DetectFlow

AI-powered detection pipelines for Kafka

Run thousands of Sigma detections on Apache Kafka & Flink streams at millisecond speed, before data reaches your SIEM.

0.005s
DETECTION LATENCY

12,000+
RULES PER PIPELINE

600K+
DETECTION LIBRARY

DetectFlow

THE PROBLEM & THE SOLUTION

🔴 THE PROBLEM

- SIEM-based detection is delayed by ingestion, indexing, and storage cycles
- Log volume growth makes "store everything" expensive and unpredictable
- ETL tools route data but aren't built to execute security detections at scale
- Most SIEMs degrade past ~500 custom rules

✅ THE OUTCOME WITH DETECTFLOW

- **Faster detection:** stream-speed tagging for real-time response
- **Lower SIEM load:** filter and route low-value noise before ingestion
- **Higher fidelity:** normalize and enrich events before rule execution
- **No rip-and-replace:** layers on existing Kafka infrastructure

How it works

Raw log events land in Kafka topics as usual. DetectFlow syncs Sigma rules from multiple sources and runs them through a five-stage Flink processing pipeline:

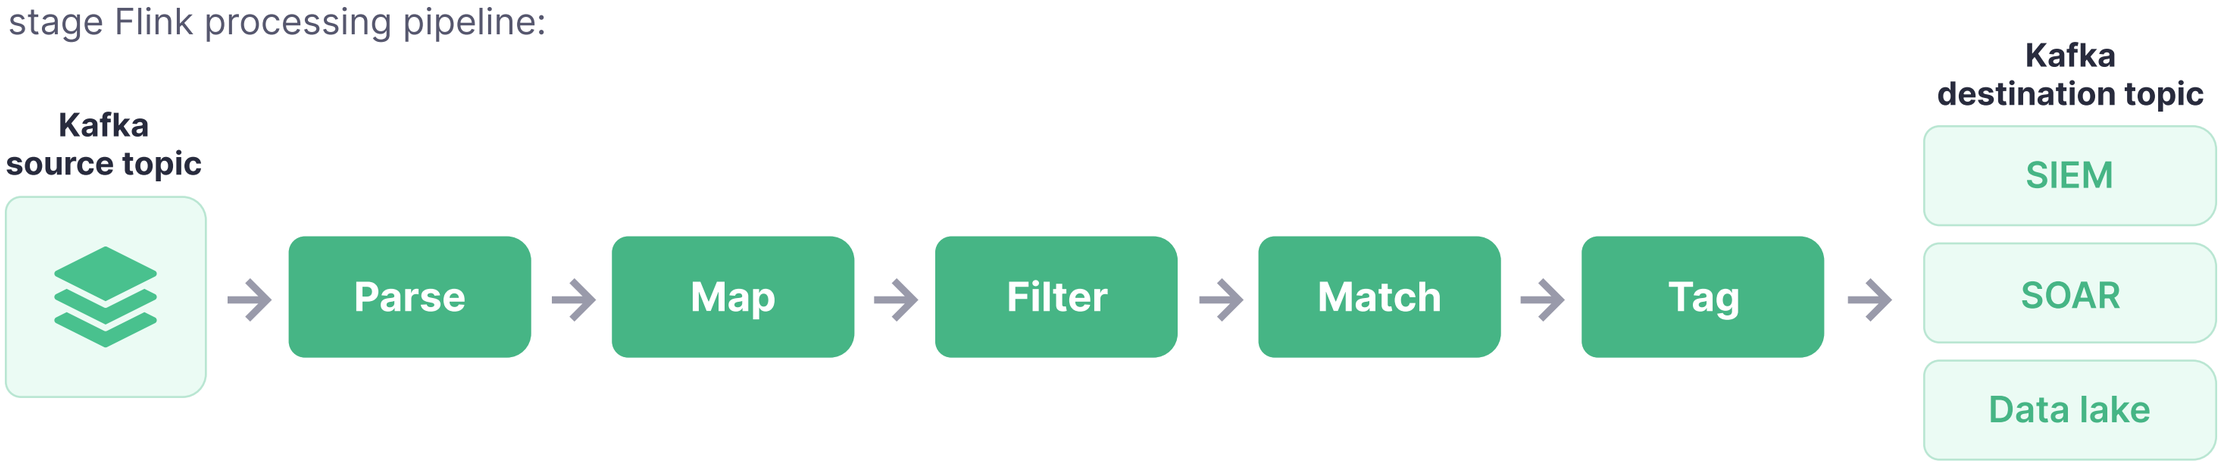

Tagged events carry detection metadata: rule ID, title, severity, and MITRE ATT&CK technique IDs.
Topic chaining lets one pipeline's output feed the next for layered detection workflows.

Rule sources

DetectFlow pulls Sigma detection rules from multiple sources — no manual import required:

- **SOC Prime cloud repo** – bidirectional sync every 5 min, 600K+ rules
- **Local repositories** – for air-gapped and classified environments
- **SigmaHQ** on GitHub
- **Splunk, Elastic, Microsoft Sentinel** GitHub repos

Key capabilities

In-stream detection
Execute thousands of Sigma rules on live streams. Sub-second latency. 12K+ rules per pipeline.

Pipeline observability
Real-time dashboard: latency, lag, throughput. RBAC and audit logs with CSV export.

AI parsing & smart filtering
AI-assisted field mapping via Uncoder AI. Filter low-value events before SIEM ingestion. Kafka replay for backtesting.

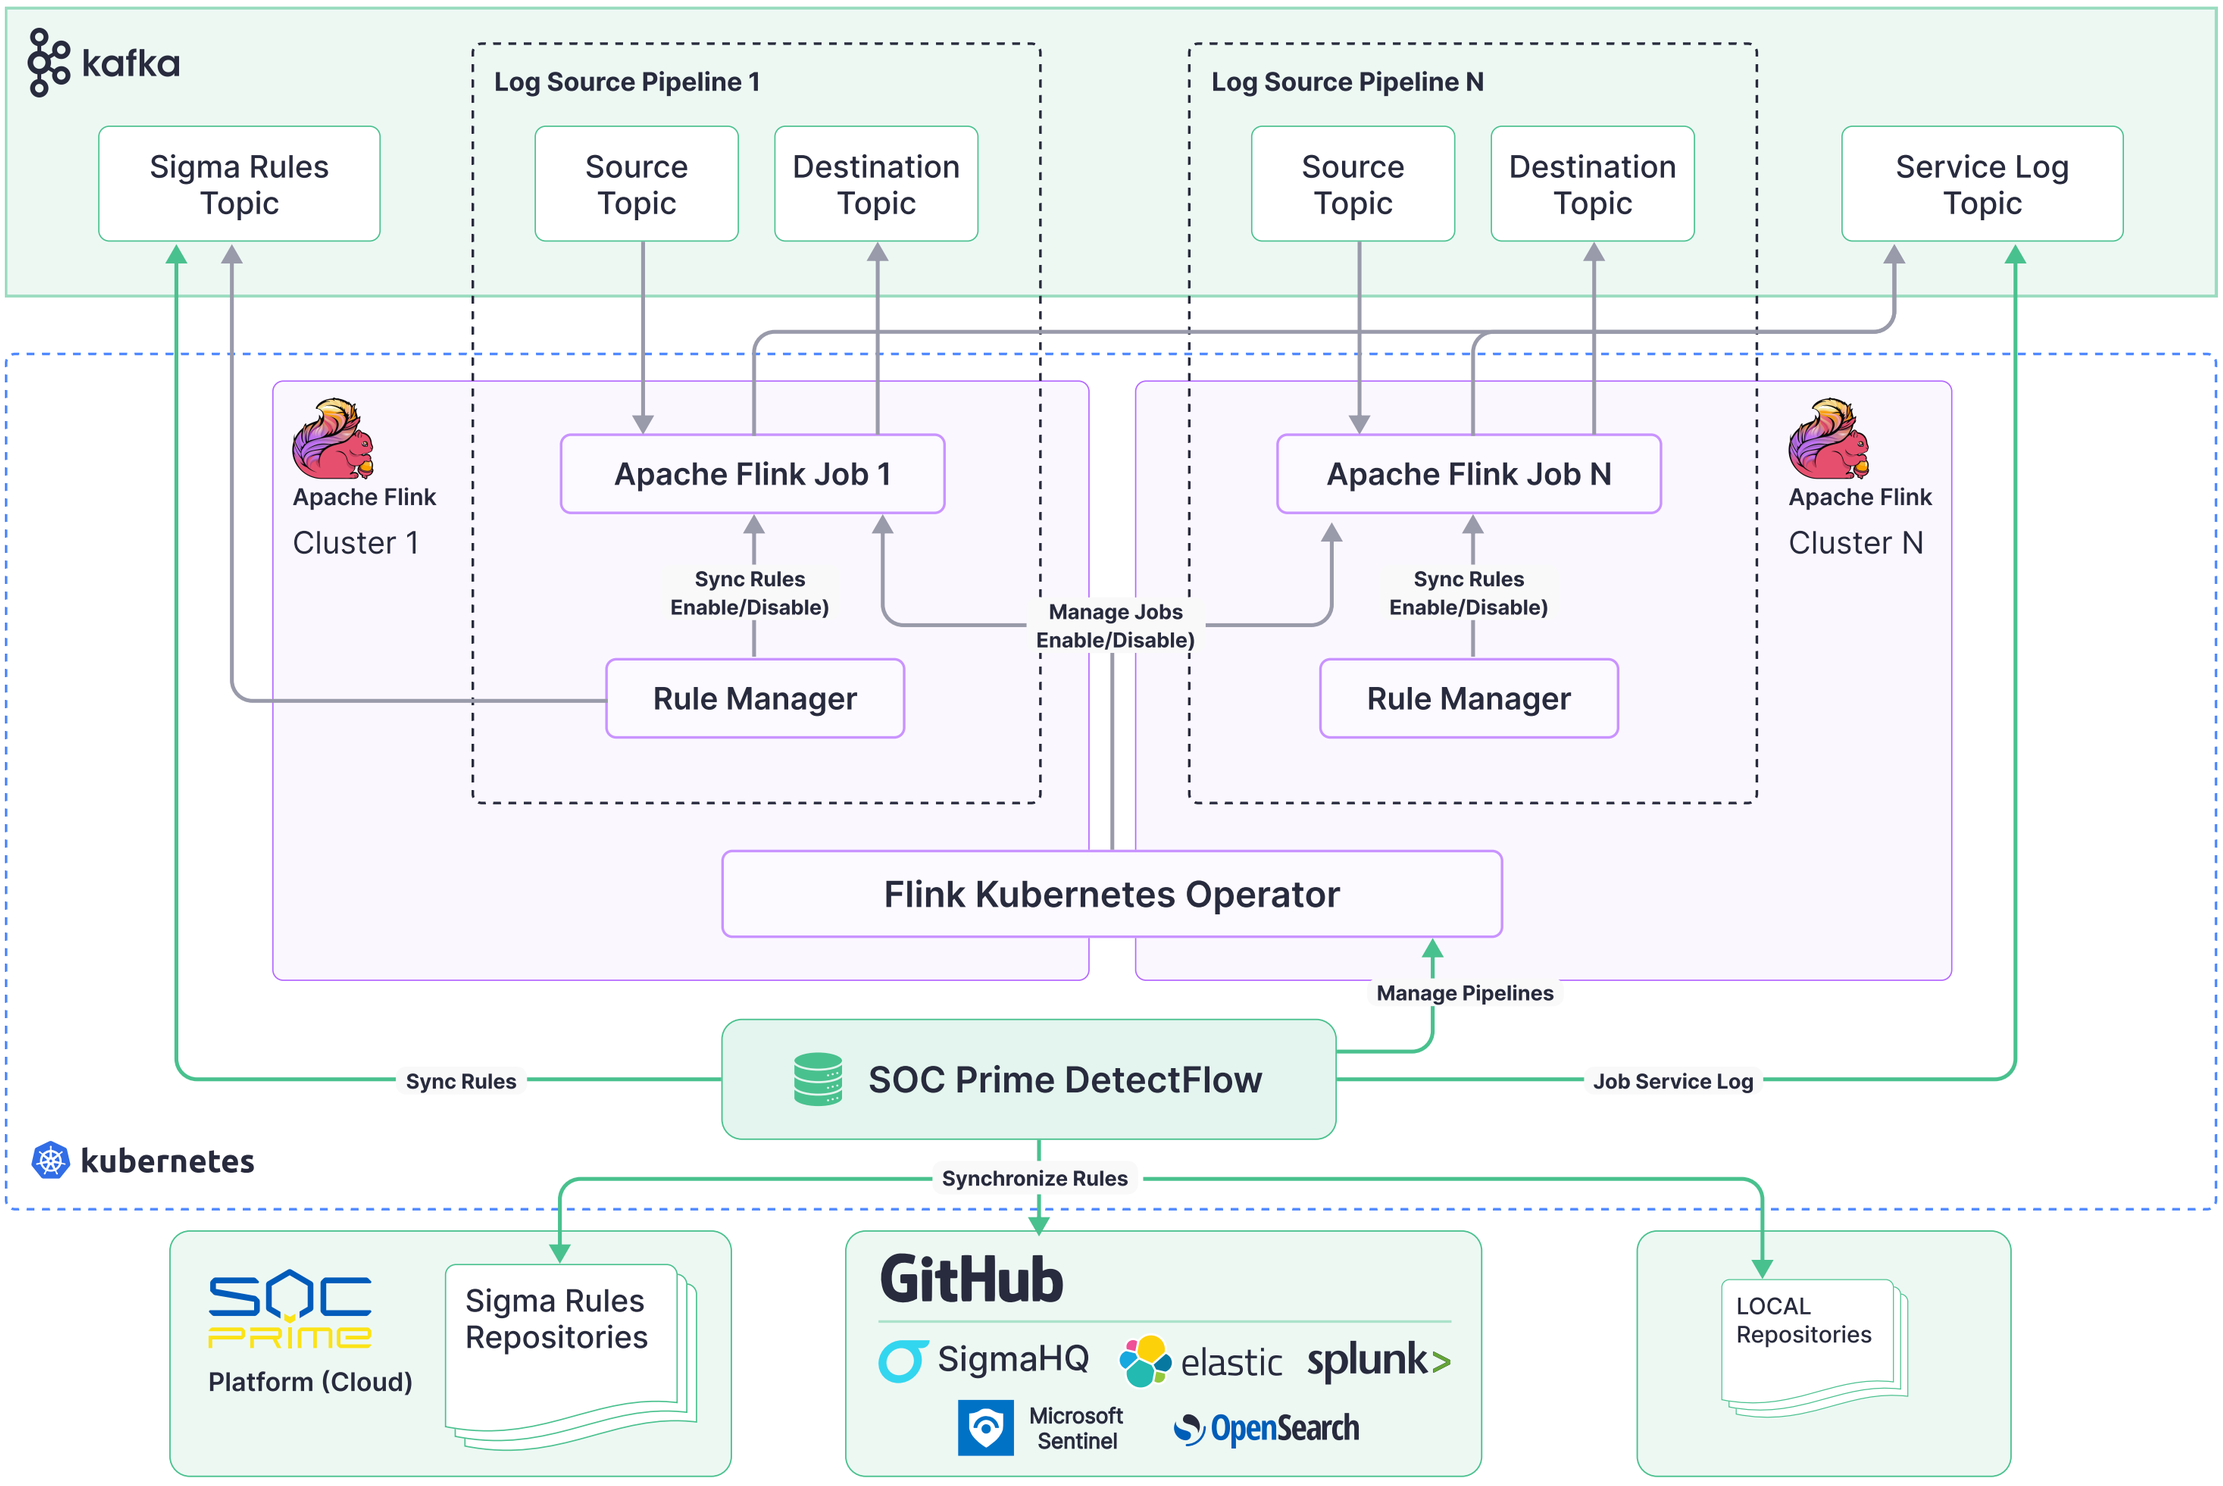

DetectFlow vs. ETL tools vs. SIEM

Dimension	ETL / Routing tools	SIEM	SOC Prime DetectFlow
Primary job	Move / transform data	Store / search / correlate	**Detect in-flight on streams**
Where detection runs	Not core (custom logic)	After ingest + index	**Before SIEM, on Kafka topics**
Speed	Batch / micro-batch	Minutes to hours	**Milliseconds**
Scale (1000s of rules)	Hard / costly	Cost grows with volume	**Built for massive rule execution**
Cost impact	Often forwards more data	Scales with ingest volume	**Cuts SIEM ingest via tag / filter**
Integration	Pipeline work required	System of record	**Layers on existing Kafka**

Get started

1. **Book a technical discovery session**
   Assess your Kafka cluster health and capacity for DetectFlow workloads.

2. **Define pilot scope**
   1–2 log event streams (Kafka topics), success metrics, timeline.

3. **Execute pilot**
   Tag and enrich events. Optional routing. Measure detection coverage improvement.

4. **Assess outcomes & expand**
   Review results. Plan phased expansion across additional topics and log sources.

Minimum requirements

KAFKA Apache Kafka 3.8+ or Confluent	STREAM PROCESSING Apache Flink 1.13+ on K8s 1.28+
METADATA STORE PostgreSQL 14+	MINIMUM CLUSTER 3 nodes, 8 vCPU / 34Gb each

Deployment options

- **One-click local deployment** — Docker + Minikube, single shell script, 7 minutes
- **Full deployment** — existing Kubernetes cluster with Helm charts
- **Air-gapped** — YAML + Docker images on physical drive, no internet required
- **Confluent Cloud / Platform** — fully compatible, identical DetectFlow behavior

SOC Prime, Inc.
socprime.com
sales@socprime.com