



Apache Druid **Success Stories**

Compiled by Imply

NETFLIX



In order to succeed in today's increasingly digital world, modern organizations need real-time, granular insights into their business and operations. As a result, a new set of analytics uses cases are emerging - including real-time observability at scale, next-gen operational visibility, revenue-impacting insights and recommendations, and for extending analytics to external customers

To achieve this, leading companies including [Salesforce](#), [Netflix](#), and [Confluent](#) are turning to their developers to build analytics applications that deliver interactive data experiences from streaming data, and deliver real-time insights to both internal and external users.

At each of these, developers chose Apache Druid to power their analytics application.

Why?

Apache Druid is a high-performance, real-time analytics database that uniquely enables:



Interactive Analytics at Any Scale

Create a dynamic experience to rapidly visualize and explore data with sub-second query response times. Quickly slice and dice data and interact with live data to rapidly understand a situation or investigate a problem.



High Concurrency at Best Value

Power applications for everyone with a highly efficient architecture that supports 100s to 100K+ queries per second by any number of people. Give everyone access to data driven insights - from product managers to ops to data scientists as well as your customers.



Insights on Real-time and Historical Data

Power applications that deliver continuous intelligence with native Kafka and AWS Kinesis integrations built for scale, low latency, and data quality. Get the full value of real-time data with real-time insights.

Read on to discover more top use cases, with real-world examples and insights into how your organization can make the leap. All information is sourced and referenced from publicly available talks and blog posts from developers from these organizations.



NETFLIX

Leading streaming entertainment service

“Druid is our choice for anything where you need sub second latency, any user interactive dashboarding, any reporting where you expect somebody on the other end to actually be waiting for a response. If you want super fast, low latency, less than a second, that’s when we recommend Druid.”

- **Parth Brahmhatt**,
Senior Software Engineer, Netflix ¹



Summary

To ensure a consistently great experience to more than 100 million members in more than 190 countries enjoying 125 million hours of TV shows and movies each day, Netflix built an analytics application powered by Apache Druid. By turning log streams into real-time metrics, Netflix is able to see how over 300 million devices (across 4 major UIs) are performing at all times in the field. Netflix chose Druid because it uniquely meets their high ingestion rate of data, high cardinality, and fast query requirements. By ingesting over 2 million events per second and querying over 1.5 trillion rows, Netflix engineers are able to pinpoint anomalies within their infrastructure, endpoint activity, and content flow.

Challenge

An ongoing challenge for Netflix is to consistently deliver a great streaming entertainment experience while continuously pushing innovative technology updates.

“How can we be confident that updates are not harming our users? And that we’re actually making measurable improvements when we intend to?”

- **Ben Sykes**,
Senior Software Engineer, Netflix ²

As Netflix’s adoption has skyrocketed, this challenge has grown more complex. With over 300 million devices spanning four major UIs including IOS, Android, Smart TVs and their own website, Netflix has a constant need to identify and isolate issues that may only affect a certain group, such as a version of the app, certain types of devices, or particular countries.

“With this data arriving at over 2 million events per second, getting it into a database that can be queried quickly is formidable. We need sufficient dimensionality for the data to be useful in isolating issues and as such we generate over 115 billion rows per day.”

- **Ben Sykes**,
Senior Software Engineer, Netflix ³


¹ [Unlocking Self-Service & Real-Time Analytics Across Netflix](#)

^{2,3} [How Netflix uses Druid for Real-time Insights to Ensure a High-Quality Experience](#)

Solution

Netflix chose Apache Druid as their real-time database to power their analytics application because it's uniquely capable of high ingestion rate of event data, with high cardinality and fast query requirements.

To quantify how seamlessly users' devices are handling browsing and playback, Netflix derives measurements using real-time logs from playback devices as a source of events.



Once they have these measures, Netflix feeds them into Druid. Every measure is tagged with anonymized details about the kind of device being used, for example, whether the device is a Smart TV, an iPad or an Android Phone. This enables Netflix to classify devices and view the data according to various aspects. With Druid, this aggregated data is available immediately for querying, either via dashboards or ad-hoc queries.

“We’re currently ingesting at over 2 million events per second, and querying over 1.5 trillion rows to get detailed insights into how our users are experiencing the service. All this helps us maintain a high-quality Netflix experience, while enabling constant innovation.”

- Ben Sykes,

Senior Software Engineer, Netflix ⁴

Netflix leverages Druid to employ A/B testing to assess how updates and changes impact various user groups. It uses the results to compare how the new version performs against the older version to tell whether users on different systems should get the update or not.

“Druid can make some optimizations in how it stores, distributes, and queries data such that we’re able to scale the datasource to trillions of rows and still achieve query response times in the 10s of milliseconds.”

- Ben Sykes,

Senior Software Engineer, Netflix ⁵

The ultimate benefit is speed, which is essential for a service that needs to react to a massive number of users in near real time.

⁴ ⁵ [How Netflix uses Druid for Real-time Insights to Ensure a High-Quality Experience](#)



Image sharing and social media service

“As we had more and more advertisers each of them were requesting different slices and cuts of their data. Everybody wanted to see their metrics a little bit differently and [HBase] didn’t really support that. We said what’s next for us? And we chose Apache Druid because it met all our key requirements.”

- Filip Jaros,

Software Engineer, Pinterest ⁶



Summary

To enable advertisers to effectively reach over 400 million people who use Pinterest every month, Pinterest chose Apache Druid to build Archmage, their real-time analytics application. This replaced usage of Apache Hbase, which was plagued with limitations and could not keep up with their growth. Once Pinterest hit millions of active users, they found it could take longer than 24 hours to give user data to their advertisers. Now, with Druid, they are able to keep up with their rising active user base as well as add additional metrics and data driven insights to help drive more revenue for their advertisers.

Challenge

When Pinterest first launched paid advertising on their platform they chose Apache HBase as their database to provide transparency to advertisers on how users engage with their ads. Metrics were precomputed in an hourly or daily batch job, transformed into a key value data model, and stored in Hbase. This approach was sufficient in the beginning, but once Pinterest hit millions of active users, they found it could take longer than 24 hours to deliver user data to their advertisers.

The challenges of the Hbase-based precomputed key value look up system became very clear:

Inability to deliver interactivity at scale

Key Value databases are not built for fast grouping and aggregation. Pinterest’s work-around was to pre-build these data views, but as the features needed for their reporting expanded, it was no longer possible to store so many different cuts.

Cardinality explodes with more columns

With 30 columns it is easy to generate a key that includes combinations of all the dimensions. Pinterest found once you expand to over 100 columns it quickly becomes unbearable because of all the possible combinations.

High maintenance cost

HBase was a complicated system, and it had high barriers to entry. As a result, it took a lot of effort to keep HBase up and running

High infrastructure cost

Infrastructure cost was high because Pinterest used primary-standby cluster pairs to achieve high availability, and they needed to use six replicas for a unique dataset.

⁶ [Pinterest: Powering Ad Analytics with Apache Druid](#)

Solution

With the help of Apache Druid's realtime analytics database, Pinterest delivers instant and full transparency to advertisers on how users engage with their ads. Druid allowed the Pinterest teams to bypass all of Hbase's complicated data slicing ingestion logic while also supporting the following key features:

Real-time ingestion via Kafka	Automatic versioning of ingested data	Data pre-aggregation based on user-set granularity
Approximate algorithms for count-distinct questions	A SQL interface	Easy to understand code and a very supportive community

As a result of Apache Druid advanced capabilities, Pinterest onboarded many new critical use cases including reporting partner and advertiser business metrics, organic pin stats, experiment metrics, spam metrics analysis, and more.

"We use the streaming to basically do ad counting as well on how many ads will be shown for each campaign. Then we also have a real-time ads campaign manager where you can look at the performance of your campaign and also start or stop and change that campaign, and a lot of that comes through going with streaming, going through Flink, and then into Druid."

- Dave Burgess,
Head of Data Engineering, Pinterest ⁷

The Druid infrastructure consists of more than 2,000 nodes in a multi-cluster setup. Overall, Pinterest is running over 500,000 events per second into Druid, and it's happening within one minute of the event happening.

High concurrency is supported with numerous teams within Pinterest running queries - the largest running over 1,000 queries per second with 99% of these queries are answered within 250ms.

"We had so much data within our HBase cluster that it was getting really slow and very difficult to manage overhead for our engineers. We decided to migrate that to Druid. The maintenance cost has gone way down, the latency has gone way down, the actual cost of running the infrastructure way down. Really happy about it."

- Dave Burgess,
Head of Data Engineering, Pinterest ⁸

⁷ ⁸ [In Conversation with Dave Burgess, Head of Data Engineering, Pinterest](#)



Multinational retail corporation

“Druid has been designed to solve the problems [of relational and NoSQL Key Value Data databases] i.e. enable exploration of real-time data and historical data while providing low latencies and high availability.”

- **Kartik Khare,**

Software Engineer at Walmart Labs ⁹



Summary

To compete with Amazon and other retailers, Walmart chose Apache Druid as part of their technology stack to track the pricing of their competitors in real-time. Druid replaced Hive and then Presto because both of these solutions impacted Walmart’s ability to make rapid decisions. Even though their data was arriving in real-time, queries could take hours to complete. After the switch to Druid, Walmart now accesses, analyzes, and makes decisions on data in near real-time. Query latencies dropped to near sub-second levels while easily scaling to nearly 1 billion events per day.

Challenge

Many data sets at Walmart are generated by the digital business, and are modeled as streams of events ranging from server logs, application metrics, to product purchases. The WalmartLab team’s goal is to make it easy for the right people across Walmart to access this data, analyze it, and make decisions in the least amount of time possible.

Walmart’s first attempt to provide low latency analytics was to leverage the Hadoop ecosystem, namely Hive, and then Presto.

“The problem we faced with both of these SQL-on-Hadoop solutions was that queries would sometimes take hours to complete, which significantly impacted our ability to make rapid decisions. Although our data was arriving in real-time, our queries quickly became a bottleneck in our decision-making cycle as our data volumes grew. We quickly realized that the workflow we were aiming to optimize was one where we could look at our event streams (both real-time and historical events) and slice and dice the data to look at specific subsections, determine trends, find root causes, and take actions accordingly.”

- **Amaresh Nayak,**

Distinguished SW Engineer at Walmart Labs ¹⁰

The team at Walmart knew they needed to make a change. The types of queries they run at Walmart require the need to get aggregates on some columns which involves scanning a lot of rows across multiple shards. Walmart knew a relational database (RDBMS) was ill equipped for this because it can’t efficiently enable data exploration in real-time. The team also ruled out a NoSQL Key Value database because of the need to query multiple partitions across a number of nodes, causing inefficient aggregation calculations. Also, a Key Value database would exponentially increase storage requirements due to a requirement of storing aggregates for all possible column combinations

⁹ [What Makes Apache Druid Great for Realtime Analytics?](#)

¹⁰ [Event Stream Analytics at Walmart with Druid](#)

Solution

After some searching, Walmart labs replaced Hive and Presto with Apache Druid to power their realtime event streaming analytics.

“Druid is an OLAP engine that is highly optimized for low latency data ingestion (streaming ingestion), as well as extremely fast aggregations. It integrates natively with Kafka and Storm, so it was relatively easy for us to get started.”

- Amaresh Nayak,
Distinguished SW Engineer at Walmart Labs ¹¹

The team at Walmart can now get instant insights on pricing because of Druid's out of the box support and integration with Kafka and Storm. With no set-up and tuning, Druid can pre-aggregate records as they are being ingested. Instead of getting a single price for an item for a specific moment in time, with Druid, Walmart can now understand in near real-time the aggregated price of that item over any span of time. This combined level of depth and speed to insights is essential to enabling Walmart to make critical pricing decisions in the least amount of time possible.

“After we switched to Druid, our query latencies also dropped to near sub-second and in general, the project fulfilled most of our requirements. Today, our cluster ingests nearly 1B+ events per day (2TB of raw data), and Druid has scaled quite well for us.”

- Amaresh Nayak,
Distinguished SW Engineer at Walmart Labs ¹²





Global leader in customer relationship management software

“We chose Druid because it gives us the flexibility to define pre-aggregations, the ability to easily manage ingestion tasks, the ability to query data effectively, and the means to create a highly scalable architecture.”

- **Dun Lu,**

Lead Software Engineer at Salesforce ¹³

Summary

To ensure a consistently great experience to more than 150,000 customers around the globe, Salesforce built an observability application powered by Apache Druid. By turning log lines into real-time application performance metrics, Salesforce is able to obtain data driven insights such as performance analysis, trend analysis, release comparison, issue triage, and troubleshooting. Salesforce chose Druid because it uniquely enables finer, more granular insights by querying increasingly distributed data sources, delivering ad hoc interactive analytics, and processes non-aggregated data at massive scale.



Challenge

With over 150,000 customers across the world, Salesforce faces the challenge of delivering a 24x7 great user experience. The Salesforce engineering team must keep their eye on the health of more than 2,700 production services - to understand performance, behavior, capacity, and more.

At Salesforce, data is massive and widely distributed. There are dozens of petabytes of data in a transactional store, over 5PB of logs data accumulates every month in data centers, and nearly 200PB accumulated in Hadoop storages. They also need to ingest 5 billion events into their system every day.

“We need to support batch, stream, and interactive ad hoc query processing on these data sets.”

- **Ram Sangireddy,**

Sr Director @ Salesforce, Platforms and Apps, Big Data, Analytics, ML/AI ¹⁴

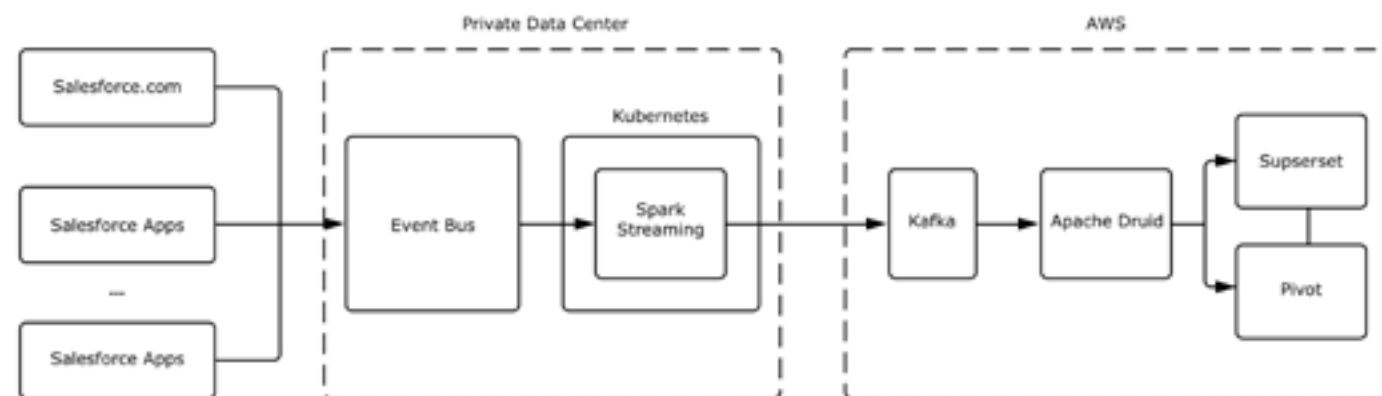
¹³ [Delivering High-Quality Insights Interactively Using Apache Druid at Salesforce](#)

¹⁴ [Big Data, Big Decisions: Finding the Right Technology for Interactive Analytics at Salesforce](#)

Solution

Salesforce chose Apache Druid as their real-time analytical database to store application performance metrics extracted from log lines. Druid uniquely meet's Salesforce core requirements:

Ingest billions to trillions of log lines per day.	Allow users to define a combination of dimensions and filters on dimensions and different types of aggregation on metrics.	Return query results of real-time data interactively within seconds.
--	--	--



“We chose Druid because it gives us the flexibility to define pre-aggregations, the ability to easily manage ingestion tasks, the ability to query data effectively, and the means to create a highly scalable architecture.”

- **Dun Lu,**
Lead Software Engineer at Salesforce ¹⁵

Many business units within Salesforce including engineers, product owners, customer service representatives, and more use this service to obtain insights including performance analysis, trend analysis, release comparison, issue triage, and troubleshooting. These teams have access to finer, more granular insights across distributed data sources.

Druid uniquely enables complex group-by aggregations (e.g. quantile group-by aggregation using data sketches), interactive queries (including SQL support) and rich analytical dashboards (including drill-down pivot tables, top-K views, histograms etc.)

Since Salesforce’s data volume is large with moderate cardinality, they are able to leverage Apache Druid’s compaction and rollup capabilities to further reduce storage footprint and drive better performance.

“For one day of data, we end up saving 82% in the total number of rows stored in Druid, which translated to 47% total savings in the storage footprint and about 30% improvement in query time performance.”

- **Dun Lu,**
Lead Software Engineer at Salesforce ¹⁶



Data streaming platform

“Leveraging Druid as part of our stack means we don’t shy away from high-cardinality data which means we can find the needle in the haystack. As a result, our teams can detect problems before they emerge and quickly troubleshoot issues to improve the overall customer experience. The flexibility we have with Druid also means we can expose the same data we use internally also to our customers, giving them detailed insights into how their own applications are behaving.”

- **Xavier Leaute and Zohreh Karimi,**
Lead Engineers at Confluent ¹⁷



Summary

To meet their strict performance and availability SLAs for Confluent Cloud, a fully managed data streaming data, Confluent built an internal-facing observability application powered by Apache Druid. Confluent turned to Druid because their existing NoSQL database could not keep up with their data growth. With Druid, they now handle 100x of the data volume and query load, deliver sub-second query latencies on high cardinality metrics, and get native Kafka ingestion. By ingesting over 3.5 million events per second and handling hundreds of queries on top of that, Confluent has real-time insights into the operations of thousands of clusters within Confluent Cloud. Confluent also leveraged Druid to build an externally-facing application, Confluent Health+, which extends performance and health insights to their customers.

Challenge

With Confluent Cloud, Confluent manages a large number of multi-tenant Apache Kafka® clusters across multiple cloud providers, including Azure, GCP, and AWS. To deliver the best user experience running on their infrastructure, it’s critical Confluent has real-time monitoring and operational visibility for their customer’s mission-critical workloads.

Confluent’s first attempt at building an observability pipeline was with a NoSQL database to store and query telemetry data. As the volume of data grew, Confluent’s legacy pipeline struggled to keep up with their data ingestion and query loads. Next, they looked into off the shelf observability solutions but quickly determined these could not handle their requirements.

“Operating multi-tenant services requires fine-grained visibility down to the individual user, tenant, or application behavior, where most traditional monitoring stacks fail to scale or become cost-prohibitive.”

- **Xavier Léauté and Zohreh Karimi,**
Lead Engineers at Confluent

To keep up with their data growth, Confluent determined their next-generation observability pipeline needed to support:

- Substantially increased data and query load (100x)
- Sub-second query latencies on high-cardinality metrics
- Inherent time-series data support

¹⁷ [How leading organizations leverage Confluent/Kafka together with Imply/Druid](#)

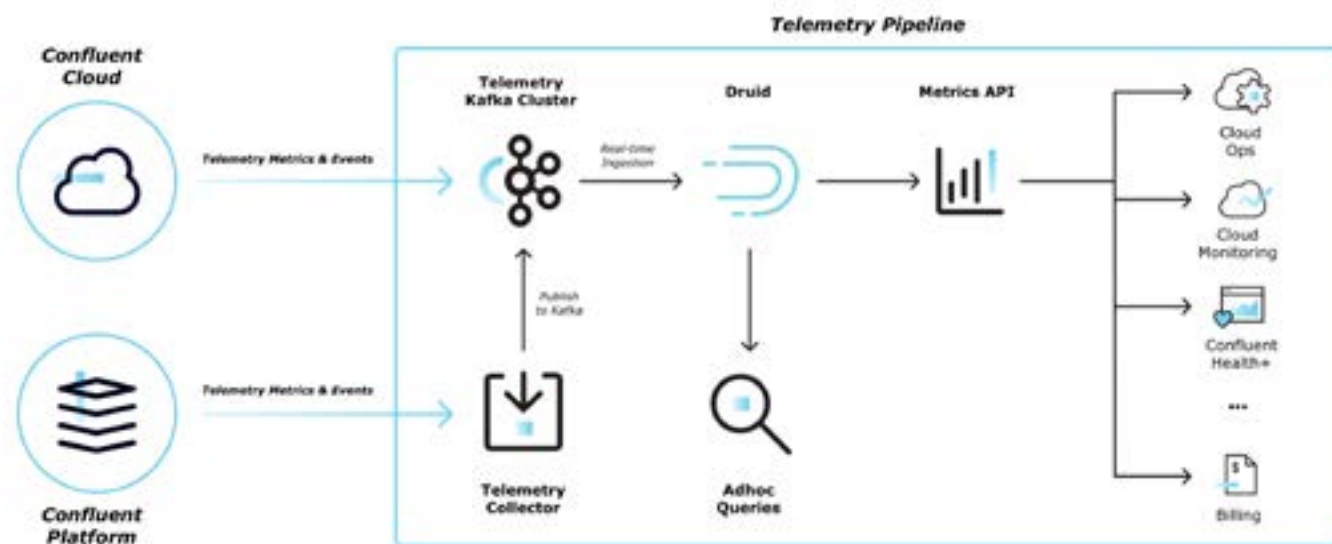
Solution

Confluent evaluated both Apache Druid and Clickhouse. They quickly ruled out Clickhouse because they found it required writing C++ plugins if they wanted to read custom format data from Kafka.

Ultimately, they chose Druid as the database to power their next-gen observability platform.

“Druid plays an important role in our monitoring and analytics platform and powers many internal and external use cases at Confluent. Among customer-facing solutions, we use Druid for Confluent Cloud monitoring dashboards, Confluent Health+ (an alerting and notification solution), Confluent Stream Lineage, and Confluent Cloud Metrics API. Inside Confluent, we use Druid to provide usage data for cloud billing and to perform ad hoc diagnostics queries.”

- **Zoe Karimi and Harini Rajendran,**
Software Engineers at Confluent ¹⁸



¹⁸ [Scaling Apache Druid for Real-Time Cloud Analytics at Confluent](#)

All Confluent Cloud clusters, as well as customer-managed, Health+-enabled clusters, publish metrics data to Confluent’s telemetry pipeline. Under the hood, the telemetry pipeline uses a Confluent Cloud Kafka cluster to transport data to Druid. From there, Druid’s real-time ingestion consumes high volumes of data from the Kafka cluster and ultimately provides low-level queries.

“Druid also has this native Kafka integration out of the box and that is one of the big advantages... we don’t even need any sort of connector or anything to make Apache Kafka and Apache Druid work together. It just works.”

- **Harini Rajendran,**
Software Engineers at Confluent ¹⁹

Confluent engineers often execute ad hoc queries directly against Druid for debugging purposes in their cloud infrastructure. This has helped Confluent tune configurations such as memory size, disk size, load-balancing thresholds, throttling thresholds, and more. With Druid, Confluent is able to have self-balancing enabled on Confluent Cloud to continuously monitor resource utilization of their cluster and load load-balance when necessary.

“We ingest over three million events per second and respond to over 250 queries per second. We keep seven days of queryable data in Druid Historical nodes and two years of data retention in S3 deep storage.”

- **Zoe Karimi and Harini Rajendran,**
Software Engineers at Confluent ²⁰

¹⁹ [Project Shapeshift Milestone 1](#)

²⁰ [Scaling Apache Druid for Real-Time Cloud Analytics at Confluent](#)

Why Developers Turn to Apache Druid

Apache Druid is the right choice when powering an analytics application at any scale, for any number of users, and across streaming and batch data. Unique capabilities include:

Sub-second at scale.

From TBs to PBs or 100s to 1000s of concurrent queries, Druid's unique distributed architecture can deliver consistent sub-second query response times - and do it without breaking the bank. It combines the performance of shared-nothing with the elasticity of shared-everything.

True stream ingestion.

Druid was built for streaming data with native integration with Kafka and AWS Kinesis. It supports massive scale ingestion to millions of events per second, query-on-arrival to analyze events instantly, and guaranteed consistency to ensure data quality.

Non-stop reliability.

For always-on applications, Druid is designed to never go down and never lose data. Its design center is built for high availability and no data loss for streams via continuous backup and automatic recovery and rebalancing.

This ebook of Apache Druid success stories was compiled by Imply.

About Imply

At Imply, we are on a mission to help developers become the new heroes of analytics. Our unique database, built from Apache Druid, enables them to develop the next generation of analytics applications. With Imply, developers can build without constraints as our database lets them create interactive data experiences on streaming and batch data with limitless scale and at the best economics.

To learn more, please visit: <https://imply.io/>.

© 2022 Imply. All rights reserved. Imply, the Imply logo, and Imply Polaris are either registered trademarks or trademarks of Imply Data, Inc. in the U.S. and/or other countries. Apache, Apache Druid, Druid and the Druid logo are either registered trademarks or trademarks of the Apache Software Foundation in the USA and/or other countries. All other marks are the property of their respective owners.