

The Apache Software Foundation Announces the 5th Anniversary of Apache™ Cassandra™

Wednesday 9 April, 2014

Highest performing "NoSQL" distributed Big Data database in use at hundreds of organizations including Adobe, CERN, Comcast, Disney, eBay, GE, GitHub, GoDaddy, HP, Hulu, IBM, Instagram, Intuit, Netflix, Plaxo, Polyvore, Sony, and The Weather Channel

Forest Hill, MD –09 April 2014– The Apache Software Foundation (ASF), the all-volunteer developers, stewards, and incubators of more than 170 Open Source projects and initiatives, announced today the 5th Anniversary of Apache™ Cassandra™, the highly-performant Big Data distributed database.

"I am so proud to see what the Apache Cassandra community has been able to achieve in five short years," said Jonathan Ellis, Vice President of Apache Cassandra and DataStax CTO. "We've come such a long way since the early days, and it is a testament to Cassandra's rapid maturation that it has been deployed in over 1,500 global critical production environments."

Apache Cassandra is an Open Source, "NoSQL" distributed database management system designed to handle large amounts of data across many commodity servers quickly and reliably without compromising performance, whether running in the Cloud or in a hybrid data store. Cassandra offers robust support for clusters spanning multiple datacenters, and provides high availability with no single point of failure.

Originally developed at Facebook in 2008 to power their Inbox Search feature, Cassandra entered the Apache Incubator in 2009 and graduated as an Apache Top-level Project in February 2010.

Apache Cassandra has consistently led the NoSQL market in performance: its fully-distributed architecture provides unparalleled fault tolerance to ensure applications will not go offline, and its linear scalability allows them to reach massive sizes while successfully handling thousands of requests per second.

Cassandra Evolution 5 Years On: Under the Hood

Over the past five years, Apache Cassandra has had 6,000 JIRA issues and 250 contributors after its initial release, making today's Cassandra significantly more performant, resilient, feature-complete, and easier to both operate and develop against.

Apache Cassandra's improvements include security features, performance, and ease of use with the implementation of Cassandra Query Language (CQL) that presents a data model familiar to relational database users. Unlike many other systems, Cassandra is ideal for read-heavy workloads, and also offers scalable write performance. During the past five years, Apache Cassandra has resolved over 6,000 JIRA issues and added more than 250 contributors, making today's Cassandra significantly more performant, resilient, feature-complete, and easier to both operate and develop against. Some of the milestones along the way include:

- The Cassandra Query Language, which offers a more intuitive data model and a performant native protocol while retaining backwards compatibility with data created under the old Apache Thrift API;
- Lightweight transactions, an industry first that allows users and applications to opt into a linearly consistent world view as necessary;
- An innovative virtual node design that allows expanding a cluster in increments as small as a single machine, and across heterogeneous hardware;
- A powerful log-structured storage engine featuring advanced compaction, compression, and SSD support;
- Thousands of enhancements from running the world's most demanding applications at scale, informing better performance, better drivers, and better management tools.

Improved Performance in Real-world Situations

Apache Cassandra powers hundreds of applications across dozens of industries that demand high performance at scale. By addressing the needs of different workloads, Cassandra has evolved beyond its initial niche in social media into a truly general purpose solution.

Apache Cassandra is used by many highly-visible organizations including: Adobe, Comcast, Disney, eBay, Eventbrite, GE, GoDaddy, HP, IBM, Instagram, Intuit, Netflix, Pearson, Safeway, Sky, Sony, Spotify, Travelocity, The Weather Channel, and Zoosk, among others. Additional organizations using

Related Sectors:

Computing & Telecoms ::

Related Keywords:

The Apache Software Foundation ::

Scan Me:



Apache Cassandra can be found at <http://planetcassandra.org/companies/>

"We knew Apache Cassandra could perform linear scaling of reads and writes with consistent performance."

--David Weinstein, Director of Software Development, Adobe

"Before adopting Cassandra, we could not monitor every malicious site and IP forever – the data volumes were just too great. No other database was ready for what we needed to do."

--Michael Kjellman, Software Engineer, Barracuda Networks

"Apache Cassandra provides us with an easy to use backend and lets us focus on our implementation and features."

--Andreas Wagner, Lead Developer, CumulusRDF at the Karlsruhe Institute of Technology (KIT)

"Apache Cassandra is critical for being able to look up historical behavior data quickly, so that we can do these model updates with low latency."

--Thomas Pinckney, Senior Director of Engineering, eBay

"Apache Cassandra provides us an easy way of providing a highly available store and not have to worry about things like sharding, multi-datacenter support and things like that. Cassandra is an amazing store. Some of the features that you get out of the box are pretty incredible. I think one of the reasons why Cassandra has taken off and is doing so well is because of its awesome community."

--Vipul Sharma, Director of Data Engineering, Eventbrite

"Apache Cassandra embodies in its core the resilience and availability we need to continue serving our enterprise and internal customers even in the face of transient outages. Most of the time we forget about Cassandra and it keeps on running."

--Michael Rose, Senior Platform Engineer, FullContact

"The Google Cloud Platform was able to sustain one million Cassandra writes per second at a cost of \$0.05 USD per million writes."

-- Ivan Santa Maria Filho, Performance Engineering Lead, Google Cloud Performance Team

"Apache Cassandra offers good performance, near linear scalability for our data model, and geo-replication all with minimal maintenance requirements."

--Andres Rangel, Senior Software Engineer, Hulu

"We're helping the world save water, which is a very laudable target for an innovative and commercial business. We're using quite a lot of innovative technology to help us do that, of which Cassandra plays a major role."

--Mike Williams, Software Director, i2O Water

"At Instagram we've either replaced or are replacing every use case of Redis with Apache Cassandra. Its operational robustness and ability to exploit solid-state disks are the primary drivers for these efforts. We've also used Cassandra as part of major new feature efforts for its ability to scale up, scale down, and transparently work in a multi-region environment."

--Rick Branson, Infrastructure Software Engineer, Instagram

"The data was stored originally in Oracle in such a way that it does not scale at all ... We chose Cassandra because of its ability to scale easily, the operational simplicity, but also because it met our requirements."

--DuyHai Doan, Senior Developer, Libon, a Division of Orange

"We absolutely love the data model and scalability that Cassandra offers us. Coupled with ease of use and the ability to get it up and running quickly in development means that we can experiment quickly and have code flow from development into production in a short amount of time."

--Hisham Mardam Bey, CTO, Mate1

"During Hurricane Sandy, we lost an entire data center. Completely. Lost. It. Our application fail-over resulted in us losing just a few moments of serving requests for a particular region of the country, but our data in Cassandra never went offline."

--Nathan Milford, US Operations Manager, Outbrain

"Portugal Telecom is now expanding to Brazil, which is a market that is roughly 25 times larger than ours, so we believe with the help of Cassandra we'll store huge loads of information."

--Ivo Jesus, Tech Lead, Portugal Telecom

"We are the number one source for weather on all platforms: TV, web, mobile, and API. Our digital group provides backend services across all our platforms —billions of requests per day— and is responsible for delivering scheduled and severe weather alerts. We make extensive use of Cassandra under the hood for our high-throughput use cases, and we support nearly every imaginable type of content. We get about 100M transactions per day on average against our busiest Cassandra-backed service, with a heavy day seeing more like 180-200M transactions. Apache Cassandra has never failed us."

--Robbie Strickland, Software Development Manager, The Weather Channel

Availability and Oversight

As with all Apache products, Apache Cassandra software is released under the Apache License v2.0, and is overseen by a self-selected team of active contributors to the project. A Project Management Committee (PMC) guides the Project's day-to-day operations, including community development and product releases. For documentation and ways to become involved with Apache Cassandra, visit <http://cassandra.apache.org/>

About The Apache Software Foundation (ASF)

Established in 1999, the all-volunteer Foundation oversees more than one hundred and seventy leading Open Source projects, including Apache HTTP Server --the world's most popular Web server software. Through the ASF's meritocratic process known as "The Apache Way," more than 400 individual Members and 3,500 Committers successfully collaborate to develop freely available enterprise-grade software, benefiting millions of users worldwide: thousands of software solutions are distributed under the Apache License; and the community actively participates in ASF mailing lists, mentoring initiatives, and ApacheCon, the Foundation's official user conference, trainings, and expo. The ASF is a US 501(c)(3) charitable organization, funded by individual donations and corporate sponsors including Budget Direct, Citrix, Cloudera, Comcast, Facebook, Google, Hortonworks, HP, Huawei, IBM, InMotion Hosting, Matt Mullenweg, Microsoft, Pivotal, Produban, WANdisco, and Yahoo.

For more information, visit <http://www.apache.org/> or follow @TheASF on Twitter.

"Apache", "Apache Cassandra", "Cassandra", and "ApacheCon" are trademarks of The Apache Software Foundation. All other brands and trademarks are the property of their respective owners.

###

Company Contact:

[The Apache Software Foundation](#)

T. 16179218656

E. press@apache.org

W. <http://www.apache.org>

[View Online](#)

Newsroom: Visit our Newsroom for all the latest stories:

<https://www.apache.pressat.co.uk>