SparkLogs

DATA SHEET

A managed data platform for telemetry analytics

SparkLogs combines petabyte-scale telemetry ingestion, flexible querying, and purpose-built MCP tools for AI-driven data exploration and analysis. Built on BigQuery, it manages evolving event data, queryable history, and portable archives, with optional Private Cloud access for custom analytics in your own Google Cloud Platform (GCP) project.

From event ingestion to custom analytics

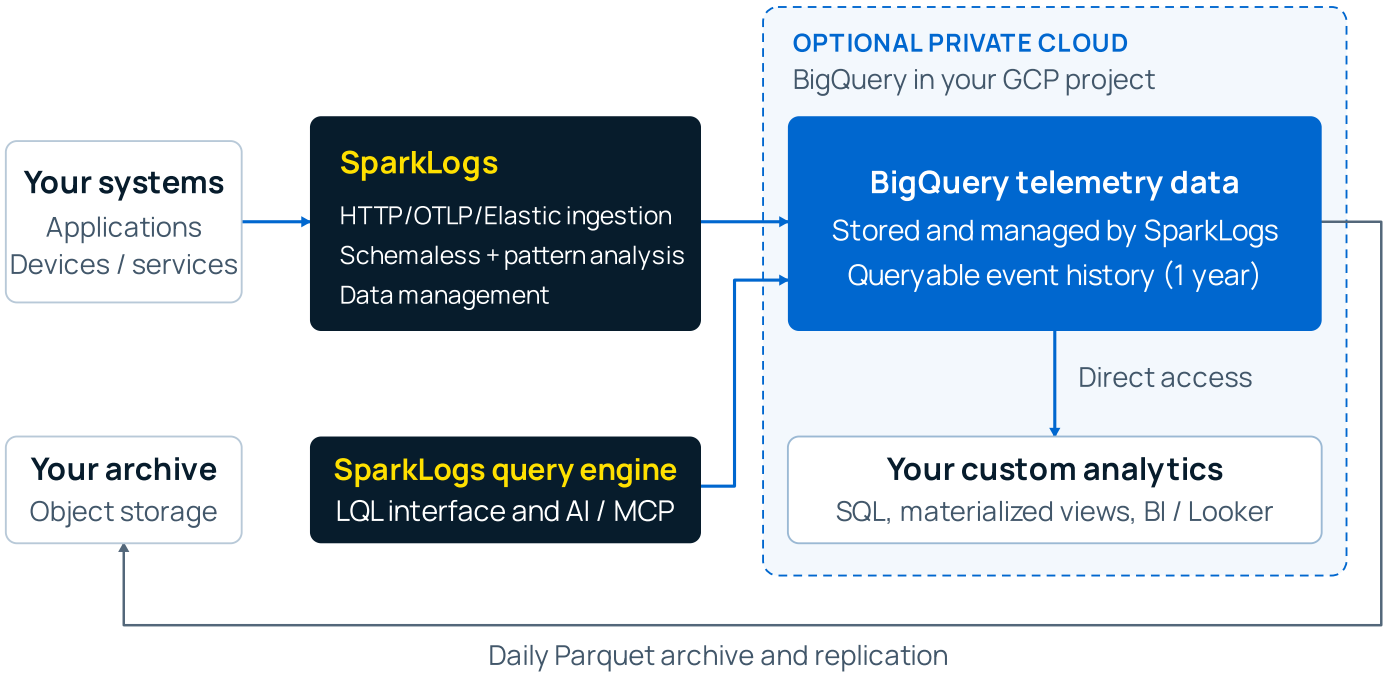

Ingest evolving event data

SparkLogs automatically scales ingestion to accommodate sharp changes in event volume, including sudden bursts from idle. Send nested JSON and custom fields through open protocols without defining an ingestion schema in advance. SparkLogs handles field extraction and mapping in realtime. The fully managed service requires no customer-managed ingestion servers or clusters.

Explore data with purpose-built AI tools

SparkLogs MCP tools help AI agents discover data, construct queries, and run aggregate and conditional analysis. Teams can also use LQL, the SQL-like query language, and full-text search in the SparkLogs interface. LQL is specialized for semi-structured data, with powerful operators, multivariate typing, and array expressions.

Custom data workflows and dashboards

With the Private Cloud option, you store your telemetry in BigQuery in your own GCP project. Your team can build SQL models, materialized views, and reporting workflows directly on that data, using BigQuery and compatible BI tools. SparkLogs provides ingestion, data management, query tools, and archive replication. Your team develops the business-specific analytics.

Private Cloud is optional. SparkLogs Cloud provides the managed service with SparkLogs-managed data storage and query processing.

Customer Case Study: Big Time Studios

Big Time Studios uses SparkLogs to ingest and analyze game telemetry to optimize player experience and game economy. The studio combined logging and analytics in one platform, retiring 3 fragmented data silos and reducing costs by 80%.

Read the customer story

Standards-based ingestion for application and operational telemetry

CAPABILITY	DETAIL
Native API	HTTPS JSON ingestion for structured and semi-structured events.
OpenTelemetry	OTLP/HTTP logs ingestion with JSON and Protobuf encodings.
Other ingestion APIs	Elasticsearch bulk ingestion and Loki push API support.
Open-source shippers	Vector, Fluent Bit, OpenTelemetry Collector, and Grafana Alloy.
Event structure	Nested JSON and custom fields without a predefined ingestion schema. Automatic standard-field mapping, storage, and querying of any arbitrarily complex JSON.

Query through SparkLogs, from your own AI systems, or directly in BigQuery

- **LQL query language:** SQL-like, type-aware queries over event fields, with full-text search and interactive exploration.
- **AI through MCP:** Purpose-built tools for data discovery, aggregate and conditional queries, and event exploration within each user's permissions.
- **Adaptive-scale querying:** enables interactive exploration of 100+ billion event datasets, dynamically zooming in on areas of interest.
- **Private Cloud:** Direct BigQuery access for customer-built SQL, materialized views, BI, and analytical workflows.

Query billions of events in seconds

SparkLogs supports analysis across tens of billions of events, with demonstrated at-scale query times under 10 seconds without sampling.

For queries in your private-cloud project, BigQuery fluid scaling adjusts query capacity each second and compute bills by the second. Capacity scales up from and back to zero within seconds. Query times depend on workload and configuration.

Queryable history and a portable archive

SparkLogs Cloud plans retain data for live querying for up to 1 year. Private Cloud plans support longer retention according to plan and configuration.

Daily archives store ingested data in compressed, Hive-partitioned Parquet. Replicate them to your own AWS S3, Google Cloud Storage, Azure Blob Storage, or S3-compatible bucket for separately managed retention and downstream analysis.

Hierarchical organization and access control

Organize data into hierarchical scopes and control access by role. Configure users for aggregate-only queries or permit access to individual event data, including through MCP.

SparkLogs permissions govern access through SparkLogs. Direct BigQuery access follows your GCP permissions. Data is encrypted in transit and at rest. Visit the Security and Trust Center

Choose where your data lives

	SPARKLOGS CLOUD	SPARKLOGS PRIVATE CLOUD
Data storage	SparkLogs-managed	BigQuery in your GCP project
SparkLogs interface and MCP	Available	Available
Direct BigQuery access	Not provided	Available
Parquet archive replication	Supported	Supported

Discuss your telemetry and reporting requirements

Review your data sources, event volume, retention needs, and reporting queries with the SparkLogs team to discuss requirements and assess fit.

Contact
sparklogs.com/contact-us
Documentation
sparklogs.com/docs/platform