


Günter Weber IBM , weberg@de.ibm.com

GS09: How IBM can modernize your Data Management



Data Management

3rd European Workshop for z/VSE, z/VM and Linux
on System z -26 -28 Oct. 2009, Dresden

GS09 – How IBM can modernize your Data Management

Speaker: Günter Weber, IBM

In this session you'll learn how the new functions of IBM's Infosphere Federation Server 9.7 and DB2 9.7 can help you to create modern applications with federated datasources. How they can help you to reduce costs by using features like DB2 Data Compression or how to include various DBM-Systems into your application landscape via federation and how using DB2s new Oracle compatibility features allows applications written for Oracle to run with DB2 without the need for complex porting of these applications.

Agenda

- **DB2 9.7**
- **Database Federation**
- **IBM Federation Server 9.7**

DB2 9.7

DB2 9.7 Highlights

Security: e.g. SSL client support, enhanced DB2 security model

Performance: concurrency enhancements, statement concentrator, scan sharing, WLM improvements

Compression: Index-, temporary table- and XML-compression, Large Object Inlining

XML: with DPF, in range partitions, database views and MQTs


HADR: Read on standby

Partitioned Indexes: Local Range Partition Indexes

Oracle compatibility: PL/SQL support, Oracle data dictionary-compatible view support, SQL*Plus compatible CLP, table truncate


DB2 9.7 – Compression Improvements

- Multiple algorithms for automatic index compression



*Unique in
the
industry*

- Automatic compression for temporary tables



*Unique in
the
industry*

- Intelligent compression of large objects and XML



Enabling an Oracle Environment in DB2 !

Oracle Enablement Step	Comments
<p>#1 - Set registry variable:</p> <pre>db2set db2_compatibility_vector=FFF</pre> <p>(Note: each feature has it's own bit position value and it's possible to enable any individual feature or combination of features. The value FFF enables them all.)</p>	<p>Sets compatibility features:</p> <ul style="list-style-type: none"> • Rownum • DUAL • (+) Outer Join • Connect By (hierarchical queries) • Number data type • Varchar2 data type • Date data type • Truncate table • Character literals (Graphic and Char) New! • Data dictionary compatible views • PL/SQL compilation and execution ability
<p>#2 - Restart instance:</p> <pre>db2stop</pre> <pre>db2start</pre>	<p>Force compatibility vector to take affect</p>
<p>#3 - Create Oracle "enabled" DB2 database:</p> <pre>db2 create database ora2db2...</pre>	<p>Creates the Oracle-like data dictionary views and exploits Oracle data types</p>

* This is not a new concept - it was introduced in DB2 9.1 and expanded in 9.5. DB2 9.7 is now bringing this home to a much greater extent.

CLPPLUS

A Familiar Command Line Environment for the Oracle DBA *

CLPPLUS Support Examples	Comments
<code>REM This is a remark example</code>	Remarks
<code>SET termout off</code> <code>SET linesize 250</code>	SET commands for CLPPLUS script environment settings (too many to list here)
<code>SPOOL myfile_out.txt</code>	Direct output to a file
<code>ACCEPT [substitution_variable_name]</code>	Keyboard input of substitution variable
<code>DEFINE [variable_name]</code>	Create or replace variable
<code>COLUMN [colname] format 9,999</code> <code>heading 'My heading'</code>	Format output data as well as headings for that data
<code>EXEC [procname]</code>	Executes a procedure
<code>CONNECT, DISCONNECT</code>	Connect to or disconnect from a DB2 database
<code>EXIT, QUIT</code>	Leave CLPPLUS session
<code>HELP, ?</code>	Help on CLPPLUS
<code>HOST [OS_command]</code>	Execute operating system commands
<code>DESCRIBE [tablename]</code>	Describe table, view, output from function, etc.
<code>EDIT, INPUT, APPEND, CLEAR, DEL, GET, etc.</code>	Buffer control commands

* The Oracle tool is named: SQL*Plus

PL/SQL support (1)

- **Oracle Programming Language for SQL Procedures**
- **It is a fully functional programming language that allows**
 - Anonymous Labeled blocks
 - Local variable declarations
 - Assignments
 - Control Flow Statements
 - Implicit and Explicit Loops
 - Error Handling
 - Imitated OO technology techniques

PL/SQL support (2)

- **Local Variables Declaration**
- **Values Assignments**
- **Control Flow**
 - Conditional Statements
 - Iterative Control Structures
- **Cursors Processing**
- **Error Handling**
- **Miscellaneous**
 - Collections
 - Bulk Collect
 - Packages
 - Built-in Packages

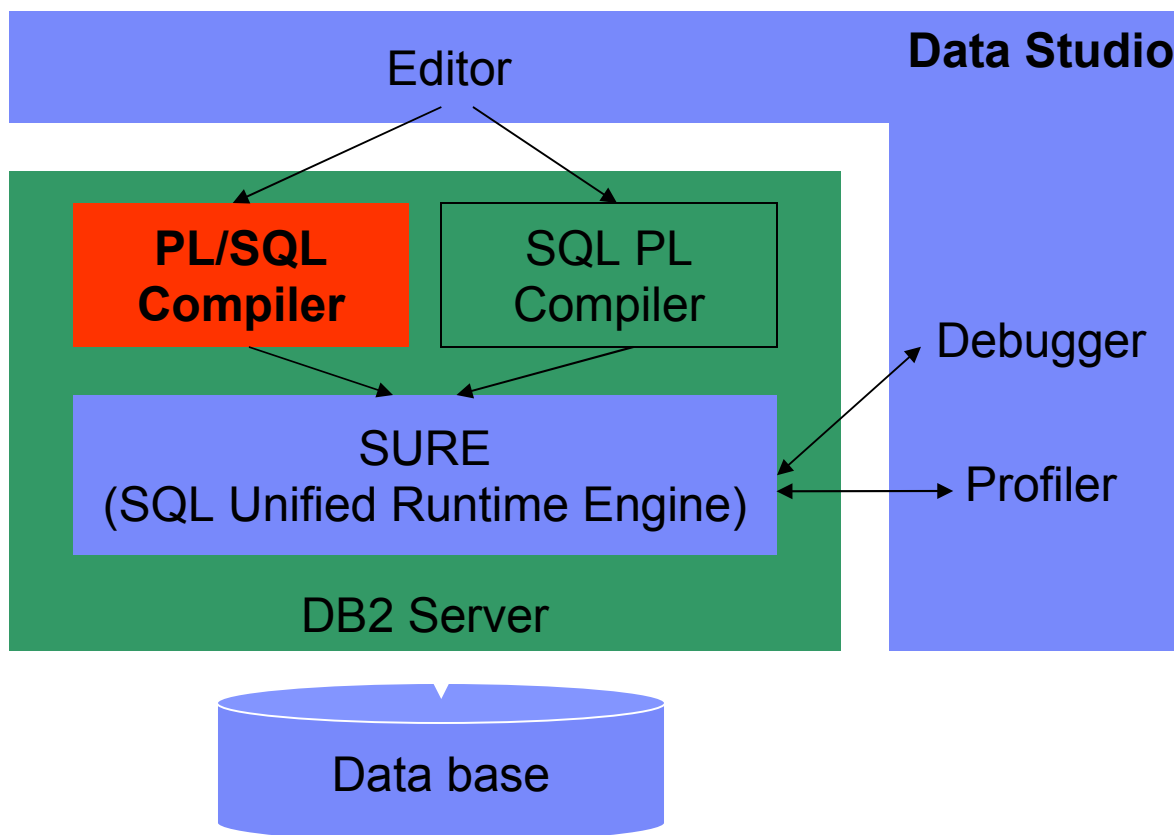
PL/SQL support (3) - Advanced Packages

- **DBMS_OUTPUT** putting and getting messages
- **UTL_FILE** reading, writing of files
- **DBMS_ALERT** sending, receiving alerts
- **DBMS_PIPE** sending messages through pipes
- **DBMS_JOB** scheduling and managing jobs
- **DBMS_LOB** routines for large object handling
- **DBMS_SQL** procedures for dynamic sql execution
- **DBMS_UTILITY** utility routines
- **UTL_MAIL** procedures for sending of emails
- **UTL_SMTP** routines for email sending via SMTP

PL/SQL support (4)

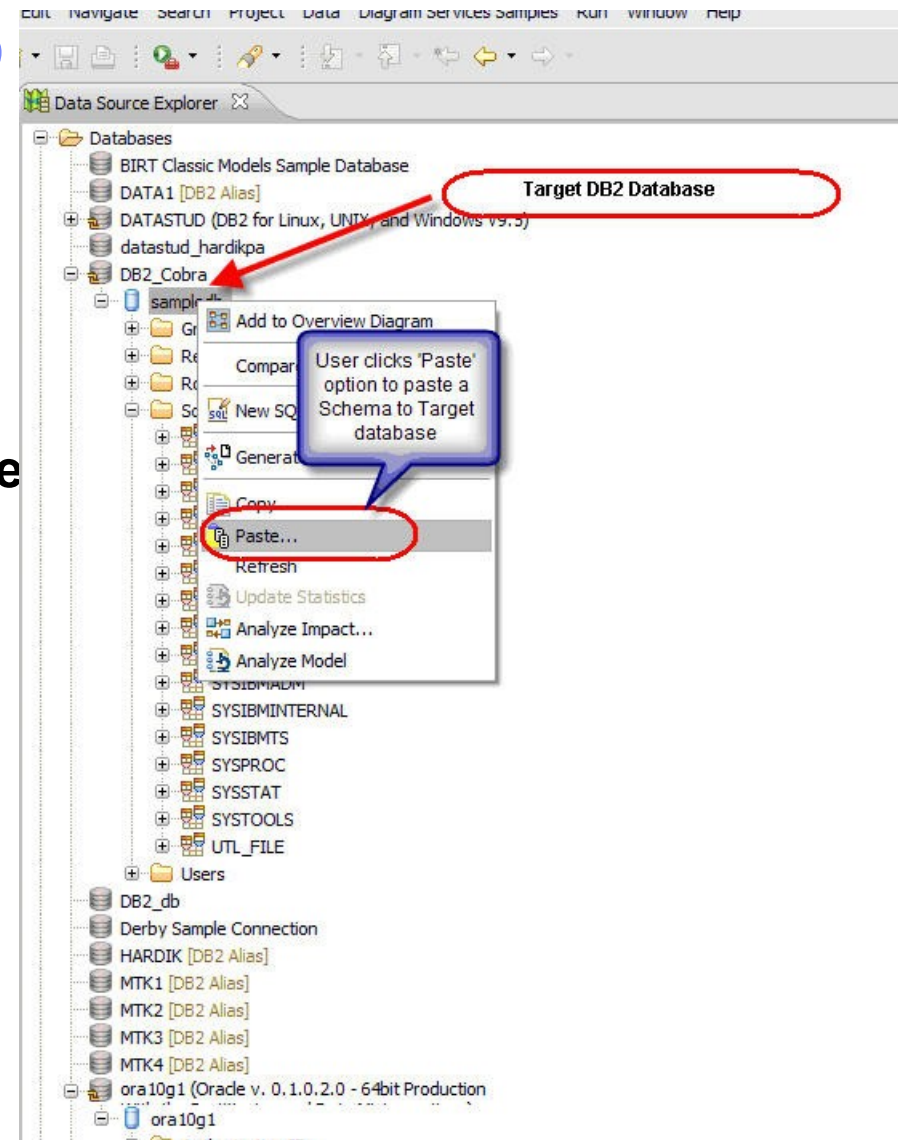
Built in PL/SQL compiler

Source level debugging and profiling



Easily Import Other Vendor's Schemas with IBM Data / OPTIM Studio

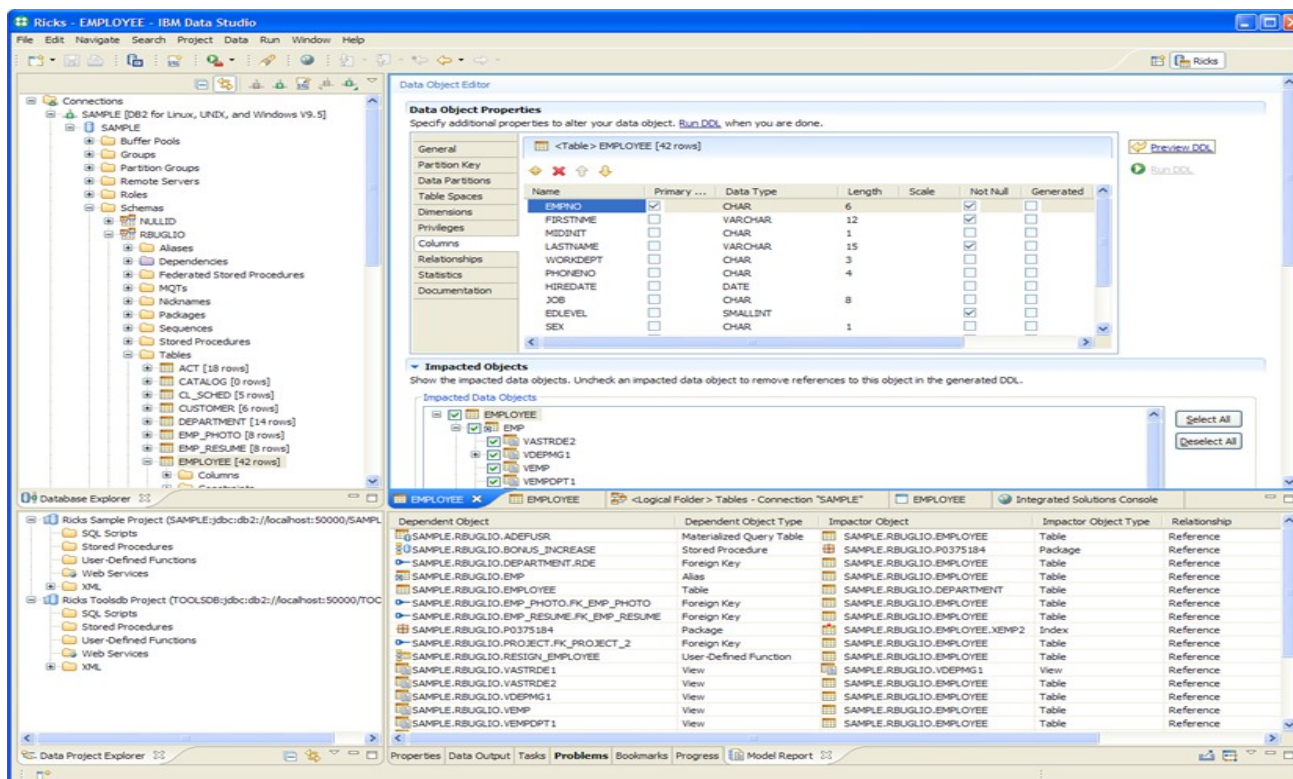
- Drag and drop databases between Oracle Database and DB2 for quick migrations and trials
- Easily map schemas and data types from Oracle Database to DB2
- Automatically map schemas and data types
- Oracle Database developers can quickly start using DB2



Administration / Development / Performance Optimization

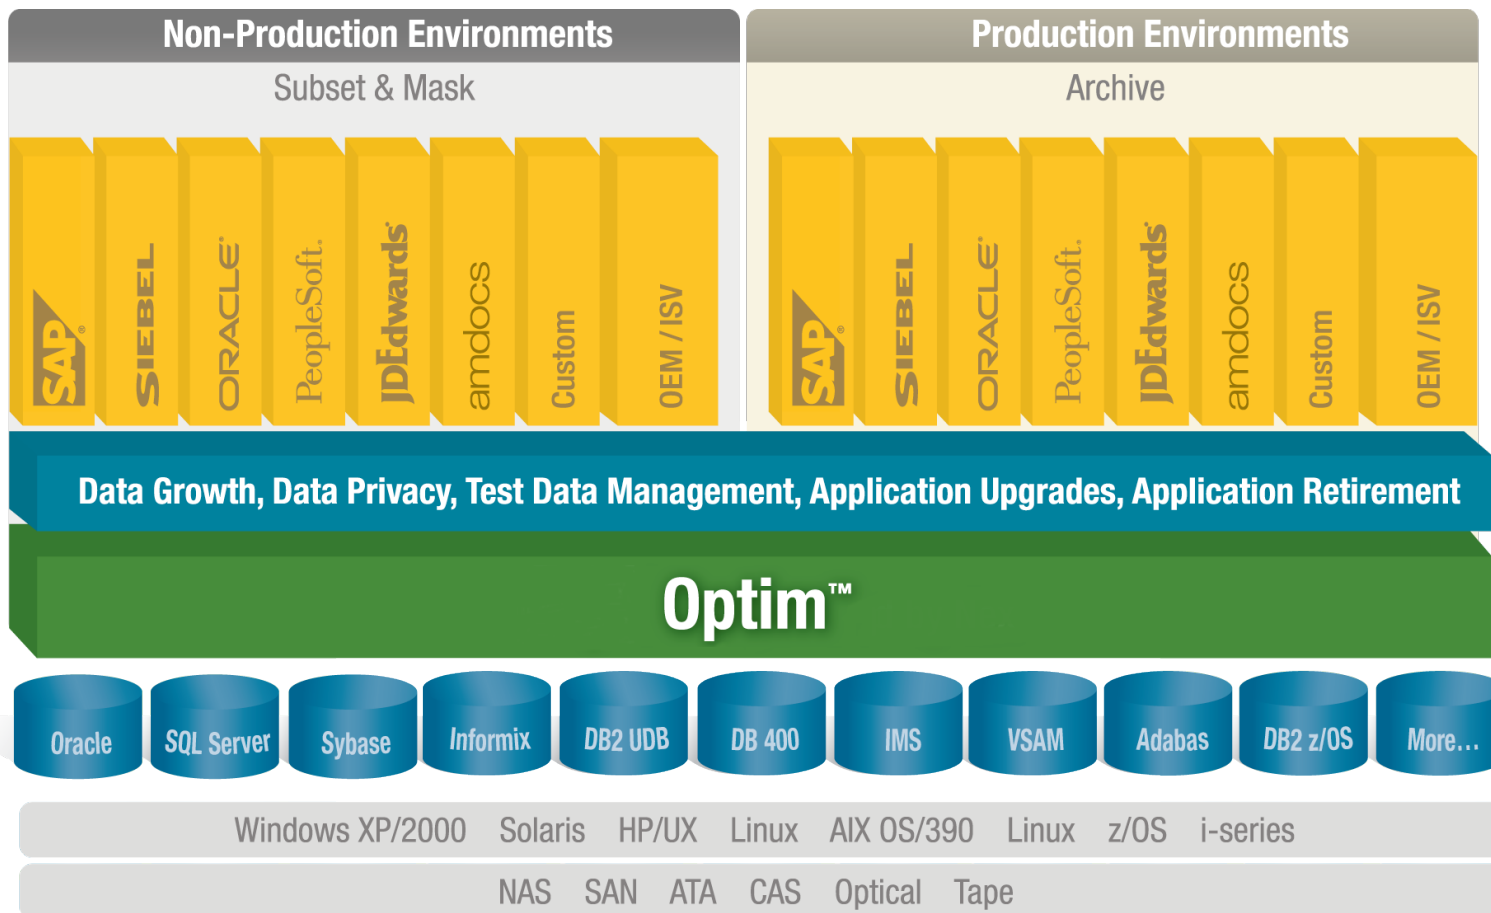
IBM Data / Optim Studio

A consistent and productive environment based on Eclipse for Data Management, Schema Management, Stored Procedure support for DB2 and ...



Optim ?

Name comes from a toolset for Testdatamanagement (incl. Data Masking) and DBMS Archiving tools...



Database Federation



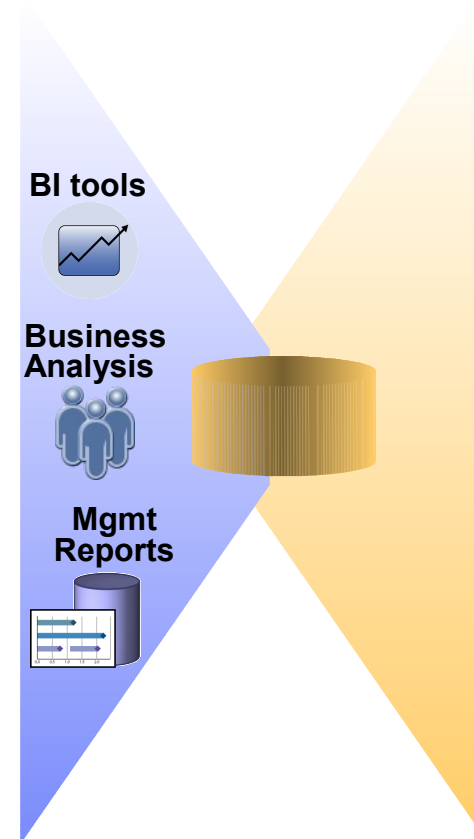
What if you could...

Access data anywhere in your enterprise

- No matter where it resides
- Regardless of what format it is in
- Regardless of vendor
- Without creating new databases and without disruptive changes to existing ones...

Using standard SQL and any tool that supports JDBC/ODBC...

all while looking to the user like a single database!



InfoSphere Federation Server

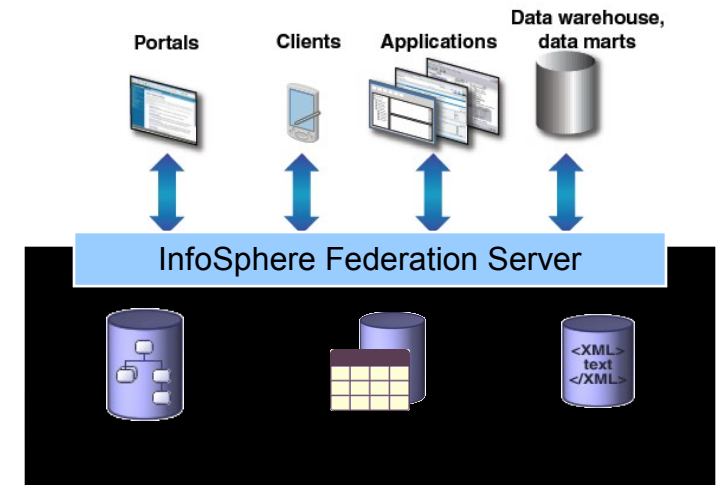
- Transparent
 - Appears to be one source
 - Independent of how and where data is stored
 - Applications continue to work despite of any change in how data is stored
- Heterogeneous
 - Accesses data from diverse sources
 - Relational, Structured, XML, messages, Web content ...
- Extensible
 - Bring together almost any data source.
 - Wrapper Development Toolkit
- High Function
 - Full query support against all data
 - Capabilities of sources as well
- Autonomous
 - Non-disruptive to data sources, existing applications, systems.
- High Performance
 - Optimization of distributed queries



IBM InfoSphere Federation Server

Access and integrate heterogeneous information across multiple sources as if they were a single source

Extend value of existing analytical applications by providing real-time access to integrated information

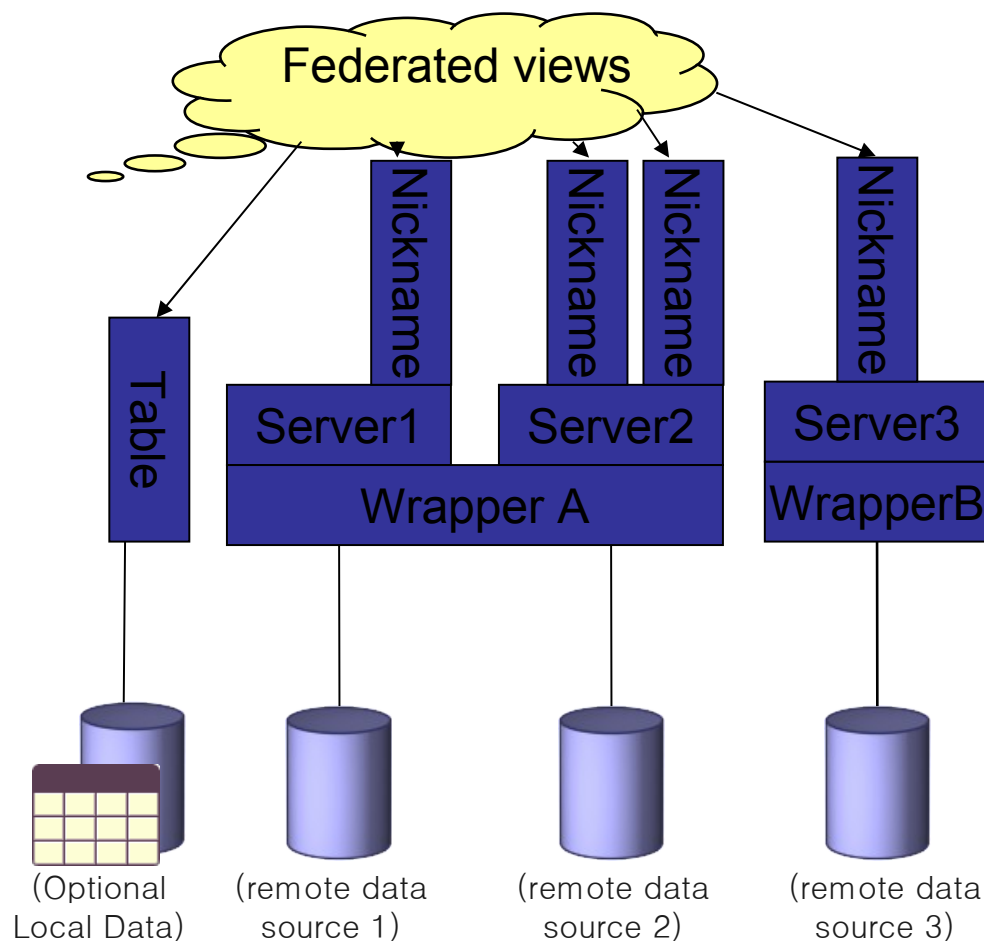


Federation Terminology

Wrapper: a library allowing access to a particular class of data sources or protocols (Net8, DRDA, CTLIB...). Gathers information about data source characteristics

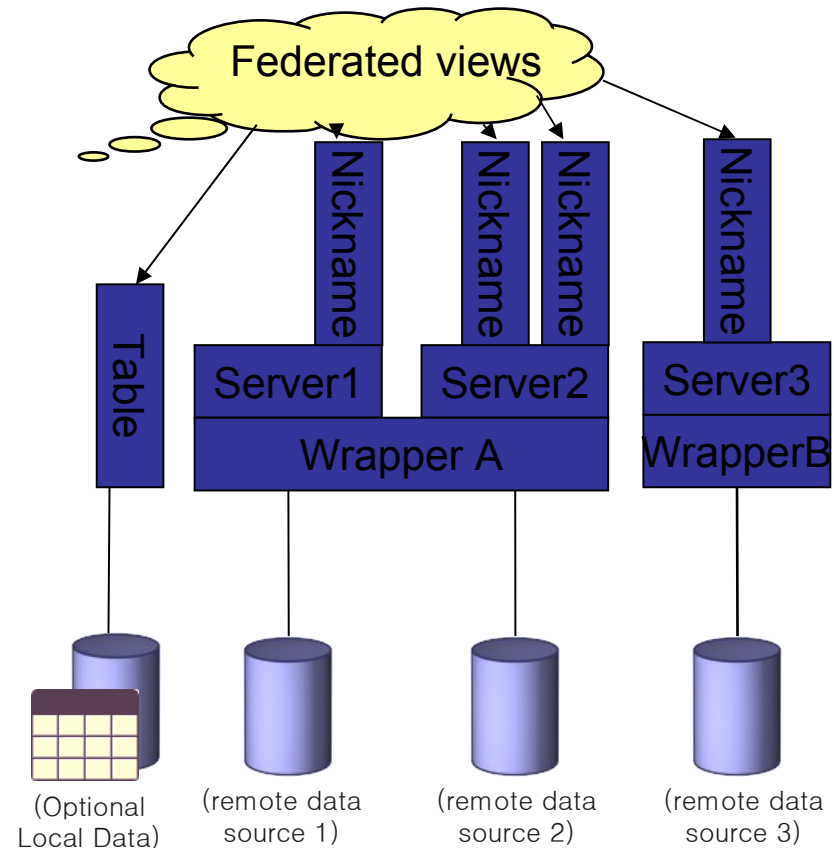
Server: represents a specific data source

Nickname: a local alias to data on a remote server (mapped to rows and columns)



How Federation Server works

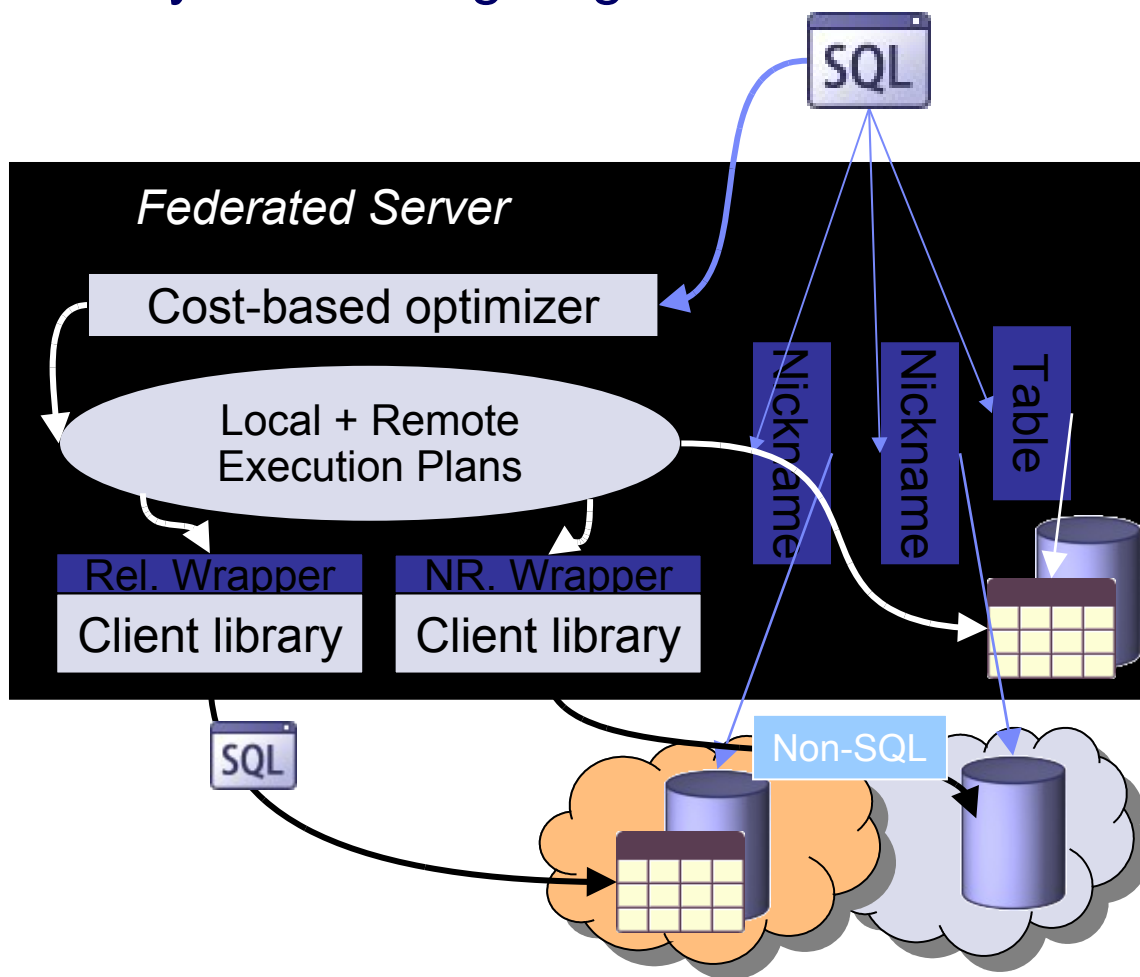
1. User's client sends a request which goes to the Federation Server interface
2. Federation Server parses the request and creates query fragments to send to the sources
3. Each source system receives its query fragment, processes it and returns the result
4. Federation Server assembles the final result (which may mean additional processing) and the result is sent back to the client



Data Federation Query Processing

Powerful Query Processing Engine

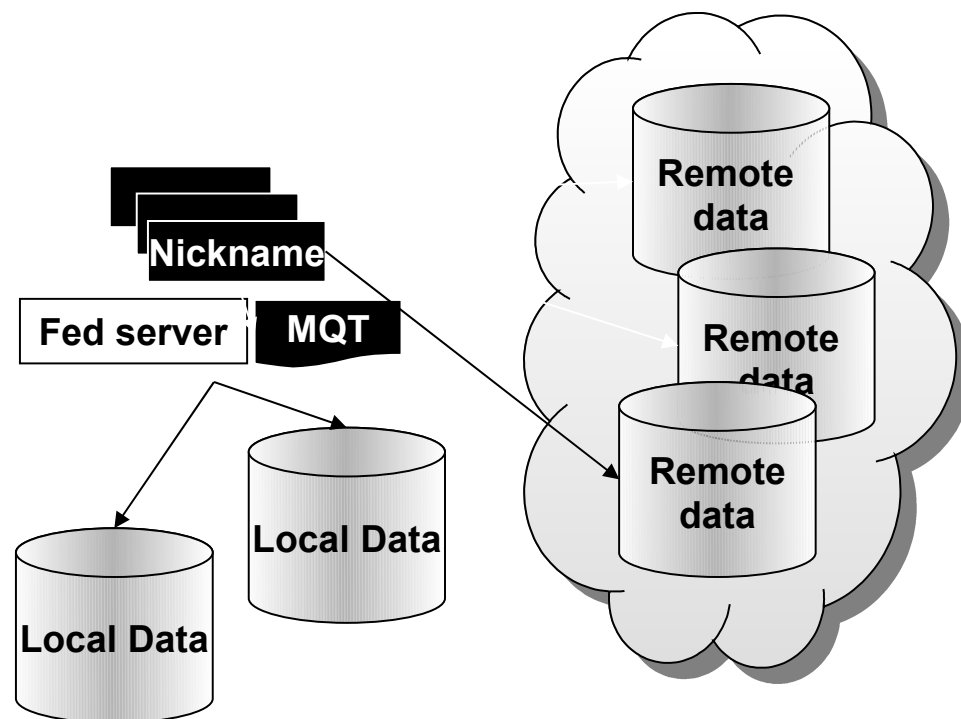
- Decomposes, rewrites and distributes queries
- Cost-based optimizer chooses query plan with pushdown as appropriate
- Query execution engine drives wrappers, combines results
- Compensates for missing function in data source
- Invokes functions at remote sources as needed



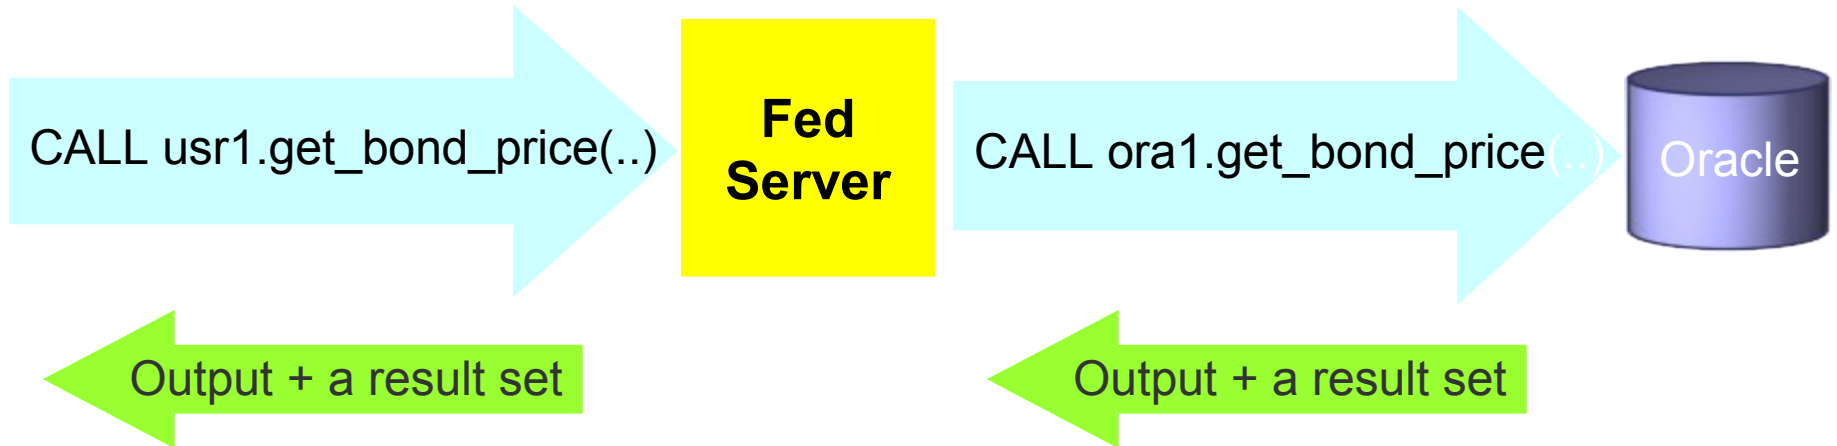
Local Caching with MQTs

Materialized Query Tables (MQTs) Over Nicknames

- MQT: local table defined by the result of a query
 - Can include joins, aggregations over multiple nicknames
 - Can be indexed, replicated in partitioned environment
 - Optimizer “routes to” them transparently as appropriate
 - Can include both local DB2 tables and nicknames
 - Can include nicknames to nonrelational objects
- Use to replace remote access with local access



Federated Procedures (1)



The Oracle procedure ora1.get_bond_price(...) is cataloged as a federated procedure called usr1.get_bond_price(...).

Federated Procedures (2)

- ▶ **What is it?:**
This feature allows the federated system to access stored procedures at remote data sources.
- ▶ **Why is it important?:**
A stored procedure is a program which is physically stored within a database. A common use is the encapsulation of commonly used tasks that might involve complex SQL queries.
- ▶ **Benefit?**
Using Federated stored procedures allows any stored procedure already developed on the remote database to be part of query. This saves a lot of work because procedures already stored in the system can be reused. This is used extensively by Sybase and Oracle.

Supported Data Sources

Relational data sources

- **IBM DB2 (for z/OS®, i, LUW)**
- **IBM Informix®**
- Oracle
- Sybase Adaptive Server Enterprise
- Microsoft® SQL Server™
- Teradata
- DB2 Server for VSE and VM
- Sources accessible by open database connectivity (ODBC) or Java database connectivity (JDBC)

Mainframe data sources¹

- VSAM, IAM, Sequential
- IMS
- Software AG Adabas
- Computer Associates CA-Datcom
- Computer Associates CA-IDMS

Extensibility

- **C++ and Java Software Development Kits**

Other sources and formats

- Web services
- WebSphere MQ message queues
- Microsoft Excel® spreadsheets
- Table-structured flat files
- XML documents
- OLE DB-accessible data sources
- Script output data (Perl, Python and others)
- BioRS

¹Via separate purchase of InfoSphere Classic Federation for z/OS


IBM Federation Server 9.7



Release Highlights


InfoSphere Federation Server 9.1

- ✓ RDA Data Modeling
- ✓ **2 Phase Commit**
- ✓ Stored Procedure Support Phase 1
- ✓ LDAP plug-in
- ✓ WISD SOA support



InfoSphere Federation Server 9.1 FPx

- ✓ LBAC
- ✓ Stored Proc SQL Server and DB2
- ✓ CCC EAL4
- ✓ HPUX IA



InfoSphere Federation Server 9.5

- ✓ XML Federation
- ✓ Trusted Contexts
- ✓ Roles Based Security
- ✓ Auto Statistics Refresh



InfoSphere Federation Server 9.7

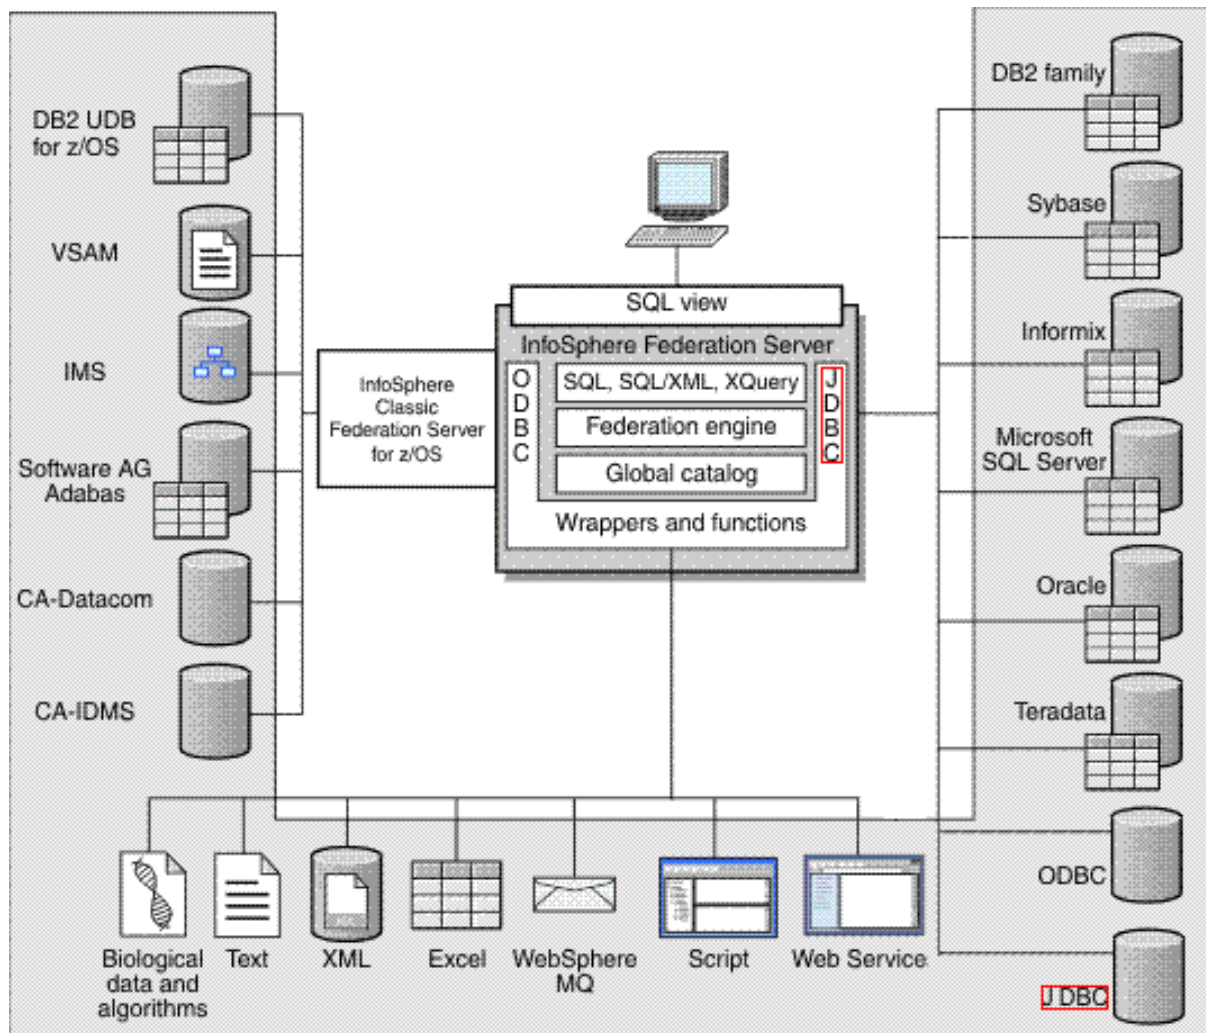
- ✓ **JDBC Wrapper**
- ✓ Federated Proc Multiple Result Set for Oracle Wrapper
- ✓ Join of Store Proc Result Set
- ✓ Remote Query Cancellation
- ✓ Installation Changes
- ✓ Database Migration Upgrade
- ✓ DECFLOAT Data Type Support
- ✓ Implicit Casting Support
- ✓ Oracle Semantic and Oracle DATE support
- ✓ Teradata LOB read
- ✓ Maintain MQ XML UDFs/SPs



JDBC Wrapper (1)

- ▶ **What is it?:**
JDBC wrapper provides transparent connectivity from the Federated Server to remote data source that supports JDBC connection.
- ▶ **Why is it important?:**
It is based on Java, which can be "Write Once, Run Anywhere". JDBC wrapper is supported at all platforms that federation server supports, and has less likelihood of programmer errors. Because JDBC standard is still developing, the latest DB technology can be added into JDBC wrapper. (e.g. F2PC will be support in the future release of JDBC wrapper) .
- ▶ **Benefit?**
It provides the connection to data sources that only support JDBC connection (e.g. Cloudscape (Derby), Universal Data, U2 product on Linux doesn't support ODBC connectivity, and DataDirect doesn't offer mssql ODBC on zlinux). Also it offers customer an alternative generous way to connect to most of remote data sources in much easier way.

JDBC Wrapper (2)



Supported Platforms

■ IFS V9.7 System and Data Source Requirements

– <http://www-01.ibm.com/support/docview.wss?uid=swg27015299>

- Linux (64 Bit) for zSeries (RHEL, SUSE) is supported

Red Hat Enterprise Linux 5.0: 64-bit kernel is required for Power and zSeries; 2.6.18 kernel level is required

SuSE Linux Enterprise Server 10.0 Service Pack 1: 64-bit kernel is required for Power and zSeries; 2.6.16 kernel level is required

SuSE Linux Enterprise Server 11.0

More Information

- **DB2 Product Page - <http://www-01.ibm.com/software/data/db2/9/>**
- **IFS Product Page - <http://www-01.ibm.com/software/data/infosphere/federation-server/>**
- **IFS Information Center - http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.swg.im.iis.db.prod.fed.nav.doc/dochome/iiypfnav_dochome.html**
- **<http://www.ibm.com/developerworks/>**

Thank You !

