

iDoctor What's New June – Oct 2010



Ron McCargar
iDoctor Development

iDoctor update resources

iDoctor e-mail list: usage tips, build updates, PTF info

iDoctor update history: embedded into the GUI. Search the list on the topic of interest.

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>

* DS6000/DS8000 support was added to 7.1 and 6.1 with Collection Services Investigator in April (links/ranks/cache statistics). See page 213 (actual page 235 within Acrobat).

Jun-Oct 2010 - Overview

Current Focus Areas

Latest builds

Collection Summary enhanced (in CSI and JW)

SQL Tables/Comparison enhancements

Create Job Summary enhanced (in CSI and JW)

Visual Studio 2008 redistributable now a GUI install prerequisite

Monitors usability enhancements

Miscellaneous

Jun-Oct 2010 – Current Focus Areas

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

Comparing data in iDoctor more effectively:

- SQL table support

- Quick view comparisons

- Additional analyses and graphs

- Goal: Reduce/eliminate users needing to go to tools like Excel

External storage

Tom McBride working on server-side VIOS data collection enablement.

Surface new metrics/features added to OS

Jun-Oct 2010 – Comparison Graph Example

The following graph was done by Helen Olson-Williams, and shows how batch workloads performed differently in a customer's application based on several different factors:

Graph label details:

1GB - memory pool size where job is running

12dsk - number of disk drives

1.0P - number of processing units, i.e. number of cores

Jcache - journal cache attribute set to *yes

ASPB - we had used trcaspbal and straspbal *hsm to move hot data to SSDs prior to this run


TRCASP - this was the run where we had trcaspbal running to collect info on the hot sectors

PStats - we were running PEX STATS at this time so run probably wasn't valid from a comparison point of view

2IOAs - we added an IOA to the config and spread the drives across 2 IOAs instead of 1.

12DSK + SSDs - we used media preference to move hot data to the SSDs prior to this run.

Jun-Oct 2010 – Comparison Graph Example



Jun-Oct 2010 – Comparison Graph Example

With the current builds, the following steps are needed:

1. Run a create job summary for each job/generic job you wish to analyze.
2. Use SQL to merge multiple create job summaries into 1 table.
`INSERT INTO TABLE1 (SELECT * FROM TABLE2)`
3. Alter the table adding a description column
`ALTER TABLE TABLE1 add column desc varchar(100)`
4. Modify the description for each desired row or set of rows using SQL.
`update TABLE1 set desc = 'COL 1 INFO' where mbr = 'COL1'`
5. Open the Create Job Summary job rankings graph over the 'merged' results table.
6. Modify the SQL to include DESC column in the output and change the legend so the label shown on graph is the DESC column.

In the future we'd like to make this easier, by adding built in options to merge SQL tables, adding description columns within the SQL tables, and add GUI options to modify the descriptions.

Jun-Oct 2010 – New Builds

Released 2 rounds of external builds in Sept, and again on Oct 13th.

Plan is to not update again for a several months unless there is a critical issue.

Generally try to do 3 or 4 major updates a year.

Jun-Oct 2010 – Collection Summary vs Create Job Summary

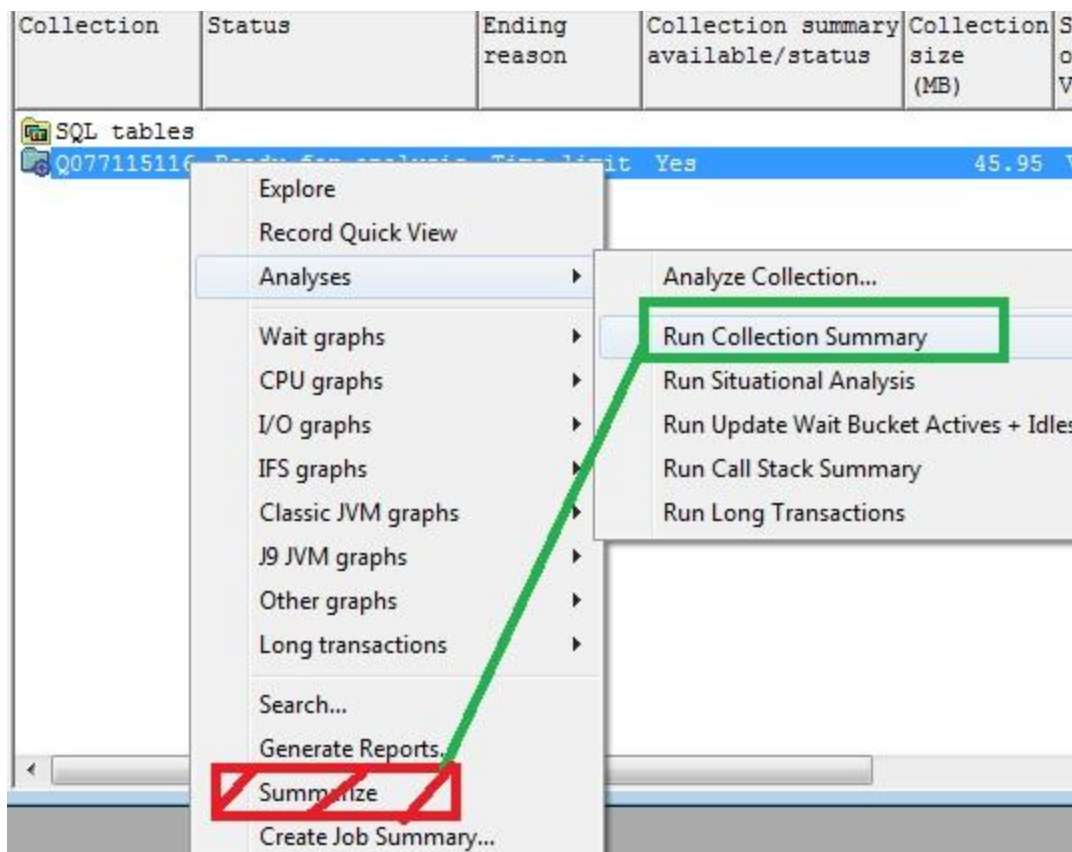
Feature	Collection Summary	Create Job Summary
SQL Tables format	By time interval	By job or thread
Includes 1 record totals SQL table	Yes (CSI only)	No
Analyze multiple collections (per table)	No	Yes
Rankings graphs	Yes, over filtered collection data. Might be slow.	Yes, over SQL tables. FAST!
By time graphs	Yes	No*
Filter by job, subsystem, time	Yes	Yes
Filter by current user	Yes	No
Filter by pool	Not yet	Not yet

* No over time graphs are available in Create Job Summary, except when drilling down in Job Watcher to graph a single job/thread over time.

Jun-Oct 2010 – Collection Summary (JW and CSI)

“Summarize” menu moved under “Analyses” -> “Run Collection Summary”.

Also accessible from “Analyses” -> “Analyze Collection” if you want to run it in batch.



Jun-Oct 2010 – Collection Summary (JW and CSI)

“Analyses” -> “Analyze Collection” menu will now show Analyze Collection(s) window.

Preferences button shows the Situational Analysis options.

The screenshot shows the 'Analyze Collection(s)' window with the following content:

This interface allows you to select which analysis functions should be performed for the selected collection(s). Additional reports will be provided after performing this option.

Functions available: Preferences... Default Toggle Selected

Description	Used by
<input checked="" type="checkbox"/> iDoctor Interval Summary	
<input checked="" type="checkbox"/> Situational Analysis	
<input type="checkbox"/> Update Wait Bucket Actives + I	
<input type="checkbox"/> Call Stack Summary	

Submit this request to a batch job instead of u

The 'Situational Analysis Preferences' dialog box is open, showing the following content:

Use these options to indicate which situations should be ran and the filters to be applied (where changing the filter values you can increase or decrease the likelihood of a situation occurring.

Situational Analysis Available: Default

ID	Situation	Filter
<input checked="" type="checkbox"/> 1	Seize/lock table	
<input checked="" type="checkbox"/> 2	Starting/ending commitment control	5
<input checked="" type="checkbox"/> 3	Poorly written/performing SQL	100
<input checked="" type="checkbox"/> 4	Missed jobs	.05
<input checked="" type="checkbox"/> 5	Seize contention due to data forced to disk	
<input checked="" type="checkbox"/> 6	Fixed length of varchar or blob too small	
<input checked="" type="checkbox"/> 7	High number of opens/closes	
<input checked="" type="checkbox"/> 8	Contention on user profile	
<input checked="" type="checkbox"/> 9	High synchronous write response time	3
<input checked="" type="checkbox"/> 10	Concurrent write support not enabled	
<input checked="" type="checkbox"/> 11	Journal cache could help performance	

Situation Filter:

Jun-Oct 2010 – Collection Summary (JW)

In order to simplify the Job Watcher summary, the “Collection Summary” analysis now **always** includes the following steps from the Summarize Window found in earlier builds:

- Create interval summary
- Create process interval summary (QAPYJWPRC stats)
- Create wait bucket gap file (i.e. actives + idles)
- Create collection aliases

The remaining steps have now become new Job Watcher “analyses”. Those are:

- Run situational analysis
- Run update wait bucket actives + idles
- Run call stack summary
- Run long transactions (at 6.1+ only)

Jun-Oct 2010 – Collection Summary (JW and CSI)

Filters are OPTIONAL but necessary if you want to do comparisons under the SQL tables folder.

If filters are NOT used then Summarized status will become “Yes” and the normal summarized graphs are used under the collection.

If filters are used then you can access graphs over the filtered Collection Summary SQL table generated. **Note:** *Currently in order for all the drill downs to work properly you must also summarize the entire collection without using any filters.*

Client 805 or higher needed

Run Collection Summary - Rchaskmb

This option will produce interval summary SQL tables for the following collection(s):

Collection name
BENFIELD/Q268000009

Selected collections shown here

NOTE: If you do not wish to filter the data just press the Submit button to continue:

If no filters are used, then the summarized tables will be utilized as the default set of iDoctor graphs.

SQL Tables filter: (OPTIONAL):

Job name: contains []

Job user name: []

Job number: []

Job current user profile: []

Subsystem name contains: []

Start time: 2010-09-25-00.00.09

End time: 2010-09-26-00.00.00

Comments: **Comment attached to SQL table**

Submit Cancel

Jun-Oct 2010 – Collection Summary (JW)

JW interval summary SQL table example:

Description	Library	Collection(s)	VRM	Comments	Job	Current user	Subs...	Start	End	SQL Table Name
Interval summary file	CRAVENS1	Q077115116	V6R1M0	QZDA jobs only	QZDA/*/*					Qaidrjwsum_q077115116
Interval summary file	CRAVENS1	Q077115116	V6R1M0							aidrjwsum_q077115116

In some cases (Create Job Summary, there will be multiple collections listed)

Optional filters

In the future there will be options to put the comments directly on the graph if doing a comparison

Actual SQL table name

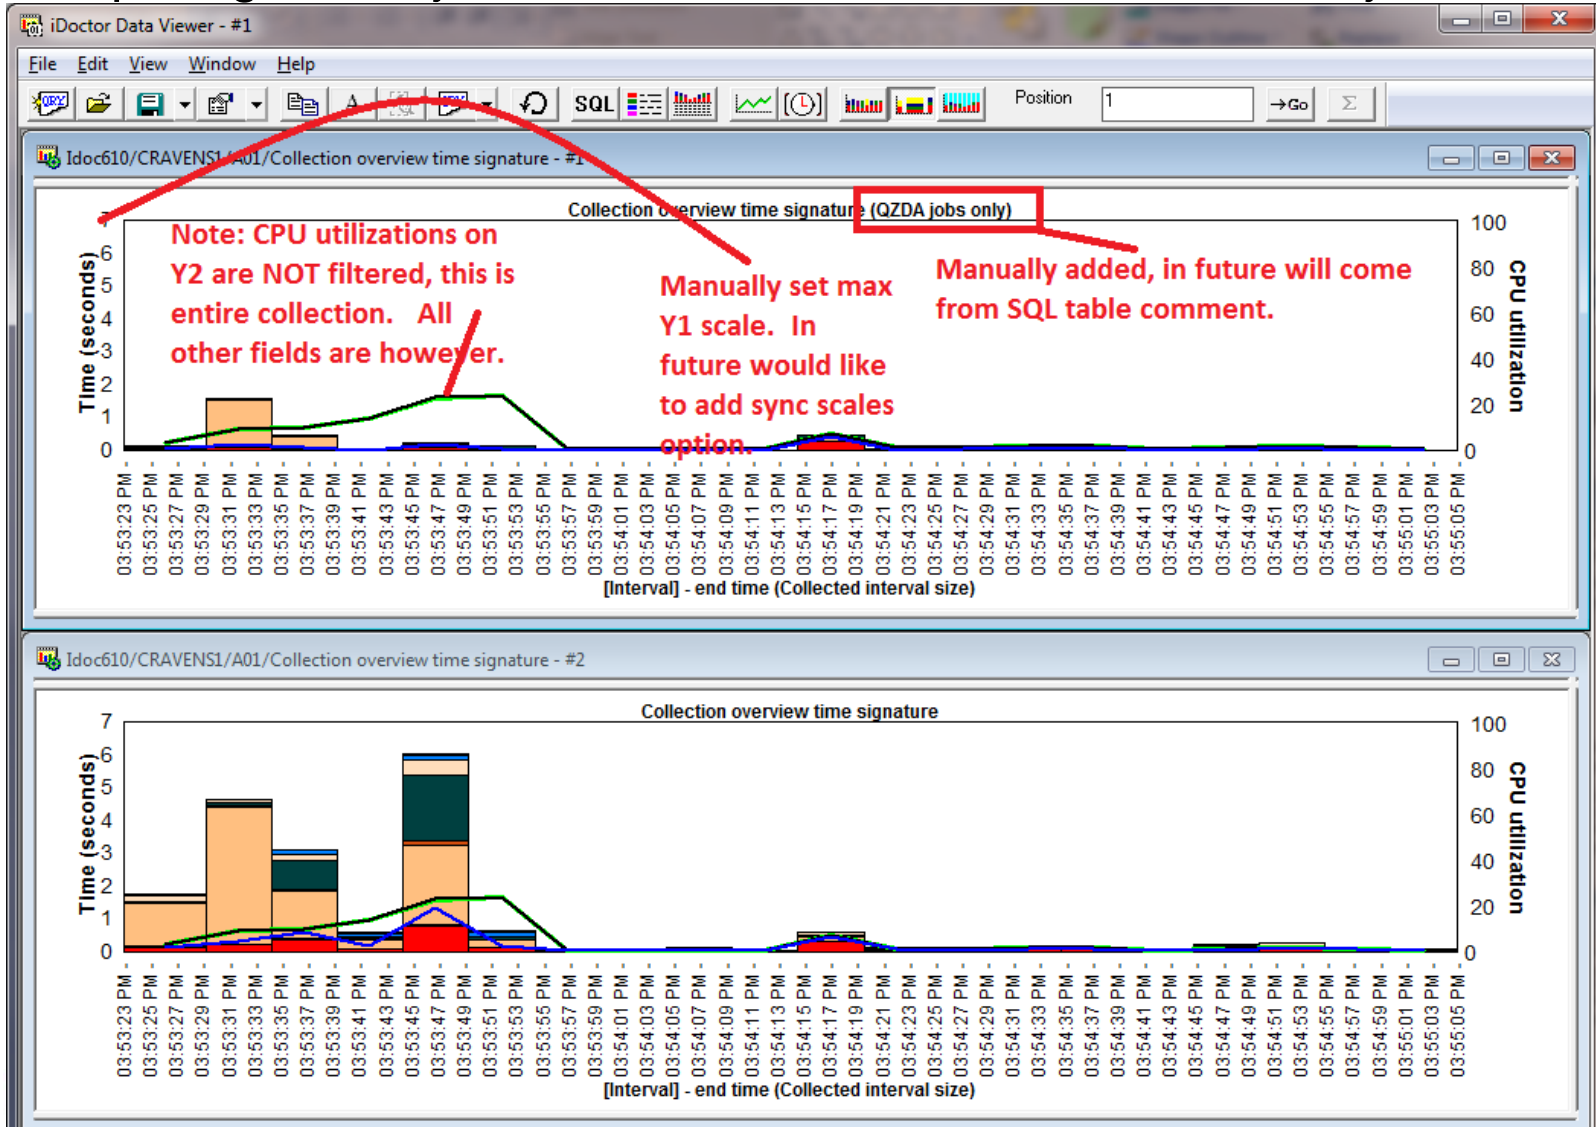
Graphs using the QZDA jobs only filter

- Open Table(s)
- Record Quick View
- Wait graphs
- CPU graphs
- I/O graphs
- IFS graphs
- Other graphs
- Edit Comment
- Delete...
- Properties

Note: The file built QAIDRJWSUM_<<MBRNAME>>* now contains totals from file QAPYJWPRC as well as QAPYJWTDE/QAPYJWSTS.

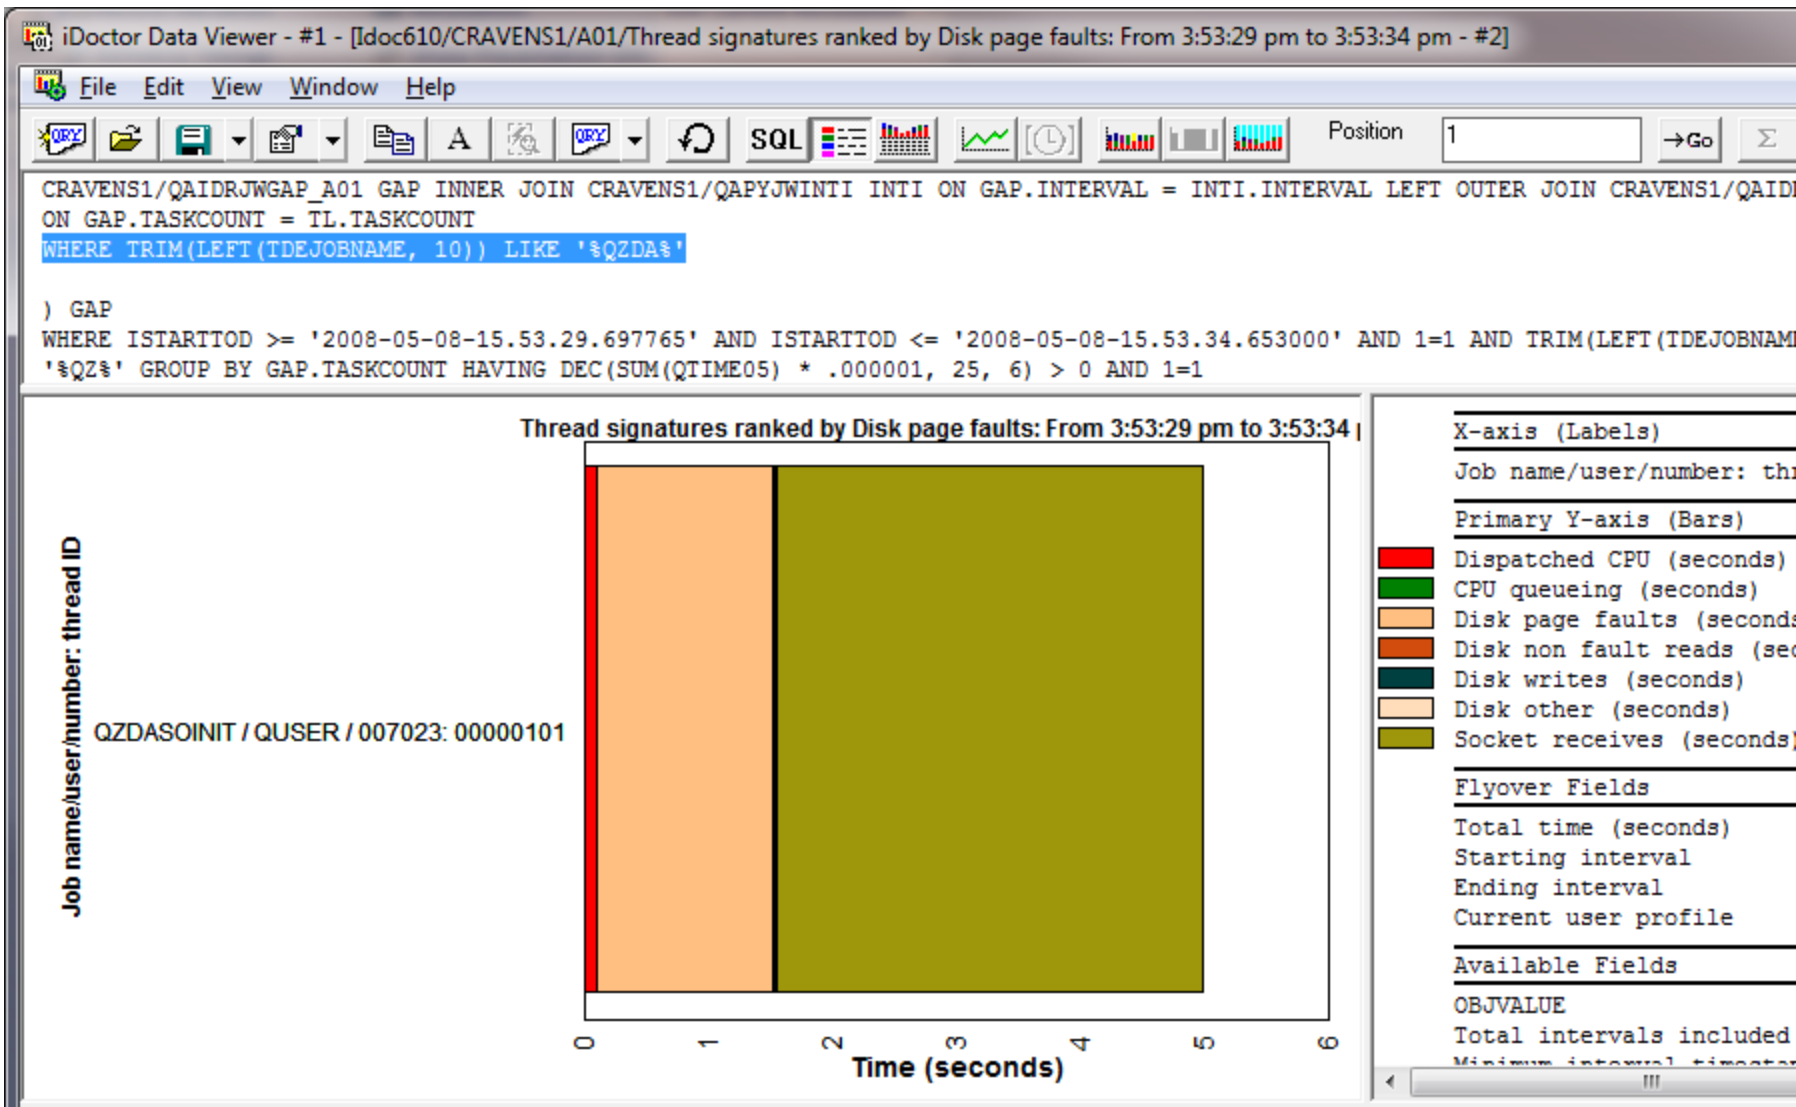
Jun-Oct 2010 – Collection Summary (JW)

Comparing QZDA jobs contributions to wait buckets vs all jobs:



Jun-Oct 2010 – Collection Summary (JW)

Drill downs from filtered overview, maintain the same filtering:



Jun-Oct 2010 – Collection Summary (JW) - Future

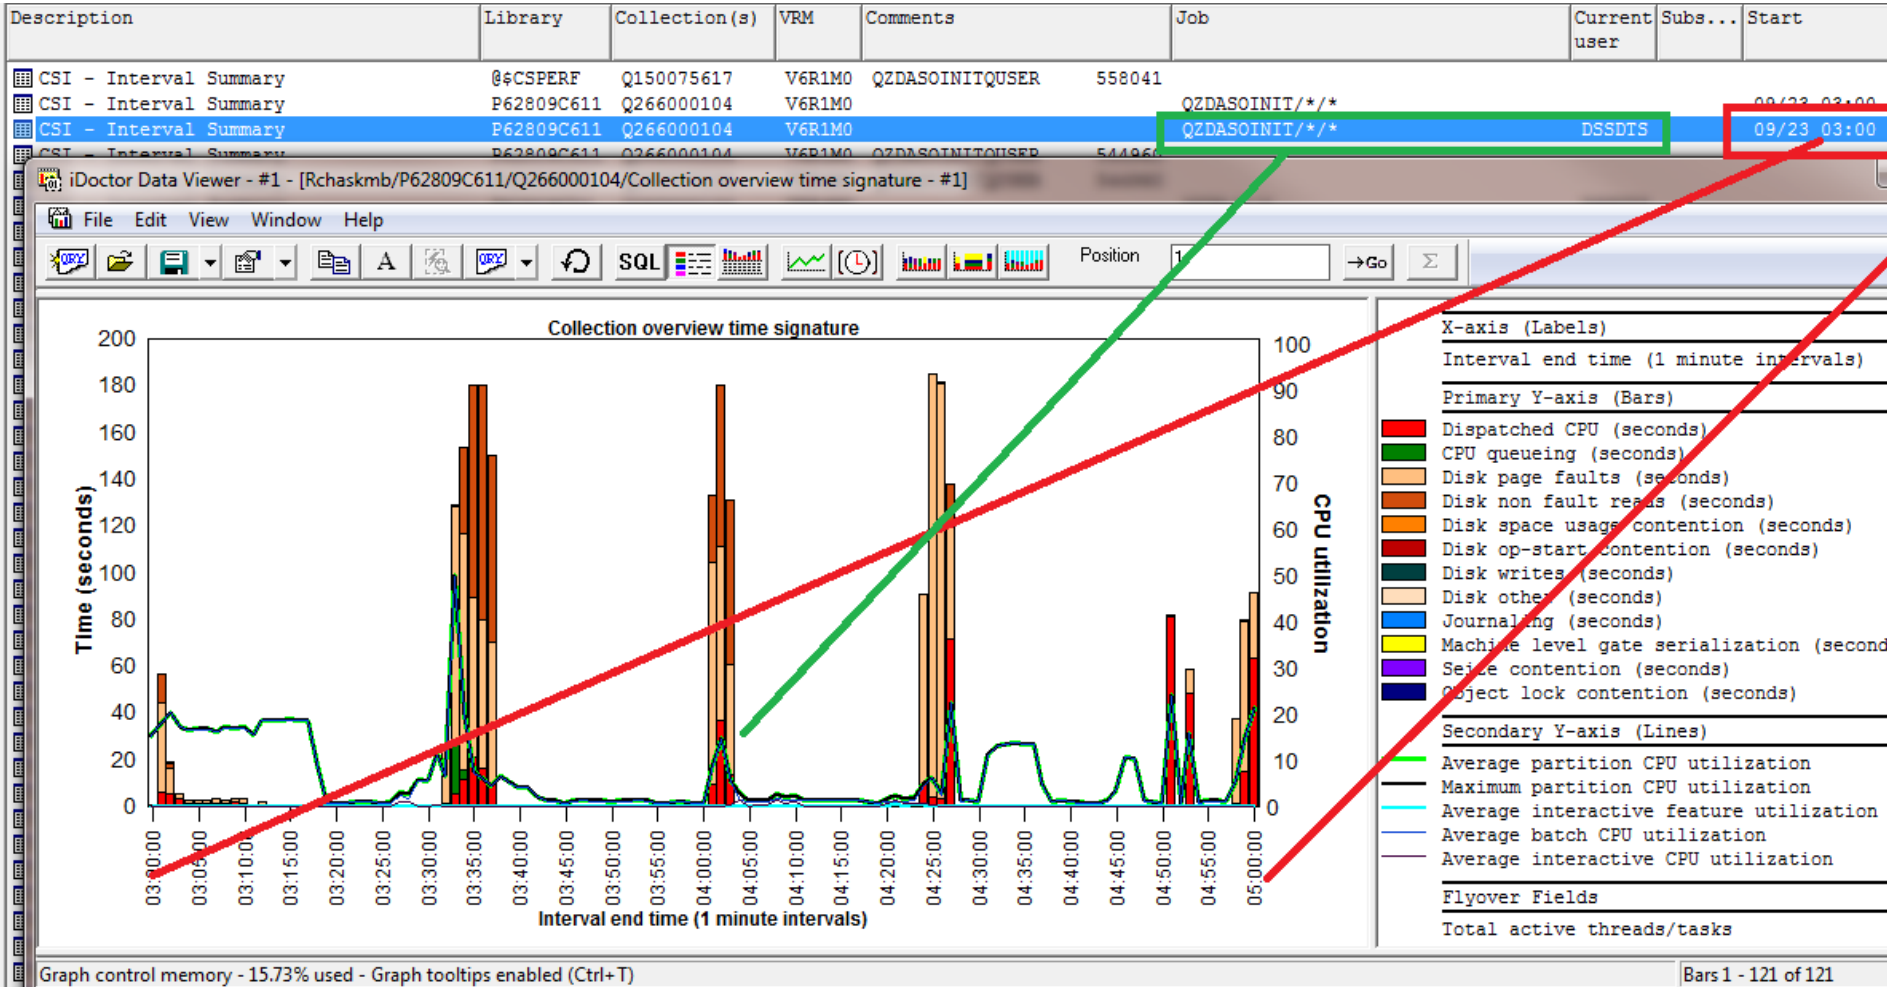
Need more flexibility in creating/showing comments on a per collection basis.

Need options to synchronize the min/max scales when comparing graphs in Data Viewer.

What other comparison graphs do you want in JW?

Jun-Oct 2010 – Collection Summary (CSI)

Just like JW, the CS Collection Summary can be filtered by time range, job, current user or subsystem to produce unique graphs.



Jun-Oct 2010 – Collection Summary (CSI)

Unlike JW, the CS Collection Summary SQL tables only provides the Wait graphs as a graphing option.

Description	Library	Collection(s)	VRM	Comments	Job	Current user	Subs...	Start	End
CSI - Interval Summary	@#CSPERF	Q150075617	V6R1M0	QZDASOINITQUSER	558041				
CSI - Interval Summary	P62809C611	Q266000104	V6R1M0			QZDASOINIT/**		09/23 03:00 am	0
CSI - Interval Summary	P62809C611	Q266000104	V6R1M0			QZDASOINIT/**	DSSDTS	09/23 03:00 am	0
CSI - Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASO					
CSI - Interval Summary	P62809C54	Q263000104	V5R4M0	QZDASO					
CSI - Interval Summary	P62809C54	Q263000104	V5R4M0			/*	DSSDTS		
CSI - Interval Summary	P62809C54	Q263000104	V5R4M0					4:00 am	0
CSI - Interval Summary	P62809C54	Q263000104	V5R4M0					3:00 am	0
CSI - Interval Summary	PMR30064AA	Q272000003	V6R1M0	QZDASO					
CSI - Interval Summary	PMR30064AA	Q271000003	V6R1M0	QZDASO					
CSI - Interval Summary	P02872C929	Q272012508	V6R1M0						
CSI - Interval Summary	PMR92635	Q270093957	V6R1M0	QZDASO					
CSI - Interval Summary	PMR92635	Q270000105	V6R1M0	QZDASO					
CSI - Interval Summary	PMR92635	Q268000106	V6R1M0	QZDASOINITQUSER	523110				
CSI - Interval Summary	PMR43853AA	Q265000025	V6R1M0		MB400				
CSI - Interval Summary	PMR43853AA	Q266000015	V6R1M0	QZDASOINITQUSER	510105				
CSI - Interval Summary	PMR43853AA	Q265000025	V6R1M0	QZDASOINITQUSER	510105				
CSI - Interval Summary	PMR02727CS	Q260000104	V5R4M0	QZDASOINITQUSER	500704				
CSI - Interval Summary	PMR72304AA	Q260000104	V6R1M0	QZDASOINITQUSER	490024				
CSI - Interval Summary	PMR72304AB	Q260000104	V6R1M0	QZDASOINITQUSER	489959				
CSI - Interval Summary	PMR03617CS	Q259000003	V6R1M0	QZDASOINITQUSER	481454				
CSI - Interval Summary	PMR03617CS	Q258000003	V6R1M0	QZDASOINITQUSER	481454				
CSI - Interval Summary	PMR79401AC	Q259000031	V6R1M0			IP_39708X/**			
CSI - Interval Summary	PMR79401AC	Q258000017	V6R1M0			IP_39708X/**			

- Open Table(s)
- Record Quick View
- Wait graphs
- Edit Comment
- Delete...
- Properties

- Collection overview time signature
- Collection overview with dispatch CPU time signature
- Seizes and locks time signature
- Contention time signature
- Disk time signature
- Classic JVM time signature
- DB record lock time signature
- Communications time signature
- Collection overview time signature comparison
- Dispatched CPU rankings
- Disk page faults rankings
- Counts

Jun-Oct 2010 – Collection Summary (CSI) Totals

The analysis also includes an iDoctor Collection Summary Totals SQL table automatically for every analysis ran.

This is a 1 record summary over the collection or filtered data.

The screenshot shows a system monitoring interface. On the left, a list of servers is displayed, including 'I2cs0112', 'Ibmcs928', 'Junk', 'Kurtz', 'Larsondr', 'Larsondr2', 'Osc0722', 'Oscar', 'P02280', 'P02872c929', 'P05697', 'P05921cs1', 'P28824', 'P37232', 'P45603cs', 'P45603pfr', 'P52758cs', 'P62809c54', 'P62809c611', 'SQL tables', 'iDoctor Interval Summary', 'iDoctor Interval Summary Totals', 'Situational Analysis', and 'Q266000104'. The 'iDoctor Interval Summary Totals' table is selected.

The table displays the following data:

CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496
CSI - Aggregated Interval Summary	P62809C611	Q266000104	V6R1M0	QZDASOINITQUSER	54496

A context menu is open over the table, showing the following options:

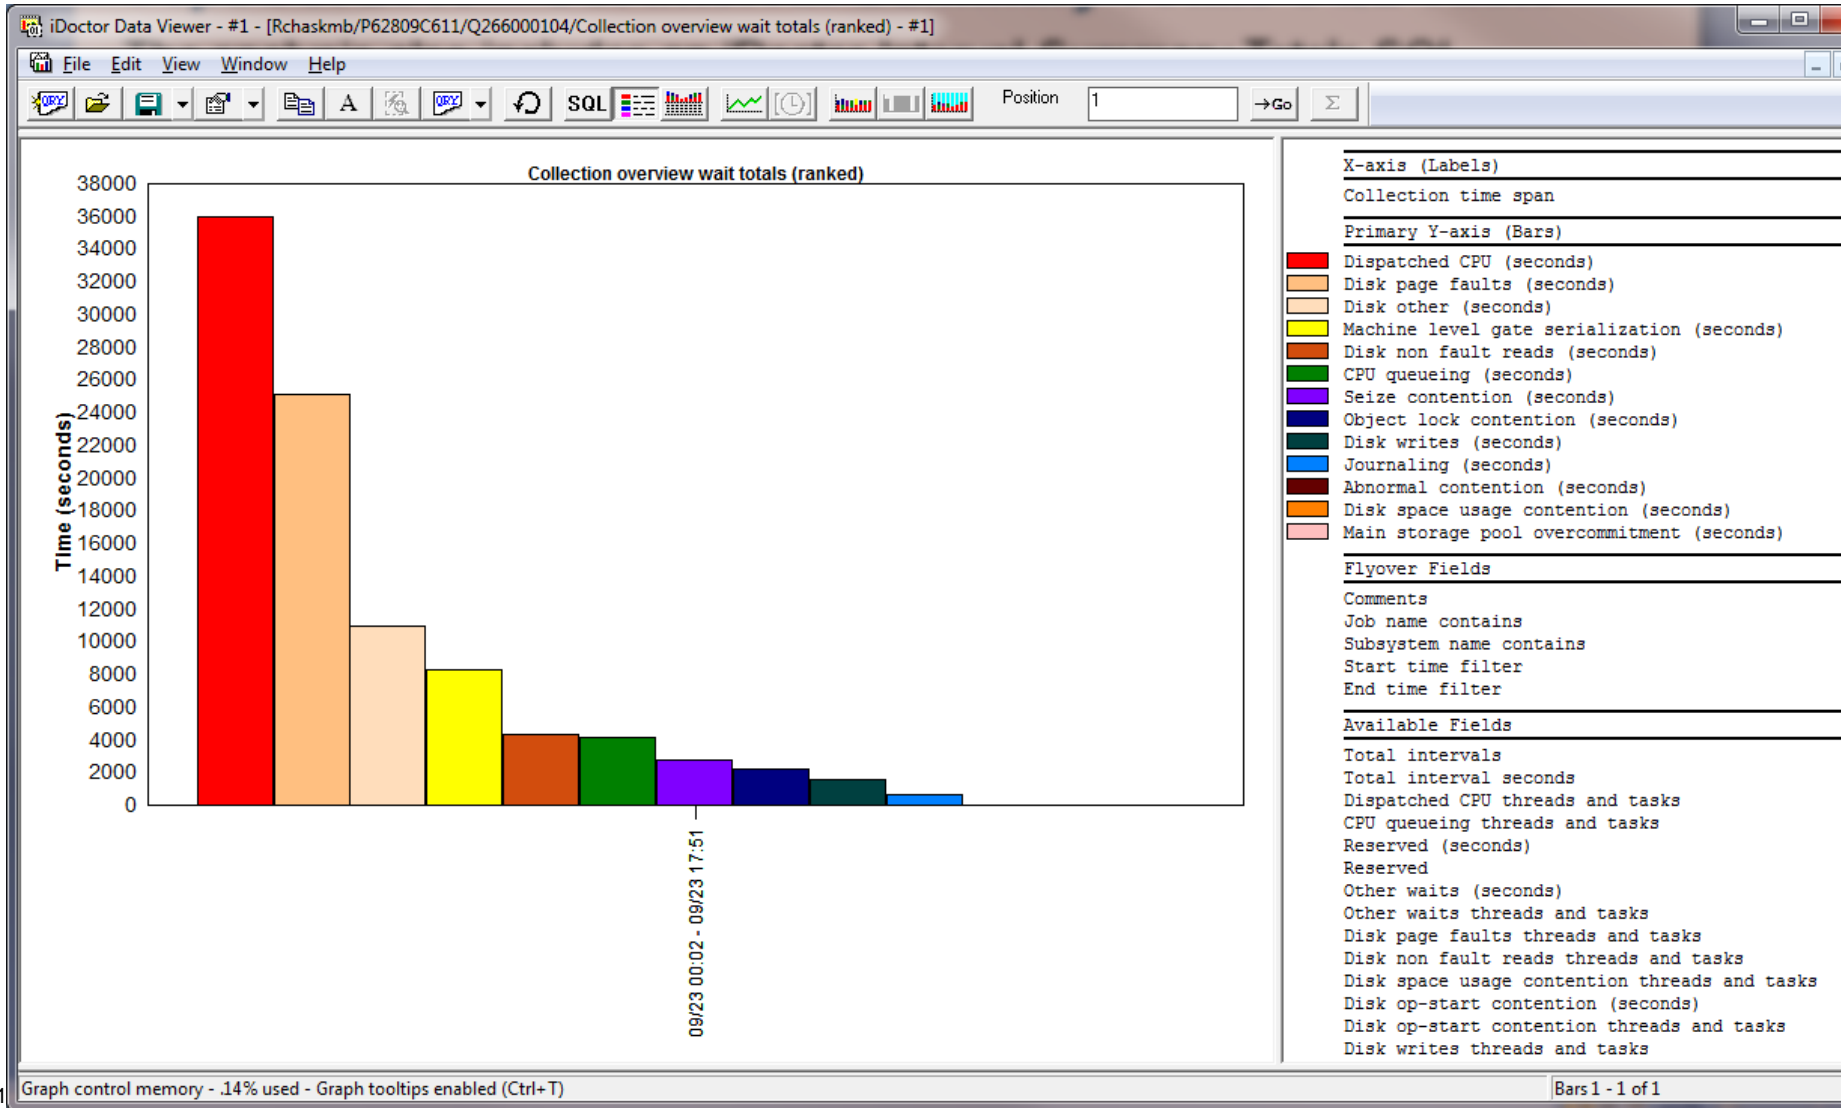
- Open Table(s)
- Edit
- Wait graphs
- IO graphs
- Delete
- Properties

The 'Wait graphs' and 'IO graphs' options have sub-menus:

- Wait graphs: Collection overview wait totals compar...
- IO graphs: Collection overview wait totals (ranked), Collection overview wait totals (pie), Counts per second, Contributing jobs


Jun-Oct 2010 – Collection Summary (CSI)

Ranking the wait buckets in the SQL table.



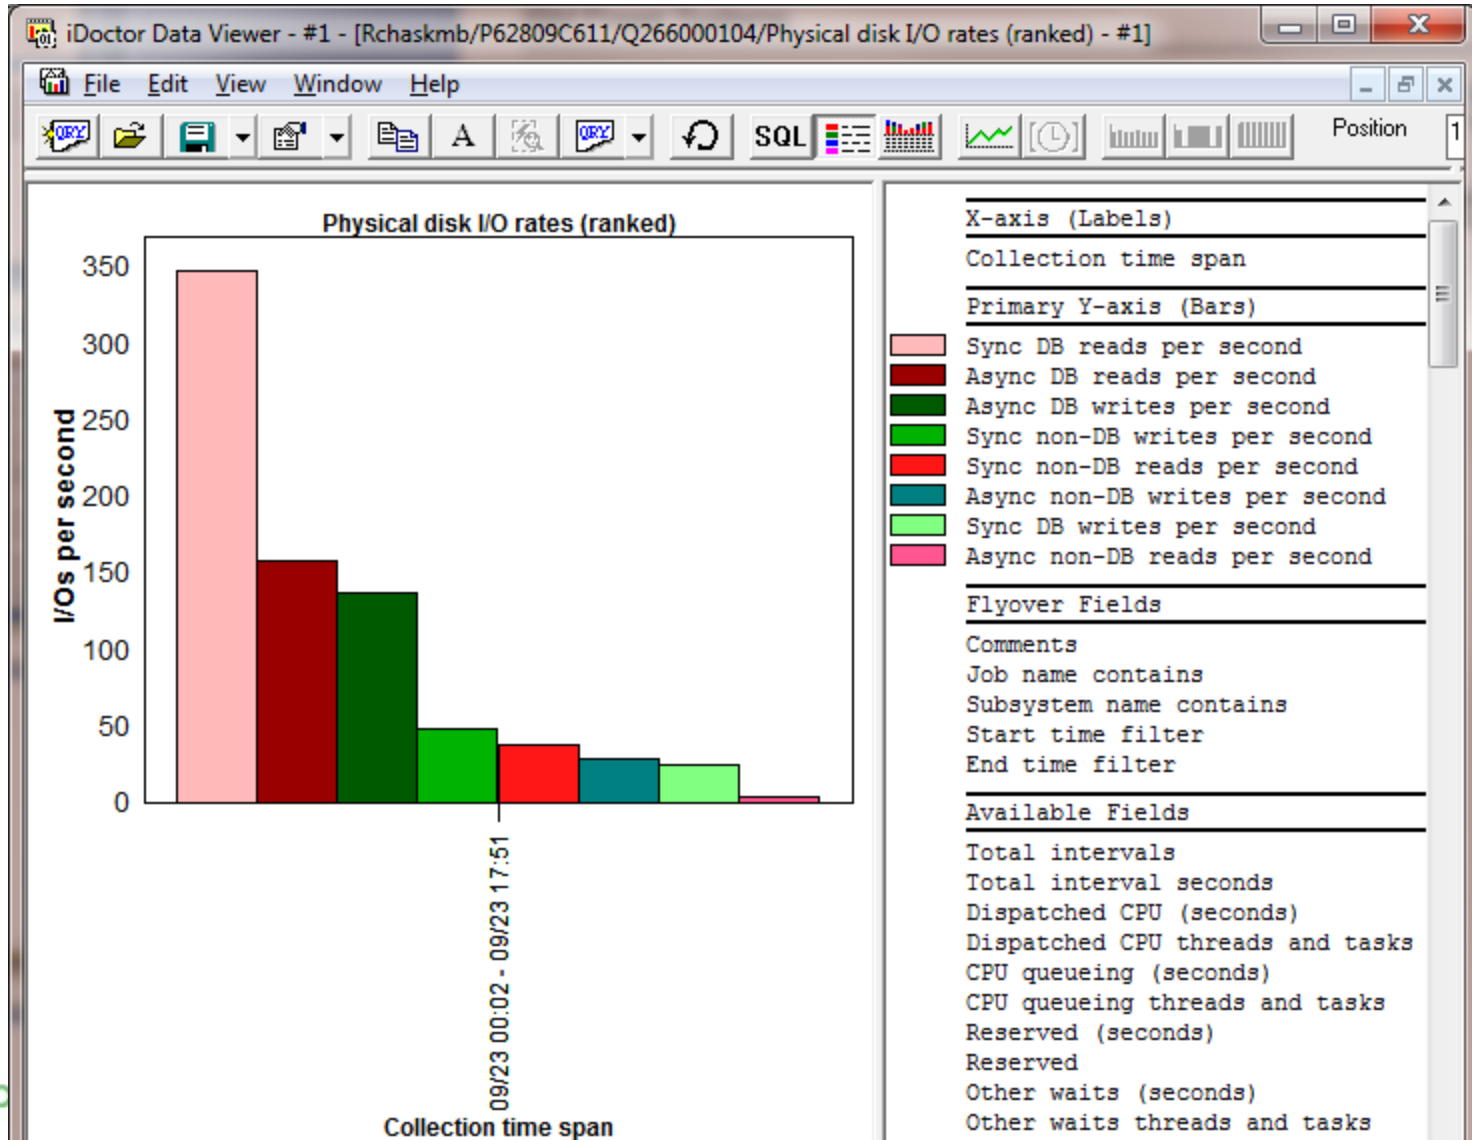
Jun-Oct 2010 – Collection Summary (CSI)

Pie chart of “interesting” wait bucket distribution also available.



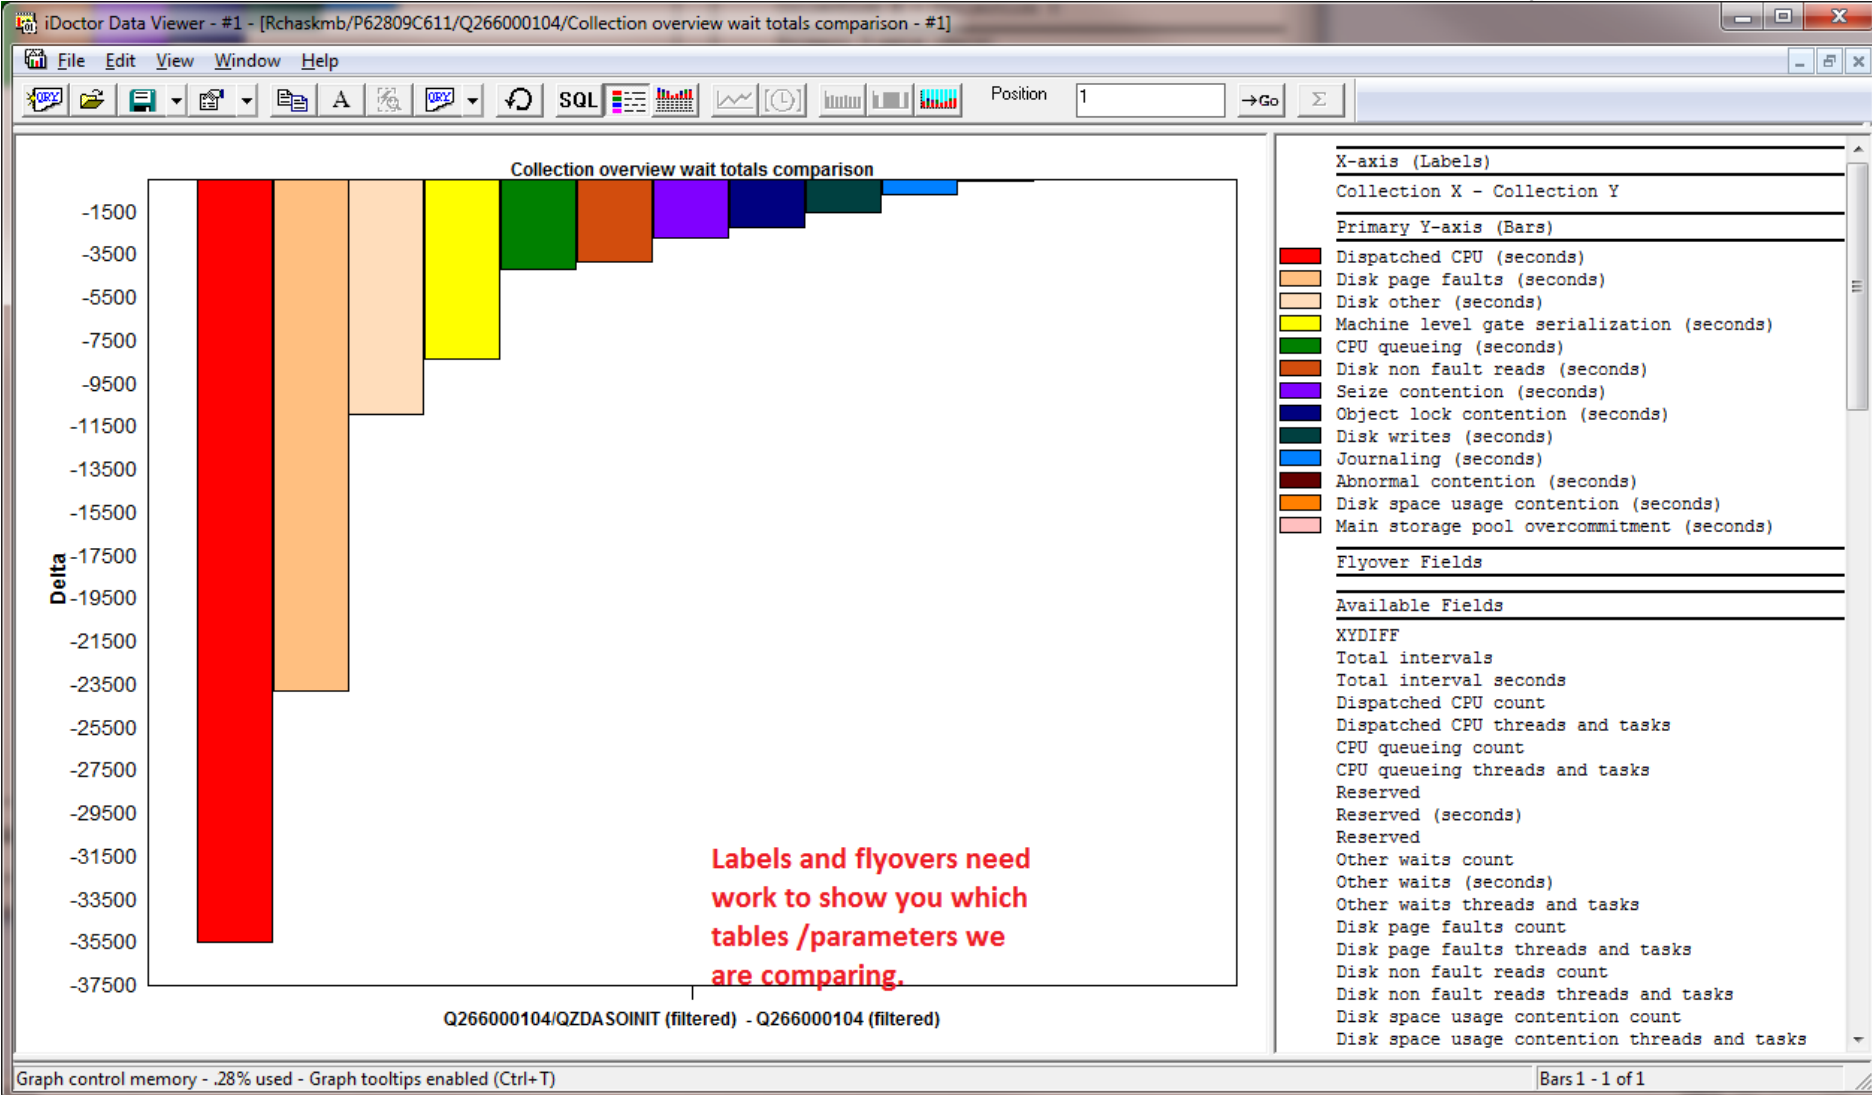
Jun-Oct 2010 – Collection Summary (CSI)

PDIO rates (ranked or in pie chart form) are included.



Jun-Oct 2010 – Collection Summary (CSI)

Delta comparison between wait buckets in 2 Collection Summary totals



Jun-Oct 2010 – Collection Summary (CSI) - Future

Need more flexibility in creating/showing comments on a per collection basis (not just per run).

SQL merge tables function

Need options to synchronize the min/max scales when comparing graphs in Data Viewer.

What other comparison graphs do you want in CSI?
(memory? disk? showing what exactly?)

Jun-Oct 2010 – SQL Tables Overview

SQL Tables folder is a repository in iDoctor for the SQL tables generated by the analyses in iDoctor.

SQL Tables folder contains a subfolder for each type of SQL table generated by the iDoctor analyses.

Allows comparisons in some instances (CSI Collection Summary) and many drill down and graphing options (CSI/JW create job summary, CSI/JW Collection Summary) by right-clicking the SQL tables.

In some cases drill downs are accessible from the SQL tables after opening them and right-clicking the records within. (stats hier for one job)

Found in JW, PEX and CSI.

Jun-Oct 2010 – SQL Tables Locations

Accessible from 2 locations (component or library level)

The top screenshot shows the 'SQL tables' component selected in the left pane. The table below shows the analysis items and their descriptions:

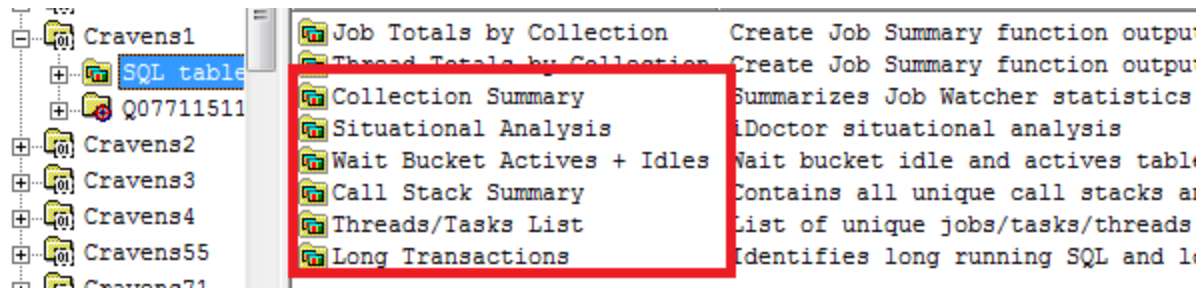
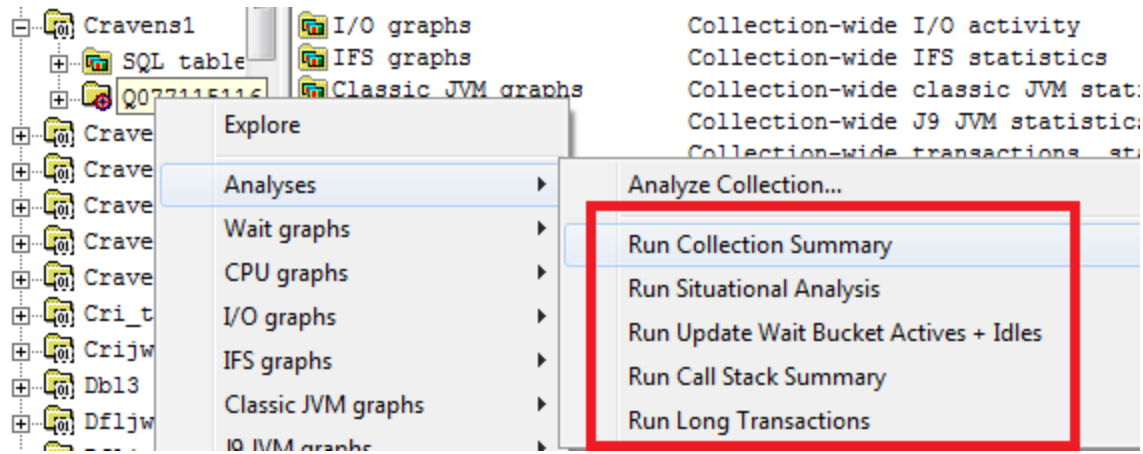
Analysis	Description	Number of tables
iDoctor Interval Summary	Statistics summarized by time interval (wait buckets, CPU, I/Os, etc) from files QAPYJWIDE/QAPYJWSTS	73
Situational Analysis	iDoctor situational analysis	74
Wait Bucket Actives + Idles	Wait bucket idle and actives table (includes all 32 buckets every interval)	72
Call Stack Summary	Contains all unique call stacks and how many there were of each	42
Threads/Tasks List	List of unique jobs/tasks/threads	73
Long Transactions	Identifies long running SQL and long periods of work where no 'idle' waits were found	41

The bottom screenshot shows the 'SQL tables' component selected in the 'Cravens1' library in the left pane. The table below shows the analysis items and their descriptions:

Analysis	Description	Number of tables
iDoctor Interval Summary	Statistics summarized by time interval (wait buckets, CPU, I/Os, etc) from files QAPYJWIDE/QAPYJWSTS	2
Situational Analysis	iDoctor situational analysis	2
Wait Bucket Actives + Idles	Wait bucket idle and actives table (includes all 32 buckets every interval)	1
Call Stack Summary	Contains all unique call stacks and how many there were of each	1
Threads/Tasks List	List of unique jobs/tasks/threads	1
Long Transactions	Identifies long running SQL and long periods of work where no 'idle' waits were found	1

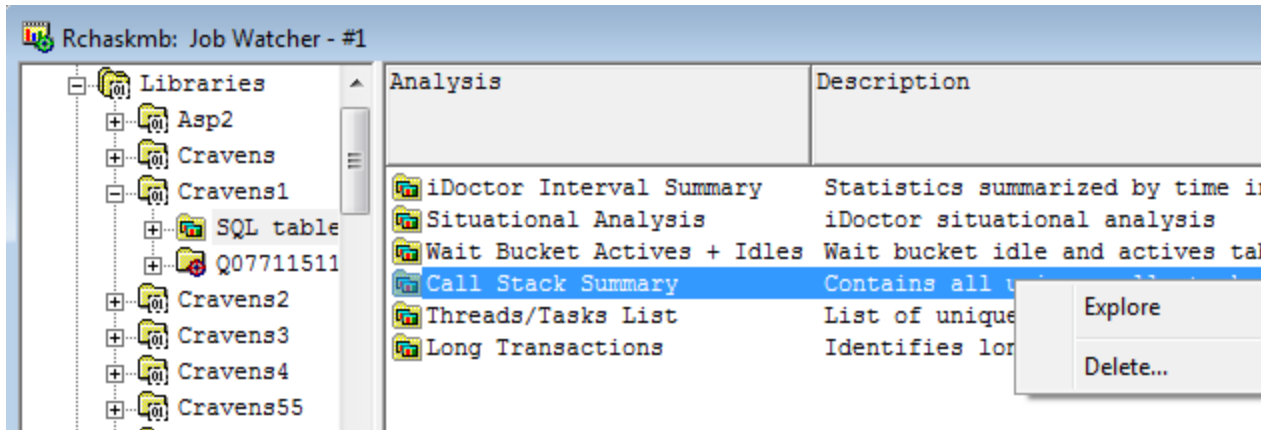

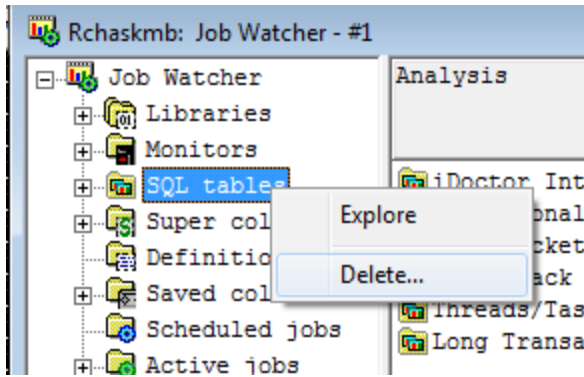
Jun-Oct 2010 – SQL Tables Creation

The Analyses menu options generally match the SQL tables subfolders and are used to create the SQL tables.



Jun-Oct 2010 – SQL Tables Deletion

For cleanup purposes, you can delete all SQL tables on the system, by library or by analysis



Jun-Oct 2010 – SQL Tables: CSI Collection Summary

CSI – Collection Summary currently only includes “Wait graphs”

Note: The same graphs are available after opening one of the SQL tables and right-clicking a row under the “Collection overview” menu.

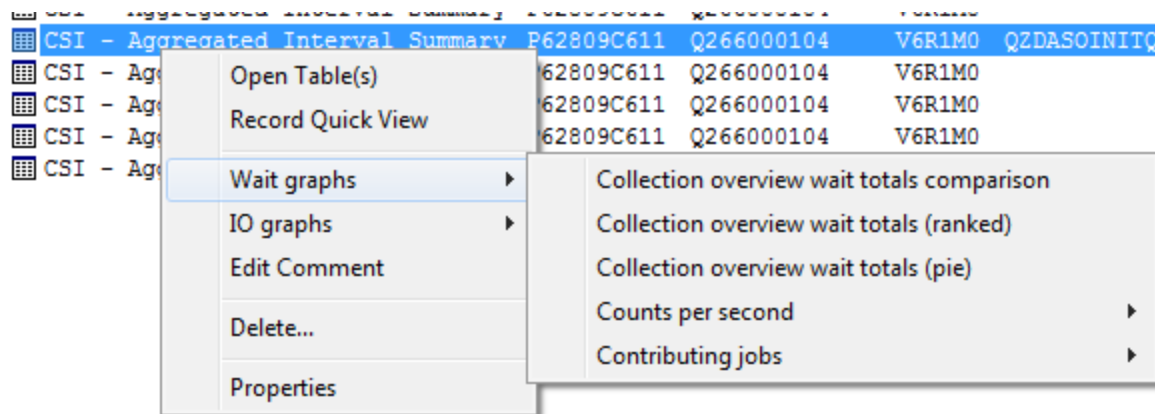
Description	Library	Collection(s)	VRM	Comments
CSI iDoctor Interval Summary				
CSI - Interval Summary	P62809C611	Q266000104	V6R1M0	
Open Table(s)	09C611	Q266000104	V6R1M0	
Record Quick View	09C611	Q266000104	V6R1M0	QZDASOINITQUSER 5.

Wait graphs	Collection overview time signature
Edit Comment	Collection overview with dispatch CPU time signature
Delete...	Seizes and locks time signature
Properties	Contention time signature
	Disk time signature
	Classic JVM time signature
	DB record lock time signature
	Communications time signature
	Collection overview time signature comparison
	Dispatched CPU rankings
	Disk page faults rankings
	Counts

Jun-Oct 2010 – SQL Tables: CSI Collection Summary Totals

CSI – Collection Summary Totals “Collection overview wait totals comparison” graph show differences in total wait buckets between 2 SQL tables

Also includes wait bucket contributing jobs and counts per second as well as PDIO and LDIO graphs.



CSI - iDoctor Interval Summary Totals

Jun-Oct 2010 – SQL Tables: CSI System Configuration

System configuration analysis includes 2 reports (show key fields or show all fields). Must select 2 SQL tables for this.

The SQL table is a field-described QAPMCONF

The screenshot shows a table with columns 'Description', 'Library', and 'Collection(s)'. Three rows are visible, all starting with 'CSI - Configuration Information'. A context menu is open over the second row, with 'Compare configurations' selected. A sub-menu is also open, showing 'Show key fields' and 'Show all fields' options.

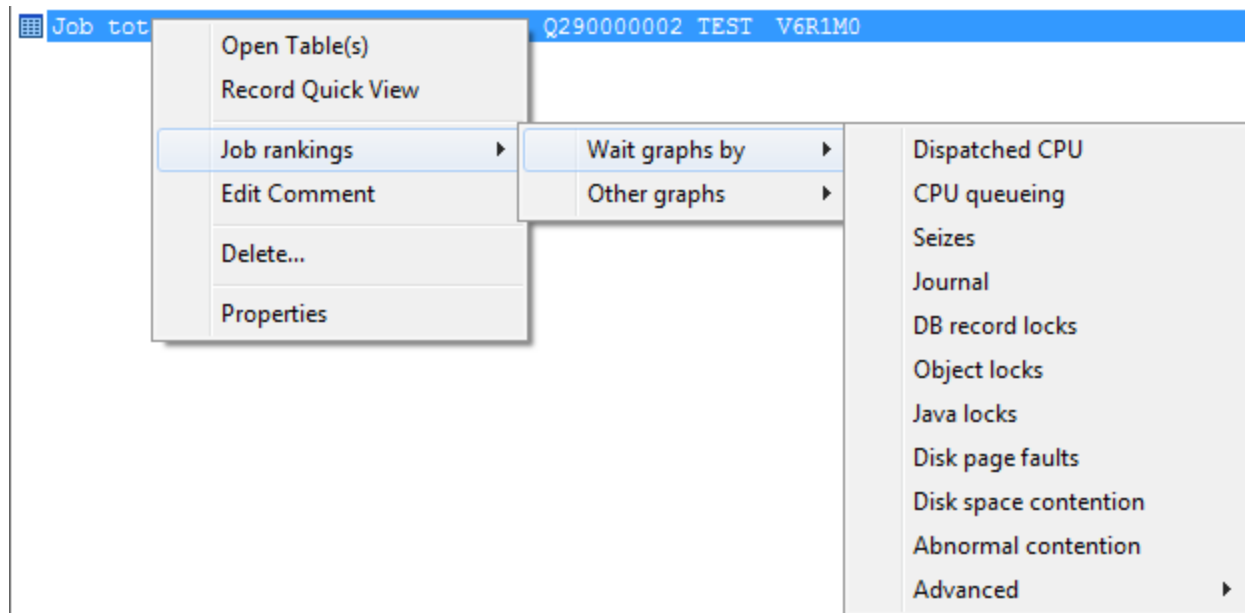
CSI - System Configuration (2 must be selected)

System name (SYS_NAME)	Collection library (COL_LIB)	Collection name (COL_NAME)	Version (VERSION)	Interval (in minutes) (INT_MINUTES)	Type (TYPE)	Model (MODEL)	Total processor cores (TOTAL_PROCS)	Installed processor cores (INSTALLED_PROCS)	Available processor cores (AVAILABLE_PROCS)	Virtual processors (VIRTUAL_PROCS)	Assigned processor cores (processor units) (PROC_UNITS)
AMWKOR	P62809C611	Q266000104	610	1	9117	MMA	16	4	12	4	4.00
INTERLG2	QPFRDATA	Q270093957	610	15	8203	E4A	4	4	0	1	1.00

Jun-Oct 2010 – SQL Tables: CSI Jobs Totals or Job Totals by Collection

Created by Create Job Summary function.

Provides Job rankings graphs for wait buckets, CPU, I/Os, and more



Jun-Oct 2010 – SQL Tables: CSI Jobs Totals or Threads Totals SQL tables

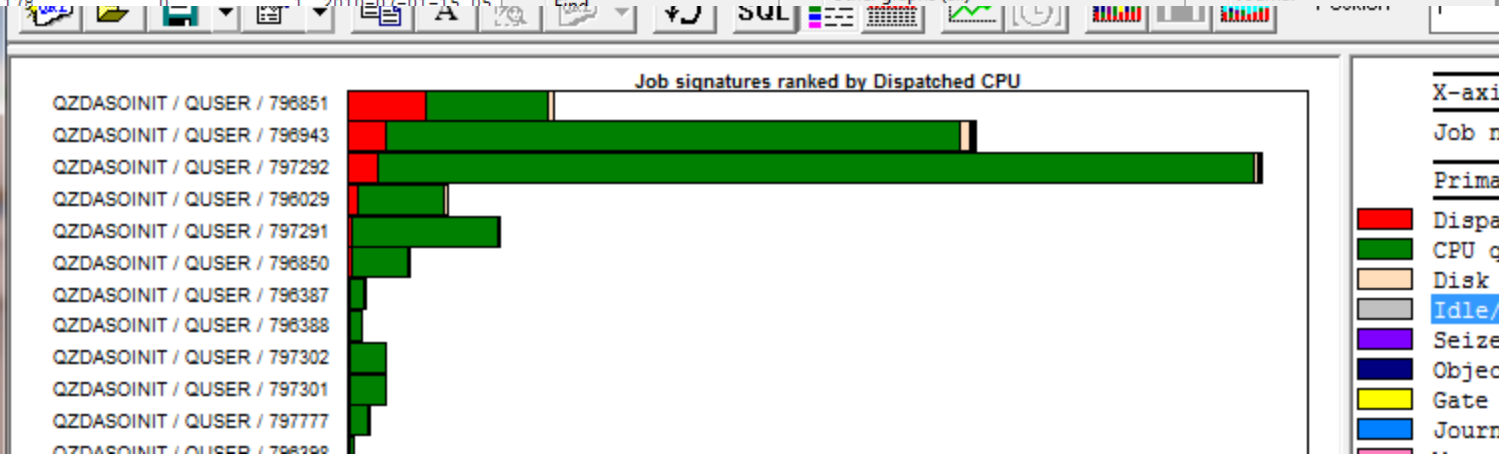
Open tables and select the desired jobs/threads to graph only the selected ones.

Contributing collections (MBRCOUNT)	Duration of job in hours (RUNTIME)	Elapsed interval seconds (TDESECS)	Start of job included 1-Yes, 0-No (JOBSTART)	End of job included 1-Yes, 0-No (JOBEND)	Start timestamp (STARTTIME)	End timestamp (ENDTIME)	Job name (JBNAME)	Job user (JBUSER)	Job number (JBNBR)	Current user profile (JBCUSR)	Job pool (JBPOOL)	CE (J
1	.650	2357	0	1	2010-07-01-09.05.00.000000	2010-07-01-09.44.17.000000	QZDASOINIT	QUSER	796017	QUSER	02	
1	.316	1159	1	1	2010-07-01-09.04.56.000000	2010-07-01-09.24.15.000000	QZDASOINIT	QUSER	796028	QUSER	02	
1	2.766	9972	1	1	2010-07-01-09.04.56.000000	2010-07-01-11.51.08.000000	QZDASOINIT	QUSER	796029	QUSER	02	
1	.650	2345	1	1	2010-07-01-09.55.14.000000	2010-07-01-10.34.19.000000	QZDASOINIT	QUSER	796387	QUSER	02	
1	.016	77	1	1	2010-07-01-09.55.15.000000	2010-07-01-09.56.32.000000	QZDASOINIT	QUSER	796388	QUSER	02	
1	.300	1106	1	1	2010-07-01-09.55.51.000000	2010-07-01-10.14.17.000000	QZDASOINIT	QUSER	796398	QUSER	02	
1	.466	1707	1	1	2010-07-01-09.55.51.000000	2010-07-01-10.24.18.000000	QZDASOINIT	QUSER	796399	QUSER	02	
1	2.900	10,468	1	1	2010-07-01-11.01.49.000000	2010-07-01-13.56.17.000000	QZDASOINIT	QUSER	796850	QUSER	02	
1	1.533	5559	1	1	2010-07-01-11.01.49.000000	2010-07-01-12.34.28.000000	QZDASOINIT	QUSER	796851	QUSER	02	
1	3.983	14,379	0	1	2010-07-01-11.25.00.000000	2010-07-01-15.24.39.000000	QZDASOINIT	QUSER	796943	QUSER	02	
1	.483	1758	1	1	2010-07-01-11.15.09.000000	2010-07-01-11.44.27.000000	QZDASOINIT	QUSER	796944	QUSER	02	
1	1.833	6648	1	1	2010-07-01-12.04.03.000000	2010-07-01-13.54.51.000000	QZDASOINIT	QUSER	797291	QUSER	02	
1	3.333	12,036	1	1	2010-07-01-12.04.03.000000	2010-07-01-15.24.39.000000	QZDASOINIT	QUSER	797292	QUSER	02	
1	.283	1056	1	1	2010-07-01-12.06.11.000000	2010-07-01-12.06.11.000000						
1	.283	1055	1	1	2010-07-01-12.06.11.000000	2010-07-01-12.06.11.000000						
1	.933	3369	1	1	2010-07-01-13.28.11.000000	2010-07-01-13.28.11.000000						
1	.433	1565	1	1	2010-07-01-13.28.11.000000	2010-07-01-13.28.11.000000						
1	3.916	14,105	0	1	2010-07-01-21.20.05.000000	2010-07-01-21.20.05.000000						
1	.316	1178	0	1	2010-07-01-15.05.00.000000	2010-07-01-15.05.00.000000						

- Job rankings
- Record Quick View
- Copy
- Find


- Wait graphs by (selected jobs only)
- Wait graphs by (all)
- Other graphs (selected jobs only)
- Other graphs (all)

- Dispatched CPU
- CPU queueing
- Seizes
- Journal



Jun-Oct 2010 – SQL Tables: CSI Jobs Totals

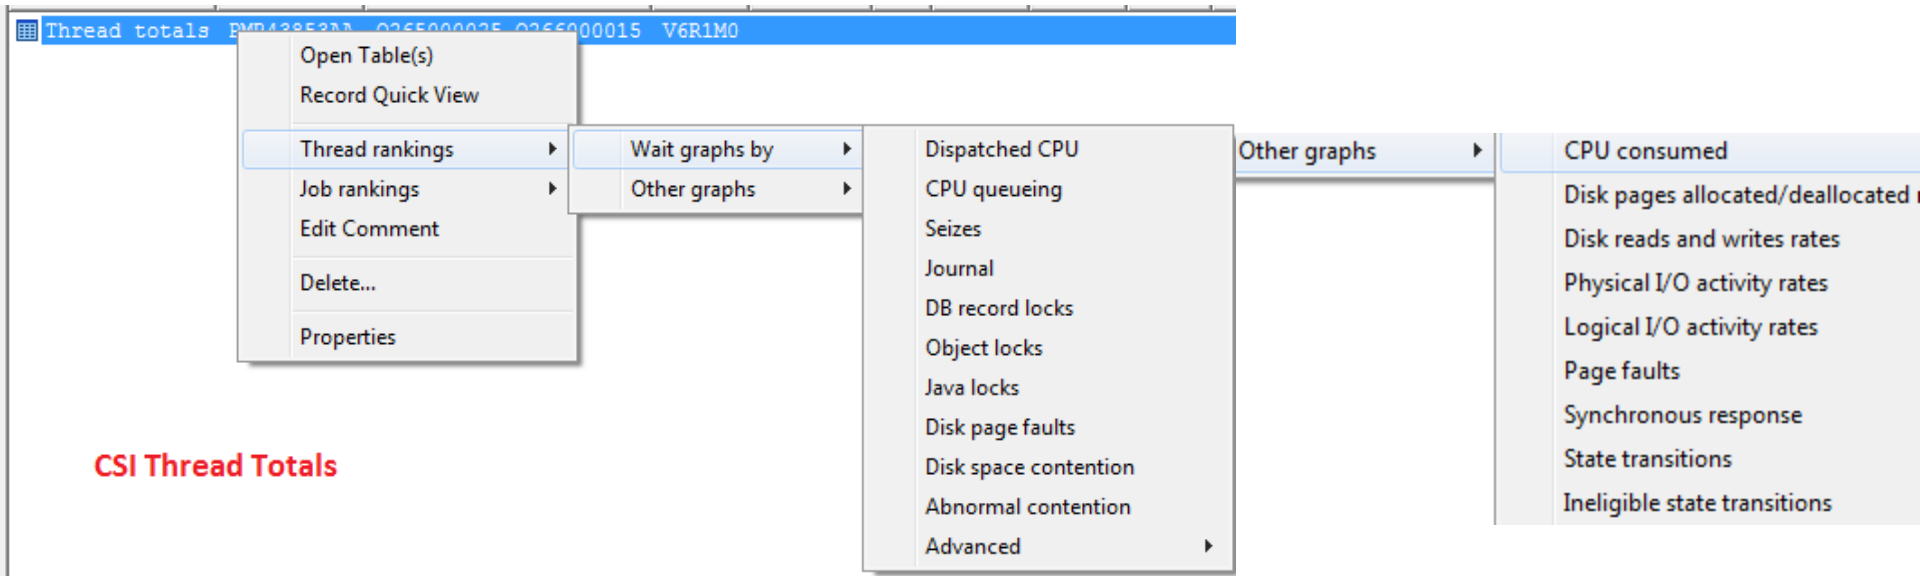
Job rankings by Dispatched CPU



Jun-Oct 2010 – SQL Tables: CSI Thread Totals or Thread Totals by Collection

Created by Create Job Summary function.

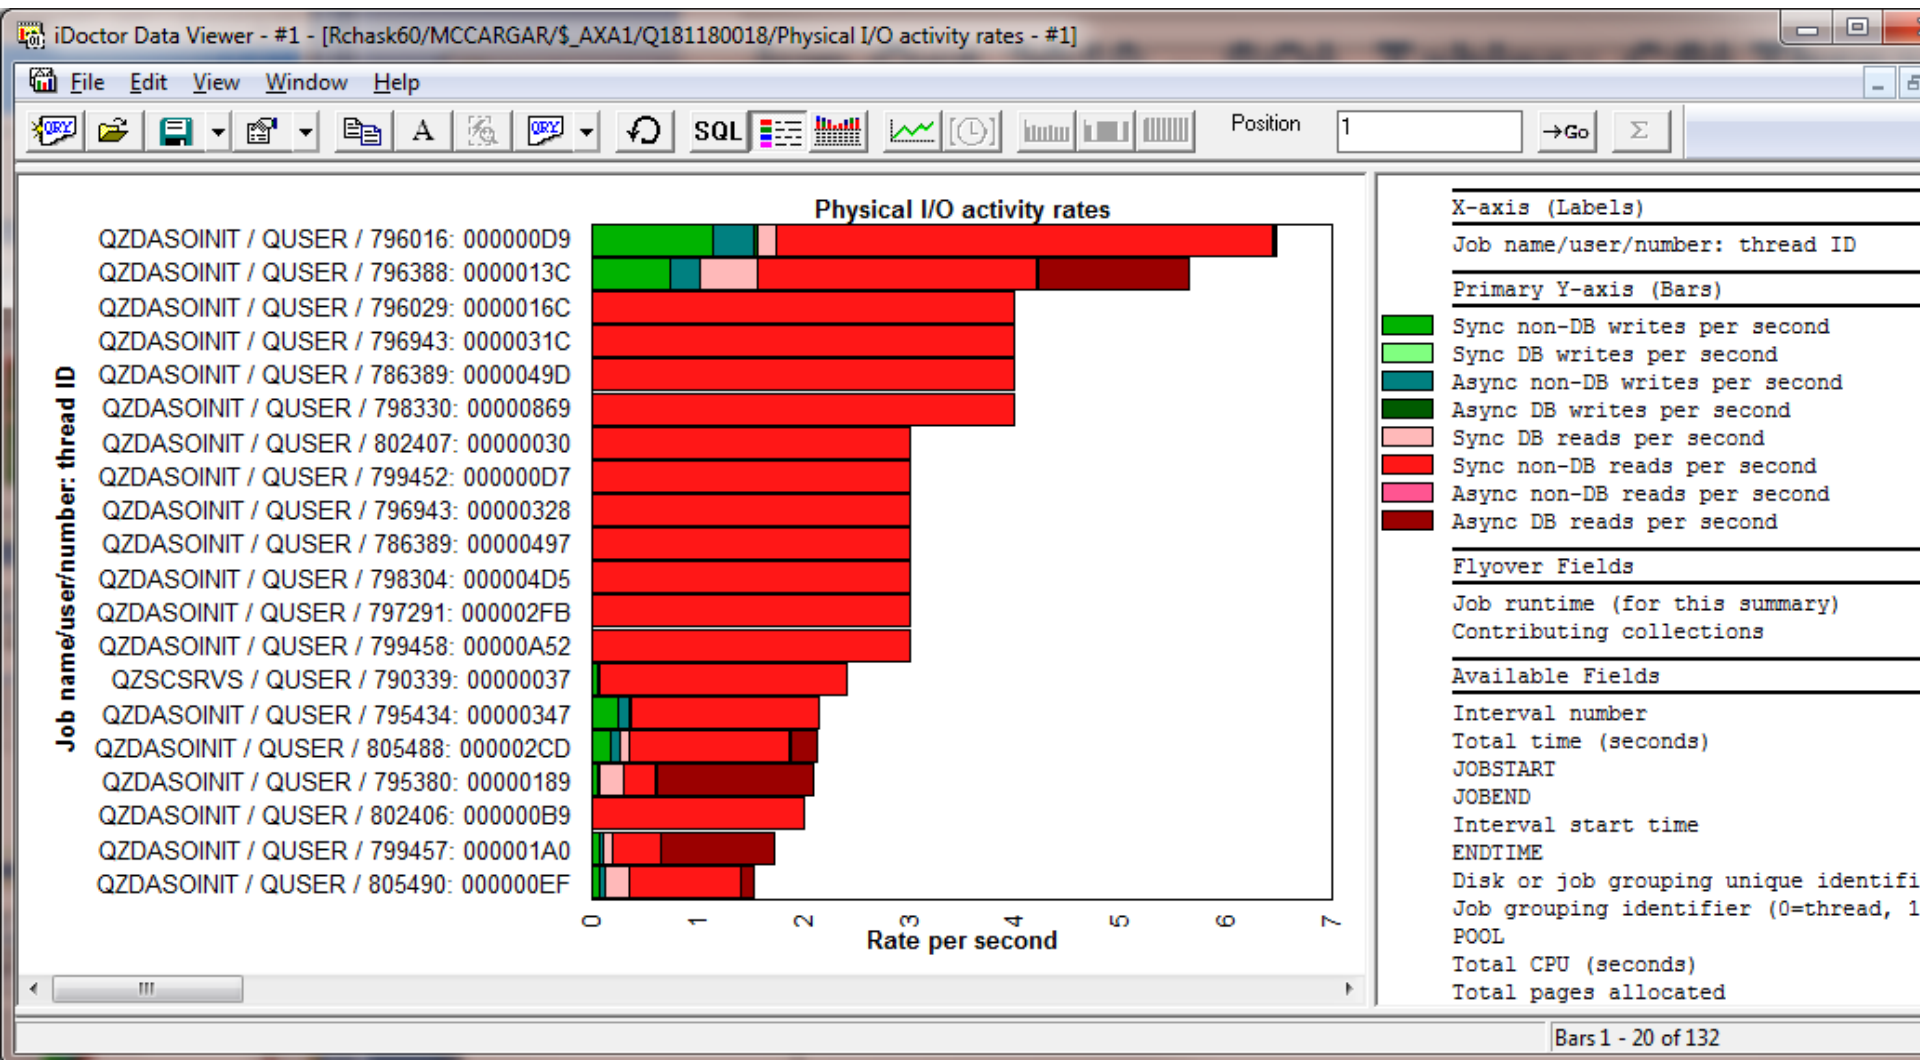
Provides Thread and Job rankings graphs for wait buckets, CPU, I/Os, and more



CSI Thread Totals

Jun-Oct 2010 – SQL Tables: CSI Thread Totals or Thread Totals by Collection

Thread rankings showing PDIO rates



Jun-Oct 2010 – SQL Tables: CSI External Storage Cache Statistics

New at 6.1+ from file QAPMXSTGD

Provides graphs over time, by disk unit and ASP over DS6000/DS8000 external storage.

The by time interval, by disk unit options will prompt you for desired ASP or give comparison option to compare 1 ASP with another.

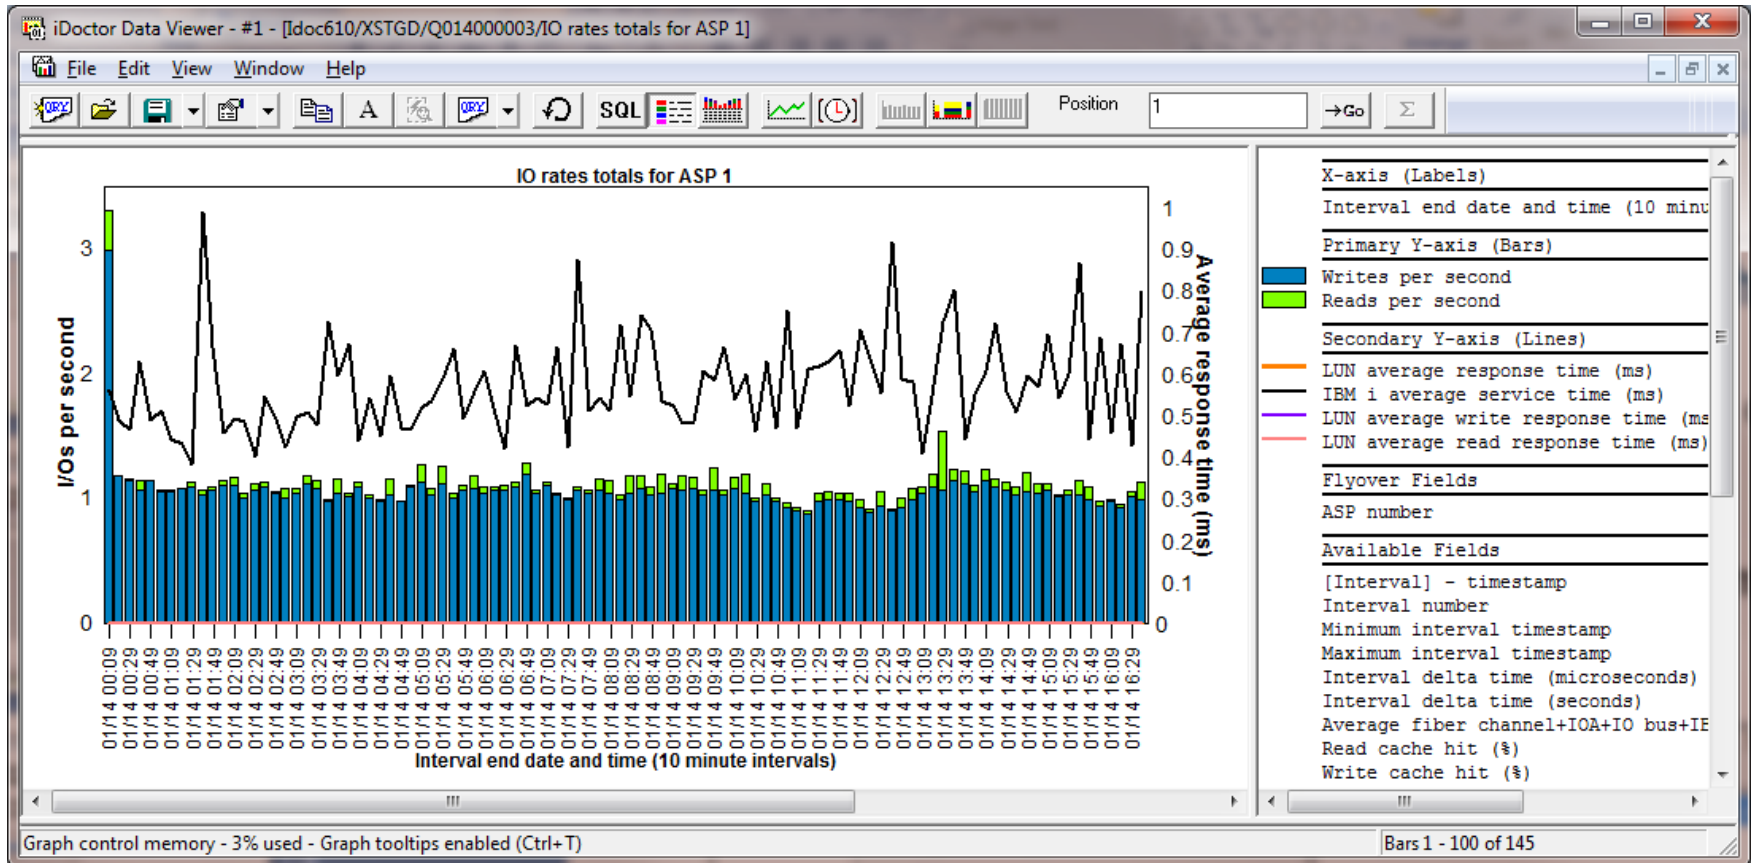
Description	Library	Collection(s)	VRM	Comments	SQL Table Name	Change date
CSI - External storage X32 by unit	XSTGD	Q014000003	V7R1M0		Qaidrcsext_units_q014000003	2010-10-05-08

- Open Table(s)
- Edit
- External storage cache statistics
 - By time interval
 - Average DS to IBM i latency for ASP <<DSASP>>
 - IO counts totals for ASP <<DSASP>>
 - IO size totals for ASP <<DSASP>>
 - IO size averages for ASP <<DSASP>>
 - IO size rates with cache hits for ASP <<DSASP>>
 - IO rates totals for ASP <<DSASP>>
 - IO rates totals with cache hits for ASP <<DSASP>>
 - Cache operation rates for ASP <<DSASP>>
 - NVS space allocation rates for ASP <<DSASP>>
 - Record mode read rates for ASP <<DSASP>>
 - By disk unit
 - By ASP
- Delete
- Properties

CSI External Storage Cache Statistics

Jun-Oct 2010 – SQL Tables: CSI External Storage Cache Statistics

Example graph



Jun-Oct 2010 – SQL Tables: CSI External Storage Links

This data is new at 7.1 from file QAPMXSTGV.

Contains fiber channel links.

Created by the External storage links and ranks analysis.

Description	Library	Collection(s)	VRM	Comments	SQL Table Name	Change date
CSI - Link deltas	XSTGD	0014000003	V7R1M0		Qaidrcsext_linkd_q014000003	2010-10-08.30

- Open Table(s)
- Record Quick View
- External storage link and rank statistics ▶
- Edit Comment
- Delete...
- Properties

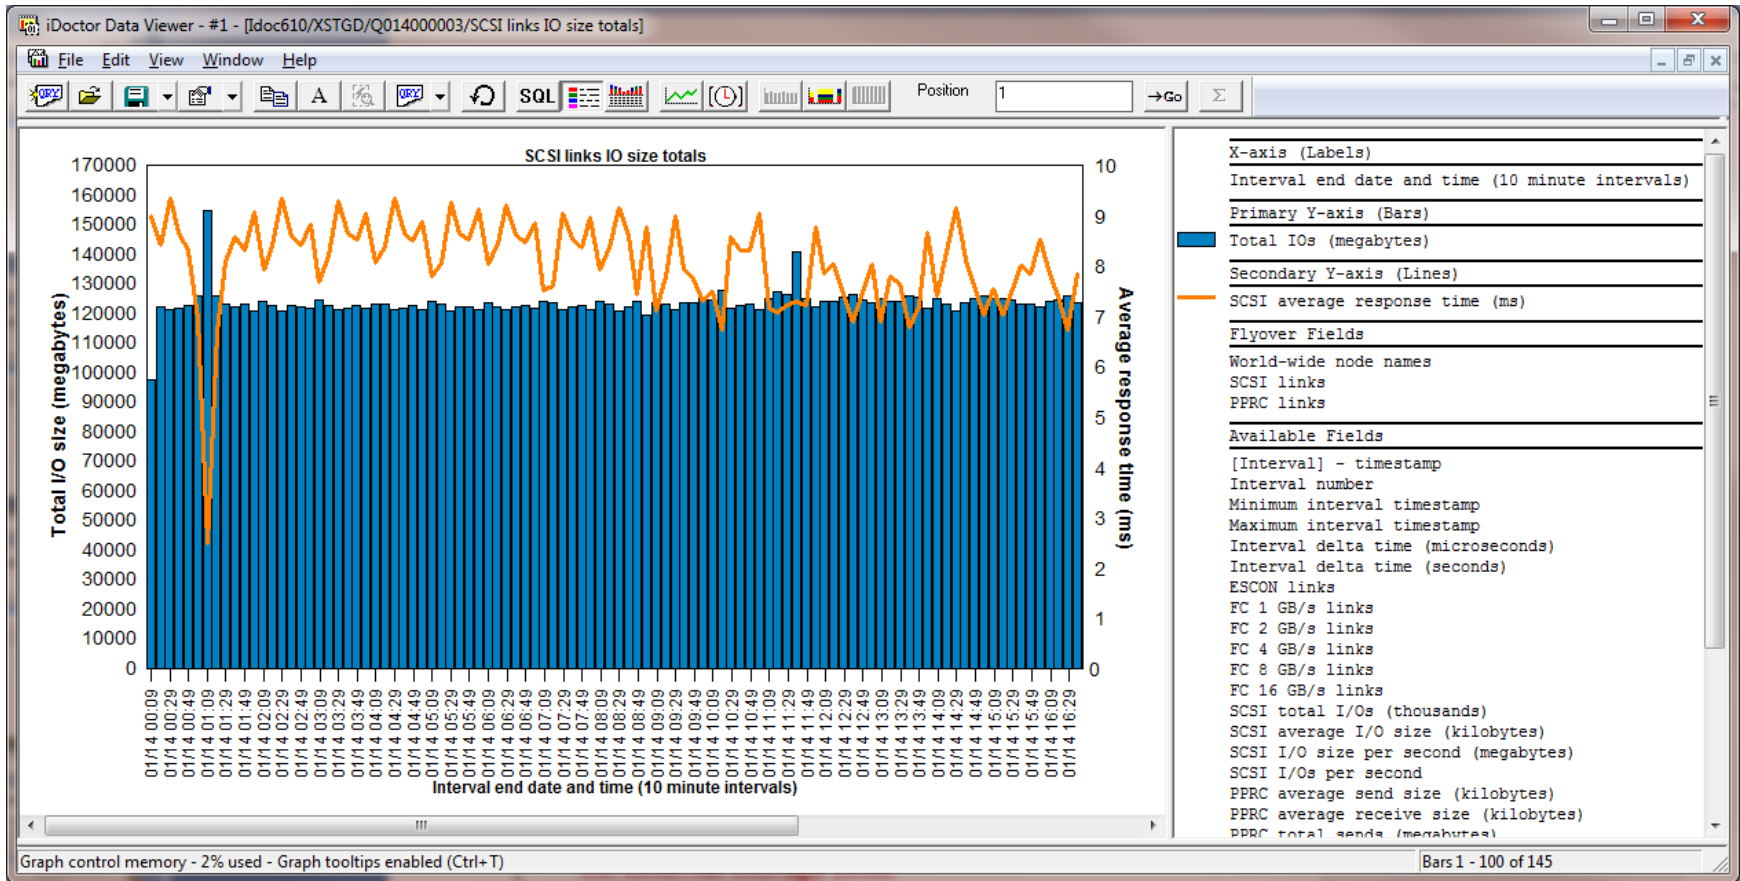
- Link graphs ▶
- Rank graphs ▶

- SCSI links IO counts totals
- SCSI links IO size totals
- SCSI links IO size averages
- SCSI links IO size rates
- SCSI links IO rates totals
- PPRC links IO counts totals
- PPRC links IO size totals
- PPRC links IO size averages
- PPRC links IO size rates
- PPRC links IO rates totals
- By interface ID ▶

CSI External Storage Links

Jun-Oct 2010 – SQL Tables: CSI External Storage Links

Example graph



Jun-Oct 2010 – SQL Tables: CSI External Storage Ranks

This data is new at 7.1 from file QAPMXSTGV.

Contains ranks (similar to an ASP).

Created by the External storage links and ranks analysis.

Description	Library	Collection(s)	VRM	Comments	SQL Table Name	Change date
CSI - Rank deltas	YSTGN	0014000003	V7R1M0		Qaidrcsext_rankd_q014000003	2010-10-08.30.23.

- Open Table(s) ▶
- Edit ▶
- External storage link and rank statistics ▶
- Delete
- Properties

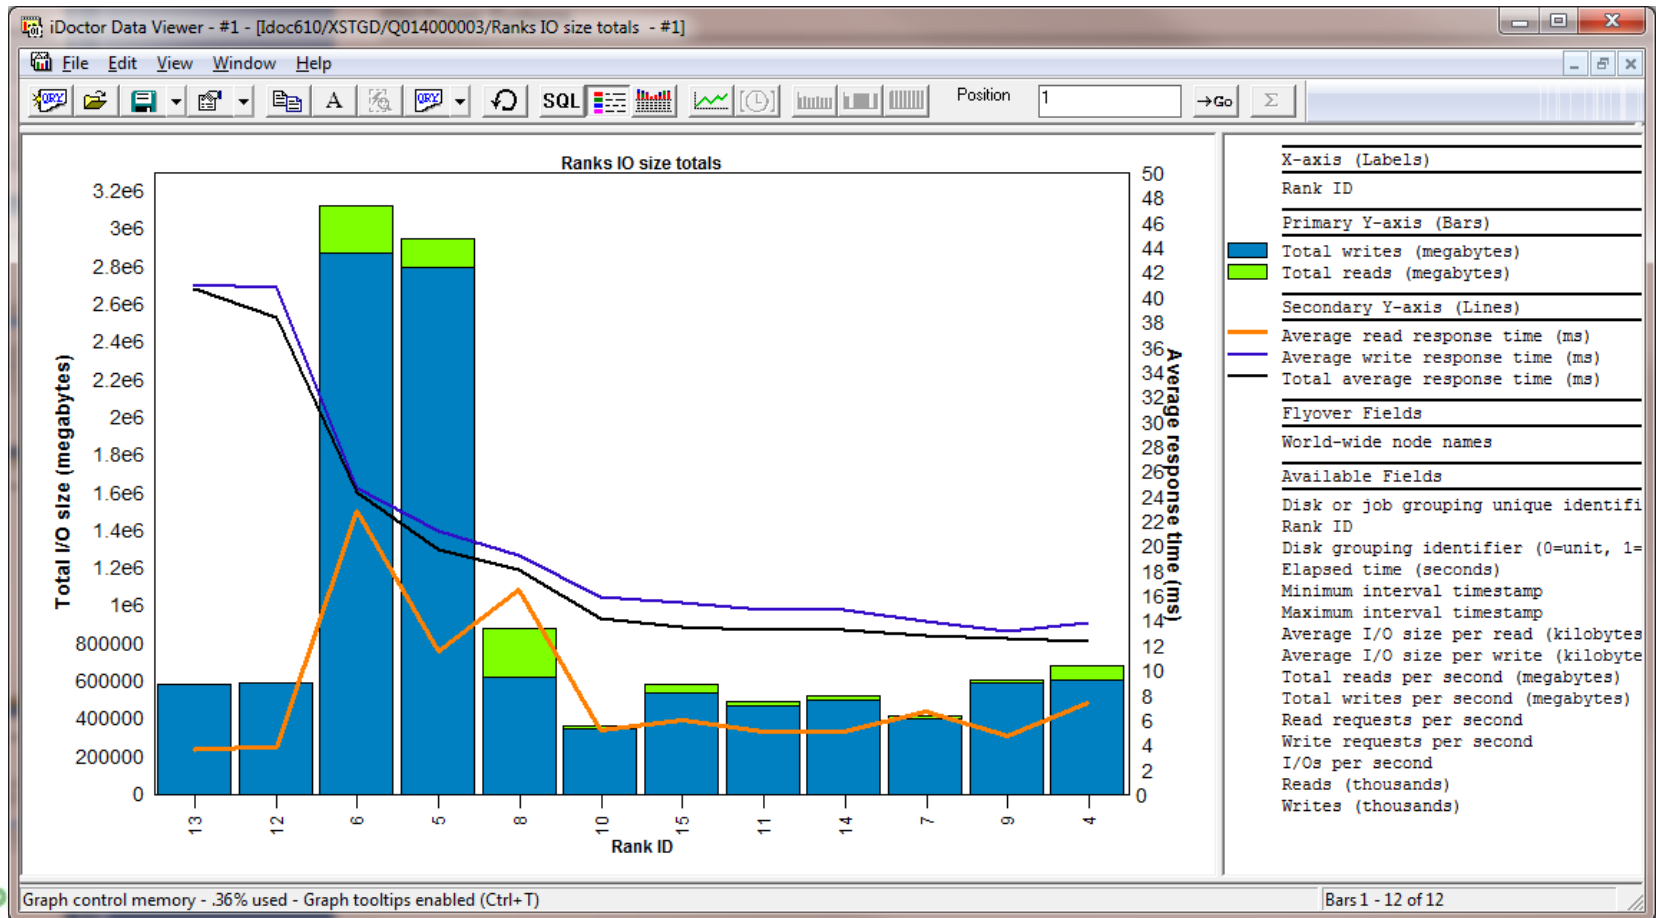
- Link graphs ▶
- Rank graphs ▶

- Ranks IO counts totals
- Ranks IO size totals
- Ranks IO size averages
- Ranks IO size rates
- Ranks IO rates totals
- By rank ID ▶

CSI External Storage Ranks

Jun-Oct 2010 – SQL Tables: CSI External Storage Ranks

This data is new at 7.1 from file QAPMXSTGV.
 Contains ranks (similar to an ASP).
 Created by the External storage links and ranks analysis.



Jun-Oct 2010 – SQL Tables: JW iDoctor Collection Summary

Graphs are filtered if they were used to create the SQL table.

Description	Library	Collection(s)	VRM	Comments	Job	Current user	Subs...	Start	End	SQL Table Name
Interval summary file	CRAVENS1	Q077115116	V6R1M0	QZDA jobs only	QZDA/**					Qaidrjwsum_q077115116_qzda_
Interval summary file	CRAVENS1	Q077115116	V6R1M0							Qaidrjwsum_q077115116

- Open Table(s)
- Record Quick View
- Wait graphs ▶
- CPU graphs ▶
- I/O graphs ▶
- IFS graphs ▶
- Other graphs ▶
- Edit Comment
- Delete...
- Properties

- Collection overview time signature
- Situational analysis overview time signature
- Collection overview with dispatched CPU time signature
- Collection overview with faulting breakdown time signature
- Current wait duration time signature
- Seizes and locks time signature
- Contention time signature
- Disk time signature
- Classic JVM time signature
- Communications time signature
- By thread
- Collection totals

JW iDoctor Interval Summary

Jun-Oct 2010 – SQL Tables: JW Long Transactions

This table contains 1 record per job/task for every time period > 1 interval that occurred where no 'idle' wait buckets were found.
These are NOT normal 5250 transactions!

Created by the Long Transactions Analysis.

Description	Library	Collection(s)	VRM	Comments	SQL Table
Job Watcher long transactions	CRAVENS1	Q077115116	V6R1M0		Qaidrjwtx

- Open Table(s)
- Edit
- Long transactions
 - Long transactions for DB server jobs
 - Long transactions for all jobs
- Delete
- Properties

Jun-Oct 2010 – SQL Tables: PEX Stats Summary for All Jobs

The drill down report from the previous slide does not allow further drill downs (where pgm used, etc). For that, use the Stats hier for selected job analysis.

Call level (CALLLVL)	Partial count status (PCSTS)	Library name (LIBNAM)	Program name (PGMNAM)	MI complex instruction (MICPXNM)	Module name (MODNAM)	Procedure name (PRCNAM)	Times called (CALLCOUNT)	Calls made (CALLMADE)	Calls to MI complex instructions (CALLMICPX)	Inline elapsed us (INELPUS)	Inline CPU us (INCPUUS)	Inline percent CPU (INPCPUUS)	Inline per c (INCE)
0	Y	QSYS	QZRCRSVS		QZRCRSVS	_C_pep	0	1	0	.9080	0	0	0
1	Y	QSYS	QZRCRSVS		QZRCRSVS	main	0	661	0	686.4640	53.7070	.0342	0
2	Y	QSYS	QZRCRSVS		QZRCRMT	RunCommand	20	85	0	243.4250	59.9970	.0382	2.99
3	Y	QSYS	QCPCMD		QCPCMD	QCPCMD	20	88	152	617.0360	199.4640	.1269	9.97
4	Y	QIDRPA	QIDRPASTCP		QIDRPASTCP	QIDRPASTCP	0	13	3	44.8690	21.5690	.0137	0
5	N			*MATQAT			1	0	0	7.8520	1.6350	.0010	1.63
5	N			*STPDCTRD			1	0	0	213.8720	6.8670	.0044	6.86
5	N			*SETACST			1	0	0	627.4920	6.4050	.0041	6.40
5	N	QSYS	QCLCLCPR		QCLCLCPR	QCLCLCPR	1	0	3	46.7450	19.3150	.0123	19.3
6	N			*MATPTR			1	0	0	3.0420	.6390	.0004	.639
6	N			*RSLVSP			2	0	0	12.3090	5.7730	.0037	2.88
5	N	QSYS	QSNDDTAQ		QSNDDTAQ	QSNDDTAQ	1	0	6	14.4290	2.9950	.0019	2.99
6	N			*LOCKSL			1	0	0	2.8320	.6610	.0004	.661
6	N			*ENQ			1	0	0	3.8900	1.2430	.0008	1.24
6	N			*MATQAT			1	0	0	3.0420	1.1520	.0007	1.15
6	N			*STPDCTRD			1	0	0	.8210	.0760	.0000	.076
6	N			*UNLOCKSL			1	0	0	1.6170	.3630	.0002	.363
6	N			*SETACST			1	0	0	440.9370	10.6880	.0068	10.6
5	N	QSYS	QCADRV		QCADRV	QCADRV	3	12	6	52.1440	13.6580	.0087	4.55
6	N			*MATPTR			6	0	0	14.1270	3.2860	.0021	.547
6	N	QSYS	QCARULE		QCARULE	QCARULE	3	0	9	25.4230	6.4270	.0041	2.14
7	N			*RSLVSP			3	0	0	28.3530	10.5210	.0067	3.50
7	N			*MATPTR			3	0	0	4.9770	1.5630	.0010	.521
7	N			*TESTAU			3	0	0	3.9080	.6210	.0004	.207
6	N	QSYS	QCAPOS		QCAPOS	QCAPOS	3	3	0	58.2260	28.5710	.0182	9.52
7	N	QSYS	QCAFSCAN		QCAFSCAN	QCAFSCAN	3	0	0	6.0450	2.3010	.0015	.767
6	N	QSYS	QCAFLD		QCAFLD	QCAFLD	3	0	0	64.8100	36.1480	.0230	12.0
6	N	QSYS	QCAIFLD		QCAIFLD	QCAIFLD	3	0	0	17.9630	9.7790	.0062	3.25
5	N	QSYS	QCATRS		QCATRS	QCATRS	3	0	3	25.0190	8.8160	.0056	2.93

Jun-Oct 2010 – SQL Tables: PEX Stats hier for selected job

After opening one of the Stats hier for selected job SQL tables, the following “stats hier classic” options are available:

- Selected program -> Where program XYZ used (up to 5 call levels)
- Selected program -> Where program XYZ used (up to 10 call levels)

- Selected procedure -> Where procedure XYZ used (up to 5 call levels)
- Selected procedure -> Where procedure XYZ used (up to 10 call levels)

- Selected MI instruction -> Where MI instruction XYZ used (up to 5 call levels)
- Selected MI instruction -> Where MI instruction XYZ used (up to 10 call levels)

Call Level (CALLLVL)	Partial Count Status (PCSTS)	Library Name (LIBNAM)	Program Name (PGMNAM)	MI Complex Instruction (MICPXNM)	Module Name (MODNAM)	Procedure Name (PRCNAM)	Times Called (CALLCOUNT)	Calls Made (CALLMADE)	Calls to MI Complex Inst (CALLMICPX)	Inline Elapsed us (INELPUS)	Inline CPU us (INCPUS)	Inline Percent CPU (INPCPUUS)	Inline us per C (INCP)
0	Y	QSYS	QZRCSRVS		QZRCSRVS	_C_pep	0	1	0	.9080	0	0	
1	Y	QSYS	QZRCSRVS		QZRCSRVS	main	0	661	0	686.4640	53.7070	.0342	
2	Y	QSYS	QZRCSRVS		QZRCRMT	RunCommand	20	85	0	243.4250	59.9970	.0382	2
3	Y	QSYS	QCPCMD		QCPCMD	QCPCMD	20	88	152	617.0360	199.4640	.1269	9
4	Y	QIDRPA	QIDRPASTCP		QIDRPASTCP	QIDRPASTCP	0	1	0	11.0000	21.5500	.0137	
5	N			*MATQAT			1						
5	N			*STPDCTRD			1						
5	N			*SETACST			1						
5	N	QSYS	QCLCLCPR		QCLCLCPR	QCLCLCPR	1			327.4920	6.4050	.0041	6
6	N			*MATPTR			1			46.7450	19.3150	.0123	19
6	N			*RSLVSP			2			3.0420	.6390	.0004	
5	N	QSYS	QSNDDTAQ		QSNDDTAQ	QSNDDTAQ	1			12.3090	5.7730	.0037	2
										14.4290	2.9950	.0019	2

Jun-Oct 2010 – Create Job Summary

Create Job Summary interface changes:

Create Job Summary - Idoc610

Use this function to query job statistics for the desired collections and produce totals for each job/thread based on the job name, subsystem filters provided.

Tip: Leave job name and subsystem name fields blank to include statistics for all jobs.

Library: Cravens1

Collection(s):

Collection name
Q290000002
Test

Add >>

Job name contains:

Subsystem name contains:

Time range (optional):

Start time: 2008-05-08-00.00.02

End time: 2009-10-18-00.00.00

Comments:

Collections to summarize:

Collection name
CRAVENS1/Q290000002(610)
CRAVENS1/TEST(610)

VRM

Can include collections from multiple libs

Remove Remove All

SQL table creation options:

Library: Cravens1

Job Totals **New output options**

Thread Totals **options**





Job Totals by Collection

Thread Totals by Collection

Submit Cancel

Jun-Oct 2010 – Create Job Summary

Depending on the desired outputs, 1 table is created into the following 4 SQL table folders:

Analysis	Description
 Job Totals	Create Job Summary function output grouped by job
 Thread Totals	Create Job Summary function output grouped by thread
 Job Totals by Collection	Create Job Summary function output grouped by job and collection
 Thread Totals by Collection	Create Job Summary function output grouped by thread and collection

Drill down options from the SQL tables was covered in previous slides.

Jun-Oct 2010 – Create Job Summary - Future

More testing needed for CSI comparison of collections of multiple releases. In JW however, you can run create job summary with multiple release collections if desired with client 807.

Add merge SQL tables support.

Add multiple (generic) job selection and give user control of wildcards. QZD* vs *QZD*

Add pool selection.

Summarize at higher groupings like generic job or pool and/or entire collection.

Anything else?

Jun-Oct 2010 – Visual Studio redistributable package

This was previously included in the install (32-bit only) and installed automatically but now first-time users will need to pick the appropriate version of Windows they have and install it manually.

This was changed to better support 64-bit users and avoiding increasing download sizes too much.

The following is displayed on all download pages on the website 5.4+. Many users have missed it ;)

New User GUI Installation Requirement

Before installing the iDoctor GUI for the first time you are required to install the Microsoft Visual Studio Redistributable package from one of the links below depending on the type of CPU and OS level you have. [How do I tell if my computer is running 32-bit or 64-bit Windows?](#)

- ➔ Microsoft Visual C++ 2008 SP1 Redistributable Package (x86, 32-bit)
- ➔ Microsoft Visual C++ 2008 SP1 Redistributable Package (x64, 64-bit)
- ➔ Microsoft Visual C++ 2008 SP1 Redistributable Package (ia64, 64-bit Itanium)

Jun-Oct 2010 – Monitors enhancements

Starting a monitor at 5.4 and higher, will now adjust the definition's ASP limit value to be the same as the one specified on the monitor.

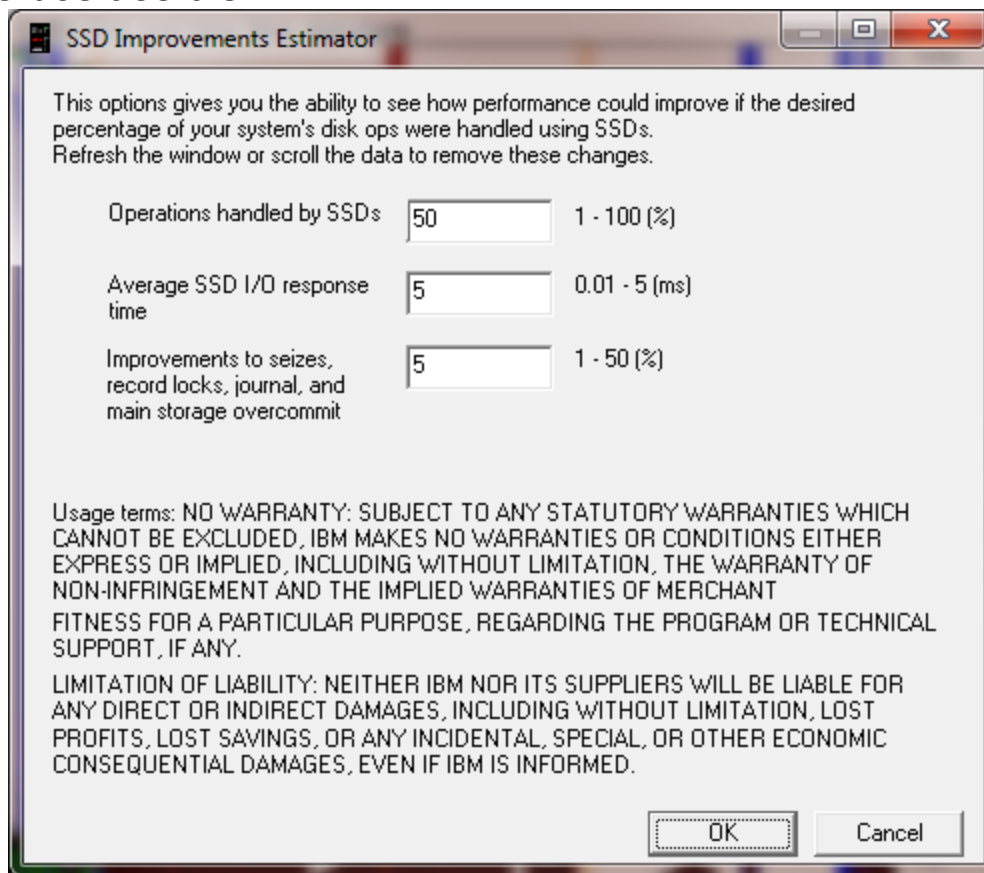
Note: This change actually caused a recent STRJW command failure if lib ASP is not 1 which requires a PTF at 6.1+ so I partially backed this out in 807. Only the system ASP limit is changed at 6.1+. I plan to change this back after PTFs are released and some time has passed.

For 5.3 and 5.4 JW, when starting a monitor the definition will be auto adjusted to make sure it runs at least as long as the collection's duration. Previously a message was shown asking the user to do this.

Jun-Oct 2010 –Misc: SSD Estimator now External

The SSD Estimator is now available for external use with client 806.

Alters wait bucket graphs in CSI/JW to show possible disk time reductions if SSDs were installed based on % of SSDs, avg SSD I/O response time and estimated improvements to other types of waits besides disk.



The screenshot shows a dialog box titled "SSD Improvements Estimator". It contains the following text and controls:

This options gives you the ability to see how performance could improve if the desired percentage of your system's disk ops were handled using SSDs.
Refresh the window or scroll the data to remove these changes.

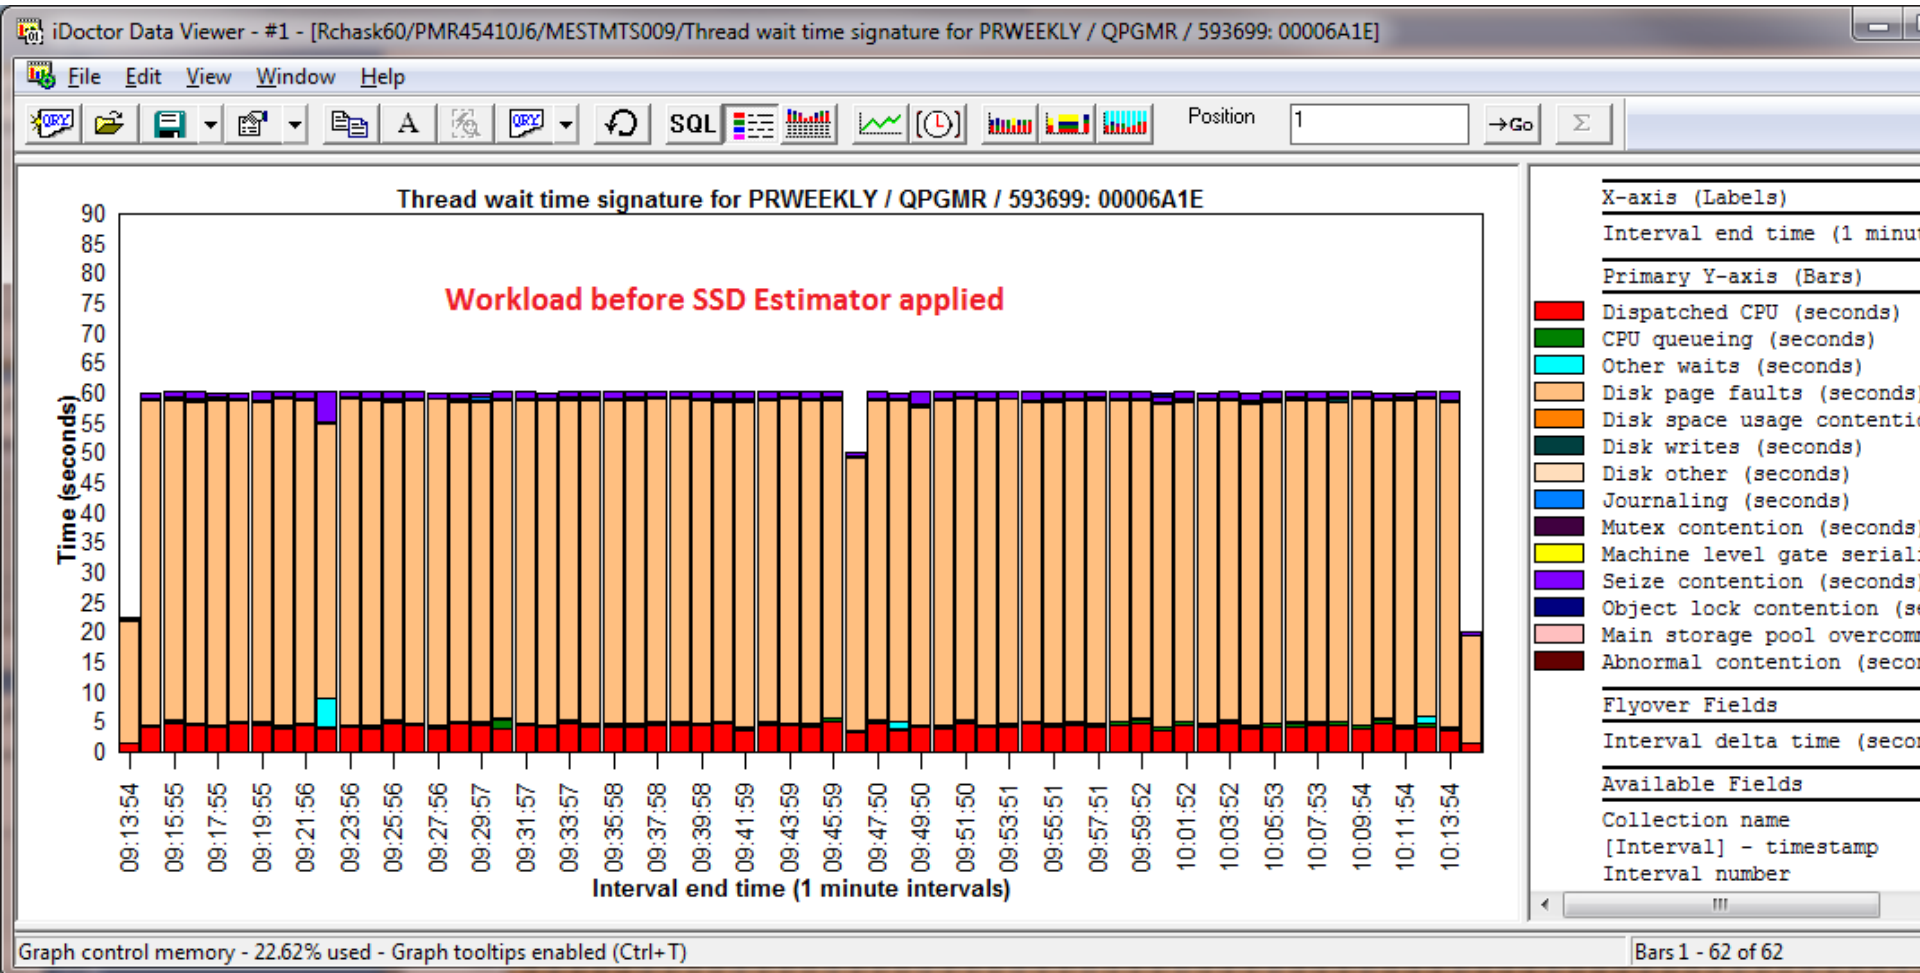
Operations handled by SSDs	<input type="text" value="50"/>	1 - 100 (%)
Average SSD I/O response time	<input type="text" value="5"/>	0.01 - 5 (ms)
Improvements to seizures, record locks, journal, and main storage overcommit	<input type="text" value="5"/>	1 - 50 (%)

Usage terms: NO WARRANTY: SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM MAKES NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, THE WARRANTY OF NON-INFRINGEMENT AND THE IMPLIED WARRANTIES OF MERCHANT FITNESS FOR A PARTICULAR PURPOSE, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

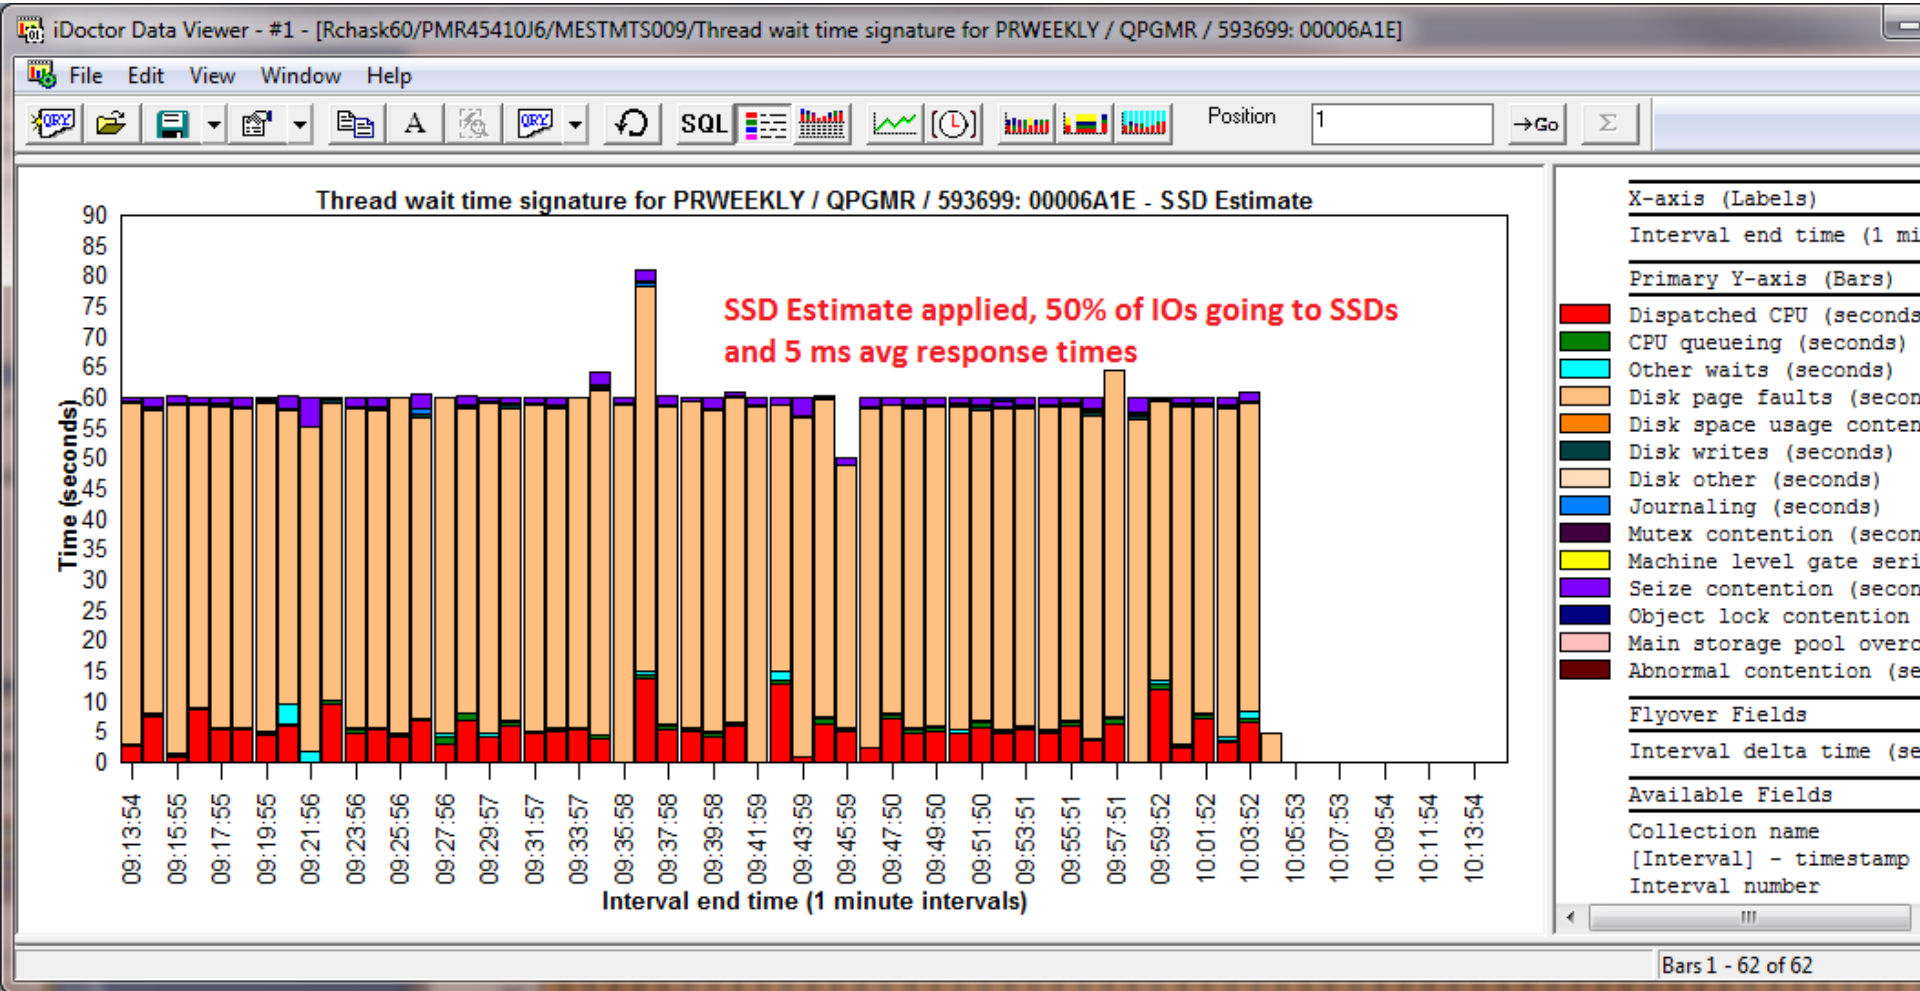
LIMITATION OF LIABILITY: NEITHER IBM NOR ITS SUPPLIERS WILL BE LIABLE FOR ANY DIRECT OR INDIRECT DAMAGES, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST SAVINGS, OR ANY INCIDENTAL, SPECIAL, OR OTHER ECONOMIC CONSEQUENTIAL DAMAGES, EVEN IF IBM IS INFORMED.

OK Cancel

Jun-Oct 2010 – Misc: SSD Estimator Example



Jun-Oct 2010 – Misc: SSD Estimator Example



Jun-Oct 2010 – Misc: JW SQL Server mode job info

For JW 6.1 (with PTFs) or 7.1 , the interval details property page now includes the SQL server mode client job if found with the option to drill down and graph the job.

At 6.1 the PTFs needed are:
 5761SS1 V6R1M0 SI39406
 5761999 V6R1M0 MF51454
 5761999 V6R1M1 MF51348

The screenshot shows the 'Interval' details page in IBM Job Watcher. The 'General' section contains the following information:

- Primary thread: QSQSRVR / QUSER / 160480: 00000051
- Interval: 6
- Job subsystem: QSYSWRK, Thread status: RUN
- Job function: Pool: 2
- Current user profile: DFL, Current state: RUN
- Priority (XPF/LIC): 20/160, Original LIC: 176
- Current or last wait: (351/Mcw) Condition wait
- Wait duration: 0 microseconds
- Object waited on: Segment type LIC HEAP (MWS) AREA DATA
- Interval duration: 5.035 seconds
- Holding job or task: None detected this interval
- Interval end: 2010-09-28-15.49.33.199000

The 'SQL client job' field is highlighted with a green box and contains: QPADEV000Q / DFL / 158948: 00000176.

Below this, the 'Wait bucket statistics' table is shown:

Bucket number	Description	Percent of Total Time	Time (seconds)	Occurrences per second	Current wait time
01	Dispatched CPU	72.76	3.66		
02	CPU queueing	3.69	.18		
30	Idle/waiting for work	23.55	1.18		

A context menu is open over the table, listing options: Wait graphs, I/O graphs, IFS graphs, Classic JVM graphs, J9 JVM graphs, and Other graphs.

Jun-Oct 2010 – Misc: PEX PDIO enhancements

PEX SQL-based PDIO has been enhanced in the following ways:

Added new rankings graphs:

1. by segment type
2. by object info (which is object name + object type/seg type)
3. by thread, object info
4. by thread, disk unit
5. by thread, object info, disk unit
6. by full object info: SID (this is object location/object name+ object type/seg type and virtual seg/object address)
7. by IFS path
8. IO adapter
9. IO adapter port

Added support for IFS. IFS path names are now surfaced in the analysis. By object groupings will now include a value IFS which is the total value for all IFS paths. The by IFS path grouping can be used to break out the statistics further.

In order for these options to work, the analysis must be reran.

Jun-Oct 2010 – Misc: PEX PDIO enhancements

PEX PDIO Rankings Graphs

Report	Folder description
I/O times by thread	
I/O times by job	
I/O times by job user	
I/O times by generic job name	
I/O times by memory pool	
I/O times by disk unit	
I/O times by ASP	
I/O times by object	
I/O times by object type	
I/O times by segment type	
I/O times by object info	
I/O times by object location	
I/O times by IO adapter	
I/O times by IO adapter port	
I/O times by thread, object info	
I/O times by thread, disk unit	
I/O times by thread, object info, disk unit	
I/O times by full object info: SID	
I/O times by IFS path	
Times	Ranking graphs showing I/O times
Rates	Ranking graphs showing I/O rates
Counts	Ranking graphs showing I/O counts
Sizes	Ranking graphs showing I/O sizes

Jun-Oct 2010 – Misc: JW System collected on column

C00800 - Job Watcher - Enhancement

Added "system collected on" to list of fields that can be selected/viewed in list of JW collections

IBM iDoctor for IBM i C00805 - [Rchaskmb: Job Watcher - #1]

File Edit View Window Help

Collection	Status	Ending reason	iDoctor summary available/status	Collection size (MB)	System collected on VRM	System collected on	Last inte collected
SQL tables							
Q077115116	Ready for analysis	Time limit	Yes	45.95	V6R1M0	RCHASKMB	

Rchaskmb: Job Watcher\Libraries\Cravens1

Jun-Oct 2010 – Misc: PEX Call Stacks Analysis

C00793 - Pex Analyzer - Enhancement

In PEX 5.4+, added new reports to the call stacks analysis that will show the top programs causing opens and closes.

The screenshot displays the PEX Analyzer interface. On the left is a tree view of analysis categories, including 'Classic Analyses' and 'SQL-based Analyses'. The 'Call stacks' folder is expanded. The main area shows a 'Report' list with several options, including 'Top programs causing full opens (program QRNXIO/_QRNX_OPEN)' and 'Top programs causing closes (program QRNXIO/_QRNX_CLOSE)'. A data viewer window titled 'iDoctor Data Viewer - #1 - [Idoc610/CRAVENS23/TEST23C/Top progra...' is open, showing a table with the following data:

Program name (PGMNAME)	Procedure (PROCEDURE)	Total call stacks (TOTAL)
QIDRDWMSTS	QIDRDWMSTS	2

The status bar at the bottom of the data viewer indicates 'Rows 1 - 1 of 1'.

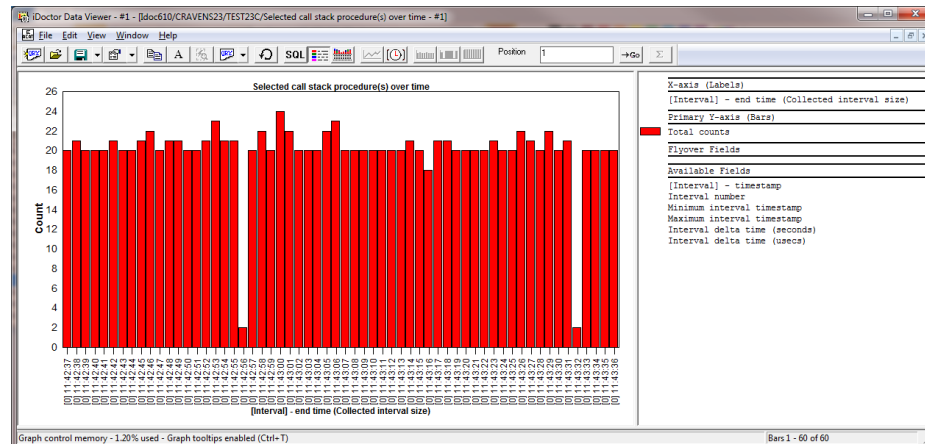
Jun-Oct 2010 – Misc: PEX Call Stacks Analysis

C00791 - Pex Analyzer - Enhancement

In PEX, the call stacks analysis now includes the trace back table addresses for each call level in the report. This allows for a new drill down option that lets you graph how many times the selected procedure(s) occurred in the collection over time.

Total call stacks (TOTAL)	Call level (LEVEL)	Program library (PGMLIB)	Program name (PGMNAME)	Module library (MODLIB)	Module name (MODNAME)	Procedure (PROCEDURE)
91479	0				#dbioest	#dbioest
91479	1				#dbioest	#dbioest
91479	2				#dbmatds	#dbmatds
91479	3				#cfmir	#cfmir
91479	4				cfscv0a	syscall_A_portal
91479	5	QSYS	QDBEXDME		QDBEXDME	QDBEXDME
91479	6	QSYS	QUSRMBRD		QUSRMBRD	QUSRMBRD
91479	7	MIMIX	LVSRV02	Z294134#	LVMBRD	lvGetMemberDescription0200
91479	8	MIMIX	LVSRV02	Z294134#	LVMBRD	
91479	9	MIMIX	DMAPPLY	Z294134#	DMASOPENF	
91479	10	MIMIX	DMAPPLY	Z294134#	DMAS2111	
91479	11	MIMIX	DMAPPLY	Z294134#	DMAS2111	

Selected procedures ▶ Selected call stack procedure(s) over time
Record Quick View
Conv



Jun-Oct 2010 – Misc: PEX TPROF Analysis

C00791 - Pex Analyzer - Enhancement

In PEX, the TPROF analysis now includes a tree report to show counts grouped by system data address register and then procedure name within.

The screenshot displays the PEX TPROF analysis interface. On the left, a tree view shows the analysis structure with folders like 'Classic Analyses' and 'SQL-based Analyses'. The main window shows a 'Report' table with various analysis options, where 'Hits by system data address register/procedure' is selected. Below this, a detailed data table is shown, listing system data address registers, their total hits, and associated component and program information.

Full name	Total hits	Component description	Library Name	Program Name	QPRMNM	Procedure Name	Traceback table Address
Total	100% - 202						
E312A7D5E308BBF8	1.49% - 3						
CF*	.99% - 2						
CFGRBLA/bla_Gennaker	.99% - 2	SLIC Common Functions	CFGRBLA	#cfgrbla		bla_Gennaker	FFFFFFF>
QX*	.50% - 1						
QXML4CINT/_BN_EXT_CALL_32	.50% - 1	MI Other	QSYS	QXML4CINT	BINDERGLUEMOD	_BN_EXT_CALL_32	215DADA>
E4AB0872AA08BBF8	.99% - 2						
D6ABAEFAA2097B0C	.99% - 2						
D2933F41F5000278	.99% - 2						
C000010B87011818	.99% - 2						
C00000A9BD0D1EA8	.99% - 2						
0000000035EF4198	.99% - 2						
0000000000004060	.99% - 2						
F25BCEFF9290000E0	.99% - 2						
FFFFFFFF804142C8	.99% - 2						