

## DataStax Enterprise – Certified Apache Cassandra™

The open source Apache Cassandra™ NoSQL database has quickly become the preferred data management platform for Web, mobile, and IOT applications that need to scale and perform in a distributed environment that consists of either single data center deployments or multiple data centers and clouds. Cassandra’s masterless, shared nothing architecture provides enterprises with constant uptime for their transactional/operational database applications as well as a flexible data model capable of storing today’s modern datatypes and operational simplicity where database management is concerned.

The following explains the differences between open source Cassandra and the version of Cassandra certified for production use in DataStax Enterprise.

### Different Development and Release Models

The open source Cassandra community is instituting an accelerated model of software development. The goal is to provide a new Cassandra release each month, with odd months introducing new features and even months restricted to bug fixes only. This model is sometimes referred to as “tick-tock” because of its new feature/patch publication cadence.

Cassandra’s open source development model works well in that it is able to quickly introduce new features to the Cassandra community for feedback and testing, which helps features mature and stabilize more quickly than waiting many months to make many new features available at the same time. Any issues found are patched in subsequent releases. Notably, there is no “long term support” branch that receives *only* bug fixes.



Figure 1 — The tick-tock development model of open source Cassandra.

While the open source development model of Cassandra is tailor made for how open source software communities’ function and is done in the full spirit of open source, it does not lend itself for use in enterprise environments. The only way to get new bug fixes is to upgrade to the latest open source version, including any new features introduced in the meantime. Because enterprises typically have to re-test/re-certify production applications for database upgrades that contain new features, open source Cassandra’s release model makes application re-certification a near certainty for each database upgrade.

By contrast, Production-Certified Cassandra is developed and released in commercial software fashion with new open source releases of Cassandra being chosen by the experts at DataStax for production certification.

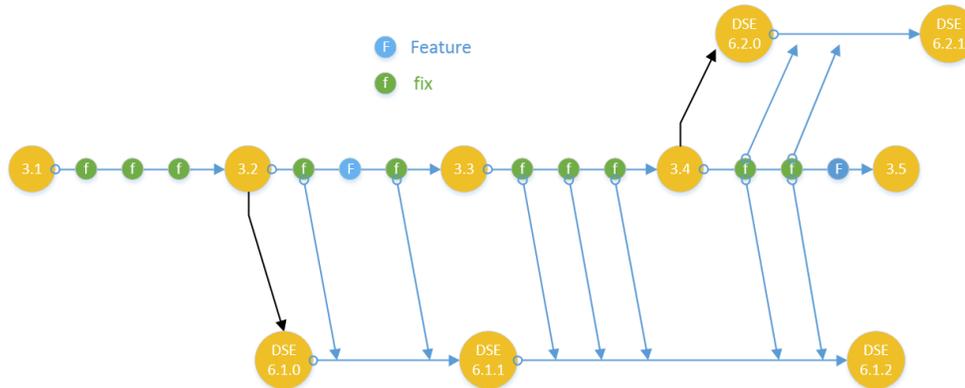


Figure 2 — The Production-Certified Cassandra development and release model by DataStax.

The production candidate selections then undergo an intensive production certification process and are managed in typical commercial software fashion, including certified software updates, long-term support, formal end-of-life policies, and more.

### DataStax Production Certification Program

To ensure a version of Cassandra is ready for enterprise-wide deployment, the staff at DataStax puts each selected version of open source Cassandra through the following.

#### Quality Assurance Testing

DataStax quality assurance (QA) testing goes far beyond typical open source community testing. Through professionally developed test harnesses that are constantly updated, the Cassandra database is rigorously tested and improved to ensure it is ready for the uptime and performance demands of enterprise environments. Any found issues are immediately included in Production-Certified Cassandra and also given back to the Cassandra open source community, although when those fixes are applied to a new open source Cassandra version is up to the community.

#### Scale Testing

Candidates for Production-Certified Cassandra are put through scalability testing that includes performance tests conducted with 1,000 node deployments on both Amazon and Google clouds. As with QA testing, any issues discovered are included in Production-Certified Cassandra and given back to the open source Cassandra community for inclusion in an upcoming version.

#### Platform Certification

Production-Certified Cassandra is certified by DataStax for production on the most popular operating systems and guaranteed to run without issue.

#### Cassandra Eco-System Software Inclusion and Validation

Cassandra enjoys a thriving software eco-system, which includes open source software that enables the running of analytics (e.g. with Apache Hadoop and Spark) and enterprise search (e.g. Apache Spark) operations on Cassandra data. The DSE platform includes built-in and tested versions of such software that are integrated with Production-Certified Cassandra and certified for production use by DataStax. This removes any guesswork associated with using popular open source software with Cassandra.

Further, the analytics and search software included in DSE enable full workload isolation on the platform where analytics, search, and transactional workloads on Cassandra do not compete with each for data or compute resources.

### Third-Party Software Validation

Each new version of Production-Certified Cassandra is validated for various third-party software offerings from DataStax partners (e.g. BI software from Tableau), with the testing being done either by DataStax or the software partner.

### Digitally Signed Code

Each version of Production-Certified Cassandra is digitally signed so that the actual software version can be validated and confirmed, a process required by many government agencies.

### **DataStax Certified Software Updates**

Certified software updates are routinely made available for Production-Certified Cassandra, the frequency of which is determined by a variety of factors (e.g. any major security issue is fixed and immediately made available). Patches from the open source community as well as those delivered by DataStax are certified for production environments and included in software updates released throughout the year.

In addition, certain software patches are backported to previous releases of Production-Certified Cassandra to ensure any bug found that threatens the production viability of earlier DSE versions is removed.

### **DataStax Hot Fixes**

DataStax customers that experience major performance or uptime issues due to a bug in Production-Certified Cassandra are able to report the problem and receive a customized, patched version of Production-Certified Cassandra that they can deploy back to production. This allows the issue to be immediately rectified vs. waiting for a new release of Production-Certified Cassandra (or open source Cassandra) that includes the patch.

### **DataStax End of Life Support**

Customers of DataStax receive full commercial software lifecycle support that includes both certified software updates and expert support for earlier versions of Production-Certified Cassandra that they still have deployed in production.

### **Differences Between Open Source Cassandra and Production-Certified Cassandra by DataStax**

The following table summarizes the differences between open source and Production-Certified Cassandra.

Feature	Open Source Cassandra	DataStax Enterprise
Commercial development and release model	No	Yes
Commercial quality assurance testing	No	Yes
Scale testing	No	Yes
Platform certification	No	Yes
Eco-system inclusion and validation	No	Yes
Third party software validation	No	Yes
Digitally signed code	No	Yes
Certified software updates (current and prior releases)	No	Yes
Hot fixes	No	Yes
Commercial software lifecycle support	No	Yes

### **Further Reading**

Production-Certified Cassandra is the only commercially developed, certified, and supported version of Apache Cassandra made available for enterprise deployments. For more resources and [downloads](#) of DataStax Enterprise, visit [www.datastax.com](http://www.datastax.com) today.