



SALES & PARTNER PLAYBOOK

IntelWatchtower

IntelWatchtower gives intelligence and security teams one accountable picture, from collection requirement to released report.

All-source fusion

Collection tasking

Cyber triage

Case files

Classified by default

PREPARED BY

Burdenoff Consultancy Services Pvt.
Ltd.

PRODUCT

intelwatchtower.com

BOOK A CALL

calendly.com/vignesh-burdenoff/1-1



00 **How to use this playbook**

Start with the opportunity signals and the one-question test; between them they will tell you inside a single conversation whether this is worth your time. Use the competitor table and the objection handling when a prospect names an incumbent, because both are written to be said out loud rather than read. When you have something live, work through the handoff section and send us the five things it lists. Everything else here is background you can dip into when a specific question comes up.

IN A HURRY

Read sections 01–03 and 18. That is enough to recognise an opportunity and open a conversation.

BEFORE A MEETING

Add 06–08, 10 and 16 — who buys, what they buy, what it costs, and what they will push back on.

HANDING OVER A LEAD

Sections 19–22 tell you exactly what to send, to whom, and what you earn for it.

CONTENTS

01	Product Overview	13	Customer Proof & Credibility
02	One-Line Pitch	14	Demo Information
03	30-Second Elevator Pitch	15	Sales Talking Points
04	Problem We Solve	16	Objection Handling
05	Value Proposition	17	Discovery / Qualification Questions
06	Ideal Customer Profile	18	How to Identify a Good Opportunity
07	Buyer Personas	19	How to Introduce Us
08	Top Use Cases	20	Lead Handoff Process
09	Product Capabilities	21	Partner / Referral Commercials
10	Pricing & Plans	22	Contact Information
11	Competitive Landscape	23	Useful Links & Sales Assets
12	Why Us / Differentiators	24	Product Mockups — Complete Visual Appendix

01 Product Overview

INTELWATCHTOWER.COM

IntelWatchtower is a platform for teams whose job is to work out what is happening and tell someone who has to act on it. It handles the collection side (registering sources and sensors, stating what you need to know, tasking sources against those needs, and tracking the feeds that deliver reporting) and the analysis side (resolving that reporting onto tracked people, places, organisations and events, maintaining threat indicators, triaging alerts, running investigations as case files, framing operations as missions, and moving finished reports through review and release). All of it sits in one data model, so a question raised on Monday and a report released on Friday are the same thread rather than five systems and a lot of copying. Every record carries a classification label and a confidence grade, and every action is attributable.

CATEGORY

All-source intelligence platform, for intelligence and fusion cells, security operations and threat-intelligence teams, investigation units, and operations centres that need one current picture.

WHERE IT STANDS TODAY

Generally available and early. The collection and fusion core works in the product today: sources, requirements, tasking orders, feeds, observations, entities, indicators, alerts, cases, missions and reports, with the shared overview dashboard, classification labelling, role-based access control and audit in place. The plotted map view, interactive link charts, an interface for writing detection rules, and most of the AI analysis beyond the page-aware assistant are still being built. There are no public reference customers yet and no completed security certifications.

HOW CUSTOMERS GET IT

Cloud-hosted and multi-tenant, with every record scoped to a workspace. A team can sign up on the free tier and start without talking to anyone. Enterprise agreements add enforced single sign-on and directory provisioning, longer audit retention, security-event streaming into the customer's own tooling, a choice of hosting region, and dedicated-tenant, self-hosted or air-gapped deployment. A small cell is productive in hours; migrating historical cases and reporting off existing tooling is a paid service with provenance and chain of custody preserved.

Web application

Desktop application for Windows, macOS and Linux

Command-line client

MCP server for AI agents

Node.js SDK

Python SDK

Browser extension

Editor extension

Mobile application (planned)

02 One-Line Pitch

SAY EXACTLY THIS

"IntelWatchtower gives intelligence and security teams one accountable picture, from collection requirement to released report."

SHORTER, FOR A CHAT MESSAGE

One accountable intelligence picture, from tasked collection to released report.

THE CATEGORY SENTENCE

IntelWatchtower is an all-source intelligence platform for analyst cells, security operations centres, investigation units and mission teams.

03 30-Second Elevator Pitch

SPOKEN, ROUGHLY 30 SECONDS

"Most intelligence teams work across five or six tools: a threat feed, a map, a spreadsheet, a case system, a chat channel. They glue it together by hand, and by the time a briefing reaches the person deciding, it is stale and nobody can say where it came from. IntelWatchtower puts the whole cycle in one place. You raise a collection requirement, task a source against it, fuse what comes back onto tracked entities, and carry it through to a released report. Every record holds a classification, a confidence level and a source trail."

IF THEY ARE TECHNICAL

IntelWatchtower is API-first: everything the web application does is available through one API, and every operation carries an access-control rule that is checked at the gateway before the request reaches the service, so the interface is never the enforcement point. Collection objects (sources, requirements, tasking orders, feeds, observations) and fusion objects (entities, indicators, alerts, cases, missions, reports) share one data model, so pivoting from an alert to the entity it implicates to the case it belongs to is a query rather than an export. Classification, confidence, and created-by, updated-by and on-behalf-of provenance are fields on every record.

IF THEY ARE A BUSINESS BUYER

Teams doing this work usually buy a threat-intelligence subscription, a mapping tool, a case system and a reporting workflow separately, then pay people to move data between them. IntelWatchtower covers that ground on one bill, one identity and one audit trail, starting at 99 US dollars a month rather than a six-figure programme. The questions an oversight or audit review asks (who saw what, how confident were we, where did it come from) are answered from the system instead of reconstructed from memory.

04 Problem We Solve

These are the complaints you will actually hear when you ask an intelligence, security or investigation team how their week went. Each one has a specific answer in the product; none of them requires you to know the tradecraft to recognise.

The picture lives in six tools and one analyst's head

Who feels it: All-source and fusion analysts, threat-intelligence leads

What it costs them: A meaningful share of an analyst's day goes on exporting, re-keying and reconciling instead of thinking, and the assembled picture only exists on the machine of whoever assembled it.

How IntelWatchtower answers it: Collection and analysis objects share one data model, so reporting from open sources, geospatial work, cyber telemetry and partner material resolves onto the same tracked entities. The picture is a record in a shared system rather than a document someone maintains.

Provenance evaporates between the source and the briefing

Who feels it: Analysts, and the commanders and executives who act on their work

What it costs them: When someone asks where a claim came from, the honest answer takes days to reconstruct and sometimes cannot be reconstructed at all.

How IntelWatchtower answers it: Observations stay attached to the entities they support, and every record carries who created it, who last changed it and on whose behalf. The trail from a released report back to the source that produced it is a query, not an archaeology project.

Nobody knows how confident to be

Who feels it: Decision-makers reading finished intelligence

What it costs them: Assessments get read as either certainty or noise, and conflicting reporting is quietly averaged away rather than surfaced for a person to resolve.

How IntelWatchtower answers it: Confidence is a graded field on every entity and indicator (low, moderate, high, confirmed) and sources carry a reliability score out of 100. Uncertainty is recorded rather than implied, and contradictions stay visible.

Collection drifts away from what was actually asked

Who feels it: Collection and tasking managers, intelligence directors

What it costs them: Requirements are raised in meetings, tracked in a spreadsheet and quietly forgotten. Nobody can say which source was tasked against which question or whether the question was ever answered.

How IntelWatchtower answers it: Priority intelligence requirements are first-class records with a priority band, a due date and a lifecycle. Tasking orders are issued against named sources with instructions and a collection window, and the requirement moves to satisfied with the answer attached.

Cyber detection sits in a console that knows nothing else

Who feels it: Security operations analysts and threat hunters

What it costs them: Every alert that matters gets its context rebuilt by hand, and the link between a technical detection and the wider picture depends on someone remembering to make it.

How IntelWatchtower answers it: Indicators, alerts, entities and cases share one model. An alert carries a severity and a controlled triage lifecycle, links to the entity or indicator it implicates, and escalates into an investigation case without leaving the platform.

The operating picture is a deck built last night

Who feels it: Watch officers, operations centres, mission commanders

What it costs them: By the morning brief the picture describes yesterday, and half the meeting is spent arguing about whose numbers are current.

How IntelWatchtower answers it: One shared overview draws mission status, the alert queue by severity, sources by discipline and the top triage items straight from the records analysts are working in, and refreshes on open and on an interval.

Classification and audit arrive after the leak

Who feels it: Security and governance leads, compliance and oversight functions

What it costs them: Handling rules live in a policy document and are enforced by habit, so proving who saw a sensitive record means asking around.

How IntelWatchtower answers it: Every record carries a classification label as a schema field, access control is checked on every operation before the request reaches the service, and the activity timeline is searchable and exportable. A read-only governance auditor role can evidence how the process ran without being able to change anything.

05 Value Proposition

Six things the product does that a prospect will care about within the first conversation. Each is backed by something they can check themselves rather than by a claim we are asking them to accept.

The whole cycle in one thread

A requirement becomes a tasking order, which produces observations, which resolve onto entities, which raise alerts, which escalate into cases, which roll up to missions and out as reports. Nothing is re-keyed at any boundary. That continuity is the product; the individual features are how it is delivered.

Evidence: Requirements, tasking orders, satisfaction transitions, entities, alerts, cases, missions and report release are all live objects in the product today, each with its own working screen.

Uncertainty is on the record

Confidence grading on entities and indicators, and a reliability score on every source, mean an assessment states how much to trust it. Conflicting reporting is surfaced for a person to resolve rather than smoothed over. Decision-makers stop guessing at the strength of what they are reading.

Evidence: Confidence and reliability are fields on the records themselves, visible in the interface and returned by the API, not free text in a note.

Governance built into the data, not the policy

Classification labels sit on every source, requirement, task, entity, indicator, alert, case, mission and report. Access control is enforced at the gateway on every operation and fails closed when a permission is missing. Workspace isolation is enforced on every write, so an identifier from another organisation cannot be read or changed.

Evidence: Six roles are seeded into every workspace, including a governance auditor that can evidence the process end to end without write access anywhere.

Cyber inside the picture, not beside it

Indicators of compromise carry confidence, classification and a validity window, and attach to the entity they implicate. Alerts move through open, triaged, escalated, resolved, false positive and suppressed with an assignee and triage notes, and escalate straight into a case. Security work participates in the intelligence picture instead of running parallel to it.

Evidence: Alert triage and the escalation of an alert into an investigation case are shipped and exercisable in the product.

A price a team can actually approve

There is a free tier with real quotas, a Starter plan at 99 US dollars a month and published capacity limits at every level. Tiers differ by capacity rather than by locking features away, so a small cell gets the same platform surface as a large one. Regional pricing means the billed price reflects the customer's market rather than a currency conversion.

Evidence: Plan prices, per-plan quotas and add-on prices are published on the website and read from the live billing catalogue.

An enterprise backbone that came free

Single sign-on, fine-grained roles, multi-tenancy, billing, file storage, comment threads, tagging, notifications, translation and accessibility are shared platform services rather than things this product had to build. That is why a young product ships with governance a much older one would still be retrofitting.

Evidence: Comments, files, tags, notifications and audit appear on intelligence records because they are the same platform services other products in the family use.

06 Ideal Customer Profile

A team of roughly five to two hundred analysts or operators who work across more than one kind of source, have to show their working to somebody, and have outgrown a spreadsheet-and-chat workflow.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
Intelligence and fusion cells in defence, homeland security and government	5 to 200 analysts and collectors, inside a larger organisation	They need the full cycle and the governance, and most of them cannot reach the budget or the procurement timeline that a programme-scale platform demands. A cell inside a large agency is a realistic entry point where the whole agency is not.
Enterprise security operations and threat-intelligence teams	20 to 500 security staff across operations, threat intelligence and governance	They already have detection tooling but nothing that fuses cyber with open-source, geospatial and physical-security signals or frames it for an executive audience. Tool sprawl and the cost of integration between them is a live budget conversation.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
Law enforcement and corporate investigation units	5 to 50 investigators and analysts	Chain of custody and defensible case files are the whole job. Cases with timestamped, attributable entries linked to the entities and alerts that produced them replace scattered folders and screenshots.
Risk, due-diligence and corporate intelligence functions	3 to 40 analysts, often inside a bank, insurer or large corporate	They fuse open-source research with internal records and partner reporting, and they are usually asked to show provenance for anything that lands in a decision. Today most of them do it in documents.
Conflict monitors, NGOs and research organisations	3 to 50 staff, budget-constrained	They need a live picture and a shared workspace without classified infrastructure. Regional pricing and a free tier make the product reachable where a defence platform is not.

GEOGRAPHIES

Sold globally with regional pricing across four bands, including native Indian rupee pricing. Hosting region is configurable for customers with jurisdiction or data-residency requirements, agreed per engagement. No United States government accreditation is held, so federal programmes requiring one are out of scope today.

NOT A FIT – DO NOT CHASE

- A buyer whose entire problem is one discipline and who wants the best specialist tool for it. A pure open-source discovery buyer, a pure threat-feed buyer or a pure mapping buyer will be better served by a specialist, and we lose those fairly.
- Anyone whose procurement requires FedRAMP, a Department of Defense impact level, CMMC or ITAR authorisation. None of these are held or in progress, and the classification labels in the product are workspace metadata the customer configures, not an accreditation.
- Very high-volume log ingestion as the core requirement. That is a job for a security information and event management platform; IntelWatchtower integrates with one rather than replacing it.
- A team that wants best-in-class visual link charting above everything else. The underlying relationships exist, the interactive chart canvas does not yet, and a dedicated link-analysis workbench will beat us on that today.
- A single individual doing casual research. The product assumes a team, roles and a governance question; a lone researcher will find it heavier than the job needs.
- A buyer who requires a long list of named reference customers before they will run a pilot. That does not exist yet, and no amount of conversation will produce it.

07 Buyer Personas

Five roles decide or block this purchase, and they care about visibly different things. The hooks below are written to be said as they are, in the first thirty seconds of a conversation.

Intelligence director or head of intelligence

WHAT THEY WANT

Deliver a current, defensible picture to whoever is deciding, shorten the distance between a question and an answer, and be able to stand behind the process when it is reviewed.

Open with: If your director asked tomorrow which of this quarter's intelligence requirements were actually answered, and by which source, how long would it take you to tell them?

They will ask: Does this cover the whole cycle, or is it another tool my team has to feed?

WHAT FRUSTRATES THEM

Analyst time going on data wrangling rather than analysis, briefings that are stale on arrival, and no way to prove that handling rules were followed other than trusting people.

Security operations or threat-intelligence lead

WHAT THEY WANT

Cut the time to work an alert, reduce the number of consoles an analyst touches, and give the executive audience something more useful than an indicator count.

Open with: When one of your analysts picks up an alert, how much of the next twenty minutes is them rebuilding context that already exists somewhere else?

They will ask: How does this sit next to the detection tooling we already pay for, and what do I switch off?

WHAT FRUSTRATES THEM

Threat intelligence sits apart from case management, which sits apart from reporting. Analysts lose shift time to context switching, and board reporting is compiled by hand.

Collection or tasking manager

WHAT THEY WANT

Turn stated intelligence gaps into tasked collection against sources that can actually satisfy them, and close the loop on what came back.

Open with: Where do your intelligence requirements live at the moment, and what tells you when one of them has been satisfied?

They will ask: Can I see, in one place, which requirements are open, which are tasked, and which source answered them?

WHAT FRUSTRATES THEM

Requirements live in a spreadsheet, source reliability is folklore held by whoever has been there longest, and a feed can go quiet for a fortnight before anyone notices.

Investigator or case lead

WHAT THEY WANT

Build a case that holds up under challenge, with every piece of evidence carrying a time, a source and an owner.

Open with: If counsel asked where a specific piece of evidence in your current case came from, could you answer inside a minute?

They will ask: Is the case record defensible, and can I control who sees which parts of it?

WHAT FRUSTRATES THEM

Case material spread across email, shared drives and screenshots, chain of custody maintained by discipline rather than by the system, and no controlled way to work a case with another agency.

Security, governance or compliance lead

WHAT THEY WANT

Prove that classification and access rules were enforced, that only the right people saw sensitive records, and that the evidence for all of it can be exported.

Open with: When your next oversight review asks who accessed a sensitive record, where does that answer come from today?

They will ask: What certifications do you hold, where does the data sit, and can I get a read-only auditor account?

WHAT FRUSTRATES THEM

Governance lives in a policy memo and gets tested for the first time during a review. Reconstructing who saw what is manual and slow.

Chief information security officer or economic buyer

WHAT THEY WANT

Consolidate overlapping tools, reduce integration debt, and get defensible numbers in front of a board without a week of preparation.


Open with: How many separate subscriptions are involved in producing the picture you take to your board, and who assembles it?

They will ask: What is the total cost against what we run today, and what is the risk of betting on a young vendor?

WHAT FRUSTRATES THEM

Several overlapping subscriptions, a growing integration bill, and reporting that is manually assembled and therefore always slightly out of date.

Twelve situations the product is built for. Each is written as the bind the team is in first, then what the product does about it, so you can recognise the situation in a conversation without knowing the tradecraft.



COLLECTION AND TASKING MANAGER

Close the loop from intelligence gap to answer


THE SITUATION

Gaps get raised in a meeting and tracked in a spreadsheet. Nobody can say which source was tasked against which question, or whether the question was ever answered.

WHAT HAPPENS WITH INTELWATCHTOWER

A priority intelligence requirement is raised with a priority band running from routine to flash, a due date and an owning mission. Tasking orders go out against named sources with instructions and a collection window, and the requirement moves through submitted, approved and tasked to satisfied as reporting arrives.

Every intelligence question has a visible path from asked to answered, with the source that answered it on the record.



COLLECTION MANAGER, INTELLIGENCE DIRECTOR

Know which sources you can actually rely on


THE SITUATION

Sources are scattered across systems and their reliability is institutional memory. The person who knows which feed is dependable is the one who has been there longest.

WHAT HAPPENS WITH INTELWATCHTOWER

Every source is registered with its discipline (human, signals, geospatial, open-source, imagery, measurement or cyber), an operational status, a region, an operator, a reliability score out of 100 and the time it was last seen.

Tasking decisions are made against a recorded reliability score rather than a corridor conversation.



WATCH OFFICER, COLLECTION MANAGER

Catch a dead feed before it costs you a fortnight


THE SITUATION

A feed quietly stops delivering. Nobody notices until an analyst asks why a region has gone silent, by which time two weeks of reporting are missing.

WHAT HAPPENS WITH INTELWATCHTOWER

Feeds attached to a source carry an ingest state (connected, stale, error, disconnected), the time of the last ingest and a running count, all shown on one health view.

Silent collection failures surface in minutes instead of at the next review.



ALL-SOURCE OR FUSION ANALYST

Fuse many partial reports into one assessment


THE SITUATION

Open-source reporting, geospatial work, cyber telemetry and partner material each describe part of the same subject, and fusing them means copying between tools until the provenance is gone.

WHAT HAPPENS WITH INTELWATCHTOWER

Reporting resolves onto tracked entities (people, organisations, vessels, aircraft, vehicles, facilities, locations, network nodes, events), each carrying aliases, attributes, a confidence grade, a classification, coordinates and a threat score, with the observations behind them still attached.

One record per subject, with an explicit confidence level and a trail back to the reporting that produced it.



THREAT ANALYST, DETECTION ENGINEER

Keep threat indicators current and attached to something


THE SITUATION

Indicators pile up in a list with no expiry and no link to what they implicate, so the list gets less useful the longer it exists.

WHAT HAPPENS WITH INTELWATCHTOWER

Indicators (addresses, domains, links, file hashes, email addresses, coordinates, signatures, behaviours) carry a confidence grade, a classification, a validity window and an active flag, and attach to the entity they implicate.

An indicator set that expires on its own terms and always points at something in the wider picture.



SECURITY OPERATIONS ANALYST

Work the alert queue without leaving the picture


THE SITUATION

Detections arrive in a console that knows nothing about the rest of the intelligence picture, so context is rebuilt by hand for every alert that matters.

WHAT HAPPENS WITH INTELWATCHTOWER

Alerts carry a severity from informational to critical and move through a controlled lifecycle (open, triaged, escalated, resolved, false positive, suppressed) with an assignee and triage notes, and each links to the entity or indicator it implicates.

A prioritised, attributable queue where every decision taken on an alert is recorded.



INVESTIGATOR, CASE LEAD

Run an investigation as one accountable file


THE SITUATION

Case material is spread across email, shared drives and screenshots. When somebody asks where a piece of evidence came from, reconstructing the answer takes days.

WHAT HAPPENS WITH INTELWATCHTOWER

A case carries a reference, a summary, a lead analyst, a status, a priority and a classification, and every note, action and piece of evidence is added as a timestamped entry attributed to its author. Alerts and entities attach to the case, and the case rolls up to the mission it supports.

Chain of custody by construction rather than by discipline.



MISSION COMMANDER, OPERATIONS LEAD

Frame an operation and keep everything underneath it


THE SITUATION

Collection, analysis and reporting for a single operation are tracked in different places, so nobody can see the whole effort at once or close it out cleanly.

WHAT HAPPENS WITH INTELWATCHTOWER

A mission holds an objective, a commander, a window and a status that moves from planning through ready and active to completed, and owns the cases and reports produced under it.

An operation can be run, reviewed and closed as one unit with its output attached.



INTELLIGENCE DIRECTOR, REPORTING LEAD

Release finished intelligence under control


THE SITUATION

Reports go out as attachments. Once one is wrong there is no way to withdraw it, and no record of who approved the version people are acting on.

WHAT HAPPENS WITH INTELWATCHTOWER

Situation reports, summaries, assessments, warnings and spot reports move through draft, in review, released, recalled and archived, with every state change recorded against the person who made it.

Finished product that is dated, attributed and withdrawable rather than loose in an inbox.



WATCH OFFICER, OPERATIONS CENTRE, COMMANDER

Give command one current picture


THE SITUATION

The operating picture is a deck someone built last night. By the morning brief it describes yesterday, and half the meeting goes on whose numbers are right.

WHAT HAPPENS WITH INTELWATCHTOWER

One shared overview brings mission status, the live alert queue by severity, sources by discipline and the top triage items together from the same records analysts are working in, refreshing on open and on an interval. The plotted map view is still being built; coordinates are captured on entities today and shown alongside them.

The room argues about the decision rather than about whose picture is current.



ALL-SOURCE ANALYST, INVESTIGATOR

Trace how people, places and events connect


THE SITUATION

The connection that matters spans three reports and two disciplines, and finding it depends on an analyst holding the whole thing in their head.

WHAT HAPPENS WITH INTELWATCHTOWER

An analyst opens an entity and pivots to the indicators and alerts attached to it, and on to the cases and missions those belong to, with the supporting record behind every hop. The interactive chart canvas is on the roadmap; the pivots and the underlying relationships work today.

Connections are traced through records rather than reconstructed from memory.



SECURITY AND GOVERNANCE LEAD, PLATFORM ENGINEER

Point your own AI agent at it, under the same rules

THE SITUATION

The team is wiring AI assistants into its tooling, and the governance question, what can it see and what can it change, has no satisfying answer.

WHAT HAPPENS WITH INTELWATCHTOWER

A page-aware assistant works alongside the analyst and can see the entity, case or alert on screen. Separately, an agent server exposes the platform's operations to AI agent tooling using scoped keys under the same access control a person gets, so an agent cannot reach a record its operator could not open. Confirmation gating on destructive operations is still being finished.

Agents work inside the same permission model as people, with the same audit trail behind them.

09 Product Capabilities

What is built, what is being built, and what is on the roadmap. Anything not marked shipped should be described to a prospect as coming, never as present. A partner who oversells one of these is the reason the next conversation goes badly.

Collection and tasking

Turning a stated intelligence need into tasked collection, and keeping the pipes that deliver reporting visible.

Source and sensor registry

Sources across seven disciplines with status, classification, region, operator, codename and a reliability score out of 100.

Collection requirements

Priority intelligence requirements and essential elements of information, with a priority band from routine to flash, an owning mission, a due date and a full lifecycle.

Tasking orders

Orders issued against a source with instructions, a collection window, dispatch and completion timestamps and a result summary.

Live feeds and ingest health

Feeds attached to a source with ingest state, last-ingest time and a running count, so a broken pipe is visible rather than silent.

Observation capture

Raw reporting recorded as discrete observations, the atomic input that everything downstream fuses from.

Open-source watchlists planned

Keyword, entity and geographic watches across news, social and public records.

Deep external connectors planned

Prebuilt connectors for commercial open-source, threat-intelligence, imagery and sensor platforms.

All-source fusion

Turning many partial reports into one picture that is honest about how much it knows.

Tracked entities

People, organisations, vessels, aircraft, vehicles, facilities, locations, network nodes and events, with aliases, attributes, coordinates, last-seen time and a threat score.

Confidence grading

Low, moderate, high or confirmed on entities and indicators, so an assessment states its own strength.

Source reliability

A recorded reliability score on every source that travels with the reporting it produced.

Provenance to the observation

Fused records keep the reporting behind them attached, so a claim traces back to what produced it.

Entity pivots

Move from an entity to its indicators and alerts, and on to the cases and missions they belong to.

Visual link charts in progress

An interactive network canvas for exploring relationships at scale.

Automated entity resolution planned

Machine-assisted matching and correlation of incoming reporting onto existing entities.

Cyber detection and triage

Security work inside the same picture as everything else, rather than in a parallel console.

Threat indicators

Addresses, domains, links, file hashes, email addresses, coordinates, signatures and behaviours with confidence, classification, a validity window and an active flag.

Alert lifecycle

Severity from informational to critical, moving through open, triaged, escalated, resolved, false positive and suppressed with an assignee and triage notes.

Escalate an alert into a case

Attach an alert to an investigation case in one step, carrying its entity and indicator context with it.

Detection rules raising alerts in progress

Rules evaluated against recorded observations that raise alerts automatically with their provenance attached. The engine exists; the interface for writing and tuning rules does not yet.

High-volume telemetry ingestion planned

Ingest at the scale a dedicated security monitoring platform handles.

Proactive threat hunting planned

Hunt workflows and adversary profiles across the whole data estate.

Investigation, missions and reporting

The containers that make a piece of work accountable from opening to release.

Investigation cases

Reference, summary, lead analyst, status, priority, classification and an owning mission on a living case file.

Chain-of-custody entries

Timestamped case entries attributed to their author, recording evidence, notes and actions as they happen.

Mission operations

An objective, commander, window and status from planning to completed, owning the cases and reports produced under it.

Report release and recall

Situation reports, summaries, assessments, warnings and spot reports moving through draft, review, release, recall and archive.

Assisted report drafting in progress

A first draft assembled from the case file and the live picture, with a person holding release authority. Drafting works; the round trip back into the report record is not finished.

Timeline reconstruction and investigation boards planned

Visual timeline and board surfaces for working a complex case.

The shared operating picture

One current view that analysts, watch officers and commanders all read from.

Shared overview

Mission status, alerts by severity, sources by discipline and the top triage items in one place, drawn from the live records.

Coordinates on entities

Latitude and longitude captured as first-class fields on entities and observations, ready for the map.

Threat scoring on entities

A score carried on the entity so the picture can be ranked by what matters most.

Plotted map with areas of interest in progress

Entities and alerts drawn on a shared map with defined areas of interest and layered context.

Live push updates planned

The picture updating itself as reporting lands, rather than refreshing on open and on an interval.

Geospatial change detection planned

Comparing a location across time to show what appeared, moved or went away.

Governance, access and audit

The part a security leader has to vouch for, enforced in the data model rather than in a policy document.

Classification on every record

Unclassified, confidential, secret or top secret as a schema field on every source, requirement, task, entity, indicator, alert, case, mission and report.

Access control at the edge

Every operation carries an access rule checked before the request reaches the service. A missing permission fails closed.

Workspace isolation

Every write is scoped to the caller's workspace with a count-checked update, so a guessed identifier from another organisation cannot be read or changed.

Provenance and activity trail

Created-by, updated-by and on-behalf-of on every record, with a searchable and exportable activity timeline.

Seeded roles including a read-only auditor

Six roles seeded into every workspace, one of which can evidence how the process ran without being able to change anything.

Need-to-know compartmentation and releasability

planned

Clearance-based compartmentation and controlled cross-organisation sharing with a disclosure record.

Security-event streaming and compliance exports

planned

Streaming the audit trail into the customer's own monitoring stack, with evidence exports for a review.

Platform, API and integrations

How the product connects to identity, to other systems and to whatever the customer builds themselves.

Single sign-on and fine-grained roles

Enterprise identity from day one, with roles scoped to workspace, resource and action.

One API behind everything

Every operation the interface performs is available programmatically, under the same access control.

Page-aware assistant

An assistant that can see the entity, case or alert on screen, so a question about the current work does not have to be re-described.

Inherited collaboration services

Comment threads, file storage, tagging, notifications, translation, accessibility and support, shared with the wider platform.

Command line, SDKs and agent server in progress

A terminal client, Node.js and Python SDKs, browser and editor extensions and an MCP server for AI agents. All are built and wired to the production API; public registry publication is held until general release.

Mobile application planned

Monitoring the picture, taking alerts and giving approvals from a phone.

Marketplace of detection content and analytic packs planned

A place to publish and buy detection rules, connectors and analytic packs.

10 Pricing & Plans

Priced per workspace on flat-fee tiers banded by the number of analyst and operator seats, with Enterprise contracted per seat annually. Tiers differ by capacity rather than by feature gating: every plan gets the whole platform surface, and a team upgrades when seats, sources, feeds, observations, entities, cases or reports run out. Four regional bands mean the billed price reflects the customer's own market rather than a currency conversion. AI-assisted analysis runs on the customer's own AI provider key, so there is no model markup inside the plan price.

PLAN	PRICE	WHO IT IS FOR	WHAT IS INCLUDED
Free	\$0 Free forever, no card required. India: zero rupees.	Evaluation, a single small cell, academic and newsroom work	3 analyst seats and 2 admin seats; 1 workspace; 15 tracked sources and 2 active feeds; 1 active mission; 1,000 observations a month, 100 tracked entities; 5 reports a month, 1 GB shared storage; Community support plus one ticket a month
Starter	\$99/mo or \$79/mo annual India: ₹3,999/mo or ₹3,199/mo annual. Billed monthly or annually.	A small analyst cell or security operations team of up to about ten analysts	10 analyst seats and 5 admin seats; 50 tracked sources and 10 active feeds; 3 active missions; 50,000 observations a month, 2,500 tracked entities; 50 reports a month, 25 GB shared storage; Email support under 48 hours; Data-processing agreement on request
Growth	\$399/mo or \$319/mo annual India: ₹14,999/mo or ₹11,999/mo annual. Includes a 14-day trial.	A scaling team ingesting at volume and running several investigations at once	40 analyst seats and 15 admin seats; 3 workspaces; 250 tracked sources and 50 active feeds; 1,000,000 observations a month, 25,000 tracked entities; 500 reports a month, 100 GB shared storage; Priority support under 24 hours; 99.5% uptime commitment
Business	\$999/mo or \$799/mo annual India: ₹39,999/mo or ₹31,999/mo annual.	An established operation of up to about 150 analysts that needs an availability commitment	150 analyst seats and 50 admin seats; 10 workspaces; 1,000 tracked sources and 250 active feeds; Unlimited active missions on fair use; 25,000,000 observations a month, 250,000 tracked entities; 5,000 reports a month, 1 TB pooled storage; Priority support under 8 hours, 99.9% uptime commitment
Enterprise	From \$12/seat/mo list; \$9/seat/mo annual floor at 250+ seats Annual contract, quoted per engagement. Regional quotes available, including a lower Indian band.	Larger, regulated or multi-agency organisations, and anyone needing a deployment choice	Custom seat, source, feed, observation and entity quotas; Enforced single sign-on and directory provisioning; Longer audit retention and a dedicated support commitment from 99.95%; Configurable hosting region for data residency; Dedicated tenant, self-hosted or air-gapped deployment; Need-to-know compartmentation and security-event streaming are committed but still being built; confirm timing in the quote; Not FedRAMP, Department of Defense impact level, CMMC or ITAR authorised

Add-ons

ADD-ON	PRICE
5-Seat Pack (+5 analyst seats)	\$45/mo
25-Seat Pack (+25 analyst seats)	\$199/mo
50-Source Pack (+50 tracked sources)	\$39/mo
250-Source Pack (+250 tracked sources)	\$149/mo

ADD-ON	PRICE
25-Feed Pack (+25 active feeds)	\$59/mo
1M Observation Pack (+1,000,000 observations a month)	\$49/mo
100-Case Pack (+100 active investigation cases)	\$49/mo
250 GB Storage	\$49/mo
AI Pro (+2,500 AI-assisted analysis operations a day, customer's own key)	\$79/mo
Support Retainer (12 hours a month, up to 6 hours rollover)	\$4,000/mo

WHAT TO TELL A PROSPECT ABOUT PRICE

- All prices are the published list in US dollars before tax. Four regional bands apply and the Indian band is materially lower; quote the customer's own band, never a currency conversion of the dollar price.
- Seats are counted on analysts and operators, the people who collect, fuse, triage, investigate and report. Admin logins are included generously and do not raise the price inside a tier.
- Add-ons are available on Starter, Growth, Business and Enterprise. A free-tier customer has to upgrade before buying one.
- AI-assisted analysis uses the customer's own AI provider key, so their intelligence data stays inside their own provider relationship and there is no inference markup in the plan price.
- The pricing page on the product website reads the live billing catalogue and is the authoritative answer. If anything here and the website disagree, quote the website and tell us.

11 Competitive Landscape

Seven real alternatives, with an honest account of where each of them is better. A comparison in which the competitor never wins is a comparison a prospect stops believing, and the ones below all win somewhere.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Palantir Gotham	The full-cycle intelligence and defence platform: ontology-driven fusion, a strong geospatial picture, deep mission planning, and deployments inside classified environments for large national programmes.	Proven at scale in classified environments, far deeper ontology and mission-planning work, and agentic AI already shipping through their AI platform. If the buyer is a large national programme with the budget, the accreditation requirement and the appetite for a services-led deployment, Gotham is the safer answer and you should say so.	Collection requirements and tasking are first-class records here; Gotham publishes no requirement-to-answer loop. Pricing is published from a free tier and 99 US dollars a month against contracts commonly quoted in the millions. Tenancy is workspace-native rather than project-centric, and cyber indicators, alerts and cases sit in the same object model instead of a separate product.
Recorded Future	The largest commercial cyber threat-intelligence provider, now part of Mastercard. Indexes an enormous external source corpus and delivers scored, curated intelligence with a human analyst team behind it.	Their data corpus, their analyst team and their finished-intelligence AI are far ahead of anything we ship, and their integration catalogue into detection tooling is deeper. If the requirement is genuinely 'give me the best cyber threat feed', they win it.	All-source rather than cyber only: human, signals, geospatial, imagery and measurement sources fuse with cyber into one picture. Intelligence is tasked against stated requirements rather than only consumed. Classification and confidence sit on every record instead of enterprise roles alone, and there is a mission and case structure around the intelligence.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
ThreatConnect, part of Dataminr	A threat-intelligence operations platform acquired by Dataminr in November 2025 and repositioned as an agentic, client-tailored cyber intelligence product with strong case and evidence workflow.	Agentic threat-intelligence operations and real-time public-signal detection at scale are live and mature, and their risk quantification ties threat to financial impact in a way executives respond to. Their security tooling integrations outnumber ours.	Multi-discipline source fusion beyond cyber, an operational mission container they do not have, classification enforced in the data model rather than through standard enterprise roles, and a requirements-to-satisfaction loop for collection.
IBM i2 / i2 Group	The long-standing link-analysis workbench family, used in law enforcement, financial crime and military investigation for more than three decades.	Visual link and network charting is their defining discipline and it is genuinely better than ours, which is still being built. Three decades of court-facing investigative work is a track record we cannot claim. For a team whose day is spent building charts, they are the right tool.	Cloud-delivered with live feeds and a shared picture rather than import-driven desktop analysis. Native collection and tasking, which the workbench does not do. Cyber indicator and alert triage in the same place. Multi-tenant workspaces and inherited identity, billing and audit rather than a per-site deployment.
Anomali	A threat-intelligence platform moving towards an agentic security-operations model, with broad feed aggregation and the ability to match intelligence against historical telemetry.	Their threat-feed marketplace and matching intelligence against historical logs at volume are real strengths, and their security-operations copilot is shipping. On pure cyber operations they are further along.	Everything outside cyber, plus requirement-driven collection, mission framing, chain-of-custody cases and classification as a schema field rather than a role setting.
Google Threat Intelligence	Mandiant's finished intelligence and VirusTotal's telemetry packaged alongside Google's security operations stack.	A very large analyst organisation, incident-derived reporting, VirusTotal telemetry and attack-surface discovery. On cyber breadth and the depth of finished reporting they are ahead of us and will stay ahead for some time.	Multi-discipline fusion rather than cyber only, a tasked collection loop, classification on every record, and no dependency on one cloud provider's security stack to get the value.
Babel Street	An open-source and web intelligence platform focused on multilingual collection, text analytics and geolocated discovery across the open web.	Multilingual open-web collection breadth and text analytics are their specialism and we do not match them. A buyer whose entire problem is open-source discovery should buy the specialist.	The full cycle around the open-source work: requirements, tasking, cases with chain of custody, mission framing and controlled report release, plus classification and confidence on every record and cyber indicators in the same picture.

THE POSITIONING LINE TO HOLD

The breadth of a full-cycle intelligence platform, on published pricing, without a multi-year procurement.

12 Why Us / Differentiators

If a prospect remembers nothing else, these are the seven things that separate IntelWatchtower from the alternatives. The first two are the ones no competitor matches.

The collection loop is a real object, not a spreadsheet

A stated intelligence need, the source that might satisfy it, the order connecting them, the pipe that delivers the reporting and the reporting itself are all records with lifecycles. None of the named competitors publishes this as a core workflow, and it is the thing collection managers react to first.

Classification and confidence are fields, not conventions

Every record carries a handling label and a graded confidence, and sources carry a reliability score. Competitors either apply enterprise roles only, or support record-level classification mainly in bespoke programme deals. Here it is in the schema, on the free tier, on day one.

Cyber shares the model rather than sitting beside it

Indicators, alerts, entities, cases and multi-discipline sources are one object model, so an analyst goes from a detection to the entity it implicates to the case it belongs to without an export. The dedicated threat-intelligence platforms handle indicators well but are cut off from geospatial, human and mission context.

One thread from question to released report

Requirement, tasking order, observation, entity, alert, case, mission, report. Nothing is re-keyed at a boundary and provenance survives every hop. That continuity is what makes the audit answer fast and the briefing traceable.

Workspace-native, with the platform backbone included

Identity, roles, tenancy, billing, files, comments, tags, notifications, translation and accessibility are shared services this product inherited rather than built. That is why an early product arrives with governance a much older one would still be adding.

Published pricing, including a free tier that does real work

Prices, per-plan quotas and add-on prices are all public, so a buyer can model their cost before speaking to anyone, and can run a genuine trial without a purchase order. In this category that is unusual enough to be a differentiator on its own.

Agents and automation run under the human rules

Everything the interface does is available through one API with the same access control, so a script, an integration or an AI agent operates inside the same permission model and leaves the same audit trail as a person. That is the answer to the governance question every security team is now asking about AI.

13 Customer Proof & Credibility

There are no customer references to lean on, so credibility has to come from three places: what a prospect can verify themselves in the product, governance that is enforced in the data model rather than promised in a policy, and being straight about what is not built. Lead with the free tier. Anything a partner asserts, a prospect can check within an hour.

A FREE TIER THAT PROVES THE CLAIM

A prospect can sign up, register a source, raise a requirement, task it, triage an alert and release a report without speaking to anyone. Nothing in this document requires them to take our word for it.

ACCESS CONTROL CHECKED BEFORE THE SERVICE SEES THE REQUEST

Every operation carries an access rule enforced at the gateway. A missing permission fails closed, and the interface is never the enforcement point.

TENANT ISOLATION BY CONSTRUCTION

Every write is scoped to the caller's workspace with a count-checked update, so an identifier guessed from another organisation's data cannot be read or changed.

PROVENANCE ON EVERY RECORD

Created-by, updated-by and on-behalf-of are recorded alongside create and update timestamps, and the activity timeline is searchable and exportable for a review.

SIX ROLES THAT MATCH REAL JOBS

Intelligence director, collection manager, all-source analyst, threat analyst, watch officer and governance auditor are seeded into every workspace, including a read-only auditor that can evidence the process without changing anything.

PUBLISHED PRICING AND QUOTAS

Plan prices, per-plan capacity and add-on prices are published rather than quoted, so a buyer can model the cost before a call and compare it against their existing stack.

CAPABILITY STATUS PUBLISHED HONESTLY

The product website marks each capability as shipped, in progress or roadmap, including the areas where a competitor is ahead. A prospect can hold us to what is written there.

COMPLIANCE POSTURE STATED ON THE PRICING PAGE

The website says in plain terms that the product is not FedRAMP, Department of Defense impact level, CMMC or ITAR authorised, and that classification labels are customer-configured workspace metadata. A buyer finds that out from us, not from their legal team.

Be straight about this. IntelWatchtower is a young product from a small company. It became generally available in 2026, it has no public reference customers, no published deployment stories, and no completed security certifications: SOC 2 Type II and ISO 27001 are on a roadmap with target dates in 2027, and it is explicitly not FedRAMP, Department of Defense impact level, CMMC or ITAR authorised. The classification labels are workspace metadata the customer configures; they are not a government accreditation. The plotted map view, interactive link charts, an interface for authoring detection rules, and most of the AI analysis beyond the page-aware assistant are still being built, and the command-line client, SDKs and agent server are complete but not yet published to public registries. What a prospect should weigh instead is whether the shipped core, the collection loop, the shared data model, classification and audit, fits how they actually work; whether the free tier proves it in a week; and whether direct access to the team that builds it is worth more to them right now than a long reference list. If the answer to that last question is no, walk away rather than push.

14 Demo Information

The most effective demonstration is not a demonstration. Get the prospect onto the free tier with one of their own real problems and let the product argue for itself. Where a guided walkthrough is wanted, keep it to the loop below, because the loop is the differentiator and individual screens are not.

HOW TO GET A DEMO

Anyone can sign up on the free tier at the product website with no card and no sales conversation, which is enough to run the whole flow below at small scale. For a guided session, or for anything involving a dedicated tenant, a hosting region or a security review, book time with Vignesh T.V. on the call link and we will run it live.

Runs about 30 to 45 minutes for a guided walkthrough. A prospect exploring the free tier alone usually needs about an hour to reach the same understanding.

WHERE TO LOOK

- **Product website and free sign-up** — <https://intelwatchtower.com>
- **The product itself** — <https://app.intelwatchtower.com>
- **Documentation** — <https://docs.intelwatchtower.com>
- **Book a guided walkthrough with Vignesh T.V.** — <https://calendly.com/vignesh-burdenoff/1-1>

The walkthrough, in order

1 The shared overview

This is what the room looks at. Mission status, the alert queue by severity, sources by discipline, the top triage items, all drawn from the records people are working in rather than a deck someone built last night.

2 The source registry

Every source you collect from, with its discipline, status, region and a reliability score out of 100. Tasking decisions are made against this rather than against whoever remembers which feed is dependable.

3 Raising a collection requirement

This is the piece nobody else does. You state what you need to know, give it a priority from routine up to flash and a due date, and it becomes a record with a lifecycle rather than a line in a spreadsheet.

4 Issuing a tasking order against a source

Now the requirement has an order against a named source, with instructions and a collection window. When reporting comes back, the requirement moves to satisfied and the answer stays attached to the question.

5 An entity built from several sources

This is where reporting lands. One record per subject, with aliases, attributes, a confidence grade, a classification, coordinates and a threat score, and the reporting behind it still attached, so you can always say where a claim came from.

6 Triaging an alert and escalating it

Severity, assignee, triage notes, and a controlled lifecycle. Notice the alert links to the entity it implicates, and escalating it opens a case with that context carried across. Your security analyst never leaves the picture.

7 A case with its entries

Every note, action and piece of evidence is a timestamped entry attributed to whoever added it. When someone asks where this came from, the answer is on the screen rather than in somebody's inbox.

8 Releasing a report, then the activity timeline

The finished product moves through review to release, and can be recalled. Then this: every action, with who did it and on whose behalf, searchable and exportable. That is the answer to the audit question, and it is the same answer in every workspace.

15 Sales Talking Points

Nine things worth saying, and why each one lands. Say them in the prospect's language, not ours; if a phrase below only makes sense to someone who does this work, drop it.

One thread from the question to the report

It is the single hardest thing to assemble from separate tools, and the thing every buyer in this category has tried and failed to do with integrations. It reframes the purchase from 'another tool' to 'the thing that removes four handoffs'.

Every intelligence question gets tracked to an answer

Collection managers have no system for this today and feel it every week. It is the capability no named competitor publishes, so it is where the conversation goes somewhere new rather than into a feature comparison.

The assessment says how confident it is

Decision-makers are used to being handed either certainty or noise. A graded confidence and a source trail change how the output gets read, and that is a boardroom argument, not a technical one.

Cyber alerts arrive with the rest of the picture attached

Security leads instantly recognise the cost of rebuilding context per alert. It is also the bridge that lets a security budget pay for something the intelligence team wanted.

Governance is in the data, not in a memo

Classification on every record, access checked on every operation, workspace isolation on every write, and a read-only auditor role. It converts the compliance stakeholder from a blocker into a sponsor.

Published pricing, starting free

In a category where everything is quoted, a public price is disarming. It also lets a champion socialise the idea internally before anyone has to ask for budget.

The customer's own AI key, no markup

For sensitive work, where the data goes matters more than what the feature does. Bring-your-own-key means their intelligence flows through their own provider relationship.

Agents work under the same rules as people

Every security team is currently being asked what happens when someone points an AI agent at their systems. Having a real answer, scoped keys and identical access control, is a differentiator right now and will not stay one for long.

We tell you what is not built

The map, the link charts, the detection-rule interface, most of the AI analysis. Naming the gaps before the prospect finds them is the fastest route to being believed about everything else.

Eleven things prospects actually say. The responses are written to be spoken; the proof line is what you point at if they push. Never invent an answer you are not sure of, and never soften an honest no.

"Ninety-nine dollars a month for an intelligence platform? What is missing?"

Capacity, not capability. Every tier gets the whole platform surface, and the plans differ by how many seats, sources, feeds, observations, entities, cases and reports you get. The reason the entry price is low is that identity, roles, billing, files, notifications and audit are shared services across a family of products rather than costs this one product has to carry alone.

PROOF · Per-plan quotas are published on the pricing page, and the free tier includes the same feature set at smaller numbers.

"You have no customers I can call."

That is right, and I would rather say so than dress it up. What you can do instead is run the free tier against a real problem of yours for a week. If the loop from requirement to report does not fit how you work, you will know by Friday, and it will have cost you nothing.

PROOF · Free sign-up at the product website, no card, no sales conversation required.

"Why would we not just buy Palantir?"

If you are a national programme with the budget, the accreditation requirement and a year for procurement, you probably should. Most teams asking me this are a cell inside a larger organisation and cannot reach that budget or that timeline. This gives you the same shape of platform, published pricing, and a team using it this month.

PROOF · The comparison page on the website sets both out side by side, including where they are ahead.

"We already pay for a threat-intelligence subscription. Why add another tool?"

Keep it. That subscription is a source; this is where sources get fused with everything else you know and turned into cases, missions and reports. The question is not which feed is better, it is where the picture lives once the feed has done its job.

PROOF · Indicators, alerts, entities and cases share one model, so a feed-derived indicator lands in the same picture as geospatial and human reporting.

"We could build this on our existing monitoring stack and a case tool."

Teams do try. What they usually end up with is two systems and a synchronisation job, and the provenance dies at the boundary. The part that is hard to build is not any single screen, it is keeping a requirement, an observation, an entity, an alert, a case and a report on one thread with classification and audit intact.

PROOF · Ask them how long their last internal integration took to build and who maintains it now.

"Are you FedRAMP authorised? We need government accreditation."

No, and it is not in progress. The product is not FedRAMP, Department of Defense impact level, CMMC or ITAR authorised, and the classification labels in it are workspace metadata you configure rather than a government accreditation. If accreditation is a procurement gate for you, we are not a fit today and I would rather tell you now.

PROOF · The pricing page states this in plain language, unprompted.

"Our data is sensitive. Where does it live and who can see it?"

Every record is scoped to your workspace and every write is checked against it, so another organisation's identifier cannot reach your data. Access is enforced on every operation before the request reaches the service, and every action records who did it and on whose behalf. Hosting region is configurable on an Enterprise agreement, and dedicated, self-hosted and air-gapped deployment are all available.

PROOF · A read-only governance auditor role is seeded in every workspace so they can inspect this themselves.

"If we put everything in here, how do we get it out again?"

Everything the interface does is available through the API, and the activity trail and analytics are exportable. There is no proprietary ontology you have to be taught and no bespoke schema you would have to unpick. It is a fair question to ask us in writing, and we will answer it in writing.

PROOF · One published API behind every surface, with command-line and SDK clients built against it.

"You are a small company. What happens if you go away?"

It is a real risk and you should price it in. Two things reduce it: your data is exportable through the same API the product runs on, and this is one product in a family that shares its identity, billing and audit services, so the platform underneath is not carried by this product alone. For a critical deployment, ask for the continuity terms in the agreement.

PROOF · Ask for the escrow and continuity language during contracting; we will put it in front of you rather than argue about it.

"The AI features look thin next to what everyone else is shipping."

They are, and that is the honest gap. What exists today is a page-aware assistant, a detection engine that raises alerts with their provenance attached, and an agent server that lets your own AI tooling operate under the same permissions your people have. Automated correlation, scoring and the autonomous watch are still being built.

PROOF · The website marks each AI capability as shipped, in progress or roadmap, so nothing here needs taking on trust.

"Who supports us when something breaks at three in the morning?"

Support scales with the plan: community on free, email within 48 hours on Starter, priority within 24 hours on Growth, within 8 hours on Business, and a dedicated commitment on Enterprise. Availability commitments start at 99.5% on Growth. Below Enterprise, be clear that this is business-hours support with a target, not a follow-the-sun operations desk.

PROOF · Support tiers and uptime commitments are published against each plan.

17 Discovery / Qualification Questions

Eleven questions that separate a real opportunity from an interesting conversation. Ask them plainly and listen for the specifics; a vague answer to any of the first four usually means there is no budget behind the pain.

ASK THIS

LISTEN FOR

Walk me through how one assessment gets produced today. How many tools does it touch?

Four or more named tools, and manual copying between at least two of them. If they name one tool and are happy with it, this is not an opportunity.

If someone challenged a claim in your last report, how long would it take to trace it back to the original source?

Hesitation, or an answer measured in days. If they say 'it is all in the case system and takes a minute', they have already solved this.

Where do your intelligence requirements live, and what tells you when one has been answered?

A spreadsheet, a shared document or 'in people's heads'. This is the question with the highest hit rate, because almost nobody has a system for it.

How does an analyst know whether a given source is reliable?

Naming a person rather than a system. That is a dependency they already worry about.

When a security alert matters, how much work is it to find out what it relates to?

Context being rebuilt per alert, or a security console that is disconnected from everything else. Quantify it in analyst minutes if you can.

Who reads your finished reporting, and how do they know how much to trust it?

No confidence grading, or confidence expressed in prose. It tells you whether the decision-maker is a stakeholder in this purchase.

What does your next audit or oversight review actually ask you to produce?

Access evidence, handling proof or a disclosure record. If a review is scheduled, you have a date to work back from.

Do you share intelligence with another team, agency or partner? How is that controlled today?

Email and trust. Be careful here: compartmentation and releasability controls are on the roadmap, not shipped, so this is a future-fit signal rather than a today answer.

How many separate subscriptions go into producing the picture your leadership sees?

A count and a total. This is the number the economic buyer will compare against.

Does anything in your procurement require a specific certification or accreditation?

FedRAMP, a Department of Defense impact level, CMMC, ITAR or a completed SOC 2 report. Any of these is a hard stop today and you should find it out early.

If this worked, who signs it off, and what is their budget cycle?

A named person and a quarter. Without both, you have interest and not an opportunity.

What have you already tried, and why did it stop?

A failed integration project or an abandoned platform evaluation. It tells you what the objection will be before they raise it.

18 How to Identify a Good Opportunity

Use this to decide where to spend your time. One green signal is worth a conversation; three is worth an introduction. A single red signal is usually enough to stop, and stopping early is a service to everyone.

STRONG SIGNAL – PURSUE

- Analysts are copying between four or more tools to produce one assessment
- An audit, oversight review or accreditation exercise is scheduled in the next two quarters
- The operating picture is a deck someone rebuilds by hand, and leadership has complained about it
- Cyber and physical or geopolitical risk are landing on the same team and nothing joins them up
- A team is running investigations in shared drives and worrying about chain of custody
- Someone has been asked what happens when an AI agent is pointed at their intelligence data

Introduce us straight away.

WORTH A CONVERSATION

- Strong interest from analysts but no named budget holder in the conversation yet
- They are mid-contract on a threat-intelligence subscription and cannot move until renewal
- The requirement leans heavily on visual link charting, which is still being built
- They want a hosting region or dedicated tenant, which means an Enterprise agreement and a longer path
- Cross-agency sharing is the driver, and compartmentation and releasability controls are on the roadmap

Qualify further before an introduction.

POOR FIT – DO NOT CHASE

- Procurement requires FedRAMP, a Department of Defense impact level, CMMC or ITAR authorisation
- They want the single best specialist tool for one discipline and nothing else
- The real requirement is high-volume log ingestion and search, which is a different category of product
- They will not run a pilot until they can call three reference customers
- One person doing occasional research, with no team, no roles and no governance question

Politely park it; suggest another product if one fits.

THE ONE-QUESTION TEST

Ask: 'Where do your intelligence requirements live today, and what tells you when one has been answered?' If the answer is a spreadsheet, a shared document or a shrug, you have a real opportunity, because that is the gap no competitor fills and it usually sits next to every other pain in this document. If they name a system that already does it well, thank them and move on.

19 How to Introduce Us

Keep the introduction short and about them. Your job is to earn twenty minutes, not to explain the platform. Say the problem back to them in their own words, name the product once, and offer the call.

IN CONVERSATION

"You know how your analysts pull a picture together from a feed, a map, a case system and a spreadsheet, and the source trail is gone by the time it reaches your director? IntelWatchtower does that whole cycle in one place, from stating what you need to know through to the released report, with a classification and a confidence level on every record. It is early and there is a free tier, so you can test it before anyone asks for budget. Want me to introduce you to the founder?"

EMAIL – SUBJECT: THE FIVE-TOOL PROBLEM YOUR ANALYSTS KEEP MENTIONING

Hi [Name],

You mentioned your team assembles its picture from a threat feed, a map, a case system and a spreadsheet, and that the source trail is usually gone by the time anything reaches leadership.

I have been looking at IntelWatchtower, which runs that whole cycle in one place: you state what you need to know, task a source against it, fuse what comes back onto tracked entities, triage the alerts, work the case and release the report, with a classification and a confidence level on every record.

It is early. Generally available this year, no long reference list. But there is a free tier you can try before any conversation, and published pricing from \$99 a month.

Worth twenty minutes with Vignesh, who runs it? Happy to introduce you.

[Your name]

SHORT MESSAGE / LINKEDIN

Saw your post about analysts spending more time moving data between tools than analysing it. IntelWatchtower puts the whole intelligence cycle on one thread, from a stated requirement through to a released report, with confidence and classification on every record. Free tier if you want to poke at it, or I can introduce you to the team. Worth a look either way.

WARM THREE-WAY INTRODUCTION

Introduce by problem, not by product. 'You two should talk. [Name] runs intelligence for [organisation] and has been fighting the tool-sprawl problem for a year. Vignesh built IntelWatchtower, which puts the whole cycle in one workspace.' Then get out of the way, copy vignesh@burdenoff.com, and let them find their own level.

Please do not: Do not claim customers, deployments, certifications or results. There are none, and one invented reference undoes everything else in this document.; Do not promise the map, the link charts, the detection-rule interface, compartmentation, releasability controls or the mobile app. All are roadmap, and a prospect who finds out later stops believing the parts that are true.; Do not quote a converted price. Quote the customer's own regional band, or say you will come back with it.; Do not imply any government accreditation or that the classification labels carry official weight. They are workspace metadata the customer configures.; Do not attempt a technical demonstration you have not rehearsed. Send them to the free tier or book the call instead.

20 Lead Handoff Process

Register first, introduce second. Registration is what protects your commission, and it has to happen before the prospect is in an active procurement conversation with us. The whole handoff should take you ten minutes.

1 Register the opportunity before you introduce

Email vignesh@burdenoff.com with the organisation, the contact, what triggered the conversation and roughly when they want to move. You will get a written confirmation, and that confirmation is the basis for your commission.

2 Send the qualifying detail

The five things in the list below. If you only have three, send three and say so; a partial picture beats a delayed one. Do not send anything the prospect told you in confidence without checking with them first.

3 Make the introduction

A short email or message introducing the prospect to Vignesh T.V. (vignesh@burdenoff.com, +91-7358445777) with one sentence on why they should talk. Copy yourself so you can see the thread move.

4 Let us qualify and demonstrate

We take the first call, run the walkthrough and answer the security and commercial questions. You are welcome on the call, and for a consulting or implementation engagement you should be.

5 Stay in the loop through the quote

We control final pricing, discounting and contract terms, and every quote is approved before it goes to the customer. If you are reselling, your quote needs written approval before you send it.

6 Agree your role at close

Referral, consulting, implementation or resale, decided before signature rather than after. Renewal commission needs an approved renewal or customer-success role, so agree that at the same time.

SEND US THIS – IT IS ALL WE NEED

- Organisation, sector and rough size of the team who would use it
- The named contact, their role, and whether they can approve spend or only recommend
- What triggered the conversation: an audit, a failed integration, a tool consolidation, a new mandate
- Which tools they use today for collection, analysis, cases and reporting
- Any certification, accreditation, hosting-region or data-residency requirement, stated early
- Rough timing and budget cycle, even if it is only a quarter
- Anything you have already told them, so we do not contradict you on the first call

WHAT HAPPENS NEXT, AND WHEN

We acknowledge a registration within one working day and confirm whether it is protected. First contact with the prospect happens within two working days of the introduction, and you get a status update after the first call and at any material change. If we go quiet, chase Vignesh T.V. directly on +91-7358445777.

Deal registration: An approved registration protects the opportunity for 90 days, with one further 90-day extension available where there is active progress. Registration must come before we are already in an active opportunity with that account, and there is no retroactive registration. Where two partners claim the same account, the first approved registration with real contribution takes priority and any split needs written approval.

21 Partner / Referral Commercials

Partners are paid on collected net subscription revenue, not on list price and not on invoiced amounts. If a customer is billed in a lower regional band or in rupees, the payout is calculated from what was actually collected in that band. Taxes, refunds, chargebacks, credits, payment fees, customer-owned cloud costs and AI provider costs are all excluded before the calculation.

PROGRAMME	YOUR ROLE	FIRST-YEAR	RENEWAL / ADD-ONS
Referral Partner	Registered introduction, no sales ownership	10% of first-year collected net subscription	0% unless separately approved · 3% of first-year collected net add-ons, capped at \$10,000
Intelligence or Security Consulting Partner	Introduction plus support through the evaluation	12% of first-year collected net subscription	0 to 3% where active in the success plan · 3% of first-year collected net add-ons, capped at \$15,000
Implementation Partner	Introduction plus delivery of the rollout, migration, integration or training	15% of first-year collected net subscription	3 to 5% where active in the renewal success plan · 5% of first-year collected net add-ons, capped at \$25,000
Authorised Reseller	Approved resale with first-line commercial ownership	Up to 15% reseller margin on the approved quote	Up to 8% where the account stays active and supported · Up to 5% on dedicated tenant or managed services, only where approved
Managed service or Strategic Partner	Operating IntelWatchtower for end customers, or including it by name in a bundled offering	Custom wholesale or managed-service terms	Custom · Custom

THE RULES THAT MATTER

- Register before you introduce. There is no retroactive registration, and no commission where we already had an active opportunity in that account unless you open a new buying centre or materially expand the scope.
- Larger deals earn more: an extra 1% on first-year net subscription between \$100,000 and \$249,999 capped at \$2,500, 2% between \$250,000 and \$499,999 capped at \$7,500, and 3% above \$500,000 capped at \$25,000. All require registration before procurement and approval of the quote before signature.
- Regional pricing is not a discount and does not reduce your eligibility. It reduces the payout only because the collected amount is lower. Extra discounts, credits and coupons do reduce the commissionable revenue.
- Dedicated tenant, self-hosted, air-gapped, premium support and managed-service revenue carry lower, quote-approved rates because of the delivery effort behind them.
- Professional services you contract directly with the customer are yours to keep, provided they do not conflict with our agreement with that customer.
- Partner commission and affiliate rewards do not stack. One customer is attributed to one programme unless we approve an exception in writing.
- Nothing is commissionable on taxes, credits, refunds, pass-through costs, the customer's own cloud spend, their AI provider spend, or marketplace publisher revenue, which follows a separate programme.
- Public-sector work needs registration and a compliance review first, and no payment, gift or referral benefit may go to a government employee or official.

GETTING PAID

Commission is paid 15 days after revenue is successfully realised, on collected net revenue only, in the currency named in your partner agreement. Multi-year contracts pay as the customer's cash is collected rather than on total contract value unless we agree otherwise. Refunds and chargebacks reverse the commission or are deducted from a future payout, and you are responsible for your own tax forms and local compliance while we withhold and report where the law requires it.

These are the standard programme terms and are shown here for orientation only. The terms that apply to you are the ones in your signed partner agreement, and the terms that apply to the customer are the ones in the quote and order form we issue to them. Where anything here and a signed document differ, the signed document governs.

22 Contact Information

YOUR CONTACT

Vignesh T.V.

vignesh@burdenoff.com
+91-7358445777

BOOK A 1-1 CALL

<https://calendly.com/vignesh-burdenoff/1-1>

Use this for anything — qualification, pricing, a joint call with your prospect, or a live demo.

EVERY BURDENOFF PRODUCT

<https://burdenoff.com/products>

If IntelWatchtower is not the right fit, the answer may well be on that page.

WHICH CHANNEL FOR WHAT

- **A live opportunity** — email with the handoff details in section 20, and copy the phone number in if it is time-sensitive.
- **A question you cannot answer** — book the 1-1 call; it is usually faster than email.
- **A prospect who wants to meet us** — send them the calendly link directly, or ask for a joint call.

23 Useful Links & Sales Assets

ASSET	WHERE	WHAT IT IS FOR
IntelWatchtower website	https://intelwatchtower.com	The public positioning, in the product's own words. Start a prospect here.
The product	https://app.intelwatchtower.com	Sign in, or sign up on the free tier without a card.
Pricing	https://intelwatchtower.com/pricing	Reads the live billing catalogue. This is the authoritative price, not this document.
Use cases	https://intelwatchtower.com/use-cases	The same scenarios as this playbook, written for the prospect rather than for you.
Competitive comparisons	https://intelwatchtower.com/vs	Side-by-side pages for each named alternative, including where they are ahead.
Security posture	https://intelwatchtower.com/security	Classification, access control and audit, for the security stakeholder.
Documentation	https://docs.intelwatchtower.com	Product guides and API reference for a technical evaluator.
Book a call with Vignesh T.V.	https://calendly.com/vignesh-burdenoff/1-1	Direct diary access for a guided walkthrough or a commercial conversation.
All Burdenoff products	https://burdenoff.com/products	The wider family this product sits in, useful when a prospect asks who is behind it.

WHAT YOU CAN SHARE, AND WITH WHOM

- Share freely: anything published on the product website, including pricing, comparison pages, use cases and the security page. It is written to be public and it is kept current.
- Prospects only, one to one: this playbook, the cheat sheet and the walkthrough flow. They are written for partners and contain positioning judgements that read badly out of context. Do not post them publicly or attach them to a mass mailing.
- Ask us first: any written comparison you author yourself, any claim about a named competitor, any statement about certifications, accreditation, hosting region or roadmap timing, and any use of the IntelWatchtower name or logo in your own marketing.
- Never share: anything a prospect told you in confidence, internal pricing floors or discount thresholds, and any figure not published on the website. Quoting an internal number to a customer is the fastest way to lose deal approval.
- Every quote we issue is approved before it reaches the customer. Do not send pricing you have constructed yourself, even where you are an approved reseller.

24 Product Mockups — Complete Visual Appendix


44 SCREENS

Every screen in the IntelWatchtower design set — 44 in total, none omitted. Use these to show a prospect what the product looks like before a live demo.


PRODUCT SCREENS • 34 SCREENS




001 Watchtower Overview




002 Landing Page




003 Common Operating Picture




004 Sources List




005 Source Detail




006 Source Register Modal




007 Requirements




008 Tasks




017 Cases List




018 Case Detail




019 Case Add Entry Modal




020 Reports List




021 Report Detail




022 Mobile Overview




023 Mobile Alert Triage




024 Light Theme Dashboard



025 Feeds Ingest Health




026 Tasking Complete



027 Raise Alert Modal




028 AI Assistant Entity Context



029 Auto Raised Alert Provenance



030 Link Analysis Graph




031 AI Report Draft



032 AI Threat Score Proposal




033 Watch Loop Handover Delta




034 MCP Agent Console




035 Payment Fraud Mule Network Case 1 Mule Network Case Graph




036 Payment Fraud Mule Network Case 2 Hold Approval On Phone




037 Critical Energy Site Security Picture 1 Site Security Picture




038 Critical Energy Site Security Picture 2 Site Alert On A Phone




039 Municipal Emergency Operations Picture 1 Flood Response Picture




040 Claims Fraud Ring Link Analysis 1 Claims Link Analysis




041 Multi Agency Fusion Conflict Board 1 Conflict Board For One V



042 Releasability Tearline Review 1 Partner Version Review



043 Satellite Tasking Requests And Pass Windows 1 Acquisition Tasking



044 Subscriber Router Botnet Watch 1 Router Botnet Watch



Found someone who needs IntelWatchtower?

You do not have to be an intelligence specialist to spot this one. If someone tells you their analysts spend more time moving data between tools than thinking about it, that is your opening. Register the opportunity, make the introduction, and we will carry it from there.

EMAIL OR CALL

Vignesh T.V.
vignesh@burdenoff.com
+91-7358445777

BOOK A 1-1 WITH VIGNESH

<https://calendly.com/vignesh-burdenoff/1-1>

ALL BURDENOFF PRODUCTS

<https://burdenoff.com/products>



Burdenoff Consultancy Services Pvt. Ltd. · IntelWatchtower is available at intelwatchtower.com.

This playbook is confidential and intended for Burdenoff partners. Figures, plans and capabilities are current at the date of issue and may change; final commercial terms are those set out in your partner agreement and in the quote issued to the customer.