

# iDoctor What's New Oct 2011 – Feb 2012

(iDoctor GUI Builds 887-912)



**Ron McCargar**  
**iDoctor Development**

## iDoctor resources

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

## Oct-Feb 2012 – New Builds

Latest builds released in November 2011 and again on March 2<sup>nd</sup>, 2012.

Generally try to do 4 major updates a year.

# Oct-Feb 2012 - Overview

## Current Focus Areas

### Enhancements to Collection Services Investigator (CSI)

(12x loops, IASP Bandwidth analysis, seizes and locks graphs, etc)

### Transfer Function

### Remote SQL/Command Status View Enhancements

### Miscellaneous:

(PEX analysis updates, collection scheduling, run default analyses option, etc)

## Oct-Feb 2012– Current Focus Areas

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

### **Enhancing Collection Services Investigator**

There always seems to be more files/data to analyze

### **External storage**

Server-side VIOS data collection enablement

(due to issues, we may transition back to a non-native solution using NMON Investigator for analysis)

### **Plan Cache Analyzer Updates**

This component might become IBM external (part of Job Watcher) soon.

## Oct-Feb 2012 – CSI - IASP Bandwidth Analysis

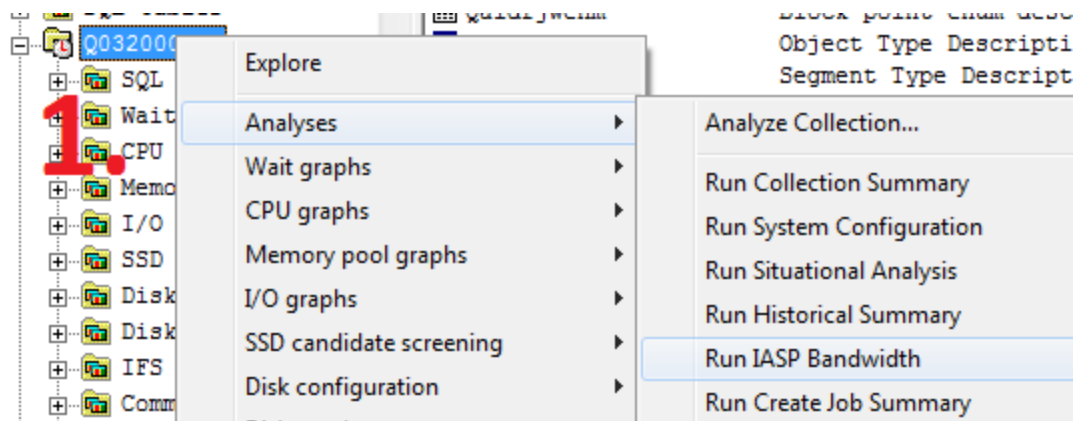
New analysis added to CSI to help a user determine if the system would be a good candidate for moving their workloads to Independent ASPs.

This is based on an Excel Spreadsheet created by David Frost in Lab Services.

Additional changes are expected soon to this analysis (waiting on feedback from Lab Services.)

# Oct-Feb 2012 – CSI - IASP Bandwidth Analysis

Right-click the collection and use the menu to kick off the analysis.



## Oct-Feb 2012 – CSI - IASP Bandwidth Analysis

When prompted, fill in the desired compression rate, bandwidth values (in megabits per second) and ASPs to include in the report and press OK.

Use these options to set the compression rate and desired bandwidth settings:

**2.** Compression rate:

Full system bandwidth:

IASP bandwidth:

ASP filtering (select ASPs to include):

ASP
1



# Oct-Feb 2012 – CSI - IASP Bandwidth Analysis

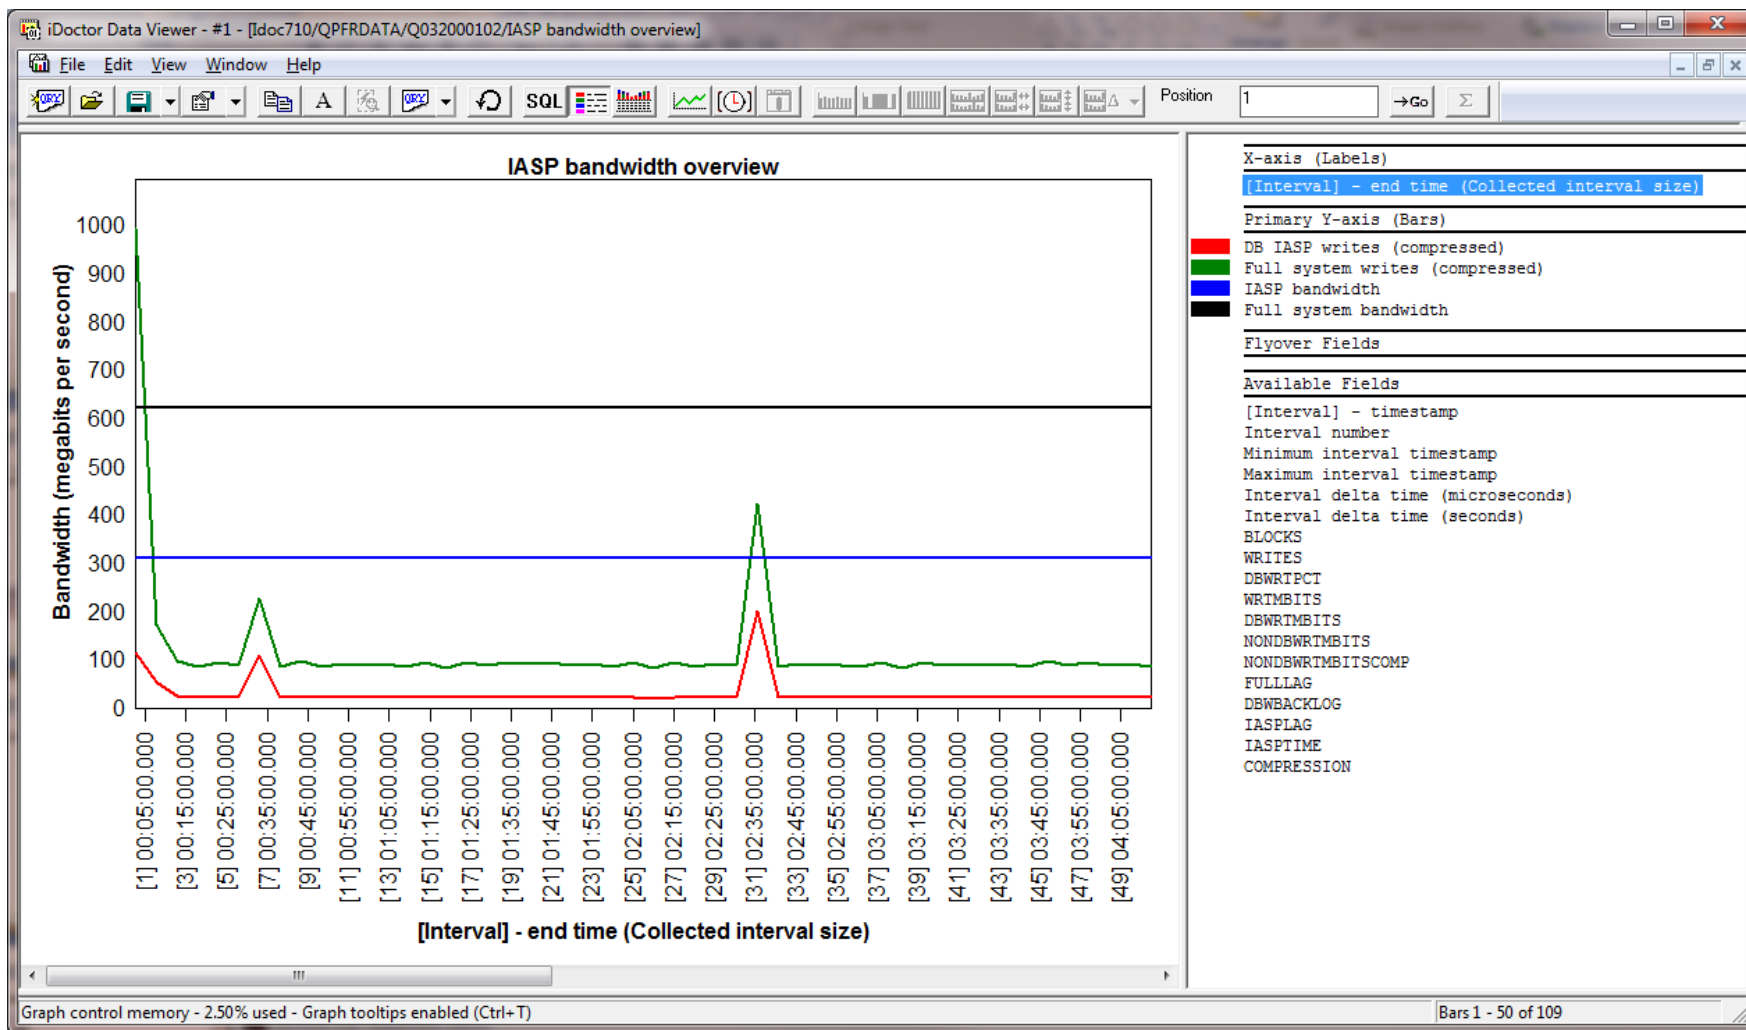
The analysis will run, refresh the collection and the new “IASP bandwidth estimations” folder will appear.

The screenshot displays the Performance Explorer interface. On the left, a tree view shows a collection named 'Qpfrdata' with various sub-folders like 'SQL tables', 'CPU graphs', and 'I/O graphs'. The 'IASP bandwidth estimations' folder is highlighted in blue. A red box encloses this folder name, and a red arrow points from it to a status message at the bottom of the window. The status message reads: 'IASP Bandwidth analysis created successfully for collection Q032000102 (1.404 seconds)'. A large red number '3.' is overlaid on the screenshot.

System	Status	SQL Statement
✓ Idoc710	IASP Bandwidth analysis created successfully for collection Q032000102 (1.404 seconds)	CALL QIDRGUI/QIDRCSIASP ('QPFRDATA', 'Q032000102', '1.5', '6

# Oct-Feb 2012 – CSI - IASP Bandwidth Analysis

Expand the new folder and open the desired reports.



## Oct-Feb 2012 – CSI – 12x loops and I/O hubs

In the Collection Services Investigator component under the Communication graphs folder.

iDoctor support to graph the 12x loop data has been available since July 2011 (client 863). On builds earlier than 887 look in the “RIO HSL loop graphs” folder.

For the latest updates client 888 or higher (released October 2011) is required. The subfolder name with this build is “12x loops and I/O hubs”.

Note: The folder will only appear if data has been collected in file QAPMBUSINT.

P.S. The location may change again in the future based on suggestions... (under “hardware graphs”?)

## Oct-Feb 2012 – CSI – 12x loops and I/O hubs

Ratsisb: Collection Services Investigator - #1

Report description

- Summarized loops/hubs traffic
- Summarized 12x loops traffic
- Summarized I/O hub traffic
- Total traffic breakdown
- Send and receive traffic breakdown
- 12x loops traffic breakdown
- I/O hub traffic breakdown
- Rankings
- Advanced


Summarizes the loop data over time.

Rankings Ranks each loop/port by throughput

Instead of summarizing the loop data, shows every loops data over time. Great for a small number of loops but gets cumbersome on 16 or more loops.


- Contains 3 styles of graphs: Summarized, rankings and advanced (shows all loops/ports over time)
- The 7 summarized graphs provide drill downs into 7 ranking graphs for the desired time period.

# Oct-Feb 2012 – CSI – Summarized loops/hubs traffic



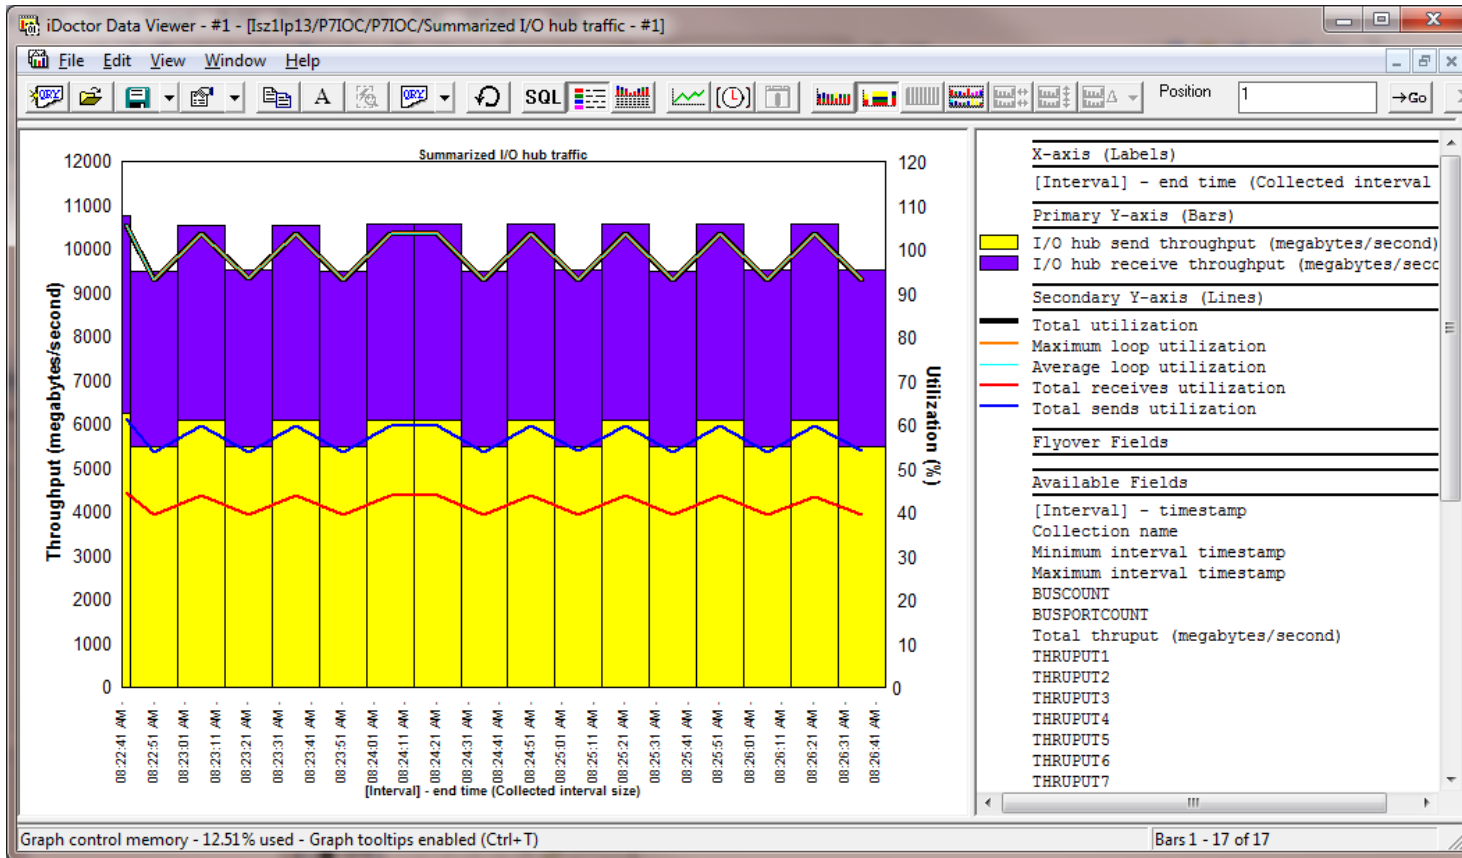
- Simply shows sends and receives throughput for loops and hubs (if available, hubs not shown in this data).
- Y2 shows utilization rates

# Oct-Feb 2012 – CSI - Summarized 12x loops traffic




- Same as previous, but does not show I/O hubs

# Oct-Feb 2012 – CSI - Summarized I/O hub traffic



- Same as Summarized loops/hubs traffic graph but only shows the I/O hubs

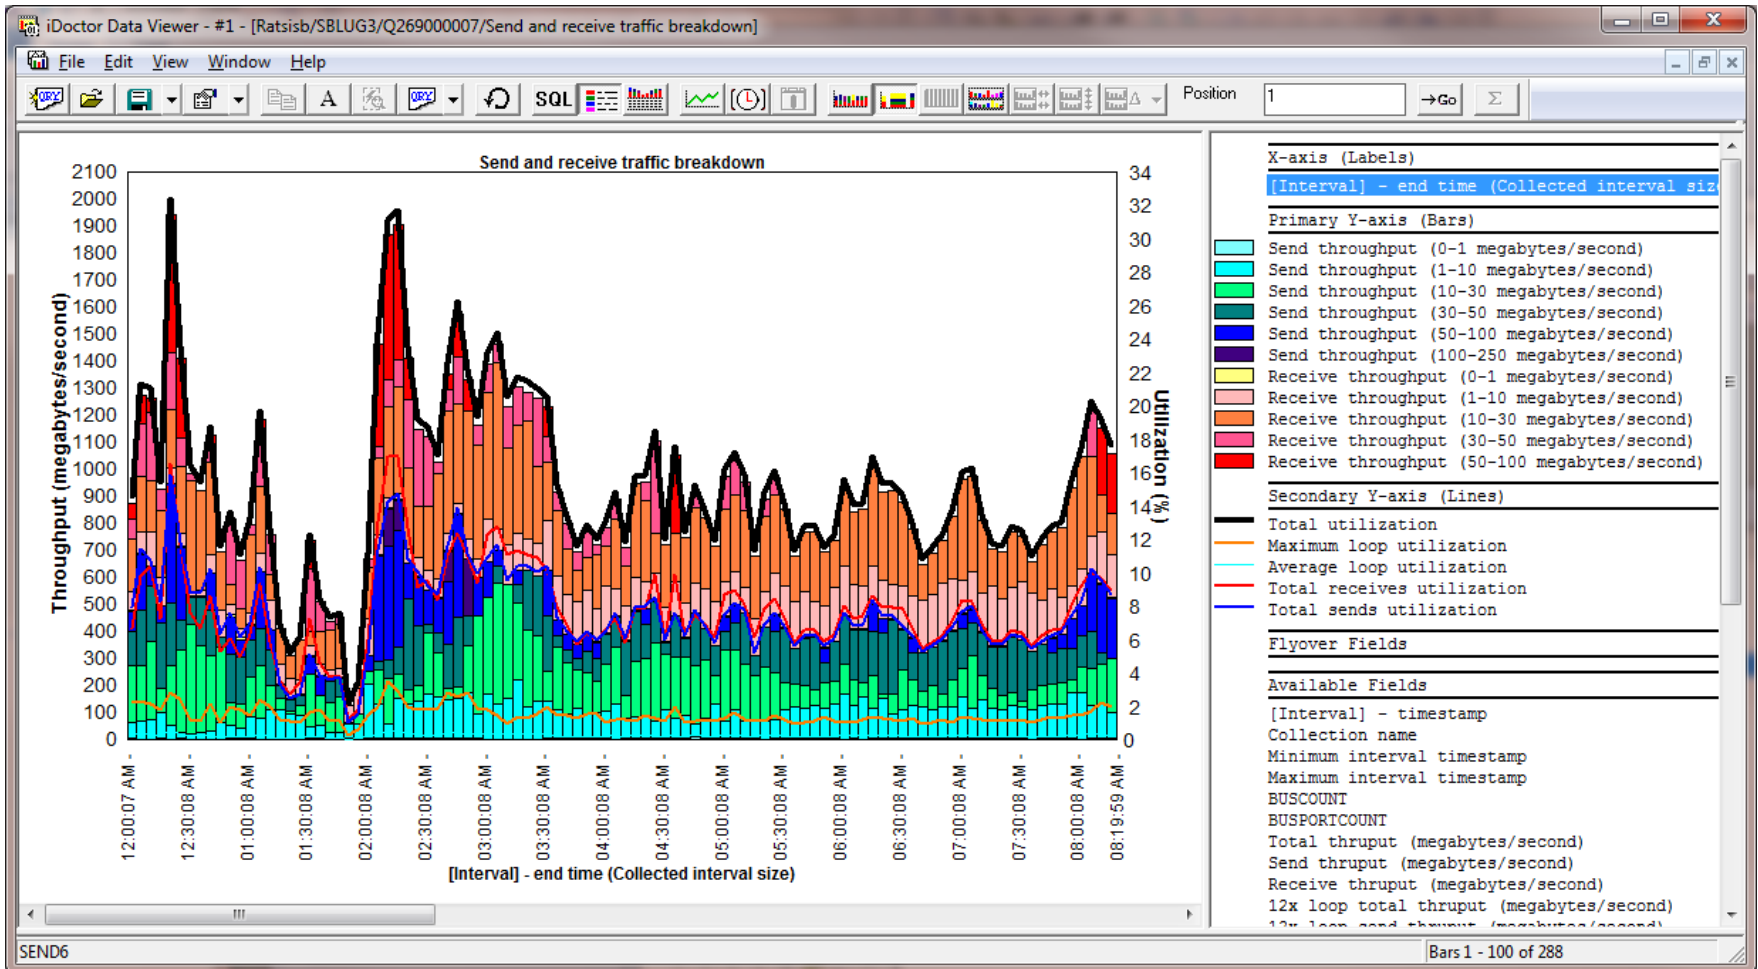
# Oct-Feb 2012 – CSI - Total traffic breakdown



- Shows sends and receives throughput for both loops and hubs broken down into 8 different “throughput” buckets.
- Useful if there are a large number of loops/hubs

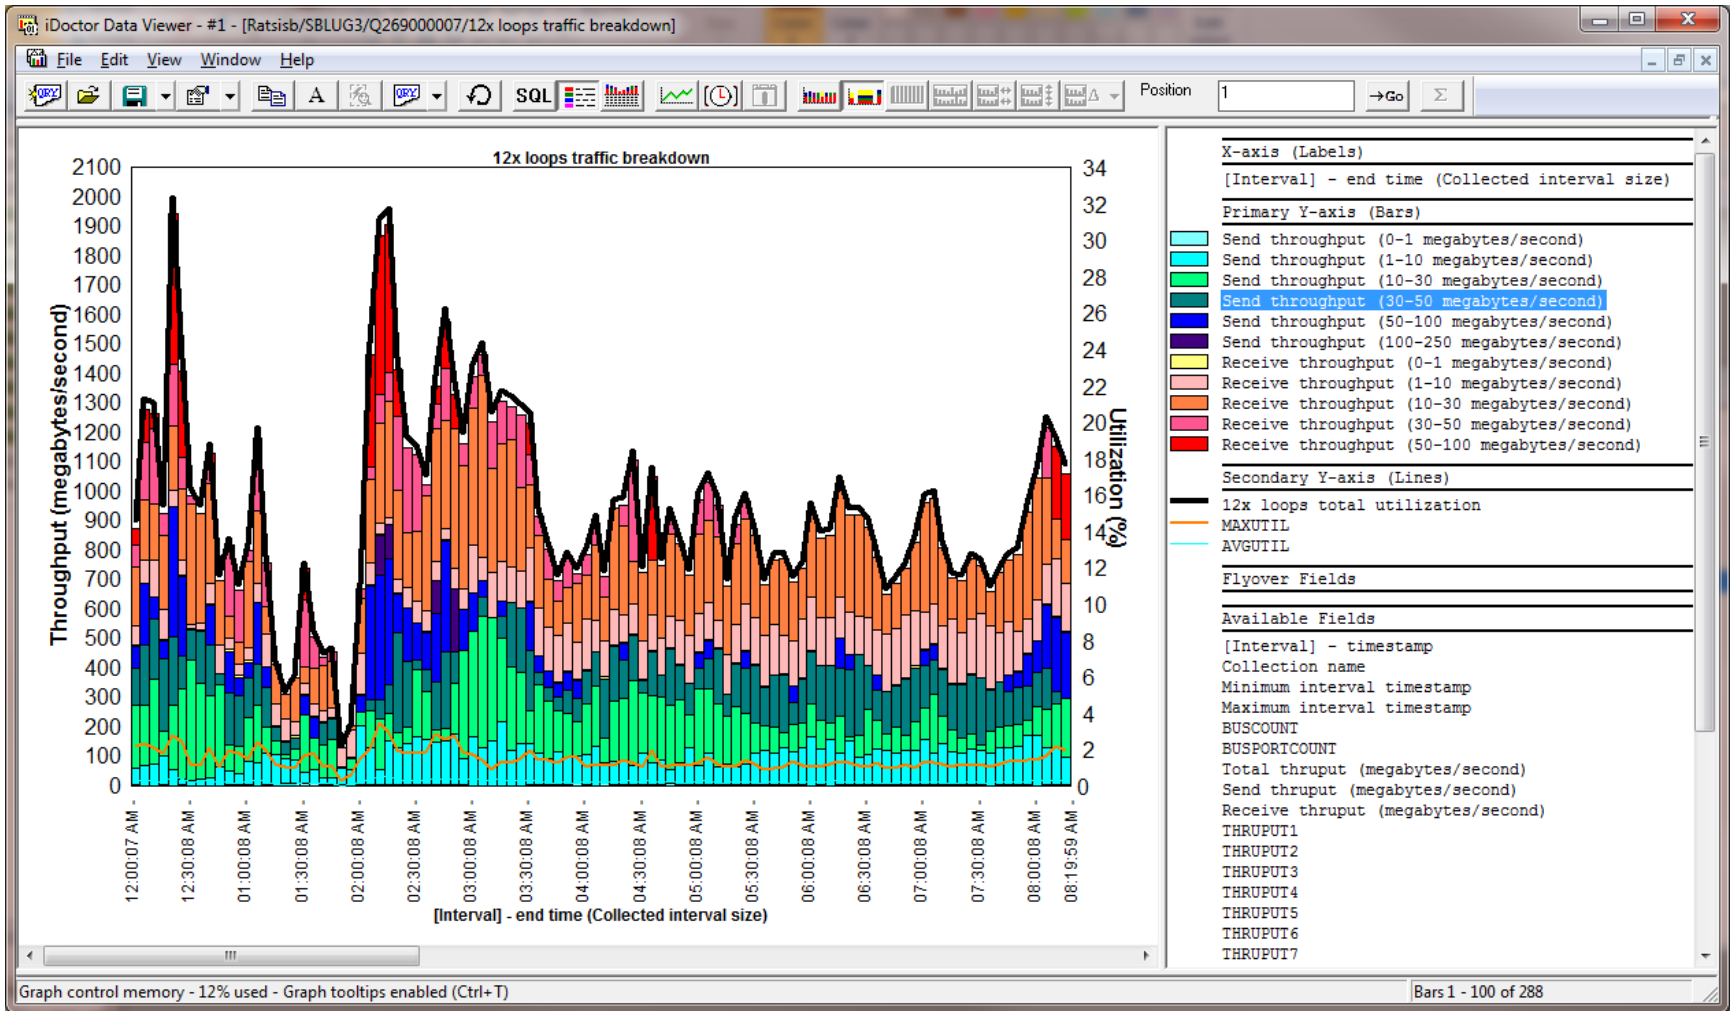


# Oct-Feb 2012 – CSI - Send and receive traffic breakdown



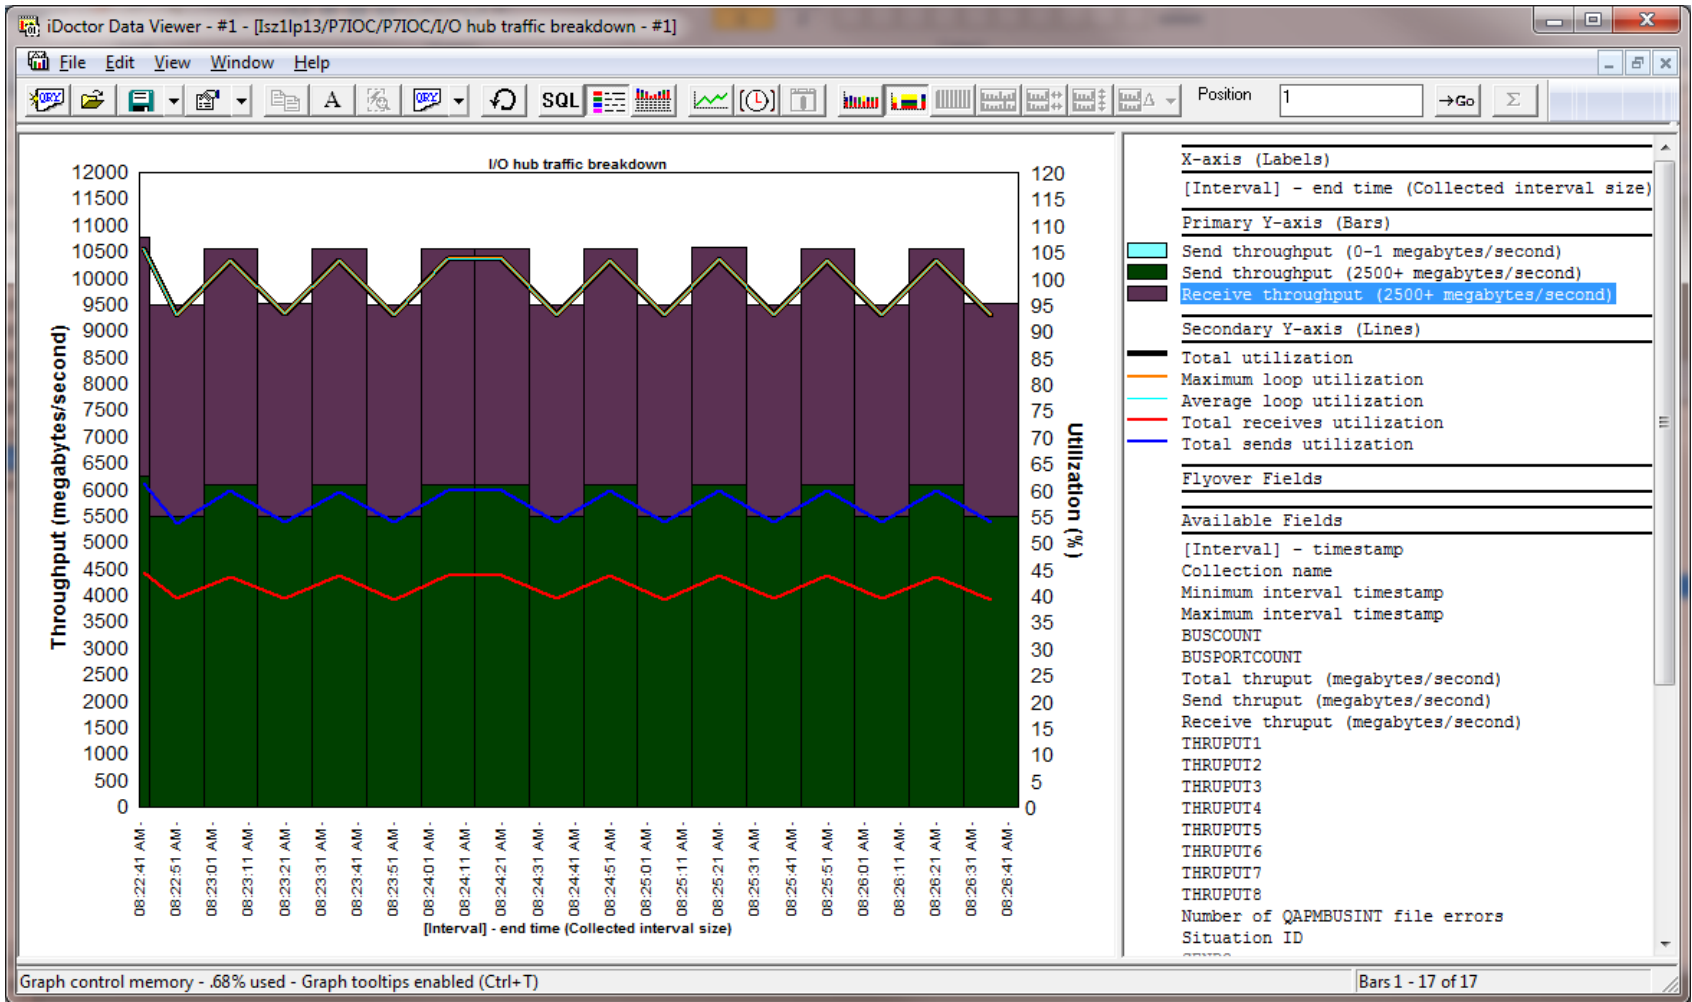
- Graphs sends and receives (for both loops and hubs) into 8 throughput buckets each.

# Oct-Feb 2012 – CSI - 12x loops traffic breakdown



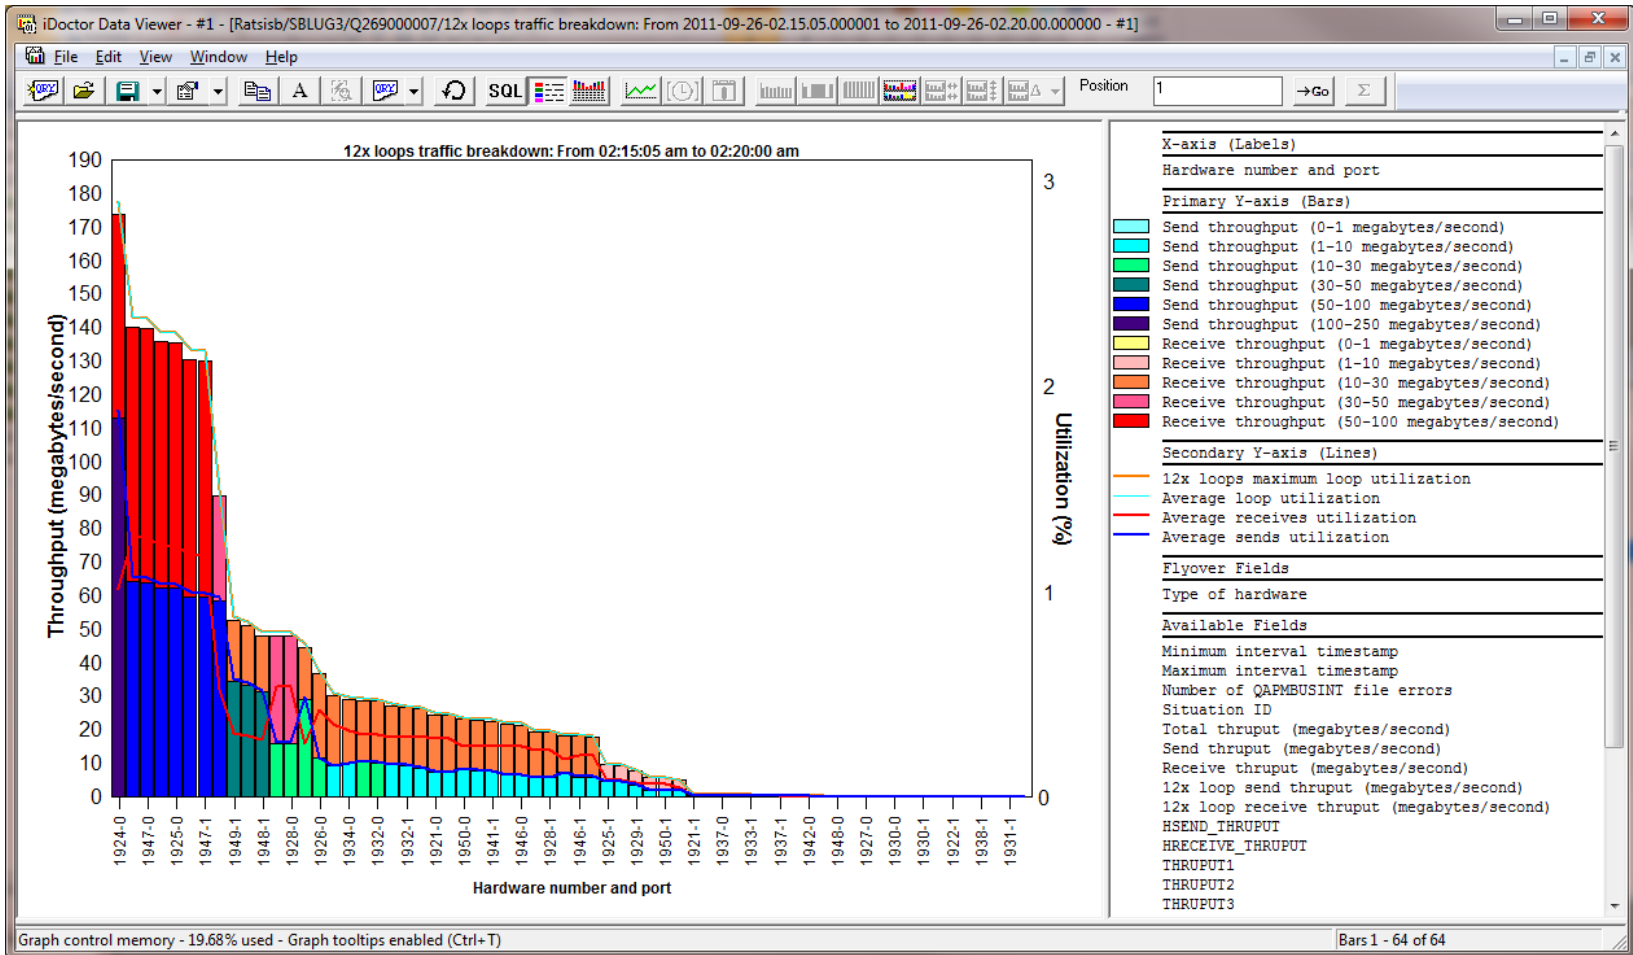
- Graphs sends and receives for 12x loops into 8 throughput buckets.

# Oct-Feb 2012 – CSI - I/O hub traffic breakdown



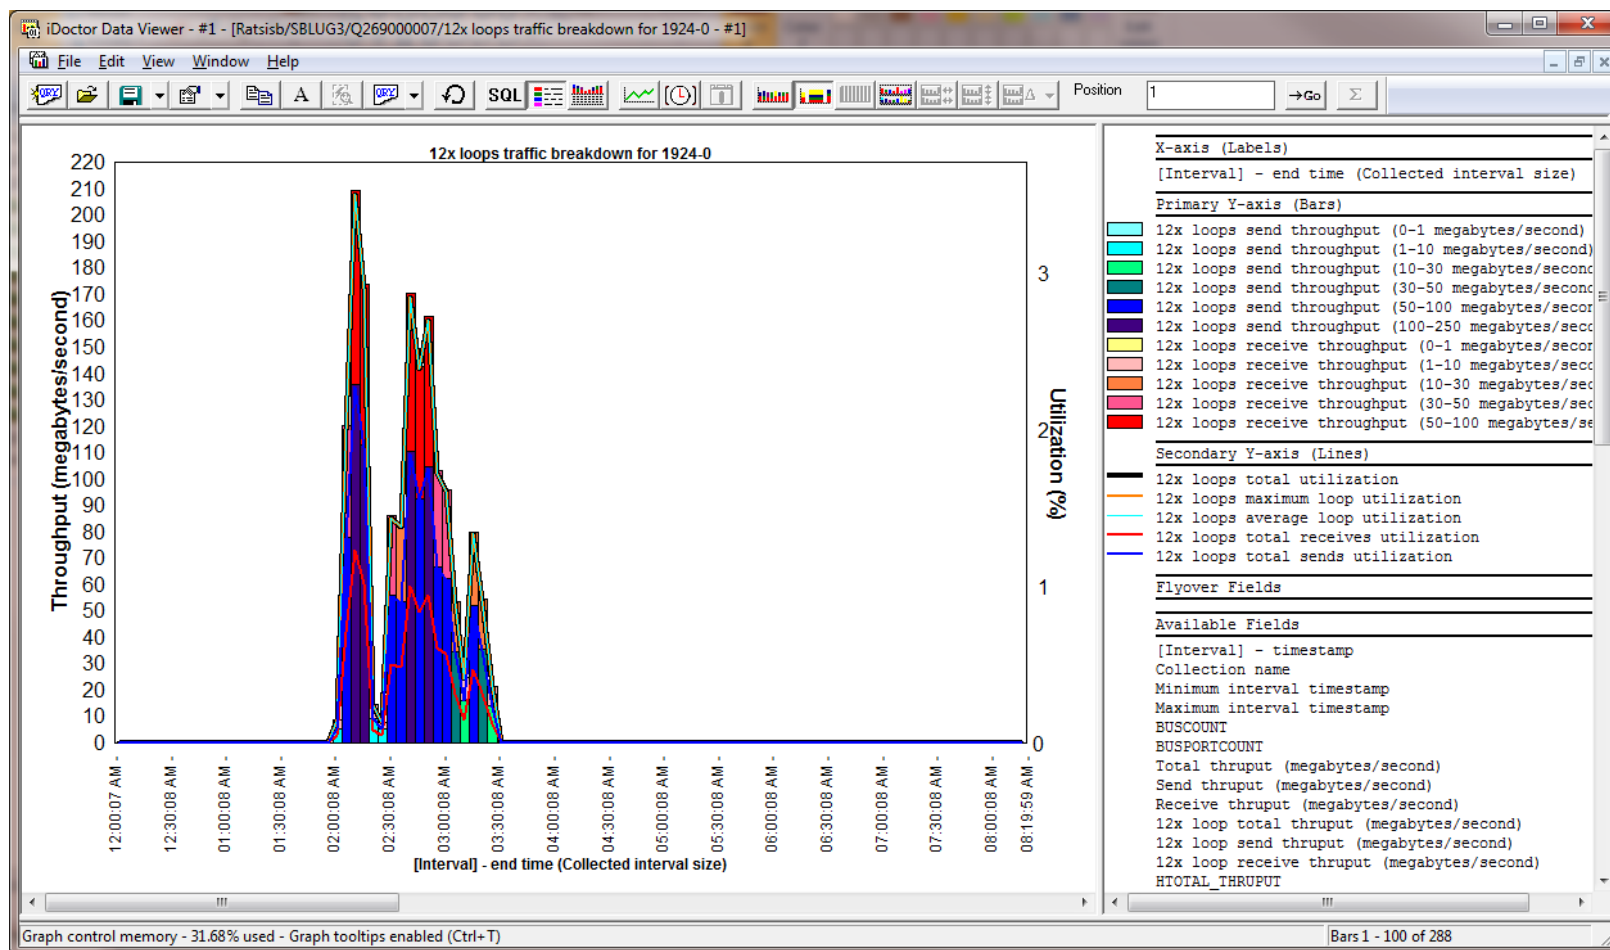
- Graphs sends and receives for I/O hubs only into 8 throughput buckets.

# Oct-Feb 2012 – CSI – 12x loops rankings graph example



- Ranks the loops by total throughput. From here a user can right-click and drill down again to see the loop data over time.


# Oct-Feb 2012 – CSI - 12x loops traffic for single loop



- Graphs the single loop over time.



# Oct-Feb 2012 – CSI – 12x loops advanced graph example



- These graphs show every loop/hub (above the filter threshold) over time.
- On some systems this works well, on others there are simply too many.

## Oct-Feb 2012 – CSI Create Job Summary Analysis

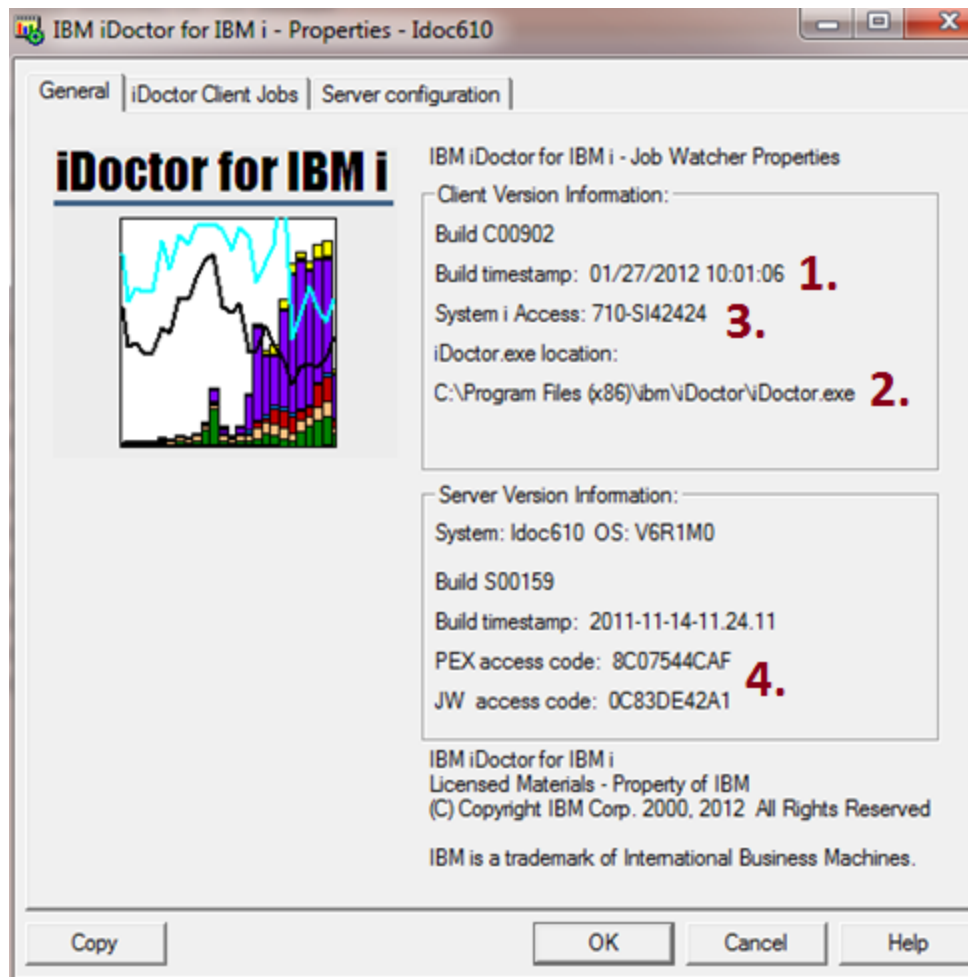
The Create Job Summary analysis will now utilize file QAPMJOBWTG to provide more accurate wait bucket times for intervals where jobs/threads did not use CPU.

The job duration calculation is also more accurate now.

## Oct-Feb 2012 – Application Properties changes

The iDoctor application properties - General tab, now provides the following additional information:

1. iDoctor.exe date and time
2. iDoctor.exe location
3. Client access VRM and service pack level
4. Access codes last applied.





## Oct-Feb 2012 – Plan Cache Analyzer Updates

Made the following changes/enhancements to Plan Cache Analyzer:

1. Plan cache snapshots can now be created using the GUI. Right-click the Plan Cache Analyzer icon and use the Create Plan Cache Snapshot menu.
2. Renamed the Start Plan Cache Collection menu/wizard to "Create Plan Cache Dump"
3. The server-side output files folder under a snapshot now correctly contains the snapshot table.
4. Added a search function over Plan Cache dumps.

## Oct-Feb 2012 – SQL Editor supports multiple statements

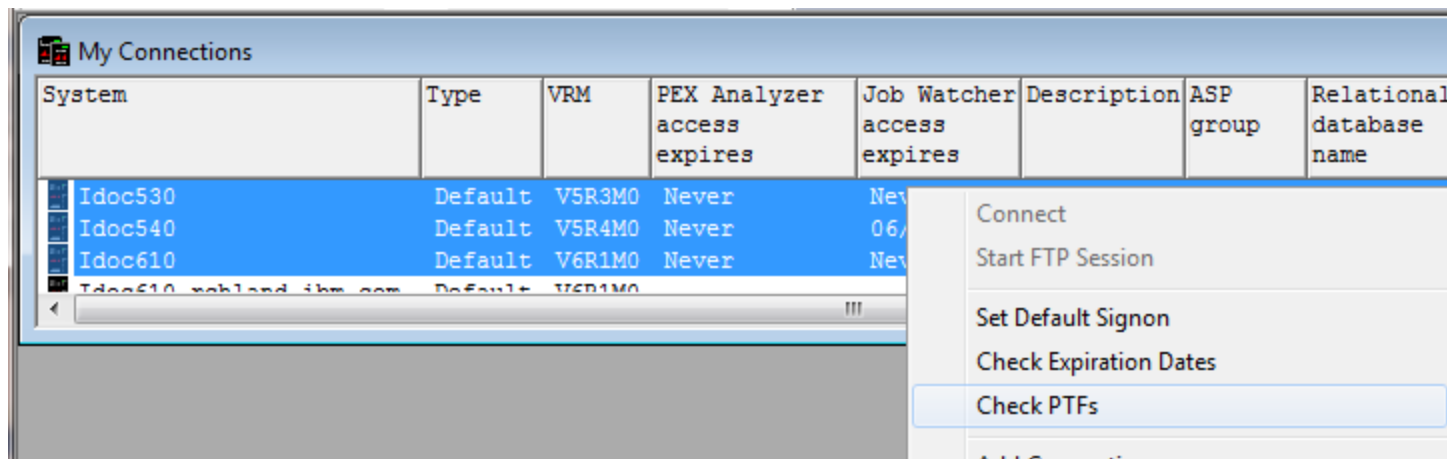
The SQL Editor can now run multiple SQL statements at once as long as the following rules are followed:

1. Each statement ends with a semicolon.
2. If results are desired to be displayed in the table or graph, then the last SQL statement should be a SELECT or WITH statement.

This can be useful if calling a stored procedure or creating aliases before executing a SELECT statement.

## Oct-Feb 2012 – PTF Checking

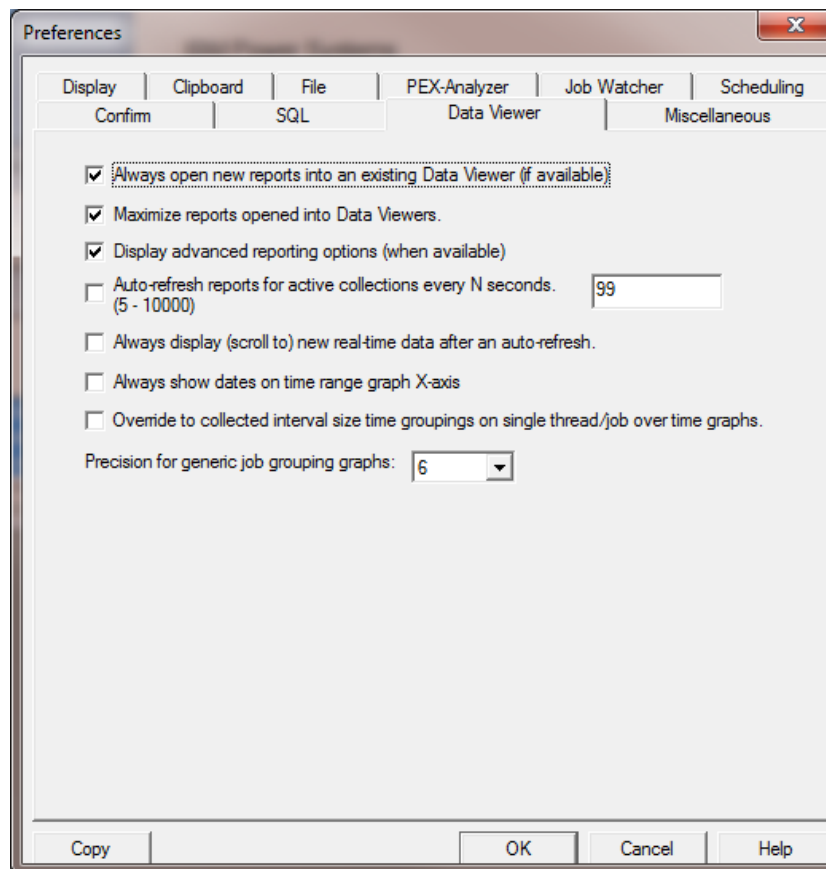
The Check PTFs menu option from the list of connections will now also check for the latest Performance Group PTF levels at 5.4 and higher systems.



## Oct-Feb 2012 – Data Viewer Preferences

Added a set of preferences on a new tab called "Data Viewer".

This contains several preferences, some being moved from the Job Watcher and Miscellaneous tabs.



## Oct-Feb 2012 – New Disk Watcher Definitions

Added additional iDoctor-supplied Disk Watcher definitions that have the OBJINF parameter set to \*ALL.

QFULLO, QFULL1MINO, QTRCO, QTRC1MINO are the new definitions.

The reload IBM-supplied definitions option must be used on systems that already have definitions in order for these to be visible.

Using these definitions will have a more significant performance impact at collection time, but will help with analysis.

## Oct-Feb 2012 – Job Watcher Definition Changes

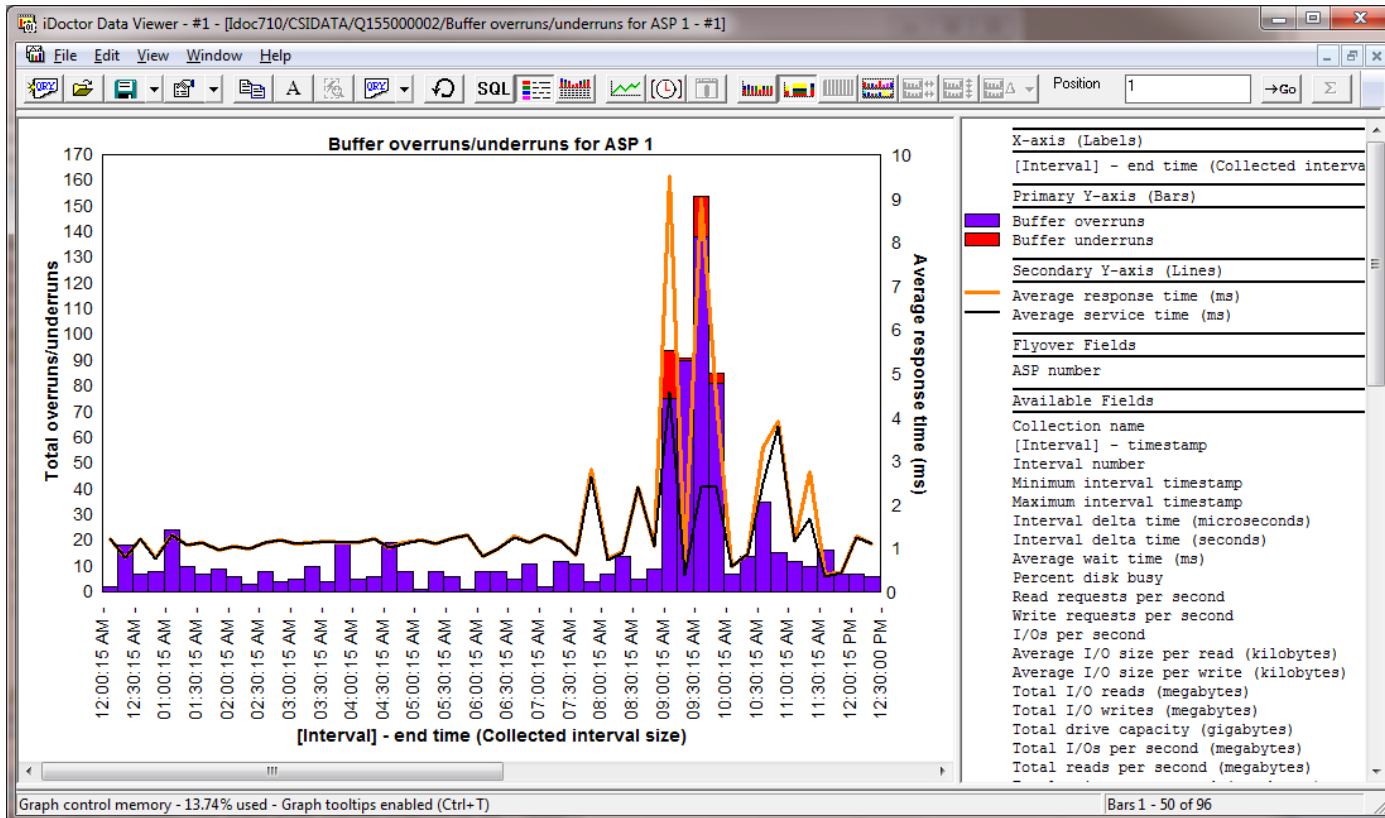
The reload Job Watcher definitions option will now change the definitions that collect SQL so the \*SQLCURSTMT option is used (collects actively running SQL statements and host variables only.) Previously the last running SQL statement was also collected.

Keep in mind that if the reload option has not been done, the original definitions created by the ADDJWDFN command do not use this option. \*SQLSTMT is used instead which captures active or last executed SQL statements.

# Oct-Feb 2012 – New CSI Disk Graph

Added a graph for buffer overruns/underruns within the CSI disk graphs.


**NOTE: These statistics are rumored to be inaccurate in some environments (external storage.)**



# Oct-Feb 2012 – CSI - New seizures and locks graphs

In CSI, added graphs over the new 7.1 lock count fields in file QAPMJOBMI. These are only available if working with 7.1+ file formats (5.4/6.1 \*MGTCOLs can be ported to a 7.1 system to create these fields).

**NOTE: You must run the Collection Summary analysis first to see these new graphs under the Wait graphs -> Seizes and locks folder.**





## Oct-Feb 2012 – Collection Summary hide filtering option

Add an option on the Run Collection Summary (filtering) window to not show the screen again.

If this filtering is disabled when running the Collection Summary analysis, you can turn it back on under the Confirm tab in Preferences.

Run Collection Summary - Idoc710

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

Collection name  
CKOUR1/CKMON002

Comments:

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.

Filters (OPTIONAL):

Job name: contains [ ]

Job user name: [ ]

Job number: [ ]

Job current user profile: [ ]

Subsystem name contains: [ ]

Start time: 2010-02-19-10.23.25

End time: 2010-02-19-10.25.35

Do not show this screen again.

Submit Cancel

Preferences

Display | Clipboard | File | PEX-Analyzer | Job Watcher | Scheduling

Confirm | SQL | Data Viewer | Miscellaneous

Confirm when closing Data Viewers

Confirm when ending iDoctor GUI

Confirm when stopping the Report Generator function

Confirm system name when opening new SQL editor

Confirm ASP selection when opening disk graphs

Confirm when closing tables or graphs that have dependent property sheets



Confirm usage of query definition if the SQL contains comments (which will be lost)

Prompt for time filtering options when running most PEX Analyzer analyses

Prompt for filtering options when running the Collection Summary Analysis in JW and CS

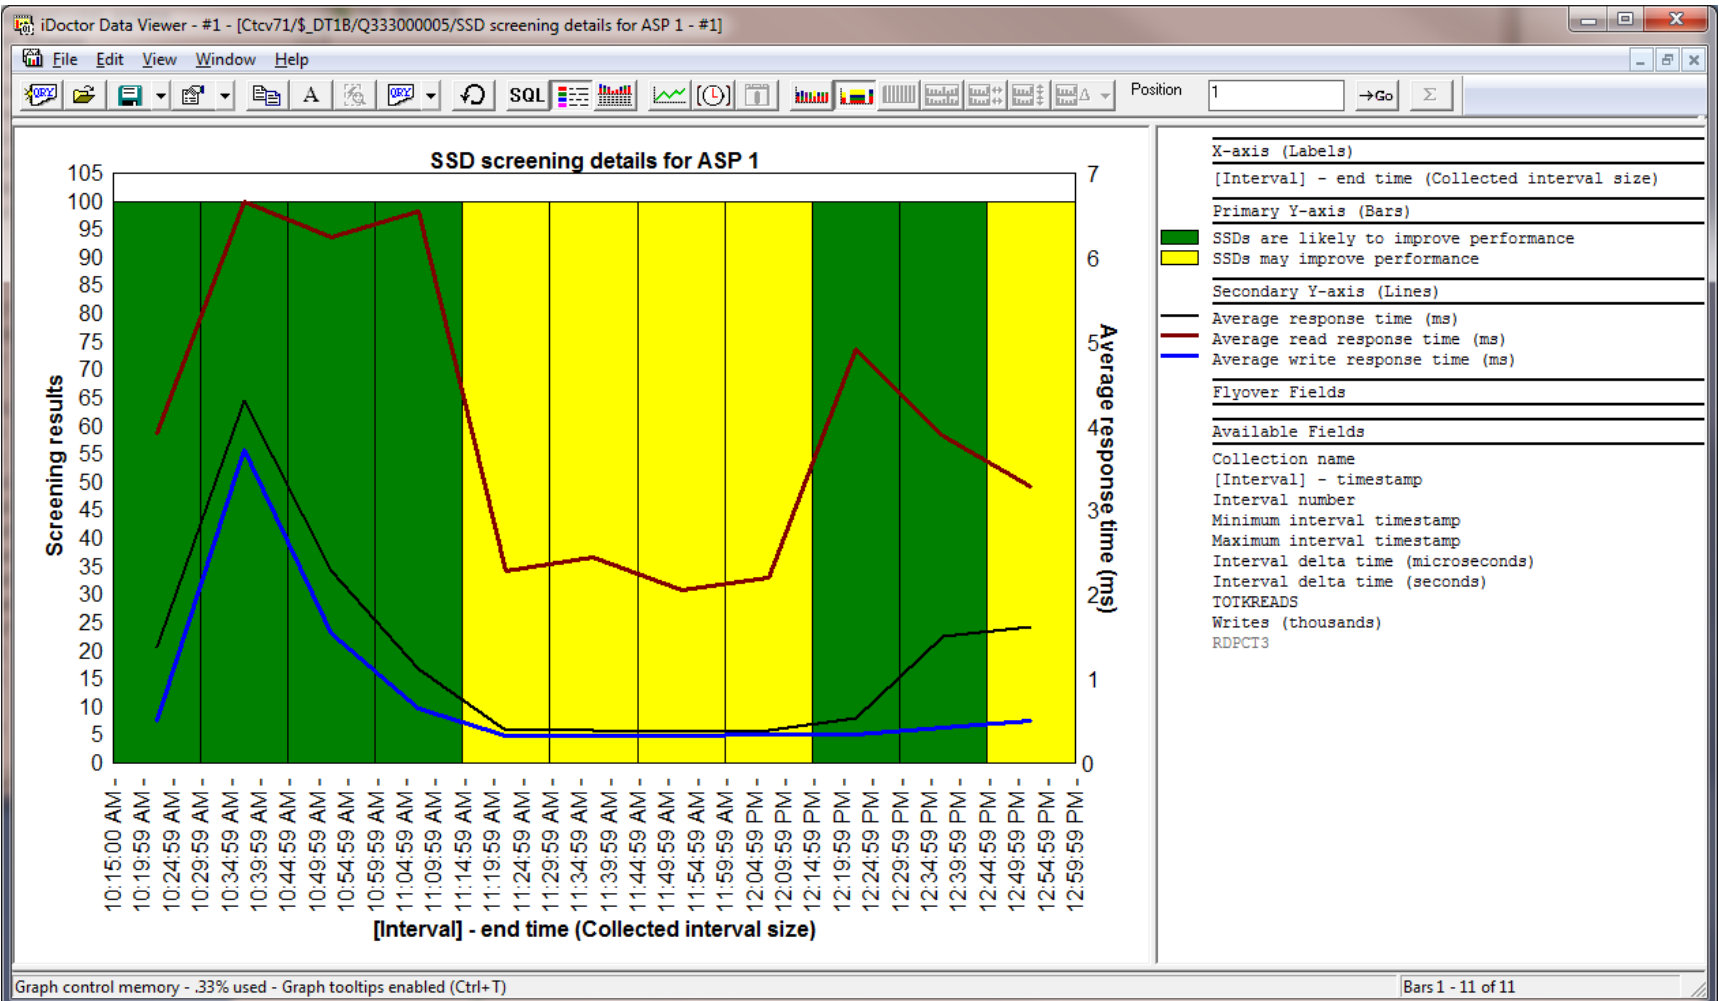
# Oct-Feb 2012 – CSI – SSD Candidate Screening

- Offered at 7.1 only, helps a user determine if SSDs could help performance. It needs data from QAPMDISKRB.



# Oct-Feb 2012 – CSI – SSD Candidate Screening

- Offered at 7.1 only, helps a user determine if SSDs could help performance.



# Oct-Feb 2012 CSI – SSD Candidate Screening job details report

▪ This is similar to the green screen SSD Analyzer tool:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS3780>

iDoctor Data Viewer - #1 - [Ctcv71/\$\_DT1B/Q333000005/SSD screening details by job - #1]

File Edit View Window Help

Position 1 →Go Σ

Job or task name (JOBNAME)	OBJNAME (JBIDE)	Total CPU time (seconds) (CPUTOT)	Disk read wait time (seconds) (READTIME)	Total DASD reads (TOTREADS)	Average read response time (ms) (RDVGRSP)	Disk read time/CPU ratio (READ_CPU_RATIO)	Job current user profile (JBCUSR)	Total contributing threads/tasks (TOTIDES)	Pool number (JBPOOL)
System Tasks	0000000000000000	406,570.24>	864.7938	127,240	6.7966	.0021		864	06
QSPLMAINT/QSYS/568603	00000000000000498	11,946.0060	73.8976	13,933	5.3038	.0062	QSYS	1	02
QWCTJOBS/QSYS/568621	000000000000004AA	546.0770	42.0807	2878	14.6215	.0771	QSYS	1	02
QPYJWCOL/MHENNINGSE/584243	00000000000005389	44,335.0070	37.1213	41,828	.8875	.0008	MHENNINGSE	1	02
QYPSFRCOL/QSYS/568661	000000000000004F1	5079.0510	34.2505	4278	8.0062	.0067	QSYS	28	02
QDBSRV03/QSYS/568593	0000000000000048E	11,662.3460	25.4500	82,719	.3077	.0022	QSYS	1	02
QDBSRV02/QSYS/568592	0000000000000048D	12,377.8530	23.5842	84,189	.2801	.0019	QSYS	1	02
RUN25000/STRSRVR/584385	000000000000005427	483.0480	18.2784	1167	15.6628	.0378	STRSRVR004	1	11
RUN25000/STRSRVR/584411	000000000000005441	464.6690	18.0947	1153	15.6936	.0389	STRSRVR001	1	11
RUN25000/STRSRVR/584425	000000000000005453	481.5900	18.0503	1181	15.2839	.0375	STRSRVR004	1	11
RUN25000/STRSRVR/584426	000000000000005454	484.0610	17.9090	1188	15.0749	.0370	STRSRVR005	1	11
RUN22000/STRSRVR/584985	0000000000000056A7	10,845.0140	17.4471	3402	5.1285	.0016	STRSRVR003	1	09
RUN25000/STRSRVR/584294	0000000000000053C9	14,094.4730	17.3294	4882	3.5497	.0012	STRSRVR003	1	11
RUN25000/STRSRVR/584516	0000000000000054B0	480.4660	17.1908	1221	14.0793	.0358	STRSRVR004	1	11
ASTRYSINV/QSYS/580091	000000000000003E3F	398.7670	17.1094	1871	9.1445	.0429	QSYS	2	02
RUN22000/STRSRVR/584981	0000000000000056A3	9853.1580	17.0874	3409	5.0124	.0017	STRSRVR004	1	09
RUN25000/STRSRVR/584610	000000000000005516	13,952.1460	16.9672	3894	4.3573	.0012	STRSRVR002	1	11
RUN25000/STRSRVR/584439	000000000000005461	508.4230	16.9495	1170	14.4868	.0333	STRSRVR005	1	11
RUN21000/STRSRVR/585132	00000000000000574C	14,511.2210	16.8581	2045	8.2436	.0012	STRSRVR004	1	10
RUN25000/STRSRVR/584507	0000000000000054A7	493.7710	16.7998	1166	14.4080	.0340	STRSRVR002	1	11
RUN25000/STRSRVR/584440	000000000000005462	469.0420	16.7316	1153	14.5114	.0357	STRSRVR005	1	11
QYPSJSVR/QYPSJSVR/568743	000000000000005BB	8824.9920	16.4559	5552	2.9640	.0019		84	02
RUN25000/STRSRVR/584512	0000000000000054AD	492.6360	16.4214	1165	14.0956	.0333	STRSRVR005	1	11
RUN21000/STRSRVR/585131	00000000000000574B	14,774.1270	16.2628	1945	8.3614	.0011	STRSRVR005	1	10
RUN21000/STRSRVR/580520	000000000000003FF1	14,849.3340	16.2335	3440	4.7190	.0011	STRSRVR002	1	10
RUN25000/STRSRVR/584438	000000000000005460	479.4850	16.1812	1152	14.0462	.0337	STRSRVR003	1	11
RUN21000/STRSRVR/585175	000000000000005778	14,414.7210	16.1243	2091	7.7113	.0011	STRSRVR004	1	10
RUN25000/STRSRVR/584511	0000000000000054AB	498.5610	16.0954	1162	13.8515	.0323	STRSRVR003	1	11
RUN25000/STRSRVR/584443	000000000000005465	417.9220	16.0600	1171	13.7147	.0384	STRSRVR003	1	11
RUN25000/STRSRVR/584475	000000000000005485	477.3270	16.0374	1181	13.5795	.0336	STRSRVR004	1	11
RUN21000/STRSRVR/585106	000000000000005731	14,760.3140	15.9373	2176	7.3241	.0011	STRSRVR002	1	10

Rows 1 - 31 of 3832

## Oct-Feb 2012 – Scheduling Window Changes

1. The current date is circled on the calendar control.
2. The system's name and current time (updating every second) is listed below the calendar.
3. If the scheduled time is before the current system's time an error will be given.
4. The default scheduled time preference is 1 day ahead now instead of 7 days.

Use this interface to schedule an action for a later time.

Schedule the collection start time

Note: Date and time values are based on the server's clock, not your PC's clock.

Frequency:

Scheduled date:

February, 2012						
Sun	Mon	Tue	Wed	Thu	Fri	Sat
29	30	31	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	1	2	3
4	5	6	7	8	9	10

**Today: 2/2/2012**

Current (system) time: 06:30:55 AM Idoc610

Scheduled collection start time:

# Oct-Feb 2012 – Transfer Collection Changes

Made several changes to the Transfer Collection function in response to problems with SAVPFRCOL/RSTPFRCOL and IASPs.

Transfer Collection(s) - Idoc610

Transfer options:

Action: Copy to temp lib, save, send and restore to IBM i library

Target system: Idoc710 - V7R1

Target library: mccargar11 ASP number or device: 1 Select ASP

Clear existing data

Disable extended passive mode (SENDEPSV 0)

Temporary library options (Recreated, then collections copied into it):

Library name: QIDOCTEMP ASP number or device: 1 Select ASP

Clear existing data  Delete when complete

Save options:

Target release: \*CURRENT Data compression: High

Collections to transfer:

Collection	Status	Ending reason	Using iDoctor collection summary	Collection size (MB)	Partition collected on VRM	Peccor
B	Ready for analysis	ASP limit	Yes	.32	V6R1M0	I

Transfer Cancel

## Oct-Feb 2012 – Transfer Collection Changes

The desired collection(s) to transfer are now copied to a temp library.

SAVLIB/RSTLIB commands are now used to save the data instead of SAVPFRCOL/RSTPFRCOL.

Added options to check disk space and select the desired ASP or device to save the data to on both the target and local systems (for the temp library.)

Added validation checks before starting the transfer.

**NOTE: PLEASE BE CAREFUL!** The data in the temp and target libs is cleared by default. See the checkboxes on the previous screen to control that behavior.

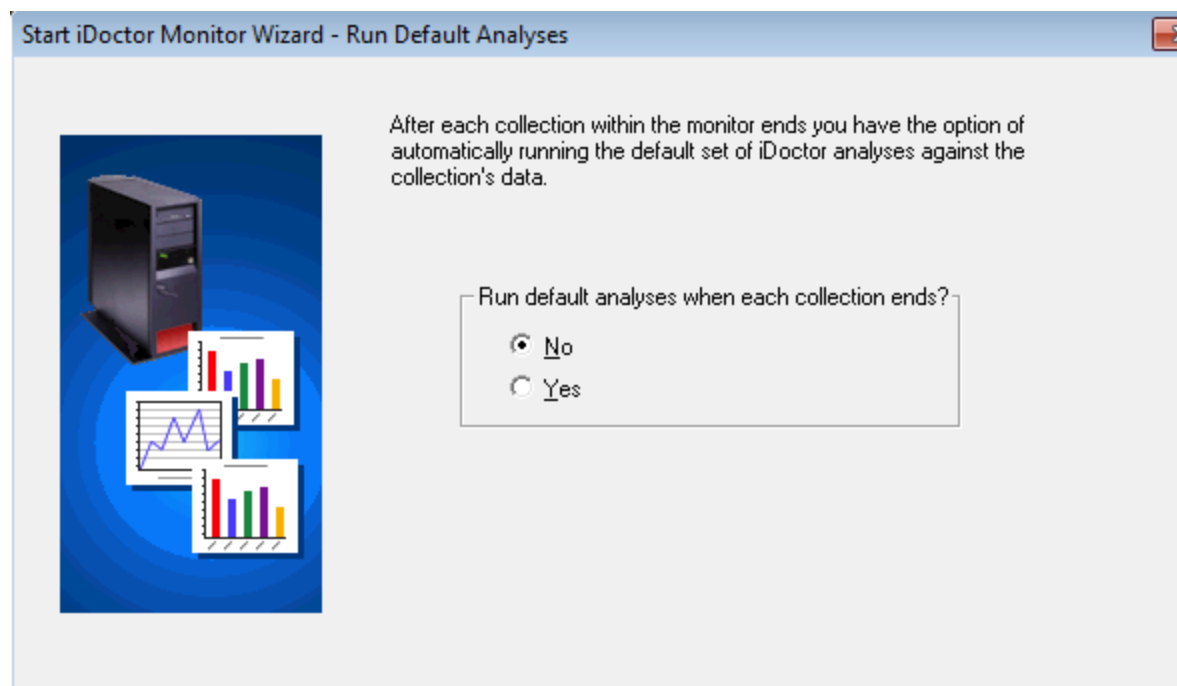
## Oct-Feb 2012 – Flattened graphs changes

For graphs that utilize the new flattening technique (CSI memory pools, CSI advanced loop graphs), add the capability to customize the graph labels, hide values on the Y1/Y2 axis, change the scaling, or use the graph definition interface to change fields used, scaling, colors and sizes.



## Oct-Feb 2012 – Run Default Analyses option

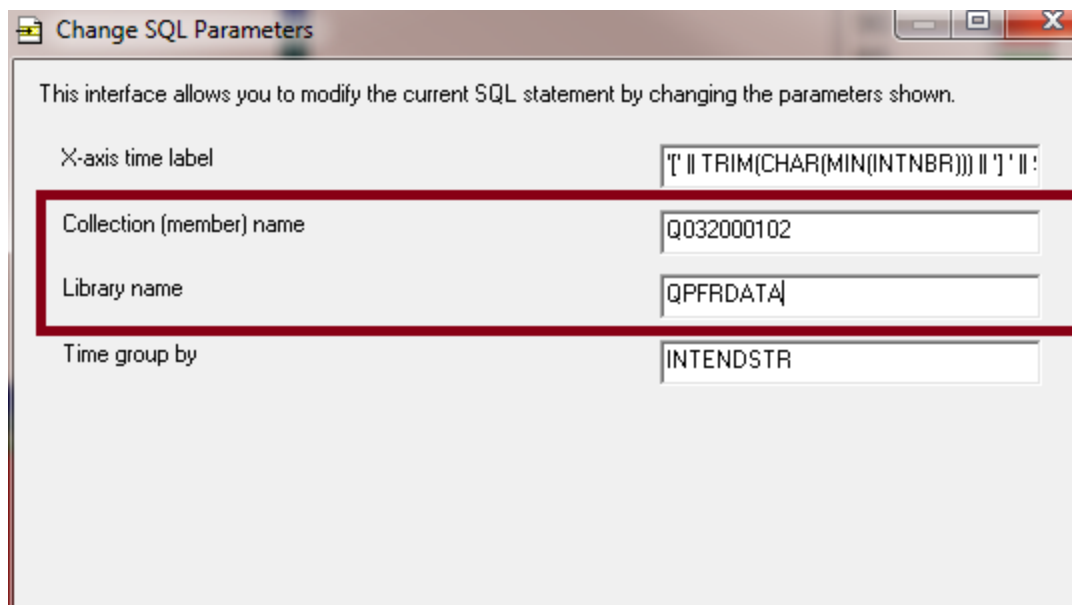
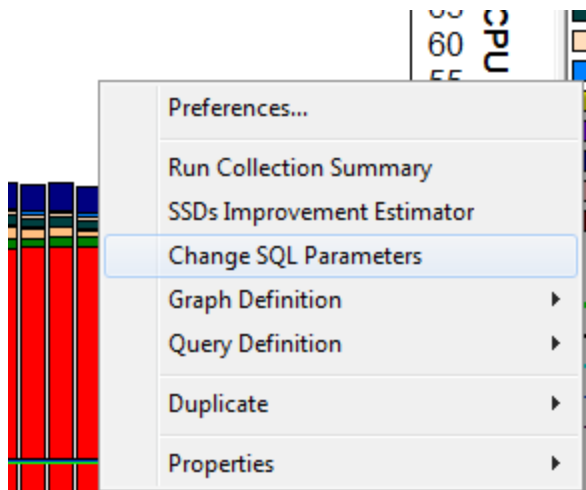
In Job Watcher at 6.1+, all the iDoctor commands used to start monitors, end monitors or end collections now include a new parameter called "Run default analyses (\*YES/\*NO)" with a default value of \*NO. The same option has been added to the Start Monitor Wizard (only appears if builds have been updated.)



## Oct-Feb 2012 – Changing the current collection


A user can now use the Change SQL Parameters menu from a graph to change the library and collection currently being used to display the graph.

**NOTE: If the collection is summarized, the new collection picked must be summarized too.**



# Oct-Feb 2012 – CPU graph updates

In Job Watcher and CSI, updated the CPU graphs so they contain a new field in the flyover “Current number of active processors.”

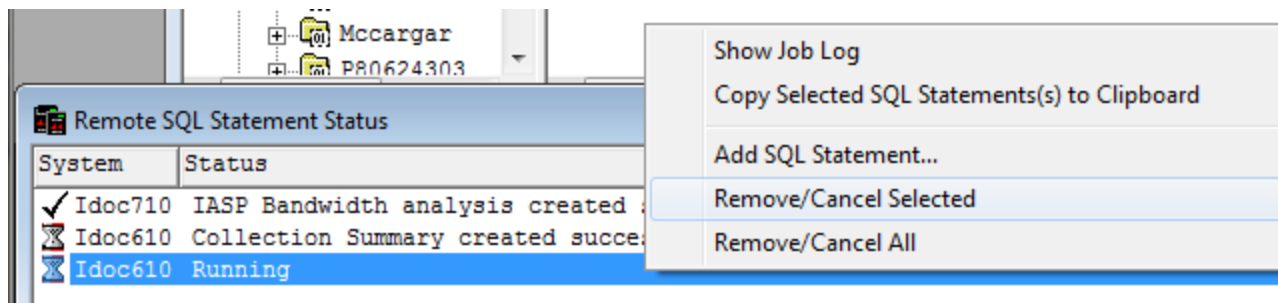


## Oct-Feb 2012 – Remote SQL Statement Status View Updates

Users can now cancel/remove entries while SQL Statements are running.

Users can use the “Copy Selected SQL Statements to Clipboard” option.

Use the new “Add SQL Statement” to define additional statement(s) to run in the view. You can also specify the system to run the statements on.

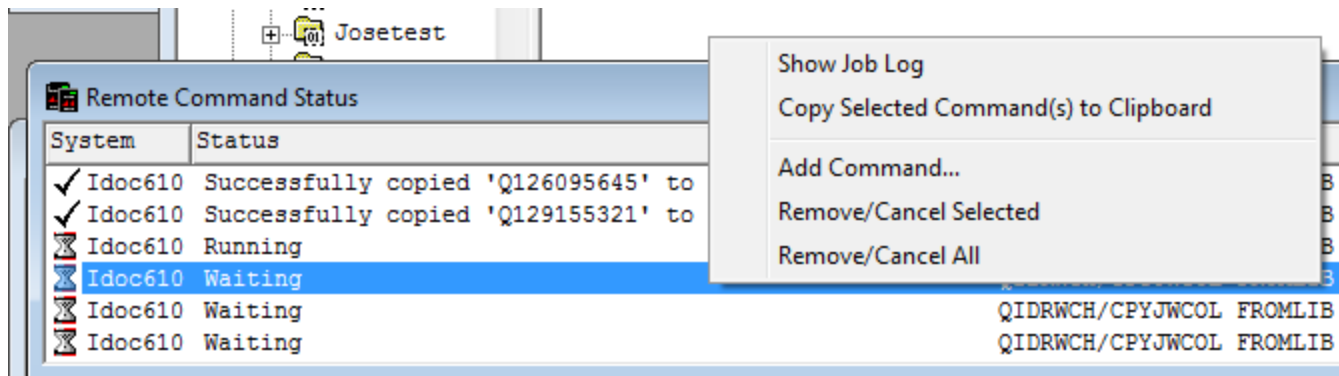


## Oct-Feb 2012 – Remote Command Status View Updates

Users can cancel/remove entries while commands are running.

Users can use the “Copy Selected Commands to Clipboard” option.

Use the new “Add Command” option to define additional command(s) to run in the view. You can also specify the system to run the commands on.

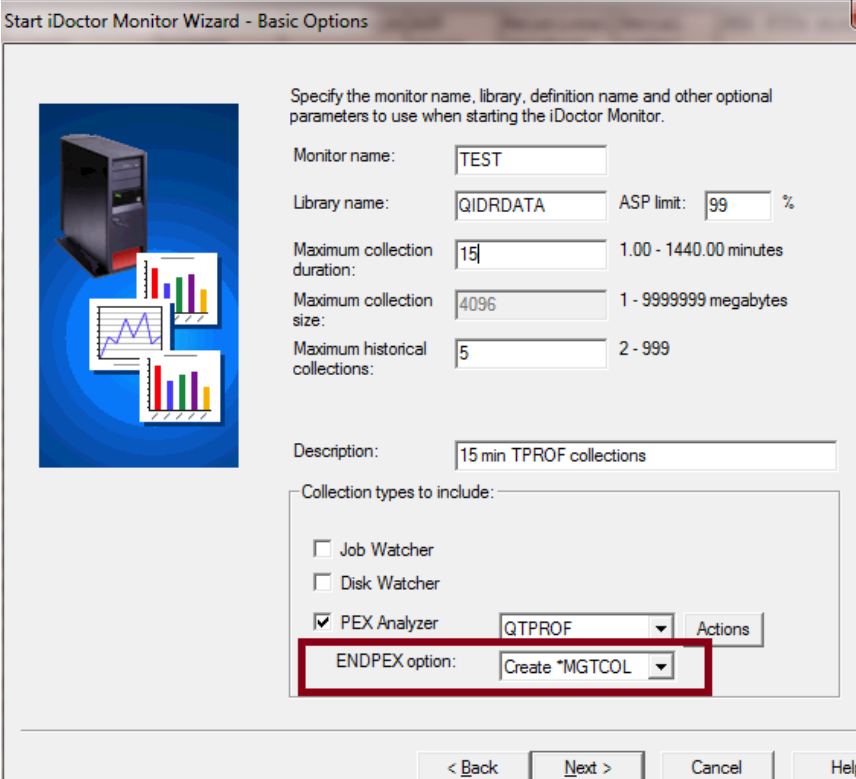


## Oct-Feb 2012 – Start PEX Monitor using \*MGTCOLs

The Start iDoctor Monitor Wizard now supports the creation of PEX monitors into \*MGTCOL objects instead of DB files.

On the basic options screen of the wizard, added an ENDPEX option parameter with 3 possible values:

- 1) Create DB files
- 2) Create \*MGTCOL
- 3) Suspend



Start iDoctor Monitor Wizard - Basic Options

Specify the monitor name, library, definition name and other optional parameters to use when starting the iDoctor Monitor.

Monitor name: TEST

Library name: QIDRDATA ASP limit: 99 %

Maximum collection duration: 15 1.00 - 1440.00 minutes

Maximum collection size: 4096 1 - 9999999 megabytes

Maximum historical collections: 5 2 - 999

Description: 15 min TPROF collections

Collection types to include:

Job Watcher

Disk Watcher

PEX Analyzer QTPROF Actions

ENDPEX option: Create \*MGTCOL

< Back Next > Cancel Help

## Oct-Feb 2012 – Start PEX Collection CPU Sample change

Updated the STRPACOL command at 5.4 to support a CPU sample value of 0. If 0 is used the PMCO event will NOT be included in the collection.

**Note:** The GUI also allows this value to be used if the server builds have been updated.

PEX Collection Wizard - Options - Idoc610

Collection Options:

Definition type:  PEX-Analyzer-supplied  User-defined

Definition: \*DB\_LDIO [Details...]

Start in standby (suspended) mode

Collection name: [ ]

Library: QPADATA

Description: [ ]

Duration: 5 1 - 1440 minutes

Maximum data to collect: 500000 1024 - 250000000 KB

CPU interval sample: 0 0.1 - 200.0 ms or 0 = no PMCO events collected

Advanced options: [Configure...]

Scheduled start time: [Configure...] Immediate

< Back Next > Cancel Help

## Oct-Feb 2012 – New Grouping Options in PEX Analyses

Added new grouping options to the following PEX Analyses:

1. CPU Profile
2. LDIO
3. Data queue
4. Data area
5. DB opens
6. Net size
7. Heap storage

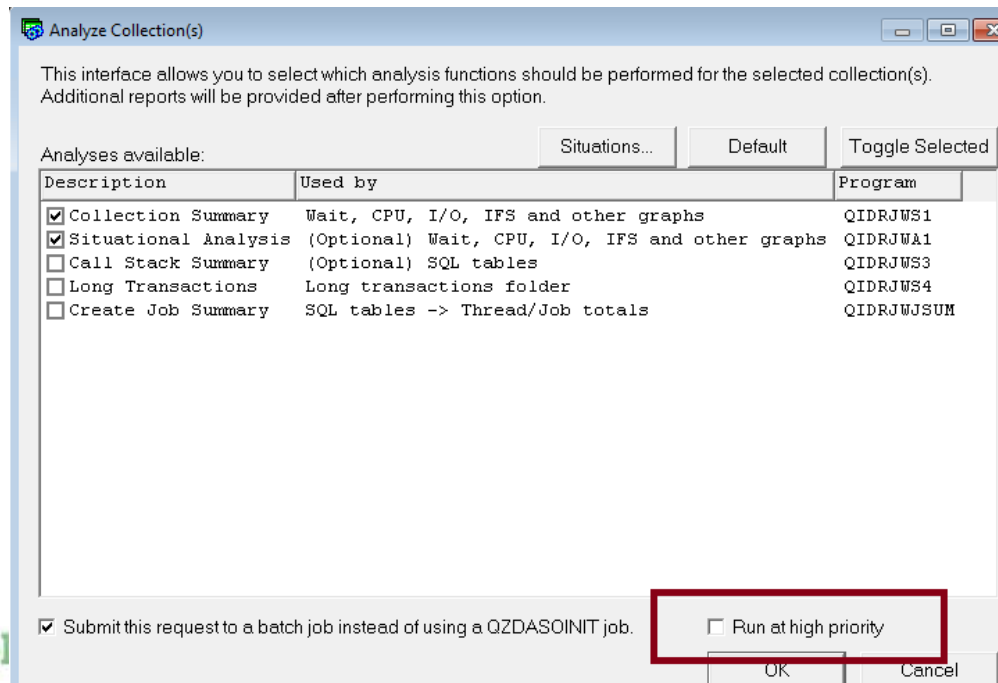
[See the update history for more information!](#)



## Oct-Feb 2012 – Analyze Collections Window – run priority option

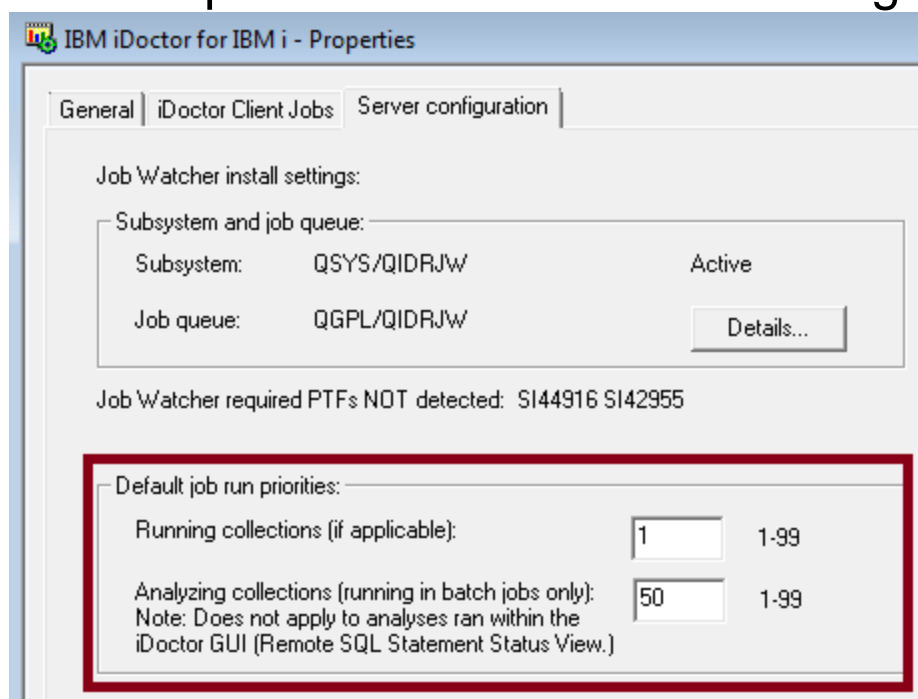
Added an option on the Analyze Collection window if submitting to batch to specify the priority level (high/low option). Running at high priority will use JOBD/CLS QIDRBCH which defaults to priority 1. Low priority will use JOBD/CLS QIDRLOWBCH which defaults to priority 50.

See the next slide for more information on changing these default values.



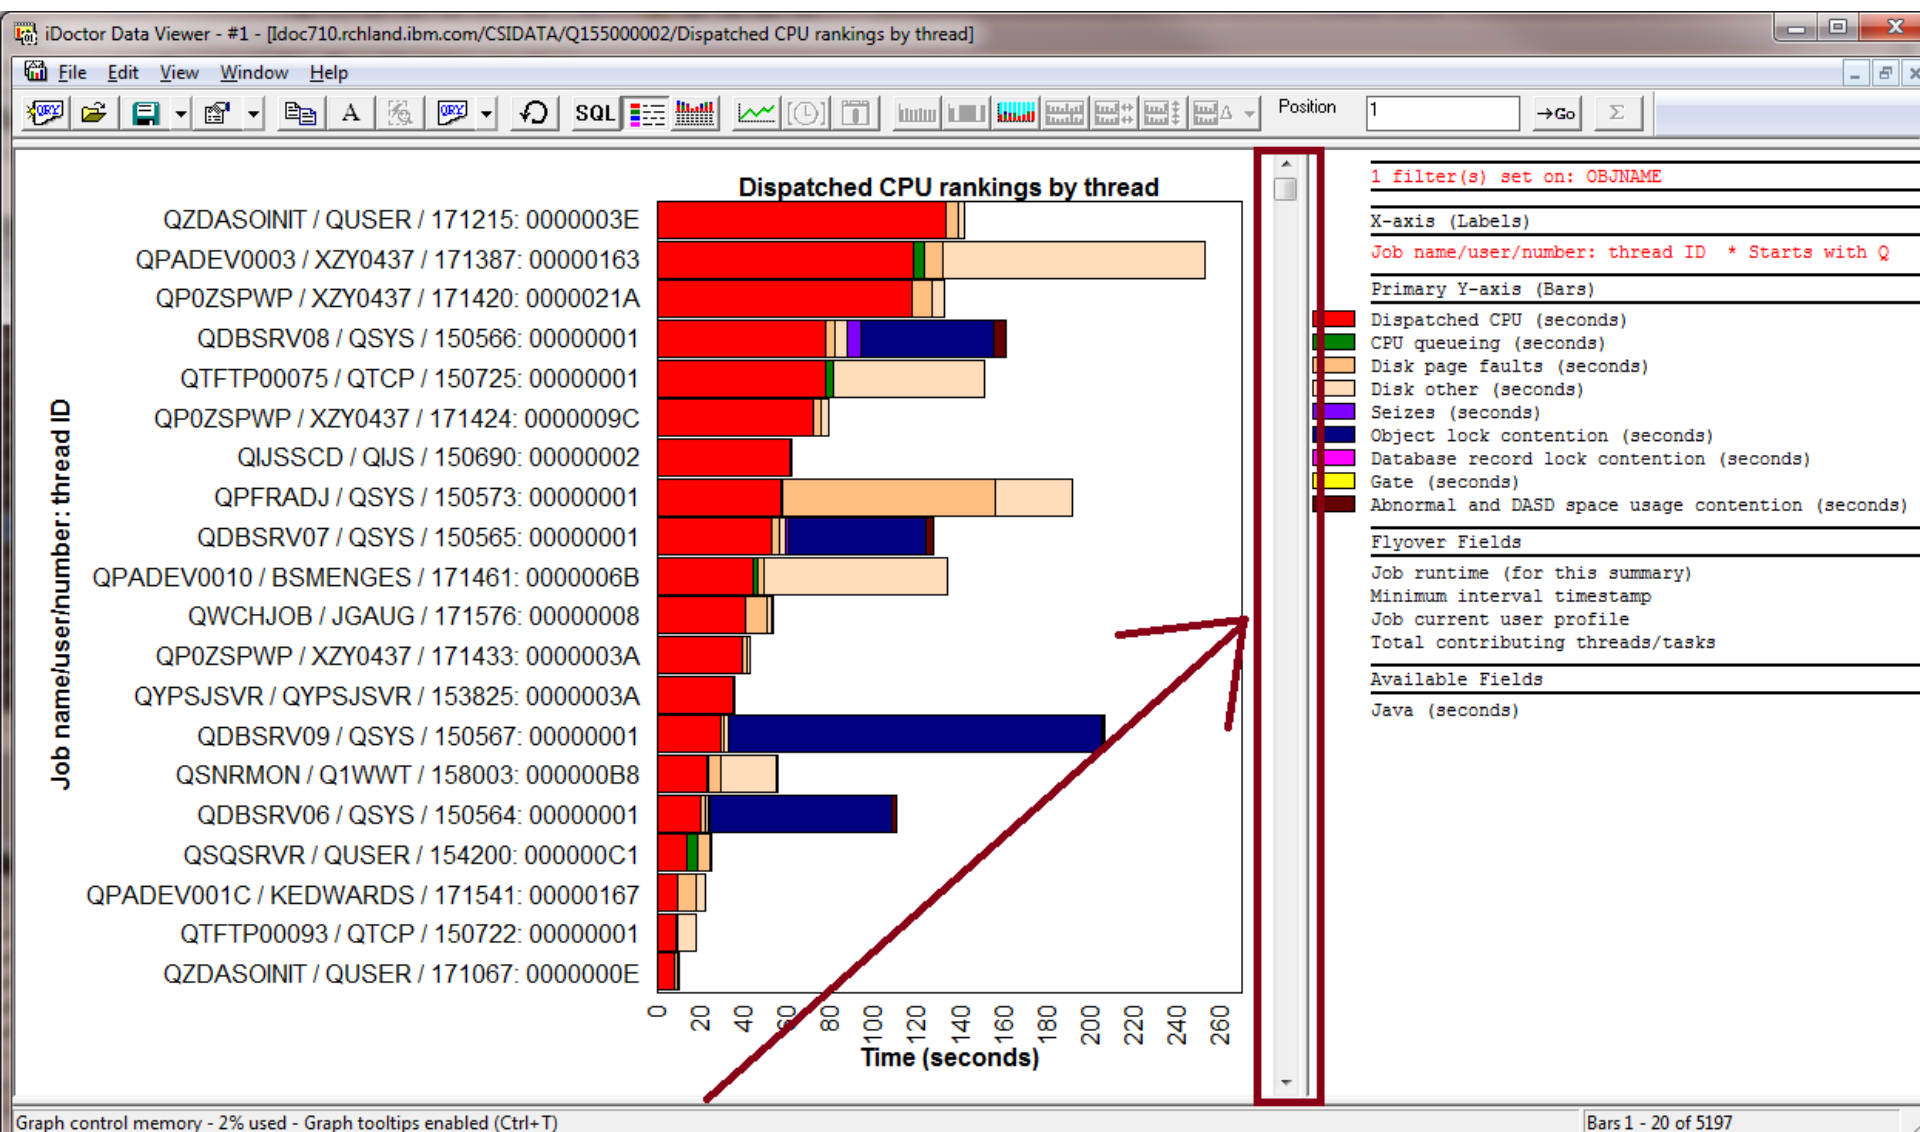
## Oct-Feb 2012 – Modifying default run priorities

In the server configuration tab of the Application Properties (right-click Job Watcher or PEX Analyzer icon and use Properties menu) you may configure the desired run priorities to use when running collections and analyses.



**NOTE: Current builds (892 and higher) will incorrectly run the analyses submitted to batch at priority 1 by default but this will change with the Feb 2012 builds.**

# Oct-Feb 2012 – Graph scrollbar moved on rankings charts



## Oct-Feb 2012 – Graph selection changes

Change selections in graphs so every left click will clear existing selections unless ctrl (add to selection) or shift (defines a range) is pressed.

Previously any existing selections made were not cleared on another left-click.

## Oct-Feb 2012 – Installation changes

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

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

## Oct-Feb 2012 – PEX miscellaneous updates

The Call Stacks analysis will now process the call stacks created by format 4 events (\*PMCO format 4).

Update the STRPACOL command to ensure that the job it submits ends abnormally if ENDPEX fails. Previously it was possible for the STRPACOL command to issue the ENDPEX command and if errors occurred the job would still end with a normal reason code.

**The latest server builds are required.**

In the Stats summary for all jobs analysis, in the report "Summarized CPU and I/O by thread/pgm/MI instruction " the JOBCPUPCT and ELPPCT fields were **incorrect**.

In PDIO, if working with data from multiple ASPs, you will now be prompted for the desired ASP (like happens with the applicable graphs in DW/CSI.)

## Oct-Feb 2012 – Job Watcher new situation

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

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

## Oct-Feb 2012 – Job Watcher single interval rankings update

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

W = has a wait object

H = holder

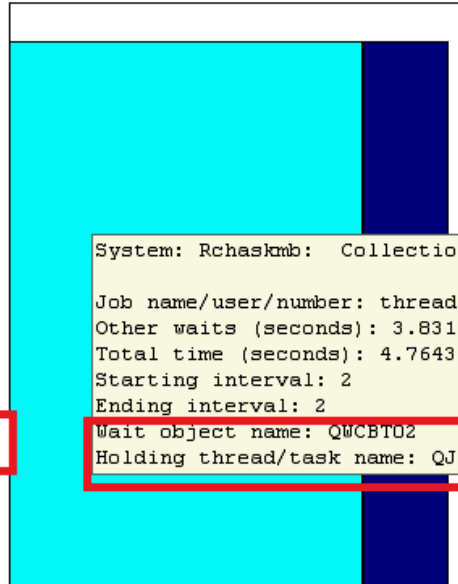
B = current wait bucket is the same as current sort /filter bucket

S = has an SQL client job (applies to 6.1 and higher only)



# Oct-Feb 2012 – Job Watcher single interval rankings update (example)

Thread signatures ranked by Object locks: From 08:30:15 am to 08:30:20 am



X-axis (Labels)

Job name/user/number: thread ID

Flags (W=wait object, H=holder, B=in current bucket, S=SQL client)

Primary Y-axis (Bars)

- Dispatched CPU (seconds)
- CPU queueing (seconds)
- Other waits (seconds)
- Disk page faults (seconds)
- Object lock contention (seconds)

System: Rchaskmb: Collection: Q047083009

Job name/user/number: thread ID: Q1WWTCMN / QSYS / 060712: 00000004 WHB

Other waits (seconds): 3.8312 (80.41% of total)

Total time (seconds): 4.764399

Starting interval: 2

Ending interval: 2

Wait object name: QWCBT02

Holding thread/task name: QJOBLOGSVRQSYS 048032

Q1WWTCMN / QSYS / 060712: 00000004 WHB

OFFSET

Total intervals included

Initial thread task count

Process initial thread task count

## Oct-Feb 2012 – Job Watcher misc updates

In Job Watcher, made changes to work around a possible issue in the collection summary analysis. Data in the qapyjwtd file, field TDEUSECS could be invalid (very large negative values) causing the analysis to fail. Changes made to the analysis will handle this possibility now.

**At 7.1, PTF SI44515 corrects the problem** with the data being invalid.

In Job Watcher, removed the "Update Wait Bucket Actives + Idles" Analysis because it is now obsolete.

When listing collections in a library or monitor, added checks that ensure that all required file/members are present. If some are missing then the graphs and analysis options will not be available for that collection and the status will show "Missing required data" and then list the missing file/members.

## Oct-Feb 2012 – CSI new memory pool graph

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

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

## Oct-Feb 2012 – CSI bug fixes

In CSI at 6.1+, fixed a bug with the External Storage links and ranks analysis where some ports may be missing from the data generated.  
**The latest server builds must be installed.**

In CSI removed use of QAPMJSUM in several graphs (I/O, IFS, other graph folders) to improve accuracy.

In CSI, when graphing multiple collections, the drill down ranking graphs from the overview graph were not including data from all collections originally selected.

In CSI at 5.4 only, fixed accuracy problems with the wait bucket graphs. Also updated Create Job Summary analysis for the same issue.

When listing collections in CSI, the DB Files VRM field was sometimes inaccurate.

## Oct-Feb 2012 – Miscellaneous bug fixes

In graph flyovers; the values shown were sometimes incorrect.

Fixed a couple of situations where the report shown in iDoctor could be from the wrong collection member. (Call stack Search in JW for example).

Fixed a problem applying access codes if the user applying it does not have \*ALLOBJ authority. **(needs latest server builds to fix)**