

Turning Data Into Business Value in the Age of AI

How data streaming platforms provide a competitive edge

Contents

- 1 Why AI Needs Real-Time Data
- 2 Every AI Problem Is a Data Problem
- 3 Data Streaming Platform:
The Data Orchestration Layer for Your AI Stack
- 4 How Confluent Can Help
- 5 Take the Next Step

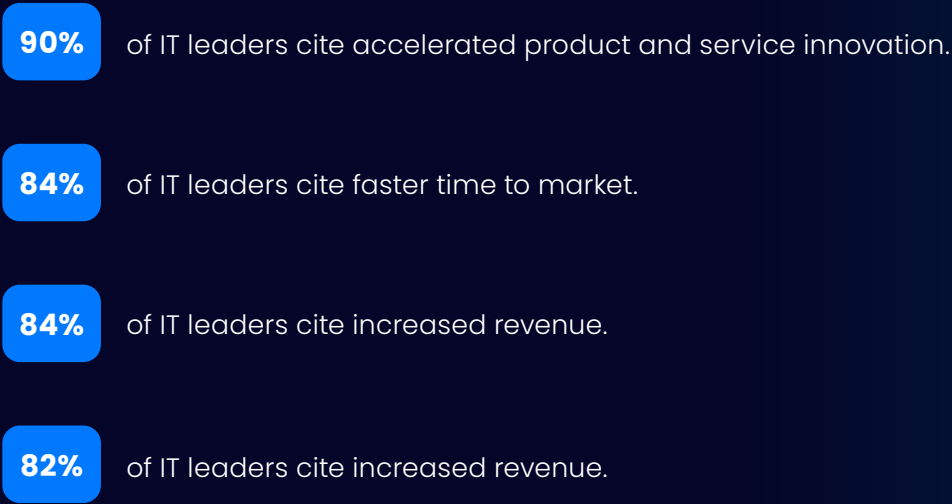
Introduction

DATA HAS LONG BEEN RECOGNIZED AS A STRATEGIC ASSET. But in today’s constantly evolving, AI-powered world, expectations for what data can and must do have changed. Businesses looking for a competitive edge must now have the ability to tap into trustworthy data in real time. As a result, many companies are shifting from using data to understand what happened to using data to drive what happens next.

For technical executives, this evolution isn’t just about adopting new tools: it’s about rethinking how data flows across their businesses. **To truly capitalize on AI, data must be fresh, trustworthy, properly contextualized, and available the moment it’s needed.**

That’s where a [data streaming platform \(DSP\)](#) comes in—powering real-time data flows that fuel AI applications, accelerate innovation, and unlock business value as fast as the world demands it.

Our [2025 Data Streaming Report](#) found that DSPs drive real business outcomes by making data instantly accessible whenever and wherever it’s needed:



LET’S DIVE IN TO FIND OUT HOW.

1 Why AI Needs Real-Time Data

FOR YEARS, BUSINESS INTELLIGENCE (BI) was the cornerstone of data-driven decision-making, offering valuable hindsight to guide strategy and operations.

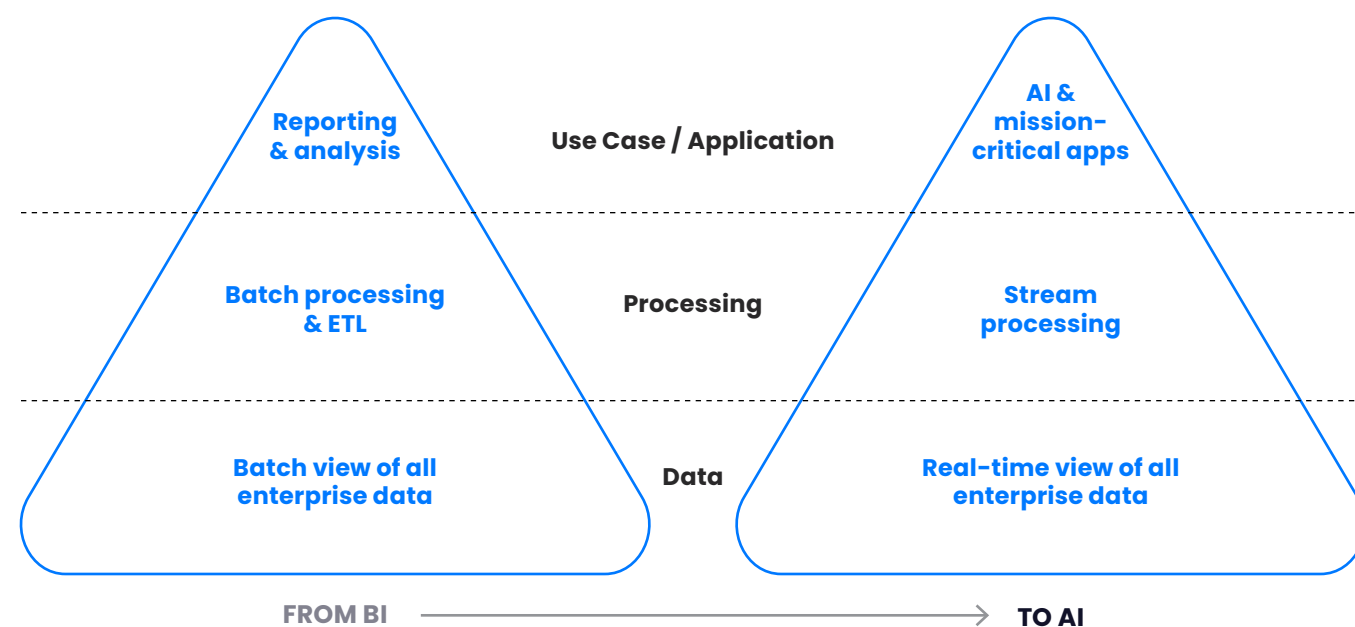
Data warehouses, central to the BI era, offered the powerful ability to bring data from across a business into one place, process it together, and serve it up through internal dashboards and reports. But here's the catch: Data warehouses are fundamentally batch-based. Data is collected, transformed, and updated on a schedule—often only once a day. That worked when the pace of business allowed for it.

The modern pace of business and rise of AI demand more than hindsight. We're now in an era—moving from BI to AI—where unlocking business value from data requires acting in the moment. Whether it's enabling chatbots to provide instant customer support, delivering

personalized recommendations on the fly, or stopping threats with fraud detection systems, real-time data is key to helping AI power this new era of business innovation.

However, AI didn't arrive in a vacuum. It's the next stage in the arc of technology evolution. As businesses have become more software-defined over time, software itself has grown more powerful, intelligent, and deeply embedded in daily business processes and operations. Data that once flowed into dashboards for manual human review is now fed into data platforms, stream processing engines, and AI applications that make decisions in milliseconds.

And as AI continues to evolve, easy and continuous access to trustworthy data in real time has never been more important.



BUSINESS IMPACT OF OUTDATED DATA

- AI models underperform due to outdated or incomplete data
- Customer experiences suffer from high latency, inconsistencies, or wrong recommendations
- Operations miss critical moments to act or adapt in real time
- Teams spend more time on data integration and fixing issues than driving innovation

The Evolution of AI

WAVE 1: PURPOSE-BUILT AI

Traditional machine learning focused on building statistical models and predictive capabilities for narrowly defined tasks. For these purpose-built models trained on a specific dataset, having access to high-quality, real-time data leads to more accurate predictions and faster decision-making.

WAVE 2: GENERATIVE AI

Generative AI (GenAI), driven by deep learning, marked a turning point. It uses reusable foundation models like large language models (LLMs) to generate content. But LLMs are trained on public data and lack domain-specific context, so when there's a knowledge gap, they may hallucinate and provide false or misleading answers. Data streaming provides LLMs with fresh, relevant enterprise data for more accurate, grounded outputs in retrieval-augmented generation (RAG) applications.

WAVE 3: AGENTIC AI

Agentic AI systems can think, adapt, and act autonomously. According to [Gartner](#)¹, “By 2028, 33% of enterprise software applications will include agentic AI, up from less than 1% in 2024, enabling 15% of work decisions to be made autonomously, and by 2029, agentic AI will autonomously resolve 80% of common customer service issues without human intervention—leading to a 30% reduction in operational costs.”

Agents bring something fundamentally new: dynamic, context-driven workflows. They're well suited for solving complex, interconnected problems in enterprise environments. Our [2025 Data Streaming Report](#) found that 40% of organizations are already in the pilot or deployment stage with agentic AI. But to power automation workflows, AI agents need constant access to what's happening now, not what happened yesterday. Real-time data gives agentic AI the context it needs to make autonomous, intelligent, timely decisions. Hardcoding won't work—agents need to perceive input, decide, act, and adapt in real time.

Real-Time AI Runs on Real-Time Data



SMARTER, FASTER DECISIONS

Real-time inputs bring enriched and complete data to fill context windows for AI to generate insights and act instantly—keeping pace with market shifts and operational demands.



GREATER EFFICIENCY

Real-time data reduces errors, improves automation, and optimizes resource allocation.



PERSONALIZED CUSTOMER EXPERIENCES

Real-time data enables personalization by analyzing customer interactions in real time to deliver tailored recommendations.

¹ Gartner Article, Intelligent Agents in AI Really Can Work Alone. Here's How, by Tom Coshov, October 1, 2024.

² Gartner Press Release, Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029, March 5, 2025.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

2 Every AI Problem Is a Data Problem

WHILE AI OFFERS tremendous opportunities for businesses—think tangible outcomes like improving productivity and creating new revenue streams—they’re too often held back by disconnected systems. The outcome? They fail to deliver the right data at the exact moment it’s needed.

Several data challenges get in the way of accelerating AI projects. At the top, according to our [2025 Data Streaming Report](#), are fragmented ownership of data across disparate systems (68%), limited ability to seamlessly integrate new data sources (65%), and insufficient infrastructure for real-time data processing (61%).

Enabling easy, continuous access to high-quality data—at the right time and in the right format—is pivotal to driving success with modern AI applications. **That’s because AI systems in a business context are only as good as the data that powers them and how quickly and reliably the data can be delivered.**

This calls for an unbroken loop of continuous data movement in real time, where operational systems feed AI applications and those applications then drive decisions and actions back into those same systems—all at enterprise scale and speed.

Failure to meet these demands isn’t just a technical issue. It’s a business risk. Security and governance issues, latency, data fragmentation, and blind spots degrade the AI experience and directly impact customer satisfaction, operational efficiency, and innovation velocity.

Creating this loop requires breaking through silos, bridging the operational and analytical data divide, and making data available when and where it’s needed. But that’s easier said than done.



Building an AI-ready organization starts with solving data infrastructure challenges.

The AI Data Mess Problem

Data has traditionally resided in two domains. The first is the operational estate, which includes databases, customer relationship management (CRM) and enterprise resource planning (ERP) applications, and billing systems for daily business operations. The second is the analytical estate, comprising data warehouses for after-the-fact analysis to inform business decisions, create reports dashboards, and more.

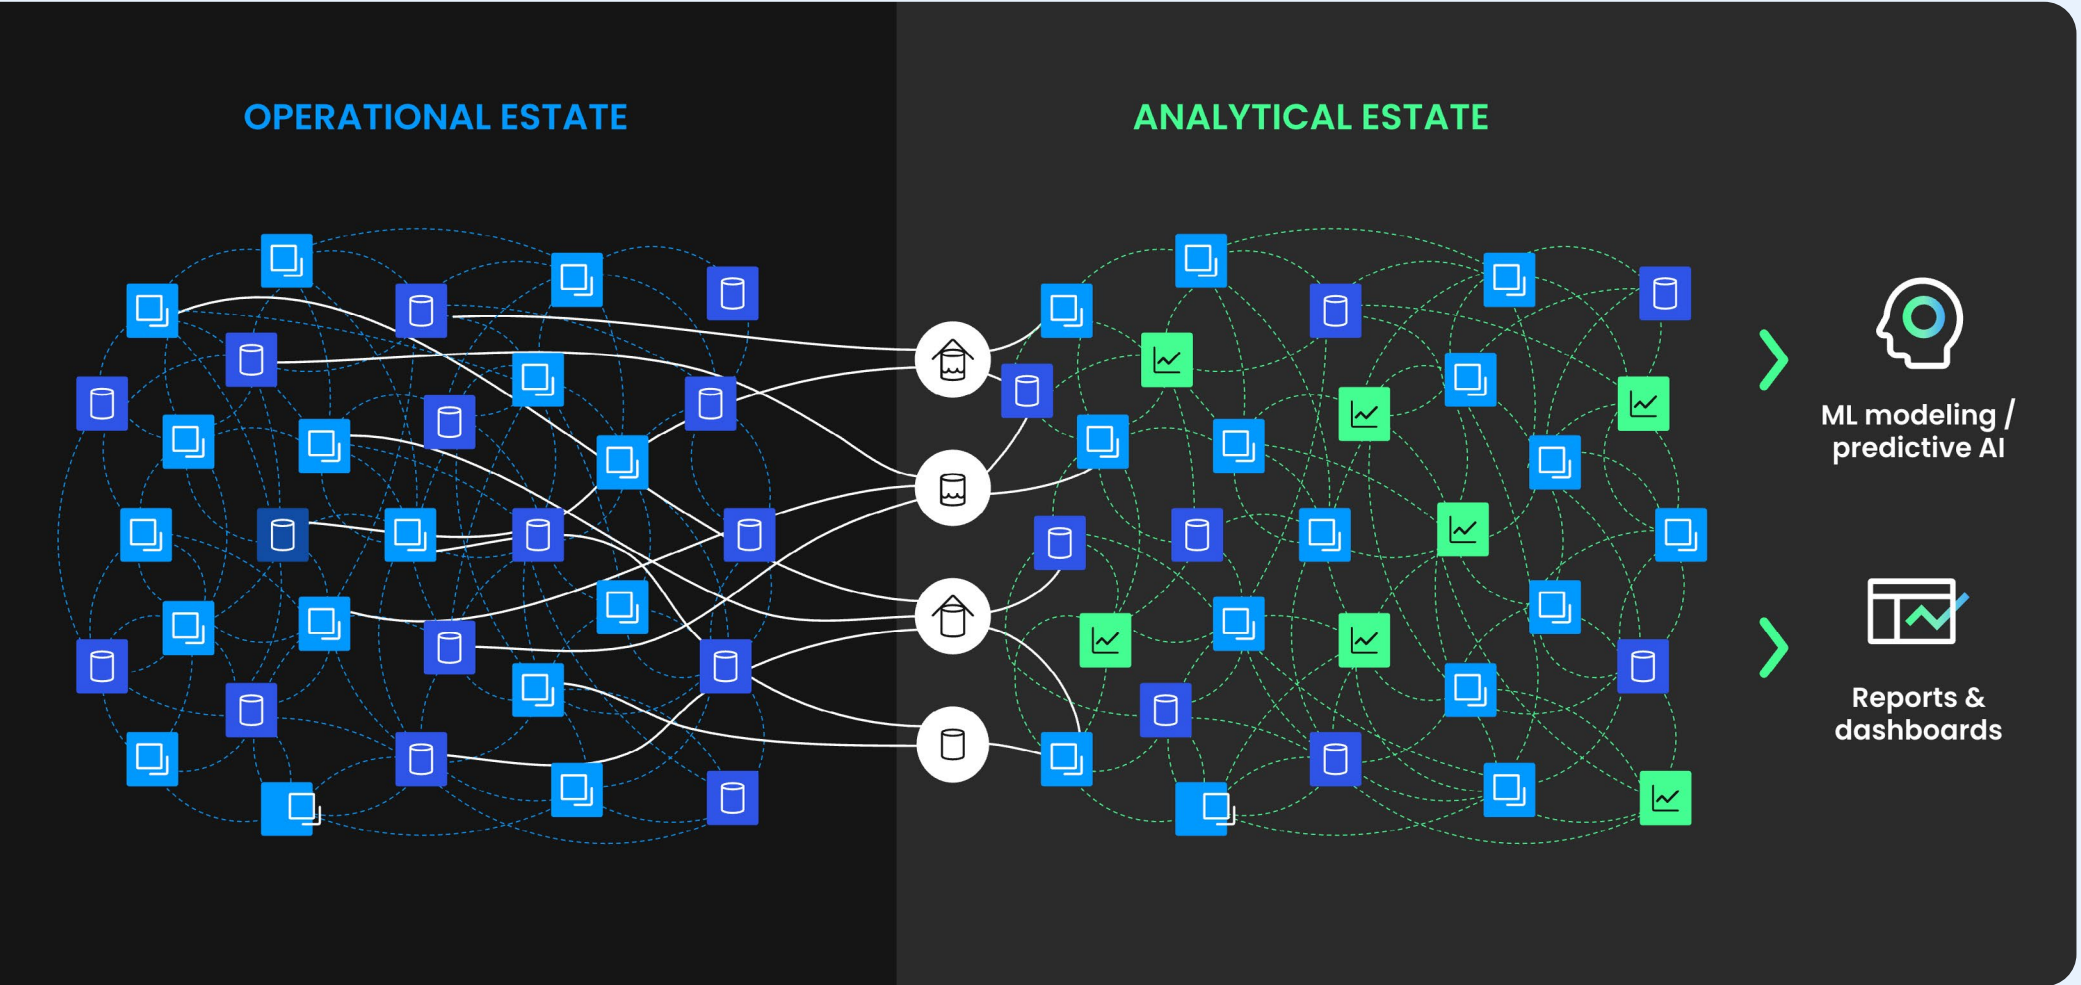
Over the years, there has been a proliferation of new business applications, software as a service (SaaS) offerings, and specialized tools alongside a hybrid operating model that integrates cloud services with on-premises systems. At the same time, development teams are transitioning from building monolithic applications to event-driven microservices, necessitating seamless integration for real-time communication.

This complexity extends to the analytical domain. Data complexity has escalated as data needs have multiplied. As organizations seek to consolidate data across silos of operational sources for various stakeholders, what used to be a handful of batch systems, messaging middleware, and APIs has evolved into a data mess—a growing sprawl of point-to-point integrations.

The results?

- Increased costs and complexity
- Hindered rate of innovation and vendor lock-in
- Stale data and integration challenges, which make it difficult to deliver real-time context for AI
- Increased latency
- Reduced data quality and trustworthiness
- Data governance and security issues

Traditional ML and AI are stuck in the analytical estate, far from real time

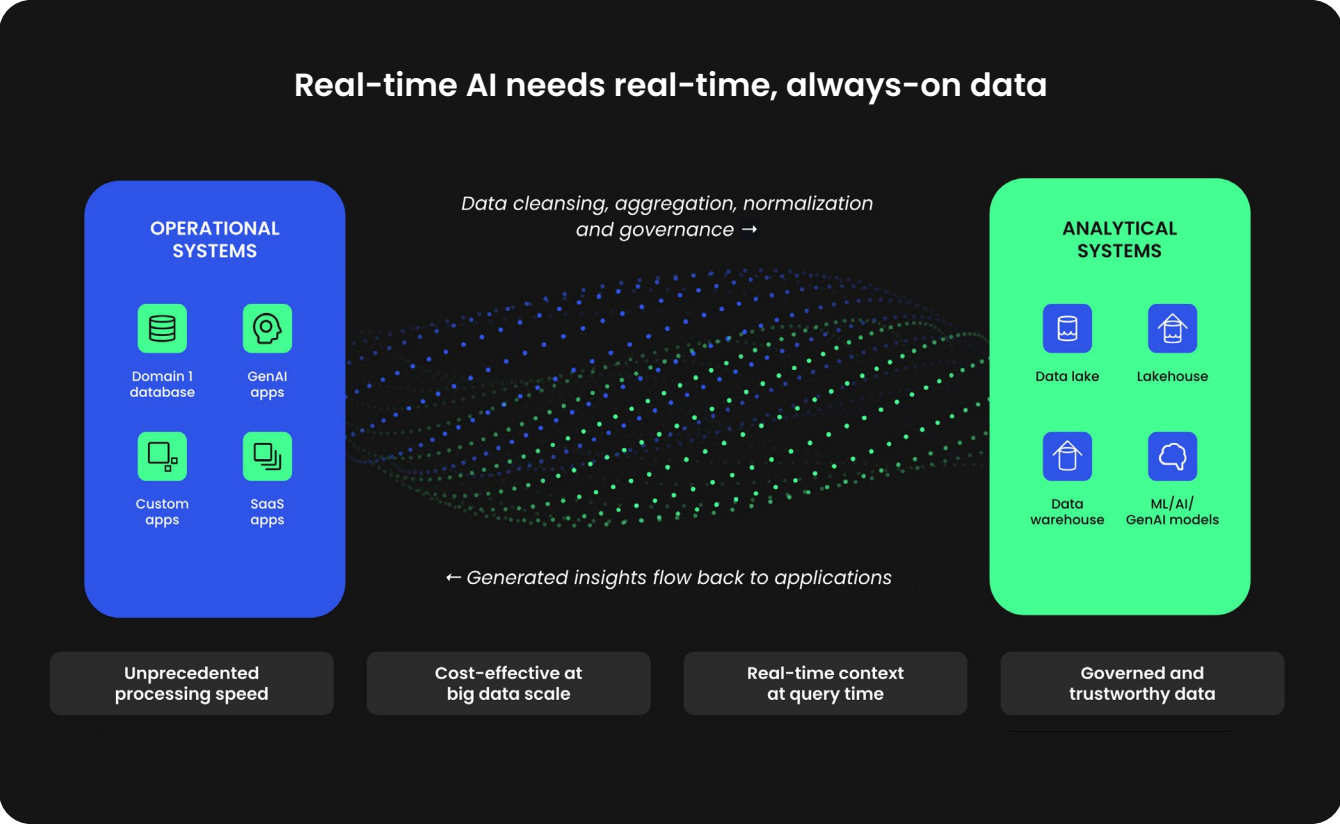


3 Data Streaming Platform: The Data Orchestration Layer for Your AI Stack

WITHOUT A PURPOSE-BUILT PLATFORM to unify and stream data continuously, your business will fail to unlock the full promise of AI. That’s where a data streaming platform comes in—DSPs make it possible for AI systems and applications to access the right data at the right time so that they can work smarter and faster.

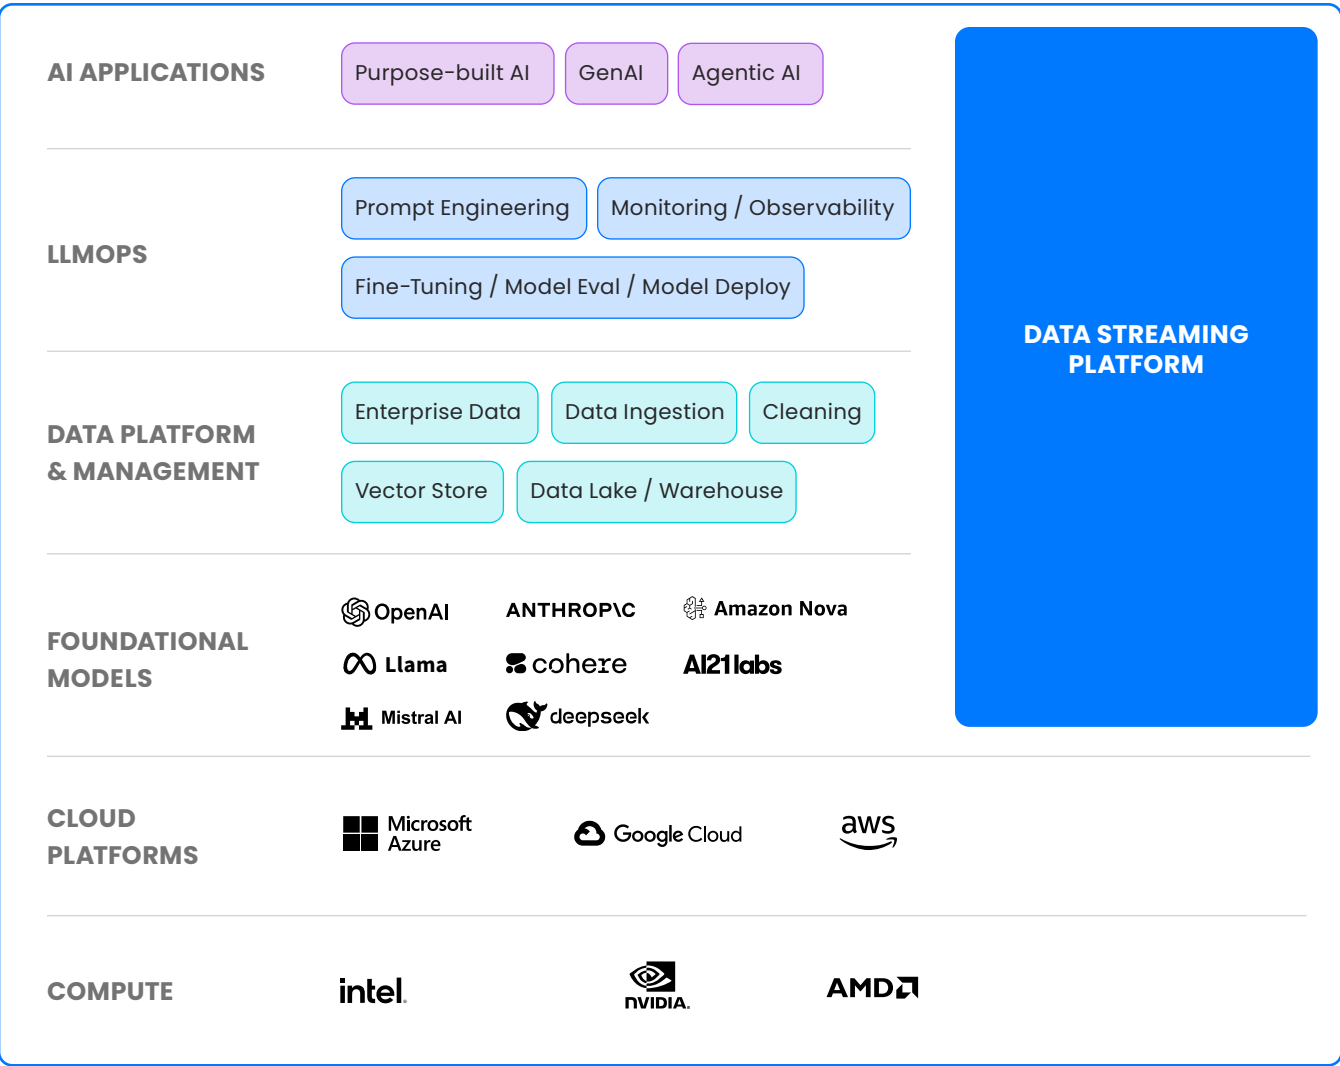
DSPs deliver the always-on, real-time foundation that powers AI applications at scale. Unlike traditional request-response architectures that introduce bottlenecks and stale data, DSPs enable continuous, low-latency access to the information that modern AI systems, models, applications, and agents need.

According to our [Data Streaming Report](#), 89% of IT leaders see DSPs easing AI adoption by helping them directly address hurdles in the areas of data access, quality assurance, and governance. And 87% say DSPs will be increasingly used to feed AI systems real-time, contextual, and trustworthy data.



At the heart of every scalable AI stack is a real-time DSP that helps:

- **Integrate** disparate operational data across an enterprise in real time for reliable, trustworthy use.
- **Organize** data into vector embeddings for RAG to help GenAI applications generate the most accurate, up-to-date information and prevent hallucinations from negatively impacting customers.
- **Provide** real-time access by eliminating batch processing delays and ensuring that AI systems make decisions based on the latest available data.
- **Enable** LLMs, agents, vector stores, etc. to be treated as modular components that can easily be substituted as technology improves—preventing vendor lock-in and allowing organizations to take advantage of the latest and greatest AI tools.
- **Ensure** secure, governed data for AI with data quality rules, access controls, and auditability at every step.



A DSP is effectively the data orchestration layer for your AI stack, unlocking data flow across cloud providers, LLMs, AI tools, and data platforms as well as customer-facing AI applications.

Business Benefits of Building a Scalable, Resilient AI Stack

Faster time to value: Deploy AI solutions quicker with an architecture that supports real-time data and streamlined integrations across systems.

Cost efficiency at scale: Optimize infrastructure and compute resources by decoupling systems and scaling only what’s needed—when needed.

Improved decision-making: Real-time data inputs ensure that AI-driven decisions reflect the latest business context.

Future-proof flexibility: Support new AI use cases (like generative and agentic AI) with an event-driven architecture that can evolve and allow you to adopt new models and tools as they become available.

Enterprise-wide integration: Break down data silos and create unified intelligence across departments.

Competitive edge: Turn AI into a strategic advantage by building a stack that powers real-time, intelligent products, services, and customer interactions.

DSPs help organizations build a range of AI applications by enabling easy access to real-time data.



PURPOSE-BUILT AI

DSPs ensure that models are continuously trained on fresh, high-quality data for accurate fraud detection, forecasting, and more. DSPs also help maintain a consistent view of data across all systems, data warehouses, and data lakes.

Use cases organizations can build:

Inventory tracking and forecasting, predictive maintenance, risk assessment, enhanced healthcare diagnosis, and much more.



GENERATIVE AI

DSPs help build a real-time, contextualized, and trustworthy knowledge base to power GenAI applications and mitigate LLM hallucinations.

Use cases organizations can build:

Customer service chatbot, coding copilot, internal knowledge assistant for employees, live sales coach, and much more.



AGENTIC AI

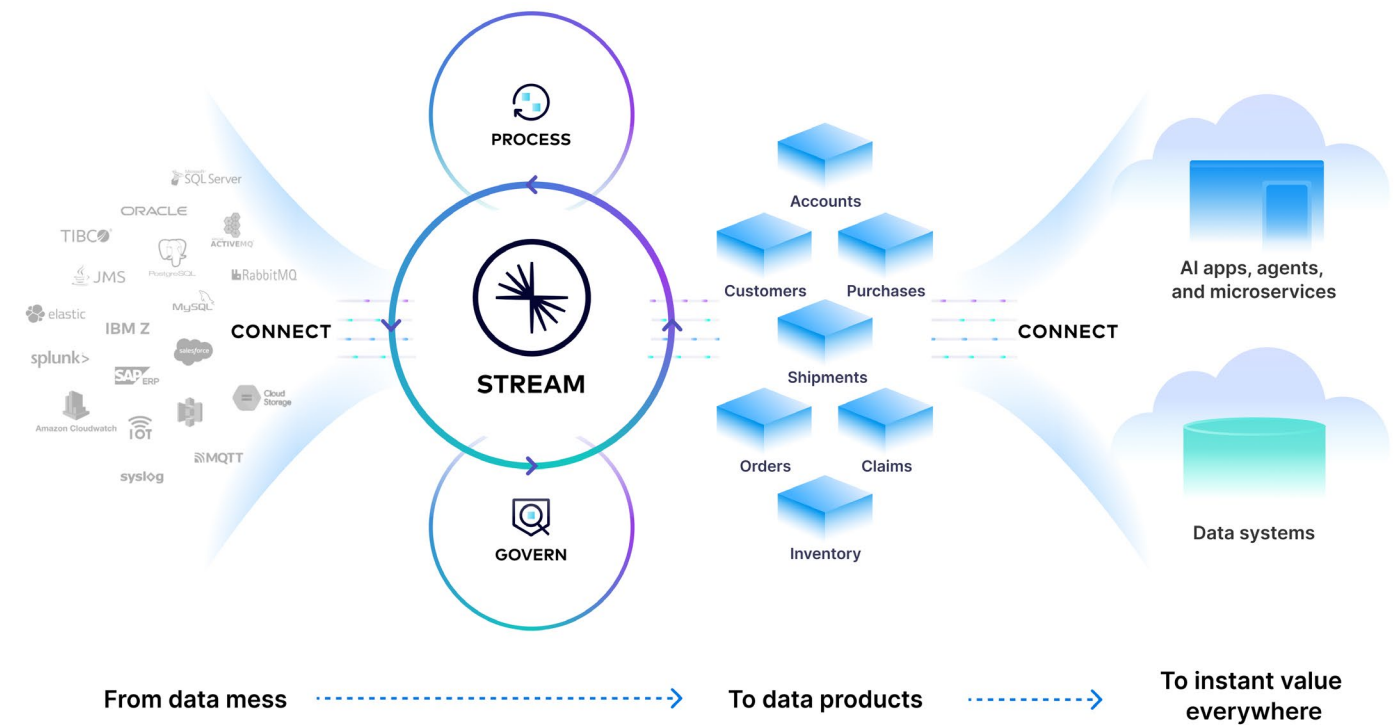
DSPs enable dynamic, context-driven workflows for agents, allowing them to access and share data effortlessly and make decisions based on real-time information.

Use cases organizations can build:

Multi-agent sales development representatives, web scraping agents, mortgage underwriting agents, onboarding agents, and much more.

4 How Confluent Can Help

HIGH-VALUE GENAI AND AGENTIC AI APPLICATIONS need high-value data. Confluent’s data streaming platform allows you to tap into continuously enriched, trustworthy data streams to quickly build and scale real-time AI applications—so that you can bring innovative AI experiences to market faster.



STREAM

Make data streams first-class citizens and share them with AI systems and applications in real time.



CONNECT

Integrate disparate data from any environment, with 120+ pre-built and custom connectors, including integrations with leading AI tools.



PROCESS

Use Apache Flink® stream processing to enrich data streams on the fly and combine publicly available datasets with proprietary business data.




GOVERN

Centrally manage, audit, and apply data policies with Stream Governance to ensure that data for AI is secure and verifiable. Secure data in transit with private networking.


Confluent turns the data mess into [data products](#) that—unlike raw data—are high-quality, reliable, and reusable data assets for your AI use cases.

Don't just take our word for it. See how customers across industries are leveraging the Confluent data streaming platform to drive business value.




“We rely on Confluent to **stream GenAI and agentic AI outputs across our architecture for our customers**, tracking token usage and enabling accurate billing. Confluent makes it easy for us to iterate quickly and build new AI features in days instead of weeks.”

Adam Watkins
Co-Founder & CTO
Reworkd



“To save our users time, write faster, and boost creativity... we use Confluent to **share new content and updates in real time**. Our product and engineering teams use data products without worrying about infrastructure. This speeds up our GenAI use cases.”


Daniel Sternberg
Head of Data & AI
Notion



Hours	Food	Clothing	Home	Other	Total
Ten-Bam	\$108,147	-	-	\$23,898	\$622,383
Bam-Bam	\$678,758	\$233,193	\$93,283	\$33,873	\$1,039,007
Bam-Bam	\$745,696	\$302,129	\$123,107	\$56,783	\$1,227,715
Bam-Bam	\$145,769	\$289,902	\$183,219	\$96,932	\$1,396,799
Bam-Bam	\$123,319	\$66,946	\$16,346	\$1,632,46	\$1,632,46
Bam-Bam	\$145,769	\$302,129	\$123,107	\$56,783	\$1,227,715


“With Confluent, we help build AI copilots that allow users to interact with data in natural language, **turning Flink jobs into agents** that continuously monitor data streams. Flink AI Model Interference simplifies our stack...providing real-time context to generate the most accurate Flink SQL queries.”

Steffen Hoellinger
Co-Founder & CEO
Airy



“Everyone wants AI—but the hard part is getting high-quality data moving in real time. The Confluent data streaming platform makes that possible for us. **It’s the foundation that gets our data moving** and gets it where it needs to be.”

Chris Kapp
Software Architect
Henry Schein One



“Confluent data streaming platform is the backbone of our multi-agent platform... Confluent **enables real-time agent orchestration, observation, and governance**... allowing us to iterate faster and ship new features weekly.”

Saul Sparber
Founder & CEO
Agent Taskflow

[See more stories](#) →

5 Take the Next Step

- Read the [2025 Data Streaming Report](#) to learn how data streaming platforms are imperative for driving business value with data.
- Read about [agentic mesh and enterprise agent ecosystems](#).
- Learn why [the future of agentic AI is event-driven](#).
- Visit the [AI hub](#) for more resources.

About Confluent

Confluent is pioneering a fundamentally new category of data infrastructure focused on data in motion. Confluent's cloud-native offering is the foundational platform for data in motion—designed to be the intelligent connective tissue enabling real-time data from multiple sources to constantly stream across an organization. With Confluent, organizations can meet the new business imperative of delivering rich digital frontend customer experiences and transitioning to sophisticated, real-time, software-driven backend operations. To learn more, please visit www.confluent.io.

