SOLUTION BRIEF

SELECTOR

Selector Foundry

Agentic NetOps, built on full-stack observability and AIOps

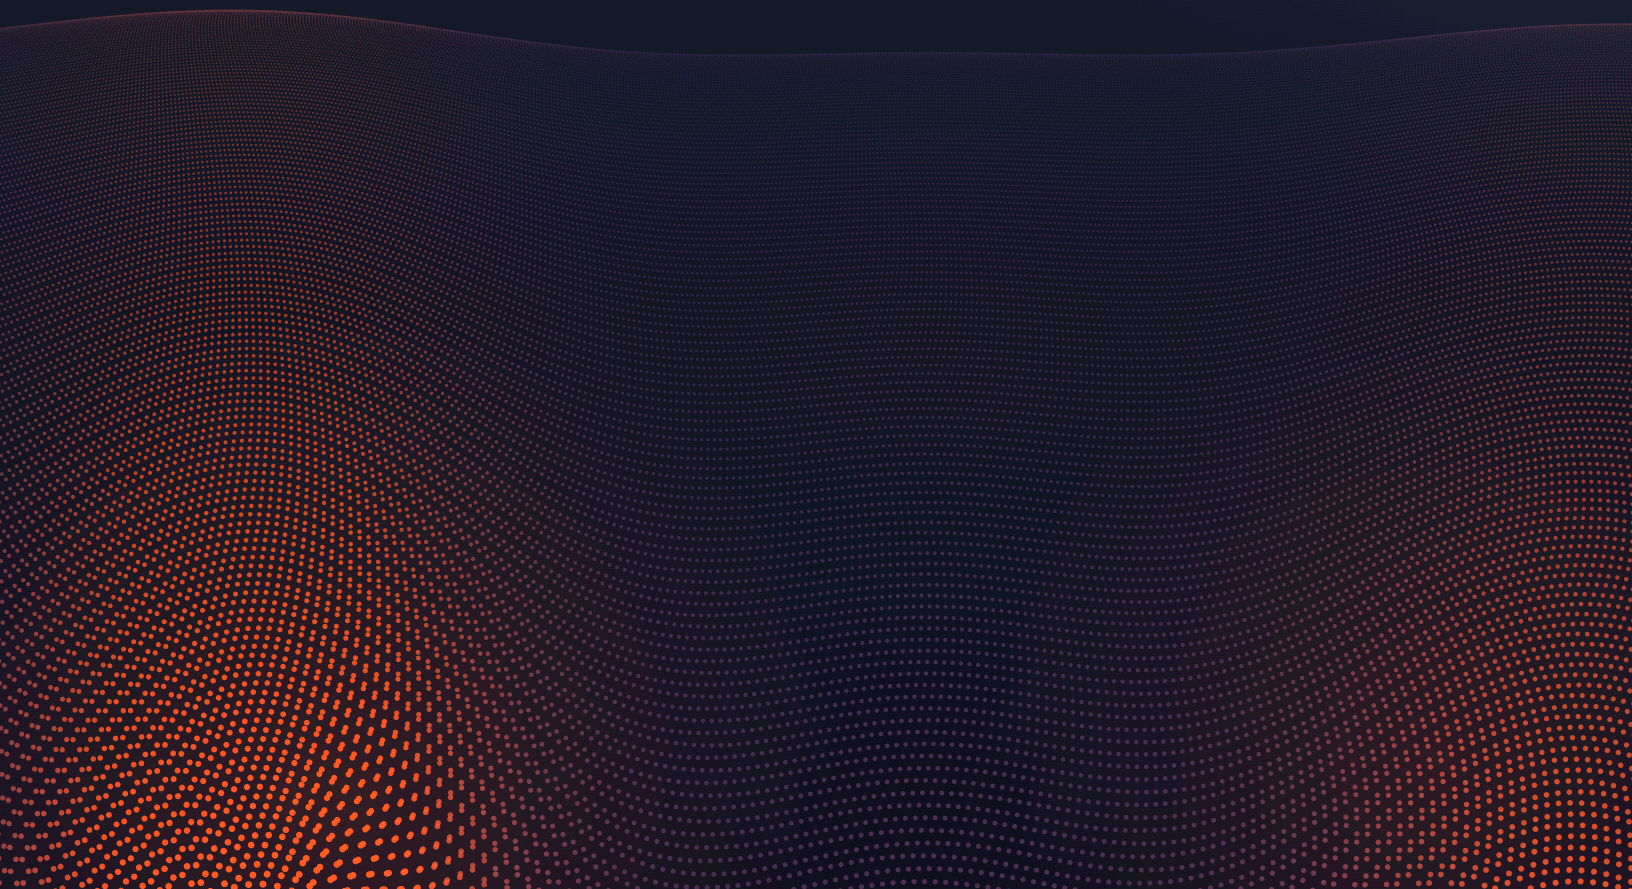

THE GAP

Network operations have changed. AI has not.

Networks no longer fail in one place. A single service disruption travels across data center fabrics, cloud regions, transit providers, firewalls, overlays, and the applications riding on top of them — and every one of those layers is instrumented by a different tool.

Yet most AI in network operations still answers one question at a time. A chat window on top of a dashboard can summarize what an operator already found, but it cannot open the investigation, follow the path across domains, or carry the work through to resolution.

The outcome is familiar. Enterprises bought AI for NetOps, and the investigation stayed human. Teams still assemble context by hand, still coordinate every handoff, and still spend most of an incident finding the problem rather than fixing it.

> **Conventional AI can answer what an operator asks. Operating a network requires something that knows what to ask next — and can do the work.**

WHAT AI DOES

DETECT
Sees the alert. Answers the question.

WHAT STAYS HUMAN

CORRELATE
connect evidence

DIAGNOSE
find root cause

ACT
prepare the fix

VERIFY
prove recovery

Prompt-driven
Waits to be asked. Never starts from the incident.

Context-limited
Sees one slice. Cannot connect the full path.

Assists, doesn't run
One question at a time. Never carries the case.

Answers without proof
Suggests a fix. Can't confirm recovery.

THE APPROACH

A new operational approach to NetOps

Meeting this challenge requires more than a better assistant. It requires an operating model where operational context, reasoning, and action live on one platform rather than in three disconnected products.

Selector delivers this through Agentic NetOps — specialized agents that reason over a single live model of the full path across network, cloud, infrastructure, and change data, then act within governed limits.

Because that context is built before an incident rather than assembled during one, work can begin the moment an event occurs, with or without an operator prompt.

NETOPS, CLOUDOPS, AND SERVICE ASSURANCE TEAMS GAIN ONE OPERATING MODEL THAT:

- Investigates across every domain, not a single slice
- Backs every finding with cited, reproducible evidence
- Carries the work from first question to closed loop

RECOGNIZED BY GARTNER®

Selector is named a Representative Vendor in the 2026 Gartner® Market Guide for Agentic NetOps Software.

"By 2030, AI agents will be the most common approach for executing network runtime activities, up from minimal adoption in late 2025."

THE PLATFORM

Agentic NetOps within a unified platform

Agents are only as good as the ground truth beneath them. Selector built that ground truth first: one horizontal collection architecture spanning network, cloud, infrastructure, and application domains, normalized into a single operational model that context travels with. Selector operationalizes this through three layers that each stand on the one beneath.

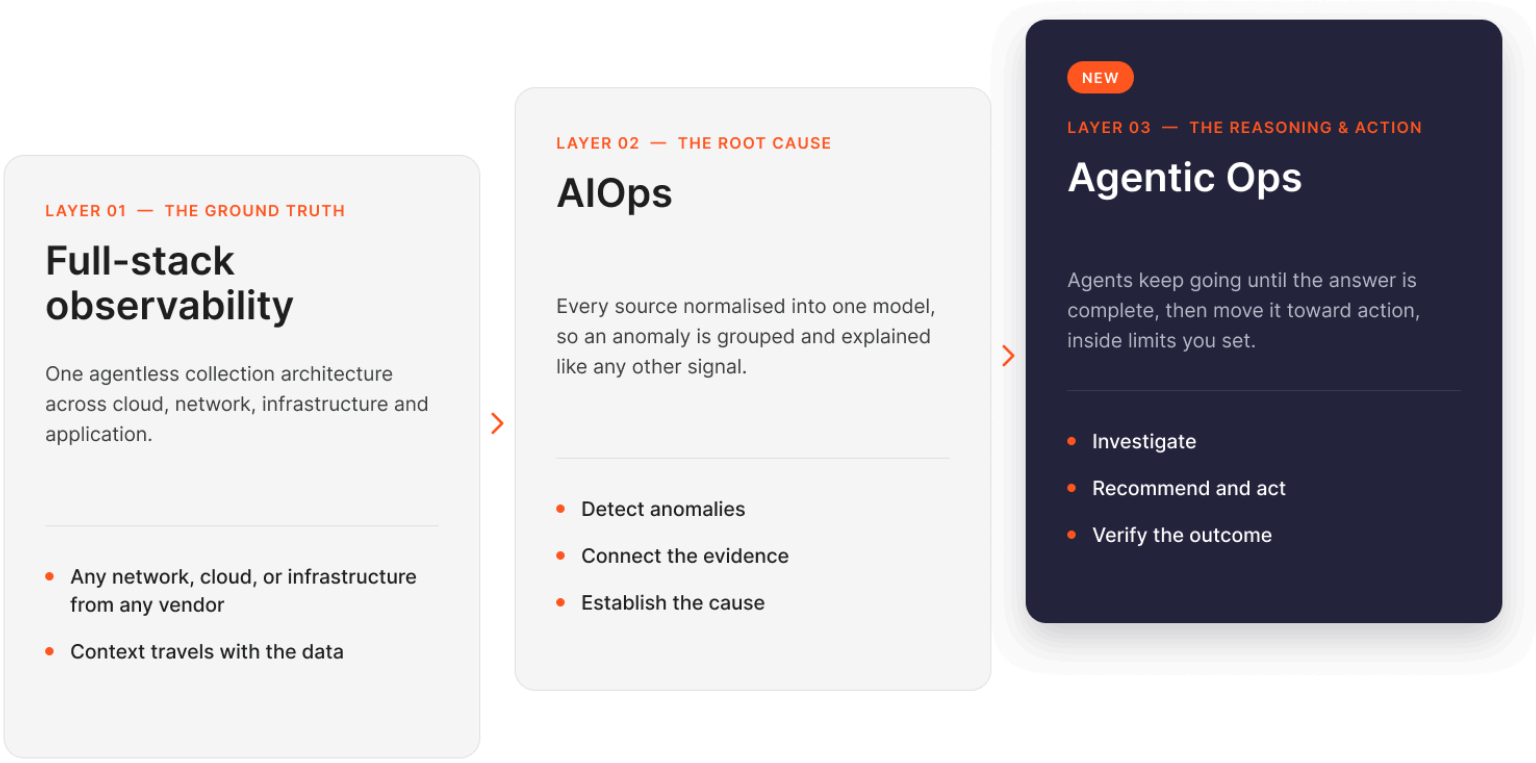

Each layer stands on the one beneath — Agentic Ops begins from narrowed, trusted context.

One live, shared model of the full path

Vendor-neutral telemetry, topology, configuration, incidents, and history in one model spanning devices, interfaces, firewalls, overlays, providers, and cloud constructs.

Specialized agents, not general-purpose AI

Network expertise encoded into agents built for defined jobs — correlation, ticketing, reporting, dashboards, and agent creation, each with approved tools and explicit limits.

Work that begins before the ask

Operator-led work runs through Rosetta. Event-led work begins automatically from an incident, schedule, API call, or another agent.

Governed action with evidence

AIOps detects anomalies, connects evidence, and establishes cause. Agents act from there through approved tools and guardrails, at the pace the organization sets.

THE IMPACT

What changes for NetOps teams

When context, reasoning, and action share one platform, the shape of the work changes. Instead of assembling evidence and coordinating specialists, teams direct investigations that agents run — and step in where judgment actually matters.

Organizations using Selector can:

01 ### Investigate in plain language across every domain

Ask why a service is failing and receive a connected answer spanning routing, cloud, firewall, and change data — with no query syntax and no tool pivots.

02 ### Collapse signal chaos into one incident story

Related anomalies, changes, and topology impact are connected into a single incident with a clear root cause and a defined blast radius.

03 ### Start the work before an operator asks

Incidents, schedules, and API triggers launch the right workflow automatically, so the investigation is already underway. An operator joins when needed.

04 ### Hand off tickets and updates that arrive complete

Tickets open with impact, evidence, owner, and next step attached. Executive updates and post-incident reports are written from the verified record.

05 ### Keep every next step explainable and controlled

Show the evidence behind conclusions, preserve approval gates, and verify outcomes against live operational state.

THE AGENTS

Applying agentic operations to real network challenges

The value of this model becomes clear in the work network teams do every day. Rosetta directs the play; five specialists carry the work from first question to closed loop.

01	Rosetta	The conversational investigation
02	Loom	Correlation and root cause analysis
03	TicketMaster	Turn findings into action
04	Artist	Operational views on demand
05	Herald	Executive-ready communication
06	Sculptor	Playbooks become governed agents

ROSETTA

The conversational investigation

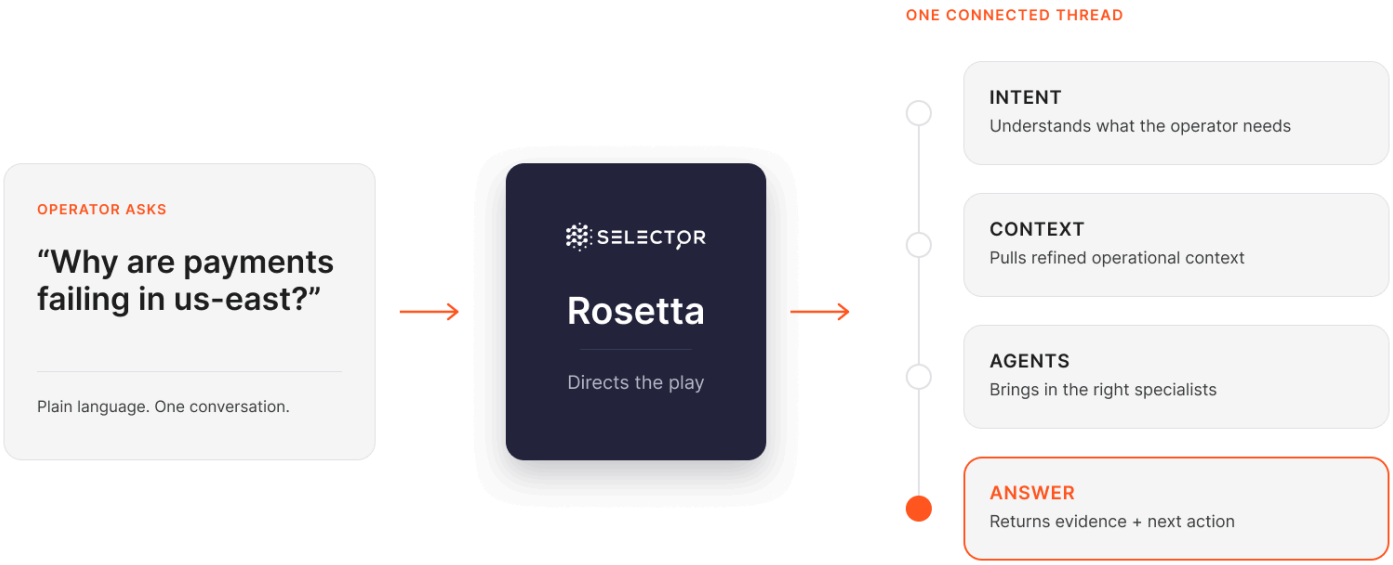

Intent, context, agents, answer — one connected thread, whether a person or an event opens it.

THE CHALLENGE

When a service degrades, the operator becomes the integration layer — pulling metrics from one tool, routing state from another, and configuration history from a third, then assembling the timeline by hand. The investigation moves at the speed of tool-switching, and the answer depends on who is on shift.

THE SELECTOR APPROACH

Rosetta interprets the question in plain language, grounds it in Selector's shared operational context, and dispatches the right agents across domains. Findings return as one connected thread with evidence, impact, and the next controlled action.

OPERATIONAL OUTCOME

One question replaces an hour of manual assembly. Every operator investigates with the reach of the team's most experienced engineer, and every answer is reproducible.

One question, one connected answer

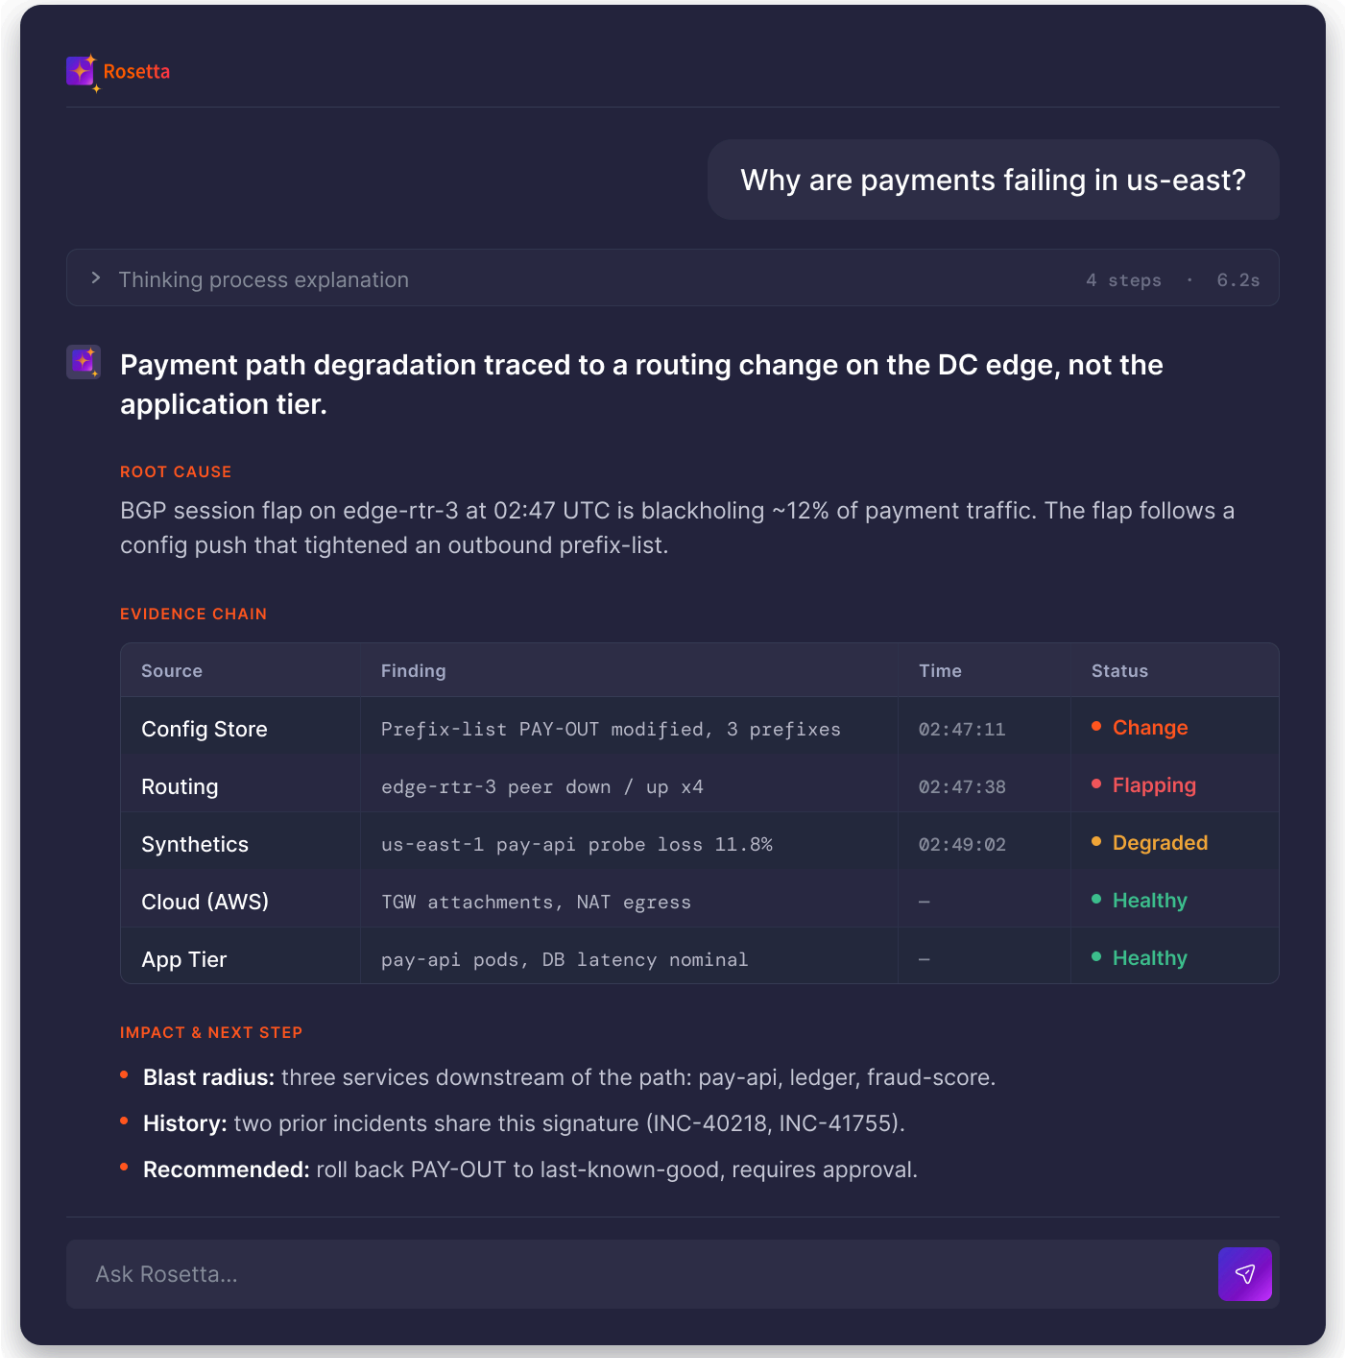

01

Asked in plain language

"Why are payments failing in us-east?" — no query syntax, one conversation.

02

Answered with cited evidence

Five sources, from the config push to a healthy app tier, timestamped in one chain.

03

Handed to a controlled action

Rollback to last-known-good is recommended, and still requires approval.

LOOM

Correlation and root cause analysis

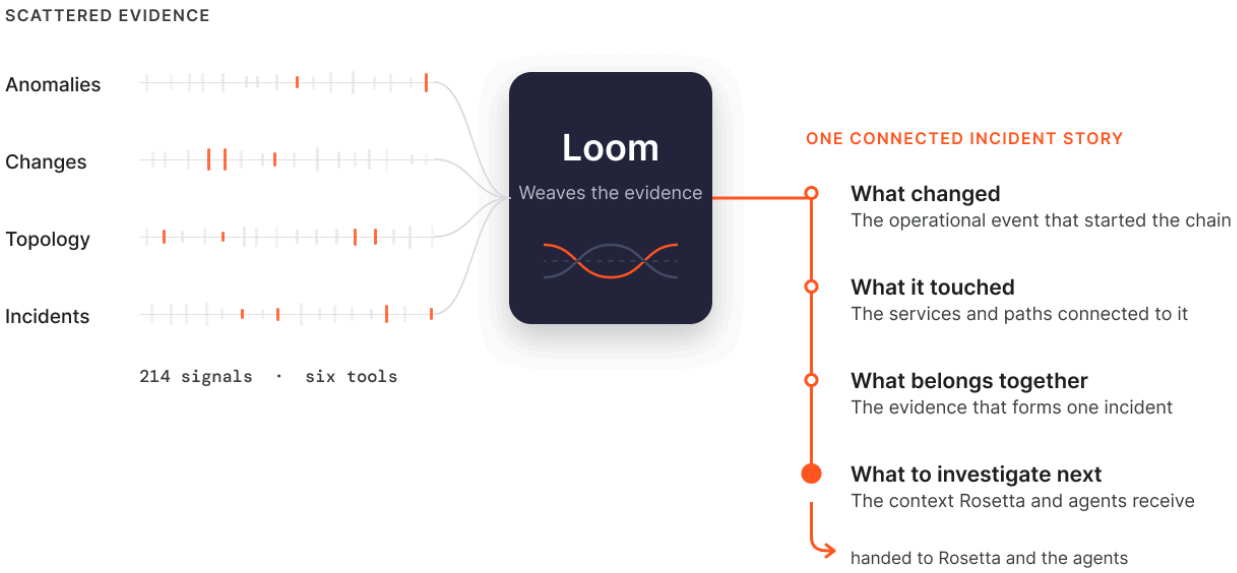

214 signals across six tools become one incident story — then hand off to Rosetta and the agents.

THE CHALLENGE

One failure produces dozens of alerts across domains, each showing a partial symptom. Teams manually reconstruct which signals belong together, and that reconstruction is slowest at exactly the moment the incident is most severe.

THE SELECTOR APPROACH

Loom connects relationships across time, topology, changes, and domains, grouping scattered evidence into one incident story: what changed, what it touched, what belongs together, and what to investigate next. Every agent downstream inherits that context.

OPERATIONAL OUTCOME

Alert floods collapse into a small number of actionable incidents, and the war room shrinks to the people who can actually fix it.

LOOM · IN THE PRODUCT

214 alerts, three incidents, one cause

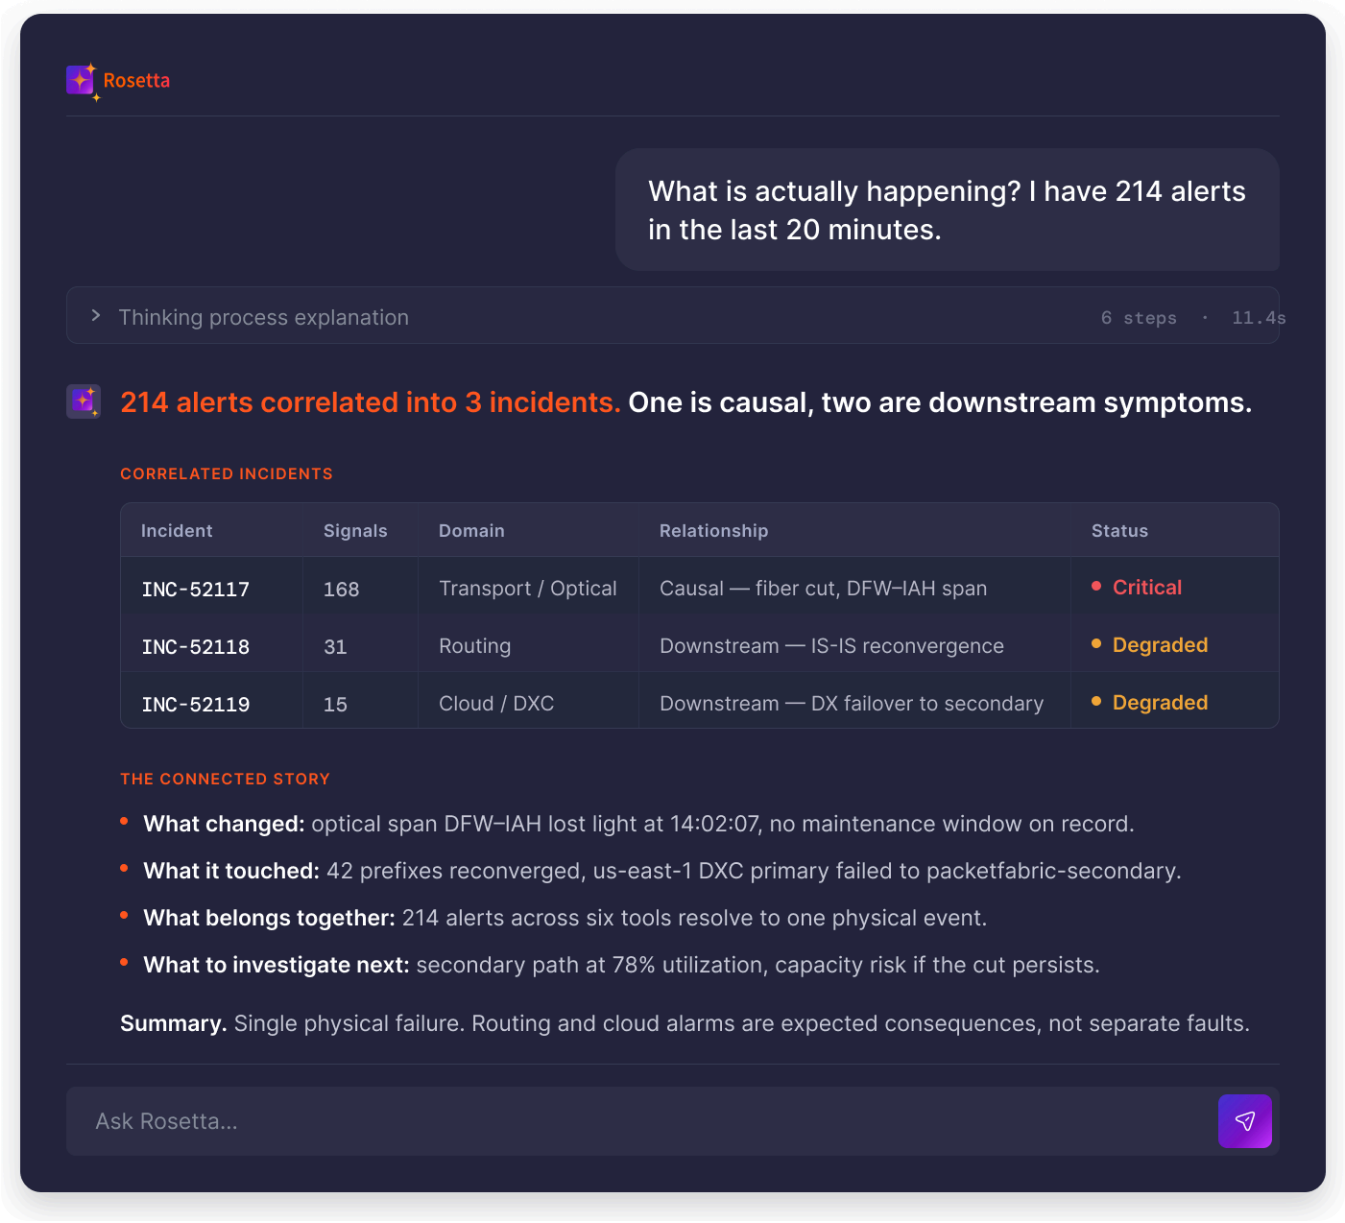

01

One causal event

214 alerts across six tools resolve to a single fiber cut on the DFW–IAH span.

02

Downstream, not separate

Routing reconvergence and cloud failover are consequences, not new faults.

03

The next risk, named

The secondary path is at 78% utilization if the cut persists.

TICKETMASTER

Turn findings into action

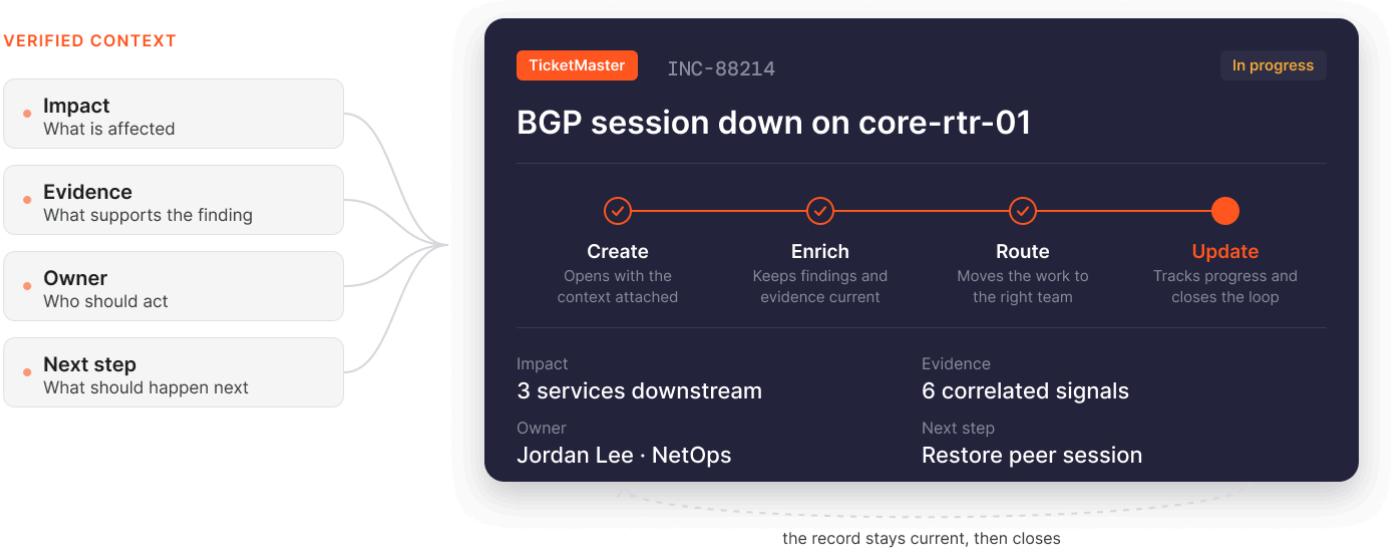

Create, enrich, route, update — impact, evidence, owner, and next step travel with the record.

THE CHALLENGE

Tickets typically arrive as a symptom and a timestamp. The assigned team repeats an investigation that has already been done, ownership is guessed at, and the record goes stale the moment the incident evolves past what was written down.

THE SELECTOR APPROACH

TicketMaster opens the record with verified context attached — impact, evidence, owner, and next step — then enriches, routes, and keeps it synchronized with the live investigation inside the ITSM platform teams already use.

OPERATIONAL OUTCOME

The ticket arrives complete. Nothing is re-investigated, and the record never falls behind the incident.

TICKETMASTER · IN THE PRODUCT

Delivered where the team already works

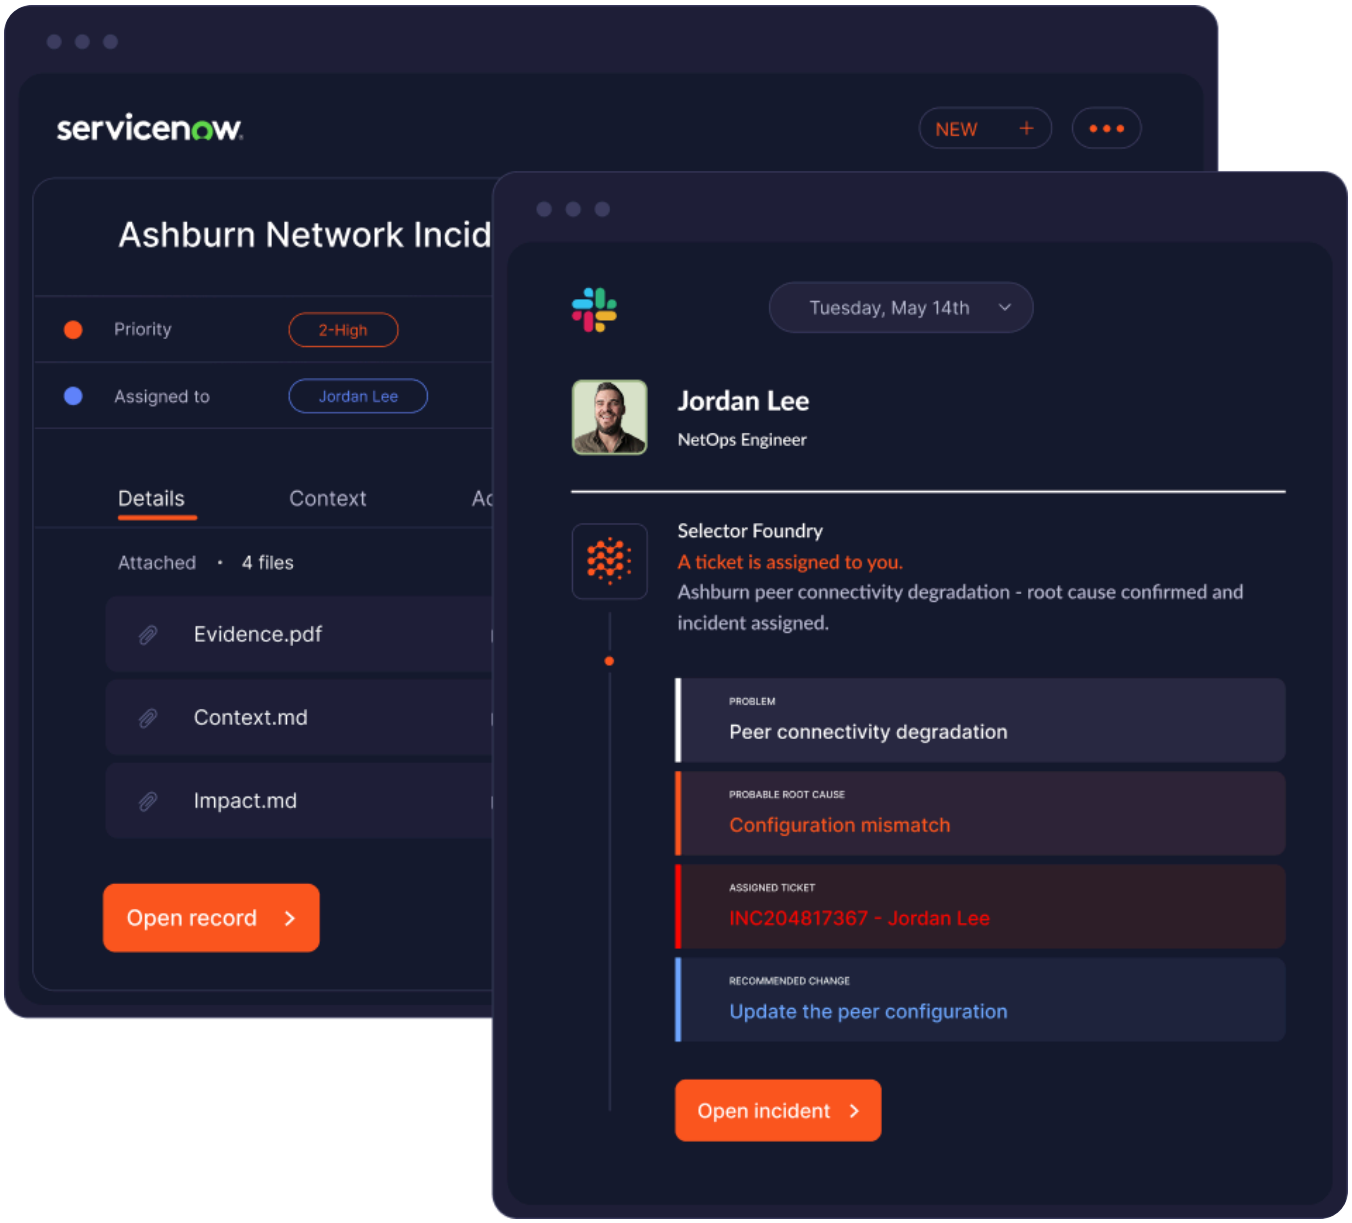

01

Arrives assigned

The ticket opens with owner, impact, and the evidence files attached.

02

Root cause stated

Configuration mismatch — not a symptom and a timestamp.

03

Next step included

The recommended peer-configuration change travels with the record.

ARTIST

Operational views on demand

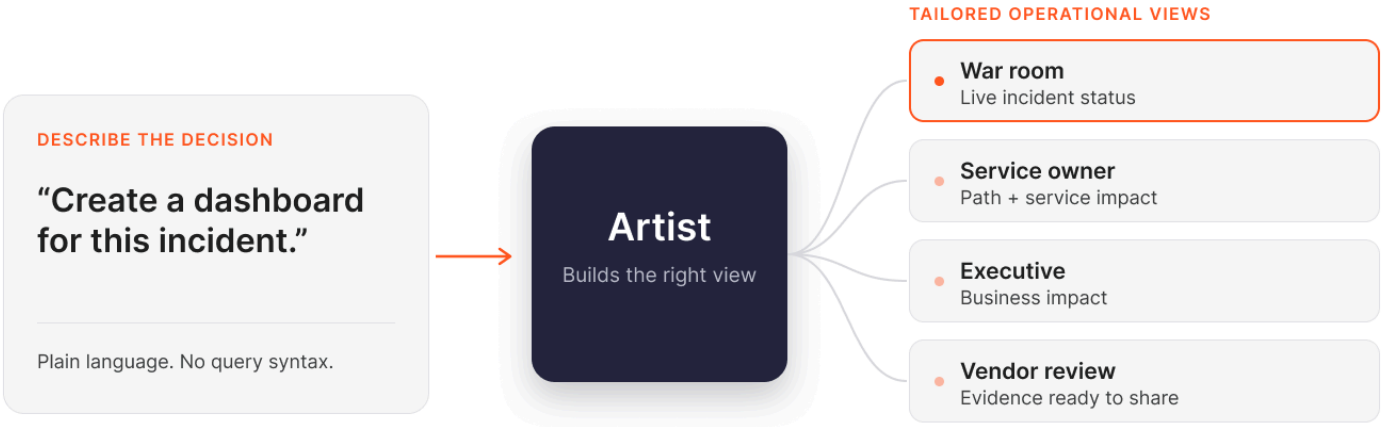

One described decision, four tailored views: war room, service owner, executive, vendor review.

THE CHALLENGE

The view a team needs during an incident rarely exists yet. Dashboard requests queue behind a platform backlog, and by the time the view ships, the decision it was built to support has already been made without it.

THE SELECTOR APPROACH

Artist turns a plain-English description of the decision into a live operational view built from current network and service data — tailored to the war room, the service owner, the executive, or a vendor review, with no query syntax involved.

OPERATIONAL OUTCOME

The dashboard becomes an output rather than a project — the right view at the moment of the decision.

ARTIST · IN THE PRODUCT

A described dashboard, built from live data

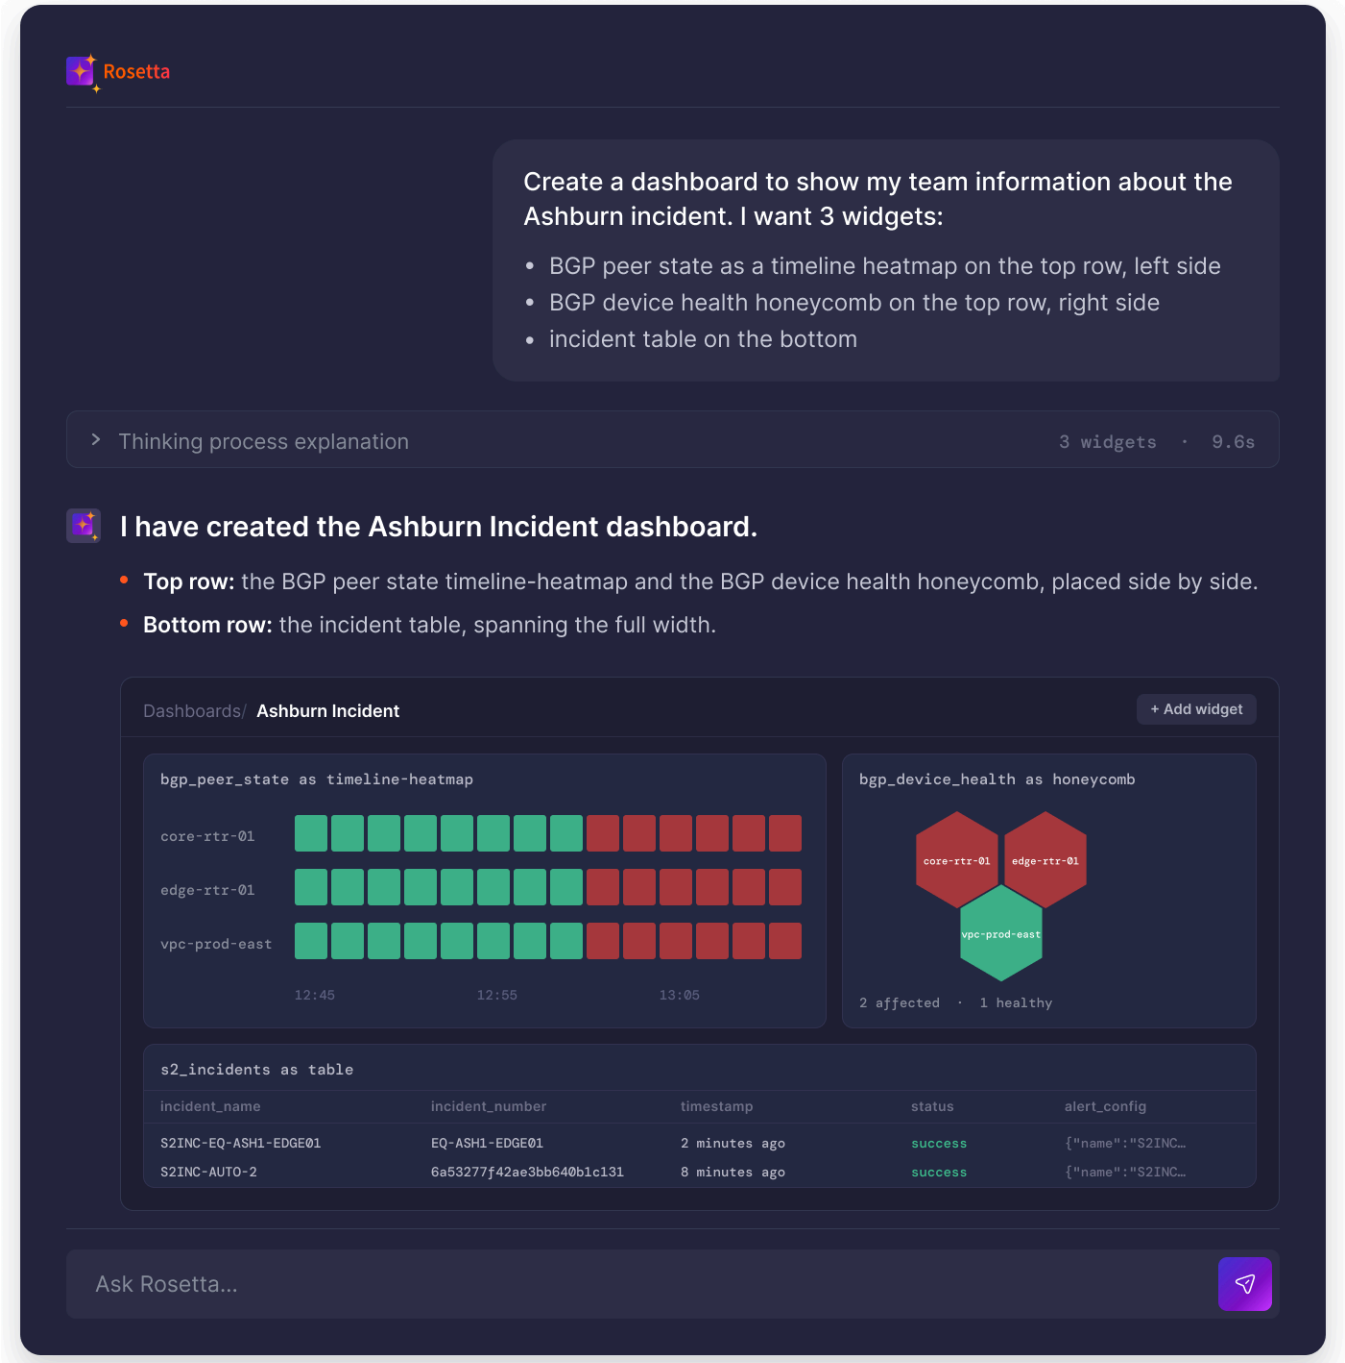

01

Described, not built

Three widgets and their placement, requested in plain English.

02

Drawn from live data

Peer-state timeline heatmap and device-health honeycomb, current at ask time.

03

Ready to share

The incident table spans the full width for the war room view.

HERALD

Executive-ready communication

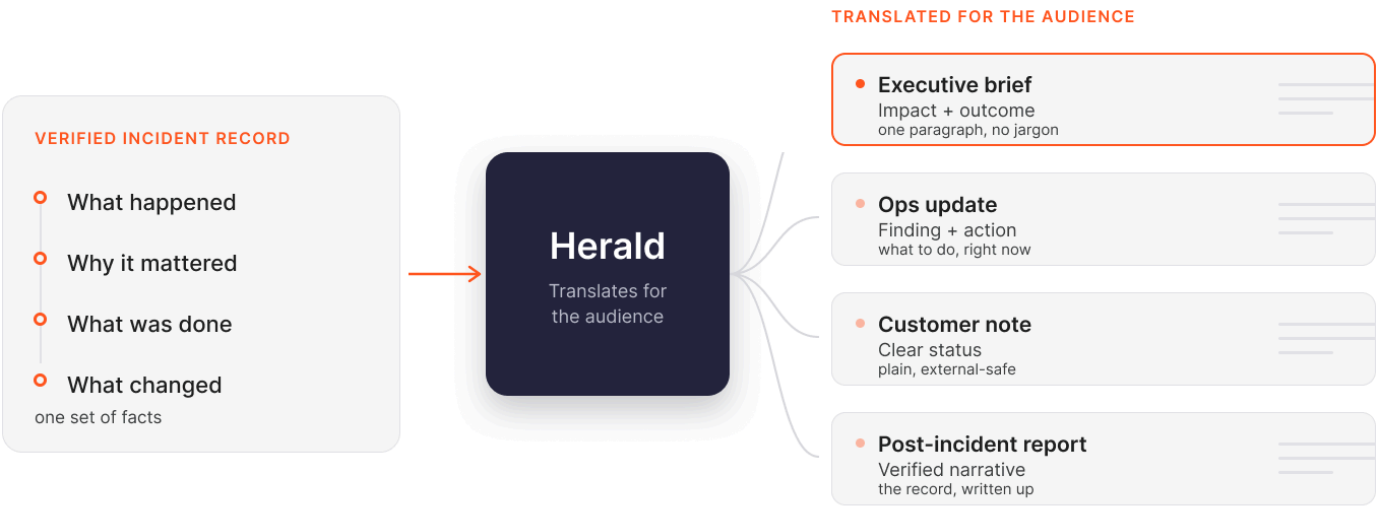

One set of facts, translated four ways — executive brief, ops update, customer note, post-incident report.

THE CHALLENGE

During a major incident, communication competes directly with resolution. Engineers stop fixing to write updates for executives, service owners, and customers — translating the same evidence three different ways, under pressure, from memory.

THE SELECTOR APPROACH

Herald reads the verified incident record and writes for the audience: what happened, why it mattered, what was done, and what changed. Executive briefs, ops updates, customer notes, and post-incident reports all come from one source of truth.

OPERATIONAL OUTCOME

Communication stops competing with resolution. The incident ends with the story already written.

HERALD · IN THE PRODUCT

The executive brief, written from the record

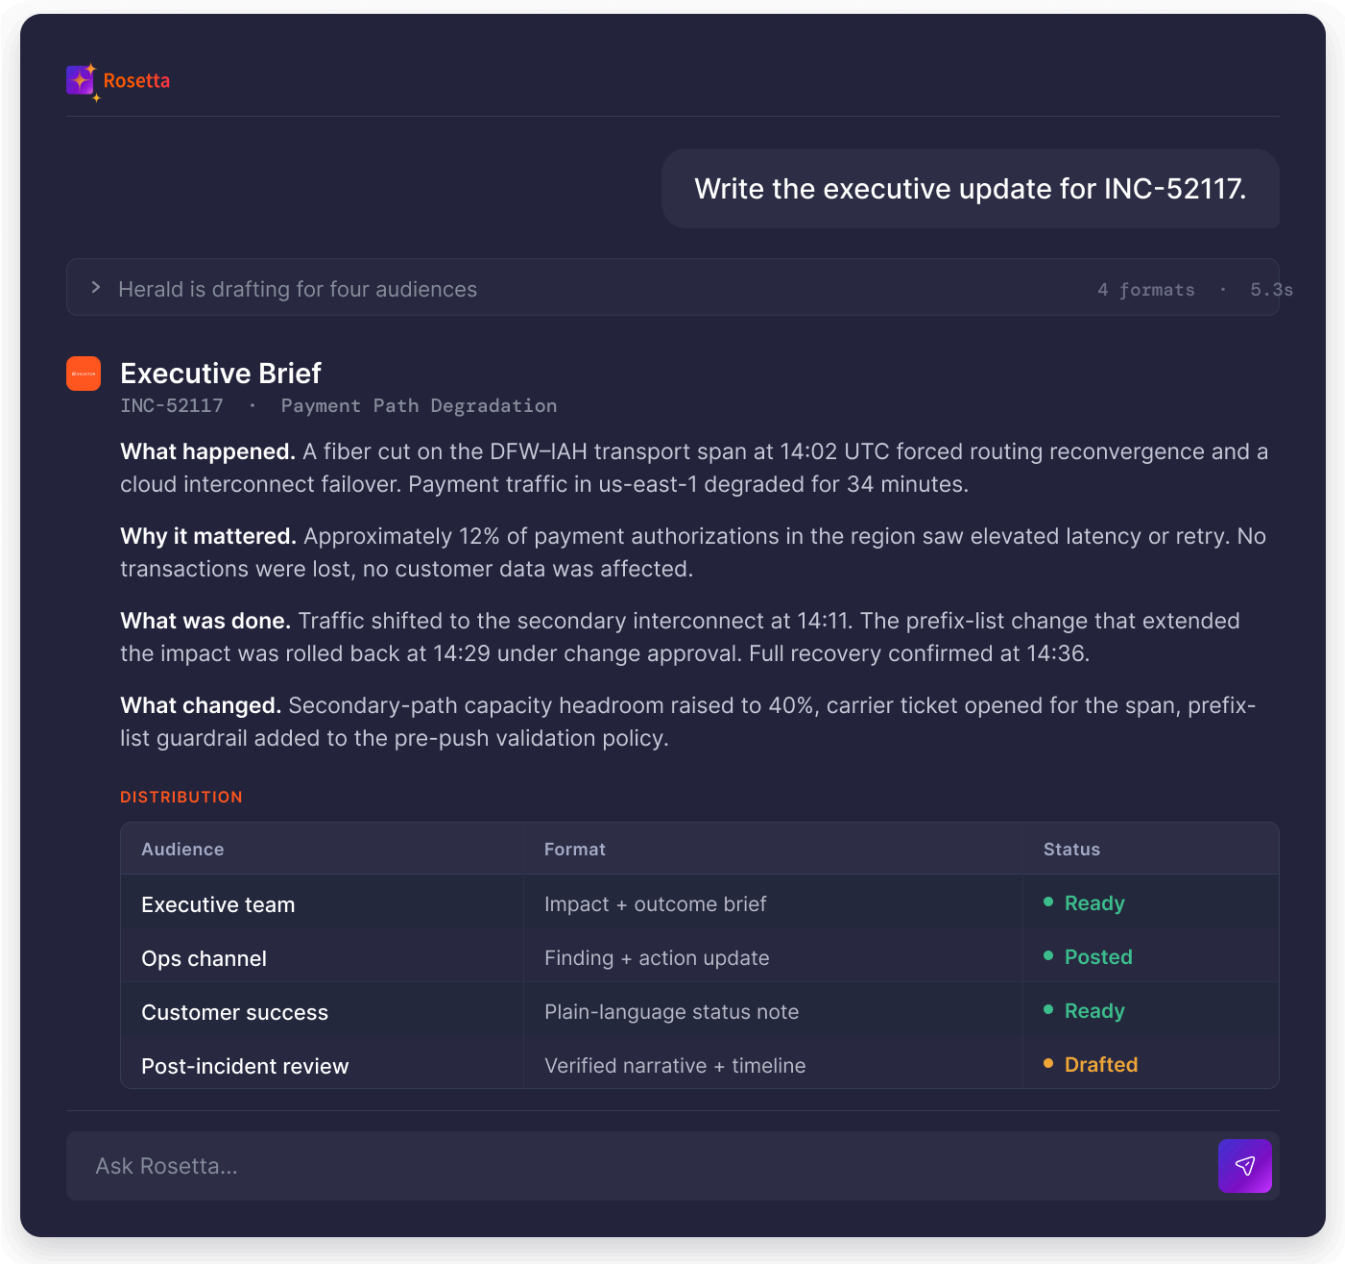

01

Written from the record

What happened, why it mattered, what was done, what changed.

02

Executive language

One paragraph per question, no jargon, safe to forward.

03

Distribution tracked

Four audiences, each with its own format and status.

SCULPTOR

Playbooks become governed agents

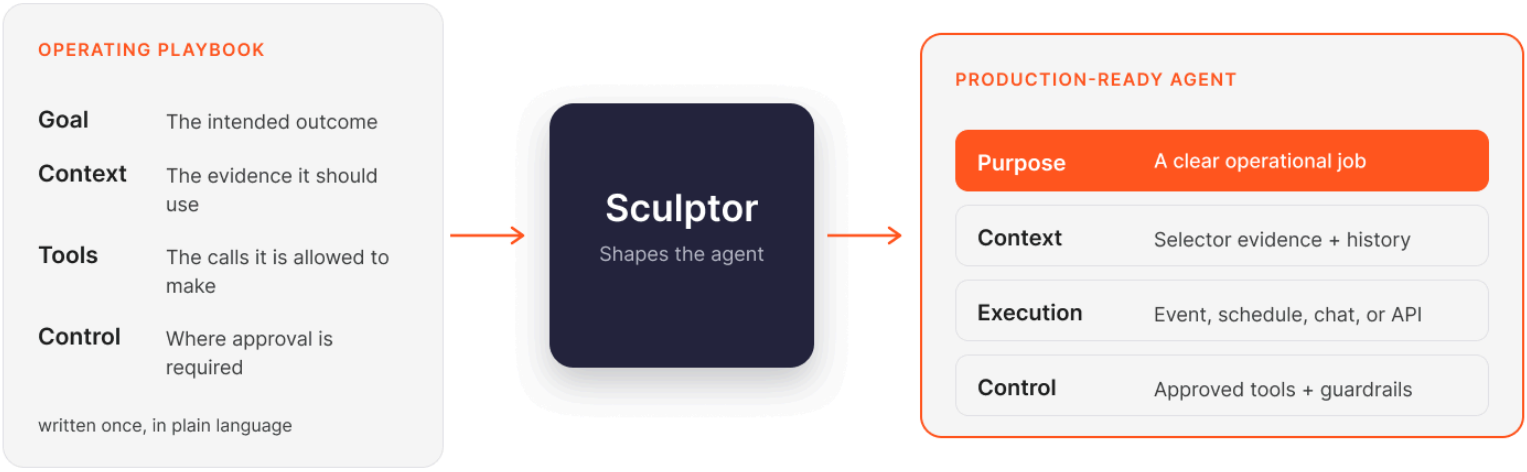

A playbook written once in plain language becomes a production-ready agent with purpose, context, execution, and control.

THE CHALLENGE

The best operational knowledge lives with a handful of engineers and in runbooks nobody has time to maintain. When those engineers are asleep, on leave, or gone, the expertise leaves with them — and the same investigation gets rebuilt from scratch.

THE SELECTOR APPROACH

Sculptor turns a described playbook into a governed agent: goal, context, approved tools, triggers, and the points where a human must approve. The agent runs from chat, an event, a schedule, or an API — with the same guardrails every time.

OPERATIONAL OUTCOME

Expertise becomes platform capability. Coverage grows without growing headcount, and every run is auditable.

A described playbook, running as an agent

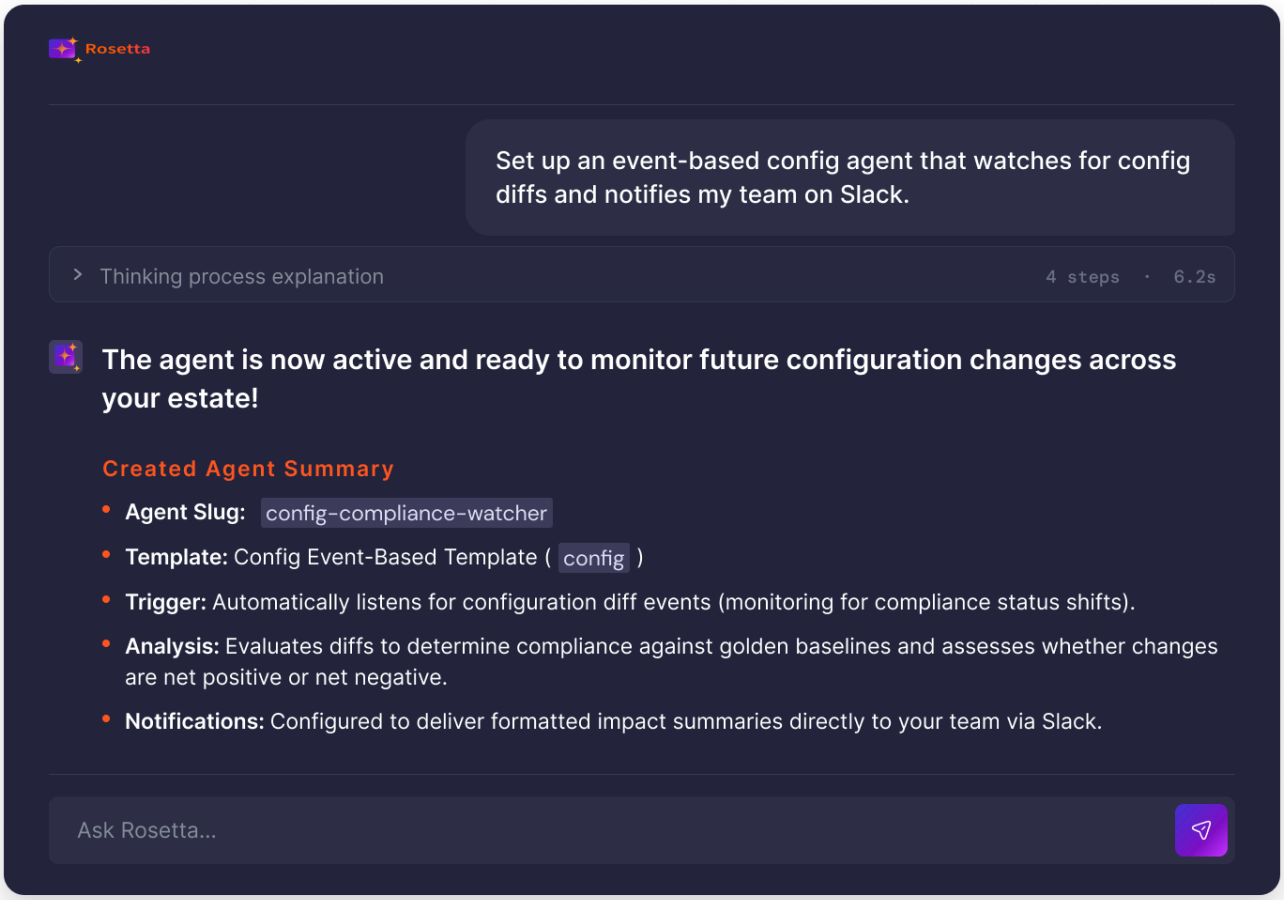

01

One sentence in chat

Watch for config diffs and notify the team on Slack.

02

Shaped into a governed agent

Template, trigger, analysis, and notification path are made explicit.

03

Active and auditable

The agent monitors the estate from creation, every run on record.

DEPLOYMENT

Designed for adoption without disruption

An agentic operating model succeeds only if it fits the environment it inherits. Selector deploys across multi-vendor networks, hybrid cloud, and existing toolchains using SNMP, streaming telemetry, syslog, flow, configuration, cloud APIs, and synthetic probes.

Agents reach outward through MCP and A2A, so established ITSM, collaboration, and monitoring investments keep working. Organizations add agents and raise autonomy incrementally, on their own operational timeline.

ADD AGENTIC OPERATIONS WITHOUT REPLACING YOUR STACK

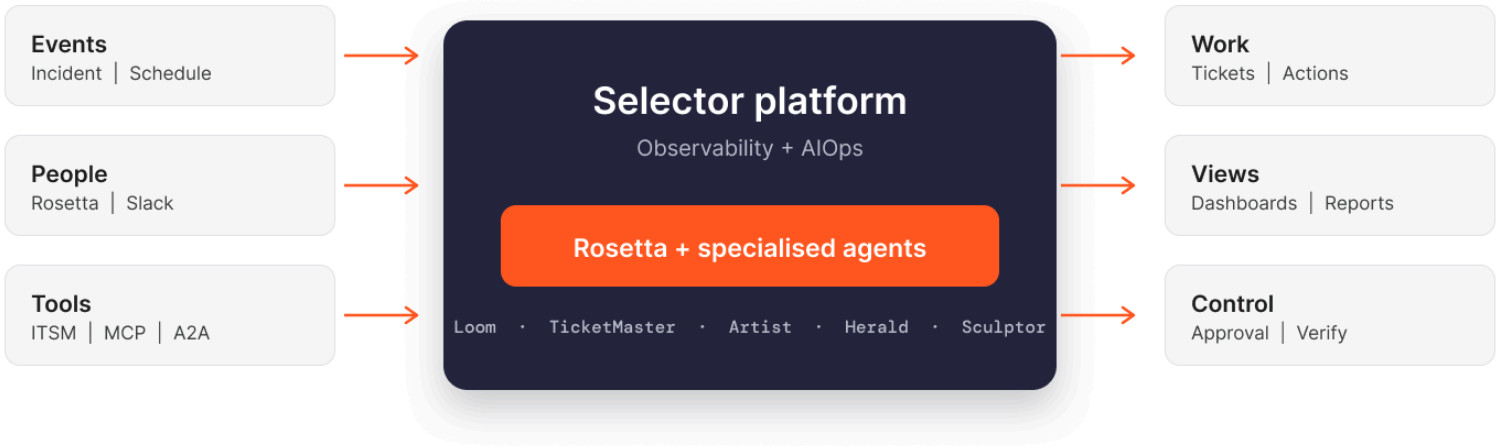

Events, people, and tools in — work, views, and control out. The existing stack stays in place.

Collect

SNMP · streaming telemetry · syslog · flow · config · cloud APIs · synthetics

Integrate

MCP · A2A · ITSM · collaboration · existing monitoring

Govern

approved tools · approval gates · verified outcomes · audit trail

THE SHIFT

From answering to operating

Network complexity will keep outgrowing the teams asked to operate it. The advantage no longer comes from seeing more — it comes from expertise that is available on every shift and evidence that stands behind every action.

Selector gives NetOps, CloudOps, and service assurance teams one intelligence layer and a coordinated set of agents that investigate, execute, and operate.

WHAT TEAMS RECOVER PER P1

250 → 5	8×	1,000	$1M+
Engineers on a P1 98% smaller	Faster resolution MTTR 4 hrs → 30 min	Engineer-hours reclaimed per P1	Recovered per P1 downtime + labor

See Selector's agents on your network

See how Selector's agents investigate, execute, and operate across your network.

selector.ai
sales@selector.ai
selector.ai/request-a-demo