iDoctor What's New Apr – Sept 2011

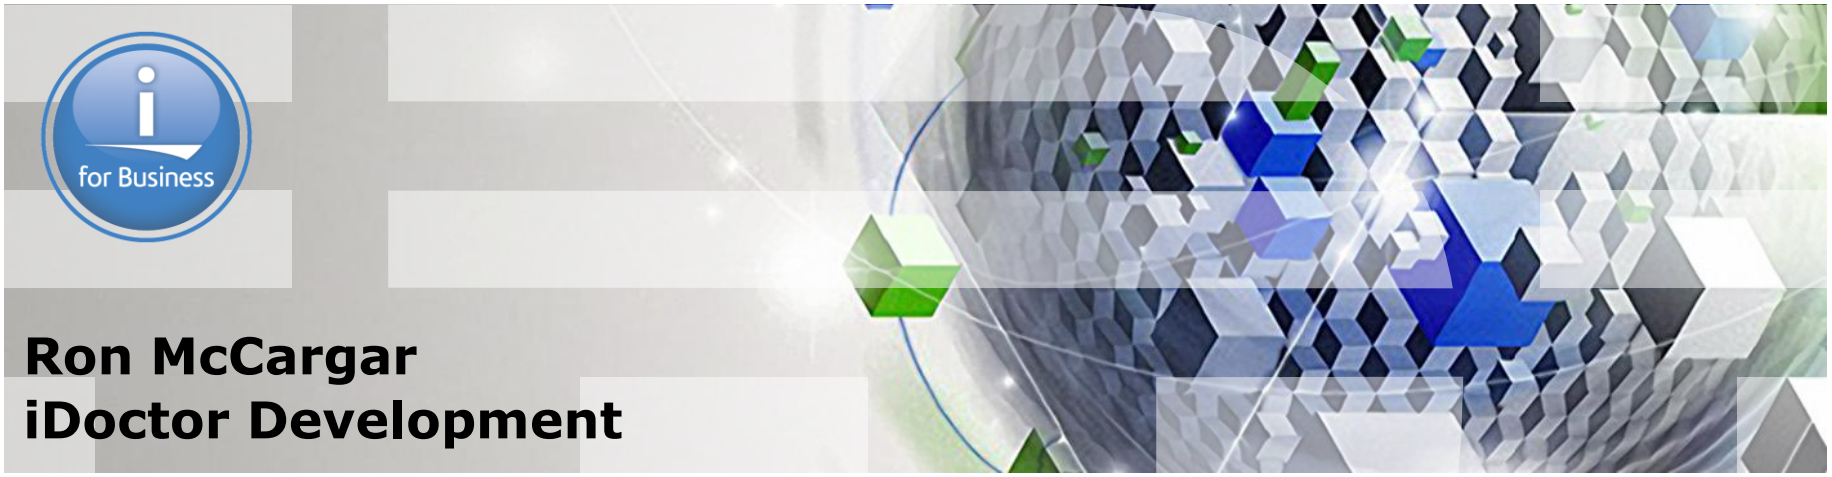

Ron McCargar
iDoctor Development

iDoctor resources

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

iDoctor update history: embedded into the GUI. *Tip: Search the update history on the area of interest to see if your question is answered there.*

YouTube Channel (20+ videos): http://www.youtube.com/user/IBMiDoctor?feature=mhum

IBM i 7.1 Technical Overview – Covers all updates from 6.1 GA to 7.1 GA: http://www.redbooks.ibm.com/redbooks.nsf/RedpieceAbstracts/sg247858.html?Open

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

Apr–Sept 2011 – New Builds

Released new external builds on May 9th, August 5th.

Note: Client update 885 was just added Sept 27th due to a bug with client 878 where a Job Watcher key was required in order to use Heap Analyzer or PEX Analyzer.

Generally try to do 4 or 5 major updates a year.

7.2 internal builds are now available

Apr–Sept 2011 - Overview

Current Focus Areas

CSI Historical Summary

Capacity planning (CPW-based CPU estimations)

Connections list enhancements

New filtering interface for tables/graphs

Installation changes

Plan Cache Analyzer updates

Job Watcher updates, CSI updates, PEX Analyzer updates

Miscellaneous

Apr–Sept 2011 – Current Focus Areas

Making users more productive (always our #1 goal!)

Historical Summaries

Graphing weeks/months of CSI data.

Collection Services

Added new or improved options for graphing HSL loops, memory pools, physical processors and LPAR data.

Plan Cache Analyzer

Made several updates based on user requests

External storage

Server-side VIOS data collection enablement still in progress.
GUI development to start when server-side is done .

Surface new metrics/features added to OS

Added new memory pool graphing options to CSI.
Added physical processor utilization graph to CSI.

Apr–Sept 2011 – CSI Historical Summaries

Allows you to collect and analyze CSI data over the long term (weeks/months)

New options in GUI and command (STRCSMON) lets you collect the data.

Most of the normal CSI collection graphs are available over the historical summary data.

Future plans: Add average day graphs, CPU estimations capability.

Note: In order for these future average day graphs to work well, data will need to be captured with hourly intervals.

Apr–Sept 2011 – CSI Historical Summaries folder

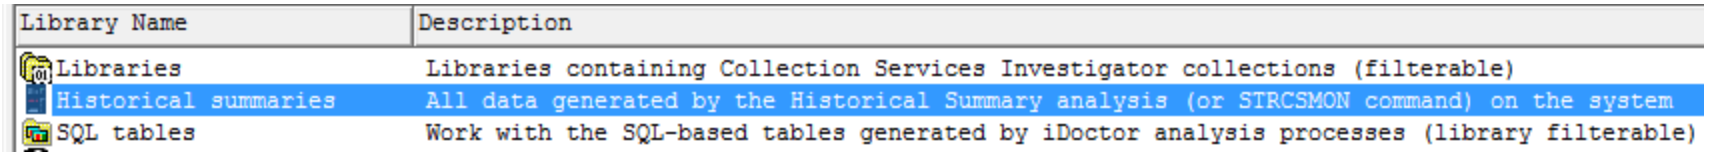

Library Name	Description
Libraries	Libraries containing Collection Services Investigator collections (filterable)
Historical summaries	All data generated by the Historical Summary analysis (or STRCSMON command) on the system
SQL tables	Work with the SQL-based tables generated by iDoctor analysis processes (library filterable)

This new folder contains all the Historical Summary data found on the system.

Historical Summaries can be created by either:

1. Running the Historical Summary analysis over 1 or more collections
2. Running the STRCSMON command (or Start Monitor GUI option) to summarize new CS data on a daily basis

Apr–Sept 2011 – CSI Historical Summaries folder contents

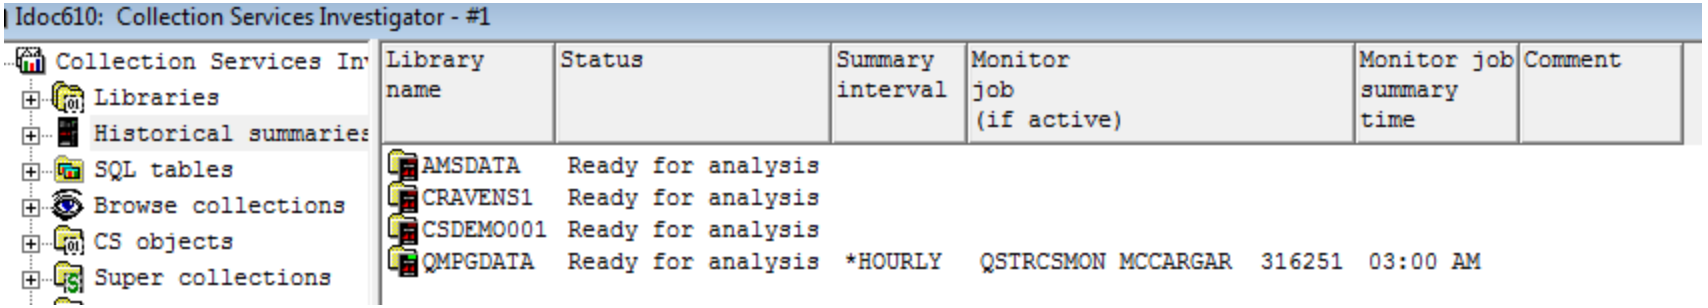

This shows an example of a list of historical summary data.
Only one historical summary “collection” can exist per library.

If the monitor is running, the job information will be displayed in the appropriate fields.

Expand any of these collections to view the historical summary graphs within.

Apr–Sept 2011 – CSI Historical Summaries graphs

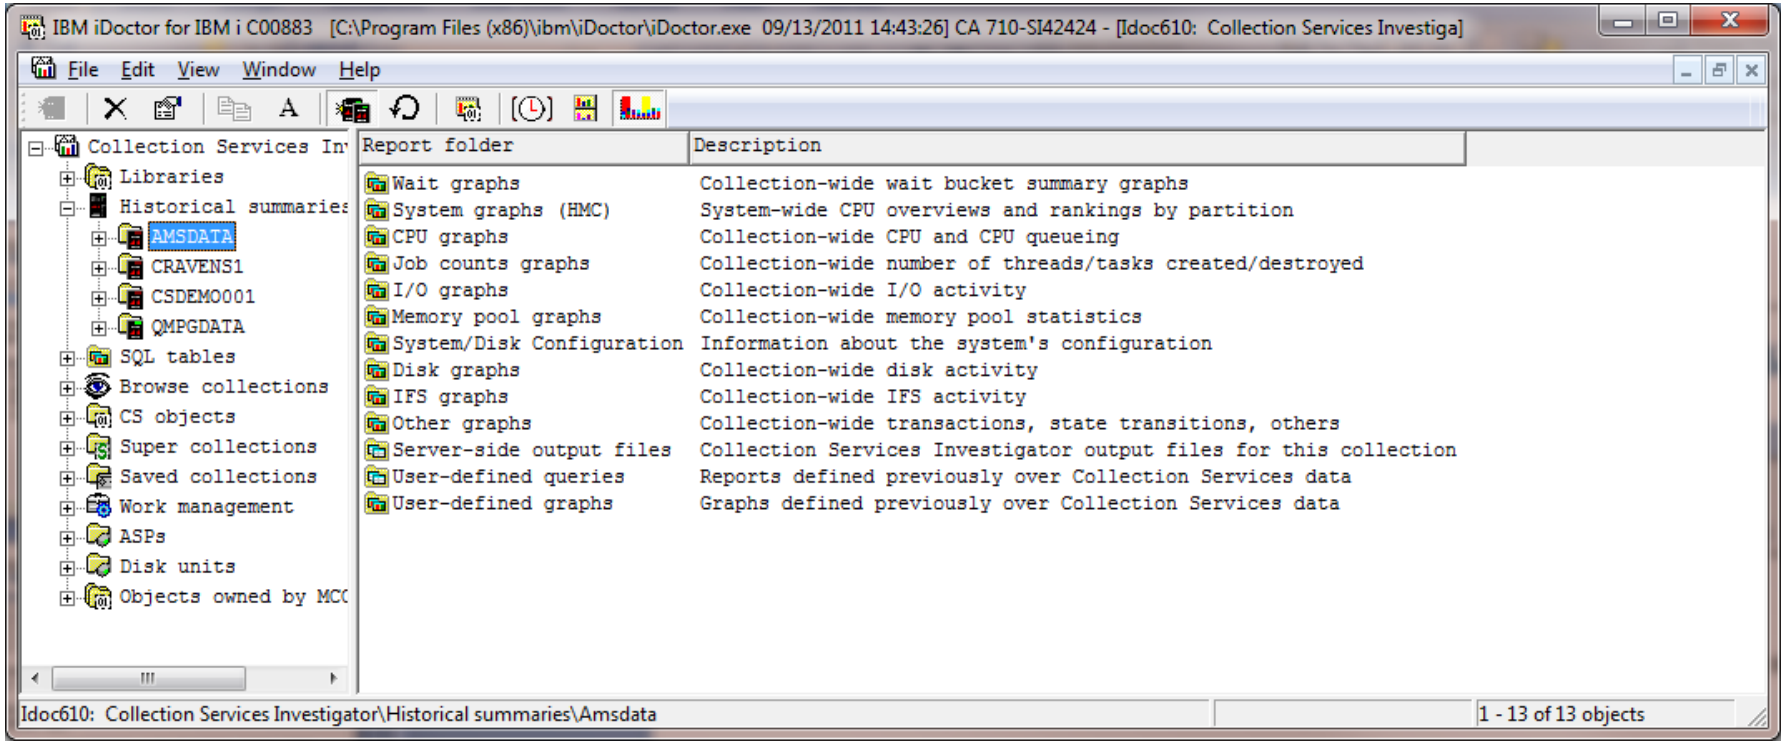

The current graphs available for a historical summary. Some options like system graphs, memory pool graphs and disk graphs may not be available depending on the CS files that were included in the original data.

Apr–Sept 2011 – CSI Historical Summary analysis tables

Data created by the historical summary can be summarized such that one record is produced every day, every 12 hours, every 8 hours, every 4 hours or every hour depending on the level of granularity desired in the data.

The Historical Summary Analysis produces the following tables:

QAIDRCSHSUM - CPU , wait buckets and job statistics.
QAIDRCSHSUMLIST - List of collections included in the summary (so far).
QAIDRCSHSYS - System configuration (or QAPMCONF output) .
QAIDRCSHDISKCFG - Disk configuration
QAIDRCSHDISK - Disk overview data (QAPMDISK , grouped by ASP).
QAIDRCSHPOOL - Memory pools
QAIDRCSHLPARH - CPU and memory for all partitions (QAPMLPARH data)
QAIDRCSSYSPRC – Physical processor statistics (QAPMSYSPRC data)
QAIDRCSHJBCNTS - Data needed for the Job counts graph.
An optionally created file if all collections have been summarized.

Apr–Sept 2011 – CSI Historical Summaries Start Monitor

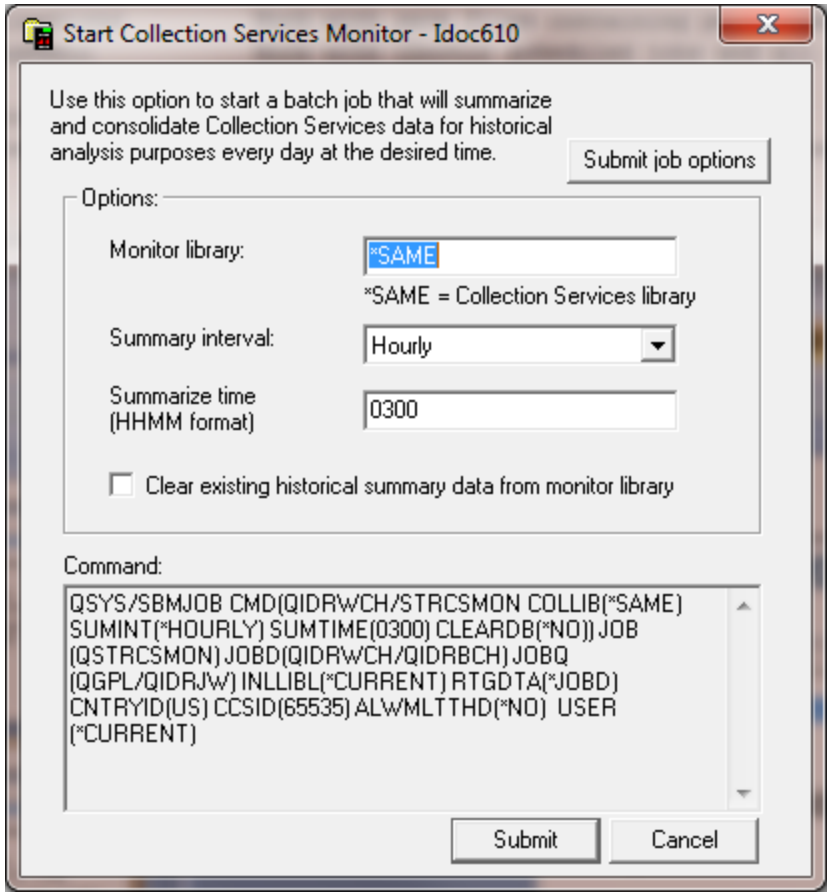

Access this screen by right-clicking the Historical Summaries folder and using the Start Monitor… menu.

Default options will create historical summary data in the same library as the Collection Services default library with 1 hour intervals.

The command/job will summarize new data every day at 3 AM until it is stopped.

Apr–Sept 2011 – CSI Historical Summaries STRCSMON

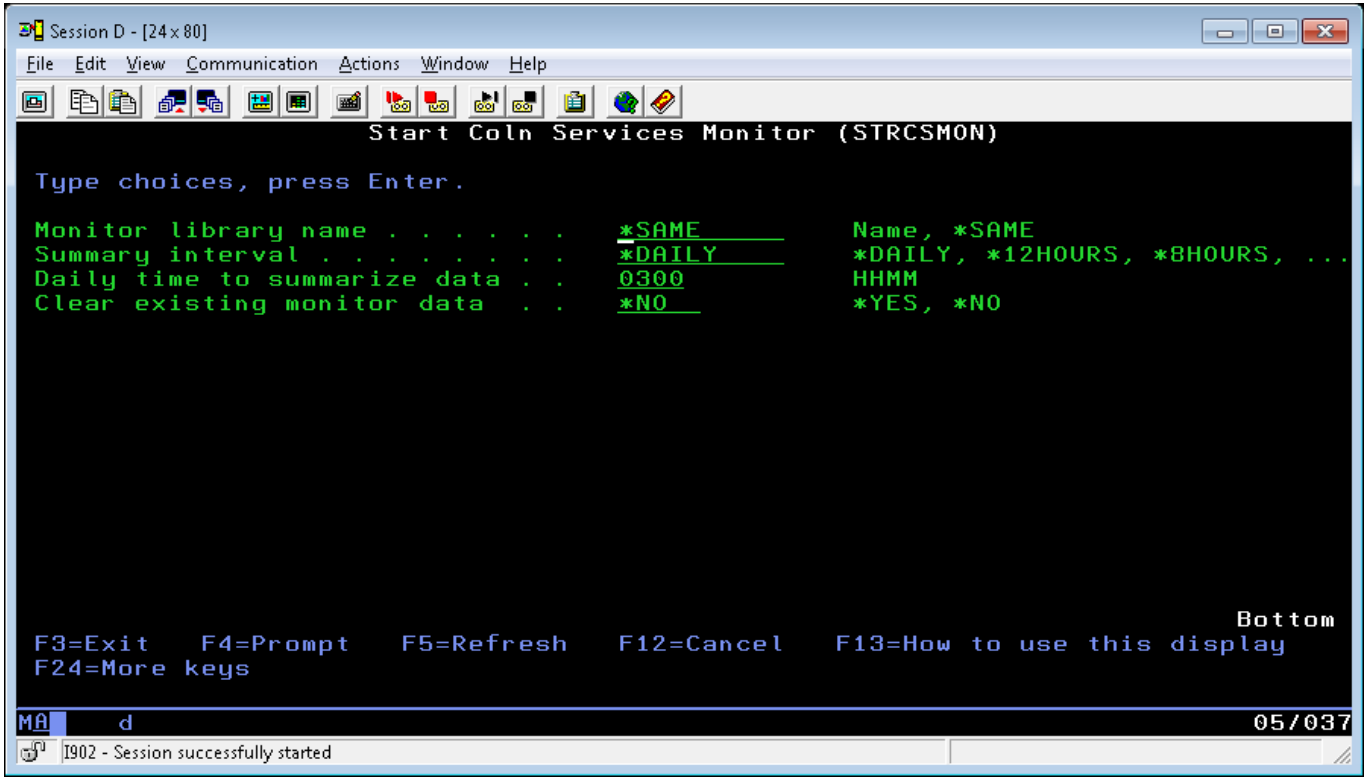

STRCSMON shipped in library QIDRWCH
Note: Use ENDCSMON command to end the monitor.

The clear option is necessary if the file formats changed between builds (watch the update history).

Apr–Sept 2011 – CSI Historical Summary Analysis

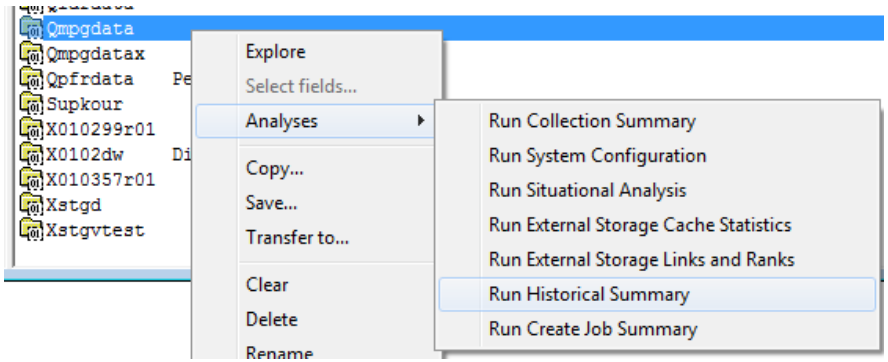

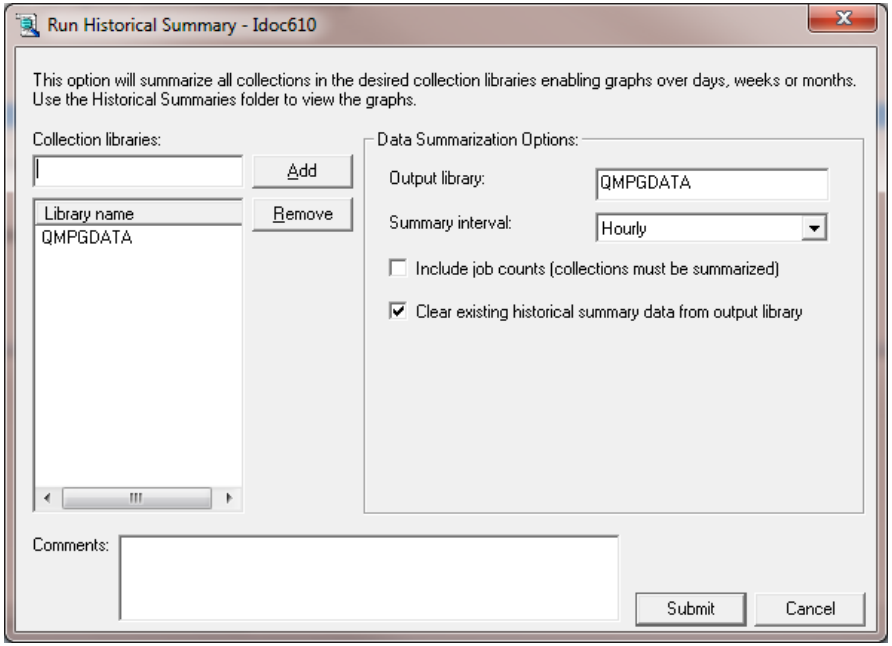

Accessible by right-clicking one or more! Libraries under the Analyses menu. All collections in all libraries included will be summarized.

Note: Most analyses in all components can now be accessed from the libraries folder in order to apply the analyses to all collections in the library.

Apr–Sept 2011 – CSI Historical Summaries Time Filtering

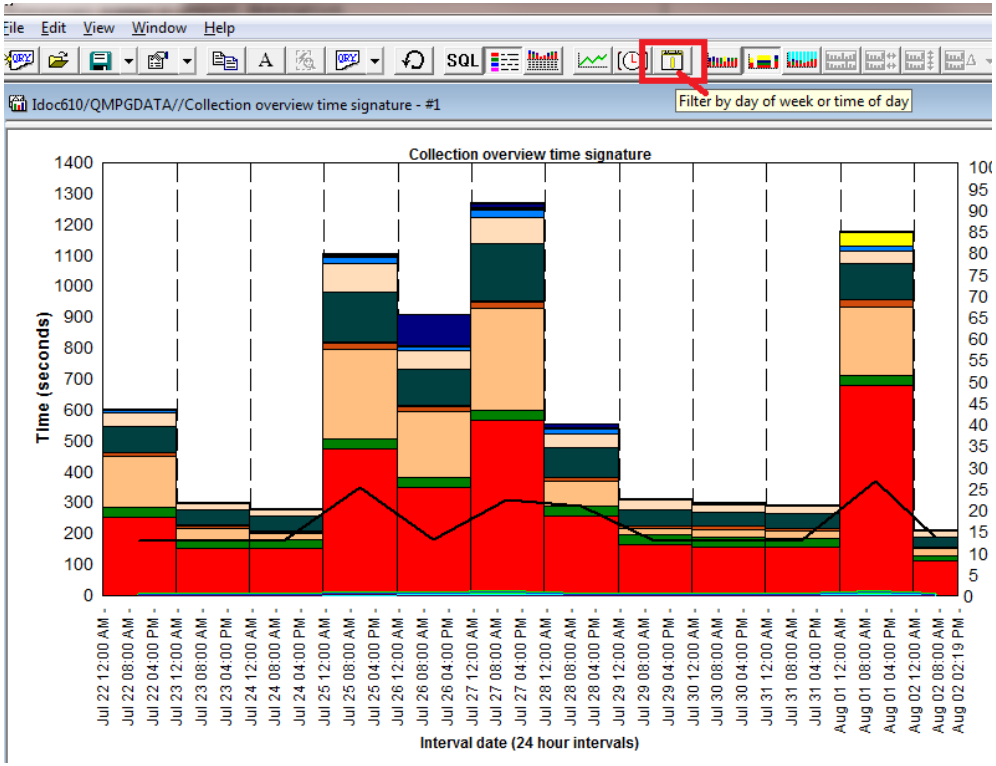

New button on toolbar lets you filter the data by time range or days of the week.

Power your planet.

Apr–Sept 2011 – CSI Historical Summaries Time Filtering

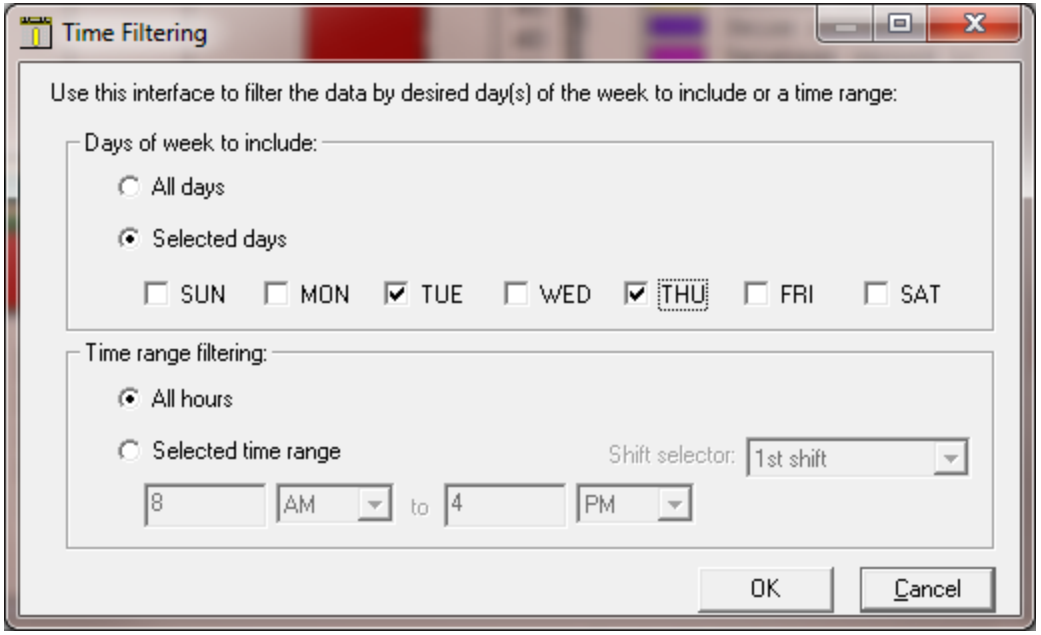

This lets you do things like only graph from 8 AM – 5 PM or exclude the weekends from the graph output.

Currently this option is only available for historical summaries, but it could be applied else where if desired.

Apr–Sept 2011 – CSI Historical Summaries Future Plans

Add drill down options into job rankings (if the original collection data is still available).

Add CPU estimations capability.

Add average day graphs.

We might also add a historical summary analysis for Job Watcher, in order to make it easier to graph lots of Job Watcher collections in a monitor at once.

Apr–Sept 2011 – Capacity planning – (CPU Estimations)

An IBM internal only component was added to allow a user to perform CPU estimations over PerfNav data from Average Day CPU graphs.

These estimations show you how the CPU utilization would likely change if the system model and number of CPUs changed to the desired new configuration.

All estimations are done using CPW values stored in IBM's System.XML file from the IFS.

Future plans:

- Add CPU estimation capability to CSI and possibly JW.
- Add disk estimation capabilities to CSI.

Apr–Sept 2011 –CPU Estimations Example

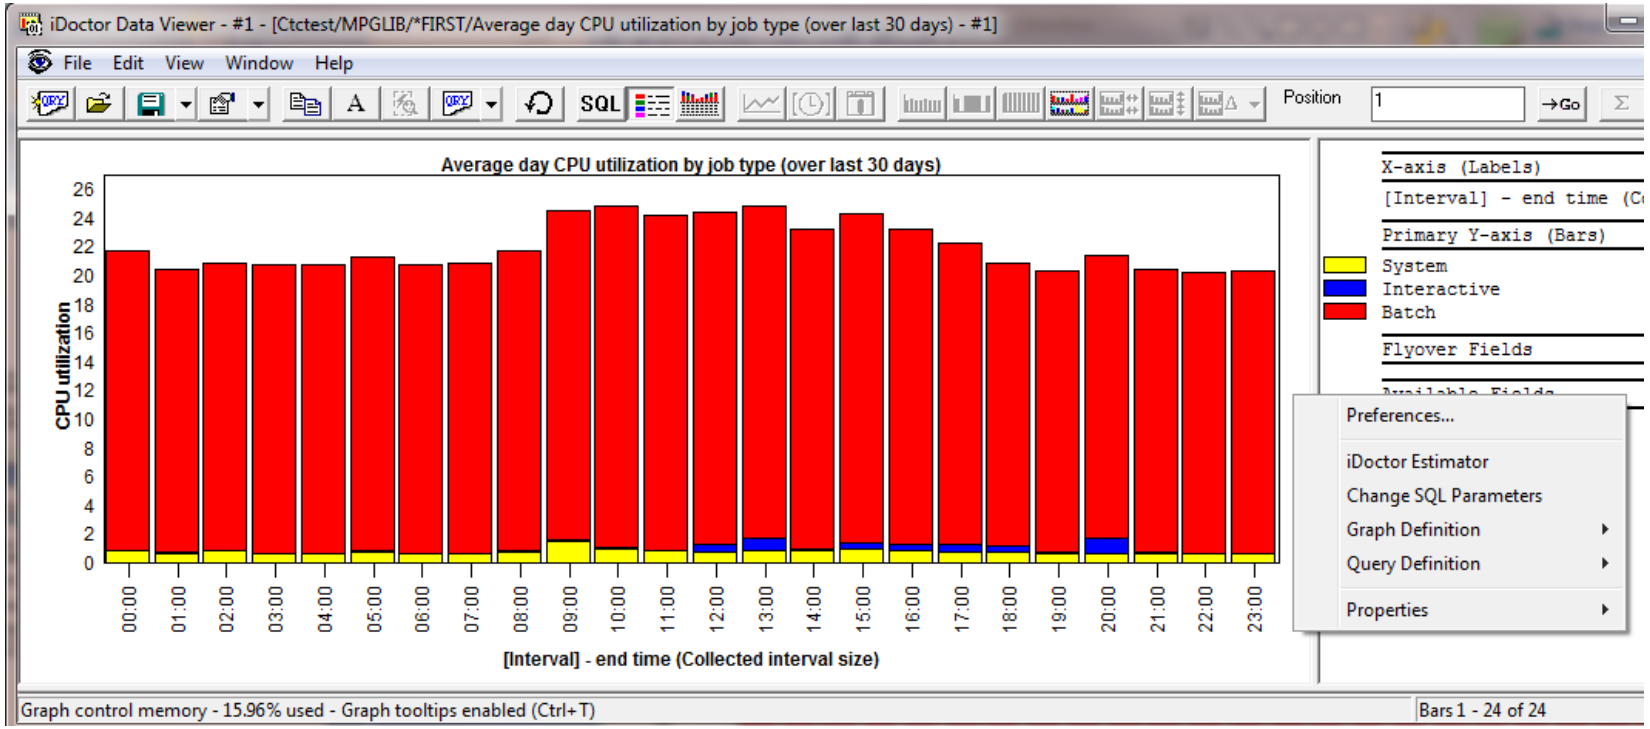

Use the iDoctor Estimator menu to launch the window to perform CPU estimations.

Apr–Sept 2011 –CPU Estimations Example (cont)

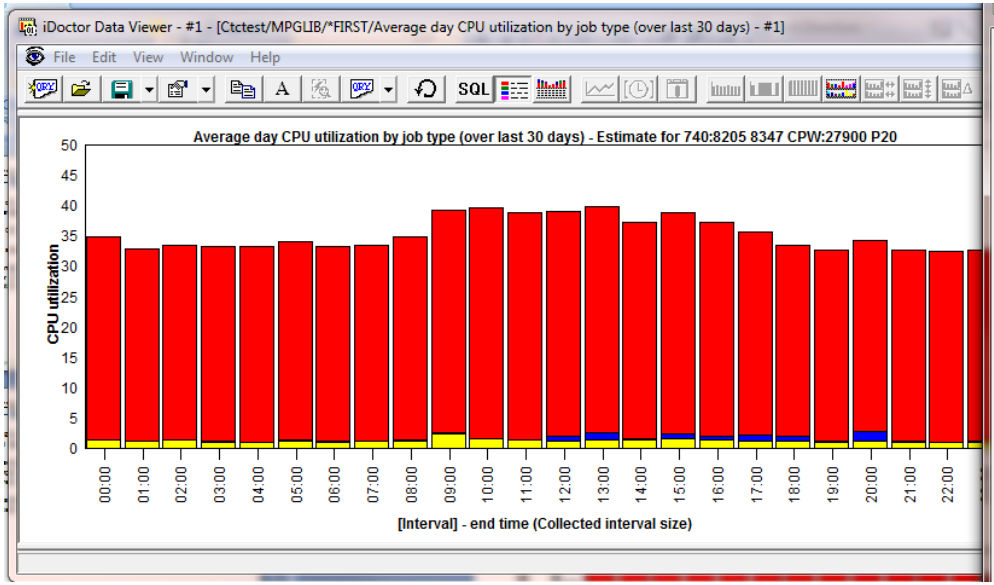

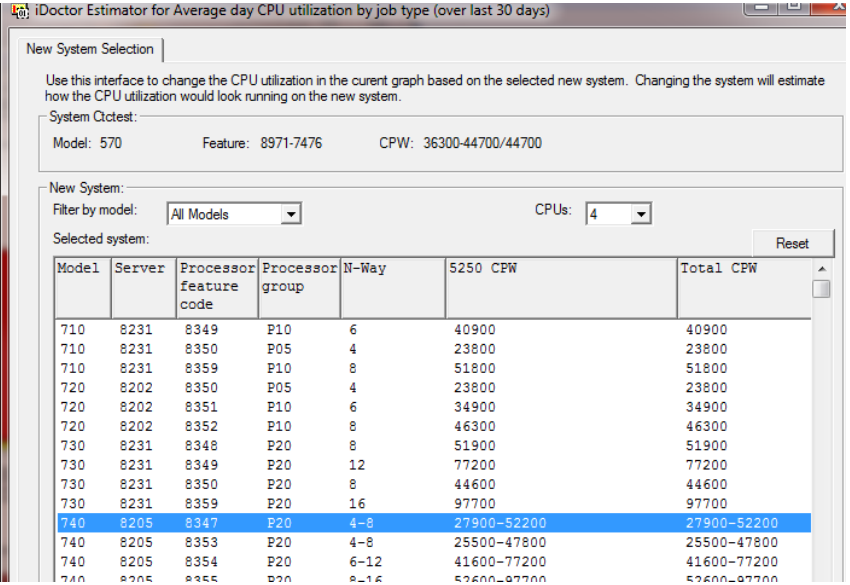

From the iDoctor Estimator window making changes to the desired CPU model and number of CPUs will update the graph accordingly.

Press the Reset button to go back to the original data.

Apr–Sept 2011 – Connections List Enhancements

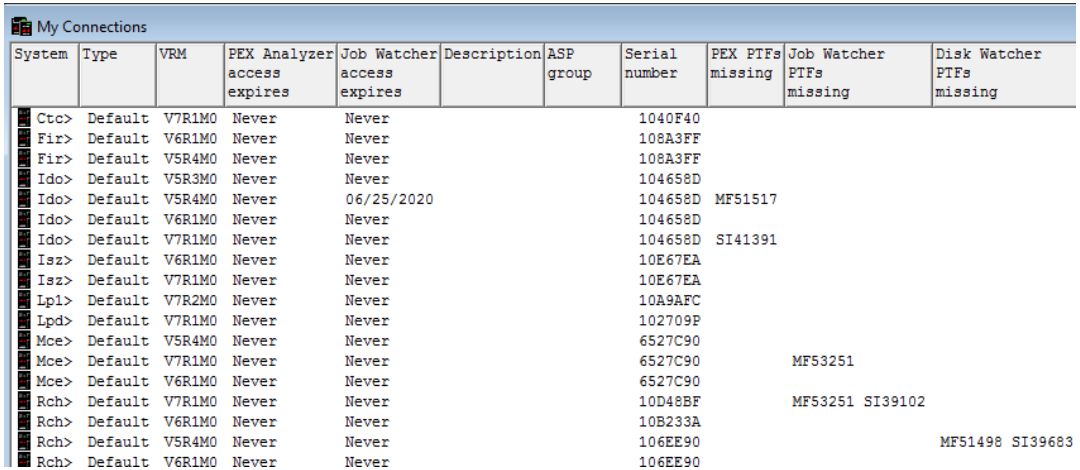
My Connections

System	Type	VRM	PEX Analyzer access expires	Job Watcher access expires	Description	ASP group	Serial number	PEX PTFs missing	Job Watcher PTFs missing	Disk Watcher PTFs missing
Ctc>	Default	V7R1M0	Never	Never			1040F40			
Fir>	Default	V6R1M0	Never	Never			108A3FF			
Fir>	Default	V5R4M0	Never	Never			108A3FF			
Ido>	Default	V5R3M0	Never	Never			104658D			
Ido>	Default	V5R4M0	Never	06/25/2020			104658D	MF51517		
Ido>	Default	V6R1M0	Never	Never			104658D			
Ido>	Default	V7R1M0	Never	Never			104658D	SI41391		
Isz>	Default	V6R1M0	Never	Never			10E67EA			
Isz>	Default	V7R1M0	Never	Never			10E67EA			
Lp1>	Default	V7R2M0	Never	Never			10A9AFC			
Lpd>	Default	V7R1M0	Never	Never			102709P			
Mce>	Default	V5R4M0	Never	Never			6527C90			
Mce>	Default	V7R1M0	Never	Never			6527C90		MF53251	
Mce>	Default	V6R1M0	Never	Never			6527C90			
Rch>	Default	V7R1M0	Never	Never			10D48BF		MF53251 SI39102	
Rch>	Default	V6R1M0	Never	Never			10B233A			
Rch>	Default	V5R4M0	Never	Never			106EE90			MF51498 SI39683
Rch>	Default	V6R1M0	Never	Never			106EE90			

Added new columns to show component access code expiration dates and missing PTFs in order to let you monitor this information across multiple LPARs more easily. The system serial number has also been added.

Added new menu "Set default signon" that will store the default user id and password to use when making new connections if one is not already available. This function gets used if selecting many connections and choosing one of the menu options that apply to multiple partitions.

Added menu "Check expiration dates" that lets you check the iDoctor access code expiration dates of the selected partitions and updates the list.

Using the Check PTFs function will now show the missing PTFs in the list of connections.

Multiple connections can be removed at once now.

The uninstall option works now against multiple partitions at once.

Apr–Sept 2011 – New filtering interface for tables

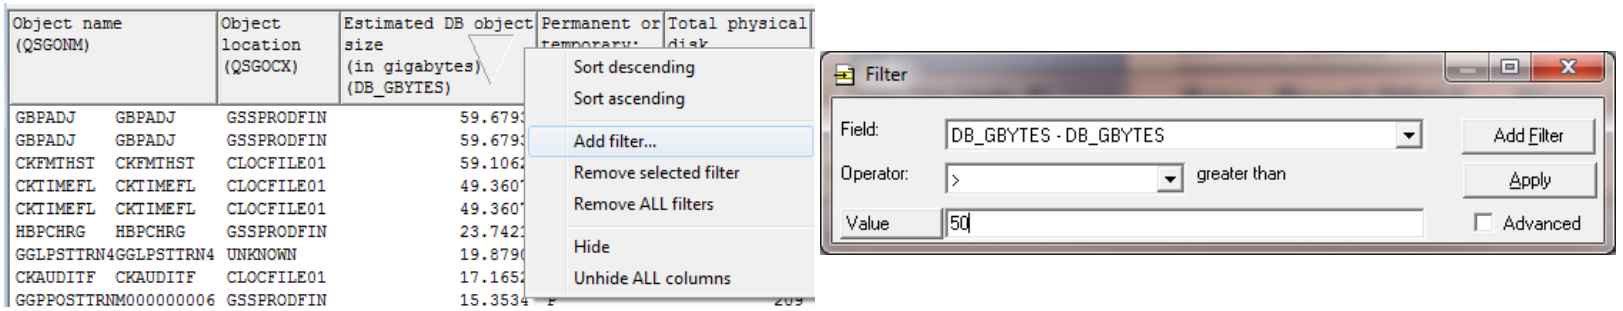

Right-click the desired column to filter on. Set the desired operator and value and press Apply to refresh immediately.

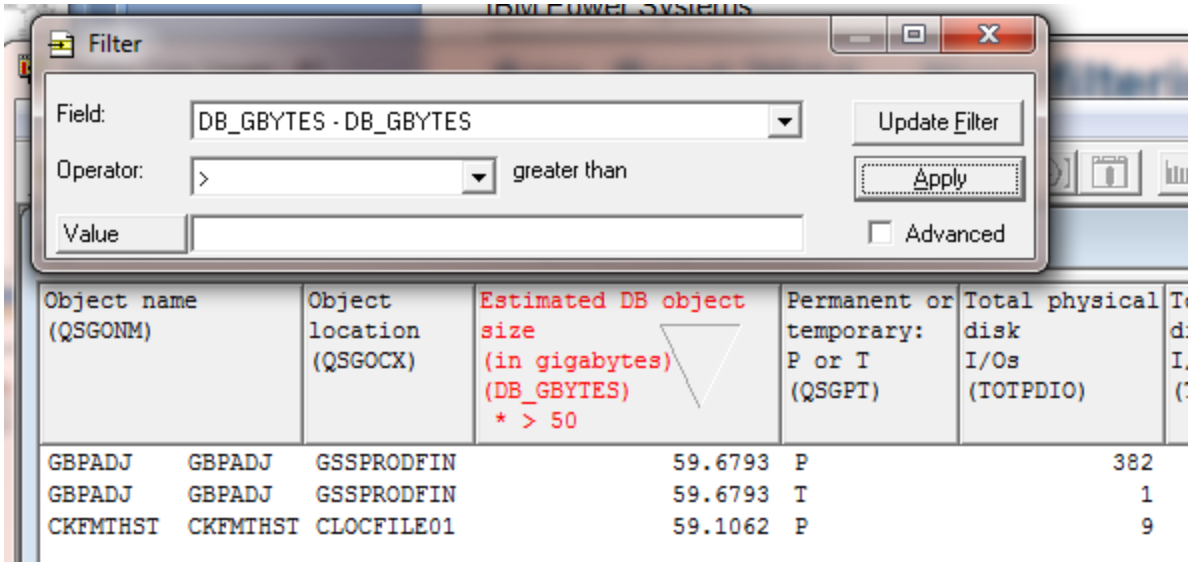

Can also right-click columns to perform other actions: Sort, Remove filter, Hide column, Unhide all columns

Apr–Sept 2011 – New filtering interface for graphs

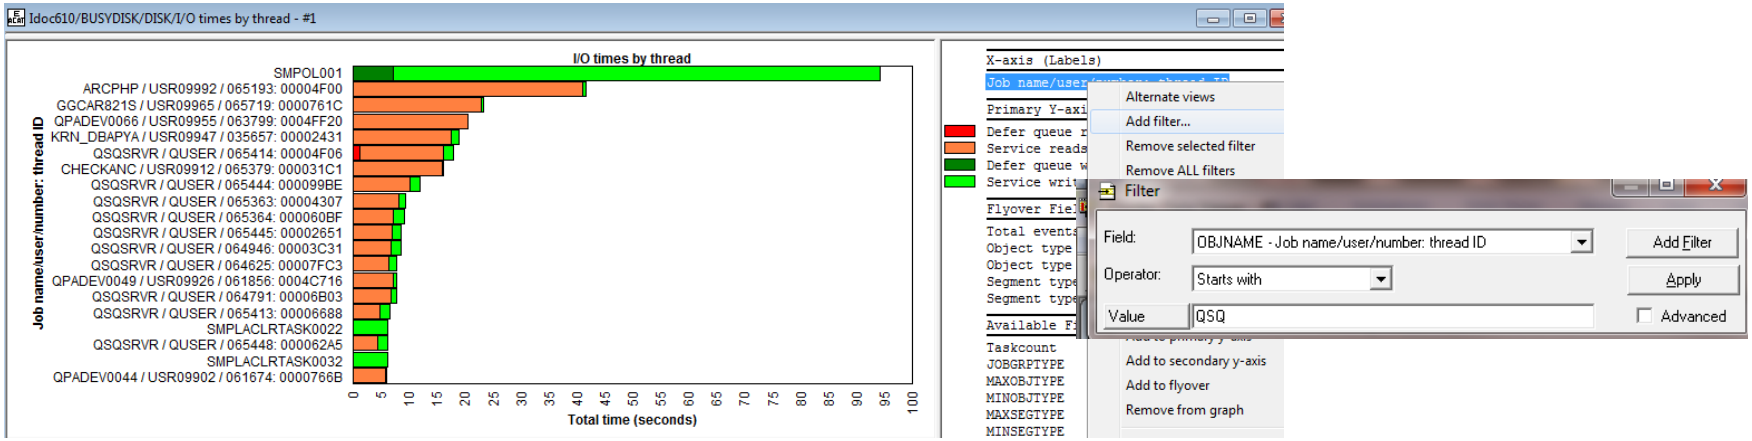

Right-click the desired column to filter on. Set the desired operator and value and press Apply to refresh immediately.

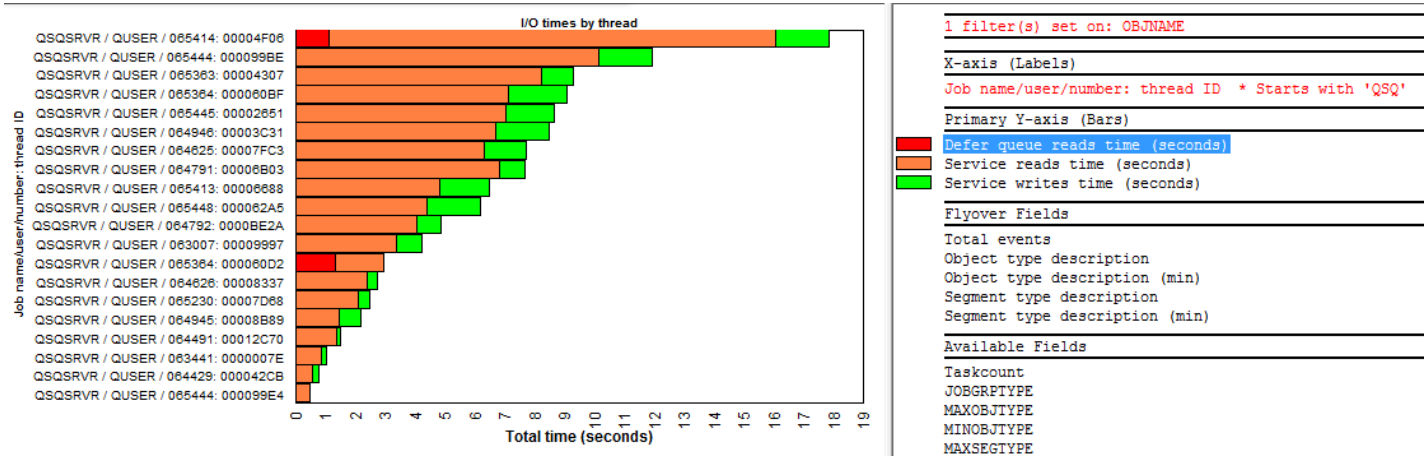

Right-click existing filters to change/remove them from the graph.

Power your planet.

Apr–Sept 2011 – Installation Enhancements

At install time of server builds, most needed stored procedures will be created and registered in QIDRGUI library to avoid long startup delays the 1st time the GUI is used after a reinstall of server builds. This only applies to 6.1+ only.

During Job Watcher installation added a check to make sure the definitions file QAPYJWDFN exists in QUSRSYS. If not it will be created and filled with the IBM-supplied definitions.

On Windows 7, the iDoctor GUI was not showing up properly in the Windows installed programs.
Also added additional information about the client build: publisher info, website, estimated size, install date, etc.

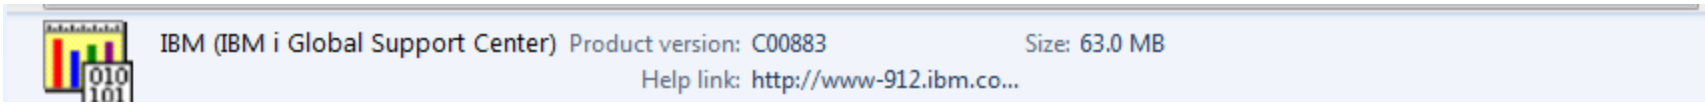

The default install directory changed twice. First it changed to C:\IBM\iDoctor and now it is C:\Program Files\IBM\iDoctor This won't change again. Be sure to save any user-defined queries/graphs .mdb databases you have created before removing the old install directory.

(Note: On 64-bit OS the default path will be C:\Program Files (x86)\IBM\iDoctor)

When the installation encounters an ODBC connection error, display an error to the user explaining the job log will not be retrievable at the end of the installation. Afterwards the installation will continue normally.
ODBC is only needed to retrieve the job log of the installation so it will no longer be a fatal error.

Apr–Sept 2011 – Installation Enhancements (cont)

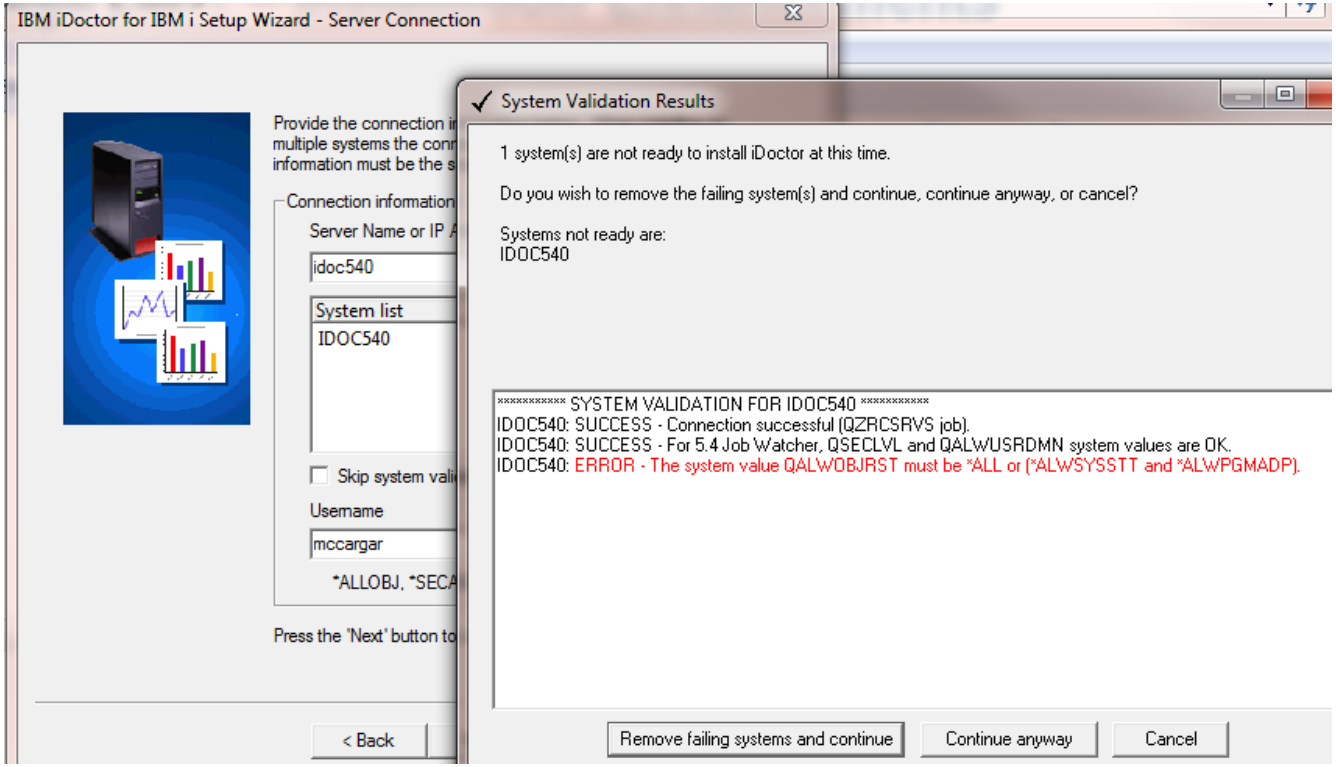

On the server information page of the installation, each system will now be validated when pressing the Next button. Though this will take some time to connect to each system it should avoid problems later. Also added a checkbox to this screen that allows you to skip the validation step.

During validation each system is tested to ensure that:
1. A valid connection can be created to the system
2. Determine the VRM of the server
3. Determine that the iDoctor server code save file exists, matching the desired VRM.
4. If 540, determines if QSECURITY is 50 and QALWUSRDMN is not *ALL then warn the user that Job Watcher won't be usable unless changed.
5. Checks if QALWOBJRST is set correctly *ALL or (*ALWSYSSTT and *ALWPGMADP)

If any of these checks fail, a message is shown allowing you to remove the system from the list and continue, not remove the system or cancel.

Apr–Sept 2011 – Plan Cache Analyzer – Snapshot support

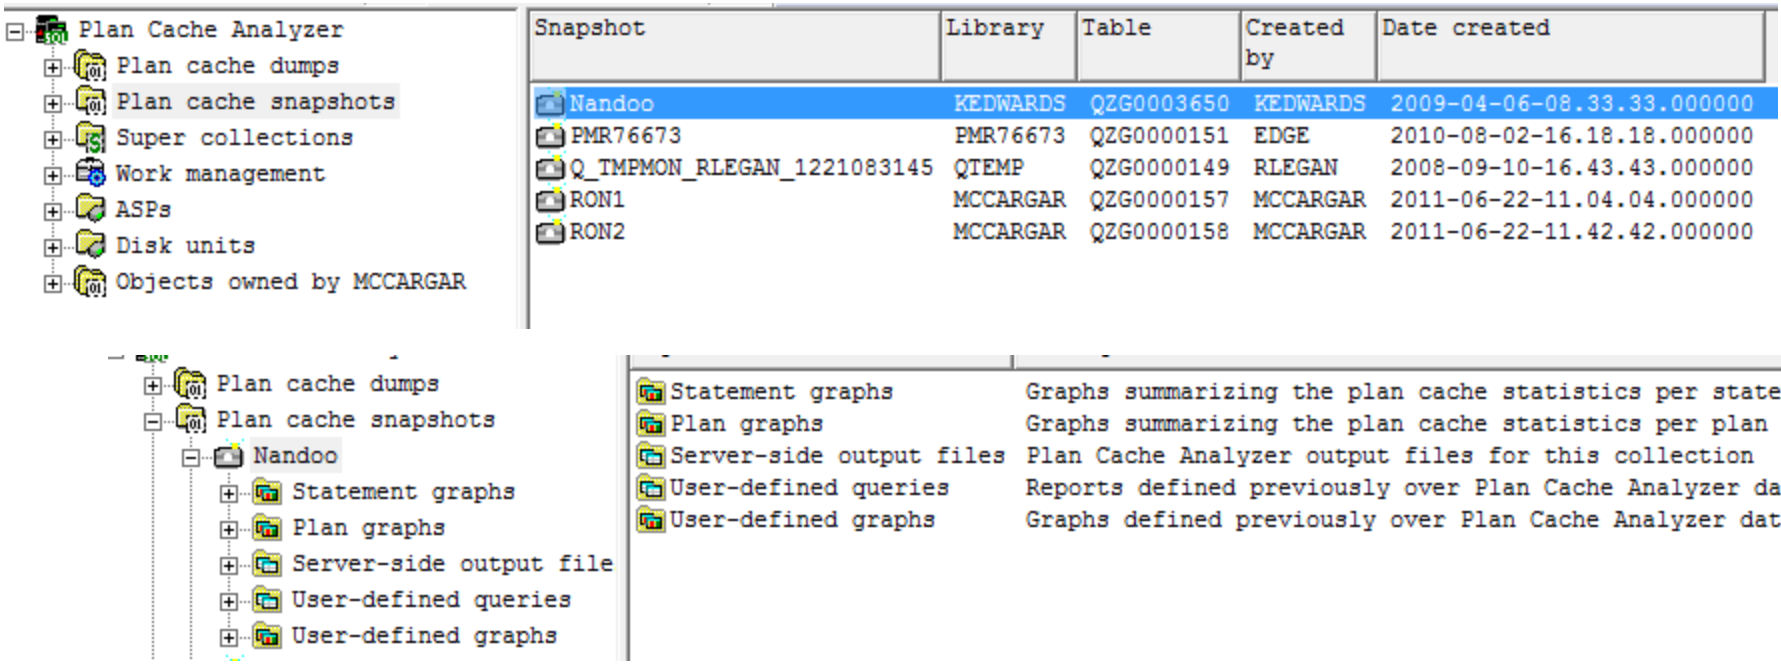

Enhanced Plan Cache Analyzer to support Plan Cache snapshots.

Separate folders now exist to work with plan cache dumps and plan cache snapshots.
Plan cache snapshots offer graphs for statements and plans.

You can select desired records from the graphs and use the Extract function to copy the data from the selected bars into a filtered table/snapshot. This option allows you to automatically import the data into the Snapshot repository in iNav to avoid having to manually import the data yourself.

Apr–Sept 2011 – Plan Cache Analyzer – Additional changes

You can now import files > 2 GB.

You can now import active query dumps (.dbops) either from the IFS, or you can create a new one using the Start Plan Cache Analyzer wizard and selecting the desired job to analyze.

You can now right-click a library under Plan Cache dumps and use the Start Analysis option to run (or rerun) the analysis data of the plan cache dump. The CDATA file must already exist (this is created after importing the data from the IFS.) Any existing analysis files will be deleted and recreated.

Apr–Sept 2011 – Job Watcher – Top threads over time graphs

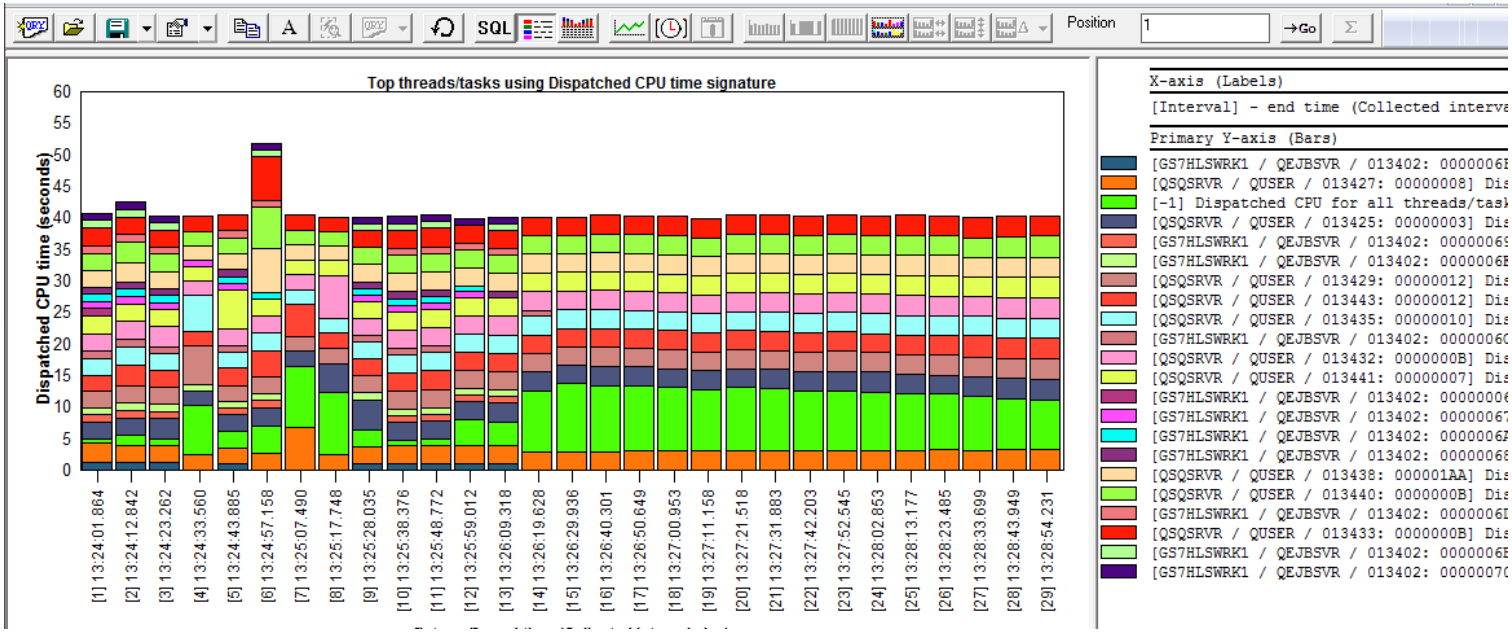

These graphs show the threads/tasks that spent the most time in the desired wait bucket over time.

A grouping filter is used with a default value of 1 second where all threads/tasks that spent less than 1 second in CPU each interval and added together. In some cases if there are a large number of threads/tasks this filter will need to be increased in order to make the graph more readable.

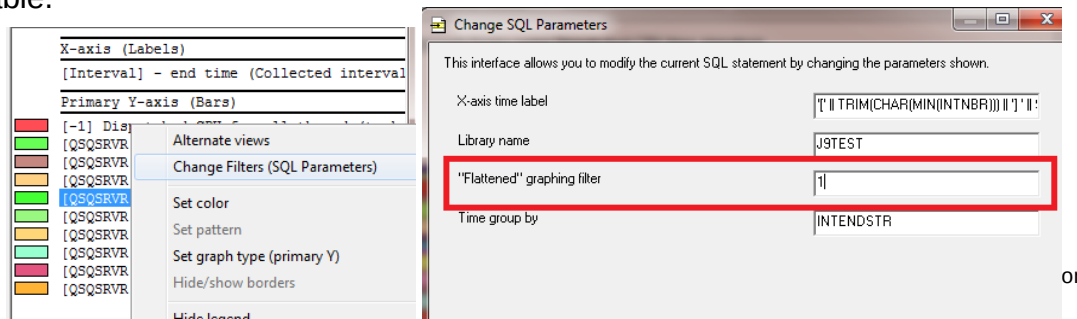

Power your planet.

Apr–Sept 2011 – Job Watcher - Objects Waited on Tab Enhancements

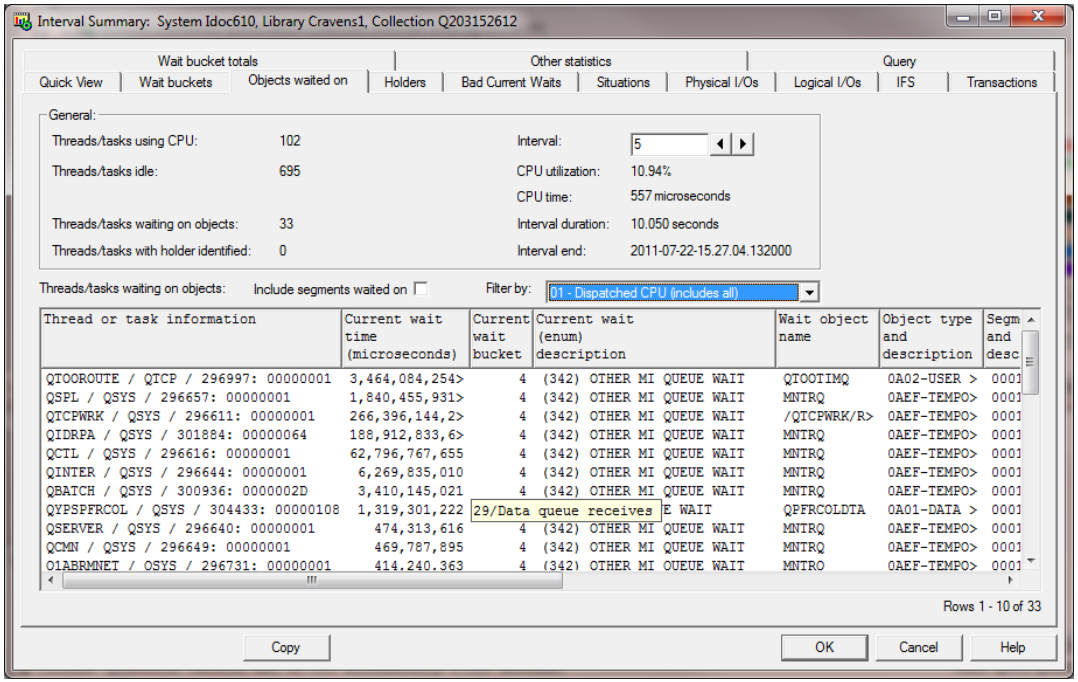

In the interval summary interface on the objects waited on tab, the list of jobs now includes jobs that are waiting on an object but did not use CPU in the interval. Previously only jobs that used CPU in the interval were shown. Also added a checkbox to show segments waited on.

For this to work and show the jobs waiting but did not use CPU, the latest server builds must be installed and the collection must be summarized.

Apr–Sept 2011 – Job Watcher – Detail reports drill down option changes

In 5.4+ JW, made the following changes to the options available under the "Detail reports" menu (from overview, rankings or single job grouping over time graph types):

1. Situation details report will filter by taskcount (thread) if coming from a rankings or single job/thread over time chart.

2. Moved the "Process -> Primary threads using CPU" report under the "Job" submenu.

3. At 6.1+, removed "Process -> Activation group details" because activation group data collection is not working properly at 6.1 and higher.

4. The reports under the following folders will now filter by the current job grouping if one exists (thread, job, user, etc) as well as time selection:
   Job, Classic JVM, TPROF, Call Stack Summary,

5. The queries under Call Stack Summary will now have their record counts calculated correctly (previously they were set to 10000 records because these queries begin with "WITH").

Apr–Sept 2011 – Job Watcher – Call Stack Reports

In JW 5.4+, added some changes to the set of call stack reports. There is a new report to group by procedure and offset, and also drill downs into these reports will filter by the offset field if it exists in the previous report.

Apr–Sept 2011 – CSI/JW – Create Job Summary

In CSI/JW from the create job summary reports, within the SQL tables interface, right-clicking one or more jobs now offers a Job runtimes graph under the 'Other graphs' menu which displays a simple graph of the run times for the selected jobs/threads (or at least the run times captured by the collection).

When working with the SQL tables generated by the JW/CSI create job summary option, changed the menus shown from Job rankings to Job rankings summarized and Thread rankings to Thread rankings summarized to indicate that all collections included in the create job summary are added together for each job or thread shown

Also added a new set of graphs called "per collection", that shows the collection name next to the job name. This should allow users to more easily compare performance of jobs from 1 collection to another.

Additionally all of the reports now use all SQL tables selected from the initial graph. So if you want to graph data from multiple create job summary SQL table files at once, you can now.

Enhanced the CSI and JW create job summary option to provide additional filtering options. You can now filter by job name, job user, job number, current user profile, subsystem or time.

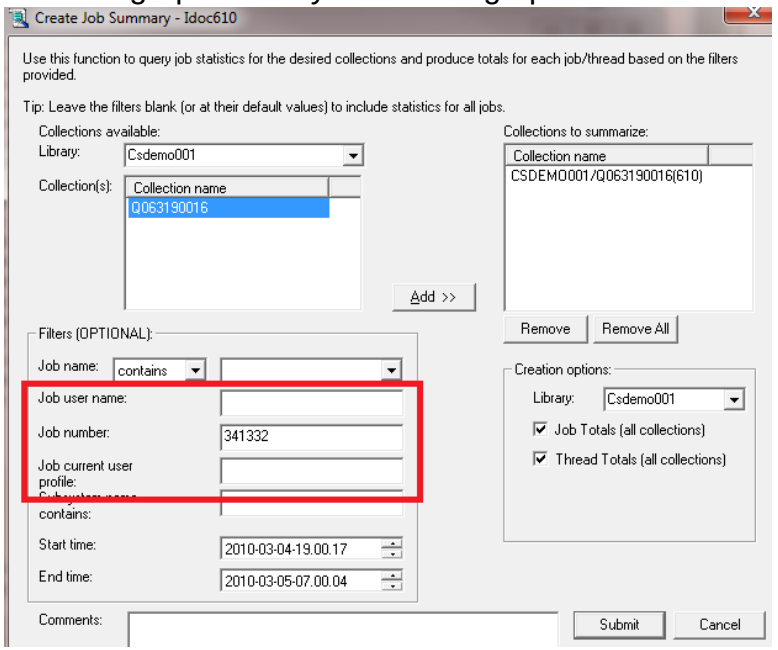

Power your planet.

Apr–Sept 2011 – CSI/JW – Wait Buckets Tab

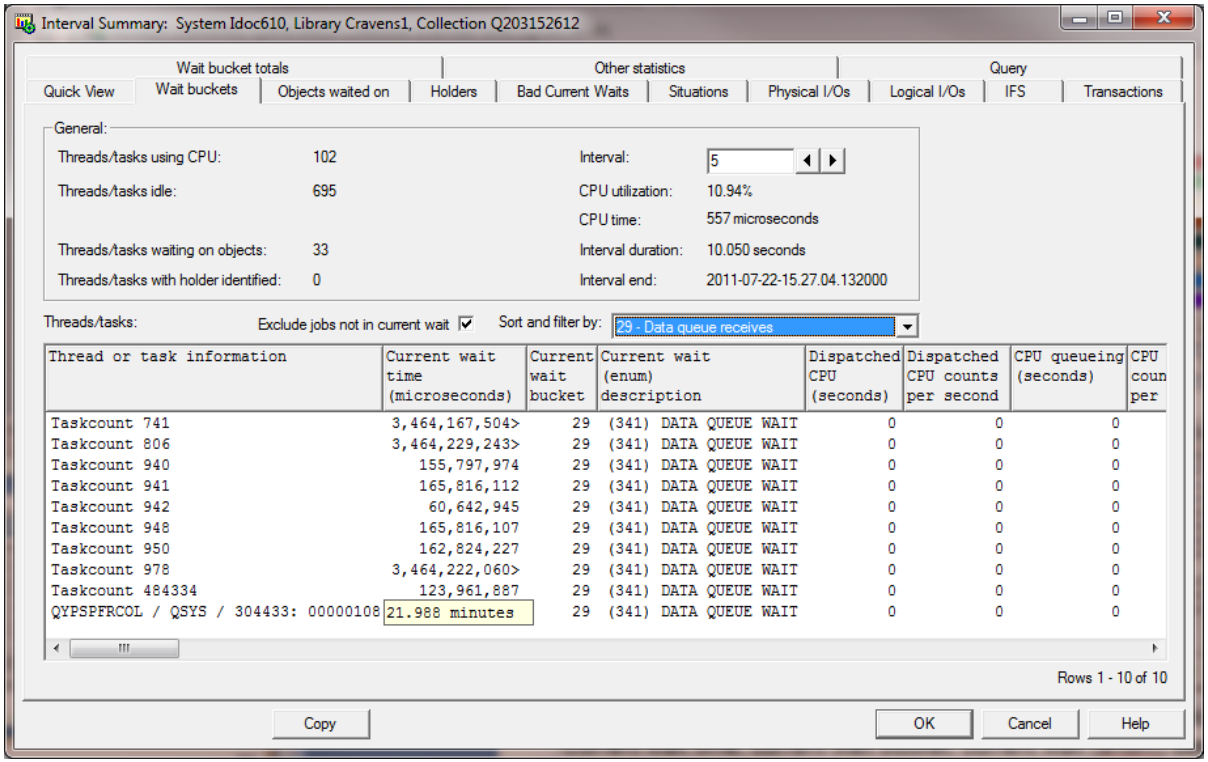

In CSI/JW, in the interval summary interface, Wait Buckets tab, added the following new fields to the list: Current wait time, current wait bucket, current wait (enum) description.

Also added a checkbox called "Exclude jobs not in current wait" that filters the list so it only shows jobs that were currently in the selected wait from the drop down list at the moment the interval ended.

The list of jobs also now contains the counts per second for each wait bucket as well as the avg sync read and write response times.

Apr–Sept 2011 – CSI - System graphs – LPAR CPU time overview graph

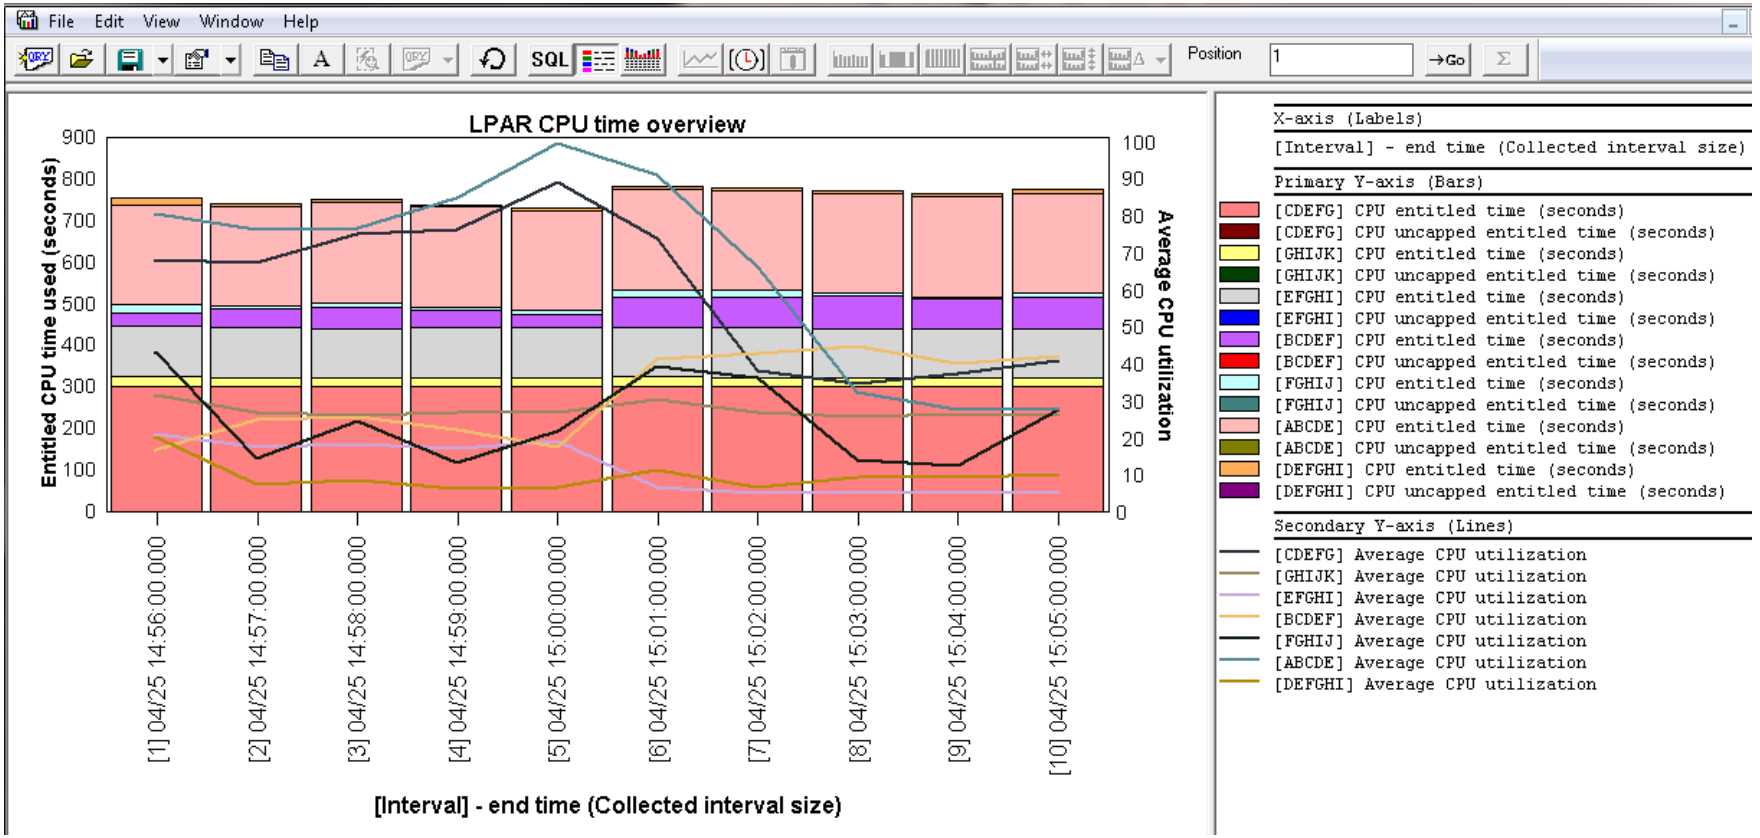

In CSI 6.1+, new graph shows CPU entitled and CPU uncapped entitled for all partitions over time along with CPU utilization for each partition on the 2nd Y axis.

Apr–Sept 2011 – CSI - System graphs – LPAR memory overview graph

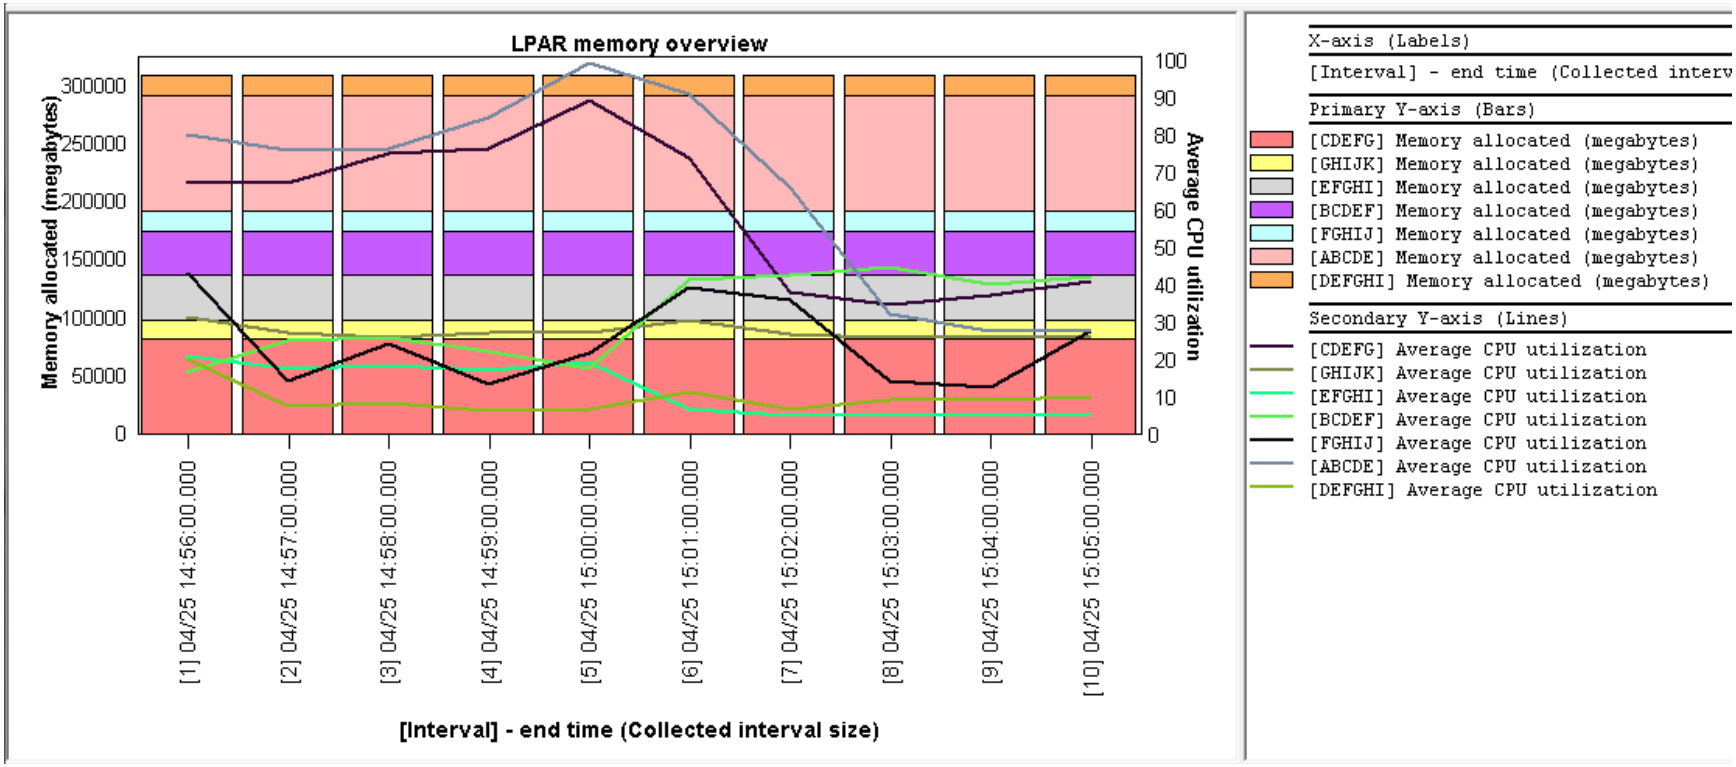

In CSI 6.1+, new graph shows the physical memory allocated to each partition over time along with CPU utilization for each partition on the 2nd Y axis.

Apr–Sept 2011 – CSI - System graphs – Physical processor utilization overview

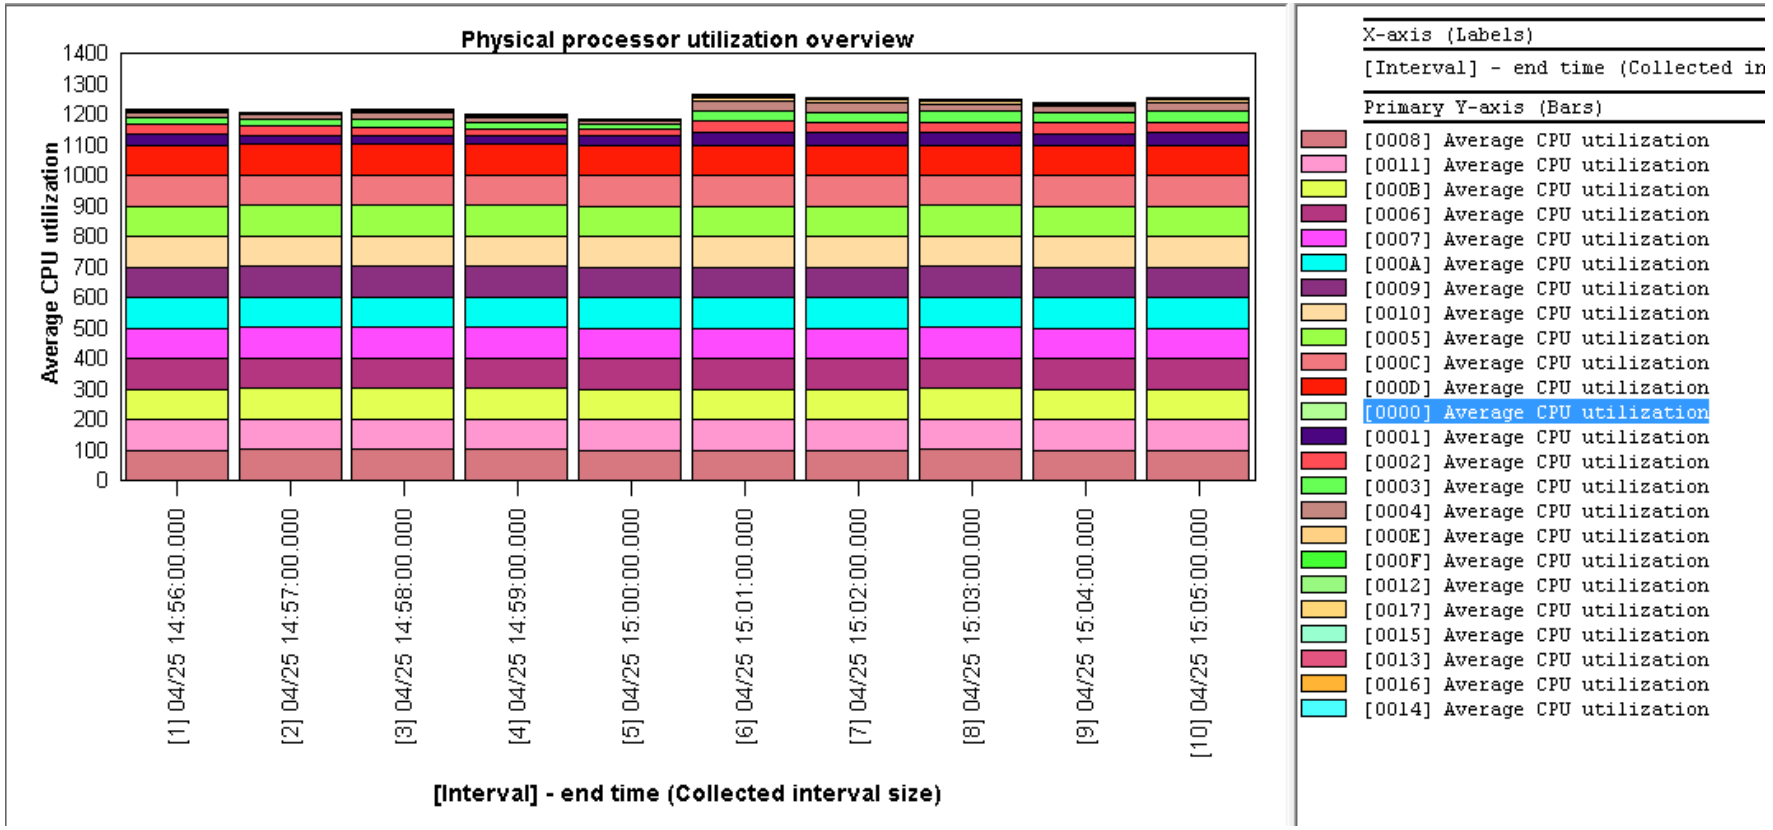

In CSI 6.1+, new graph shows the average CPU utilization for each processor ID over time.
Additional graphs are also available that filters the list of processors shown by processor state (guarded off, unlicensed, shared, borrowed and dedicated)

Apr–Sept 2011 – CSI - System graphs – LPAR dedicated processor utilization rankings

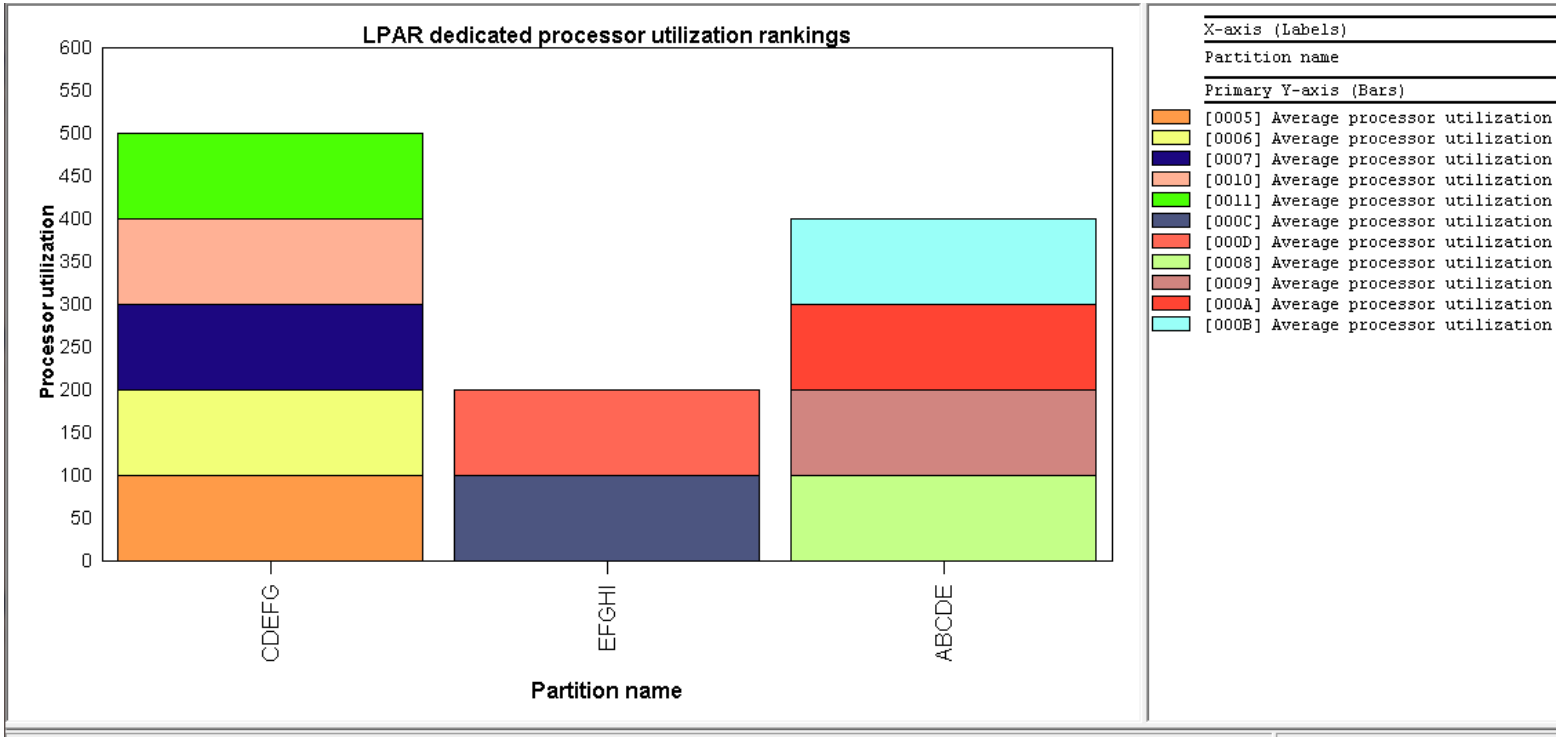

In CSI 6.1+, new graph shows a ranking of processor utilization by partition with a breakdown by processor ID.

Apr–Sept 2011 – CSI – Job counts graphs

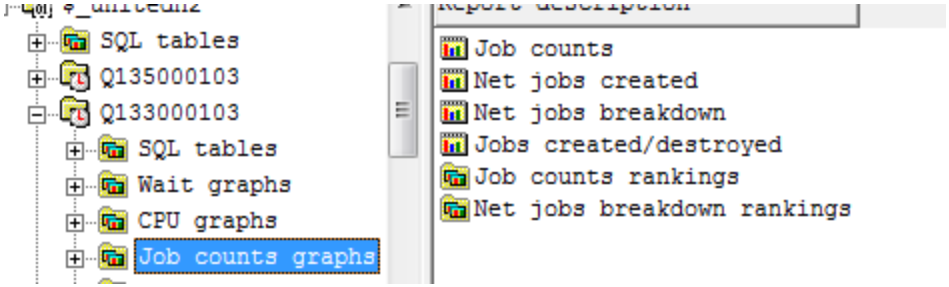

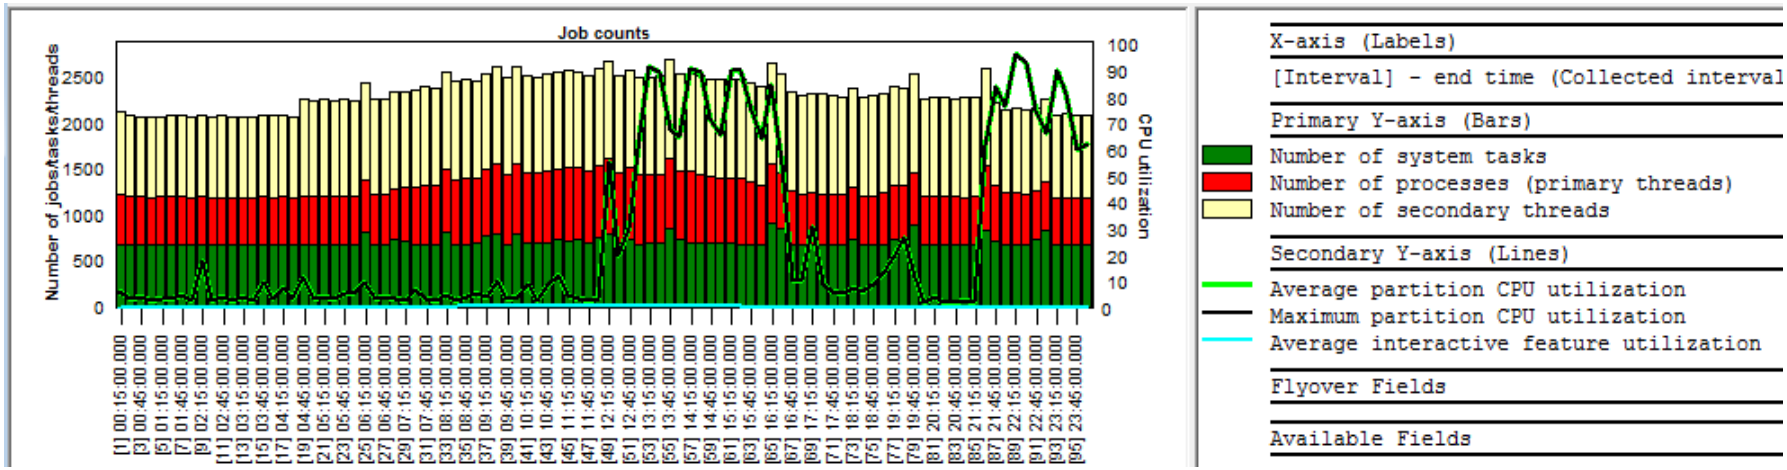

In CSI 6.1+, added a set of graphs under the "Job counts graphs" folder that shows the number of active or idle tasks/jobs/threads that occurred over time, or ranked by one of the job groupings. These graphs only appear if the collection has been summarized.

Note: The counts in these graphs are higher than those given in the CPU graphs because the CPU graphs do not include active jobs/tasks/threads that did not use CPU each interval.

Apr–Sept 2011 – CSI – Job counts by generic job name (example)

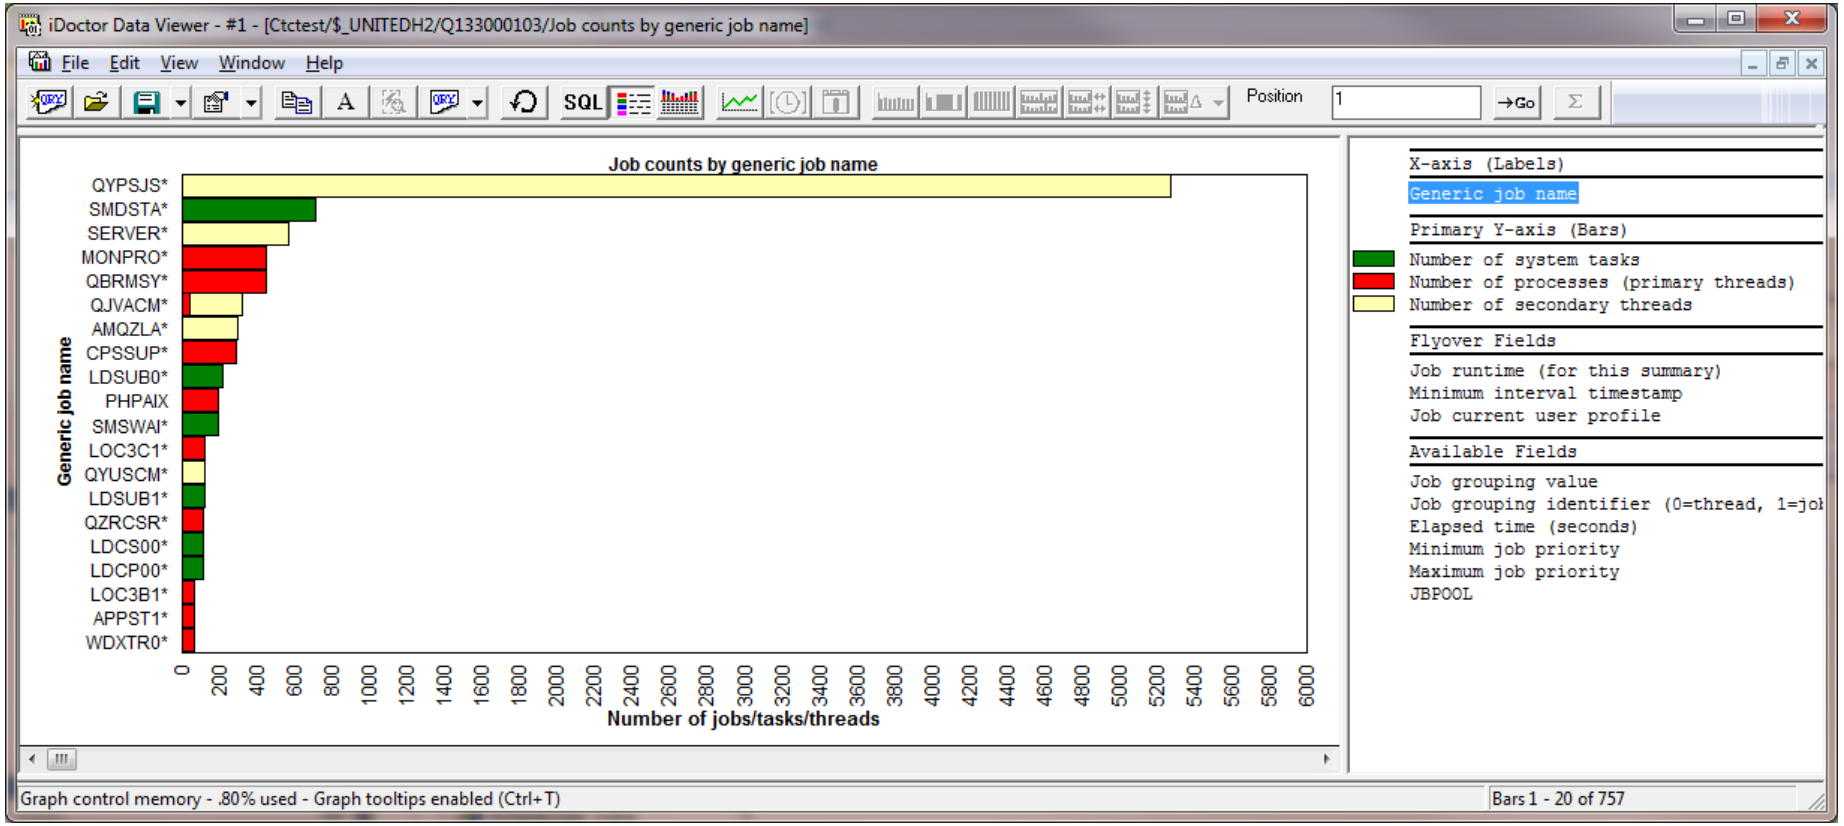

Above is an example of the Job counts by generic job name graph included under the Job counts graphs folder. Tasks are green, primary threads are red and secondary threads are yellow.

Apr–Sept 2011 – CSI – Memory pool graphs – DB vs non DB faulting rates

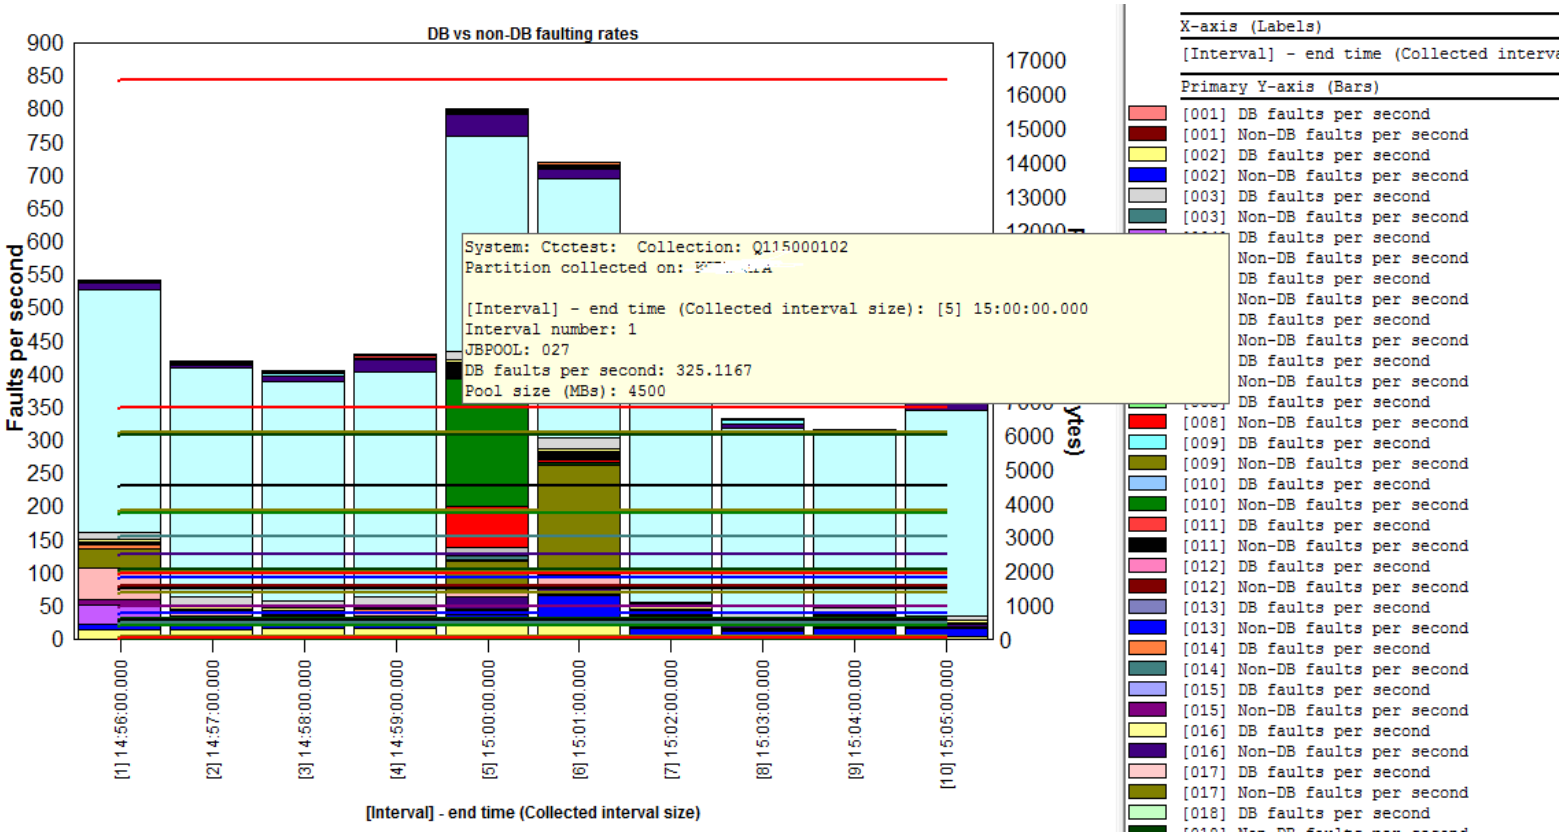

Within the CSI memory pool graphs, this graph shows the highest DB faults and non-DB faults for all memory pools on the system at once over time.

Apr–Sept 2011 – CSI – Memory pool graphs – DB vs non DB paging rates

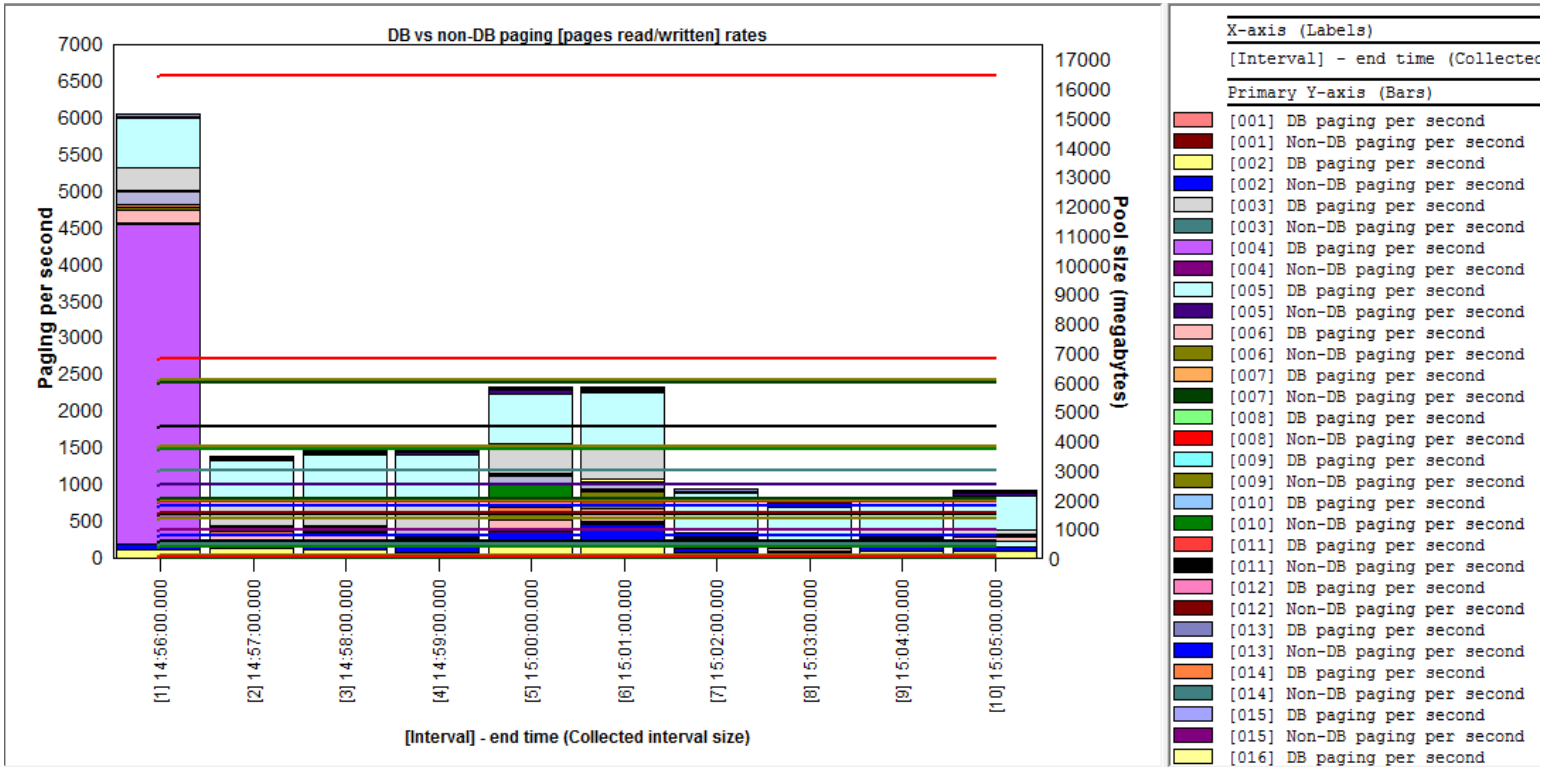

Within the CSI memory pool graphs, this graph shows the database vs non-database paging (4K pages read/written) for all memory pools on the system at once over time.

Apr–Sept 2011 – CSI – Memory pool graph drill downs

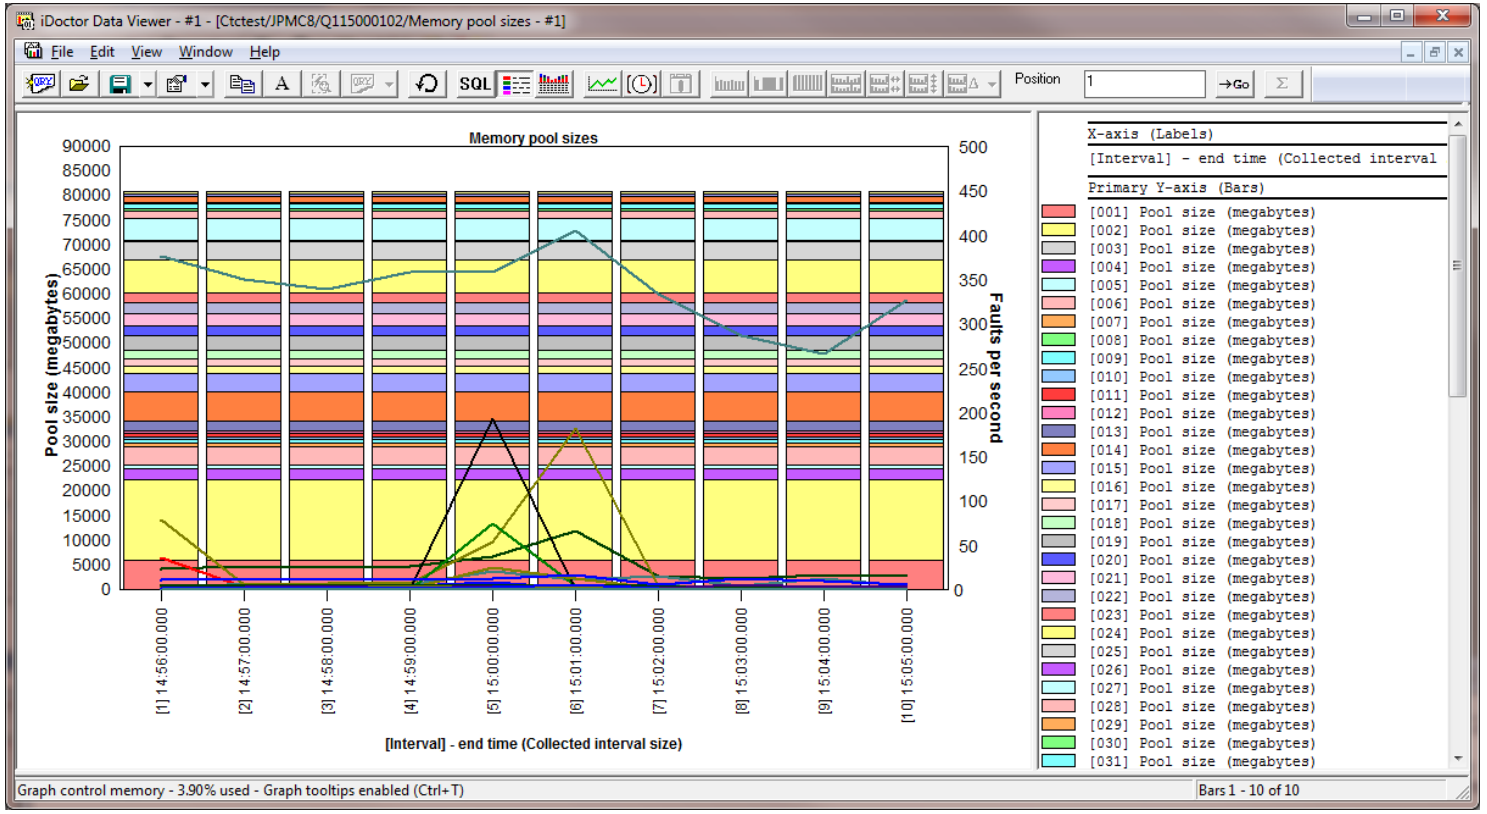

From the CSI memory pool graphs, drilling down into the rankings graphs for all graph types except disk graphs and virtual I/O graphs will now filter on the pool number that was right-clicked on.

Apr–Sept 2011 – CSI disk graphs – I/O size and Ethernet rates

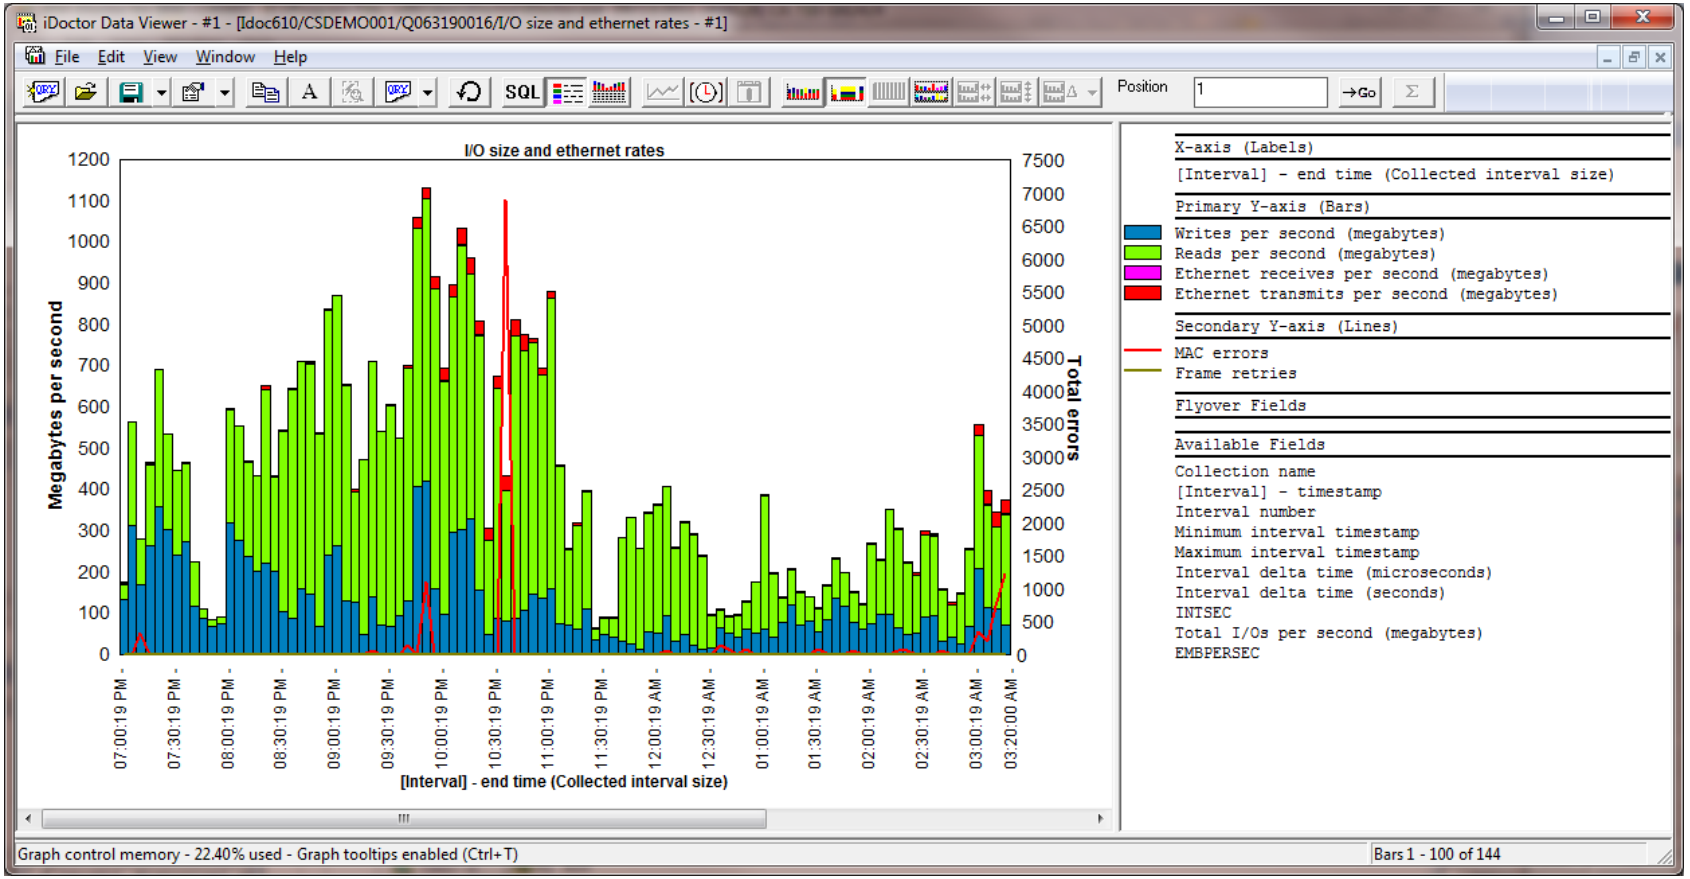

In CSI 5.4+, under the disk graphs , added a new graph called I/O size and ethernet rates that shows disk read sizes (MBs) per second, disk write sizes (MBs) per second , ethernet transfers (MBs) per second and ethernet receives (MBs) per second. The 2nd Y axis shows MAC errors and frame retries.

Apr–Sept 2011 – CSI disk graphs – I/O average service and wait time categorized totals

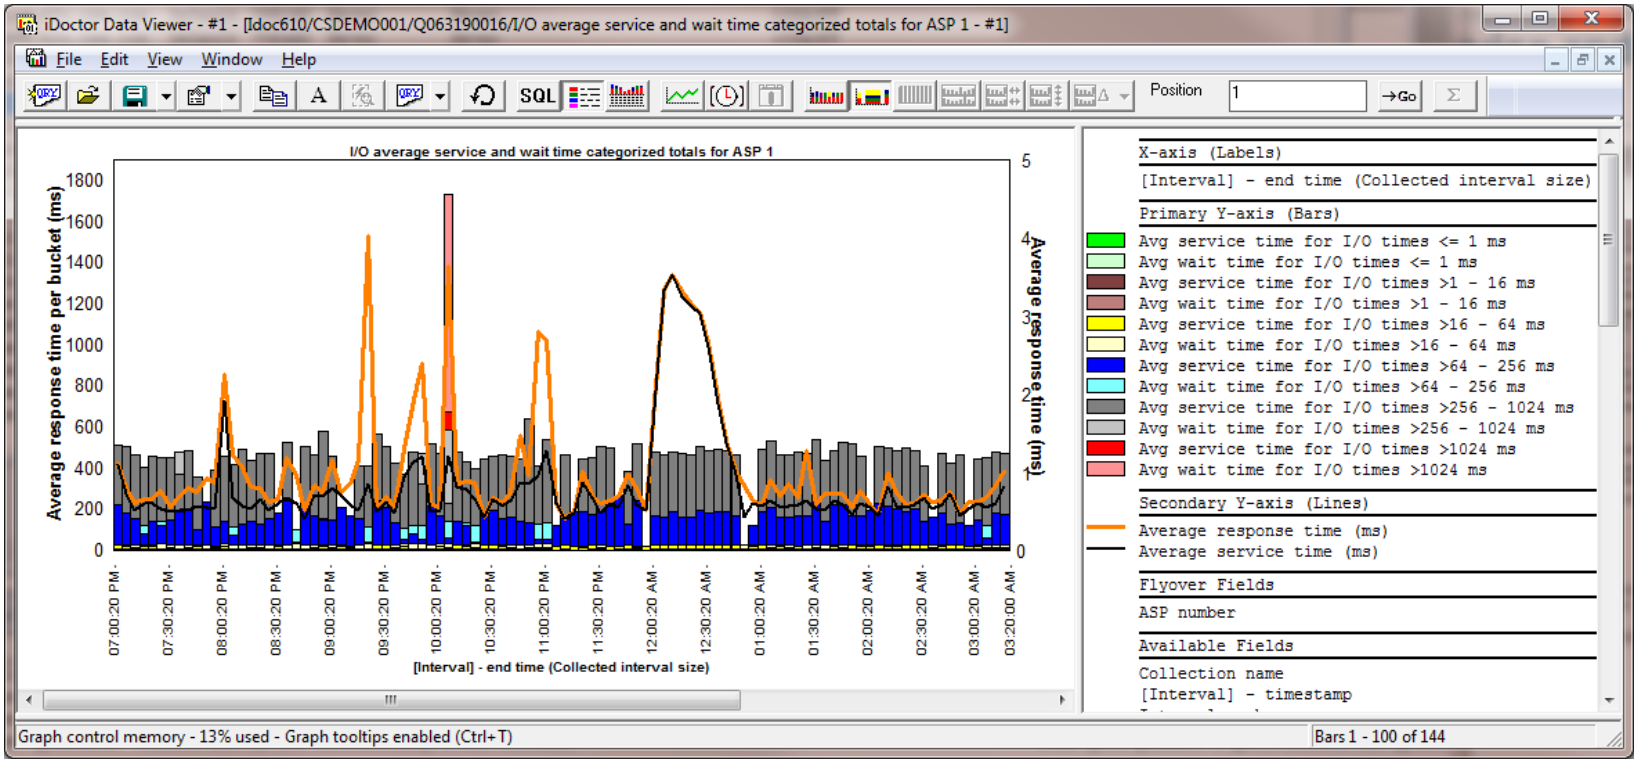

In CSI disk graphs at 5.4+ added a new graph at all 3 levels (overview, rankings, selected unit/path/etc over time) called "I/O average service and wait time categorized totals for ASP <<DSASP>>"

This graph is like the I/O average response time categorized totals graph but breaks each bucket down into wait time and service time components.

Apr–Sept 2011 – CSI communication graphs – RIO HSL 12x loops

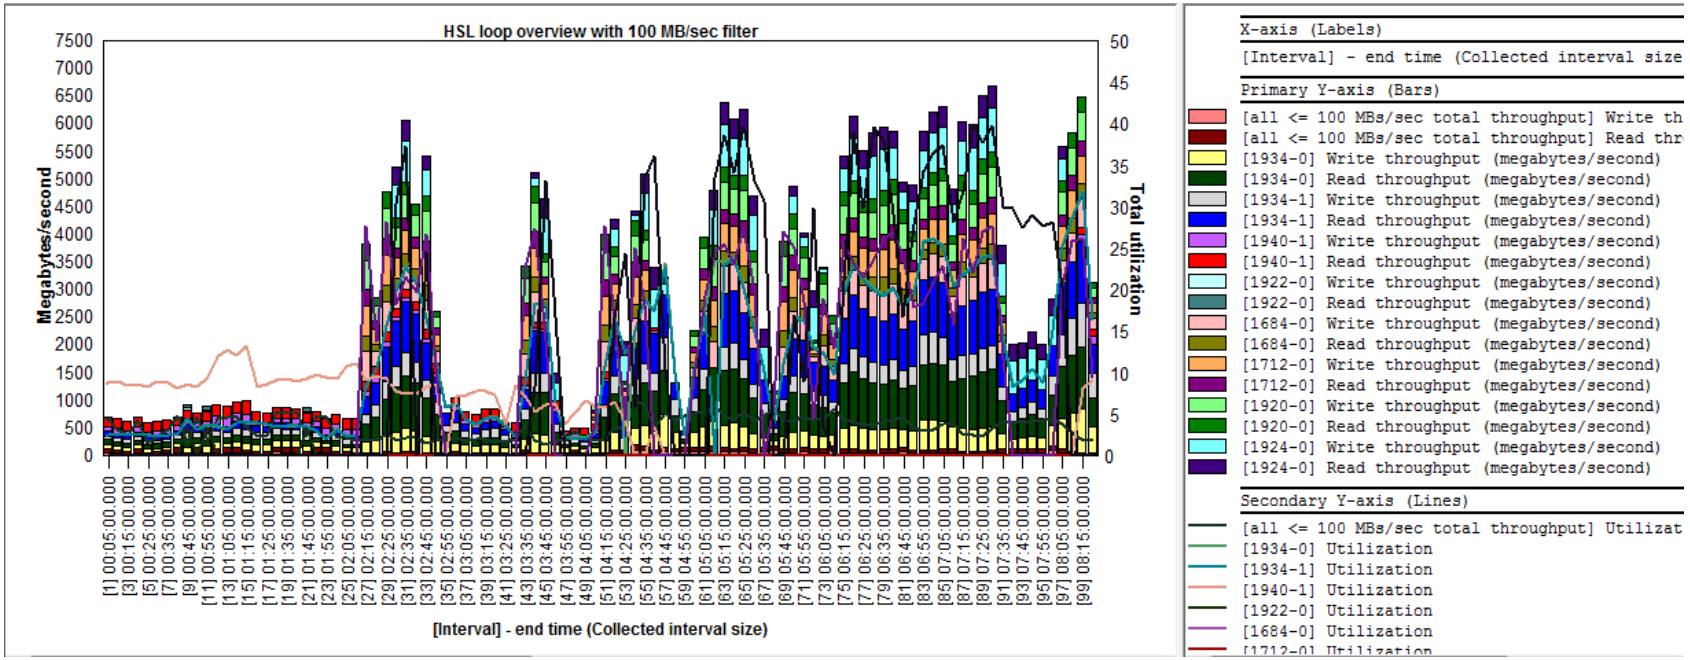

In CSI at 7.1, under the Communication graphs -> RIO HSL 12x loops folder, added a set of 12 graphs to show HSL read/write throughput in various ways. Options of using a filter of 0 MB/sec, 1 MB/sec, 100 MB/sec or 250 MB/sec are provided.

After any of these graphs are opened, the current filter can be modified by right-clicking the legend. In graphs where there are large number of loops, the usage of the filters is recommended to group smaller values together.

Apr–Sept 2011 – CSI communication graphs – SSL

In CSI, under the communication graphs folder added the graphs SSL authentication totals and SSL authentication rates to show SSL fast and full server (or server+client) handshakes. These graphs also show cpu utilization on the 2nd Y axis. If these numbers are high it can help determine the need for a 4764 Cryptographic Coprocessor

Apr–Sept 2011 – PEX – PDIO analysis

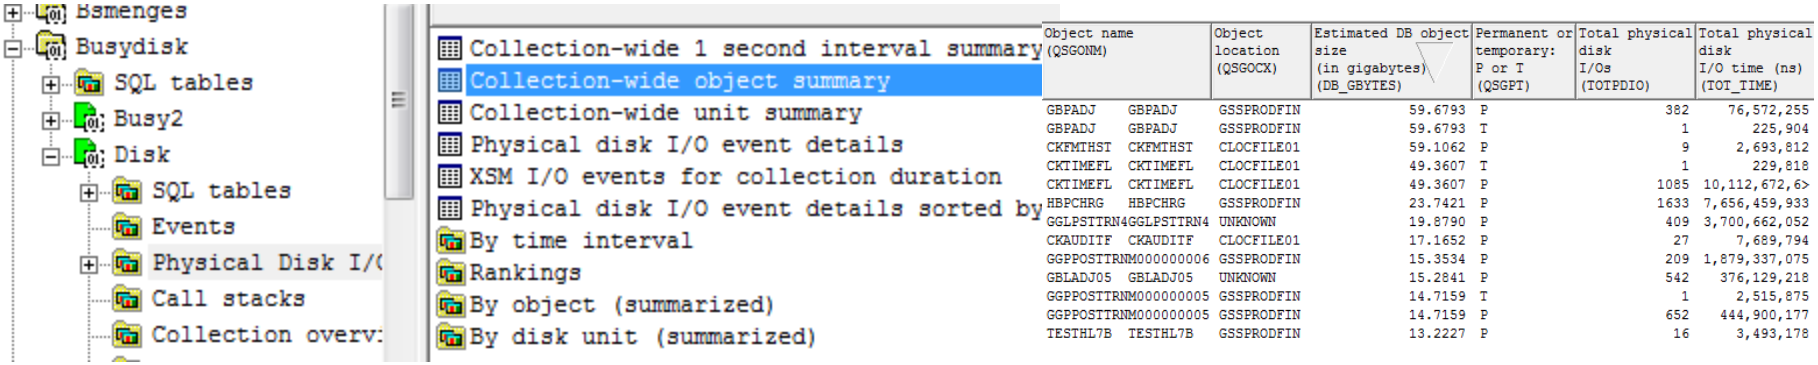

Object name (QSGONM)		Object location (QSGOCX)	Estimated DB object size (in gigabytes) (DB_GBYTES)	Permanent or temporary: P or T (QSGPT)	Total physical disk I/Os (TOTPDIO)	Total physical disk I/O time (ns) (TOT_TIME)
GBPADJ	GBPADJ	GSSPRODFIN	59.6793	P	382	76,572,255
GBPADJ	GBPADJ	GSSPRODFIN	59.6793	T	1	225,904
CKFMTHST	CKFMTHST	CLOCFILE01	59.1062	P	9	2,693,812
CKTIMEFL	CKTIMEFL	CLOCFILE01	49.3607	T	1	229,818
CKTIMEFL	CKTIMEFL	CLOCFILE01	49.3607	P	1085	10,112,672,6>
HBPCHRG	HBPCHRG	GSSPRODFIN	23.7421	P	1633	7,656,459,933
GGLPSTTRN4	GGLPSTTRN4	UNKNOWN	19.8790	P	409	3,700,662,052
CKAUDITF	CKAUDITF	CLOCFILE01	17.1652	P	27	7,689,794
GGPPOSTTRNM000000006		GSSPRODFIN	15.3534	P	209	1,879,337,075
GBLADJ05	GBLADJ05	UNKNOWN	15.2841	P	542	376,129,218
GGPPOSTTRNM000000005		GSSPRODFIN	14.7159	T	1	2,515,875
GGPPOSTTRNM000000005		GSSPRODFIN	14.7159	P	652	444,900,177
TESTHL7B	TESTHL7B	GSSPRODFIN	13.2227	P	16	3,493,178

In the collection-wide object summary report added a new field called “Estimated DB object size (in gigabytes)”

Added a new report "Physical disk i/o event details sorted by thread/task, time"

Apr–Sept 2011 – PEX – Call stacks analysis

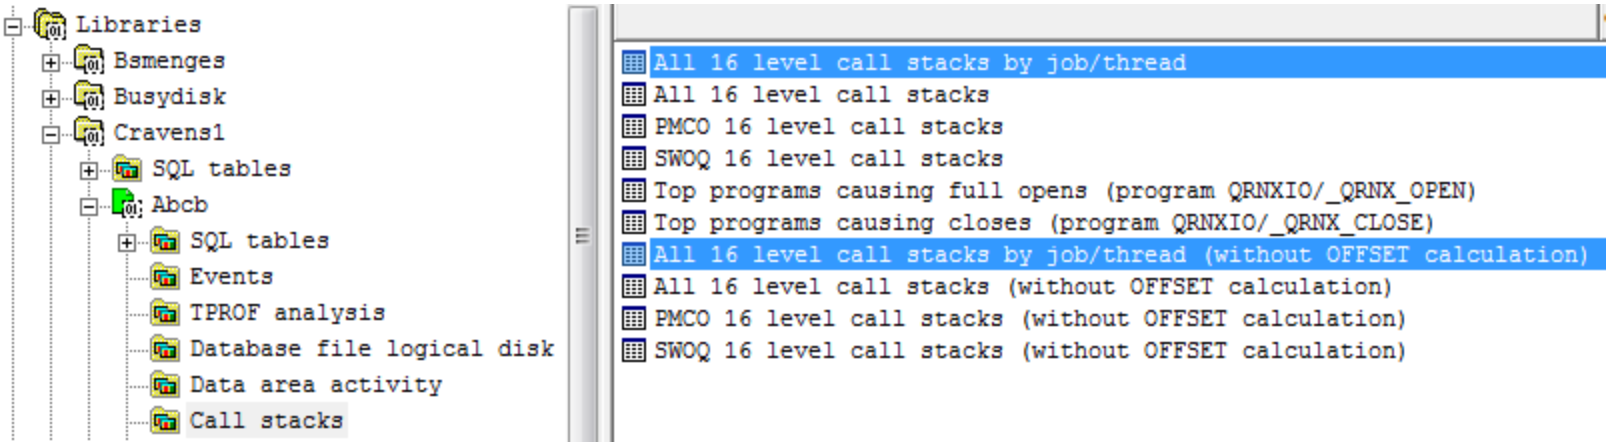

Added 2 new reports to the PEX call stacks analysis which gives the option to show the call stacks grouped by job/thread.

All 16 level call stacks by job/thread,
All 16 level call stacks by job/thread (without OFFSET calculation)

You can also right-click any of the call stack reports that do not group by job/thread and get 2 new reports to break down the selected call stack by job/thread.

Selected call stack -> Ranked by contributing jobs/threads
Selected call stack -> Ranked by contributing jobs/threads (includes selected call stack)

Apr–Sept 2011 – PEX – Trace Details analysis

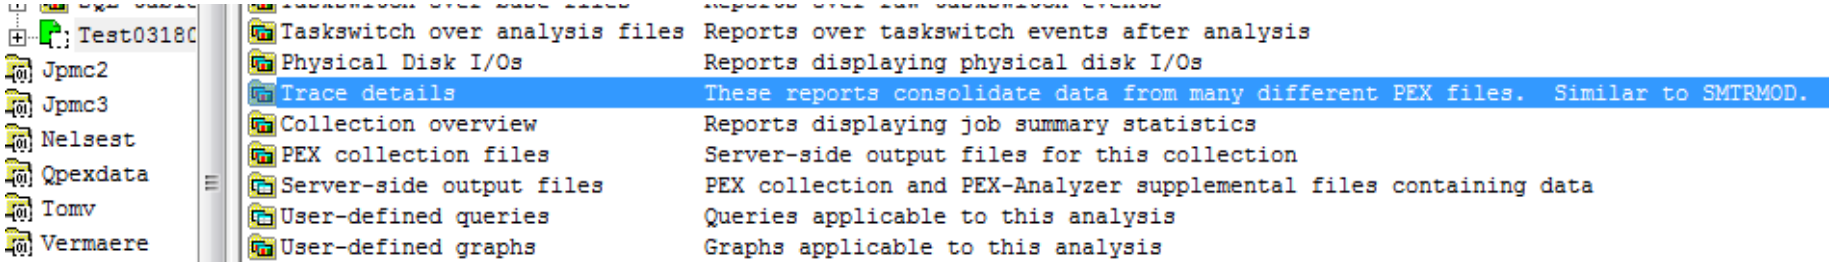

In PEX, if the trace details analysis has been ran, a new Trace Details folder will appear under the collection. This folder contains the following 2 reports over the SMTRMOD like output table:

Trace details sorted by time,
Trace details sorted by job/thread, time,

Apr–Sept 2011 – PEX – TPROF analysis

In the PEX TPROF analysis, made a fix where the OS component may not be completely filled causing the reports to show a blank component in some of the reports. Also fixed the same problem in the procedure XYZ by caller drill down.

Apr–Sept 2011 – PEX – New rate graphs

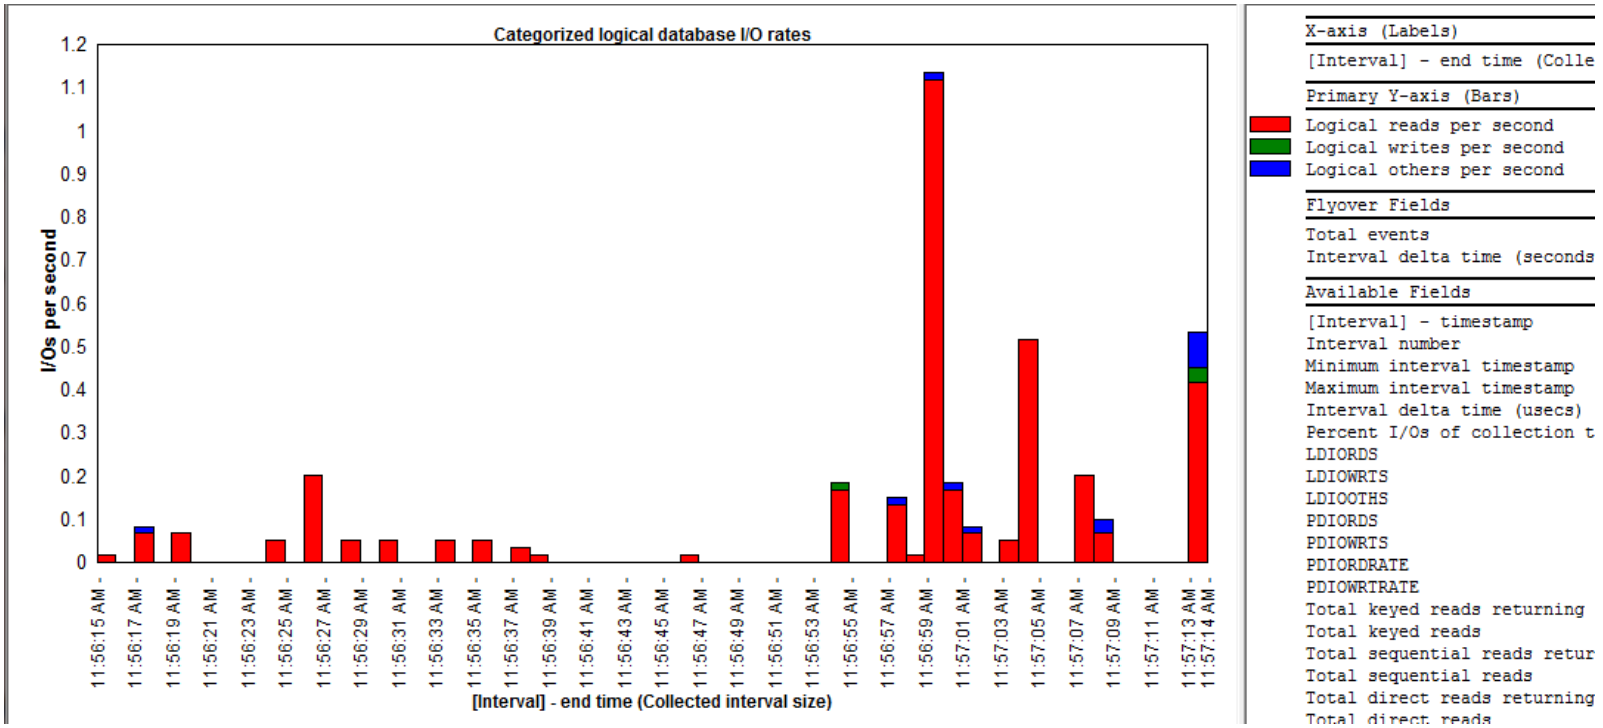

Added new graphs in PEX Analyzer to show rates in the following analyses:

1. Database opens/closes,
2. Database file logical disk IO
   (New graphs Categorized logical database I/O rates, Logical and physical I/O rates)
3. Data area activity
4. Data queue activity

Apr–Sept 2011 – PEX – Collection Properties General Tab

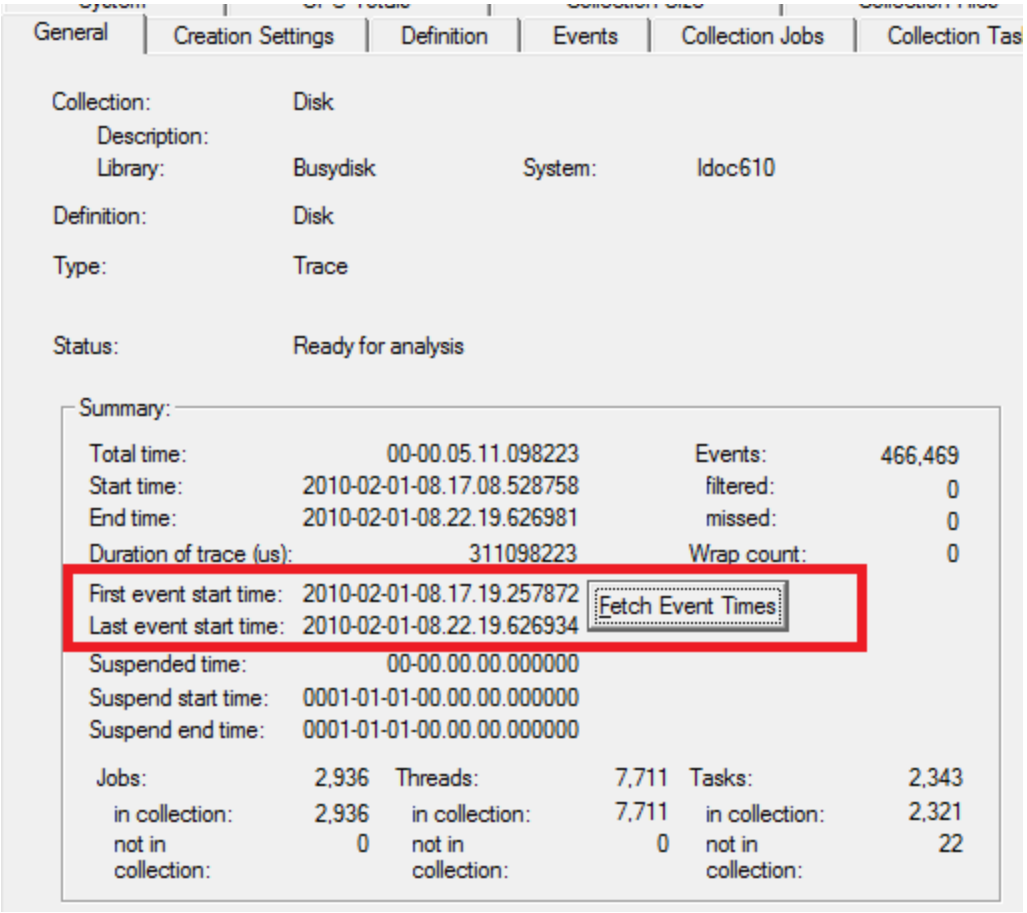

In PEX Collection properties for trace collections only, the first event time, and last event time will be listed on the general properties tab. To retrieve them, press the Fetch Event Times button next to these labels.

Apr–Sept 2011 – PEX Analysis Time Filtering

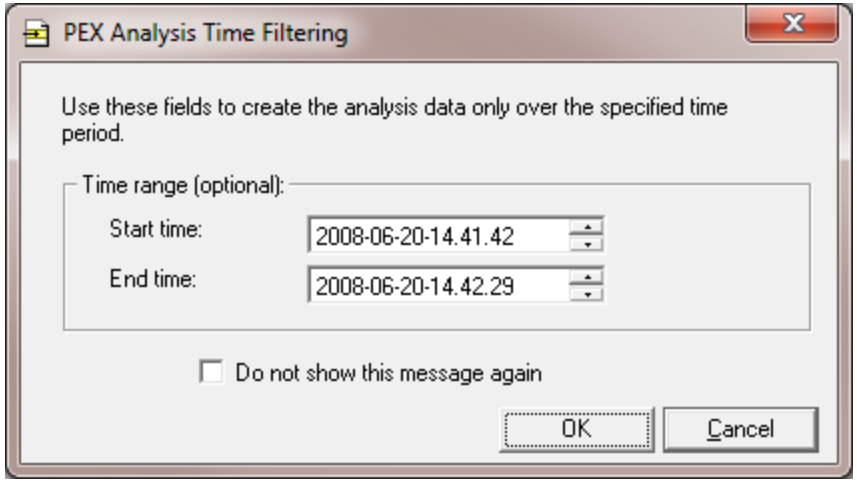

Most of the PEX Trace analyses now provide a prompt where you can filter the data by a desired time range.

If you rarely/never need to do this, you can check the box on the screen to turn this feature off. Turn it back on again under Preferences -> Confirm tab,

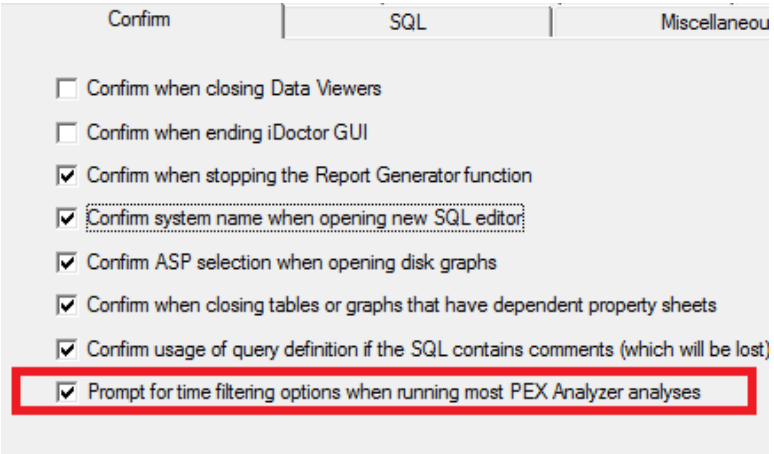

Power your planet.

Apr–Sept 2011 – PEX Monitor Updates

---> C00862 - Job Watcher - Bug Fix Request
At 5.4+, the PEX monitor included with Job Watcher (QIDRWCH/STRPAMON command) was running ENDPEX at priority 1 instead of 50.

---> C00842 - Job Watcher - Bug Fix Request
Fixed a possible "MCH0603 received by QIDRPEXCNT" error when running a PEX monitor (STRPAMON command).

Also fixed a bug with STRPAMON if using the *MGTCOL option with the latest server builds.

Latest server builds required.

Apr–Sept 2011 – Client Access ODBC Driver issues

Tested all Client Access service pack levels at 5.4 and higher and disabled iDoctor usage with those service pack levels that failed due to various ODBC driver issues.

The disabled Service pack levels are:
5.4: GA, SI20465 (errors related to being unable to add the sql default schema to library list)
7.1 GA, SI36916, SI37895 (invalid cursor state listing collections in CSI)

We recommend upgrading to 7.1 SP3 or higher.

Apr–Sept 2011 – Situational Analysis On/Off Toggle

On the main window toolbar added a button that enables/disables the display of situational analysis background colors in graphs. A simple click on the button will turn it on/off for all graphs (even open ones). Another click on the graph or legend will redraw the graph with/without the situations (if found in the data).

Apr–Sept 2011 – Performance Improvements

Creating ODBC connections is faster.

Opening a component is much faster the 2nd time because stored procedure version numbers are now cached.

If the SQL preference -> Use sql catalog tables to improve performance is enabled, opening the server-side output files folder could be slow on some systems with a large number of members. Modified the SQL so it now runs much faster.

Apr–Sept 2011 – Copy button on Property Sheets

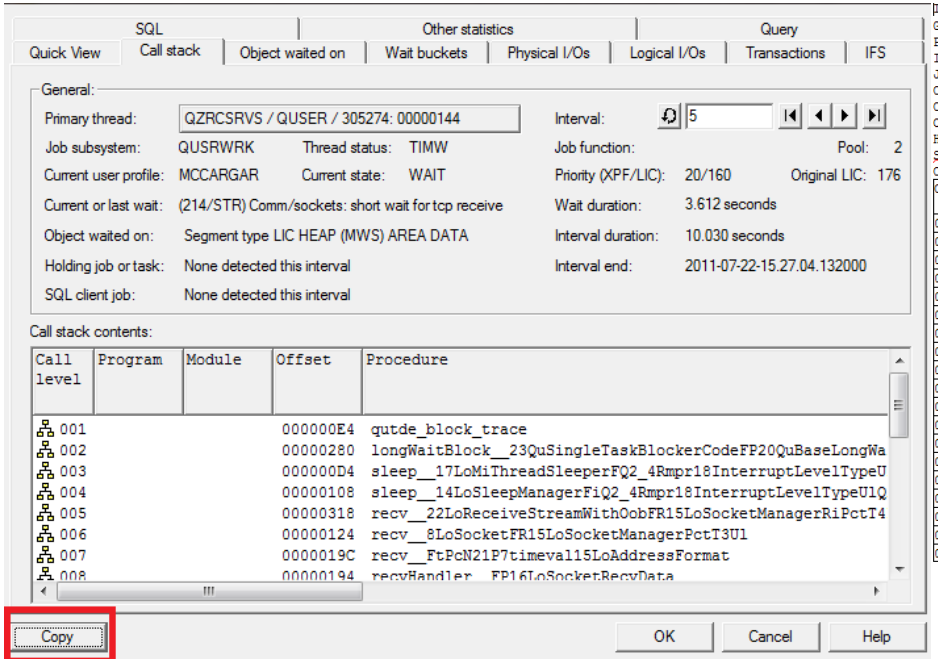


```
Interval Details:  System Idoc610, Library Cravens1, Collection Q203152612
General:
Primary thread:      QZRCSRVS / QUSER / 305274: 00000144
Interval:     IDB_REFRESH  5       0            IDB_FIND_NEXT
Job subsystem:       QUSRWRK       Thread status:       TIMW   Job function:        Pool:  2
Current user profile:       MCCARGAR       Current state:       WAIT   Priority (XPF/LIC):  20/160 Original LIC: 176
Current or last wait:       (214/STR) Comm/sockets: short wait for tcp receive      Wait duration:       3.612 seconds
Object waited on:    Segment type LIC HEAP (MWS) AREA DATA    Interval duration:   10.030 seconds
Holding job or task: None detected this interval Interval end: 2011-07-22-15.27.04.132000
SQL client job:      None detected this interval
Call stack contents:
```

Call level	Program	Module	Offset	Procedure
001			000000E4	qutde_block_trace
002			00000280	longWaitBlock__23QuSingleTaskBlockerCodeFP20QuBaseLongWaitObjectR
003			000000D4	sleep__17LoMiThreadSleeperFQ2_4Rmpr18InterruptLevelTypeUlQ2_4Rmpr
004			00000108	sleep__14LoSleepManagerFiQ2_4Rmpr18InterruptLevelTypeUlQ2_8TDQSEnu
005			00000318	recv__22LoReceiveStreamWithOobFR15LoSocketManagerRiPctT4P8sockaddr
006			00000124	recv__8LoSocketFR15LoSocketManagerPctT3Ul
007			0000019C	recv__FtPcN21P7timeval15LoAddressFormat
008			00000194	recvHandler__FP16LoSocketRecvData
009			000000E4	socketop
010			000000E8	#cfmir
011			0000012C	syscall_A_portal
012	QSOSRV1	QSOSYS	00000098	recv
013	QZBSCOMM	QZBSCOMM	00000190	QzbsReceiveClientReq
014	QZRCSRVS	QZRCSRVS	000000C0	RcvClientReq
015	QZRCSRVS	QZRCSRVS	0000040C	main
016	QZRCSRVS	QZRCSRVS	00000290	_C_pep
017			000001D0	cblabranch
018			0000005C	aimach_program_call_portal
019			000007B8	pmInitiateProcessUnderTarget__Fv

The new Copy button found on property sheets will now copy the contents of the call stack, quick view to the clipboard in Rich Text Format. This can be useful in the CPS database to enable text searching against call stack contents and more.

Apr–Sept 2011 – Main Window Title Bar Changes

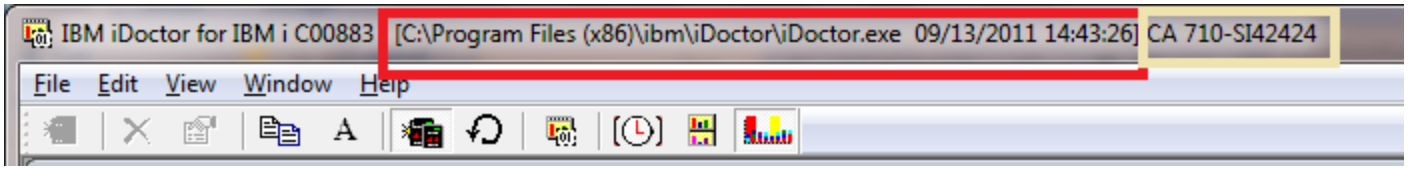

The main window title bar now shows the following information in addition to the build number:

iDoctor.exe location and change date/time (red above)
Client access VRM and fix level (yellow above)

Apr–Sept 2011 –Work Management Folder

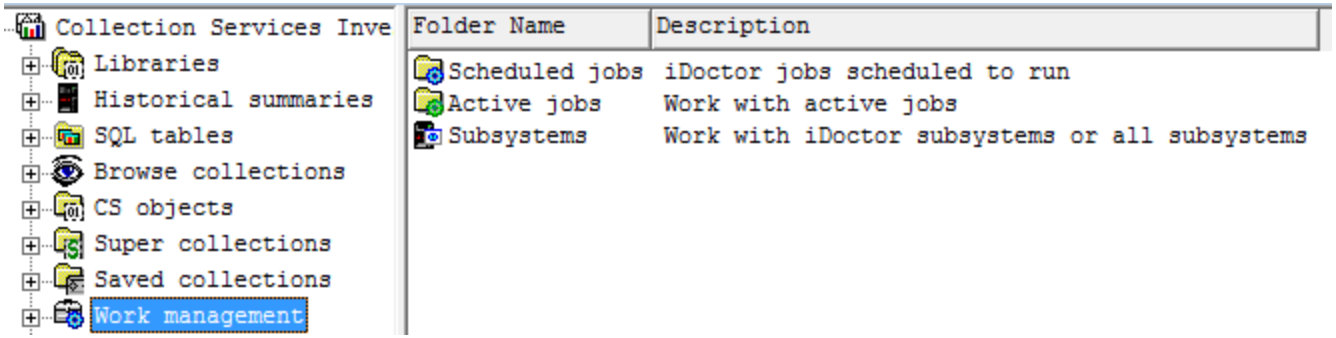

Added a new Work Management folder that contains: Scheduled Jobs, Active Jobs and Subsystems.

The Subsystems folder lets you work with either the jobs in the iDoctor subsystems or All subsystems. QIDRJW for example contains all the jobs submitted by a monitor and this interface can be used to check their status.

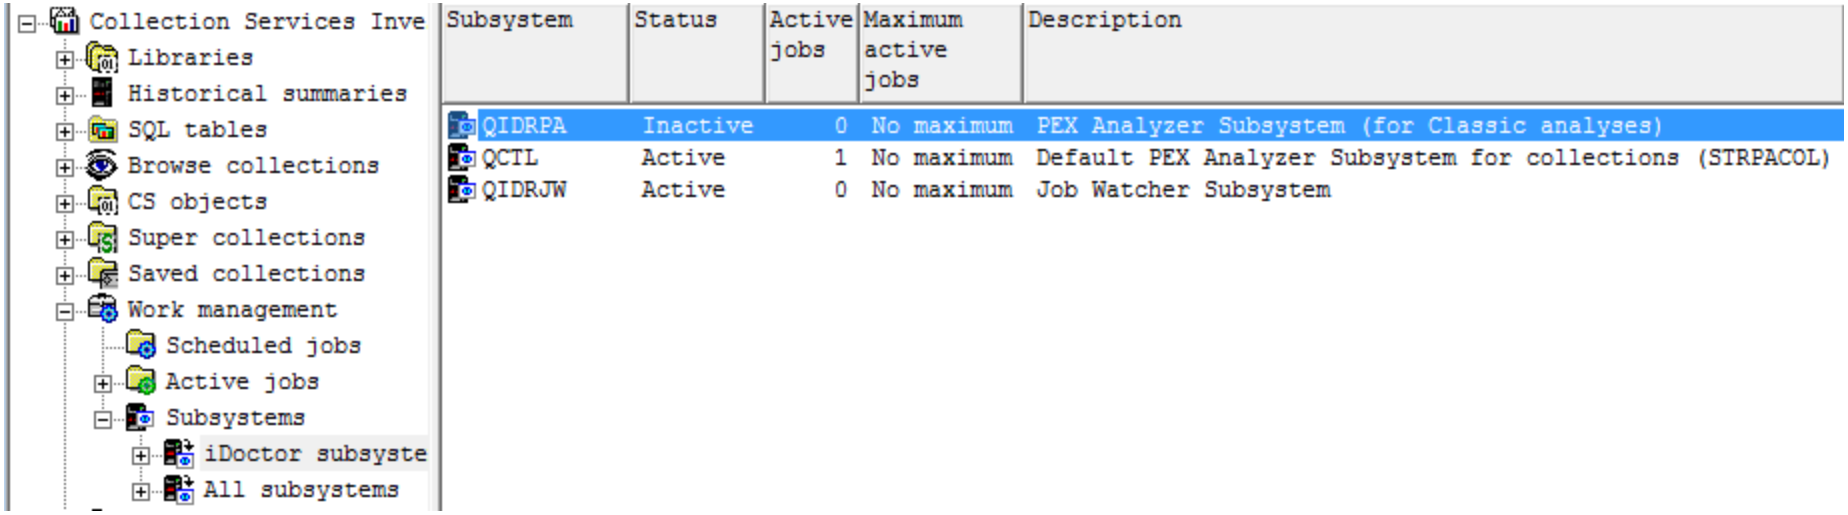

Power your planet.

Apr–Sept 2011 – Monitor Updates

---> C00862 - Job Watcher - Bug Fix Request
Previously in all start monitor commands (STRPAMON, STRJWMON, STRDWMON) at 5.4 and higher, if the monitor failed immediately then restarting the same monitor would not be possible.

---> C00845 - Disk Watcher;Job Watcher;Pex Analyzer - Bug Fix Request
At 5.4+, command QIDRGUI/ADDIDRUSR was not giving *USE authority to the following commands required in order to run monitors properly without additional authorities:

STRPEX
ENDPEX
STRDW
ENDJW (6.1+ only)
ENDDW (6.1+ only)

Also changed QIDRGUI/RMVIDRUSR to revoke the authority to the additional commands that ADDIDRUSR now provides.

Previously, users that started a monitor without authority to these commands would not be able to run or end the monitor properly depending on the monitor type.

Apr–Sept 2011 – Clock Icon changes

The clock icon now supports new additional time range groupings of: 5 seconds, 15 seconds, 8 hours, 12 hours and 24 hours.

In the future we most likely will add weekly and monthly groupings.

Also in order to eliminate redundant choices, the clock icon in the Data Viewer will now only show possible time groupings based on the collection’s interval duration and the collection's total run time (if available).

Apr–Sept 2011 – Running Analyses from Libraries

In JW, CSI and PEX, you can now right-click one or more libraries and use the Analyses menu in order to run the desired analysis against all collections in all selected libraries in one step.

Use the new preference on the Miscellaneous tab if you wish the analysis to run in a batch job.

Also related to this the Create Job Summary function now is an analysis and behaves like one. (It shows up in the Analyses menu, etc.)

Apr–Sept 2011 – App Properties – iDoctor client jobs tab

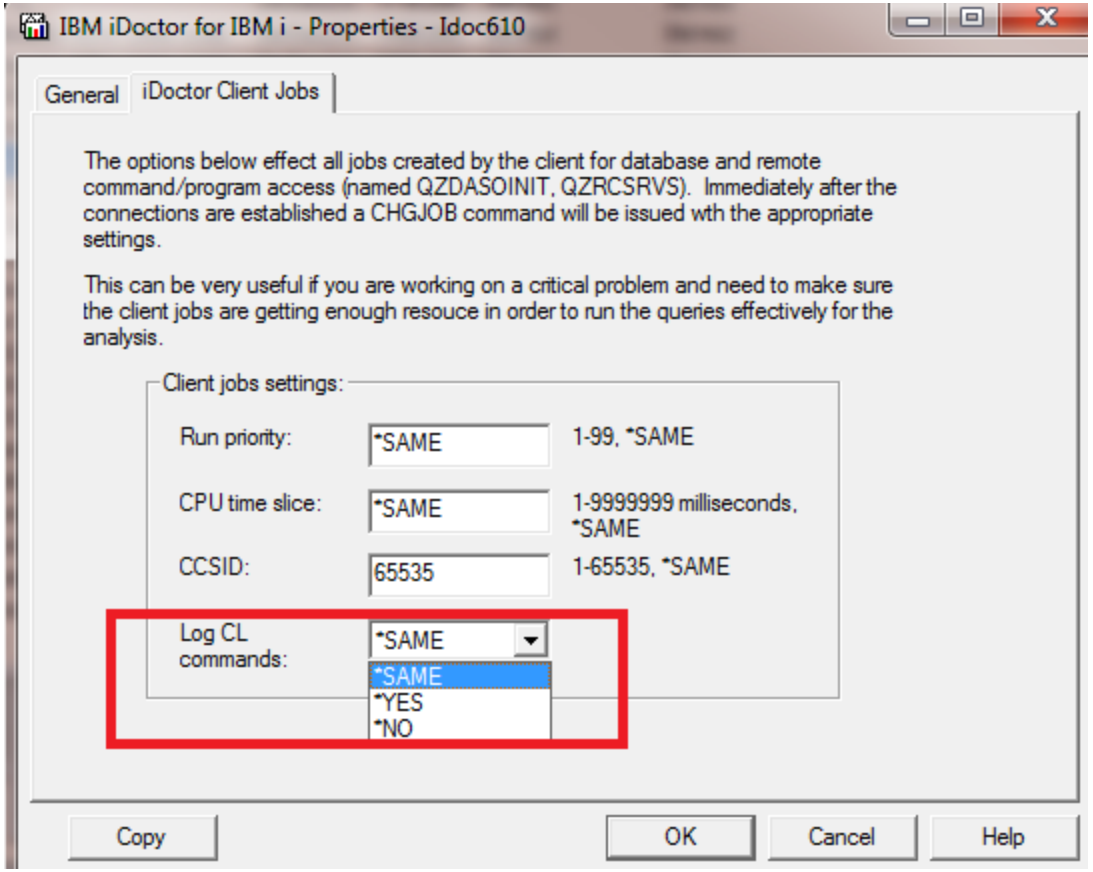

In the application properties, iDoctor client jobs tab, added a new option called “Log CL commands". If set to something other than *SAME, then the CHGJOB command parameter by the same name will be used at startup of new iDoctor jobs.