

Data Intensity and the Five Critical Components That Impact Application Performance

→ What is Data Intensity?

A measure of how rigorous data requirements are for a given application.

→ Why is it important?

High-volume, high-velocity data workloads are increasing at a fast rate. The more data intensive your application, the more difficult it's going to be to meet your SLAs.


→ How will it impact the performance of my application(s)?


A better understanding of data intensity and the five factors that contribute to it will help you determine the optimal data infrastructure you need to meet your requirements.

How Data Intensive is Your Application?

[Take Our 3-Minute Assessment to Find Out](#)

The SingleStore **Data-Intensity Index** is derived from five discrete variables:






Data Size

Modern databases must be able to effortlessly manage terabytes to petabytes of data generated by increasingly data-intensive applications.

Data sizes of **less than 1 TB** are lightweight analytics, or small-to-medium operational workloads.

Data sizes **1 – 10 TB** are very large for operational workloads, but medium-sized for analytics workloads.


Data sizes of **tens of TBs** or larger are typically big data analytics.



Data Ingest Speed

Data-intensive applications must be able to load – or ingest – data at very high rates, including hundreds of thousands or millions of rows per second.

Many databases can handle a **few hundred** to a **few thousand** rows per second. Workloads with thousands to millions of rows per second or higher are typical in data-intensive applications. This is where many data infrastructure systems struggle.



Query Complexity


Data-intensive applications must be able to handle both simple and complex queries equally, effectively and fast.

Complex data queries join multiple data sets with multiple filters and often include aggregation and other calculations. The more joins and subqueries, the greater the complexity.

Low Most operational databases (and some data warehouses) can handle low complexity queries in fractions of a second.

Med There are several databases that can manage medium complexity with simple joins, distributed databases perform scale better.

High Highly complex queries on a large data set require the use of distributed systems to run.



Query Latency

Query latency is the amount of time it takes a database to execute a query and receive a result. Data-intensive applications often have strict SLAs, or requirements around query latency.

Query latency requirements can vary dramatically depending on the use case.

- | | |
|-------------------------|--|
| Minutes/1-10 Sec | Daily reports can take minutes or hours to complete. Single-node databases or big data open-source databases are sufficient. |
| 100 ms - 1 Sec | Applications delivering interactive experiences need queries to complete in less than half a second. |
| 10 - 100/0-10ms | For operationalized Machine Learning use cases, queries need to be < 100 ms. |



Concurrency

Data-intensive applications often need to support a large number of users or concurrent queries while managing low-latency queries simultaneously.

Data infrastructures need the capacity to meet demand; however, most legacy data warehouses were built for generating reports, not handling concurrency. While many operational databases have improved their ability to handle concurrency, they do so only for simple queries.

- + **Applications** with a limited set of users (e.g., departmental applications, low traffic web applications) generate few queries.
- + Fast-growing **SaaS applications** with high numbers of concurrent users generate tens to hundreds or thousands of queries at peak times.

Data-Intensive Applications Need a Modern Data Infrastructure



Accelerate the Performance of Your Data-Intensive Applications

The first modern, relational database for cloud and on-premises, SingleStore enables organizations to simplify their data architectures while delivering speed, scale and immediate insights for modern apps and analytical systems.

SingleStore is a cloud-native database built with speed and scale to power data-intensive applications.


With SingleStore, you can:



+ Deliver real-time analytics by leveraging architecture that provides fast ingest and query performance with high concurrency.



+ Supercharge modern SaaS apps with a ludicrously fast database infrastructure that works on any data, anywhere.



+ Scale your database and ensure that you have the backend stability you need to successfully deliver your SaaS product, every time.



+ Consolidate databases with a single database that unifies your data and supports multiple, massive workloads concurrently.

How Data Intensive is Your Application?

Take Our Free 3-Minute Assessment Now

Find out how ready your data infrastructure is to deliver the user experiences you need to power your business.

Start Now →

Unleash the Power and Performance of Your Data with SingleStore

To learn how SingleStore can help you deliver low latency, sub-millisecond performance at nearly unlimited scale, visit singlestore.com

