

DDT Performance

Version 1.0.0

Sep.13,2024

Table of Contents

| | |
|--------------------------------------|----------|
| DDT Performance | 3 |
| Introduction | 3 |
| Synchronization Modes | 3 |
| Full Synchronization: | 3 |
| Real-Time Synchronization: | 4 |
| Full + Real-Time Synchronization: | 4 |
| Full + Incremental Synchronization: | 4 |
| System Requirements and Dependencies | 5 |
| Minimum System Requirements: | 5 |
| Supported Environments: | 5 |
| Performance Testing | 5 |
| Throughput: | 5 |
| Latency: | 5 |
| Resource Usage: | 6 |
| Bandwidth Efficiency: | 6 |
| Monitoring and Data Validation | 6 |
| Monitoring: | 6 |
| Data Validation: | 7 |
| Conclusion | 7 |

DDT Performance

Introduction

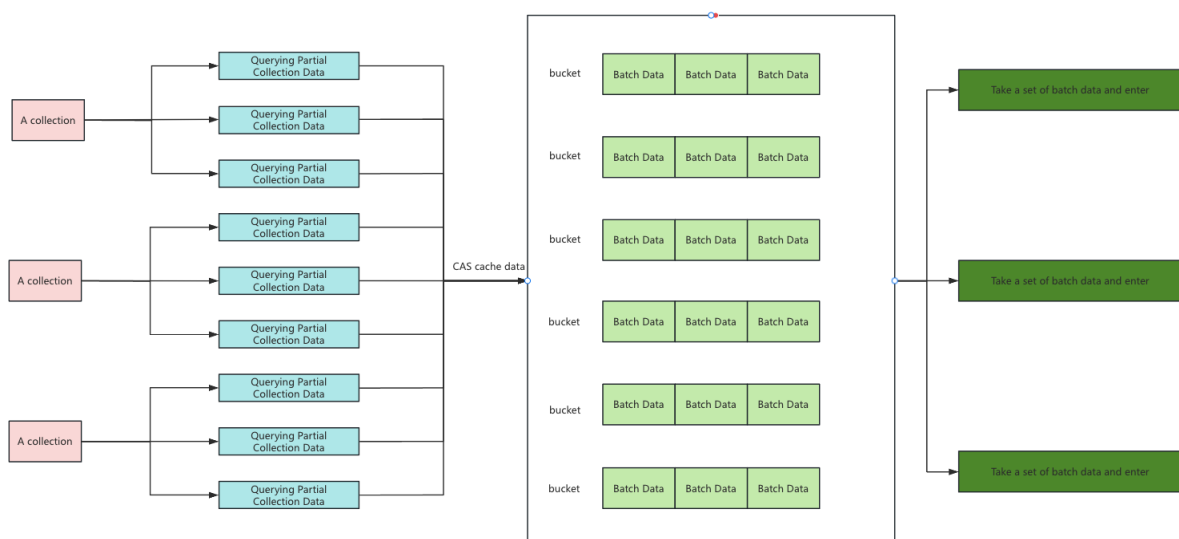
Document Data Transfer (DDT) is a cutting-edge data synchronization tool designed to meet the diverse and demanding needs of modern enterprises. As businesses increasingly rely on distributed systems and databases, the need for efficient, reliable, and scalable data synchronization solutions has never been more critical. DDT addresses these challenges with its versatile synchronization capabilities, supporting multiple synchronization modes tailored to different use cases and system environments.

Synchronization Modes

DDT offers four distinct synchronization modes, each designed to handle specific data transfer scenarios:

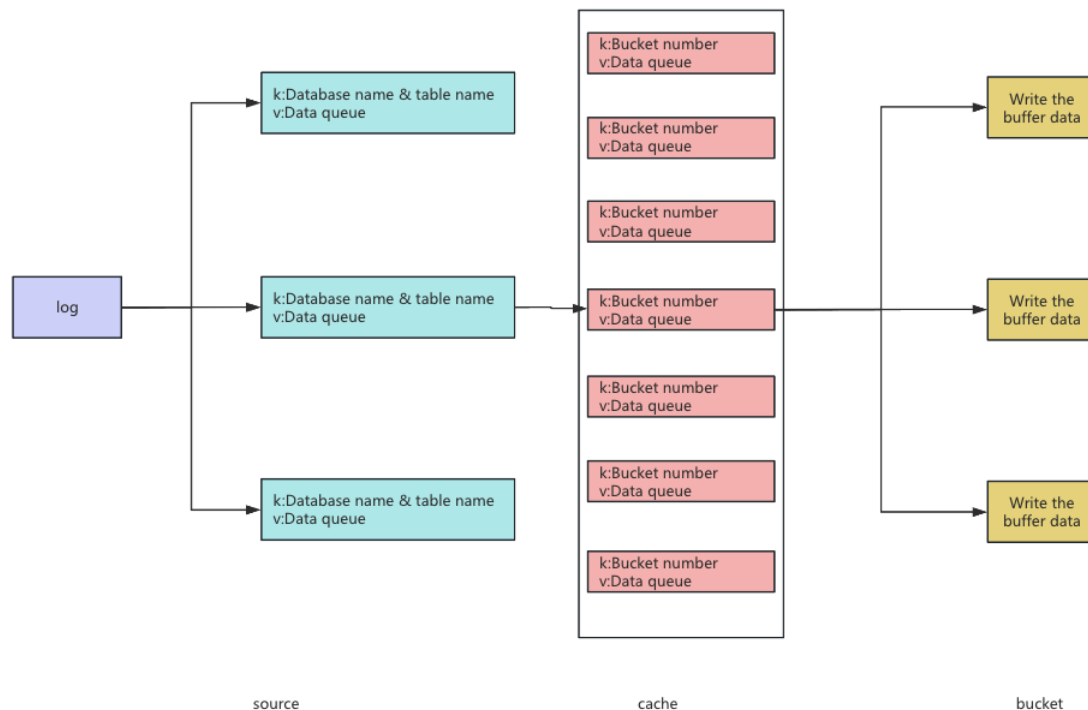
Full Synchronization:

This mode performs a complete data transfer from the source to the target, ensuring that all data is fully synchronized. It is ideal for initial data migration or when a comprehensive update of the target dataset is required. Use Case: Suitable for large-scale data migrations, backups, or when there is a need to synchronize an entire dataset.



Real-Time Synchronization:

Designed for scenarios requiring minimal latency, this mode continuously syncs data from the source to the target as changes occur. Real-Time Synchronization is essential for applications where up-to-the-minute data consistency is crucial. Use Case: Ideal for real-time applications such as financial systems, live data feeds, and any environment where data freshness is paramount.



Full + Real-Time Synchronization:

Combining the strengths of both Full and Real-Time Synchronization, this mode first performs a complete data transfer and then continues to synchronize any new changes in real-time. This ensures that the target system is always up-to-date, even as new data is generated. Use Case: Best for systems that need an initial data load followed by continuous updates, such as when setting up a new system that needs to stay in sync with an active database.

Full + Incremental Synchronization:

This mode begins with a full synchronization and then performs periodic incremental updates. It is particularly useful when the source data changes frequently but real-time synchronization is not required.

Use Case: Ideal for environments where data is updated in batches or where bandwidth and resource usage need to be optimized.

System Requirements and Dependencies

DDT is engineered to operate seamlessly across various system environments, making it a highly adaptable tool for different IT infrastructures. The tool is built on OpenJDK 11, ensuring compatibility with a wide range of operating systems including Windows, Linux, and macOS. This dependency on OpenJDK 11 allows DDT to leverage the latest advancements in Java, providing robust performance and security.

Minimum System Requirements:

- CPU: Quad-core processor or better.
- RAM: 8 GB (16 GB recommended for large-scale operations).
- Storage: SSD with at least 100 GB of free space.
- Network: High-speed internet connection or intranet with low latency for real-time synchronization.

Supported Environments:

- Operating Systems: Windows, Linux, macOS.
- Databases: DDT is database-agnostic and supports synchronization across NoSQL and SQL databases (e.g., MongoDB, Couchbase, PostgreSQL, MySQL).

Performance Testing

Performance is a critical aspect of any synchronization tool, and DDT has been extensively tested to ensure it meets the highest standards. The tool is designed to handle large datasets efficiently, even in resource-constrained environments.

Throughput:

DDT can achieve a throughput of over **500,000 documents per minute** in high-performance environments, depending on the network and hardware specifications.

Latency:

In real-time synchronization, DDT has been optimized to ensure that data is synchronized with minimal latency, with typical delays being under **100 milliseconds** for most use cases.

Resource Usage:

DDT's architecture is designed to minimize CPU and memory usage during synchronization. The tool dynamically adjusts its resource consumption based on the size and complexity of the dataset, making it scalable across different hardware setups.

Bandwidth Efficiency:

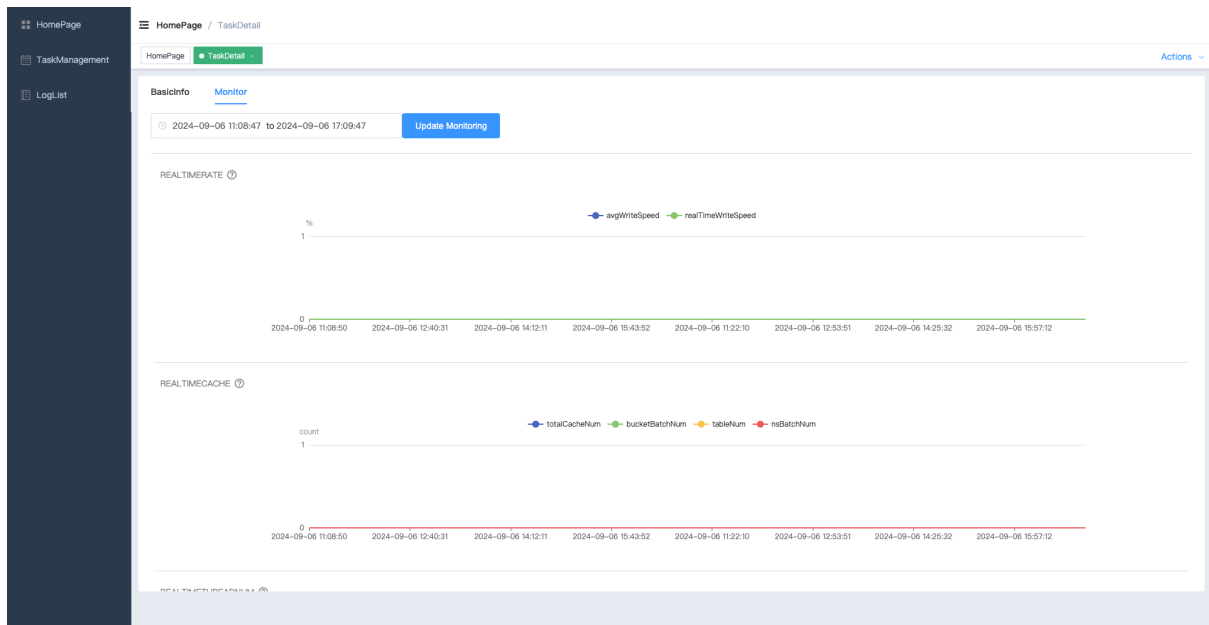
In incremental synchronization mode, DDT only transfers the changes made to the dataset, significantly reducing the amount of data transmitted over the network. This optimization makes it ideal for environments with limited bandwidth.

Monitoring and Data Validation

One of the key features that set DDT apart is its comprehensive monitoring and data validation scripts. These tools are essential for ensuring the accuracy and integrity of the synchronized data, providing users with the confidence that their data transfers are both complete and correct.

Monitoring:

DDT includes a powerful monitoring tool that enables users to track the progress of their synchronization tasks in real time. The monitoring interface displays crucial information such as the start and end times of tasks, the current status of each synchronization process, and the URLs of the source and target endpoints. This feature is invaluable for managing and optimizing multiple synchronization tasks across complex environments.



Data Validation:

After the synchronization process is complete, and assuming that no new data is being written to the source, users can utilize DDT's data validation script to verify the accuracy of the transfer. This script checks that all data has been correctly synchronized, helping to prevent discrepancies between the source and target datasets.

Conclusion

DDT is more than just a data synchronization tool—it is a comprehensive solution designed to meet the evolving needs of today's data-driven enterprises. Whether the requirement is for full data migration, real-time updates, or a combination of both, DDT provides the flexibility, reliability, and performance needed to keep your data synchronized across systems.