SELECTOR

www.selector.ai

DATASHEET

Selector Rosetta

Agentic investigation for network operations, orchestrated across Selector Foundry agents

Rosetta is the gateway to the Selector platform and the orchestrator behind Selector Foundry. Operators ask a question in plain language, and Rosetta plans the investigation, dispatches the right agents across domains, and returns one connected answer with the evidence, the impact, and the next controlled action.

Rosetta reasons over the operational context Selector builds before an incident begins: normalized telemetry, topology, configuration, change history, and correlated root cause across network, cloud, and infrastructure. Every answer traces back to that evidence, and every action waits on the approvals your team sets.

Where Investigations Slow Down

Operators become the integration layer

Metrics come from one tool, routing state from another, and configuration history from a third. The timeline gets assembled by hand while the incident is still running.

Answers depend on who is on shift

The fastest path to root cause lives with a few senior engineers. When they are unavailable, the same investigation starts over from scratch.

Single-turn AI loses the thread

A chat window that answers one question at a time cannot open an investigation, follow the path across domains, or carry the work to resolution.

Findings arrive without proof

A suggested fix with no evidence attached still has to be verified by hand before anyone is willing to act on it.

How Rosetta Works

01

Ask

The operator asks in plain language, with no query syntax and no switching between tools.

02

Plan

Rosetta interprets the intent and retrieves the relevant operational context.

03

Orchestrate

It dispatches the right agents, often in parallel, and correlates their findings across the path.

04

Resolve

It returns the finding, the evidence, the impact, and the next controlled action for approval.

ONE REQUEST

"edge-rtr-01 in Denver looks unhealthy. Investigate the impact, build me a view, and open a ticket for whoever owns it."

ROSETTA DISPATCHES IN PARALLEL

Rosetta · Loom · Artist · TicketMaster · Herald

One answer: a configuration push on edge-rtr-01 is destabilizing the Denver edge. Rosetta lists the actions taken and recommends rolling back the 02:47 push.

Conversational
Plain-language questions, with no query syntax to learn

Orchestrator
Plans the work and calls the right agents

Multi-turn
One thread holds the full case across every follow-up

Evidence-backed
Every answer traces to live telemetry and history

Delegation in Depth

Rosetta is the single way into every agent and function on the platform. It hands each request down through platform agents, domain specialists, and MCP tools, then brings one answer back.

ORCHESTRATOR **Rosetta**	Receives the question, conversation history, and customer-specific context. Interprets intent, decides which agents and domains the question needs, delegates the work, and synthesizes one evidence-backed answer.
PLATFORM AGENTS **Selector Foundry**	**Loom** — Correlation and root cause **TicketMaster** — ITSM tickets and updates **Artist** — Operational views on demand **Herald** — Briefs, updates, and reports **Sculptor** — Playbooks become agents
DOMAIN SPECIALISTS **Focused agents**	Specialists for areas such as routing, firewalls, synthetics, and cloud infrastructure. Each runs a defined job with a constrained toolset, which keeps context lean and results precise.
TOOLS **MCP layer**	MCP tools read and act on Selector data and services, such as BGP tables, latency metrics, and device health. External tools and agents connect at the same layer over MCP and A2A.

Context Engineering

- **Operational context.** Selector's AIOps platform normalizes every signal and establishes root cause before Rosetta reasons over it.
- **Organizational context.** Your language, conventions, and incident history shape how Rosetta reads each question.
- **Agent and tool descriptions.** Precise descriptions of each agent and MCP tool let Rosetta pick the right one without hardcoded logic.
- **Case memory.** Rosetta keeps the case open across follow-ups, including what was checked and ruled out.
- **Traceable reasoning.** Click any claim to see the telemetry or configuration diff behind it.

Agents You Own

- **Your data, your environment.** Agents run where your data already lives, so nothing goes to an outside service.
- **Your choice of LLM.** Rosetta is LLM-agnostic and plans and acts over the model you select.
- **Built and versioned by your team.** Author agents, commit them to your Git, and test them against past incidents before promotion.
- **Governed action.** Agents run under your RBAC, with approval gates on controlled actions and every run on record.
- **Open integration.** Existing ITSM, collaboration, and monitoring tools join through MCP and A2A.

Autonomy at Your Pace

Teams decide how much autonomy each workflow gets and raise it one workflow at a time.

OPERATOR-LED

Ask and approve

Operators direct investigations through Rosetta, and recommended actions wait for approval before they run.

EVENT-LED

Start before the ask

An incident, schedule, API call, or another agent opens the workflow, so the investigation is underway when an operator joins.

CLOSED-LOOP

Act within guardrails

Agents act through approved tools, verify outcomes against live state, and stay inside the limits the organization sets.

See Rosetta on your network

Ask a question about your own environment and watch Rosetta run the investigation and return a verified answer.

Request a demo

www.selector.ai | sales@selector.ai