



## Level Set

SQL: special-purpose programming language designed for managing data in a relational database management system (RDBMS)

SQL/PL: A set of SQL statements introduced in DB2 UDB Version 7. Provides procedural constructs necessary for implementing control flow logic around traditional SQL queries and operations.

Supports comprehensive high-level programming in SQL.

SQL PL is a subset of the SQL Persistent Stored Modules (SQL/PSM) language standard. The specification of the current SQL/PSM standard can be found in ANSI/ISO/IEC 9075-4:1999 Information Technology, Database Language SQL, Part 4: Persistent Stored Modules (SQL/PSM).

Native vs. External:

Native: Definition + Executable are stored in the DBMS

External: Definition in DBMS, Executable in File System

- zOS Library (COBOL)

- Unix file system (Java)



Today we'll talk about 3 major topics:

1. How we convinced Java developers to give SQL/PL native procedures a try.

2. What were the next steps after our DB2 access technologies throwdown, aka the Java DB2 Access Options Project

3. What is the current status of Shelter's adoption of SQL/PL Native Procedures?





1. With this much buzz in the DB2 community about SQL/PL stored procedures, how could Java developers were so skeptical?

2. And how did we convince them to give SQL/PL stored procedures a try?

- 3. This presentation is really about two things:
- Our "discovery" of SQL/PL native stored procedures at Shelter when we were looking for alternatives to JDBC that would provide more predictable, governable performance then dynamic SQL and
- The story of taking a technology much touted by IBM and DB2 experts and meeting Shelter's requirements for enterprise-quality technology adoption.



- 1. We began the Java DB2 Access Options Project in 2009.
- 2. The landscape at Shelter at the time included these factors:
- An interest in services and service-oriented architecture, including answering the question, What criteria should be used to judge the effectiveness of a particular technology used for service enablement?
- A few highly visible production issues where the performance of dynamic SQL coming from distributed Java applications surprised everyone (in a bad way) and people were put under great pressure to solve those issues in live fire .

(Raise your hand if you've been a witness to something similar!)

 These two factors influenced the business objectives for the Java DB2 Access Options Project



- 1. Why did we choose these technologies to evaluate?
- JDBC had been in use at Shelter since Shelter developed its first browser-based web application server application, around 1999
- COBOL stored procedures were adopted in the mid-2000s
  - We began exploring COBOL stored procedures because the approach chosen for the early JDBC applications was to replicate IMS/DB or denormalized DB2 structures into more normalized DB2 tables
  - The process of replication was becoming prohibitive and we wanted an alternative
- We wanted to check out the WSDL-based interface for CICS Web Services. Shelter did <u>not</u> have a large inventory of CICS transactions that it was interested in reusing.
- We heard a good deal about SQLJ in the mid- to late 2000s. We were interested in it because it allowed Java developers to stay within the Java API for SQL development (like JDBC) but with the potential advantage of the static SQL model.
- For the same reason we wanted to check out pureQuery (and of course IBM encouraged us to evaluate it).



- 1. [At this point perhaps half the room is cheering or at least laughing and the other half is hissing.]
- 2. Does it sound like I have a chip on my shoulder about Java?!
- Quite the contrary! I am a fan of Java and am even paying my own money to take online classes to learn Java! I want to be respected and look people straight in the eye and say words like "interface" "encapsulation" "inheritance" and "iterate." Oh, and don't forget "method signature," "primitive," and "enumeration." Anyway ....



- Frameworks (such as Spring and Hibernate) are extremely important to Java developers. One of the advantages Java developers see in frameworks is that they can "introspect" relational tables and map tables and columns to objects. This is called object-relationship mapping ("ORM"). Apparently, frameworks do not support ORM of stored procedure result sets. My observation is that Java developers place high value on frameworks because frameworks are perceived to enable speed of development; and that, at least for some Java developers, speed of development can be valued above efficiency of execution.
- Java developers have also objected to COBOL stored procedures because they are too "black box" – "I don't know what's going on inside it, so I don't trust it." Java developers place a high value on the effectiveness of Java methods of unit testing. Apparently JDBC works well with Java unit testing, and stored procedures are perceived to not work as easily.
- 3. With SQL/PL stored procedures, we wanted to be able to provide a stored procedure vehicle that Java developers would trust as much as they trust their own Java code.



- 1. Perhaps we cannot address all the philosophical questions, but we <u>can</u> measure performance. To do this, we built a Test Harness application to drive our DB2 Access Options and enable fair and equal comparisons.
- 2. Test scripts
  - Run this combination of DB2 Access Options this many times.
  - Record results, start / stop times. That enabled us to correlate to SMF data for performance reporting.



- 1. This diagram shows a conceptual diagram of the application we built to measure the 10 different access technologies: The Test Harness
- 2. Notice from the diagram:
  - 1. Test Harness on Web Sphere Application Server distributed platform in our case, Windows
  - 2. DB2 for z/OS
    - 1. JDBC talks directly to DB2 for z/OS tables
    - 2. JDBC talks directly to SQL/PL native stored procedures, which run within DB2
  - 3. External stored procedure address space (COBOL, JDBC, SQLJ, SQL/PL external stored procedures)
    - WLM-managed
    - Test Harness calls SP via JDBC. Stored Procedures talk to DB2 for z/OS tables
  - 4. CICS Web Services
    - 1. Test Harness speaks WSDL to CICS Web Services
    - 2. CICS region talks to DB2 for z/OS tables
  - We used different DB2 AUTHIDs depending on whether the DB2 Access Method used static or dynamic SQL
  - For dynamic SQL, the BINDer of the SQL statement requires SELECT, INSERT, UPDATE, DELETE authority to the DB2 object based on the statement it is BINDing
    - \_\_\_\_\_\_



ITCAM = IBM Tivoli Composite Application Manager



Excerpt from Test Script to put test database in a consistent state:

**Instructions** 

- 1. For each of the 9 DB2 Access Methods, execute Retrieve Test 1 followed by Retrieve Test 2.
- 2. Before each pair of tests:
  - a. Stop spaces
  - b. Start spaces
  - c. Run Access Database command to open spaces
  - d. Run RUNSTATS UPDATE NONE REPORT NO to flush DSC



- 1. We wanted to be able to compare DB2's version of Application Elapsed Time to the application's measurement.
- There was an urban legend floating around that asserted that the application elapsed time did not necessarily correlate to the "mainframe" or DB2 elapsed time. We found that these two measurements were generally correlated although there were a couple of exceptions.

<ul> <li>clientApplicationName property</li> </ul>	Value of Client Information ApplicationName property set by Wrapper	
<ul> <li>Test Harness set unique name for each DB2</li> </ul>	CICSWebServiceDelete COBOLSPDelete	
Access Option it invoked	SQLPLFencedDelete	
<ul> <li>IBM Data Server property</li> </ul>	SQLPLNonFencedDelete	
<ul> <li>Subsequently incorporated into JDBC API</li> </ul>	JDBCDataAccessMethodDelete SQLJInterpretiveDelete	
	SQLJStaticDelete	
	JDBCSPDelete	
	SQLJSPDelete	

- 1. At our shop, the RACF administrators tend to prefer having fewer rather than more RACF userids. RACF userids tend to be at the application level rather than a more granular level.
- 2. This makes it difficult to measure, monitor, and isolate a particular function within a large application.
- 3. What we were trying to accomplish in this project was what we get easily with statically bound applications, where it's very easy to measure/monitor at the DB2 package level.
- 4. We discovered the clientApplicationName property of the IBM data server driver. A Java application can set this value which is stored in DB2's accounting data. In turn, DB2 Accounting Reports can then be used to group reporting by this property. The information how to set this property is documented in the DB2 Application Programming Guide for Java.
- 5. The ability to set this property was subsequently incorporated into many of Shelter's Java DB2 applications since it was so useful in allowing us to report on more granular pieces of our applications.
- 6. In a subsequent project, one of our developers discovered that this (formerly IBM proprietary) feature was moved into JDBC API.



- 1. We measured 3 components of Elapsed Time
  - 1. Test Harness, DB2 Class 1, and DB2 Class 2
  - 2. See DB2 Elapsed Time chart on next slide
- 2. CPU time
  - 1. We wanted to see how DB2 CPU time compared among the three options
  - 2. We were not interested in comparing the different application languages
  - or transaction managers.
  - 3. We were interested in seeing which of our applications used the General Purpose Processor the least,
  - so we computed the percentage use of general processor for each DB2 Access Option



- 1. This is the definitive IBM diagram explaining the components to user response time for DB2 applications
- 2. Class 1 (Application Elapsed Time) is shown in green
  - 1. This includes the time from when the application connects to DB2 until it disconnects.
  - 2. Class 1 time includes Class 2 time
- 3. Class 2 (DB2 Elapsed Time) is show in blue
  - 1. This is the time spent executing SQL statements



	etrics		2	~		
	Total (lower better)	rethene	a tagreed	Capsed DB2 Eager	Junerore	GP-USAF
SQL/PL Native	6360	370	284	275	5135	296
SQL/PL External	6497	272	226	217	5249	533
SQLI Interpretive	6557	354	319	288	5281	316
COBOL SP	6604	280	238	225	5329	532
SQLI SP	6710	396	269	226	5285	534
JDBC	7081	331	370	344	5757	279
CICS Web Services	7120	516	260	236	5508	600
JDBC SP	7561	429	346	255	5985	546
Sector and a sector secto	76.77	4.8.4	201	267	COEC	224

## Elapsed Time Metrics

Test Harness Elapsed Time

Average elapsed time in milliseconds (1/1,000 of a second) for the execution of a transaction

Calculated by the Test Harness application using the Java API and averaged using Microsoft Excel functions

Mainframe Elapsed Time (DB2 Accounting Class 1 Elapsed)

Average elapsed time in milliseconds

Time application is connected to DB2

DB2 Elapsed Time (DB2 Accounting Class 2 Elapsed)

Average elapsed time in milliseconds

Time spent performing SQL statements



			2	7	1	1
	Total (lower is better)	Testhane	starsed	capsed	o Canada	GR UPAT
SQL/PL Native	6360	370	284	275	5135	296
SQL/PL External	6497	272	226	217	5249	533
SQLI Interpretive	6557	354	319	288	5281	316
COBOL SP	6604	280	238	225	5329	532
SQLI SP	6710	396	269	226	5285	534
IDBC	7081	331	370	344	5757	279
CICS Web Services	7120	516	260	236	5508	600
IDBC SP	7561	429	346	255	5985	546
COLL Statia	7577	441	391	357	6056	331

## **CPU** Consumption Metric

Mainframe Service Units

A measure of the CPU used performing SQL statements in DB2

Service units are a way of measuring CPU usage in z/OS in a consistent way that is independent of processor model or CPU speed.

CPU time varies across processor model, while number of service units should remain constant across process models for a given workload

DB2 Accounting Long Report, DB2 Class 2 Service Units

The lower the Mainframe Service Units, the less CPU the application uses

Therefore, Mainframe Services Units are a measure of <u>efficiency</u> (CPU resource usage) or <u>performance</u>



1. Use of specialty engines can be important to enterprises using z/OS because of their cost model.

2. An enterprise pays for a specialty engine once. There is no capacitybased licensing model for specialty engines as there is for general purpose processes.

3. A capacity-based licensing model can be thought of as a variable cost because as the enterprise increases (or decreases)

its available mainframe computing capacity, software licensing charges will also increase (or decrease).

4. On the other hand, work running on the specialty engine incurs a fixed cost, the one-time cost of procuring the specialty engine processor.

*Note:* I have found there is a tendency to confuse <u>efficiency of execution</u> with <u>cost of computing</u>.

• Remember, regardless of whether work runs on the specialty engine or the general purpose processor, the lower the total Mainframe Service Units, the less CPU the work uses, and the more efficient it is.

complined Score I	Vietrics		-			-
	Total (lower is better)	Testhame	stansed .	Dat teos	A Harrene	Service Unit
SQL/PL Native	6360	370	284	275	5135	296
SQL/PL External	6497	272	226	217	5249	533
SQLI Interpretive	6557	354	319	288	5281	316
COBOL SP	6604	280	238	225	5329	532
SQLI SP	6710	396	269	226	5285	534
JDBC	7081	331	370	344	5757	279
CICS Web Services	7120	516	260	236	5508	600
JDBC SP	7561	429	346	255	5985	546
SOLI Static	7577	441	391	357	6056	331

- 1. Highest possible score = 600 (6 tests \* 100% = 600%)
- 2. A lower score can be considered "better" if your goal is to exploit the zIIP
- 3. Notice:
  - a. CICS had the highest GP usage all its workload ran on the GP Processor.
  - b. JDBC had the lowest GP usage followed by SQL/PL native procedures

ombined Score I	Vietrics					-
	Total (lower is	Testhane	a tagoed	tagood page	ed summerine	Service Unit
SQL/PL Native	6360	370	284	275	5135	296
SQL/PL External	6497	272	226	217	5249	533
SQU Interpretive	6557	354	319	288	5281	316
COBOL SP	6604	280	238	225	5329	532
SQLI SP	6710	396	269	226	5285	534
JDBC	7081	331	370	344	5757	279
CICS Web Services	7120	516	260	236	5508	600
JDBC SP	7561	429	346	255	5985	546
SOLI Static	7577	441	391	357	6056	331

- 1. The idea for a "Combined Score Metric" came from one of my coworkers who had worked for the public schools in our town.
- 2. She told me that teachers would use a similar approach as an alternative to assigning percentages for assignments. The idea is that every assignment receives a certain number of points which are summed to provide an overall student score for a set of assignments for the grading period.
- 3. Here we are combining dissimilar measurements to provide a combined score.
- 4. The overall goal is to be able to provide a measure of "overall goodness"
- 5. Lower is better:
  - Elapsed time
  - Service Units
  - GP Usage

Performance Difference	1	
Compare JDBC to all SPs	Mainframe Service Units	% Better (Worse) than JDBC
SQL/PL Native	5135	11%
SQL/PL External	5249	9%
SQLI SP	5285	8%
COBOL SP	5329	7%
JDBC SP	5985	(4%)

- 1. The Service Units reported here are the DB2 Class 2 Service Units.
- 2. This chart reflects DB2 CPU Usage for JDBC compared to stored procedures.
- 3. The points I want to emphasize are:
  - SQL/PL Native Procedures in our test workload used 11% less CPU than JDBC.
  - COBOL stored procedures in our test workload used 7% less CPU than JDBC.
  - For our test workload where we made very sure to control for all the factors that could sway the results, this is very convincing evidence of the performance advantages of both SQL/PL and COBOL stored procedures over JDBC direct access to tables.
  - JDBC using direct access to tables is highly efficient, but the evidence from our project is that SQL/PL Native Procedures offer a considerable opportunity for CPU savings.



- 1. Here are some factors over and above metrics that influenced our decision:
  - Our research indicated that pureXML would be supported in SQL/PL native procedures.
- 2. SQL/PL external vs native WLM effect
  - There is a WLM queuing effect for external stored procedures.
  - This is discussed in the Peggy Zagelow blog post cited in References.
- Some Java developers were familiar with use of PL/SQL procedures from working in Oracle shops and were receptive to the idea of using a similar technology in DB2.
- 4. Recognition that IBM was not pushing SQLJ as they had 3-5 years earlier
  - Perception that SQLJ was a dead-end technology
- 5. Consideration of pureQuery
  - Shelter does not invest lightly in a new technology
  - Too many hands in the deployment pie
    - DBAs and WebSphere Administrators would need to be involved in code deployment as well as Java developers
    - Specialized skill required not a typical DBA skill
    - "Recording" on WAS; XML file imported into Data Studio, then put back into WAS; custom properties in WAS

















- We initially assumed we would manage SQL/PL source code with Serena Dimensions like we do for Java source code, and that we would use separate products and processes to manage source code and deployment for SQL/PL like we do for Java. In that scenario, we would manage SQL/PL source code with Serena Dimensions, initially deploy SQL/PL stored procedures into UNIT using Data Studio, then DEPLOY into higher testing levels and production using Serena Changeman ZMF.
- 2. In September 2012 Serena came onsite at Shelter and demonstrated the Serena Client Pack. Following the demo, the project team reconsidered the initial assumptions.
- 3. We looked at integration considerations and the licensing costs for Serena Dimensions seat licenses vs. Client Pack licenses. We looked at our target developer audience. Although our initial audience is Java developers, we are also targeting COBOL developers as SQL/PL procedure developers. It made sense to have management of SQL/PL stored procedures entirely under Serena Changeman. The price of a Serena Client Pack seat license is approximately ¼ the cost of a Dimensions seat license, so it makes sense to consider the Serena Client Pack if our target developers are COBOL as well as Java developers. (All Java Developers have Serena Dimensions Seat licenses).
- 4. Note, we plan to completely implement Serena Changeman ZMF SQL/PL support before embarking on the Client Pack.







- 1. This example illustrates the use of BIND PACKAGE ... DEPLOY for a stored procedure called SQPDDR01.
- 2. It is compiled / initially created in the UNIT testing level using DB2 subsystem DSNT.
- 3. DB2 object qualifiers correspond to the testing level or Production.
- UNIT uses a qualifier of SHELTRU
- TEST uses a qualifier of SHELTRT
- PROD uses a qualifier of SHELTRP
- 4. DSNT is Shelter's main development DB2 subsystem, and

DSN is Shelter's main production DB2 subsystem.



- 1. Despite our Changeman administrator's initial preference for the Compile Once Model, there are issues with the use of this model for SQL/PL stored procedures.
- 2. The Security team does not really like production and DB2 subsystems to be able to communicate with each other.
- This is still possible at Shelter because we use some DB2 replication techniques between production and test that require this communication (the DB2 Crossloader for one).
- So we were able to successfully test BIND PACKAGE ... DEPLOY from the test DB2 subsystem to production.
- We found out that Changeman ZMF 7.1.2 (the first release Shelter has used that supports SQL/PL native procedures) uses the Compile Many model. It performs a DROP PROCEDURE followed by a CREATE PROCEDURE when promoting or installing.
- Since our administrators also prefer as little customization as possible, that sealed the deal for the Compile Many Model.



- 1. I had hoped to have the project completely wrapped by the time I needed to have the presentation ready for SHARE in December 2012!
- 2. Oh well!







1. SQLPLGRP is the RACF group we designated for SQL/PL stored procedure development.



















