# iDoctor Updates (April – August 2018)

Ron McCargar
iDoctor development
IBM i Global Support Center

# Agenda

- Overview
- Cover the latest GUI updates and enhancements for this period
- Questions?

# Overview

- iDoctor is a suite of **dynamic** performance tools offered by the Global Support Center.
- We cover all areas of performance but historically focused more on low-level details.
- Started in V4R5 with the PEX GUI plug-in for Operations Navigator it now consists of 8 external and 2 IBM internal components.

# Dynamic on demand

- iDoctor was originally created to help IBM support solve performance problems.
- Working on unique customer problems requires a high-level of flexibility and ability to add or change functions quickly and easily.
- Feedback from iDoctor users continues to help shape the enhancements added on an ongoing basis.

# Support

- Supported currently at IBM i V6R1 or higher.
  - iDoctor support for V6R1 and 7.1 IBM i will end in early 2019.

# August 2018 (1325) – Fixed several broken graphs in CSI

- Under Job counts graphs -> net jobs breakdown rankings
  - Net jobs breakdown by subystem

- Under IFS graphs -> IFS read total rankings
  - IFS read totals by thread, current user

- Under CPU Graphs folder:
  - CPU utilization breakdown by core vs total CPU utilization
  - CPU utilization breakdown by core

# July 2018 (1323) – Primary keys removed from JW Job Summary

- A user was experiencing check constraint errors on the tables built by this analysis, so removed these keys for now in order to resolve and investigate further.

- CPF502F occurred on the SQL insert of data on the 1st table.

# July 2018 (1322) – VIOS columns flipped in directories

- When connected to a VIOS using Power Connections and looking at directories, the File type and Modified date column data were flipped around.

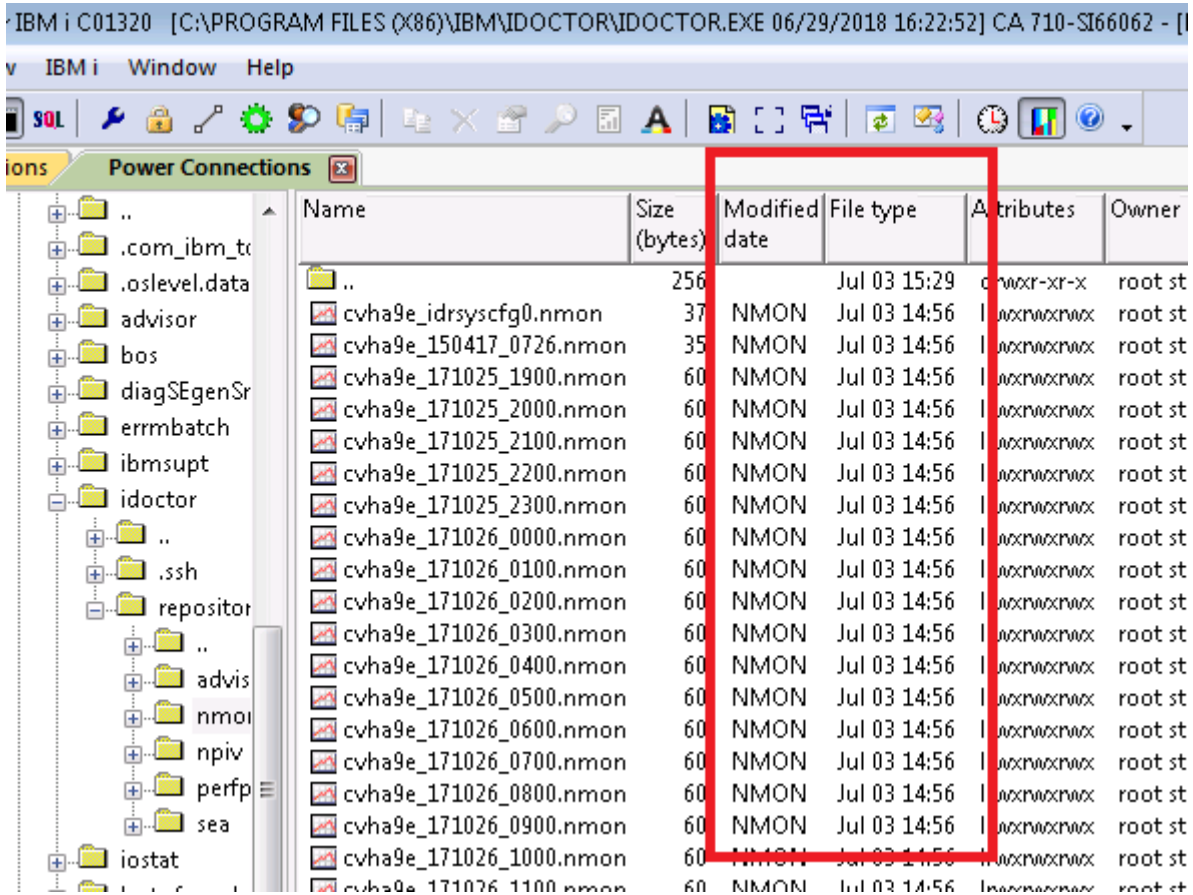

# July 2018 (1322) – VIOS symbolic links were incorrect

- When listing files on a VIOS that are symbolic links, the symbolic link path was incorrect. It was missing the 1st and last characters. This caused the inability to open files (using the Open local file option.)

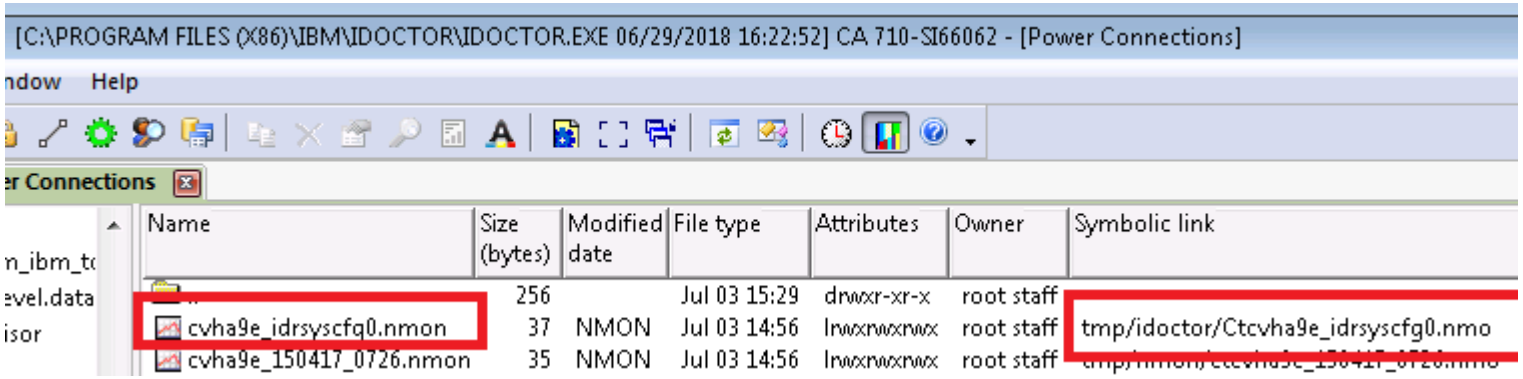

# July 2018 (1322) – Power – Display details preference changed

- In Preferences -> Power, renamed the preference
  "Display details about collections in the Import folder while connected to the IFS (very slow)" to
  "Display details about collections in the Import folder (possibly very slow!)"

  This option applies to nmon, npiv, sea import folders when connected to a VIOS or when connected to an IBM i.  If there are thousands of files to list then this option can mean the difference between seeing results in a few seconds vs a few minutes!

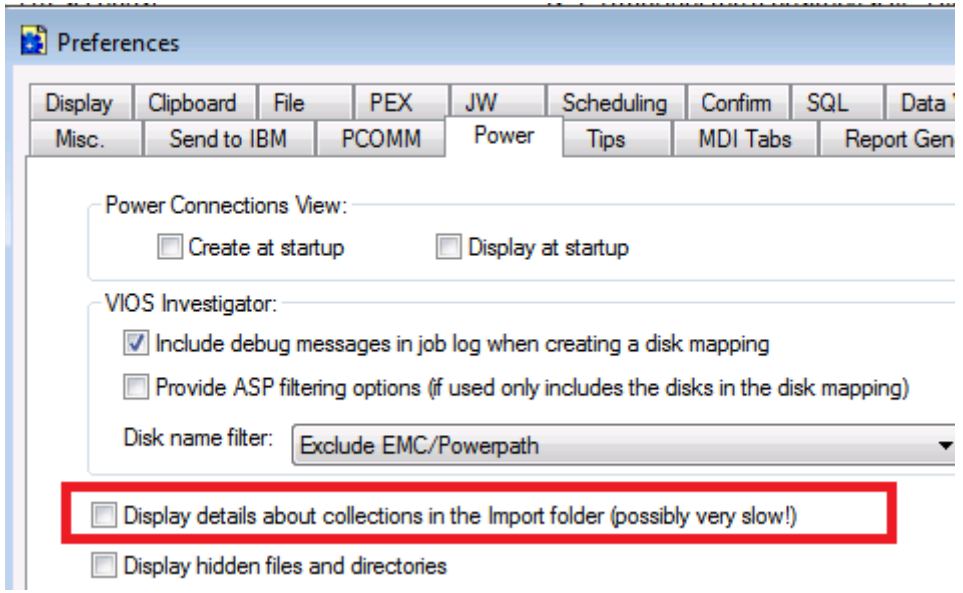

# July 2018 (1321) – CSI workload capping delay rankings SQL error

- In CSI fixed an SQL error on TL.JBPOOL column in the graph "Workload capping delay rankings for <XYZ> by thread“

- Work around is replace TL.JBPOOL with JBPOOL in the SQL statement.

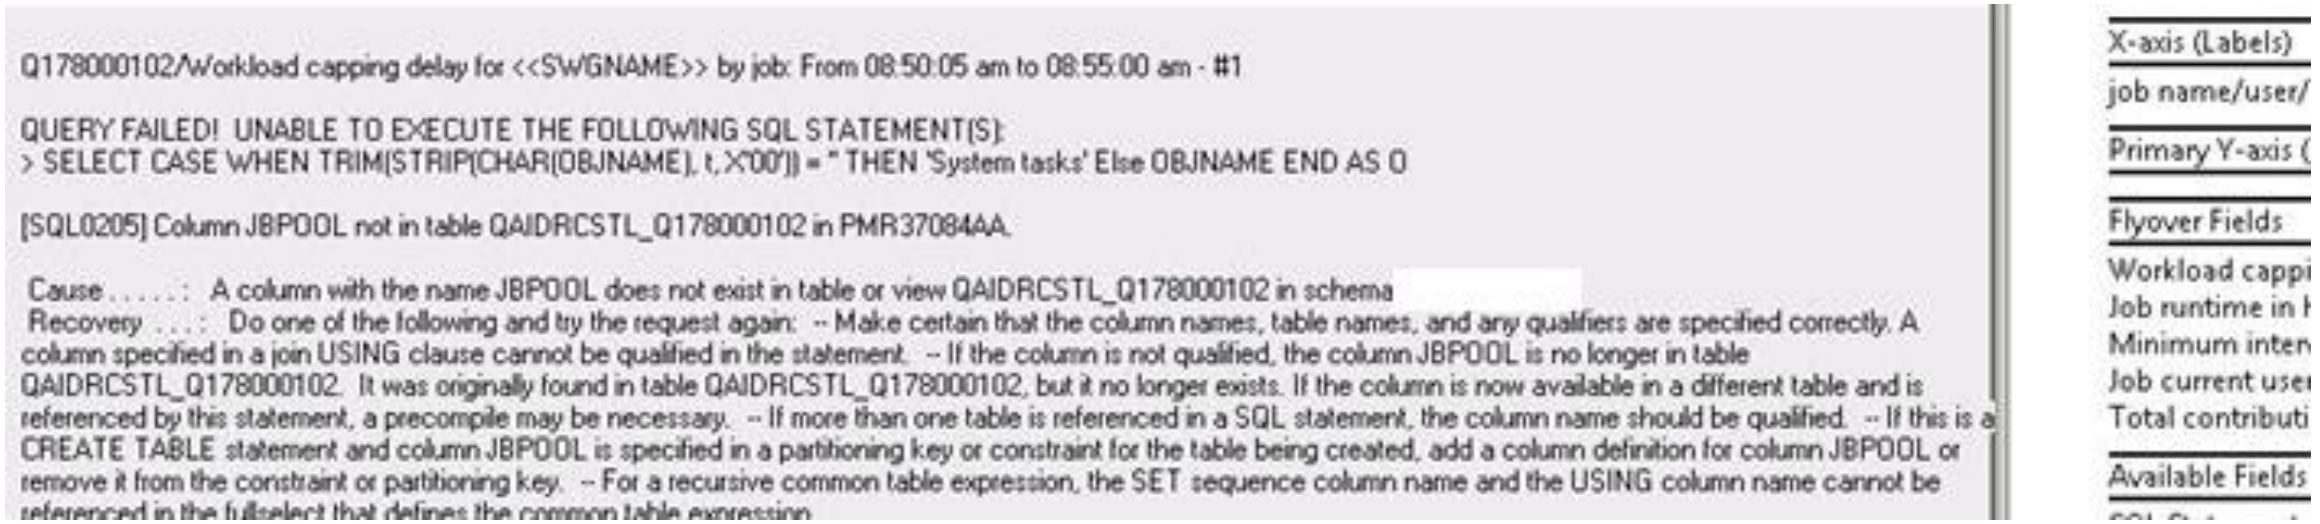

# June 2018 (1320) – CSI workload capping delay rankings prompt

- If drilling down in CSI from the Collection Overview to "Workload capping delay rankings for <XYZ> by thread“ the user will be prompted for the value of <<SWGNAME>>.

- Added the following text on the Change SQL Parameters window (that this is the):
  - “Workload capping group name from file QAPMSYSWLC”

# June 2018 (1320) – Configure Collection Services cycle time

- Fixed some bugs when using the Configure Collection Services window.  The cycle time could not be edited without an error message "The parameter is incorrect" being shown.

  Also the cycle time when loaded from the saved value on the system into the interface was not being shown correctly.

# June 2018 (1320) – Select analysis system window text change

- Changed the message text when selecting an analysis DB (IBM i) on the System Selection window that appears on a PC when first using Power Connections (since “Power Doctor” is really “Power Connections” now)

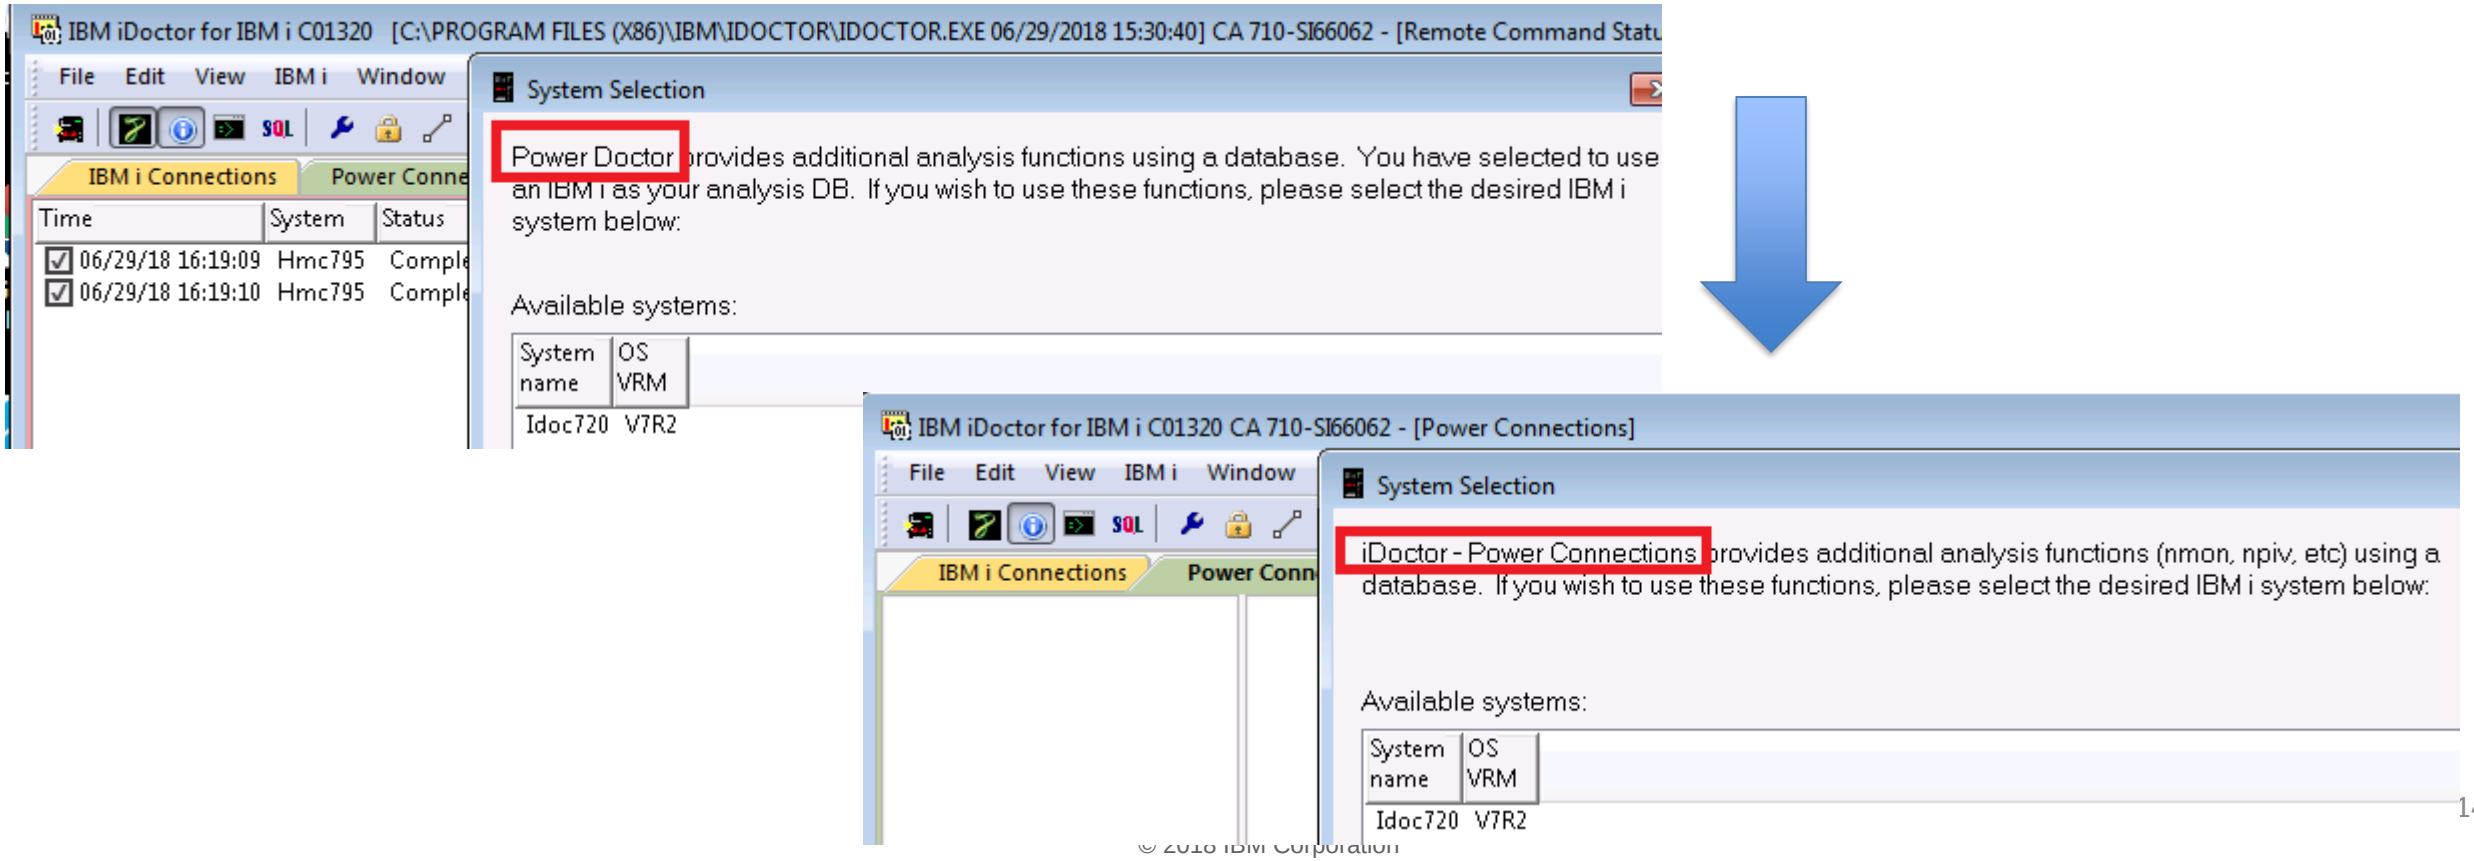

# June 2018 (1320) – Connection fails when using SSH due to no Java

- If connecting to a non-IBM i system (HMC/VIOS/etc) using SSH but Java is not installed or could not be found in the user's path, improved the error message shown to make the resolution more clear.
- Also the user's password is no longer returned in the error message window.

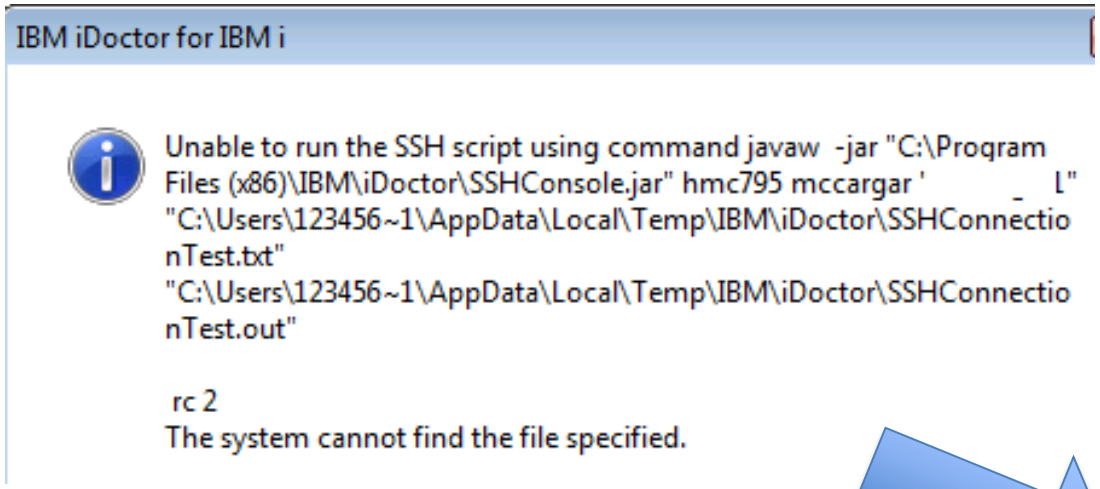


# June 2018 (1320) – PEX Col Wizard – events option hidden if scheduling

- In the PEX Collection Wizard in the Basic Options screen the “Maximum events to collect” option is now hidden if a scheduled collection start time is set.
- **Note:** This option does not apply to scheduled PEX collections.

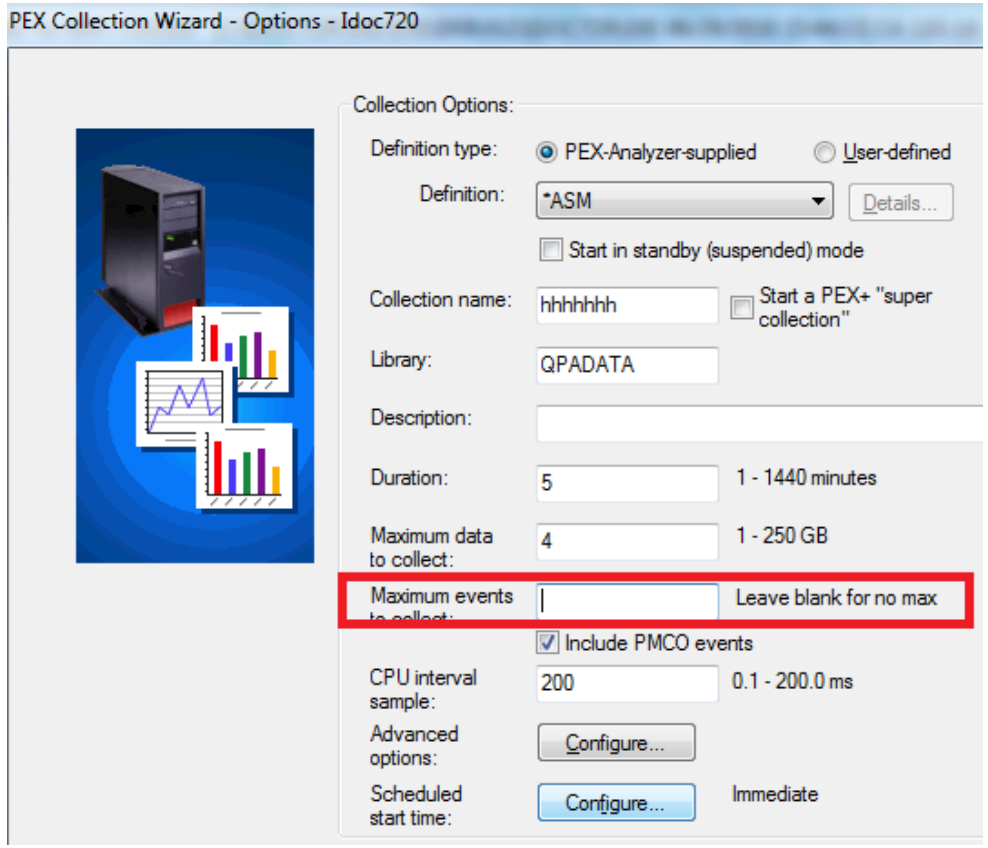


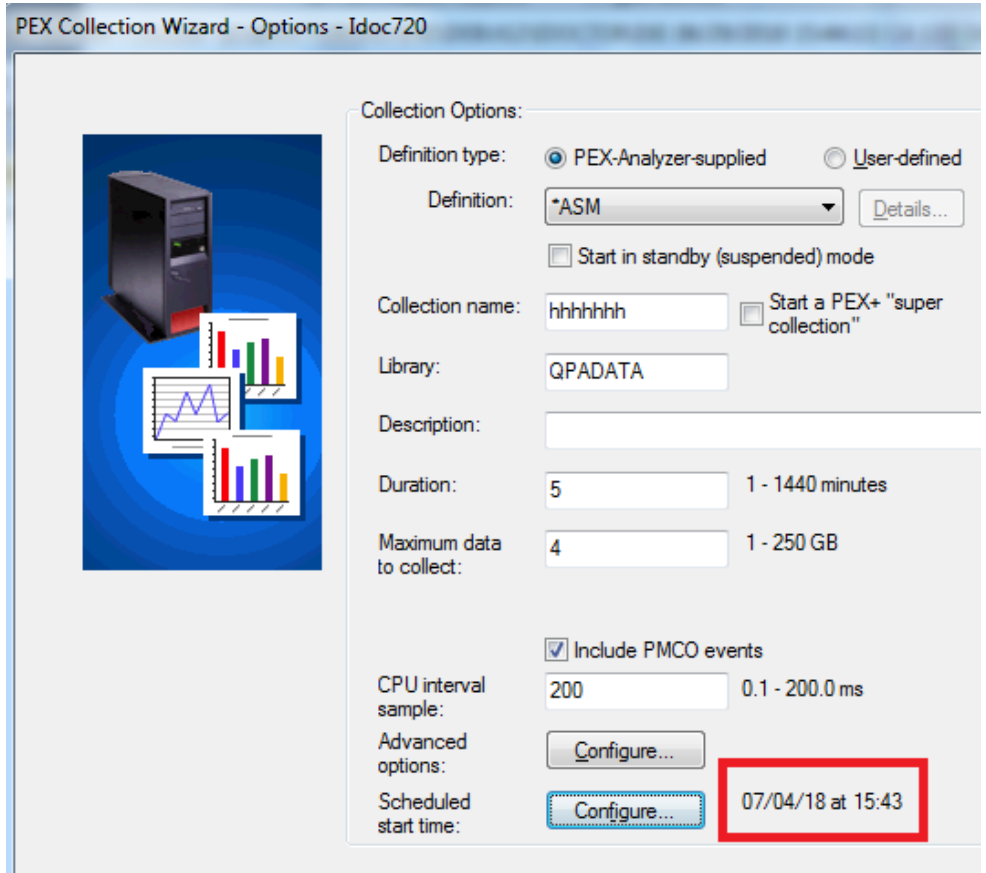

# June 2018 (1320) – GUI crash if drilling down from comparison graphs

- The iDoctor GUI will crash if you produce a comparison graph (i.e. such as a CSI Disk graphs comparison of ASP 1 vs ASP 33) then trying to drill down from there into another graph.

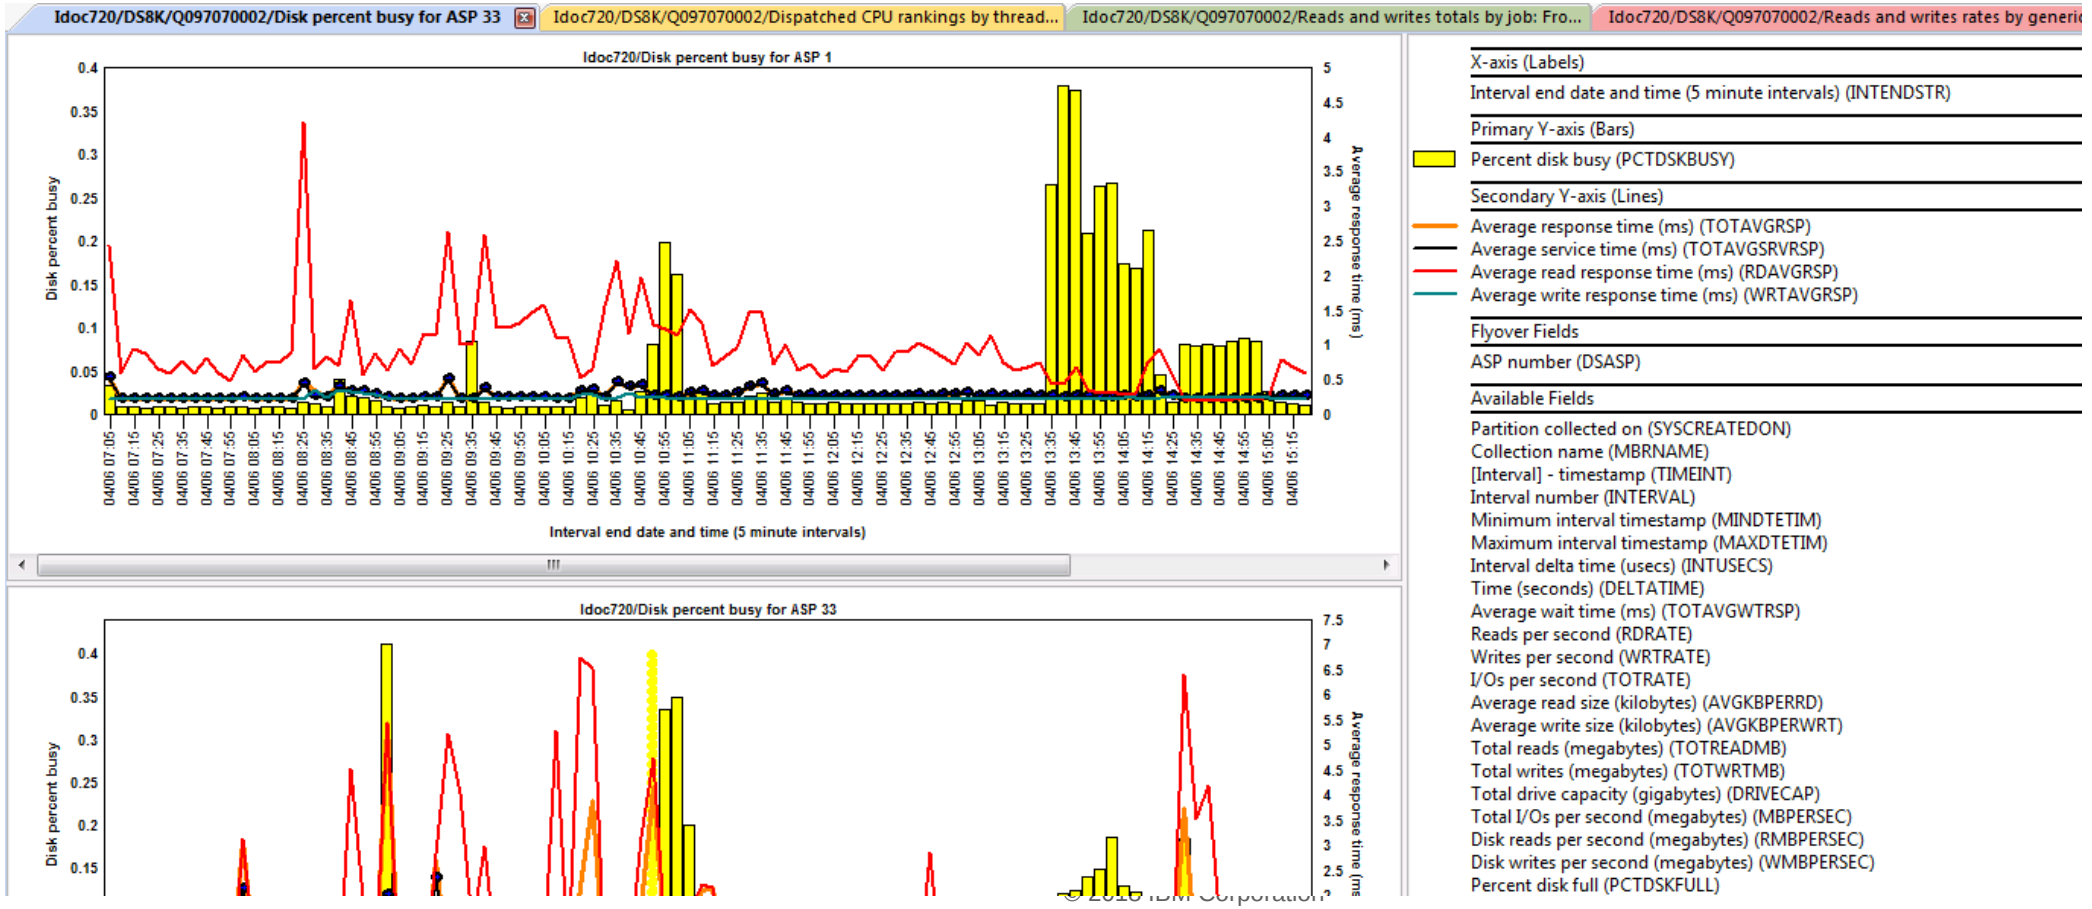

# June 2018 (1320) – Job Watcher objects waited on reports

- In Job Watcher under the Wait graphs -> Objects waited on folder made the following updates:

  The folder will now appear even if the collection summary analysis has NOT be ran.

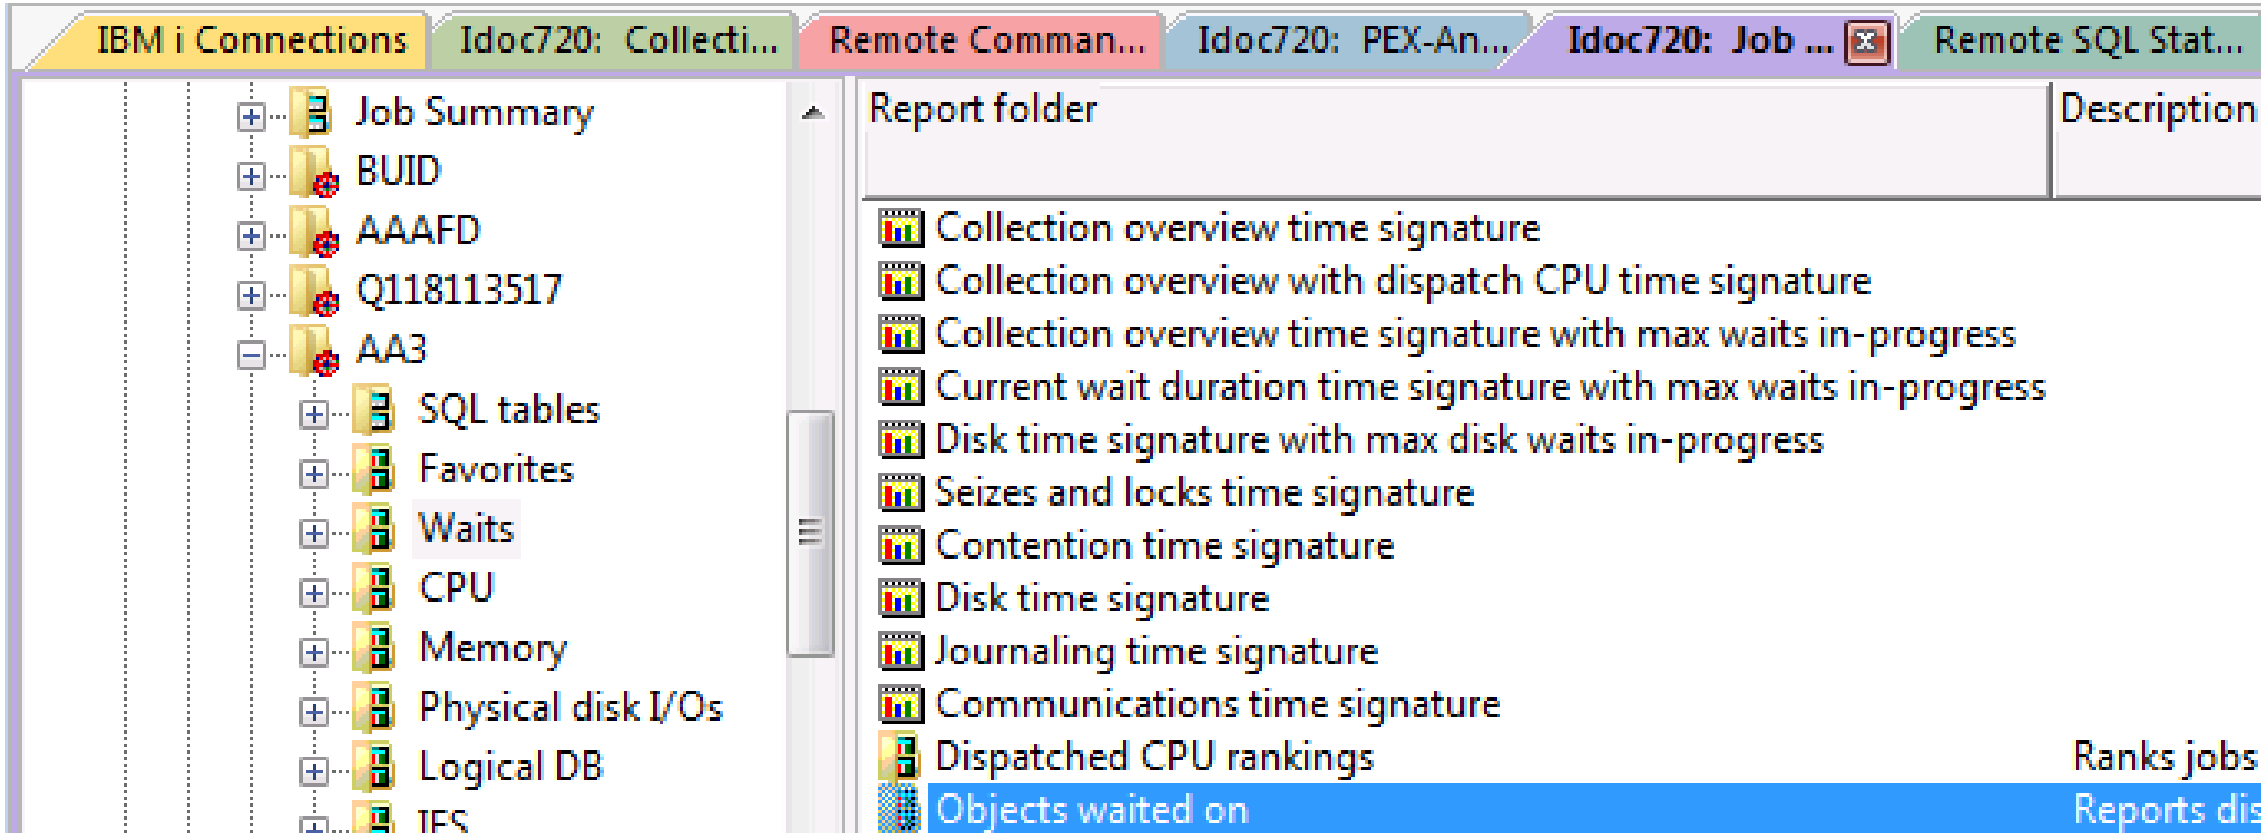

# June 2018 (1320) – Job Watcher objects waited on reports

- In Job Watcher under the Wait graphs -> Objects waited on folder (or Detail reports -> Objects waited on) made the following updates:

  The reports that show "all waits of interest" will now show "Total waits" in the 1st column instead of incorrectly labeling these as "Total page faults".

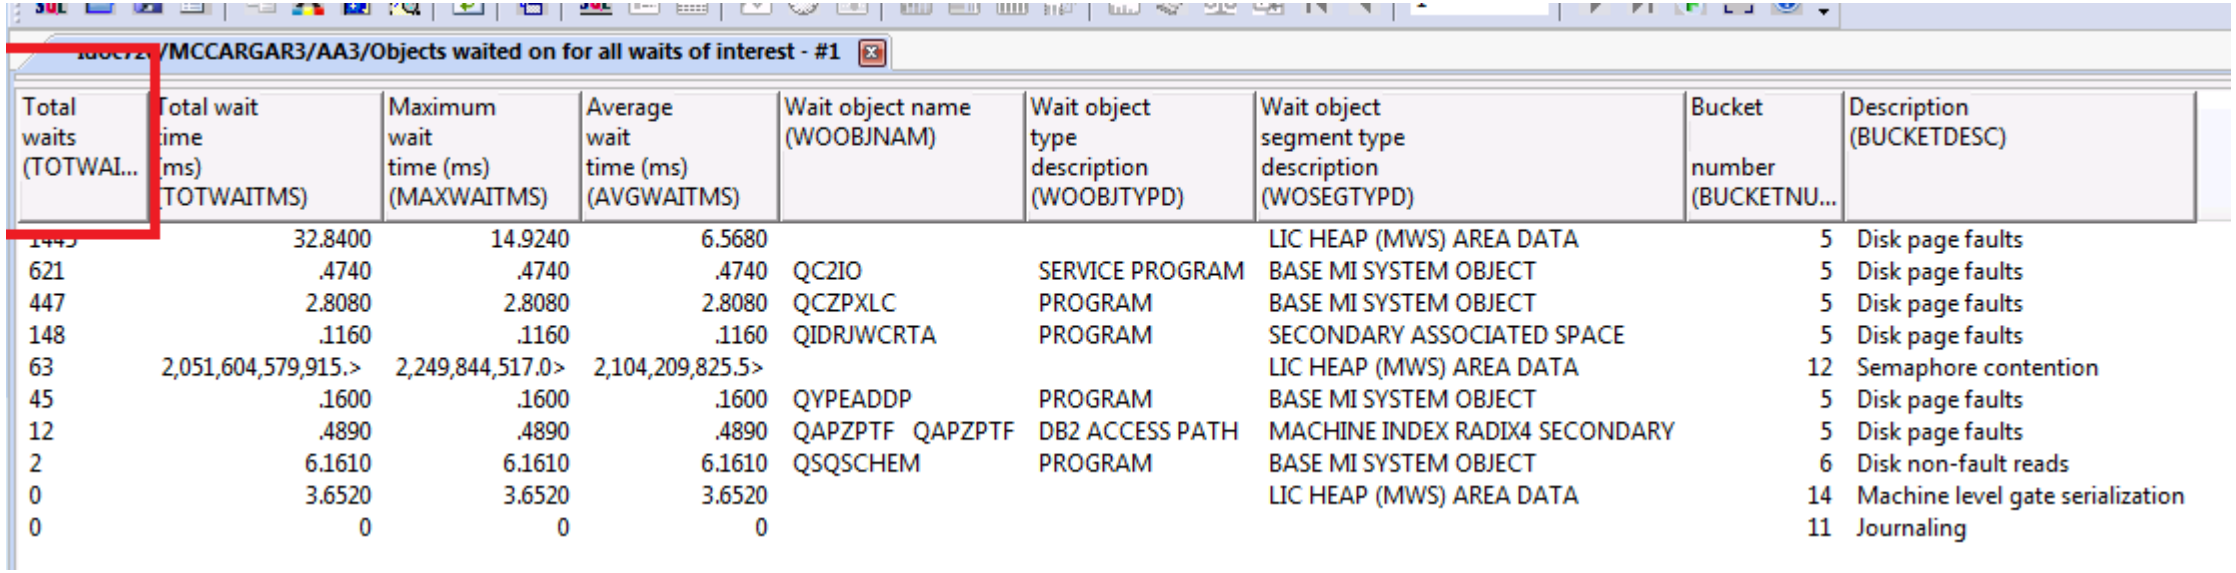

/MCCARGAR3/AA3/Objects waited on for all waits of interest - #1

| Total waits (TOTWAI... | Total wait time (ms) (TOTWAITMS) | Maximum wait time (ms) (MAXWAITMS) | Average wait time (ms) (AVGWAITMS) | Wait object name (WOOBJNAM) | Wait object type description (WOOBJTYPD) | Wait object segment type description (WOSEGTYPD) | Bucket number (BUCKETNU... | Description (BUCKETDESC) |
|---|---|---|---|---|---|---|---|---|
| 1443 | 32.8400 | 14.9240 | 6.5680 | | | LIC HEAP (MWS) AREA DATA | 5 | Disk page faults |
| 621 | .4740 | .4740 | .4740 | QC2IO | SERVICE PROGRAM | BASE MI SYSTEM OBJECT | 5 | Disk page faults |
| 447 | 2.8080 | 2.8080 | 2.8080 | QCZPXLC | PROGRAM | BASE MI SYSTEM OBJECT | 5 | Disk page faults |
| 148 | .1160 | .1160 | .1160 | QIDRJWCRTA | PROGRAM | SECONDARY ASSOCIATED SPACE | 5 | Disk page faults |
| 63 | 2,051,604,579,915.> | 2,249,844,517.0> | 2,104,209,825.5> | | | LIC HEAP (MWS) AREA DATA | 12 | Semaphore contention |
| 45 | .1600 | .1600 | .1600 | QYPEADDP | PROGRAM | BASE MI SYSTEM OBJECT | 5 | Disk page faults |
| 12 | .4890 | .4890 | .4890 | QAPZPTF QAPZPTF | DB2 ACCESS PATH | MACHINE INDEX RADIX4 SECONDARY | 5 | Disk page faults |
| 2 | 6.1610 | 6.1610 | 6.1610 | QSQSCHEM | PROGRAM | BASE MI SYSTEM OBJECT | 6 | Disk non-fault reads |
| 0 | 3.6520 | 3.6520 | 3.6520 | | | LIC HEAP (MWS) AREA DATA | 14 | Machine level gate serialization |
| 0 | 0 | 0 | 0 | | | | 11 | Journaling |

# June 2018 (1320) – Job Watcher objects waited on reports

- In Job Watcher under the Wait graphs -> Objects waited on folder (or Detail reports -> Objects waited on) made the following updates:

  The reports that showed total page faults in the 1st column reported incorrect results. The value before was the total records in the QAPYJWTDE file matching the current selection.

Idoc720/MCCARGAR3/AA3/Collection overview time signature | Idoc720/MCCARGAR3/AA3/Objects waited on due to p

| Total page faults | Total wait time (ms) | Maximum wait time (ms) | Average wait time (ms) | Wait object name | Wait object type description | Wait object segment type description |
|---|---|---|---|---|---|---|
| 5 | 32.8400 | 14.9240 | 6.5680 | | | LIC HEAP (MWS) AREA DATA |
| 1 | 2.8080 | 2.8080 | 2.8080 | QCZPXLC | PROGRAM | BASE MI SYSTEM OBJECT |
| 1 | .4890 | .4890 | .4890 | QAPZPTF QAPZPTF | DB2 ACCESS PATH | MACHINE INDEX RADIX4 SECONDARY |
| 1 | .4740 | .4740 | .4740 | QC2IO | SERVICE PROGRAM | BASE MI SYSTEM OBJECT |
| 1 | .1600 | .1600 | .1600 | QYPEADDP | PROGRAM | |
| 1 | .1160 | .1160 | .1160 | QIDRJWCRTA | PROGRAM | |

Idoc720/MCCARGAR3/AA3/Objects waited on due to page faulting - #1

| Total page faults (FLTS) | Total wait time (ms) (TOTWAIT... | Maximum wait time (ms) (MAXWAIT... | Average wait time (ms) (AVGWAIT... | Wait object name (WOOBJNAM) | Wait object type description (WOOBJTYPD) | Wait object segment type description (WOSEGTYPD) |
|---|---|---|---|---|---|---|
| 1445 | 32.8400 | 14.9240 | 6.5680 | | | LIC HEAP (MWS) AREA DATA |
| 621 | .4740 | .4740 | .4740 | QC2IO | SERVICE PROGRAM | BASE MI SYSTEM OBJECT |
| 447 | 2.8080 | 2.8080 | 2.8080 | QCZPXLC | PROGRAM | BASE MI SYSTEM OBJECT |
| 148 | .1160 | .1160 | .1160 | QIDRJWCRTA | PROGRAM | SECONDARY ASSOCIATED SPACE |
| 45 | .1600 | .1600 | .1600 | QYPEADDP | PROGRAM | BASE MI SYSTEM OBJECT |
| 12 | .4890 | .4890 | .4890 | QAPZPTF QAPZPTF | DB2 ACCESS PATH | MACHINE INDEX RADIX4 SECONDARY |

# June 2018 (1320) – Job Summary analysis broken at 6.1

- In CSI/JW the Job Summary analysis did not work on release 6.1 and has been fixed.

# June 2018 (1320) – PEX Definition Wizard – Add Jobs updates

- The subsystem field is no longer case sensitive.

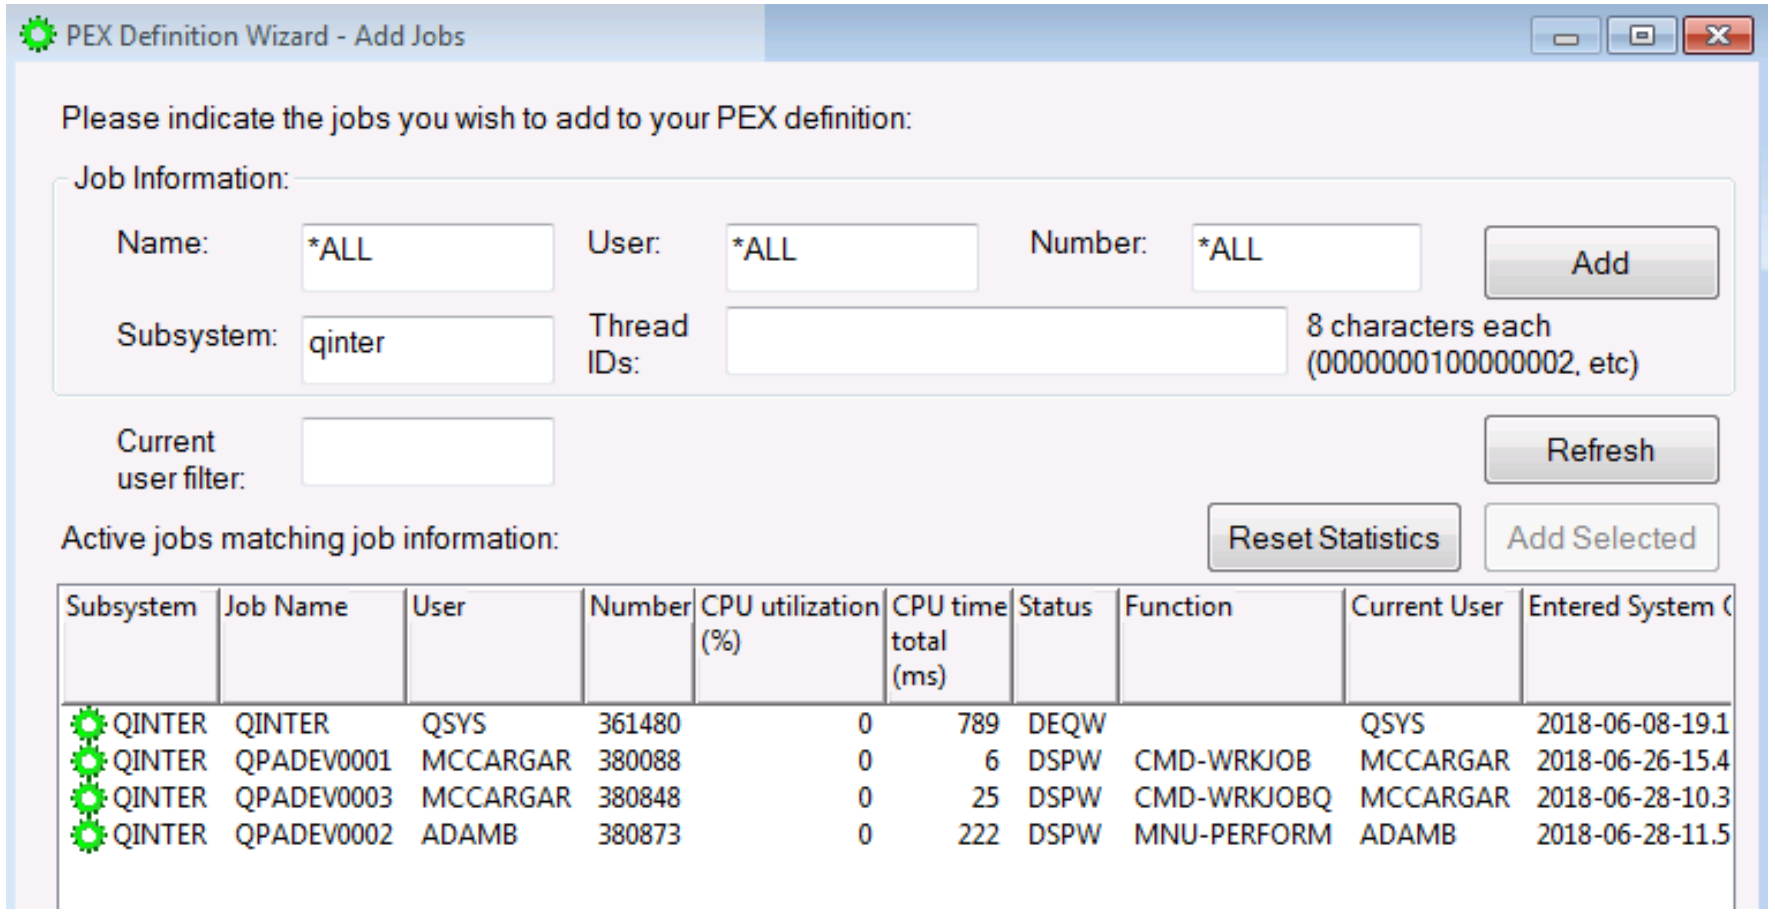

# June 2018 (1320) – PEX Definition Wizard – Add Jobs updates

- The list will now show “No rows found” if no results were found when during a search.

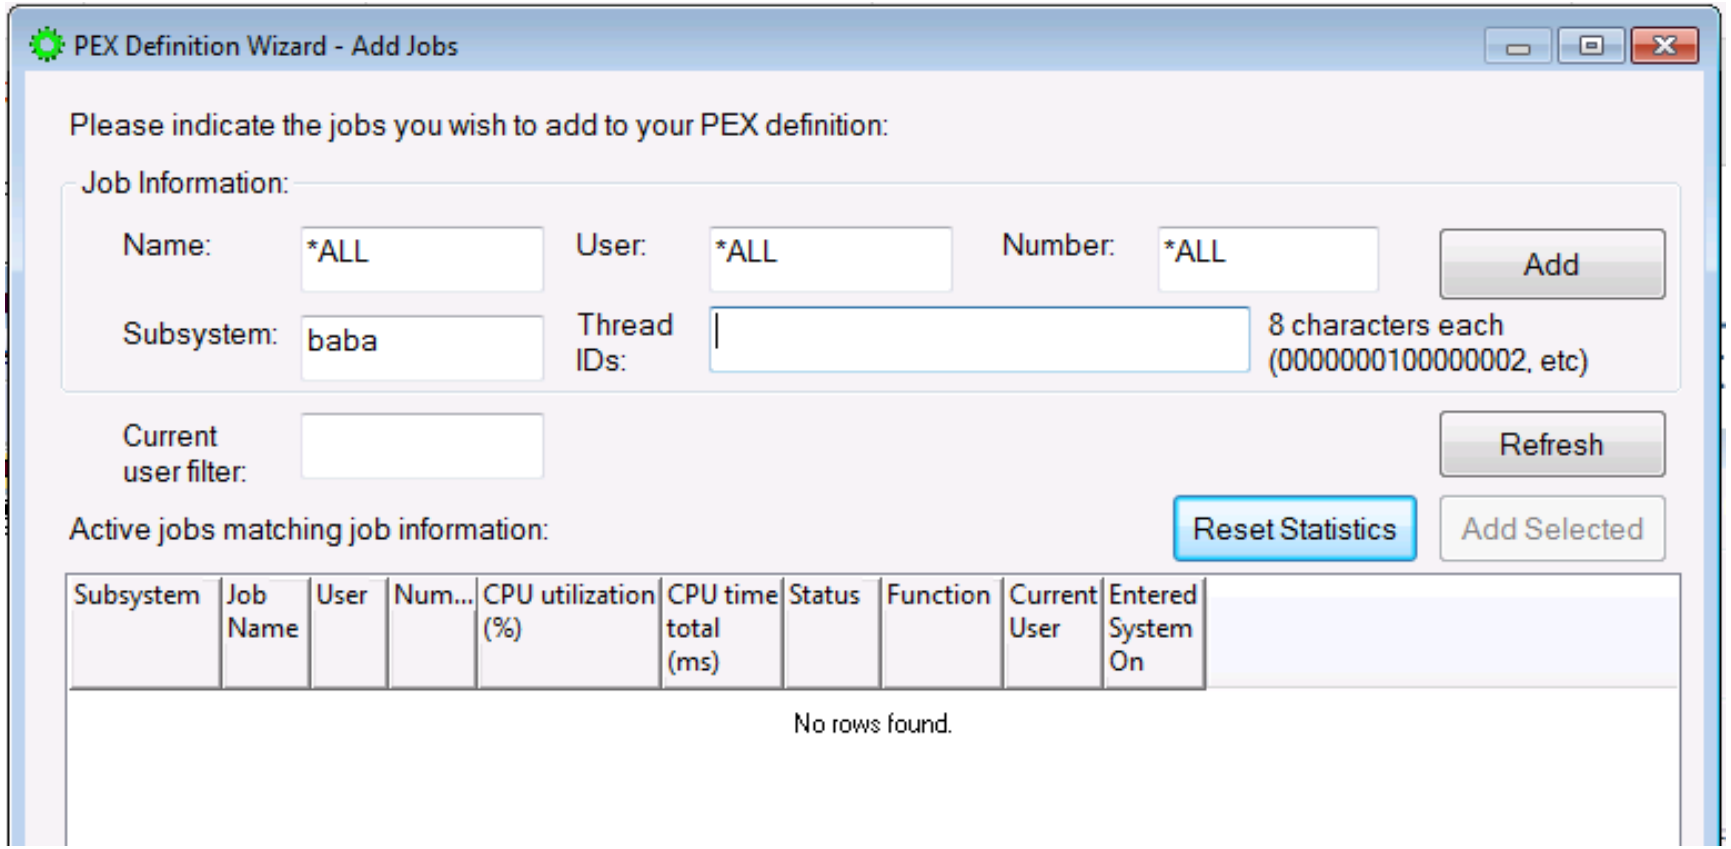

# June 2018 (1320) – CSI HTTP graph fixes

- In CSI under the Communication graphs -> HTTP (Apache) folder, made the following updates:

  1. The rates were calculated incorrectly in some cases.
  2. Fixed a potential divide by zero error in the SQL statements.
  3. The Y2 axis is now always labeled from 0 to 100%. In some cases no line was drawn at all if all values were zero.
  4. Instead of hiding time intervals where no HTTP requests were sent/received these blank intervals will now be shown. This behavior is consistent with the rest of the time-based graphs in CSI.

# June 2018 (1319) – JW Job Summary drill down menu fixes

- In Job Watcher made the following fixes and improvements to the menu options that show up under the Job Summary analysis reports:

  The "Selected Thread" option appeared twice (as "Selected Thread" and "Selected thread").

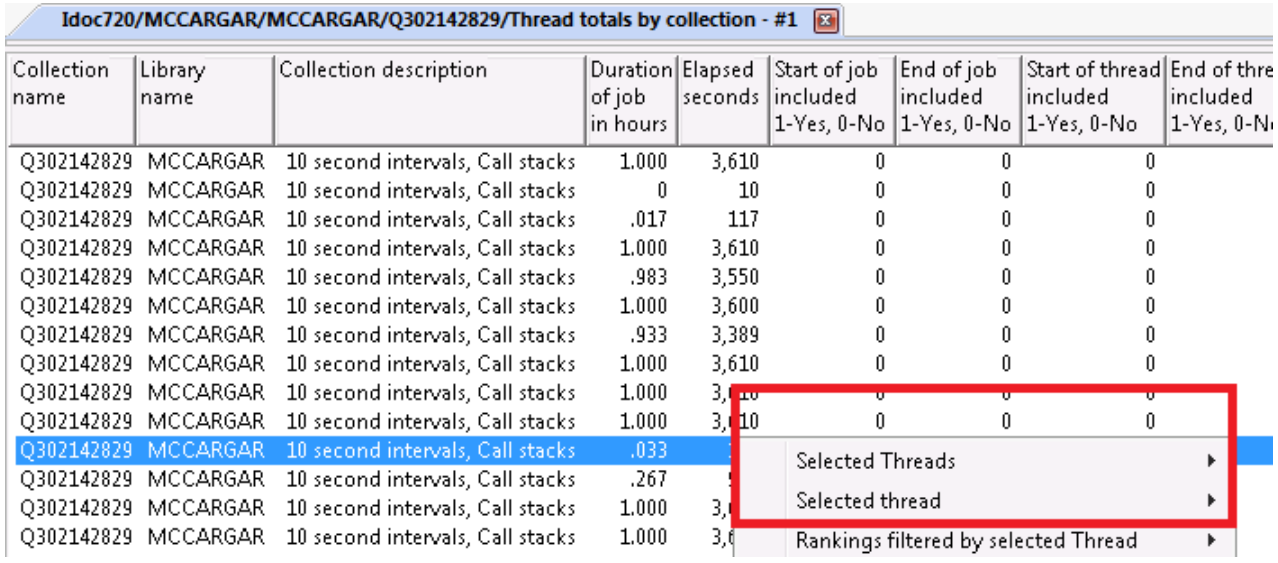


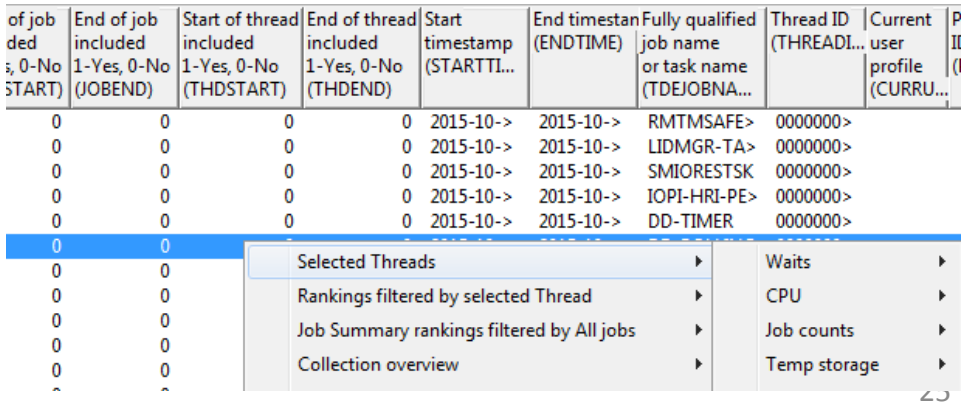

# June 2018 (1319) – JW Job Summary drill down menu fixes

- A 2nd "Rankings" folder appeared in the list but was redundant with the other options (Rankings filtered by selected thread and Collection overview) and has been removed.  It also contained non ranking graphs so was unintuitive the way it was.

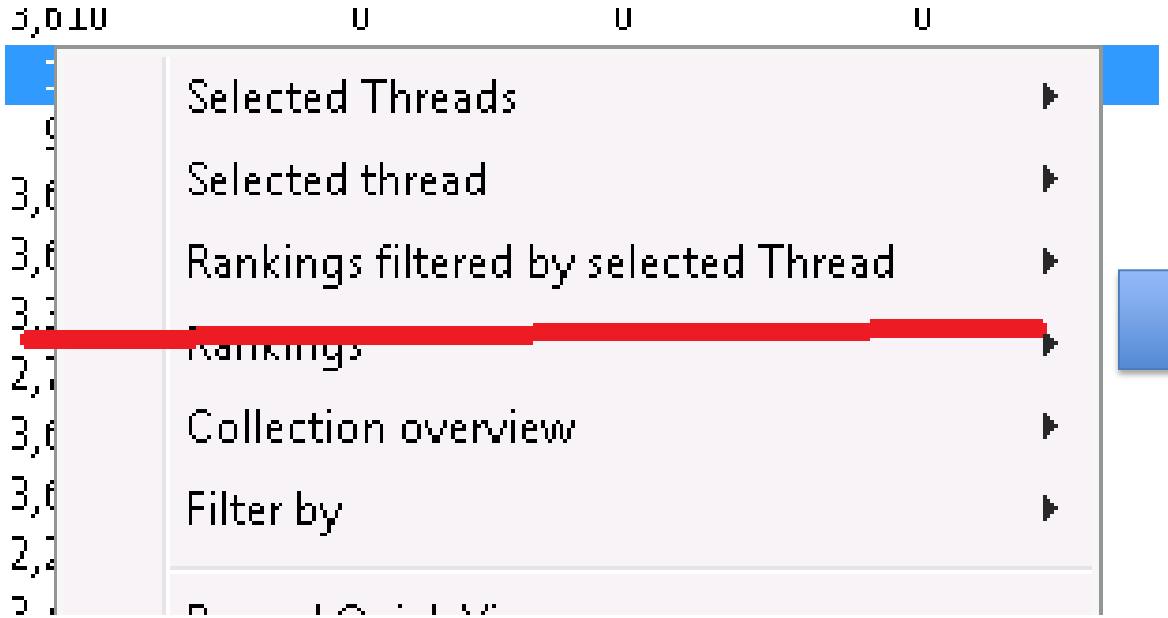


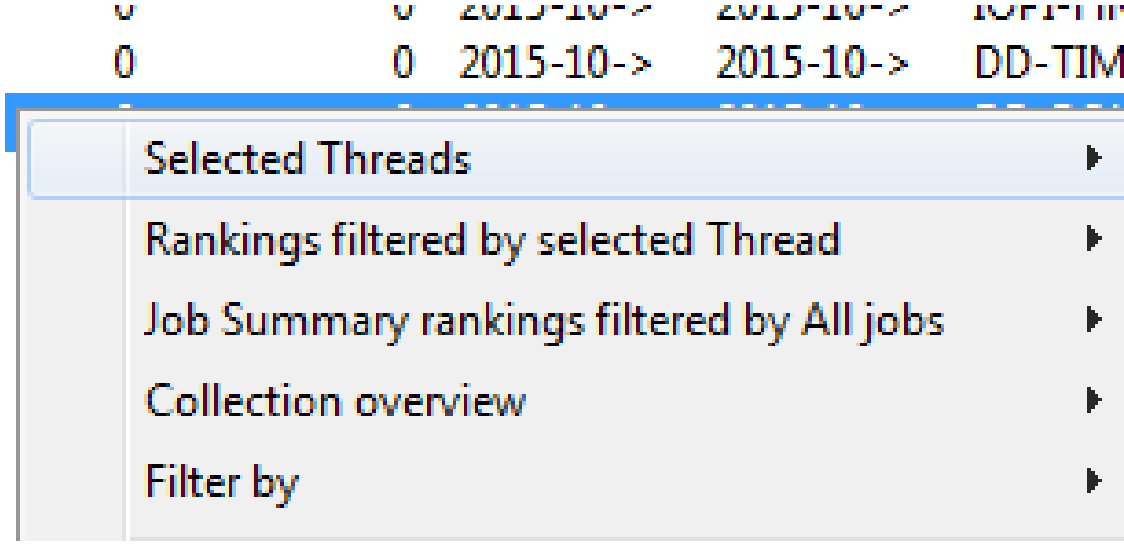

# June 2018 (1319) – JW Job Summary drill down menu fixes

- After drilling down to the Thread signature for a specific job/thread, the "Detail reports" drill down menu option did not appear like it should have, but has been fixed.

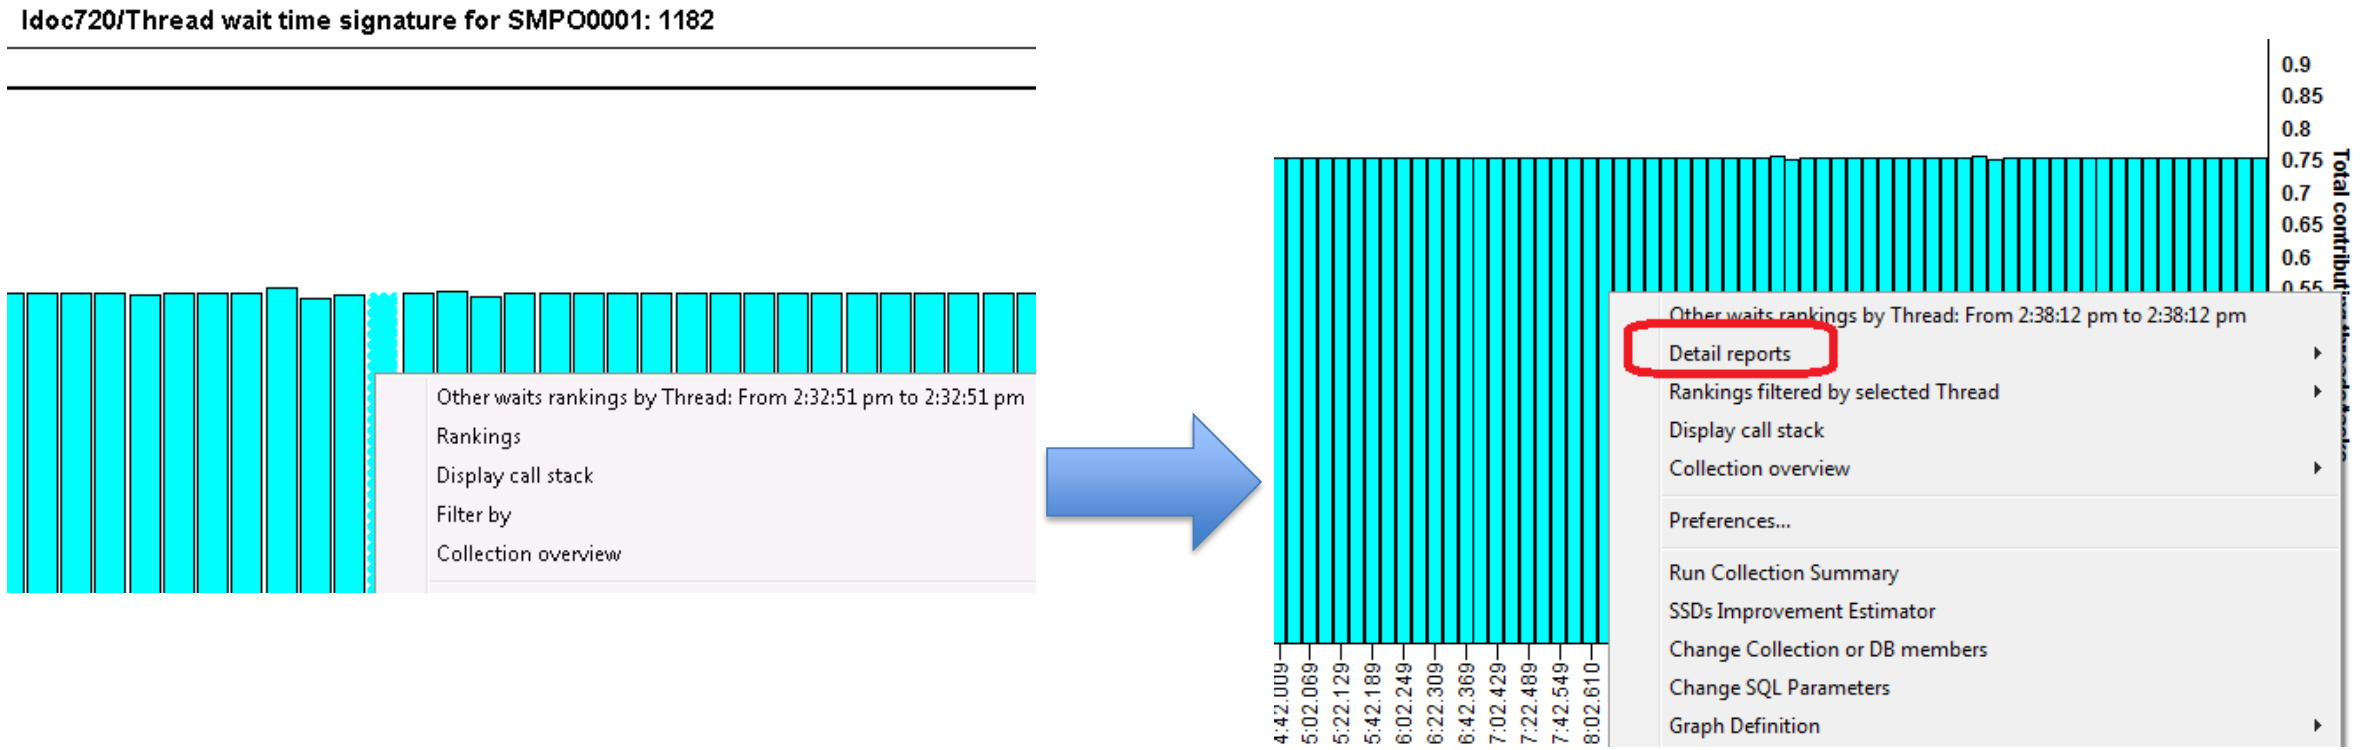

# June 2018 (1319) – JW Job Summary drill down menu fixes

- Moved the menu options that are specific to the Job Summary analysis (SQL results over the job summary sql tables instead of the collection data) from "Rankings -> Wait graph rankings <type>" menu to "Job Summary rankings filtered by <type>".
  - Also renamed these graphs so they match the current naming convention for the wait bucket ranking graphs. (Example: from "Thread signatures ranked by Dispatched CPU" to "Dispatched CPU rankings by thread".)

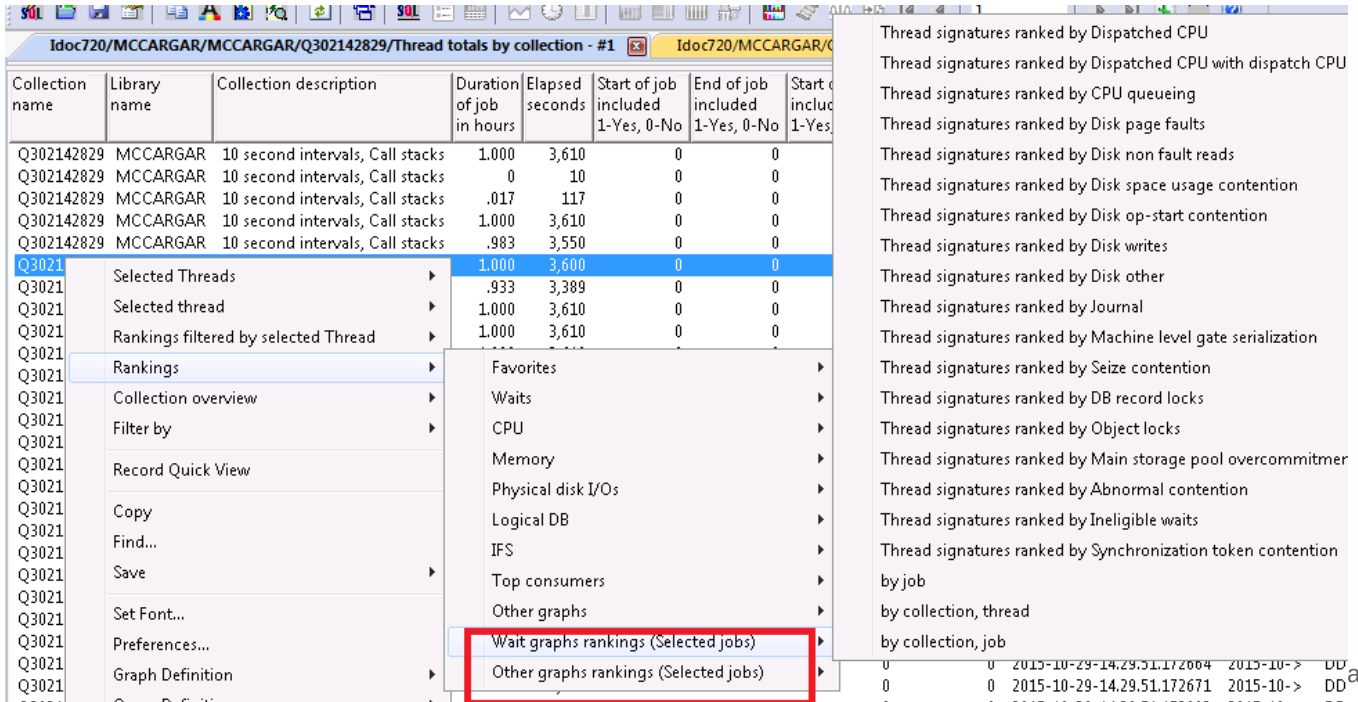


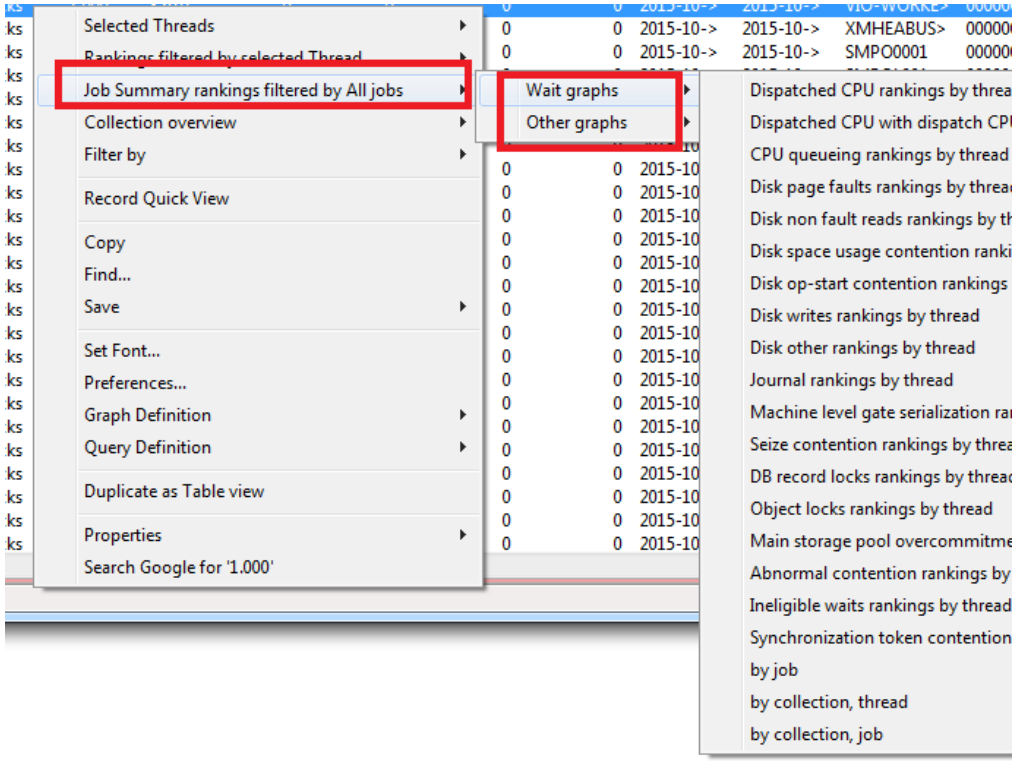

# June 2018 (1319) – JW Job Summary drill down menu fixes

- In some of the drill down graphs from Job Summary analysis the "Filter by" (All jobs, selected job, prompt for generic job) option appeared in the menu where it did not work and should not have appeared.
  - This feature is only intended to apply to the initial SQL table reports in the analysis.
  - It is linked to the Job Summary rankings filtered by <type> option ONLY.

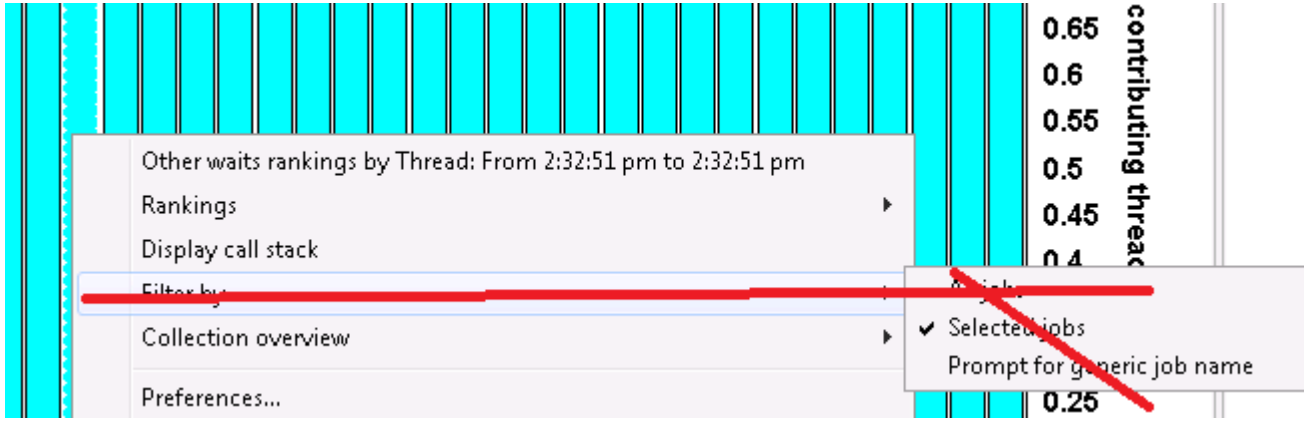

# June 2018 (1319) – JW Job Summary drill down menu fixes

- Support to drill down from one of the **generic job totals** SQL tables was NOT implemented properly (user would get SQL errors). Added by generic job options.

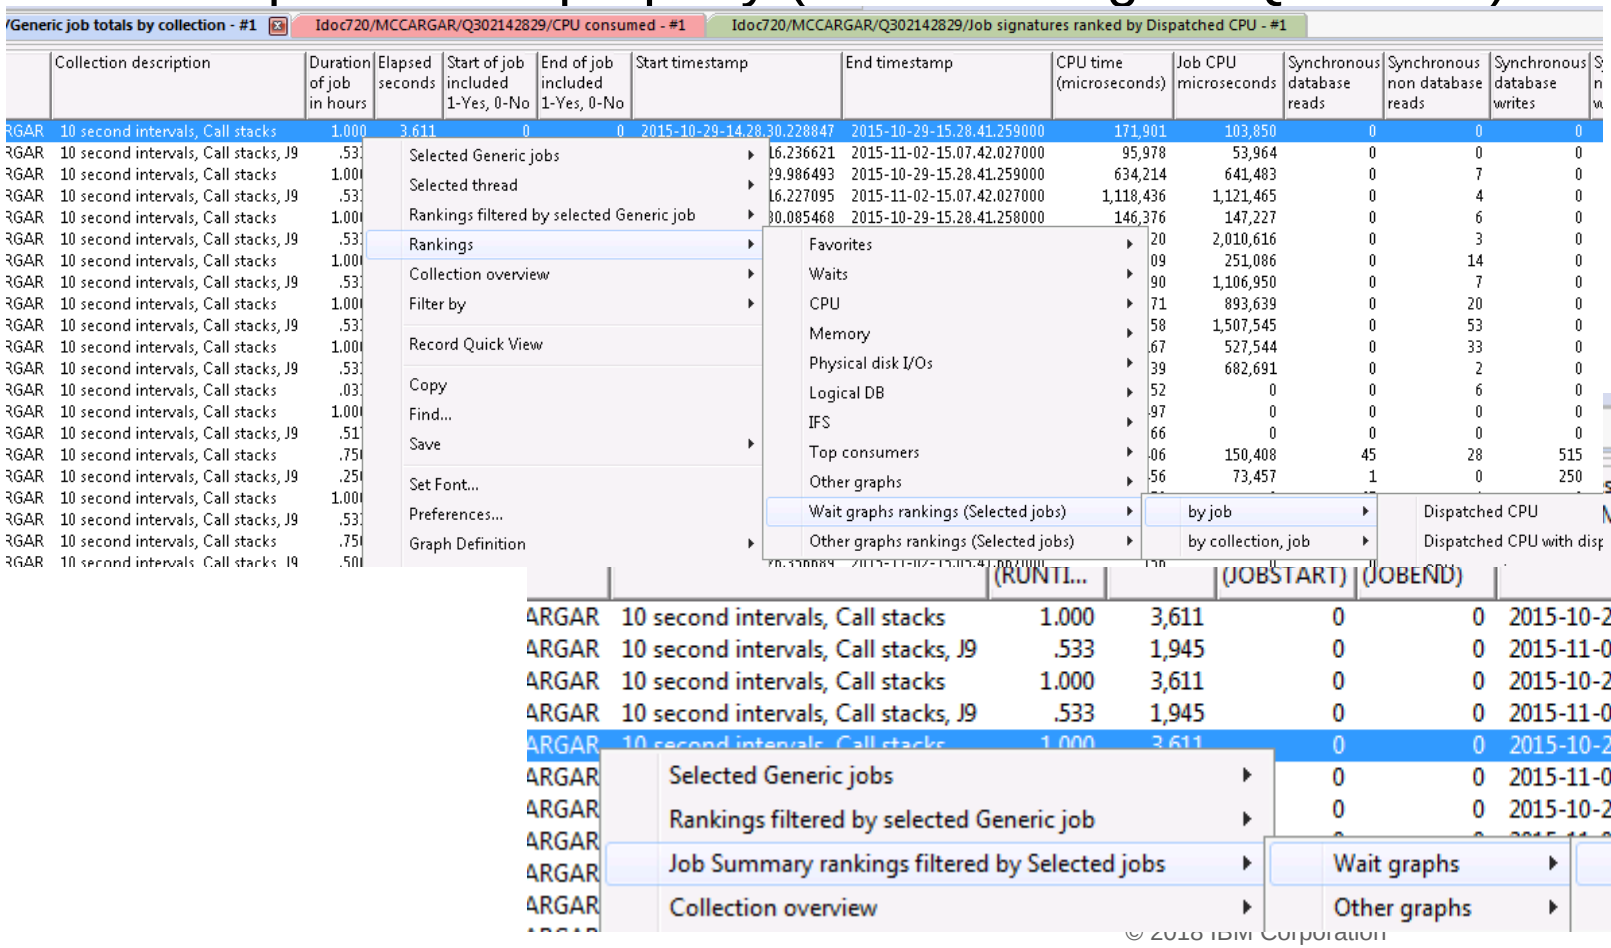


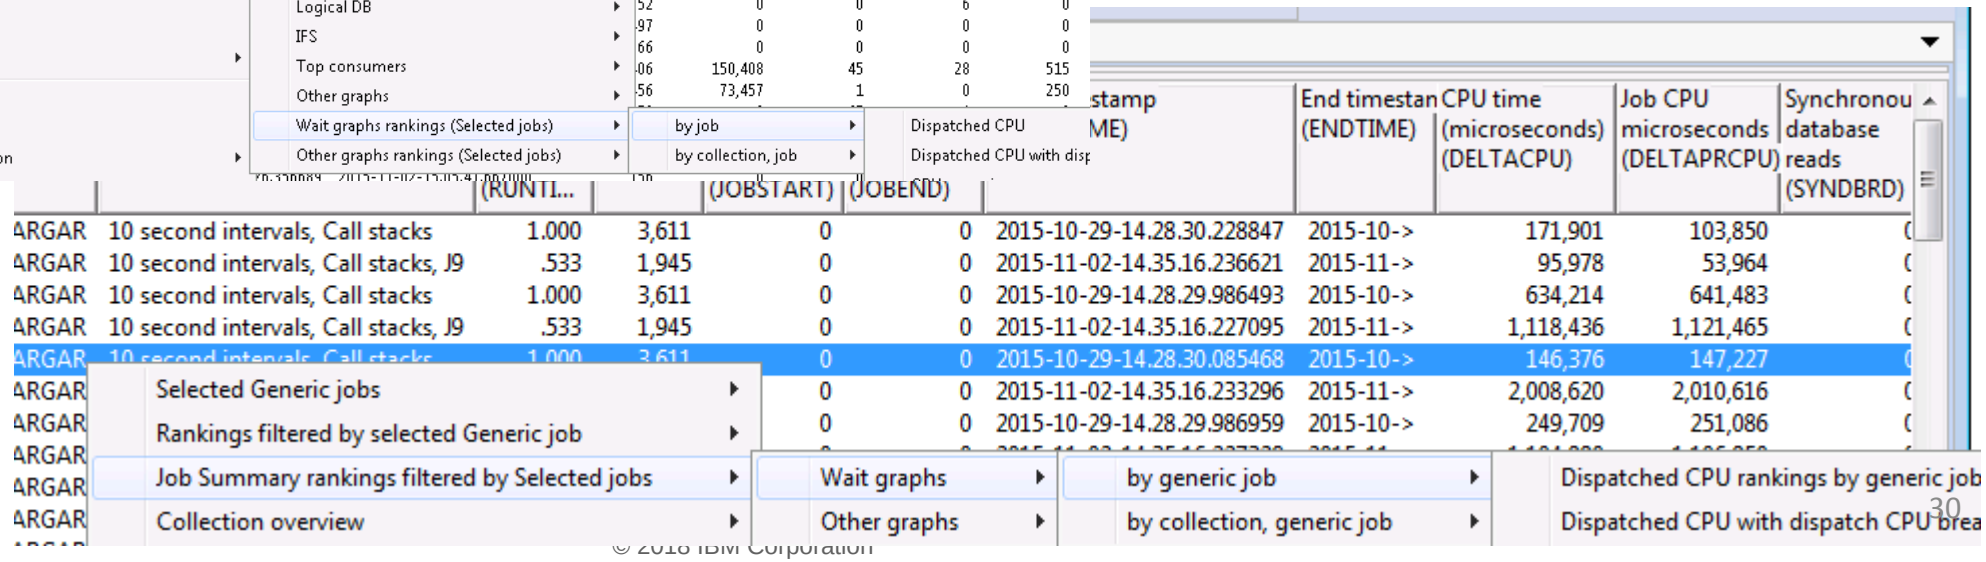

# June 2018 (1319) – JW Job Summary analysis

- The value shown in the Job Summary graph drilldowns for Contributing collections (MBRCOUNT) was wrong in some cases depending on which SQL table you started from.

- All SQL tables now contain a MBRCOUNT column.

# June 2018 (1319) – CSI Job Summary analysis

- Made updates / fixes to CSI Job Summary analysis so the drill down options are consistent with these recent updates to JW Job Summary.

- Added support for  new columns JBHSQLSTMT, JBTICC, JBTICU, JBTTMBU and JBPICC  to be included in the job summary results in CSI at 7.2 and higher.

# June 2018 (1319) – Dispatched CPU waiting vs sharing

- In wait bucket graphs, renamed the column:
  "Dispatched CPU sharing/waiting" to "Dispatched CPU sharing"

  In tables, renamed the column
  "Dispatched CPU waiting" to "Dispatched CPU sharing“

- **Note:** I was not aware “Dispatched CPU waiting” still existed in iDoctor (I thought we had renamed this years ago), but it still remained in table views and there was a CPS about this topic.

# June 2018 (1319) – IBM i connection error message window change

- When connecting to an IBM i and there is a connection failure changed the window shown from a message box to a window that allows copy and paste of the information shown.

  Also improved the contents of the messages shown for some of the most common errors.

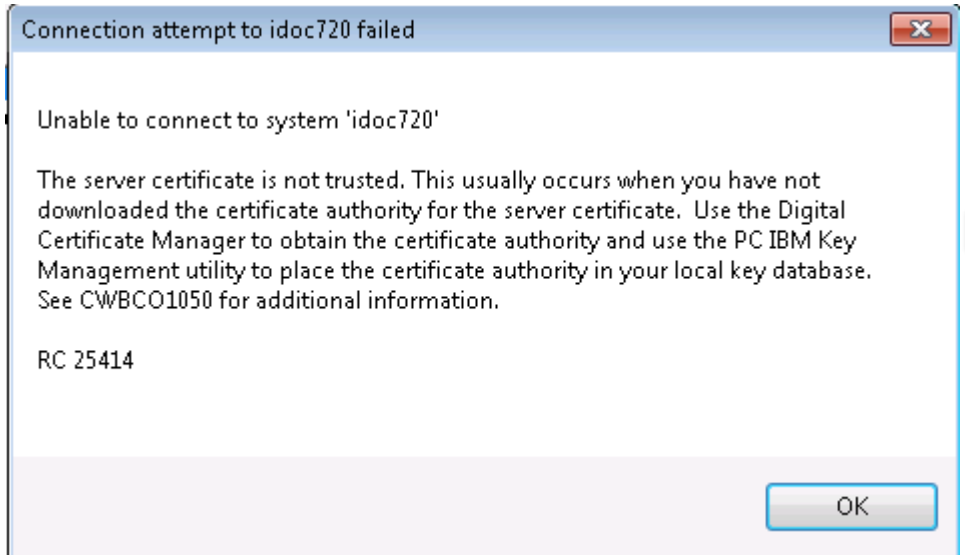


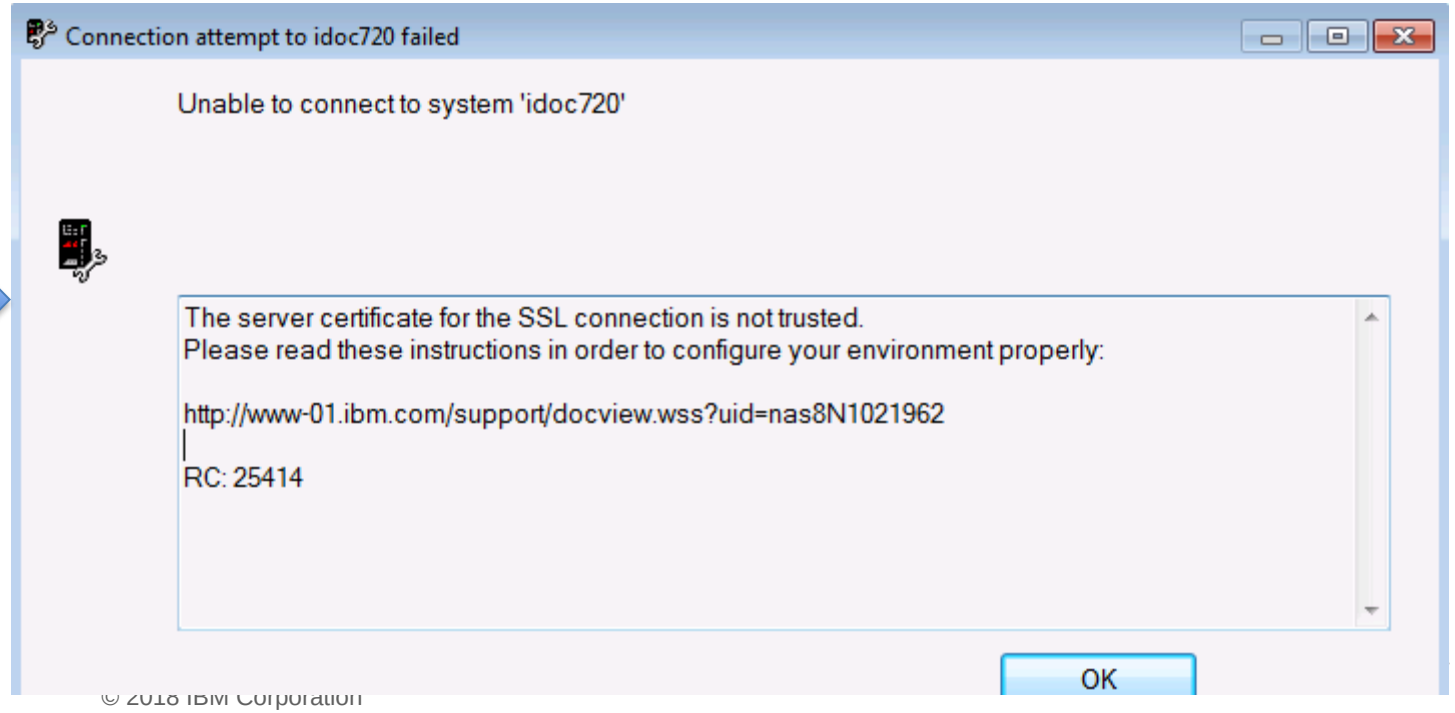

# June 2018 (1318) – Table search now includes # of columns

- The table search function will now include the (number of) columns in the results shown.

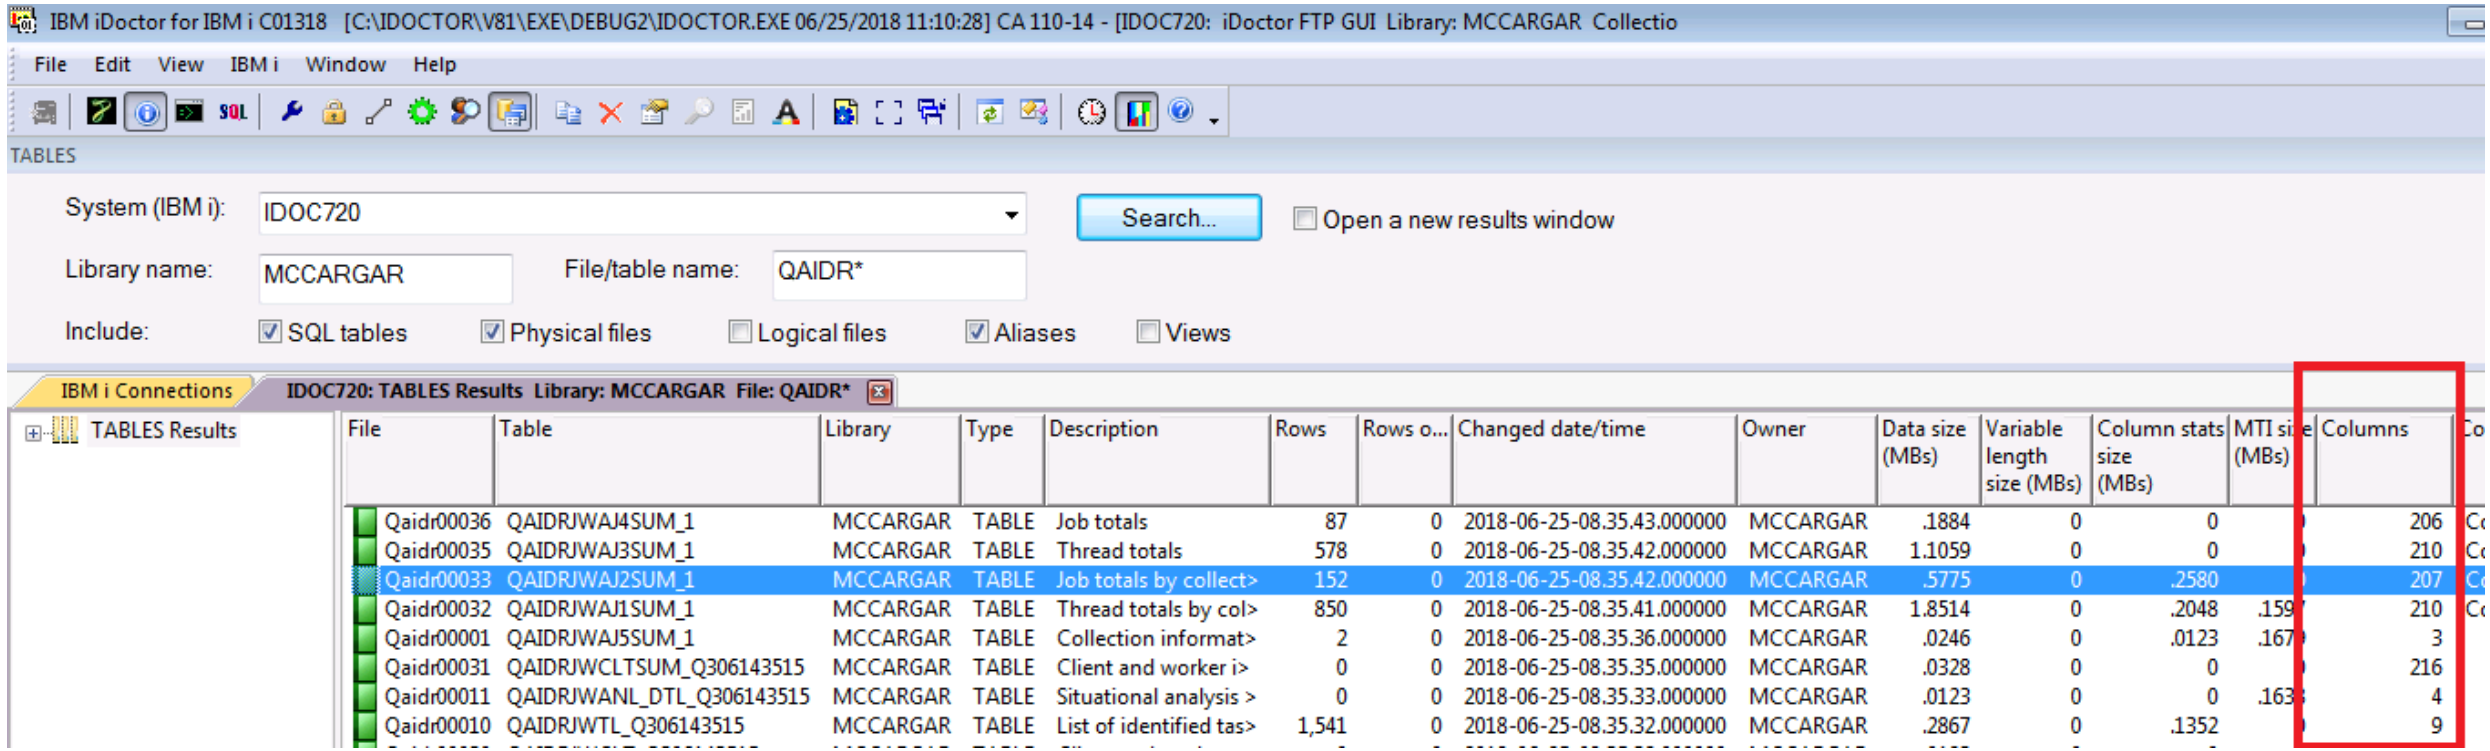

# May 2018 (1317) – CSI disk graph labels for bucket mappings

- In CSI, in the disk graphs at 7.1+, the collection's QAPMCONF response times given in fields G1-GA are now used to provide the labels shown on any of the graphs that show disk response time buckets.

  Note:  The default graphs do not use the B1-B5 mapping at all (because the categorized fields in QAPMDISK do not break down reads vs writes and the ones in QAPMDISKRB do)

  Instead the default graphs use a combination of buckets from the advanced mapping to determine the "simplified" bucket mappings shown in the default disk graphs.  For example
  Bucket 1 = advanced buckets 1-3  (< 1 ms)
  Bucket 2 = advanced buckets 4-6  (> 1-16 ms)
  Bucket 3 = advanced bucket 7  (> 16-64 ms)
  Bucket 4 = advanced bucket 8  (> 64-256 ms)
  Bucket 5 = advanced bucket 9 + 10 (> 256 -1024ms)
  Bucket 6 = advanced bucket 11(> 1024 ms)

# May 2018 (1317) – CSI disk graph labels for bucket mappings

- The primary difference you will notice in the advanced graphs legend, the metrics are no longer all in microseconds.
- In the default graphs there is no noticeable change.

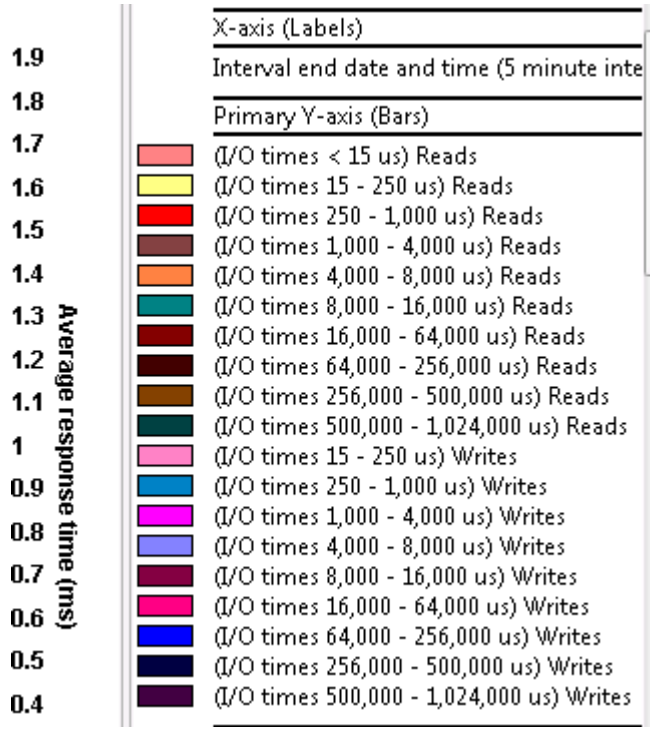


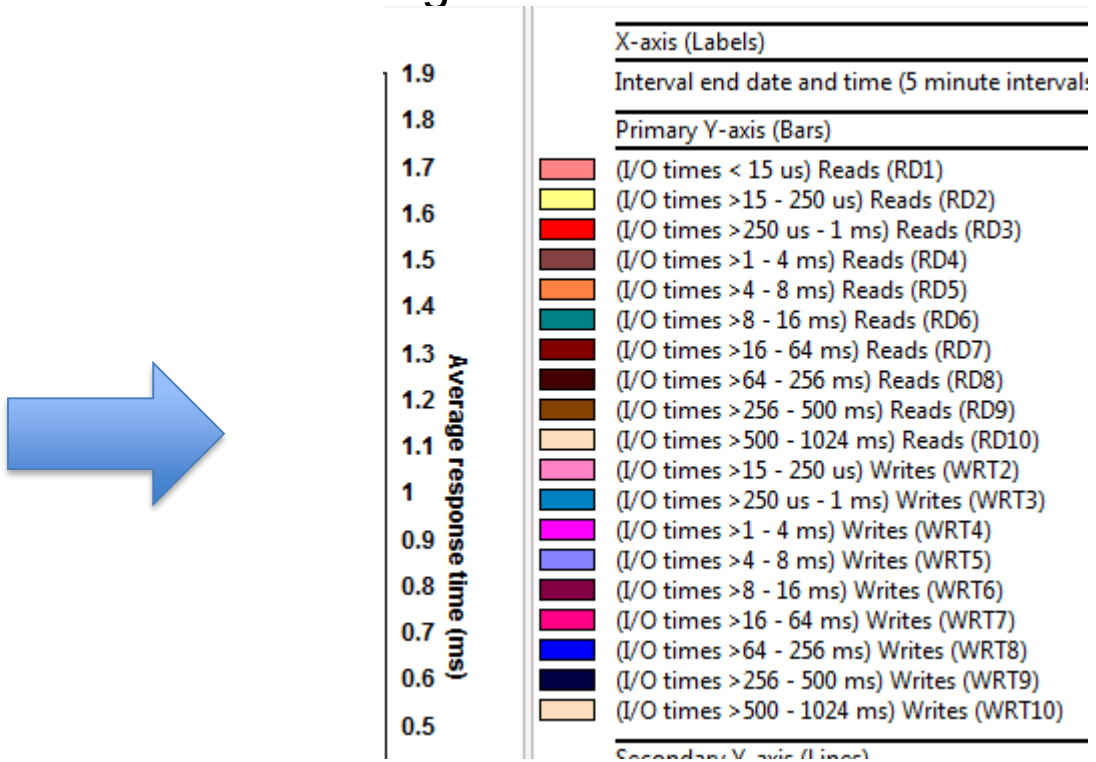

# May 2018 (1317) – CSI system graph updates

- In CSI under the System graphs folder the graph "Physical processor utilization overview [guarded off state only]" has been removed since it would never return any data.

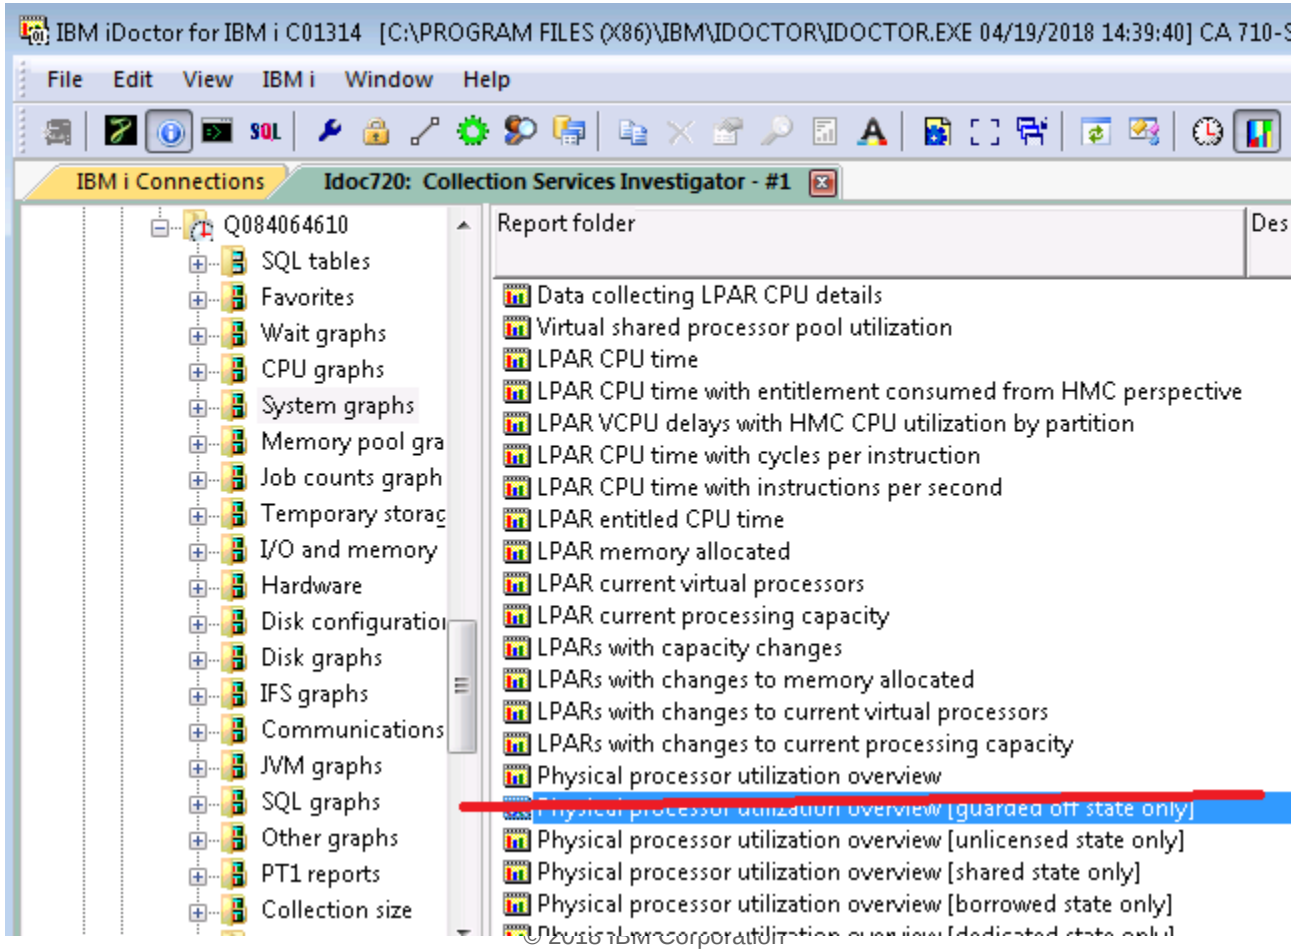

# May 2018 (1317) – CSI system graph updates

- In CSI under the System graphs folder added a new graph "Total physical processors by state" that simply breaks down the total processors for each state occurring over time.

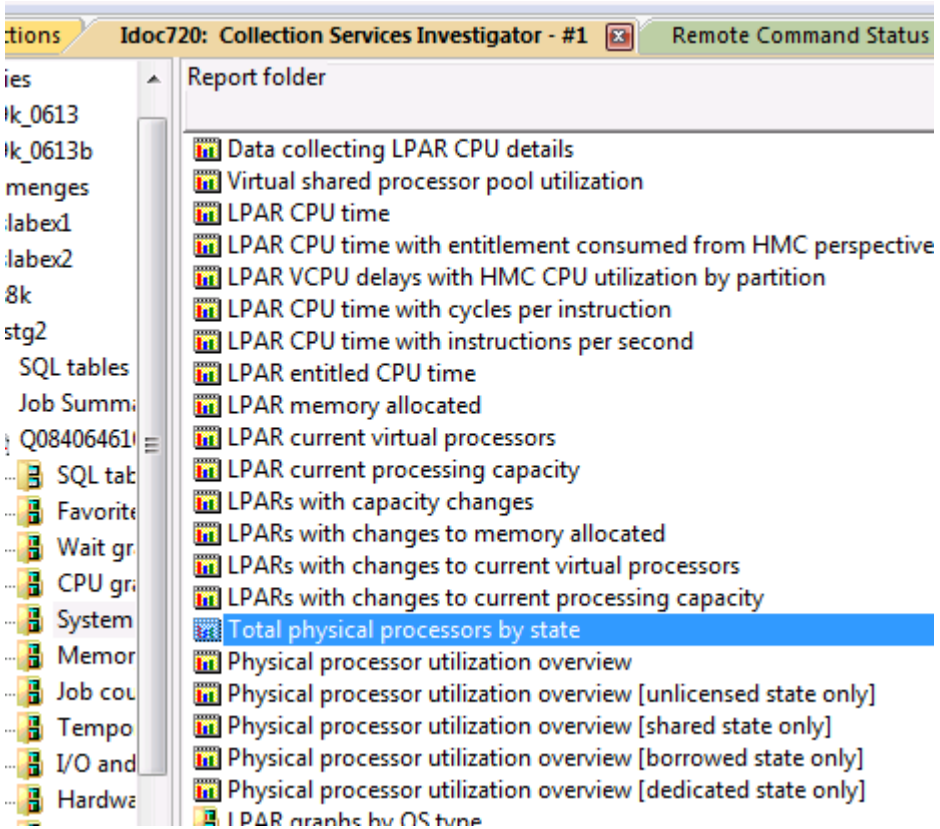


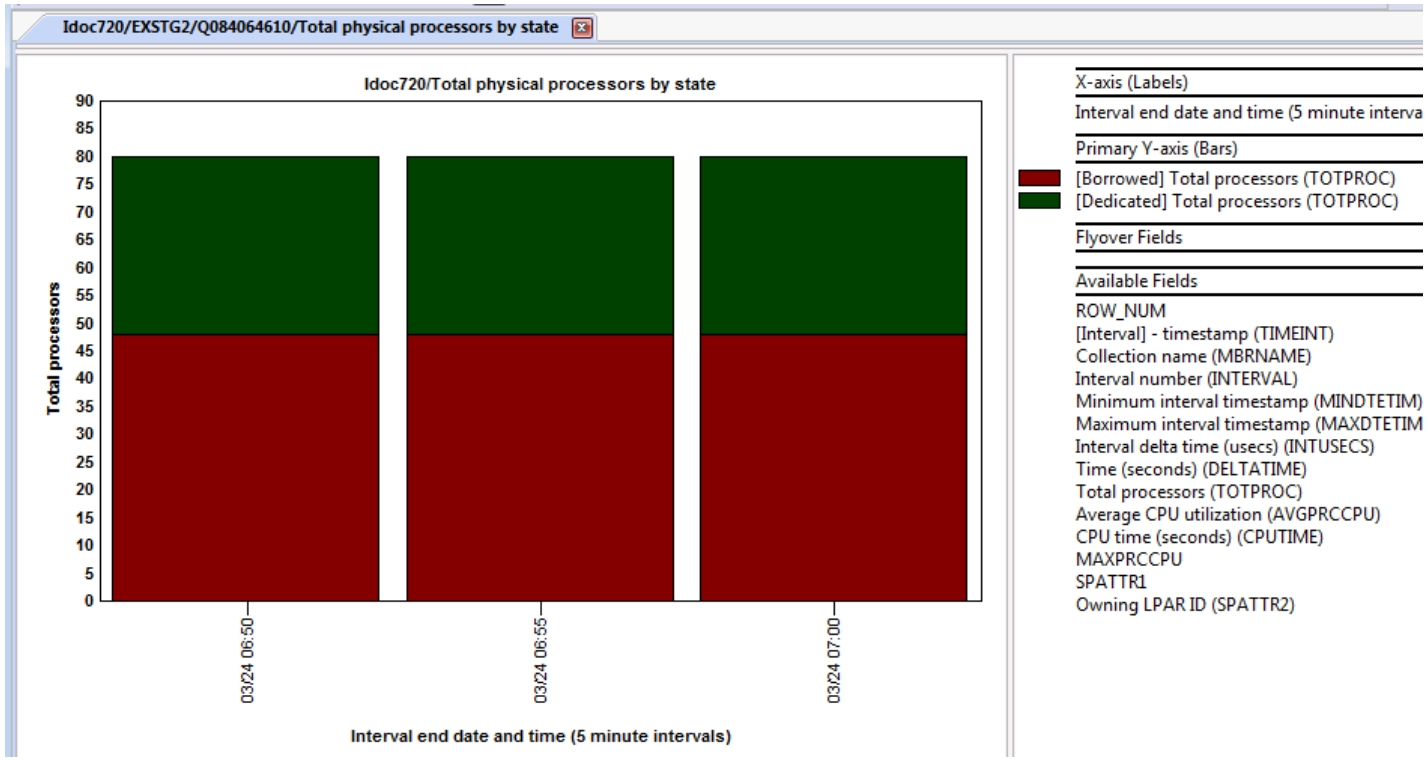

# May 2018 (1316) – PEX Physical Disk I/O graphs redesigned

- We (Brad/Kristie and Ron) decided to make these graphs look and feel much more like CSI disk graphs.
  - We corrected inconsistencies with the graph names, column names, colors, etc.
  - Drill down mechanism looks more consistent as well.

- You can now define your own response time buckets (advanced or default/basic) if you want to.
  - This could be useful if working with newer hardware with very small response times.
  - Existing mapping used in CS at 7.1-7.3 is becoming obsolete. (> 1 second response times useless?!)
  - SQL statements all changed so response time bucket mapping is no longer “hard coded.”

# May 2018 (1316) – PEX Physical Disk I/O graph example

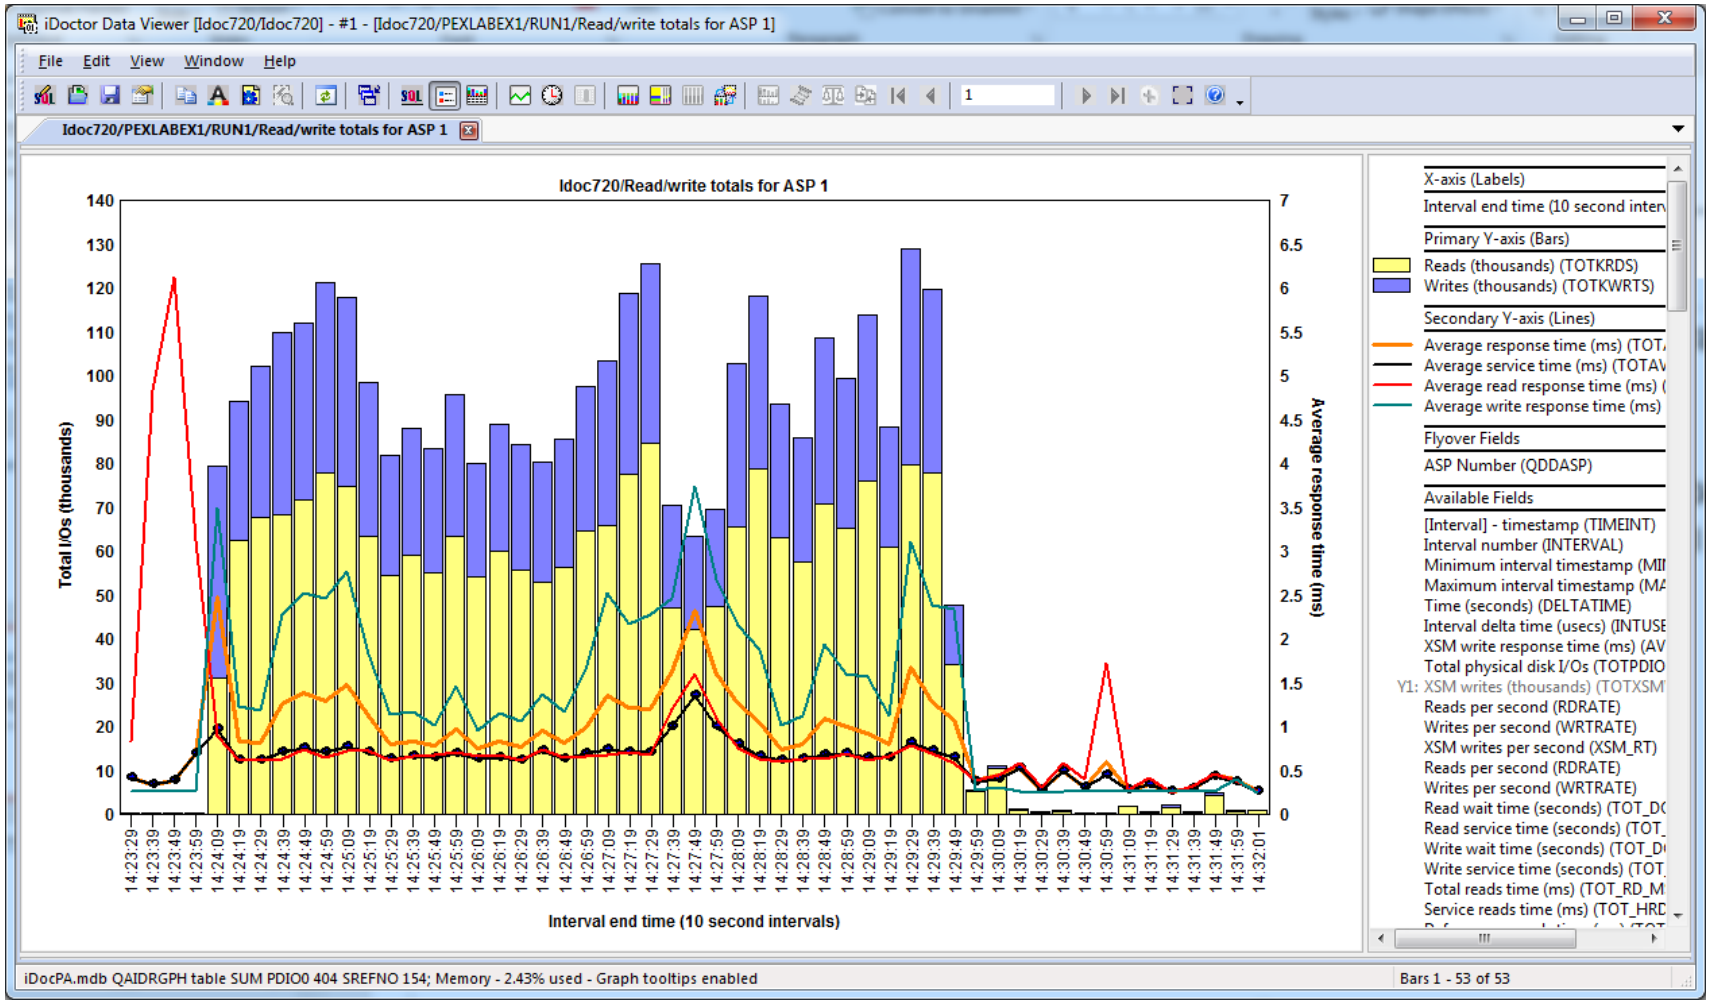

# May 2018 (1316) – PEX Physical Disk I/O example

- 6 read buckets, 6 write buckets

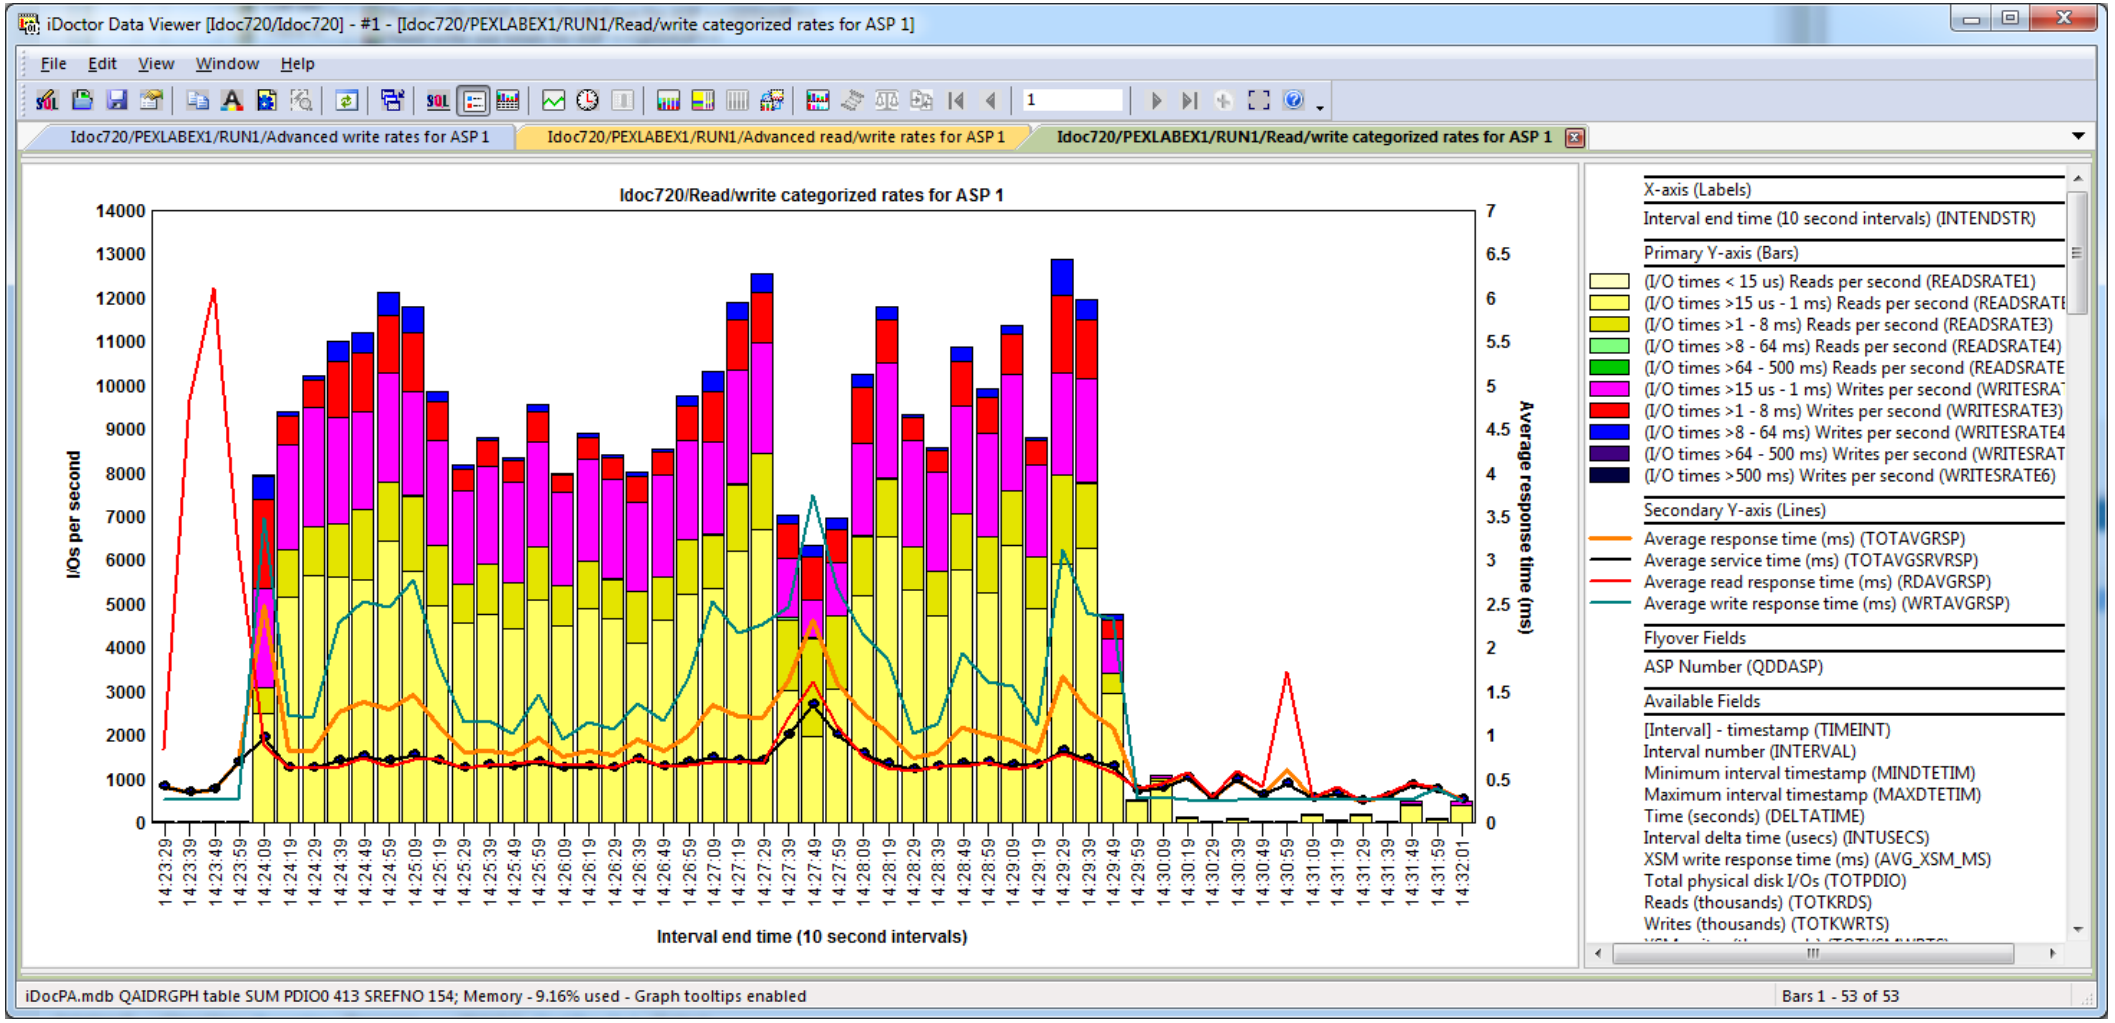

# May 2018 (1316) – PEX Physical Disk I/O advanced example

- 11 read buckets, 11 write buckets

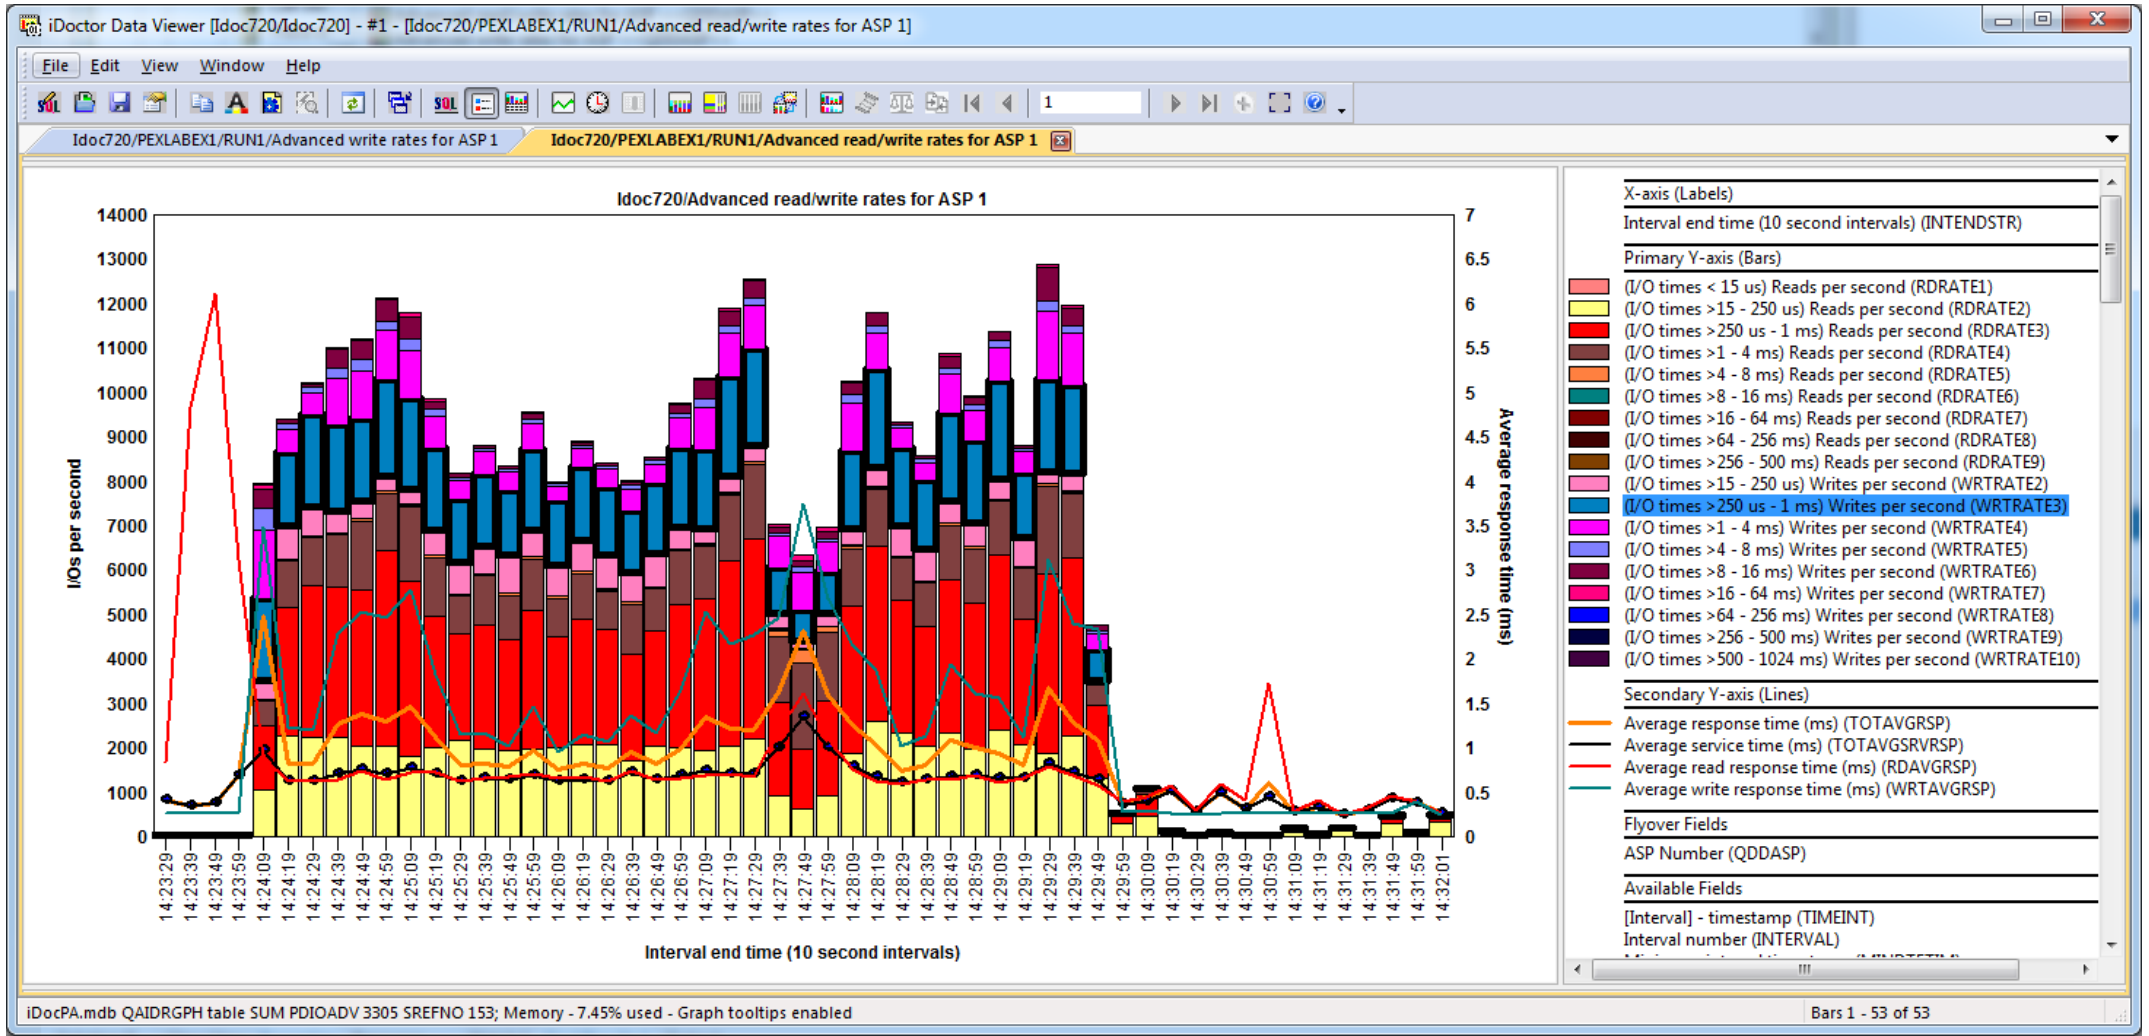

# May 2018 (1315) – Disk response time mapping changes (default)

- The disk response time default mapping used in Collection Services is changing in the next IBM i release (7.4). These are identified in file QAPMCONF and cannot be changed.
- QAPMDISK response time fields are not used, these will likely be removed at 7.4 (set to reserved.)
- Changes necessary since disks are getting faster!

| Bucket | 7.1-7.3 mapping | 7.4 mapping |
|---|---|---|
| 1 | 0 to <= 1 ms | 0 to <= 8 us |
| 2 | > 1 to 16 ms | > 8 to 64 us |
| 3 | > 16 to 64 ms | > 64 to 256 us |
| 4 | > 64 to 255 ms | > 256 us to 1 ms |
| 5 | > 256 to 1024 ms | > 1 to 16 ms |
| 6 | > 1024 ms | > 16 ms |

**Note:** The 7.4 mapping is subject to change but will apply to CS at 7.4 or optionally in PEX/DW at 7.1+ if using the iDoctor mapping V740. IBM Confidential

# May 2018 (1315) – Disk response time mapping changes (advanced)

- The advanced mapping applies to file QAPMDISKRB (in CS) or used by iDoctor in PEX or DW.

| Bucket | 7.1-7.3 mapping | 7.4 mapping |
|---|---|---|
| 1 | 0 to <= 15 us | 0 to <= 8 us |
| 2 | > 15 to 250 us | > 8 to 16 us |
| 3 | > 250 us to 1 ms | > 16 to 64 us |
| 4 | > 1 to 4 ms | > 64 to 128 us |
| 5 | > 4 to 8 ms | > 128 to 256 us |
| 6 | > 8 to 16 ms | > 256 to 512 us |
| 7 | > 16 to 64 ms | > 512 us to 1 ms |
| 8 | > 64 to 256 ms | > 1 to 8 ms |
| 9 | > 256 to 500 ms | > 8 to 16 ms |
| 10 | > 500 to 1024 ms | > 16 to 64 ms |
| 11 | > 1024 ms | > 64 ms |

**Note:** The 7.4 mapping is subject to change but will apply to CS at 7.4 or optionally in PEX/DW at 7.1+ if using the iDoctor mapping V740. IBM Confidential

# May 2018 (1315) – Disk response time mapping file

- iDoctor has added a file QUSRSYS/QAIDRDBKT to keep track of these disk response time mappings.
- This applies to PEX currently only to allow you to choose the mapping to use when analyzing data.
- This could apply to Disk Watcher as well, but will NOT apply in Collection Services.
  - No mechanism exists to change the mapping in CS.
- Run the PEX analysis “Rebuild the disk response times mapping” to create this file (if desired.)
  - Otherwise iDoctor will create it when needed.

# May 2018 (1315) – PEX Physical Disk I/O analysis prompt

- When running PEX PDIO you will be prompted for the disk response time mapping to use.
- You can use an existing iDoctor defined mapping or create a new one.

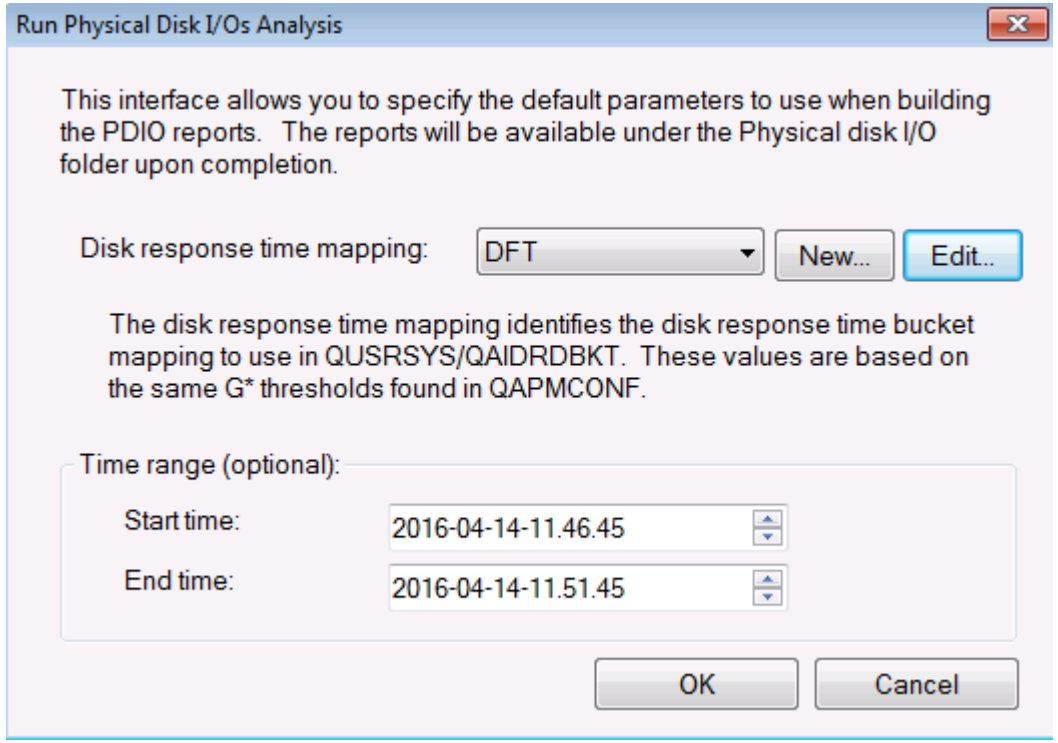

# May 2018 (1315) – Disk response time mapping example 1

- For example these boundaries results in the following **default** 6 bucket graph mapping:

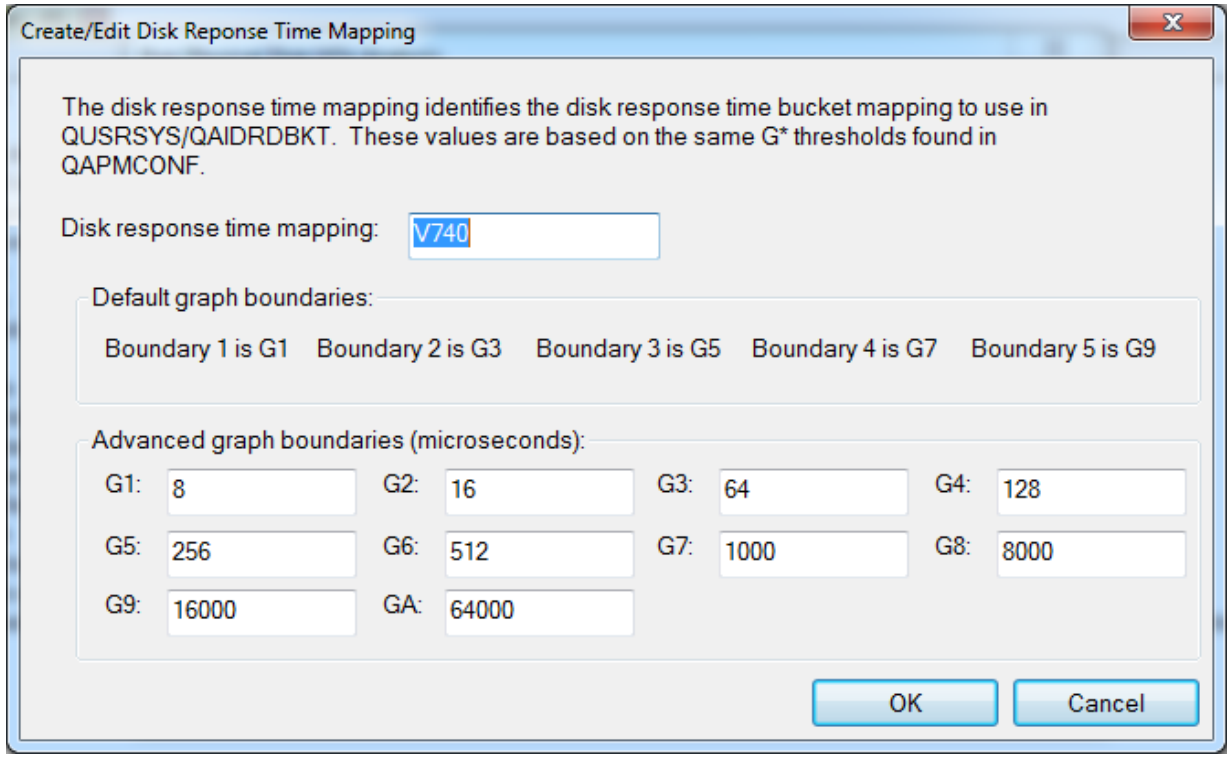


1: 0 to <= 8 microseconds

2: > 8 to 64 microseconds

3: > 64 to 256 microseconds

4: > 256 microseconds to 1 millisecond

5: > 1 millisecond to 16 milliseconds

6: > 16 milliseconds

**Note:** The 7.4 mapping is subject to change but will apply to CS at 7.4 or optionally in PEX/DW at 7.1+ if using the iDoctor mapping V740. IBM Confidential

# May 2018 (1315) – Disk response time mapping example 2

- Or these for the **advanced** 11 bucket graph mapping:

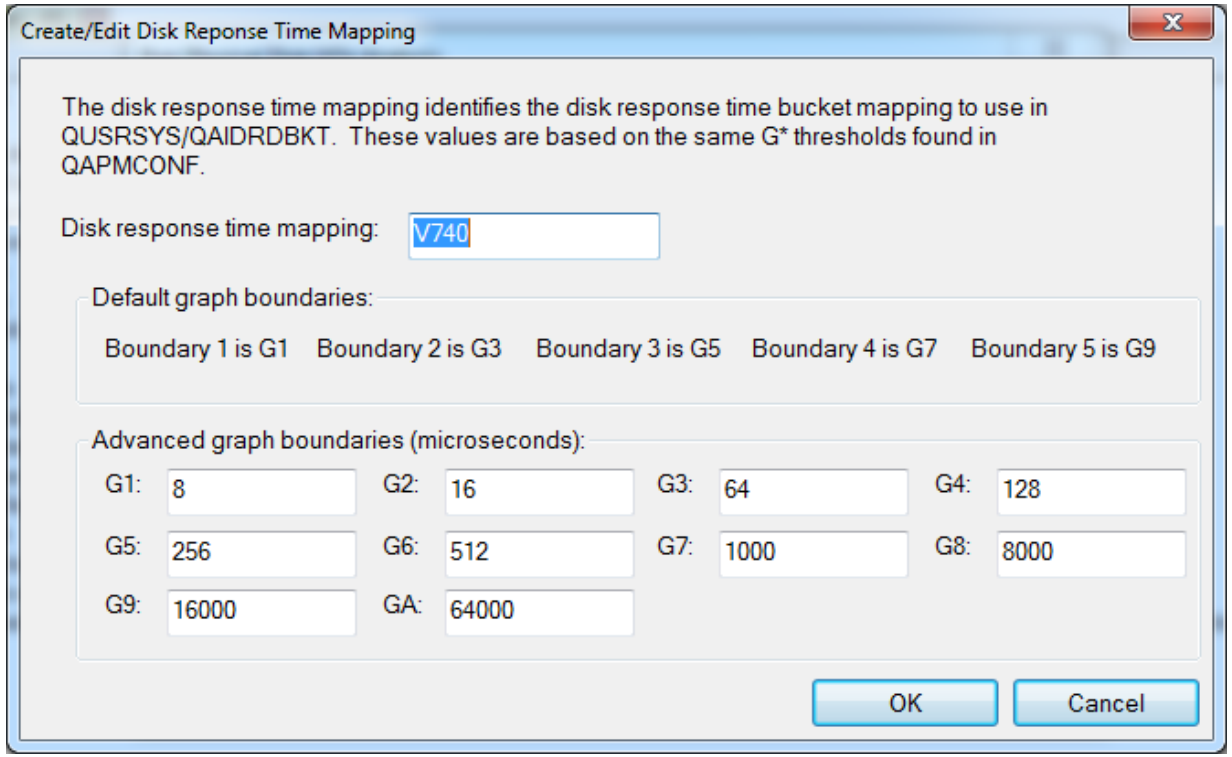


1: 0 to <= 8 microseconds
2: > 8 to 16 microseconds
3: > 16 to 64 microseconds
4: > 64 microseconds to 128 microseconds
5: > 128 microseconds to 256 microseconds
6: > 256 microseconds to 512 microseconds
7: > 512 microseconds to 1 millisecond
8: > 1 millisecond to 8 milliseconds
9: > 8 milliseconds to 16 milliseconds
10: > 16 milliseconds to 64 milliseconds
11: > 64 milliseconds

**Note:** The 7.4 mapping is subject to change but will apply to CS at 7.4 or optionally in PEX/DW at 7.1+ if using the iDoctor mapping V740. IBM Confidential

# April 2018 (1314) – Create Indexes analysis added in CSI/JW

- **<u>Note: These were already added externally in client 1314, but leaving here for reference.</u>**

- To speed up some of the graphs you can now run the Create indexes analysis.
  - Run the Collection Summary analysis first before using this.
  - It does NOT run automatically as a “default” analysis.

- Do NOT use this on actively running collections.

- It primarily speeds up the SQL for several flavors of rankings graphs.

- You can delete these indexes under the SQL Tables -> Create Indexes folder (right-click -> Delete… menu)