

WhereScape®

WHEREscape
RED TECHNICAL
ARCHITECTURE

White Paper

WhereScape®

| | |
|---|-----------|
| Architectural Overview | 01 |
| WhereScape RED Clients | 02 |
| Target Data Warehouse and WhereScape RED Repository™ | 03 |
| Scheduler | 03 |
| Dependent Objects | 04 |
| System Architecture | 04 |

WhereScape®

WS RED TECHNICAL ARCHITECTURE

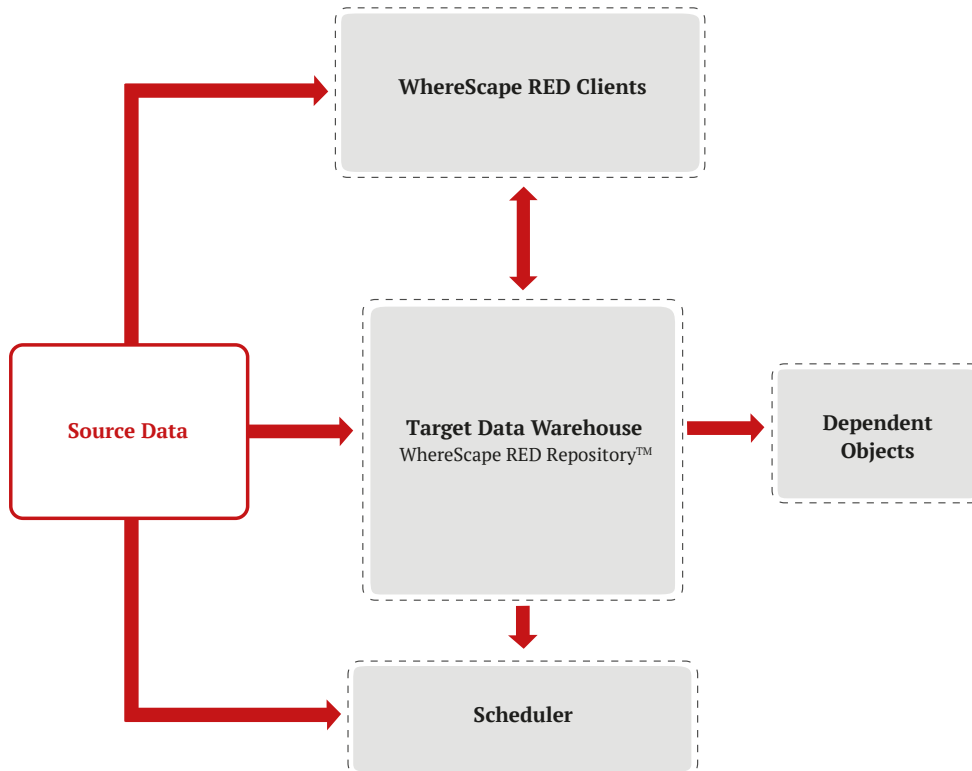
Architectural Overview

WhereScape RED is an integrated development environment (IDE) for developing and managing data warehouses. It can be used to build SQL Server, Teradata, Oracle or DB2 data warehouses. The output from a WhereScape RED project is a data warehouse built in a supported database platform – the Target Data Warehouse. The WhereScape RED Desktop reads from and writes to a set of database metadata tables (WhereScape

RED Repository™) that are stored in the Target Data Warehouse.

A WhereScape RED Repository™ is installed in each target data warehouse environment – development, test, production etc.

WhereScape RED generates database specific objects (such as tables, indexes and cubes) as well as procedural code and scripts to move and process data.



WhereScape®

All the objects are native database objects that can be viewed and modified by standard third party tools. Code can be changed, tables can be altered, and indexes can be added, changed or removed. They can also be scheduled through the integrated WhereScape RED scheduler and edited, versioned and backed up through the WhereScape RED Desktop. The integrated metadata can also be used to automatically generate data warehouse documentation, diagrams and lineage information.

The WhereScape RED Desktop is the main interface into the WhereScape RED Repository™. From the WhereScape RED Desktop connections can be set up to a WhereScape RED Repository™ (for metadata) and between external data sources and the Target Data Warehouse (for data loading).

The WhereScape RED Repository™ contains metadata about the objects created in the Target Data Warehouse and metadata about Dependent Objects. Metadata is created implicitly as operations are performed in the WhereScape RED Desktop and can be maintained manually. Metadata is used to generate DDL and DML and to drive the data warehouse wizards, scheduler and automated documentation. Existing data warehouse tables can be identified (retrofitted) to the WhereScape RED Repository™.

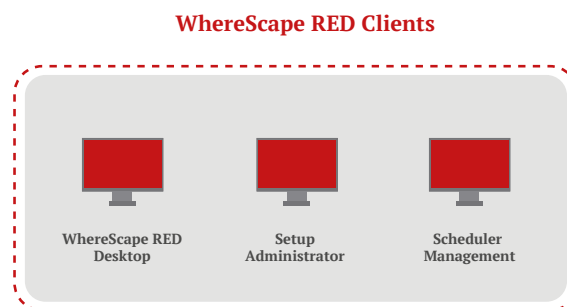
Each WhereScape RED Repository™ can have one or more schedulers associated with it, and the schedulers can either be Unix, Linux or Windows based. A scheduler polls the metadata looking for tasks that are available to run. Typically tasks update the metadata, which is the basis for the workflow. The workflow routines are standard generated code and can be easily added to manually created code.

Dependent Objects are external to the Target Data Warehouse, but still have metadata about them stored in the WhereScape RED Repository™.

Associating metadata with the Dependent Objects allows them to be created, managed and documented within the WhereScape environment. Dependent Objects include Microsoft Analysis Services cubes, MicroStrategy Projects and export objects used for downstream feeds.

WhereScape RED Clients

There are three WhereScape RED client components.



The WhereScape RED Desktop is the main development interface for the WhereScape RED Repository™. A copy of the WhereScape RED Desktop is installed on each developer's PC. Development carried out within the WhereScape RED Desktop automatically creates metadata in the WhereScape RED Repository™.

The Setup Administrator is used to setup repositories, maintain Windows schedulers and load applications.

The Scheduler Management client is optional (the functionality is included in the WhereScape RED Desktop) and provides an interface into the WhereScape RED Repository™ specifically for scheduling. It is used primarily by operations personnel for job status information.

WhereScape®

Target Data Warehouse and WhereScape RED Repository™

The Target Data Warehouse consists of data warehouse objects and the WhereScape RED Repository™. WhereScape RED supports Oracle, SQL Server, DB2, and Teradata databases as Target Data Warehouses.

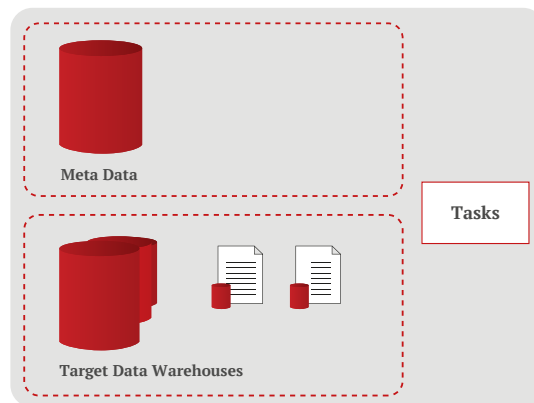
The data warehouse objects created by WhereScape RED are standard database objects such as tables, scripts, procedures and indexes. They are specific to the target database e.g. Integration Services packages, Analysis Services cubes and T-SQL in SQL Server, SQL Loader control files and PL/SQL procedures in Oracle, and MultiLoad, FastLoad, TPT or BTEQ scripts in Teradata.

The WhereScape RED Repository™ is a set of metadata tables that reside in the Target Data Warehouse. It consists of information about each of the data warehouse objects. The metadata tables are standard database tables that are accessible by off the shelf query tools.

Multiple versions of objects can be retained within the WhereScape RED Repository™, and the WhereScape RED Desktop includes version control functionality.

WhereScape RED provides a number of mechanisms for promoting between repositories (e.g. from a development to a test repository). These include REFRESH and IMPORT (available from the WhereScape RED Desktop) or APPLICATIONS (built in the WhereScape RED Desktop and loaded via the WhereScape Setup Administrator). Both metadata and data warehouse objects are promoted.


Tasks are created by the scheduler, and run within the Target Data Warehouse database.



Scheduler

A WhereScape RED Repository™ can have one or more Unix, Linux or Windows schedulers associated with it. Typically these are on the same server as the target data warehouse, although this is not a requirement. The WhereScape RED schedulers poll the metadata looking for tasks that can be run. Tasks can be organized into jobs which can have dependencies defined.

Typically tasks run in the target database. For external loading of source data (e.g. using Oracle's SQL Loader) tasks may run on the scheduler platform. Typically scheduler tasks initiate loads (which are database specific SQL) or procedures (which are database specific procedural code). The integrated WhereScape RED schedulers can be called from an enterprise scheduler if required.



WhereScape®

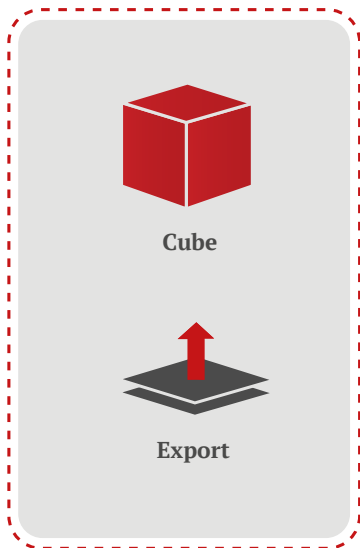
Dependent Objects

Dependent Objects are external to the Target Data Warehouse, but still have metadata about them stored in the WhereScape RED Repository™. Dependent Objects include Analysis Services cubes, Exports and MicroStrategy Projects.

Analysis Services cubes for example, are created by WhereScape RED from metadata stored in the WhereScape RED Repository™. They are designed to be identical to manually created cubes, and can be managed by the WhereScape RED scheduler.

In addition, as they have been created through WhereScape RED, they have associated metadata which is available to the generated documentation. Exports are created from data warehouse tables and are used to feed downstream systems.

Dependent Objects



System Architecture

WhereScape RED Client applications are PC based client applications. The WhereScape RED Desktop is installed on developer's PCs, The Setup Administrator resides on an administrator's PC and, if required, Scheduler Management clients can be installed on operator's PCs. All the WhereScape RED clients can reside on the same PC.

Using the WhereScape RED Desktop, connections are defined to data sources and to the WhereScape RED Repository™. Data sources can be flat files, XML or any database that is accessed via ODBC. and can reside on any server accessible to the WhereScape RED Desktop (for development) and the Target Data Warehouse (for data loading). Generally ODBC is only used for loading meta data (e.g. column names). Data is usually loaded using the database platforms' own bulk loading utilities which are managed by WhereScape RED.

The WhereScape RED Repository™ resides in the Target Data Warehouse database, and consists of a series of open database tables containing metadata. There is one WhereScape RED Repository™ installed in each environment – development, test, production etc. WhereScape RED supports Oracle, SQL Server, DB2, and Teradata Target Data Warehouses.

Each WhereScape RED Repository™ can have one or more schedulers associated with it. Schedulers reside on Unix, Linux or Windows platforms, and poll the WhereScape RED Repository™ looking for tasks to run. They can reside on the same platform as the Target Data Warehouse. Connection Objects define how the data will be loaded from data sources into the Target Data Warehouse. Data movement within the Target Data Warehouse is accomplished through database procedures that can be generated, modified or manually created using the WhereScape RED Desktop.

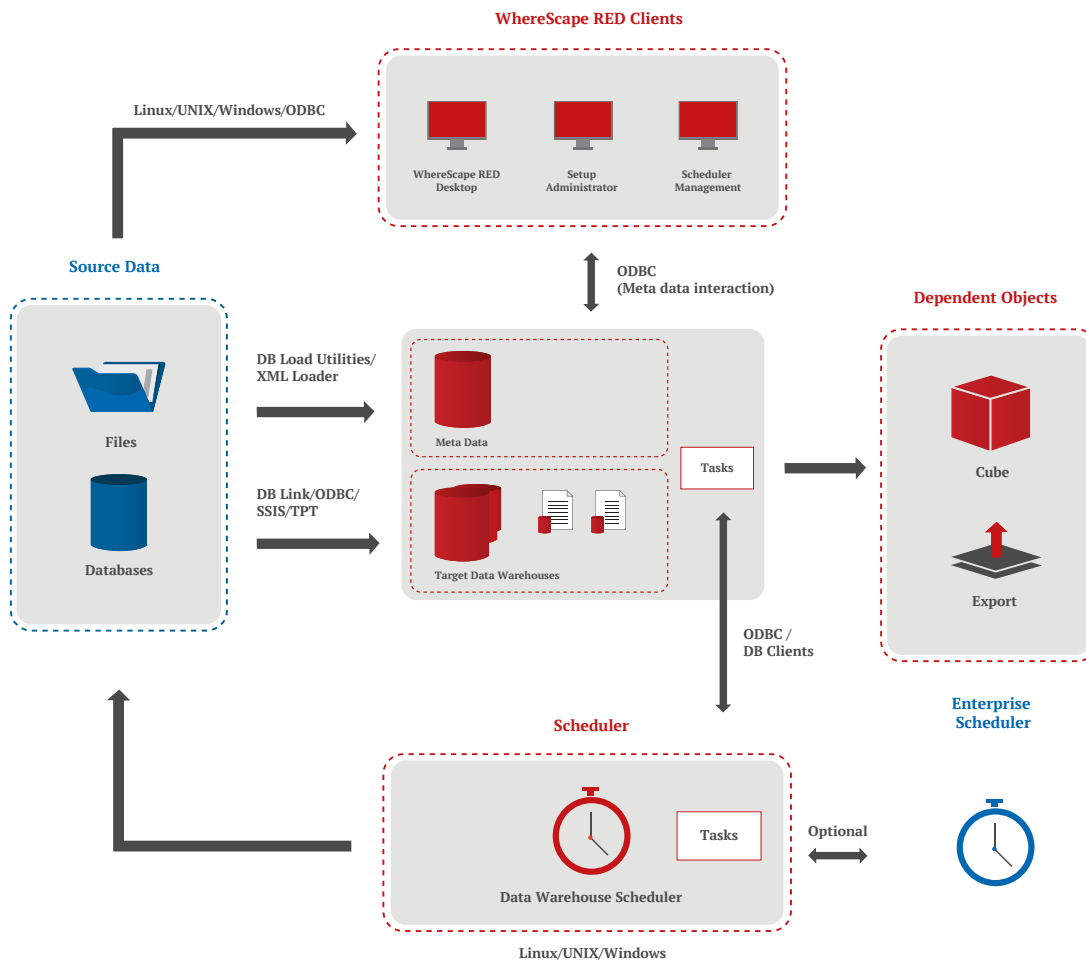
WhereScape®

Analysis Services cubes reside on a Windows Server, which can be the same platform as the Target Data Warehouse.

The WhereScape RED Desktop is used for data warehouse development. Outside the development the only WhereScape component running is the WhereScape RED scheduler – the data warehouse

objects created are all standard database objects that can be viewed and queried through any standard database query tool.

For more information on how WhereScape RED operates, or to arrange an online demonstration, please visit: www.wherescape.com



About WhereScape

The pioneer in data warehouse automation software, WhereScape empowers organizations constrained by time, money or lack of resources, to deliver business value from their decision support infrastructure – including enterprise data warehouses, business facing data marts, and big data solutions. WhereScape has global operations in the USA, UK, Singapore, and New Zealand. www.wherescape.com