ACTIAN™

Avalanche
Hybrid Cloud Data Warehouse

vs.

snowflake

Avalanche is designed for the most efficient use of resources on-premises and in the cloud

Actian Avalanche is a fully managed hybrid cloud data warehouse designed from the ground up to deliver unrivaled performance on commodity infrastructure. Initially developed for on-premises deployment, Avalanche was reimagined for the cloud. Avalanche delivers all of the features that one expects of a cloud data warehouse service such as elasticity and separation of compute from storage which enables the scaling of resources to meet changing business needs.

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."

Snowflake is built for the abundant resources of the cloud

Snowflake Cloud Data Warehouse was architected specifically for the cloud and does not have an on-premises equivalent. Its stated design goal was to deliver an elastic solution that enabled the independent scaling of compute and storage resources as a managed service and at an attractive price point.

But there is a down-side to their "built for the cloud" mission. The cloud provides limitless resources, so Snowflake is not designed to make efficient use of resources and cost often spirals out of control when expanded to enterprise scale. Performance also is not their focus, as evidenced by their use of slow storage that results in lengthy execution times when reading data from disk.

Avalanche delivers 6x superior price-performance over Snowflake. Substantially boost your performance or lower your cost: the choice is yours.

ACTIAN

Avalanche
Hybrid Cloud Data Warehouse

vs.

snowflake

Don't get caught with sticker shock that hurts your eyes—and your wallet

Both Avalanche and Snowflake are sold in terms of resource units. That is where the similarities end. Pricing may be similar, but the compute resource units are not created equal.

Due to differences in design goals, Avalanche completes 6x as many queries per resource unit as Snowflake. Avalanche also supports 8x more concurrent users than Snowflake – an important distinction in today's world that includes many business users and decision-makers who need access to data.

Together, this means more work done faster and at a lower cost.

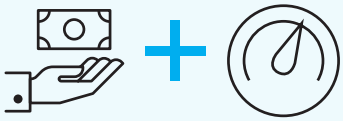

"Price and performance are critical points of interest when it comes to selecting an analytics platform…Our analysis reveals Avalanche to be the industry leader on this criterion."

GigaOm, "High Performance Cloud Data Warehouse Performance Testing," 2020

What causes Snowflake's runaway cost?

- **Slow performance:** Because resource efficiency wasn't a design goal for Snowflake, it takes a lot more resources to answer business questions and the results are produced slowly, which translates to higher cost.

  Avalanche's approach: *Avalanche is designed to deliver high performance. Its cloud cost meter runs 6x slower—meaning you can do a lot more with the same resources.*

- **Concurrent users:** If more than eight users need to access Snowflake concurrently, it will spin up a new compute warehouse with its autoscaling feature for each additional group of up to 8 users.

  Avalanche's approach: *Avalanche allows up to 64 concurrent users out-of-the-box, which keeps costs down when deploying at production scale. Remember that concurrent means queries running at exactly the same instance in time. Combined with Avalanche's superior performance, this translates into support for hundreds, even thousands, of active users.*

ACTIAN

Avalanche
Hybrid Cloud Data Warehouse

vs.

snowflake

Actian Avalanche Delivers the Best of Both Worlds – Fast Insights at Low Cost

GigaOm TPC-H Benchmark 30TB data, 5 concurrent users

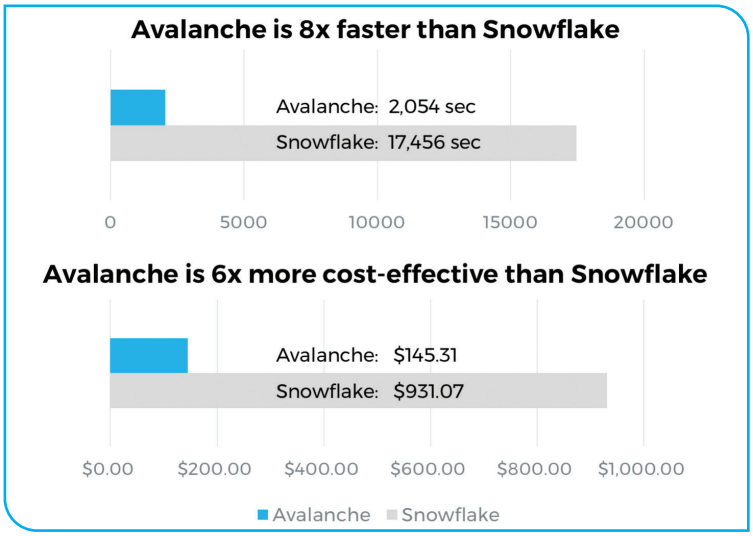

Figure 1. Avalanche outperforms Snowflake by 6x in the Oct 2020 GigaOm benchmark study. Its cost per query advantage increases as concurrency and query complexity rises. You can choose to improve performance and/or cost to meet your needs.

Actian Avalanche delivers the best price-performance in the industry because its highly efficient design takes advantage of the performance features in modern advanced CPUs to maximize compute, storage, and memory resources. Scale up or down as your needs change, but at a much lower and predictable cost than Snowflake.

Avalanche uses patented features and industry best practices:

- **Vector processing** – Operate on hundreds of tuples of data by exploiting SIMD support in x86 CPUs
- **CPU cache maximization** – Use private CPU core and caches as execution memory – 100x faster than RAM
- **Separation of compute and storage** – Scale compute resources to meet the needs of the business
- **Zero penalty updates** – Perform analytical queries as the data warehouse is being updated without noticeable impact
- **Pure columnar** – Pure columnar implementation not just at the storage layer
- **Advanced Compression** – Maximize the efficiency of decompression; but also deliver 4-6x compression ratio i.e. reduced disk footprint
- **MPP architecture** – Parallelize query execution within and across nodes to power through business workloads regardless of size and complexity
- **High Concurrency** – Support high volume of concurrent users, allowing up to 64 concurrent users out-of-the-box
- **Modern analytics** – Incorporate Hadoop, Kafka, streaming, mobile and IoT data, and connect to your favorite AI/ML applications

ACTIAN

Avalanche
Hybrid Cloud Data Warehouse

vs.

snowflake

Avalanche comes with native integration

Avalanche is the industry's first and only cloud data warehouse to offer integration capabilities natively built into the product. Avalanche provides connectors and templates for easily sourcing and moving data from SaaS applications to Avalanche data warehouse at scale—no special ETL is required. It also comes with Universal Connect patented technology that enables you to connect any data source or application.

Unique Advantages of the Avalanche Hybrid Model

The same Avalanche data warehouse service can be delivered on-premises or in multiple clouds such as AWS, Azure and Google Cloud. Snowflake does not provide an option to deploy on-premises. 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-premises, you will work with a single data model, consistent ETL integration, and have 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 as needed
- **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

Avalanche Hybrid Cloud Data Warehouse vs. snowflake

Category	Detail	Actian Avalanche	Snowflake
Data Store	Physical Storage Design	Columnar	Columnar
Scale Architecture	ACID	Yes	Yes
	CPU usage	Vectorized	Vectorized
	Indexing	Automatic storage indexes	Auto column indexes
		User created secondary indexes	No user defined indexes supported
	Performance with updates	Negligible drop (~10%)	~40% drop (10TB workload)
	Designed for concurrency	Yes – default to 64	No - defaults to 8
Cloud Architecture	Elasticity – Dynamic Scalability	Full on Azure (AWS and Google Cloud coming soon)	Full
	Elasticity - Turn On/Off Service	Yes	Yes
	Elasticity – Auto suspend	Yes	Yes
Hybrid Architecture	Federated Query	Federated query across Avalanche deployments	No
	Deployment Options	On-Prem, AWS, Azure, Google Cloud, VPC, Hadoop	AWS, Azure, Google Cloud
Data-types and functions	ANSI SQL support	Yes	Yes
	Stored procedure support	Yes	Yes
	External table support	Yes	Yes
	UDF support	SQL, JavaScript & Python	SQL, JavaScript & Java
	Data type support - JSON	Yes	Yes
Security	Data at rest encryption	All data encrypted by default; column level encryption available	All data encrypted
	Data masking	Yes	Yes
	SOC 2 Type II compliance	Yes	Yes
Data Ingestion	APIs and access method	CLI/JDBC/ODBC/.NET	CLI / JDBC/ODBC/.Net
	Data loading	Parallel and real-time (via Avalanche Connect)	Parallel and real-time (via SnowPipe)

ACTIAN™

www.actian.com | **Toll Free** +1 888 446 4737 | **Tel** +1 650 587 5500
2300 Geng Rd., Suite 150, Palo Alto, CA 94303