Selector Collection Service

A horizontal ingestion architecture for full-stack operational context

Selector Collection Service eliminates a structural limitation at the core of traditional observability: the separation of telemetry from operational context. In most platforms, telemetry and the context needed to interpret it are collected in different systems, processed through separate pipelines, and correlated only after an incident has already begun. Selector takes a different approach. Its collection layer sits horizontally at the bottom of the platform, allowing organizations to ingest raw telemetry and contextual data from across network, cloud, infrastructure, application, and operational systems into a shared model for cross-domain correlation, root-cause analysis, and action. The result is a collection architecture that scales by workload, adapts to hybrid environments without re-architecting, and begins delivering value from existing tools on day one.

Why Traditional Collection Breaks Down

Traditional monitoring and observability platforms were architected around a vertical data model. Logs are routed into one pipeline, metrics into another, and traces into a third, while topology, CMDB, configuration state, and operational events often remain in separate tools altogether. Each pipeline operates with its own schema, storage format, and query engine—optimized for depth within a single data type, not breadth across types. This architecture may simplify point deployment, but it creates a major operational gap: the data needed to understand an incident is collected in pieces, stored in different formats, and analyzed in isolation.

- **Type-Specific Pipelines Preserve Data Silos**
  Correlation usually happens later through predefined rules, tags, or manual investigation, rather than during ingestion.
- **Operational Context is Structurally Absent**
  Logs, metrics, configurations, topology, and change signals often live in different tools, slowing root-cause analysis and increasing operator effort.
- **Scaling one Workload Forces Overbuilding the Entire Stack**
  SNMP polling, high-frequency streaming telemetry, and log volume often require scaling beyond the specific collection tier under load, increasing infrastructure overhead and reducing operational efficiency.

How Selector Collection Service Works

Selector Collection Service was built to address that limitation at the foundation. Instead of treating each telemetry type as a separate operational world, Selector uses a horizontal collection model powered by self-supervised and unsupervised learning at the ingestion layer. The collection layer spans the bottom of the platform so ingested signals—regardless of source, type, or domain—flow into a shared model where structure, relationships, and context are learned automatically, without predefined schemas or manual tagging.

Horizontal placement: Collection engines are deployed close to the data source—in data centers, branch locations, cloud regions, or network segments—reducing latency between signal generation and ingestion and maintaining collection continuity in distributed and hybrid environments.

Push and pull ingestion: Selector supports direct device collection via SNMP, gNMI, syslog, and CLI alongside upstream-tool ingestion from platforms such as Splunk, Prometheus, and Kafka. Both push-based streaming and pull-based polling are supported, allowing teams to consolidate existing tool investments into a single collection pipeline without forklift replacement

Context-rich ingestion: Collection is not limited to time-series data and logs. Selector simultaneously ingests topology and dependency graphs, CMDB and inventory records, configuration state, routing policy, interface aliases, and change signals. This ensures that every telemetry record enters the shared intelligence layer with the operational context required to correlate symptoms to causation — not just to detect that something changed, but to understand what it affected and why.

Scale by workload: Because the collection layer is horizontal, each ingestion workload scales independently. If SNMP polling volume increases, capacity is added to that tier without provisioning additional infrastructure across the rest of the platform. The same applies to high-frequency gNMI streaming, NetFlow export, syslog ingestion, or API-based pull collection — each scale on its own terms.

What can Selector Ingest

Data type	Common Ingestion Sources
Metrics / time-series	CPU, memory, interface counters, and custom device and service metrics collected directly via SNMP, gNMI, and streaming telemetry, or pulled from existing monitoring systems including Prometheus-compatible sources and REST APIs. Supports both high-frequency polling and model-driven telemetry using YANG data models.
Logs / syslog	Syslog streams from network devices, infrastructure, and application layers, including ingestion from upstream aggregation platforms such as Splunk. ML-based clustering groups similar log patterns to surface anomalies and reduce noise before signals reach the correlation layer.
Events / alerts / anomalies	Operational events, SNMP traps, syslog-derived events, generated incidents, and anomaly signals from monitoring, ITSM, and observability tools. Events are normalized on ingestion and enriched with topology and configuration context for correlation and root-cause analysis.
Flow and network telemetry	Flow records including NetFlow, sFlow, and IPFIX for traffic analysis and anomaly detection, alongside high-frequency streaming network telemetry via gNMI for near-real-time KPI monitoring. Flow data is integrated into the unified analytics layer alongside metrics, logs, and topology.
Configuration state	Running and intended device configuration ingested via Netconf/CLI, used to track configuration state, detect drift, and correlate change events directly with performance degradation or incidents. Configuration context enables the platform to answer not just what changed, but what that change affected across the operational graph.
Topology / dependency context	Device relationships, service dependency graphs, and digital twin context built and continuously updated from SNMP, configuration data, and telemetry. Topology is stored in a Knowledge Graph and used to enrich every ingested signal with dependency relationships — enabling the platform to trace a symptom to its origin across the full service and infrastructure chain
Metadata and CMDB context	ServiceNow CMDB records, inventory data, ownership attributes, site and region labels, interface aliases, routing policy context, and knowledge base entries. This metadata enriches every telemetry record with the operational relationships needed to identify service impact and assign causation accurately.
Third-party tool telemetry	APM, log, metric, and trace data pulled from existing monitoring platforms including Splunk, NetBox, InfluxDB, Kafka, ThousandEyes, Prometheus, SolarWinds, LogicMonitor, and ServiceNow. Selector connects to 300+ pre-built integrations, enabling teams to unify signals from their current toolchain without forklift replacement or schema re-instrumentation.

Operational Benefits

By changing collection at the architectural level, Selector improves how teams scale, investigate, and operationalize observability across hybrid environments. The result is a collection foundation that reduces fragmentation, accelerates time to value, and supports more effective cross-domain operations.

- **One collection architecture across domains:** A single horizontal collection layer replaces the need for separate vertical stacks per domain. Network, infrastructure, cloud, application, and edge signals all feed the same shared intelligence layer, eliminating the domain silos that force manual correlation during incidents.
- **Faster time to value:** Selector ingests directly from tools already in production — Splunk, Prometheus, SolarWinds, ServiceNow, and 300+ others — without requiring schema changes, re-instrumentation, or replacement of existing infrastructure. Cross-domain analysis begins from existing data sources on day one.
- **Better root-cause analysis:** Because telemetry and operational context — topology, CMDB, configuration state, and change signals — are ingested together into a shared intelligence layer, correlation happens during ingestion rather than after. The platform connects symptoms to causation across domains, not just within a single telemetry type.
- **Operational flexibility:** The collection layer is source-agnostic and schema-agnostic. Teams can expand telemetry coverage, swap upstream tools, or migrate to new protocols without redesigning the ingestion architecture.

Built for Hybrid, High-Throughput Operation

Selector Collection Service gives operations teams a structurally different starting point: a single, horizontal ingestion layer that unifies raw telemetry with operational context across network, cloud, infrastructure, application, and edge domains simultaneously. By positioning collection at the foundation of the platform rather than as a series of domain-specific pipelines, Selector ensures that every signal entering the system — whether a syslog from a core router, a CMDB record from ServiceNow, a flow export from a peering edge, or a configuration change from a CI/CD pipeline — is immediately available for cross-domain correlation, causal root-cause analysis, and AI-driven action. The result is an ingestion architecture that is operationally resilient, scales to hybrid environments without re-instrumentation, and forms the data foundation on which every Selector capability depends.

Build a stronger foundation for operations

Explore how Selector enables teams to break down silos, gain broader context, and act faster across modern hybrid environments. Book a demo to learn more today.

Book a Demo

www.selector.ai | sales@selector.ai