What's New with iDoctor Feb-July 2012

(iDoctor GUI Builds 903-949)

Ron McCargar
IBM i Global Support Center

Agenda

- **Overview**
  - Goals, components, unique features and website.
- **What's new with iDoctor (Feb 2012 – July 2012):**
  - Note: As of August 3rd, 2012 some of these functions described are not yet available to customers but will be once the next builds are released.
- **Questions?**
  - Contact idoctor@us.ibm.com or mccargar@us.ibm.com

Overview

- **iDoctor is a suite of dynamic performance tools offered by the Global Support Center.**
- **We cover all areas of performance but historically focused more on low-level details.**
- **Started in V4R5 with the PEX GUI plug-in for Operations Navigator it now consists of 5 external and 5+ IBM internal components.**

Goals of iDoctor

- **Broaden the user base for Performance Investigation**
  - enable Operators, Programmers, IS Management
  - as well as Performance Specialists, Consultants
- **Simplify and automate processes**
- **Provide quick, immediate access to collected data**
- **Provide more analysis options**
- **Reduce the dependency on PEX traces**

iDoctor Functional Areas

Functional Area	Performance Tool
High-level system/job monitoring	**iDoctor – Collection Services Investigator** IBM i Performance Tools (PT1) Management Central Monitors
Medium-level system/job monitoring	**iDoctor – Job Watcher** WRKSYSACT IBM i Job Watcher (PT1) / STRJW
Low-level system/job tracing, stats, profiling	**iDoctor – PEX Analyzer** IBM i Performance EXplorer / PRTPEXRPT
Disk stats/tracing	**iDoctor – Disk Watcher** **iDoctor – Collection Services Investigator** **iDoctor – PEX Analyzer (PDIO analysis)** **iDoctor – VIOS Investigator (new in 2012)** IBM i Disk Watcher (PT1) / STRDW
Plan cache analysis	**iDoctor – Plan Cache Analyzer (new in 2012)**
Determine if SSDs could help performance	**iDoctor – Collection Services Investigator** SSD Analyzer Tool for IBM i

iDoctor GUI

- **It's a Windows client offering superior flexibility and functionality**
- **All components offer a similar user experience**
- **The latest GUI builds provide access for iDoctor components installed on servers running IBM i V5R4 or higher.**
  - Note: Latest GUI builds probably won't work well with V5R3 systems.
- **Requirements:**
  - System i Access for Windows
  - Trial or License keys for Job Watcher (includes DW, CSI) and PEX Analyzer component usage.

iDoctor Resources

- **iDoctor e-mail list: usage tips, build updates, PTF info, etc**
  Send join requests to mccargar@us.ibm.com

- **iDoctor Website:**
  http://www-912.ibm.com/i_dir/idoctor.nsf/

- **Presentations (What's New, etc):**
  http://www-912.ibm.com/i_dir/idoctor.nsf/downloadsDemos.html

- **YouTube Channel (20+ videos):**
  - http://www.youtube.com/user/IBMiDoctor?feature=mhum
    - These videos are also available on IBM.COM if your company blocks YouTube.
    - Just click the appropriate links titled "Video name on IBM.COM" from the Video Library pages on our website:
      - https://www-912.ibm.com/i_dir/idoctor.nsf/videos.html

- **iDoctor Forum:**
  http://www.ibm.com/developerworks/forums/forum.jspa?forumID=871

- **Documentation:**
  https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf

iDoctor Team

- **Larry Cravens**
  - Performance Expert, Taskswitch, Wait Bucket Analysis
- **Tom McBride**
  - VIOS/External Storage, Strategic Direction
- **Ron McCargar**
  - GUI, Website, Monitors, Builds, Documentation, Videos
- **Brad Menges**
  - Education (on-site visits), Sales, Project Manager
- **Chris Pilcher (Business Manager)**
- **Lab Services (Sales Contacts)**
  - Karen Anderson, Mark Even (USA)
  - Jin-Ming Liu (AP)
  - Virginie Cohen, Gerard Barneaud (EMEA)

What's New with iDoctor: Overview

- **General GUI Enhancements**
- **Job Watcher**
- **Collection Services Investigator**
- **PEX Analyzer**

General – New Components Added in May 2012

Plan Cache Analyzer is now included with a Job Watcher license.
Note: Only snapshot support within this component is provided with the customer version.

VIOS Investigator (previously known as NMON Investigator) is now available and is presently a free option.

The Documentation also has new chapters available on these components.

If connected to the IBM i, the iDoctor FTP GUI component is now also included in the list of components. Additionally the functionality within the IBM internal component formerly known as Object Explorer is now included in the iDoctor FTP GUI under the Libraries folder within the iDoctor FTP GUI.

General – VIOS Investigator Notes

The initial release of VIOS Investigator features a GUI that analyses the data using DB2 on the IBM i.

A future release will allow several functions for collecting data (and analyzing) directly against one (or more) VIOS.

We intend to support PerfPMR data collection (and send to IBM), NMON collection, and NPIV/FC statistics collection using fcstat –n.

These latest changes will not be covered in this set of slides and will instead be covered once these functions have been released.

General – Shortcuts

- **You can now save a shortcut file (*.idr) from a Data Viewer graph or table using the Create Shortcut menu.**
- **You can send this to other users or save it on your PC to revisit this report later.**
- **Double-click the file to launch iDoctor and the report(s) saved.**
- **Especially useful for customized drill-down reports to get you back to exactly where you were before!**

General – iDoctor URLs

- **Allow you to save a URL for things in iDoctor, and revisit it later or send it to another user (libraries, collections, call stacks, etc)**
- **Use the Copy URL menu or button to copy the URL to the clipboard.**
- **Then paste this into a web browser or save it for future use to revisit the library/collection/call stack, etc.**
- **Example (opens library COMMON within Job Watcher on system idoc610):**
  - idoctor:///viewinfo1[type=CFolderLib,sys=Idoc610,lib=COMMON,comp=JW]

General – Batch Window Views

Jobs in ranking graphs can now be graphed in a “batch window” type of view using Gantt Chart support.
You can find additional details on this in the Update History.

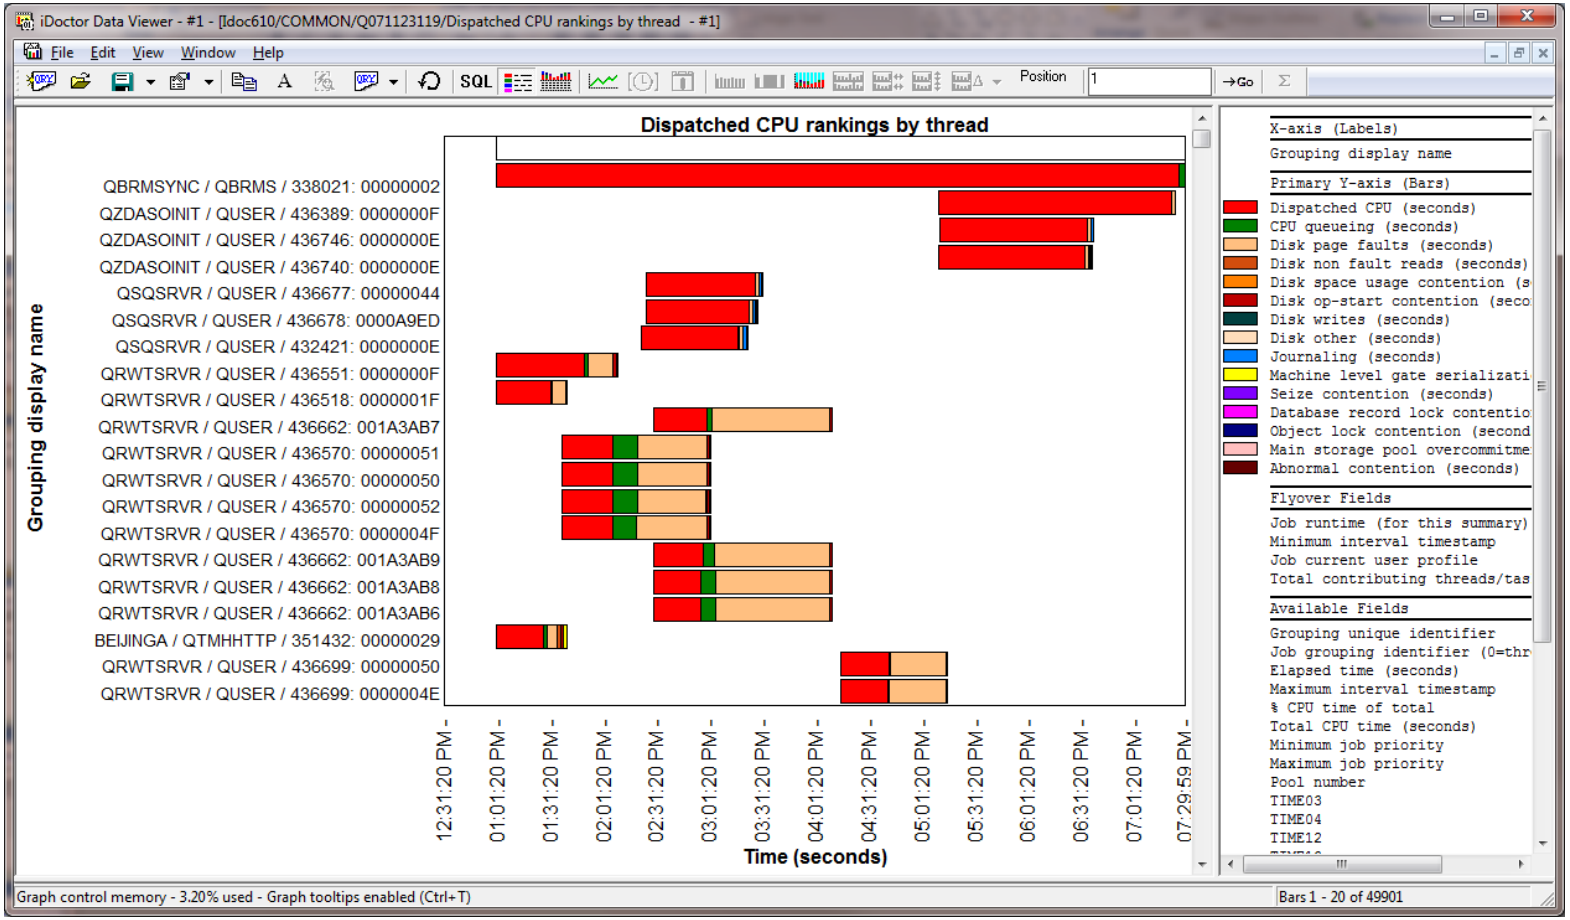

General – Updated all analyses to generate primary keys on tables

All iDoctor analyses were updated so that each SQL table generated contains a primary key. This will ensure tables have a valid sort order if opened with using an order by.

General – SQL Editor will now show Stored procedure result sets

Calling a stored procedure in an iDoctor SQL editor that returns a result set will now display the result set in the table view.

An example of doing this is

CALL QIDRCSGRI('LIBNAME')
where LIBNAME is a library that contains collection services data.

General – Copy collection

In all components that support copying a collection, you can now specify the collection name in the target library. This also lets you copy a collection to a new name in the same library.

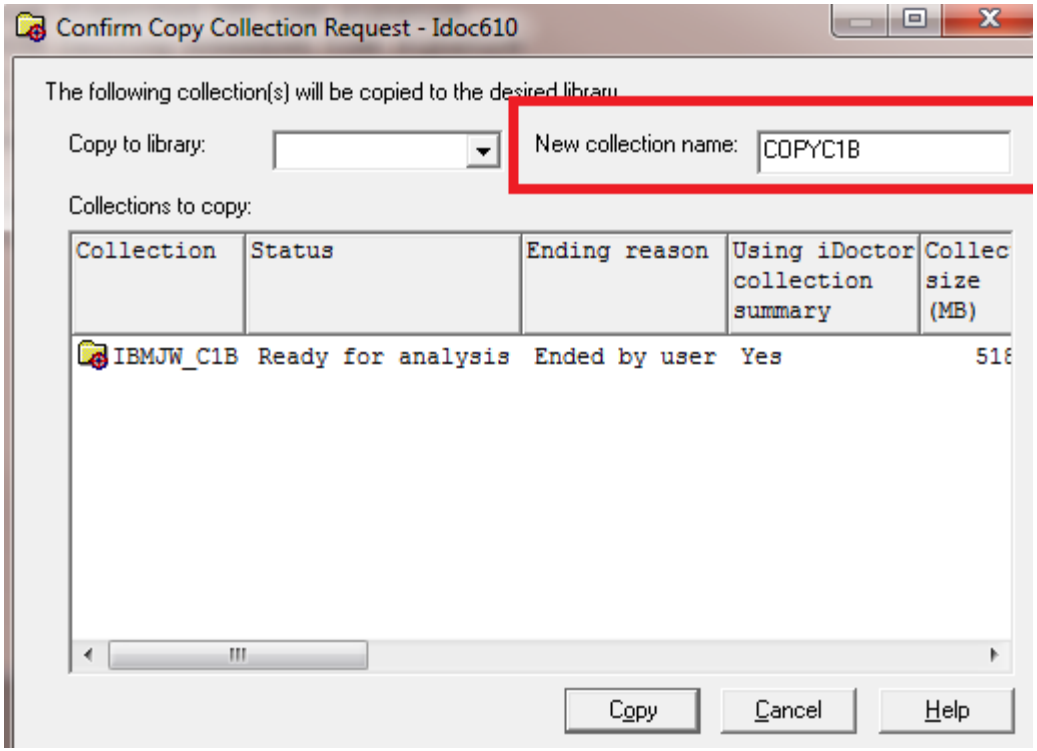

General – Report Generator Changes

1. Added a debug option (only shows reports with SQL errors.)

2. Added checks to ensure no popup messages are shown when running. If graphs normally provide popup selections, default values will be taken.

3. Added additional information about the collection at the top of the report: created on system and VRM and the collection start and end times.

General – Firmware level check for Power 7

In PEX Analyzer and VIOS Investigator added a check for P7 firmware levels being at least Ax730_078 to fix a potentially serious problem when collecting a TPROF.

General – Install option for QIDOCTOR

In the installation, added an option on the Component Selection window called "Create user profile QIDOCTOR (applies to Base support only)" with the default of unchecked.

This will cause the QIDOCTOR profile to be created if checked only. Some users prefer to run collections under this profile (via the Submit job options window) , but others may not want additional profiles added to their systems.

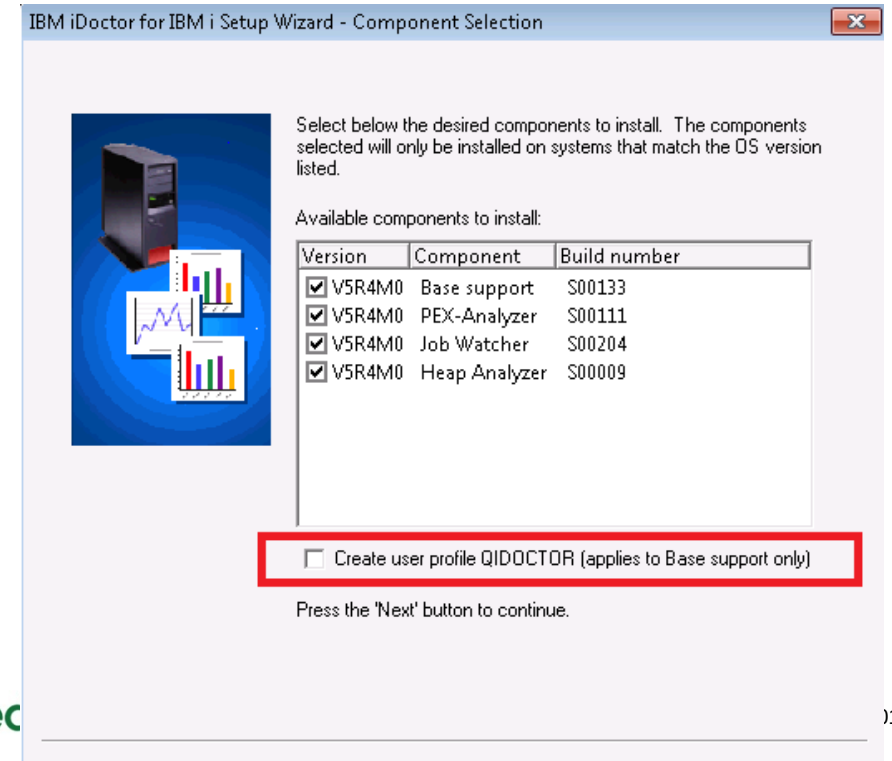

General - Change SQL Parameters – Parameter Prompting

If opening a graph where a parameter marker has not yet been resolved in the SQL Statement, you will now be prompted using the Change SQL Parameters Window for the desired value.

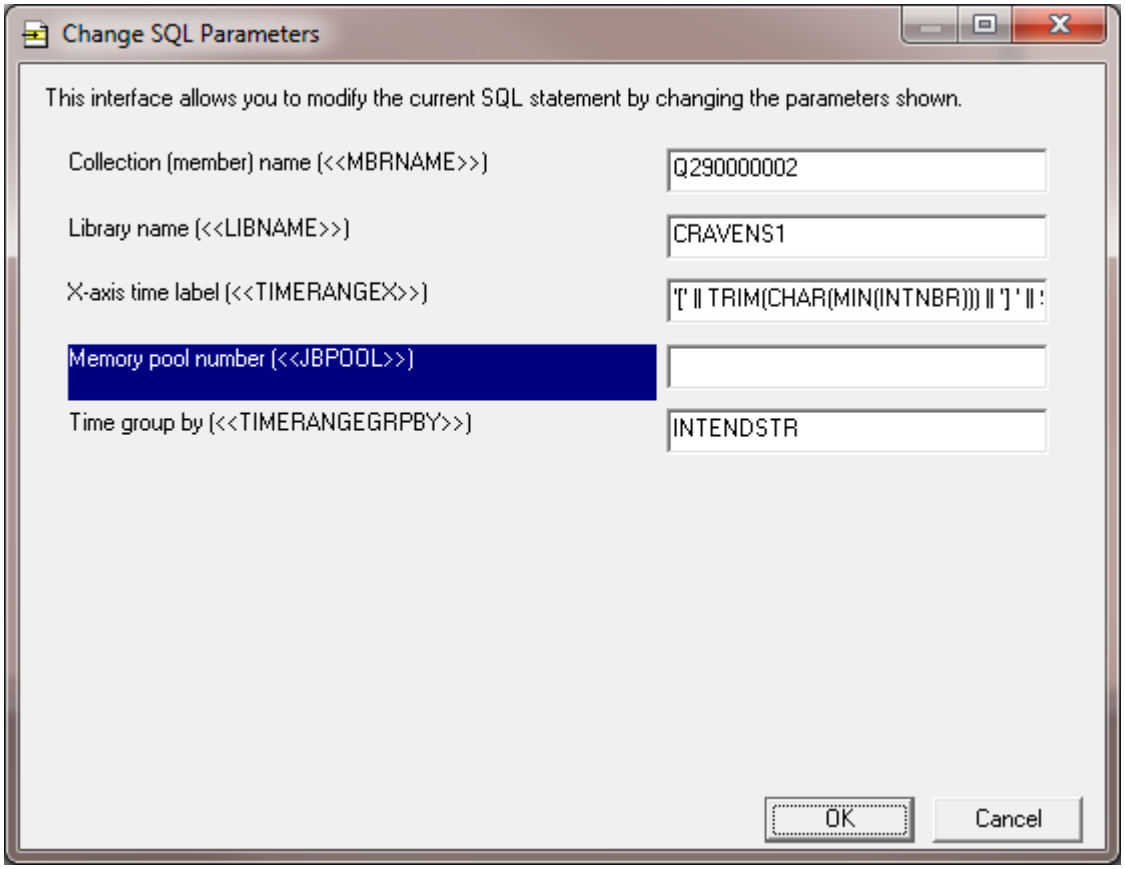

General – Install failure of “Base Support” or QIDRGUI library

In the installation of library QIDRGUI (base support), on some non-English systems the call to program CRTSQLPGMs would fail.

Ultimately this seemed to because of the SQL being unable to compile stored procedures containing this line:
DECLARE CMDLEN DECIMAL(15,5);

Changed several stored procedures to work around this issue.

Please Note:

1. The iDoctor tools are still usable to despite the error. The Stored Procedures will be loaded instead the 1st time the GUI is used.

2. This has already been fixed but as of Aug 3rd, 2012 not yet released on our website.

General – Database Redesign

The iDoctor databases used to store query and graph definitions underwent a major redesign in order to reduce redundancy and improve maintenance times. As a result some colors and descriptions shown in the reports may have changed.

Please report anything that looks incorrect to us for review.

See the Update History for more information.

Job Watcher Updates

JW – Interval Details – Holder’s client job

In Job Watcher at 6.1+, in the interval details interface (call stack/etc), if the holder job has an SQL client job identified, then it will now be shown as the "Holder's client job" saving the user an extra step to determine that.

JW – Selected job grouping over time graphs

Added the Total Contributing threads/tasks on secondary Y axis for the selected job grouping over time wait bucket graphs in Job Watcher. (for example "Generic job wait time signature for QZDA*")

JW – Interval Details - Stepping through intervals

In Job Watcher in the interval details interface (call stack/etc), when moving forward through the intervals if the job goes into an idle wait then all previous information shown (holder, wait object, etc) will be retained while stepping forward through the idle intervals.

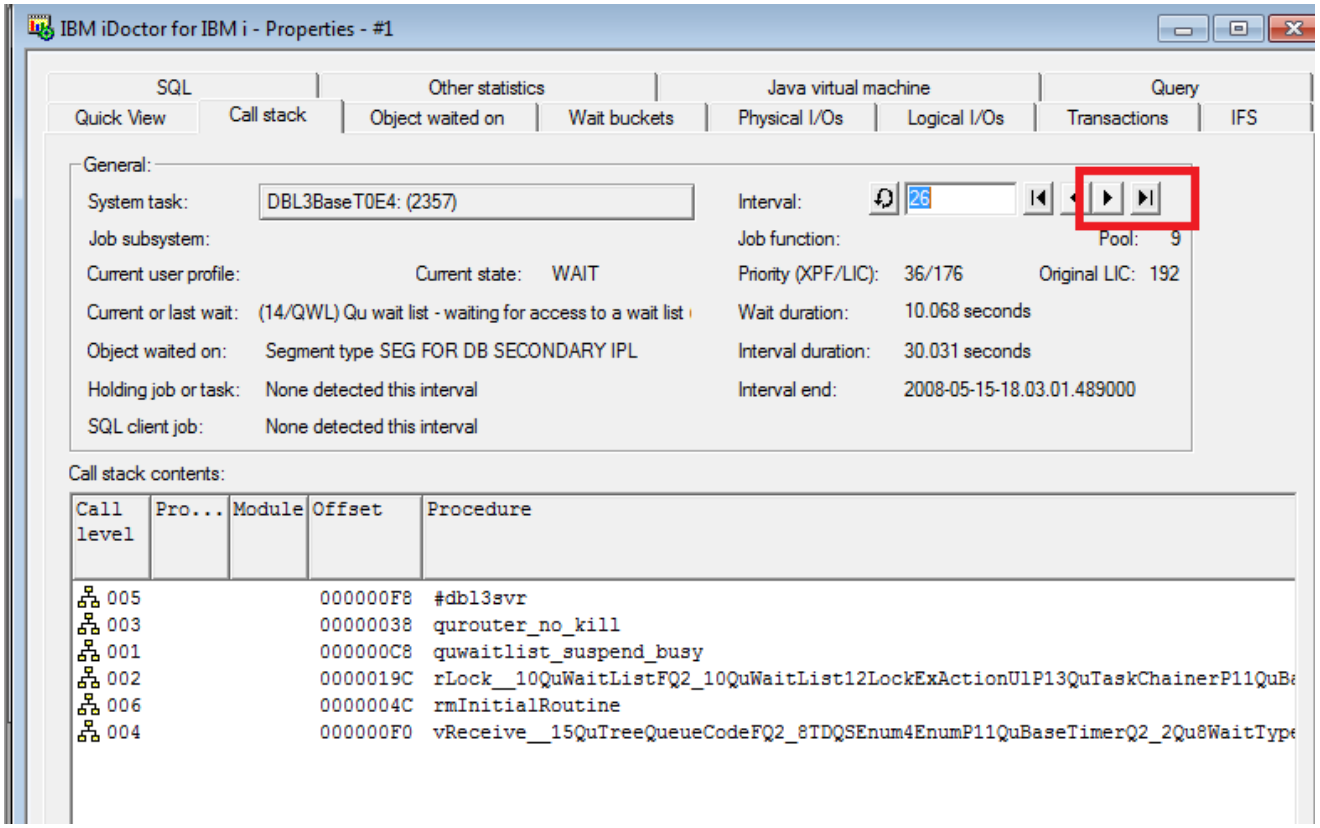

JW – Include segments preference

Added a new Job Watcher preference "Include segments when viewing objects waited on graphs and reports" with a default value of unchecked.

If checked then segments will be included by default in the objects waited on graphs and the interval summary -> objects waited on interface.

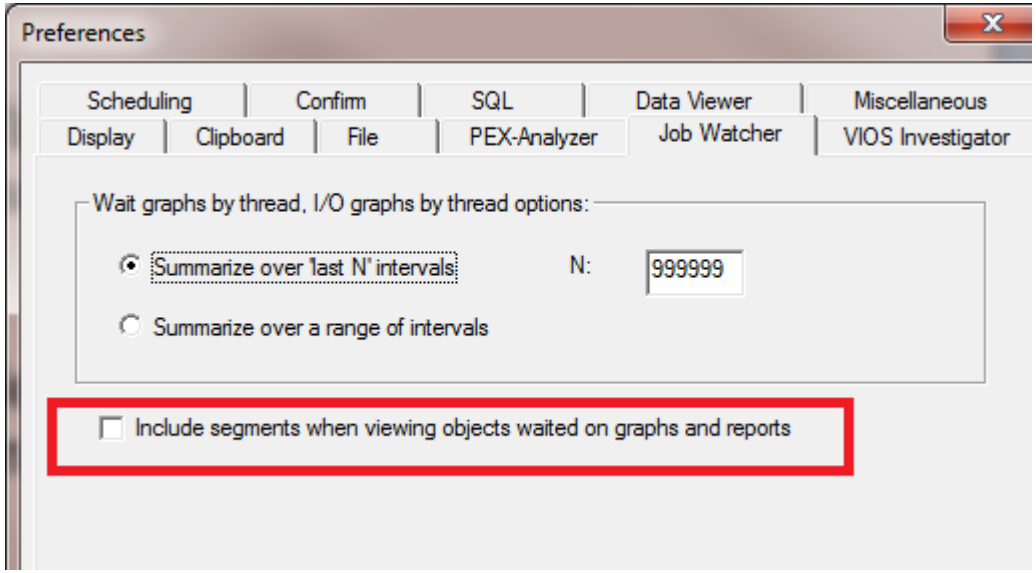

JW - Single interval rankings identifying flags

In Job Watcher, if drilling down into a wait bucket rankings graph on a single (job/thread) interval added a new flags field at the end of the job/thread name that will contain the following possible values and meanings:

W = has a wait object
H = holder
B = current wait bucket is the same as current sort /filter bucket
S = has an SQL client job (applies to 6.1 and higher only)

JW - Single interval rankings identifying flags (example)

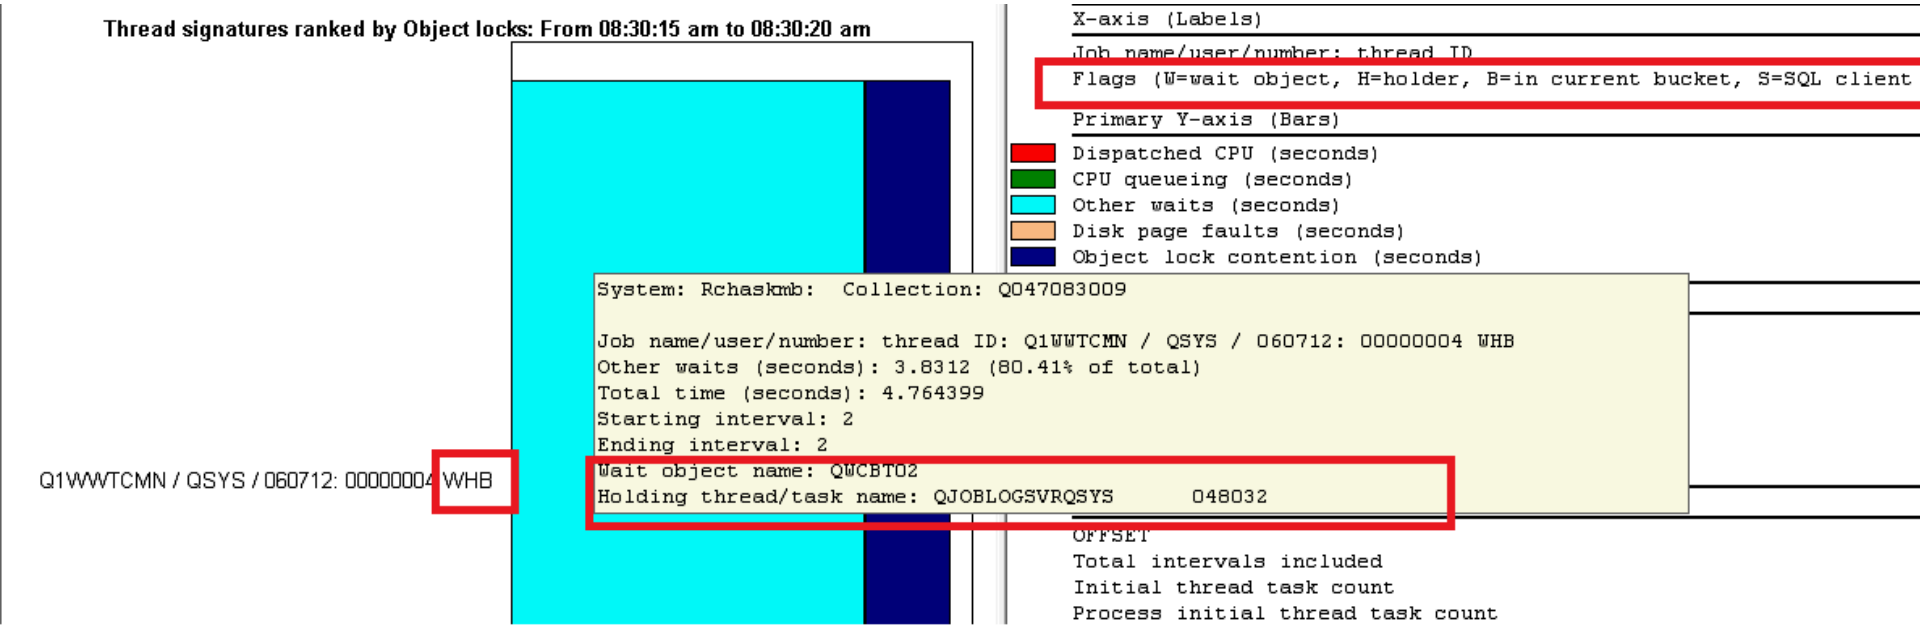

Power is performance redefined

JW – Objects waited on having <<BUCKEDESC>> current wait for thread XYZ

Added a new Objects waited on graph for a single job/thread that filters the results based on the selected wait bucket.

A 2nd option focusing on just disk reads was also added called "Objects waited on having disk reads current wait".

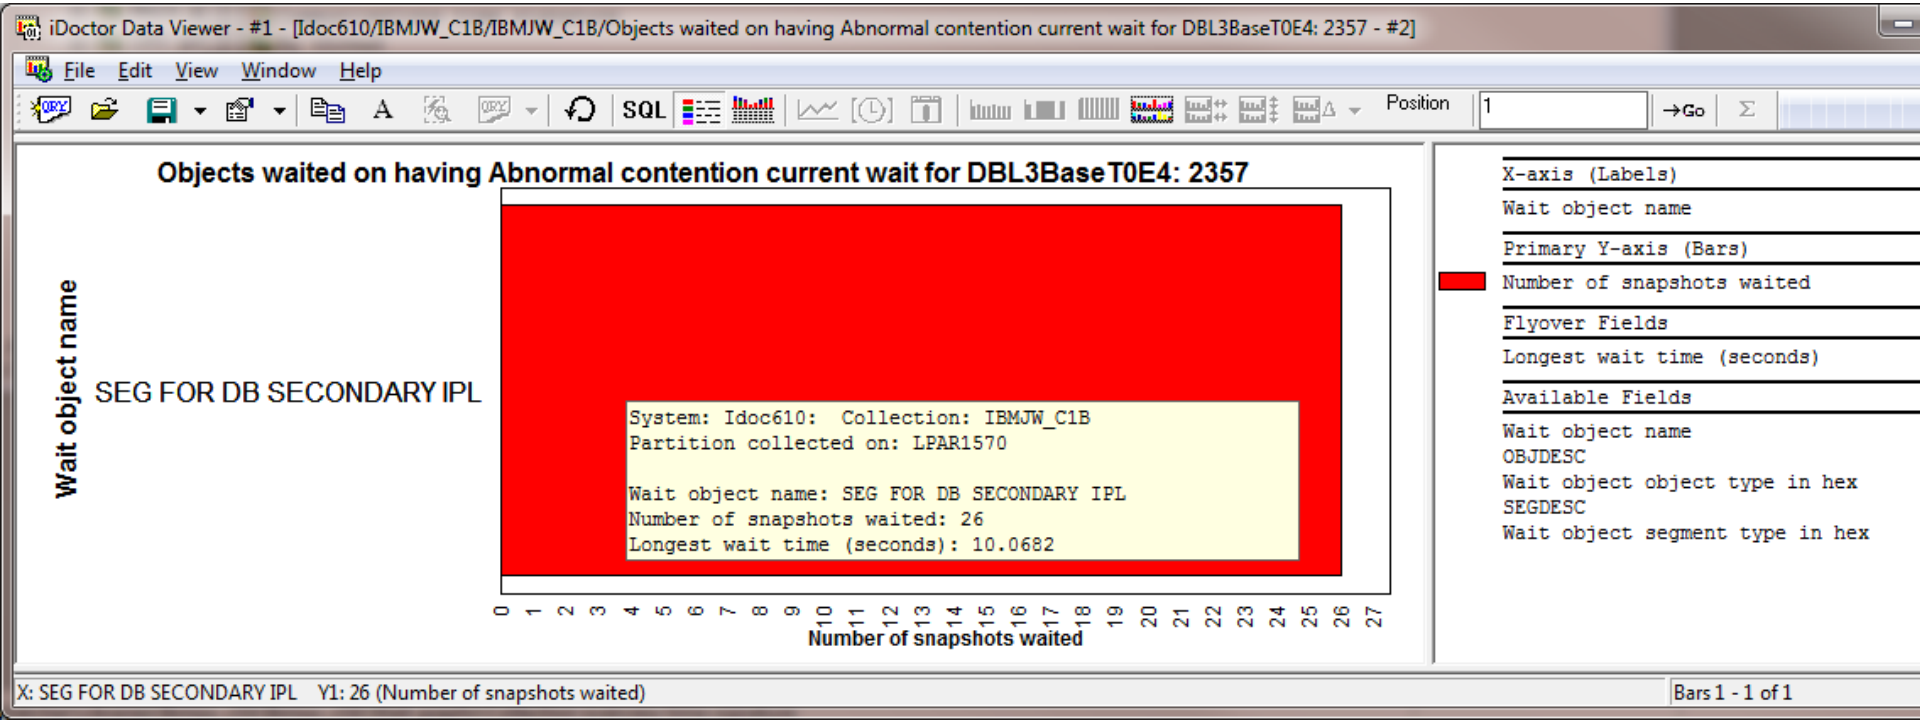

JW – New Situation – Large number of locks

Added a new situation to Job Watcher called "Potentially large number of locks.“

If this is found the recommendation is to review the “Wait graphs -> seizes and locks graphs” in Collection Services Investigator for the same time period to see the jobs and threads experiencing these locks.

JW – New Situation – Deadlock due to DB record locks

Added a new situation to Job Watcher to look for a possible two job deadlock due to DB record locks.

If this is found the recommendation is to use DBMON for further investigation.

Collection Services Investigator Updates

CSI – Collection Description

You can now view and set a Collection Services' collection description using iDoctor. Some analyses will also have options to surface this in their reports to help when comparing data.

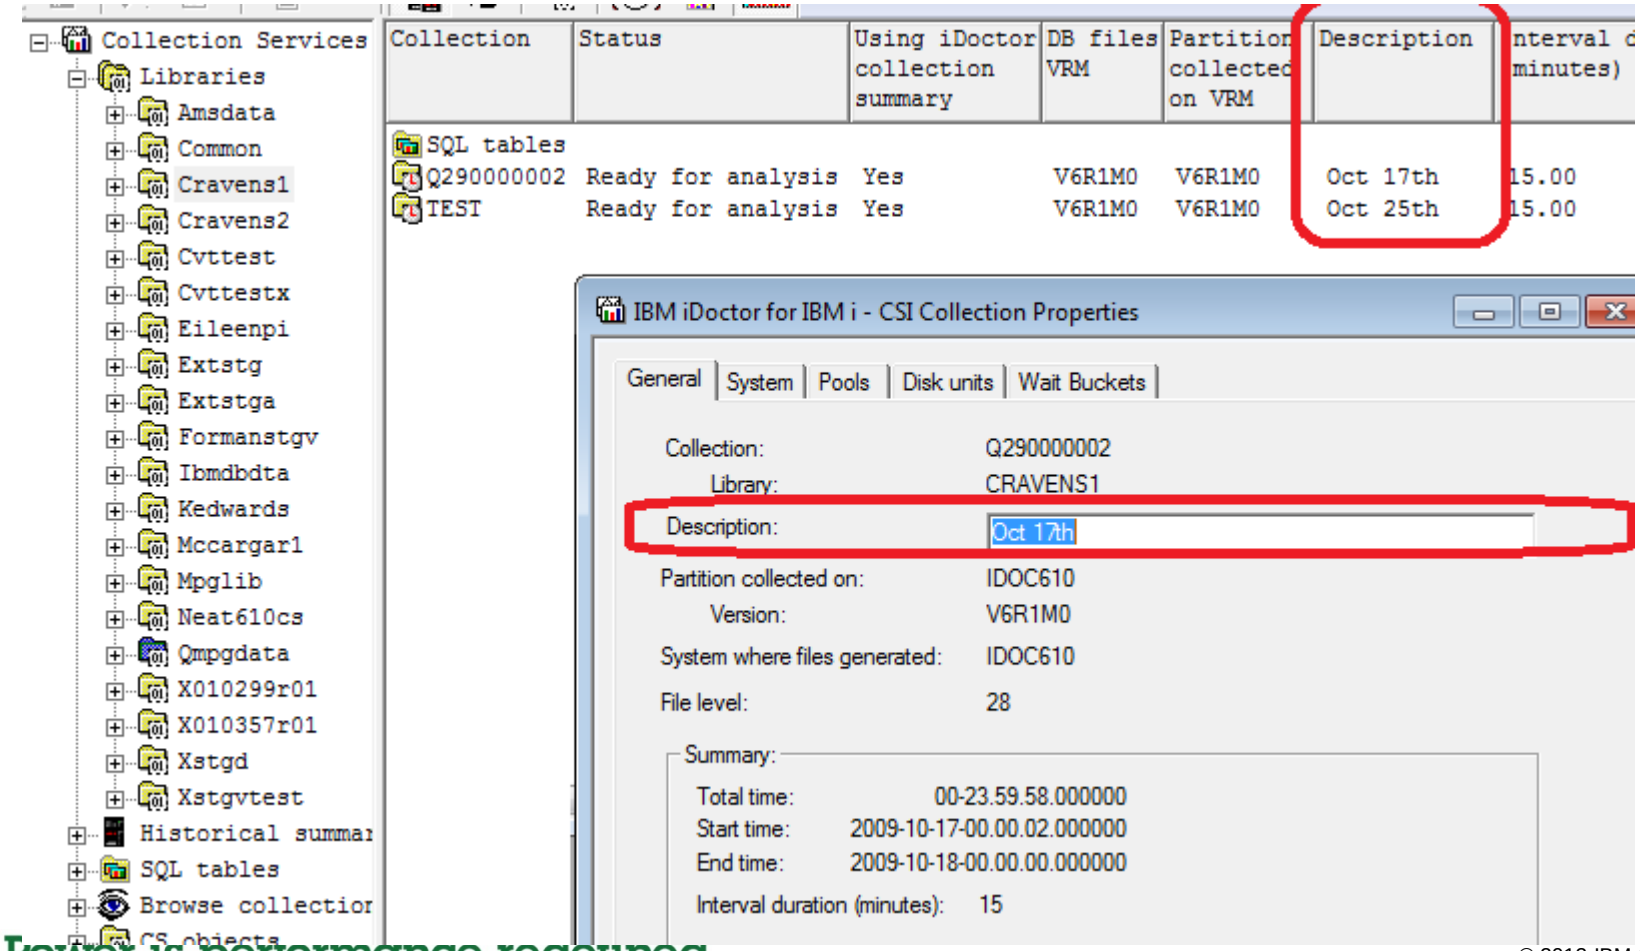

CSI – Advanced Job Summary

Created a new analysis to replace the Create Job Summary called Advanced Job Summary.

It allows multiple parameters for the job name, current user profile and subsystem. You can also filter by CPU time or job run time.

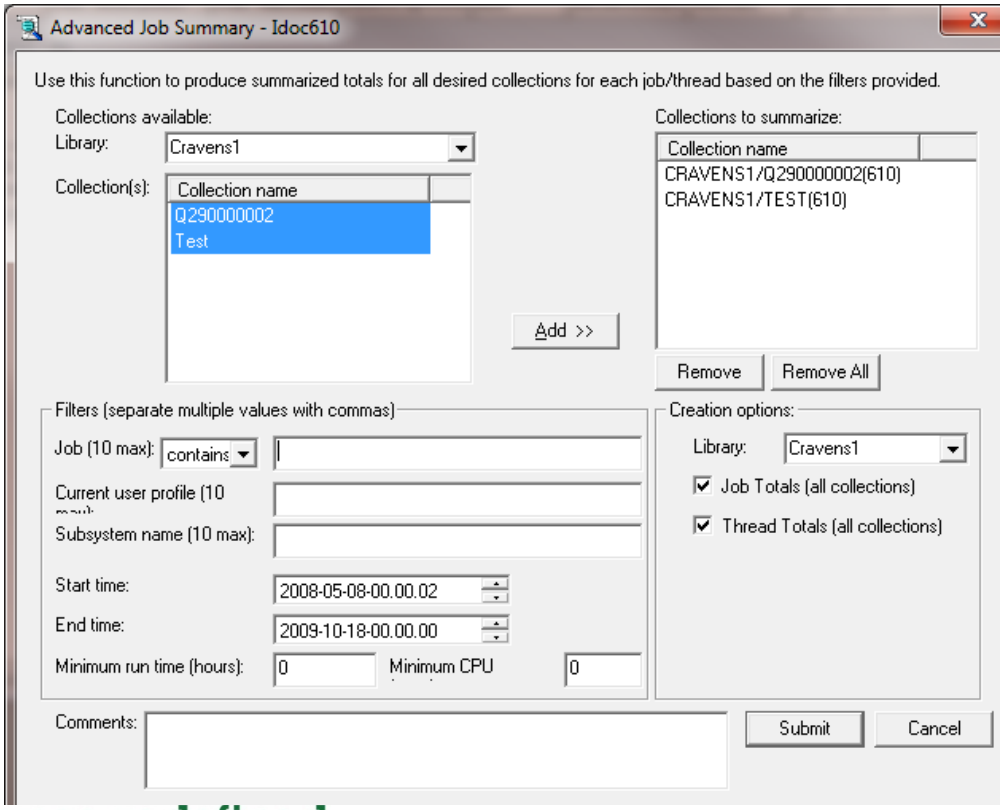

CSI – Advanced Job Summary

Graphing options now include a function to graph the data by generic job name. An option was also added to filter the data by generic job name.

Also changed the drill down graphs under the "Other graphs" menu to show counts/totals instead of rates in all applicable cases.

The graphs that show job/thread rankings by collection now includes the collection description in the list of available fields. This can be moved to the x-axis label if desired.

Note: In the future these same enhancements will be moved to Job Watcher.

CSI – 64K versus 4K page faults graphs

In CSI 6.1+,. under the memory pool graphs added a new graph called "64K versus 4K page faults for pool <<JBPOOL>>". When you open this graph you will be prompted to enter the desired pool number.

The graph can help with page thrashing issues on newer hardware related to 64K memory pages.

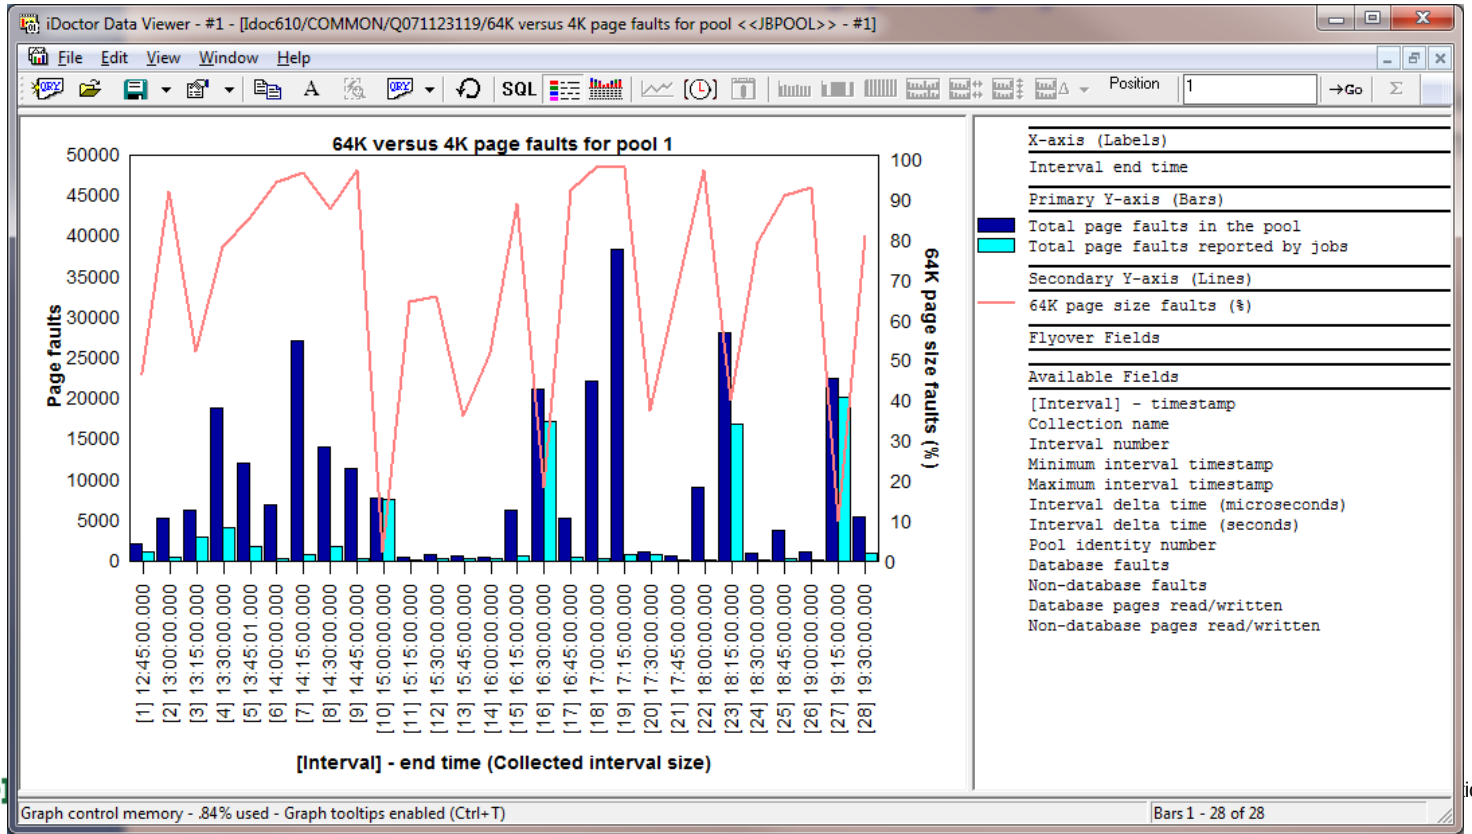

CSI – External Storage Links and Ranks graph changes

1. Now showing avg response time, avg read response time and avg write response time in all the graphs (some only showed avg response time.)

2. The selected links or ranks over time graphs now show side-by-side vertical bars instead of a "step" chart.

3. The selected links or ranks over time graphs are now flattened graphs (on INTERFACE_ID or RANKID).

CSI – Showing the default Collection Services Library

The library the collection services data is being created in is now shown in a different color.

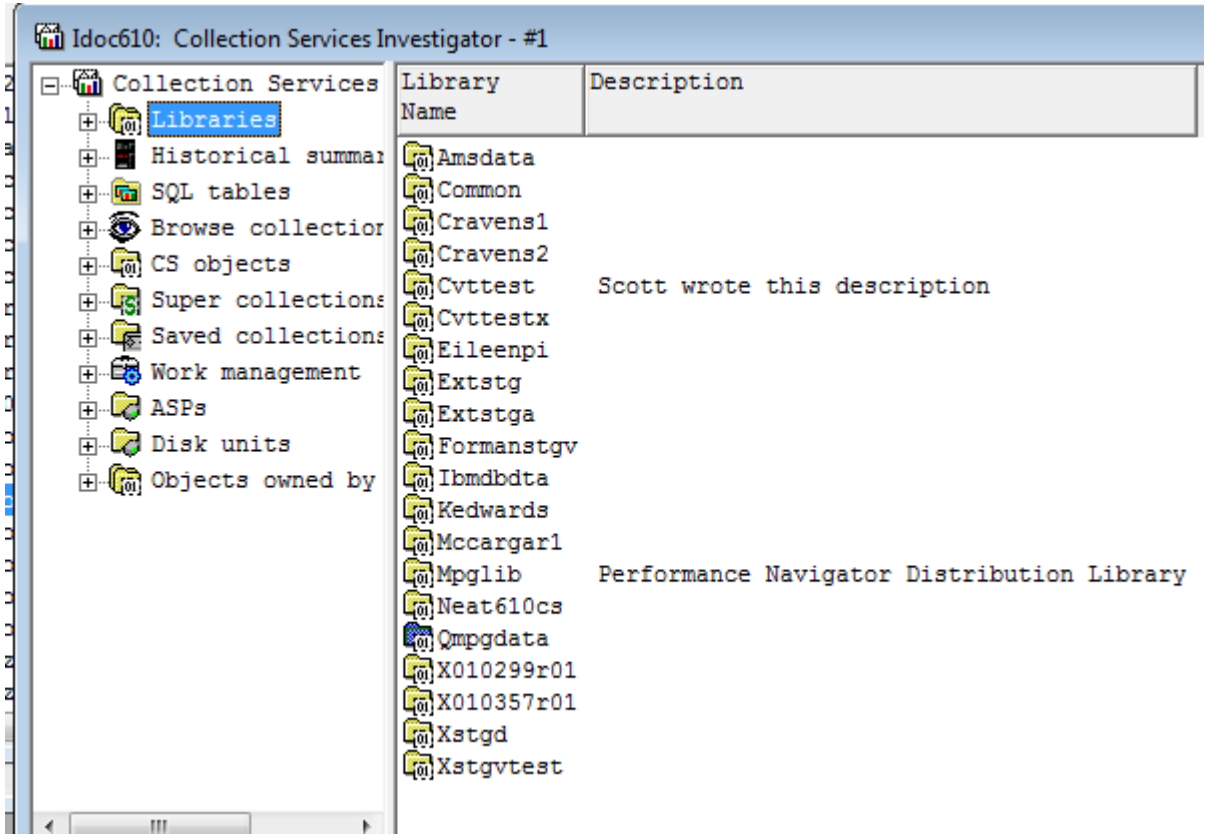

PEX Analyzer Updates

PEX Collection Split

Allows you to divide up a huge PEX collection into a more manageable size (if SQL queries are too slow against it.)

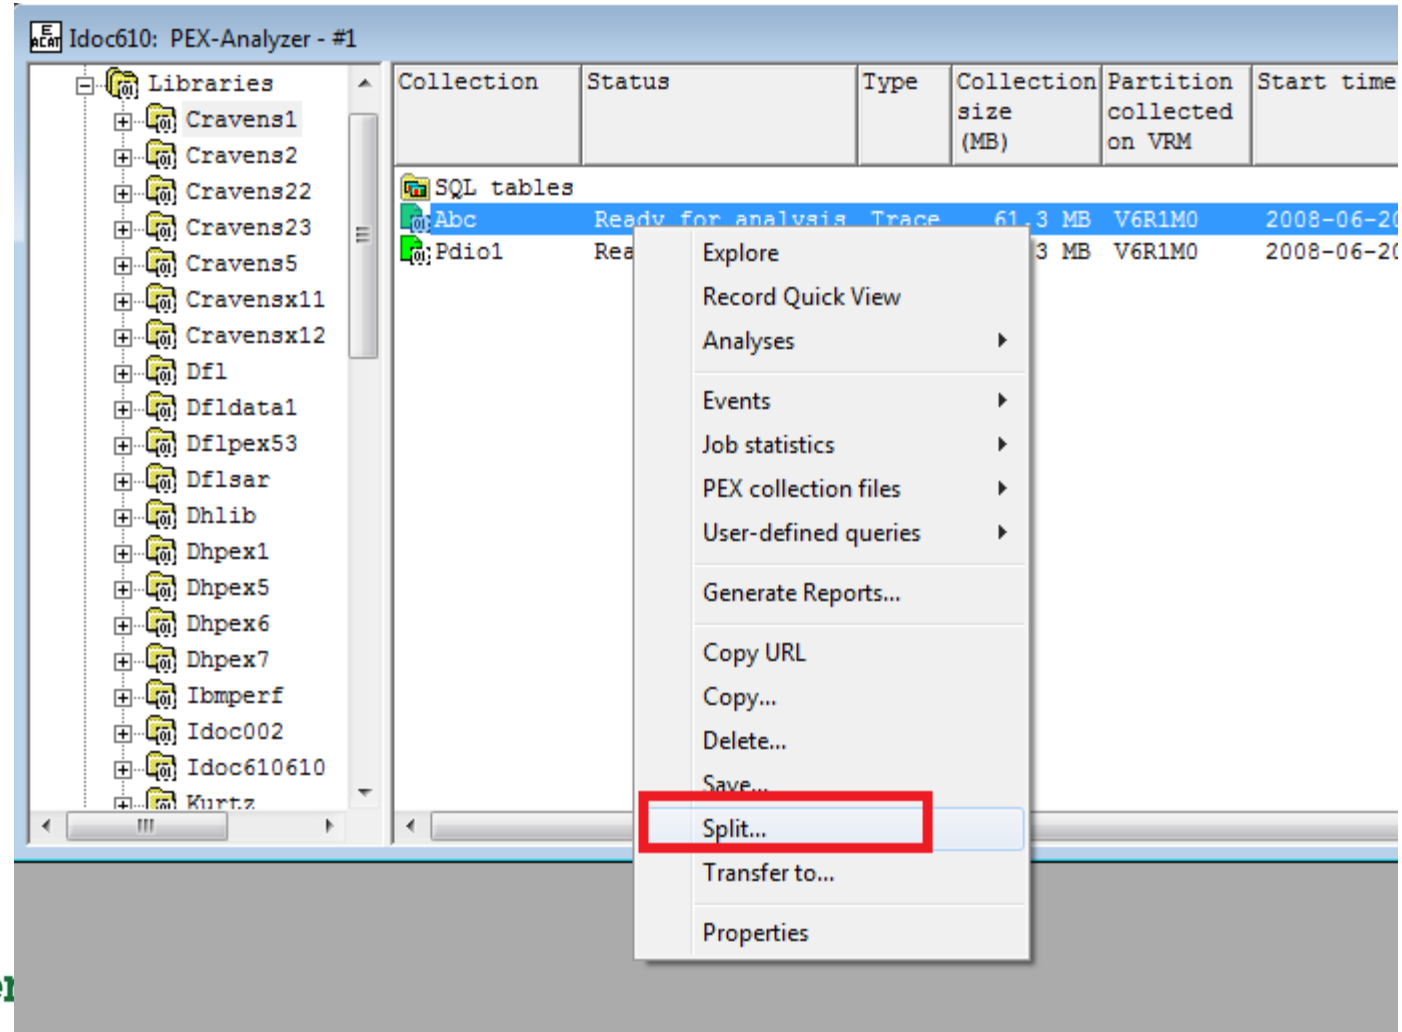

PEX – Stats Hier Analysis Changes

The Stats Hier Analysis was changed so that the SQL table names no longer contain the Taskcount parameter. Parameters are now in the SQL comment field for the table instead.

PEX – PDIO support ASP prompting / filtering

In PEX Analyzer, in the PDIO analysis added support for ASP prompting/comparisons when opening a graph that contains data from more than 1 ASP. Previously this support only existed in CSI and DW.

PEX – Call Stacks Analysis Changes

In PEX Analyzer, running the Call Stacks analysis will now prompt the user with a window asking if the call stacks by job SQL table should be created or not. The default now is to NOT create this table. Specifying to create this extra table will likely double the execution time of the analysis (at least).

Also modified the analysis so the QRECN (max) and MINQRECN are no longer included in the output.

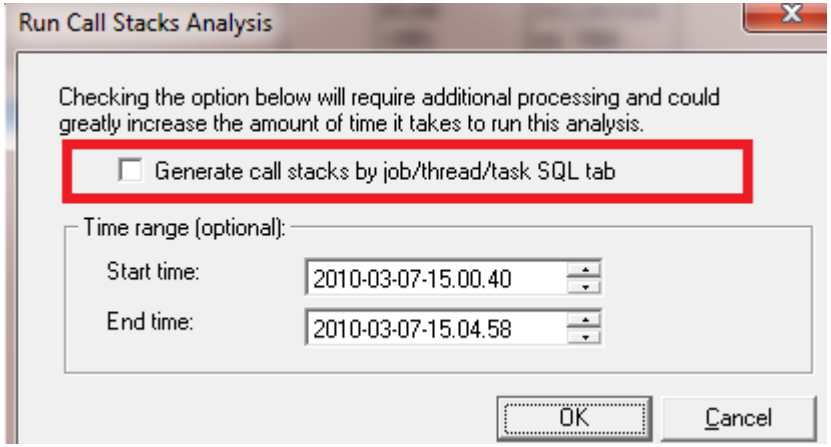

PEX – MCLI analysis options added to TPROF analysis

The 5.4 and higher PEX TPROF analysis folder now contains additional reports to support MCLI analysis if format 4 PMCO events were collected.

At 5.4 the new report is called

Hits by SDAR by Procedure with Start Address

At 6.1+ the reports are
Hits by SDAR by Procedure with Start Address
Hits by SDAR by Procedure with Start Address and Object Name and Context
Hits by SDAR by Procedure with Start Address and Object Name and Context and Job/User

From these you can drill down into the "selected address" report to show the possible call stacks

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.

Revised September 26, 2006

Special notices (cont.)

Revised December 2, 2010

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
	Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- **rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.**

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

==========

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007