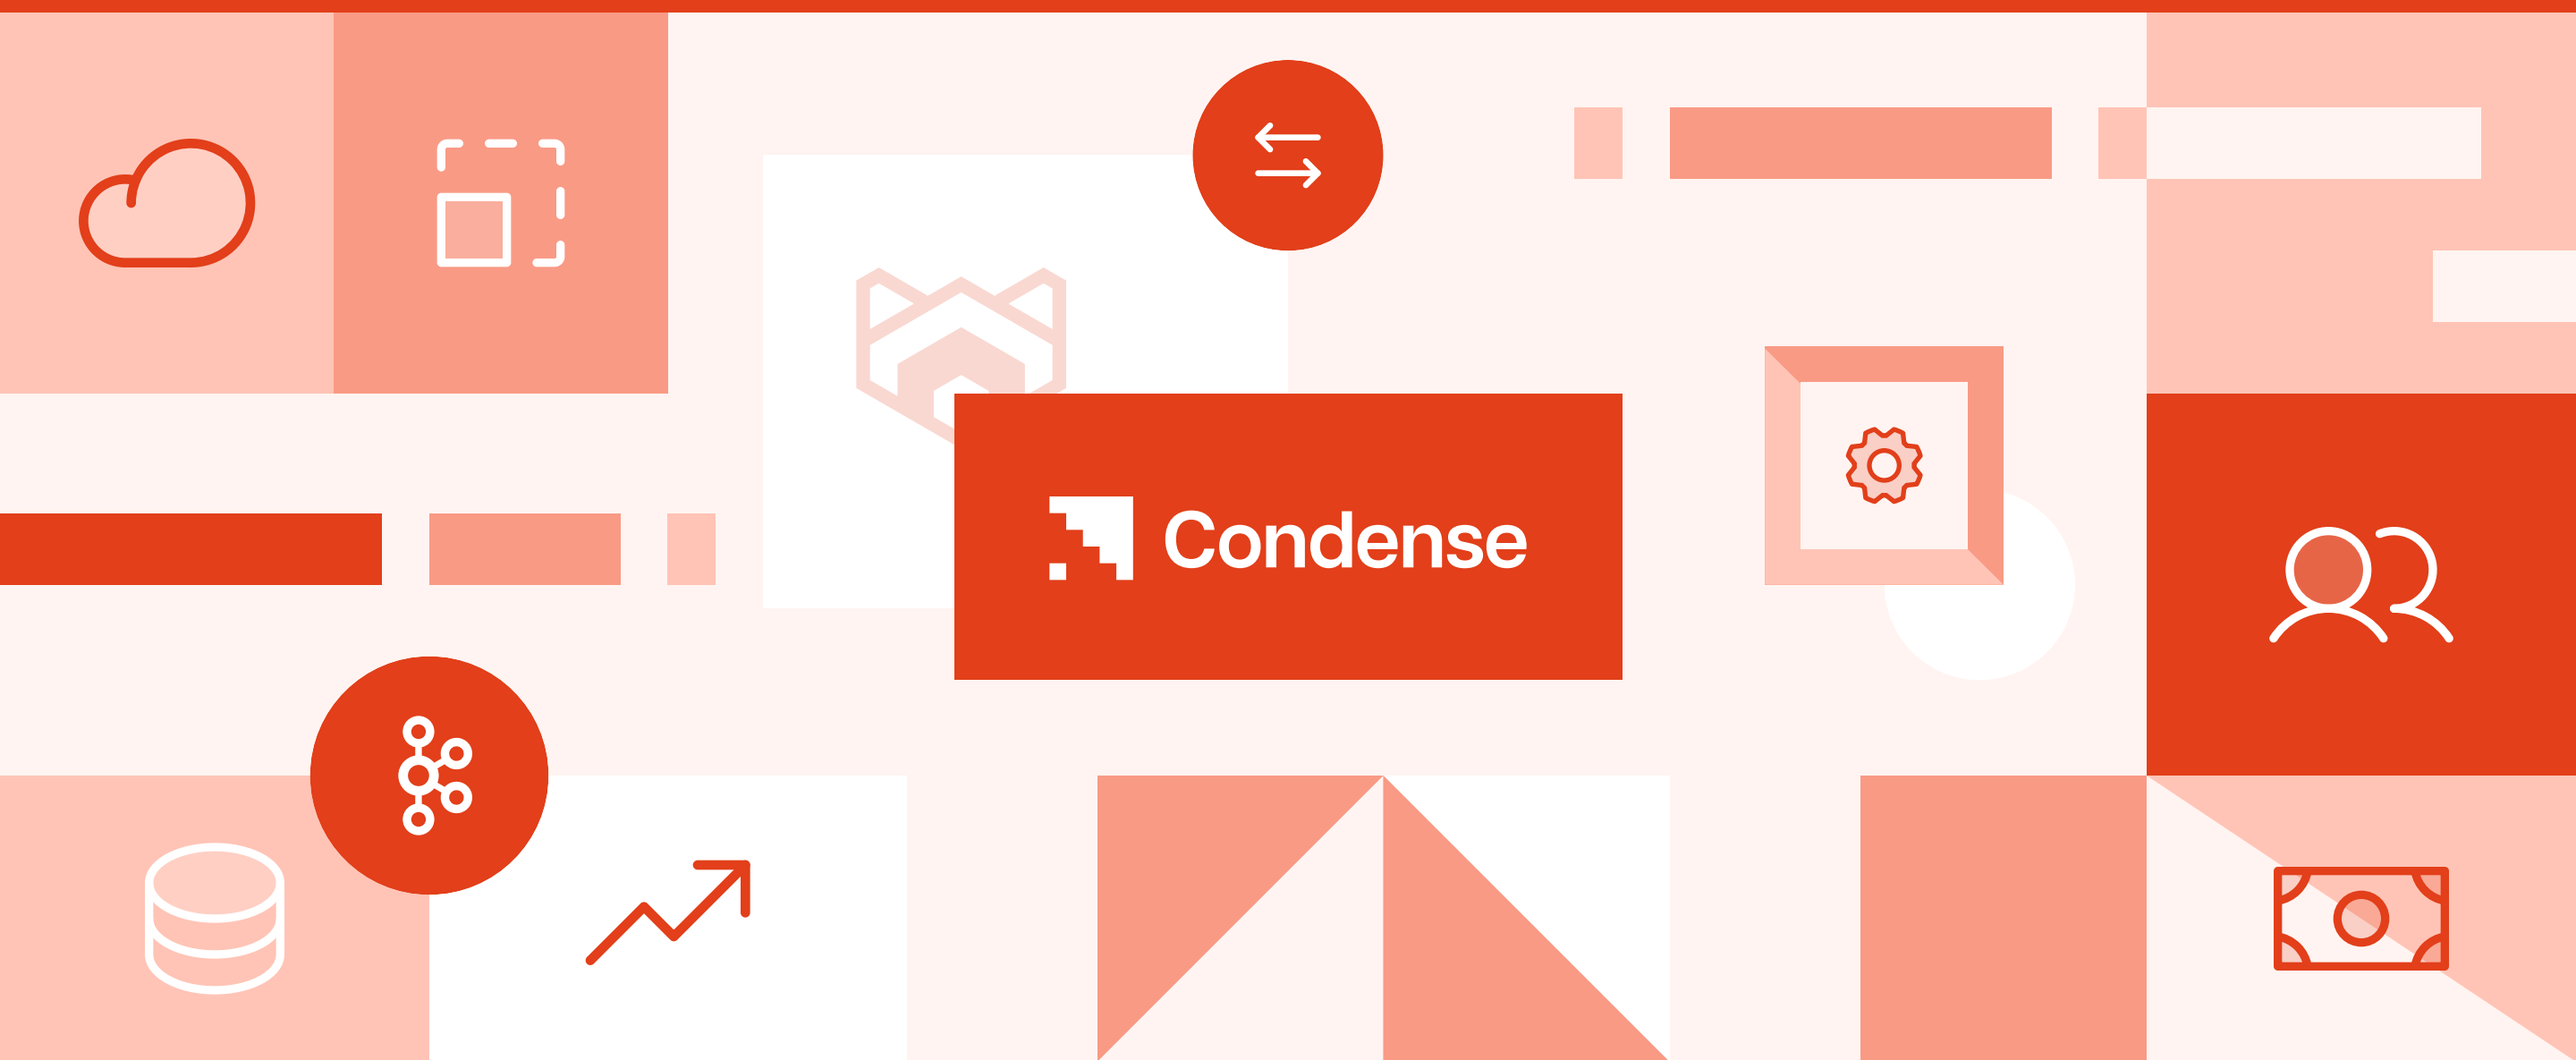


 Comparison

# Data in Motion, Simplified at Scale

Evaluating Condense vs Redpanda  
for Enterprise Real-Time Data Streaming



 **Condense**

ZELIOT

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

 [in](#) | [@](#) | [M](#) | [X](#) | [m](#)

All Rights Reserved 2026 @ Zeliot

# How Does **Condense** Compare vs Redpanda?

Condense and Redpanda both handle GBps-level throughput, but Condense provides the complete application environment that Redpanda lacks. Redpanda is a storage engine that requires you to manage separate clusters like Flink or Spark to actually use your data. Condense solves the entire lifecycle integrating high-performance 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 manual work of managing it

## > Architecture

FEATURE	CONDENSE	REDPANDA
System Design	A unified environment that bundles Kafka, Kubernetes, and the App Runtime into one integrated stack	Designed as a native C++ engine that consolidates the broker, proxy, and storage into a high-performance, JVM-free execution model
Unified Platform	Unified Application Fabric: Merges the Kafka engine with a native event-driven microservice runtime. It is an "Execution Fabric"	Broker-Centric: A high-performance storage engine; requires external tools for logic
Scaling Scope	Full-Stack Autonomy: Automatically scales brokers, connectors, and custom-code transforms (Java/Python/Go) based on real-time consumer lag	Infra-Centric and Broker-Level Scaling: Efficiently scales cluster nodes, but scaling the processing power for your code 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	Tiered Storage: Integrated S3 offloading; Requires Enterprise License for production
Enterprise BYOC	100% BYOC Native: Specifically engineered to run in your VPC to eliminate "SaaS" networking taxes and maintain 100% data sovereignty	Flexible Managed/BYOC: Offers SaaS, Serverless, and BYOC; however, Azure BYOC typically requires more manual configuration
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	Ultra-Low Latency: C++ engine optimized for sub-millisecond tail latencies

> Ecosystem

FEATURE	CONDENSE	REDPANDA
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 complexity (no JVM/KRaft), but leaves the "Application Toil" like managing the external apps that process the data to your team
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. Any industry-specific schema or logic must be engineered from scratch by your team
Apache Kafka API	100% Native: Built on upstream Kafka 4.0	API Compatible: C++ reimplement of the protocol; 100% API compatible
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 inline WASM for simple tasks or piped to external clusters (Flink/Spark/K8s) that you must 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	WASM & Console: Inline transforms via WebAssembly for simple tasks. Excellent UI (Console) for topic management, but lacks integrated logic runtimes
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 cluster health and throughput, but application-level tracing usually requires external APM tools
Flink Integration	Native Runtime Alternative: Replaces the need for Flink for most use cases	External: Heavily relies on external Flink or Spark for stateful logic

> Support & Compliance

FEATURE	CONDENSE	REDPANDA
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 limited in BYOC)
Security	Own Cloud -Sovereign: Inherits your IAM and security groups; 100% data localization	Hardened: Strong security features and enterprise RBAC
SLAs	99.95% Full-Stack: Covers the brokers, the processing logic, and the connectors under one SLA	99.99% Broker: Industry-leading uptime for the ingestion engine
Pricing Model / TCO	Compute-Based: Simple, predictable vCPU/hr billing. Zero "Connector " or per-partition fees	Usage/Performance: Pricing scales with throughput and volume. While hardware efficient (fewer nodes), costs can fluctuate significantly with data spikes., but often requires Enterprise licenses for Tiered Storage or BYOC support

## Why Switch to Condense?

### PRIVATE CLOUD SOVEREIGNTY (BYOC)

Many managed services require data to leave your environment to reside in a vendor-owned account, which can complicate security audits and increase network costs

- Condense is built for Bring Your Own Cloud (BYOC). It deploys directly into your private VPC (AWS, Azure, or GCP), keeping 100% of your data within your control. This allows you to inherit your existing security policies and encryption keys while utilizing your existing cloud enterprise credits for the underlying compute and storage

### VERTICAL SOLNS VS GENERAL-PURPOSE PIPES

Redpanda is a horizontal tool designed for raw speed across any use case, meaning all industry-specific logic must be built from scratch

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

# Why Switch to Condense?

## AUTONOMOUS FULL-STACK SCALING

In a typical Redpanda architecture, scaling is restricted to the infrastructure. If your processing logic (WASM transforms or external apps) starts to lag during a traffic spike, you must manually intervene to adjust compute resources

- 💡 Condense introduces Autonomous Scaling for the entire pipeline. It monitors real-time consumer lag and throughput at the event level. When data volume surges, the platform automatically provisions compute for your custom Java, Python, or Go transforms. It ensures your processing power stays in sync with your data volume and automatically scales back down to optimize costs once the surge passes

## UNIFIED DEVELOPMENT ENVIRONMENT

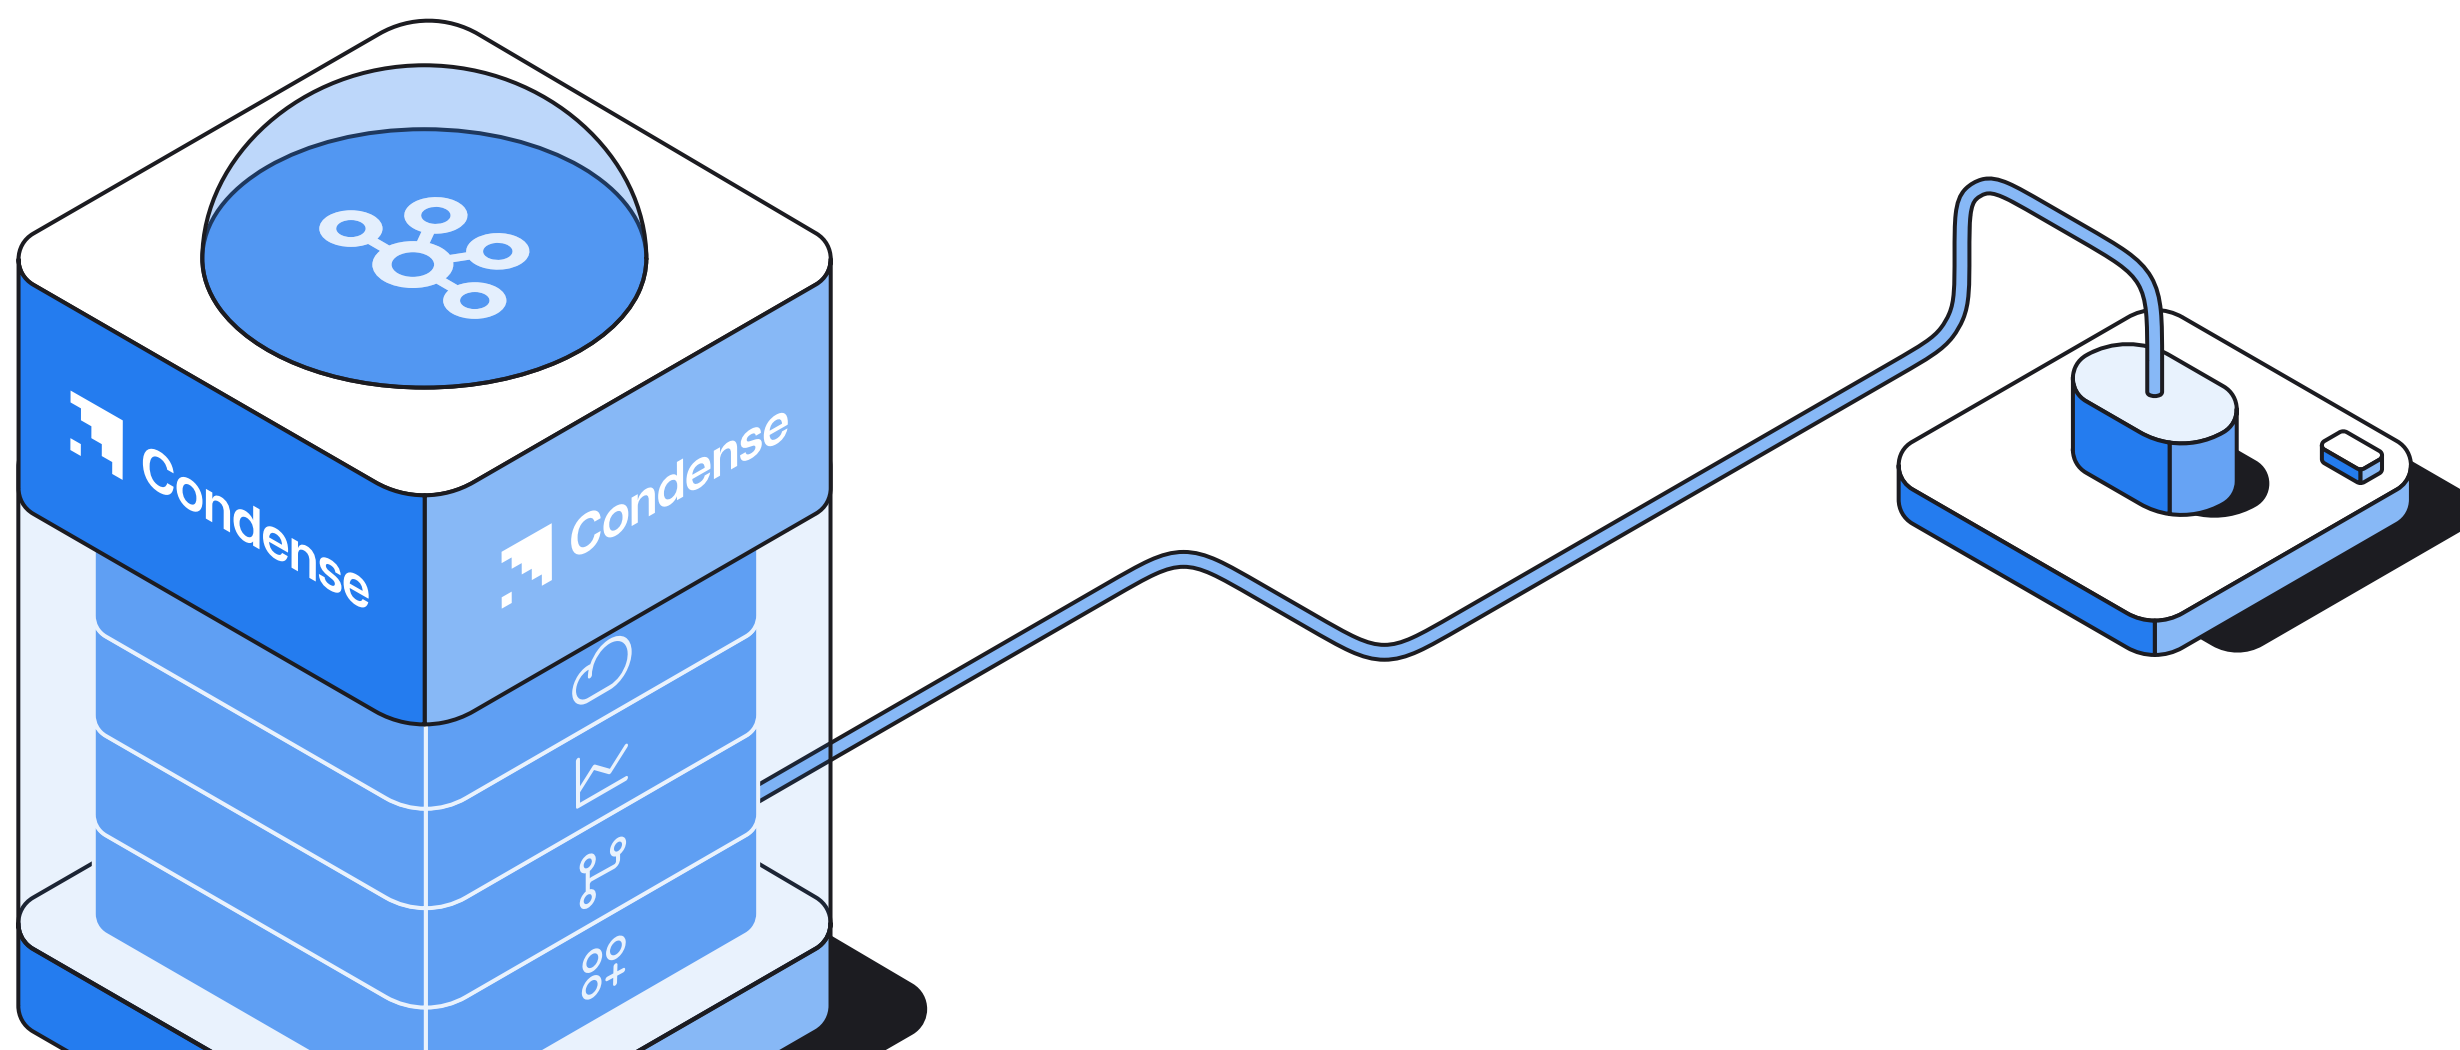
While Redpanda supports WebAssembly (WASM) for simple tasks like data filtering, complex stateful business logic often requires moving code to an external system

- 💡 Condense embeds a Full-Code IDE directly into the platform. Developers can write, test, and deploy production transforms in standard languages with built-in state management and version control. This environment is supported by an AI-assisted layer that handles rebalancing and root-cause analysis, significantly reducing the manual effort required to keep complex pipelines healthy

## TRANSITION FROM MESSAGE BROKER TO APPLICATION ENVIRONMENT





Redpanda simplifies the broker by removing the JVM, but it remains a storage and messaging layer. To actually use the data, teams still face the overhead of managing external processing clusters, such as Flink or Kubernetes microservices


- 💡 Condense merges these layers. It provides a managed environment that hosts both the data and the business logic as a single entity. This allows developers to ship production-grade code directly on the stream without the burden of building and maintaining a separate processing tier



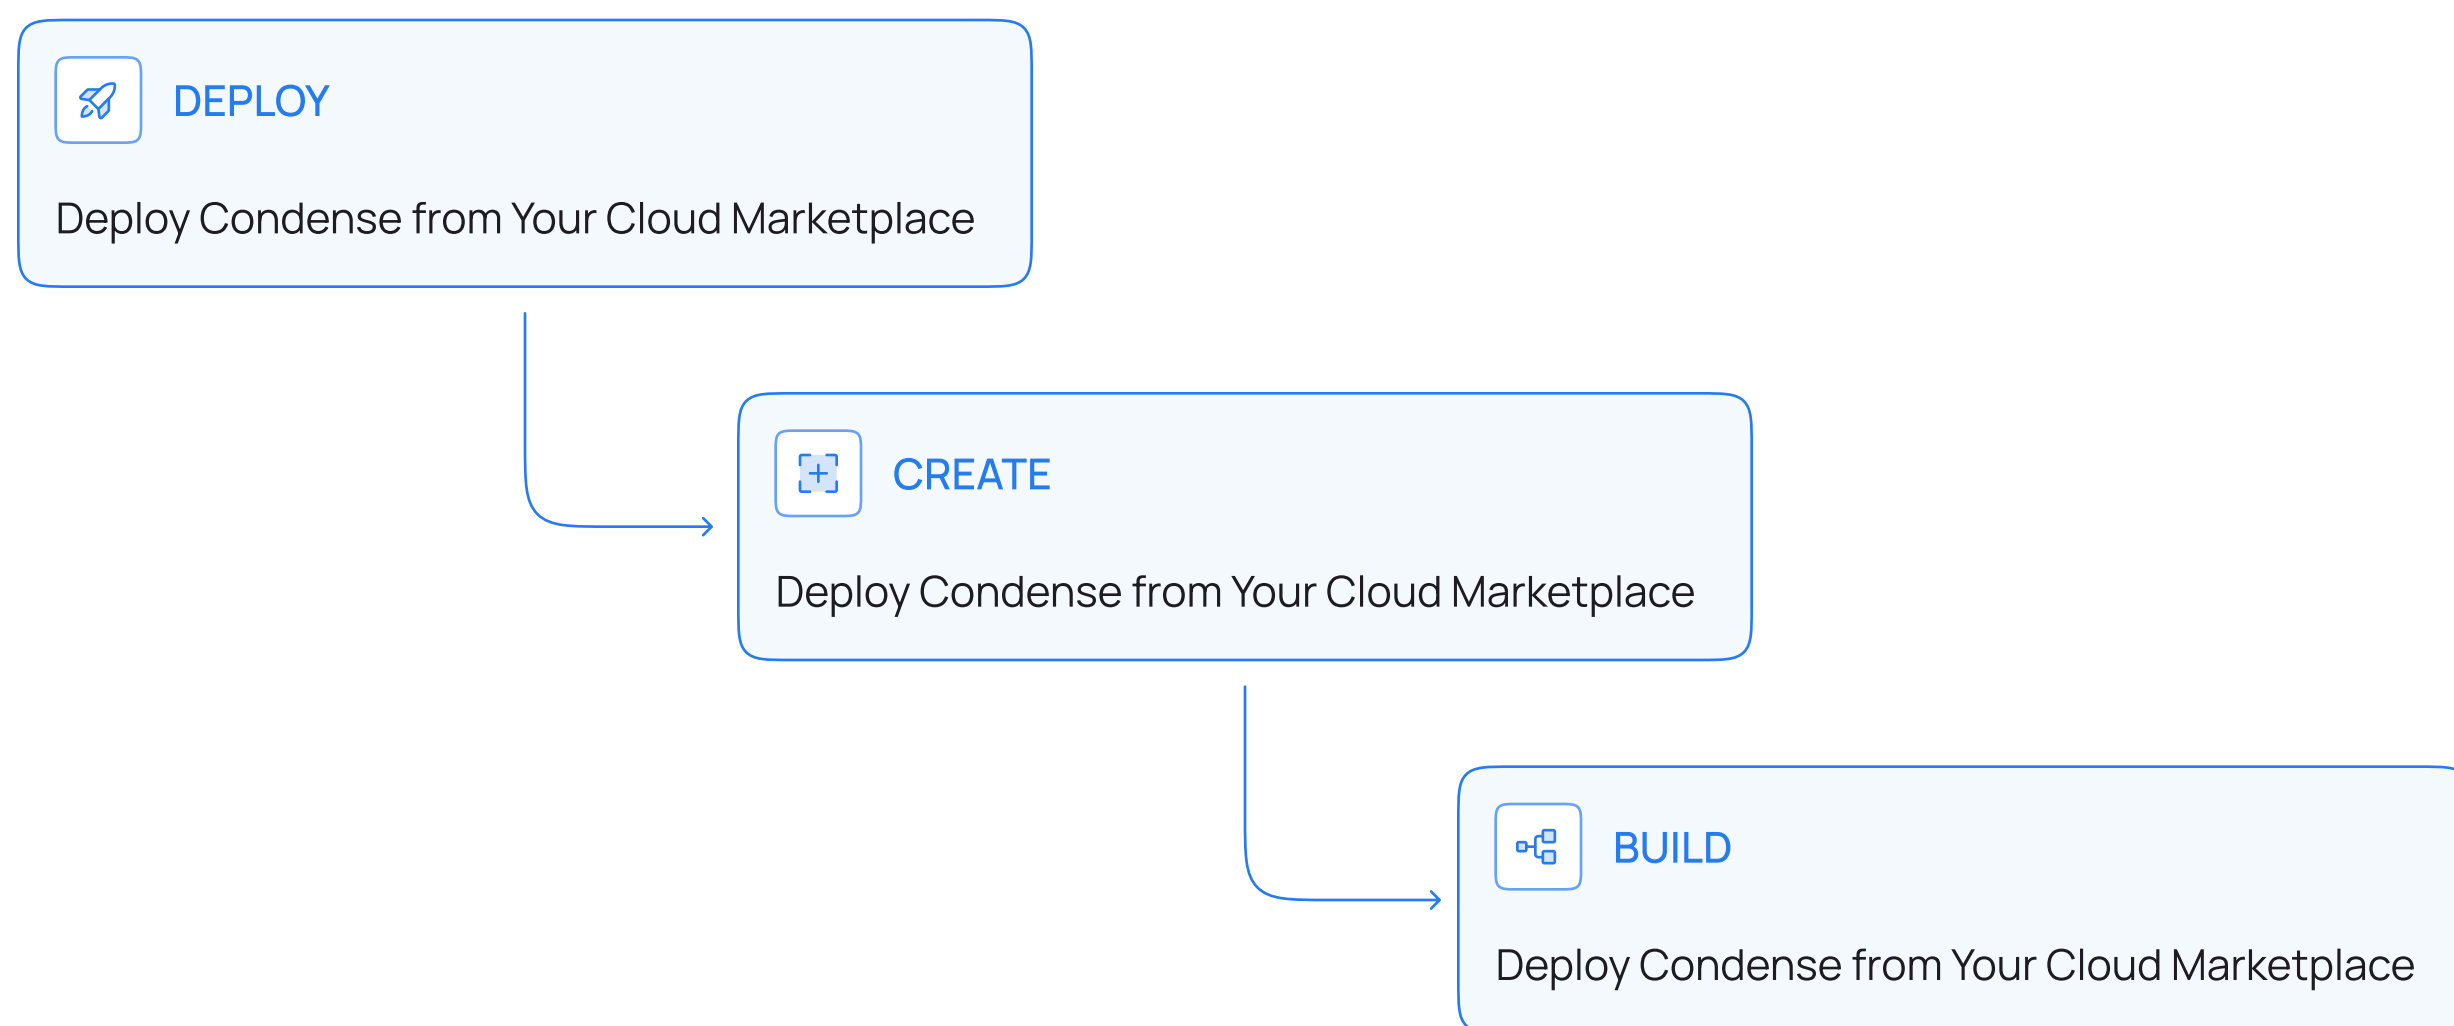
# Vertical Data Platform for Real-Time Businesses

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



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



# 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

