



Diffusion

Intelligent Data Mesh

Turbocharge Real-Time Applications

“We saw a 15% increase in revenue the first quarter after going live with Diffusion.”
Betsson

Introduction

Today’s technology and business needs are focused on information management for distributed data sources. Big corporations and small businesses alike experience the challenges of dealing with constantly expanding and evolving data sources and the complex systems around them, while driving revenue from existing data. As event-based applications and real-time systems become fundamental to new business opportunities, there is a clear need and challenge yet to be addressed.

Traditional streaming data pipeline tools provide ways to help unify and normalize distributed data sources, but these tools are fundamentally built around polling-based resources (e.g. REST, SOAP) – an approach that is incompatible with the requirements for delivering truly live data. In the diverse ecosystems of today’s digital business environment, architectures can include any combination of polling-based, event-based, and bespoke infrastructures – often with challenging integration requirements. Businesses need a platform that delivers the operational benefits of data pipelines, and is also built to handle the unique challenges and interactions of real time – a platform to manage live data that is:

| Fast | Reliable | Secure | Scalable |

Diffusion Intelligent Data Mesh provides the ability to manage, optimize, secure, and distribute live data, no matter the source. Use cases exist across all industries, wherever there is business value in immediate data distribution: sports and trading feeds, geo-location data, IoT sensor streams, and many others.

At the forefront of this new solution space, Diffusion combines the benefits of traditional data pipeline and pub-sub solutions with the power of Real-Time, tying together both polling & event-based back-ends to provide a single unified platform for managing and distributing live data. Diffusion removes the traditional constraints of data management, allowing businesses to create new revenue channels from pre-existing infrastructure while greatly simplifying development of new and innovative applications. With a focus on efficiency, scalability, and cost-effectiveness, Diffusion brings the power of real-time to you.



“We needed to **build a single piece of code** for a variety of web transports – Diffusion does that, plus **we got >5 times latency reduction!**”

Derivco

“**Diffusion is a secure, robust, fast, and flexible data platform** which supports our Electronic Trading Platform & the features within our GUI.”

Tradition

Extending data

The fundamental responsibility of any streaming data pipeline solution is to integrate with existing data sources and extend them to new consumers over edge networks such as web, mobile, and satellite. The Diffusion Intelligent Data Mesh provides out-of-the-box adapters to back-end products including Kafka, JMS, and Apache Camel, empowering businesses to easily capitalize on existing event stream data to drive new applications. Diffusion also provides adapters for non-streaming resources such as REST and SQL, for easy transformation of traditional polling-based APIs into real-time streams.

With Diffusion, businesses can complement and enhance in-place APIs, building more sophisticated real-time features on top of existing infrastructure, without having to rip and replace everything. In addition, for use cases where data is accessed via proprietary or bespoke systems, Diffusion includes SDKs and application adapters (e.g. Excel) for the publication and consumption of data with custom logic and sophisticated control mechanisms.

The extension of real-time data is not solely based on the concept of making internal streams accessible over the Internet. For many mobile or IoT applications, there is a need for real-time data to be sent from remote devices back to a central event processing system. The Diffusion Intelligent Data Mesh provides fully bi-directional communication -- the same integration points allow data ingestion at scale, enabling businesses to both distribute internal feeds while receiving remote data streams, for collection and processing within back-end systems.

Workflows

For most businesses, the data coming from internal sources is unlikely to be suitably formatted for external application consumption. This problem is exacerbated when considering the potential and likelihood for multiple applications to be built on top of the same common data feeds. Normalizing and structuring data on a per-application basis can be time-consuming, error prone, and complex, especially when it is incumbent upon the development teams of each front-end application to communicate their specific requirements to the owners of the back-end systems. With Diffusion, applying structure and ordering on top of streams is simple. Diffusion features such as categorization, aliases, and dynamic views, reduce the complexity of how applications access and consume data, resulting in less development work and faster time-to-market.

“We made the right choice - Diffusion integration only took 2 weeks and the performance is great!”

Brenock

Ease of Development

An important benefit for Diffusion is the way in which it simplifies development of applications. Modern applications often require multiple types of data interactions -- event streams, direct messaging, and time series data are all leveraged to deliver rich application experiences and functionality. Diffusion is a unified platform to support multiple mechanisms for distributing data among systems, applications, and devices; with security and management layers applied above all data flows. Development teams who use Diffusion can focus on core business features, instead of low-level transport concerns, which allows their organizations to quickly and easily capitalize on existing or new data sources with significantly lower development costs and go-to-market time.

Security

Access control is a core concern for any business moving data. Exposing real-time data feeds can introduce additional operational risk, if not properly managed. At the same time, the ability to expose real-time feeds enables substantial opportunity for revenue growth. As a unified platform, Diffusion acts as a single access point to deliver centralized security control over all real-time data. Diffusion has a pluggable authentication system that allows businesses to enforce identity control across all data using whatever mechanisms are required -- SSO, LDAP, or custom authentication -- with fine-grained permissions and dynamic authorization for granular access control to easily grant or revoke privileges, as required.

Quality of Service

A key differentiator between conventional management of data streams and Diffusion is the Quality of Service, which affects the performance and viability of all aspects of the data flow. When dealing with polling-based solutions, the ability to explicitly handle connection quality is limited by the fundamental delivery mechanisms. By comparison, because Diffusion operates on streams of data, it provides significantly more sophisticated mechanisms for maintaining consistent and reliable data delivery. Automatic optimization of data assures extremely efficient performance, even for large numbers of remote devices -- ideal for IoT scenarios where the cost of bandwidth can be a significant issue. Regardless of scale or throughput, Diffusion guarantees that all real-time data will be delivered quickly, efficiently, and reliably, under even the most adverse conditions.

“Due to Diffusion’s reliability, we reduced our off-hours, on-site support personnel by 90%.”

OddsChecker

Scale

The value of a data streaming solutions such as Diffusion is not only the ability to manage real-time data, but also how efficiently and effectively the platform handles the scale of distribution. The wide array of corporate applications requires different types of scale including the abilities to: serve large and variable client volume, to handle tens of thousands of unique data streams, and to provide high throughput of data across geographically dispersed and/or remote regions. REST-based approaches often require large numbers of server instances to support heavy traffic loads, plus the associated operational complexity of coordinating data and monitoring systems.

Diffusion is purpose-built to deliver optimal performance and reduce operational risk across both axes of scale - traffic volume and data throughput - regardless of fluctuating conditions.

“With Diffusion, we can **scale our servers** up and down **to match demand** each day, and even by the hour, using the cloud infrastructure **in the most cost effective way.**”

Racing Post

Visit www.pushtechnology.com or contact sales@pushtechnology.com today!



Rethink Real-Time™

London
+44 (0) 20 3588 0900

Silicon Valley
1 (408) 780-0720

Belfast
+44 (0) 20 3588 0900

© 2020 Push Technology Ltd

All trademarks are the property of their respective owners