

## How to Build Your Own DB2 Performance Database

Joe Borges  
Sr. Technical Specialist  
BMO Financial Group

**IDUG-NA Regional Event – California  
October 28<sup>th</sup>, 2004**



## Agenda

- Bank of Montreal (**BMO Financial Group**)
- Introduction to DB2's Instrumentation Facility
- Building a DB2 Performance Database
- Reporting and Exception Monitoring
- Monitoring Key Performance Metrics
- Summary



International DB2 Users Group

“**How to Build Your Own DB2 Performance Database**” is a practical implementation of a DB2 performance database using the DB2 Instrumentation Facility and DB2PM.

Monitor:

- User activity (accounting)
- System thresholds (statistics)
- Lock suspension/timeouts/deadlocks (locking)
- Audit table access (audit)
- Dynamic SQL, Sort information, Rid Pool details and Buffer Pool Statistics at the dataset level (performance traces)
- Set and monitor DB2 thresholds (exceptions)
- ... AND much more!

- **BMO Financial Group**
  - ❖ Bank of Montreal (*Canada's Oldest Bank!*)
  - ❖ Harris Bank (Chicago)
  - ❖ Nesbitt Burns (Investment Firm)
- "A Diversified Financial Services Institution"
- CDN\$247 billion in assets
- 1,100+ Branches & 2,200+ ABMs  
Coast to Coast ... 5 time zones
- 24 \* 7 \* 365 service
- Virtual banking ... Internet, PC, Telephone
- Over 33,000 employees



International DB2 Users Group

#### About Us

**BMO FINANCIAL GROUP** <http://www.bmo.com>

Beginning in 1817, BMO Financial Group (NYSE, TSX: BMO) is one of the largest financial services providers in North America. With average assets of \$247 billion as at July 31, 2002 and more than 33,000 employees, BMO provides a broad range of retail banking, wealth management and investment banking products and solutions. BMO serves clients across Canada and in the United States through its Canadian retail arm BMO Bank of Montreal, Chicago-based Harris Bank, a major U.S. mid-west financial services organization which also has wealth management offices and branches across the United States and BMO Nesbitt Burns, one of North America's leading full-service investment firms.

#### The History of Bank of Montreal

Bank of Montreal, Canada's first chartered bank, opened for business on November 3<sup>rd</sup>, 1817. It provided Canada's first sound and plentiful currency and has played a major and continuing role in the development of the country, taking part in the financing of the first transcontinental railway in the 1880s. The Bank took part in the creation of Canadian Confederation in 1867 and served as Canada's central bank until 1935. The first Canadian bank to open branches abroad, it has long been active in European, Latin American, and East Asian markets as well as in the United States. Today it continues as one of Canada's pre-eminent financial institutions and has a significant presence in United States and world markets.

**About Harris Bank:** visit <http://www.harrisbank.com/facts.html>

**About BMO Nesbitt Burns:** visit <http://www.bmonesbitburns.com>

### System Environment

- **z/OS Environment**
  - 6 mainframes
  - IBM zSeries
  - 8,000+ MIPS, 100+ GB memory
  - 4 Coupling Facilities
  - EMC, IBM, Hitachi, StorageTek DASD
  - StorageTek Automated Cartridge Silo
- **Open Systems**
  - IBM AIX/SP2
  - Production Applications
  - Information Warehouse and Datamarts
  - DB2 Connect Servers

### Software Configuration

- **DBMSs**
  - DB2 UDB for z/OS V7
  - DB2 UDB EEE (AIX/SP2) V7.2 and V8.1
  - DB2 UDB for Windows/NT
  - IDMS
  - Oracle, Sybase
  - MAAS (homegrown!)
  - ... and LOTS of VSAM
- **DB2 for z/OS Production Environment**
  - 5 DB2 Data Sharing Groups: 4-way or 5-way
  - 5 DB2 Subsystems
  - 8 GigaBytes of Bufferpools
  - 2 TeraBytes of DB2 Data
  - 70+ applications
  - 12 (24\*7\*365) mission critical applications  
e.g. MasterCard, Online Banking



International DB2 Users Group

The six mainframes are carved out into multiple logical sites. All multi-LPAR sites have been defined as SYSPLEXes.

In the Open Systems environment, IBM's SP2/AIX operating system is the predominant and the Bank's strategic platform. The Bank is currently running production applications, the Information Warehouse, some Datamarts, and the DB2 Connect servers on the SP2 environment.

Today, the Bank's network is a combination of SNA and TCP/IP.

The "Development" environment is also the Bank's contingency "Disaster Recovery" site. Most of the offsite data for DR is transmitted electronically to the remote site using NDM (Network Data Mover) software and T3 network links. The remaining DR data is remotely mirrored using Hitachi's TrueCopy technology.

The Bank's standard DBMS platform is DB2 UDB for both z/OS and the open systems SP2/AIX environments.

Today approximately 50% of the production "databases" on z/OS are DB2... the majority of the remaining "databases" are VSAM. The Bank has a very large database using a homegrown DBMS (called MAAS) that has been around for 25years! And remnants of IMS and IDMS databases.

In the Open Systems environment the majority of the production databases are DB2 UDB EEE. Other DBMS in this environment include, DB2 UDB for Windows/NT, Oracle and Sybase.

The Bank's standard Transaction Manager is CICS, and it's standard messaging software is MQSeries.

There are 25+ DB2 subsystems and DB2 Data Sharing members at the Bank, supporting over 70 DB2 production applications. Many of these applications are mission critical... 12 have 24\*7\*365 availability today (or close to it!) and most have a requirement for continuous availability.

There are 5 production DB2 Data Sharing Groups at the Bank. The strategy is to have a DB2 data sharing group member on each LPAR (Logical Partition). Most of the Bank's DB2 applications are on, or are migrating to DB2 Data Sharing groups. The applications not scheduled for DB2 data sharing are on 5 DB2 Subsystems (primarily because they reside on single image LPARS). The only exception to this rule is our PeopleSoft DB2 subsystem (data sharing for this subsystem is being considered).

The DB2 production environment consists of approximately 2 TeraBytes of DB2 data and utilizes approximately 8 GigaBytes of memory for bufferpools (all virtual pools).

The Development environment consists of 3 DB2 Data Sharing Groups. All groups consist of two members most of the time, but have the capability of running with four members... to simulate the production environment during QA or volume testing. There are also 5 DB2 Subsystems in the development environment. All DB2 development subsystems and members utilize approximately 3 GigaBytes of memory for bufferpools (all virtual pools).

## Introduction to... DB2's Instrumentation Facility

- Overview
- What is an IFCID?
- DB2 Trace Types and Classes
- Reporting methods
- How to format an IFCID



International DB2 Users Group

This section will cover DB2's Instrumentation Facility (IF).

After an overview and definition of DB2's Instrumentation Facility, this section will... define what an IFCID is; it will describe the various traces (and it's types and classes) ... and finally, it will discuss various methods for reporting the trace data, including how to format an IFCID!

## Instrumentation Facility Overview

- Instrumentation Facility (IF) is a component of DB2's System Services Address Space [DSNMSTR]
- Instrumentation Facilities component services include:
  - DUMP services
  - **TRACE services**
    - Collects Accounting and Statistical information
    - Provides Performance and Serviceability TRACE data
    - Controls TRACE execution
    - Manages the Instrumentation Facility Interface (IFI)
    - TRACE can be activated automatically at DB2 startup or manually by the DB2 -START TRACE command



International DB2 Users Group

DB2's Instrumentation Facility (IF) is a component of DB2's System Services Address Space [DSNMSTR]. The Instrumentation Facilities (IF) subcomponent provides TRACE and DUMP services. DUMP services will not be covered in this presentation.

TRACE services provides a facility to record DB2 data and events ...

- Collects Accounting and Statistical information
- Provides Performance and Serviceability trace data
- Controls TRACE execution
- Manages the Instrumentation Facility Interface (IFI). The IFI is designed for programs needing online trace information (i.e. Online Performance Monitor)
- TRACE can be activated automatically at DB2 startup or manually by the DB2 -START TRACE command

## Instrumentation Facility What is an IFCID?

IFCID = Instrumentation Facility Component Identifier

- Each TRACE event has a specific format and Instrumentation Facility Component Identifier - **IFCID**
- An IFCID uniquely identifies a DB2 TRACE event
- IFCIDs can be activated by either the -START TRACE command or automatically at DB2 startup
- There are 220+ IFCIDs
- DB2's DSNNSAMP(DSNWMSG) library contains descriptions of all trace IFCIDs



International DB2 Users Group

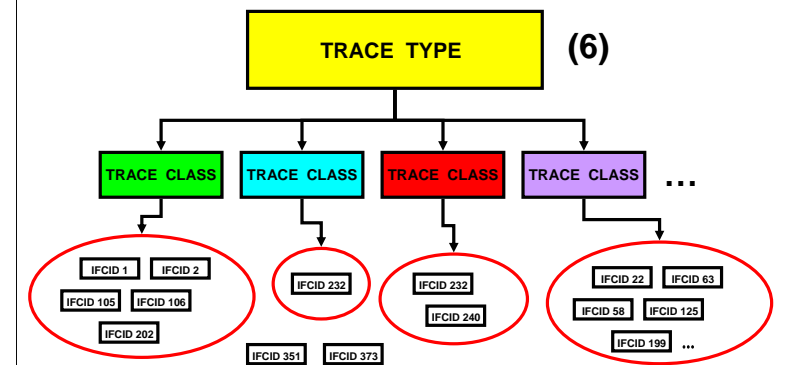
**What is an IFCID?** An IFCID is an "Instrumentation Facility Component Identifier" and it uniquely identifies a DB2 TRACE event. Each TRACE event has a specific format and Instrumentation Facility Component Identifier (IFCID).

IFCID can be activated by either the -START TRACE command or automatically at DB2 startup.

Today (DB2 V7) there are over 220 IFCIDs ... and growing with every release!

IFCID trace descriptions can be found in *prefix.DSNNSAMP(DSNWMSG)*; IFCIDs are also described by the comments in their mapping macros, contained in *prefix.SDSNMACS*. Both datasets are shipped with DB2. Check with your System Programmer if you can't find these datasets, or you don't have access to them.

## Instrumentation Facility DB2 Trace Types and Classes



International DB2 Users Group

DB2 Instrumentation Facility TRACE events and data are grouped into trace **TYPES** and trace **CLASSES**. There are 6 trace **TYPES**:

- > **Accounting** - application program information
- > **Statistics** - DB2 system services information
- > **Audit** - DB2 security controls
- > **Performance** - variety of DB2 events
- > **Monitor** - data for online monitoring
- > **Global** - IBM serviceability data

Trace **CLASSES**:

- > A CLASS defines a certain group of events or data within a trace TYPE
- > A trace CLASS contains at least 1 IFCID, but usually contains multiple IFCIDs
- > Most IFCIDs are grouped into CLASSES ... and some IFCIDs may be found in multiple CLASSES

## Instrumentation Facility DB2 Trace Types and Classes

Trace Type	Trace Classes
Accounting [SMF 101]	1-3, 5, 7-8; 4 [installation defined]; 6, 9-29 [reserved]; 30-32 [available for local use]
Statistics [SMF 100 + 102]	1, 3-6, 8; 2 [installation defined]; 7, 9-29 [reserved]; 30-32 [available for local use]
Performance [SMF 102]	1-14, 16-17, 20-22; 15 [installation defined]; 18-19, 23-29 [reserved]; 30-32 [available for local use]
Audit [SMF 102]	1-8; 9 [installation defined]; 10-29 [reserved]; 30-32 [available for local use]
Monitor	1-3, 5-8; 4 [installation defined]; 9-29 [reserved]; 30-32 [available for local use]



International DB2 Users Group

## Instrumentation Facility DB2 Trace Types and Classes

Trace Type	TRACE CLASS	DESCRIPTION	IFCID
Accounting	1	STANDARD ACCOUNTING DATA	3 106, 239
	2	ENTRY OR EXIT FROM DB2 EVENT SIGNALLING – "in-DB2" Time	232
Statistics	3	ELAPSED WAIT TIME IN DB2	6-9, 32-33, 44-45, 117-118, 127-128, 170-171, 174-175, 213-216, 226-227, 242-243, 321-322,329
Performance	5	TIME SPENT PROCESSING IFI REQUESTS	187
Audit	7	ENTRY OR EXIT FROM DB2 EVENT SIGNALLING FOR PACKAGE AND DBRM ACCOUNTING – "in-DB2" Time	232, 240
Monitor	8	WAIT TIME FOR A PACKAGE	6-9, 32-33, 44-45, 117-118, 127-128, 170-171, 174-175, 213-216, 226-227, 241-243, 321-322



International DB2 Users Group

This chart shows the available trace TYPES (except for the GLOBAL trace) and their corresponding trace CLASSES. For a description of each trace class ... refer to the –START TRACE command in the [DB2 for z/OS Command Reference](#) manual.

Note that trace classes 30-32 for all trace types are available for local use ... more on this later!

This chart also shows the SMF record type for the various trace types when DESTINATION is SMF:

- > **Accounting** - **SMF 101:** IFCIDs: 3 (Agent Accounting) \*  
239 (Agent Accounting Overflow)
- > **Statistics** - **SMF 100:** IFCIDs: 1 (System Services Statistics)  
2 (Database Services Statistics)  
202 (dynamic System Parameters)  
230 (Data Sharing Global Statistics)
- **SMF 102:** All other IFCIDs
- > **Performance** - **SMF 102:** All IFCIDs
- > **Audit** - **SMF 102:** All IFCIDs

\* Accounting data for class 1 is accumulated by several DB2 components during normal execution. This data is then collected at the end of the accounting period (thread termination); it does not involve as much overhead as individual event tracing.

On the other hand, when you start class 2,3,7, or 8, many additional trace points are activated. Every occurrence of these events (IFCIDs) is traced internally by DB2 trace, however, these traces are not written to any external destination. Rather, the accounting facility use these traces to compute the additional total statistics that appear in the accounting record ... IFCID 003, when class 2 or class 3 is activated. Accounting class 1 must be active to externalize the information.

[Reference:](#) DB2 for z/OS Administration Guide

## Instrumentation Facility DB2 Trace Types and Classes

Trace Type	TRACE CLASS	DESCRIPTION	IFCID
	1	STATISTICS DATA	1, 2, 105-106, 202
Accounting	3	DEADLOCK, GROUP BUFFER POOL, DATA SET EXTENSION INFORMATION, AND INDICATIONS OF LONG-RUNNING URS AND ACTIVE LOG SPACE SHORTAGES	172, 196, 250, 258, 261-262, 313, 330, 337
<b>Statistics</b>	4	DB2 EXCEPTIONAL CONDITIONS	191-195, 203-210, 235-236, 238, 267, 268
Performance	5	DB2 DATA SHARING STATISTICS RECORD	230
Audit	6	STORAGE USAGE DETAILS	225
Monitor	8	DATA SET I/O STATISTICS	199



International DB2 Users Group

## Instrumentation Facility DB2 Trace Types and Classes

Trace Type	TRACE CLASS	DESCRIPTION	IFCID
	3	SQL EVENTS	22, 53, 55, 58, 59-62, 63, 64-66, 92, 95-96, 106, 112, 177, 233, 237, 250, 272-273, 325
Accounting	4	READS TO AND WRITES FROM THE BUFFER AND EDM POOLS	6-10, 29-30, 105-107, 127-128, 226-227, 321-322
Statistics	8	DATA SCANNING DETAIL	13-17, 18, 105-107, 125, 221-222, 231, 305, 311
<b>Performance</b>	9	SORT DETAIL	26-28, 95, 96, 106
Audit	16	DISTRIBUTED PROCESSING	157-163, 167, 183
Monitor	20	DATA SHARING COHERENCY SUMMARY	249-251, 256-257, 261-262, 267-268
	21	DATA SHARING COHERENCY DETAIL	255, 259, 263



International DB2 Users Group

## Instrumentation Facility Reporting Methods

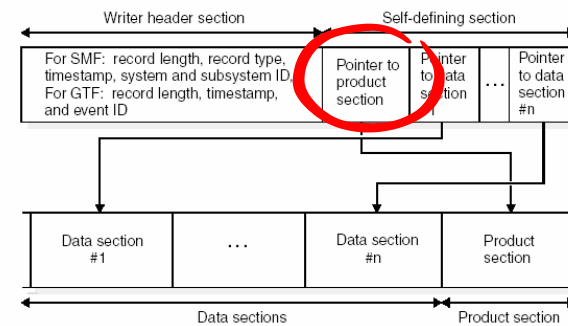
- DB2 Traces are activated by the `--START TRACE` command ...  
`--START TRACE(PERFM) CLASS(30) IFCID(22,63,96,125) DEST(SMF)  
 PLAN(QMFPLAN) AUTHID(JBORGES) ...`
- Trace data can be written to SMF, GTF, MVS OS dataset or Instrumentation Facility Interface (usually used by online monitors)
- IF component provides a trace facility to record DB2 data and events, however analysis and reporting of the trace records must take place outside of DB2
- Use DB2PM (or similar product) to ...
  - View trace data online
  - Format, print and interpret DB2 trace output
  - Generate files [ which may be loaded into DB2 tables ] ✓
- Or, format trace data using your own program...



International DB2 Users Group

## Instrumentation Facility How to format an IFCID

- Formatting trace data using your own program!



International DB2 Users Group

Trace records contain up to four basic sections:

- An SMF or GTF writer header section
- A self-defining section
- A product section
- Zero or more data sections

The writer header section begins at the first byte of the record and continues for a fixed length. Writer headers for "statistics" records (SMF type 100) are mapped by macro DSNDQWST, for "accounting" records (SMF type 101) by macro DSNDQWAS, and for "performance", "audit" and "monitor" (SMF type 102) by macro DSNDQWSP. When these macros are assembled, they include the other macros necessary to map the remainder of the trace records sent to SMF.

The self-defining section follows the writer header section. The first self-defining section always points to a special data section called the product section. Among other things, the product section contains an instrumentation facility component identifier (IFCID). Descriptions of the records differ for each IFCID... the table on the following foil shows the mapping macros for the various IFCIDs.

## Instrumentation Facility How to format an IFCID

- Use assembler mapping macros in SDSNMACS to interpret DB2 trace records

IFCID	Mapped by Macro
0001	DSNDQWST, subtype=0
0002	DSNDQWST, subtype=1
0003	DSNDQWAS
0004 – 0057	DSNDQW00
0058 – 0139 (except 0106)	DSNDQW01
0106	DSNDQWPZ
0140 – 196, 198, 199	DSNDQW02
0201 – 0249 (except 202, 230 and 239)	DSNDQW03
0202	DSNDQWS2, subtype=2
0230	DSNDQWST, subtype=3
0239	DSNDQWAS and SNDQWA1
0250 – 0330	DSNDQW04

See “Interpreting DB2 Trace Output”  
Appendix D - DB2 Administration Guide Volume 2




International DB2 Users Group

## Building a DB2 Performance Database (PDB)



- Benefits
- Data Sources
- Data Extraction
- Database Structure
  - Accounting
  - Statistics
  - Performance trace (IFCIDs) (Rid Pool, Sort, Locking, etc)
  - Exception
  - Business Metrics
- Data Summarization/Aggregation
- Data Retention



International DB2 Users Group

To interpret a record ... find its description (by IFCID) in one of the provided mapping macros.  
For example: Macro **DSNDQWST** (subtype 1) maps out IFCID 2 - Database Services Statistics



## Building a DB2 PDB ... Benefits

- Pro-Active Monitoring
  - Alerts
  - Triggers
  - Prevent problems before they happen!
- Automation
  - Increases efficiency
- Data Available to everyone!
  - Resulting in better performing applications
- A MUST for Capacity Planning
  - DB2 subsystem
  - Application



International DB2 Users Group



## Building a DB2 PDB ... Data Sources

- SMF
  - Accounting (Classes 1,2,3,7,8)
  - Statistics (Classes 1,2,4,5,6,8)
  - Performance trace (various IFCIDs)
- DB2PM Batch Reporter
  - Exceptions
- Miscellaneous Sources
  - Other PDBs (e.g. CICS, RMF, MQSeries, ...)
  - Application/Business metrics




International DB2 Users Group




## Building a DB2 PDB ... Data Extraction

- Use DB2PM (or similar program) to extract DB2 trace data into files
  - Accounting
  - Statistics
  - Performance trace
- Use DB2PM to create DB2 “Exceptions” file
- Load extracted data into DB2 PDB



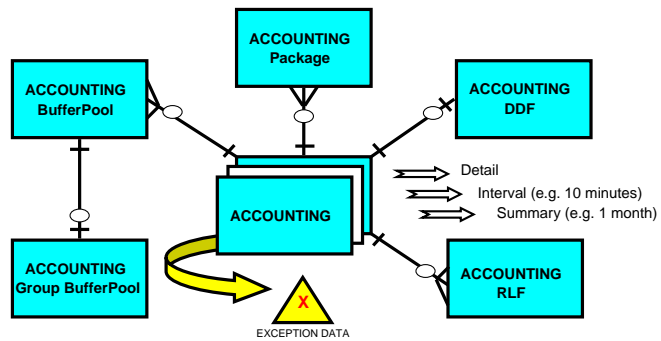
## Building a DB2 PDB ... Database Structure

- DB2 PDB is made up of 3 main components



## Building a DB2 PDB ... Database Structure

- DB2 PDB – Accounting Data component

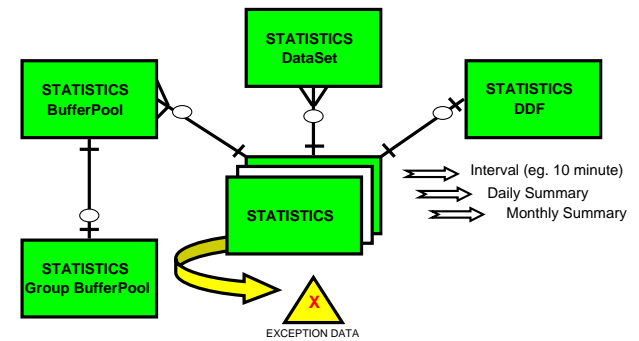


International DB2 Users Group



## Building a DB2 PDB ... Database Structure

- DB2 PDB – Statistics Data component




International DB2 Users Group



## Building a DB2 PDB ... Database Structure

- DB2 PDB – Performance Data component




International DB2 Users Group



## Building a DB2 PDB ... Data Summarization/Aggregation

- Summarization may occur at EXTRACTION time
  - DB2PM interval summarization (e.g. Accounting\_Interval)
- Otherwise ... Summarization and/or Aggregation occurs after the data is loaded into the DB2 PDB

Two Methods:  
1) SQL queries  
2) Triggers



International DB2 Users Group



## Building a DB2 PDB ... Data Retention

- Detail data
  - **Accounting**
    - as needed ... then delete
  - **Statistics**
    - 550 days (6 quarters)
  - **Performance**
    - as needed ... then delete
  - **Exceptions**
    - 365 days (1 year)
- Summary data
  - **Accounting**
    - 10 minute interval - 550 days (6 quarters)
    - monthly summary/aggregation - 10 years +
  - **Statistics**
    - daily summary - 5 years
    - monthly summary/aggregation - 10 + years



International DB2 Users Group



## Reporting and Exception Monitoring



- Build Your Own Reports
- Historical Trending
- Exception Monitoring
  - Set Your Thresholds
  - Set off Alerts and Trigger Actions
- Adhoc Reporting
- Sample Reports



International DB2 Users Group

## Reporting & Exception Monitoring Build Your Own Reports

- Now that your DB2 PDB is a relational database, you have many choices of tools to “Build Your Own Reports”!
  - QMF, spreadsheets, GUI tools, Web ... or your favourite report writer
- Tailor reports to suit your needs and requirements
- Report on...
  - Key performance metrics
  - Historical trends
  - Exceptions
  - Bottlenecks
  - Adhoc
  - Etc...



International DB2 Users Group



## Reporting & Exception Monitoring Historical Trending

### 1) Statistics: DBM1 Storage Trend ... (IFCID 225) STATS CLASS(6)

```
QUERY:  Select date, interval, available_storage, getmained_storage,
        virtual_bufferpool
        from w1z.db2_stat_intrvl
        where local_location = 'dbga' and member = 'dba1'
        and date between '2002-07-10' and '2002-07-16'
```

DATE	INTERVAL	AVAI LABLE STORAGE	GETMAI NED STORAGE	VI RTUAL BUFFERPOOL
2002-07-10	14. 31. 44	902848512	607201096	514048000
2002-07-11	14. 34. 03	897490944	610728552	514048000
2002-07-12	14. 31. 44	892841984	612335784	514048000
2002-07-13	14. 38. 03	885747712	616972456	514048000
2002-07-14	14. 35. 52	887836672	613703848	518144000
2002-07-15	14. 35. 21	882028544	619028648	518144000
2002-07-16	14. 33. 39	850505728	642920616	542720000



International DB2 Users Group



```

SELECT A.DATE
, A.PRIMAUTH
, SUM (CASE WHEN SUBSYSTEM_ID='DBA1'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.42
      WHEN SUBSYSTEM_ID='DBA2'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.67
      WHEN SUBSYSTEM_ID='DBA3'
      THEN (CLASS1_CPU_NNESTED)/60 * 14.02
      END)
      AS CLASS1_CPU
, SUM (CASE WHEN SUBSYSTEM_ID='DBA1'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.42
      WHEN SUBSYSTEM_ID='DBA2'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.67
      WHEN SUBSYSTEM_ID='DBA3'
      THEN (CLASS1_CPU_NNESTED)/60 * 14.02
      END) / SUM(COMMIT)
      AS CPU_PER_COMMIT
, SUM (COMMIT)
      AS COMMIT
, SUM (SELECT + INSERT + UPDATE + DELETE + DESCRIBE_TABLE
      + DESCRIBE + PREPARE + OPEN + FETCH + CLOSE )
      AS DML_SQL
-----
FROM WIZ.DB2_ACCT_INTRVL A
-----
WHERE A.DATE BETWEEN '2002-01-01' AND '2003-01-31'
AND A.LOCAL_LOCATION = 'DBGA'
AND A.SUBSYSTEM_ID IN ('DBA1','DBA2','DBA3')
AND A.CONNECT_ID IN ('CICS006','CICS006B','CICS006C')
AND A.PRIMAUTH IN ('ENRQ')
-----
GROUP BY A.DATE, A.LOCAL_LOCATION,
A.PRIMAUTH

```

**NOTE:** CPU converted to Model 200E

```

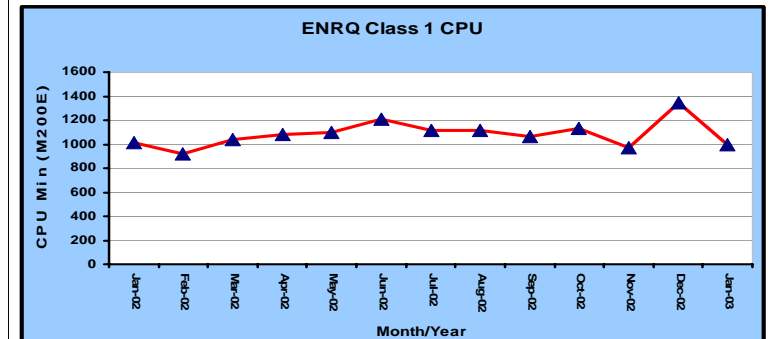
SELECT A.DATE
, A.PRIMAUTH
, SUM (CASE WHEN SUBSYSTEM_ID='DBA1'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.42
      WHEN SUBSYSTEM_ID='DBA2'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.67
      WHEN SUBSYSTEM_ID='DBA3'
      THEN (CLASS1_CPU_NNESTED)/60 * 14.02
      END)
      AS CLASS1_CPU
, SUM (CASE WHEN SUBSYSTEM_ID='DBA1'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.42
      WHEN SUBSYSTEM_ID='DBA2'
      THEN (CLASS1_CPU_NNESTED)/60 * 13.67
      WHEN SUBSYSTEM_ID='DBA3'
      THEN (CLASS1_CPU_NNESTED)/60 * 14.02
      END) / SUM(COMMIT)
      AS CPU_PER_COMMIT
, SUM (COMMIT)
      AS COMMIT
, SUM (SELECT + INSERT + UPDATE + DELETE + DESCRIBE_TABLE
      + DESCRIBE + PREPARE + OPEN + FETCH + CLOSE )
      AS DML_SQL
-----
FROM WIZ.DB2_ACCT_INTRVL A
-----
WHERE A.DATE BETWEEN '2002-01-01' AND '2003-01-31'
AND A.LOCAL_LOCATION = 'DBGA'
AND A.SUBSYSTEM_ID IN ('DBA1','DBA2','DBA3')
AND A.CONNECT_ID IN ('CICS006','CICS006B','CICS006C')
AND A.PRIMAUTH IN ('ENRQ')
-----
GROUP BY A.DATE, A.LOCAL_LOCATION,
A.PRIMAUTH

```

**NOTE:** CPU converted to Model 200E

## Reporting & Exception Monitoring Historical Trending

### 3) Accounting: Online Application – Workload Trend (graph)




International DB2 Users Group



## Reporting & Exception Monitoring Exception Monitoring

- Set Your Threshold Values and Build Queries to monitor ... it's that easy!
- Automate monitoring process
- When threshold values are exceeded ... fire off an alert and trigger corrective action!!



## Reporting & Exception Monitoring Exception Monitoring

1) Accounting I/O Response Exception: Average I/O > 20 ms

```

QUERY: select Interval, plan, prlmauth as txid, db_io,
        class3_db_io / db_io as avg_io_resp
from wtz.db2_acct_intrvl
Where date = '2002-12-23' and connect_id = 'clcs006'
and (class3_db_io / db_io) > 0.020 and db_io > 0
  
```

INTERVAL	PLAN	TXID	DB IO	AVG IO RESP
21. 30. 00	PLNCC	NCAL	13	0.0784
03. 20. 00	PLNCC	NSIV	16	0.0484
14. 00. 00	PLNCC	NA1D	46	0.0363
23. 20. 00	PLNCC	SAPR	5	0.0313
14. 30. 00	PLNCC	NEFL	2	0.0258





## Reporting & Exception Monitoring

### Exception Monitoring

#### 2) EDM Pool Statistics Exception: EDM Pool Full

```
QUERY: select date, interval, local_location, member_name, fail_pool_full
from w.z.db2_stat_intrl
where fail_pool_full > 0 and date >= '2002-07-01'
```

DATE	INTERVAL	LOCAL LOCATION	MEMBER NAME	FAIL POOL-FULL
2002-07-15	20. 32. 20	DBG0	DB02	17
2002-07-18	23. 05. 43	DBG0	DB01	10
2002-08-09	10. 50. 04	DBGB	DBB4	5
2002-08-30	11. 58. 09	DBGB	DBB2	35
2002-08-30	12. 08. 08	DBGB	DBB2	147
2002-08-30	12. 18. 09	DBGB	DBB2	8



International DB2 Users Group



## Reporting & Exception Monitoring

### Exception Monitoring

#### 3) EXCEPTION Table: Check Exceptions Daily!

```
QUERY: select db2id, instance, field_event, description,
compare_basis, op, excp_value, thresh_value, excp_lvl
from w.z.db2_exceptions
where local_location = 'dbga' and date = '2002-12-23'
```

DB2ID	INSTANCE	FIELD EVENT DESCRIPTION	COMPARE BASIS	OP	EXCP VALUE	THRESH VALUE	EXCP LVL
DBA1	B888B4F8E50B	PLAN TOTAL CLASS 3 SUSPENSIONS	M	>	215	180	W
DBA1	B888B4F9E8AE	PLAN TOTAL CLASS 3 SUSPENSIONS	M	>	583	240	P
DBA1	B888B4D04017	PLAN GETPAGES PER SYNCHRONOUS READ I	V	<	1	2	W
DBA3	B879B857D7D	SYST ROWS IN MESSAGE BUFFER	V	<	0	1	P
DBA1	B888B4FE7424	PLAN SERV.TASK SWITCH SUSP. TIME (C	M	>	22.929666	7	P
DBA1	B888B4FFB059	PLAN SERV.TASK SWITCH SUSP. TIME (C	M	>	5.279576	5	W
DBA1	B888B51F2F15	PLAN RID LIST TERMINATED - ANY REAS	V	>	2	0	P
DBA3	B888AEA4C598	PLAN GETPAGES	V	>	3714849	200000	P
DBA1	B89D7A78CCD2	SYST TOTAL SQL DDL STATEMENTS	V	>	3	0	P
DBA1	B89D7A78CCD2	SYST RID LIST NOT USED - ANY REASON	V	>	1	0	P



International DB2 Users Group



## Reporting & Exception Monitoring Adhoc Reporting

"Who has used QMF in the last 3 months, and how frequently?"

QUERY: `Select plan, pri mauth, count(distinct date) as frequency_days`

**Adhoc Report**

PLAN NAME	PRI MAUTH	FREQUENCY (DAYS)	PRI MAUTH	FREQUENCY (DAYS)
QMFPLAN	C5889DR	90	P8094UL	41
	DTOMPKI	88	M8007D1	39
	AJWHI TE	47	JHONG	37
	JANI CEA	45	C4064PC	35
	T5936JR	43	DBABA	35



International DB2 Users Group



## Monitoring KEY Performance Metrics



- Key Performance Metrics
  - Accounting
  - Statistics
- Key Performance IFCIDs
  - Dynamic SQL
  - Explain
  - RID Pool
  - Sort
  - and many more...



International DB2 Users Group

# Monitoring KEY Performance Metrics

## Key Performance Metrics

- Accounting: IFCID003

- Most commonly used DB2 trace record
  - Perform program related tuning
  - Charge DB2 costs
- Trace record written at thread termination
- Accounting traces provide information at the application level...

Elapsed time, CPU time, Wait time and counts, SQL events, Buffer Pool metrics, Locking, Data Sharing information, Package level metrics, etc...



International DB2 Users Group



IDENTIFICATION		CLASS 2 TIME DISTRIBUTION		CLASS 3 SUSPENSIONS		EVENTS		HIGHLIGHTS	
APPL	10%	DB2	8%	SUSP	82%	NOTACC	2%	SUSP	91%
TYPE	EVENTS	APPL (CL. 1)	DB2 (CL. 2)	IFI (CL. 5)	CLASS 3 SUSPENSIONS	ELAPSED TIME	EVENTS	HIGHLIGHTS	
2905	LOCK/LATCH(DB2-IRLM)	2:42	147920	N/P		2905	THREAD TYPE	ALLIED	
11808	SYNCHRON I/O	1:37	970261	N/A		11808	TERM. CONDITION	NORMAL	
11645	DATABASE I/O	1:33	620410	N/A		11645	INVOKE REASON	DEALLOC	
163	LOG WRITE I/O	4:349851	163	N/A		163	COMMIT S	1672	
97	OTHER READ I/O	1:836027	97	N/A		97	ROLLBACK	0	
0	OTHER WRITE I/O	0:000000	0	N/A		0	SVPT REQUESTS	0	
51	SER. TASK SWITCH	51:070366	1666	N/P		51	SVPT RELEASE	0	
0	UPDATE COMMIT	0:000000	0	N/A		0	SVPT ROLLBACK	0	
0	OPEN/CLOSE	0:000000	0	N/P		0	INCRM BINDS	0	
0	SYSLOGRG REC	0:000000	0	N/A		0	UPDATE/COMMIT	8.76	
0	EXT/DEL/DEF	0:000000	0	N/A		0	SYNCH I/O AVG.	0.008297	
0	OTHER SERVICE	0:000000	0	N/A		0	PROGRAMS	9	
0	ARC LOG (QUE) ES	0:000000	0	N/A		0	MAX CASCADE	0	
0	ARC LOG READ	0:000000	0	N/A		0	PARALLELISM	NO	
0	STOR. PRC SCHED	0:000000	0	N/A		0			
0	UDF SCHEDULE	0:000000	0	N/A		0			
0	DRAIN LOCK	0:000000	0	N/A		0			
0	CLAIM RELEASE	0:000000	0	N/A		0			
90	PAGE LATCH	0:709714	90	N/P		90			
0	NOTIFY MSGS	0:000000	0	N/A		0			
1191	GLOBAL CONTENTION	9:13	014067	N/A		1191			
0	COMMIT PH1 WRITE I/O	0:000000	0	N/A		0			
0	ASYNCH IXL REQUESTS	0:000000	0	N/P		0			
14:26	TOTAL CLASS 3	14:26	748354	N/P		17757			



ACCOUNTING: IFCID003 – Sample Report (partial)



Class3 Wait: Global contention chart ...

GLOBAL	CONTENTION	L-LOCKS	ELAPSED TIME	EVENTS
L-LOCKS			9: 12. 924287	1183
PARENT	(DB, TS, TAB, PART)		0. 313377	104
CHILD	(PAGE, ROW)		9: 12. 610910	1079
OTHER			0. 000000	0
GLOBAL	CONTENTION	P-LOCKS	ELAPSED TIME	EVENTS
P-LOCKS			0. 089780	8
PAGESET/PARTITION			0. 000000	0
PAGE			0. 000000	0
OTHER			0. 089780	8

## Monitoring KEY Performance Metrics Key Performance Metrics

1) Accounting: Threads with CLASS1\_CPU > 30 minutes

QUERY: (see NOTES page ...)

```

SELECT DATE, LOCAL_LOCATION, PRMAUTH, CORRNAME
      , INSTANCE_NBR
      , CLASS1_CPU_NNESTED AS CLASS1_CPU
      , CLASS1_CPU_NNESTED AS CLASS2_CPU
FROM WIZ.DB2_ACCT_DETAIL
WHERE LOCAL_LOCATION = 'DBG1'
      AND DATE BETWEEN '2002-05-01' and '2002-06-30'
      AND CLASS1_CPU_NNESTED / 60 > 30
ORDER BY 6 DESC
    
```



International DB2 Users Group



### Query:

```

SELECT DATE, LOCAL_LOCATION, PRMAUTH, CORRNAME, INSTANCE_NBR
      , CLASS1_CPU_NNESTED AS CLASS1_CPU
      , CLASS1_CPU_NNESTED AS CLASS2_CPU
FROM WIZ.DB2_ACCT_DETAIL
WHERE LOCAL_LOCATION = 'DBG1'
      AND DATE BETWEEN '2002-05-01' and '2002-06-30'
      AND CLASS1_CPU_NNESTED / 60 > 30
ORDER BY 6 DESC
    
```

## Monitoring KEY Performance Metrics Key Performance Metrics

2) Accounting: Threads with CLASS3\_WAIT\_OTHER > 25% of Total Class3 Wait

QUERY: (see NOTES page ...)

DB2ID	PRMAUTH	PLAN NAME	WAIT LOCK	WAIT SYNCH/I/O	WAIT OTHER	WAIT OTHER %	# OCCURS
DBA2	EPOI	PLNCC	5.64	66.12	84.36	54.04	1444
DBA1	MI SCONS	FOCUS709	0.03	35.48	69.07	66.04	18
DBC2	LS2SPRP	SQLACCES	0.76	48.69	54.46	52.41	167130
DBB2	AI PI	AI FI PL16	0.07	53.54	53.09	49.75	10778
DBB2	AI PI	AI FI PLNC	0.21	41.36	43.85	51.33	9410
DBC4	C5889DR	LSI	0.01	6.47	32.89	83.55	158
DSN6	HRMSPPS	PSQED. EX	0.01	12.57	32.73	72.23	471
DSN6	HRMSPPS	PSNVS. EX	0.18	9.93	29.49	74.47	614
DSN6	HRMSPPS	PSAPI SRV	3.27	1.33	26.10	85.00	77061
DSNC	PERHRM7	PSAPI SRV	0.17	1.90	25.35	92.48	49322
DBC2	SAMR	SAMSMR03	0.01	0.01	6.12	99.67	25608
DBA2	MPRD012	PLFDC	0.02	0.11	6.04	97.93	21



International DB2 Users Group



CLASS3\_WAIT\_OTHER is all CLASS3\_WAIT time except for WAIT\_LOCK\_LATCH and WAIT\_SYNCH\_IO.

### Query:

## Monitoring KEY Performance Metrics Key Performance Metrics

- **Statistics:** IFCID002 (Database Statistics)
  - Trace record written at end of interval
  - Statistics traces provide information at the DB2 subsystem level...

**SQL (DML,DDL,DCL),  
 Subsystem Services, EDM Pool, LOG Activity,  
 RID LIST processing,  
 Buffer Pool, Group Buffer Pool,  
 Locking, Data Sharing Locking,  
 Query Parallelism, Global DDF Activity, Dynamic SQL,  
 DBM1 Storage usage,  
 DB2 Address Space CPU times,  
 And much more...**



International DB2 Users Group



SQL DML	QUANTITY	/SECOND	/THREAD	/COMMIT
SELECT	48497.3K	565.97	182.67	24.17
INSERT	3306.4K	38.59	12.45	1.65
UPDATE	4034.7K	47.09	15.20	2.01
DELETE	11775.00	0.14	0.04	0.01
PREPARE	2576.00	0.03	0.01	0.00
DESCRIBE	55.00	0.00	0.00	0.00
DESCRIBE TABLE	0.00	0.00	0.00	0.00
OPEN	16000.3K	186.72	60.27	7.98
CLOSE	15986.2K	186.56	60.21	7.97
FETCH	79153.9K	923.73	298.14	39.45
TOTAL	167.0M	1948.82	628.99	83.24



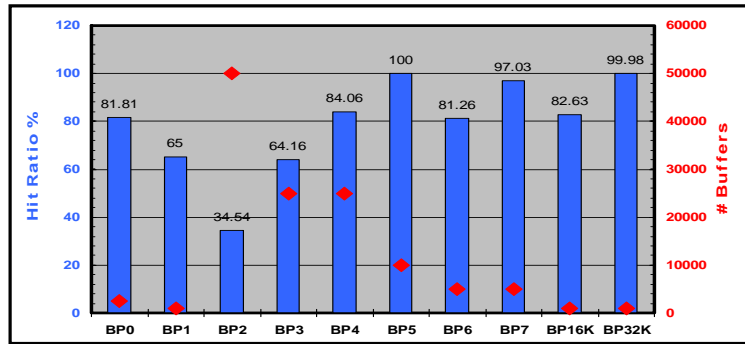
STATISTICS: IFCID001 & IFCID002 – Sample Report (partial)



# Monitoring KEY Performance Metrics

## Key Performance Metrics

### 1) Statistics: Bufferpool Metrics ... [ Bufferpool Size & Hit Ratio ]




International DB2 Users Group



# Monitoring KEY Performance Metrics

## Key Performance Metrics

### 2) Statistics: Bufferpool Metrics ... [ Bufferpool Hit Ratio Trend ]



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance Metrics

### 3) Statistics: Bufferpool Data Set metrics ... (IFCID 199) STATS CLASS(8)

```
QUERY: select date, interval, bp_id, syn_io_avg_delay, sync_io,
        asyn_io_avg_delay, asyn_io
from wtz.db2_statsn_intrvl
Where local_location = 'dbga' and space_name = 'xnccacm1' and date ...
```

DATE	INTERVAL	SYN_IO		ASYN_IO		
		BP ID	AVG DELAY	SYNC IO	AVG DELAY	ASYN IO
2002-12-12	14. 10. 05	4	4	1006428	1	45087
2002-12-13	14. 10. 21	4	3	1150367	0	51802
2002-12-14	14. 10. 23	4	7	25704	2	301993
2002-12-15	14. 10. 27	4	5	14026	1	407
2002-12-16	14. 10. 33	4	3	1259252	0	89641
2002-12-17	14. 10. 41	4	3	1107219	1	49272
2002-12-18	14. 10. 49	4	4	1093357	1	52508



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

- Dynamic SQL – IFCID 63
- Explain Information – IFCID 22
- RID Pool – IFCID 125
- Sort – IFCID 96
- And many more ...



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

- Dynamic SQL performance analysis is difficult ... especially after the fact!
  - Which query used all the CPU?
  - We know **who** did it ... but which query was it?
- Using **IFCIDs 003, 022** and **063** we can get our answers!
  - IFCID 003 will identify the **EXPENSIVE** thread
  - IFCID 022 will provide the **EXPLAIN** data
  - IFCID 063 will identify the **SQL STATEMENT**



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

- RID Pool – IFCID 125 (Exception: RID Pool Failures)

LOCAL\_LOCATION: **DSN2WCB** SUBSYSTEM\_ID: **DSN2** DATE: **2002-07-11**

LW I N S T A N C E	P R I M A U T H	T I M E S T A M P	R I D U S A G E	R I D N O T _ U S E D * S T O R *	R I D N O T _ U S E D * M A X *	R I D S I N _ F I N A L L I S T
B7E96DB89551	JBORGES	2002-07-11-17.18.26.287192	N	Y		-2
B7E9A0BB32E7	KULKER	2002-07-11-21.09.00.645262	N		Y	-2

LW I N S A N C E	P L A N N A M E	P R O G N A M E	S T A T E M E N T	D B I D	O B I D	I N D E X R I D S	I N D E X T H R E S H O L D
B7E96DB89551	QMFPLAN	DSOCFSQL	2771	471	6	-2 (NO STORAGE)	327680
B7E9A0BB32E7	ACT711DM	ACTBOT	628	632	6	-3 (MAX LIMIT)	65536



International DB2 Users Group





## Monitoring KEY Performance Metrics Key Performance IFCIDs

- Sort – IFCID 96 (Sort End Statistics)

LOCAL_LOCATION: DSN2WCB SUBSYSTEM_ID: DSN2 PRI MAUTH: KULKER					
PLAN_NAME: J458YPLN PROGRAM_NAME: J458Y244 LUW_INSTANCE: B7EA2546F9ED					
STMT NO	TIMESTAMP	RETURN CODE	RECORDS SORTED	SORT_DATA AREA	SORT_KEY SIZE
117	2002-07-12-13.32.01.517807	0	1207630	20	19
117	2002-07-12-13.55.46.044012	0	1030156	20	19
117	2002-07-12-14.02.57.666741	0	813664	20	19
SORT_RECORD SIZE	WORK FILES	INIT_WORK FILES	MERGE PASSES	SORT TYPE	
39	29	28	1	ESA	
39	28	27	1	ESA	
39	20	19	1	ESA	



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

- And many more ...
  - Data sharing
    - IFCID 230 – Group Buffer Pool Attributes
    - IFCID 254 – Coupling Facility Cache Structure Statistics
    - IFCID 261 – Group Buffer Pool Checkpoint
    - IFCID 262 – GBPOOLT Castout Threshold Processing
    - IFCID 263 – PAGE Set and Partition Castout Detail
  - Parallelism
    - IFCID 221 – Parallel Group Execution
    - IFCID 222 – Parallel Group Elapsed Time
    - IFCID 231 – Parallel Group Task Time



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

- And many more ...
  - Locking
    - IFCID 172 – Deadlock Data
    - IFCID 196 – Timeout Data
    - “Various” Lock Suspend/Resume – IFCID 44-45, 213-216, 226-227
  - Scan End
    - IFCID 18 – Scan End
  - End SQL
    - IFCID 58 – End SQL
  - Long Running Units of Work
    - IFCID 313 – Uncommitted Unit of Recovery
  - Etc ...



International DB2 Users Group



## Monitoring KEY Performance Metrics Key Performance IFCIDs

There are 220+ IFCIDs ...  
and growing with every release!  
A wealth of information  
**Ready to be mined!**




International DB2 Users Group



## Summary ...

**NO**  
**DB2 Performance Monitoring**  
**Is complete**  
**Without a**  
**DB2 Performance Database.**  
**Start Mining Your**  
**DB2 PDB!**



International DB2 Users Group

How to Build Your Own DB2 Performance Database  
Session: A6

# Thank-You!

**Joe Borges**  
**BMO Financial Group**  
**joe.borges@bmo.com**

