Université IBM i 2018

16 et 17 mai

IBM Client Center Paris

S48 – Best Practices for IBM i Memory Tuning for Performance

Stacy L. Benfield
IBM i Performance Consultant - Lab Services Power Systems Delivery Practice
IBM i Large User Group (LUG) Program Manager
IBM, Rochester MN
stacylb@us.ibm.com

Agenda

- Memory Concepts and page faulting
- IBM i memory tuning features
- Tools used to identify memory issues
- Best practices for memory tuning

Importance of Memory

- The CPU(s) can't do anything if the required data is not in memory
- Having to go to disk is slow, in relative terms

How to improve performance:

- Tuning can help your system use its memory better
- IBM i has technology built in to help manage memory automatically
- IBM i has excellent performance tools for analyzing memory performance
- Adding hardware such as more memory or SSDs can reduce I/O wait times
- Application changes can greatly reduce how often your jobs wait on disk I/O

Bringing data into memory

- Implicit memory transfers – ***"page fault reads"***
  - Page faults – Synchronous memory transfer where jobs wait until I/O completes
  - I/O Pending faults – waiting on your/or someone else's I/O request to complete
  - Collection Services & Job Watcher have a "wait bucket" to track this time

- Explicit memory transfers– ***"non-fault reads"***
  - Asynchronous memory transfer where OS pre-fetches data before application needs it (jobs don't have to wait)
    - DB2 "read-ahead" or "asynchronous brings"
  - SETOBJACC command puts an object into a memory pool
  - Read part of CRTDUPOBJ and CPYF operations
  - Collection Services & Job Watcher have a "wait bucket" to track this time

Types of I/O Requests for Reads and Writes

- Synchronous I/Os
  - Processing waits until I/O completes
  - Contributes to response/runtime
  - Examples (reads):

    SETOBJACC, CRTDUPOBJ, CPYF, Faults, …
- Asynchronous I/Os
  - Processing concurrent with I/O
  - Can turn synchronous
  - Examples (reads):

    DB2 asynchronous brings

Page Fault basics

The primary focus for memory analysis is to reduce the amount of time waiting on disk faults.

- Page Faults are normal and expected
  - It is the mechanism used to bring most things into memory
  - Some things need to be brought into memory to be cleaned up / deleted
- A single page fault can bring multiple pages into memory
  - Faults per second < pages per second
  - We don’t typically care about pages per second
- Requested pages are placed into the memory pool the job is running in
  - Pages can be shared between jobs
  - A job can access a page in a different pool

Why do we care if page fault wait times are high

- High page fault wait times indicates pool tuning issues, unnecessary memory usage, and/or lack of memory that will affect performance
  - Longer interactive response times for I/O intensive transactions
  - Longer batch run times
  - Poor disk response times
  - Less efficient query implementation methods

Page Faults – Faults Per Second vs. Wait Time - PDI

- Focus on time being spent waiting on faults, **not** the rate of faults
  - Large memory pools can have a high number of faults, but little impact to wait times
- However, rates can be useful for monitoring and when you can't get wait time easily
  - Dashboard, System Monitors, WRKSYSSTS, etc.
  - Also can indicate new workload has started up
- Which chart below is more helpful in determining if faulting might be a problem?

faults/sec

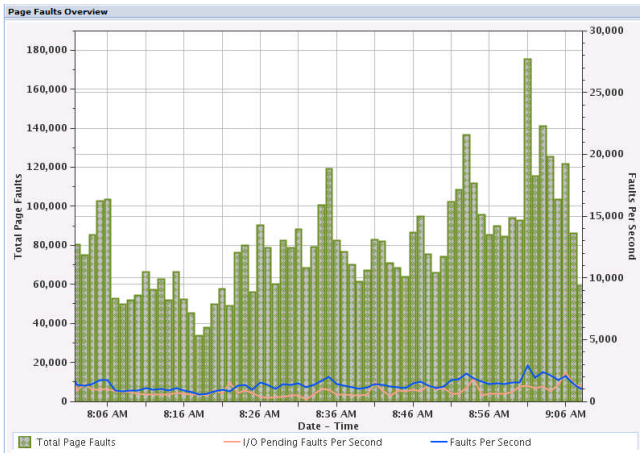

wait time

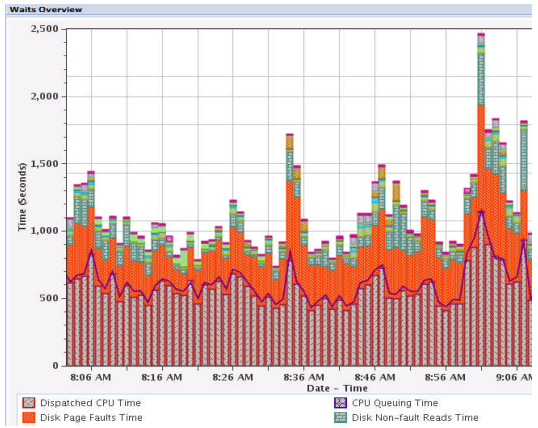

Page Faults Wait Time at a Job Level

- Collection Services (and Job Watcher) collect disk page fault wait time both at a **system** level, as well as an **individual job/thread/task level**
- ??? Is it more interesting to know that JOBX did 120 faults per second, or that it waited 370 seconds (out of 5 minutes) on disk page faults

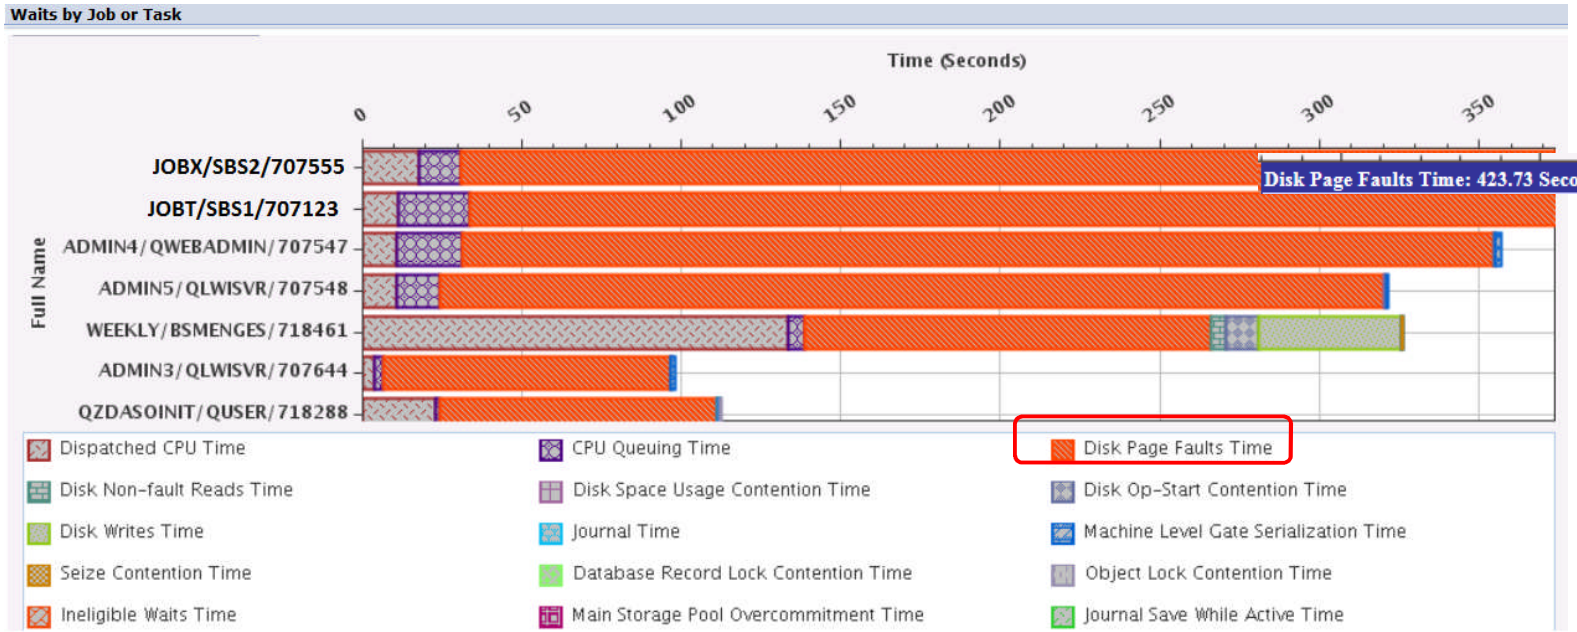

IBM tasks associated with removing data from memory

System Controlled

- There are two tasks that are responsible for writing changed pages of memory out to disk to make room for data being paged into memory.
  - SMPOL001: low priority page out task
  - SMPO0001: high priority page out task
    - becomes active if the low priority task cannot keep up with demand
- High activity in these tasks, especially the high priority task, can indicate a need for additional memory.

User Program Controlled

- Changed pages are written out when a job ends

Single-level Storage

- Unique to IBM i
- Main memory, traditional spinning disks, and solid-state drives all appear as **one** address space
- Objects automatically managed by the OS, moving between memory and disk
- One copy of file/data can be shared by many users, jobs, programs
- Objects spread across drives for parallelism
- No need to create “table spaces”, “buffer pools”, etc. required by other platforms
- **Think of memory as a giant cache for disk**

IBM i Memory / Storage Hierarchy

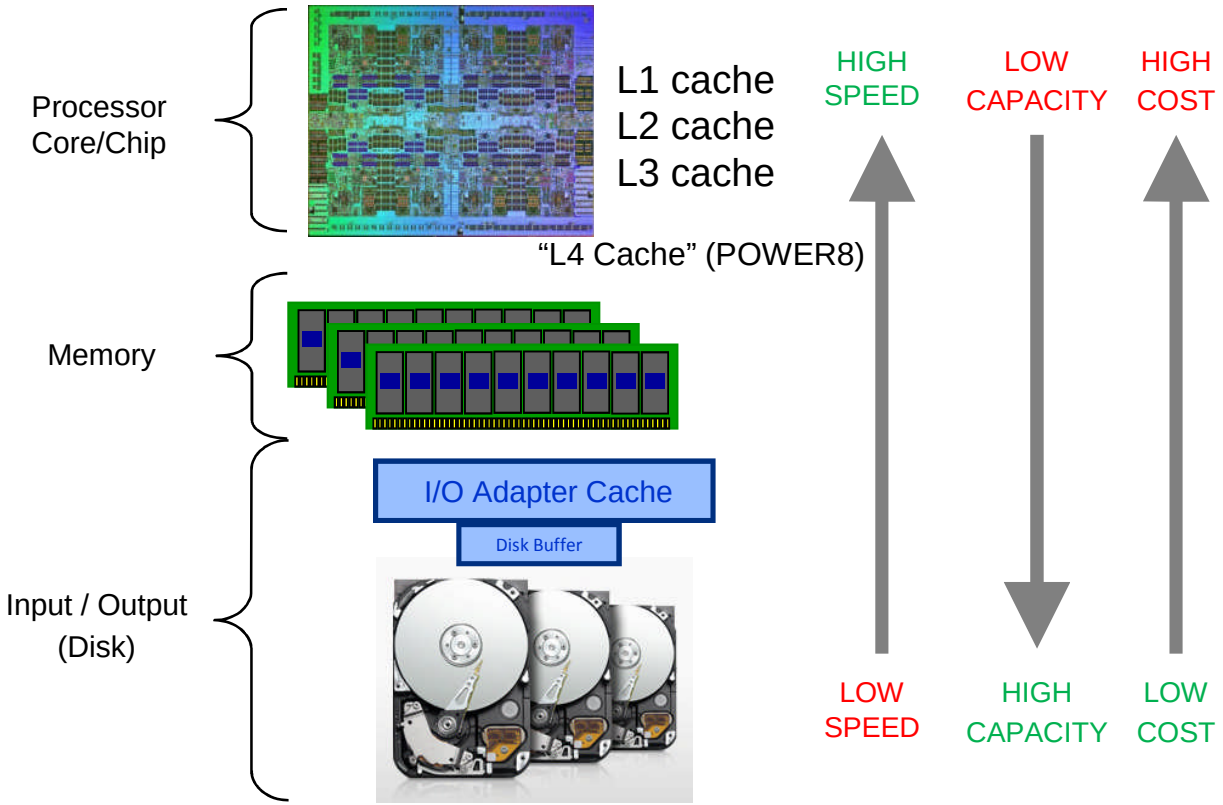

POWER Memory / Storage Hierarchy

	Size	Speed / cycles
POWER8 L1 cache	32K instr + 64K data per core	~3 cycles
POWER8 L2 cache	512 KB per core	~10 cycles
POWER8 L3 cache	96 MB shared per chip	~30 cycles
POWER8 “L4 cache”	16 MB per memory chip*	~200 cycles
Internal Memory	10s of GB per core	~300-800 cycles (~100ns)
Solid State Drives	100s of GB per drive	<1 ms
Hard Disk Drives	Can be TBs per drive	1-5 ms

100x

10,000x

Range due to “distance” of access.

1 ms = 1000 µs = 1,000,000 ns
400 cycles is about 0.1 µs if 4.0 GHz

Memory Configuration – Verify there's enough memory first

- A general rule of thumb for memory (minimum requirements) based off internal benchmarks
  - 32 GB/core for POWER8
  - 24 GB/core for POWER7/7+
  - 16 GB/core for POWER6
- Partitioning considerations
  - If dynamically adding cores, also typically want to add memory (DLPAR)
  - Uncapping can lead to an imbalance between CPU and memory
- Can utilize SSDs/Flash technology to improve page fault wait times

Memory Pools

- Memory pools are logical subdivisions of physical memory
- Used with subsystems to isolate memory usage by different applications
- Two types – shared and private

	Shared	Private
Subsystems	single or multiple	single
QPFRADJ	yes	no
Expert Cache	yes	no

Memory Pools

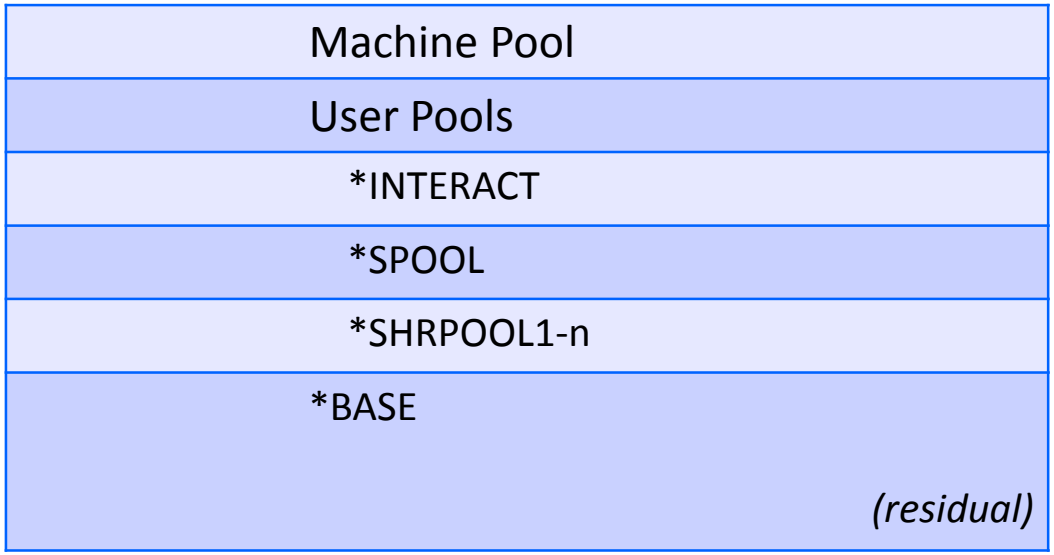

Pool Maximum Activity Level

- The maximum number of threads in the pool that can use the CPU(s) concurrently
  - Threads without an activity level are ineligible to run
- Does not apply to the machine pool
  - No jobs run in the machine pool
- Can be adjusted + or - by the Performance Adjustor
  - Adjuster conservative on decreasing
- Considerations
  - Generally want high enough to avoid transitions to ineligible
    - There is an “Ineligible Waits time” bucket in CS and JW
  - Setting too low can lead to severe performance problems
  - Setting too high can lead to more faulting

IBM i memory tuning features

- Expert Cache
- Automatic performance adjust system value (QPFRADJ)
- SETOBJACC command
- DB2 keep in memory

Expert Cache (*CALC)

- Paging parameter for shared storage pools
  - *FIXED
  - ***CALC** – enables expert cache
- What it does:
  - Monitors the I/O reference pattern for database files
  - Reduces I/O operations by adjusting the size and type of I/Os
- Cannot be used by:
  - Private pools (can be used by running an API – QWCCHGTN)
  - Machine pool
- Activity you may see:
  - SMXCSPRVSR (Expert Cache supervisor) task
  - SMXCAGERnn tasks (1 per *CALC pool, 01 => pool 2)
- Expert Cache almost always provides benefit and should be enabled

Expert Cache Enablement

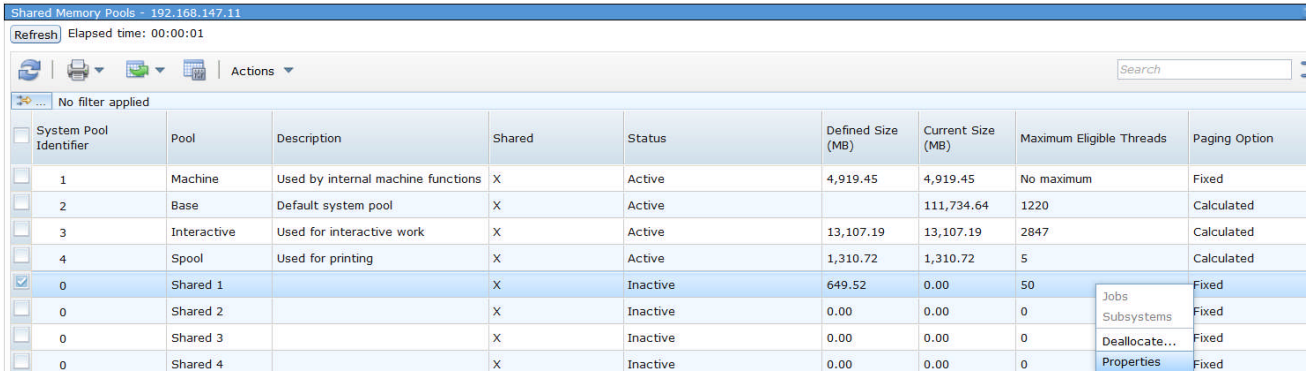

System Pool Identifier	Pool	Description	Shared	Status	Defined Size (MB)	Current Size (MB)	Maximum Eligible Threads	Paging Option
1	Machine	Used by internal machine functions	X	Active	4,919.45	4,919.45	No maximum	Fixed
2	Base	Default system pool	X	Active		111,734.64	1220	Calculated
3	Interactive	Used for interactive work	X	Active	13,107.19	13,107.19	2847	Calculated
4	Spool	Used for printing	X	Active	1,310.72	1,310.72	5	Calculated
0	Shared 1		X	Inactive	649.52	0.00	50	Fixed
0	Shared 2		X	Inactive	0.00	0.00	0	Fixed
0	Shared 3		X	Inactive	0.00	0.00	0	Fixed
0	Shared 4		X	Inactive	0.00	0.00	0	Fixed

```
                         Work with System Status                CTCLWPRF
                                                        03/13/15  10:19:17
% CPU used . . . . . :          .5    Auxiliary storage:
Elapsed time . . . . :   00:00:00      System ASP . . . . . :     9080 G
Jobs in system . . . :          670    % system ASP used  . . :   87.4248
% perm addresses . . :         .016    Total  . . . . . . . :     9080 G
% temp addresses . . :         .263    Current unprotect used :   184855 M
                                       Maximum unprotect  . . :   312991 M

Type changes (if allowed), press Enter.

System     Pool     Reserved     Max   Paging
 Pool   Size (M)   Size (M)   Active  Option
   1     1632.78     843.36   +++++  *FIXED
   2    28677.61       7.51     300  *CALC
   3      163.83      <.01      41  *CALC
   4      327.67        .00       5  *CALC
   5      327.67        .00      50  *CALC
                                                              More...
Command
===>
F3=Exit    F4=Prompt    F5=Refresh    F9=Retrieve    F10=Restart    F12=Cancel
F19=Extended system status          F24=More keys
```

```
                              Work with Shared Pools
                                                              System:   CTCLWPRF
Main storage size (M)  . :       32768.00

Type changes (if allowed), press Enter.

              Defined    Max   Allocated    Pool  -Paging Option--
Pool        Size (M)  Active  Size (M)      ID   Defined  Current
*MACHINE     1632.78   +++++     1632.78     1   *FIXED   *FIXED
*BASE       28677.61     300    28677.61     2   *CALC    *CALC
*INTERACT    1638.39      57     1638.39     6   *CALC    *CALC
*SPOOL        327.67       5      327.67     4   *CALC    *CALC
*SHRPOOL1     163.83      41      163.83     3   *CALC    *CALC
*SHRPOOL2   10485.75      50                     *CALC
*SHRPOOL3     327.67      50      327.67     5   *CALC    *CALC
*SHRPOOL4    3276.79     820                     *FIXED
*SHRPOOL5   26099.82    2048                     *CALC
*SHRPOOL6        .00       0                     *CALC
                                                                  More...
Command
===>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F11=Display tuning data
F12=Cancel
```

The Performance Adjuster

- Enabled using the QPFRADJ system value.
- Will manage the size of the shared memory pools for you.
  - Will also adjust the maximum activity level.
- Uses complex algorithms to ensure your pools are operating at peak efficiency.
- Refer to *The Performance Adjuster (QPFRADJ)* experience report on the IBM i Information Center.
  http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=%2Fexperience%2Fwork3abstract.htm

Automatic Changes in Pool Sizes - PDI

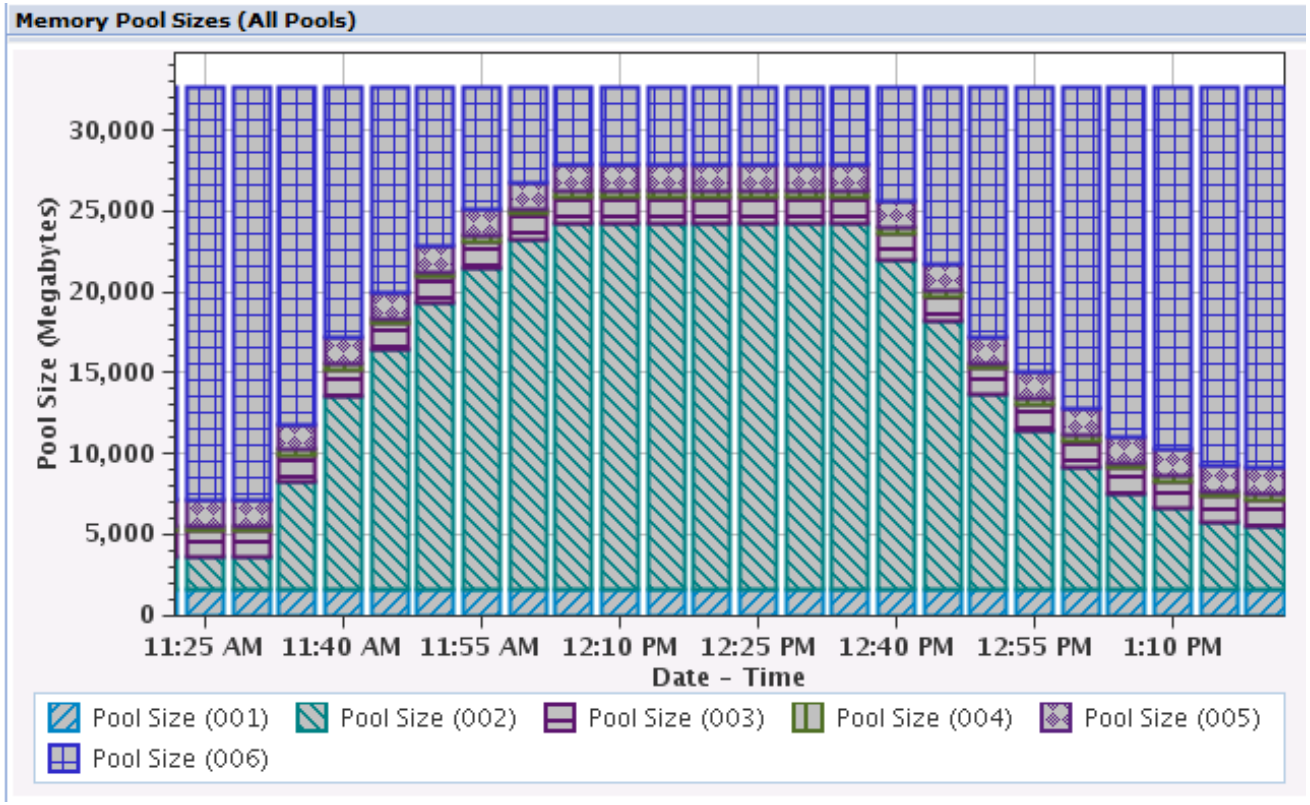

Turning QPFRADJ On or Off

The Performance Adjuster is controlled by the System Value QPFRADJ

- 0=No adjustment ← User managed!
- 1=At IPL only based on static information
- 2=At IPL and automatic (default)
- 3=Automatic only ← Recommended

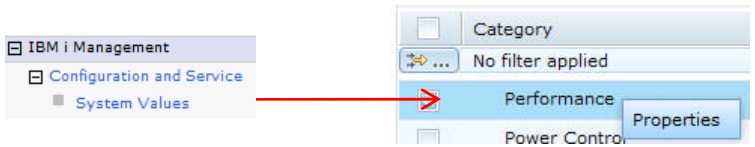

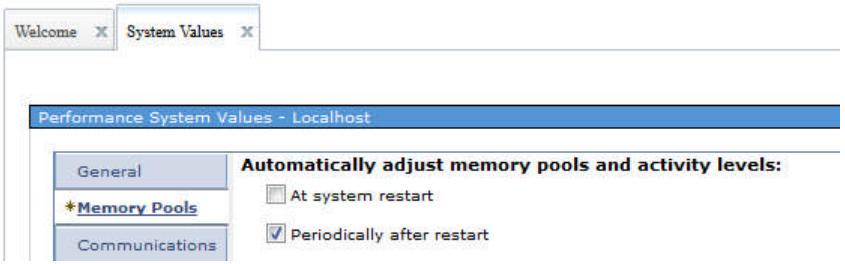

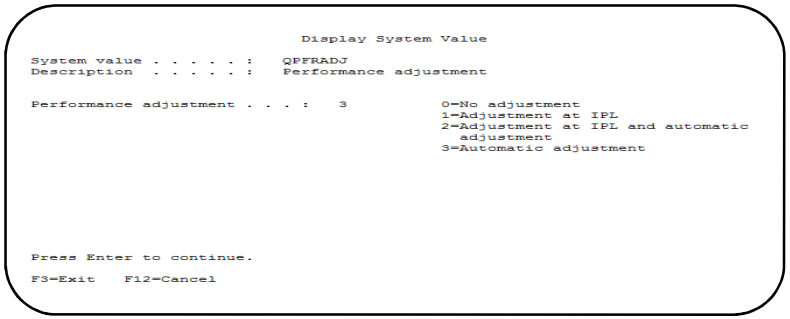

Tuning QPFRADJ (WRKSHRPOOL)

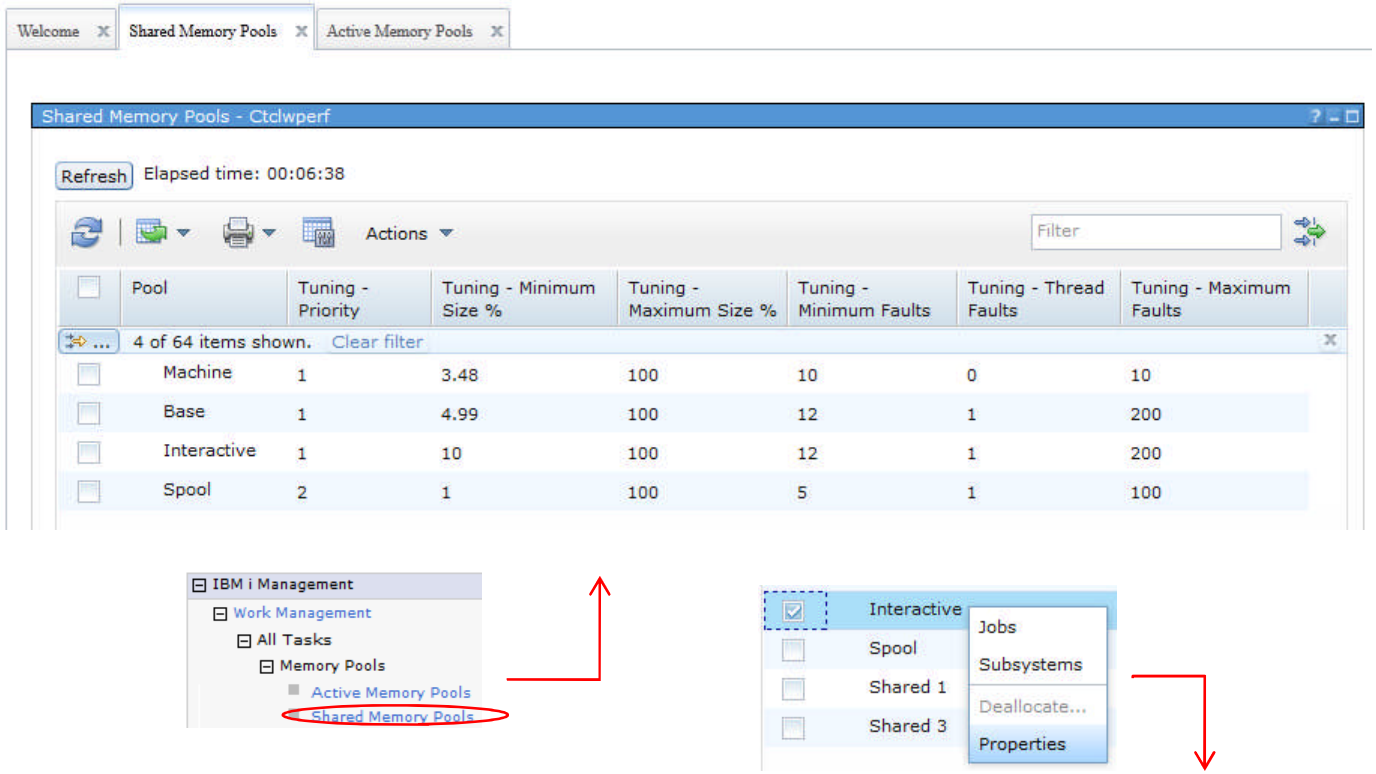

Tuning QPFRADJ

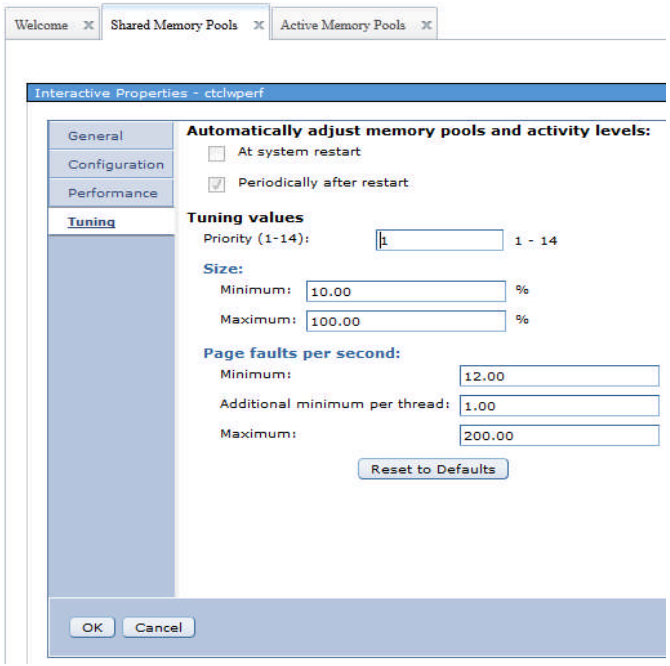

WRKSHRPOOL:

```
                          Work with Shared Pools
                                                                  System:
Main storage size (M)  . :       32476.00

Type changes (if allowed), press Enter.

                          -----Size %-----      -----Faults/Second------
Pool          Priority  Minimum  Maximum   Minimum  Thread  Maximum
*MACHINE             1     5.00      100     10.00     .00    10.00
*BASE                1     4.99      100     12.00    1.00      200
*INTERACT            1    10.00      100     12.00    1.00      200
*SPOOL               2     1.00      100      5.00    1.00      100
*SHRPOOL1            2     1.00      100     10.00    2.00      100
*SHRPOOL2            2     1.00      100     10.00    2.00      100
*SHRPOOL3            2     1.00      100     10.00    2.00      100
*SHRPOOL4            2     1.00      100     10.00    2.00      100
*SHRPOOL5            2     1.00      100     10.00    2.00      100
*SHRPOOL6            2     1.00      100     10.00    2.00      100
                                                                  More...
Command
===>
F3=Exit    F4=Prompt    F5=Refresh    F9=Retrieve    F11=Display text
F12=Cancel
```

Tuning shared pools when QPFRADJ is ON

- Use WRKSHRPOOL command to tune shared memory pools
- Determine min/max size for critical pools
  - Monitor max active settings
- Set a range of pool priorities.
- **Java/WebSphere pools minimum size should always be enough to contain all JVMs in the pool**
- Set max on memory intensive pools to limit impact to other jobs on the system
- Large changes to the size of the memory pool can cause the query optimizer to rebuild access plans which can contribute to poor performance
- **If see pools sitting at min size a lot, consider decreasing its min size**
- **If see pools hitting max size, consider increasing max size.**
- **Keep total of minimum sizes < 70% of memory to allow QPFRADJ some flexibility**
- When adding memory, may need to adjust min/max
  - When memory is for a particular workload, rather than a general upgrade
  - Values are percentages of total

Refer to *The Performance Adjuster (QPFRADJ)* experience report on the IBM i Information Center.
http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp?topic=%2Fexperience%2Fwork3abstract.htm

SETOBJACC to ‘pin’ objects in memory

- Allows selected objects (database file, index or program) to be "pinned" in memory to reduce I/O
  - Typically, define a private pool where no jobs run
  - Can use a shared pool. Set Min/Max to keep QPFRADJ from shrinking
- May increase memory requirements
- Protects objects from “demand paging” activity
- Objects loaded into memory very quickly by a single thread
- Typically used to improve performance of batch jobs
- Run a CLRPOOL command before loading objects in pool
- Re-run SETOBJACC periodically to pin changed pages (updates/inserts)
- Can load and purge objects programmatically during batch processing

SETOBJACC/CLRPOOL Command Technote:

http://www.ibm.com/support/docview.wss?uid=nas1dc0a2297bdaefddb86256d6c0069907f

DB2 KEEPINMEM to 'pin' objects in memory

- New in IBM i 7.1
- Can be used **with tables and indexes**
- Objects are brought into memory when first accessed by SQE queries
- Brought in asynchronously and using parallel I/O
  (vs. SETOBJACC which uses 1 thread and synch IO)
- Usage:
  - CHGPF FILE(library/table) KEEPINMEM(*YES)
  - CHGLF FILE(library/index) KEEPINMEM(*YES)
- The MEMORY_POOL_PREFERENCE parameter in the QAQQINI file determines which pool the objects will be held in unless the object is already in memory
- In IBM i 7.2. the KEEP IN MEMORY memory-preference support has been extended to the DDL statements: ALTER TABLE, CREATE INDEX, CREATE TABLE and DECLARE GLOBAL TEMPORARY TABLE

Memory Pool and Subsystem Creation Recommendations

- Minimize user work in *BASE (in general)
- Put disparate workloads into separate pools
- Often want to run SQL in its own pool
- Heavy memory use when not critical
  - Often see software replication jobs in own pool with lots of faulting, but still keeping up
- Java / WebSphere should run in its own memory pool
  - These environments cannot tolerate faulting
- Controlling potentially inefficient work
  - QZDA jobs in own pool by IP for controlling ad-hoc queries
- Max active settings
  - If not using QPFRADJ set max active high enough to avoid transitions to ineligible
  - If running 6.1 don’t set too high for pools running SQL
- Settings related to QPFRADJ (min/max size, etc.)
- If see high priority page out task (SMPO0001), pool needs more memory.

Tools Used for Memory Analysis

- Real Time Tools
  - Green screen commands / IBM Navigator for i
    - WRKSYSSTS, WRKACTJOB
  - IBM Navigator for i Dashboard
  - IBM Navigator for i Monitors (New in IBM i 7.2)
- Collector based
  - Collection Services – System and Job level memory related wait times, fault rates, and pool sizes
  - Job Watcher – Object waited on, Call Stacks, and SQL statement
- GUI Tools
  - **Performance Data Investigator** (“PDI” Web based, part of IBM Navigator for i) – Collection Services, Job Watcher, Disk Watcher, limited PEX
  - **IBM iDoctor for i** (Windows based) – Collection Services, Job Watcher, Disk Watcher, PEX Analyzer
  - System i Navigator

WRKSYSSTS – Faulting by Memory pool

- Can be invoked from green screen or web based IBM navigator for i
- Faults per second column doesn't show real fault wait times
- Monitor the max active and ineligibles

```
                          Work with System Status
                                                          02/25/15  10:24:54
% CPU used . . . . . . . . :       24.9     System ASP . . . . . . . . :     6209 G
Elapsed time . . . . . . . :   00:00:01     % system ASP used  . . . . :    66.4656
Jobs in system . . . . . . :      83397     Total aux stg  . . . . . . :    20322 G
% perm addresses . . . . . :      1.699     Current unprotect used . . :     1100 G
% temp addresses . . . . . :     51.451     Maximum unprotect  . . . . :     1489 G


Sys       Pool    Reserved     Max  ----DB-----  --Non-DB---  Act-   Wait-  Act-
Pool    Size M    Size M       Act  Fault Pages  Fault Pages  Wait   Inel   Inel
  1     103923     25891   +++++    .0    .0     .0    .0   850.6    .0     .0
  2     256397        70     850    .0    .0     3.7  26.4  10718     .0     .0
  3      81920         3    1100   16.0 148.3    13.2  32.1   1928     .0     .0
  4      12288         0     150    .0    .0     .0    .0     .0     .0     .0
  5     163840        <1    1200  116.2  2120  174.8 602.0   5330     .0     .0
  6      16384        <1      90   68.9 441.3    3.7   3.7  510.3     .0     .0
  7      40960        26    3000    .0    .0     3.7  17.0  +++++     .0     .0
  8      20480         0     500    .0    .0     .0    .0   964.0     .0     .0
  9     122880        <1     114    9.4  77.5     .0    .0   226.8     .0     .0
                                                                    More...
===>
F21=Select assistance level
```

Memory Pools (WRKSYSSTS)

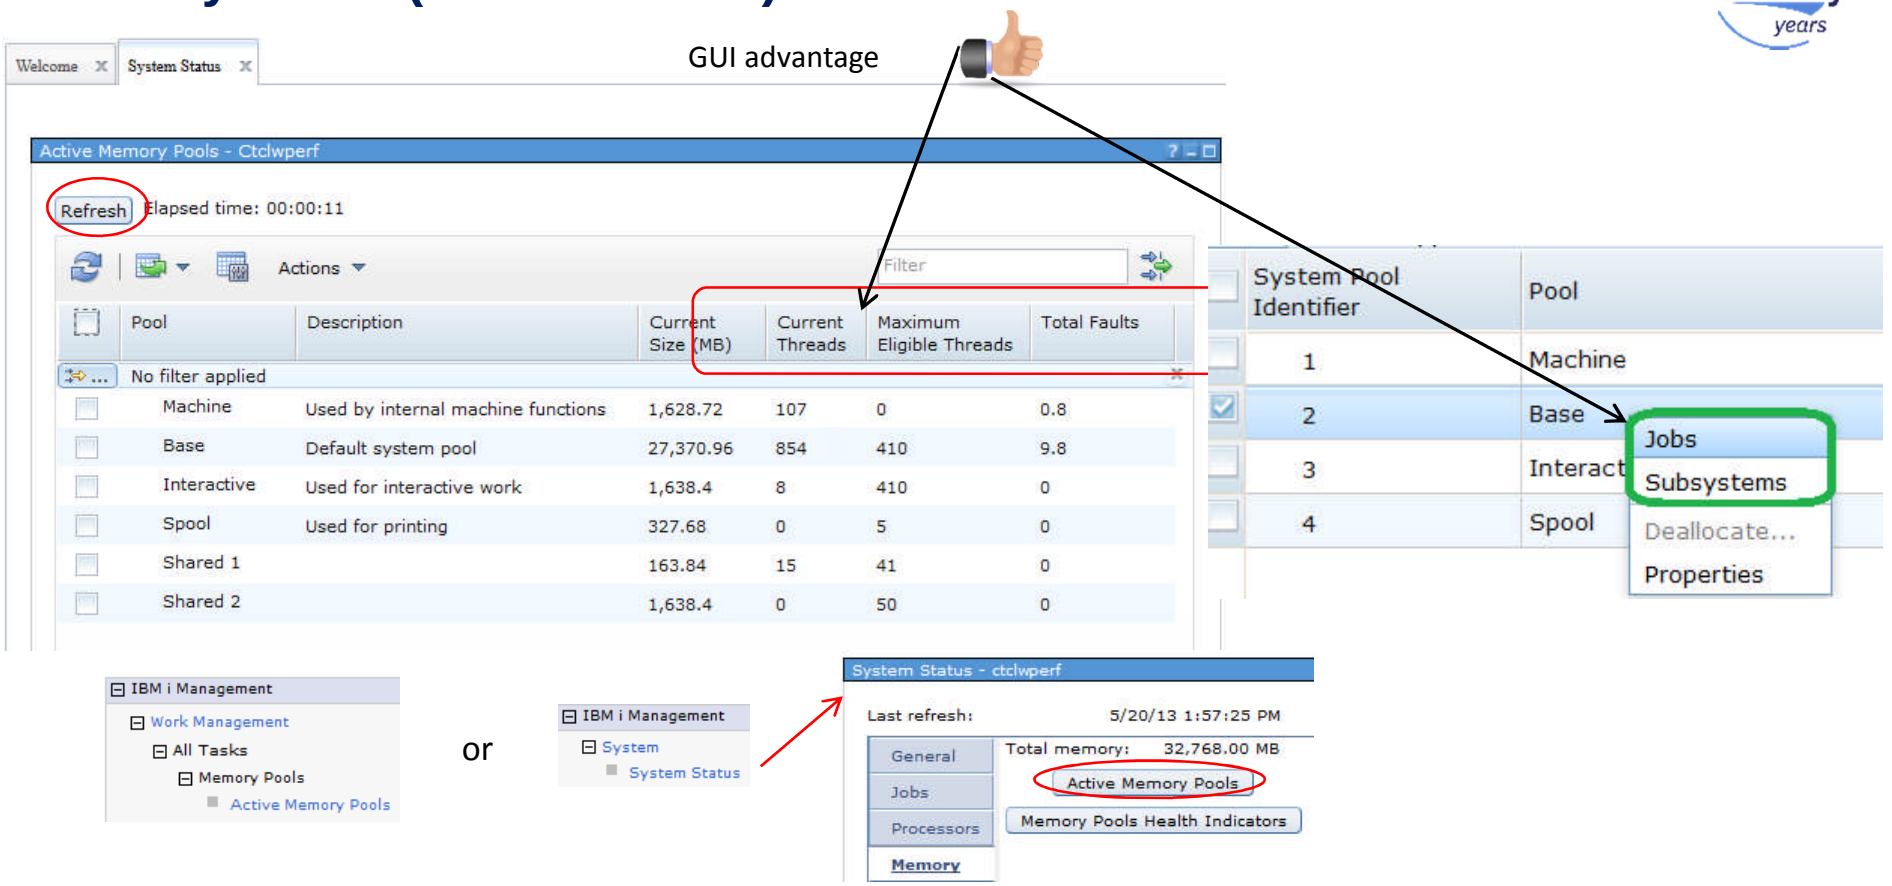

WRKACTJOB – Page faulting by job

The Page fault rate column is available of the IBM Navigator for i interface

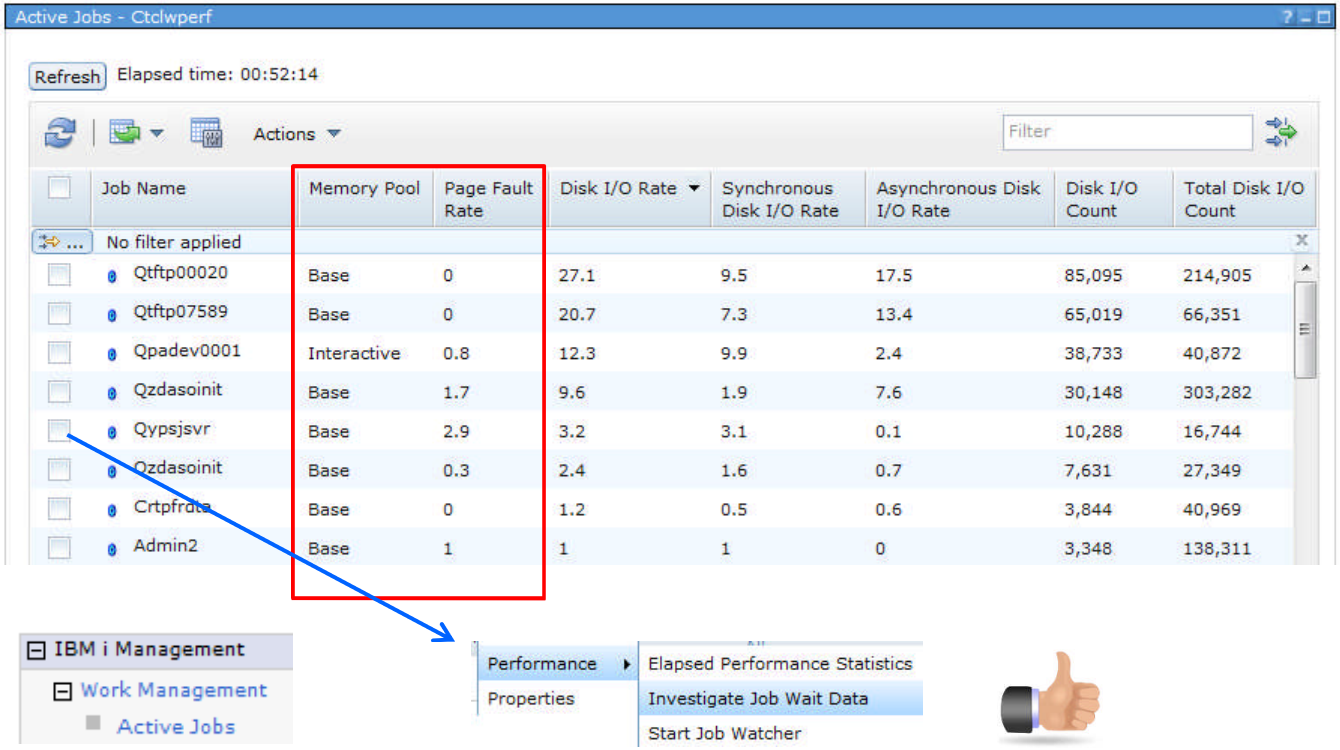

Navigator Dashboard

Also available in iAccess Mobile: http://system.name:2001/iamobile

Navigator for i System Monitors (7.2)

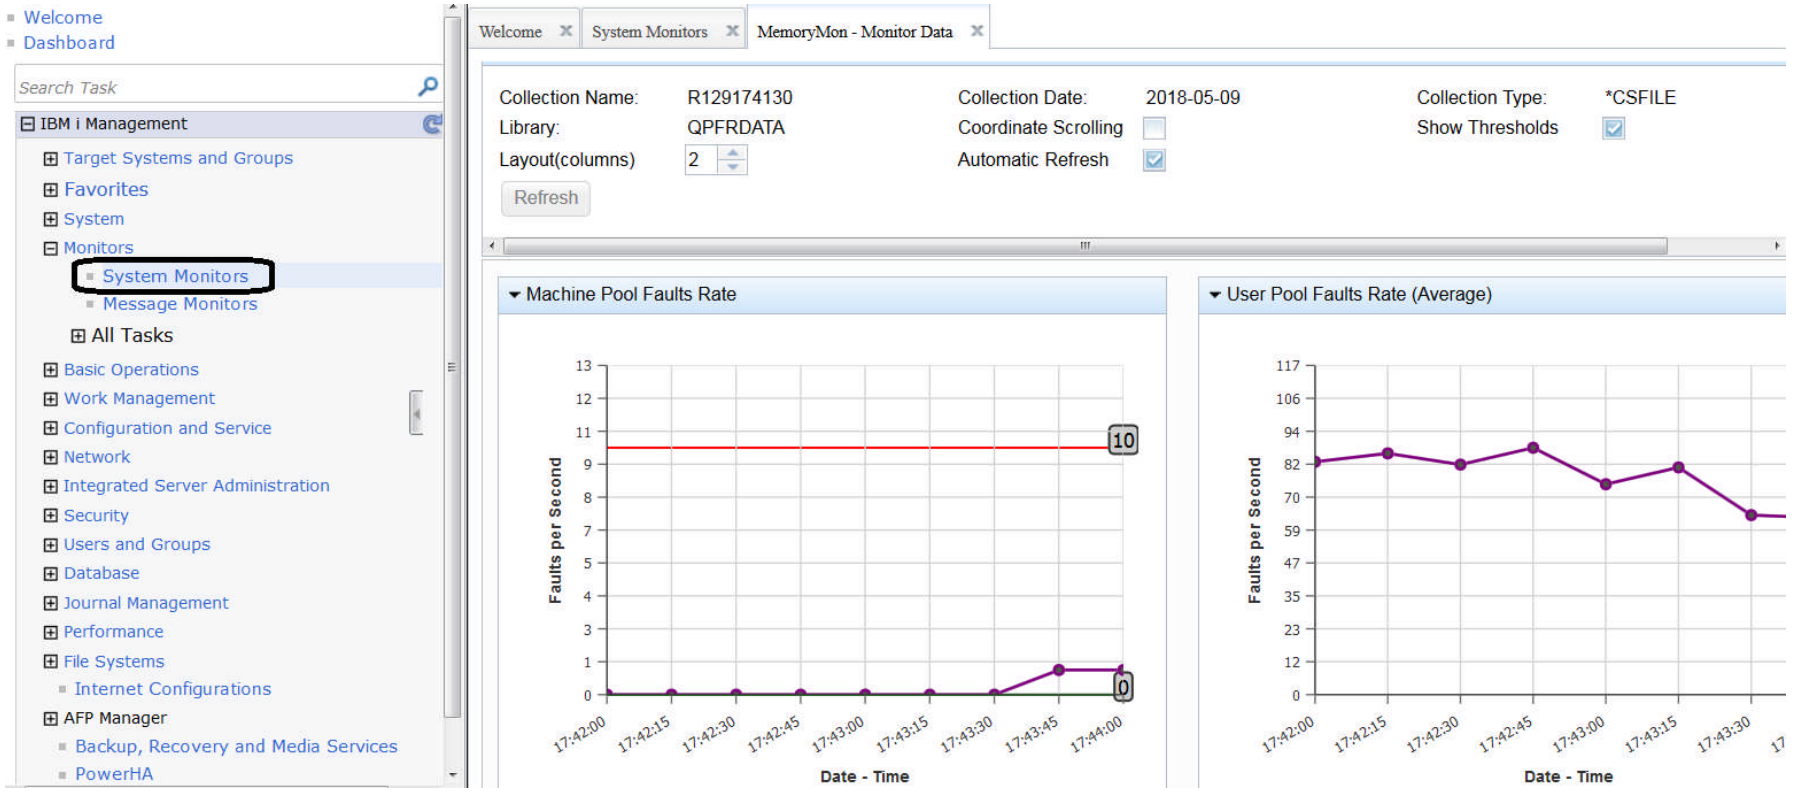

IBM Graphical Analysis Tools

IBM provides two powerful tools to aid in making your analysis **more efficient and productive:**

- **Performance Data Investigator**
- **IBM iDoctor for IBM i**

Both solutions support data analysis (varying degrees) for the 4 collectors:

- Collection Services
- Job Watcher
- Disk Watcher
- Performance Explorer (PEX)

Performance Data Investigator (PDI)

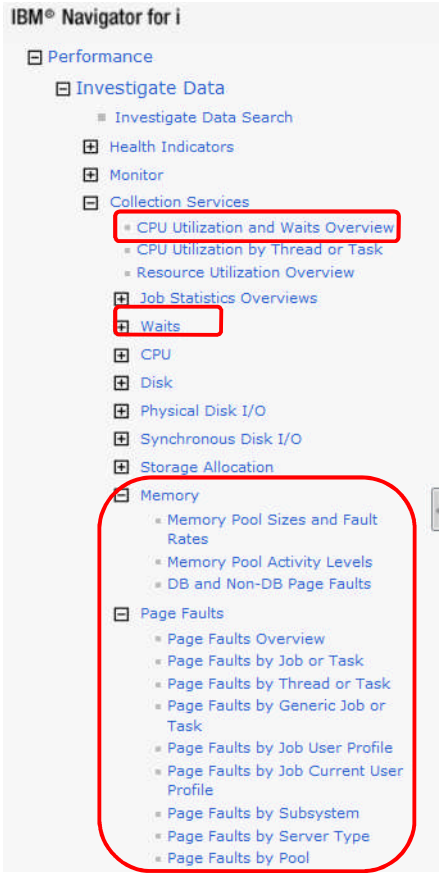

- Browser (web) based solution
- Integrated as part of IBM i OS
- Included in IBM Navigator for i

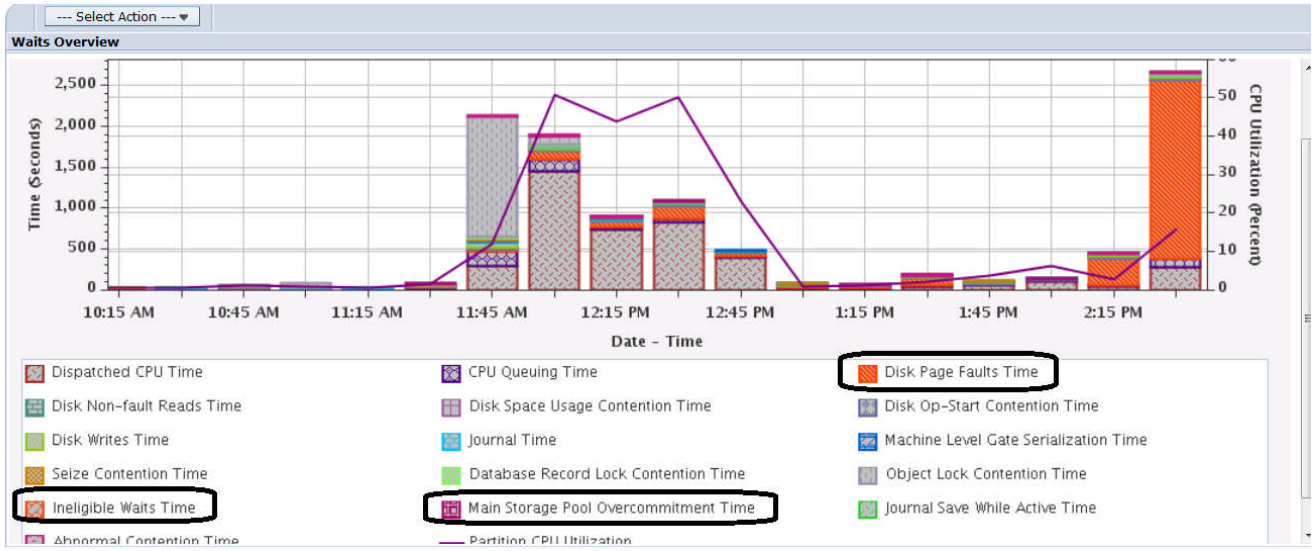

IBM iDoctor for IBM i

- Microsoft Windows based client
- Service/Support offering
- Deep Job Watcher and PEX analysis capability

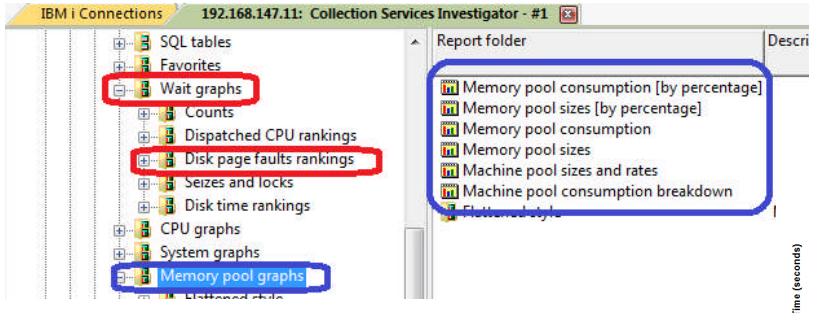

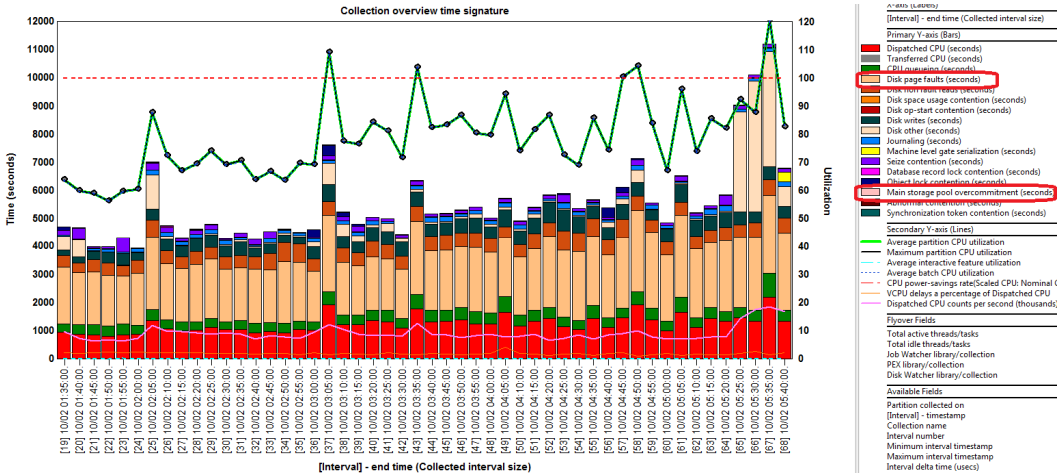

Graphical Analysis Tools

- You have two graphical interfaces for performance data analysis…
  - Which is right for you?

Feature	iDoctor	PDI
Interface	**Windows client**	**Browser**
Wait Analysis	Yes	Yes
Collection Services	Yes	Yes
Job Watcher	Yes (In-depth)	Yes
Disk Watcher	Yes	Yes
Performance Explorer	Yes (In-depth)	Profile collections only
Level of analysis provided	**Deep**	**Basic to Medium**
Job Watcher Monitors (Built –in)	Yes	No
User Defined graphs and queries	Yes	Yes
Update Frequency	Quarterly	Twice Yearly
Support	Email idoctor@us.ibm.com	Standard SWMA
Chargeable	**Yearly license for each component (by serial number)** **▪Job Watcher** **–Includes Job Watcher, Collection Services Investigator, and Disk Watcher** **▪PEX Analyzer**	**▪Collection Services & Health Indicators at no additional charge with i** **▪Disk Watcher, Database, and Performance Explorer included with base PT1 (Performance Tools LPP) product – Option 1 Manager feature** **▪Job Watcher is an additional option of PT1 and has an additional charge - Option 3 Job Watcher**
DS8K graphs & VIOS Investigator	Yes	No
Multinational language support	No	Yes

Collection Services - Memory Related Data

- What will be analyzed using CS?
  - Page fault wait time for system, job, thread, memory pool etc.
    - At a job level, can be useful to know if the wait time occurred primarily during start-up time, or consistently while active
  - Rate of faulting in each memory pool
  - Memory pool size/config of the jobs with high page faults

Memory Analysis - CSI

- The first step is to identify page fault wait times in a collection services member

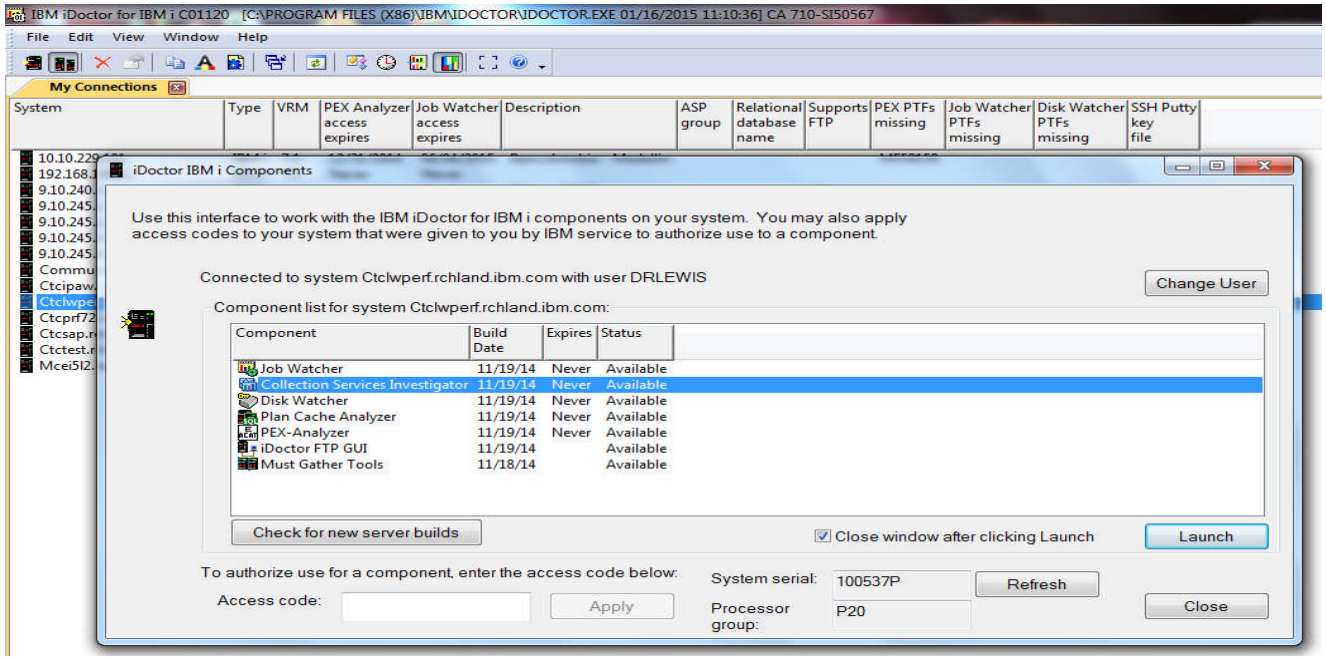

Memory Analysis - CSI

- A good place to start is the Collection overview time signature.
- We want to see the relative amount of time being spent waiting on disk faults.

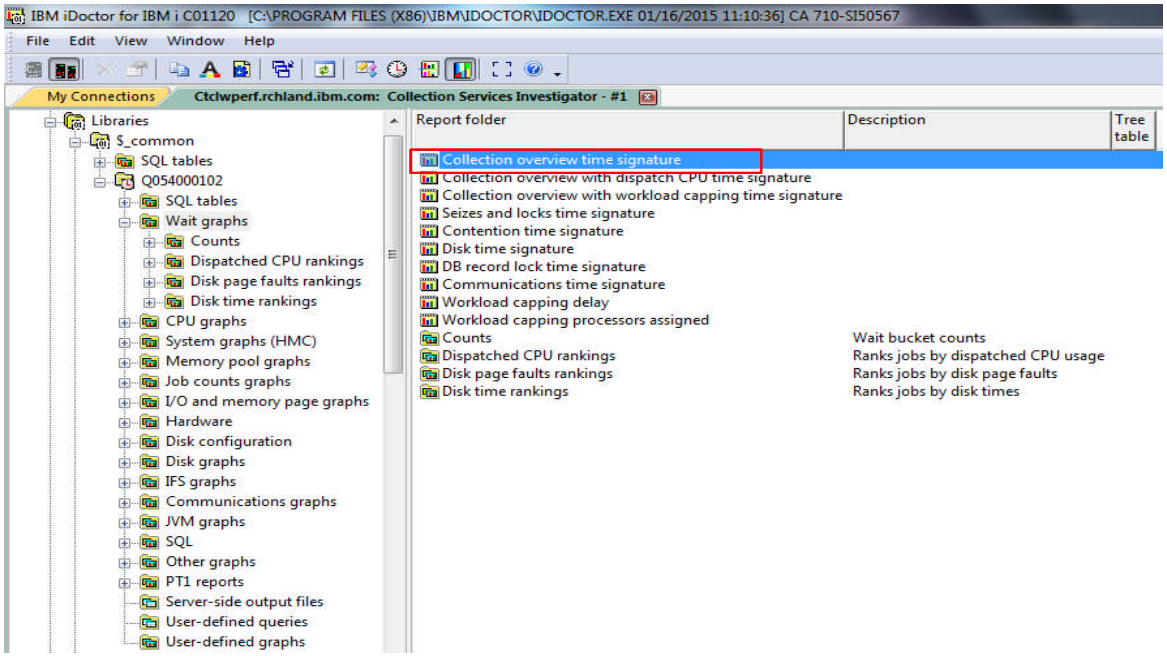

Memory Analysis - CSI

- Identify intervals with high page fault wait times (tan)
- Wait times could be skewed some by large collection services intervals

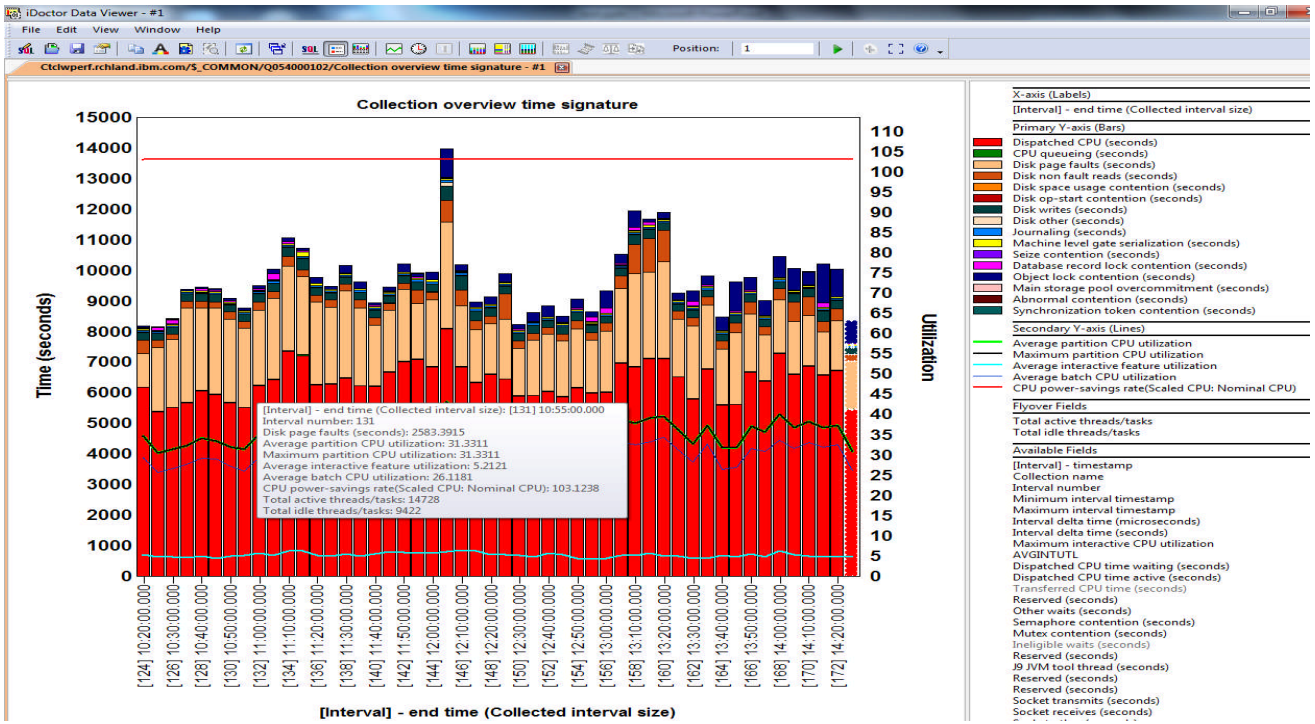

Memory Analysis – Memory Pool Graphs in CSI

- No change in pool size indicates that performance adjuster system value (QPFRADJ) was not on during this collection

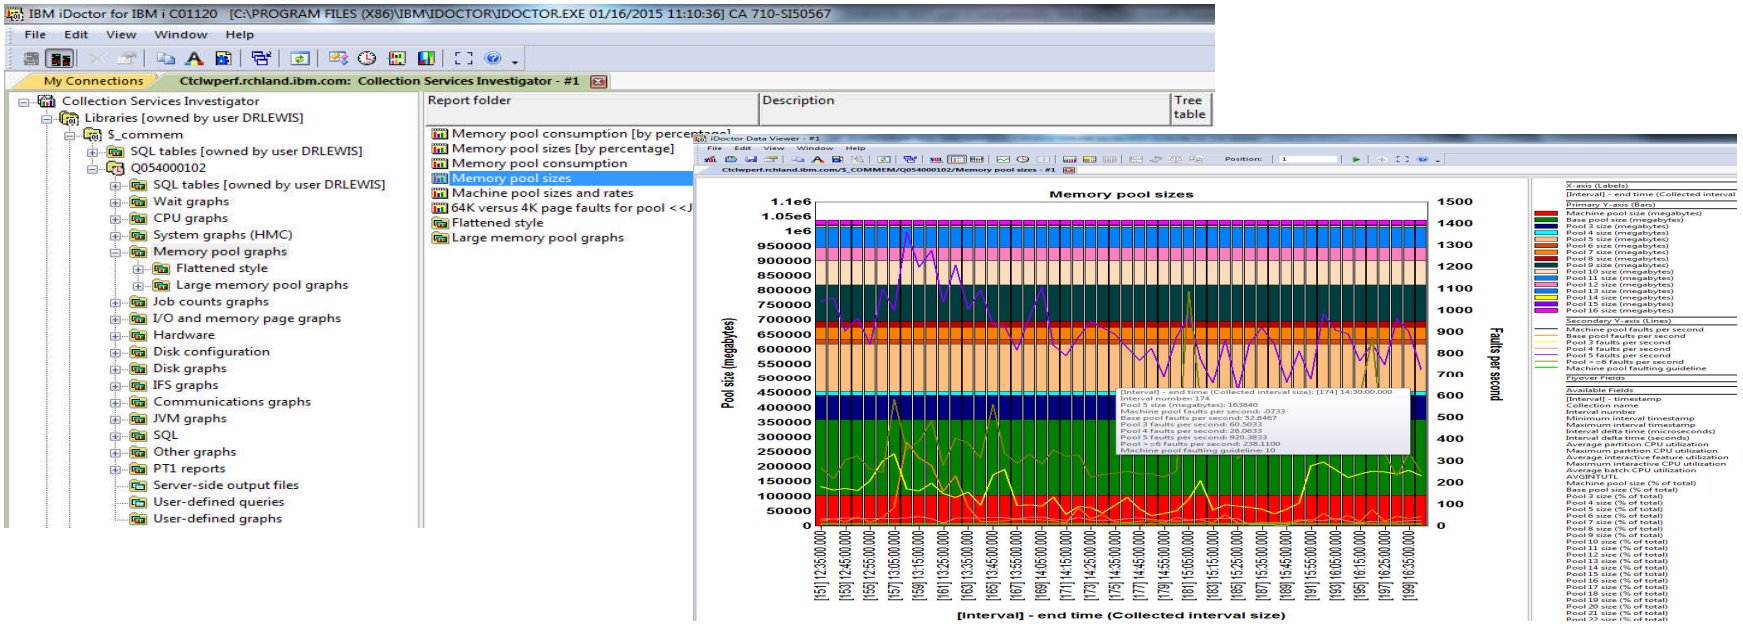

Memory Analysis - CSI

- Highlight intervals with high page fault waits and right click to rank by memory pool

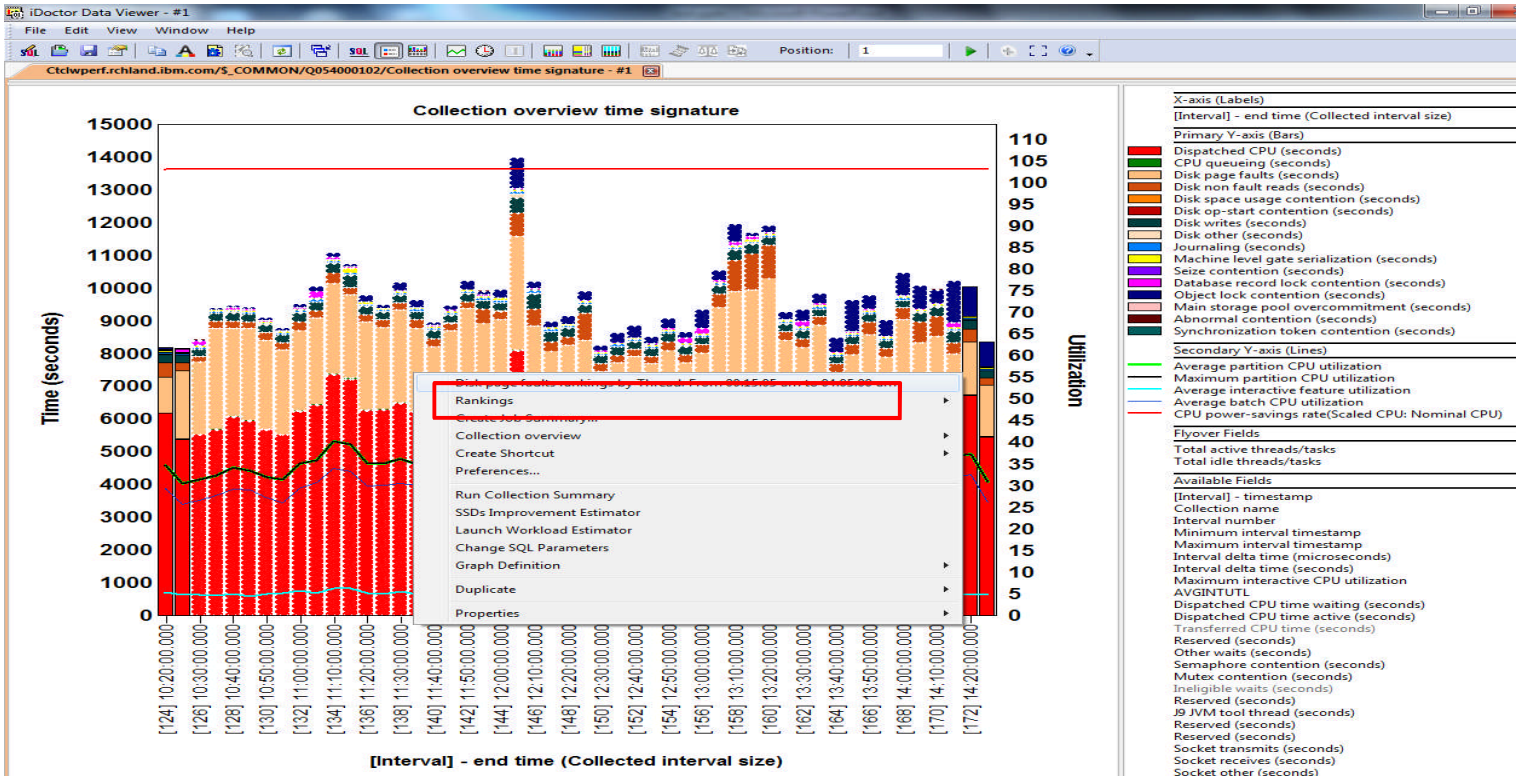

Identify top jobs in the pool with most page faults

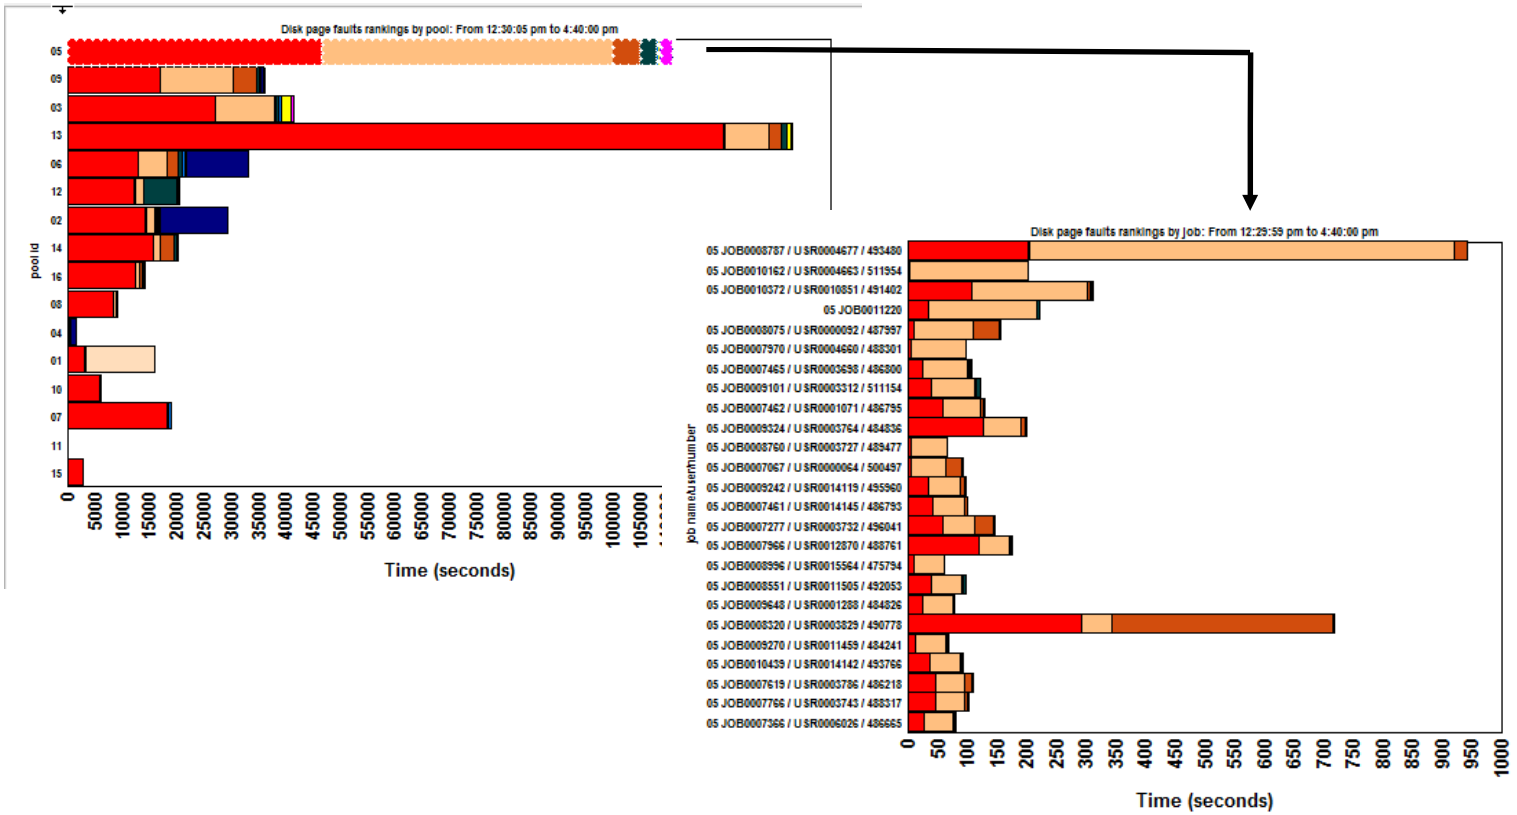

Worst page faulting jobs for all pools

- The following chart shows worst offending jobs and the pools they run in
- We added the Pool No. to the X-Axis label.

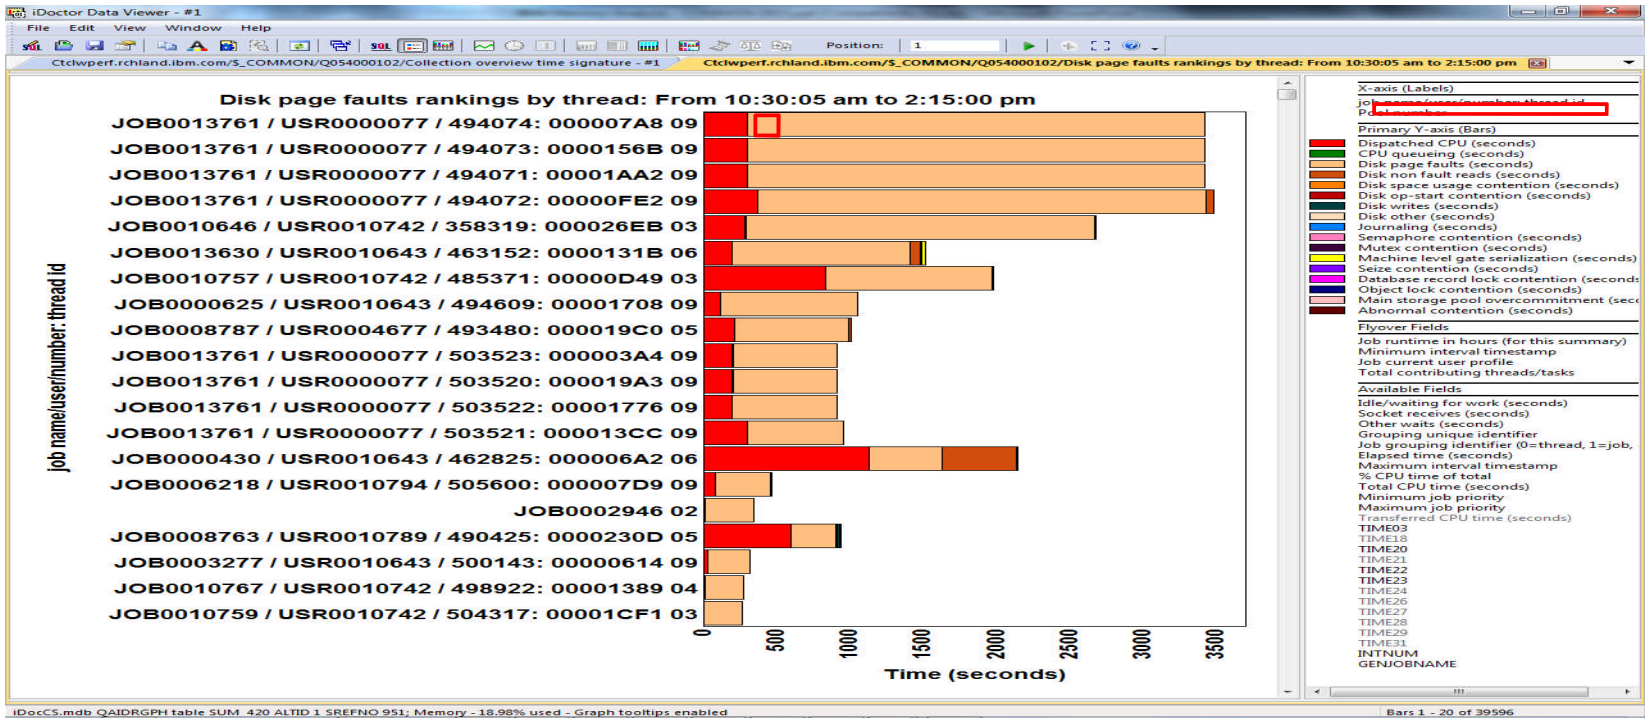

Memory Analysis - CSI

- Right click on jobs of interest to get wait time signature
- It is often useful to group threads in different ways such as by generic name, current user id etc.

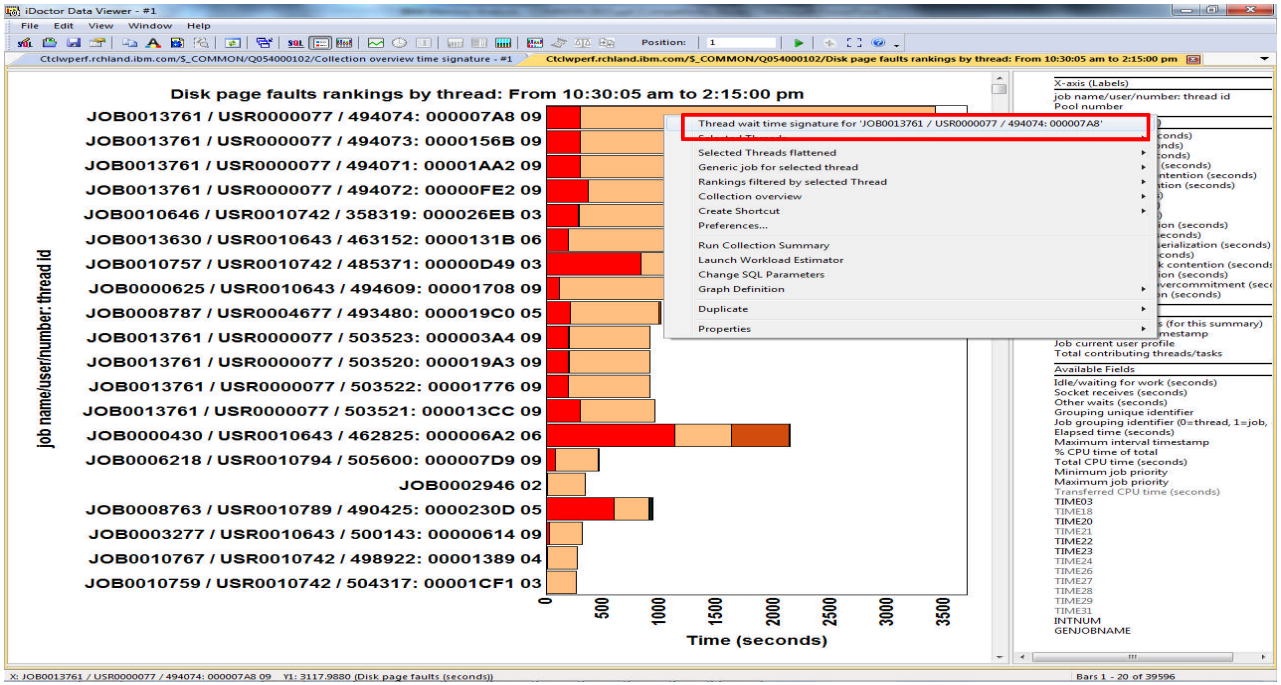

Memory Analysis - CSI

- Click on any interval to get the specific wait statistics

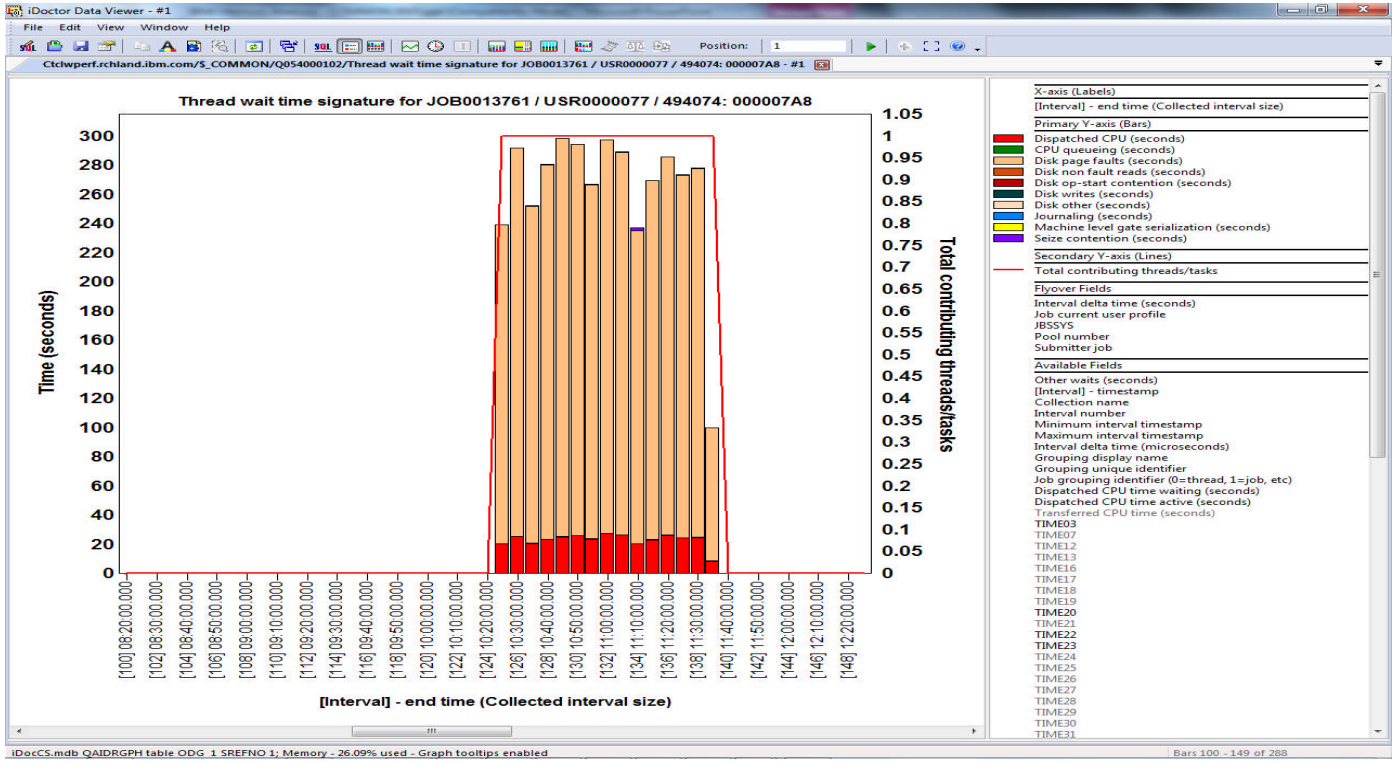

Memory Analysis - CSI

- The wait time details are shown in the tabs below

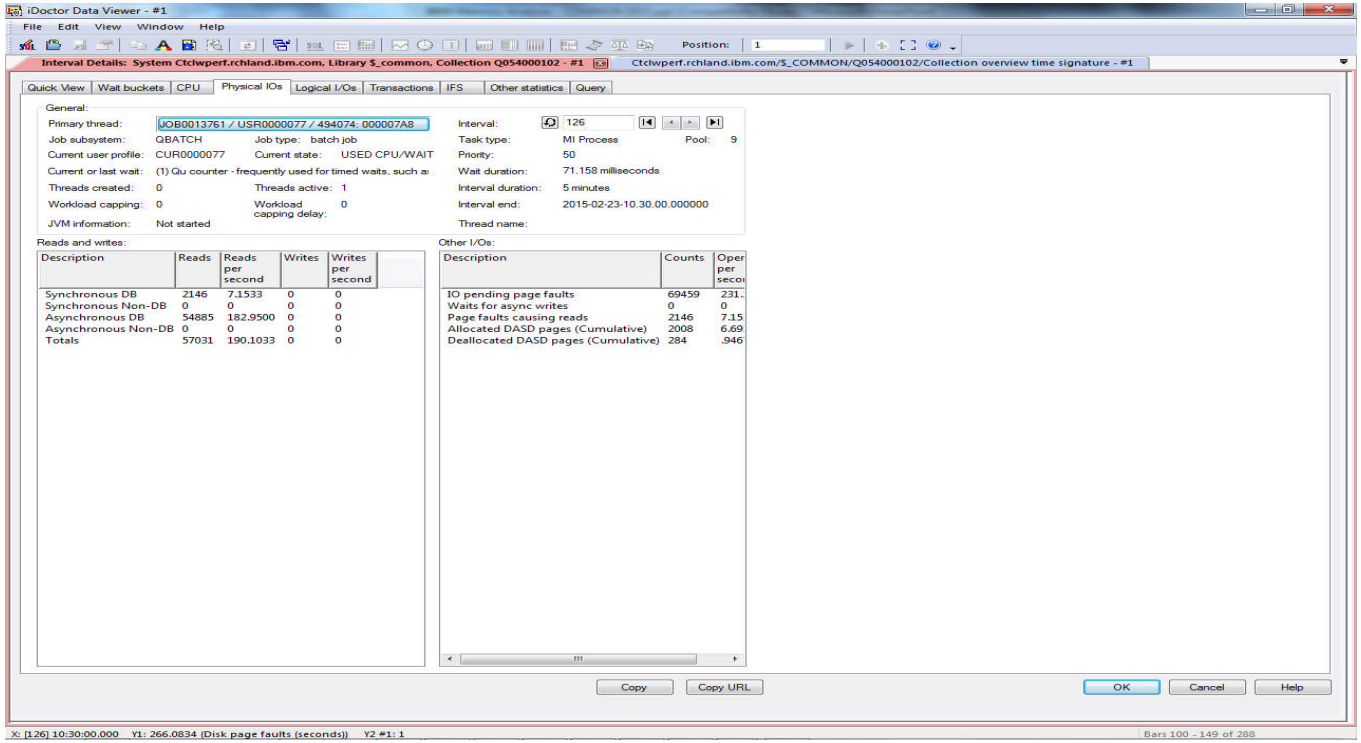

50

Collection Services Memory analysis summary

- We can tell how much faulting is happening at a system level
  - And at a job or thread or task level
- We can tell what pools have the most faulting
- We can see if memory pools are changing in size
- We can't tell what the jobs are doing that is causing faulting
  - Or what they are faulting on

- For additional details….
  - Job Watcher data can provide answers to many questions.

Job Watcher Memory Related Data

- Job Watcher has similar information as Collections Services. Key differences are:
  - Intervals typically much shorter (5-10 seconds vs. 5-15 minutes)
  - Additional data is collected
    - Call stacks are collected
    - Objects being waited on are collected
    - SQL being run is collected

- Things are different once we start looking at interval details

Launching Job Watcher

- Select Job Watcher and Click Launch

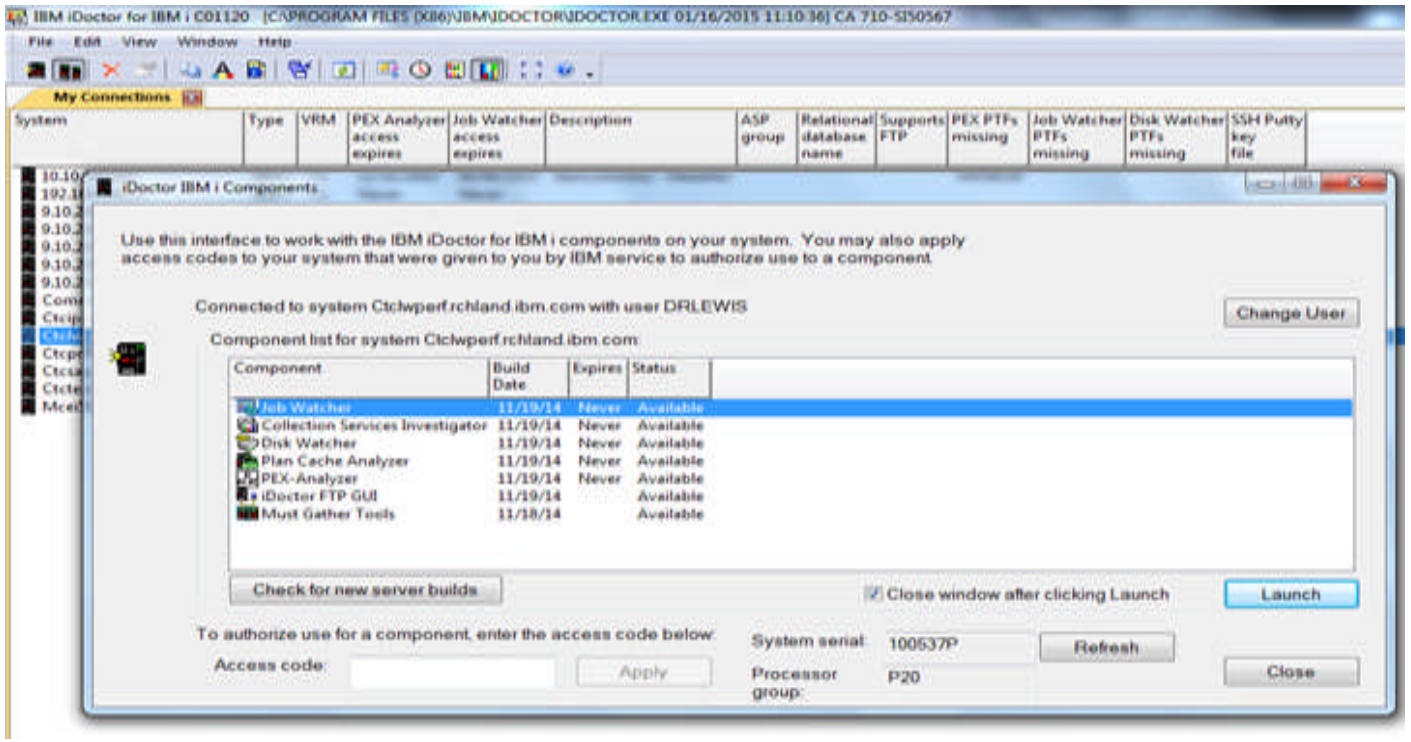

Memory Analysis - Job Watcher

- Entry Point is the same as CSI – Collection Overview time Signature, but the chart will show shorter intervals for a more obvious impact of all wait times

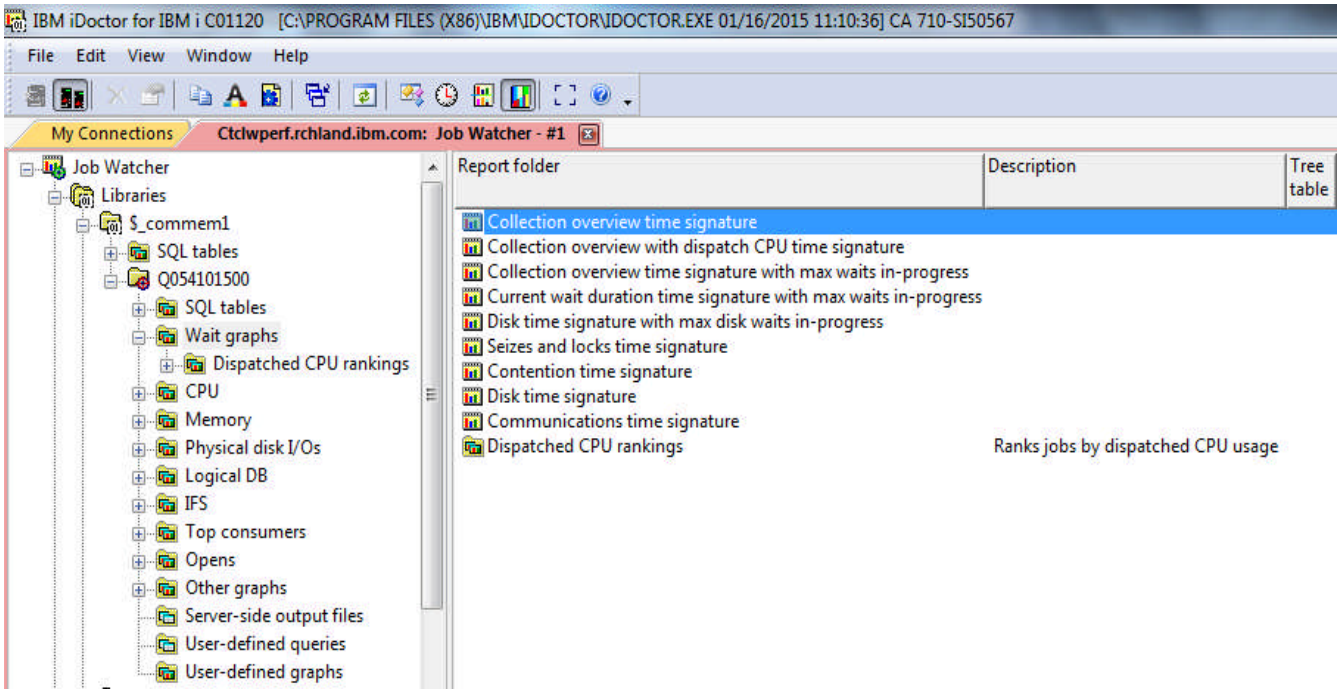

Job Watcher drill down for memory analysis

- The same drill down process can be used here as it was in CSI

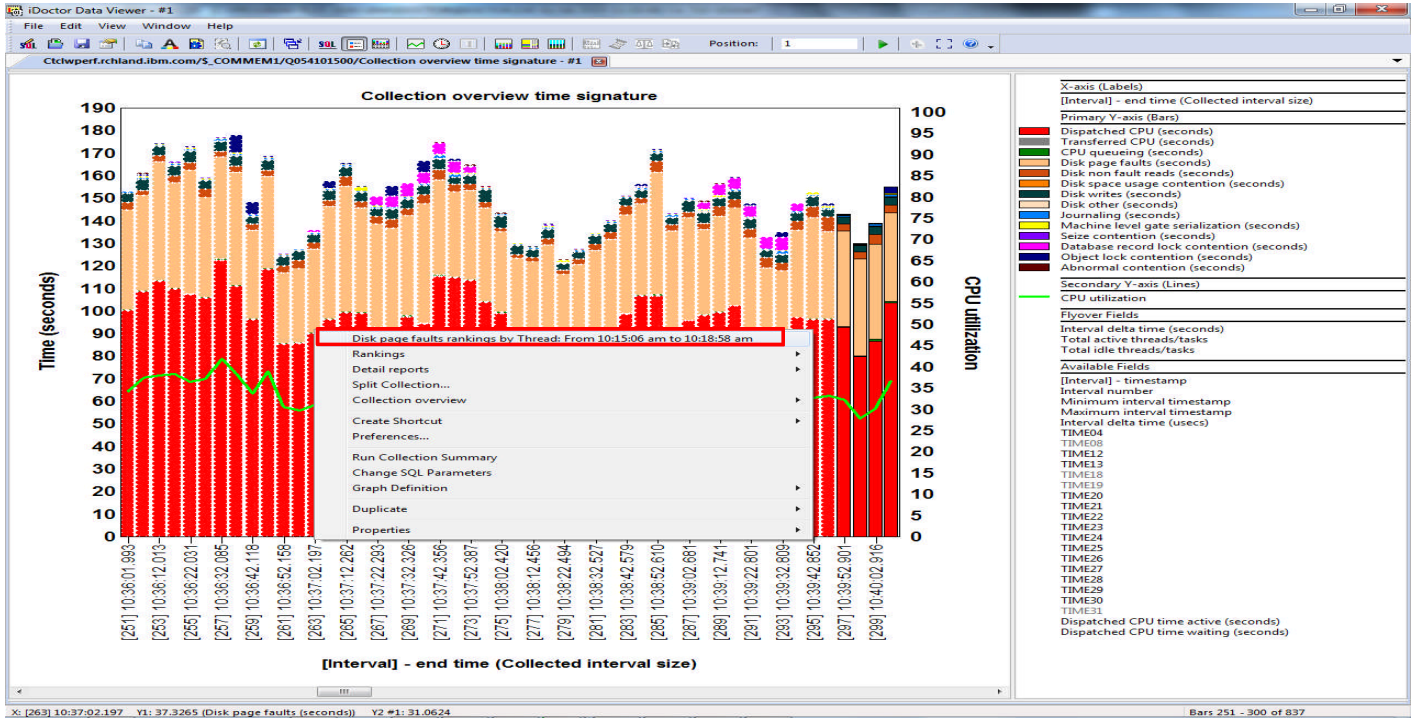

Analyzing job-threads with high page faulting

- Threads ranked by longest page fault wait times

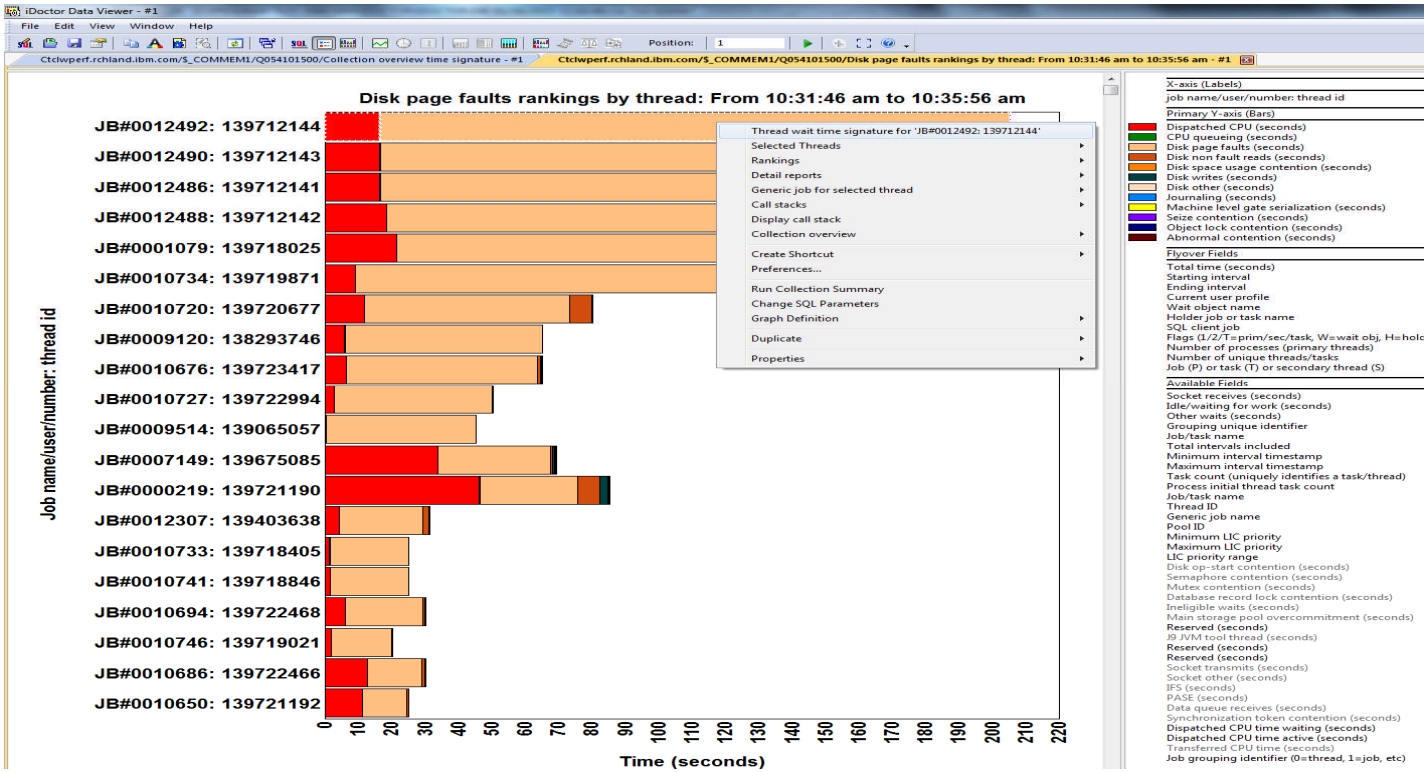

Investigating wait signature over time

- Use your mouse to fly over the intervals to see the CPU Utilization and page fault wait times in seconds

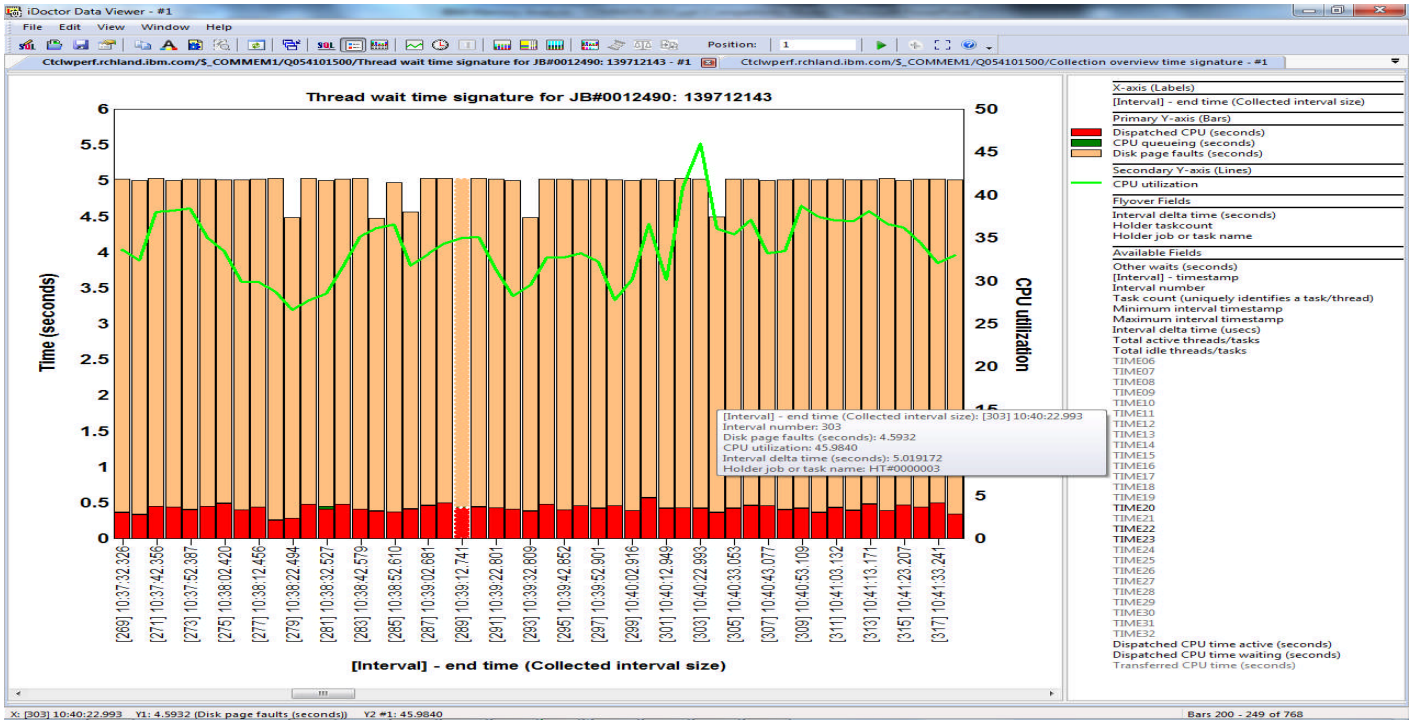

Reviewing Call stacks

- Right click on the intervals of interest and click on “Display call stack”

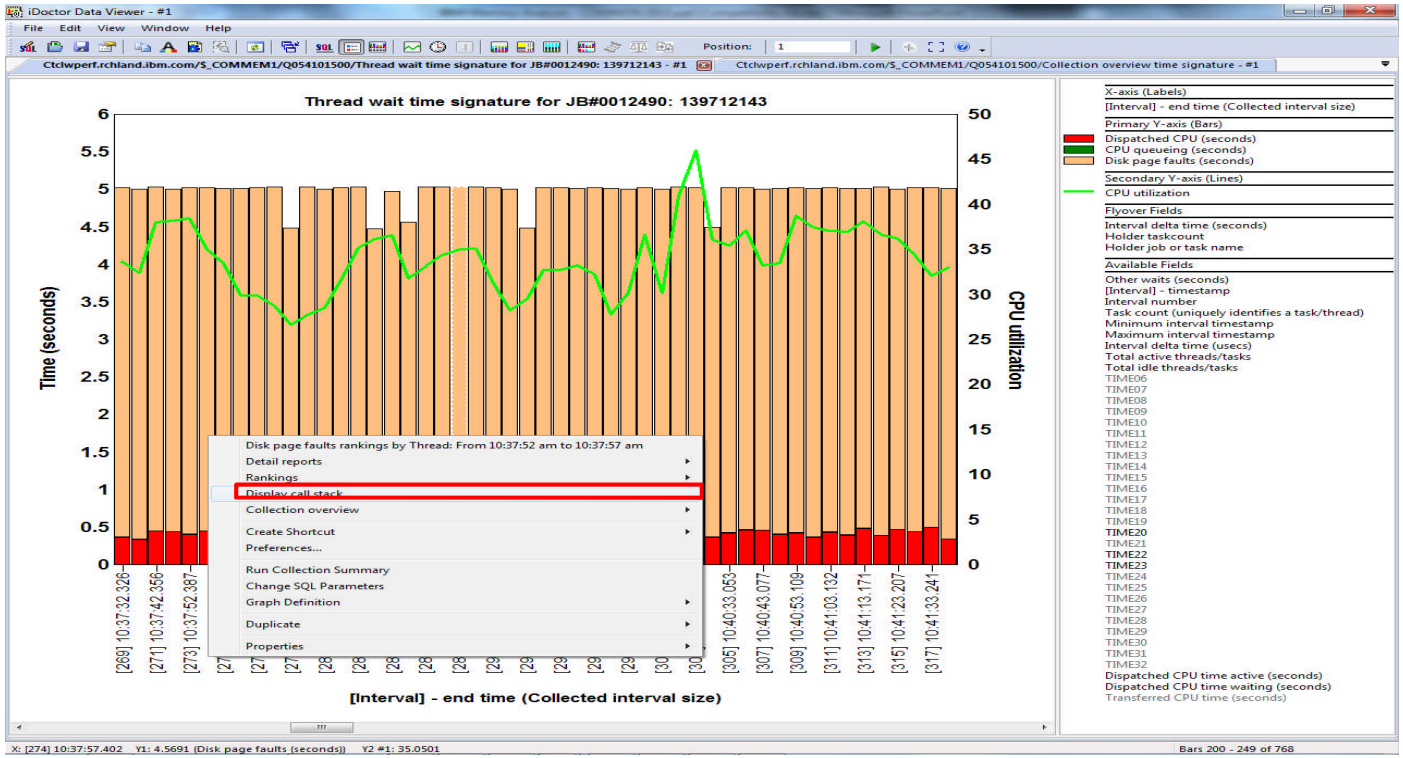

Call Stack information

- The call stack will show the object faulted on, pool, and program driving it

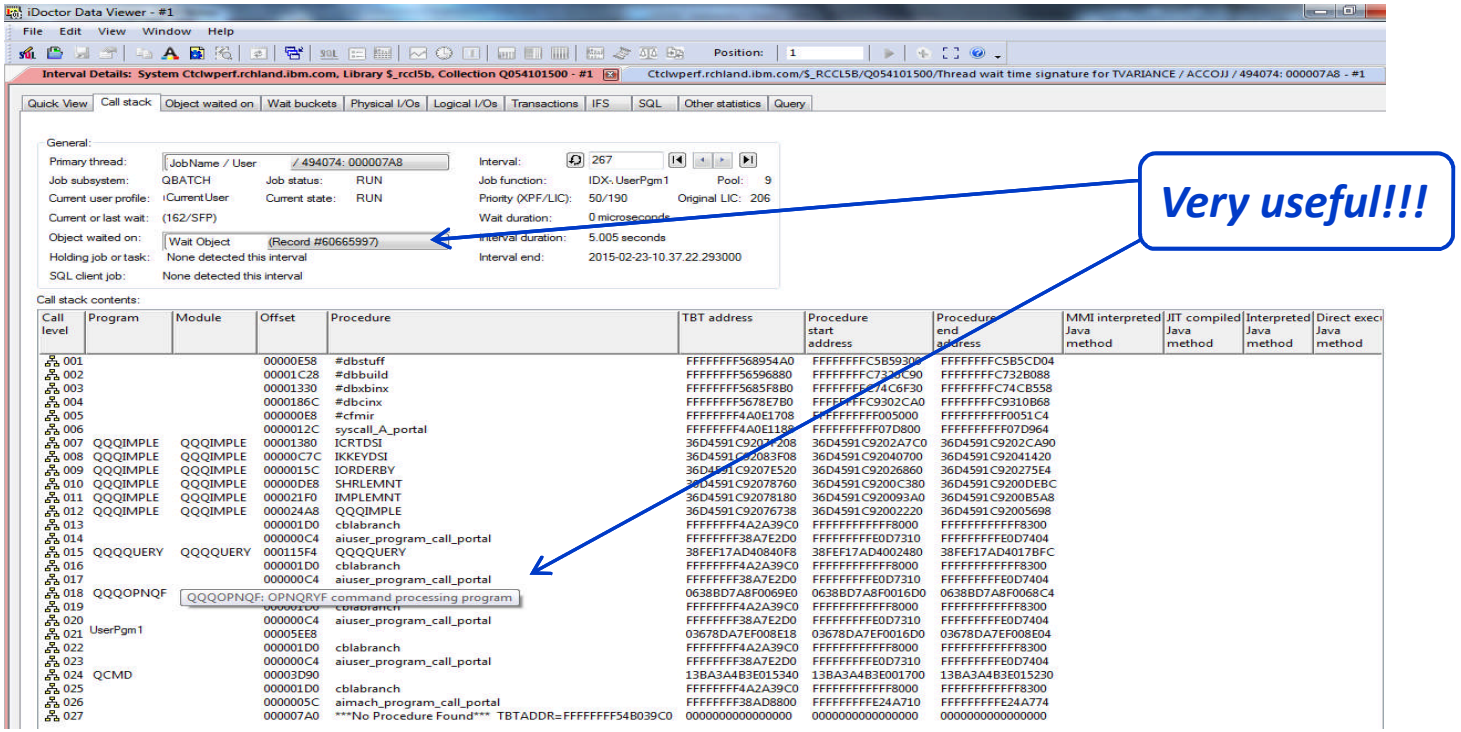

Example: Finding cause for intermittent page faults in a job

- In this example, a job shows intermittent faults

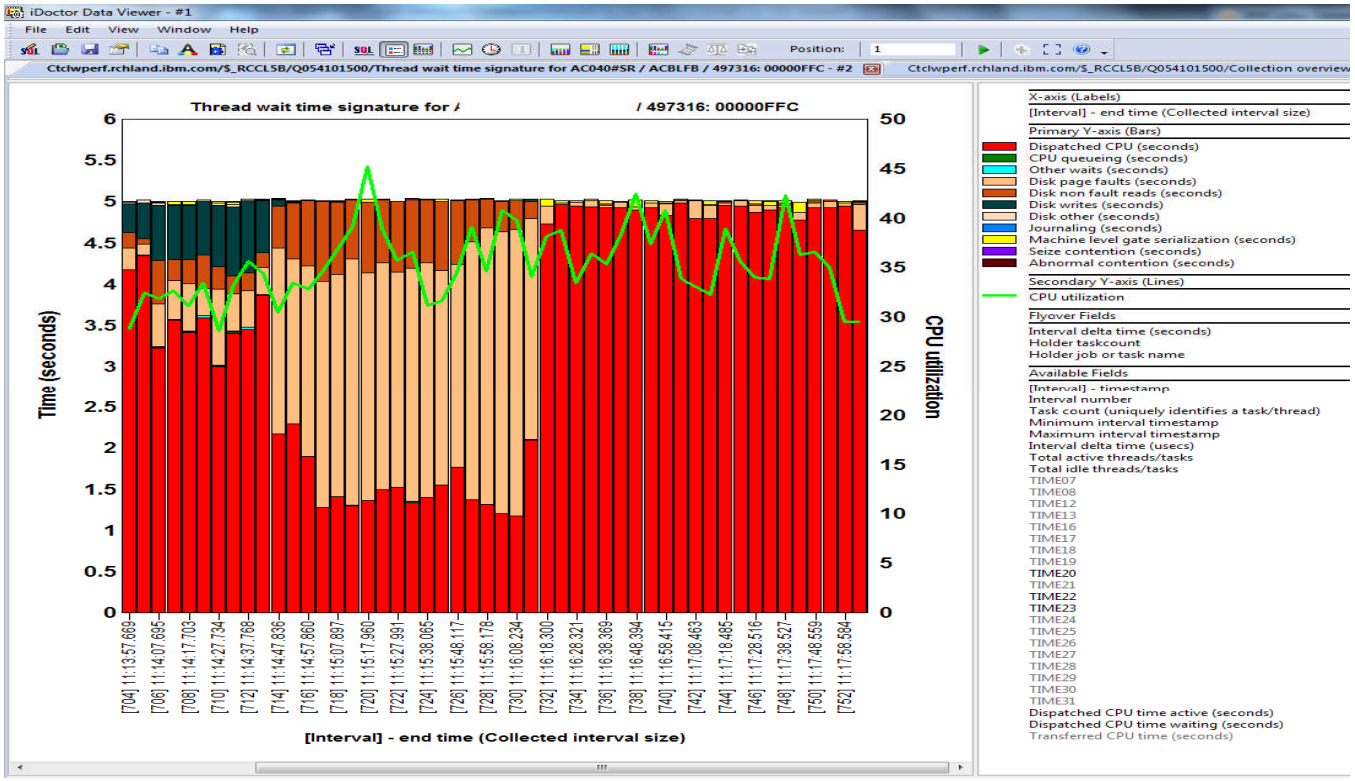

Identify program causing page faults

- The user program here is calling an SQL statement that then calls the QDBGETMQO (table or index scan) MI program

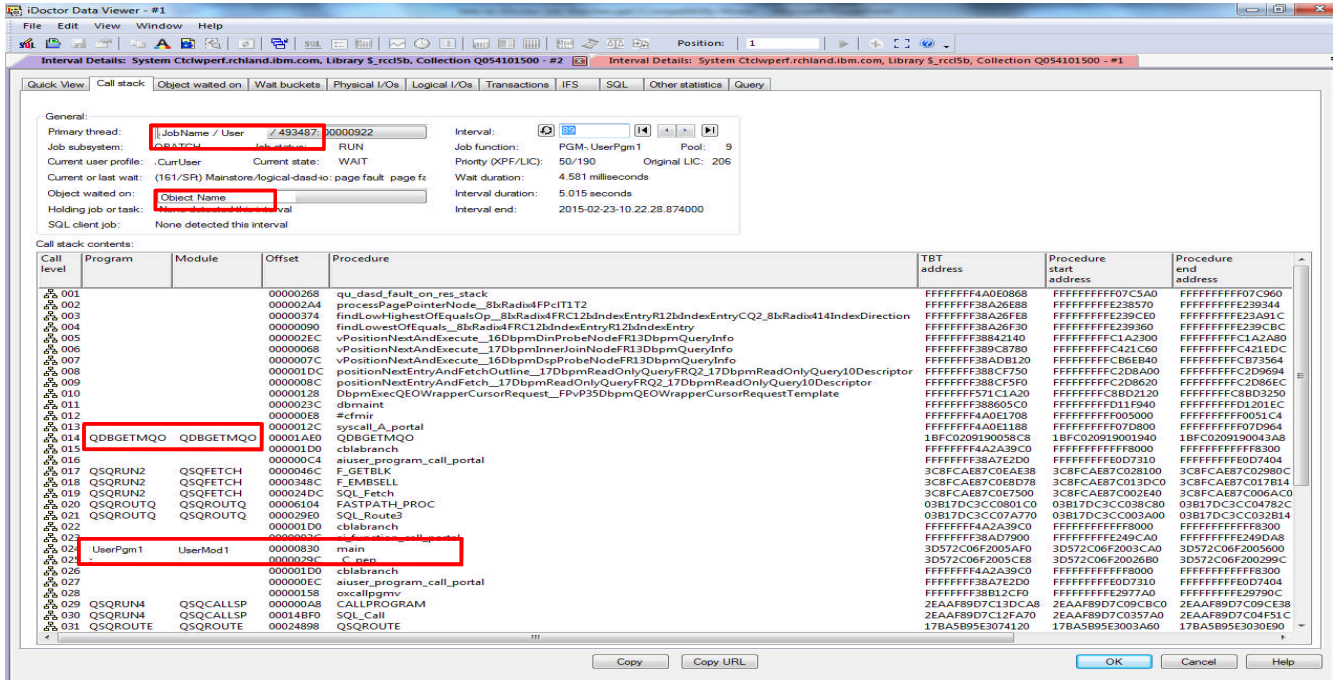

Identify the SQL statement causing high page faults

- The SQL tab will often show which statement was run to cause the faulting

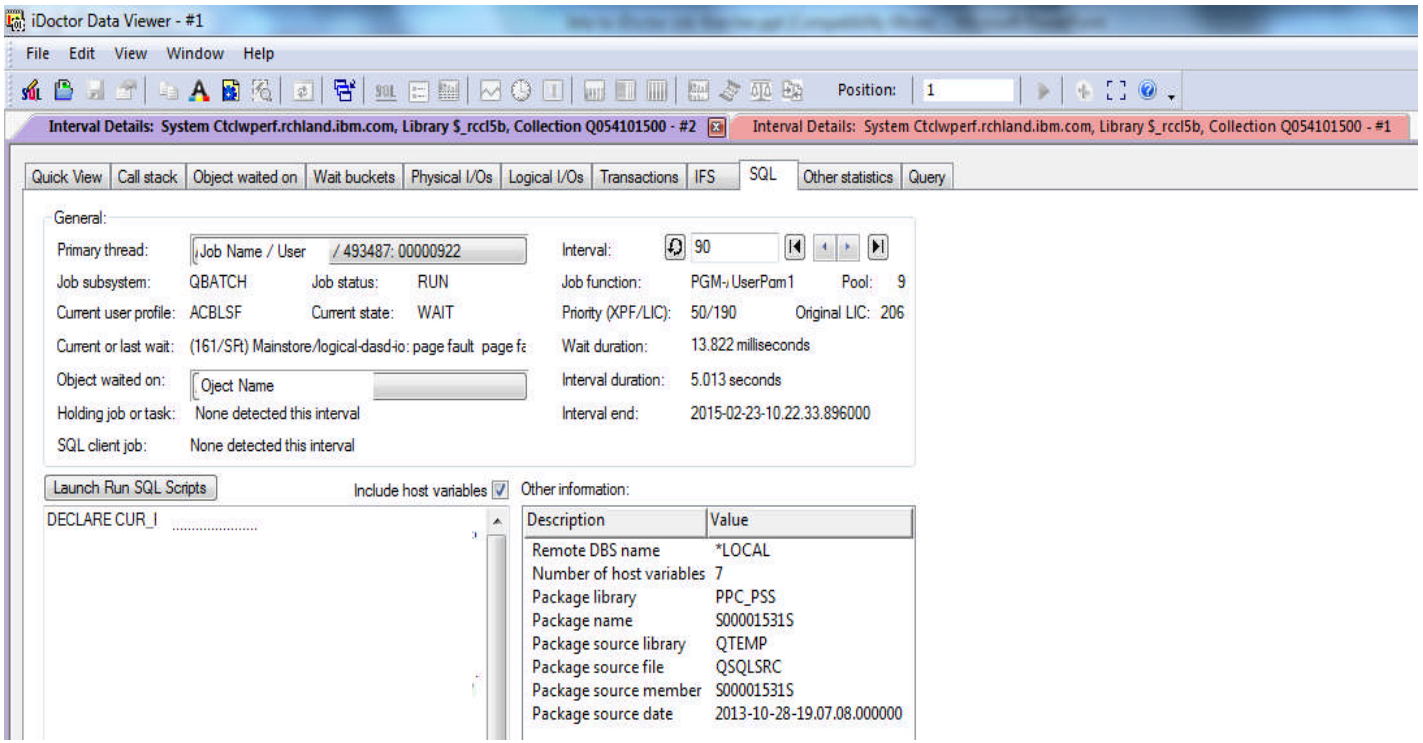

Job Watcher Memory Analysis Summary

- We can tell with more granularity what CS data showed us:
  - The job waiting the most on disk faults
  - The type of faults (DB vs. non-DB)
  - The average wait time on faults
  - The pool the job is running in
  - Etc.
- Additionally, we now know:
  - The object being faulted on
  - The program running
  - The OS operation causing the faulting
  - May also get SQL statement and host variables if running SQL

Navigator SQL Plan Cache Snapshot Analysis

- Run a plan cache snapshot against the tables being faulted in the most

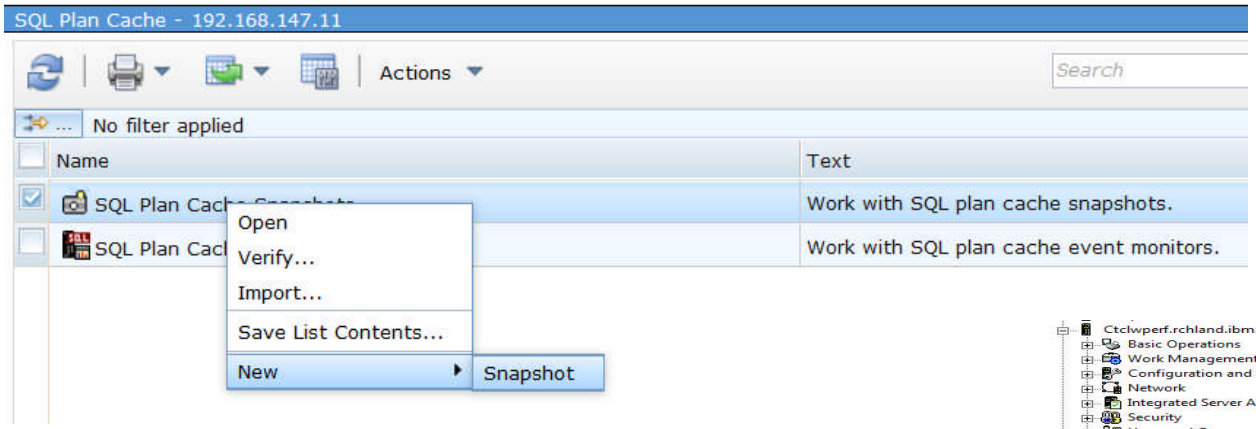

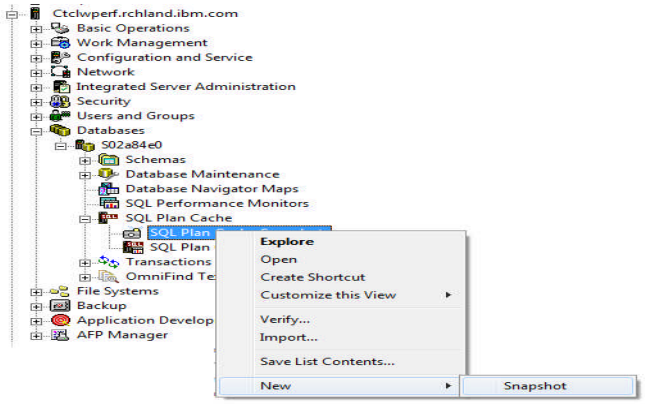

Plan Cache Snapshot Analysis

- Filter the plan cache snapshot by objects

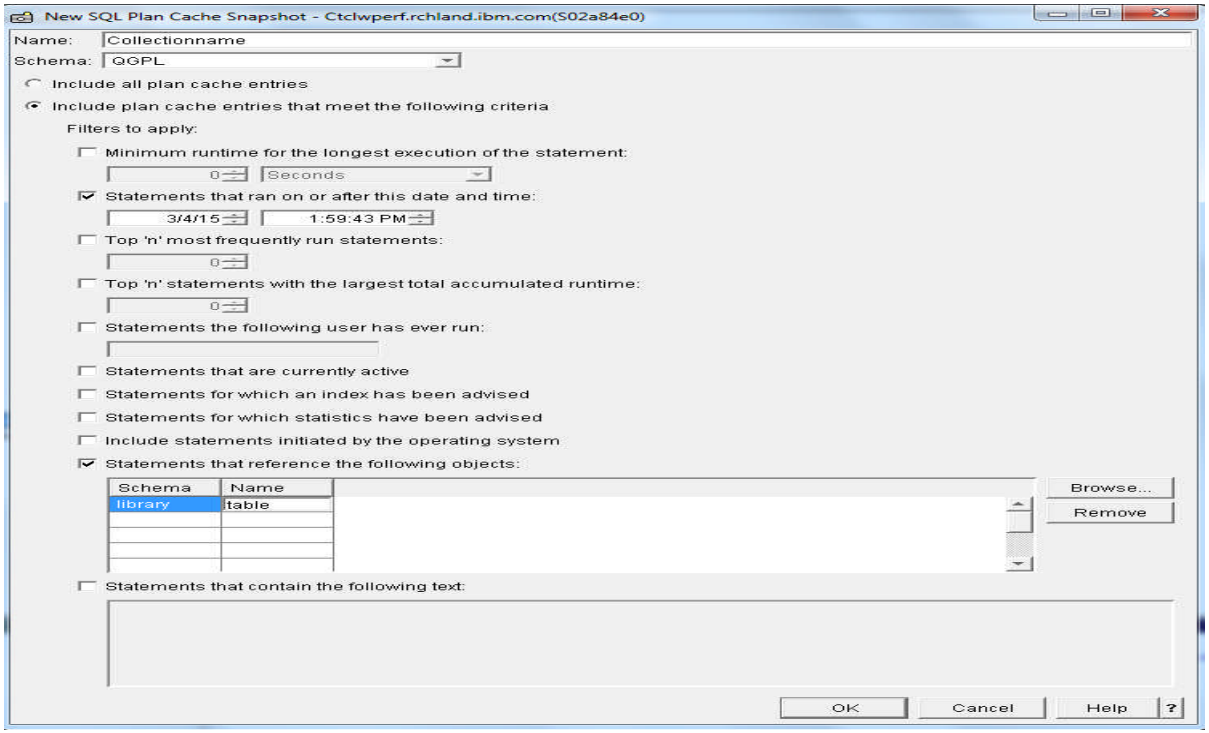

Data Management Best Practices for Memory

- Reduce file sizes
- Implement an index strategy
- Review index page size
- Remove deleted records (RGZPFM command)
- Purge historical data
- Consider adding SSD arms for those objects that have to be faulted in
- Avoid making any pool with activity too small – the faulting caused additional I/O, which in turn can affect response times for other jobs

IBM i Memory Tuning Best Practices - Summary

- Verify the memory configuration
  - Memory per processor
  - Processor affinity score
  - Reduce IO time by removing I/O bottlenecks
- Verify the pools are tuned correctly
  - Min/max pool size set up
  - Jobs running in the right pool
  - IBM i performance features are being used
- Tune application data access
- Optimize SQL queries
- Verify files size are right for an OLTP environment
  - Remove deleted records
  - Purge historical data

Questions?

References

IBM i Performance FAQ a MUST read!

October 2017 update *(watch for a Spring 2018 soon!):*

https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=POW03102USEN

IBM Power Systems Performance

IBM i on Power - Performance FAQ

October 9, 2017

IBM i Web Sites with Performance Information

- IBM Knowledge Center:
  - 7.2 Performance
  - 7.3 Performance
- IBM i Performance Management:
  i Performance Management
- developerWorks:
  - IBM i Performance Tools: developerWorks Performance Tools
  - IBM i Performance Data Investigator: developerWorks PDI
- IBM iDoctor for IBM i: iDoctor
- IBM i Wait Accounting information:
  - Job Waits Whitepaper
  - KnowledgeCenter: The basics of Wait Accounting
  - developerWorks: IBM i Wait Accounting

A Redbooks publication!

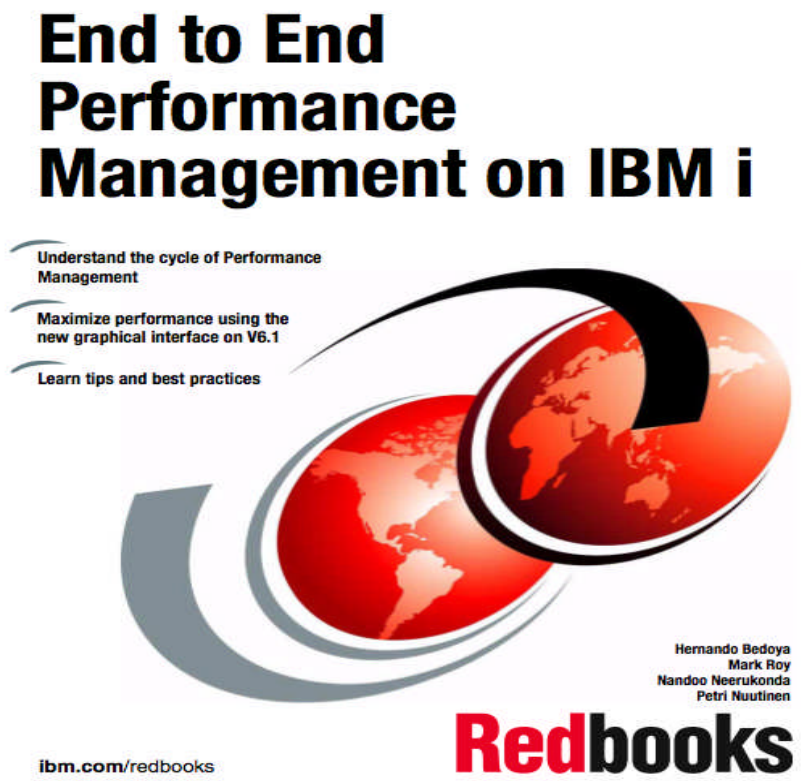

http://www.redbooks.ibm.com/redbooks/pdfs/sg247808.pdf

IBM i 7.2 Technology Refresh Updates

Covers the 7.2 content through Technology Refresh 1

Section 2.8 – Performance

Section 8.6.7 – Job level SQL stats in Collection Services

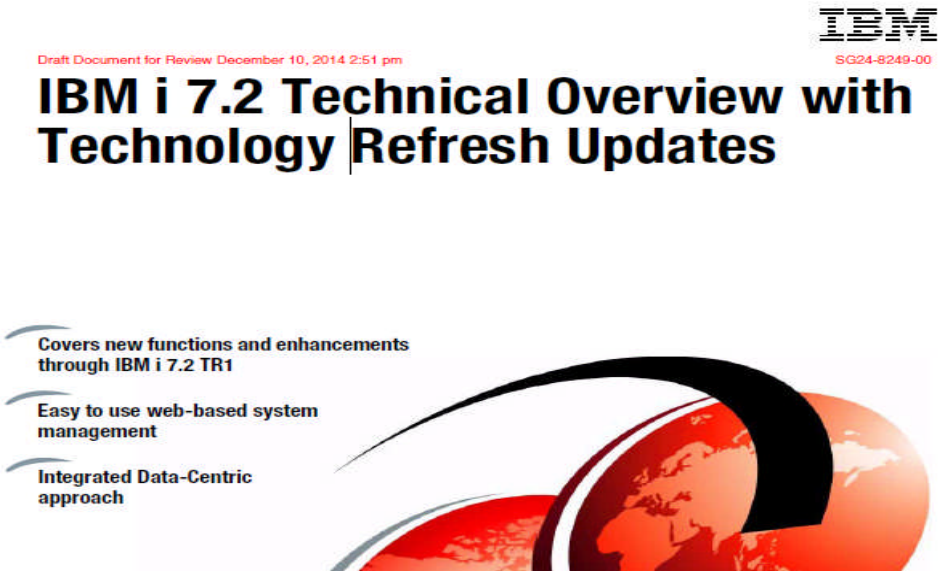

IBM i Performance Analysis Workshop

Learn the science and art of performance analysis, methodology and problem solving

Managing and analyzing the data can be quite complex. During this workshop, the IBM Systems Lab Services IBM i team will share useful techniques for analyzing performance data on key IBM i resources, and will cover strategies for solving performance problems. It will aid in building a future foundation of performance methodology you can apply in your environment.

Overview:

- Topics covered include:
  - Key performance analysis concepts
  - Performance tools
  - Performance data collectors (Collection Services, Job Watcher, Disk Watcher, and Performance Explorer)
  - Wait accounting
- Core methodology and analysis of:
  - Locks
  - Memory
  - I/O subsystem
  - CPU
- Concept reinforcement through case studies and lab exercises
- Discussions on theory, problem solving, prevention and best practices

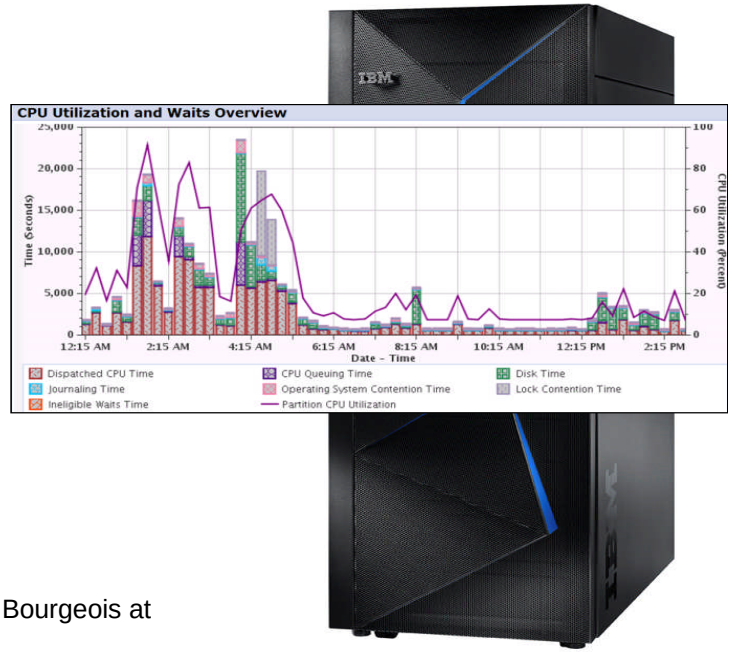

Workshop details:

- Intermediate IBM i skill level
- 3-4 day workshop, public or private (on-site)
  - For general public workshop availability and enrollment: IBM i Performance Analysis Workshop
  - For public workshop availability and enrollment in France, please contact Philippe Bourgeois at pbourgeois@fr.ibm.com or Françoise Laurens at f_laurens@fr.ibm.com
- For additional information, including private workshops, please contact Eric Barsness at ericbar@us.ibm.com or Stacy Benfield at stacylb@us.ibm.com, members of Systems Lab Services

IBM Systems Lab Services Power Systems Delivery Practice - **ibm**.com/systems/services/labservices - ibmsls@us.ibm.com

IBM i Performance and Optimization Services

The IBM i Performance and Optimization team specializes in resolving a wide variety of performance problems. Our team of experts can help you tune your partition and applications, including:

- Reducing batch processing times
- Resolving SQL query and native IO performance problems
- Tuning RPG, COBOL, C, and Java (including WebSphere Application Server) programs
- Removing bottlenecks, resolving intermittent issues
- Resolving memory leaks, temporary storage growth problems, etc.
- Tuning memory pools, disk subsystems, system values, and LPAR settings for best performance
- Optimizing Solid State Drive (SSD) performance
- Tuning client interfaces such as ODBC, JDBC, .Net and more

Skills transfer and training for performance tools and analysis also available!

Contact Eric Barsness at ericbar@us.ibm.com for more details.

www.ibm.com/systems/services/labservices

And finally…..

Thank you

Don't forget to fill-in the feedback form!

ithankyou

www.ibm.com/power/i

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)

End of Presentation material…..