Comparison

Data in Motion, Simplified at Scale

Evaluating Condense vs Warpstream for Enterprise Real-Time Data Streaming

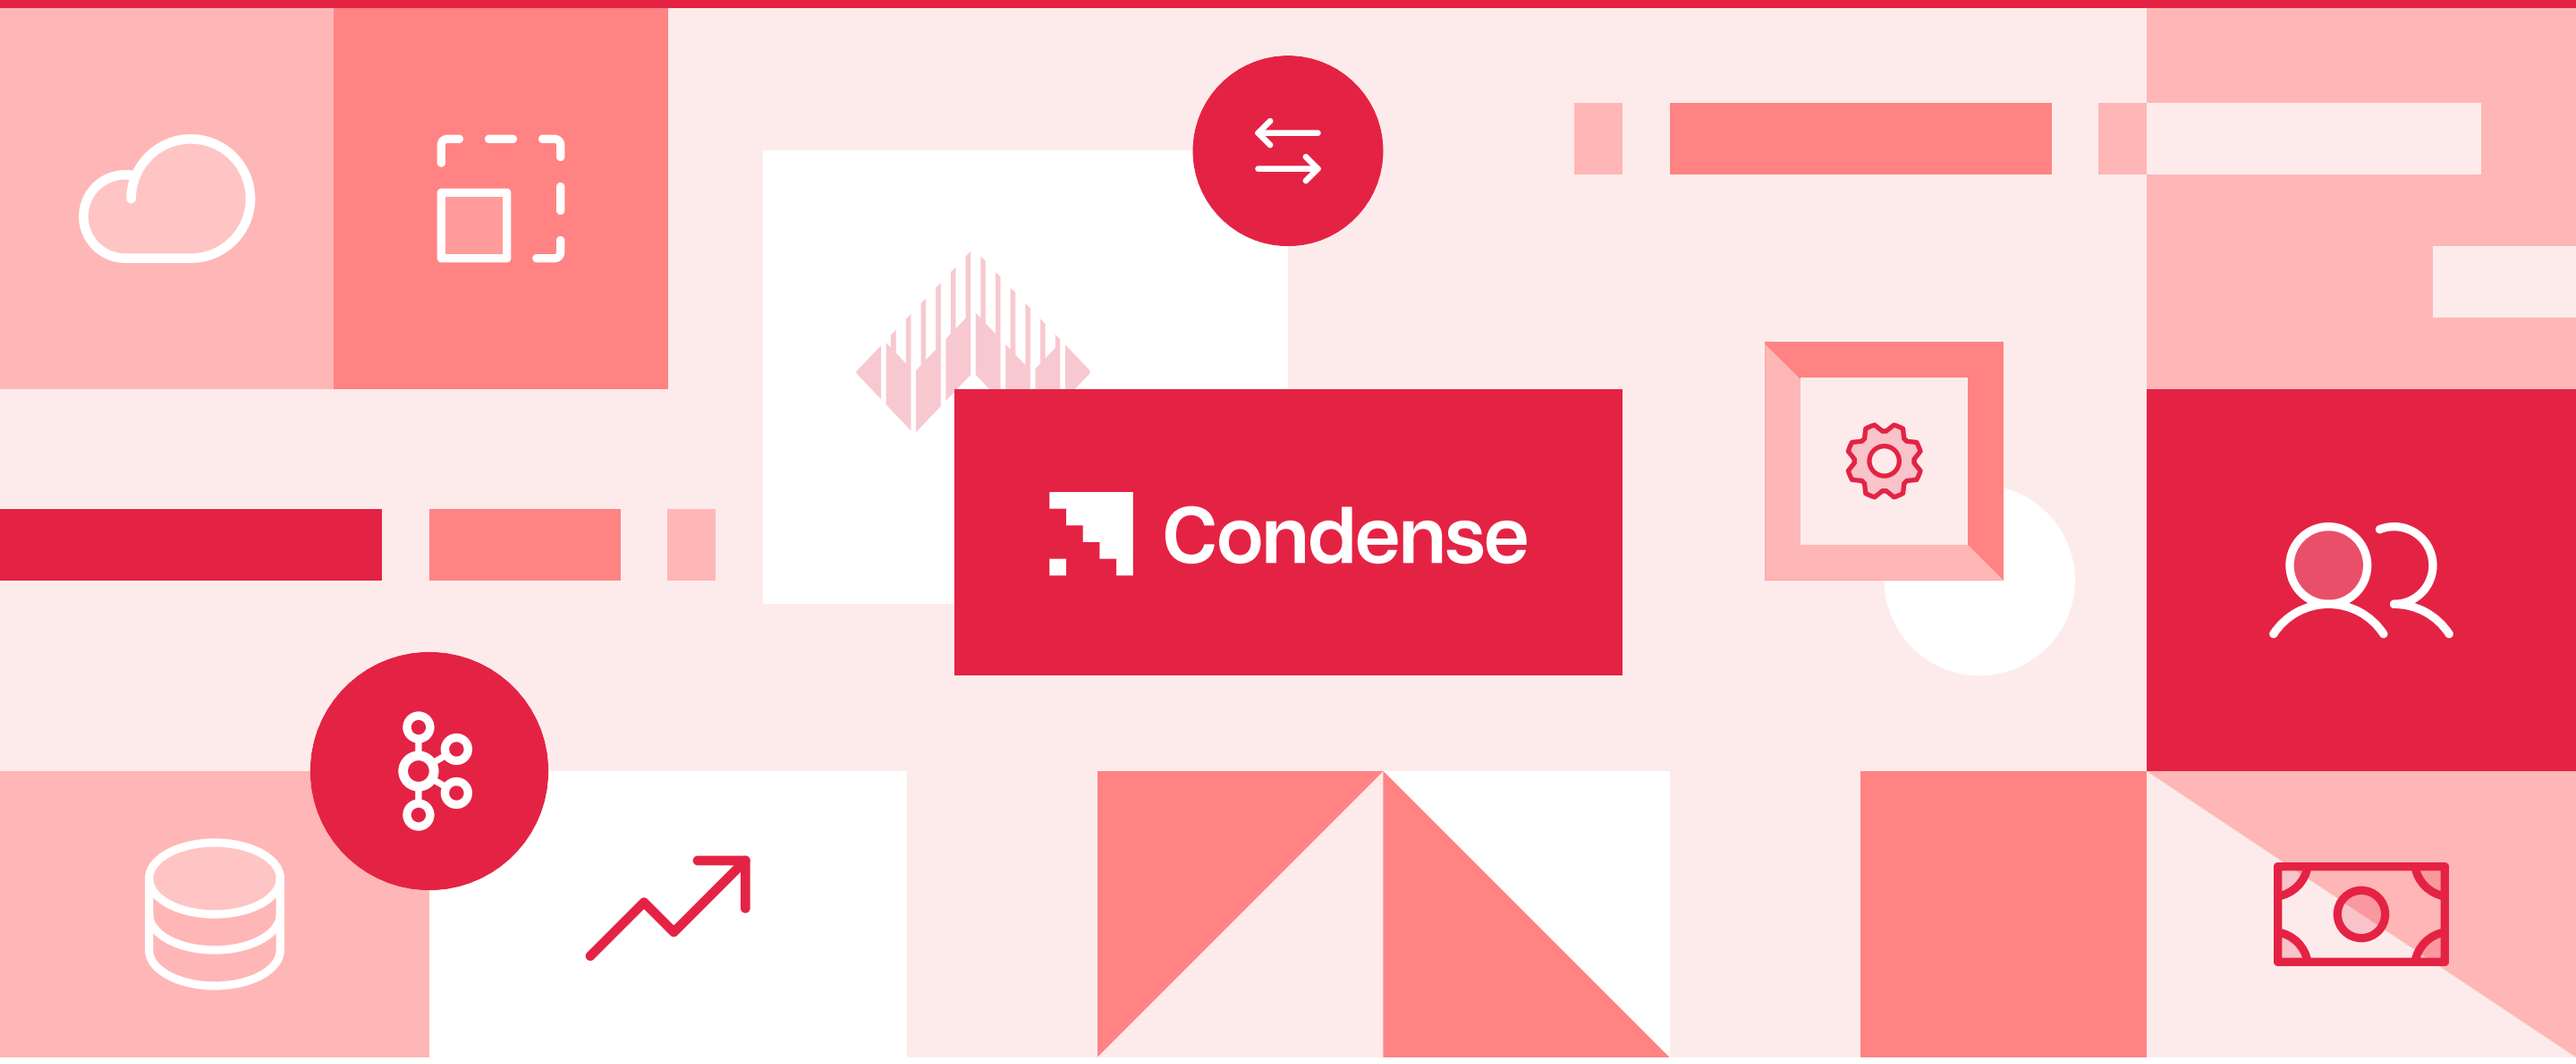

Condense

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

ZELIOT

How Does Condense Compare vs Warpstream?

Condense and WarpStream both address the complexities of cloud-native streaming, but they solve for entirely different outcomes. WarpStream is an infrastructure-focused, stateless reimplementation of the Kafka protocol designed to slash networking and storage costs by writing data directly to S3, but it remains a transport layer that introduces significant "latency ceilings"often 400ms–600ms and requires you to manage separate external clusters like Flink or Spark to actually process your data. Condense solves the entire application lifecycle by integrating high-performance, native Kafka 4.0 brokers with a native execution layer for your Java, Python, or Go code. It autonomously scales your brokers and your logic in lockstep, eliminating the need for separate processing infrastructure and the "latency tax" of moving data between decoupled storage and external compute

› Architecture

FEATURE	CONDENSE	WARPSTREAM
System Design	A unified environment that bundles Kafka, Kubernetes, and the App Runtime into one integrated stack	A stateless reimplementation of the Kafka protocol that replaces local disks with a 100% S3-direct storage model
Unified Platform	Unified Application Fabric: Merges the Kafka engine with a native event-driven microservice runtime. It is an "Execution Fabric”	Storage-Centric: A high-performance transport layer; requires external clusters (Flink/Spark) for any business logic
Scaling Scope	Full-Stack Autonomy: Automatically scales brokers, connectors, AND your custom-code transforms (Java/Python/Go) based on real-time consumer lag	Stateless Scaling: Scaling agents is instant because they hold no data, but scaling the application logic is a separate manual task
Cloud Storage	Native Object Storage: Direct offloading to S3/GCS/Azure Blob; data stays in your buckets for infinite retention	Direct-to-S3: Writes data directly to object storage. Eliminates EBS cost but relies entirely on S3 availability and performance
Enterprise BYOC	100% BYOC Native: Specifically engineered to run in your VPC to eliminate "SaaS" networking taxes and maintain 100% data sovereignty	BYOC (Zero-Data SaaS): Agents run in your VPC, but the control plane and metadata are managed by WarpStream
Performance	Serverless Scale and Zero-Hop: Running GBps+ enterprise data workload and autonomous scale for spikes and variations. Application logic runs "local" to the broker, eliminating the network latency of external processors	Elastic Performance: Highly elastic for throughput, but subject to S3-bound tail latencies (400ms+) for primary writes

› Ecosystem

FEATURE	CONDENSE	WARPSTREAM
Operational Effort	Zero-Ops Application: Eliminates "Microservice Sprawl." You simply write the logic; the platform autonomously runs, scales, and monitors the entire application pipeline	Simplified Infrastructure: Removes Kafka disk management and rebalancing, but leaves you to manage the external apps/clusters that process the data
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	Horizontal Generalist: Optimized for raw byte-moving at scale. Any industry-specific logic must be engineered from scratch by your team
Apache Kafka API	100% Native: Built on upstream Kafka 4.0/ KRaft	Protocol-Compatible: A custom re-implementation in Go. Supports core produce/ consume, but misses some Admin APIs and internal protocol fields
Logic Execution	Managed App Runtime: Custom transforms run in a managed, serverless-style environment. Condense handles containerization, state management, and orchestration	External Execution: Business logic is either limited to simple inline transforms or must be piped to external clusters (Flink/K8s) that you manage
Dev Experience	Integrated AI IDE: Build, test, and deploy production Java/Python logic natively. Includes Agentic AI for root-cause analysis and code generation	Standard API: Works with existing Kafka clients and tools, but lacks an integrated environment for building and deploying production code
Observability	Inbuilt Full-Stack: Trace an event from ingestion through your custom code to the final sink in a single, unified dashboard	Broker Observability: Strong metrics for agent health and S3 throughput, but application-level tracing requires external APM tools
Flink Integration	Native Runtime Alternative: Replaces the need for Flink for most use cases	Dependent (External): Heavily relies on external Flink or Spark for stateful logic, adding network complexity and "hops"

› Support & Compliance

FEATURE	CONDENSE	WARPSTREAM
Access Experts 24/7	Direct Engineering: Built in support widget to raise ticket. Condense offers support priority based TAT for resolution	Enterprise Support: Mature support model for 24/7 cluster operations
Cloud Availability	AWS, GCP, Azure (Customer Subscription)	AWS, GCP, Azure (customer Cloud)
Security	Own Cloud -Sovereign: Inherits your IAM and security groups; 100% data localization	Private: Data stays in your bucket, but metadata is hosted externally in the WarpStream Control Plane
SLAs	99.95%: Covers the brokers, the processing logic, and the connectors under one SLA	Broker HA: Industry-standard uptime for the ingestion engine
TCO	Compute-Based: Simple, predictable vCPU/hr billing. Zero "Connector " or per-partition fees	Usage-Based: Pricing scales with data volume; eliminates EBS and Inter-AZ fees, but costs rise with throughput

Why Switch to Condense?

TRANSITION FROM STORAGE TIER TO APPLICATION ENVIRONMENT

WarpStream is an infrastructure-focused reimplementation of Kafka designed to slash storage costs by writing directly to S3. However, it remains a transport layer; to actually process data, teams must still manage external clusters like Flink or Kubernetes

Condense unifies these layers. It integrates the message broker with a native execution fabric, hosting both your data and your business logic (Java, Python, or Go) as a single, managed entity. This allows you to ship production-ready applications directly on the stream without the operational burden of building a separate processing tier

FULL-STACK AUTONOMOUS SCALING VS BROKER SCALING

WarpStream scales its "agents" quickly because they are stateless, but it does not scale the applications processing that data. If your external logic falls behind during a spike, the pipeline lags until you manually scale your compute clusters

Condense provides Full-Stack Autonomous Scaling. It monitors real-time consumer lag and automatically provisions compute for the brokers, connectors, and your custom transforms simultaneously. This ensures your processing power stays in lockstep with data volume, scaling back down automatically to optimize costs when the surge passes

Why Switch to Condense?

ZERO-HOP PERFORMANCE VS S3 LATENCY CEILINGS

By writing directly to object storage, WarpStream introduces an inherent "latency ceiling," with P99 tail latencies often reaching 400ms–600ms. Furthermore, every time data moves to an external processing cluster, it incurs a "network hop" that adds further delay

Condense utilizes a Zero-Hop Architecture. Because your logic runs local to the broker, it eliminates the network delay of external processing. Combined with optimized local caching, Condense provides the sub-millisecond end-to-end performance required for mission-critical Mobility and IoT use cases

DOMAIN-AWARE LOGIC VS GENERAL-PURPOSE PIPES

WarpStream is a horizontal "byte-mover" designed for raw infrastructure efficiency; any industry-specific logic must be engineered from scratch

Condense offers a Verticalized Ecosystem. It includes pre-built, domain-aware transforms for industries like Mobility (e.g., VIN parsing, trip decoding) and IoT. By leveraging these pre-tuned assets, teams can bypass months of custom development and transition from prototype to production significantly faster

COMPLETE DATA SOVEREIGNTY (100% BYOC)

While WarpStream agents run in your VPC, the metadata and control plane are hosted externally by the vendor, which can create friction during strict security audits

Condense is 100% BYOC Native. It deploys the entire stack: brokers, runtime, and metadata, directly into your private cloud (AWS, Azure, or GCP). This ensures 100% data and execution localization, allowing you to inherit your existing security policies and utilize your cloud enterprise credits for the entire data lifecycle

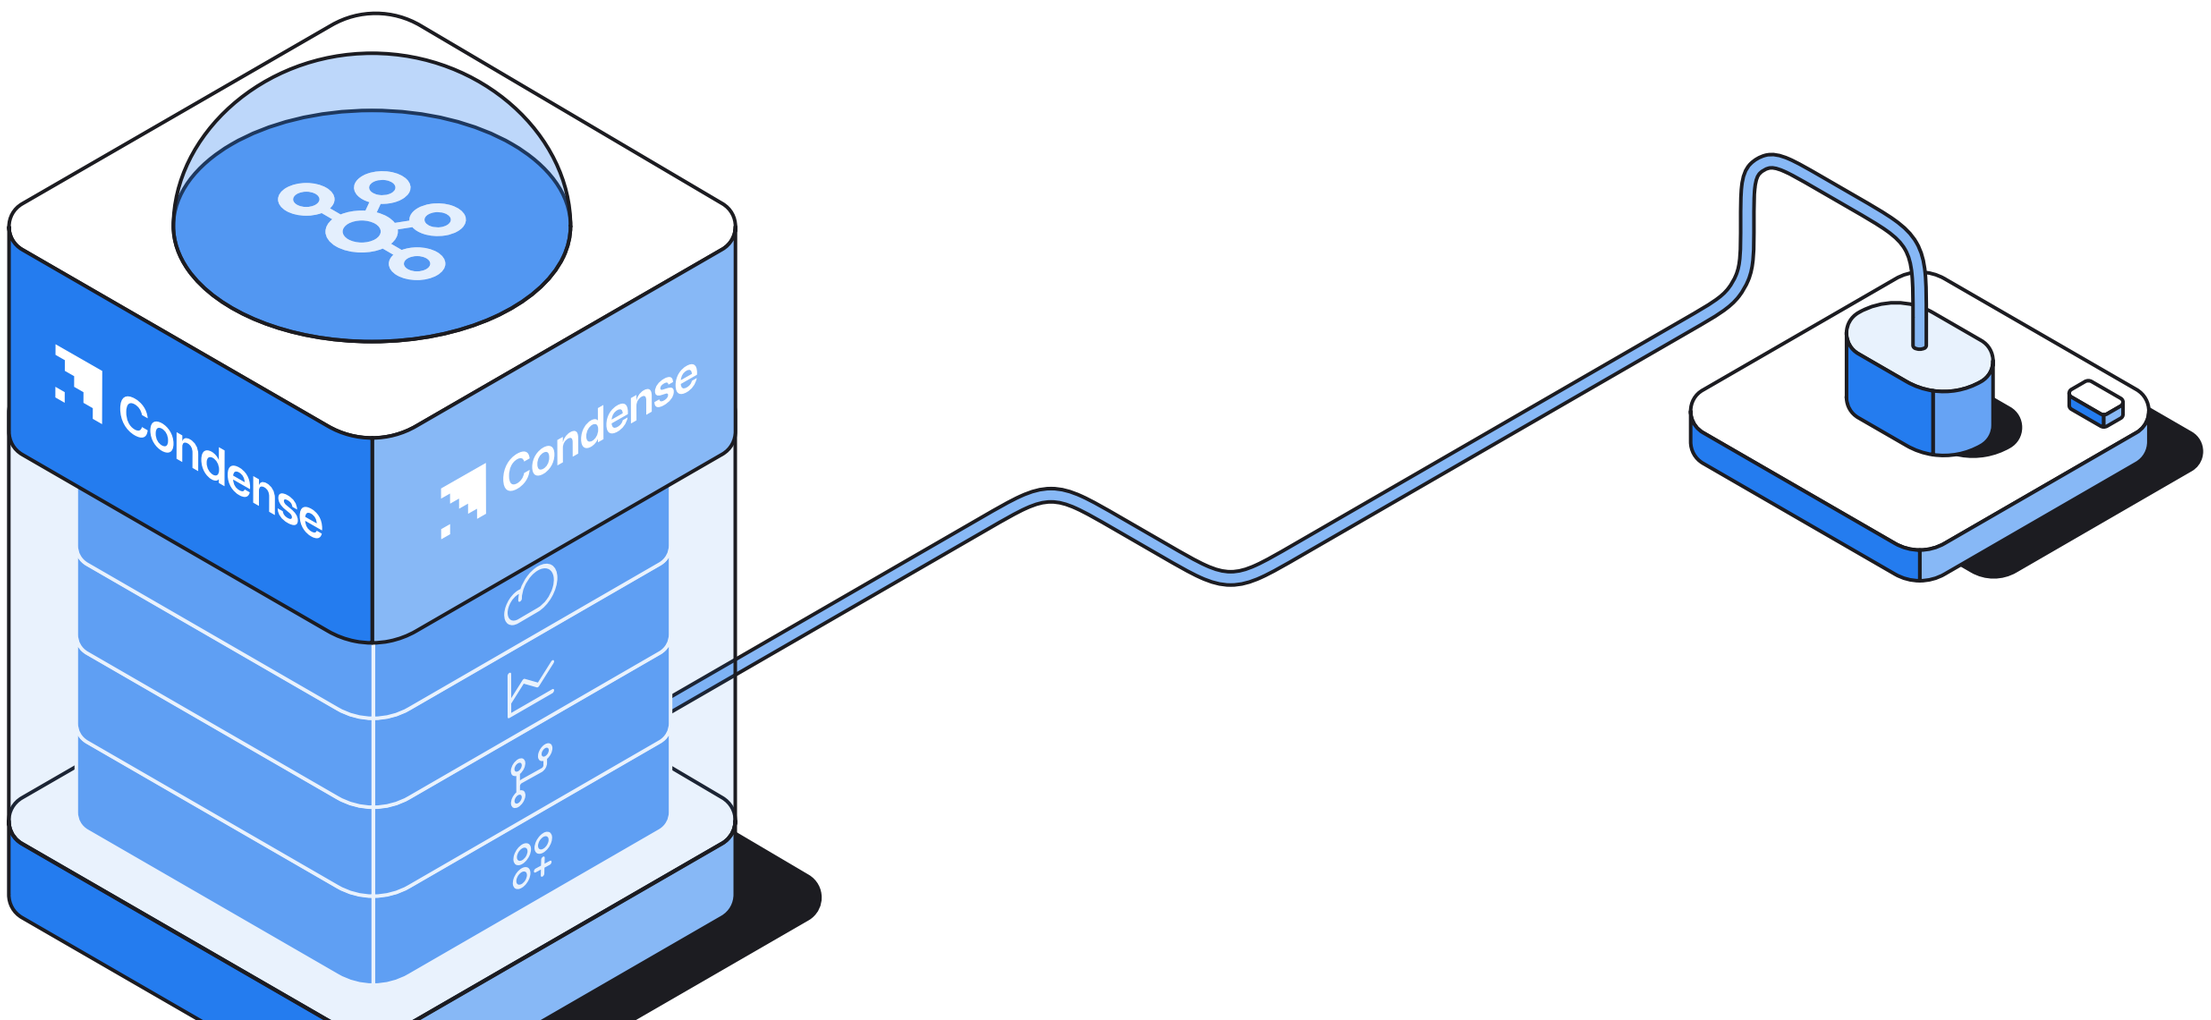

Vertical Data Platform for Real-Time Businesses

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

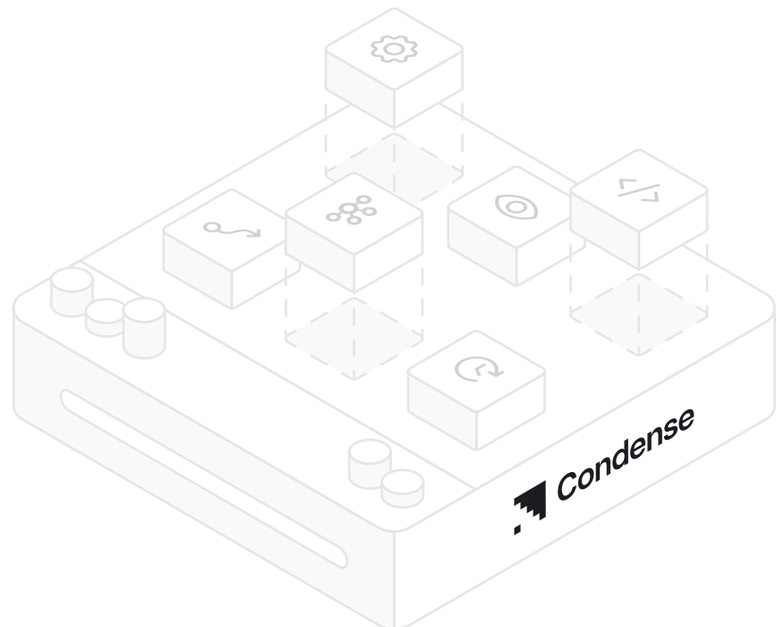

Get Started with Condense in 3 Simple Steps

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