




10 Things to Know About InfluxDB Clustered





01 WHY READ THIS?

Large enterprises who want real-time monitoring and control of their mission-critical systems and applications, are facing enormous challenges managing the time series data that gets generated in the process due to its unique characteristics, including the volume, frequency, variety, and cardinality of the data. InfluxDB is purpose-built to take the complexity out of time series data management including easily collecting, storing, and analyzing the massive volumes of time-stamped data generated. InfluxData addresses these challenges with three flavors of its distributed, scalable architecture, Cloud Serverless, Cloud Dedicated and Clustered. Let's take a closer look at Clustered.

InfluxDB Clustered is the solution designed and built for organizations that not only face the challenges of volume, frequency, variety, and cardinality above but also have stringent requirements around data residency and privacy that require them to run InfluxDB on-premises or in self-managed private clouds.

This guide explores the benefits of InfluxDB Clustered and enables you to make an informed decision on whether InfluxDB Clustered is the right solution for you and your business.

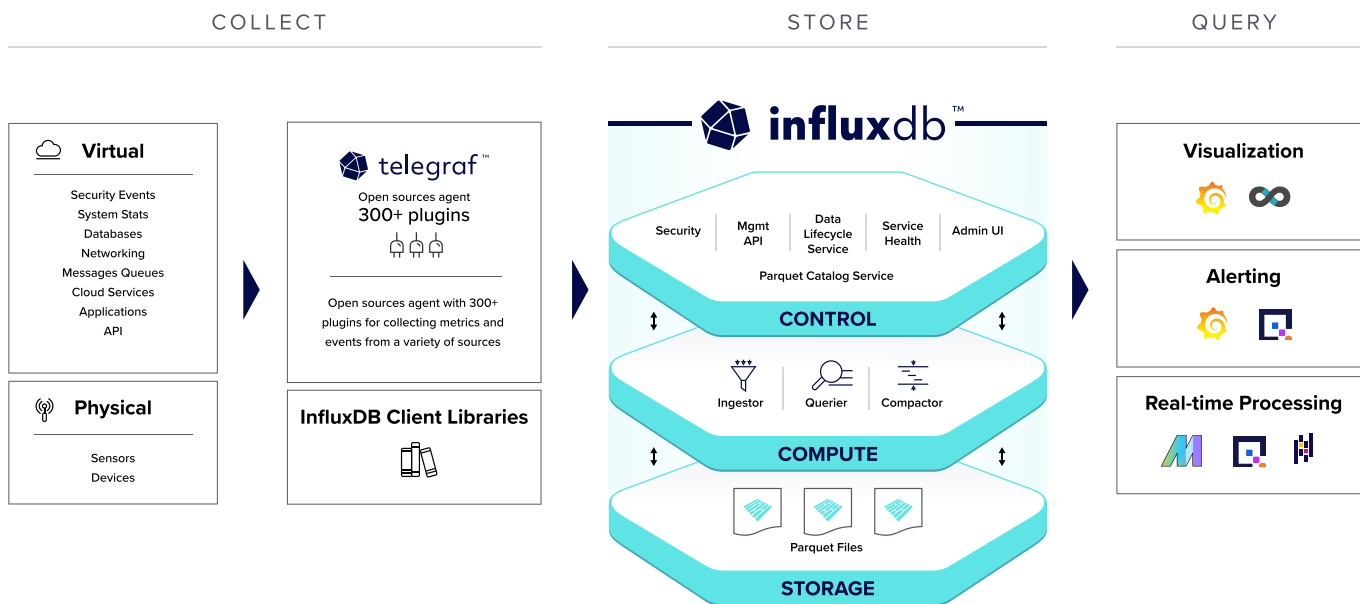


02

WHAT IS INFLUXDB CLUSTERED?

influxdb™ clustered

InfluxDB Clustered is a distributed, scalable and highly performant edition of the InfluxDB 3.0 engine built on the Apache FDAP stack. It is designed to be deployed and managed as a collection of Kubernetes-based containers and services. This enables you to run and operate InfluxDB at scale within on-premise or self-managed private cloud environments and allows you to customize resource allocations to meet your specific workload demands.



InfluxDB Clustered is architected to provide the highest levels of availability, performance, security, and cost optimization. It is also the next-generation evolution of our InfluxDB Enterprise offering.

03

WHY INFLUXDB 3.0?

(THE ENGINE POWERING INFLUXDB CLUSTERED)

Purpose-Built for Time Series Data

Designed specifically for time series data, InfluxDB ensures efficient data ingestion and compression, ideal for sensor-generated data in manufacturing.

High Performance Data Handling

InfluxDB handles large volumes of time series data continuously with high throughput. InfluxDB's columnar analytics based on [Apache Arrow](#) and in-memory data handling unlocks real-time querying

Designed for Unlimited Cardinality

Collect granular data to gain richer, more in-depth insights. InfluxDB has no storage or cardinality limits, so you can ingest billions of data points in real-time.

Reduced Total Cost of Ownership

Best-in-class data compression with [Apache Parquet](#) and persistence on commoditized storage means InfluxDB saves more data using less space, reducing the cost of keeping your data.

InfluxDB 3.0 vs OSS

45x

Better ingest performance

100x

Faster queries for high cardinality data

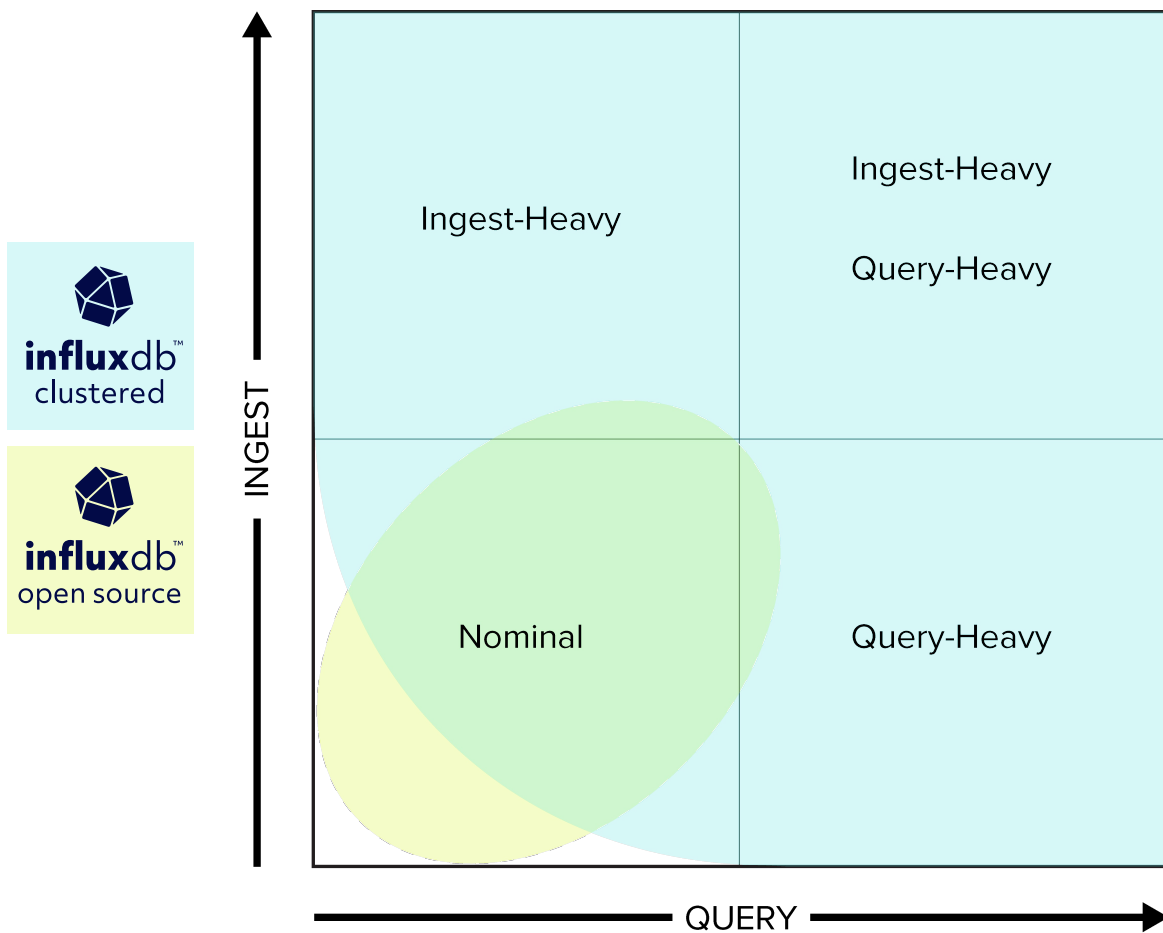
04 WHY INFLUXDB CLUSTERED?

Self-Managed

- Meet specific data residency and privacy requirements.
- Have complete control over the underlying infrastructure.

Custom Scaling to Specific Workloads

- Choose the cluster configuration that is customized to your data workload.
e.g. add higher number of ingestors for ingest-heavy workload type.

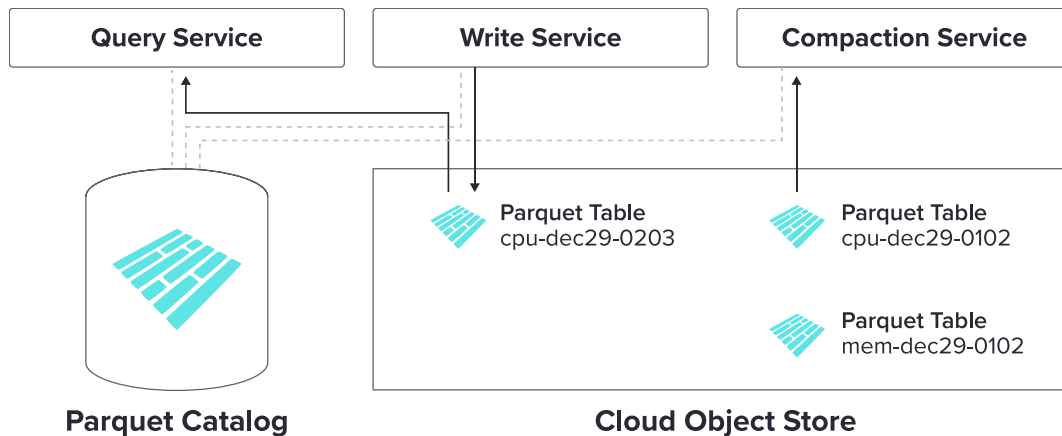


05

HOW DOES INFLUXDB CLUSTERED ARCHITECTURALLY COMPARE TO INFLUXDB OPEN SOURCE?

InfluxDB Clustered: 10x better cost optimization and operational efficiency

InfluxDB Clustered is built as a distributed, cloud native architecture with the storage and compute tiers being separate, allowing each to scale independently based on demand. By isolating compute and storage, InfluxDB Clustered is able to reduce resource contention, improving overall system performance and reliability. Compute nodes can be added independently to ingestor or querier services to scale out to handle increased loads without affecting storage performance.



InfluxDB Open Source: Storage and compute are bundled together

In contrast, in a self-managed InfluxDB OSS instance running on a VM, the storage and compute resources are bundled together, often residing on the same host. While this provides some advantage in terms of reducing latency associated with data transfer between storage and processing units, it quickly presents challenges with real world production workloads. For example, intensive compute operations (ingest heavy or query heavy operations) can impact the performance of storage operations and vice versa, leading to potential bottlenecks and degraded performance. The only alternative is to add CPU and memory which in turn increases overall costs.

06

WHAT SECURITY FEATURES ARE AVAILABLE IN INFLUXDB CLUSTERED?

InfluxDB Clustered is designed with a strong focus on security, suited for self-managed environments where control and compliance are priorities. Here are the enterprise-grade security features available in InfluxDB Clustered:



End-to-end Encryption:

Clustered can be configured for complete encryption of data at-rest and in-flight.



Single Sign-On (SSO):

Enhances the ease of use and security by allowing users to log in with a single ID across multiple systems.



Private Networking Options:

Offers capabilities to configure networks that are isolated from public access, enhancing data protection.



Attribute-Based Access Control (ABAC):

Provides an access control mechanism that grants permission based on policies applied to user attributes.



Support for Fully Air-Gapped Deployments:

Ensures that systems can be run completely isolated from the internet, which is critical for highly secure environments.

These features are part of the broader suite of enterprise-grade capabilities intended to meet various regulatory requirements on data privacy and residency while providing the flexibility to run on-premises or in private and public cloud environments.

07

HOW IS INFLUXDB CLUSTERED PRICED?

Annual Contract

Capacity-based Pricing

- Priced by CPU / RAM configuration
- Flexible scaling to support unplanned growth
- Round-the-clock with global coverage across 6 continents

Optional Add-ons

Enterprise Requirements

- Enterprise SSO (Single Sign-On)
- Non-production environments

08


WHAT ARE THE KEY REQUIREMENTS FOR RUNNING AND MANAGING INFLUXDB CLUSTERED?



Built to be deployed and managed using Kubernetes. Helm chart is available to streamline the deployment and upgrade process.



Requires S3 compatible storage for persisting data in Parquet format. Supported storage includes MinIO, Ceph.



Dedicated PostgreSQL or compatible database instance for data catalog service.



OAuth 2.0 provider for Device Authorization Flow. Supported providers include Microsoft Entra ID, Keycloak, AuthO.

09

IS INFLUXDB CLUSTERED THE RIGHT SOLUTION FOR ME?

I have expertise managing Kubernetes based applications. I need a highly performant InfluxDB solution for on-premise or private cloud environment.



influxdata.com/products/influxdb-clustered

I do not have expertise managing Kubernetes environments



influxdata.com/products/influxdb-cloud/dedicated

I do not have expertise managing Kubernetes environments and I need an on-premise InfluxDB solution



influxdata.com/products/start-on-enterprise

I do not have high cardinality data and I need an on-premise InfluxDB solution



influxdata.com/products/start-on-enterprise

10

WHAT ADDITIONAL RESOURCES DO YOU HAVE FOR INFLUXDB CLUSTERED?

Quickstart guide to InfluxDB Clustered Document

InfluxDB Clustered is deployed and managed using Kubernetes. This multi-page guide walks through setting up prerequisites and configuring your InfluxDB cluster deployment.

[View Quickstart Guide](#) 

Sign up for self-paced course in InfluxDB University Course

This course covers the fundamentals of InfluxDB 3.0. It walks through the various components of the platform, and introduces the InfluxDB data model and how to map raw data onto it.

[View Course](#) 

Check out InfluxDB's storage engine internals Document

The InfluxDB v3 storage engine is a real-time, columnar database optimized for time series data built in Rust on top of Apache Arrow and DataFusion.

[Open Document](#) 

Schema best practices Course

Just because you don't need to define a schema before using InfluxDB, doesn't mean you can't optimize your schema for your data and use case.

[Open Course](#) 



TAKE A TEST DRIVE

INFLUXDB CLUSTERED | [REQUEST A PROOF OF CONCEPT](#)

