

Apache Cassandra Changelog #8

Author(s):	Community
Reviewer(s):	Community
Outlet:	C* blog
Date:	Week of 06/28
Excerpt:	4.0-rc2 released, say hello to our Google Summer of Code intern and new community intro to Cassandra videos



Our monthly roundup of key activities and knowledge to keep the community informed.

Release Notes

Released

The Apache Cassandra team is pleased to announce the release of Apache Cassandra [4.0-rc2](#) ([pgp](#), [sha256](#) and [sha512](#)). Please read the [release notes](#) and let us [know](#) if you encounter any problems.

This version is a release candidate on the 4.0 series.

Note: The bintray location is now replaced with the [ASF's JFrog Artifactory location](#).

Downloads of source and binary distributions for the latest stable and older supported version are listed in our [download section](#):

For the latest status of Apache Cassandra 4.0 GA, check the Cassandra 4.0GAScope board [here](#).

We recommend joining the Cassandra [mailing list](#) to stay up-to-date.

Community Notes

Updates on Cassandra Enhancement Proposals (CEPs), how to contribute, and other community activities.

Added

Netflix's internal teams create easy-to-use tooling that helps streamline user experience and incorporates best practices. It has open-sourced one of the tools, [Data Explorer](#), which provides fast, safe access to data stored in popular databases, including Apache Cassandra. Gim Mahasintunan has written an explanation [here](#). The tool includes a drag and drop Schema Designer for writing CREATE TABLE statements, support for binary data, and multiple cluster access, along with a powerful CQL IDE in Query mode.

Added

[CEP-10](#) is under discussion. This is a proposal for a mechanism to evaluate whole clusters, or individual classes, with a deterministically pseudorandom ordering of all thread and message events. The goal will be to simulate a cluster and actions on it (or simpler unit tests) so that the behavior is deterministic, repeatable but pseudo-random.

Added

Yifan Cai has noted that JMXTimer exposes its attributes in inconsistent time units: nanos and micros. By correcting this inconsistency and using micros, we can reduce the timer memory footprint by 50%. You can read about the details [here](#).

Added

We are pleased to welcome Abuli Palagashvili, who was selected to participate in Google Summer of Code (GSoC) on Apache Cassandra this year. He will be working primarily on [CASSANDRA-16451](#).

Added

As the project looks to the future, we are adding a roadmap page to the new [Apache Cassandra project website](#). The community is also looking towards new features, and you can take a look at the features that developers plan to work on [here](#).

Added

We've released new Apache Cassandra community videos: [What is Apache Cassandra?](#) and [What's New in Apache Cassandra 4.0?](#). A big thank you goes to Ben Bromhead, Ekaterina Dimitrova, and Patrick McFadin for presenting and helping put these together! 🙌 Remember, if you weren't able to attend Cassandra World Party, we've also uploaded all the recorded lightning talks to this [playlist](#).

```
<iframe width="560" height="315" src="https://www.youtube.com/watch?v=kOf29M8qFts"
frameborder="0" allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope;
picture-in-picture" allowfullscreen></iframe>
```

Added

The DROP COMPACT STORAGE feature will be marked as [‘experimental’](#), but this will be until the following points are solved:

- We have a similar set of unit tests to the one we have for COMPACT storage for after DROP COMPACT and another set where we mix writes performed before and after DROP COMPACT.
- We solve the issues linked to the lack of primary key liveness ([CASSANDRA-16675](#)).
- We have a way to prevent users from having issues with their tables by using ALTER DROP statements.
- We find a solution for [CASSANDRA-1606](#).

The project has generated an umbrella ticket to have DROP COMPACT STORAGE out of experimental as soon as possible and allow users to be able to use it without taking any risk.

User Space

Paidy

Paidy is a financial services company that provides a cardless payment method using only your email address and a phone number and is currently one of the largest online payment businesses in Japan. To realize the company's mission of creating a world to removes barriers and embraces simplicity for conducting online payments, the company needed a robust and scalable infrastructure for the technology platform:

"Our platform is designed with an event sourcing and reactive architecture, and we need to be capable of storing the data in multiple data centers with redundancy and high availability being critical to our services. Early on, we identified Apache Cassandra as being the right database technology for Paidy's payment services. We also understand the complexity of managing this type of databases to maintain the highest level of availability and the dedication and expertise required to tune our environment for continued low latency and high performance." - [Saravanan Karuppia](#)

Hornet

Hornet is a social networking community and app for the gay community with 30 million members worldwide. As the company is focused on community development, the Hornet team knew they would have to scale up the existing messaging clusters. Hornet had implemented Apache Cassandra previously, but its current version required an update and Hornet has found the benefits have been significant. For example, it has been able to reduce the messaging cluster volumes by 50% and improved the content service archive from two weeks to one month of messages:

"This is probably going to be the most engineering non-answer ever, which is simply that I haven't really had to care about Cassandra since we made the changes and upgrades. Usually, if I was getting paged in the middle of the night, it probably had something to do with a brief Cassandra blip that was causing an increased response time. That has just gone away

completely. We are also able to accept far more users at this point, before we had to look into doing any sort of infrastructure upgrades - infrastructure meaning disk size, CPU, memory etc."
- [Nate Mitchell](#)

Do you have a Cassandra case study to share? Email cassandra@constantia.io.

In the News

JAXEnter: [Apache Cassandra 4.0: Taming Tail Latencies with Java 16 ZGC](#)

TechRepublic: [You Don't Have to Love Facebook to Love the Code it Open Sources](#)

Container Journal: [Why a Cloud-Native Database Must Run on K8s](#)

Cassandra Tutorials & More

[Change Data Capture for Distributed Databases @Netflix](#) - Raghuram Onti Srinivasan

[Running your Database on OpenShift and CodeReady Containers](#) - Mick Semb Weaver

[Apache Cassandra Lunch #50: Machine Learning with Spark + Cassandra](#) - Rahul Singh

[Apache Cassandra Data Model: Components And Statements](#) - SimpliLearn



JOIN THE COMMUNITY

Cassandra Changelog is curated by the community. Please send submissions to cassandra@constantia.io.