A Deep Dive into Cube Security

Cube Cloud has robust security measures designed to protect the confidentiality, integrity, and availability of your valuable data so you don't have to. We comply with stringent security controls and regulations. For up-to-date and detailed information, please refer to our online security statement.

Data Confidentiality

Secure Infrastructure

Cube Cloud is deployed on trusted cloud providers like Amazon Web Services, Google Cloud Platform, and Microsoft Azure. If you choose the Enterprise or above subscription plan, you have the option to utilize a virtual private cloud (VPC) with dedicated infrastructure and enhanced networking security, ensuring data confidentiality both at rest and in motion.

There are also options for a "bring your own cloud" (BYOC) deployment.

Secure Storage

Cube makes use of a separate storage layer for storing metadata as well as for persisting pre-aggregations as Parquet files. Cube Store can be configured to use either AWS S3 or Google Cloud Storage giving you the choice to use external storage for data cache.

Cube Store requires strong consistency guarantees from underlying distributed storage and recommends AWS S3, Google Cloud Storage, and Azure Blob Storage.

API Security

Cube Cloud gives you complete control over API exposure. Configure cross-origin resource sharing (CORS) and employ industry-standard JSON Web Token (JWT) based authentication flow. Have full control over the JWT authentication behavior by directly writing authorization code, not just changing parameters. You can also customize API security through checkAuth and checkSqlAuth configuration options, guaranteeing the confidentiality of data.

Semantic Layer Access Control

With Cube Cloud, you have full control over which semantic layer entities (cubes, views, measures, dimensions) are exposed to specific users and groups. You can define query permissions and control direct access to upstream data sources, ensuring data confidentiality.

Multitenancy

Cube Cloud supports flexible multi-tenant configurations. This enables each tenant to have their own semantic layer configuration, promoting data segregation and confidentiality.

Data Integrity

Development Mode

Cube Cloud provides isolated environments for developing and validating the semantic layer configuration. This prevents any unintended changes from compromising the integrity of your data.

Continuous Deployment (CD)

By seamlessly integrating with popular version control systems like Git and hosting services such as GitHub, GitLab, and BitBucket, Cube Cloud aligns with your continuous deployment process. This ensures the integrity of the semantic layer throughout the development lifecycle.

Single Sign-On (SSO)

Cube Cloud offers secure access to the semantic layer configuration through integration with external authentication providers. In addition to options to login using email and password, a GitHub account, or a Google account, Cube Cloud provides single sign-on (SSO) via identity providers supporting industry-proven SAML 2.0 protocol, e.g., Okta, Google Workspace, Azure AD, etc. This enhances the integrity of the semantic layer by restricting access to authorized personnel. Learn more about SSO here.

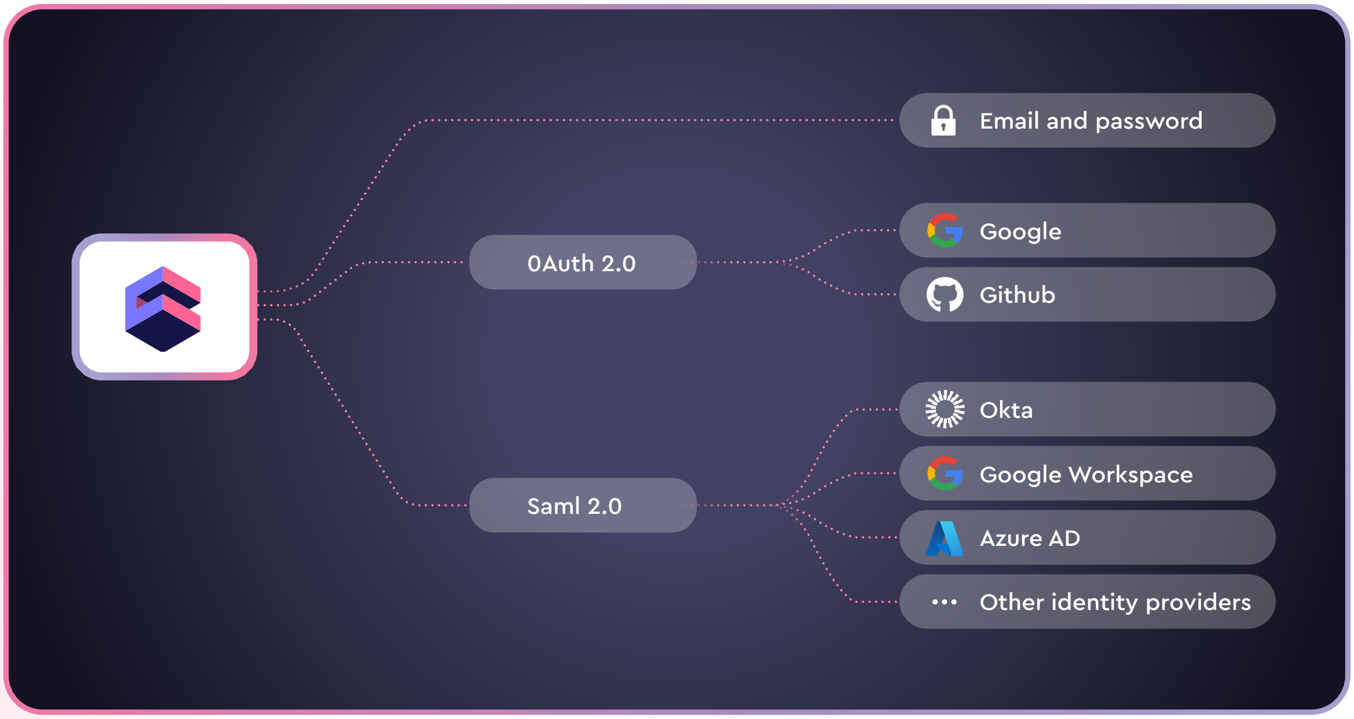

Role-Based Access Control (RBAC)

Granular access control is enforced by Cube Cloud, allowing you to assign roles to responsible personnel. This ensures that only authorized individuals can modify the semantic layer, maintaining data integrity.

CUBE SECURITY DEEP DIVE

Data Availability

Blue-Green Deployments

Cube Cloud uses the Blue-green deployment methodology. At any given time, only one server is handling requests (e.g., being pointed to by the DNS). Changes are installed on the non-live server, which is then tested through the private network to verify the changes work as expected. Once verified, the non-live server is swapped with the live server, effectively making the deployed changes live.

This prevents faulty changes from impacting production environments during semantic layer updates and ensures uninterrupted service availability for downstream data applications.

Cache Warm-Up

Before going live, Cube Cloud populates data cache based on its configuration. This prevents cache misses, ensuring continuous service availability for downstream data applications.

High Availability

Cube Cloud provides fully managed infrastructure with production cluster and multi-cluster deployment options. Redundant components, including API instances and Cube Store nodes, are designed to eliminate single points of failure and ensure fault tolerance and high availability.

Auto-Scaling

Our fully-managed infrastructure supports auto-scaling of components like API instances and Cube Store nodes. This enables Cube Cloud to handle load bursts, prevent denial of service, and ensure continuous service availability for your data applications.

Monitoring

Cube Cloud offers configurable alerts and monitoring integrations to notify on-call personnel about upstream data source timeouts and API outages. This proactive approach helps ensure continuous service availability for downstream data applications.

Compliance

Cube Inc. is committed to meeting rigorous compliance standards. We continuously undergo audits and annual pentests. The Report can be shared upon request.

Cube complies with the following controls and regulations.

In regard to GDPR:

- For the account data we collect we act as data controller and adhere to our Privacy Policy.
- For the data you upload, we act as a data processor and adhere to our DPA which we sign with enterprise customers.

Learn more in our Documentation or reach out to your Account Exec or anyone on the Cube team for a deeper dive into your specific security needs.

AICPA SOC 2 TYPE 1

AICPA SOC 2 TYPE 2

HIPAA

GDPR