



MongoDB is an open-source cross-platform document-oriented database program. Classified as a NoSQL database program, MongoDB uses JSON-like documents with optional schemas.



Apache Kafka is a framework implementation of a software bus using stream processing. The project aims to provide a unified, high-throughput, low-latency platform for handling real-time data feeds.



A relational database (RDBMS) is a digital database based on the relational model of data. Many relational database systems have an option of using the SQL for querying and maintaining the database.



Axon Server complements the Axon Framework by providing robust, unified messaging services such as event storage, clustering, load-balancing, message routing, service discovery, monitoring, and zero-configuration.



Axon Server with Enterprise functionality contains all of the features of the Axon Server Open Source. Additionally, it contains features to make it suitable for use in mission-critical, enterprise use cases.

Messaging

<p>Event-driven architecture Traditional approaches gather data in discrete areas to be called on only when needed. An event-driven approach can be more efficient by immediately reacting to new published information.</p>	No, MongoDB is a data store, not a messaging component.	Yes	No, a RDBMS is a data store, not a messaging component.	Axon Server provides a purpose-built event bus and event store, used to stream and store events between a multitude of applications.	Axon Server provides a built-in event bus and event store, used to stream and store events between a multitude of applications. On top of that, Axon Server provides automatic fail-over through its Raft-based clustering, ensuring events can always flow.
<p>Event streaming Event streams can create notifications about new events, and by subscribing to the stream, business can react to these changes in real-time. Event Streaming requires the ability to push (rather than pull) messages from the sender to the receiver(s).</p>	No, MongoDB is a data store, not a messaging component.	Yes, but the client is required to pull the messages. This can be solved by a dedicated client library.	No, a RDBMS is a data store, not a messaging component.	Axon Server provides a purpose-built Event Bus and Event Store, used to stream and store events between a multitude of applications.	Axon Server provides a built-in Event Bus and Event Store, used to stream and store events between a multitude of applications. On top of that, Axon Server provides automatic fail-over through its Raft-based clustering, ensuring events can always flow.
<p>Commands, Queries, and Events Different types of messages require different communication patterns. It is a common misunderstanding that all communication can be served well using asynchronous, one-way communication. Commands, which request changes in state, require a reply to indicate acceptance. Queries are requests that could benefit from ScatterGather or Streaming approaches, while Events are one-way and generally allow for Pub-Sub.</p>	No, MongoDB is a data store, not a messaging component.	No, Kafka is a purpose-built Message (Event) streaming bus and does not provide for other communication patterns.	No, a RDBMS is a data store, not a messaging component.	Yes, Axon Server can understand the different message-handler discovery and message routing patterns, and doesn't need boiler-plate code to support the different message types.	Yes, Axon Server can understand the different message-handler discovery and message routing patterns, and doesn't need boiler-plate code to support the different message types.

Event store and event sourcing

<p>Audit trail An event-driven system with event storage stores your data as a series of immutable events over time, providing one of the strongest audit trail options available.</p>	No, a database could be used to track changes, but falls short as a reliable audit trail.	Not reliably, since Kafka does not support the Event Sourcing features to make the event store the reliable source of truth.	Some RDBMSes allow you to keep and query a transaction log, which stores any committed changes, but they are never meant to be kept indefinitely.	Yes, Axon Server provides an event store, showing all auditable changes from your system.	Yes, Axon Server provides an event store, showing all auditable changes from your system.
<p>Time travel Since all state changes are kept as events in the event store, it is possible to move systems backward and forwards in time which is extremely valuable for audits and "what if" analyses.</p>	No. Add-ons can keep track of the past state of databases, but not the changes.	Not reliably, since the Kafka event store does not offer a reliable source of truth and the retention time of messages effectively prevents travelling back further than that time.	No, going back in time is only supported as a disaster recovery procedure and will lose you any changes after that point.	Yes, Axon Server provides an Event Store, showing all changes in your system and allowing you to move to any point in history.	Yes, Axon Server provides an Event Store, showing all changes in your system and allowing you to move to any point in history.
<p>Event sourcing The storage mechanism supports all essential features to implement event sourcing reliably. For example, it is optimized for appending events and streaming them. It must also be able to efficiently retrieve events for a specific aggregate and be able to append them using optimistic locking.</p>	MongoDB can be used as an Event Store, but the performance will degrade as it grows over time.	Kafka can be used as the central event bus for a CQRS implementation, but comes with limitations in practice. As it doesn't support Event Sourcing, there is always a challenge to keep the command model's state in sync with the events published on Kafka	RDBMS can be used as an Event Store, but the performance will degrade as it grows over time.	Yes, Axon Server is purpose built for the append-only pattern of the Event Store and optimized for event replays.	Yes, Axon Server is purpose built for the append-only pattern of the Event Store and optimized for event replays. On top of that, Axon Server with Enterprise functionality does this with Raft-based fault tolerance, ensuring your commands, events and queries always flow.
<p>Backups Event streams are fundamentally just logs with strong backup and recovery characteristics. Writing just the core "source of record" data to the event stream enables the rebuilding of downstream projections.</p>	MongoDB supports creating backups, but it isn't aware of the fact that this data is append-only. Backups will therefore be increasingly large and will overlap with any previous backups created.	Kafka doesn't offer backup tools out of the box. It is possible to use external tools.	Yes, all RDBMSes have support for online or offline backups.	Yes, Axon Server is a single source of truth for any model within an application, at any point in time.	Yes, Axon Server is a single source of truth for any model within an application, at any point in time.
<p>Event replay Event streams are fundamentally logs with strong backup and recovery characteristics. The current state of an aggregate can be reconstructed using a replay of events, optionally using snapshots to reduce the required replay interval. Rebuilds of query models can use full replays if they start empty, or from a point in time for a partial rebuild.</p>	Replays are possible only if the appropriate indices are available.	Yes, given the data retention configured. Replays are possible within the configured timeframe. Setting the timeframe to "indefinite" is possible, but not recommended.	Replays are possible only if the appropriate indices are available.	Yes, Axon Server is a single source of truth for any model within an application, at any point in time.	Yes, Axon Server is a single source of truth for any model within an application, at any point in time.
<p>Multi-context architecture An arbitrary number of logical contexts on a single cluster. This allows for multitenancy, access control, and differentiated retention policies.</p>	Yes	Possible by using topics	Yes, most RDBMS will support multiple schemas.	No, Axon Server only supports a single context.	Yes, Axon Server with Enterprise functionality enables multi-context support out of the box. This means messages from different contexts or tenants have a dedicated message bus and storage location.
<p>Data protection and GDPR compliance Field-level encryption and cryptographic erasure.</p>	Yes, using the AxonIQ Data Protection module in Axon Framework applications.	Yes, using the AxonIQ Data Protection module in Axon Framework applications.	Yes, using the AxonIQ Data Protection module in Axon Framework applications.	Yes, using the AxonIQ Data Protection module.	Yes, the AxonIQ Data Protection module allows personally identifiable information in any message to be encrypted. By removing the key the encrypted information is no longer readable (also called "cryptographic erasure").

High availability and performance

<p>Asynchronous first Event sourced systems strive for the minimum amount of synchronous interaction; consistency boundaries are consciously chosen so that business requirements are met, and everything else is eventually consistent. This results in responsive, high performing, scalable systems.</p>	While clients may be asynchronous, the database itself doesn't stream updates in a user-friendly way. This means clients need to resort to (synchronous) polling to retrieve new data.	Appending and reading events is done asynchronously, but there are no consistency checks when appending data.	While clients may be asynchronous, the database itself doesn't stream updates in a user-friendly way. This means clients need to resort to (synchronous) polling to retrieve new data.	Axon Server is a message routing solution at its core. It complies to the routing requirements of commands, events, and queries, enabling completely asynchronous communication between your applications.	Axon Server is a message routing solution at its core. It complies to the routing requirements of commands, events, and queries, enabling completely asynchronous communication between your applications.
<p>Constant performance Remain highly performant, regardless of the data.</p>	No, indices required for event sourcing will grow and become slower over time.	Only for full sequential reads of the events. When using Kafka for event sourcing, the performance degrades quickly as more data is added.	No, indices required for event sourcing will grow and become slower over time.	Yes, the Axon Server event store is purpose-built to support the append-only model, using segmented file-stores and sparse indices to facilitate back-in-time searches.	Yes, the Axon Server event store is purpose-built to support the append-only model, using segmented file-stores and sparse indices to facilitate back-in-time searches.
<p>High availability clustering The ability of the system to continue functioning properly regardless of the failure of some components</p>	Yes	Yes	Yes, provided you have configured the database in HA mode. This may require additional components on client or server side.	No	Yes, using a distributed consensus algorithm.
<p>Automatic geographical clustering backup The ability of the system to provide backups in other regions, so the data is still available in the case of a disastrous loss of the entire region.</p>	Yes	Yes	No, this needs to be configured outside of the database, or using a managed service that adds this functionality.	No, unless you use periodic backups of the filesystem and copy those to another region. This must be configured outside of Axon Server.	Yes, because Axon Server can be deployed between several zones or regions. A node (or group of nodes) in a different region can be configured to receive a copy of the event store, ensuring storage of the backup.
<p>No data loss guarantee The guarantee that at least one backup node contains all the committed events.</p>	Yes	Yes	Yes, provided you have configured the database in HA mode. This may require additional components on client or server side.	No, because you only have a single copy of the data.	Yes, because the distributed consensus algorithm used by Axon Server ensures the majority of nodes contains the data at all times.
<p>Fast manual and assisted disaster recovery The ability to quickly rebuild an alternate group of primary nodes, in case of complete loss of the main region.</p>	Yes	No, separate plugins are required to assist in disaster recovery.	Yes	No	Yes, by using backup nodes in a different region. You can reconfigure these to become primary nodes so they start at the same point in time as the old one, ensuring quick recovery.
<p>Tiered storage The ability to use faster storage for data more often used, while keeping a full copy on slower storage.</p>	No	Kafka supports partitioning and data retention policies, but doesn't support storing older data in separate tiers.	Yes, using partitioning. This requires the chosen product to support partitioning on age of the data, and will migrate/clean data as it ages.	No, only a single Event Store is provided, with no automated partitioning/cleaning of data.	Yes, Axon Server allows for the addition of secondary nodes to keep the full event store, where the primary nodes will remove older events. This significantly reduces the size of the event store on primary nodes.