

**IBM Information**

>>> On Demand

**2007**



# Managing Dynamic SQL Performance

*Thanikachalam “Billy” Sundarrajan*

*Lead DBA, Fifth Third Bank, Cincinnati*

*[billy.sundarrajan@53.com](mailto:billy.sundarrajan@53.com)*

*1819A Data Servers - System z - DB2 and Tools*



***Act.Right.Now.***

**IBM INFORMATION ON DEMAND 2007**

**October 14 - 19, 2007**

**Mandalay Bay**

**Las Vegas, Nevada**

# Agenda

- Fifth Third Bank Environment
- Performance tuning challenges
- DB2 V8 Enhancements
- DB2 9 Enhancements
- SQL Performance Warehouse
- Summary
- Questions?



# Fifth Third Bank Environment

- Fourth largest merchant processor in US
- Data Sharing (2-way in Dev/Test/Load, 2-way in Prod-DW, 6-way in PROD-OLTP)
- DB2 V8 NFM, 22TB DASD, Several multi terabyte tables
- Data Marts, Data Warehouse (z Star schema), Reporting, OLTP
- Stored Procedures (SQL and Cobol)
- Majority of SQL workload is dynamic SQL (~22 Million dynamic SQLs)
- Majority of large tables populated by LOAD SHRLEVEL CHANGE or MQ (24 x 7)





# Dynamic SQL Performance Tuning Challenges

*Act.Right.Now.*



# Dynamic SQL performance tuning challenges

- Identifying the “actual” end user performing a dynamic SQL – Most dynamic SQL get executed by an application server/common authorization ID
- Ability to drill down accounting information to a more granular level than just the authorization ID
  - Middleware such as Business Objects, Websphere perform the connect to DB2 using common authorization/RACF ID
- Identifying the contents of the dynamic SQL cache and the access paths being used
- Reducing the number of accounting (SMF 101) records for DISTSERV



# Dynamic SQL performance tuning challenges (addressed by DB2 9 Enhancements)

- Using Resource Limit Facility to control the service units for dynamic SQL – at a more granular level than AUTHID/PLAN/PACKAGE
- Using DB2 traces for dynamic SQL applications – How to control the unit of work being traced
- Optimizing prepared dynamic SQL where the host variables change
- Securing the application server user ID



# Identifying the “actual” user in a 3-tier architecture

*Act.Right.Now.*



## Identifying the “actual” end user:

- Most dynamic SQL intensive applications use a functional ID/ common Authorization ID
  - Single ID avoids the creation and maintenance of a large number of RACF IDs in a web-based application
  - DBA does not have additional information about the “actual” end-user
  - In most instances, the IP address in the DB2 monitor is the IP address of the DB2 Connect gateway





## Identifying the “actual” end user:

- DB2 v7 provided ability to pass “client accounting” / “client identification” information to SMF 101 records.
- V7 –Enhancements to supply client accounting information to the SMF 101 records and to the monitor
  - SQLESETI
  - RRSAF
  - JDBC Universal driver
- What does V8 provide?



## Client Identification registers –V8:

- V8 - provides new registers
  - CURRENT CLIENT\_ACCTNG
  - CURRENT CLIENT\_APPLNAME
  - CURRENT CLIENT\_USERID
  - CURRENT CLIENT\_WRKSTNNAME
- Allows applications to intercept and use the information supplied/set by a web interface



## Client Accounting Information register –V8:

- **CURRENT CLIENT\_ACCTNG**  
**QMDASUFFIX – VARCHAR(200)**  
contains the value of the accounting string from the client information that is specified for the connection. Set using:
  - `sqleseti (SQLE_CLIENT_INFO_ACCTSTR)`
  - `RRS DSNRLI`
  - `DB2Connection.setDB2ClientAccountingInformation(String info)`
- In a distributed environment, if the value set by the API exceeds 200 bytes, it is truncated to 200 bytes.



## Client Userid register – V8:

- CURRENT CLIENT\_USERID  
QWHCEUID – CHAR(16)

contains the value of the client user ID from the client information that is specified for the connection. Set using:

- sqleseti (SQLE\_CLIENT\_INFO\_USERID)
- RRS DSNRLI
- DB2Connection.setDB2ClientUser(String info)
- In a distributed environment, if the value set by the API exceeds 16 bytes, it is truncated to 16 bytes.



## Client Workstation register – V8:

- CURRENT\_CLIENT\_WRKSTNNAME  
QWHCEUWN – CHAR(18)

contains the value of the workstation name from the client information that is specified for the connection. Set using:

- sqleseti (SQLE\_CLIENT\_INFO\_WRKSTNNAME)
- RRS DSNRLI
- DB2Connection.setDB2ClientWorkstation(String info)
- In a distributed environment, if the value set by the API exceeds 18 bytes, it is truncated to 18 bytes.



## Client Application Identifier register – V8:

- CURRENT CLIENT\_APPLNAME

QWHCEUTX - CHAR(32)

contains the value of the application name from the client information that is specified for the connection. Set using:

- sqleseti (SQLE\_CLIENT\_INFO\_APPLNAME)
- RRS DSNRLI
- DB2Connection.

setDB2ClientApplicationInformation(String info)

- In a distributed environment, if the value set by the API exceeds 32 bytes, it is truncated to 32 bytes.



## EXCSQLSET – DRDA Applications – V8:

- DRDA Applications can use EXCSQLSET
  - SET CLIENT USERID 'my\_userid'
  - SET CLIENT WRKSTNNAME 'my\_wrkstn'
  - SET CLIENT APPLNAME 'my\_applname'
  - SET CLIENT USERID 'my\_userid'
- Additional information allows the ability to store client identifier in tables (such as who last updated by data)
- EXCSAT EXTNAM/RRS is used for passing the correlation ID – stored in QWHCCV
- DSNDQWAS SUBTYPE=ALL provides layout of SMF101 record



## Drilling down on accounting information

- V7 – No easy mechanism to drill down information at a more detailed level than the authorization ID
- V8 –
  - With additional CLIENT registers that are available to the application, additional client information can be easily externalized into a DB2 table
  - Reduces the need to wait until SMF records are available







# What is in my Dynamic SQL cache?

*Act.Right.Now.*



# What's in my Dynamic SQL cache?

- Dynamic SQL cache (DSC) contains statements that are prepared and inserted into the global cache (except SQLs using GTTs)
  - CACHEDYN=YES
- With IFCID 318 (Monitor trace) turned on, cache contains information about SQL execution, CPU time, GETPAGES etc
- Until V8, information on cached statements and access paths could not be easily externalized
  - Required a tool such as DB2 PE, Omegamon etc
  - Reading IFCID 22 (mini plan)/63 (SQL text)
- What does V8 provide?



# What's in my Dynamic SQL cache –V8

- Unlike static SQL, DBAs do not have easy access to the dynamic SQL statements run by an application
- V8 provides additional enhancements to EXPLAIN to view information in the DSC
- New EXPLAIN options:
  - EXPLAIN STMTCACHE ALL
  - EXPLAIN STMTCACHE STMTID
  - EXPLAIN STMTCACHE STMTTOKEN
    - Differs from traditional EXPLAIN – Does not re-explain to obtain access path
    - Can be run using SPUFI, DSNTEP2, Visual Explain



## New EXPLAIN options – V8

- **EXPLAIN STMTCACHE ALL**
  - Dumps the SQL statement cache –
  - Populates a table <current sqlid>. DSN\_STATEMENT\_CACHE \_TABLE
  - DB2 V8 NFM required
  - Uses LOBs
  - Member specific – In a Data sharing environment should be run on all members



## New EXPLAIN options – V8

- **EXPLAIN STMTCACHE ALL**
  - Information identical to information from IFCID 316 and IFCID 317
  - Collection and reset of statistics is controlled by IFCID 318
  - Visual Explain/OSC provides easy to use interface to perform EXPLAIN command



# EXPLAIN STMTCACHE ALL using DB2 OSC/OE

The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The main window is titled "SQL Cache" and has a tabbed interface with "View Queries" selected. The interface is divided into two main sections: "Source" and "Query text".

**Source Section:**

- Source:** Specify the source of the query that you want to tune and then, if applicable, select a view, customize, and save it. After the query is identified, tune the query in the Query text section.
- Query source:** A dropdown menu is set to "Statement cache". To its right are two buttons: "Enable Cache Trace" and "Disable Cache Trace".
- View name:** A dropdown menu is set to "<Select a view or create a new one>". A tooltip box is visible over this dropdown, containing the text: "Select a source from which you want to display the queries. After you select the source, you must select a view before the queries are displayed. For more detail on each option, press F1 to open the help for this panel." Below the dropdown is a list of view names: "ACCU\_M\_CPU\_DESC", "ACCU\_ELAP\_DESC", "EXECUTIONS\_DESC", "GETPAGES\_DESC", and "<New view...>".

**Query text Section:**

- Query text:** There are several options to tune the selected query. Format or categorize selected query text, analyze the query, or use additional tools for more analysis.
- Tools:** A dropdown menu is set to "EXPLAIN timestamp".
- EXPLAIN options:** Three radio buttons are present: "Run EXPLAIN again" (selected), "Use subsystem EXPLAIN information", and "Use local EXPLAIN information".
- A large empty text area is provided for entering or editing the query text.

The bottom of the window shows a taskbar with the Start button, several application icons, and the text "DB2 Performance Expert..." and "IBM DB2 Optimizatio...". The system tray on the right shows the time "4:58 PM".



# EXPLAIN STMTCACHE ALL using DB2 OSC/OE

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Configure Subsystems View Queries View Workloads View Monitors \* SQL Cache x

**Source:**

Specify the source of the query that you want to tune and then, if applicable, select a view, customize, and save it. After the query is identified, tune the query in the Query text section.

Query source:

View name:

All of the rows are displayed. The number of rows is 4929.

STMT_ID	STAT_EXEC	STAT_CPU	STAT_ELAP	STAT_GPAG	STMT_TEXT
7331	1	1.6035937370878516E-4	1.7643749380630313E-4	12	SELECT A.* FROM "SYSIBM"."SYSC
8010	1	1.603125019898346E-4	1.6745263202727644E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
7995	1	1.6020311524306296E-4	1.612499945765069E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8683	1	1.601562435241124E-4	1.6129686629545747E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8935	1	1.6014842914563476E-4	1.843124955627775E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8300	1	1.5977344084211497E-4	0.022092171948330075	2	SELECT COUNT(*) FROM "CID1"."C
8446	1	1.5971874019277154E-4	1.6107812190574466E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8016	1	1.594921814245811E-4	1.6057811809849797E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8332	1	1.5939062360820233E-4	2.1226561877091287E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8650	1	1.5933593751077412E-4	1.6065624733135908E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8216	1	1.5929686561838596E-4	1.60421874184691E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
9089	1	1.5923436514248012E-4	1.6070311905030965E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA
8041	1	1.5923436514248012E-4	1.6070311905030965E-4	3	SELECT A.* FROM "SYSIBM"."SYSTA

**Query text**

There are several options to tune the selected query. Format or categorize selected query text, analyze the query, or use additional tools for more analysis.

Query   EXPLAIN timestamp:

EXPLAIN options:  Run EXPLAIN again  Use subsystem EXPLAIN information  Use local EXPLAIN information

Project Query x

Enable Cache Trace



## New EXPLAIN options – V8

- EXPLAIN STMTCACHE STMTID
  - Specifies that the cached statement associated with the statement ID (STMT\_ID column in DSN\_STATEMENT\_CACHE\_TABLE) is to be explained.
  - Statement ID is an integer that uniquely determines a statement that has been cached in dynamic statement cache. The statement ID can be retrieved through IFI monitor facilities from IFCID 316 or 124 and is shown in some diagnostic IFCID trace records such as 172, 196, and 337.





## New EXPLAIN options – V8

- EXPLAIN STMTCACHE STMTID
  - Column QUERYNO is given the value of the statement ID in every row inserted into the plan table, statement table, or function table by the EXPLAIN statement. (i.e., STMT\_ID column copied to QUERYNO column)
  - BIND\_TIME is given the value of CACHED\_TS
  - Writes information to the DSN\_STATEMENT\_CACHE\_TABLE and PLAN\_TABLE



## New EXPLAIN options – V8

- **EXPLAIN STMTCACHE STMTTOKEN**
  - Explains all cached statements associated with a statement token (statement token up to 240 bytes)
  - STMTTOKEN is associated with the cached statement by the application program that originally prepares and inserts the statement into the cache.



## New EXPLAIN options – V8

- EXPLAIN STMTCACHE STMTTOKEN
  - Application programs use RRSAF SET\_ID function, or sqleseti API from a remotely-connected program.
  - Column STMTTOKEN is given the value of the statement token, and the column QUERYNO is given the value of the statement ID for the cached statement with the statement token, BIND\_TIME is given the value of CACHED\_TS.
  - Writes information to the DSN\_STATEMENT\_CACHE\_TABLE and PLAN\_TABLE



## EXPLAIN versus EXPLAIN STMTCACHE

- Unlike EXPLAIN, EXPLAIN STMTCACHE does not go through the EXPLAIN process
- DB2 writes the current access path to the PLAN\_TABLE; re-explaining the statement can provide different results
- SQLCODE -20248 if statement no longer in cache



## PLAN\_TABLE – How do I identify a cached statement?

- COLLID Column in PLAN\_TABLE, and DSN\_STATEMENT\_CACHE\_TABLE will contain the value DSNDYNAMICSQLCACHE for cached statements
- New column STMTTOKEN in PLAN\_TABLE contains the statement token value for the Statement being explained



# PLAN\_TABLE – How do I identify a cached statement?

DSN\_STATEMENT\_CACHE\_TABLE

(Key columns)

STMT_ID
STMT_TOKEN
COLLID
CACHED_TS
EXPLAIN_TS
STMT_TEXT

STMT\_ID – Token identifying a cached statement

STMT\_TOKEN – Identification token set by user (null if not populated)

COLLID – DSN DYNAMIC SQL CACHE to indicate cached statement

CACHED\_TS – Timestamp when statement was inserted to SQL cache

EXPLAIN\_TS – Timestamp when EXPLAIN STMT CACHE ALL/STMTID was run

(Multiple EXPLAIN\_TS possible for the same Statement)

STMT\_TEXT – SQL Text

PLAN\_TABLE

(Key Columns)

QUERYNO
STMT_TOKEN
COLLID
BIND_TIME

QUERYNO – Contains STMT\_ID from the dynamic SQL Cache

STMT\_TOKEN – Contains STMT\_TOKEN from the dynamic SQL cache

COLLID – DSN DYNAMIC SQL CACHE to indicate statement being explained is from the dynamic SQL cache

BIND\_TIME – CACHED\_TS from dynamic SQL cache (Duplicate rows possible when EXPLAIN STMT CACHE STMTID executed multiple times)





# How do I reduce the volume of DISTSERV SMF101s

*Act.Right.Now.*



# Too many SMF101 records for DISTSERV?

- Most installation of DB2 use type 2 inactive connection support (i.e., CMTSTAT=INACTIVE)
- In DB2 v7, with CMTSTAT=INACTIVE, SMF101 record is cut every time a thread becomes Inactive
- The number of SMF records could exceed several million (We cut 22 million+ DB2 SMF101 records in a day for DISTSERV)





## Too many SMF101 records for DISTSERV?

- Large volumes inhibit the ability to extract useful information for DDF/RRSAF threads
- Large volumes result in increased batch run times (for processing SMF records)
- What does V8 provide?



## SMF 101 records rollup- V8

- New dynamic ZPARM ACCUMACC, ACCUMUID rolls-up accounting information from DDF/RRSAF threads
- Can be turned on/off dynamically
- Allows the DBA to obtain detailed information if needed without disruption
- Accounting data accumulated by either, or combination of
  - User ID
  - Transaction name
  - Workstation name



## SMF 101 records rollup- V8

- How do I set the ZPARMS and register used for roll-up?
  - ACCUMACC
  - ACCUMID
  - User ID
  - Transaction name
  - Workstation name



## SMF 101 records rollup- V8

- ZPARM– “DDF/RRSAF Accum” ACCUMAC (DSNTIPN)
- Controls whether DB2 accounting data should be accumulated for DDF and RRSAF threads. The type of accumulation is controlled by ZPARM ACCUMUID.



## SMF 101 records rollup- V8

- ZPARM– “DDF/RRSAF Accum” ACCUMAC (DSNTIPN)
  - If “10” is specified (the default), then DB2 continues to write an accounting record for '10' occurrences of the “Aggregation field”
  - If “NO” is specified, an accounting record is cut when a DDF thread is made inactive, or when signon occurs for an RRSF thread.
  - If 2-65535 is specified, then DB2 writes an accounting record every 'n' occurrences of the “Aggregation field” on the thread, where 'n' is the number specified for this parameter.
  - “Aggregation field” defined by ACCUMUID



## SMF 101 records rollup- V8

- ACCUMUID – ZPARM – Defines the “Aggregation Fields” to be used for DDF and RRS accounting record rollup.
- Aggregation is based on the following three fields that are provided to the application:
  - ID of the end user (QWHCEUID, 16 bytes)
  - End user transaction name (QWHCEUTX, 32 bytes) contains the name of the executable
  - End user workstation name (QWHCEUWN, 18 bytes)



## SMF 101 records rollup- V8

- Client accounting fields can be set using
  - “Server Connect” and “Set Client” (SQLESETI) calls
  - RRSAF threads via the RRSAF SIGN, AUTH SIGNON, and CONTEXT SIGNON functions
  - Java programs when using the new Java Universal Driver



## SMF 101 records rollup- V8

- **ACCUMUID values:**
  - 0 : End user ID, AND end user transaction name, AND end user workstation name
  - 1 : End user ID
  - 2 : End user transaction name
  - 3 : End user workstation name
  - 4 : End user ID AND end user transaction name
  - 5 : End user ID AND workstation name
  - 6 : End user transaction name AND workstation name
- The default value is 0 (zero). The ACCUMUID value is ignored if ACCUMACC=NO (no DDF/RRS rollup).





## SMF 101 records rollup- V8

- DB2 can override ACCUMACC in certain situations:
  - Storage threshold is reached for the accounting rollup blocks
  - No updates have been made to the rollup block for 10 minutes
  - Not all of the fields specified in the aggregation criteria are supplied
- May provide relief to the “SMF flooding” issue
- May want to check CPU changes



# SMF 101 records rollup

- Enhances the ability to use accounting information
- Example:
  - Significantly reduces the amount of time taken to run performance reports (we saw a 90% reduction in time for DISTSERV plans)
  - Rolling up information by “end user” allows the DBA to drill down and obtain CPU utilization information at a more granular level than AuthID
- IFCID 239 provides significant amount of package level data (buffer manager, lock manager, and SQL statistics accumulated at the package level )
  - Overhead associated with collection
  - Accounting Class(10) controls collection of additional data (PK28561)





# DB2 v8 and DB2 9 Enhancements

*Act.Right.Now.*



## Other DB2 v8 Enhancements

- Dynamic SQL statements – like static SQL was limited to 32KB – Version 8 – limit changed to 2 MB
- “Execute Immediate” can use a LOB/CLOB
- Entire statement information is provided by IFCID 350 –similar to IFCID 63
- Invalidate dynamic SQL cache - RUNSTATS UPDATE NONE REPORT NO, CREATE INDEX
- New ZPARM EDMSTMTC – Controls size of dynamic SQL cache



## DB2 9 Enhancements

- Enhancement to DSNRRLST based on client variables
- REOPT AUTO enhancement
- -START TRACE enhancement
- Trusted Context – Prevent the application server ID from getting used outside of the app server



## Using RLF enhancements – DB2 9

- DB2 9 allows the RLF to be set at a more granular level than AUTHID/PLAN/PACKAGE
  - For dynamic SQL PLAN is always DISTSERV which does not provide adequate granularity
  - Utilizes a table DSNRLMTxx – Middleware resource limit facility
    - Allows the option to specify
      - RLFEUAN – End User application name
      - RLFEUWN – End user workstation name
      - RLFEUID – End User ID
      - RLFIP – IP from which dynamic SQL originated
    - START RLIMIT starts RLST as well as the RLMT tables
    - Only one RLMT table can be active at any point in time



## REOPT(AUTO) – DB2 9

- Enhancement to REOPT – REOPT(AUTO)
  - DB2 determines if a new access path based on the parameter values
  - DB2 saves the access path in the dynamic SQL Cache
  - Combines the advantages of REOPT(ALWAYS) and REOPT(ONCE)



## -START TRACE enhancements (DB2 9)

- -START TRACE has new keywords
  - USERID - End Client UserID
  - APPNAME - Application Name
  - WRKSTN - Client Workstation Name
  - CONNID - Connection
  - CORRID - Correlation ID
- Constraint block modified to include specific end client User ID, Application name, workstation name, connection, or correlation ID
- Filtering block allows XUSERID, XAPPNAME, XWRKSTN, XCONNID, XCORRID – To exclude threads matching specific end clients
- X option allowed with wildcards but a x(\*) is not allowed





# TRUSTED CONTEXT

- Addresses security/audit challenges using a common middleware/functional ID to connect to DB2
  - Database entity which allows for a unique set of interaction based on Authorization ID and connection attributes
  - Two types of trusted context
    - Implicit – No application changes are needed
    - Explicit – Application changes are needed
  - Allows Authorization IDs to connect from a specific IP address without authentication and assign default roles
  - Explicit trusted context allows for user ID switching
  - Switching User IDs allows the actual user to be used for logging all database changes



# TRUSTED CONNECTION

- Addresses challenges using a common middleware/functional ID to connect to DB2
  - ROLE is a Database entity to which authorities can be granted
  - TRUSTED CONNECTION is a database connection established in a TRUSTED CONTEXT
  - Misuse of application IDs can be alleviated by granting security on DB2 objects to roles
  - Common application IDs can be granted the ROLE under a TRUSTED CONTEXT





# Performance Warehouse

***Act.Right.Now.***



## Performance Warehouse

- Performance warehouse is a collection of DB2 performance data
  - Performance information (Summary and exception) at a Plan, Package, SQL level stored in custom tables serves as a tuning repository
  - Plan/Package Information can be extracted from SMF 101 (IFCID 3 and IFCID 239 records)
  - Information from DB2 augmented with user specific information
  - Tables for Top “n” dynamic and static SQL
  - SQL level data extracted from SQL level monitoring tools such as IBM DB2 Query Monitor, Dynamic SQL cache, DB2/OSC monitor profiles (DB2 9 onwards)



## Performance Warehouse

- Using Enhancements in DB2 v8/9 for collecting SQL level performance data
  - DB2 v8 – Top ‘n’ SQL (CPU, elapsed time, GETPAGES) - EXPLAIN STMTCACHE ALL output can be parsed to populate custom built SQL cache tables
  - DB2 9 – Using profile monitor - Statements which leave the cache or statements (static/dynamic) which exceed thresholds are written to DSN\_STATEMENT\_RUNTIME\_INFO table.



## Summary

- V8/9 – Provides a plethora of tools to “tame” dynamic SQLs
  - Special registers
  - SQLESETI/JDBC/RRS Signon to set client variables
  - ACCUMAC/ACCUMID to reduce SMF records
  - IFCID 350
  - RUNSTATS REPORT NO/UPDATE NONE to invalidate dynamic SQL cache
  - ZPARM EDMSTMTC
  - -START TRACE, DSNRLMT, REOPT(AUTO), Profile Monitor



# Questions?

