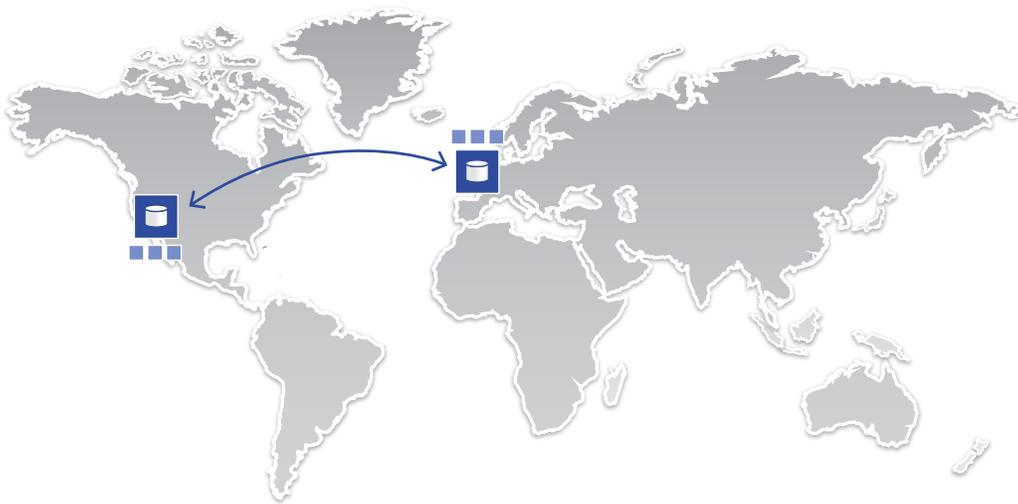


Flexible Replication

Getting data when and where you need it can be challenging—and even risky or impossible. Global organizations often run into latency issues trying to replicate data from centralized data centers to remote areas. And, for military operations in particular, replicating data to the field is often problematic because of cybersecurity concerns and technical roadblocks created by adversaries armed with signal jamming equipment.

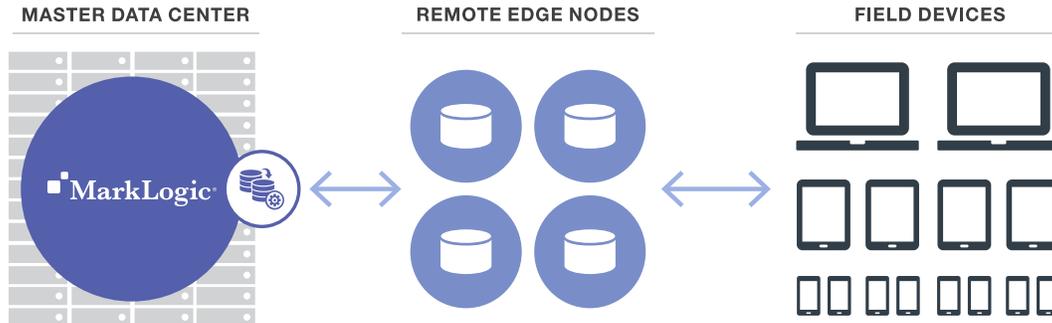
Flexible Replication solves these problems by allowing you to easily distribute portions of your data without worrying about security vulnerabilities or disconnected, intermittent, and latent networks.



Agility Around the Globe

Flexible replication gives you the ability to leverage the full benefits of a distributed database platform, distributing your data where you need it around the world while maintaining all of the enterprise features of MarkLogic®. MarkLogic has full database replication, which organizations use for disaster recovery. But, flexible replication is an ideal choice when organizations need the capability to:

- Enable field agents and remote workers to access data in austere environments
- Replicate filtered data from a centralized database to thousands of edge nodes around the world
- Collect data on edge nodes and have it automatically added to the main database when reconnected
- Provide improved performance by conducting queries and updates locally rather than over a network
- Address issues with disconnected, intermittent, and latent networks, often using radio or satellite
- Prioritize security and maintain the ability to use Role Based Access Control (RBAC)



Flexible and Robust Replication Model

- **Master-slave architecture** – MarkLogic clusters generally employ a shared-nothing architecture but flexible replication employs a master-slave architecture with centralized master nodes and distributed edge nodes that act as slaves. With conflicts, masters win, though MarkLogic can be customized so that slaves win
- **Document-based** – MarkLogic stores data as documents, which is the unit of replication. Any time a document changes it is transferred using an HTTP or HTTPS protocol. Documents can be pushed to the replica or pulled by the replica
- **Asynchronous** – Replication occurs after Master database transactions commit. This means that flexible replication is non-transactional, but this feature helps support the data needs of detached users with intermittent or latent networks
- **Query-based** – Use queries to specify data to be replicated based on a collection, URI path, or serialized query—whichever method you want to use to specify the data you want to replicate
- **Transformation filters** – Use filters to modify a document’s content, URI, properties, collections, permissions, or more as it is replicated. You can split a single document into multiple documents, or transform the content from one schema to another
- **Entity enrichment** – Send data to a data layer where it is enriched—including tagging and highlighting people, places, modes of transportation, equipment, and events—and then send the data back out to the field operators
- **Safe updates** – Updates are staged at higher echelons and are initiated by detached users only when they have the time and place to make safe updates
- **Transitive replication** – Replicate data through one node to others, allowing for easy tiered distribution across geographies

About MarkLogic

For more than a decade, MarkLogic has delivered a powerful, agile, and trusted Enterprise NoSQL database platform that enables organizations to turn all data into valuable and actionable information.

For more information, please visit www.marklogic.com.