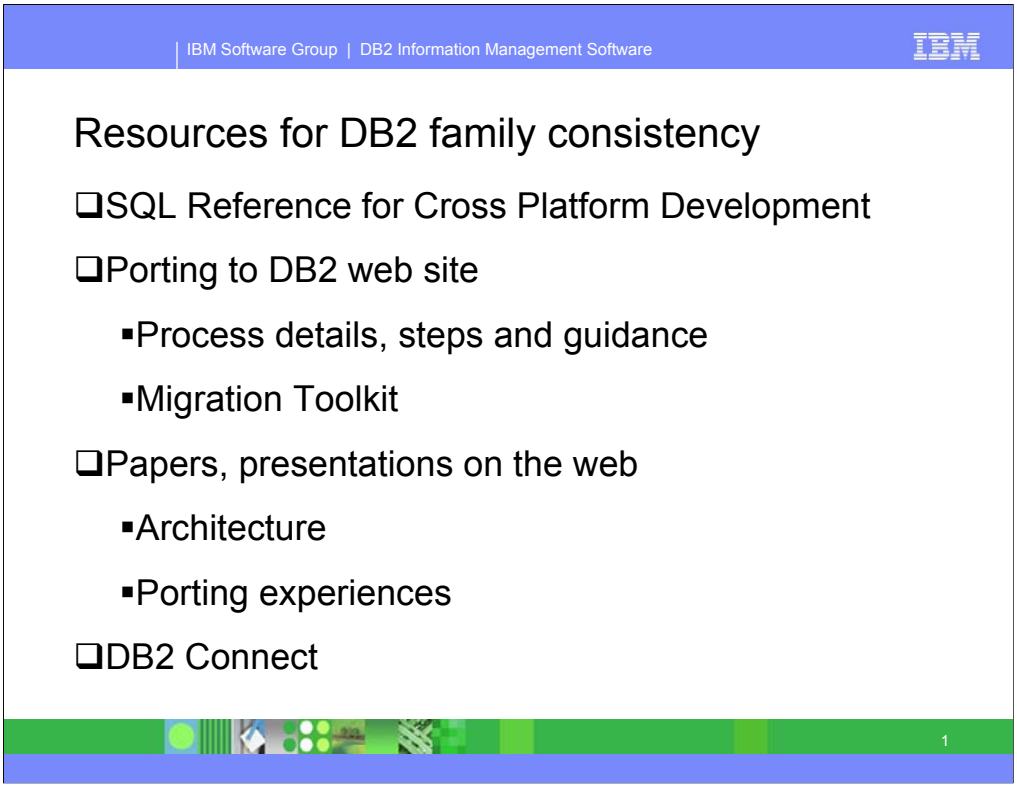

Notes on Resources for migration from DB2 to DB2:

The details of SQL across the DB2 family are provided in the 934 page book, in the SQL References for each version and platform and on the web pages. This presentation has a quick overview, showing how much SQL has become similar, but not the same. So one of the key choices is generally having the target DB2 be as current as possible.

There are a number of presentations showing the architecture, terminology and some porting experiences.

For DB2 Connect, see this information
http://www.ibm.com/software/data/db2/db2connect/

Some resources may come from groups who port, or perhaps from vendors and consultants.

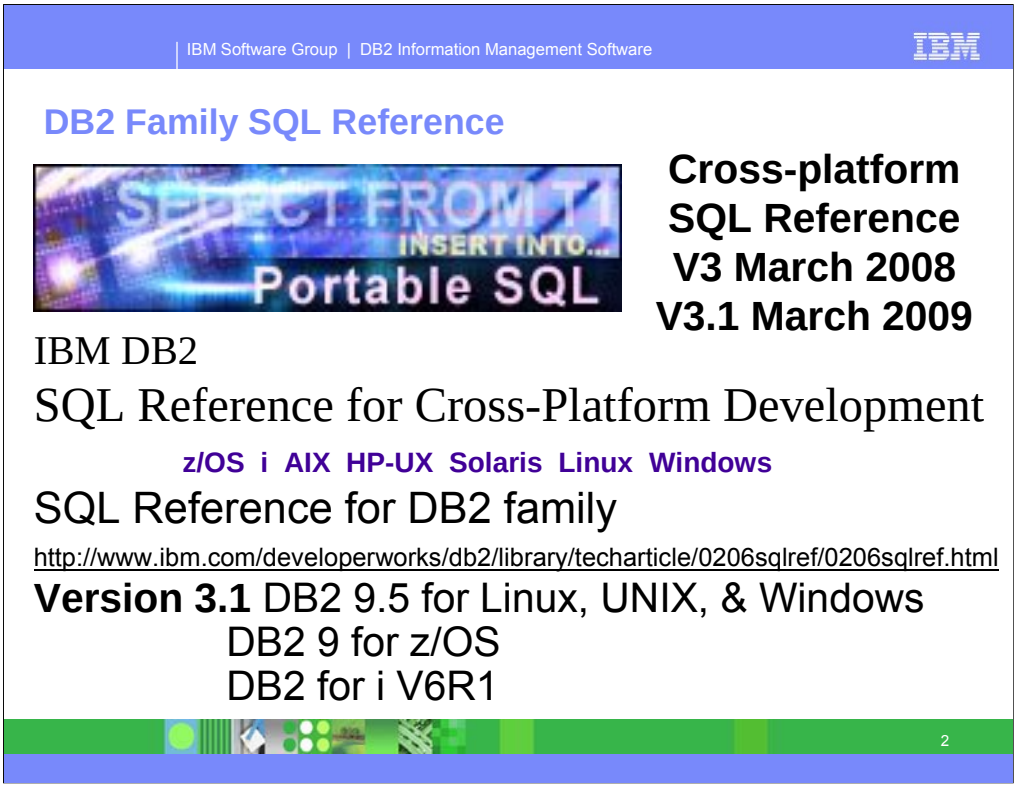

This is the web page for the SQL Reference for Cross-Platform Development. It has four versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well. If you want to write programs that work well across the DB2 family, then this is the book to download and use.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions. The latest book is Version 3.1 for platforms DB2 9 for z/OS, DB2 for I V6.1, and DB2 9.5 for Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

- DECFLOAT data type
- ROW CHANGE expressions
- order-by-clause and fetch-first-clause in a subselect
- SELECT FROM INSERT
- IMPLICITLY HIDDEN columns
- RESTRICT on DROP of functions and procedures
- CURRENT DECFLOAT ROUNDING MODE special register & SET statement
- ALTER FUNCTION
- Additional timestamp format
- New built-in functions
- Full outer join
- row-change-timestamp columns
- Other portability enhancements

Other books are for DB2 for z/OS Version 7 and Version 8.
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv31.pdf
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
http://www.ibm.com/developerworks/data/library/techarticle/db2common/?S_TACT=105AGX11&S_CMP=LP

www.ibm.com/developerworks/db2/zones/porting/

The best web site for application developers who work with DB2 is DeveloperWorks.

There is extensive help for porting to DB2 UDB on the web, with suggestions for the process and many resources, including a migration toolkit. The dramatically improved DB2 family compatibility in DB2 for z/OS V7 and V8 makes porting to DB2 UDB for z/OS from Unix or Windows platforms much easier.

http://www.ibm.com/developerworks/db2/zones/porting/

Here is the URL for the main DeveloperWorks page

http://www.ibm.com/developerworks/db2/

There is also a DeveloperWorks page about DB2 for z/OS.

http://www.ibm.com/developerworks/db2/products/db2zos/

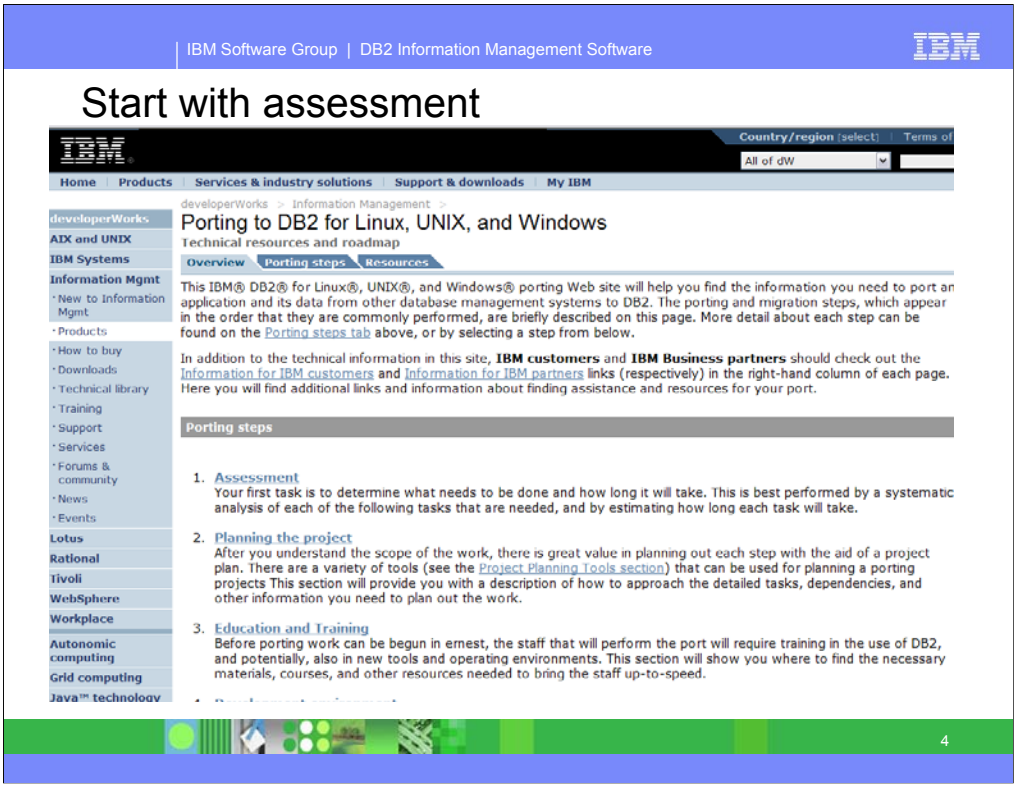

The questions at the bottom of the web page are crucial to communicating and understanding the customer objectives. With the key questions answered, many of the uncertainties are avoided. These

- Primary reason for porting to DB2
- Primary function of the application, Workload type (OLTP, OLAP/DSS, both)
- DBMS, platform, version that will be ported from and to
- Languages, APIs & web technologies (EJB, ASP, JSP, etc.)
- Availability & performance requirements
- Application SQL characteristics, know differences
- Database structure characteristics
- Time frame, People and resources available, experience

http://www.ibm.com/software/data/db2/migration/mtk/

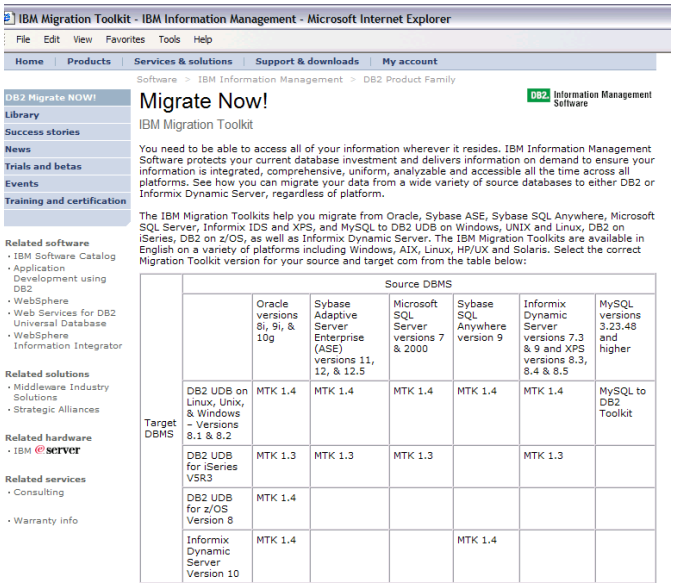

IBM Migration Toolkit - IBM Information Management - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Home | Products | Services & solutions | Support & downloads | My account

Software > IBM Information Management > DB2 Product Family

DB2 Migrate NOW!
Library
Success stories
News
Trials and betas
Events
Training and certification

Related software
- IBM Software Catalog
- Application Development using DB2
- WebSphere
- Web Services for DB2 Universal Database
- WebSphere Information Integrator

Related solutions
- Middleware Industry Solutions
- Strategic Alliances

Related hardware
- IBM @server

Related services
- Consulting
- Warranty info

Migrate Now!

IBM Migration Toolkit

DB2 Information Management Software

You need to be able to access all of your information wherever it resides. IBM Information Management Software protects your current database investment and delivers information on demand to ensure your information is integrated, comprehensive, uniform, analyzable and accessible all the time across all platforms. See how you can migrate your data from a wide variety of source databases to either DB2 or Informix Dynamic Server, regardless of platform.

The IBM Migration Toolkits help you migrate from Oracle, Sybase ASE, Sybase SQL Anywhere, Microsoft SQL Server, Informix IDS and XPS, and MySQL to DB2 UDB on Windows, UNIX and Linux, DB2 on iSeries, DB2 on z/OS, as well as Informix Dynamic Server. The IBM Migration Toolkits are available in English on a variety of platforms including Windows, AIX, Linux, HP/UX and Solaris. Select the correct Migration Toolkit version for your source and target com from the table below:

		Source DBMS					
		Oracle versions 8i, 9i, & 10g	Sybase Adaptive Server Enterprise (ASE) versions 11, 12, & 12.5	Microsoft SQL Server versions 7 & 2000	Sybase SQL Anywhere version 9	Informix Dynamic Server versions 7.3 & 9 and XPS versions 8.3, 8.4 & 8.5	MySQL versions 3.23.48 and higher
Target DBMS	DB2 UDB on Linux, Unix, & Windows – Versions 8.1 & 8.2	MTK 1.4	MTK 1.4	MTK 1.4	MTK 1.4	MTK 1.4	MySQL to DB2 Toolkit
	DB2 UDB for iSeries V5R3	MTK 1.3	MTK 1.3	MTK 1.3		MTK 1.3	
	DB2 UDB for z/OS Version 8	MTK 1.4					
	Informix Dynamic Server Version 10	MTK 1.4			MTK 1.4		

This is the migration toolkit page, which includes a capability to migrate from many other DBMS to DB2 with the 1.4 Migration Tool Kit (MTK). This is not generally needed for migration from DB2 to DB2. Here is an article on using the MTK.

http://www.ibm.com/developerworks/db2/library/techarticle/0209jarzebowicz/0209jarzebowicz.html

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS and OS/390 Version 7, comparing a March 2001 z/OS version with an October 2004 LUW version. V7 has almost no unique function, there is a small set of common function, and a larger set of SQL unique to LUW.

The next step in the process is DB2 for z/OS Version 8. There are three sets of SQL noted above, with none that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

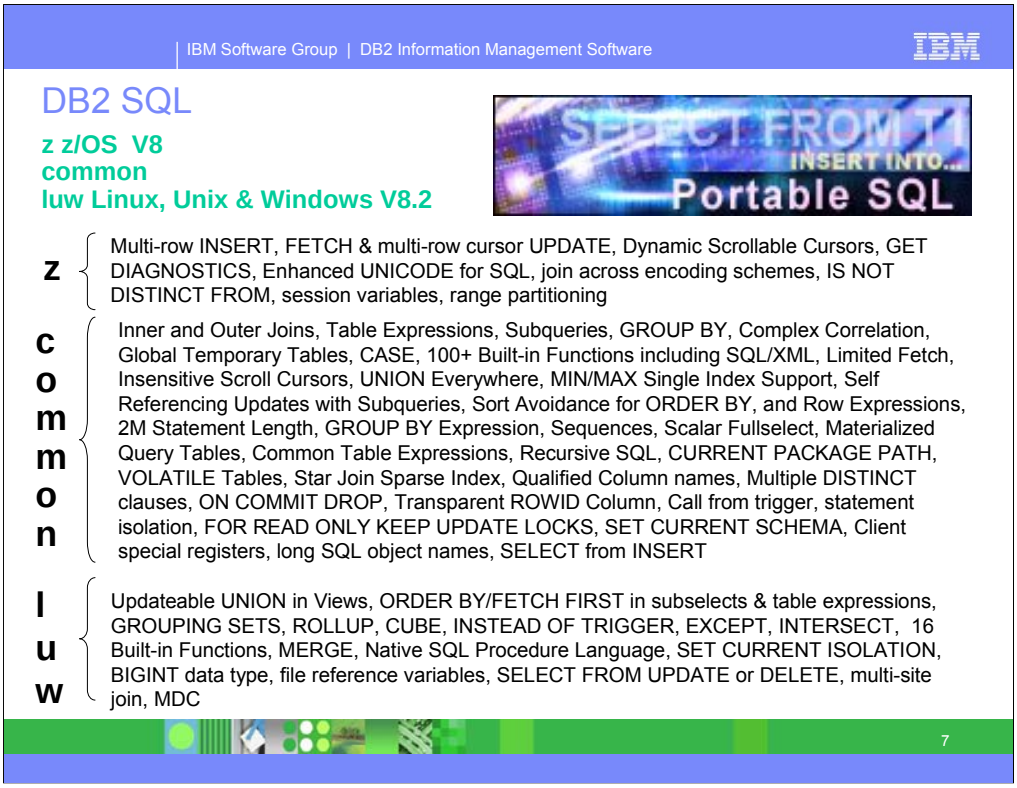

This chart shows the relationship of SQL in the DB2 family comparing DB2 for Linux, Unix & Windows with DB2 for z/OS for key language constructs. This chart compares the z/OS Version 8 from March 2004 with the LUW version from October 2004.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. Sheryl Larsen provided the base for this information, but the mistakes are mine.

If you want to improve DB2 family consistency, then DB2 for z/OS Version 8 is a big step, changing the game from one of catch up to one of leapfrog.

If you want to have a book for SQL across platforms, see the 2004 Cross-Platform SQL Reference.

- Cross-Platform Development, http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

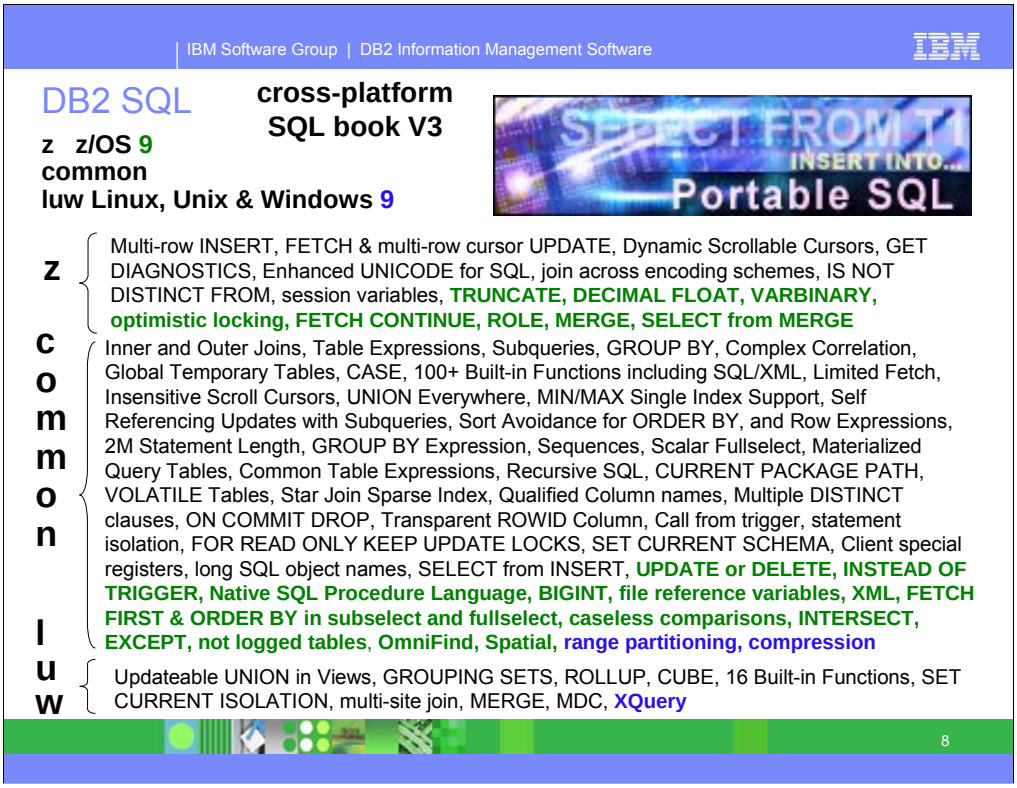

This chart shows the 2007 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS and DB2 9 for LUW. DB2 9 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. We are able to move more from the z list to the common list with Viper. There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

The Cross-Platform SQL Reference Version 3 documents this combination, with DB2 for i5/OS V5R4.

Cross-Platform Development Version 3,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf

This chart shows the 2008 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS, (DB2 9). DB2 9 moved about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, delivered in 2008. We are able to move more from the unique z list to the common list with DB2 9.5 for LUW.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. The changes in a specific version are not consistent. As we introduce new function, sometimes it will be on one platform first, but movement from unique lists into the common list continues to be the strongest trend.

The Cross-Platform SQL Reference Version 3.1 documents this combination, with DB2 for i V6R1.

Cross-Platform Development Version 3.1,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

This text just shows the projected relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. The key at the top has the shorthand, the platform and the level. This step in the process is DB2 X for z/OS and DB2 for luw Y. The projected DB2 X for z/OS and DB2 for luw code named Y are being developed, not announced or generally available for some time. We expect to be able to move more from the z and the luw list to the common list with these changes to DB2.

Three sets of SQL are noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. The changes for a specific version are not consistent. As we introduce new function, often one platform will be first, but movement from unique lists into the common list continues to be the strongest trend.

DB2 family consistency and productivity for those who use SQL have improved a lot in the past few versions. We use the Cross-platform SQL Reference to compare DB2 across the Linux, UNIX, Windows, i5/OS and z/OS platforms. See the 900+ page SQL Reference for Cross-Platform Development, Version 3,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
Watch for a new version of this book to show new SQL being delivered.

DB2 for z/OS handles many more languages than most people think. There are many different interfaces used for the languages to fit the style appropriate to the language. The languages on this chart are not all of those which are supported by or to DB2 for z/OS.

The DB2 for z/OS precompiler works with assembler, C, C++, COBOL, PL/I and Fortran. The DB2 for z/OS coprocessor works with C, C++, COBOL, and PL/I. ODBC or CLI APIs are provided for C and C++. Java has JDBC and SQLJ, with JLINQ on the way. REXX and APL2 interfaces are provided. Other languages and application generators connect to DB2 using call attach or RRS attach for local connections, DRDA or private protocols for remote connections. Other languages use these of APIs or those implemented in DB2 Connect and other client deliveries. Some examples include WebSphere Developer for z, Enterprise Generation Language. Both Microsoft .NET and open source languages use this variety of connections.

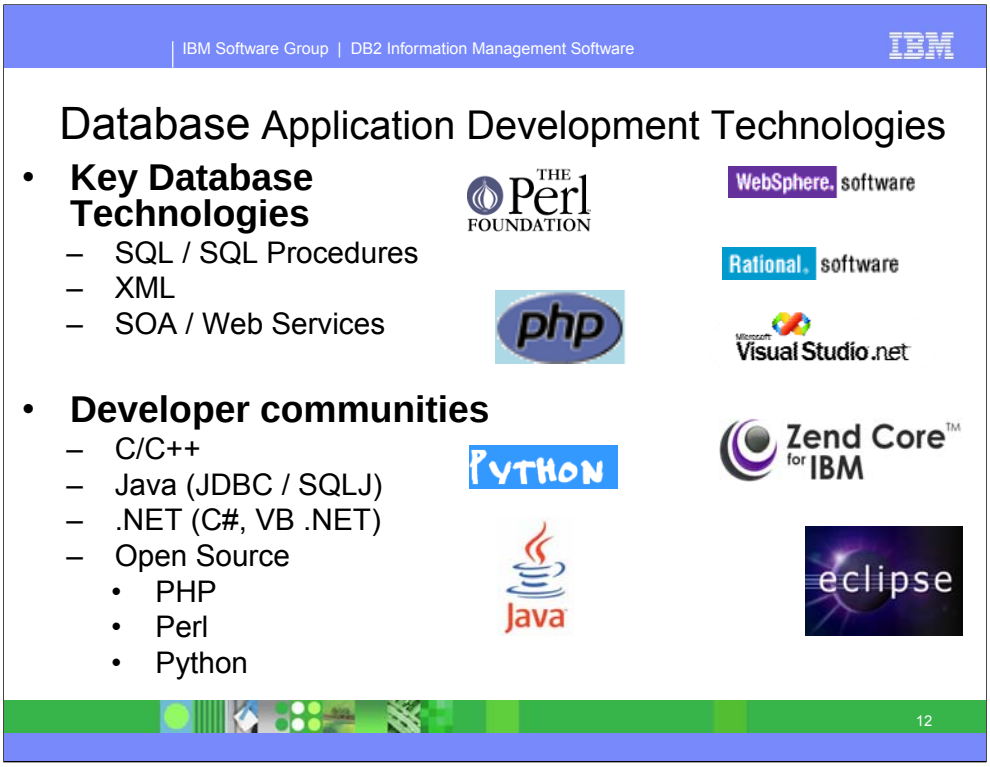

Application programming is using a wider range of tools, environments and languages. The Eclipse framework is growing strongly. We need to connect the new languages and environments to the scale and value of the existing infrastructure. The Rational and WebSphere product lines provide part of the connection, with products like Rational Data Architect and WebSphere Information Integration. DB2 clients provide more support for new environments and new languages.

Integrated Data Management Core Values

An integrated, modular, data management environment designed to increase organizational productivity and effectiveness while improving the quality of service, cost of ownership, and governance of diverse data, databases, and data-driven applications

- Providing end-to-end data lifecycle management
- Facilitating cross-organizational collaboration for business alignment
- Flexibility to provide the ease of use small businesses require with the scalability to manage the large enterprises

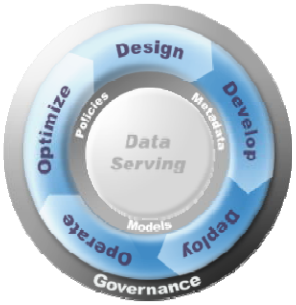

Core Users of the IBM Data Studio

14

Data Studio Packaging

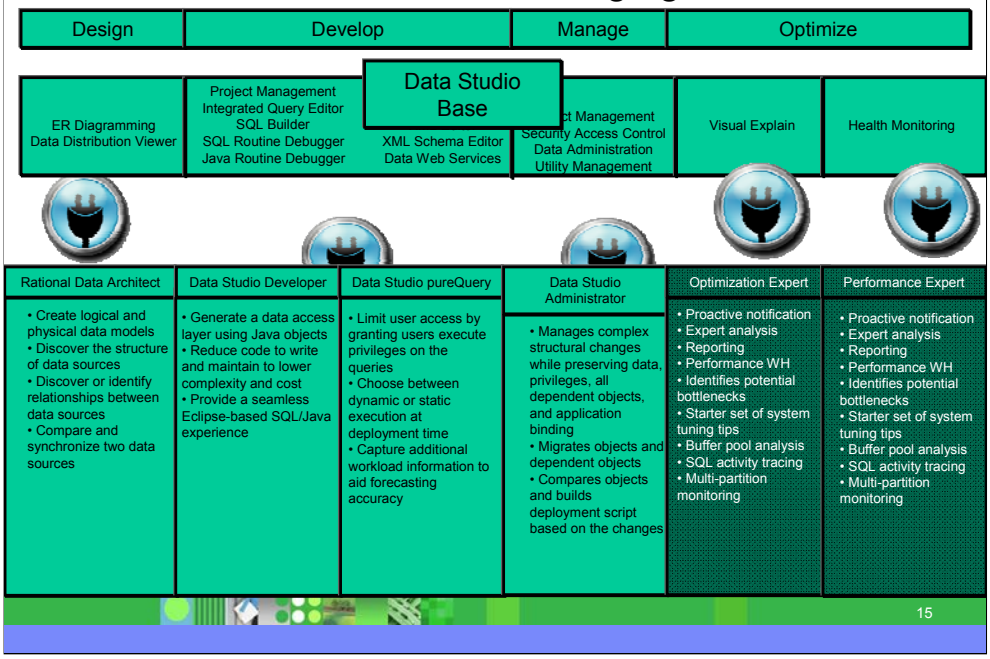

Rational Data Architect	Data Studio Developer	Data Studio pureQuery	Data Studio Administrator	Optimization Expert	Performance Expert
• Create logical and physical data models • Discover the structure of data sources • Discover or identify relationships between data sources • Compare and synchronize two data sources	• Generate a data access layer using Java objects • Reduce code to write and maintain to lower complexity and cost • Provide a seamless Eclipse-based SQL/Java experience	• Limit user access by granting users execute privileges on the queries • Choose between dynamic or static execution at deployment time • Capture additional workload information to aid forecasting accuracy	• Manages complex structural changes while preserving data, privileges, all dependent objects, and application binding • Migrates objects and dependent objects • Compares objects and builds deployment script based on the changes	• Proactive notification • Expert analysis • Reporting • Performance WH • Identifies potential bottlenecks • Starter set of system tuning tips • Buffer pool analysis • SQL activity tracing • Multi-partition monitoring	• Proactive notification • Expert analysis • Reporting • Performance WH • Identifies potential bottlenecks • Starter set of system tuning tips • Buffer pool analysis • SQL activity tracing • Multi-partition monitoring

15

DB2 Developer Workbench vs. Data Studio

before

IBM DB2 Developer Workbench V9.1

- SQL Query Editor
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

Data Studio is a full replacement of DB2 Developer Workbench plus much more

- ***DB2 for Linux, Unix, Windows v8.x, v9.1.x, v9.5***
- ***DB2 for z/OS v7, v8, 9***
- ***DB2 for i v5r2, v5r3, v5r4, v6***
- ***Informix Dynamic Server (IDS) v9.x, v10.x, v11***

now

IBM Data Studio

- Integrated Query Editor – SQL + XQuery
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

- ER Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support
- Health Monitoring DB2 for LUW 9.5 and DB2 z/OS v9

16

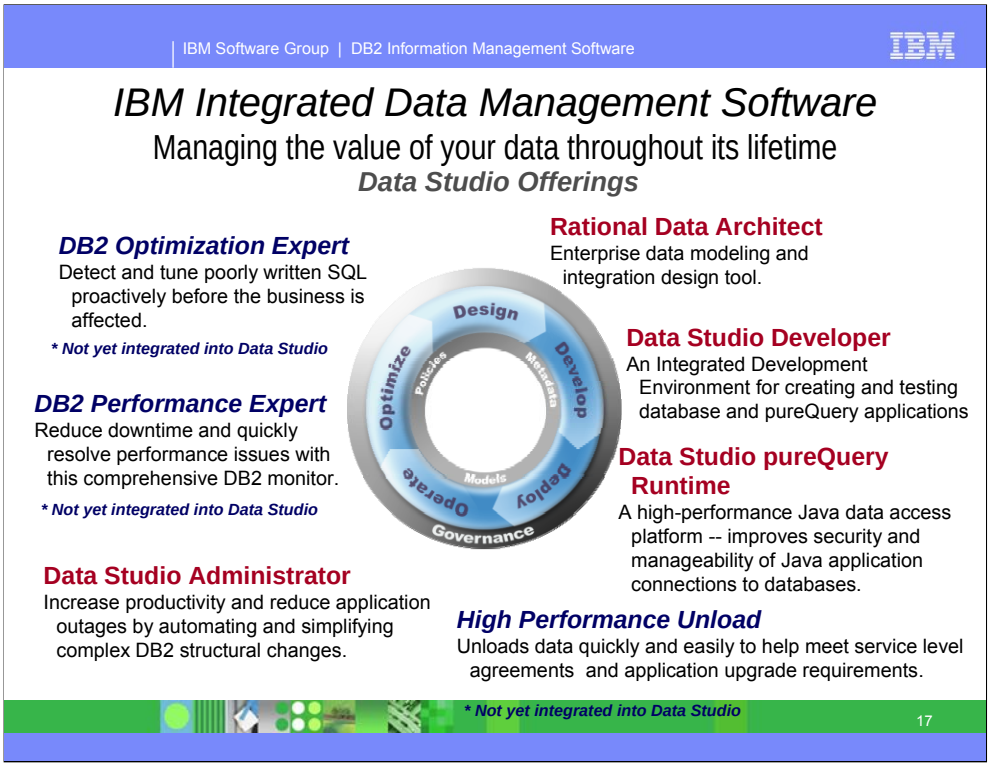

Recent and upcoming Integrated Data Management Innovations include:

Java Database Connectivity

Transaction affinity routing in a Shared Queue environment
This feature gives IMS users more control over where a transaction or group of transactions should be processed. The end result is that IMS users have direct control over shared message queue workload balance and can better manage availability.

Integration between dashboard and threshold processing
Users can set threshold values based on dashboard definitions. Threshold exceptions are logged in history database for after the fact viewing.

More status data for RM structure entries
Users can make more informed decisions about when it is safe to delete entries. Multiple delete capability is added for mass removal of structure entries.

- **DB2 Accessories Suite for z/OS 1.3**
  Data Studio Administration Console (DSAC)

  The Data Studio Administration Console (DSAC) is a component of Data Studio available for Web-download only, that provides:

  A rich Web interface for

  Database health monitoring and troubleshooting

  Q replication and event publishing monitoring and management

  - Support for DB2 on z/OS, Linux, UNIX and Windows

Archive file maintenance and

Audit security functions

17

Rational Data Architect

Rational Data Architect is a collaborative, data design solution to discover, model, relate, and standardize diverse data assets.

- Increase Data Quality and Integrity
  - Discover, explore, and visualize the structure of data sources
  - Analyze and enforce compliance to enterprise standards
  - Support business and IT collaboration via a common business glossary
  - Use with IBM Industry Models for industry-specific best practices
  - Facilitate model-driven development via seamless integration with Rational Software Delivery Platform
  - Automate transformations between the application model and the data model

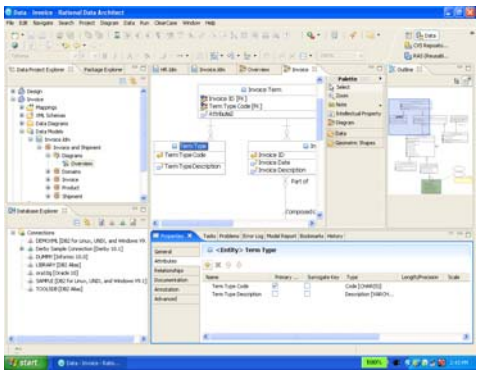

18

IBM Data Studio Developer and pureQuery Runtime

IBM Data Studio Developer is an integrated database development environment that speeds application design, development, and deployment while increasing data access performance and manageability.

- Ease JAVA coding
  - pureQuery code assistance
- Improve predictability and manageability with static SQL
  - Switch dynamic to static SQL without changing a line of code
- Visualization of JAVA request to SQL code
  - SQL Outline correlates SQL to JAVA code

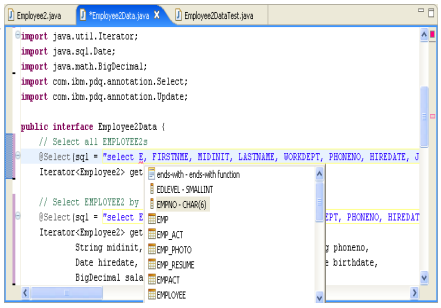

Java Data Access – many forms

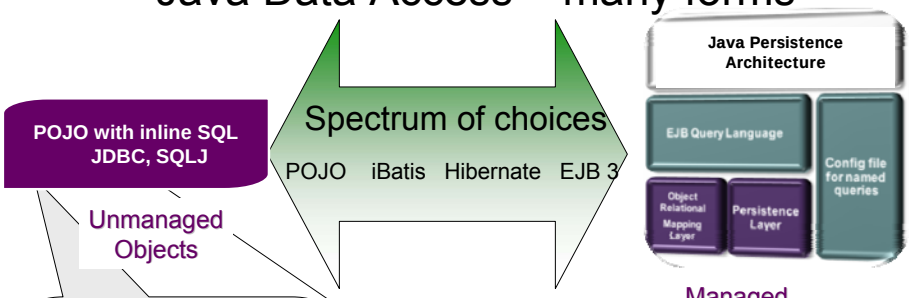

Unmanaged Objects

Pro's:
- Simplicity
- Easy to control SQL
- Good performance
- Good monitoring (SQLJ)

Con's:
- Not tied to object model
- More work for app pgmr

Managed Objects

Pro's:
- Less work for app pgmr
- Access via OO business objects

Con's:
- Complexity
- Less control over SQL issued
- Performance can suffer
- Very difficult to monitor or diagnose problems

20

pureQuery Outline

Speed up problem isolation for developers – even when using frameworks

- Capture application-SQL-data object correlation (with or without the source code)
- Trace SQL statements to using code for faster problem isolation
- Enhance impact analysis identifying application code impacted due to database changes
- Answer “Where used” questions like “Where is this column used within the application?”
- Use with modern Java frameworks e.g. Hibernate, Spring, iBatis, OpenJPA

Data Studio pureQuery Runtime for z/OS

- In-house testing shows double-digit reduction in CPU costs over dynamic JDBC

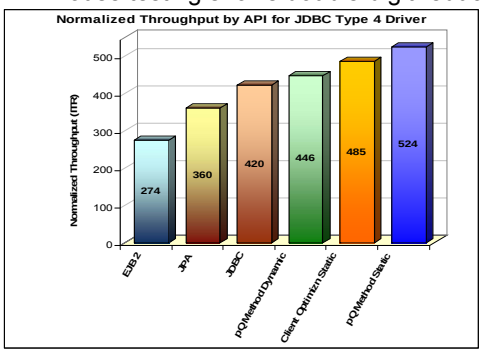

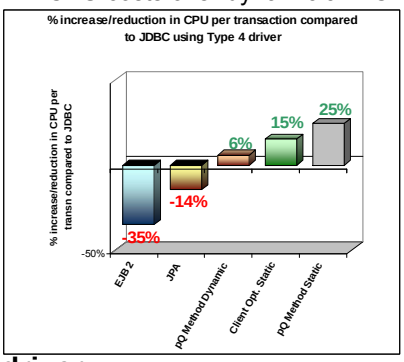

- IRWW – an OLTP workload, **Type 4 driver**
- Cache hit ratio between 70 and 85%
- **15% - 25% reduction on CPU per transaction** over dynamic JDBC

IBM Data Studio Administrator

IBM Data Studio Administrator improves DBA productivity and reduces application outages by automating and simplifying complex DB2 structural changes.

- Manages changes while preserving data, privileges, all dependent objects, and application binding
- Synchronizes, copies, clones, or merges database schema definitions from the source to the target
  - Change in place and object migration
- Documents changes for collaboration and audit

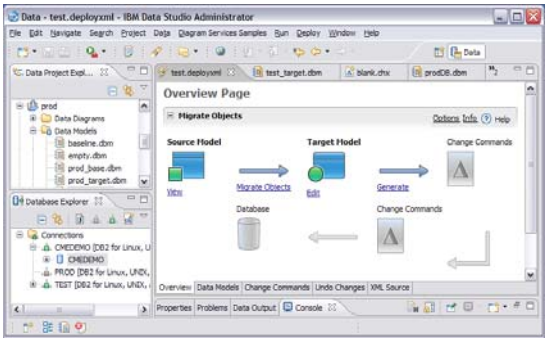

- Integrates with Rational V7 software for easy collaboration

Differentiated features:

-We can customize changes and migrate in a single change

-We have the most flexible data preservation options

-We can incorporate RDA physical data models

23

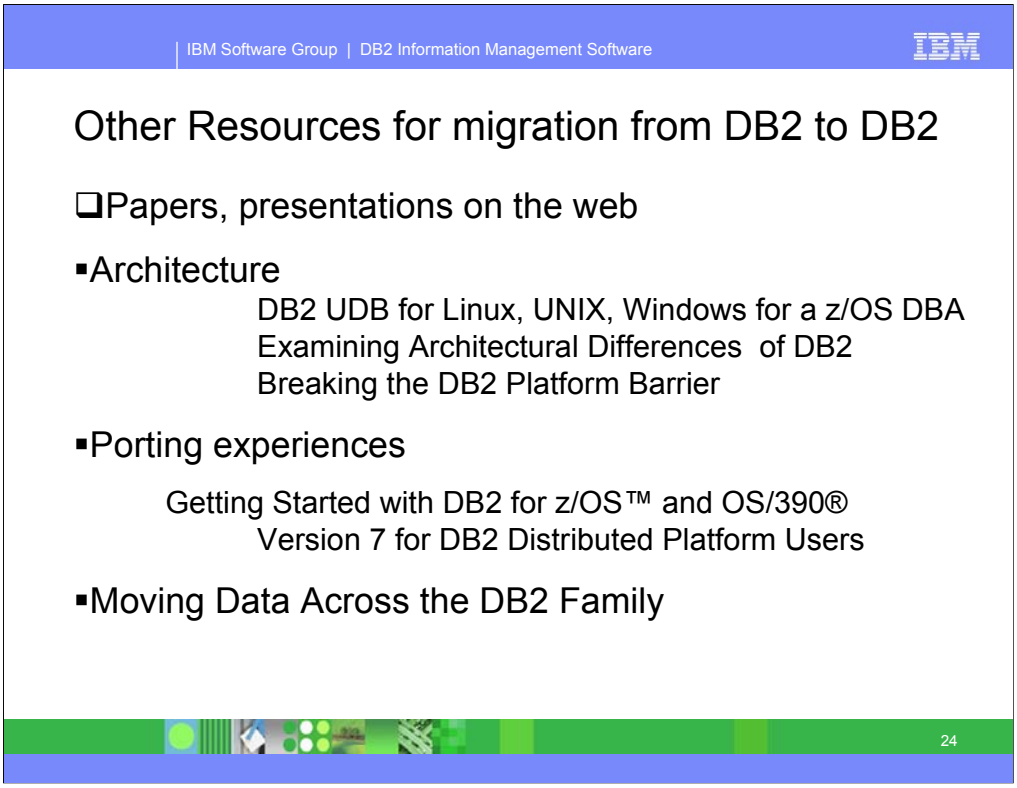

•DB2 UDB for Linux, UNIX, Windows for a z/OS DBA: Chris Eaton, IDUG Berlin 2005

•Examining Architectural Differences of DB2 on OS/390 vs. UNIX/NT, Jim Wankowski, Quest

•Breaking the DB2 Platform Barrier, Jim Wankowski, Quest

•Getting Started with DB2 for z/OS™ and OS/390® Version 7 for DB2 Distributed Platform Users, Raul Chong

http://www.ibm.com/developerworks/db2/library/techarticle/0207chong/0207chong2.html

http://www.ibm.com/developerworks/db2/library/techarticle/0207chong/0207chong.html

•A colorful introduction to DB2 UDB, Version 8 for UNIX, Linux, and Windows: Visualize fundamental DB2 concepts, Raul Chong

http://www.ibm.com/developerworks/db2/library/techarticle/0301chong/0301chong.html

Moving Data Across the DB2 Family redbook, SG24-6905

http://www.redbooks.ibm.com/abstracts/SG246905.html?Open

http://www.redbooks.ibm.com/redbooks/pdfs/sg246905.pdf

DB2 Cross-Platform SQL Reference

updated March 2009

ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

V3.1 DB2 for LUW 9.5
DB2 for i V6.1
DB2 9 for z/OS

New SQL features now included

- DECFLOAT data type
- Additional timestamp format
- ROW CHANGE expressions
- New built-in functions
- order-by-clause and fetch-first-clause in a subselect
- SELECT FROM INSERT
- Full outer join
- IMPLICITLY HIDDEN columns
- row-change-timestamp columns
- RESTRICT on DROP of functions and procedures
- CURRENT DECFLOAT ROUNDING MODE special register & SET
- ALTER FUNCTION
- Other portability enhancements

This is the web page for the SQL Reference for Cross-Platform Development. It has four versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well. If you want to write programs that work well across the DB2 family, then this is the book to download and use.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions. The latest book is Version 3.1 for platforms DB2 9 for z/OS, DB2 for I V6.1, and DB2 9.5 for Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

- DECFLOAT data type
- Additional timestamp format
- ROW CHANGE expressions
- New built-in functions
- order-by-clause and fetch-first-clause in a subselect
- SELECT FROM INSERT
- Full outer join
- IMPLICITLY HIDDEN columns
- row-change-timestamp columns
- RESTRICT on DROP of functions and procedures
- CURRENT DECFLOAT ROUNDING MODE special register & SET statement
- ALTER FUNCTION
- Other portability enhancements

Other books are for DB2 for z/OS Version 7 and Version 8.
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv31.pdf
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
http://www.ibm.com/developerworks/data/library/techarticle/db2common/?S_TACT=105AGX11&S_CMP=LP

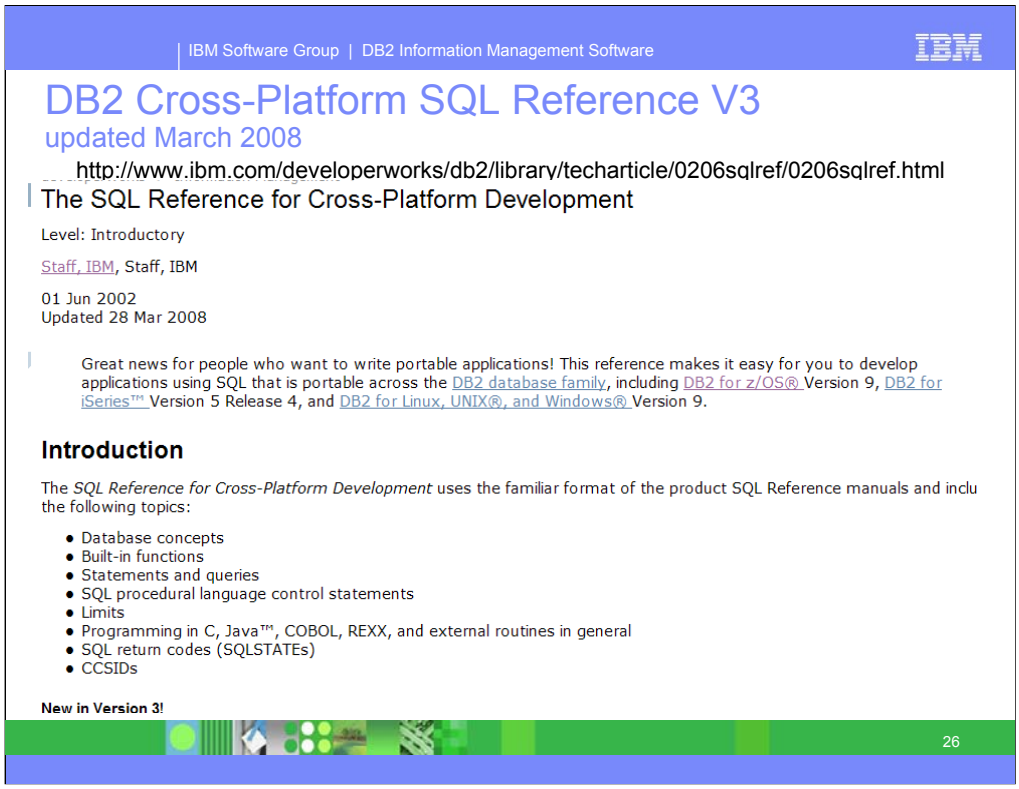

This is the web page for the SQL Reference for Cross-Platform Development. It has three versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well. If you want to write programs that work well across the DB2 family, then this is the book to download and use.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions. The current book is Version 3 for platforms Version 9 of DB2 for z/OS, Version 5 Release 4 of DB2 for i5/OS, and Version 9 of DB2 for the Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

- BIGINT data type
- New built-in functions
- Recursive common table expressions
- Additional common features for ALTER COLUMN on ALTER TABLE
- Table partitions
- SQL PL extensions including the FOR statement and support for nested ATOMIC compound statements
- Other portability enhancements
- OLAP specifications
- INTERSECT and EXCEPT
- INSTEAD OF triggers

Other books are for DB2 for z/OS Version 7 and Version 8.

http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf

Get updated books **latest April 2009**

http://publib.boulder.ibm.com/infocenter/imzic/

- ✓Program Directories
- ✓Installation Guide
- ✓Redbooks
- ❑Data Sharing: Planning and Administration
- ❑Application Programming Guide & Reference for JAVA™
- ❑RACF Access Control Module Guide
- ❑ODBC, Spatial, …
- ✓Messages
- ✓Codes
- ✓Administration Guide
- ✓Performance Guide
- ✓Utility Guide and Reference
- ✓Internationalization Guide (Unicode)

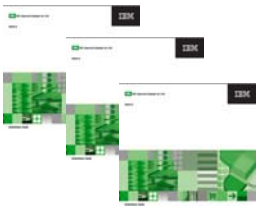

For installation, you need many books. Some are optional, for example the data sharing book is not needed if you don't use data sharing. If you use data sharing, then part of the install process is in that book. Part of the Java install process is in the Java book. Part of the ODBC install process is in the ODBC book. If you use RACF access control, then you need the RACF Access Control Module Guide. You can get most of the books from the DB2 Library web page. The books were updated in December 2007 and February, March, June, and August 2008, with some coming later, so get the latest ones. Some of the Redbooks may be helpful. You may need books from the z/OS Library as well.

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011658

http://www.ibm.com/systems/z/os/zos/bkserv/r9pdf/

Be sure to use the latest information to save time and problems. Some of the IBM Redbooks publications have always been updated and added lately (next page).

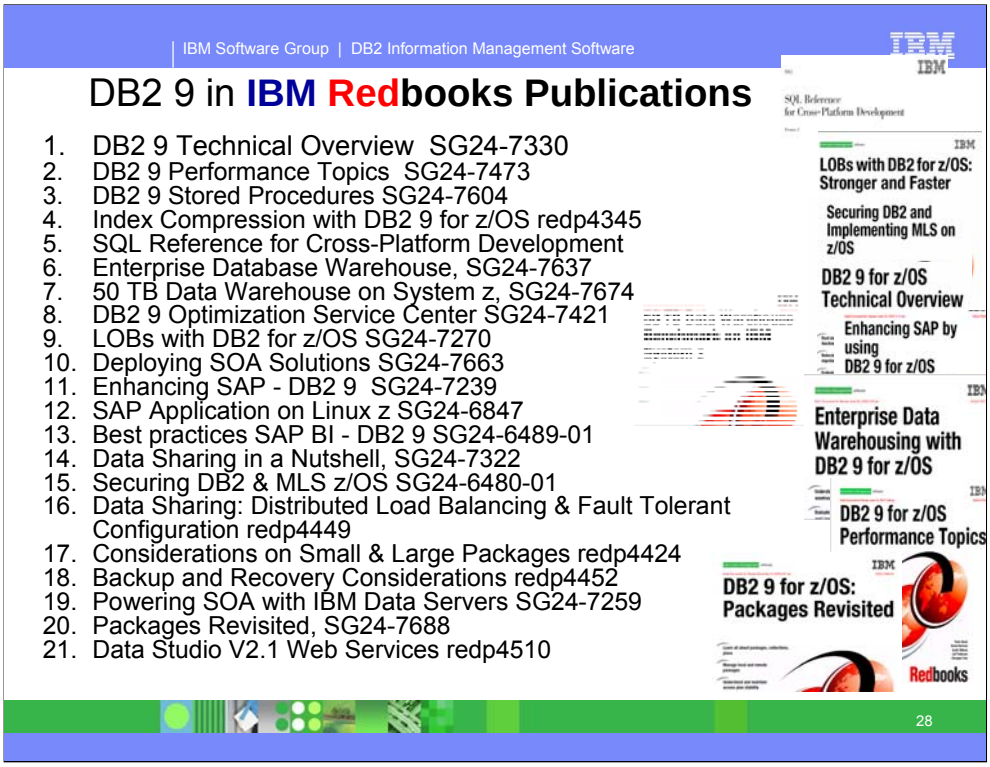

DB2 library more information http://www.ibm.com/software/data/db2/zos/library.html

Many IBM Redbooks publications, Redpapers and one cross-platform book on DB2 9 are published, in addition to the standard library, with more in the works. Check for updates.

1. DB2 9 Technical Overview, SG24-7330 http://www.redbooks.ibm.com/abstracts/SG247330.html
2. DB2 9 Performance Topics, SG24-7473, http://www.redbooks.ibm.com/abstracts/SG247473.html
3. DB2 9 Stored Procedures, SG24-7604, http://www.redbooks.ibm.com/abstracts/SG247604.html
4. Index Compression DB2 9, REDP4345, http://www.redbooks.ibm.com/abstracts/redp4345.html
5. Deploying SOA Solutions SG24-7663, http://www.redbooks.ibm.com/abstracts/SG247259.html
6. Cross-Platform Development Version 3, http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
7. Enterprise Data Warehousing, SG24-7637, http://www.redbooks.ibm.com/abstracts/sg247637.html
8. LOBs: Stronger & Faster SG24-7270, http://www.redbooks.ibm.com/abstracts/SG247270.html
9. Securing DB2 & MLS z/OS, SG24-6480-01, http://www.redbooks.ibm.com/abstracts/sg246480.html
10. Enhancing SAP, SG24-7239, http://www.redbooks.ibm.com/abstracts/SG247239.html
11. Best practices SAP BI, SG24-6489-01, http://www.redbooks.ibm.com/abstracts/sg246489.html
12. Optimization Service Center, SG24-7421, http://www.redbooks.ibm.com/abstracts/sg247421.html
13. Data Sharing in a Nutshell, SG24-7322, http://www.redbooks.ibm.com/abstracts/sg247421.html
14. DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration http://www.redbooks.ibm.com/abstracts/redp4449.html
15. DB2 for z/OS: Considerations on Small and Large Packages redp4424 http://www.redbooks.ibm.com/abstracts/redp4424.html
16. DB2 9 for z/OS: Backup and Recovery Considerations redp4452 http://www.redbooks.ibm.com/abstracts/redp4452.html
17. Powering SOA IBM Data Servers, SG24-7259 http://www.redbooks.ibm.com/abstracts/SG247259.html
18. DB2 9 for z/OS Packages Revisited, SG24-7688 http://www.redbooks.ibm.com/abstracts/SG247688.html
19. 50 TB Data Warehouse Benchmark on IBM System z http://www.redbooks.ibm.com/redpieces/abstracts/sg247674.html
20. SAP on DB2 9 for z/OS: Implementing Application Servers on Linux for System z http://www.redbooks.ibm.com/redpieces/abstracts/sg246847.html
21. IBM Data Studio V2.1: Getting Started with Web Services on DB2 for z/OS http://www.redbooks.ibm.com/redpieces/abstracts/redp4510.html
22. Parallel Sysplex Operational Scenarios http://www.redbooks.ibm.com/redpieces/abstracts/sg242079.html
23. Watch for titles on DB2 distributed; serialization & concurrency; utilities

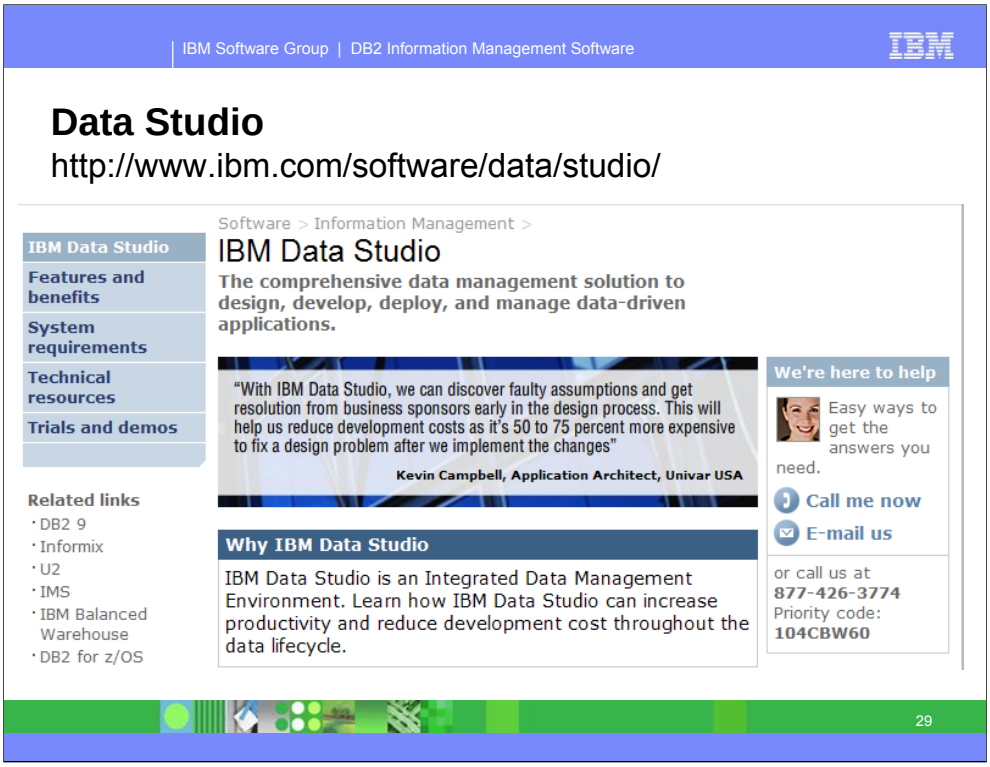

See this page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast. IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio Developer: An Integrated Development Environment for creating and testing database and pureQuery applications.

IBM Data Studio pureQuery Runtime: A high-performance Java data access platform -- improves security and manageability of Java application connections to databases.

http://www.ibm.com/software/data/studio/

ftp://ftp.software.ibm.com/software/data/db2zos/IOD1298_ADadminTrendsDirections_CotnerOct2007.pdf

DB2 Connect makes your company's host data directly available to your personal computer and LAN-based workstations. It connects desktop and palm-top applications to your company's mainframe and minicomputer host databases for access to your enterprise information no matter where it is. **DB2 Connect** provides the application enablement and robust, highly scalable communication infrastructure for connecting Web, Windows, UNIX, Linux and mobile applications to z/OS and AS/400 data. **DB2 Connect** is included in many of the DB2 products.

http://www.ibm.com/software/data/db2/db2connect/

More Notes

- See the detailed notes below for other resources and web pointers.
- http://www.redbooks.ibm.com/abstracts/sg247048.html
- http://www.redbooks.ibm.com/abstracts/tips0674.html

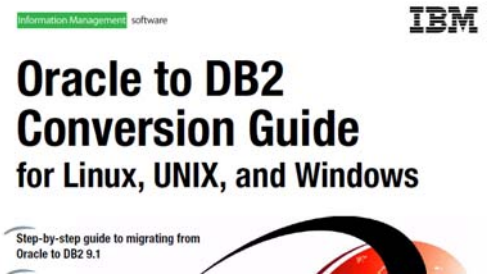

In general, if you are looking for the best portability, then you should be using the one year old DB2 9, rather than the four year old DB2 V8. Here are some additional resources:

Presentation on porting:

SQL Reference for Cross-Platform

ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv2.pdf

Porting web sites:

http://www.ibm.com/developerworks/db2/zones/porting/index.html

http://www.ibm.com/developerworks/db2/products/db2zos/index.html

Key DB2 for z/OS platform differences http://publib.boulder.ibm.com/epubs/pdf/dsnitk11.pdf

DeveloperWorks papers:

http://www.ibm.com/developerworks/db2/library/techarticle/0207chong/0207chong2.html

http://www.ibm.com/developerworks/db2/library/techarticle/0207chong/0207chong.html

http://www.ibm.com/developerworks/db2/library/techarticle/dm-0712gosavi/index.html

http://www.ibm.com/developerworks/db2/library/techarticle/dm-0408whitlark/index.html

Jim Wankowski papers, presentations and articles

http://www.quest-pipelines.com/newsletter-v6/db2_0305.pdf

http://www.quest.com/quest_central_for_db2/pdfs/Crossing_over_DB2.pdf

http://www.quest.com/Quest_Site_Assets/WhitePapers/WPD_BreakingDB2PlatfBarrier_F.pdf