 2026.

Canadian (CAF / DND) references

[CAF-1] Department of National Defence (Canada). Canadian Defence Policy (2022). DND/CAF

modernization and technology priorities.

[CAF-2] Department of National Defence (Canada). DND/CAF Artificial Intelligence Strategy (2023). Strategic direction for AI-enabled defence capabilities.

[CAF-3] Department of National Defence (Canada). DND Instruction 2004/4, Classification, Handling, and Declassification of Defence Information. Defines Protected A/B/C, Secret, and Top Secret and handling procedures.

[CAF-4] Department of National Defence (Canada). CFSO.04.035, Communications Security Orders. Encryption, key management, and secure-communications protocols.

[CAF-5] Department of National Defence (Canada). CFAO 11-001, Canadian Forces Administrative Orders Organization. Defines CF Int (Intelligence Branch) and command structures.

[CAF-6] Department of National Defence (Canada). CFAO 20-204, DND Intelligence Policy. Tasking, collection management, all-source fusion, and Priority Intelligence Requirements (PIRs).

[CAF-7] Department of National Defence (Canada). CFAO 11-025, Remotely Piloted Aircraft Systems (RPAS) Operations. Operating constraints, airspace integration, and tasking for Canadian RPAS.

[CAF-8] Department of National Defence (Canada). DND Information Technology Security Requirements (DITSR). Security-control baselines for DND systems.

[CAF-9] Canadian Centre for Cyber Security (CCCS). Common Criteria Evaluation (EAL 2 4 guidance for software-module accreditation).

[CAF-10] Government of Canada. ITSP.40.111, Government of Canada Security Control Profile for Cloud. Cloud-security baseline for central-fusion-centre hosting.

[CAF-11] Department of Justice (Canada). Personal Information Protection and Electronic Documents Act (PIPEDA); Access to Information Act (ATIA). Govern personal-data handling and disclosure.

[CAF-12] Communications Security Establishment (CSE) (Canada). National Security Policy on Responsible Artificial Intelligence (2024). Principles for AI use in Canadian security and defence.

[CAF-13] Canadian Armed Forces Cyber Operations Centre (CAFCOC). Organizational mandate for CAF cyber operations and defence.

[CAF-14] NORAD (North American Aerospace Defence Command). NORAD Modernization Strategy. Technological and operational priorities for Canada-US air and aerospace defence.

[CAF-15] Innovation, Science and Economic Development Canada (ISED). Radio Spectrum Management. Frequency allocation and licensing for tactical MANET radios.

[CAF-16] Government of Canada. Security of Information Act (R.S.C., 1985, c. O-5). Special Operational Information (SOI) and Persons Permanently Bound to Secrecy (PPBS).

[CAF-17] NATO Security Committee. NATO security classifications COSMIC Top Secret, NATO Secret, NATO NATO Restricted €24 and the ATOMAL marking for atomic information.

[CAF-18] S. Jobin, CWO (Ret). Enlil Conceptual Architecture Advisory Document v1.1. Independent subject-matter-expert review for MIOVSKI GROUP, 1 July 2026 (internal). The CAF-framework analysis, tactical scenarios, and sensor inventory that Sections 1.4, 1.5, 6, 9, and 11 of this draft respond to.

" Stage Control applied Chain-of-evidence tie

" 1 - Acquire Pull only from the authoritative curator s canonical source; record origin, version, and retrieval time. Provenance entity (origin)

" 2 - Verify Check publisher signature or checksum; confirm the licence permits the intended use; reject on any mismatch. Signature + licence gate

- " 3 - Scan Run statistical-anomaly, duplicate, label-flip, and backdoor-trigger detection; quarantine outliers for review. Measurable outcome
- " 4 - Record Content-hash the admitted artifact, hash-link it to its source, sign it, and set an expiry date. Signed provenance record
- " 5 - Admit Add to a versioned, access-controlled corpus with least-privilege write; no anonymous contribution. Least privilege (AC-5 / AC-6)
- " 6 - Re-validate After each fine-tune, regression-test on a trusted held-out benchmark; re-check drift-prone modalities on a cycle. Freshness / expiry rule

Table 18 - The training-data ingestion gate. Each stage maps onto a control in the Trust Assurance Framework, so the corpus carries the same provenance guarantees as live evidence.

The effect is that the training corpus becomes auditable in the same way Kittu makes inference auditable. For any fielded model, Enkidu can answer a single question what data trained this, from where, verified how, and when does it expire? which is precisely the question the Trust Assurance Framework answers for every live decision. The open items this section surfaces the scarcity of seismic and UGS data, the central curation of INT text, and the build-out of the red-team corpus are carried into Section 13.

12.6 Audit log schemas and export formats

Accreditation and disclosure both depend on machine-readable evidence. Kittu therefore emits its chain of custody as structured, exportable audit logs. A compliance audit entry, in JSON Lines, looks like:

```
{
  "timestamp": "2026-06-23T14:32:00Z",
  "source_sensor": "node-12-eoir-camera",
  "source_classification": "Unclassified",
  "data_type": "IMINT/EO-IR",
  "ingestion_stage": "Authenticate, inspect, and admit",
  "origi_decision": "PERMIT",
  "tangram_zone": "Sector-4-IMINT",
  "trust_zone_required": "Z1",
  "trust_zone_current": "Z2",
  "provenance_hash": "sha256:abc123def456...",
  "signature": "[cryptographic signature of entry]",
  "audit_trail": "Signed by node-12@FIB-cluster-3, timestamp, nonce"
}
```

An exportable compliance report aggregates these entries for DND accreditation and ATIA review, and carries: data lineage from ingestion through fusion output; every policy rule applied (classification, Protected B handling, coalition-release restrictions); enforcement actions taken (permit, restrict, downgrade, segregate); trust-zone transitions and their justifications; chain-of-custody signatures and hash links; and any operator overrides with the human-approval evidence behind them. The report is exportable as JSON, XML, or CSV to external compliance-review tools DND accreditation support systems and Privacy Impact Assessment tooling. [CAF-3][CAF-11]

13. Open Questions and Follow-On Work

As a first draft, this architecture deliberately leaves the following to be resolved in detailed design. They are the spine of the needs-assessment.

Sensor selection and integration. Confirm specific field sensors per discipline and tier (SDR/DF receivers, EO/IR gimbals, SAR payloads, UGS, MASINT and CBRN detectors) and their physical and data integration into each FIB the build-out of the Section 6 sensor requirements.

Central-feed connectors and the common schema. The secure reachback connectors for Military Intelligence direction, GEOINT, OSINT, and Allied/Partner feeds, and the formal, machine-readable schema of the Enlil Observation whose concept model Section 6.4 fixes (Table 7A).

Jetson-tier HMI capability research. Confirm, per tier, the display pipelines, audio I/O, and concurrent model headroom available for HUD and on-device speech the explicit open item behind Section 9.

Model selection, fine-tuning, and quantization budgets. Confirm the specific base models per data type and per agent (Section 7), the fine-tuning datasets and LoRA/QLoRA recipes, and the quantization that fits each tier's power and memory envelope.

Bearer profiles and RF down-selection. Validation of the two deployment profiles of Section 11.4 LCSS-integrated and infrastructure-denied and, for the latter, the specific MANET radio and waveform (range, throughput, EMCON posture), and how Kittu's transport security layers over each bearer.

Hardware bill of materials and SWaP validation. Final enclosures, carrier boards, connector sets, and battery counts per tier, validating the estimated power, runtime, and weight of Section 11 against measured loads.

FOB-IFC hardware down-selection. Confirm the FOB-IFC platform (DGX Station GB300 single vs dual, or a GB200-class shelter rack), its ruggedization and power/cooling at the Forward Operating Base, and the ~80B correlation models it runs across dozens of FIB hubs (Section 3.3). [98][99]

Anshar zone and continuum schema. A formal, machine-readable schema for zones and continuums, their transition events, and the relevance rules that scope notification.

Kittu <-> NeMo Guardrails integration. The precise interface by which Kittu's trust-zone checks compose with NeMo Guardrails' execution rails on each tier, and the deterministic fallback when a guard model is unavailable.

Chain-of-custody and key management. Signing-key provisioning anchored in each Jetson's hardware root of trust, revocation on capture, and the confidence-scoring function over the evidence chain.

ATAK / CoT mapping. The exact mapping of Enlil tracks, zones, and alerts to CoT message types, and plugin vs. server-relay integration. [13][14]

Reconciliation semantics. How divergent post-partition pictures are merged, and how conflicts are surfaced rather than silently resolved.

Accreditation and sovereignty. Mapping the design to the relevant security accreditation, and ensuring the build, models, and provenance are controllable end-to-end.

Formal accreditation pathway. Map the design to the DND Trusted IT Product Evaluation Program (TIPTOP) for hardware and to Canadian Centre for Cyber Security (CCCS) Common Criteria evaluation (EAL 2-4) for the software modules (Kittu, Anshar agents), against DITSR security-control baselines. [CAF-8][CAF-9]

Audit-log export APIs and formats. Formalize the JSON/XML/CSV schemas for compliance audit logs (Section 12.6) and a REST API for exporting them to DND accreditation systems and ATIA/Privacy Act review tools, keeping the hash-linked chain of custody immutable and tamper-evident on export. [CAF-11]

Cross-domain solution (CDS) integration. Selection and accreditation pathway for the CDS brokered at accredited domain boundaries (Section 6.4, stage 3), and the custody record that spans both sides of each transfer.

Joint data-link gateway integration. The concrete gateway interfaces by which the hub and FOB-IFC tiers produce and consume J-series traffic (Link 16 via JREAP-C) and AdatP-3/APP-11 messages through existing certified gateways at the TF HQ integration point (Section 10.3) including message-set subsetting and the releasability profile per network. [CAF-18]

14. Conclusion

Enlil moves fusion to the edge without giving up coherence or accountability. A blended mesh and distributed hub-and-spoke fabric keeps the picture alive through the loss of any unit; Anshar bounds the fabric into physical zones and the logical continuums that run over them, and makes the movement of threats observable and notifiable across the fabric; and Kittu wraps every agent in a trust harness that, layered over NeMo Guardrails, decides what each agent may sense or command, secures every link, and signs an end-to-end chain of custody that turns evidence into calibrated threat confidence. Each FIB becomes a virtual combat assistant that informs, observes, records, and shares seamlessly, and within policy.

This first draft is intended to anchor the detailed design and the needs-assessment that follows. The open questions in Section 13 are the next increment of work; the principles in Section 2 are the standard against which that work should be judged.

Infer forward. Survive partition. Earn trust. Evidence before confidence. Serve the operator. Everything in Enlil follows from those five commitments.

References

- [1] Department of National Defence (Canada), IDEaS Program. Reliable AI sensor fusion for real-world missions. Competitive Projects challenge, 2026. canada.ca.
- [2] Enkidu Enterprises. Anshar Overview Zones, Policies, Conops & Reporting (companion document).
- [3] Enkidu Enterprises. Trust Assurance Framework v1.6 and Trust Assurance Overview (companion documents).
- [4] Enkidu Enterprises. Intelligence Fusion Center Operations: Detailed Overview (companion document).
- [5] S. Rose et al. Zero Trust Architecture, NIST Special Publication 800-207, 2020 (PDP/PEP; never trust, always verify).
- [6] NVIDIA. Jetson Orin module family Orin Nano (up to 67 TOPS, 7.25 W), Orin NX (up to 157 TOPS, 10.40 W), AGX Orin (up to 275 TOPS, 15.60 W). [nvidia.com / developer.nvidia.com](https://nvidia.com/developer.nvidia.com), 2024-2026.
- [7] NVIDIA. Jetson Thor / AGX Thor T5000 up to ~2070 FP4 TFLOPS, 128 GB, 40-130 W. nvidia.com and NVIDIA Technical Blog, 2025-2026.
- [8] Forecr. NVIDIA Jetson Comparison: Orin Nano to AGX Thor (module specifications and carrier boards), 2026.
- [9] ProventusNova. Jetson Orin AGX vs Orin NX vs Orin Nano (TOPS, memory bandwidth, and concurrency trade-offs), 2026.
- [10] ThinkRobotics. NVIDIA Jetson Orin NX Module Review and Jetson AGX Thor Developer Kit Review, 2026.
- [11] T. Rebedea et al. NeMo Guardrails: A Toolkit for Controllable and Safe LLM Applications with Programmable Rails. EMNLP (System Demonstrations), 2023; [arXiv:2310.10501](https://arxiv.org/abs/2310.10501).
- [12] NVIDIA. NeMo Guardrails documentation input, dialog, retrieval, execution, and output rails; agentic/tool-call validation. docs.nvidia.com/nemo/guardrails, 2026.
- [13] Android Team Awareness Kit (ATAK) and Cursor-on-Target (CoT) TAK Product Center / tak.gov; ATAK-CIV open-source release. 2020-2026.
- [14] CoT over mesh bearers e.g., Meshtastic ATAK Plugin (CoT <-> mesh; PLI, chat, off-grid). github.com/meshtastic/ATAK-Plugin, 2025-2026.
- [15] U.S. Army. FM 2-0, Intelligence (all-source intelligence; the G-2/J-2 staff and the Analysis and Control Element). globalsecurity.org; irp.fas.org.
- [16] U.S. Army. FM 34-25-3, Analysis and Control Element / all-source production and dissemination (collection management; multi-INT integration; coalition feeds).
- [17] Signals intelligence COMINT and ELINT overview. en.wikipedia.org/wiki/Signals_intelligence; MAG Aerospace, What Is Signals Intelligence (SIGINT), 2026.
- [18] Electronic Support Measures and direction finding (emitter parameters, DF, signature libraries). radartutorial.eu; Sierra Nevada Corp., Full-Spectrum ESM/SIGINT.
- [19] Tactical SDR-based COMINT/ELINT and direction-finding systems (man-portable receivers, multichannel DF). Novator Solutions, EW/SIGINT products, 2026.
- [20] U.S. Army. FM 2-22.3, Human Intelligence Collector Operations (HUMINT collection methods; SALUTE reporting; source reliability matrix). 2006.

- [21] U.S. Army CALL. Commander's Guide to Human Intelligence (HUMINT) (tactical reporting flow; combat information; PIR-driven collection).
- [22] Tactical Questioning: Soldier's Handbook (SALUTE report format; point-of-collection reporting). publicintelligence.net.
- [23] Insitu / Inside Unmanned Systems. Sensors: A Force Multiplier (EO/IR full-motion video; wide-area search radar; GMTI/target tracking on UAS). 2021.
- [24] Unmanned Systems Technology. ISR Sensors / Payloads EO/IR gimbals (LWIR/MWIR/SWIR), SAR, LiDAR, GMTI; SWaP and platform constraints. 2025.
- [25] Unmanned Systems Technology / Skyfish / JOUAV. Tactical ISR UAV payloads EO/IR zoom, SAR, AIS, GNSS-aided INS (anti-jam/anti-spoof); backpackable ISR. 2024 2026.
- [26] National Centre for Geospatial Intelligence (UK). GEOINT and OSINT delivery (foundation geospatial data; mapping, targeting, navigation). gov.uk, 2025.
- [27] U.S. Army. FM 2-0, Chapter 9, Measurement and Signatures Intelligence (MASINT subdisciplines; radar, E-O, IR, acoustic, RF, nuclear, seismic sensors). globalsecurity.org.
- [28] IC21 / GovInfo. MASINT: Measurement and Signatures Intelligence (sensor classes: radar, optical, IR, acoustic, nuclear/radiation, spectroradiometric, seismic, material sampling).
- [29] Measurement and signature intelligence REMBASS / Improved REMBASS unattended ground sensors (seismic, magnetic, IR, acoustic). en.wikipedia.org/wiki/Measurement_and_signature_intelligence.
- [30] Geophysical MASINT unattended ground sensors for personnel/vehicle detection (ADSID heritage; seismic/acoustic). en.wikipedia.org/wiki/Geophysical_MASINT.
- [31] Social Links / The Defense Post. OSINT in Military Operations (social media, civilian video, commercial satellite, real-time feeds at the tactical level). 2026.
- [32] Venntel. OSINT Data Sources & Geolocation Intelligence (seven source categories incl. AIS/ADS-B vessel and flight tracking; limitations). 2026.
- [33] Payload / CETaS (Alan Turing Institute). Commercial-satellite and social-media OSINT in the Ukraine conflict (tactical geolocation; ML-scaled collection). 2022 2023.
- [34] CallSphere. Small Language Models That Beat GPT-4: Phi-4, Gemma-3, and SmolLM-3 Benchmarks (sub-10B models; distillation from frontier teachers; deployment tiers). 2026.
- [35] Intuz. Top 10 Small Language Models (SLMs) (Phi, Gemma, Mistral 7B, Llama 3.2 1B/3B, Qwen 2.5; on-device vs. cloud trade-offs). 2026.
- [36] Renard Digital. Small Language Models on Edge Devices (Qwen3-4B as strongest fine-tuning base; AWQ/GGUF/INT4 quantization; Ollama/llama.cpp/MLX). 2026.
- [37] Meta-Intelligence. Small Language Models: Phi-4 vs Gemma 3 vs Llama 3.3 Enterprise Edge AI (LoRA fine-tuning; 4-bit Qwen 2.5-7B reaching 92% on a vertical task with ~3k labels). 2025 2026.
- [38] Digital Applied / KDnuggets / Local AI Master. SLM business and ranking guides (Gemma 3/4 E2B/E4B effective params; Phi-4-mini 3.8B; Qwen 3; SmolLM-3). 2026.
- [39] E. J. Hu et al. LoRA: Low-Rank Adaptation of Large Language Models ; T. Dettmers et al. QLoRA: Efficient Finetuning of Quantized LLMs (parameter-efficient fine-tuning). 2021 2023.
- [40] H. Wang et al. LMW-YOLO: Lightweight model for small-object detection in remote-sensing images (~2.6M params; vs. YOLOv8n/v11n/v12n/YOLO26n). Scientific Reports, 2026.
- [41] Designing Object Detection Models for TinyML (YOLO/SSD lineage; PPYOLOtiny, Trident-YOLO depthwise-separable for mobile). [arXiv:2508.08352](https://arxiv.org/abs/2508.08352), 2025.

- [42] TinyChirp: Bird-song recognition with TinyML on low-power acoustic sensors (tiny CNNs on raw waveforms). arXiv:2407.21453, 2024.
- [43] M. M. Farag. A Tiny Matched-Filter-Based CNN for ECG Classification at the Edge (tiny 1-D-signal CNN exemplar). Sensors 23(3):1365, 2023.
- [44] Label Your Data. VLM: How Vision-Language Models Work (Phi-4-Multimodal on Jetson Orin sub-100ms; small VLMs 1B 10B for edge; VLA models). 2026.
- [45] A. Marafioti et al. SmoVLM: Redefining small and efficient multimodal models (256M 3B); PaliGemma 3B (SigLIP 400M + Gemma 2B). arXiv:2504.05299, 2025.
- [46] Dextralabs / BentoML. Top vision-language models (Phi-4-Multimodal, Pixtral for drones, Qwen2.5-VL/Qwen3-VL 3B/7B/30B-A3B, Gemma 3). 2025 2026.
- [47] T. J. O Shea et al. Over-the-Air Deep Learning-Based Radio Signal Classification and Convolutional Radio Modulation Recognition Networks (RadioML 2016.10a; CNN on raw IQ). 2016 2018.
- [48] Deep Learning-Based Automatic Modulation Classification Using Robust CNN (raw IQ; ~99.8% at 10 dB SNR). PMC10708862, 2023.
- [49] DL-AMC: Deep Learning for Automatic Modulation Classification (ResNet-18/50, MobileNetv2 on eye-diagram representations). arXiv:2504.08011, 2025.
- [50] Seeing Radio: Explainable Modulation Recognition with Vision-Language Models (QLoRA-fine-tuned Qwen2.5-VL-7B on RF spectrograms/IQ panels). arXiv:2601.13157, 2026.
- [51] NVIDIA. NeMo Guardrails agentic security: execution rails for tool-call validation; reasoning content-safety models (e.g. Nemotron-Content-Safety). docs.nvidia.com/nemo/guardrails, 2026.
- [52] Llama Guard / ShieldGemma-class guard models (1B 27B) for input/output content moderation beneath trust-zone checks. Meta / Google, 2024 2026.
- [53] Things Embedded. Rugged NVIDIA Jetson Computers (AGX Orin; 4× GbE, USB 3.1, RS-232/422, CAN, M.2; 18 32 V DC; 40&+70 °C; IP67; MIL-STD-1275/704/461/810). things-embedded.com, 2026.
- [54] Bren-Tronics / Grokipedia. BT-70791JV BB-2590/U lithium-ion battery (298 Wh, ~1.4 kg, >750 cycles, MIL-PRF-32052/1, SMBus). 2024 2026.
- [55] Connect Tech. Sentry-X2 MIL-Rugged System (Jetson AGX Orin Industrial; 4× SDI in / 2× out; 10/2.5/1 GbE; USB 3.2; IP67 D38999 connectors; MIL-STD-810H). connecttech.com, 2024.
- [56] NEXCOM. Jetson AGX Orin edge AI solution (8× MIPI/GMSL2 + 2.5 GbE for cameras/LiDAR; expansible GNSS, LTE/5G, Wi-Fi; 9 36 V; MIL-STD-810H). nexcom.com, 2023.
- [57] Forecr. Military Computer Solutions / MILBOX-AGX, MILBOX-THR, MILBOX-ORNX (Orin & Thor; MIL-STD-1275/704/810/461; IP67; MIL-DTL-38999 circular connectors; fanless). forecr.io, 2026.
- [58] Forecr. DSBOARD-AGX carrier board and Things Embedded I/O notes (GbE, USB, HDMI/DP, GPIO, serial, CAN, MIPI CSI/GMSL; M.2 LTE/5G; 12 32 V DC). 2024 2026.
- [59] Rugged Science. RAC-1000 Jetson AGX Orin IP67 rugged AI computer (9 50 V DC, ignition control, MIL-STD-810G, EN50155). ruggedscience.com, 2026.
- [60] Neosys Technology. NRU / SEMIL rugged Jetson AGX Orin / Orin NX systems (fanless wide-temperature; MIL-STD ruggedization). neosys-tech.com, 2026.
- [61] ADS Inc. / Bren-Tronics. BB-2590/U battery and portable power systems incl. 6-Pack APU (six BB-2590; ~1.2 1.75 kWh; UPS/standalone). equipment.adsinc.com, 2024 2026.
- [62] Inventus Power. Conformal Wearable Battery (CWB) and Hi-Performance BB-2590 (298 Wh); MIL-PRF-32383/4A warfighter-worn power. inventuspower.com, 2022 2026.

- [63] Epsilor / Denchi / Power-Time / LiTech. BB-2590/U-compatible battery family (14.4 V / 28.8 V dual-mode; SMBus smart battery; tactical/sensor/robot power). 2024 2026.
- [64] Silvus Technologies (Motorola Solutions). StreamCaster MANET radios & MN-MIMO waveform (hundreds of nodes; handheld/mounted/OEM; IP67/IP68, MIL-STD-810G; ATAK/WinTAK interop). silvustechnologies.com, 2022 2026.
- [65] Silvus Technologies. StreamCaster NEXUS / SM5200 (up to ~100 Mbps; self-forming adaptive mesh; EW resilience; dismounted SWaP). Motorola Solutions press release, 2025 2026.
- [66] Persistent Systems. Wave Relay MANET (MPU5; large-scale mobile ad-hoc networking; UK MoD Project CAIN). defenseadvancement.com, 2026.
- [67] Silvus Technologies. StreamCaster 4400/4200 enhanced MANET radios; DoD certification for secure UAS operations (drone-to-ground mesh; high-bandwidth video/voice/sensor). 2026.
- [68] Spectra Group / general tactical comms. Beyond-line-of-sight reachback bridging MANET to SATCOM for deployed forces (GENSS). defenseadvancement.com, 2025.
- [69] SupplyNet / military battery vendors. BB-2590/U (~298 Wh) and BB-2557/U (99 Wh) capacities, weights, and charger ecosystem. thesupplynet.com, 2026.
- [70] NVIDIA. Jetson Thor Series Modules Data Sheet DS-11945-001 (T5000 configurable 40 130 W; 120 W default; 7 20 V input; 25&+115 °C Tj). developer.nvidia.com, June 2026.
- [71] NVIDIA Jetson Orin module power profiles Orin Nano (7 25 W), Orin NX (10 40 W), AGX Orin (15 60 W). developer.nvidia.com, 2024 2026.
- [72] PowerFilm Solar / Energy Technologies Inc. Government & DoD foldable thin-film tactical solar (dismounted, packable, non-glass; solar briefcases/suitcases; charges standard batteries). 2026.
- [73] Global Solar / Bren-Tronics. Sunling foldable 30 W military solar charger and solar-capable BB-2590 charge controllers (CIGS thin-film; field battery recharge). 2024 2026.
- [74] Survey & guidance on ML data poisoning: targeted, availability, and backdoor attacks; small poisoned fractions, dormancy through ordinary validation, and the ineffectiveness of machine unlearning. NIST/industry and academic syntheses, 2023-2026.
- [75] DIU (Defense Innovation Unit) & National Geospatial-Intelligence Agency. xView (overhead object detection, 60 classes, WorldView-3) and xView3-SAR (Sentinel-1 dark-vessel detection, ~243k objects; DIU + Global Fishing Watch). iuu.xview.us; DIUx-xView. 2018-2022.
- [76] Air Force Research Laboratory & DARPA. MSTAR (Moving and Stationary Target Acquisition and Recognition) public release SAR ATR dataset, 10 ground-vehicle classes, X-band 1-ft. sdms.afrl.af.mil.
- [77] G.-S. Xia et al. DOTA: A Large-scale Dataset for Object Detection in Aerial Images (v1.0-v2.0; 15-18 categories, oriented bounding boxes). captain-whu.github.io/DOTA.
- [78] Teledyne FLIR. Free ADAS Thermal Dataset v2 (~26k aligned RGB-thermal frames, 15 classes); and X. Jia et al., LLVIP: A Visible-Infrared Paired Dataset for Low-light Vision (~15k pairs). flir.com/oem/adas; 2020-2021.
- [79] DeepSig Inc. RadioML 2018.01A (24 modulations, 2.56 M I/Q frames; CC BY-NC-SA 4.0); and TorchSig / Sig53 open RF signals toolkit and synthetic dataset. deepsig.ai/datasets; torchsig.com.
- [80] Military Audio Dataset (MAD), Scientific Data, 2024 (7 classes incl. gunshot, vehicle, helicopter, shelling); K. Piczak, ESC-50; J. Salamon et al., UrbanSound8K; E. Fonseca et al., FSD50K; J. Gemmeke et al., AudioSet.
- [81] J. Liu et al. (Salesforce). APIGen / xLAM-function-calling-60k (execution-verified function-calling data); S. Patil/Yan et al., Berkeley Function-Calling Leaderboard (BFCL). huggingface.co/Salesforce;

gorilla.cs.berkeley.edu.

[82] S. Ghosh et al. (NVIDIA). Shedu / Nemotron Content Safety Dataset v1 (CC-BY-4.0) and v2 (Aegis2.0), curated by the NeMo Guardrails team for content-safety guardrails. arXiv:2404.05993; arXiv:2501.09004.

[83] S. Han et al. WildGuard / WildGuardMix and WildJailbreak (Allen AI); PKU-Alignment BeaverTails; ToxicChat. Allen Institute for AI / academic releases, 2024.

[84] deepset/prompt-injections and safe-guard-prompt-injection (Hugging Face); M. Mazeika et al., HarmBench (arXiv:2402.04249); P. Chao et al., JailbreakBench (NeurIPS D&B 2024).

[85] E. Debenedetti et al., AgentDojo: A Dynamic Environment to Evaluate Prompt Injection Attacks on LLM Agents (arXiv:2406.13352); Q. Zhan et al., InjecAgent (arXiv:2403.02691) - indirect prompt injection at agent boundaries.

[86] ai4privacy. open-pii-masking / pii-masking series (~1.4 M samples, 23 languages, 19 PII classes; GDPR / EU-AI-Act tagged). huggingface.co/datasets/ai4privacy.

[87] C. Niu et al., RAGTruth: human-annotated hallucination spans for retrieval-augmented generation; J. Li et al., HaluEval. 2023-2024.

[88] N. Baracaldo et al. and related work on data-provenance-based mitigation of poisoning (data lineage to detect and filter malicious inputs); cryptographic dataset signing, checksums, and held-out trusted-benchmark validation as layered defences. 2018-2026.