Toronto Users Group - March 19 2014
DB2 Performance: Utilizing DB2 for i More Efficiently

Speaker Name Linda M Swan
lmswan@us.ibm.com

New Wiki for DB2 Enhancements via PTF

- **Regularly check (or Subscribe) to the DB2 for i Updates Wiki!**
  - Contains details on new PTFs that deliver new DB2 capabilities
  - Examples:
    - CANCEL_SQL system stored procedure
    - PROGRAM NAME keyword for controlling SQL Triggers Program Name
    - SQL Query Engine 6.1 support for Logical File on FROM clause
    - New RUNSQL CL command
  - Wiki URL:

    https://www.ibm.com/developerworks/systems/ibmi/techupdates/db2

- **The wiki is part of the IBM i zone in IBM developerWorks launched in April 2011**

  https://www.ibm.com/developerworks/systems/ibmi/

- **New DB2 for i Blog too! - http://db2fori.blogspot.com/**

Agenda

- **Why care about DB performance**
- **Causes of poor DB performance**
- **What tools are useful**
- **Things to look for**
- **DB performance tips**

Why care about DB performance

- **Key objective is to provide *Information***
  - Which means processing data
    - Sometimes very large quantities of data
  - In timely manner
  - Using resources efficiently
    - Both systems and personnel
  - To accommodate growth – SQL set at a time
    - At the Transactions
    - At the Data volume level
    - **Complexity** of **requests**

Typical causes of DB performance problems

1. **Accessing more data than needed**
2. **Physical disk operations are consecutive (versus concurrent)**
3. **Trying to consume more resource than is available, or more than your share.**
4. **Under utilizing hardware or not using it efficiently**
5. **High full opens or access plan rebuilds**
6. **Inefficient access plans – does the optimizer have what it needs?**

What tools would I use

- **Collection Services**
- **Job Watcher**
- **Disk Watcher**
- **Performance Explorer**
- **Plan Cache Analyzer**

  – Available through
    - IBM Navigator for i - Performance Data Investigator
    - iDoctor

- **iNavigator – SQL and DBA specific tools**
  – Plan Cache, Visual Explain, Index advisor….

What tools would I use …

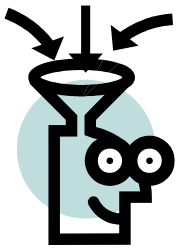

- **Application knowledge**
- **Data knowledge**

Monitor, Analyze and Tune

- **Needs to be someone's Job**
  - DB engineer position
  - OR, cross team (Operations, Developers)
- **Get the right training**
  - **iDoctor workshop**
  - **DB2 for i SQL Performance Workshop**
    - http://www-03.ibm.com/systems/i/software/db2/db2performance.html
- ***Need to establish best practices.***
  - *One being to capture performance data expectations before major changes*

What to look for ?

Look at System Resources

–Memory
- WRKSYSSTS

–I/O Disk Opts
- WRKDSKSTS

–High CPU
- WRKSYSACT

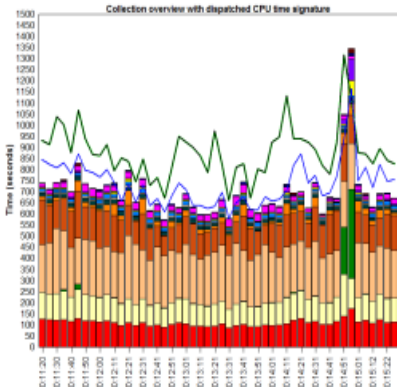

Collect Based on What You Find/*Suspect*

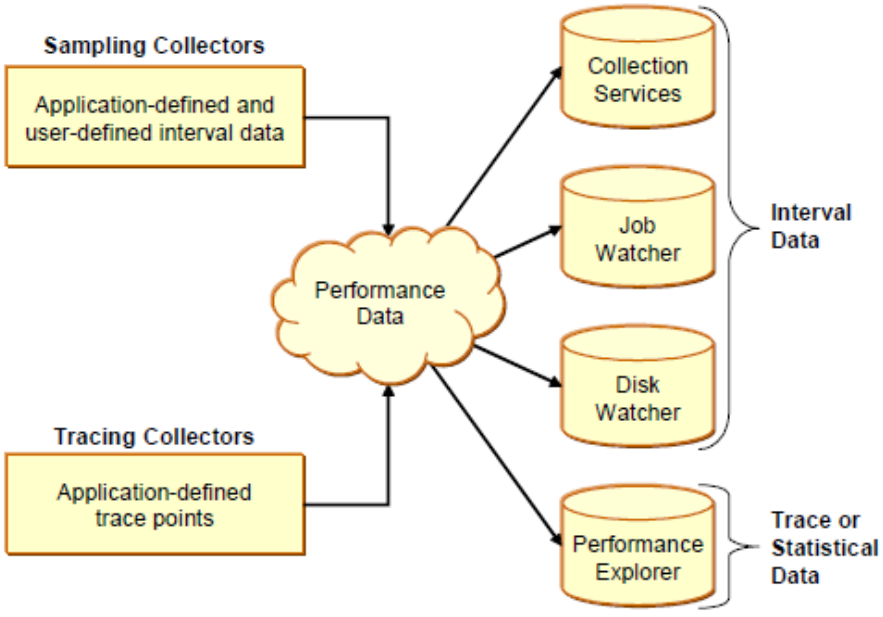

The collections at a glance

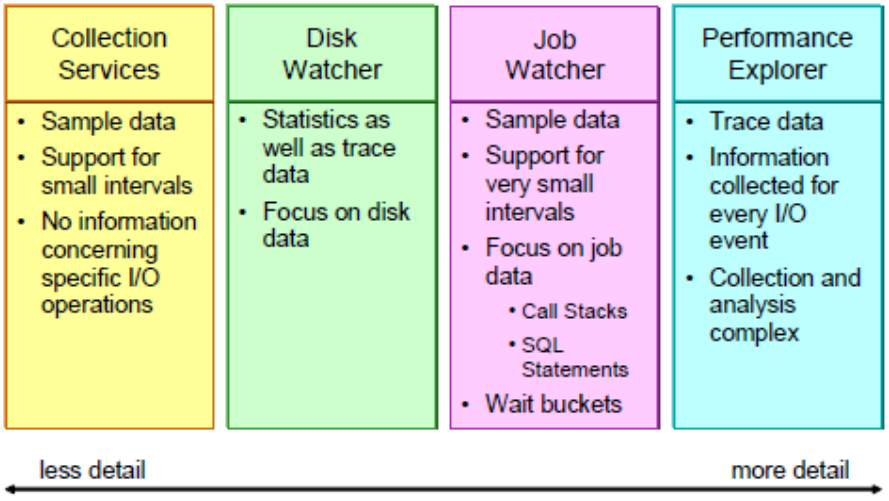

Power is performance redefined

The collections at a glance

Example: Job are typically
- using CPU
- queued up waiting to use CPU
- in an idle wait
- waiting for something specific…JW wait buckets

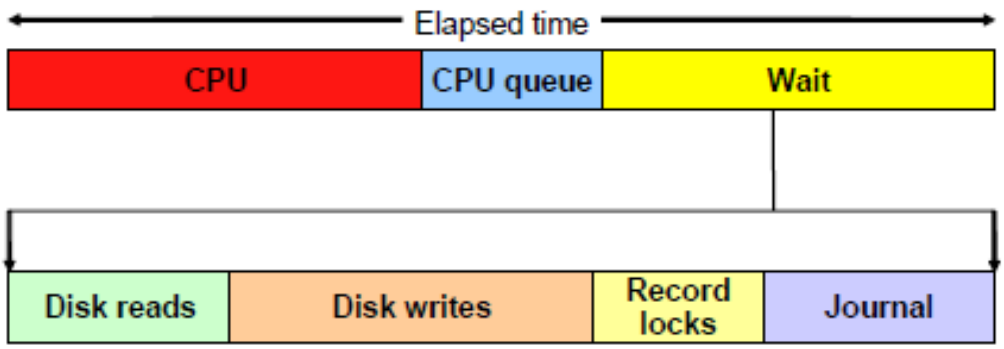

iDoctor

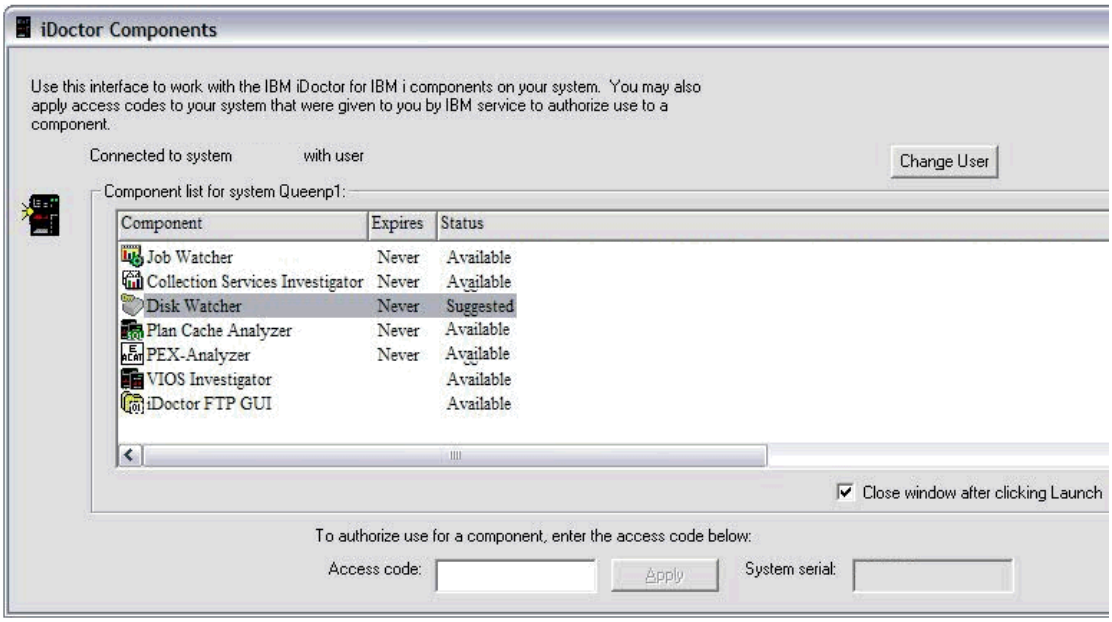

Symptom Example – High CPU
Job Watcher→ CPU Graphs→ dispatched CPU/CPUQ by high/Low Prior..

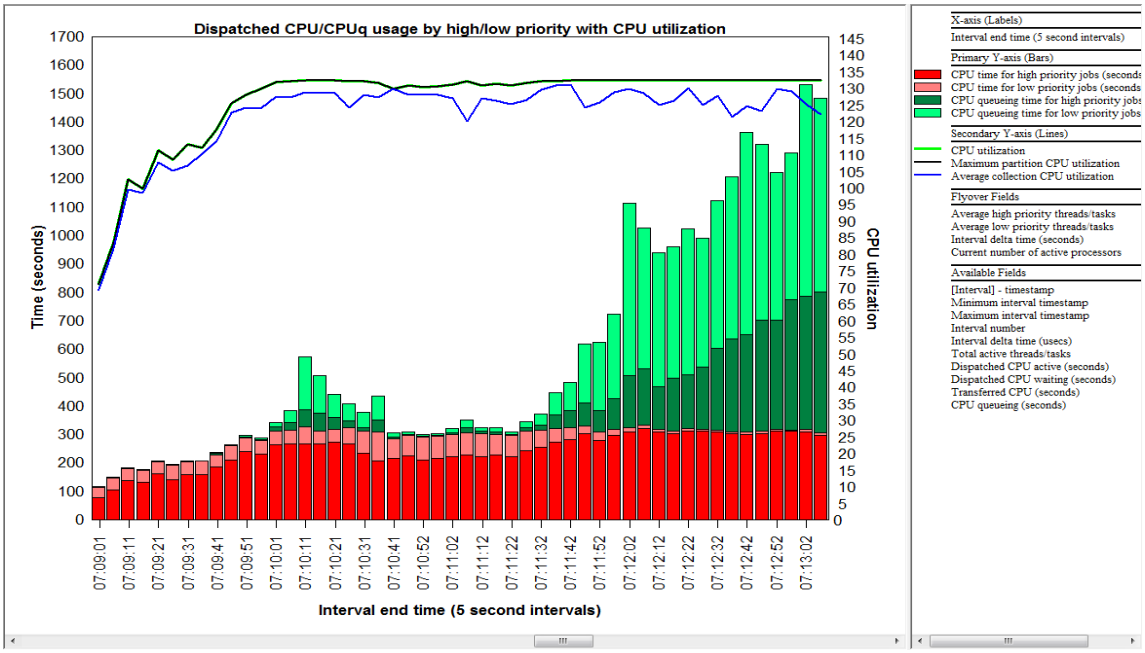

Symptom Example – High CPU
Drill into an interval or set of intervals

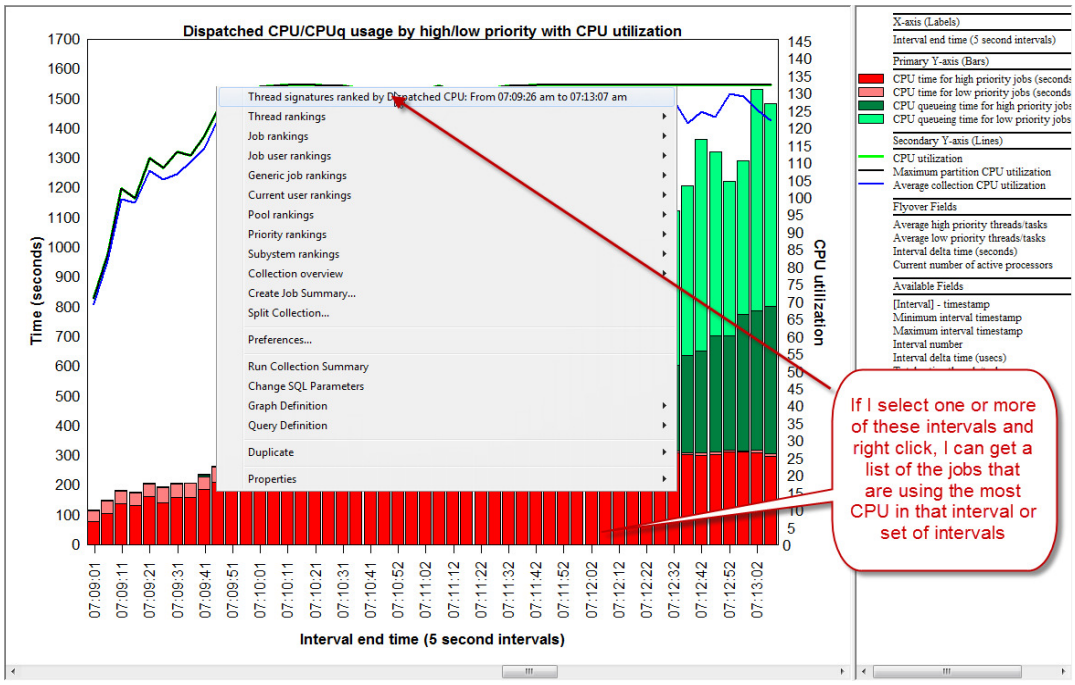

Symptom Example – High CPU
(JW) From list of jobs I can see call stacks to get an idea of What is running

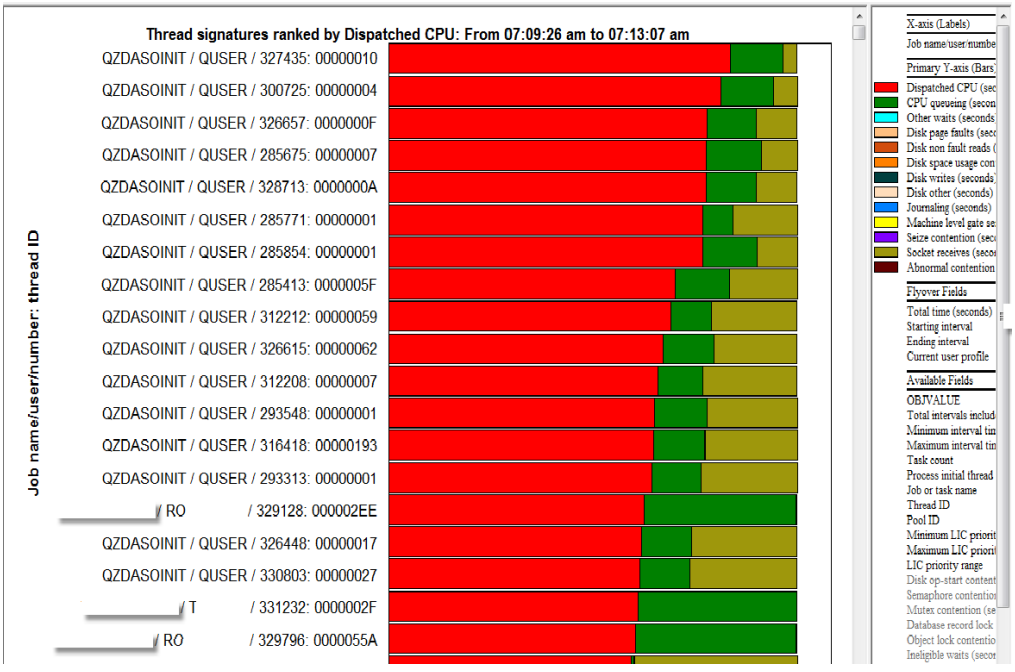

Symptom Example – High CPU
I can have job watcher collect current SQL too

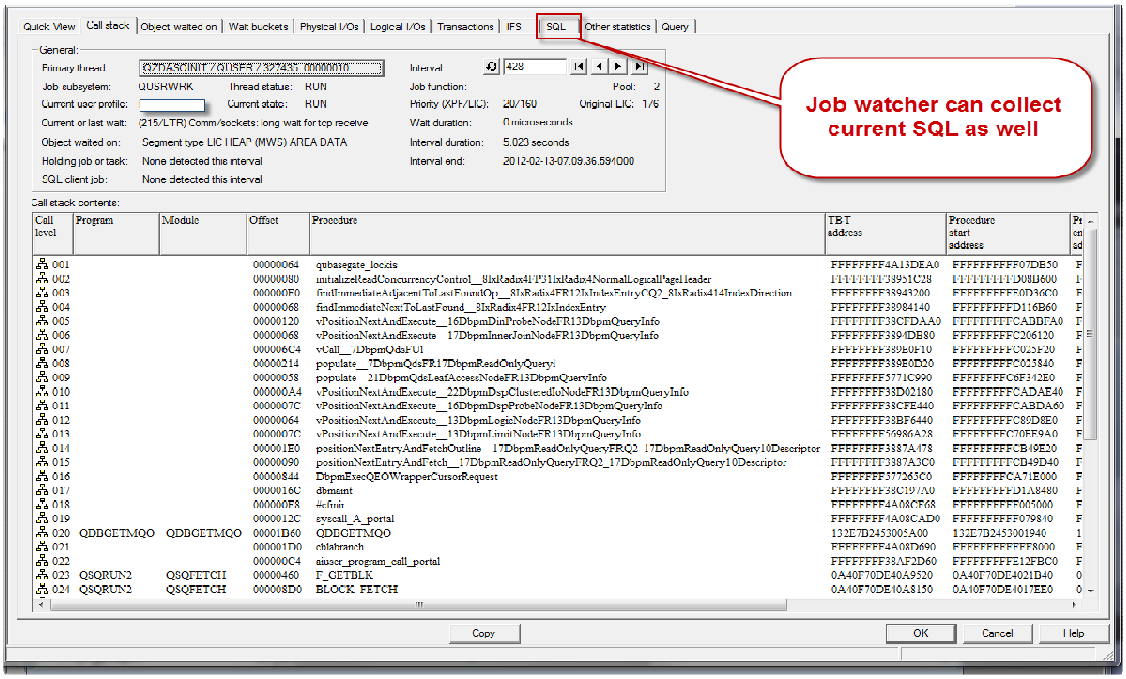

Performance Data Investigator – New Database Perspective

- Robust graphical interface enabled by Collection Services support.
- New breakdown makes it easier to notice “what changed”?

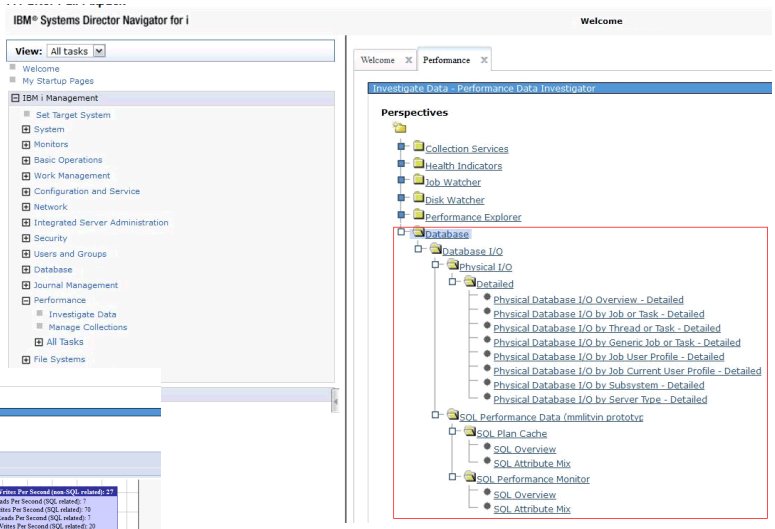

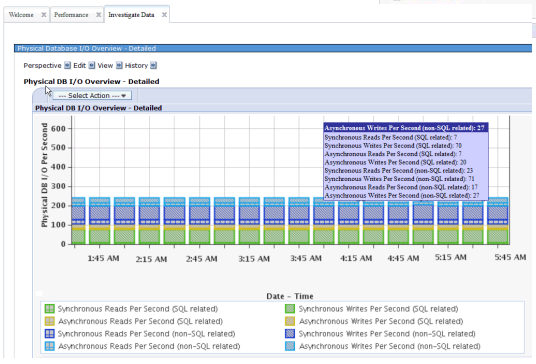

Getting started with CS, JW and iDoctor

- **Get the redbook for**
  - Job Watcher and
  - End To End Performance Management on IBM i
- **Know about iDoctor links**
  - www.youtube.com/user/IBMiDoctor
- **Use the GUI interface, easier to drill and understand data/labels, contains some custom profiles for capturing the right data**
- **v7.1 PTFs more SQL metrics, instrumentation in Collection Services**
- **Experience makes the best teacher**

What to Look for - How to Analyze and Tune

- **Time Consuming SQL or queries -**
  - Where to start… SQE Plan Cache
    - Long running SQL
    - High usage (often extremely short running)
    - Those scoped to job only tables…QTEMP awareness
    - Temporary index use

Proactive vs Reactive

SQL Plan Cache properties – Greater insight & control

- **Autosizing enhancement**
  - The default hit ratio changes from 70%→90%
  - The maximum size allowed for autosizing, which was previously set to a hard coded value, is now calculated based on partition size.
- **Slowest runs information. Plan Cache Activity Thresholds**

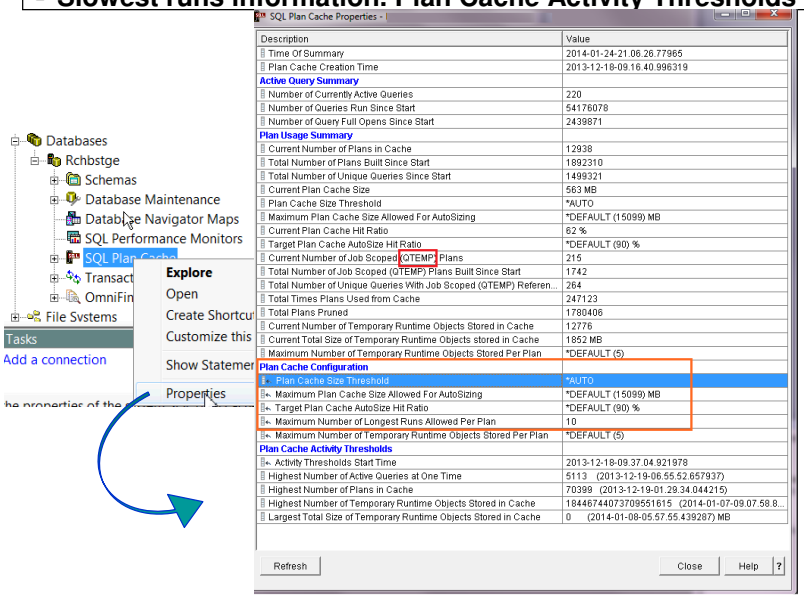

Customizable

- **Maximum Autosize**
- **Target Hit Ratio**
- **# of longest runs pre plan to keep in the cache**

Plan Cache Example – High use statements
Filter and ask for top N most time consuming statements

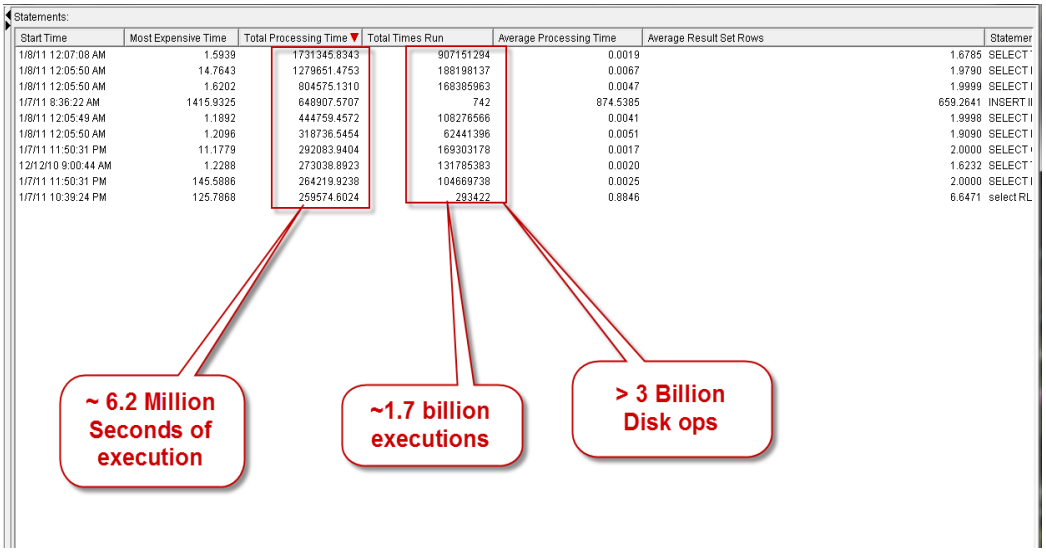

Tuning these can have a major impact on the system !!

Long running SQL – Examine the access plans

- **What to look for**
  - Full table scans
  - Full index scans
  - Temp index builds
  - High use statements needing better indexes
    - Index Only Access

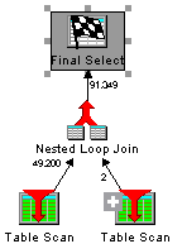

Look for expensive table scan

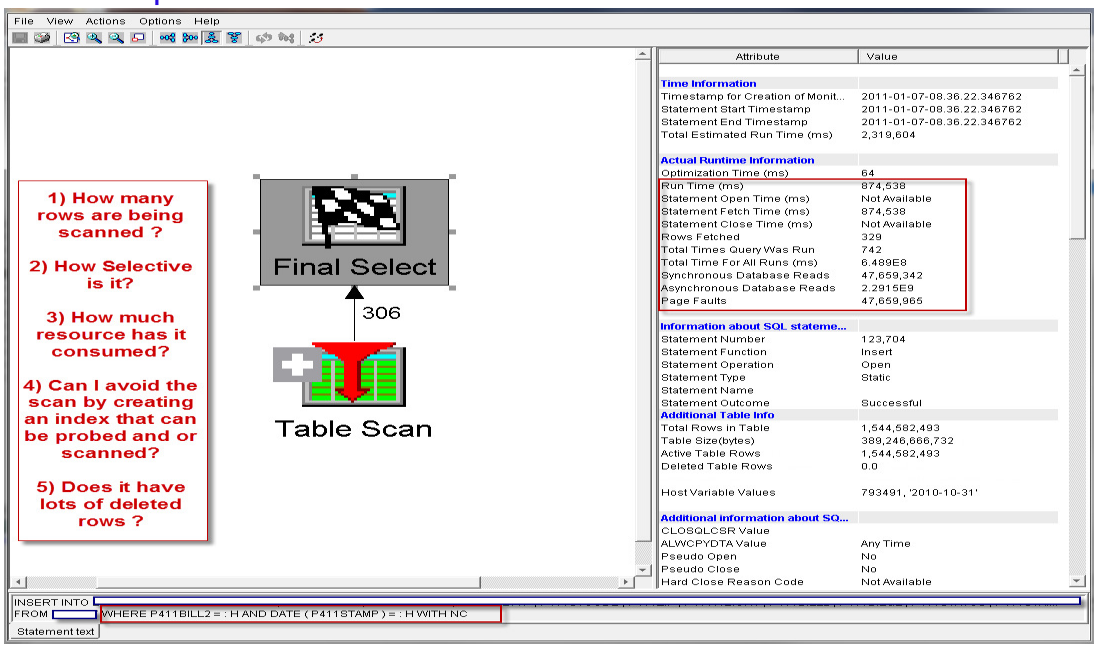

CREATE INDEX DTALIB.TABLENAME_IDX ON DTALIB.TABLENAME (P411BILL2 ASC, DATE(P411STAMP) as p411Date);

Replace a high use Index scan with a Index probe

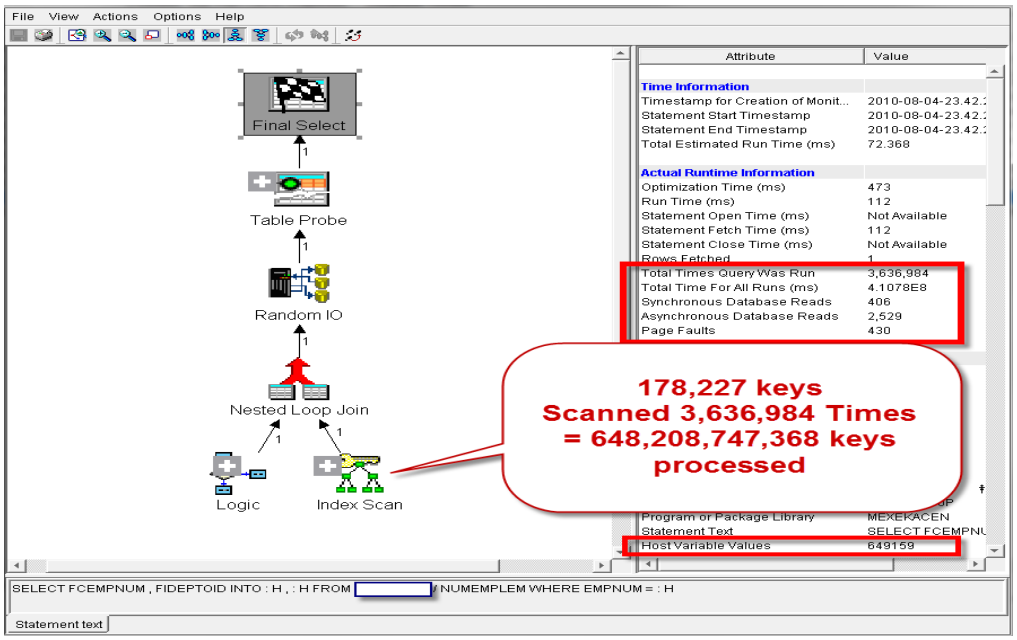

Pass matching Host variables, or alter the column to be an integer.

Replace high use temporary index with permanent **SQL INDEX**

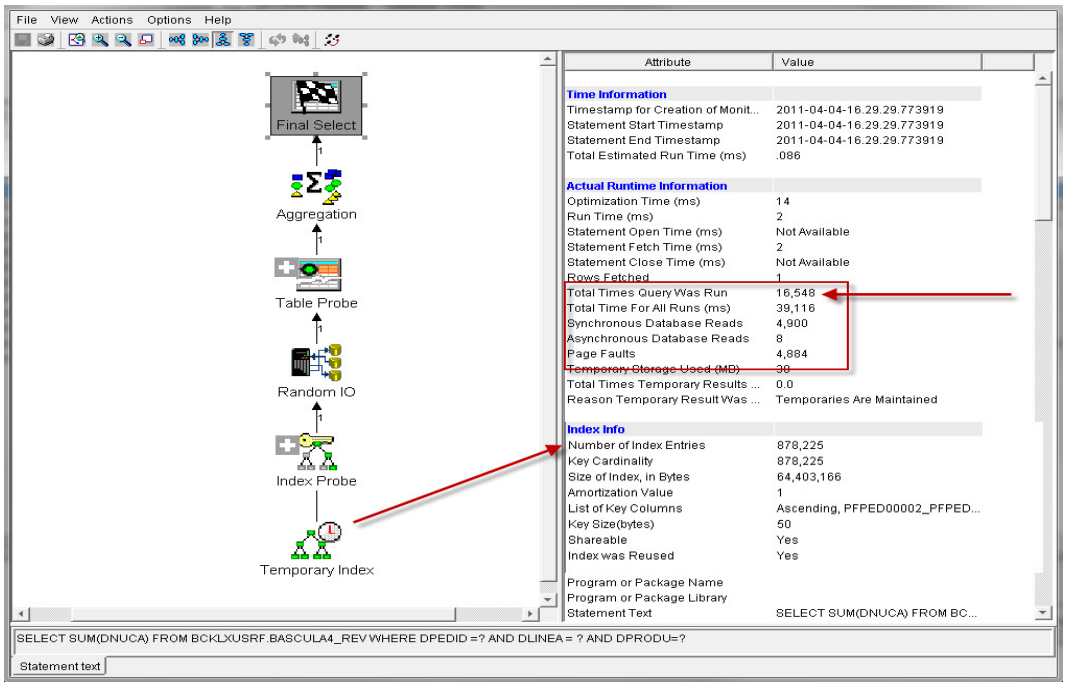

Create a better index – Avoids touching the table

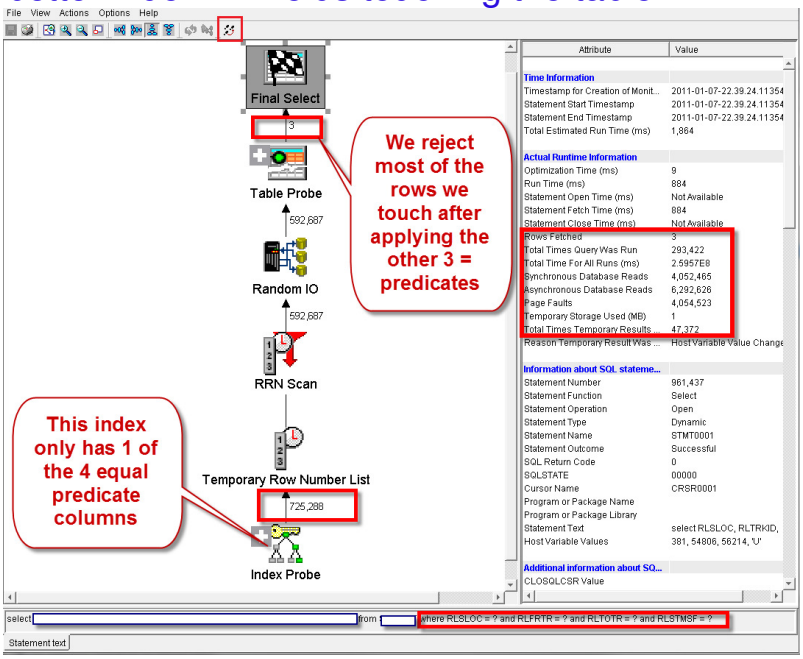

CREATE INDEX DTALIB.TABLENAM_IDX ON DTALIB.TABLENAME (RLSTMSF ASC, RLSLOC ASC, RLFRTR ASC, RLTOTR ASC);

Identify SQL running in CQE

- **CQE Only Monitor**

  STRDBMON OUTFILE(LIB/MonFile) OUTMBR(*FIRST *ADD) JOB(*ALL) TYPE(*DETAIL) COMMENT('**WANT_CQE_ONLY**')

- **Summarize CQE Reasons**
  - See appendix for query to run against the DB monitor
- **Move them to SQE**
  - 6.1 PTF to support SQL that references simple LFs
  - Migrate to i 7.1 – All SQL goes to SQE
  - Replace non- SQL queries with SQL

Profile the data – Look for BIG tables

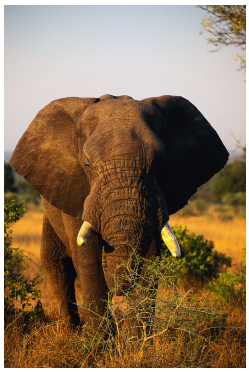

- **How big are your tables ?**
  - The bigger they are the more important it may be to pay attention
  - Use Qsys2 views to understand your biggest files. Or Files with largest numbers of deleted rows

```
Select Table_schema, Table_name, Table_partition, Number_rows,
Number_Deleted_ROWS, DATA_SIZE,
Variable_length_size,  Number_Distinct_Indexes
FROM QSYS2.SYSPARTITIONSTAT
order by DATA_SIZE DESC
FETCH FIRST 100 rows only;
```


Power is performance redefined

Look to see how the big files are used in SQE SQL Plan Cache

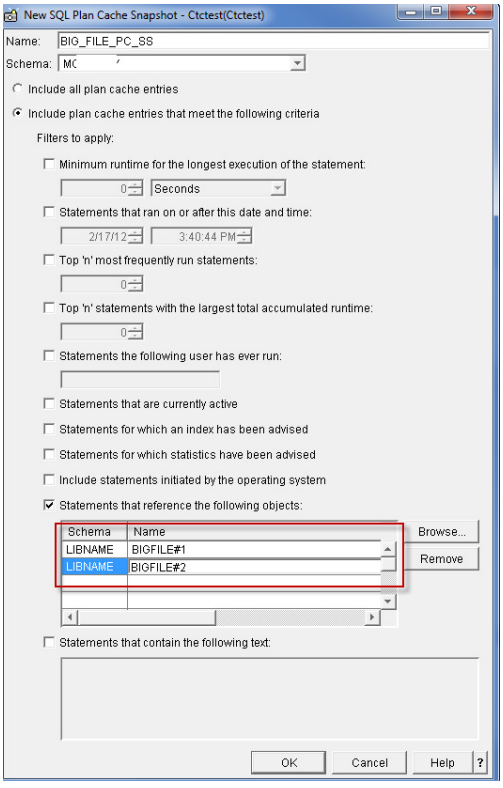

Look for

1. **Table scans**
2. **Index scans**
3. **Index improvements**
4. **Temp index builds**

Profile the data – What kind of indexes do I have

- **What kinds of indexes do I have?**
  - DDS versus SQL index
  - Page Size or max index size

Select INDEX_TYPE, SUM(CASE WHEN accpth_type =0 then 1 else 0 end) as ...

INDEX_TYPE	RADIXMAX1TB	RADIXMAX4GB	SQLEVI
FOREIGN KEY	2420	0	0
PHYSICAL	1366	522	0
PRIMARY KEY	2846	1	0
UNIQUE	125	0	0
LOGICAL	401	634	0
INDEX	8686	0	7

* See Appendix for Index Profile queries

Power is performance redefined

Profile the data – What kind of indexes do I have

Select INDEX_TYPE, LOGICAL_PAGE_SIZE, SUM(CASE WHEN accpth_type =0 then 1 else 0 end) as ... - Ctctest(C...

INDEX_TYPE	LOGICAL_PAGE_SIZE	RADIXMAX1TB	RADIXMAX4GB	SQLEVI
FOREIGN KEY	8192	5	0	0
FOREIGN KEY	65536	2415	0	0
INDEX	65536	8686	0	0
INDEX	-	0	0	7
LOGICAL	4096	0	634	0
LOGICAL	8192	397	0	0
LOGICAL	65536	4	0	0
PHYSICAL	2048	0	15	0
PHYSICAL	4096	0	507	0
PHYSICAL	8192	1359	0	0
PHYSICAL	16384	6	0	0
PHYSICAL	65536	1	0	0
PRIMARY KEY	2048	0	1	0
PRIMARY KEY	8192	5	0	0
PRIMARY KEY	65536	2841	0	0
UNIQUE	65536	125	0	0

* See Appendix for Index Profile queries

What tools should I use to look at the files

- **Use the following to get a file profile**

DSPFD FILE(*ALLUSR/*ALL) TYPE(*ATR) OUTPUT(*OUTFILE) FILEATR(*PF) OUTFILE(<yourlib>/DSPFD_ATR)

```
Select count(*) as Total_Files_Tables, Sum( case WHEN PHFTYP='R' then 1 ELSE 0 end) as
DDM_files, SUM(CASE WHEN PHFATR ='PF38' then 1 else 0 end) as Sys_38_PF, -- Reuse
Setting SUM(CASE WHEN PHRUSE ='Y' then 1 else 0 end) as Reuse_Delt_Rows,
SUM(CASE WHEN PHRUSE ='N' then 1 else 0 end) as NOT_Reuse_Delt, -- Partitioned tables
sum(case when PHPRBY <>0 then 1 else 0 end) as Partitioned_tab, -- SQL Tables,  If not a T
or M, its a PF SUM(CASE WHEN PHSQLT ='T' then 1 else 0 end) as SQL_Tab,  -- MQTs
SUM(CASE WHEN PHSQLT ='M' then 1 else 0 end) as MQT, SUM(CASE when PHFRCR<>0
then 1 else 0 end) as With_Frc_WRT_RATIO, -- Program Described sum(case when
PHFLS='N' then 1 else 0 end) as NOT_Ext_Described

from <yourlib>.dspfd_ATR ;
```

TOTAL_FILES_TAB...	DDM_FILES	SYS_38_PF	REUSE_DELT_ROWS	NOT_REUSE_DELT	PARTITIONED...	SQL_TAB	MQT	WITH_FRC_WRT_RATIO	NOT_EXT_DESCRIBED
21977	0	0	7989	13436	4	7525	3	9	755

Look for

- **High number of deleted rows**
- **Max 4GB Indexes**
- **High use indexes with small page (8K) size**
- **Indexes that are not used**
  - How do I find these?
    - iNav **show indexes** OR
    - QUSRMBRD API,
      - **Data space index last query statistics use.** The date and time the last time this data space index was used for statistics when optimizing a query. This is in the YYYYMMDDHHMMSS format.
      - **Data space index last query use.** The date and time the last time this data space index was used in a query full open. This is in the YYYYMMDDHHMMSS format.

Interview the application developers

- **What data are you using?**
- **How do you access it?**
  - One row at a time?
  - Set based processing
  - Do you make copies of the data?
- **How much data do you process?**
  - What are the most expensive components
- **Have you done performance testing on production-like data?**
  - Size and content needs to be comparable
  - Environment is key as well
  - What analysis did you do on that test?

Power is performance redefined

10 DB2 for i Performance Tips

© 2012 IBM Corporation

QUSER 4/2012

What should I do about it – Tip #1 - **Improve indexing Strategy**

- **Minimize accessing data that you don't need to**
  - More Selection columns in the index
  - Utilize Index Only Access when possible
- **Utilize new indexing technology in 6.1 and 7.1**
  - **Derived indexes**
  - EVI Aggregates
  - Memory resident indexes
  - Indexes on SSDs (tables too)
- **Get rid of MAX 4G indexes**
- **High use indexes should have bigger page size**

Read the newly updated Indexing white paper

https://www-304.ibm.com/partnerworld/wps/servlet/ContentHandler/stg_ast_sys_wp_db2_i_indexing_methods_strategies

What should I do about it – Tip #2 - **Identify critical, inefficient jobs, applications and programs. Make those more efficient!!**

- **Move from Record at a time to Set based SQL processing**
  - Minimize the data you access in your programs
  - Rewrite to use SQL set at a time (don't do SQL record at a time)
    - Focus on the expensive, not everything!!
  - Potentially use ***RPG Open Access*** to intercept native I/O and do SQL set processing in the handler.
    - http://www.ibm.com/developerworks/ibmi/library/i-roaforsql/index.html?ca=drs-

What should I do about it – Tip #2 CONT **Native Record Level Access to SQL Set based processing**

- **Take more of a Data Centric Approach versus applications centric**
  - get the database management system to do more on your behalf
- **Drive as much work down into the DBMS as possible**

  With traditional record level access:

  You tell DB2 what you want, and how to do it.

  With SQL:

  You tell DB2 what you want, but not how to do it.
- **Key attributes of Data centric Approach and set based SQL**
  - More sophisticated access
  - More Concurrent I/O
  - More Easily multithreaded
  - Centralized data integrity

What should I do about it – Tip #3 **Ensure that Jobs Memory share is not too Small**

- Optimizer calculates Jobs fair share and considers Max Active job setting in the pool. Smallest active number of jobs that will be used to calculate the Share is 10% of max
- Avoid setting Max active higher than it needs to be

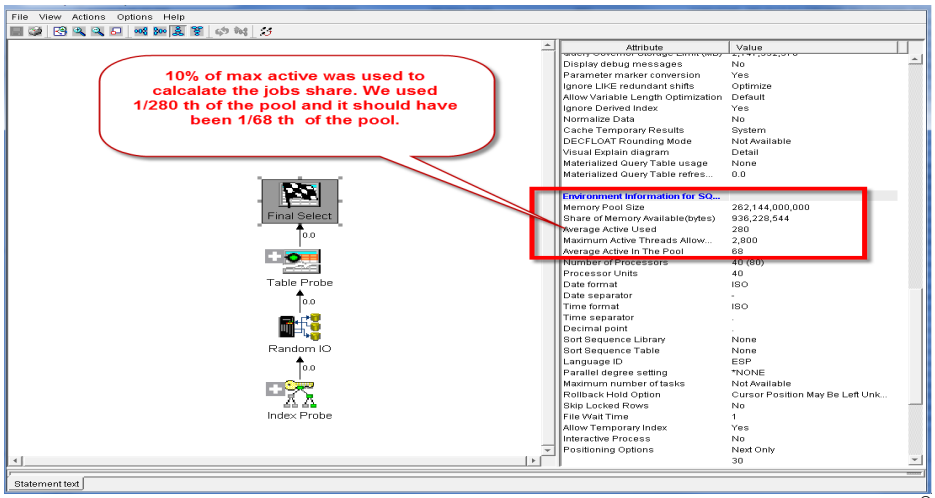

What should I do about it – Tip #4 **Set Optimization Goal to match behavior**

- Set via optional SQL statement clause
  - –OPTIMIZE FOR n ROWS
  - –OPTIMIZE FOR ALL ROWS
- Set via QAQQINI options file -or- ODBC/JDBC connection attributes
  - –*FIRSTIO
  - –*ALLIO
  - –(chgqrya qryoptlib(<ini lib>))
- Default for dynamic interfaces is **First I/O**
  - –ODBC, JDBC, STRSQL, run SQL scripts (iNav), dynamic SQL in programs
  - –**CQE - 3% of expected result set**
  - –**SQE - 30 rows**
- Otherwise default is **ALL I/O**
  - –Extended dynamic, RUNSQLSTM, INSERT + subSELECT, CLI, static SQL in programs
  - –**All of expected excepted result set**
- Optimization goal will affect the optimizer's decisions
  - –Use of indexes, SMP, temporary intermediate results like hash tables
  - –**Tell the optimizer as much information as possible**
  - –**If the application fetches the entire result set, use *ALLIO**

Power is performance redefined

What should I do about it – Tip #5 **Reorganize high use files with lots of deleted rows**

- **We can reduce disk space, main store foot print, CPU and disk ops by compressing out deleted rows**
  - – Rule of thumb says > 10% should be reorganized ??
- **Look for the really bad ones first**

-- Top 25 tables that Have > 10,000,000 deleted rows

```
Select Table_schema, Table_name, Table_partition, Number_rows, Number_Deleted_ROWS,
DATA_SIZE, Variable_length_size,  Number_Distinct_Indexes

FROM QSYS2.SYSPARTITIONSTAT WHERE NUMBER_DELETED_ROWS > 10000000
order by Number_Deleted_ROWS DESC FETCH FIRST 25 rows only;
```

- Use Reorganize to compress out deleted rows

```
RGZPFM FILE(FLGHT400/ORDERS) MBR(ORDERS) KEYFILE(*RPLDLTRCD)
ALWCANCEL(*YES) LOCK(*SHRUPD) ; – This example allows concurrent updates and replaces
deleted records at the start of the file with valid records from the end of the file.
```

What should I do about it – Tip #6 **Control adhoc queries**

- **Control significant resource usage**
  - Ad hoc queries
    - Limit access to certain interfaces
    - Use Query Governor (see the QAQQINI file)
  - Look for users that are doing expensive queries
    - Should they be doing that off peak?
    - Should they be doing it at all??

What should I do about it – Tip #7 **Don't ask for more data that is needed**

- **Limit result set to something reasonable**
  - Ask yourself, will a end user really scroll through millions of rows in a report?
  - Use FETCH N rows
  - Search for SQL Pagination
    - http://www.itjungle.com/fhg/fhg111109-story01.html

What should I do about it – Tip #8 **Avoid making copies of data**

- **Replace creation of work files or temp table with a virtual table**
  - Optimizer is often given no implementation tools with these temp tables AND the plans cannot be shared.
  - Create a view that defines that data set and access the view instead of the work file (if reuse of view is likely in other queries and no hostvar requirements)
  - Otherwise, Use **Common Table Expression** or derived table
    - See CTE examples in the appendix

What should I do about it – Tip #9 **Use separate environment for Query and reporting**

- **Mixing core applications with analytical type queries/reports can cause problems**
  - Query and reporting may be resource intensive
  - Needs different set of indexes than OLTP apps
  - May need different DB design (OLAP vs OLTP)
- **Move the query and reporting to separate system or LPAR**
  - Maybe the HA target
  - Separate box with an Operational Data Store
    - Appliance like i for BI Optimized Solution
      - http://www-03.ibm.com/systems/resources/systems_i_software_db2webquery_whyibmiforbi.pdf

What should I do about it – Tip # 10 Reduce Full opens

- **First determine if it's a problem**
  - Analyze plan cache snapshot or DB monitor
  - Use System Tools
- **Common Causes**
  - Frequently getting new connection from remote app
  - In RPG setting LR on exit
  - Programs compiled incorrectly

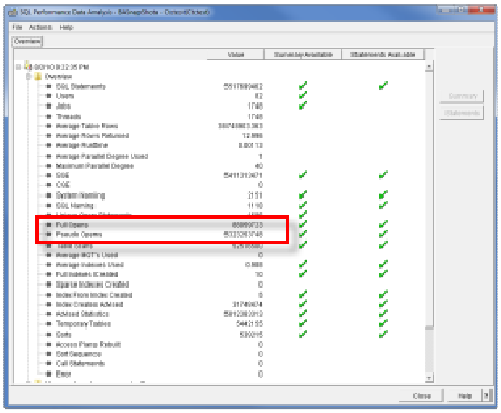

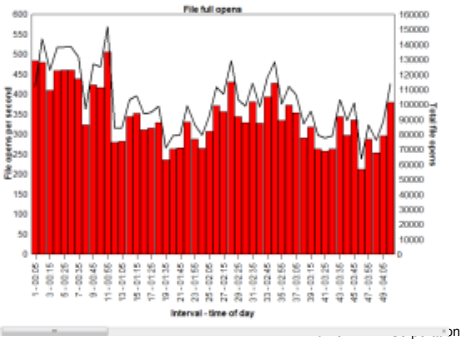

See Appendix for more info

Miscellaneous Tips

- **For Varchar columns set the allocated length so it covers most occurrences**
- **Utilize the DB2 for i Grouping sets support**
  - Allows grouping at many levels/dimensions within one query
- **Avoid use of Force write ratio**
- **Rebuild EVIs with high over flow count**
- **Utilize Tuned Views – DB engineer creates, tests and tunes before making them available**
- **Improve DB design (3rd normal form recommended)**
- **Use Referential integrity – Avoid data integrity checking**
- **Ensure you have a balanced configuration**
  - Avoid bottleneck on CPU, Memory Disk (do analysis and tuning before adding HDW)
- **You can't cheat the SQE Stats collector – leave it on!**
- **Use Performance Explorer to identify** expensive programs **and focus attention on those – histogram of the expensive calls**
- **Utilize SMP to take advantage of CPU power**
  - Do proper tuning first
  - Careful about SMP scope (do NOT set it system wide unless its a dedicated query box)
  - Utilize % capability in the QAQQINI file (*OPTIMIZE NN – Set degree to NN % of what the optimizer calculates) to throttle back SMP

Miscellaneous Tips –continued

- **Be aware of need (or NOT) for Live Data Mode or Allow Copy Data *NO**
  - This restricts optimizer choices
- **There are three possible options for cursor sensitivity. The default on a create cursor is ASENSITIVE (the most flexible) but be aware while debugging or developing an application that different environments may have this set to other values.**
  - ASENSITIVE - Specifies that the cursor may behave as SENSITIVE or INSENSITIVE depending on how the select-statement is optimized (optimizers choice!)
  - INSENSITIVE -Specifies that once the cursor is opened, it does not have sensitivity to inserts, updates, or deletes performed by this or any other activation group (tends towards plan with more cached results in the query)
  - SENSITIVE – Specifies that changes made to the database after the cursor is opened are visible in the result table (results in a 'live data' mode access plan, focuses on indexes and direct table access methods)
- **Beware of the DEFAULTS for Creating UDFs (and UDTFs)!**
  Although defaults on UDF are the 'safe' **choice, they are NOT the most performance friendly!**
  - NOT DETERMINISTIC - function might not return the same result each time that the function is invoked with the same input arguments (prevents caching)
  - EXTERNAL ACTION - function can take an action that changes the state of an object that the database manager does not manage (prevents caching)
  - FENCED – Function runs in a separate thread (additional overhead for thread)
  - READS SQL DATA – function can execute SQL but not statements that modify SQL
  - **IF** the nature of the UDF permits, the following options can result in a more optimal performing access plan: DETERMINISTIC, NO EXTERNAL ACTION, UNFENCED, CONTAINS SQL
  - Ensure that UDFs are NOT FENCED unless they need to run on a different thread

Miscellaneous tips -continued

- **Inline UDFs (available in PTF)**
  - In cases of very simple SQL functions, instead of invoking the function as part of a query, the expression in the RETURN statement of the function will be in-lined to the query. To take advantage an existing UDF needs to be altered or recreated with enablement PTFs.
    - A function is an inline function IF:
      - It is deterministic
      - SQL-routine-body contains only a RETURN statement.
      - RETURN statement does not contain a scalar subselect or fullselect.
      - SQL-routine-body does not contain BEGIN ATOMIC.
      - Example Tip. If you do an if/then/else with a RETURN on each leg, it cannot be inlined…but …if you return a CASE statement .. it can …

Additional Education

- **IBM i Database Performance and Query Optimization Guide (infocenter pub)**
- **SQL Performance tuning workshop**
  - http://www-03.ibm.com/systems/i/software/db2/db2performance.html
- **DB2 for IBM i Website**
  - http://www-03.ibm.com/systems/i/software/db2/
- **DB2 Performance tuning Redbook (old)**
  - http://www.redbooks.ibm.com/abstracts/sg247326.html?Open
- **iDoctor**
  - https://www-912.ibm.com/I_dir/idoctor.nsf
- **Navigator for i (browser)**
  - http://www-03.ibm.com/systems/i/software/navigator/directornavigator.html

Power is performance redefined

Additional Education - IBM i Navigator database tasks

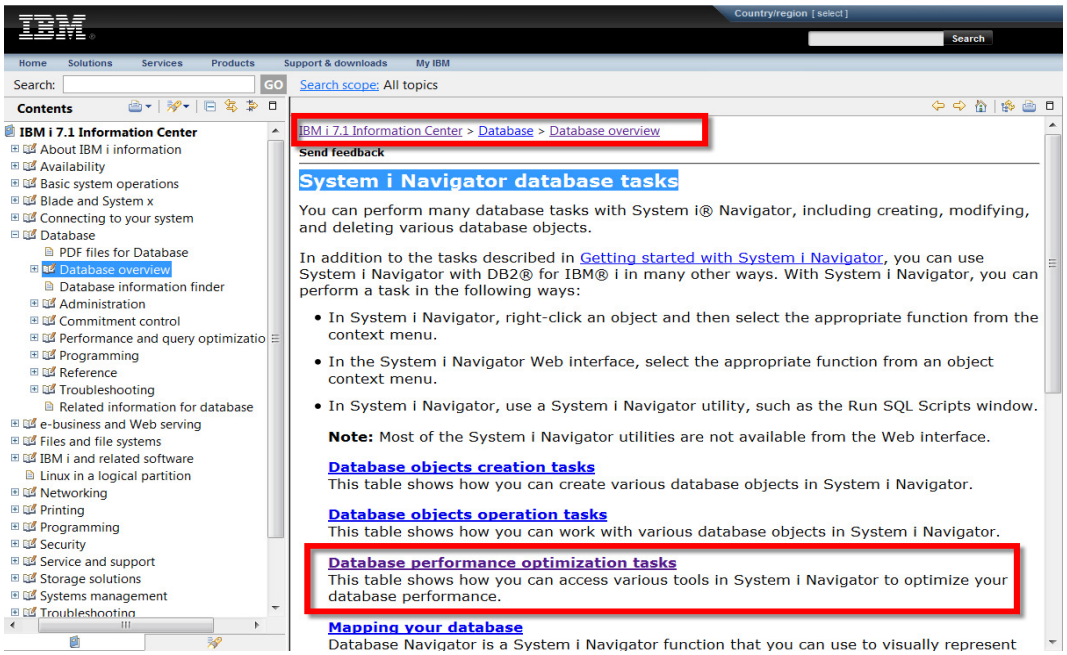

Summary

- **The more data being processed the more important efficiency becomes.**
  - Not just lots of data per request. But large number of small requests as well.
- **Learn to use the tools**
  - Collection Services, Job Watcher and iNavigator
  - Experience is the best teacher
- **Focus tuning effort on what is expensive**
  - low hanging fruit
- **Think Data Centric, not Application Centric**
  - Learn to utilize the Power of SQL - set at a time processing!

OnDemand Performance Center – User Authority Simplification

***JOBCTL (Job Control Authority)**

- Whatever worked with *JOBCTL in IBM i 6.1 will continue to work

QIBM_DB_SQLADM – Database Administrator

- This is a database specific alternative to *JOBCTL. It is a superset of the function authorized by QIBM_DB_SYSMON.
- Examples:
  - Change parallel degree for DB2 SMP feature
  - Work with Plan Cache
  - Work with OmniFind Text Search Server

QIBM_DB_SYSMON – Database Information

- This allows a user to view some system level details, but not specifics about operations or anything related to changing or controlling the database.
- Examples:
  - QUSRJOBI for SQL information
  - Show SQL Information for Jobs

User Authorization Commands:

CHGFCNUSG FCNID(QIBM_DB_SQLADM) USER(userid) USAGE(*ALLOWED)

CHGFCNUSG FCNID(QIBM_DB_SYSMON) USER(userid) USAGE(*ALLOWED)

No Special Authority required when using OnDemand Performance Center with own job

- Starting/ending SQL Performance Monitors on your own job
- Analysis of SQL Monitor/Plan cache data
- Visual Explain in Run SQL Scripts

Questions ?

Appendix

Index Profile queries

```
Select  INDEX_TYPE, SUM(CASE WHEN accpth_type =0 then 1 else 0 end) as
  RADIXMAX1TB,
        SUM(CASE WHEN accpth_type = 1 then 1 else 0 end) as RADIXMAX4GB,
  SUM(CASE WHEN ACCPTH_TYPE=2 then 1 else 0 end) as SQLEVI

  FROM qsys2.syspartitionindexstat
  where table_NAME not like 'Q%' and table_SCHEMA NOT LIKE 'Q%'
  group by index_TYPE;

Select  INDEX_TYPE, LOGICAL_PAGE_SIZE, SUM(CASE WHEN accpth_type =0
  then 1 else 0 end) as RADIXMAX1TB,
        SUM(CASE WHEN accpth_type = 1 then 1 else 0 end) as RADIXMAX4GB,
  SUM(CASE WHEN ACCPTH_TYPE=2 then 1 else 0 end) as SQLEVI

  FROM qsys2.syspartitionindexstat
  where table_NAME not like 'Q%' and table_SCHEMA NOT LIKE 'Q%'
  group by index_TYPE, LOGICAL_PAGE_SIZE
  order by index_type, logical_page_size;
```


CTE EXAMPLE – Logical Step-by-Step

- **Break query into logical steps when query requires multiple SQL statements**
  - Improved readability
  - Removes SQL view management, host variables also supported
  - Can be used to avoid **temp** files

```
WITH staff (deptno, empcount) AS
(SELECT deptno, COUNT(*) FROM employee
 WHERE division = :div_var GROUP BY deptno)
SELECT deptno, empcount FROM staff
WHERE empcount =
        (SELECT MAX(empcount) FROM staff)
```

CTE EXAMPLE – Logical Step-by-Step & Work Tables

- **CTE logical step processing can eliminate the need for temporary work tables**

```
DECLARE GLOBAL TEMPORARY TABLE t1 AS
(SELECT shipdate, customer, phone, orderkey, linenumber
  FROM item_fact i, cust_dim c
  WHERE c.custkey=i.custkey AND discount=0.08) WITH DATA;

DECLARE GLOBAL TEMPORARY TABLE t1 AS
(SELECT customer, phone, orderkey, linenumber, year, quarter
  FROM t1, star1g.time_dim t
  WHERE t.datekey=shipdate ) WITH DATA;

...
```


CTE EXAMPLE – Logical Step-by-Step & Work Tables (Continued)

- **CTE provides the same step-by-step approach without the overhead of populating physical tables**

```
WITH t1 AS
 (SELECT shipdate, customer, phone, orderkey, linenumber
  FROM item_fact i, cust_dim c
  WHERE c.custkey = i.custkey AND discount=0.08),
    t2 AS
  (SELECT customer, phone, orderkey, linenumber, year, quarter
   FROM t1, star1g.time_dim t
   WHERE t.datekey = shipdate)
SELECT * FROM t2;
```

Finding EVIs with large over flow

```
Select Index_schema, Index_name, Column_names, Number_keys,
   Last_query_use, Query_use_count, Query_Statistics_count, index_size,
   Overflow_values, EVI_CODE_SIZE
   FROM QSYS2.syspartitionindexstat
   where ACCPTH_TYPE=2 and overflow_values > 0 and table_NAME not like
   'Q%' and table_SCHEMA NOT LIKE 'Q%'
   order by Overflow_values desc
```

Refresh the symbol table using CHGLF

CHGLF FILE(LIBNAME/INDEXNAME) FRCRBDAP(*YES)

More Info on Tip #1 – Find and replace max 4GB indexes

Replace MAX 4G indexes with MAX 1TB

```
Select *
from qsys2.syspartitionindexstat
where (table_schema not like 'Q%' or table_schema = 'QUSRSYS') and
ACCPTH_TYPE=1
Order by INDEX_SIZE DESC;
```

- **For LFs**
  - Option 1 Change these to MAX 1TB **USE THIS OPTION**
    CHGLF FILE(DTALIB/PF1) ACCPTHSIZ(*MAX1TB)
  - Option 2 – Recreate Simple Keyed LFs using SQL CREATE INDEX and new 6.1 syntax. You can access these as LFs from Native I/O Commands

- **For PFs**
  CHGPF FILE(DTALIB/PF1) **ACCPTHSIZ(*MAX1TB)**

Tip #2 – Monitor, Analyze and Tune strategy

- **Monitor**
  - Keep archive of historical performance data
    - For a typical Day or week
      - Collection services
      - Plan cache snapshot
      - Health center file information
      - DB monitor sample
    - For Peak periods (end of month, peak season…)
    - Batch window as well (maybe a separate collection of data)
  - If problem occurs capture data and save to allow analysis of what just happened
  - Automate collection process (i.e. schedule job to monitor or collect data)
  - Here is article showing how to automate a Plan cache snapshot collection
    - http://www.mcpressonline.com/tips-techniques/system-administration/techtip-automatically-generate-an-sqe-plan-cache-snapshot.html

Tip #2 – Monitor, Analyze and Tune strategy…..

- **Analyze**
  - Periodically analyze top N most time consuming SQL from Plan cache (Save those to a new snapshot and compare to old one to see if anything is slower)
    - Look for Table scans, Index scans, Temp index builds, queries that access more data than they need to. Apply the tips (i.e. copies of data, large result sets…)
    - Are there statements that can be written more efficiently
  - Look for expensive executions (sort descending by longest run time)
    - Should these even be run?
    - Look for Users that are doing expensive queries and investigate
  - Assess the system
    - Are any resources a bottleneck. Don't just look at system CPU %
  - Look for Indexes that are advised a large number of times
    - Find the statement(s) that are causing that advice and assess those

Tip #2 – Monitor, Analyze and Tune strategy…..

- **Tune**
  - Proactive analysis for new or changed applications
    - Part of testing needs to include data collection and analysis
    - **A good indexing strategy is KEY**!!
  - Apply the tuning Tips given earlier
  - Use Visual Explain on the statements that are generating index advise for those IXs that have a high advise count.

Reducing Full Opens

- **Use DBMON analyze program summary to see program with high SQL full Opens**
  - CLOSQLCSR setting needs to be explored for those programs that show all full opens
- **Use iDoctor Pex full open trace to identify programs with high native full opens**
- **In addition the Information given, the SQL Query and performance optimization guide should be used. The following link can be used to access that manual:**
  - http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/topic/rzajq/rzajq.pdf

Reducing Full Opens

- **Reducing Full opens in native applications**
  - we typically see high rate of full opens and closes is that programs issue SETON LR and then Return or end the program. This has a consequence of closing the open files. In some cases the SETON LR (or *INLR = *ON) can be removed avoiding the close and subsequent open on the next call to the program. You would want to use RETURN (without SETON LR) if you are calling this program multiple times from the same calling program, in order to reduce overhead of opening and closing the files each time. For more info on use of SETON LR See the following link:
    - IBM i 7.1 Information Center > Programming > Programming languages > RPG

Power is performance redefined

Reducing Full Opens

- **Reducing Full Opens in JDBC Access**
  - The key to reducing full opens in the JDBC access is to use **connection pooling**. Based on the data we see, you are not using connection pooling successfully. You need to avoid getting a new connection for each external request. Java has classes that will support the creation and reuse of connections. IBM Info Center has information on JDBC Object pooling, which also includes ConnectionPoolDataSource properties. Go to IBM InfoCenter to see additional info
    - IBM i 7.1 Information Center > Programming > Java > IBM Developer Kit for Java > Database access from Java programs > Accessing your IBM i database with the Java JDBC driver
  - You may also build your own connection pooling. See that topic at the following info center location:
    - IBM i 7.1 Information Center > Programming > Java > IBM Developer Kit for Java > Database access from Java programs > Accessing your IBM i database with the Java JDBC driver > JDBC object pooling

OnDemand Performance Center & Sensitive Data – SECURE columns

- **Prevents sensitive data values from being displayed in DB2 performance tools – Database Monitor & Plan Cache (5.4 & 6.1 PTFs)**
  - Only security officer will be able to see sensitive values, '*SECURE' value presented to *normal* users (… WHERE cardnumber=:hostvar1)
  - User must register sensitive columns with DB2 tooling

- **Registration interface is system stored procedure: SET_COLUMN_ATTRIBUTE**
  - Procedure parameter descriptions
    - Table_Schema - System name of a table's schema
    - Table_Name - System name of a table
    - Column_Name - System column name being secured.
    - Attribute - Secure attribute setting for column
      - SECURE NO
      - SECURE YES
  - Example:

```
CALL SYSPROC.SET_COLUMN_ATTRIBUTE
  ('MYLIB1', 'ORDERS', 'CCNBR', 'SECURE YES');
```


Power is performance redefined

Trademarks and Disclaimers