Comparison

Data in Motion, Simplified at Scale

Evaluating Condense vs Confluent for Enterprise Real-Time Data Streaming

Condense

Presented By :
Zeliot, for their platform Condense: a Kafka- native, fully managed streaming platform deployed in your cloud

ZELIOT

How Does Condense Compare vs Confluent?

Condense and Confluent Cloud both handle GBps-level throughput, but Condense provides a complete autonomous environment that Confluent lacks. Confluent Cloud is a managed infrastructure service that still leaves the heavy lifting to the user llike scaling external Flink clusters and paying high egress fees to move data into a vendor network. Condense solves the entire lifecycle by merging high-performance brokers with a native execution layer for Java, Python, or Go code. Since it runs on customer cloud, the platform scales brokers and logic automatically as one, finally cutting out the need for separate processing clusters and the hidden costs of moving data between networks

› Architecture

FEATURE	CONDENSE	CONFLUENT
System Design	A unified environment that bundles Kafka, Kubernetes, and the App Runtime into one integrated stack.App-Fabric: Not a single binary, but a unified modular stack that bundles Kafka, Kubernetes, and streaming App Runtime	A cloud-native tiered architecture that uses the proprietary Kora engine to disaggregate compute from storage for elastic, multi-tenant scaling
Platform Model	Unified Application Fabric: Merges the Kafka engine with a native event-driven microservice runtime. It is an "Execution Fabric"	Infrastructure Suite: A collection of separate distributed services (Kafka, Flink, Registry) that must be integrated by the user
Scaling Scope	Full-Stack Autonomy: Automatically scales brokers, connectors, and custom-code transforms (Java/Python/Go) based on real-time consumer lag	Infrastructure-Centric and Broker-Level Scaling: Efficiently scales cluster nodes, but scaling the processing power for custom code or microservices remains a manual or external engineering task
Cloud Storage	Native Object Storage: Direct offloading to S3/GCS/Azure Blob; data stays in your buckets for infinite retention	Kora Engine (Proprietary): Proprietary serverless storage engine with high-performance tiered abstraction
Enterprise BYOC	Native First: 100% of the data plane, including the app logic, lives in your private cloud account	Split Control Plane: WarpStream (BYOC) exists but depends on Confluent's external control plane for metadata
Performance	Application logic runs "local" to the broker, eliminating the network latency of external processors	Best-in-class for GBps+ global throughput and multi-region synchronization but subject to variability caused by network overhead in multi-tenant SaaS

› Ecosystem

FEATURE	CONDENSE	CONFLUENT
Industry Readiness	Dual-Layer Ecosystem: Access a vast library of generic connectors alongside industry specific domain-aware transforms for Mobility, IoT, and FinTech (e.g., VIN parsing, Geofencing) out of the box	The industry gold standard for general enterprise IT (SAP, Oracle, Snowflake)
Apache Kafka API	100% Native: Built on upstream Kafka 4.0; no protocol emulation or compatibility layers	Native Support: Generally strong, but proprietary engines like WarpStream have specific API limitations and behavioral differences
Operational Logic	Fully Managed & Autonomous: Hands-off Kafka operations with automated scaling for brokers, connectors, and custom-code transforms	Managed Services: Reliable but reactive scaling (eCKUs); requires manual intervention for custom application scaling and tuning
Stream Processing	Simplified Full-Code: No-code visual blocks + native built in IDE with git link support with auto-managed scaling	High Complexity: Simple tasks often require a full Flink or ksqlDB deployment, increasing overhead
Dev Experience	Integrated AI IDE: Build, test, and deploy production Java/Python/Go logic natively. Includes Agentic AI for root-cause analysis and code generation	Decoupled Workflow: Fragmented development requiring external management for Flink clusters or standalone ksqlDB instances
Observability	Inbuilt Full-Stack: One view for broker health, connector throughput, and custom transform execution traces	Siloed Monitoring: Basic metrics included; deep application-level tracing requires external APM (Datadog/Splunk) integration
Flink Integration	Native Runtime Alternative: Provides a built-in event-driven microservice runtime that replaces the need for Flink	Managed Flink: Fully integrated, production-grade serverless Flink for stateful enterprise logic

› Support & Compliance

FEATURE	CONDENSE	CONFLUENT
Access Experts 24/7	Direct Engineering: Built in support widget to raise ticket. Condense offers support priority based TAT for resolution	Global Tiered: Thousands of experts across every timezone; follow-the-sun enterprise support
Cloud Availability	AWS, GCP, Azure (Customer Subscription)	AWS, GCP, Azure (Confluent controlled)
Security	Own Cloud -Sovereign: Inherits your IAM and security groups; 100% data localization	Extensive certifications (SOC2, HIPAA, PCI, FedRAMP) for multi-tenant SaaS
SLAs	99.95% Full-Stack: Covers the brokers, the processing logic, and the connectors under one SLA	99.99%, Higher uptime guarantee for the broker specifically (Kora engine)
Pricing Model / TCO	Compute-Based: Simple, predictable vCPU/hr billing. Zero "Connector " or per-partition fees	Usage-Based: Multi-dimensional pricing that fluctuates based on data volume, partition counts, and connector usage

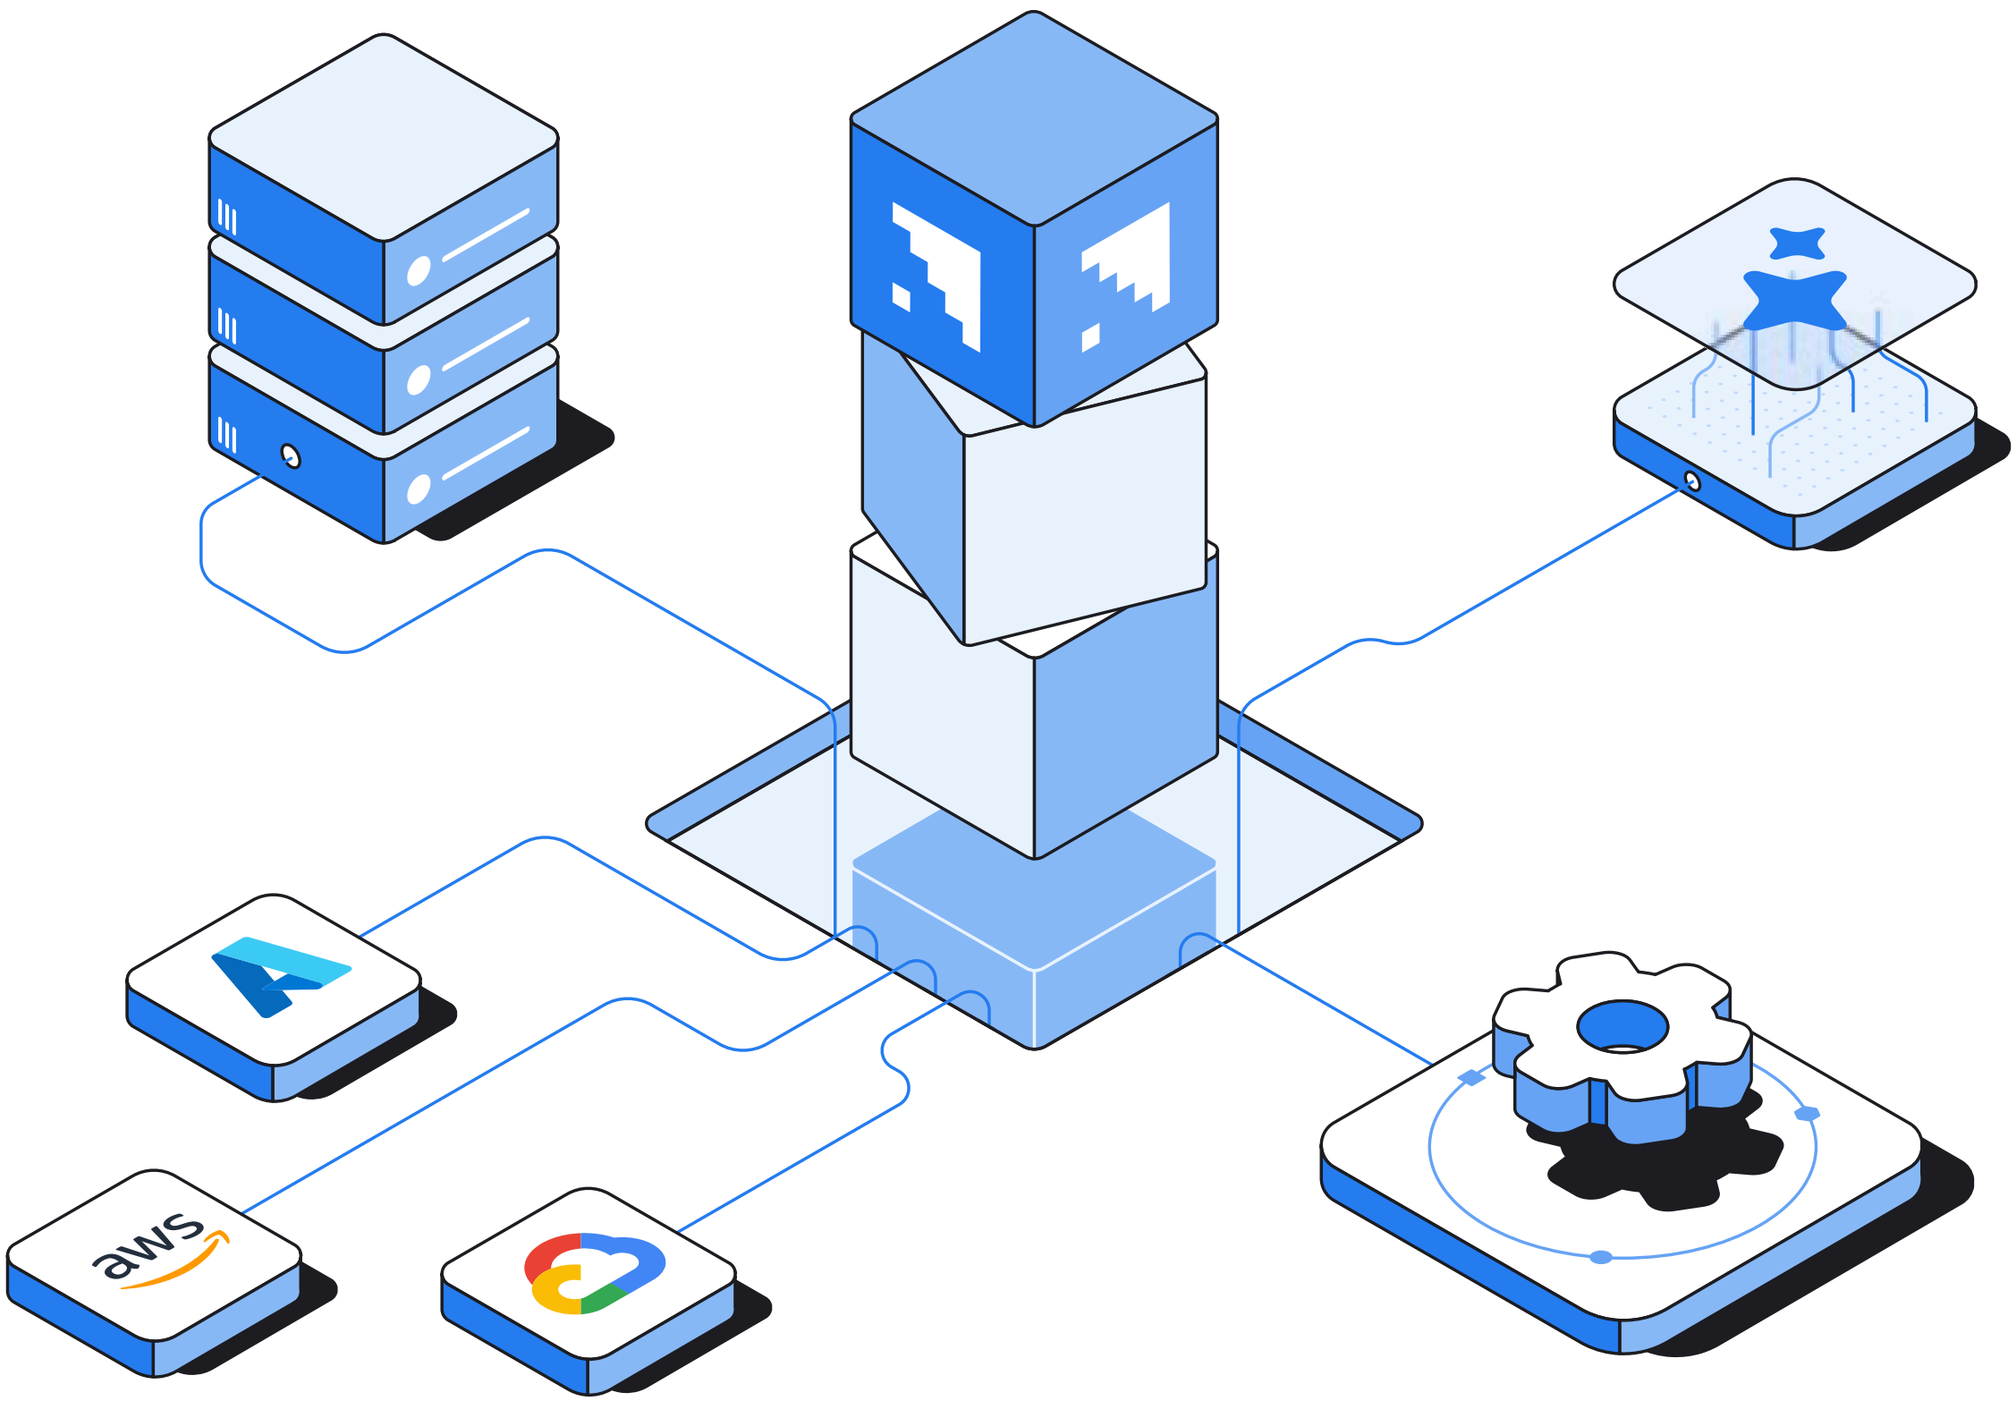

Why Switch to Condense?

NATIVE DATA SOVEREIGNTY (BYOC ARCHITECTURE)

Legacy SaaS models require data to exit your secure perimeter to reside in a vendor-owned account, introducing egress costs and compliance risks

- Condense is architected for Bring Your Own Cloud (BYOC). It deploys directly from cloud marketplaces into your private VPC (AWS, Azure, or GCP). This ensures that 100% of your data residency remains within your control, inherits your existing IAM/KMS security policies, and eliminates the "SaaS Tax" of cross-network data transfer fees

AUTONOMOUS FULL-STACK SCALING

In a standard streaming architecture, scaling is reactive and siloed brokers are scaled based on disk/CPU, while processing apps are scaled via Kubernetes HPA

- Condense introduces Autonomous Scaling for the entire lifecycle. Through its Custom Transform Framework (CTF), the platform monitors throughput and lag at the event level. It automatically provisions compute resources for your custom Java or Python transforms and connectors, ensuring that your application logic scales in perfect lockstep with your Kafka brokers without manual intervention

VERTICALIZED VS HORIZONTAL ECOSYSTEM

While general-purpose platforms offer generic connectors, they leave the domain-specific logic to the user

- Condense collapses these silos into a single Unified Application Fabric. It merges the Kafka engine with a native, event-driven microservice runtime. You don't "glue" tools together; the tools are built into the fabric

INTEGRATED DEVELOPER WORKFLOW (THE AI IDE)

Condense removes "Integration Sprawl" by embedding a Full-Code IDE directly into the platform

- Developers can write, test, and publish production-grade transforms in their preferred languages (Java, Python, Go) with built-in GitOps support. This environment is augmented by an Agentic AI Layer that assists in code generation and provides real-time root-cause analysis, reducing the engineering effort required to maintain complex streaming pipelines by up to 90%

EVOLUTION FROM INFRASTRUCTURE TO APPLICATION PLATFORM

The primary limitation of traditional managed Kafka is the "operational gap." While legacy providers host the brokers, engineering teams are still responsible for the secondary layers: scaling processing clusters, managing containerized microservices for transforms, and wiring up external observability

- Condense collapses these layers into a unified Streaming Application Substrate. It provides an event-driven runtime that treats both the infrastructure and the application logic as a single, managed entity, allowing developers to focus purely on business logic rather than pipeline scaffolding

Vertical Data Platform for Real-Time Businesses

Condense understands vertical data by design and connects directly to real time data sources like vehicles or GDS platforms. Its purpose-built industry specific connectors and transformations accelerate the development of real-time vertical use case

- Accelerate the path from idea to production
- Rapidly realize and deploy industry use cases
- Enterprise scale with native governance
- TCO ↓ through vertical industry-first approach

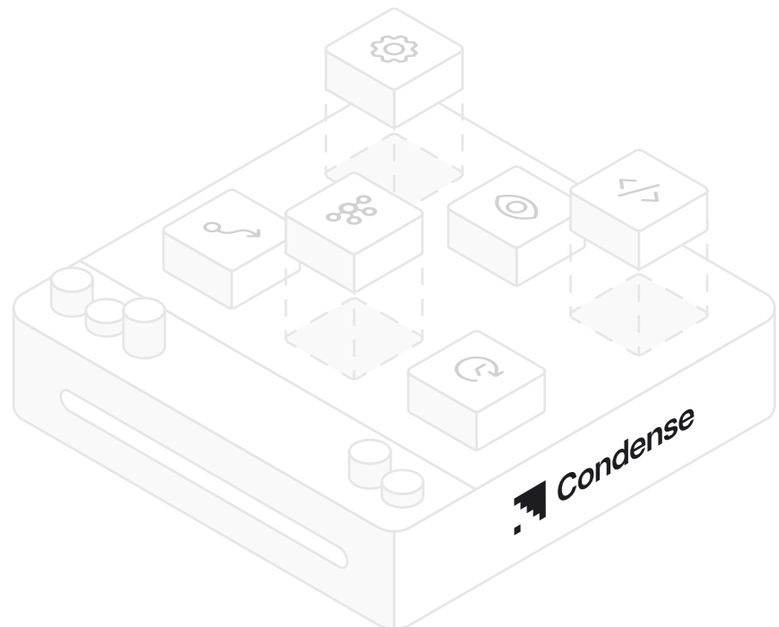

Get Started with Condense in 3 Simple Steps

DEPLOY

Deploy Condense from Your Cloud Marketplace

CREATE

Deploy Condense from Your Cloud Marketplace

BUILD

Deploy Condense from Your Cloud Marketplace

Most platforms focus on streaming data, Condense focuses on understanding it. It's not just managed Kafka; it's Kafka with context. Infused with domain intelligence, **Condense transforms raw streams into decisions, automations, and impact,** because the right Kafka platform should understand what it's moving

ZELIOT

Hope this analysis sparked a new vision for your data architecture!

Stick Around, It Gets Better

Read Our Blogs →

Read Our Documentation →

Read Our eBooks →

We're Ready When You Are. Let's Talk.

Schedule a Meeting

Don't Miss a Thing - Come Along