

DATA VIRTUALITY LOGICAL DATA WAREHOUSE

The Data Virtuality Logical Data Warehouse (LDW) is a step up in the evolution of data warehousing, combining the best of ETL and data virtualization in a single, comprehensive data management solution. With more than 200 ready-to-use-connectors, the LDW connects to almost any database and API, enabling you to access your data in real time and gain instant business insights. An agile yet highly scalable data querying tool, the Logical Data Warehouse unites structured and unstructured as well as real-time and historical data in one virtual data layer. In short, the LDW is the only solution that both enables and accelerates complex data analyses with minimum effort and maximum flexibility. And all that by using only one language: SQL.

DATA VIRTUALITY BENEFITS

CONNECT IN 5 MINUTES




Data Virtuality provides a simple wizard for typing in your data source credentials to establish a connection. It only takes minutes and you see your data in SQL queryable tables instantly.

SIMPLY USE SQL



Query, join and integrate data from Google Analytics, Salesforce, Oracle, MySQL and many other data sources using SQL commands.

FAST & AGILE



Data Virtuality provides a single platform to connect, transform, query and join data from multiple data sources immediately – without depending on the IT department.

READY TO USE



Data Virtuality provides +200 ready to use connectors and is implemented in just one day. You can access your data from any database and API within minutes.

REAL-TIME & PERSISTED



Data Virtuality unifies all connected data in one virtual layer. This enables live-data access and persistent processing of queries and thus real-time reporting and data triggered task automation.

DATA VIRTUALITY

LOGICAL DATA WAREHOUSE 2.1

The release of the Data Virtuality Logical Data Warehouse 2.1 takes data integration to the next level. The LDW 2.1 introduces significant advancements that will make data management even easier for you. It bridges the gap between technological complexity and data consumers by separating the logical and business data layer. Thus, the LDW 2.1 permanently optimizes the backend data structure without influencing the business view. In other words, it simplifies and accelerates the access to the information your business users require and exposes the data in a business-friendly manner. With the Data Virtuality Logical Data Warehouse 2.1 you can integrate and model all your data with unprecedented flexibility and scalability – in every big data environment.

LOGICAL DATA WAREHOUSE 2.1

NEW KEY FEATURES

1. Data Virtuality Web

The Data Virtuality Web client enables all business users (depending on their permissions) to search metadata and query SQL from the webbrowser without using the DV Studio.

- Inspect available data sources and their contents
- Create simple or advanced SQL statements
- Generate SQL queries directly from the data source
- Download results as .csv-files to process further and/or import into e.g. Excel

2. Data Virtuality Synchronization Tool

The Data Virtuality Sync Tool synchronize data model changes between multiple LDW instances (staging/development and production).

- Work with multiple instances (procedures, jobs/schemas, views etc.)
- No manual synchronization necessary: automatically transfer your data model changes to the production instance
- Compare instances and identify changes

3. Modular Connectors

Modular Connectors are smart, easy-to-deploy connectors that contain both business logic as well as replication logic. Modular Connectors provide templates that cover most of the use cases working with APIs.

- Connect data sources without using an editor
- Immediately query important data/insights
- Create statements/procedures to connect to modular connectors
- Drastically reduce the workflow of getting business relevant data out of the source system

DATA VIRTUALITY HOW IT WORKS

Connect to Your Data Sources

The Data Virtuality Logical Data Warehouse connects to multiple data sources and allows querying data from there by using SQL. Data sources can either be relational or non-relational.

Create a Central Data Logic

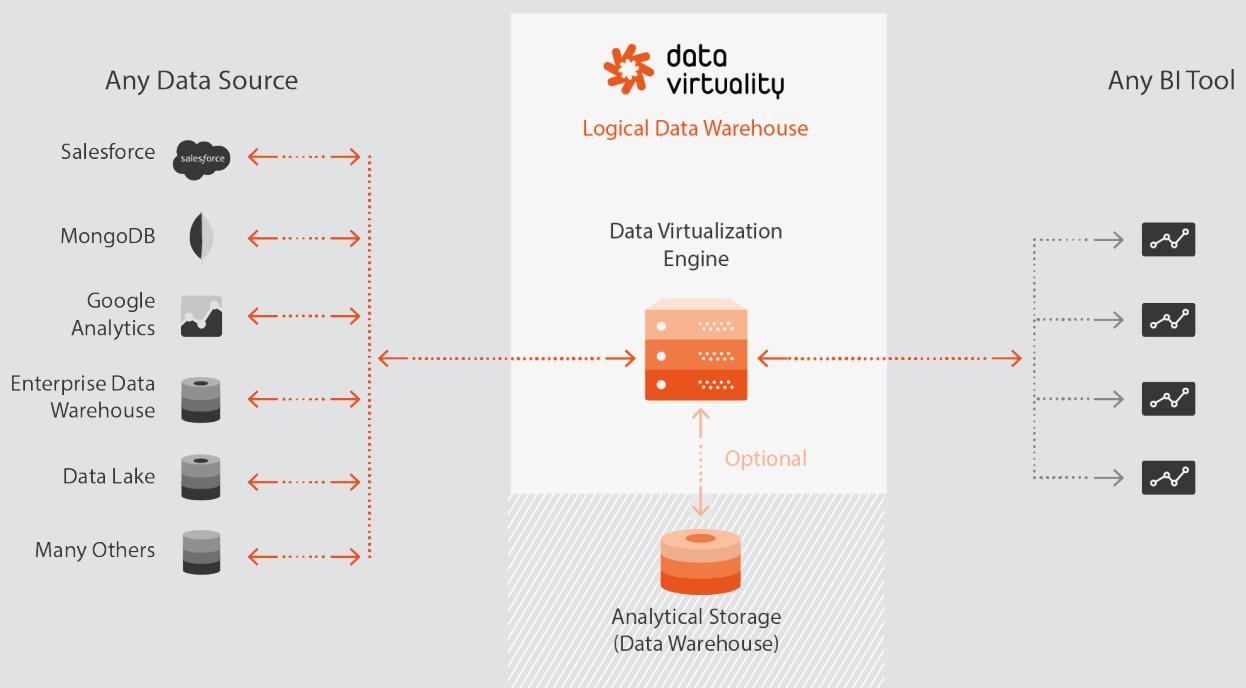
The LDW also allows you to create a central data logic that covers the business logic as well as the logical connections between the different systems. This layer can easily be implemented by using SQL views and/or stored procedures.

Get Your Data in a Data Warehouse

The Logical Data Warehouse copies/replicates data for you using SQL commands. It transforms these SQL commands into the respective native data source commands, such as SQL commands, API calls, access to various filesystems, or access to CDC (Change Data Capture) interfaces of different systems and databases. Data Virtuality provides predefined jobs for the most common transformations. The LDW gathers statistics about the usage of the data and uses these statistics to create a data warehouse in a special database, on a database server of your choice. This is done through automatic recommendations for the internal structure of the data warehouse that are then proposed to the database administrator. The data is then updated in your data warehouse with different time schedules.

Make Your Data Accessible

Finally, the LDW supports the standard interfaces (JDBC, ODBC, REST) to deliver data to the data consumers. This could be, for example, reporting tools, advanced analytics tools, or custom programs in various programming languages.



DATA VIRTUALITY LOGICAL DATA WAREHOUSE MAIN COMPONENTS

WHAT'S NEW?

- Data Virtuality Web
- Data Virtuality Sync Tool
- Modular Connectors

DATA VIRTUALITY SERVER

- Windows Server 64bit
- Linux 64bit (Redhat, CentOS, Ubuntu and many more)

DATA VIRTUALITY STUDIO

- Windows 64bit
- Linux 64bit
- Mac 64bit
- SQL Editor code completion on column level
- Metadata dependency viewer (Data Lineage)
- Metadata catalog and search
- Graphical view builder
- Wizards for easily connecting generic data (files/(S)FTP/S3/Web-services) using formats XML, JSON, CSV, xSV.
- Wizard for querying Google Analytics APIs

DATA VIRTUALITY WEB

- Search and browse metadata
- Write and run queries
- Download data

DATA FEDERATION

- Cross-database joins
- Nested loop
- Merge join
- Dependent semi-joins
- Cross-database unions
- Cross-database SELECT INTO, INSERT INTO
- Dynamic cost-based query optimization

SQL DIALECT

- ANSI-92 with extensions
- DDL, DML, procedural SQL
- Nested subqueries
- Common Table Expressions (CTEs)
- Window functions/Framing clauses
- XML/JSON parsing
- Web service access
- Scripting languages (server-side javascript)

DATABASES AND CONNECTORS

- More than 200 ready-to-use connectors. All our connectors can be found [here](#)

ACCESS DATA

- Via JDBC
- Via ODBC
- Windows (32bit/64bit)
- Linux (unixODBC 32bit/64bit)
- Mac (unixODBC 32bit/64bit)
- Via REST API (REST-JSON)

MOVE AND EXPORT DATA

- One query language: SQL
- Permission-based INSERT, UPDATE, DELETE statements on all relational databases, Salesforce, SAS
- Push-export via FTP, SFTP, SCP, email, S3, Azure Blob storage, web services (REST, SOAP, plain HTTP), file system and many more
- Export data using Data Virtuality Studio/SQL

SECURITY, AUTHENTICATION, AUDIT

- Built-in user/role based permission system
- Permission granularity on schema, table, column level
- LDAP authentication

- Active Directory authentication
- History of changes (versioning) for all custom metadata
- Access to audit information and usage statistics using SQL from external tools
- Security protocols: SSL/TLS, HTTPS

STRUCTURE OPTIMIZATION

- Materialized source tables and (virtual) views
- Precalculated joins
- Precalculated aggregations
- Automatic index creation

MATERIALIZATION ALGORITHMS

- Full copy (used with materialized tables, views, joins, aggregations)
- Incremental replication based on timestamp/id fields (used with materialized tables and views)

JOB TYPES

- Full copy with different cleanup options
- Batch update (optionally with overlap cleanup)
- History update (slowly changing dimension type 2)
- Upsert with optional surrogate keys
- Custom SQL jobs
- External programs and script

SCHEDULE TYPES

- Once with optional delay
- On time interval (every X minutes, hours etc.)
- Daily at certain times of day
- Weekly on certain weekdays
- Monthly
- Using custom cron expressions
- Dependant on other job or schedule (on success/failure/always)

IN-MEMORY CACHING FOR EVEN FASTER RESPONSES

- Session scope
- User scope
- Virtual database scope

WHAT ELSE?

- Mail notification on job and replication status
- Multi-tenancy
- Graphical web-based performance monitoring
- Password encryption
- Smart data movement approaches (Redshift S3 load, Azure DWH Blob storage load, Salesforce Bulk API)
- Programmatic access to all server functionality using Data Virtuality Management API

Visit: datavirtuality.com

Message: info@datavirtuality.com

Request Demo: demo@datavirtuality.com

US offices
1355 Market Street, Suite #488
San Francisco, CA 94103
Phone: +1 650 898 0227

1507 21st Street,
Set 200
Sacramento, CA 95811
Phone: +1 916 282 1679

German offices
Katharinenstrasse 15
04109 Leipzig
Phone: +49 341 26437217

c/o WeWork
Neue Rothofstraße 13-19
60313 Frankfurt am Main
Phone: +49 341 26362258

