CONFLUENT

Data Streaming Platform

The Key to an Evolutionary, High-Velocity Organization

JON MCCULLAGH-VINES

Contents

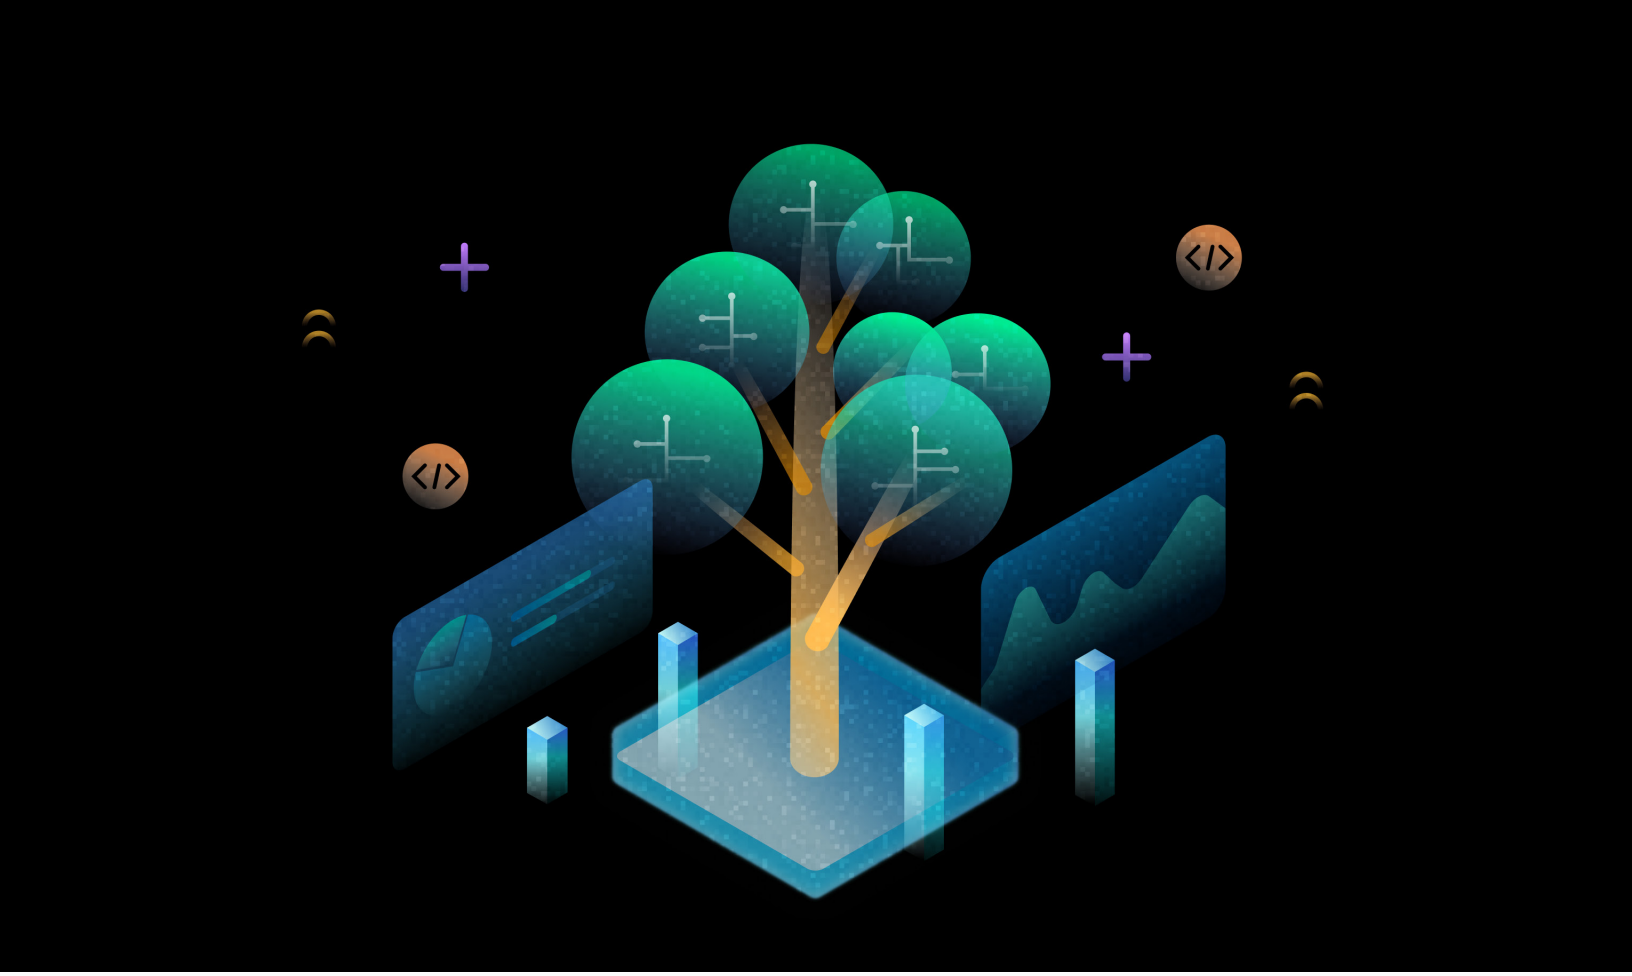

TO THRIVE IN TODAY'S FIERCELY COMPETITIVE LANDSCAPE, businesses require the ability to make intelligent, data-driven decisions in the moments that matter. However, traditional technology and data architectures create systemic friction that delays critical insights, stifles innovation, and increases operational risk. This technical debt prevents organizations from keeping pace with modern business demands and realizing their full potential.

A strategic shift towards an evolutionary, resilient architecture, supported by a data streaming platform, is essential. This architectural evolution transcends the operational and analytical data planes and powers the potential of AI to create an intelligent, data-driven enterprise. Providing this data backbone through a data streaming platform enables tangible business value through accelerated speed to market, enhanced operational resilience, real-time analytical data availability, and a superior developer experience.

This paper provides a north star for technology leaders to champion this critical modern streaming data backbone. Guiding this cultural and mindset-shift is the key to enabling enterprises to transition from a brittle, reactive state to a proactive, data-driven future.

The Strategic Imperative: Building a High-Velocity Enterprise

The Data Mess: An Inhibitor to Enterprise Agility

For decades, complex point-to-point integrations created a brittle system defined by latency bottlenecks and reliance on slow, scheduled batch processes. As business cycles have accelerated, accepting latency and batch-only is no longer an option. The cost of this data latency has become clear, resulting in lost customers, missed opportunities, and exposure to real-time threats.

This web of integrations between applications creates a tightly coupled architecture. A failure in one service can cascade and may lead to system-wide outages. This inherent brittleness makes the entire technology estate fragile and challenging to change, directly opposing the business's need for agility. Of concern is the ever-growing demand for data to fuel AI initiatives, with businesses expecting their data footprint to grow by over 40%[1], being built on this shaky foundation. This tangled, expensive, and brittle web of integrations can be viewed as the Data Mess. Alongside organizational culture, lack of executive buy-in, or insufficient funding, the data mess is one of the biggest inhibitors to enterprise agility.

The Four Pillars of a Modern Technology Strategy

To thrive, organizations must build a modern technology strategy based on four essential pillars:

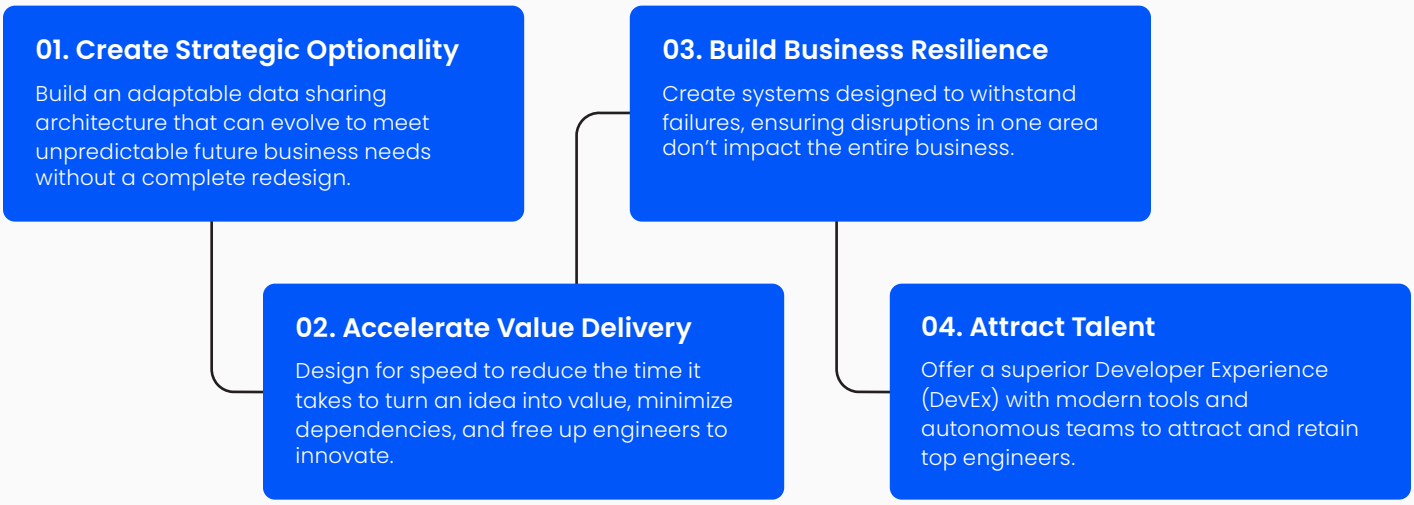

A core component to realizing these four pillars is through building a data architecture that provides ubiquitous, real-time access to the right data at the right time.

1 New NetApp Research Identifies Trade-off Between UK Business' Sustainability and AI Ambitions

Building the Streaming Data Backbone

Architecting the data-driven enterprise requires a fundamental mindset shift from viewing data as something static and stored (data at rest) to something dynamic and flowing (data in motion). Creating a data streaming backbone becomes the bedrock for enabling innovation through data within the enterprise, where every business action, be it a customer click, new order, or sensor reading — is an event. This stream of events provides an immutable, chronological log of everything that has happened in the business. With this approach, data becomes available for immediate action and decision-making.

Three Foundational Capabilities of a Modern Data-Driven Enterprise

To achieve the strategic goals outlined in the four pillars, a modern technology strategy must be built upon three foundational capabilities that guide every architectural and organizational decision. These are:

1. **Create Data as a Product:** Democratize data by treating it as a first-class product, owned and managed by the business domains that produce it. By adopting a "data as a product" mindset, data becomes trustworthy, discoverable, and equally accessible across the entire organization. Through this process, we enable organizational agility, foster innovation and dramatically improve time to market of new use cases driving business value.
2. **Empower Teams to Process and Govern Data Closer to the Source:** Empower engineering teams with the tools to process and govern data as it is created. Governance is an automated function that is "shifted left," to the source, ensuring quality and security by design. Data is a first class citizen of the organization regardless of where it is created.
3. **Deliver through a Self-Serve Platform:** Offer all the capabilities of the data streaming platform through a unified, self-service internal platform. Platform thinking enables us to provide a "golden path" for engineers. This abstracts away the underlying complexity of the technology and allows platform consumers to build and innovate with speed and autonomy using trusted, quality data.

A Data as a Product Mindset Drives Business Value

Adopting an event-driven architecture underpinned by a data streaming platform represents a crucial evolution in data strategy. If you haven't already, the time is now to introduce this organizational-wide "data as a product" mindset. This fundamental change is a catalyst for transformation, pivoting the focus from the transfer mechanism to the data itself with data producers taking full ownership of the data. The shift of ownership means the producing team now builds and delivers the data as a reliable and reusable product.

> Businesses that base decisions on data — not instincts or experience — are 19 times more likely to be profitable.
>
> **Forbes, Why Businesses Need Data To Make Better Decisions**

The value from data is measured through reuse and its availability for consumption by the whole of the organization. This reframing of the value of data results in accelerated speed to market, enhanced operational resilience, real-time analytical data availability, and a superior developer experience. Considering where to start? McKinsey asserts that *the lion's share of the potential value to a company comes from only 5 to 15 data products*[2] so focus on assessing those with the highest potential ROI.

Transcend the Operational and Analytical Estates

A profound impact of a data streaming platform is its ability to dissolve the wall between operational and analytical data. Operational systems are built to run the business, and analytical systems are built to understand it. Historically, these two planes have been entirely separate with a slow, brittle ETL wall existing between the two. Today, that slow, brittle ETL wall is no longer sufficient in meeting an organization's needs.

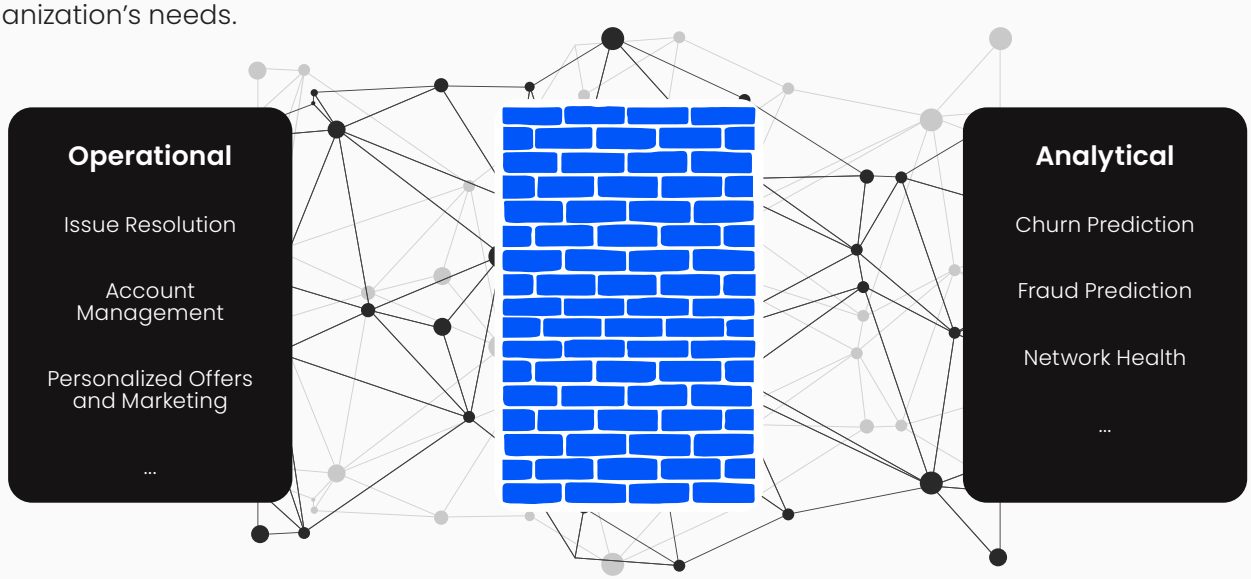

Operations and analytics connected by a web of complex point-to-point integrations and a slow, brittle ETL wall between the two.

With a data streaming platform, the same stream of events can drive operational microservices *and* analytical systems. The "OrderPlaced" event can trigger a shipping process (operational). This same event can also be streamed directly into a lakehouse to update a sales dashboard (analytical). By unifying the enterprise's data flow, business analytics are based on fresh, immediately available data, not stale, day-old batches.

2 The Missing Data Link: Five Practical Lessons to Scale Your Data Products

Consider what it means to really “shift-left”. Rather than being at the point it enters your analytical estates, view proximity to source as the entry point of data into the organization, whether that’s through an application, api endpoint, or 3rd party integration. This viewpoint will have a profound impact on your approach to data management and its impact on your customer experiences.

Power the Future of Enterprise AI

In addition to modernizing data infrastructure, adopting a data streaming platform creates the essential foundation for building an intelligent, data-driven enterprise that directly addresses some of the primary obstacles hindering most corporate AI initiatives.

Traditionally, AI models have been trained on stale, historical data through batch systems. A data streaming platform allows AI/ML models to be trained and perform inference on up-to-the-second data streams. Access to these data streams enables “Operational AI,” where we not only report but also immediately act upon insights within live business processes, such as in the moment fraud detection or instant customer personalization.

The “data as a product” mindset solves the “garbage in, garbage out” problem. When AI teams have self-serve access to a catalog of trusted, high-quality, and well-documented data streams, they can spend less time on data wrangling and more time building accurate, reliable models.

Data Streaming Backbone

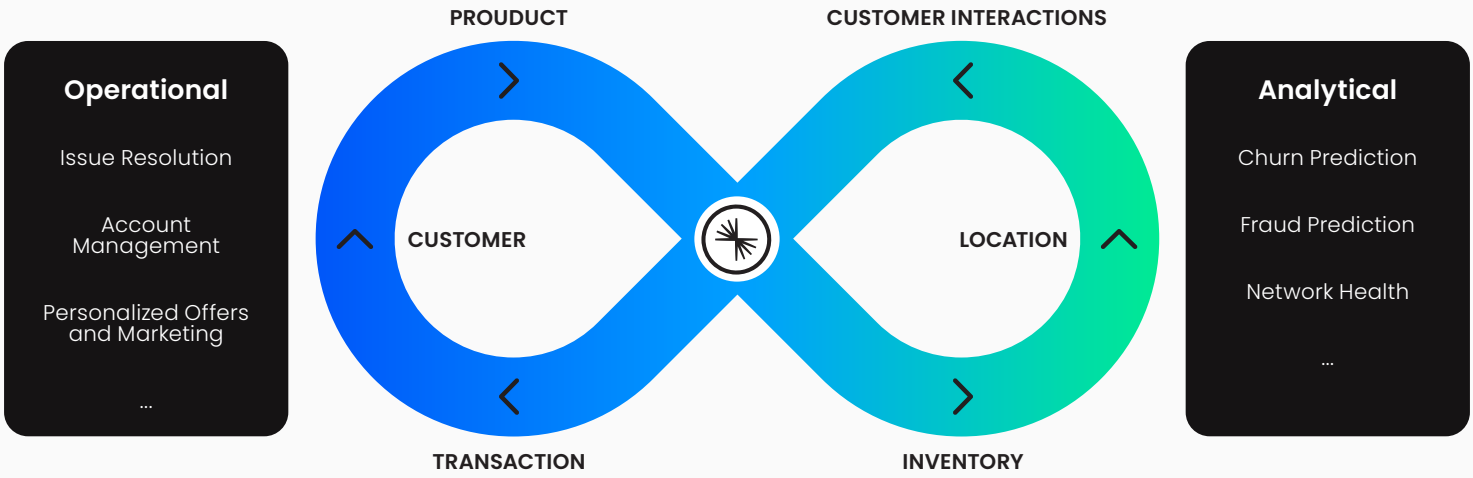

Meet in-the-moment customer needs and drive immediate business insight by delivering foundational data through a data streaming platform.

By transcending the operational and analytical planes, the data streaming platform creates a seamless feedback loop. An event can trigger a model that generates a prediction or a generated response. We then publish the response as a new event back into the platform. This new event triggers an immediate operational response, transforming AI from an analytical tool into an active participant in the business.

Learning from Past Transformations

The shift to a data-driven transformation is not an isolated phenomenon. It follows a proven path that mirrors other successful transformations that have moved organizations from slow, centralized, and brittle models to fast, decentralized, and resilient ones.

The Agile Parallel: From Batch Value to Continuous Flow: The move from overnight batch ETL to data streaming is the data equivalent of the Agile transformation. Traditional batch processing is like Waterfall development: value is delivered infrequently in one large, slow-moving lump. Data streaming is similar to Agile development: value is delivered continuously, allowing businesses to react and adapt in the moment.

The Cloud Parallel: From Manual Provisioning to Self-Service: Adopting a data streaming platform does for data what the cloud did for computing. Before the cloud, provisioning a server was a slow, manual process controlled by a central team. Today, engineers provision infrastructure via self-service APIs. Similarly, a self-service data streaming platform replaces the slow, ticket-based process of creating new resources or requesting access to data with an autonomous, self-service model, empowering engineers with autonomy and speed.

The DevOps Parallel: From Siloed Teams to Shared Ownership: The journey to data streaming mirrors the DevOps transformation. Just as a CI/CD pipeline automates the delivery of code, a data stream automates the flow of business value. Treating a data stream's schema as a version-controlled "Schema as Contract" is the new "Infrastructure as Code," bringing discipline and reliability to data integration.

Avoiding the Mistakes of Past Initiatives

Learning from past transformation helps to provide a strategic framework for success by highlighting common pitfalls to avoid and best practices to embrace for creating long-term value.

	DO	DON'T
Strategic Vision	Define a clear, value-driven strategy. Aim to create new opportunities and fundamentally change how the business operates.	Focus solely on cost reduction or confuse optimizing with transforming. Lack a clear, overarching strategy and vision for what you want to achieve.
Investment Strategy	View the investment as a long-term driver of value. Start with a single, high-impact value stream to prove value, build momentum, and scale incrementally.	Let adoption be driven by short-term budgets or financial engineering. Attempt a risky "big bang" rollout across the entire organization.

	DO	DON'T
Technology Approach	Bridge the gap between operational and analytical data to meet the insatiable demands of customers and AI agents.	Define your transformation by legacy technology or through "lift and shift" of a messaging platform, operational data store, or Kafka topics.
Implementation	Embrace a "buy" strategy for the core platform to enable teams to focus on building customer-centric solutions, creating an ecosystem designed for speed, ownership, and decoupling.	Build commoditized platform capabilities in-house unless there is a specific need.
Data	Invest in data governance and quality processes to ensure the data flowing through your new platform is trustworthy and accurate.	Ignore data quality issues and assume that migrating bad data will fix it.
People	Champion a culture of change and data literacy. Empower your teams with the necessary skills and resources to effectively use and manage the new platform.	Neglect the human side of the transformation, fail to provide training, secure buy-in, or address cultural resistance.

A Data-Driven Backbone for the Future

> The best time to plant a tree was twenty years ago. The next best time is now.
>
> **Chinese Proverb**

The journey ahead is a strategic awakening. We stand at a pivotal moment, poised to move beyond the limitations of the past and embrace a future where data is the lifeblood of innovation. The strategic shift to a modern, resilient architecture, powered by a data streaming platform, is the next step in our enterprise evolution. It is a transformation that echoes the paradigm shifts of Agile, Cloud, and DevOps, where we moved from slow, siloed operations to a culture of speed, autonomy, and shared ownership.

Adopting a "data as a product" mindset, powered through a data streaming backbone, shatters the walls between the operational and analytical data plane. This unlocks a continuous, trusted flow of information, empowering a high-velocity enterprise to proactively shape its future rather than simply react to change. It's the essential foundation for leading the charge in an age of AI and instant customer demands.

How Confluent Can Help

The time to act is now. The journey to a modern data-driven architecture is not just critical for survival, it is the pathway to competitive dominance. To help your organization take this first powerful step, Confluent has developed a Data Readiness Assessment to baseline your current maturity. By completing this assessment with a Confluent Executive Advisor, you can gain a crystal-clear understanding of where you are today and chart a tailored, ambitious roadmap to where you need to be. The future is waiting, and getting help on your journey is closer than you think.