

VECTOR GRAPH - AUTOMATICALLY GENERATE GRAPH DATA AND IMPROVE QUERY RELEVANCY

[Back to Technology](#)

TECHNOLOGY • JULY 11, 2024

DataStax Hyper-Converged Database: The Future of Data Infrastructure Is Here

Your cloud native data solution.



Patrick McFadin

VP, Developer Relations & Cassandra Committer



Cloud native is now mainstream in infrastructure. In a recent [CNCF survey](#), 84% of organizations used or evaluated Kubernetes (and that figure has been growing steadily). Every cloud provider has a managed Kubernetes environment, and on-prem organizations are creating their own self-managed clouds. DataStax has been on the leading edge of the Data on Kubernetes movement since the release of [K8ssandra](#) and the K8ssandra operator nearly four years ago. And probably one of the strongest endorsements of this trend in technology, our [Apache Cassandra® as a Service, DataStax Astra DB](#), runs completely on Kubernetes and cloud native technology. We successfully deploy thousands of clusters a week reliably and quickly. As you'll see, we aren't just sitting on all that knowledge.

Our journey and experience

With over twelve years of bringing Cassandra to the enterprise and some of the best partnerships in the industry, we've developed a wealth of experience working with your needs. We listen and respond to your challenges, always keeping your needs at the forefront of our product development. A common request: you want to be more cloud native when it comes to self-managing your data infrastructure. Astra DB is a popular offering for those enterprises migrating to a managed Cassandra in the cloud. Not all of you can take advantage of this service, but you've been asking for something similar as a self-managed offering.

Introducing Hyper-Converged Database (HCD)

[DataStax Enterprise \(DSE\)](#) is close, but we can't just put it in a container and call it cloud native. So we took your input, dug deep into our experience with DataStax Astra and K8ssandra, and released our new self-managed option, [Hyper-Converged Database \(HCD\)](#)—a green-field approach to delivering data infrastructure for the future. HCD includes [Mission Control](#) to manage the deployment and ongoing operations while providing

integration with existing cloud native Kubernetes installations. The initial release will support Cassandra workloads with notable enhancements such as the Data API, which eliminates data modeling with Cassandra and has been adopted by developers in Astra. For generative AI workloads, we have an add-on for the state-of-the-art vector search engine, JVector.

How HCD differs from DataStax Enterprise

The first question you might ask is how this differs from DSE. DSE 6.9 was released today and is still a product line, but it is available in the classic download format of tar, package, or Docker image. DSE remains an option for teams using advanced workloads such as Solr, Spark, or Graph. If you're an OpsCenter user, it's available as a download for an easy in-place upgrade. Mission Control can manage DSE and HCD clusters if you want a single control plane for all your Cassandra clusters. These systems complement each other based on your use case and deployment profile.

The future of your cloud native infrastructure

The real difference is what this means for the future of your cloud native infrastructure. To be truly "cloud native," your applications should exhibit properties like scalability, elasticity, self-healing, and observability. HCD encapsulates our expertise in delivering these properties, enhancing your experience as you grow and strive for efficiency. We're committed to delivering a future-proof solution that will keep you ahead of the curve. Experienced users will no doubt notice that we're resetting some things to avoid technical debt, and that's true. The result is that you'll have an enterprise data option that will move fast and have a much more interesting roadmap based on the agility we'll now have. Roadmap items that we've already announced include the addition of streaming and OpenSearch fully integrated and optimized for developer experience.

What we're releasing today is version 1.0 of HCD, which covers core Cassandra workloads and includes the Data API; you can add [vector search](#) for next-generation application needs. If you want to build cloud native GenAI applications, this is a great place to start. You also get Mission Control to help deploy and administer your infrastructure in one or more datacenters. Backup, restore, repairs, migrations, security—all much easier than anything you've ever done before with Cassandra clusters. We can say that confidently because we do this every day on Astra DB, and now you can in your own self-managed systems.

To get started, we have some great resources for you. You'll begin with [Mission Control](#), which enables you to create and maintain an HCD cluster. You can follow the "[Getting Started](#)" section in the documentation or view our curated [YouTube playlist](#). Once Mission Control is up and running, deploying your first HCD cluster only takes a few minutes. You'll be amazed at how easy it is, and congratulations, welcome to the club! This is how our engineers running Astra DB feel every day.

[We've also made DataStax Enterprise 6.9 generally available; read about it!](#)

SHARE



More Technology

[View All](#) >



TECHNOLOGY

JULY 18, 2024

The Hitchhiker's Guide to Vector Embeddings