

Actian Avalanche

Hybrid Cloud Data Warehouse

(AWS, Azure, Google Cloud, On-prem, VPC)

Highlights

Do more with less: Industry's best price-performance means our cloud cost meter runs 6x slower than alternative solutions. Our advantage gets even better with more users and data volume.

Future-proof with hybrid multi-cloud architecture: Avalanche can be deployed on-premises, virtual private cloud, AWS, Azure, and Google Cloud. Only one technology to learn, deploy and manage.

Integration is included: Avalanche comes with enterprise-class integration that includes over 200 application connectors so you can get up and running quickly

Pay only for what you use: Avalanche takes advantage of modern cloud innovations such as separation of storage and compute as well as elastic scalability

Lower administration costs with fully managed service: Reduce implementation time from weeks to minutes

Actian Avalanche is a fully managed hybrid cloud data warehouse service designed from the ground up to deliver high performance and scale across all dimensions – data volume, concurrent user, and query complexity – on commodity infrastructure. An MPP columnar SQL data warehouse, Avalanche is a true hybrid platform that can be deployed on-premises as well as on multiple clouds, including AWS, Azure, and Google Cloud, enabling you to migrate or offload applications and data to the cloud at your own pace.

“Actian packs a strong punch with Avalanche as a hybrid data warehouse offering capable of cloud, on-premises and/or hybrid deployments that also includes built-in data ingestion functionality while striking a favorable price-to-performance ratio.”

– Jim Curtis, Senior Analyst, 451 Research

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

Actian Avalanche delivers the best price-performance in the industry out-of-the-box without DBA tuning and optimization techniques. For the same cost as alternative solutions, you can benefit from substantially better performance or chose the same performance for significantly lower cost. For example, Avalanche provides up to 6x the price-performance advantage over Snowflake as measured by GigaOm's TPC-H industry standard benchmark 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 or thousands 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

Key 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




Figure 1. Avalanche is a high-performance platform that is optimized for compute, memory and storage, and the only platform built from the ground up for intelligent multi-cloud hybrid environments

Hybrid and Multi-Cloud Deployment Options

The same Avalanche data warehouse service can be delivered on-premises or in multiple clouds such as AWS, Azure and Google Cloud. With Avalanche, 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-premise, 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-premise 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



Figure 2. Avalanche offers a flexible consumption model.

Built for Real-Time Performance

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 freshest 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 Built-In Connectivity

Avalanche is the industry’s first and only 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 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. And 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 Power BI, Tableau, MicroStrategy and Looker.

The Avalanche Cloud Advantage

Avalanche delivers operational elasticity, enabling you to scale your data warehouse up and down to meet changing business needs. For example, Avalanche first uses its automatic indexing feature to load data into partitions equally distributed across the baseline nodes, say four. But you can increase to eight nodes for the business users during each day, and to 32 nodes on the last day of each month to handle peak demand. Every night, Avalanche would be ramped back down to four nodes.

Avalanche takes advantage of a logical connection between each compute node and storage where logical partitions get reassigned with no data redistribution or movement. Compare this to an “appliance in the cloud” approach, which relies on physical connections between compute and storage where the data must be redistributed physically to balance across nodes to take advantage of MPP.