

GOOGLE CLOUD PLATFORM

STORAGE OVERVIEW





STORAGE NEEDS

- Almost all applications need persistent durable storage
- User accounts, images, events, documents, logs, etc
- Different applications have different storage needs
- Google Cloud Platform offers seven storage options - Bigtable, Cloud Datastore, Cloud SQL, Cloud Storage, Cloud Spanner, Persistent Disk and BigQuery
- BigQuery is both a storage service and a powerful analysis tool





KEY FEATURES

- Single API across storage classes
- Scalable to exabytes of data
- Designed for 99.9999999999% durability
- Very high availability across all storage classes
- Time to first byte in milliseconds
- Strongly consistent listing





STRUCTURED AND UNSTRUCTURED

- We can divide these seven storage services into two categories: **structured** and **unstructured**.
- If the data you want to store can be organized into a table structure with columns and rows, then it is **structured** data.
- Examples: user profile information, event logs, sensor measurements, sales records, or stock trade data.
- **Structured** data comes in many different shapes, sizes, and usage patterns, and there is a great diversity of ways to store it and interact with it.
- Cloud SQL, Cloud Datastore, Cloud Bigtable, Cloud Spanner, Persistent Disk and BigQuery all store structured data.





STRUCTURED AND UNSTRUCTURED

- Google Cloud Storage stores **unstructured** data
- Provide Cloud Storage a sequence of bytes to store, a place you want to store it, and a name to identify that sequence of bytes, and it stores them
- Cloud Storage provides no insight into the internal structure
- Data is stored in Buckets as Objects





PERSISTENT DISK

Description

- Fully-managed, price-performant block storage that is suitable for virtual machines and containers.

Good for

- Block storage for Google Compute Engine and Google Kubernetes Engine
- Snapshots for data backup

Common Workloads

- Disks for virtual machines
- Sharing read-only data across multiple virtual machines
- Rapid, durable backups of running virtual machines





CLOUD BIGTABLE

Description

- A scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for

- Low-latency read/write access
- High-throughput analytics
- Native time series support

Common Workloads

- IoT, finance, adtech
- Personalization, recommendations
- Monitoring
- Geospatial datasets
- Graphs





CLOUD DATASTORE

Description

- A scalable, fully-managed NoSQL document database for your web and mobile applications.

Good for

- Semi-structured application data
- Hierarchical data
- Durable key-value data

Common Workloads

- User profiles
- Product catalogs
- Game state





CLOUD SQL

Description

- A fully-managed MySQL and PostgreSQL database service that is built on the strength and reliability of Google's infrastructure.

Good for

- Web frameworks
- Structured data
- OLTP workloads

Common Workloads

- Websites, blogs, and content management systems (CMS)
- Business Intelligence (BI) applications
- ERP, CRM, and eCommerce applications
- Geospatial applications





CLOUD SPANNER

Description

- Mission-critical, relational database service with transactional consistency, global scale and high availability.

Good for

- Mission-critical applications
- High transactions
- Scale + Consistency requirements

Common Workloads

- Adtech
- Financial services
- Global supply chain
- Retail





CLOUD BIGQUERY

Description

- A scalable, fully-managed Enterprise Data Warehouse (EDW) with SQL and fast response times.

Good for

- OLAP workloads up to petabyte-scale
- Big Data exploration and processing
- Reporting via Business Intelligence (BI) tools

Common Workloads

- Analytical reporting on large data
- Data Science and advanced analyses
- Big Data processing using SQL





CLOUD STORAGE

Description

- A scalable, fully-managed, highly reliable, and cost-efficient object / blob store.

Good for

- Images, pictures, and videos
- Objects and blobs
- Unstructured data

Common Workloads

- Storing and streaming multimedia
- Storage for custom data analytics pipelines
- Archive, backup, and disaster recovery





MOBILE STORAGE OPTIONS

- **Cloud Storage for Firebase:** Mobile and web access to Google Cloud Storage with serverless third party authentication and authorization.
- **Firebase Realtime Database:** A realtime, NoSQL JSON database for your web and mobile applications.
- **Firebase Hosting:** Production-grade web and mobile content hosting for developers.
- **Cloud Firestore for Firebase:** A NoSQL document database that simplifies storing, syncing, and querying data for mobile and web apps at global scale.





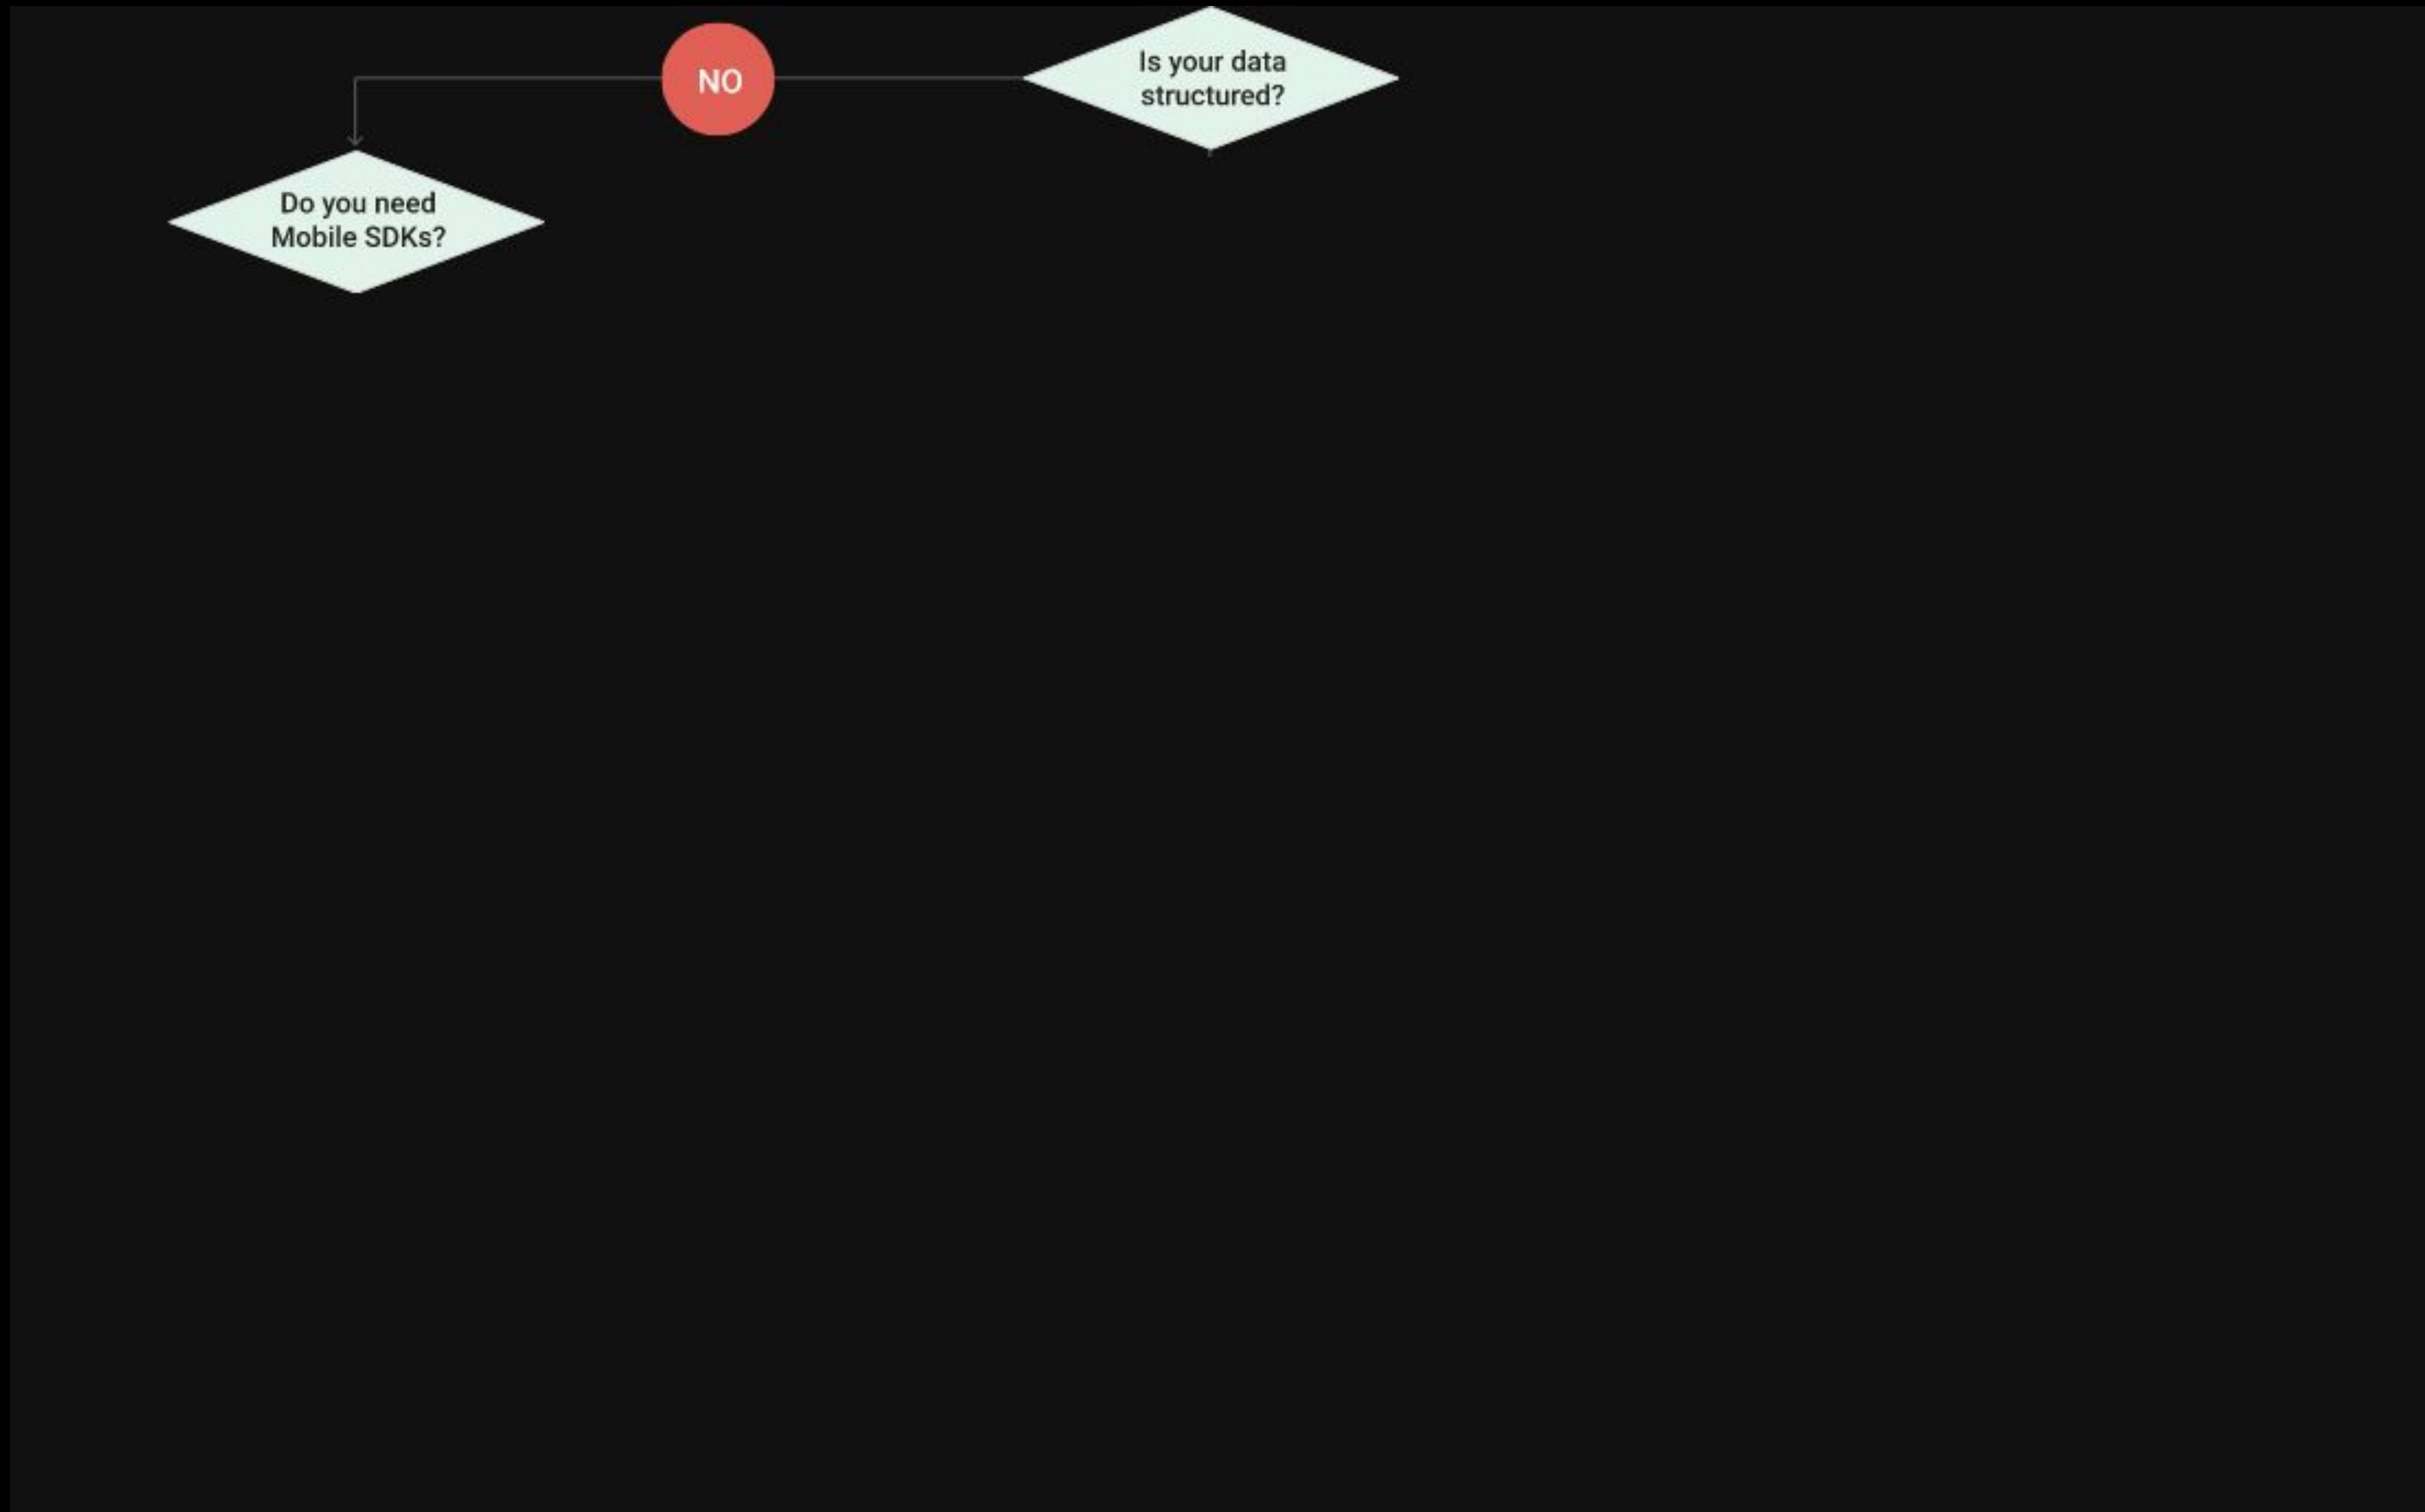
HOW TO CHOOSE

Is your data structured?



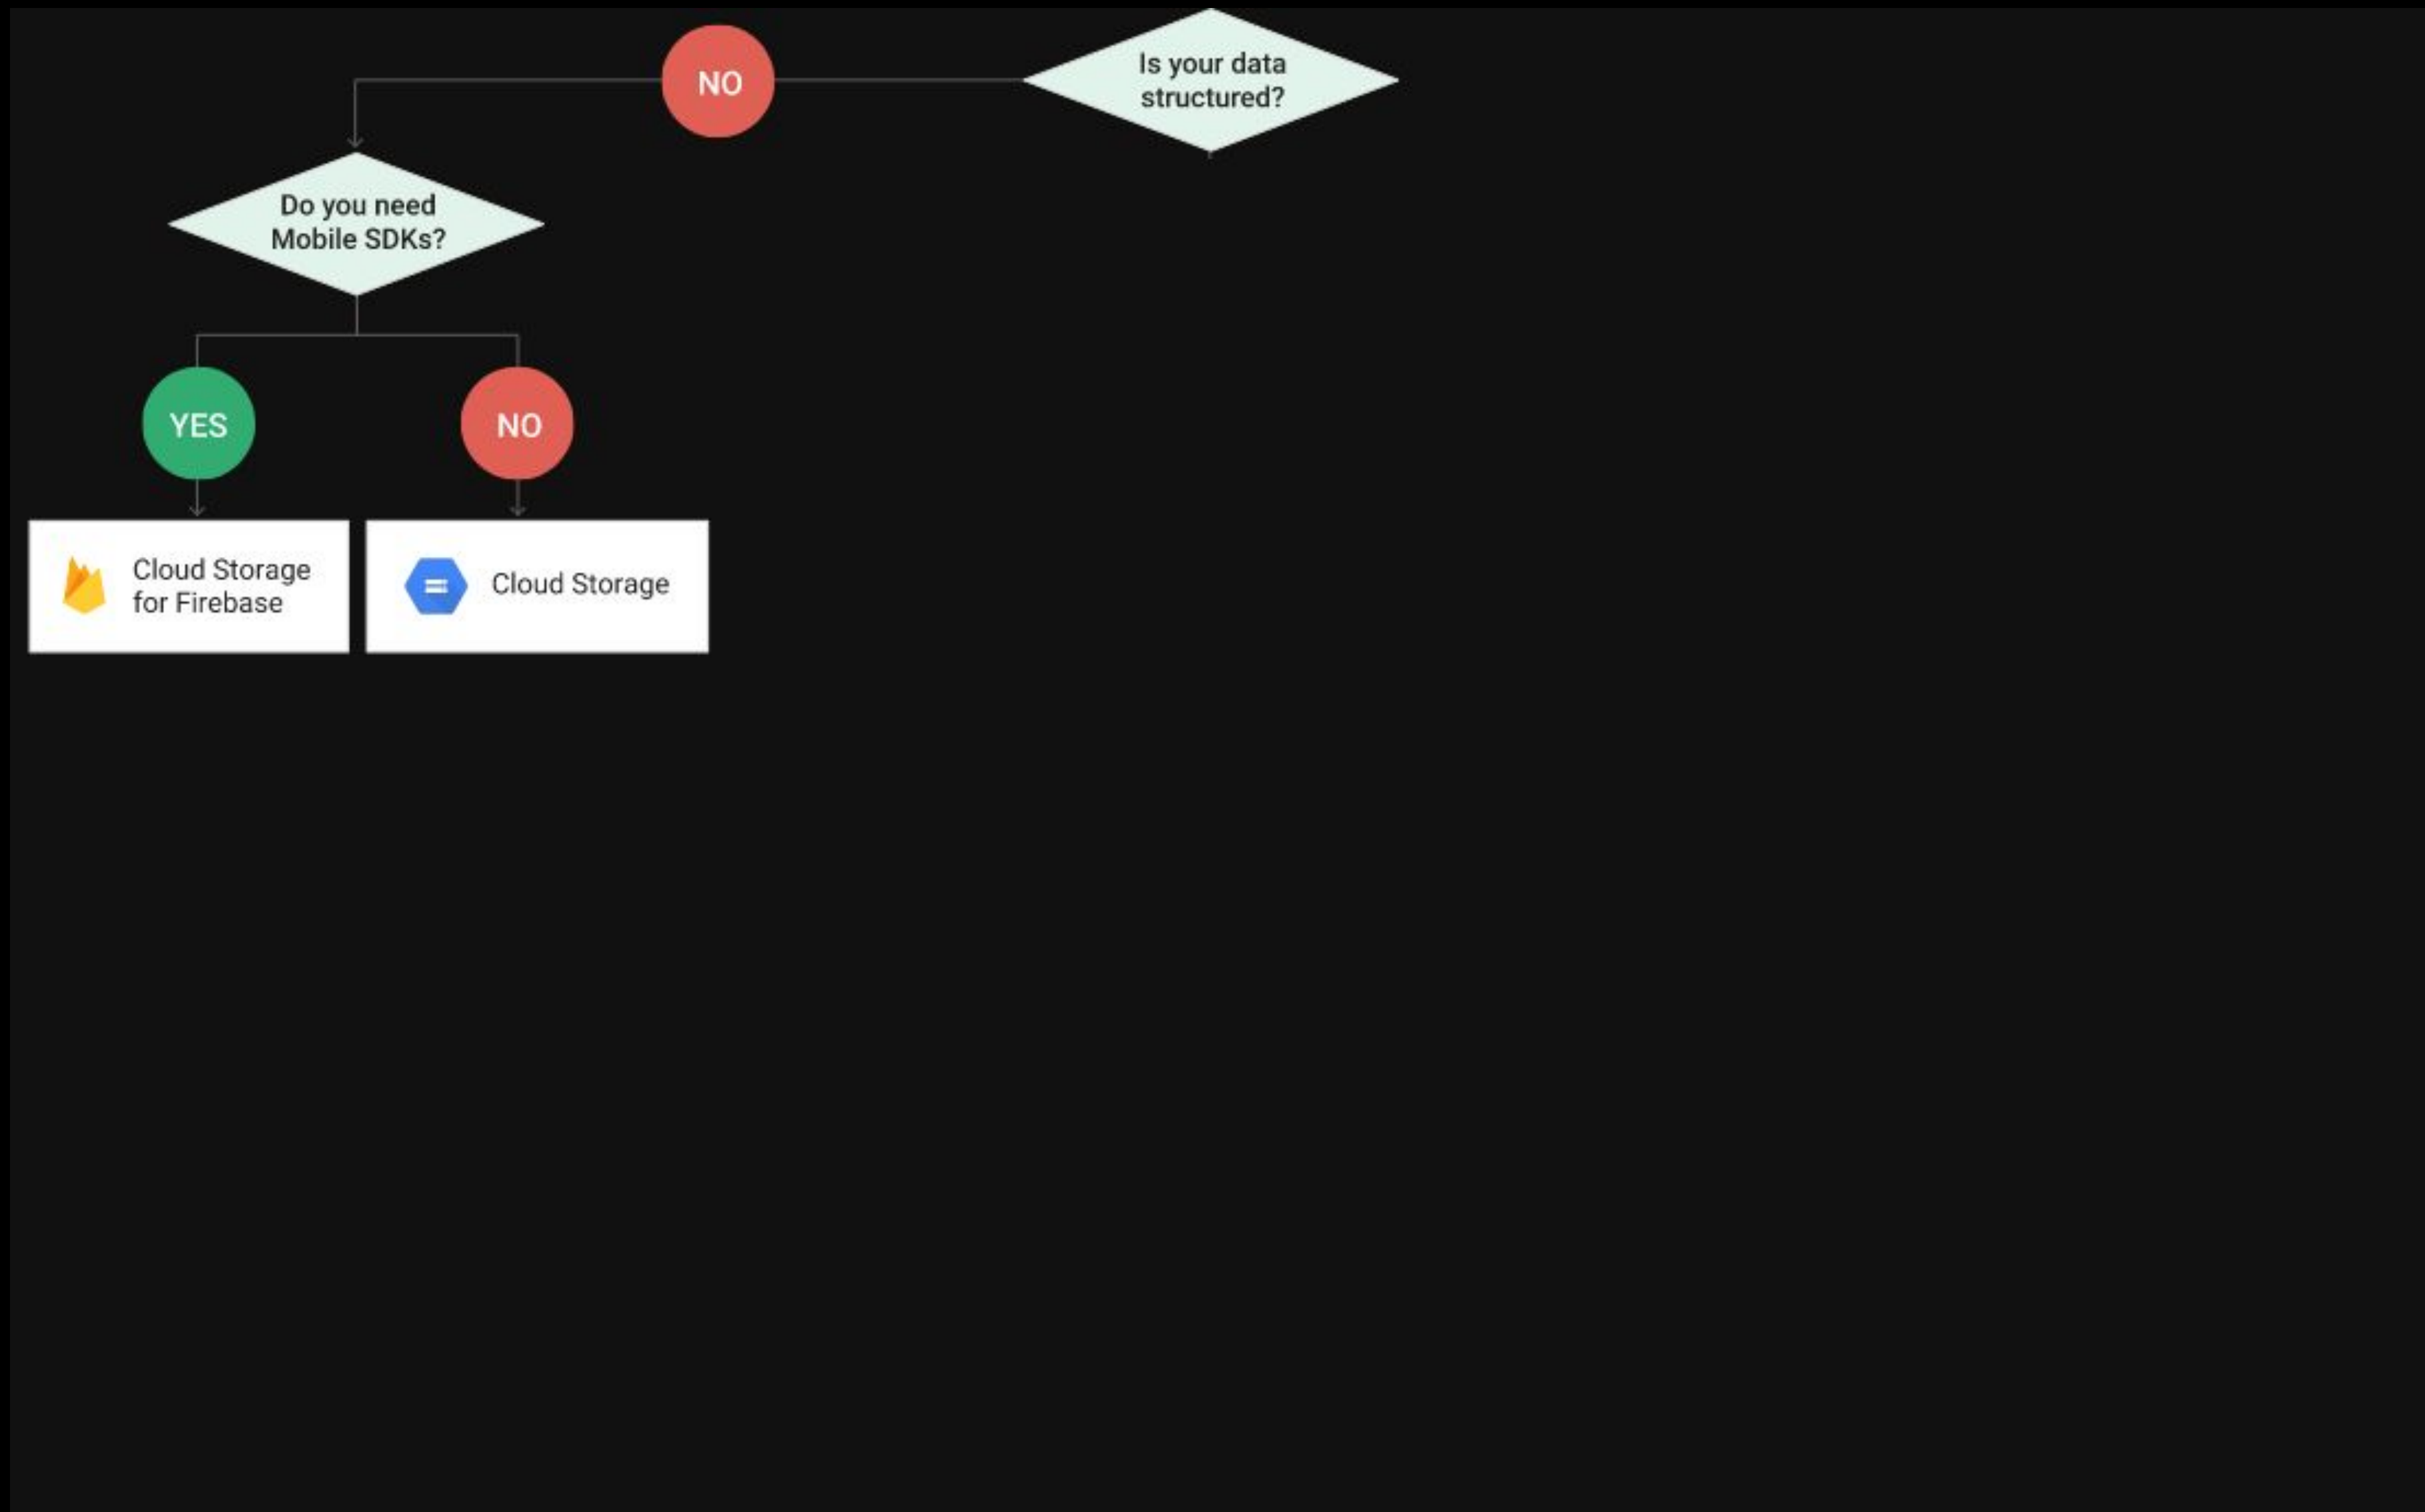


HOW TO CHOOSE



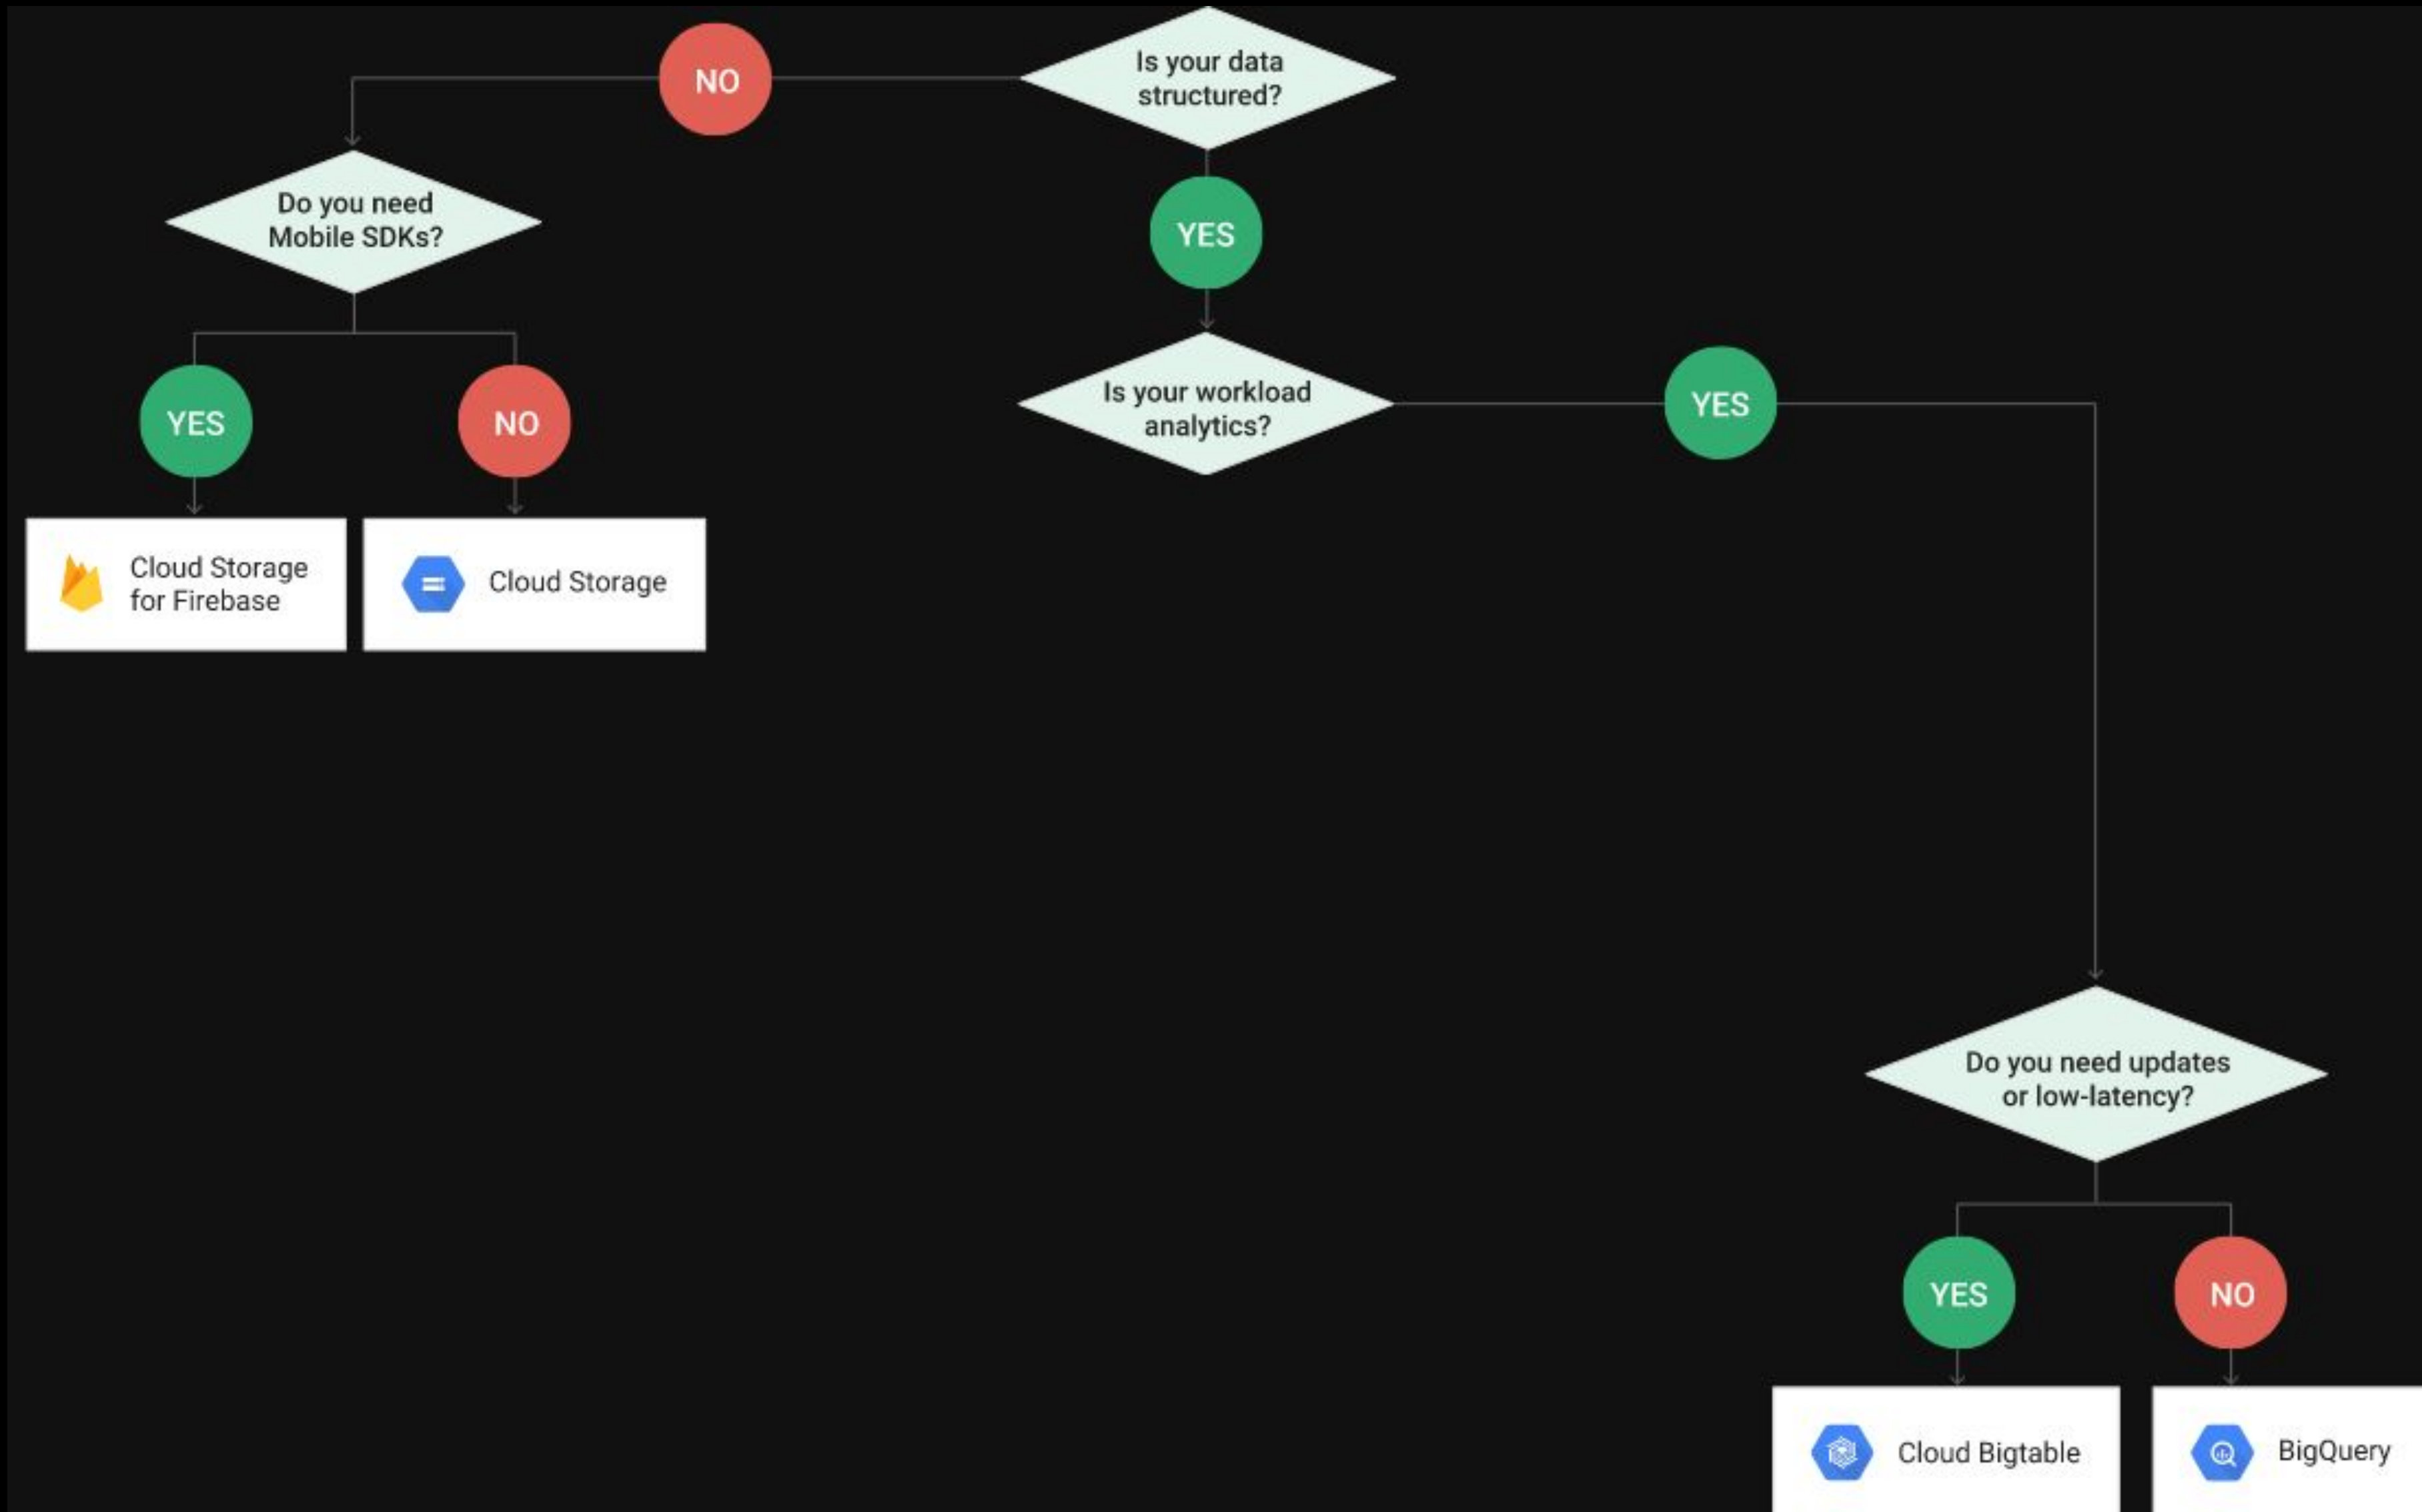


HOW TO CHOOSE





HOW TO CHOOSE



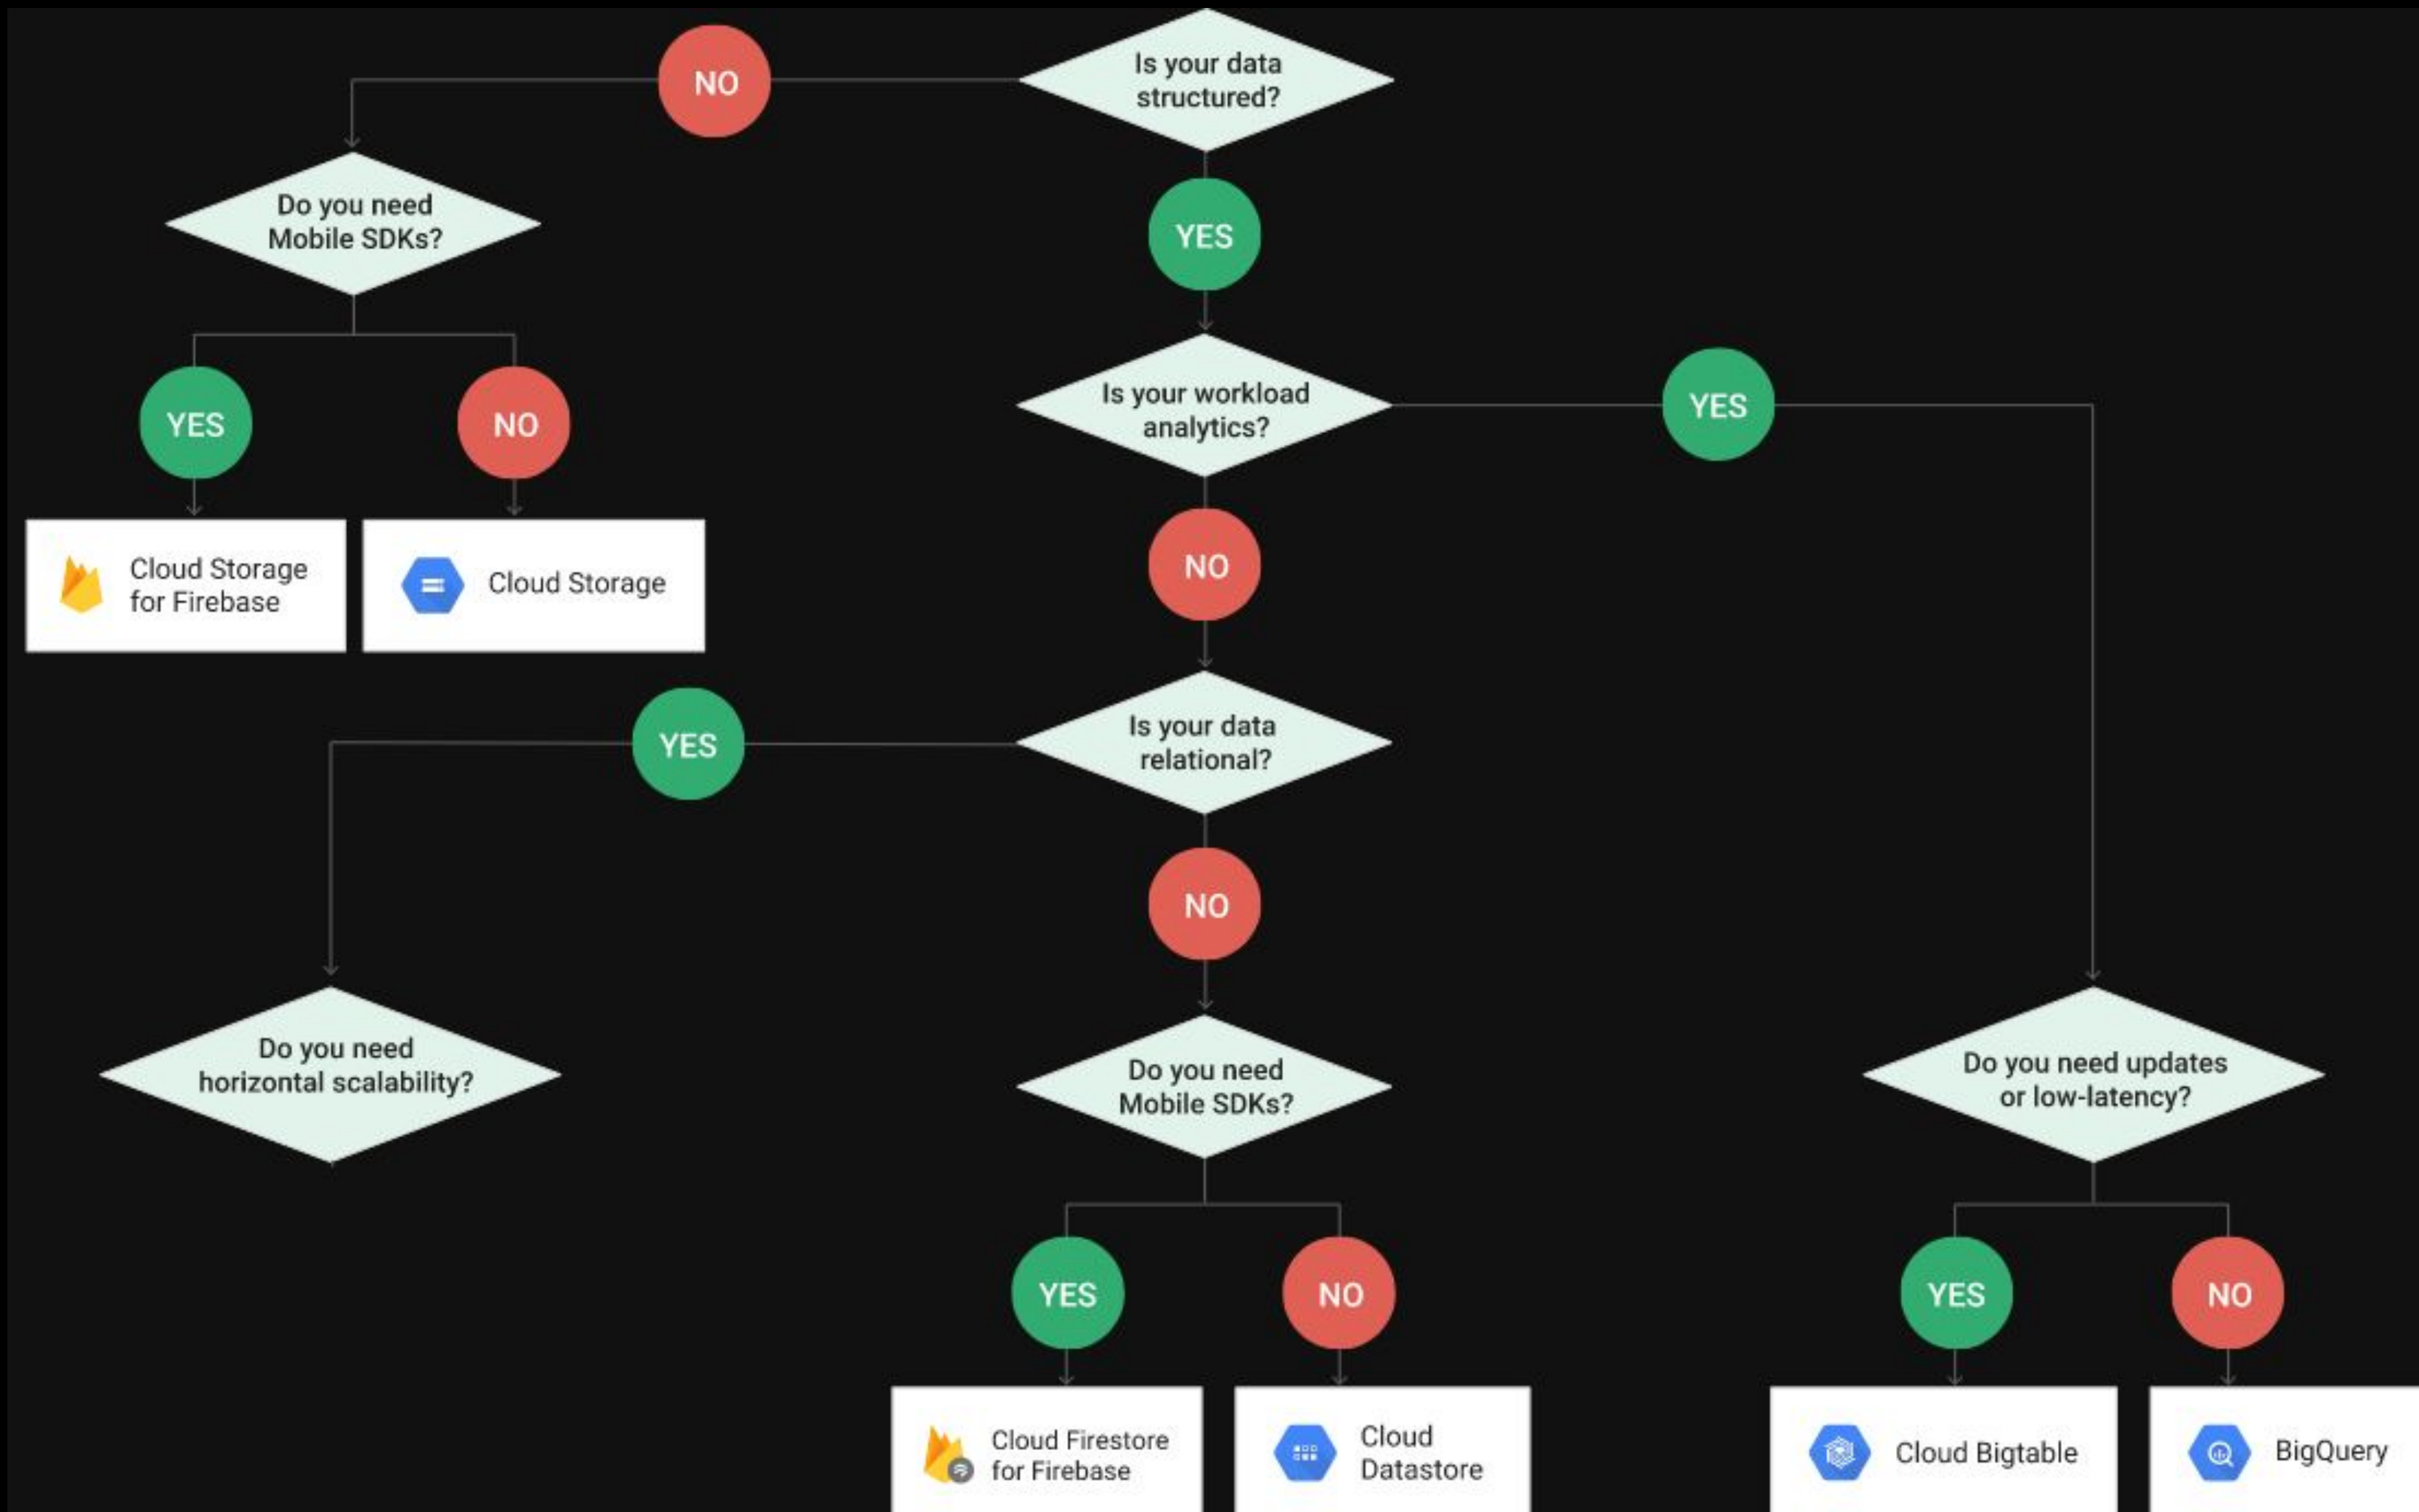


HOW TO CHOOSE





HOW TO CHOOSE





HOW TO CHOOSE

