

DQbroker®

- **True distributed query**
- **All data looks and acts like one database**
- **Global metadata management**
- **Native access to multiple sources**



DQbroker is the foundation of *DQpowersuite*, Decision Support's enterprise data access and extraction, transformation and loading (ETL) tools. *DQbroker* lets applications access all needed data, regardless of database vendor, hardware platform or geographic location. In essence, *DQbroker* is a single virtual relational database. This give users and developer simultaneous access to all available data in real time. It enables movement of notably less data across the network when processing queries that join distributed data, and interacts natively with most major relational databases. *DQbroker* is significantly more efficient and effective at accessing and joining enterprise data than the three-tiered, ODBC-based or hub design solutions widely available today.

Simple Installation and Configuration

Any system administrator can accomplish the basic installation and configuration of *DQbroker* in less than an hour. It is easy. Other software with similar capabilities often requires a week of training just to learn how to install and configure it.

True Distributed Query Processing

DQbroker has an n-tiered, thin-client architecture that distributes the processing of queries as close to the source data as possible. Queries that access and join data from multiple heterogeneous data sources are processed on multiple servers simultaneously. Only the data

necessary to resolve the query crosses the network. This differs from the fat-server, three-tier approach common in the marketplace today, which requires entire data sets to move across the network to be processed by a single hub server. Specifically, *DQbroker* converts a single SQL query that joins distributed data into multiple SQL queries. Each sub-query is targeted at a single data source. The sub-queries contain as much selection, filtering, joining, and sorting as possible. This allows *DQbroker* to leverage the capabilities of each database management system (DBMS), minimizing the amount of data returned and the time needed to return it. The processing happens as efficiently as possible because *DQbroker* communicates with all major relational database management systems (RDBMS) using native access routines.

Global Metadata Management

A global metadata cache is valuable because it enables every *DQbroker* server to know the current state of all data in a distributed domain. Caching metadata also makes the retrieval of database properties faster. This accelerates query processing in a distributed environment because *DQbroker* has all the information it needs to distribute sub-queries.

Local Control of Access & Security

Each DBMS controls access to its own data. *DQbroker* provides a layer of security in addition to that already provided by the DBMS. *DQbroker* only grants access or update privileges as allowed by the DBMS. With *DQbroker*, the goals of local control and universal views are both achieved

Enterprise View Builder GUI

DQbroker includes a graphical user interface (GUI) that simplifies the creation of distributed enterprise views. It enables the creation of joins between tables in different data sources and on different platforms.

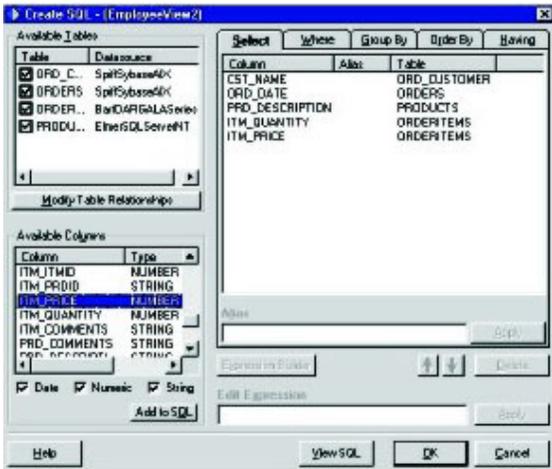


Figure 1: Fields from four different tables, in three different databases (Sybase, Unisys A-Series hierarchical, SQL Server), on three different platforms (IBM RS6000, Unisys A-Series, Windows NT) are selected for this view. The end-user or developer is shielded from the location and type of data. All available data looks like it is in one database.

Enterprise Metadata Cache

DQbroker can request metadata dynamically, or collect it in a global metadata cache. A global metadata cache allows each DQbroker server to be universally aware of the location and structure of all available distributed data. This lets DQbroker make the best possible decisions about how to distribute processing for a query that joins data from multiple data sources.

Additional Access, Not Additional Training

With DQbroker, end-users can perform simultaneous, real-time analysis on distributed enterprise data in warehouses, marts and operational data stores (ODS) with the ODBC-compliant tools already familiar to them. Developers can develop web-based and client/server applications for the entire enterprise with the development tools they already know. With DQbroker, end-users and developers get more data access, not more application training. They become more productive immediately.

Updates at the Source

DQbroker can add, delete or modify data in any available table that supports data update. DQbroker updates operational data directly. The next person to access it has the absolute most up-to-date information available. Rather than data that is only as recent as the last update or synchronization process.

Operating System Specifications

- UNIX
- Windows NT Server or Workstation
- AIX
- Linux
- Solaris
- Clearpath (Unisys)

Hardware Environments

- Data General
- DEC
- HP
- IBM
- Intel compatible
- NCR
- Sun
- Unisys

Database Management System Access

- Informix
- ODBC (32-bit)
- Oracle
- DB2
- Sybase
- SQL Server
- DMSII
- VSAM
- IMS
- DMS (2200)
- Progress
- ADABAS
- IDMS