



# Imply Helps Atlassian Provide Product Usage Analytics for Their Customers

## Background

[Atlassian's](#) Confluence Analytics Experience team builds user-facing analytics features for Confluence, an industry-leading team workspace cloud application. The features surface user behavior insights to the end-user resulting in increased adoption and higher engagement on the Confluence product.

## Challenge

The development team building Confluence Analytics features needed a high-powered database that could meet the needs of the analytics features embedded within Confluence. The database had to deliver an interactive experience at high concurrency to support the large number of users of their cloud application, and serve as an underlying database that was flexible and extensible for new features and data and user growth without compromising performance.

To meet these critical requirements, the Confluence Analytics Experience Team chose to deploy Imply Enterprise Hybrid, a complete, real-time database built from Apache Druid® that runs in Atlassian's VPC with Imply's management control plane. This replaced usage of the Postgres database for event-based data, which was plagued with limitations.

The previous solution required the team to write a lot of custom code for handling large aggregations because it was unable to scale beyond two years of data retention at varying granularities. This meant queries that were error-prone and maintenance-heavy.

In addition, pre-aggregations of data in PostgreSQL provided aged data to users resulting in sub-optimal user experience. The application's response time was several seconds, whereas the application was required to return results within milliseconds.

## Customer Quote

***"We started with PostgreSQL, but its scale limitation became really painful to deal with. Very happy we switched to Apache Druid with Imply."***

- Puneet Arora, Engineering Manager, Atlassian



Atlassian is a leading provider of team collaboration and productivity software, including Confluence, Jira and Trello.

## Challenge

Atlassian needed a scalable, multi-tenant external-facing analytics experience for a growing number of end-users.

## Solution

Imply provided the database to build the interactive, scalable, and multi-tenant analytics features within the Confluence cloud application.

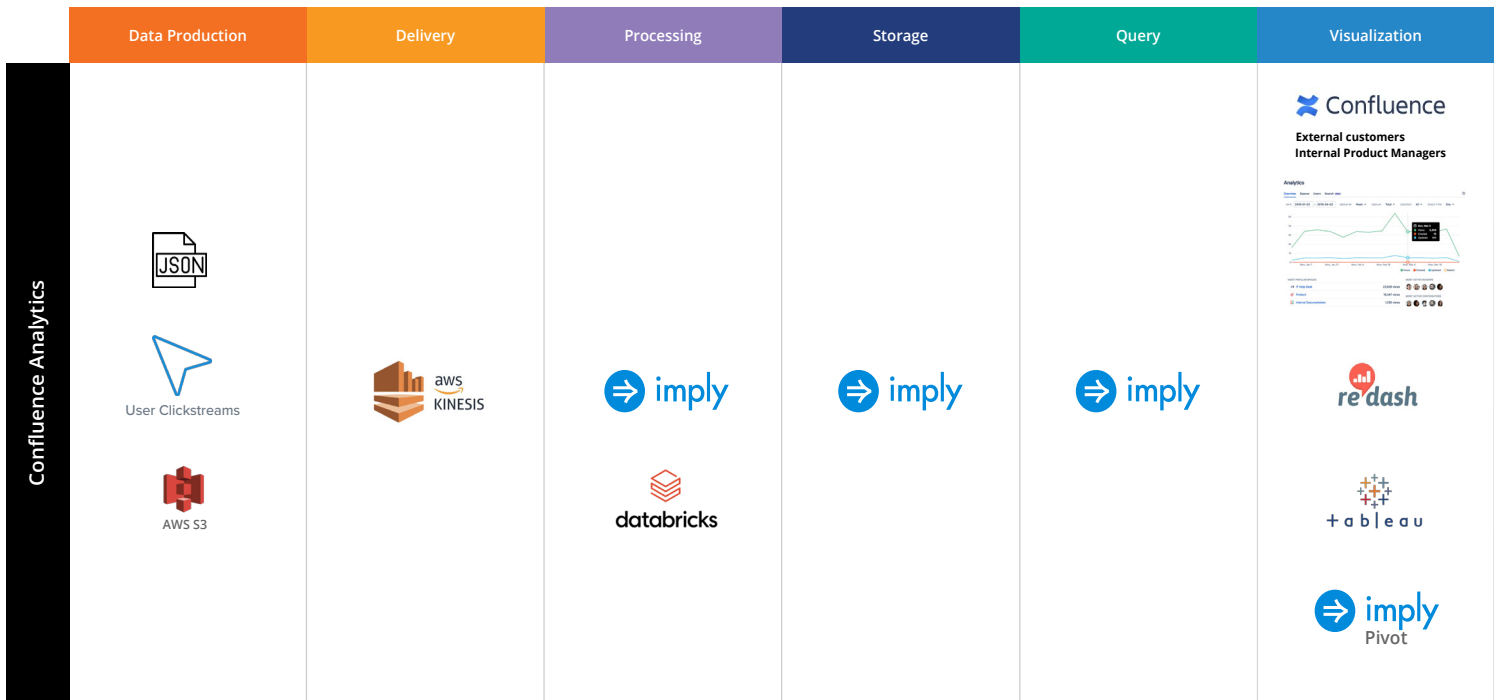
## Highlights

- 5X performance improvements
- Five years+ of data retention
- Ability to do funnel analysis on up to two years of data

## Solution

Atlassian's Confluence Analytics Experience Team needed a partner who could assist them in their rapid and continuous growth. They found that partner with Implied. Implied delivered the enterprise distribution of Apache Druid, which was not only blazingly fast and reliable, but was also effortless to manage.

## Customer Environment



## Results

With Implied, the Atlassian Confluence Analytics Experience Team has achieved more than they anticipated. They have seen massive performance improvements of up to 5X previous results with low latency of 100 msec or less. They're now able to increase data retention up to five years or more. They also have the ability to do funnel analysis on up to two years of data with no issues on real-time data ingestion and aggregation. Better still, they can now easily manage cluster configurations using Implied Manager.

### More Insights

Whether you are a Druid expert or data analytics driven, Implied Customer Stories will give you insights into how to build and leverage better data-driven applications.

Checkout more Implied Customer Stories at [implied.io/blog](https://implied.io/blog)



For more information visit us at [implied.io](https://implied.io)