WHITEPAPER

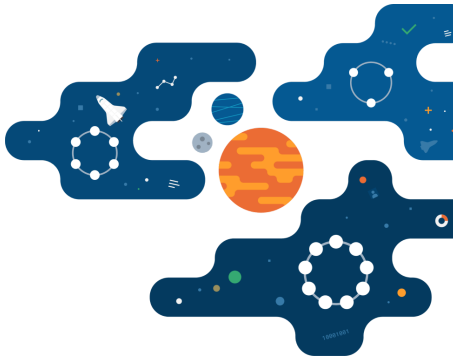

DataStax Astra Security Overview

Enterprise-level Security for Serverless Database-as-a-Service

DataStax Astra is the only open, multi-cloud, scalable, and serverless Database-as-a-Service. Built on Apache Cassandra, DataStax Astra simplifies cloud-native application development, and reduces time to install, deploy and scale from weeks to minutes. Some of the benefits of using DataStax Astra:

- **Zero Operations.** DataStax Astra eliminates operational overhead, the biggest obstacle to using Apache Cassandra, the open-source NoSQL database behind the largest applications in the world, including Netflix and Instagram. With DataStax Astra, enterprises can build secure, multi-cloud, multi-region applications with zero cloud vendor lock-in, zero-ops, and massively scalable database-as-a-service.
- **3x-5x Lower TCO.** In addition to savings due to less operational headcount, the benefits of the very scalable and highly available database can be achieved at significantly reduced costs since you only pay for the reads, writes, and storage that you use, instead of sizing for peak and paying for unnecessary capacity that sits idle most of the time.
- **High Availability.** The Astra service is designed to be resilient and highly available to minimize both downtime and the need for site-reliability engineering. Replication is done across nodes in multiple zones to ensure the highest possible availability. Astra relies on the DataStax-built and open-sourced Cassandra Kubernetes operator.

- **100% Open.** Astra is built on multiple Cloud Native Computing Foundation (CNCF) projects: Kubernetes, Prometheus, and Envoy and uses native GKE, AKS and EKS control and management planes with built-in management sidecars, metrics collectors, and configuration builders.
- **Developer Ready.** DataStax Astra makes it easy for developers to add persistence and statefulness to modern applications. Native support for Stargate Data APIs enables your organization to get productive immediately with Cassandra using REST, GraphQL, and schemaless Document (JSON), without having to master CQL or NoSQL database concepts.

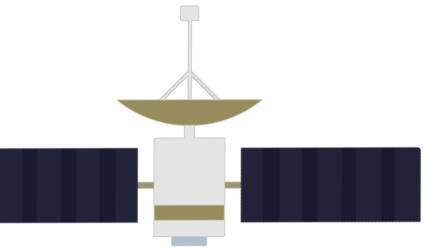

01

Trusted in the Cloud

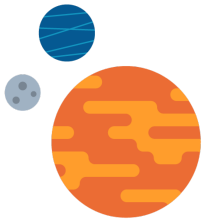

DataStax Astra employs various strategies to ensure security and availability.

Secure Shared Infrastructure

Within Astra, customer resources are segregated using Kubernetes namespace and networking policies. Customer clusters are deployed in subordinate accounts and projects across public clouds. Astra as a platform uses native Kubernetes Role Based Access Control to govern pod and service account actions to enforce security at the cloud platform level. This is not to be confused with Cassandra RBAC for database users.

As a database-as-a-service offering, all Astra infrastructure—the control plane, the Astra Cloud Console UI, and the database clusters are configured and operated within DataStax environments. The Astra control plane is built within AWS, Azure or GCP accounts, purpose-hardened to reduce security vulnerabilities.

Intrusion Detection and Prevention

Astra uses a cloud-native Intrusion Detection System to monitor and protect against malicious and unauthorized access through custom log-based alerting and container scanning.

VPC Peering

Select tiers of Astra support VPC Peering, which allows a user's personal or corporate-owned VPC to peer with the Astra VPC where the database runs. This ensures that the network connection between application clients and the database is entirely private. VPC peering is available for all cloud providers where DataStax Astra is offered, i.e., AWS, GCP and Azure.

Service Availability

DataStax commits to at least 99.9% Service Availability for a database during each calendar month. In the event DataStax does not meet the Service Availability, the Customer may be eligible to receive a Service Credit as described in the Service Level Agreement. This Service Level Agreement (“SLA”) applies only to DataStax Astra Databases at any paid Compute Tier that has been up for a minimum of 24 hours, and does not apply to any other products offered by DataStax.

02

Security Programs

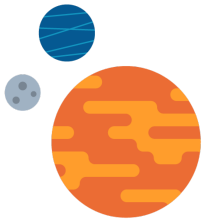

The DataStax organization follows global security standards and engineering best practices to weave security and privacy risk management traits into our DNA.

Industry Guidelines and Compliance

In the effort to infuse security and privacy best practices into our development and operational routines, DataStax is informed by the following industry guidelines.

- **NIST 8000-53.** The National Institute of Standards and Technology (NIST) 800-53 CyberSecurity framework forms the base of the program.
- **ISO 27001.** Information security management is informed by ISO 27001 best practices.
- **GDPR and CCPA.** Core privacy principles are based on the EU's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

DataStax is certified under the AICPA SOC2 framework. Reviews by external auditors affirm that controls are in place as of May 2020. The Type1 report is available under NDA as of late-June 2020 and the Type2 period ended on 31 January 2021, also available for review under NDA.

Organization Service Control Policies (SCPs)

The majority of operations performed on customer clusters are performed by software rather than staff. To the extent that manual operations are necessary, Service Control Policies (SCPs) provide guardrails against potential breaches:

- Staff access to production environments is granted by role with distinct privileges, following the principles of *least privileged access*.
- Access to compute instances running customer clusters is restricted.
- Actions are logged and used for security monitoring.
- Logs are audited and used for automated alerts to detect unauthorized activities.

Engineering Best Practices

In order to further ensure the highest level of security risk minimization, DataStax implements the following practices.

- The DataStax Security & Compliance team is made an integral part of the Engineering organization. They ensure that information security and privacy protections are built in by design.
- Third-party penetration testing is conducted at least three times per year, and a bug bounty program is always ongoing.
- All infrastructure is written as code and is peer-reviewed and tested as part of the SDLC (Software Development Lifecycle).

03

Data Platform Security

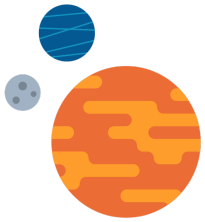

The DataStax Astra data platform includes capabilities that ensure confidentiality and data integrity.

Encryption

Users connect to Astra via a secure endpoint that provides in-transit encryption via industry-standard mutually authenticated TLS (mTLS)—where both sides of the TLS connection use information in DataStax-generated certificates to verify the connection—with unique certificates created for each database cluster.

At rest, data is protected with the use of cloud native encryption for ephemeral instance storage and persistent backup stores. Astra backups are stored in persistent storage volumes and are encrypted in each cloud provider's respective blob store: AWS Simple Storage Service (S3), GCP Cloud Storage, etc.

Identity and Access Management

JWT Based Token Authentication

Access to the Astra cloud portal UI is secured using ***KeyCloak.*** DataStax uses KeyCloak to implement third-party authentication with Google and Github for administrators and developers that need to manage Astra databases through the UI.

Role Based Access Control

DataStax uses ***Open Policy Agent*** to implement role based access control in Astra, where authorized access is granted based on role and the level of access is configured based on need. Default roles are available to give varied levels of permissions within an organization, and custom roles can also be configured.

Permission Log Audit

Audits can be made to review changes to the roles and permissions associated with each user, up to the past 90 days.

Developer Endpoints

Access using REST, GraphQL, DevOps, and Documentation API Endpoints
Astra offers developer endpoints for REST, GraphQL, and Documentation APIs that can be used to perform CRUD operations against the database. Authorization for the developer endpoints is achieved by the generation of an application token.

TLS is used for all interactions with the Astra developer endpoints (see Encryption). Specific instructions on how to use the developer endpoints can be found in the developer endpoint documentation.

Access to CQL endpoint using Secure Connect Bundle
The natively supported secure connect bundle streamlines the process for developers, administrators and service users to access the Astra database through a CQL secure endpoint. The secure connect bundle contains the keys and certificates required for mutual authentication (mTLS), avoiding the need for the user to interact with them directly. It also contains the CQL configuration (cqlshrc) needed to connect to Astra from the CQL shell (cqlsh).

The secure connect bundle makes it 10x easier and more efficient to connect with Astra compared to a traditional, encrypted connection against Cassandra.

Client ID and Client Secret
When an application token is first generated, Client ID and Client Secret are also generated and displayed. The permissions granted to these credentials are determined by the role selected. Client ID and Secret are available to copy or to download as a CSV only at the time of token generation.