



IBM's System z Forum

Make the most of your mainframe.

“Ten Things You May Not Have Known”

Neale Armstrong
Geoff Nichols

Session Premise

Australian users of IMS and DB2 have good version currency.



But, are you aware of all the useful capabilities within
The latest releases of IMS and DB2 ?

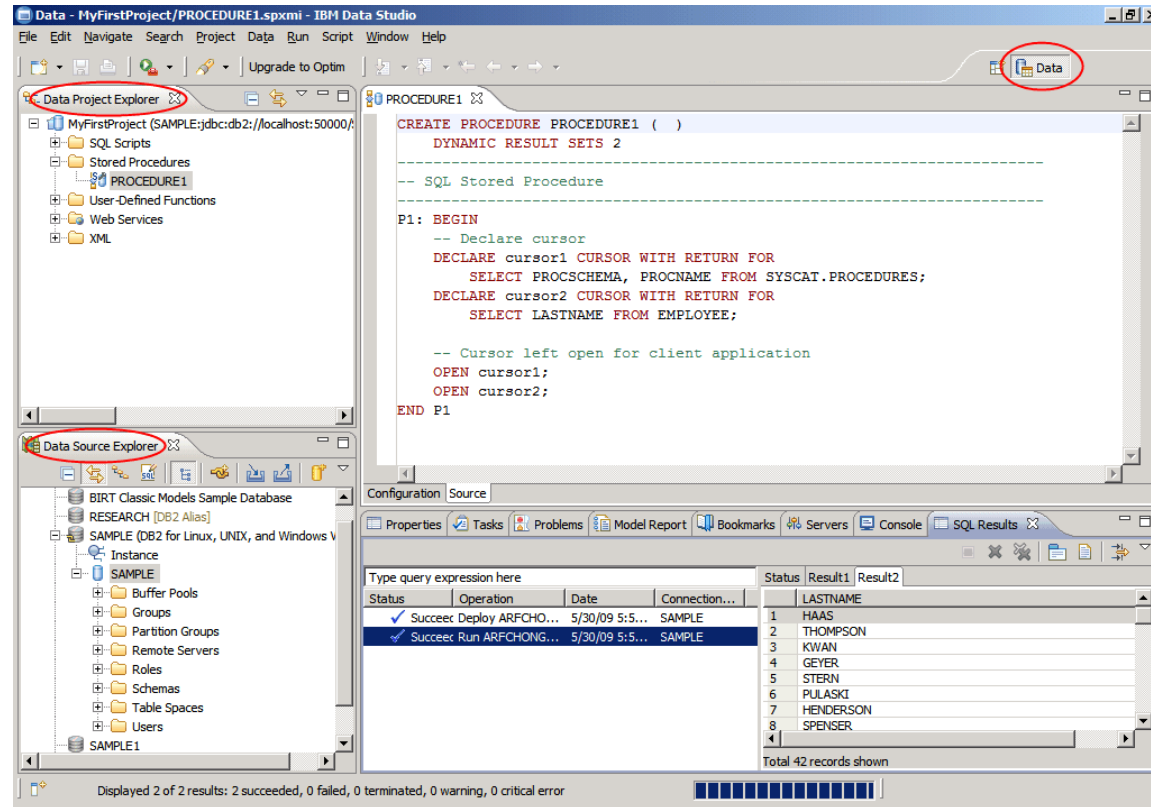


Guided Wine Tasting Session of what's in your own cellar



1. Data Studio

Every Vintage since 2008 has grown more full bodied
2012: V3.1.1



Best Practices for **DB2 z/OS**

<http://www.ibm.com/developerworks/data/library/techarticle/dm-1006datastudiooptintips/index.html>

Download:

<http://www.ibm.com/developerworks/downloads/im/data/>



Reasons to download and use Data Studio

No Charge (& V3.1.1 includes the DRDA license file)

Productivity UI for DBA

The Primary Vehicle for Visual Explain and SQL tuning

No charge : single query tuning

Charge: Optim Query Workload Tuner

The Primary Vehicle for development, debug, & deployment of DB2 Routines (Stored Procs, UDFs etc...)

Support for pureXML, Data Model diagrams etc...

... no threat to 3270 based DB2 tools
... it is a powerful additional toolset



2. Rational Developer & Test V8.5

A very lively number !



System Z Test Environment



Reasons to consider RD&T

Flexibility

- Change Management Practices is a great strength of the System Z Platform
- Change Management Practices is a weakness of the System Z Platform

Productivity

- Ability to Prototype & Test in a sandpit without worries
- Rational Development & Testing

Costs

- Low cost environment for Development & Test purposes



3. IMS Enterprise Suite

A Blend of Classic Californian Grapes and New World Production Methods



Here you can create SQL scripts with Select, Update, Delete, Insert statements

SQL Builder with content assistance to build a SQL statement

Connection
PSB = Catalog
DB PCB = Schema

Database Segments

The screenshot shows the Eclipse IDE interface for database development. The main window displays an SQL script: `SELECT PCB01.HOSPITAL.HOSPNAME, PCB01.PATIENT.PATNAME, PCB01.HOSPITAL.HOSPCODE FROM PCB01.HOSPITAL, PCB01.PATIENT`. Below the script is the SQL Builder, which shows two tables: HOSPITAL and PATIENT. The HOSPITAL table has columns HOSPCODE, HOSPLL, and HOSPNAME. The PATIENT table has columns HOSPITAL_HOSPCODE, WARD_WARDNO, PATNUM, PATLL, and PATNAME. The SQL Builder also shows a table with columns: Column, Alias, Output, Sort Type, and Sort Order. The SQL Results window shows a table with columns: Status, Operation, Date, Connection, HOSPLL, HOSPCODE, HOSPNAME, HOSPITAL_HOSPCODE, and WAR. The results table contains 10 rows of data.

View the SQL results



Reasons to download IMS Enterprise Suite

IMS is opening up

IMS Catalog

SQL Language

IMS Explorer

Eclipse Bedfellow with Data Studio for DB2

Etc...

Platform for Development activities now

... platform for Administration activities in future

A cornerstone of IMS integration strategy



4. pureXML


A wine with a complex structure



XML data type
XML Schema Repository
XML Schema Validation
XML Indexes
SQL/XML Language

XMLQUERY to SELECT specific elements or attributes from a document
XMLEXISTS to apply a WHERE predicate against elements or attributes
XMLSERIALIZE to construct the document in a external string format
XMLPARSE to parse an XML document into the internal node format
XMLTABLE to externalize an XML document in relational row / column format with SQL types

```
select deptno
       ,row_timestamp
       ,xmlpot_sequence
       ,xmlserialize(xmlquery('/DEPT/PROJ[@PROJNO="AD3110"]'
                              passing xml_time_account) as clob) as xml_time_account_proj
       ,xmlserialize(xml_audit as clob) as xml_audit
from time_accounting_part
where xmlexists('/DEPT[@DEPTNO="D21"]' passing xml_time_account)
      AND XMLEXISTS('/DEPT/PROJ/EMP[STARTDATE>="1983-01-01"]'
                    PASSING XML_TIME_ACCOUNT)
```



Reasons to use pureXML

Your organisation almost certainly DOES use XML extensively

Storing XML in files or LOBs is wasteful

- Hides away the essence of the information
- A lot of processing costs to look inside the XML document

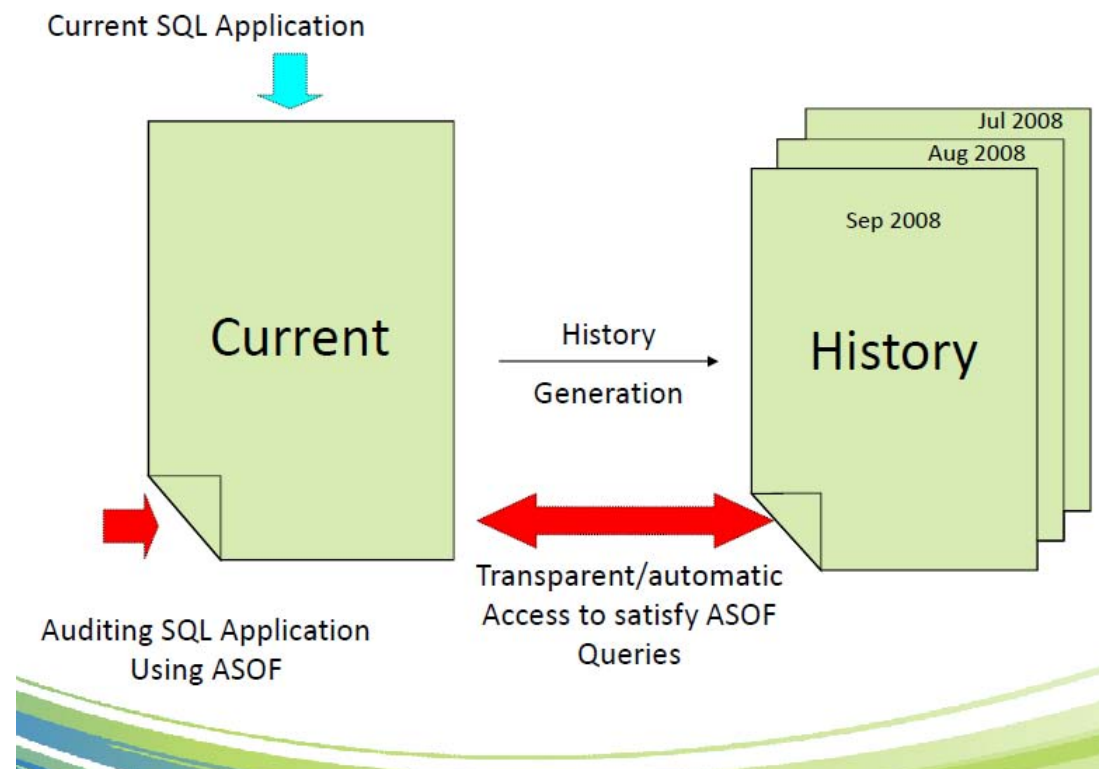
Store it natively in DB2

- Index It
- Process is alongside Relational Data with ease
- Integrate with your enterprise service bus
- Treat it like any other first class datatype
- zIIP offload all XML validation



5. Bi-temporal Tables

Know when to drink your wine !



Reasons to use bi-temporal data

Productivity:

Don't wait for Release 2 or 3 of your new application to implement Archival or History Functions.

Let DB2 look after the logic required to manage historical data From Day 1.



6. Tools Customiser for Z

Wine for Sharing



Goal: A more consistent, usable and simple solution for the simultaneous and ongoing configuration of multiple IBM Tools

- Assists in the post-SMP/e configuration, tailoring and customization
 - **Consumability** – Faster up and running time to Tool usage
 - Product templates are customized by TCz
 - Provides Job execution sequence
 - Due to automatic discovery of previous release customization parameters, there is less manual entry
 - **Easy customization of multiple tools simultaneously**
 - Step by step, with HELP text, ISPF panel-driven dialog allows specification of multiple DB2 Tool products customization
 - **Easier upgrades**
 - Parameters from previous customization are saved for future new product releases and DB2 upgrades

TCz is a no-charge prerequisite for every enabled DB2 & IMS Tool and ships as part of the IBM Tools Base 5655-V93



Reasons to use TCz

1. You have to
2. Productivity



7. DB2 10 Security

... A fortified wine



Separation of Duties: Move over SYSADM...

SECADM – security administration
DBADM ON SYSTEM – database administration

with/without DATAACCESS
with/without ACCESCTRL

SQLADM – performance administration

Audit Traces – audit policy catalog tables

Row and Column Level Permissions



Reasons to use DB2 10 Security features

Comply with today's growing security demands

And devise a security design that saves DBA & Security workload.



8. Package Management

Store your wine carefully:



- Plan Management Policies (PLANMGMT)
 - DB2 9 delivered Access Path / BIND Stability (PLANMGMT) via REBIND
Delivered in maintenance stream via PK52523 via REBIND
PLANMGMT Rebind keyword or System Parameter
 - BASIC: Keeps 2 copies of a package (Current and Previous)
 - EXTENDED: Keeps 3 copies (Current, Previous, and Original)
 - Use the Profile tables to set up automatic Plan Management Policies
 - DB2 10 evolves this capability
 - More formal Catalog implementation for PLANMGMT (V9 delivery)
 - SYSPACKCOPY is populated as part of the ENFM process.
 - Old copies that are from before V6 will still be marked invalid



Reasons to use package versioning

Eliminate the risk of performance regression after rebind.

One or two bad access path experiences

Leads some DB2 users to avoid rebinds at all costs

... so they will never realise the many access path improvements

Use Package Versioning:

Controlled programme of rebinds

With procedures to fall back if necessary



9. Indexing Options

There's an App for that
Pocket Sommelier



Index Include Columns?


```
CREATE UNIQUE INDEX ON t1(c1,c2,c3) INCLUDE (c4,c5)
```

Or No Index at all !

```
SELECT * ..WHERE  
ITEMNO = 'W0133-1662996 '
```

1 data page access

Unique Key Value



- Key finds the row without index
- Reduced:
 - Page visits
 - CPU time
 - Elapsed time

- Page in buffer pool
- Page is read from disk

1 data page disk I/O
(Possibly in buffer pool)

-Trade-off: extra space used

Rid



Consider New Physical Design Options

Each new release of DB2 contains many new physical design options

Take a moment to consider them for new and existing databases.





10. Tivoli Omegamon XE for DB2 Performance Expert V5.1


Monitor your wine consumption...



Where's My Problem ?



Reasons to use OMPE



Top SQL statements executed by Java or CLI applications like SAP, Cognos, DataStage or WebSphere
 ■ Zoom in on a selected SQL

The screenshot shows the Optim Performance Manager (OMPE) interface. The main dashboard is titled 'Extended Insight Analysis Dashboard: OMP1D911'. It includes a 'Graph' section showing 'Average End-to-End Response Time' and a 'SQL Statements' table. A red box highlights the 'SQL Statements' table, and another red box highlights the 'Detail Area for Average End-to-End Response Time'.

Statement Text	Statement Executions	Average Data Server Time (sec)
SELECT 'PVT_40K' AS WKLID, '...	1	0.504
SELECT 'PVT_40K' AS WKLID, '...	1	0.474
SELECT 'PVT_40K' AS WKLID, '...	1	0.518
SELECT 'PVT_40K' AS WKLID, '...	1	1.393
N/P	1	1.023

Detail Area for Average End-to-End Response Time

Overall average response time per transaction:	0.075 sec
Maximum response time:	15.282 sec
Maximum Time of running transactions:	10.688 sec
Number of transactions:	61,245
Statements:	65,344

Time Distribution (%)

Client time	32.00%
Network time	6.67%
Data server time	61.33%

Detailed End-to-End Response Time

Ten Things you may not have known that you already have

1. Data Studio
2. Rational Developer & Test (extra cost)
3. IMS Enterprise Suite
4. pureXML
5. BiTemporal Tables
6. Tools Customiser for Z
7. DB2 Security
8. Package Versioning
9. Indexing Options
10. Tivoli Omegamon XE for DB2 PE (extra cost)





© Copyright IBM Australia Limited 2012. ABN 79 000 024 733. © Copyright IBM Corporation 2012. All Rights Reserved.

TRADEMARKS: IBM, the IBM logos, ibm.com, smarter planet and the planet icon are trademarks of IBM Corporation registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

IMPORTANT PRIVACY INFORMATION: If you or your organisation would prefer not to receive further information on IBM products, please advise us on 132 426 (Australia) or 0800 444 714 (New Zealand). If you would like IBM Australia Limited to refrain from sending you commercial electronic messages you may send an unsubscribe message to contact@au1.ibm.com. The sending of this message was authorised by IBM Australia Limited, and IBM Australia Limited can be contacted at rlm@au1.ibm.com or on 132 426 (Australia) or 0800 801 800 (New Zealand). IBM may store data on international servers used by it.
GL_13679