

Actian Avalanche on Google Cloud

Hybrid Cloud Data Warehouse

Highlights

Run blazing fast ad-hoc queries:

The fastest cloud data warehouse gets even faster—Avalanche takes advantage of Google Cloud’s advanced infrastructure to run 20% faster than on other cloud platforms

Enjoy unbeatable cost savings:

Provides unbeatable price-performance compared to any other cloud data warehouse offering, including a 6x advantage over Snowflake

Future-proof with hybrid multi-cloud architecture: Avalanche can be deployed on Google Cloud, AWS, Azure, virtual private cloud, and on-premises using the same technology

Rapidly deploy with native integration: Avalanche comes with pre-built integration to various cloud storage options (e.g. Google Cloud Store, S3, ADLS) as well as 200+ connectors to SaaS applications such as Salesforce

Improve insights for all with advanced analytics: Avalanche supports a broad spectrum of users, from business analysts to data scientists, with familiar tools and next-gen ML powered by high performance model training

Work seamlessly within the Google Cloud ecosystem: Avalanche is tightly integrated with the core Google Cloud ecosystem and maximizes the most of Google innovation

Actian Avalanche hybrid cloud data warehouse leverages Google Cloud’s scalable infrastructure and global network to deliver breakthrough performance, concurrency and cost savings for data-driven enterprises. Avalanche features native integration with Google Cloud’s Looker business intelligence and analytics platform, and hundreds of popular SaaS and enterprise applications, to provide a comprehensive solution that is easy to deploy and consume.



Figure 1. Avalanche is integrating with Google Cloud ecosystem to make the user experience seamless.

Avalanche on Google Cloud Delivers the Best of Both Worlds – High Performance at Low Cost

Actian Avalanche is designed from the ground up to deliver unrivaled performance on commodity infrastructure. With this purpose-built architecture, Actian Avalanche delivers the best price-performance in the industry without the need for DBA tuning and optimization techniques.

When combined with Google Cloud technology, Avalanche becomes even faster. Avalanche intelligently uses Google Cloud’s superior backhaul and high-bandwidth networking infrastructure to run 20% faster than on other cloud platforms –while providing some of the lowest TCO available.

Features

- Highly scalable MPP relational analytics database
- One-touch deployment of clusters across multiple clouds from a single User Interface
- Industry's fastest performance at scale with no tuning required
- Perform queries on live data with updates enabled through patented algorithms that do not impact read performance
- Fully SQL-2016 compliant
- Python, SQL and JavaScript UDF support
- Support for diverse data types, including JSON, with auto-schema recognition and ingestion via easy user interface
- Open, industry standard data access, including Spark, ODBC, JDBC, .NET
- Extensive interoperability with native cloud ecosystems
- 200+ pre-built enterprise application and SaaS connectors included
- Comprehensive security features, including:
 - SOC-II compliance
 - Built-in data masking
 - Customer managed keys
 - Cloud Security Alliance member
 - Enterprise firewalls
 - Intrusion detection
 - SIEM logging
 - AES 256-bit data encryption
 - Authentication integration
 - Security patching
 - Google Cloud's security infrastructure and features

For the same cost as alternative solutions, you can benefit from substantially better performance or choose the same performance for significantly lower cost. For example, Avalanche provides on average **6X price-performance advantage over Snowflake** as measured by GigaOm's TPC-H performance benchmark study and even more against many of the appliance vendors. As you increase workload complexity and number of concurrent users, Avalanche outperforms by even wider margins.

With its roots in an on-premises deployment, Avalanche was designed for a scarcity of resources and therefore makes the most of every CPU clock cycle, every byte of RAM and every I/O operation. This means you benefit from the most efficient use of these resources while also enjoying cloud economics such as "pay only for what you use."

Avalanche takes advantage of patented features and industry best practices such as the following:

- **Vectorized processing:** Avalanche executes an operator such as "add" on hundreds of data elements during a single CPU cycle
- **CPU cache maximization:** Avalanche uses private CPU core and caches as execution memory, which is 100x faster than using RAM
- **Optimized I/O:** Tuning algorithms optimize the cache just before and after operations to further enhance performance.
- **Advanced compression:** Avalanche maximizes the efficiency of decompression but also delivers 4-6x compression ratio
- **High concurrency:** Avalanche supports a high volume of concurrent users, allowing up to 64 concurrent users out-of-the-box
- **Columnar storage:** Avalanche stores data together in data blocks on column-by-column basis
- **Automated storage indexes:** Avalanche assigns indexes automatically as it applies positional changes on load to/from disk
- **Trickle updates:** Avalanche applies updates to the database based on system load to always maximize performance




Figure 2. Actian Avalanche is available on Google Cloud, AWS and Microsoft Azure, enabling multi-cloud deployment options from a single pane.

Hybrid, Multi-Cloud Deployment Opportunities for Maximum Flexibility

Avalanche is the industry's first hybrid data warehouse solution that can be deployed on all three leading cloud platforms (Google Cloud, AWS, Azure) as well as on-premises, with native integration with Google Kubernetes Engine (GKE) and Google Anthos.

Because Avalanche can be used in a hybrid configuration, organizations can configure Avalanche to run on-premises and in the cloud as well as in multi-cloud environments (public/public as well as public/private). This enables organizations to realize the true potential of hybrid by bringing the compute power of Avalanche to the place where their data resides, rather than moving their data to the data warehouse.

Whether you're running Avalanche in the cloud or on-premises, you can leverage the same database engine, the same physical data model, the same ETL/ELT tools, and the same BI tools of your choosing. You gain all the benefits of a modern cloud deployment such as elastic and independent scaling of compute and storage while retaining the flexibility to keep on-premises applications for as long as you need.

What this means for you:

- **Lower cost and improved performance:** Processing and querying data where applications are producing it, whether on-premises or the cloud, often produces the best results in terms of cost and performance since data movement from the cloud is expensive and slow
- **Simplified, futureproof architecture:** Since Avalanche relies on the same patented vectorized database engine both in the cloud and on-premises, you will work with a single data model, consistent ETL integration, and one technology to learn
- **Stronger compliance and security:** You have the option to retain complete control over sensitive datasets. Highly secure workloads can remain in the data center if the CISO demands it, while the cloud is used for other workloads.
- **Amortized on-premises investments:** Workloads that are optimal for the cloud can move immediately to the cloud while those that can be handled on-premises can run on infrastructure that has already been paid for and may be cheaper
- **Flexible CapEx/OpEx pricing:** Share your expenses between CapEx and OpEx as your needs dictate
- **Phased, non-disruptive migrations:** Workloads can stay on-premises until it is ready to move to the cloud

Built for Real-Time Performance

With Avalanche, you can achieve blazing fast response on ad-hoc BI queries using the freshest data. Unlike other cloud data warehouses, Avalanche enables real-time updates to data in the data warehouse without adding latency. You can be assured of accessing the latest data to power your analytics without paying a performance penalty.

Its advanced columnar implementation enables the least I/O performed while retrieving data from disk. Avalanche's vectorized compute leverages CPU Single Instruction, Multiple Data (SIMD) and processes data in the L1/L2 CPU cache instead of RAM, leading to significantly faster performance.

Rapidly Deploy with Native Integration for Enterprise BI & Analytics

Avalanche is the industry's first cloud data warehouse to offer integration capabilities natively built into the product. Avalanche has direct loaders to pull in bulk data from popular data sources such as Google Cloud Storage, S3, ADLS and Hadoop. Avalanche also includes over 200 connectors and templates for easily sourcing and moving data from SaaS applications to Avalanche data warehouse at scale—no special ETL is required. You can load structured and semi-structured JSON data, including event-based and streaming, without coding.

Through ODBC, JDBC, .Net and Python, Avalanche enables business analysts, data scientists and data engineers to use their favorite tools for visualization, reporting, and advanced analytics. Examples include Looker, Power BI, Tableau, and MicroStrategy.

Work seamlessly with the Google Cloud Ecosystem

Action's close collaboration with Google Cloud's cloud architect and storage teams has resulted in a data warehouse experience that is easy to use, fully integrated, and one that maximizes the most of Google innovation.

Avalanche is deployed in the form of containers and micro services on the latest compute nodes of Google Kubernetes Engine (GKE), Google Cloud Dataproc, and Google Cloud Storage (GCS). Avalanche supports agile data processing by enabling the rapid ingestion of data into the data warehouse—or it can be used to query data stored in Google data lakes via external tables. You can use Google Cloud's Looker business intelligence and analytics platform to visualize your results.

Action is actively working with the Google Cloud team to deliver an Edge-to-Cloud experience via the Google Anthos framework. We are also working on full integration with core Google Cloud services such as Dataprep, Dataflow and Kubeflow to facilitate access to the real-time insights evolving organizations need.




Figure 3. Actian Avalanche is a high performance platform offering seamless native integration with hundreds of SaaS and enterprise applications

Avalanche with Google Modernizes Data Analytics

Avalanche delivers a modern analytics platform that supports a broad spectrum of users, from business analysts to data scientists. It excels in performing complex queries needed to enable diagnostic, predictive and prescriptive analytics.

Avalanche provides next-generation ML powered by high performance model training using vectorized processing. Its highly scalable atomic data analysis is exposed via Python, Java or C++ library which can be plugged into tools such as Jupyter notebook. Avalanche also provides an optimized native Spark integration and KNIME plugin that helps data scientists define pipelines for advanced analytics work.