SELECTOR

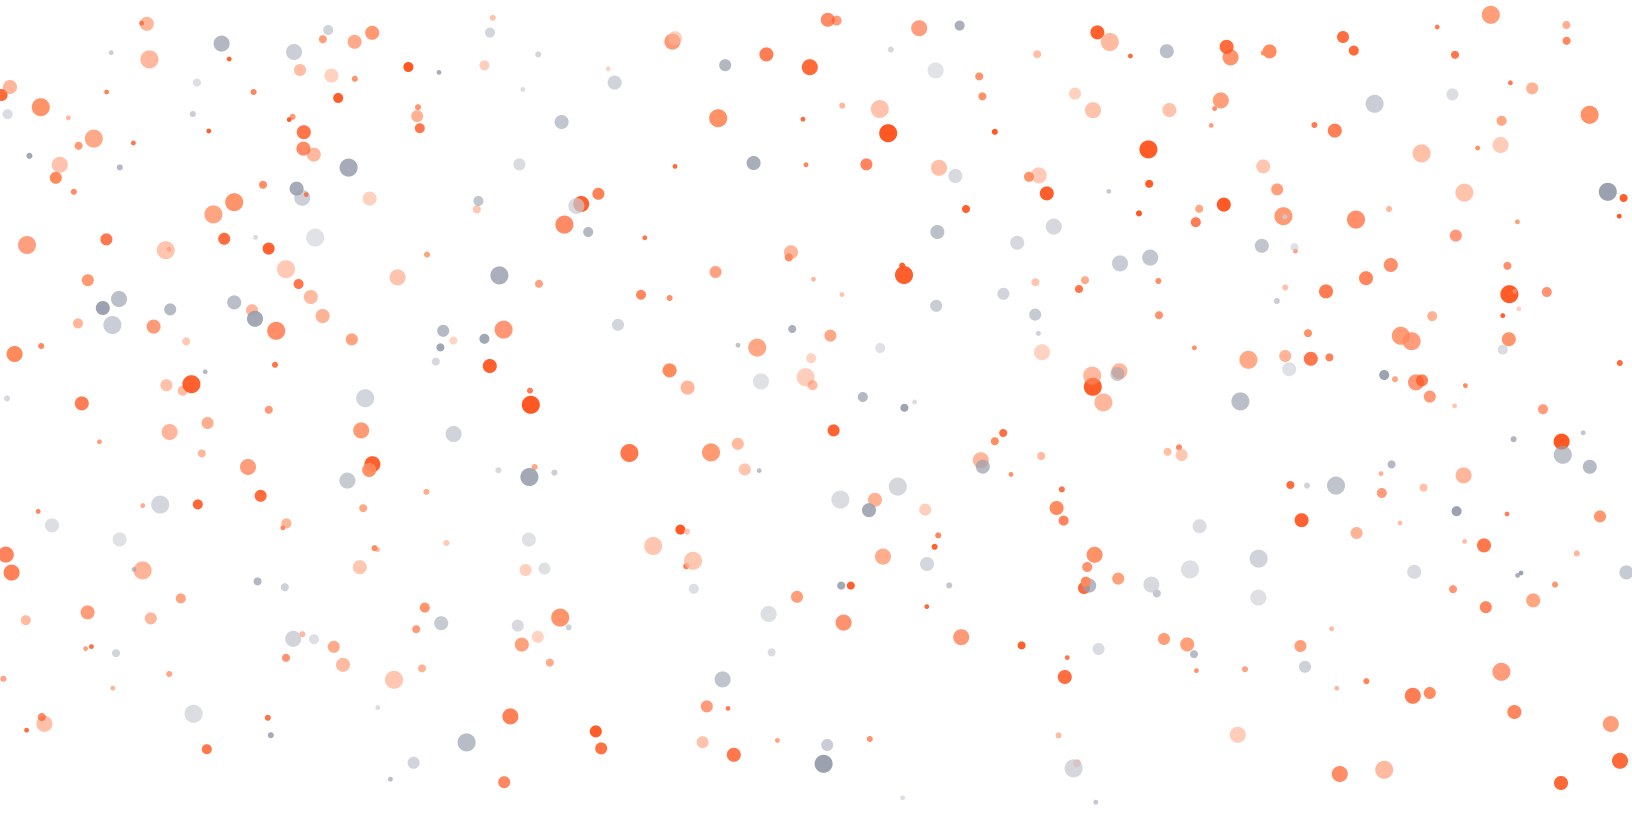
WWW.SELECTOR.AI

Build vs. Buy

Selector **vs.** Do It Yourself

Build vs. Buy for Observability and Event Intelligence

October 2026



How Each Platform Builds Operational Context

Selector collects raw telemetry and operational context together from every vendor, normalizes and enriches it before any model runs, and establishes causation with purpose-built ML before agents act. A typical DIY approach combines open-source and cloud components, such as Telegraf or gNMIc collectors, Kafka, Prometheus or InfluxDB, OpenSearch or Elasticsearch, Grafana, custom Python correlation logic, and LLM APIs, integrated and maintained by an internal team.

Selector		DIY Stack	
Foundry agents	Rosetta, Loom, TicketMaster, Artist, Herald, Sculptor; MCP and A2A.	**LLM scripts**	Custom agents or chat interfaces built on LLM APIs.
Collaboration	Natural language, dashboards, ChatOps, ITSM workflows, query-backed APIs.	**Dashboards and alerts**	Grafana dashboards and Alertmanager rules.
Knowledge	Anomaly detection, correlation, causation, forecasting, knowledge graph.	**Custom correlation**	Python rules, scripts, and models written in-house.
Data Hypervisor	ML schema inference, enrichment, ELT raw preservation, routing.	**Data stores**	Prometheus, InfluxDB, OpenSearch, and Kafka.
Collection	Horizontal engines near the source; push and pull; independent scaling.	**Collectors**	Telegraf, gNMIc, Logstash, flow collectors, custom scripts.
Sources	Any vendor's devices, cloud, CMDB, IPAM, ITSM, tools, and files.	**Sources**	Whatever the team integrates and keeps working.

Where a DIY stack falls short architecturally:

- **Every layer is the team's to build and maintain:** Collection, normalization, storage, correlation, and agents each need engineering, upgrades, and on-call ownership. Selector delivers all of them as one maintained platform.
- **Causation is the hardest part to build:** Dashboards and alert rules are straightforward; deterministic, topology-aware causation across vendors is not. Selector's patented models deliver it.
- **LLM-first shortcuts are fragile:** LLM scripts over raw data can hallucinate and give different answers to the same question. Selector uses LLMs only to explain deterministic findings.

Copyright Selector Software, Inc.

USE CASE 1: OBSERVABILITY

Multivendor Data Collection

Open-source collectors can gather most data types, but each needs integration and upkeep. Selector collects everything into one pipeline as a maintained product.

DATA TYPE	SELECTOR	DIY STACK
Metrics	SNMP, gNMI and model-driven streaming telemetry, Prometheus-compatible sources, REST APIs.	Telegraf, gNMIc, and exporters; each integration built and maintained in-house.
Logs	Syslog from any device, infrastructure, or application, structured by ML; Splunk, Logstash, TACACS.	Logstash or Fluent Bit with authored parsers.
Flow	NetFlow, sFlow, IPFIX from devices or collectors, enriched with application, owner, and criticality.	Open-source flow collectors such as pmacct or GoFlow.
Routing	BGP, BMP/OBMP, and IGP state ingested from customer devices.	BMP collectors such as OpenBMP, integrated in-house.
Configuration	Running and intended config via NETCONF and CLI; drift correlated with impact.	Oxidized or scripts for backups. **Drift correlated with impact must be built.**
Topology	From SNMP, CDP/LLDP, configuration, and telemetry, across all vendors; digital twin.	**Topology modeling must be built.**
Context	ServiceNow CMDB, NetBox, Nautobot, Infoblox; owner, site, circuit, criticality.	Custom CMDB and NetBox integrations.
Hybrid & multicloud	AWS, Azure, Google Cloud APIs, flow logs, transit gateways, VPCs, correlated with on-prem.	Cloud exporters and APIs per provider.
Digital experience	ICMP and UDP probes, ThousandEyes, traceroute, custom tests.	Blackbox exporters and custom probes.
Files / streams	Kafka, CSV, Excel, Google Sheets.	Kafka and scripts.

DIY gaps

- Every collector, integration, and upgrade is owned by the internal team.
- Topology and dependency modeling must be built from scratch.
- Coverage drifts as vendors, firmware, and APIs change.

Selector value

- Streaming telemetry, flow, routing, configuration, and context from any vendor in one pipeline.
- Routing changes tied to the services they affect, from BMP and IGP state on the customer's own devices.
- Ingestion from files and streams for data with no API.

Copyright Selector Software, Inc.

USE CASE 1: OBSERVABILITY

Data Normalization and Enrichment

Selector normalizes and enriches every signal before any model runs. In a DIY stack, every parser, schema, and enrichment rule is written and maintained by the team.

CAPABILITY	SELECTOR DATA HYPERVISOR	DIY STACK
Schema inference	ML infers structure from free-form logs, syslog, traps, and CLI output; no parsing templates.	**Hand-written parsers per source.**
Entity resolution	Named-entity and clustering models group events across vendor and format variations.	**Must be built.**
Enrichment	Topology, dependency, owner, site, maintenance state, and criticality added before analysis.	Custom enrichment scripts.
Data fidelity	ELT: raw data lands with original timestamps; transformation follows.	Depends on storage design and retention budget.
Noise suppression	Duplicate, flapping, and irrelevant signals removed before the intelligence layer.	Alertmanager rules and inhibition.
Business logic	Maintenance windows, provider notices, suppression, and criticality in YAML via DevOps workflows.	Rules files and scripts.
Customization	S2ML tunes enrichment, labeling, noise rules, prioritization, and LLM output.	Unlimited, at the cost of engineering time.
Data access	Every view is a query and can be exposed as an API endpoint; MCP server.	Multiple query languages across stores.

DIY gaps

- Parsers, schemas, and enrichment logic are written and maintained by hand.
- Data spreads across several stores with different query languages.
- Raw-data replay of network state must be designed and built.

Selector value

- Learned normalization: no parsing templates or rules to maintain.
- Raw records retained with original timestamps for replay and re-analysis.
- One enrichment layer applies topology, ownership, maintenance, and criticality to every signal.

Copyright Selector Software, Inc.

USE CASE 1: OBSERVABILITY

AI-Powered Operational Analytics

Selector models every time series individually and correlates across domains automatically. A DIY stack starts with thresholds and requires custom data science for anything more.

CAPABILITY	SELECTOR	DIY STACK
Anomaly detection	**Self-supervised, statistical, and regression models.** A model per time series, retrained every 30 minutes, with seasonality.	Static thresholds or custom models.
Log analysis	**Small language models.** Log semantics interpreted, clustered across vendors, and turned into structured events.	Search and authored rules.
Prediction	**Forecasting and predictive models.** Capacity forecasting and optical transceiver failure prediction.	**Requires custom models.**
Cross-domain analysis	**One model across domains.** Correlation runs automatically across network, cloud, and application signals.	**Requires custom correlation logic.**
Natural language	**Rosetta.** Plain-language questions answered with every claim linked to source data.	LLM scripts over data, built in-house.

DIY gaps

- Anomaly detection must be built, tuned, and retrained.
- Cross-domain correlation requires custom models.
- Prediction requires data science investment.

Selector value

- Per-series models retrained every 30 minutes detect anomalies no threshold would catch.
- Automatic correlation across network, cloud, and application signals in one model.
- Optical failure prediction and capacity forecasting before impact.

Copyright Selector Software, Inc.

USE CASE 2: EVENT INTELLIGENCE

Event Correlation and Root Cause

Correlation and causation are the hardest parts of a DIY stack. Selector delivers deterministic, topology-aware causation as a product.

CAPABILITY	SELECTOR	DIY STACK
Inputs	Alerts and the raw telemetry, logs, flow, and configuration behind them, correlated together.	Alerts from Alertmanager and other tools.
Correlation method	Contextual embeddings and clustering by shared devices, interfaces, IP addresses, topology, and time.	Rules and scripts written in-house.
Noise reduction	Up to 99.9% of irrelevant noise removed before correlation.	Inhibition rules and deduplication.
Root cause	Deterministic causation from topology and domain logic: reboot, then link failure, then BGP flap.	**Requires building causal models.**
Symptom vs. cause	Downstream effects labeled as consequences, not separate faults, across every domain.	**Must be engineered.**
Change correlation	Configuration diffs and change events correlated directly with degradation.	**Must be built.**
Evidence and replay	Every finding links to source data; Network DVR replays past network state.	**Depends on in-house design.**

Selector example

The Loom agent resolved 214 alerts across six tools to a single fiber cut on one long-haul span, labeled routing reconvergence and cloud failover as consequences, and flagged the secondary path at 78% utilization as the next risk.

The DIY reality

A Grafana alert storm during a fiber cut produces many alerts with no named cause. Building logic that names the cut and labels the rest takes sustained engineering.

DIY gaps

- Correlation is rules and scripts unless causal models are built.
- Symptom-versus-cause labeling must be engineered.
- LLM-based analysis can return different answers for the same data.

Selector value

- Alert floods collapse into one incident with a named cause, consequences, and next risk.
- Telemetry, logs, flow, and configuration correlated with alerts in one model.
- Every finding is traceable and reproducible with Network DVR.

Copyright Selector Software, Inc.

USE CASE 2: EVENT INTELLIGENCE

Agents, ITSM, and Governed Action

FUNCTION	SELECTOR FOUNDRY	DIY STACK
Investigation	**Rosetta.** Interprets plain-language intent, dispatches agents across domains, returns one cited thread.	LLM chat scripts over data.
Correlation / RCA	**Loom.** Connects time, topology, changes, and domains into one incident with a root cause.	**Must be built.**
ITSM	**TicketMaster.** Bi-directional ServiceNow sync across the lifecycle, with per-customer field mapping and incident-aware deduplication.	Custom ServiceNow integration.
Communications	**Herald.** Executive briefs, ops updates, and post-incident reports from the verified record.	Custom notification scripts.
Custom agents	**Sculptor.** Turns a described playbook into a governed agent with goal, tools, triggers, and approval points.	LLM agent frameworks. **Governance, approvals, and audit must be built.**
Action	Agents investigate, act, and verify within operator-set limits on high-confidence findings.	Ansible or scripts. **Verification and guardrails must be built.**
MCP	Production MCP server with OAuth2, granular RBAC over insights, recommendations, and actions.	**MCP servers and security must be engineered.**
Deployment	SaaS, on-premises, customer cloud, or air-gapped.	Wherever the team runs it.

DIY gaps

- Agent governance, approvals, and audit must be built.
- MCP servers, RBAC, and security must be engineered.
- Ongoing ownership competes with operations work.

Selector value

- Sculptor builds governed agents from the team's own playbooks, acting only on verified findings.
- MCP server shares insights and governed actions with other agent platforms.
- Every Foundry agent is available now, in every deployment model.

Copyright Selector Software, Inc.

SELECTOR

Functionality Gained with Selector

Capabilities a team gets on day one with Selector instead of building and maintaining them in-house.

1. One maintained platform

Collection, normalization, storage, correlation, and agents delivered and maintained as one product, instead of a stack the team must own.

2. Deterministic diagnosis

Root cause comes from deterministic topology and domain models, so the same data yields the same answer.

3. Detection beyond thresholds

Per-series models retrained every 30 minutes find conditions no static threshold or rule would catch.

4. Streaming and routing depth

gNMI streaming telemetry and BMP/OBMP and IGP analytics from the customer's own devices, plus optical failure prediction.

5. Full-fidelity history

Raw data preserved with original timestamps; Network DVR replays past network state so findings can be verified and re-run.

6. No parsers to maintain

Learned normalization infers structure across vendors and formats, replacing hand-written parsers and enrichment scripts.

7. Your playbooks as agents

Sculptor turns the team's own runbooks into governed agents with approved tools, triggers, approval points, and audit.

8. Governed action through MCP

Selector's MCP server exposes insights, recommendations, and actions to other agent platforms under OAuth2 and RBAC.

9. Custom outcomes with engineers

Forward deployed engineers build integrations, data models, views, and agents on the customer's own data.

10. Reuse what you have

Selector ingests from existing collectors, Kafka, Prometheus, and InfluxDB, so prior DIY investment carries forward.

Copyright Selector Software, Inc.

Selector Patents Driving Innovation

The techniques that define the Selector platform are covered by Selector's U.S. patent portfolio. Patent records are available on request.

CAPABILITY	WHAT SELECTOR'S PATENTED APPROACH DELIVERS
Root Cause Analysis	Deterministic causation across every domain.
Reproducible Investigations	Network DVR replays network state so an investigation can be re-run.
Log Intelligence	Structured signal from any vendor's syslog without parsing templates.
Conversational Investigation	Learns each customer's names for sites, services, and devices without manual labeling.
Custom Outcomes on Demand	Declarative query chains for drill-downs and customer-specific views.
Operational Context	Maintenance-window-aware alerting and reporting.
Optical Failure Prediction	Flags failing optics before they cause an outage.
Capacity Planning	Forecasts capacity limits before they affect services.
Scale	Efficient storage and querying of millions of data points per minute.

Copyright Selector Software, Inc.

Selector vs. Do It Yourself at a Glance

CAPABILITY	SELECTOR	DIY STACK
OBSERVABILITY		
Data Collection	Any vendor, including gNMI, BMP/IGP, flow, config, and file sources, in one pipeline.	Possible, but every integration is owned in-house.
Normalization	Learned schema inference; ELT raw preservation.	Hand-written parsers.
Detection	Per-series models retrained every 30 minutes.	Thresholds or custom models.
Cross-Domain Analysis	Automatic correlation in one model.	Must be built.
EVENT INTELLIGENCE		
Correlation	Alerts plus the raw telemetry behind them.	Rules and scripts.
Root Cause	Deterministic causation with symptom labeling.	Must be built.
Reproducibility	Network DVR replay; source-linked evidence.	Depends on design.
Custom Agents	Sculptor builds governed agents from playbooks.	Built in-house without governance.
MCP	Server for insights, recommendations, and governed actions.	Must be engineered.
Deployment	SaaS, on-premises, or air-gapped.	Self-managed.

The Bottom Line

Building in-house gives full control and works well for dashboards and alerting, but collection upkeep, normalization, correlation, causation, and agent governance become long-term engineering commitments.

Selector delivers those capabilities as a maintained platform, with patented causation, raw-data replay, governed agents, and forward deployed engineers, so the team spends its time operating the network instead of building tools.

30 min	99.9%	8x
Model retraining cycle	Noise reduction	Faster resolution per P1

See Selector on your own data

Request a Demo

www.selector.ai | sales@selector.ai

Copyright Selector Software, Inc.