

THE ALL-SOURCE INTELLIGENCE PLATFORM

IntelWatchtower

The whole intelligence cycle —
collect, task, fuse, detect, investigate,
report — in one governed workspace
instead of a dozen disconnected tools.

COLLECTION AND TASKING

ALL-SOURCE FUSION

COMMON OPERATING PICTURE

CYBER DETECTION

CLASSIFICATION BY DEFAULT



Pre-launch. Launching 2026, waitlist open. Everything described here is marked built, part-built or planned.

[INTELWATCHTOWER.COM](https://intelwatchtower.com)

01 The short version

IntelWatchtower is one place to gather what you know, work out what it means, and act on it — with a record of where every claim came from and how sure you are of it.

A question you need answered becomes a requirement. The requirement becomes an order issued to a source that can answer it. The source delivers reporting, which resolves onto the people, places, vessels and events you are tracking.

Two labels travel with every record. A classification marking says how it must be handled. A confidence grade says how much to trust it. Both are fields in the data, not conventions in a document.

THE PROMISE

One current picture, honest about its own uncertainty, with the trail back to the source still intact.

One

Model for the whole cycle

Eleven record types form one connected chain rather than eleven separate products.

7

Kinds of source

Human, signals, geospatial, imagery, measurement, open-source and cyber, each with a reliability score.

4

Classification levels

A first-class field on every record, set by a person and never inferred by a machine.

Free

Tier that runs real work

Full feature surface at activation-sized limits, with no card and no expiry.

Who it is for

Intelligence analysts and fusion cells

One entity model, one search and one link trail, so you stop re-deriving what the system already knows.

Collection and tasking managers

Requirements, tasking, source reliability and feed health are records, so an untasked gap is visible.

Investigators and case teams

A living case file with timestamped entries and attributed evidence that survives a handover.

Smaller teams priced out of programme-scale platforms

Newsrooms, NGOs, conflict monitors, corporate investigators and small fusion cells get full-cycle tooling on published pricing that starts free.

Security operations and threat teams

Cyber lives inside the same picture. An alert names what fired, pivots to the entity, escalates into a case.

Commanders and operations centres

One shared view of map, entities, alerts and mission status, so the brief reflects the current situation.

Security and governance leads

Classification, role-based access and an attributable trail are enforced, so a review is a query.

03

What breaks today

Ordinary, daily friction from running an intelligence function on tools that were never designed to talk to each other. Each one costs either time, accuracy or accountability.



Where the picture fails

The picture lives in a dozen places

Analysts export from one tool and re-key into another, and provenance evaporates in the middle.

- One tasked pipeline normalises everything and tags it with source, confidence and time.

Nobody can say how sure we are

Conflicting reporting is quietly averaged away, and a corroborated fact reads like a single claim.

- Confidence and reliability are schema fields; conflicts surface for a person to resolve.

Collection drifts from what was asked

A requirement raised in a meeting is tasked informally, and months later the gap is still open.

- Requirements carry priority, lifecycle, owner and due date. The loop closes on the record.

A feed goes silent and nobody notices

A pipe stops delivering at three in the morning and stays broken through the weekend.

- Every feed carries an ingest state, a last-ingest timestamp and a running count.

Where the system around it fails

Cyber sits in its own console

Detections are worked in one tool while the entities and cases they relate to live somewhere else.

- Indicators, alerts, entities and cases share a single model, so a pivot needs no second tool.

The briefing ages the moment it is built

By the time the decision is taken, the picture underneath it has already moved.

- The operating picture refreshes as collection arrives; products move through draft, review, release and recall.

Governance is bolted on after the leak

Classification lives in a filename convention, and proving who saw a record means reading the wrong logs.

- Access is checked on every read and write before the request lands, and every change is attributed.

Full-cycle tooling is priced out of reach

Multi-year procurement, seven-figure contracts and embedded engineers put the alternatives out of reach.

- Published tiers from free to per-seat enterprise, priced by analyst seat, in four regional bands.

04 What you get

A small number of record types that each mean something specific in intelligence work.

COLLECTION

Sources and sensors

A collection asset across seven disciplines, with status, region, operator and a reliability score.

REQUIREMENTS

The questions you need answered

An intelligence need with a priority from routine to flash, an owner, a due date and a lifecycle.

TASKING

Orders issued against a source

Instructions, a collection window and a result summary. This is how a gap becomes action.

FEEDS

Live pipes, with their health on show

Structured, threat-exchange, system-log or geospatial reporting, with ingest state visible.

FUSION

Tracked entities

People, organisations, vessels, facilities, nodes and events with confidence, coordinates and threat score.

INDICATORS

The observables that tie cyber in

Addresses, domains, hashes, coordinates, signatures and behaviours with validity windows.

ALERTS

The queue a watch actually works

Severity, a triage lifecycle, an assignee, detection and resolution timestamps and triage notes.

CASES

An investigation as one thread

Timestamped, attributed entries of evidence, notes and actions — chain of custody by construction.

PRODUCTS

Missions and finished reporting


A mission owns the cases and reports under it; reports move through draft, review, release and recall.

05

Put to work

The shapes the work takes, written from the point of view of the person doing it — with an honest note where a capability is still being built.

Collection and tasking




COLLECTION MANAGER

Bring every source into one tasked pipeline

Sources are registered across all seven disciplines. Feeds attach and stream reporting in. Material is normalised, de-duplicated and tagged with source, confidence and arrival time.

- One current stream fusion can rely on, with the source of every record still attached.




COLLECTION MANAGER

Turn an intelligence gap into a tasked order, and close it

The gap is written down as a requirement with a priority and a due date, then broken into tasking orders against the sources qualified to satisfy it. When reporting answers it, it moves to satisfied.

- Every question has a visible path from asked to answered.

Fusion and detection




FUSION ANALYST

Fuse many partial signals into one honest assessment

Open-source, geospatial, cyber, sensor and partner reporting correlate against the same entity and event model. Each claim carries a confidence grade and a trail back to the observations behind it.

- One coherent picture, with the reasoning and the uncertainty both still visible.




SECURITY OPERATIONS ANALYST

Work one prioritised alert queue, not five consoles

Every observation is compared against active indicators and enabled rules before the save completes. Each match raises its own alert naming exactly what fired, and pivots to the entity, indicator or case.

- Triage starts from provenance rather than a hunch.

Casework and command




INVESTIGATOR

Run an inquiry as a living case file

Every note, action and piece of evidence goes in as a timestamped case entry attributed to whoever added it, so chain of custody is a property of how the file is built rather than a discipline someone maintains.

- An investigation that is auditable, defensible and genuinely handover-ready.



COMMANDER

Give command one shared, current picture

The operating picture fuses the map, the entities that matter now, active alerts and mission status into one view. It refreshes when opened and on an interval, which is why the honest description is near-real-time.

- A distributed team stays aligned because there is only one picture to be aligned to.

06 Built AI-native

Linked as you write it, so an answer can follow the trail

A chat box bolted onto disconnected tables can only answer table-shaped questions. Eleven kinds of record and ten named relationships between them are written in the same moment as the record itself, so the corpus is a traversable graph from the first save.

Relationships mean something in the domain. An indicator indicates an entity. An alert concerns an entity, was raised by a rule and is part of a case. Traversals are capped by design, because on a dense adversary network an unbounded trace is an outage.


The second commitment is provenance over scores. A false positive here is a matched string an analyst can read and argue with, not a model score.



06 Four design commitments


Linked the moment you record it

Sources, observations, entities, indicators, alerts, cases, reports, missions, tasking and rules connect as you save them. Nothing waits for a nightly rebuild.




One front door for people and agents

Screens, the assistant and your own external agents ask through the same interface. An agent inherits the operator's rights and not one action more.




Provenance, not opaque scores

A machine-raised alert names the exact indicator or rule that fired, and the rule's owner is notified — so a noisy rule shows up as a rule problem.



Classification travels with the record

Work stays inside your workspace, and the activity trail hides sensitive values before it stores them.



06 What the AI does today

BUILT

- Assistant on all twelve product screens, answering from your live workspace as you
- Screen context resolved before your question is sent — this case, not an identifier
- Detection on every recorded observation, one alert per match
- Typed graph projection on write, and search over the live picture
- Credits checked before each turn, consumption and model reported after

PART-BUILT

- Bounded multi-hop traversal is callable; the visual link chart is not built
- Case-grounded drafting returns text; draft-into-report is planned
- Thirty-six operations for your own AI client; the confirm-before-destroy gate is off
- Rule catalogue runs behind the interface; the screens are the missing half
- Six durable events publish; the consumer that would drive a loop does not exist

PLANNED

- Search by meaning — there are no embeddings over the corpus today
- Proposed threat scores and near-duplicate entity merging
- A ranked triage queue and requirement-to-source matching
- Tradecraft coaching: classification discipline and source weighting
- An unattended watch loop, once the ledger, gate and spend cap exist

06 What keeps it in bounds

Agents get your permissions, nothing more

Every request is checked at the same enforcement point the screens go through. An agent is a faster front door, never a privileged one.

No privileged credential behind the agent surface

The agent server holds no credential of its own. The operator authenticates with their own identity.

Release authority stays with a person

Drafting a product can be delegated; publishing it cannot. Release and recall are separate, gated human actions.

Classification is set, never inferred

No model proposes, guesses or writes a marking, because a marking that arrived from a machine is one nobody can defend.

Proposals attach, they never overwrite

Scores and merges will arrive with the features that produced them visible, and be rejectable without ceremony.

Unattended runs stay off until the guardrails exist

The scheduled loop is blocked on a run ledger, a confirmation gate and a hard spend cap, by choice.

07 Who does what

Six roles seeded into every workspace. What each deliberately cannot do is the access control.

Intelligence director

Programme ownership, final authority on release

Owns the programme end to end. The only role with administrative reach across both halves, and therefore the only one that can retire a source or delete a detection rule.

Collection manager

The collection register and the tasking loop

Registers sources and feeds, owns requirements, issues and closes tasking. Holds nothing on entities, indicators, alerts, cases or reports — the fusion half is outside the remit by design.

All-source analyst

Fusion, casework and finished reporting

Resolves reporting onto entities, grades source reliability, raises requirements, works cases and drafts products. Holds no administrative level anywhere, so nothing it touches can be deleted.

07 Who does what

Detection engineering, the duty watch, and the read-only governance role.

Threat analyst

Detection engineering and the alert queue

Owns detection rules and indicators, triages alerts, runs cyber cases. Given no report permission at all, because detection engineering should not carry release authority.

Watch officer

The live picture during a shift

Works the live picture and escalates, but holds no source, indicator, case or report permission — escalating an alert into a case is enforced as an action on the alert.

Governance auditor

Evidencing how the process ran


Read-only and confined to the planning and tasking record. It can evidence what was asked, tasked and in force — but not read the sensitive content underneath.

08 Under the hood

IntelWatchtower is interface-first. Everything the screens can do is available over a single governed interface, and every client — web, assistant, external agent, partner integration — goes through the same door under the same permissions.

Product state is one domain in two halves that share a data model: collection, covering sources, requirements, tasking, feeds and observations; and fusion, covering missions, entities, indicators, alerts, cases and reports.

Identity, billing, quotas, files, notifications and activity history are inherited from the underlying platform rather than rebuilt, so a workspace is tenant-isolated from the moment it is created.



THE SPLIT THAT MATTERS

Authorisation is enforced at the edge, never inside the product. The gateway checks every query and mutation before it reaches the service, and a missing permission fails closed — exactly one place to audit, and exactly one place a mistake can be made.

08 Standards and integrations

OPEN STANDARDS IT SPEAKS

- OpenID Connect single sign-on for enterprise identity
- OAuth 2.0, including device-code flow for headless clients
- SAML and SCIM provisioning on the enterprise tier
- Structured threat-exchange formats for indicator ingest
- System-log ingest for network and endpoint telemetry
- Keyhole markup and coordinate formats for geospatial layers
- Adversary technique mappings on threat indicators
- Model Context Protocol for external AI agent clients

PLUGS INTO

- Live data feeds on any registered source, with visible ingest health
- Webhook endpoints for outbound delivery into your own systems
- Your existing identity provider, enforced on the enterprise tier
- Your own AI client over the open agent protocol, signed in as the operator
- Security-log platforms via export on enterprise, rather than rip-and-replace
- Shared file storage for documents and imagery, referenced not duplicated
- In-app, email, push and webhook notifications with per-user preferences
- Planned: command-line client, typed Node and Python libraries, mobile and extensions

08 How it deploys

Multi-tenant cloud

The default. Strict per-workspace isolation and the fastest path to a working picture — hours rather than the six-to-eighteen-month cycles legacy platforms need.

Configurable hosting region

Choose the jurisdiction your data sits in. Customer-selected region is an enterprise capability; other tiers run in the shared default region.

Dedicated tenant

Enterprise. Your own isolated deployment, with customer-managed keys, private networking and address allow-lists alongside it.

Self-hosted

Enterprise. Run the platform inside your own environment where the sensitivity of the work or a procurement rule requires it.

Air-gapped

On the enterprise roadmap and not available today. Listed because buyers ask, and because saying planned is more useful than saying nothing.

QUOTAS

Seats, sources, feeds, observations, entities, cases, reports and interface calls are metered at the action layer — warned at eighty per cent, blocked at the limit, never billed as silent overage.

09 Security and trust

Classification on every record

Four levels sit as a first-class field on all twelve record types, set explicitly by a person and never inferred.

Delete is administrator-only, everywhere

No editor-level role can destroy a record, on all twelve types — so future cleanup arrives governed rather than ungoverned.

Sensitive values hidden before storage

The audit trail proves what happened without becoming a second, less protected copy of the sensitive content.

Access checked before the request lands

The gateway checks resource, action and scope before the request reaches the product. A missing permission fails closed.

Tenant isolation by construction

Every write is scoped to the calling workspace and count-checked, with row-level isolation behind it at the storage layer.

Encryption in transit and at rest

Strict transport security on public endpoints, AES-256 at rest, and application-layer encryption for keys and tenant secrets.

Sensitive records withhold viewer reads

Eight of the twelve record types refuse read at viewer level, because of the codenames, coordinates and payloads they carry.

Attributable, exportable activity trail

Who acted, on whose behalf, against what and when — searchable and exportable, with retention that scales by plan.

Operator access is gated and time-bound

Mandatory multi-factor, no shared accounts, just-in-time elevation, separation of duties, access revoked within twenty-four hours of departure.

STATED PLAINLY

IntelWatchtower is pre-launch and claims no certification it has not earned

ISO 27001 and SOC 2 Type II are on the roadmap and not in force: readiness assessment is targeted for late 2026, a Type II report for 2027.

Classification markings are customer-configured workspace metadata used for in-product handling and access control. They do not imply accreditation of the hosting environment, personnel or facilities.

There is no FedRAMP authorisation and none in progress, no Department of Defense Impact Level authorisation at any level, and no CMMC assessment.

Clearance-based need-to-know compartmentation and releasability controls are on the enterprise roadmap and are not available today.

One product, many front doors

AVAILABLE

Governed application interface

Federated, permission-checked on every field, callable by people, the assistant and external agents under identical rules.

AVAILABLE

Analyst assistant

On every product screen, aware of the record you have open, answering with your own identity and metered against workspace credits.

IN PROGRESS

Web experience

The full platform in the browser across twelve screens, from collection and tasking through to missions and reporting.

IN PROGRESS

Composed product experience

Intelligence alongside the other products a team already uses, under one identity and one set of permissions.

PART-BUILT

Agent protocol server

Thirty-six operations, four resources and two guided prompts, authenticated as the operator. The confirm-before-destroy gate is not switched on.

IN PROGRESS

Documentation

Product guides, interface reference and integration tutorials, published openly as the product moves toward launch.

PLANNED

Command-line client · Typed Node and Python libraries · Mobile monitoring and approvals · Browser and editor extensions · Watchfloor and desktop console · Community hub

11 Plans

Priced on analyst seats. Tiers differ by capacity — not by hiding functionality — so every tier carries the full feature surface.

FREE \$0 forever A single analyst proving the platform on real work. No card, no expiry. 3 analyst seats · 15 sources · 2 feeds · 1,000 observations a month · 100 entities · 5 cases	STARTER \$99/mo A cell or security operations team of up to about ten analysts. 10 analyst seats · 50 sources · 10 feeds · 50,000 observations · 2,500 entities · 50 cases	GROWTH \$399/mo A scaling team of up to about forty analysts, or a busy operations centre. 40 analyst seats · 250 sources · 50 feeds · 1M observations · 25,000 entities · 250 cases	BUSINESS \$999/mo An established function of up to about a hundred and fifty analysts. 150 analyst seats · 1,000 sources · 250 feeds · 25M observations · 1 TB · 99.9% service level	ENTERPRISE From \$12/seat/mo Agencies and multi-organisation groups, contracted annually. Enforced sign-on · residency · customer-managed keys · dedicated or self-hosted · 99.95%+
--	---	---	---	--

No silent overage

Warned at eighty per cent, blocked at a hundred, then offered an upgrade or an add-on. Data above quota becomes read-only rather than deleted.

Regional pricing, not currency conversion

Four public bands reflect purchasing power. Band D covers India — Starter ₹3,999 a month, Growth ₹14,999, Business ₹39,999.

Why teams choose it

The collection loop is a first-class workflow

Requirement, tasking, source, feed, observation and satisfaction are records with lifecycles, not a spreadsheet beside the real system.

Cyber lives inside the intelligence picture

A cyber signal can corroborate reporting from a completely different discipline, without changing tools.

All-source, not cyber-only or open-source-only

Seven disciplines fuse into one entity and indicator picture. Most of the field is excellent at indicators and blind to everything else.

Governance you cannot forget to apply

Access checked on every read and write, delete administrator-only on all twelve record types, and documented omissions in the role catalogue.

Classification and confidence are schema

A handling marking and an uncertainty grade sit on every record as queryable fields, not conventions in a document.

Published pricing that starts at zero

Against a category where full-cycle alternatives commonly start in the high six or seven figures with a multi-year procurement attached.

12 Honest answers to hard questions

You are pre-launch. Why would we bet on you?

You should weight that heavily. No customers, no deployments, no revenue, no certifications. What exists is a working collection and fusion system of record. The sensible move is the free tier on a real workflow, not a signature.

Is the operating picture actually real-time?

No. It refreshes when opened and on an interval rather than streaming continuously, so near-real-time is the accurate description. The durable event stream already publishes; the consumer side is the missing half.

Where is the visual link chart?

Not built. Relationships exist as records and the bounded traversal is callable today, but the interactive network canvas is on the roadmap. If visual charting is the whole job, buy the specialist.

Can we put genuinely classified material in this?

The markings are workspace metadata for handling and access control, not a government accreditation. Whether real classified material may lawfully be entered is a determination you must make.

How locked in are we?

Everything the screens do is available over the same interface, exports run to structured formats, and files are referenced rather than swallowed. No proprietary ontology you must adopt.

12 Where it sits against the alternatives

VS PROGRAMME-SCALE DEFENCE PLATFORMS

The same breadth of full-cycle fusion with a first-class tasking loop, on tiers that start free. They lead decisively on proven deployments at the highest classification levels and on agentic workflows at scale.

VS CYBER THREAT-INTELLIGENCE PLATFORMS

Indicators and triage sit inside a multi-discipline picture with a mission container and a tasked loop. Their curated data corpora, analyst teams and enrichment scale are far beyond anything here at launch.

VS LINK-ANALYSIS WORKBENCHES

Cloud-native and live, with feeds, an operating picture, native casework and reporting rather than an import-driven desktop session. The established workbench remains the benchmark for visual link analysis; ours is not built.

VS OPEN-SOURCE INTELLIGENCE PLATFORMS

Open-source is one of seven disciplines feeding a full cycle rather than the entire scope. Their multilingual collection breadth, text analytics and geolocated discovery are a toolset we do not have.

VS ASSEMBLING IT YOURSELF FROM POINT TOOLS

One platform, one bill, one identity, one audit trail and one classification-aware model, against a stack that commonly totals a six-figure annual spend before integration and training. A stack of specialists will beat this feature-for-feature in each individual discipline — and if one discipline is the whole job, that is the better buy.

START TODAY

Prove it on a real workflow before you commit to anything

IntelWatchtower launches in 2026 and is in waitlist mode today. The fastest way to judge it is to run one collection loop end to end.

- Join the waitlist and take the free tier — three analyst seats, fifteen sources, no card and no expiry
- Run one loop: raise a requirement, task a source, fuse an entity, triage an alert, open a case, release a product
- Bring your own AI client and check that it inherits exactly the permissions you gave the operator
- Then talk to us about capacity, governance depth and the deployment shape your jurisdiction needs

JOIN THE WAITLIST

intelwatchtower.com

BOOK A WALKTHROUGH

intelwatchtower.com/contact

PRICING IN FULL

intelwatchtower.com/pricing

TALK TO SALES

sales@intelwatchtower.com



IntelWatchtower is a product of Algoshred Technologies Private Limited. Pre-launch: no customers, deployments, revenue or certifications are claimed. Capabilities marked planned or roadmap are not available today.