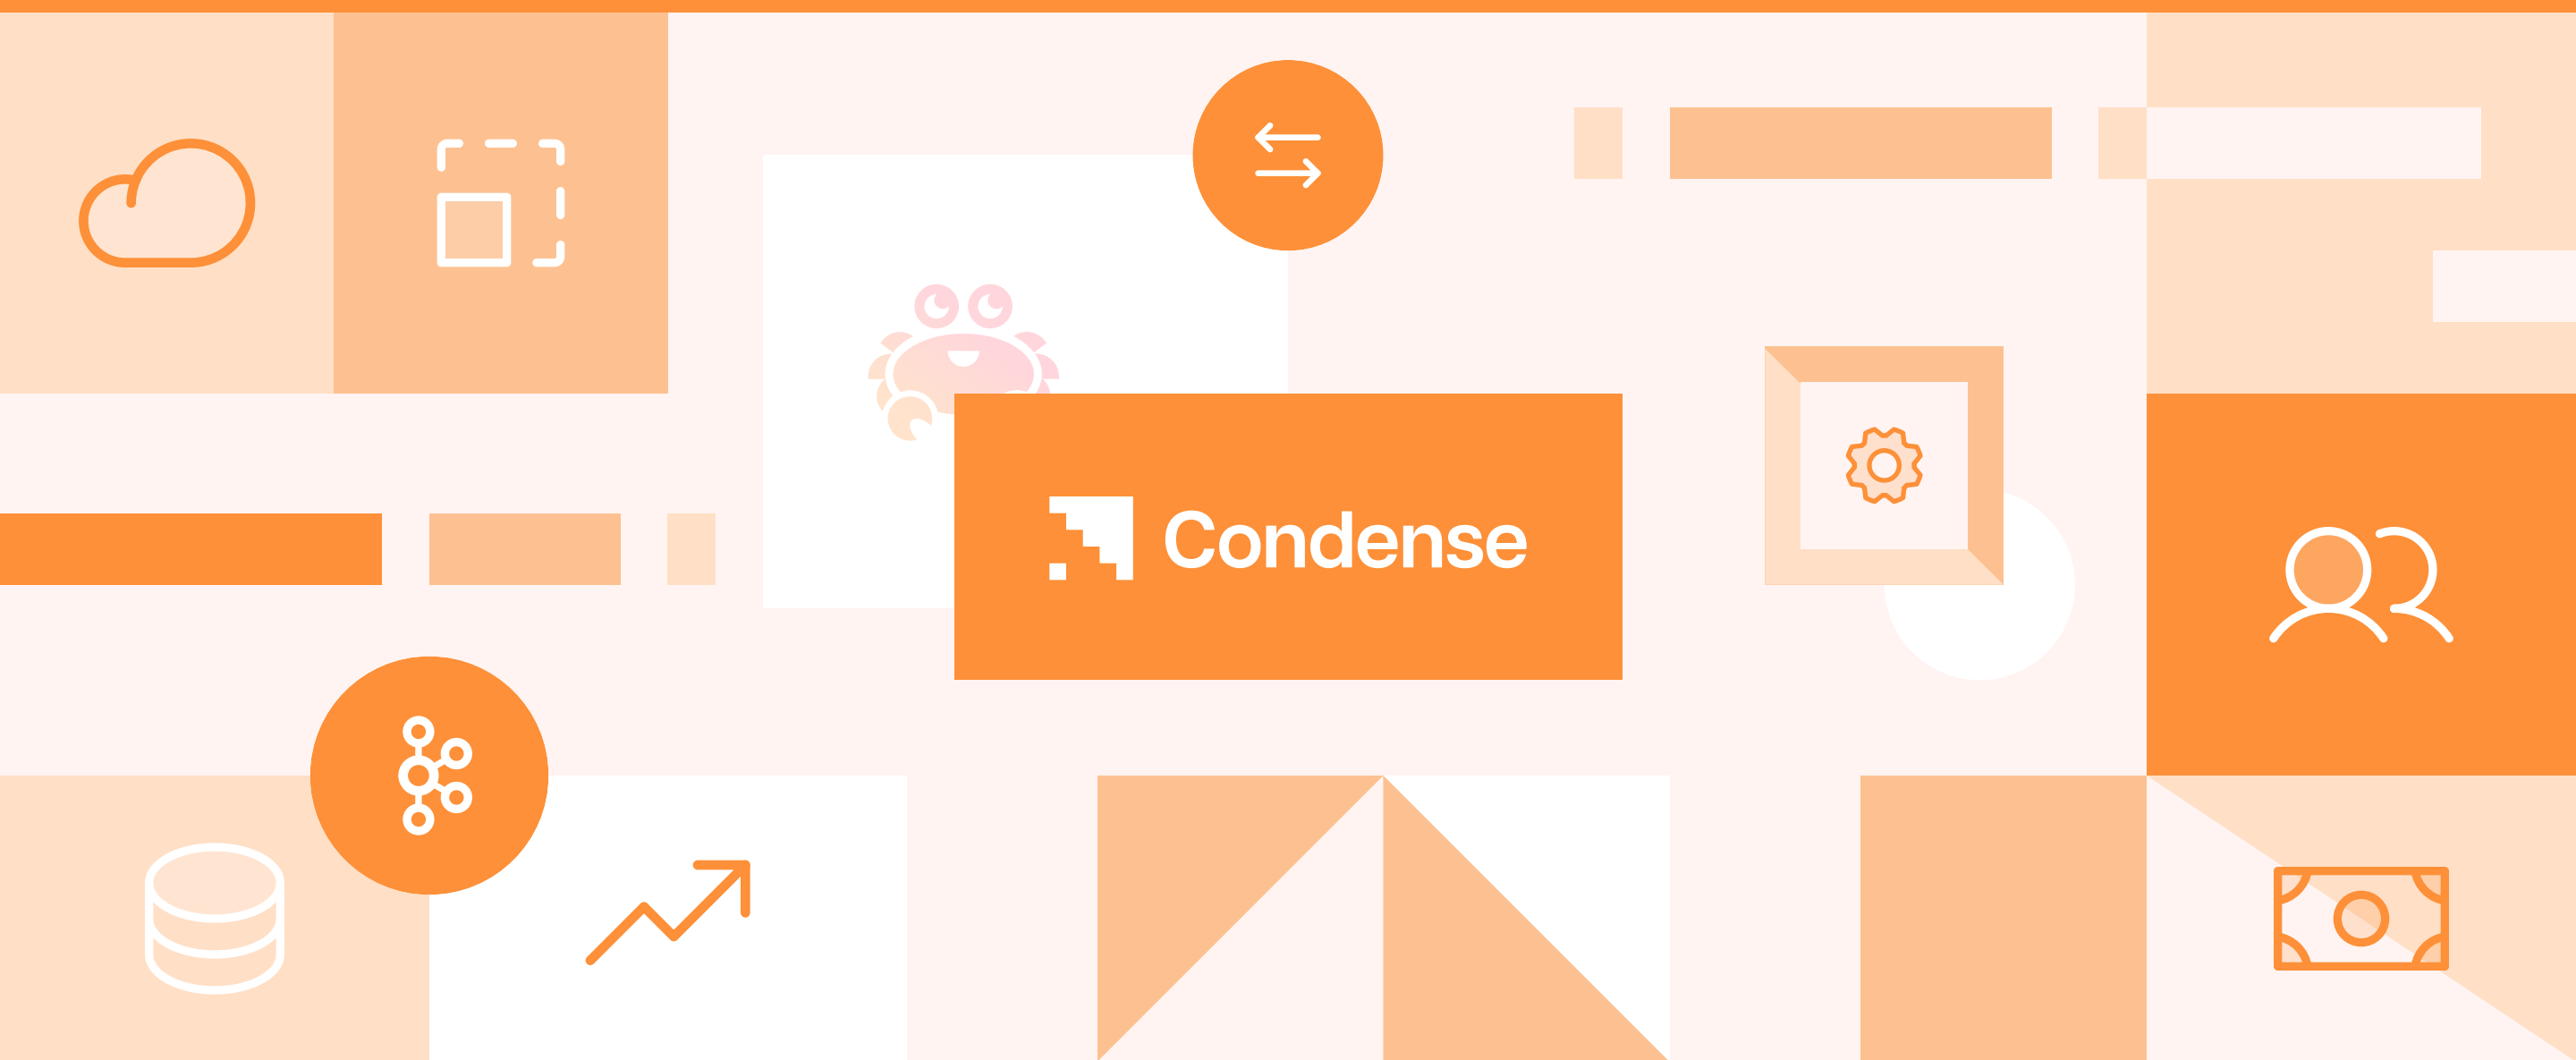


 Comparison

Data in Motion, Simplified at Scale

Evaluating Condense vs Aiven
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

 |  |  |  | 

All Rights Reserved 2026 @ Zeliot

How Does **Condense** Compare vs Aiven?

Condense and Aiven both provide managed Kafka, but they serve different operational goals. Aiven is a general-purpose managed service that excels at hosting standard open-source Kafka brokers; however, it leaves the "Application Gap" to the user, requiring separate clusters like Flink or custom microservices for data processing. Condense is a complete streaming platform that merges high-performance brokers with a native execution layer for Java, Python, or Go. By running entirely within a private VPC, Condense scales the brokers and the processing logic as a single unit, removing the need for separate infrastructure and eliminating the network costs of moving data between different vendor services

> Architecture

FEATURE	CONDENSE	AIVEN
System Design	A unified environment that bundles Kafka, Kubernetes, and the App Runtime into one integrated stack	Platform design that provides a consistent management layer to orchestrate and automate a wide range of standard open-source technologies (Kafka, Flink, Redis, etc.) across multiple clouds
Platform Model	Managed App Runtime: Custom transforms run in a managed, serverless-style environment. Condense handles containerization, state management, and orchestration	Decoupled: Requires separate services (like Aiven for Flink) or managing external apps to process data streams
Scaling Scope	Full-Stack Autonomy: Automatically scales brokers, connectors, and custom-code transforms (Java/Python/Go) based on real-time load	Infrastructure-Centric: Scales broker nodes and disk, but scaling the processing power for code remains a manual or external task
Cloud Storage	Native Object Storage: Direct, optimized offloading to S3/GCS/Azure Blob for infinite data retention	Tiered Storage: Available as an add-on; supports S3 but requires specific plan tiers for production
Enterprise BYOC	100% BYOC Native: Specifically engineered to run in your VPC to eliminate "SaaS" networking taxes and maintain 100% data sovereignty	Tiered BYOC: Primarily a SaaS-based model; BYOC is typically restricted to higher enterprise tiers
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	Vertical Scaling: High-performance VMs with plan-based IOPS and memory scaling

> Ecosystem

FEATURE	CONDENSE	AIVEN
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	30+ General: Robust, horizontal Kafka Connectors for general enterprise databases. Any domain-specific logic must be engineered from scratch
Operational Effort	Zero-Ops Application: Eliminates "Microservice Sprawl." You simply write the logic; the platform autonomously runs, scales, and monitors the entire application pipeline	Low-Ops Infrastructure: Removes Kafka complexity, but leaves the "Application Toil" of managing the external systems that process data
Apache Kafka API	100% Native: Built on upstream Kafka 4.0/ KRaft	100% Native: Strong commitment to open-source purity with zero forks
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 DevOps: Focuses on CLI and Terraform. Lacks an integrated environment for developing stream logic
Apache Flink Integration	Native Runtime Alt: Replaces Flink for most cases with a native, event-driven microservice runner	Managed Service: Offers a distinct, managed Flink service for complex stateful processing
Simplified Processing	Low-Code / No-Code/Full code): Visual Transform builder using utilities and prebuilt domain transforms (e.g., Geofence, trip formation). Also inbuilt full code git backed IDE	Infrastructure Only: Users must build and host their own transformation logic separately
Observability	Inbuilt Full-Stack: Trace an event from ingestion through custom code to the final destination in one unified dashboard	Broker Observability: Strong metrics for cluster health, but application-level tracing usually requires external APM tools


> Support & Compliance

FEATURE	CONDENSE	AIVEN
Access Experts 24/7	Direct Engineering: Built in support widget to raise ticket. Condense offers support priority based TAT for resolution	Tiered Support: Traditional ticket-based support; response times vary by plan level
Cloud Provider Availability	AWS, GCP, Azure (Customer Subscription)	AWS, GCP, Azure, DigitalOcean, UpCloud
Security	Own Cloud -Sovereign: Inherits your IAM and security groups; 100% data localization	Hardened SaaS: Fully certified (SOC2, HIPAA, GDPR) with VPC peering/PrivateLink options
SLAs	99.95% Full-Stack: Covers the brokers, the processing logic, and the connectors under one SLA	99.99% Infrastructure: Excellent uptime guarantee for managed service availability
TCO	Predictable BYOC: Pay for compute (vCPU/hr) inside your VPC. No throughput or egress fees; uses your existing cloud credits (EDP)	Usage-Based: Pricing scales with disk and volume. Often incurs network costs for data moving into the vendor's account

Why Switch to Condense?


TRANSITION FROM INFRASTRUCTURE HOSTING TO APPLICATION RUNTIME

Aiven is a managed hosting service that provides stable brokers but stops at the transport layer. To build an actual application, teams must still provision and manage separate services like Aiven for Flink or external microservices to process the data

 Condense merges these layers into a single runtime. It hosts both the high-performance messaging and the business logic in one environment. This allows developers to deploy production-grade code directly on the stream, removing the need to build, wire, and maintain a separate processing tier

AUTONOMOUS FULL-STACK SCALING

Scaling in Aiven is typically an infrastructure-level task where users manually change service plans or use APIs to add nodes. However, scaling the processing power for the code that uses the data remains the user's responsibility

 Condense provides Autonomous Scaling for the entire stack. The platform monitors real-time consumer lag and event-level throughput to automatically provision compute for custom Java, Python, or Go transforms. Your processing power scales in lockstep with your data volume and automatically scales back down during quiet periods to keep costs efficient

Why Switch to Condense?

PRIVATE CLOUD SOVEREIGNTY (BYOC)

While Aiven offers a BYOC model, it is often a complex, advanced-tier feature, with many users still residing in Aiven-managed accounts where data leaves their network

- Condense is engineered as a BYOC-native platform. It deploys 100% of the data plane including brokers, logic, and state inside your private VPC (AWS, GCP, or Azure). This eliminates "network taxes" (egress fees), keeps data sovereign, and lets you apply your existing cloud security policies and enterprise discounts to all streaming resources

VERTICALIZED ECOSYSTEM VS GENERIC CONNECTORS


Aiven provides a robust but horizontal platform, offering generic Kafka Connectors for moving raw bytes between systems. Any industry-specific logic must be engineered from scratch

- Condense offers a Vertical Ecosystem with pre-built, domain-aware transforms specifically for industries like Mobility, IoT, and FinTech. Whether it is parsing VINs, decoding telemetry, or geofencing, these ready-made assets allow teams to move from prototype to production months faster than building everything on a general-purpose pipe

UNIFIED DEVELOPMENT ENVIRONMENT

In an Aiven setup, development and observability are fragmented; you manage brokers in one console, code in an IDE, and processing health in an external APM tool


- Condense embeds a Full-Code IDE and AI-Assisted Layer directly into the platform. Developers can write, test, and deploy production transforms with built-in version control and state management. The unified dashboard provides full-stack observability, allowing an event to be traced from its raw ingestion point through your custom code to its final destination in a single view



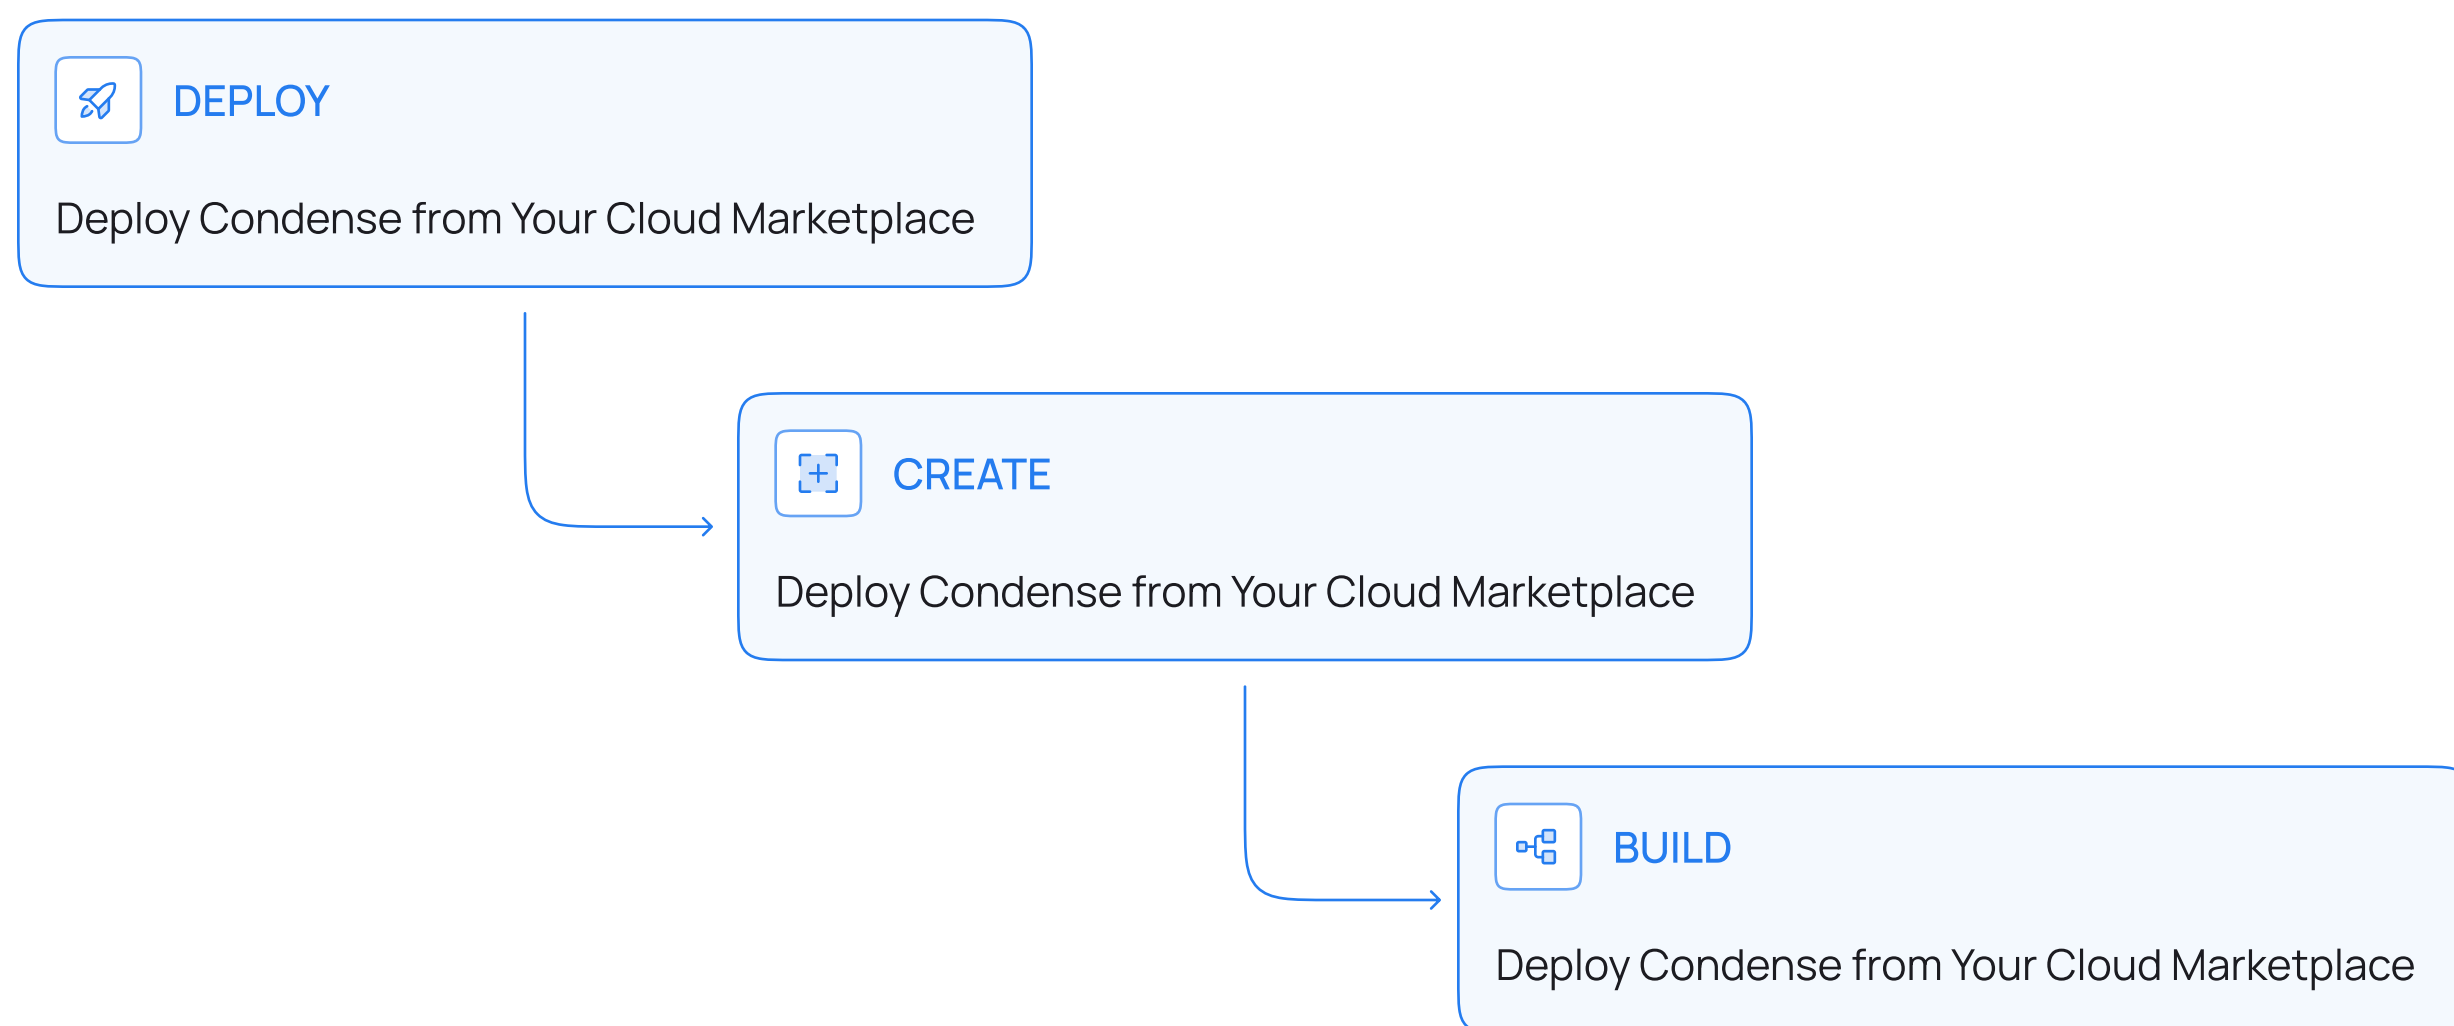
Vertical Data Platform for Real-Time Businesses

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

