

TECHNICAL DATA SHEET

Stelo is an enterprise-class tool that dynamically delivers data from anywhere to anywhere for analysis, reporting and prediction or for managing business operations, B2B interactions and supply chains. Stelo moves data easily among your core relational and No-SQL databases in real-time across firewalls, to other teams, or to the cloud.

SET IT AND FORGET IT

Simple installation with GUI interfaces, configuration wizards, and advanced tools make product setup and operation straightforward, with no programming needed. Once running, Stelo reliably operates in the background without needing dedicated engineering support to maintain and manage. A single Stelo license is needed for source and destination databases, regardless of how many tables are replicated, delivering extremely low TCO.

SPEEDY AND AGILE

The Stelo solution provides reliable, high-speed, affordable replication for any relational database accessible via ODBC and non-relational databases via Stelo Spark Connector or KAFKA/AVRO. Stelo leverages native data loading functions, and exploits multithreaded processing to provide fast, reliable performance for replicating multiple tables concurrently.

POWERFUL AND SECURE

Safe data transformation on your own terms Stelo gives you greater control by allowing you to choose replication of the data you want with both vertical and selected horizontal partitioning. Replace existing tables or merge on a row-by-row basis. Stelo also supports derived destination columns and data transformations, including automatic and customized conversion and selection of dates, views, indexes, and constraints. While flexible, Stelo fully enforces the enterprise security mechanisms by integrating with both SQL Server/Windows security and the native security of your Db2 system.

EXPERT SUPPORT

- Extensive diagnostic expertise and support for related systems Supported by experts who know your system, not generalists who may not understand data replication
- Mission critical support you can count on to keep your systems running and data available

FLEXIBLE LICENSING

A single instance can support multiple sources and destinations without additional licensing. The Stelo model is independent of the number of cores of either the source or destination- so you only pay for the capacity required to support the transaction volume. Development/QA and Disaster Recovery (DR) pricing discounts are also provided.

AVOID VENDOR LOCK-IN

Stelo takes full advantage of open standards such as DRDA, SQL, ODBC, and JDBC. to maximize compatibility and interoperability within an enterprise network. We are an active member of The Open Group software industry consortium, which was responsible for the adoption of DRDA as an industry standard for database interoperability.

THIN SERVER, ULTRA-LOW FOOTPRINT

Database-neutral solution at less than 1% resource consumption Stelo can process millions of changes per hour at a near zero footprint, avoiding any burden on your production or host databases. With no need to install or upgrade software on the source or target systems, deployment is fast, and downtime is zero for migrations or upgrades. Stelo gives you the highest performance at the lowest impact.

EXTENDED SUPPORT AGREEMENTS

- Defined SLAs and 24x7 support
- Proactive monitoring and maintenance
- Performance tuning
- Upgrade service performed by Stelo
- Consistent application of time-tested best practices to ensure secure, reliable and efficient operation

STELO FEATURE SUMMARY

Functionality

- Incremental and Snapshot replication (mirroring)
- No-SQL Data Streaming with Continuous Apply or Scheduled Batch Apply
- Bidirectional replication
- Flexible replication scheduling and filtering
- Change Data Apply preserves transactional integrity

Performance

- Uses high-performance bulk-loading technology featuring restart checkpoints and multithreaded operation
- Subscription grouping and multi-path apply of change data
- MERGE SQL functionality to deliver data using micro-batches

Management

- Centralized “publish” control of tables published for replication
- Audit trail features (change timestamp, job identification, unique identifier, Delete as Update options)
- No-maintenance window: Refresh uses shared locks – operates concurrently with production
- Transaction replay
- Performance statistics
- User Role-based Authentication Control (RBAC)

Monitoring

- Schema change notifications
- Event log and email notifications

Configuration

- Replication of views, indexes, and constraints
- Automatic conversion of data types, including up-conversion to Unicode
- Customizable conversion of data types
- Automatic Data Modeling (dynamic creation of destination objects including automatic mapping of column attributes while permitting user-defined customizations).
- Vertical (column picking) & Horizontal (row selection criteria) definitions
- Ability to define columns using built-in macros to access change data audit information.
- Support for standard SQL functions to transform existing data
- Import of table and column mapping definitions

Programming

- SQL processing before or after full-refresh
- Comprehensive PowerShell scripting tools
- Push changes to message buses using XML formatted messages

SUPPORTED SOURCE DATABASES

- Actian Pervasive PSQL*
- Amazon RDS
- Amazon Redshift*
- Apache Derby*
- AWS Aurora
- Db2 for z/OS*
- Db2 LUW
- Exasol
- Google Cloud SQL for MySQL
- IBM Db2 for i
- Informix
- Kognitio WX2*
- MariaDB
- Microsoft Azure SQL
- Microsoft SQL Server
- MySQL
- Netezza*
- Oracle 10g*
- Oracle 11g to 19c
- Oracle 21
- PostgreSQL
- Salesforce*
- SAP HANA
- Snowflake*
- Tibero*
- TmaxSoft

*snapshot replication only

SUPPORTED DESTINATION DATABASES

- Any ODBC Database (including all supported sources)
- Azure Synapse
- Confluent
- DataBricks
- Google BigQuery
- Snowflake

SUPPORTED CLIENTS

- Windows 8.1, 10, 11 64-bit Edition
- Windows Server 2012R2, 2016, 2019, 2022
- Linux x86-64 (AMD64 & Intel EM64T)