iDoctor Updates
(August 10th 2018 to present)

Ron McCargar
iDoctor development
IBM i Global Support Center

Currently recommended “stable” builds

- IBM internal:
  - Latest client is 1334 (Oct 9th)
  - Latest "stable" client is 1328 (Sept 5th)

- IBM external:
  - Latest client is 1334 (Oct 9th on ftp site only)
  - Latest "stable/announced" client is 1325 (Aug 7th)

Oct 2018 (1334) – Multiple column sort from graph legend

- Sort option from graph legend now supports multiple columns.
  Use this to reorder the data based on the summation of the columns you have selected.

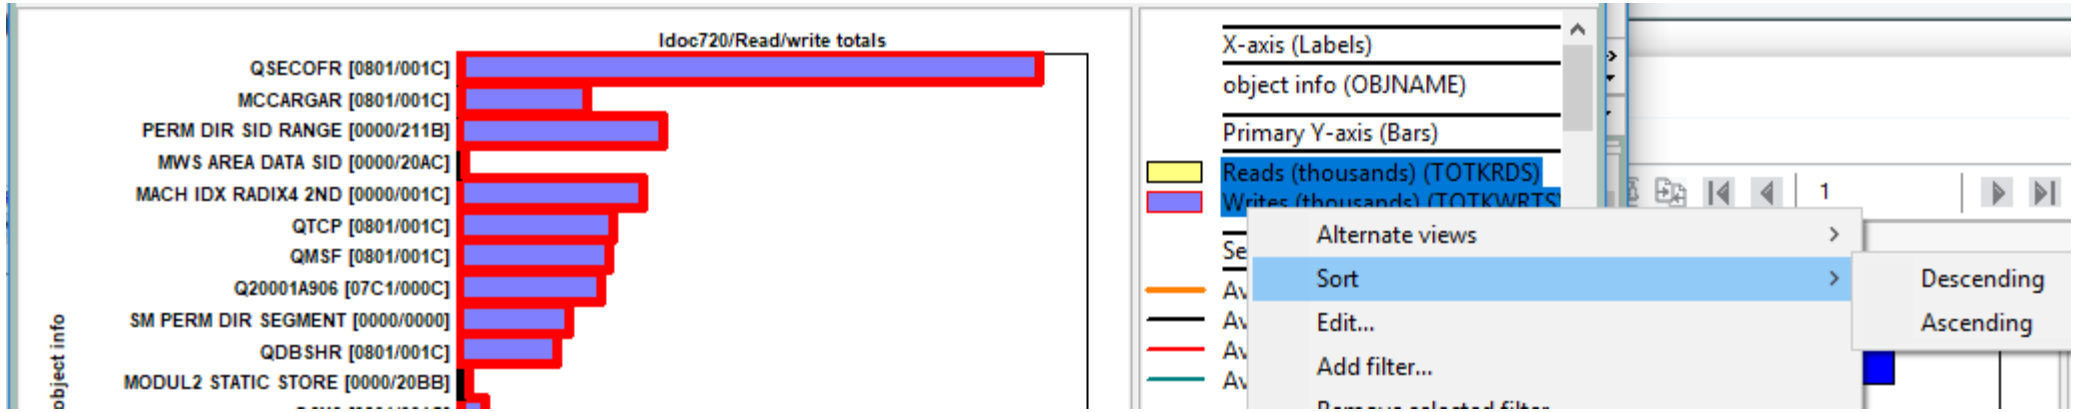

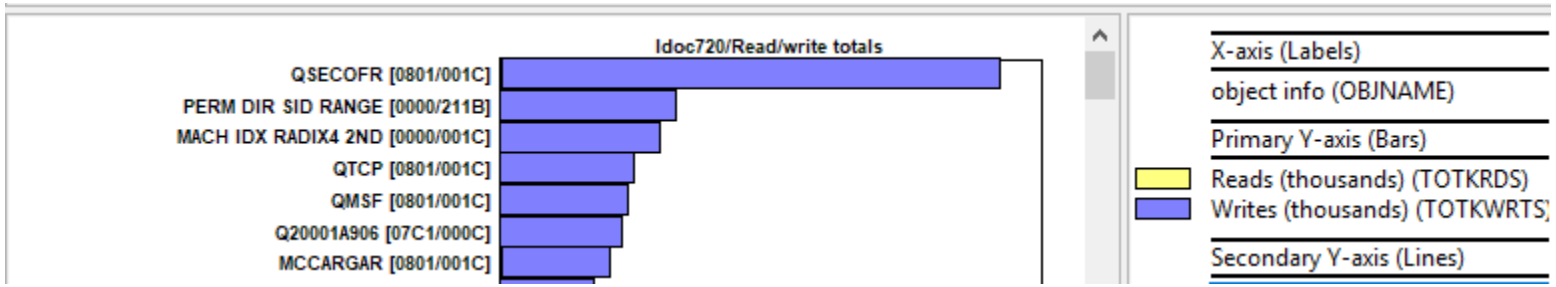

Oct 2018 (1333) – Toggle graph format button

- Button on toolbar can now be used to quickly convert ranking graphs from vertical to horizontal bars and vice versa.

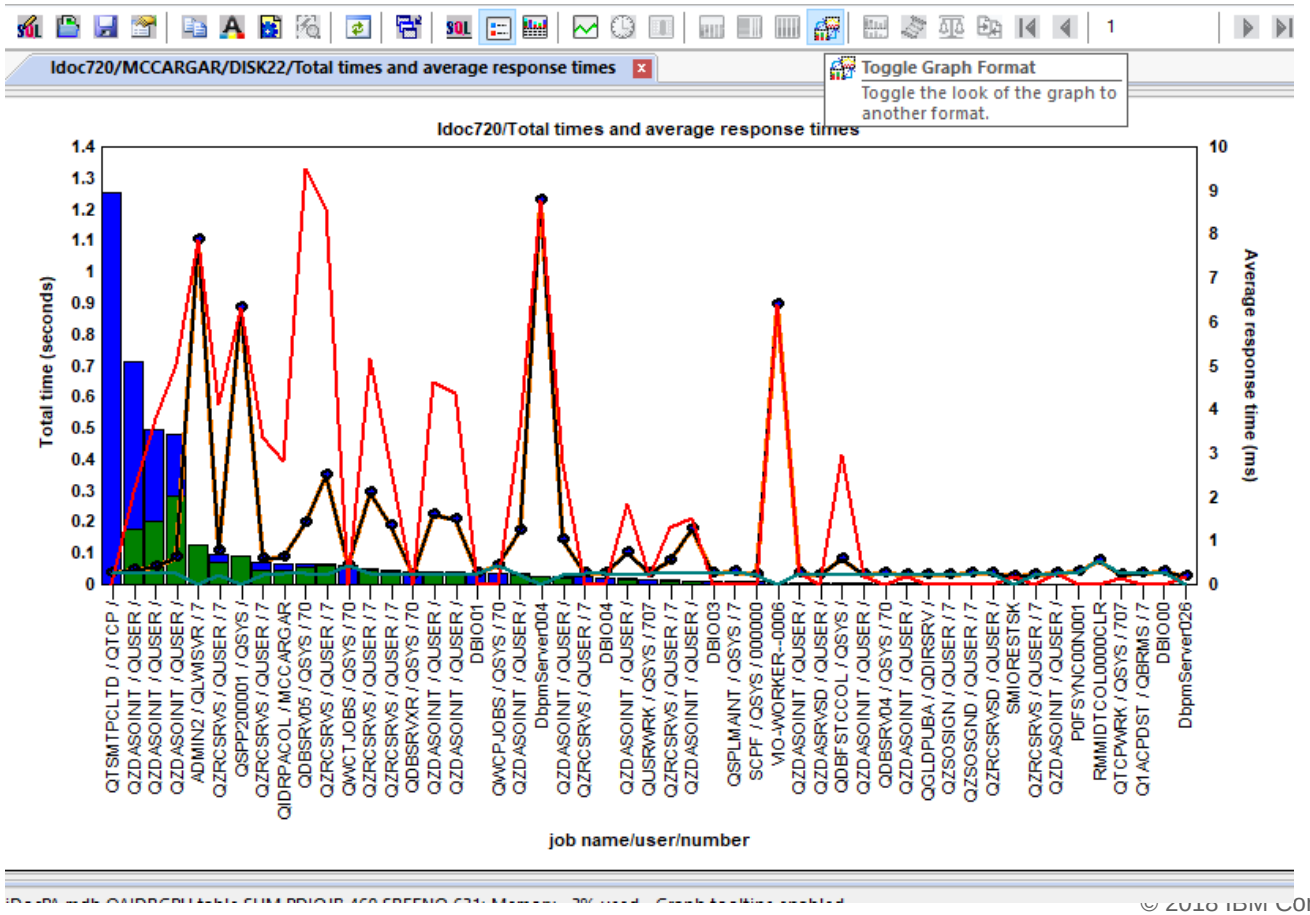

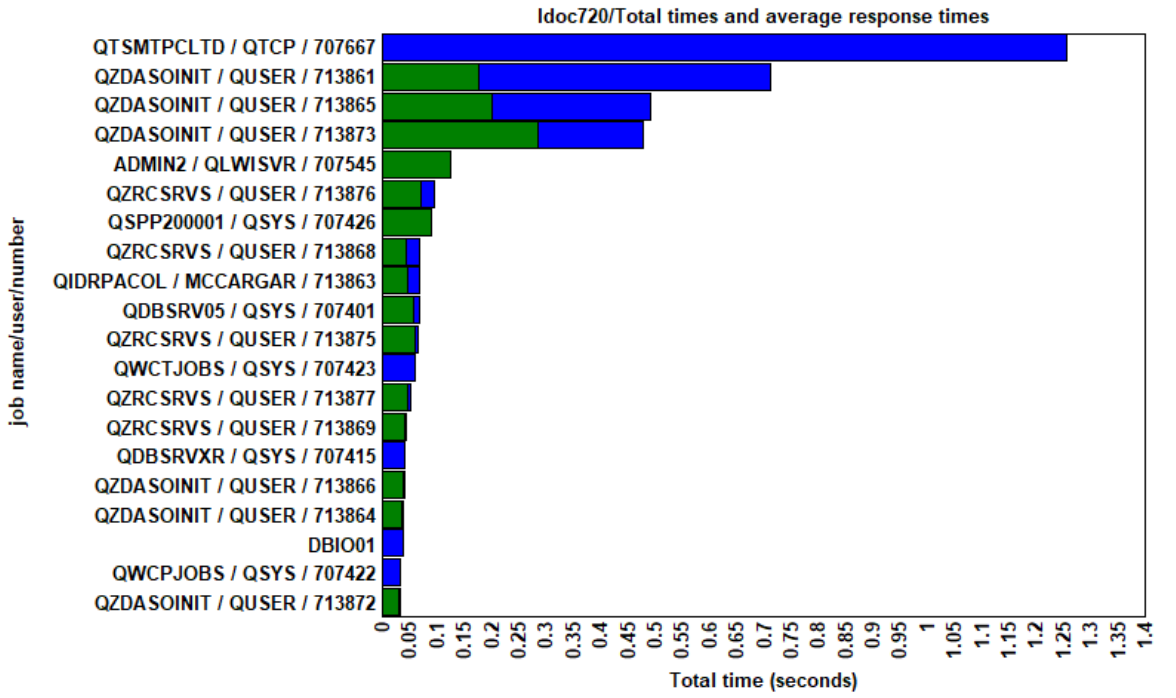

Oct 2018 (1333) – Clock icon use did not change SQL statement

- Using the clock icon on a graph to change the time grouping and then viewing the SQL statement, would not show the correct/updated SQL statement and has been fixed.

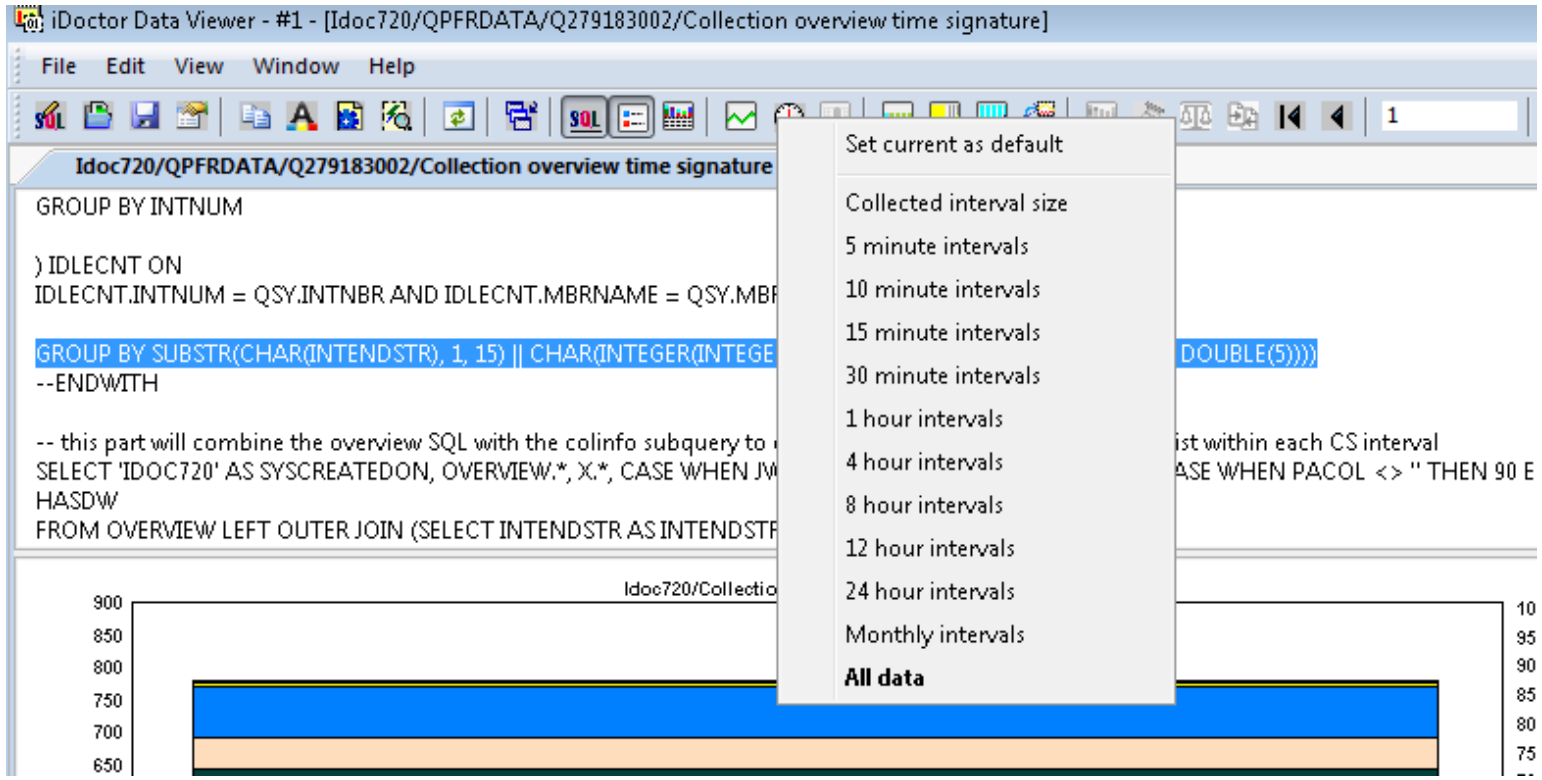

Oct 2018 (1333) – CSI default drill down not showing up

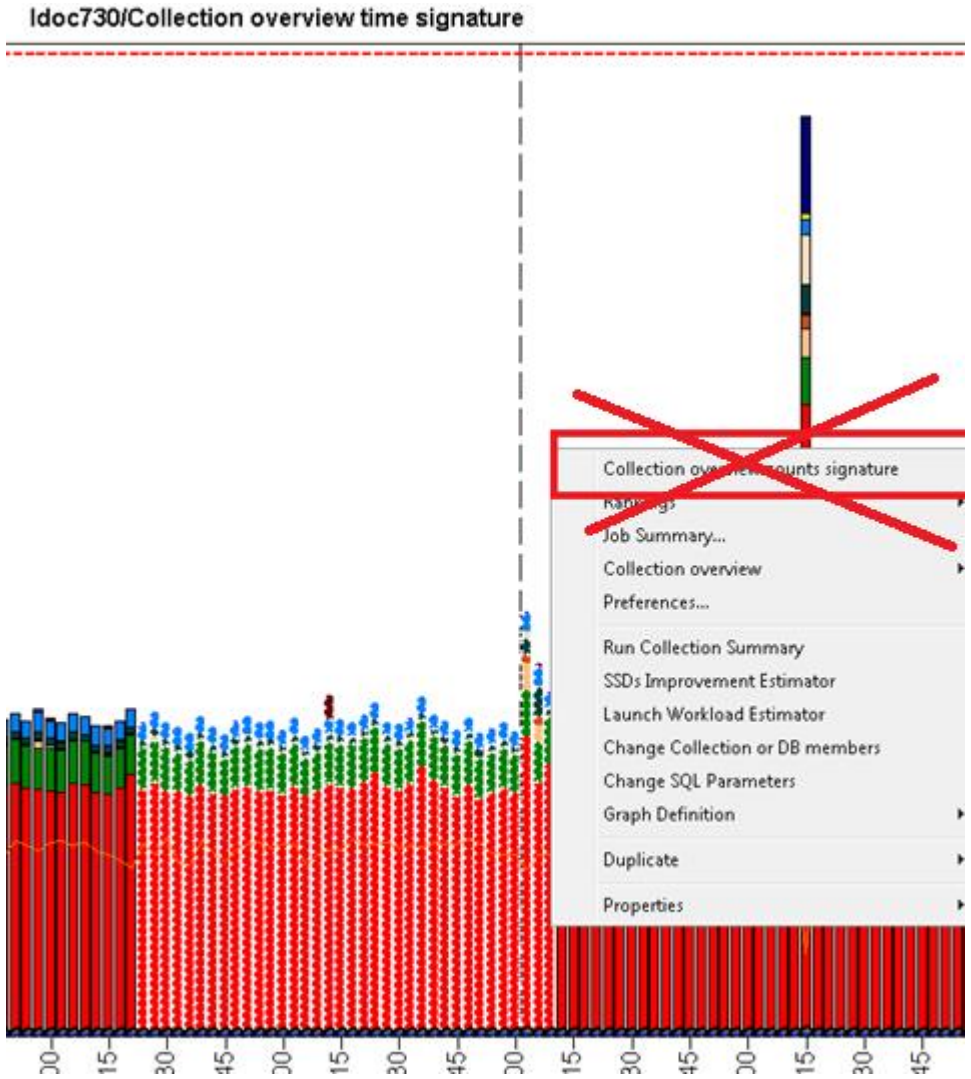

- When drilling down from the Collection overview time signature graph (with 1332), fixed an issue where the default drill down option would appear as the "Collection overview counts signature" graph instead of the Dispatched CPU wait bucket rankings by thread graph.

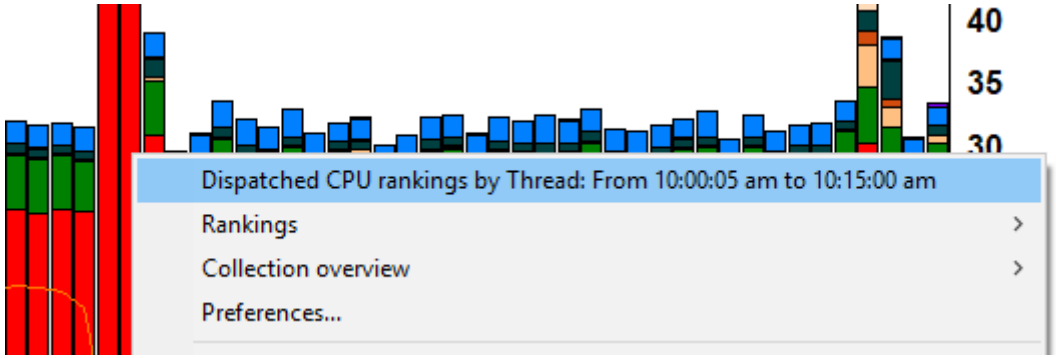

Oct 2018 (1333) – Time groupings will now show YYYY-MM-DD format

- In all iDoctor time-based graphs that support use of the clock icon, the following 3 groupings will now always show the year month and day in the label on the X-axis. (YYYY-MM/DD format)

  24 hour intervals
  Monthly intervals
  All data

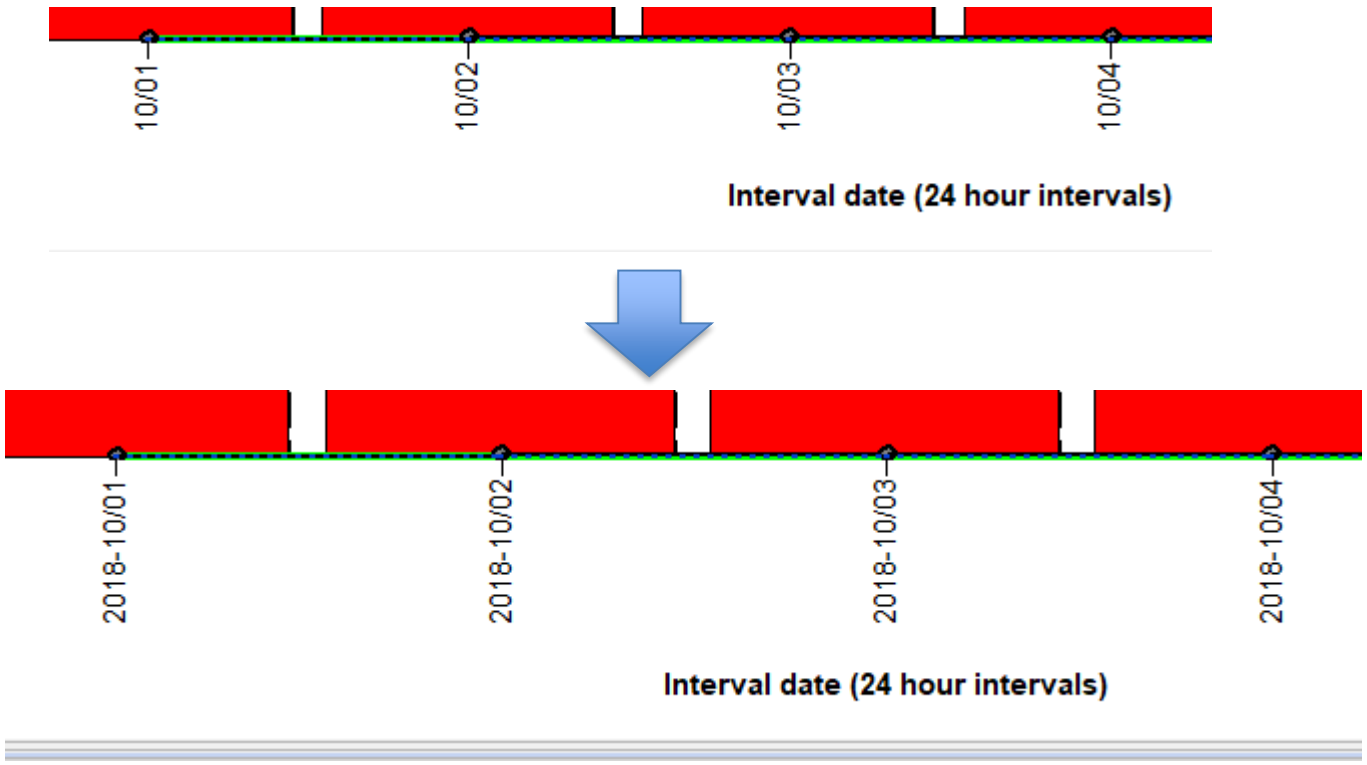

Oct 2018 (1333) – On IBM i the nmon import procedure no longer worked

- The nmon import stored procedure on IBM i no longer worked on some systems due to apparent changes to SQL syntax and use of SQL reserved words.

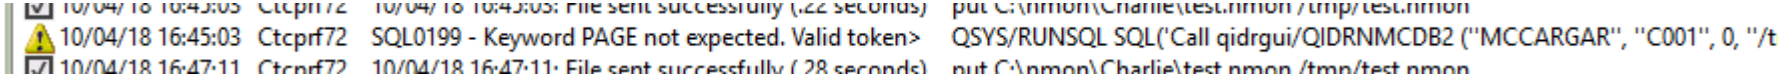

- Also corrected an issue where the GUI would incorrectly load the 6.1 version instead of the 7.1 version of the stored procedure on 7.1+ systems.

Oct 2018 (1333) – QMGTOOLS install change

- When installing QMGTOOLS, updated the installation so the QMGTOOLS library SAVF is transferred BEFORE clearing the QMGTOOLS library, just in case it fails to FTP.

Oct 2018 (1333) – Y2 field changes in JW graphs

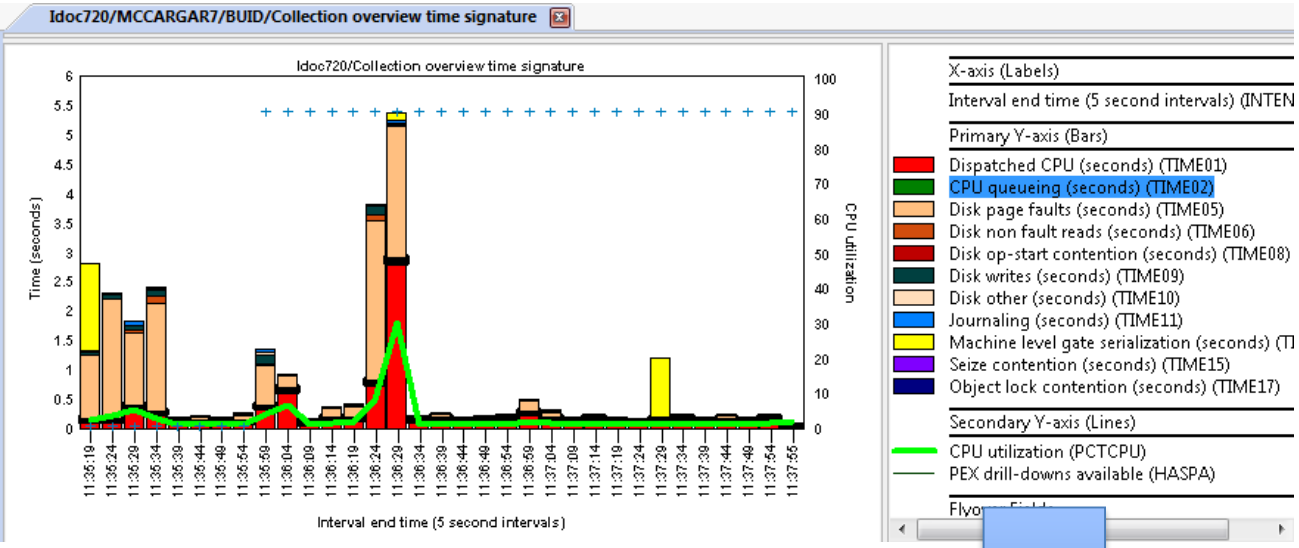

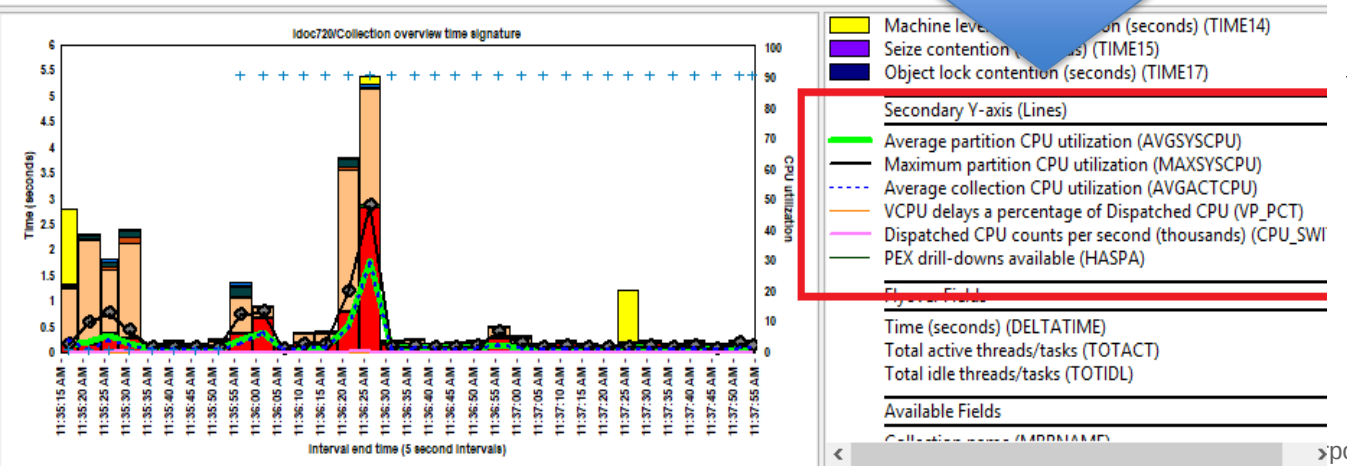

Added extra fields to the Y2 axis whether or not collection summary analysis done.

AVGSYSCPU - average system CPU utilization
MAXSYSCPU - maximum interval average system CPU utilization
AVGACTCPU - average collected CPU utilzation
VP_PCT - VCPU delays as a percentage of dispatched CPU
CPU_SWITCH - Dispatched CPU counts per second (thousands)

Oct 2018 (1333) – Added Virtual CPU delay graph to JW

- At 7.1+ added the Virtual CPU delays graph to Job Watcher in the Wait graphs folder.

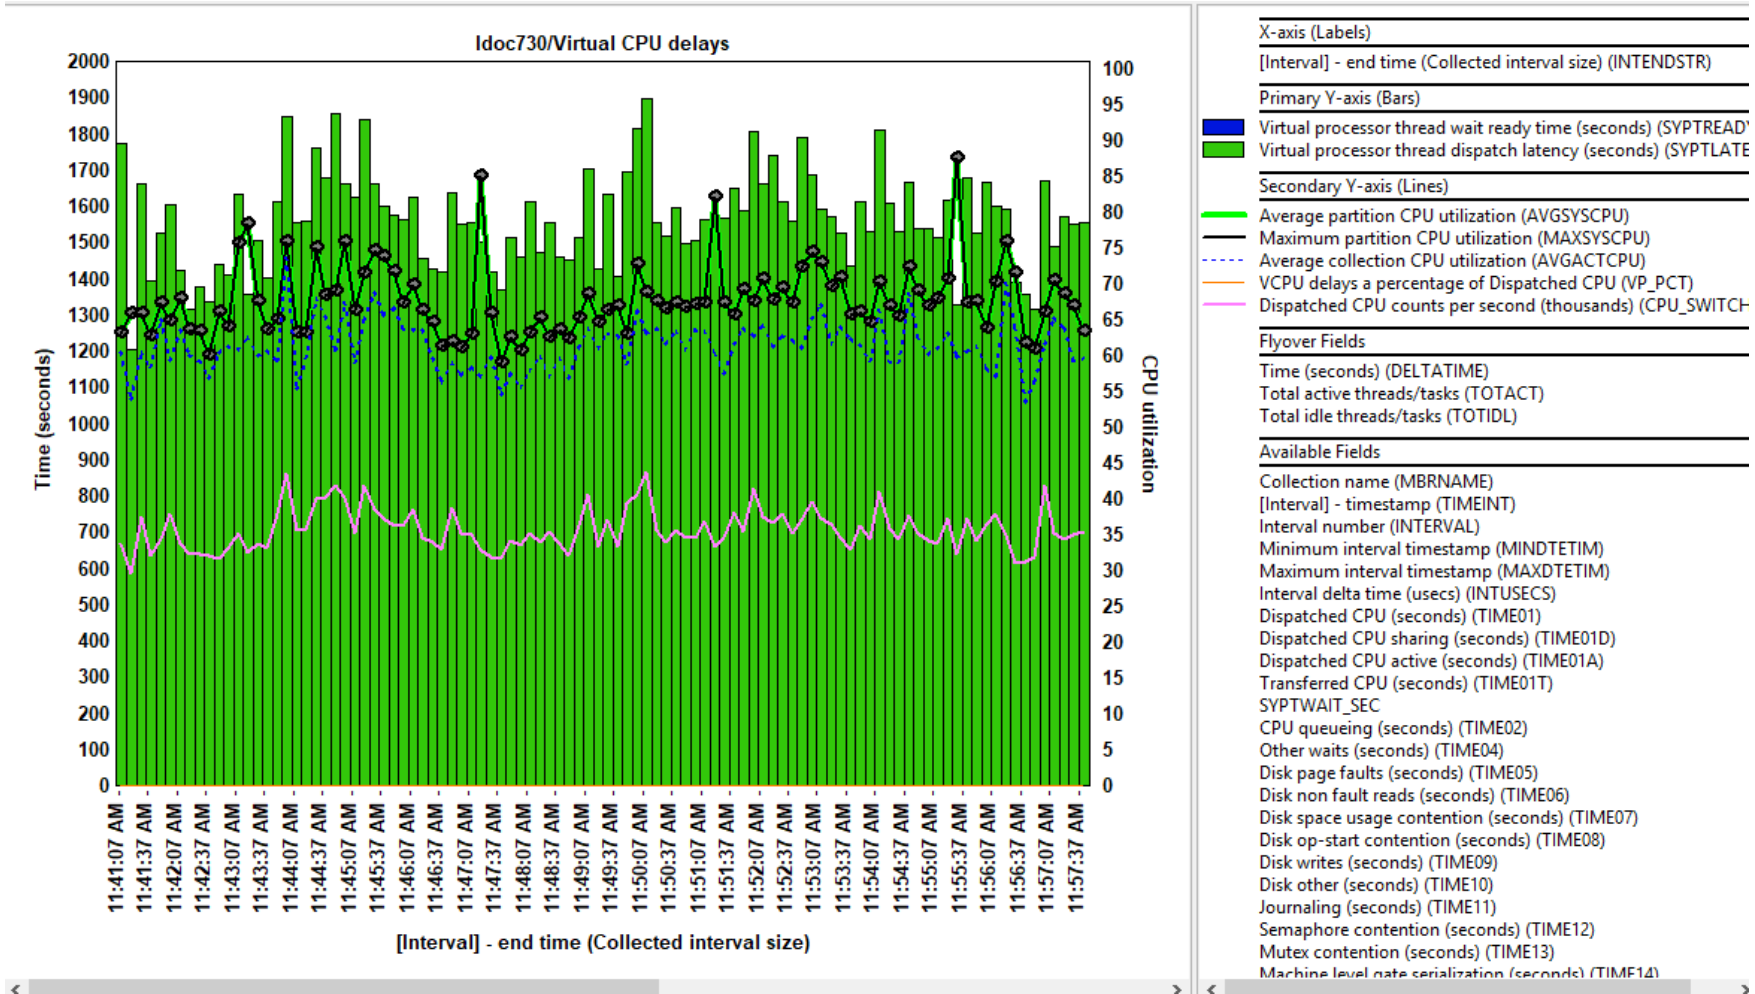

Oct 2018 (1333) – New google search options

- Added support to do a google search from the following locations in iDoctor:

  1. From the graph legend by right-clicking on a field description (and short name).

  2. From the table below a graph by right-clicking on a column header. 2 options are given: search on the the field description (and short name) or search on the value last selected in the table in the curent column.

  3. From any table in the Data Viewer you can right-click, any column header or cell value and do a search in the same ways as #2 above.

Oct 2018 (1332) – Find option – SQL tab

- In the SQL tab of a report definition's properties (either from graph search or user-defined reports), a new Find (text) option has been added to let you find something of interest in the SQL statement.

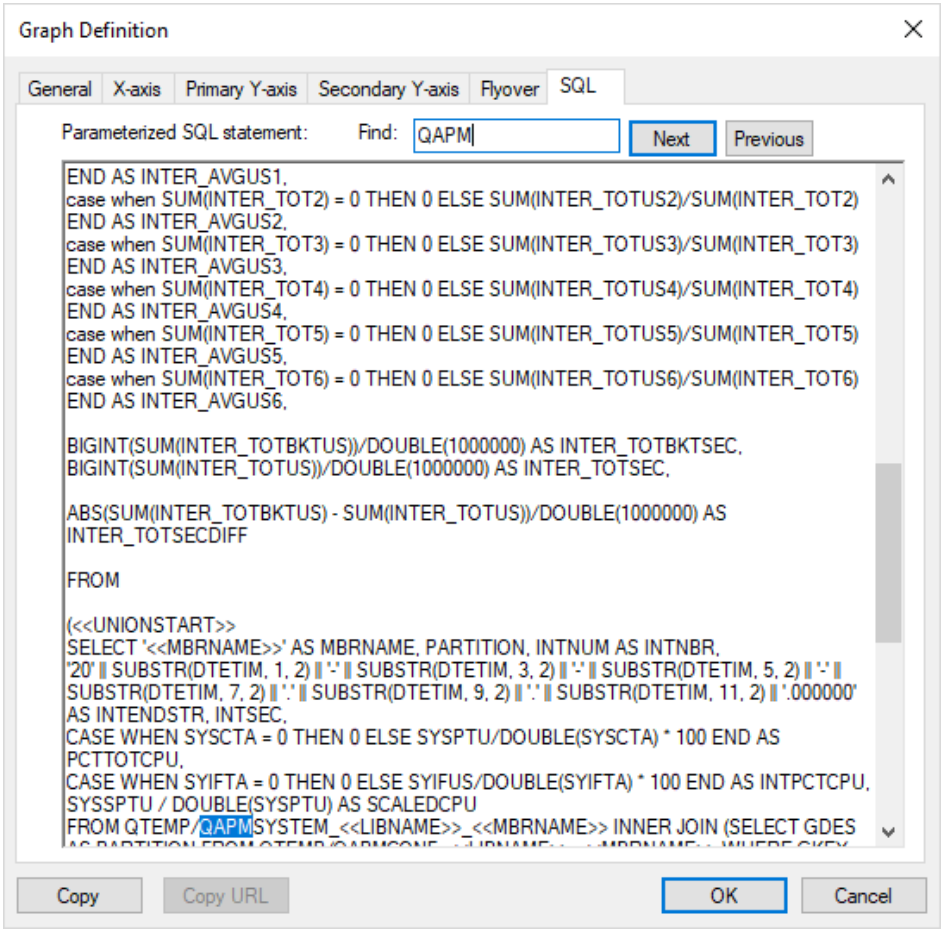

Oct 2018 (1332) – Graph search pane fixes

- Opening, closing and reopening the graph search pane would crash the GUI.
- Double-clicking graph search results was not showing the properties correctly.
- Gave the graph search results a proper title for the Pane/Tab including the values used on the search (if any)
- Updated the graph search results for CSI graphs related to TLBIEs, Interrupts and Fragmentation (free space map) to include PTFs needed in the Comments column.

Oct 2018 (1331) – Graph search pane

- Added a graph search function /pane that lets you browse the iDoctor report databases for reports of interest in the IBM i components CSI, Job Watcher and PEX Analyzer.
- This tells you required files, analyses, PTFs and VRMs needed in order for the report to appear and/or show data.

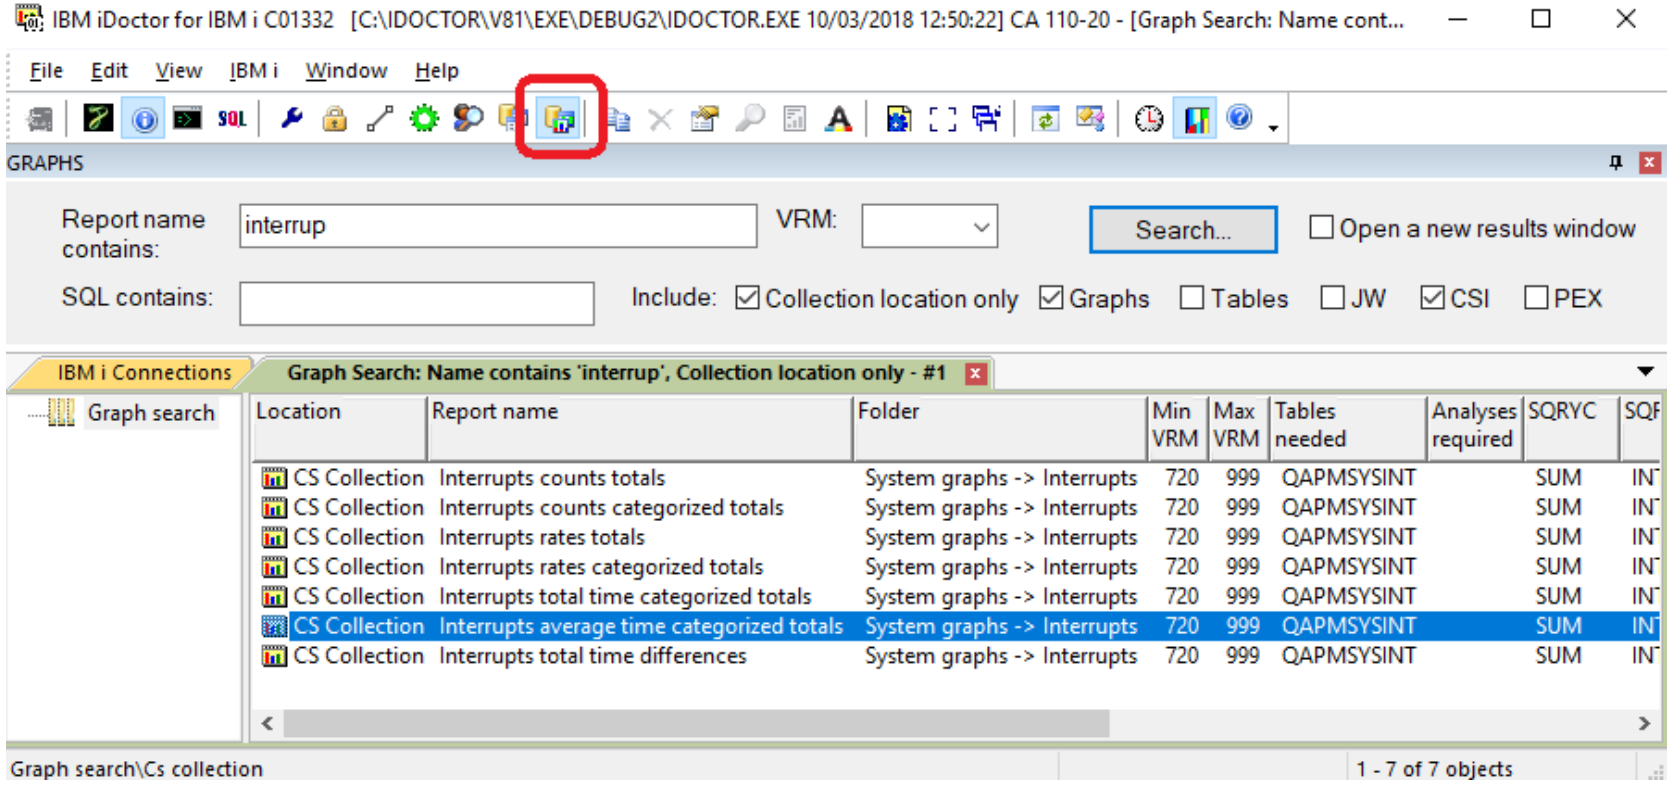

Oct 2018 (1331) – Graph search pane (continued)

- When viewing the properties by double-clicking one of the results you can view the graph settings used as well as the parameterized SQL statement for the report.
  Note: The properties interface is read only.

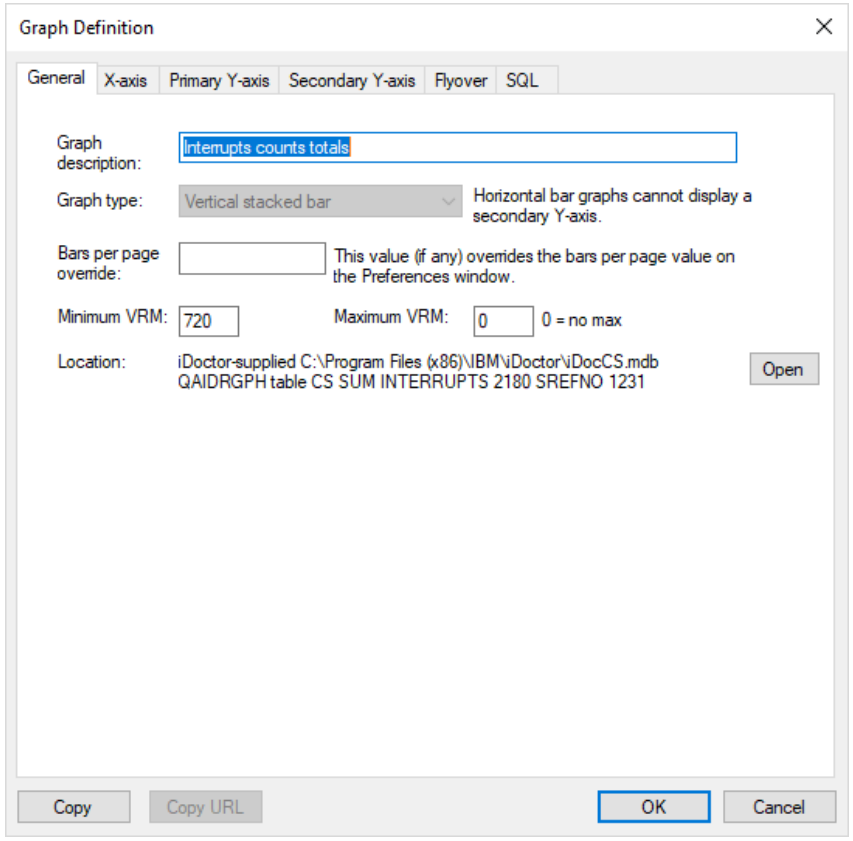

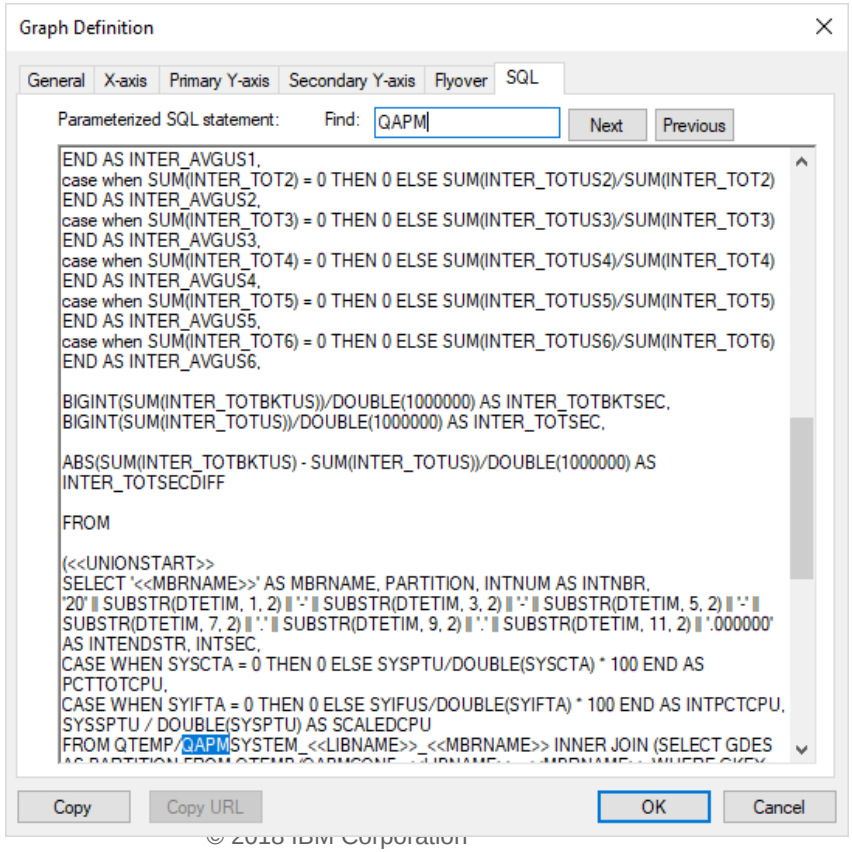

Oct 2018 (1331) – JW - Most drill downs via SQL tables removed.

- In JW, removed graph drill down support previously found when right-clicking individual SQL tables except ones related to the Job Summary analysis.
  - These did not work in most cases and in some cases did not make sense to appear. This includes subfolders like "Collection Summary, Call Stack Summary, Long Transactions", etc.

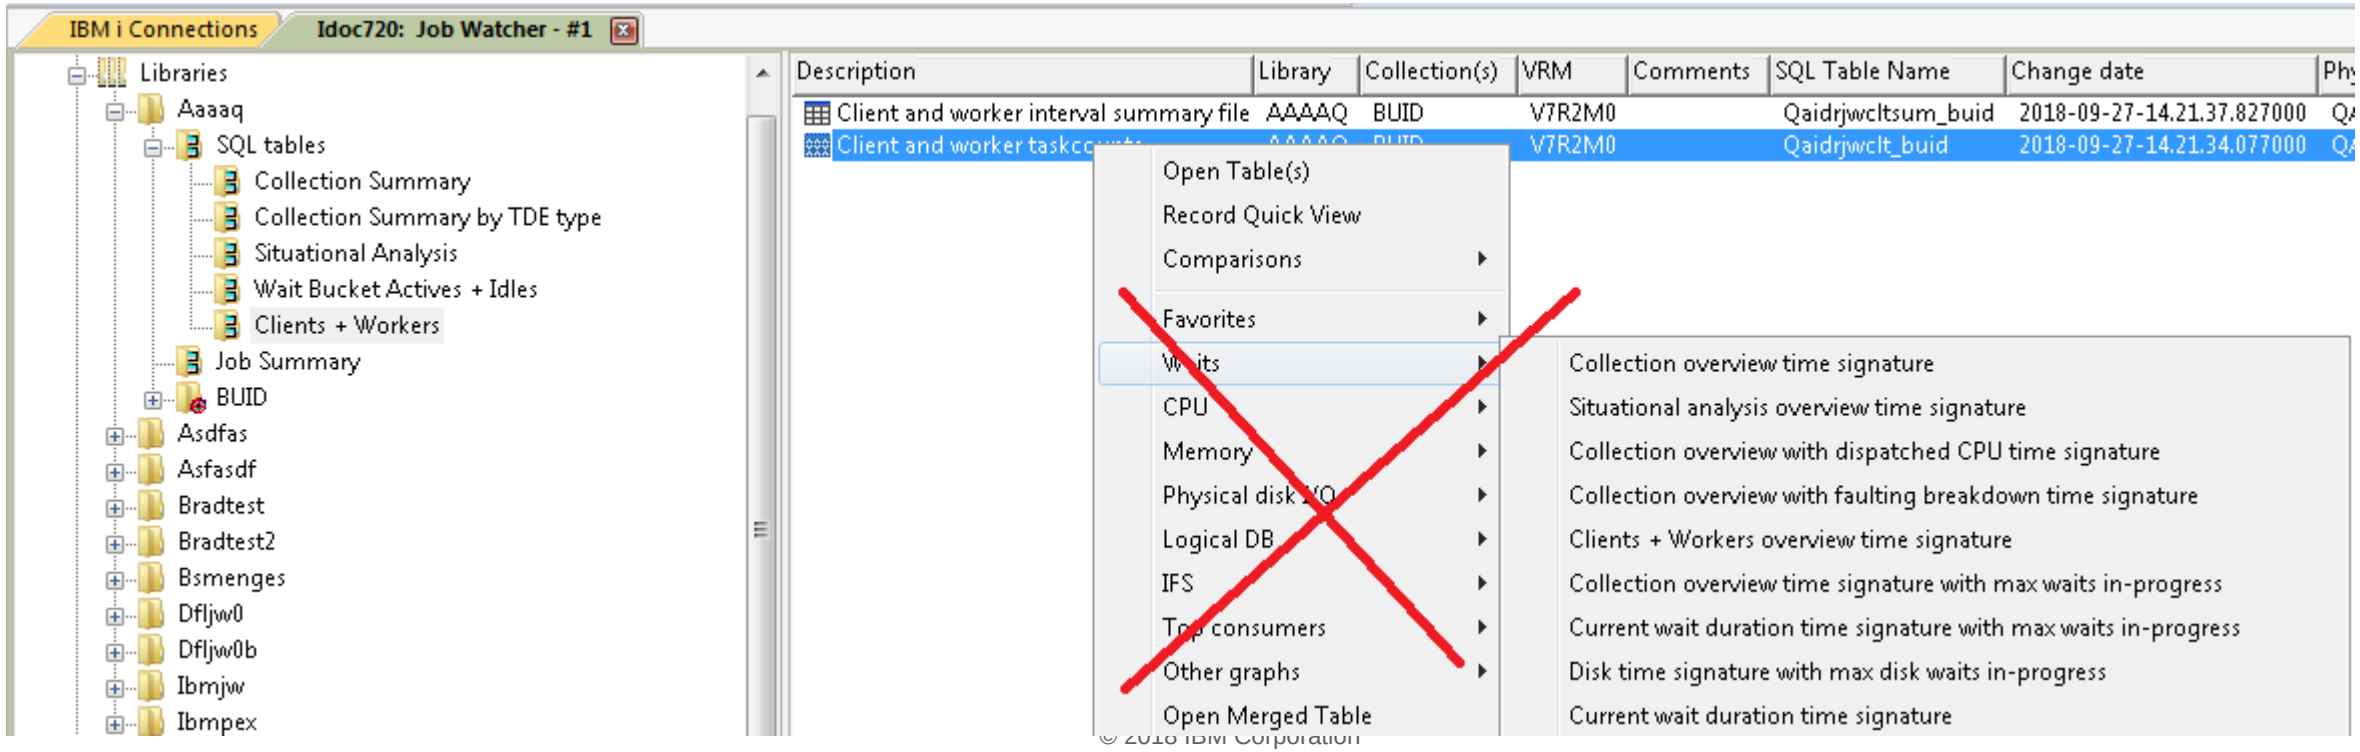

Oct 2018 (1331) – NMON – SQLite – CPU utilization graph

- Corrected SQL errors with the nmon graph CPU utilization for <CPU #N> when using SQLite as the analysis database.

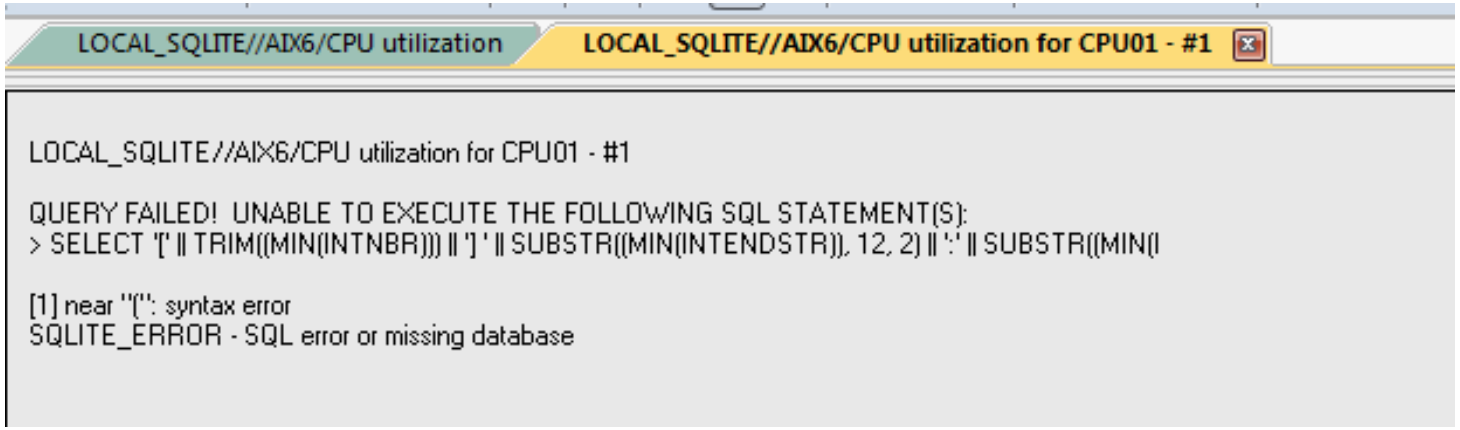

Oct 2018 (1331) – NMON – SQLite – Delete collection issues

- Fixed an issue when using nmon and SQLite as the analysis database. Deleting collections created using the Merge option would not delete all files in the collection. This causes the user to be unable to create new collections (using the merge option) with the same name.

- When deleting nmon collections created using SQLite as the analysis database, improved the error handling and process in the remote command status view.

  If any file cannot be deleted this will be listed in the results column (and history log). Previously the GUI would indicate the delete was successful when certain files may not have been deleted.

Oct 2018 (1331) – JW – Holder chase broken at 7.4

- In JW, the Holder chase drill down sql syntax no longer worked at the next ibm i release due to changes to SQL syntax. Made a work around to handle this.

Oct 2018 (1331) – CSI – CPU utilization with SMT Context

- Fixes for the CPU utilization with SMT Context graph

  1. At 7.1 the calculation for SMT Context was likely incorrect.

  2. A 7.2 the extra fields included in the table for total virtual processors and active virtual processors previously would not be able to exceed 32 and has been fixed so the true values are shown.

  3. Added total virtual processors and active virtual processors to the graph flyover.

Sept 2018 (1330) – CSI – CPU utilization for <XYZ> graphs

- Fixed the SQL statement for the CSI drill down graphs CPU utilization for <XYZ>. When graphing multiple collections the results were incorrect.

Sept 2018 (1330) – PEX PDIO By objects (summarized) folder back

- By request, the By object (summarized) folder is back and the 6 graphs within it.

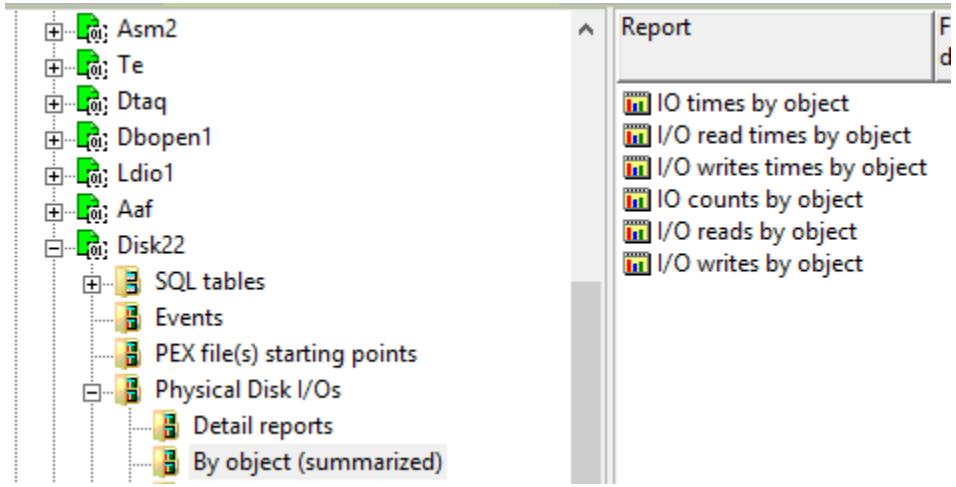

Sept 2018 (1330) – PEX PDIO sort option fixed

- In Preferences -> PEX the Sort PDIO ranking graphs option was broken and has been updated to add several new options (or type in your own choice will work too.)

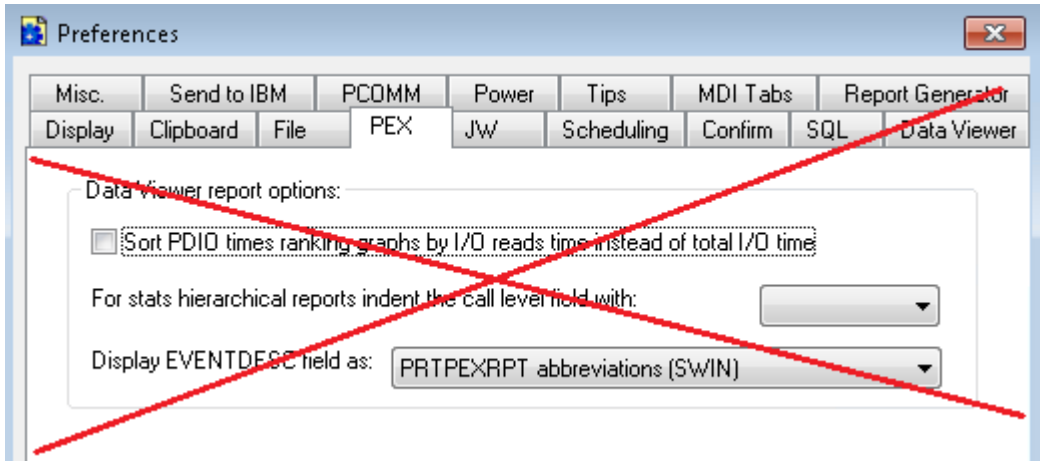

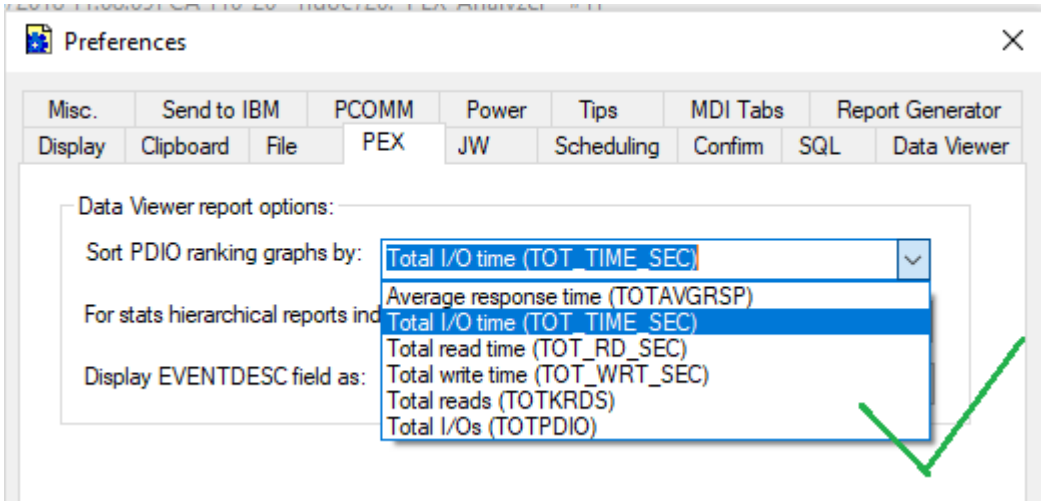

Sept 2018 (1329) – JW – Call stack – Program details option

- In JW the call stack window now has an option called “Program details” in the drop down list that lets you see many additional details about the programs found in the call stack. **Note:** This will only work if the “Retrieve program/module details” analysis has been ran on at least one collection in the library.

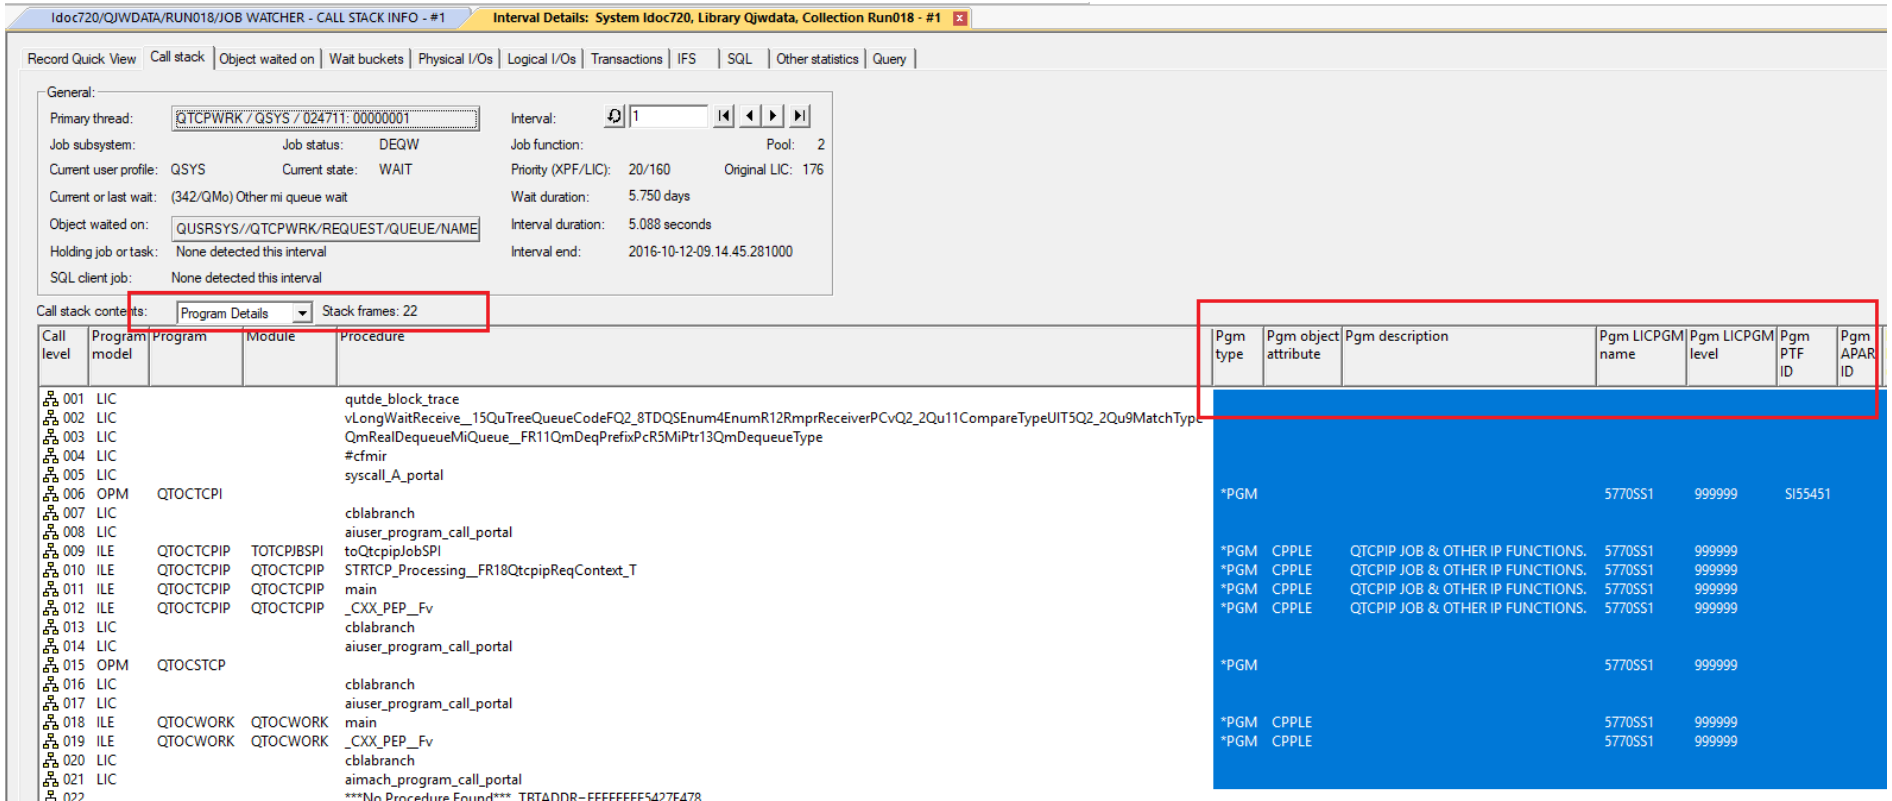

Sept 2018 (1329) – JW – Retrieve program/module details analysis

- In JW renamed the analysis "Retrieve program/module information" to "Retrieve program/module details“. This was done to provide consistency with the new “Program details” option on the JW call stack.

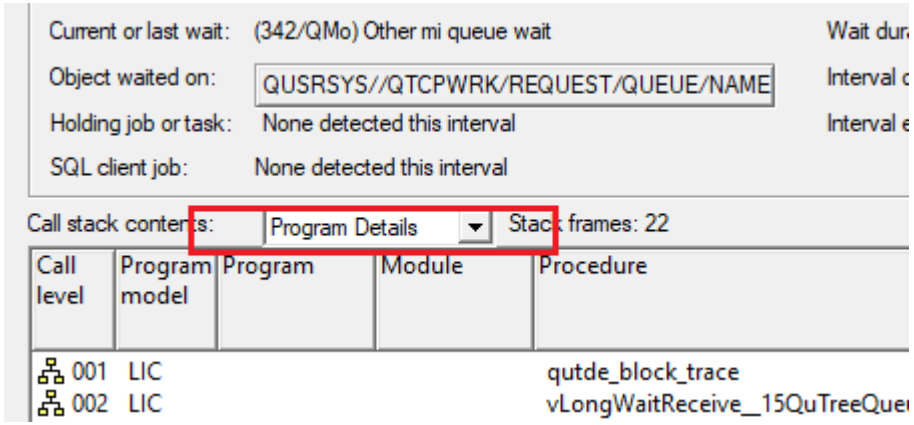

Sept 2018 (1329) – JW – New call stack reports options

- When drilling down from a call stack on a single entry, added two new options:
  - Program details
  - Module details

Note: This provides additional information about the program or module selected in the call stack.

Sept 2018 (1329) – Power Connections – Create Disk Mapping Check

- When running the Create Disk Mapping function the check for OpenSSH has changed so it should now work with any version installed on the IBM i. Previously it would fail if the version of OpenSSH installed did not match the version of IBM i.

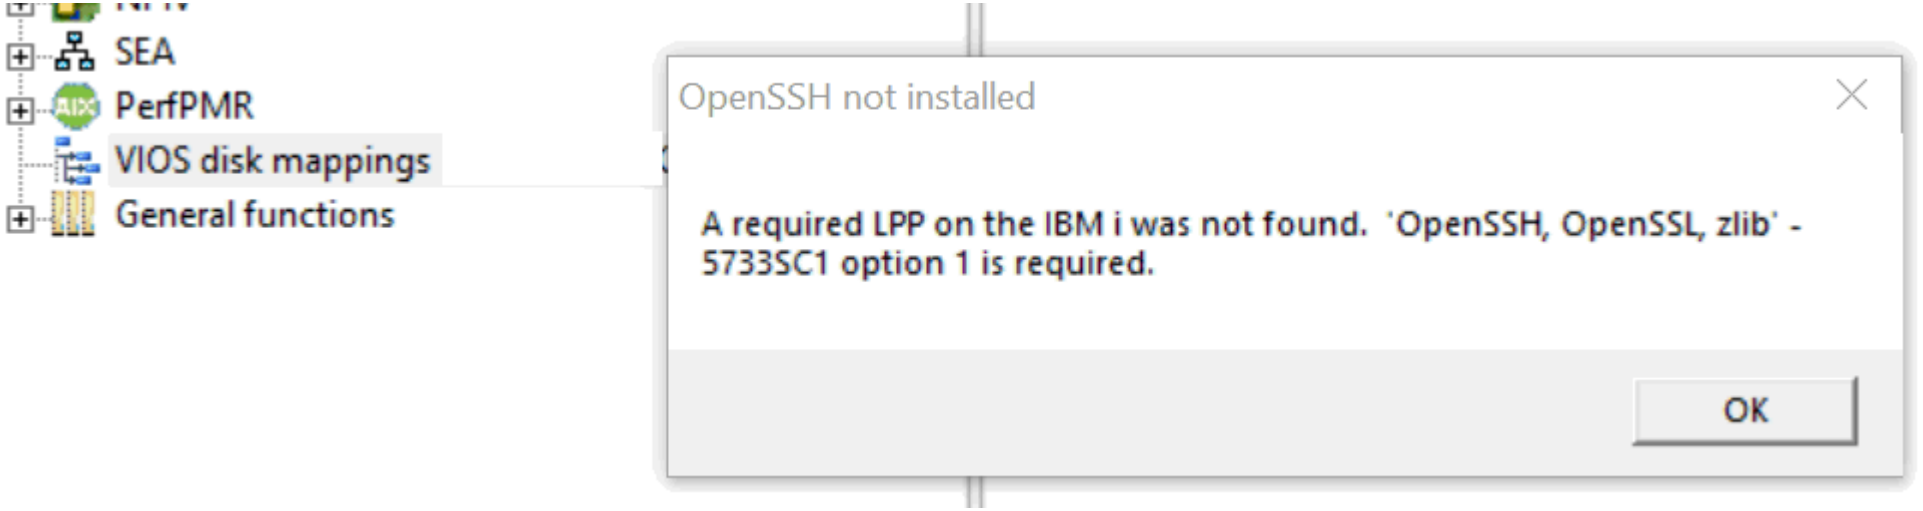

Sept 2018 (1329) – Power Connections – Create Disk Mapping Authentication Failure

- Fixed a timeout issue with the SSH connection used by MkAuthKeys.jar. This is part of the Create Disk Mapping function and would cause the user to think their user/password was incorrect when this wasn't true. The log provided clues to the true reason for the failure.

```
java.net.SocketTimeoutException: Read timed out
java.net.SocketTimeoutException: Read timed out
        at java.net.SocketInputStream.socketRead0(Native Method)
        at java.net.SocketInputStream.socketRead(Unknown Source)
        at java.net.SocketInputStream.read(Unknown Source)
        at java.net.SocketInputStream.read(Unknown Source)
        at trans.Transport.readPacket(Transport.java:290)
        at trans.Transport.read(Transport.java:154)
        at trans.Transport.read(Transport.java:134)
        at ssh.Connection.authenticateUser(Connection.java:606)
        at ssh.Connection.<init>(Connection.java:176)
```

Sept 2018 (1328) – CSI – Firmware truncated in Collection properties

- In CSI the firmware level in the System tab was truncated to 7 characters when it should have been 9.

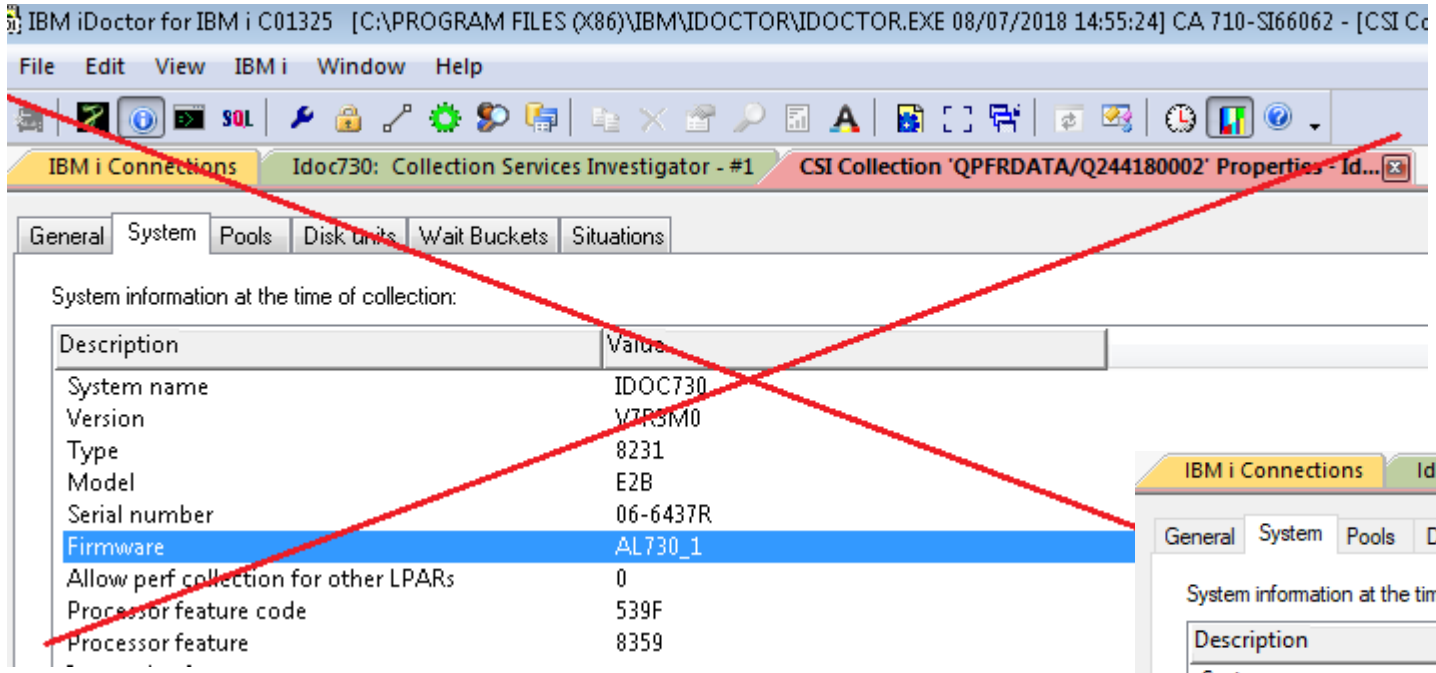

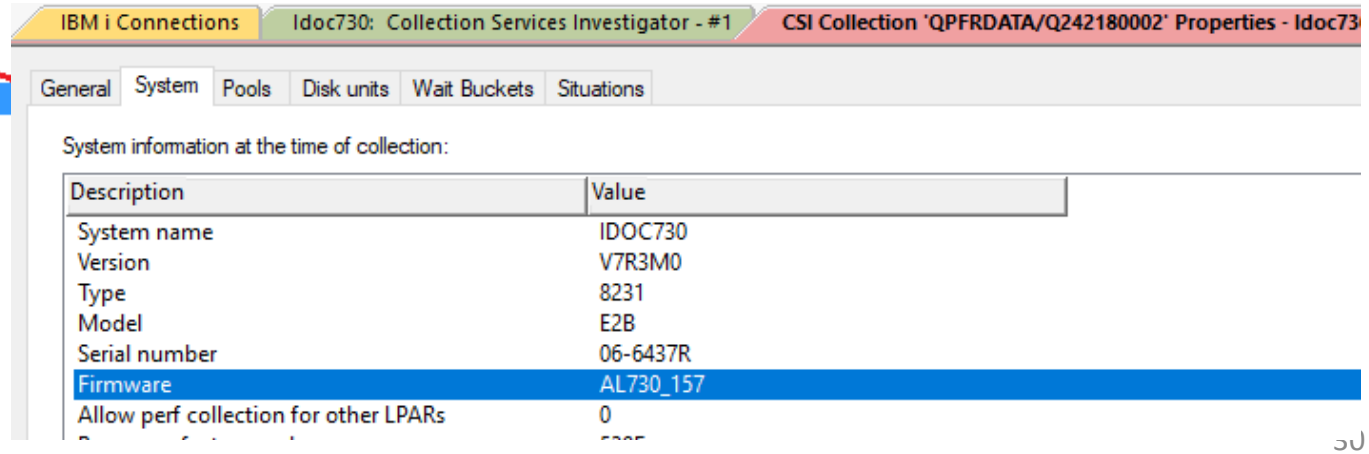

August 2018 (1327) – CSI – Generic job X wait time signature bad data

- In CSI the wait bucket times for the graph Generic job <XYZ> wait time signature were being calculated incorrectly for periods where the jobs were idle and has been fixed.

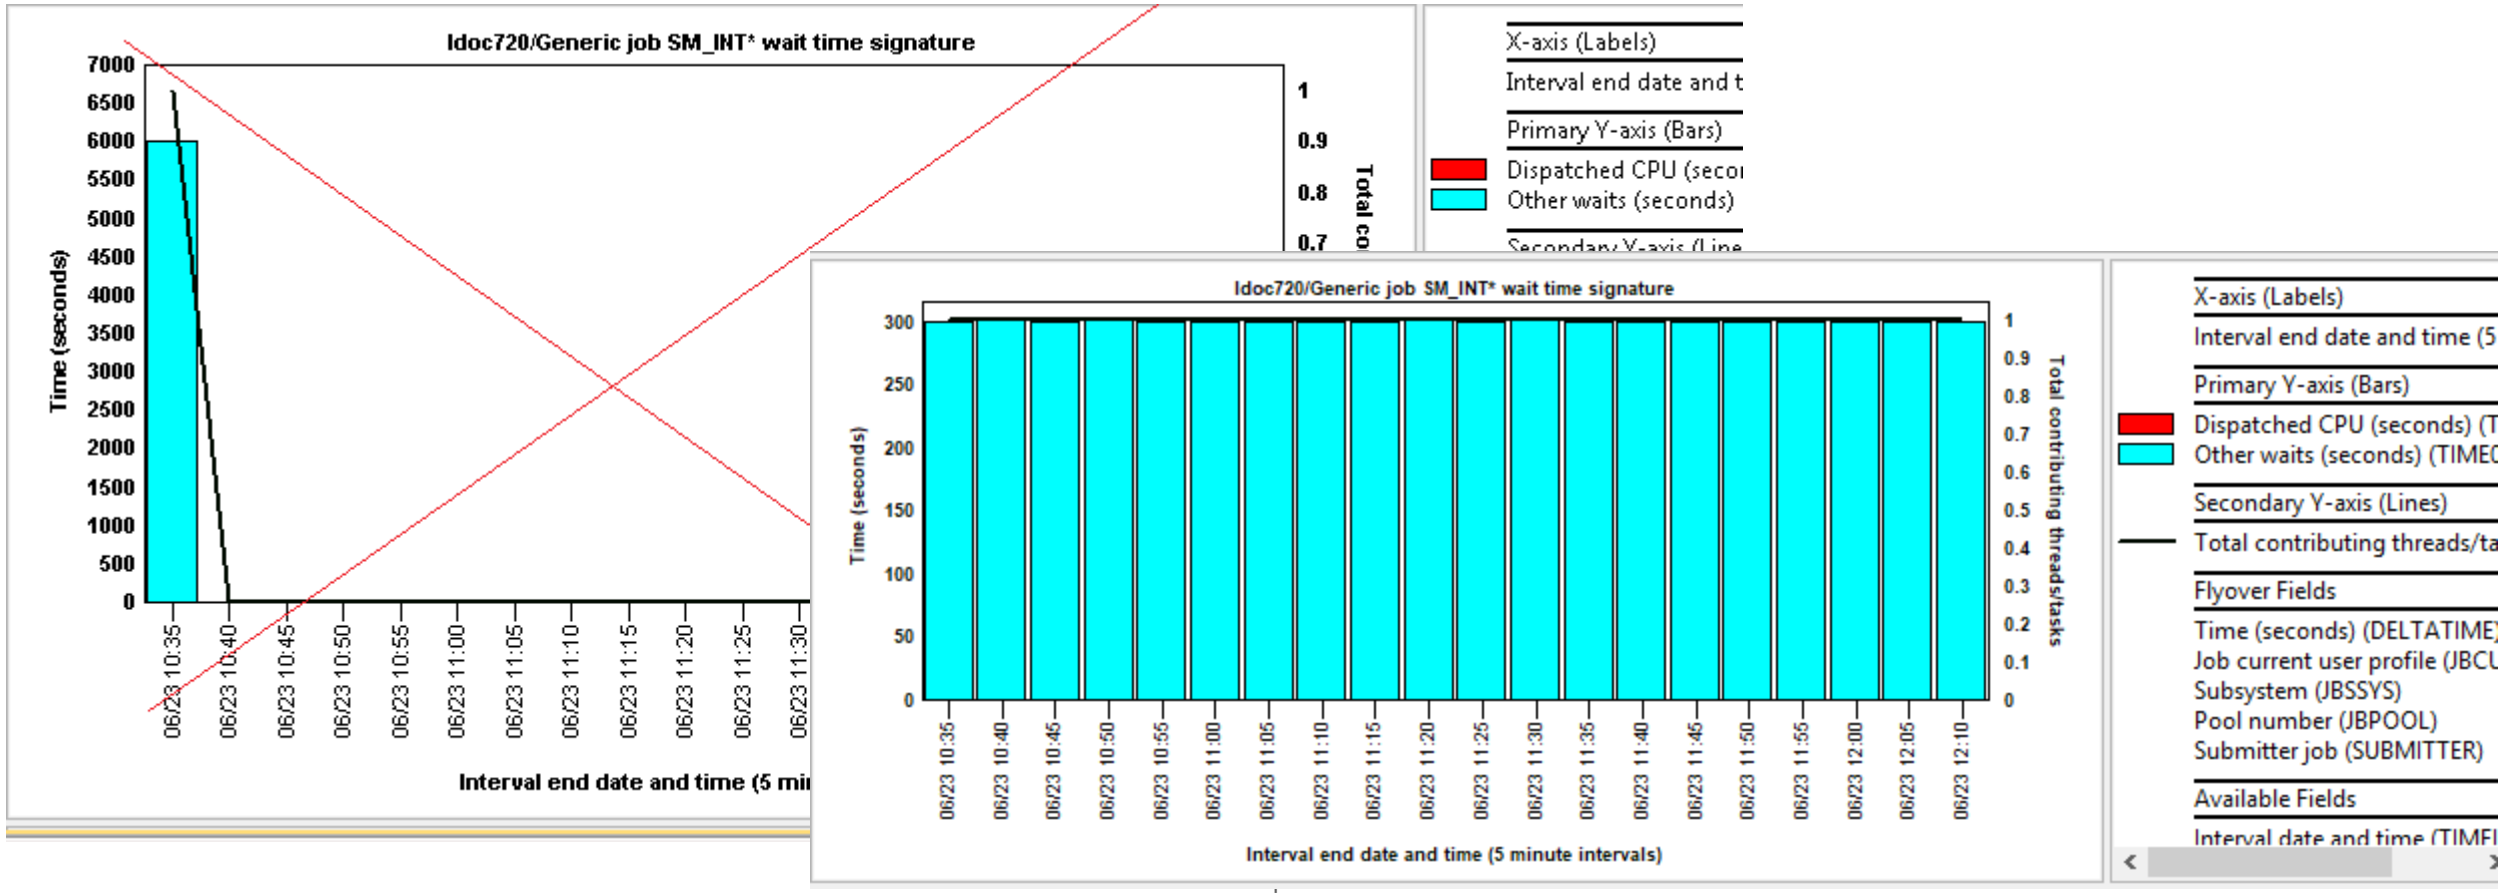

August 2018 (1327) – CSI – Generic job drill downs broken

- The Generic job drill down graphs would fail to produce any data (in some cases) because the previous SQL statement value for field GENJOBNAME did not contain a '*' at the end of the name if the job name length was the same as the generic job name length preference on the Data Viewer tab.

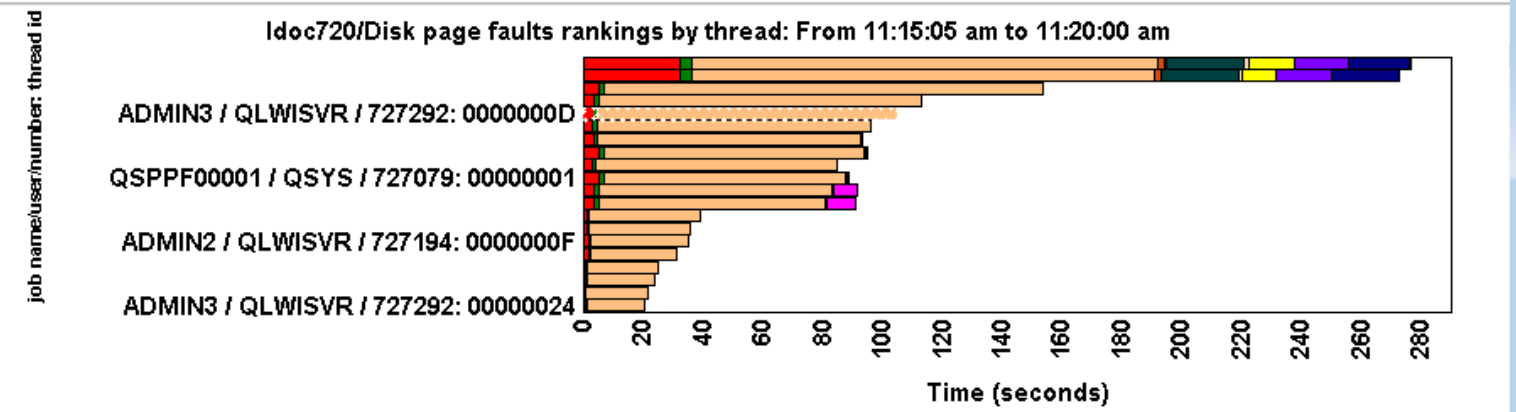

ved nds) 23)	Socket transmits (seconds) (TIME24)	Socket receives (seconds) (TIME25)	Socket other (seconds) (TIME26)	IFS (seconds) (TIME27)	PASE (seconds) (TIME28)	Data queue receives (seconds) (TIME29)	Idle/waiting for work (seconds) (TIME30)	Synchronization token contention (seconds) (TIME31)	Abnormal contention (seconds) (TIME32)	Interval number (INTNUM)	GENJOBNAME
0	0	0	0	0	0	0	0	0	.0336	10	BAC062*
0	0	0	0	0	0	0	0	0	.0135	10	BAC062*
0	0	0	0	0	146.7125	0	0	0	0	10	ADMIN2
0	0	0	0	0	186.9516	0	0	0	0	10	ADMIN5
0	0	0	0	0	194.9752	0	0	0	0	10	ADMIN3
0	0	0	0	0	203.9746	0	0	0	0	10	QSRVMO*
0	0	0	0	0	207.2050	0	0	0	0	10	ADMIN2

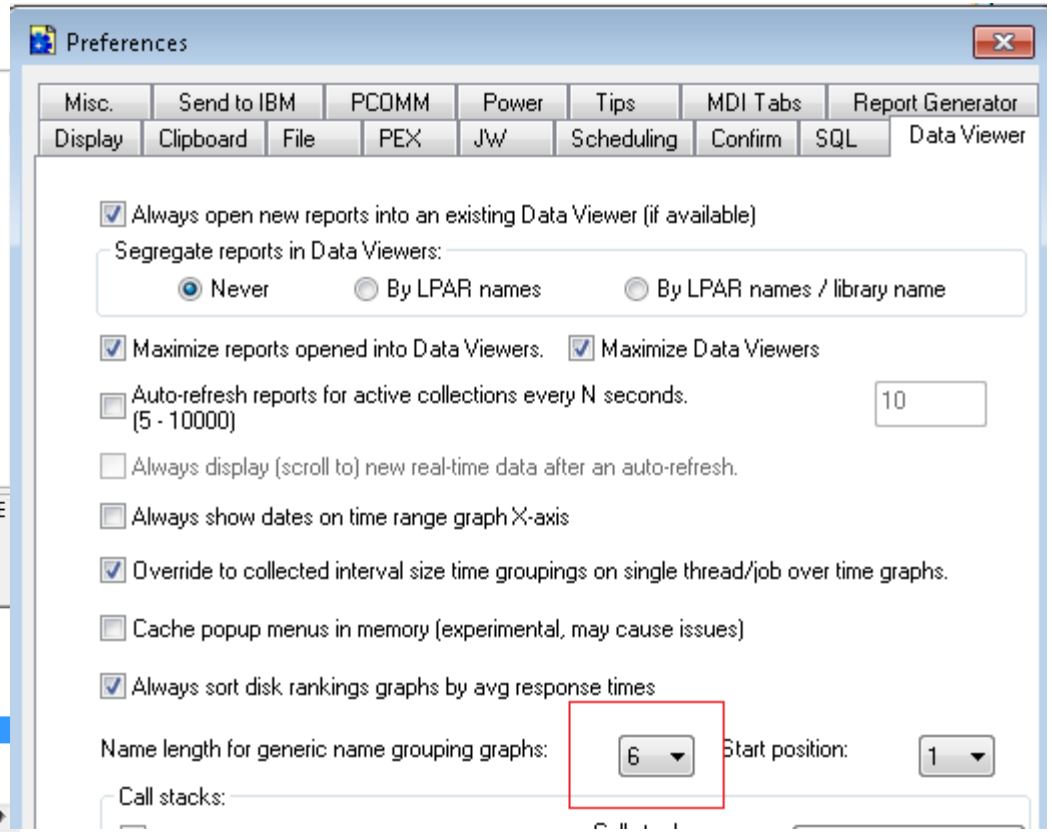

August 2018 (1327) – CSI – Generic job drill downs broken (page 2)

- The Generic job drill down graphs would fail to produce any data (in some cases) because the previous SQL statement value for field GENJOBNAME did not contain a '*' at the end of the name if the job name length was the same as the generic job name length preference on the Data Viewer tab.

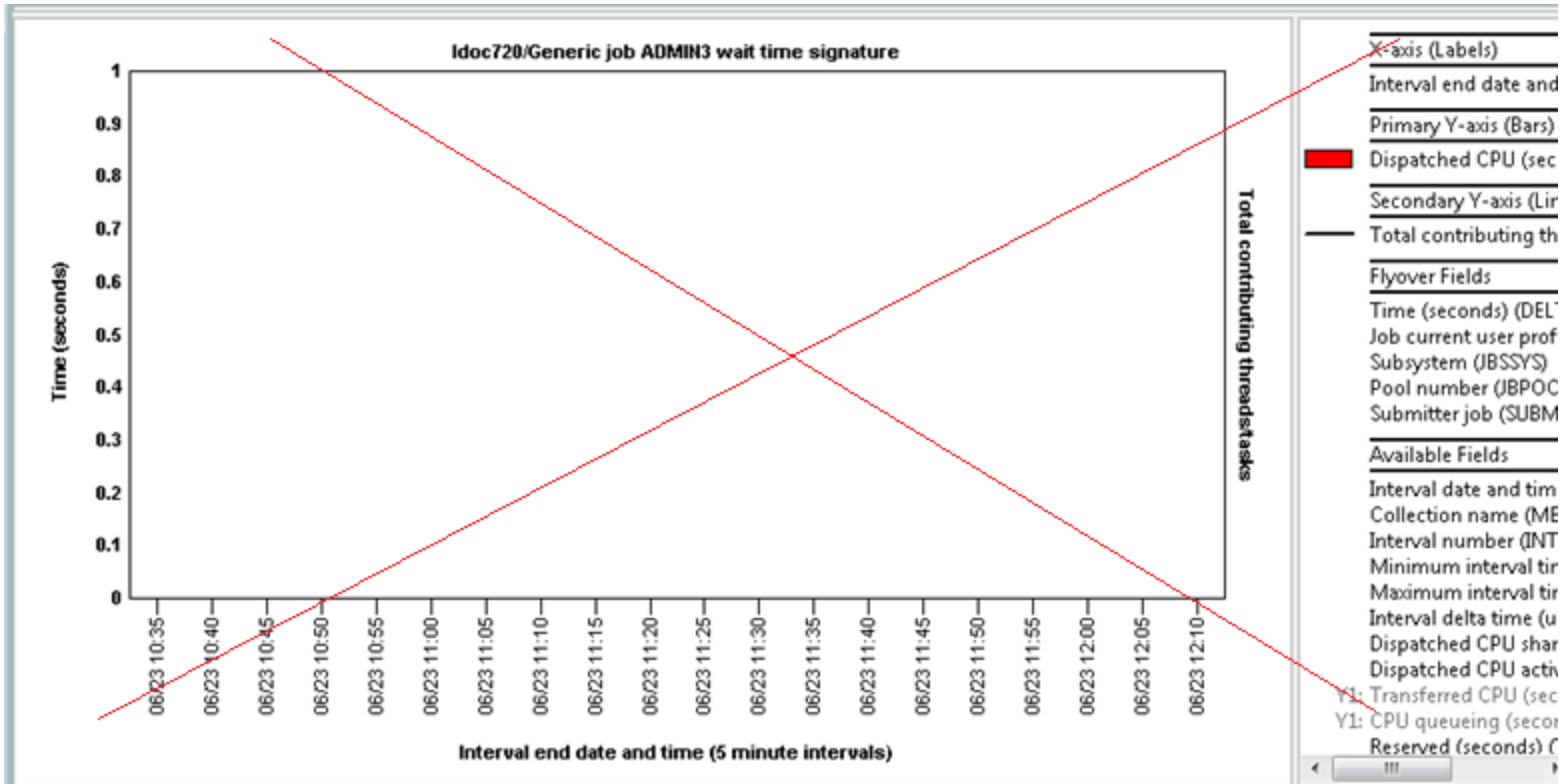

August 2018 (1327) – CSI – Underlines in job names breaking drill downs

- In CSI the generic job for selected thread drill down graphs would return incorrect data if the generic job name selected contains underlines (such as SM_IN)

  This fix causes the underlines to be treated correctly which have special meaning in SQL LIKE statements

August 2018 (1327) – CSI – Collection Search underline support

- If the search value contains an underline '_' , the search results for job, subsystem or current user did not handle it correctly. More results would be returned than desired.

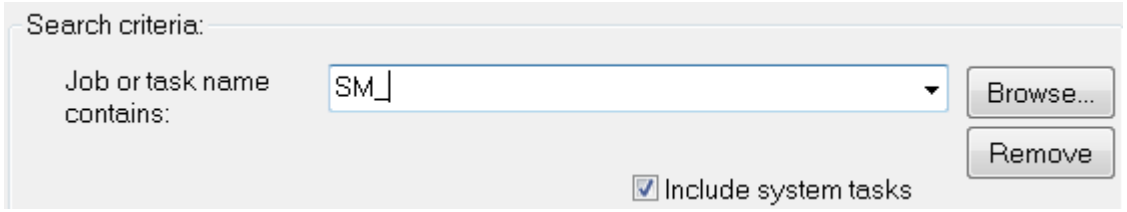

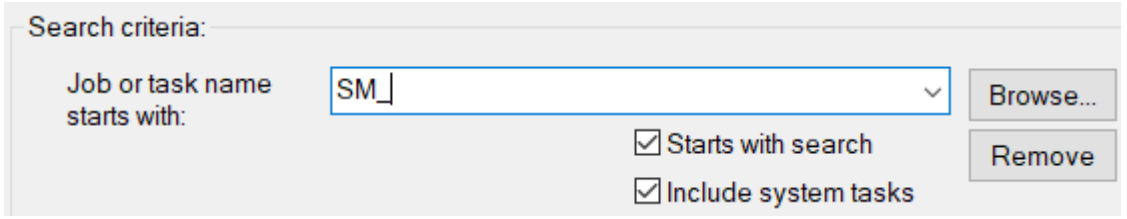

Idoc720/CSLABEX1/Q175102853/Job or task name search - #1

Collection name (MBRNAME)	iDoctor grouping name (OBJNAME)	iDoctor grouping value (OBJVALUE)
Q175102853	AM-ANSMGR	0000000000000
Q175102853	BAC0622AA / BSMENGES / 729238: 0000004B	00000000000Et
Q175102853	BAC0622AA / BSMENGES / 729238: 0000004B	00000000000Et
Q175102853	BAC0622AA / BSMENGES / 729238: 0000004B	00000000000Et
Q175102853	BAC0622AA / BSMENGES / 729238: 0000004B	00000000000Et
Q175102853	BAC0622AB / BSMENGES / 729239: 0000002A	00000000000Et
Q175102853	BAC0622AB / BSMENGES / 729239: 0000002A	00000000000Et
Q175102853	BAC0622AB / BSMENGES / 729239: 0000002A	00000000000Et
Q175102853	BAC0622AC / BSMENGES / 729221: 00000064	00000000000Et
Q175102853	BAC0622AC / BSMENGES / 729221: 00000064	00000000000Et
Q175102853	BAC0622AC / BSMENGES / 729221: 00000064	00000000000Et
Q175102853	BAC0622AD / BSMENGES / 729222: 00000004	00000000000Et
Q175102853	BAC0622AD / BSMENGES / 729222: 00000004	00000000000Et
Q175102853	BAC0622AD / BSMENGES / 729222: 00000004	00000000000Et
Q175102853	BAC0622AD / BSMENGES / 729223: 0000000C	00000000000Et
Q175102853	BAC0622AD / BSMENGES / 729223: 0000000C	00000000000Et

Idoc720/CSLABEX1/Q175102853/Job or task name search - #1 Idoc72

Collection name (MBRNAME)	iDoctor grouping name (OBJNAME)	iDoctor grouping value (OBJVALUE)	Interval number (INTNU	Interval date time (DTETIM)
Q175102853	SM_INTERFACE_SRC	0000000000000328	1	16062310

August 2018 (1327) – JW – Collection Search underline support

- If the search value contains an underline '_' , the search results for job, subsystem, current user, call stack (pgm, module, procedure), sql statement, wait object name or J9 call stack procedure did not handle it correctly. More results would be returned than desired.

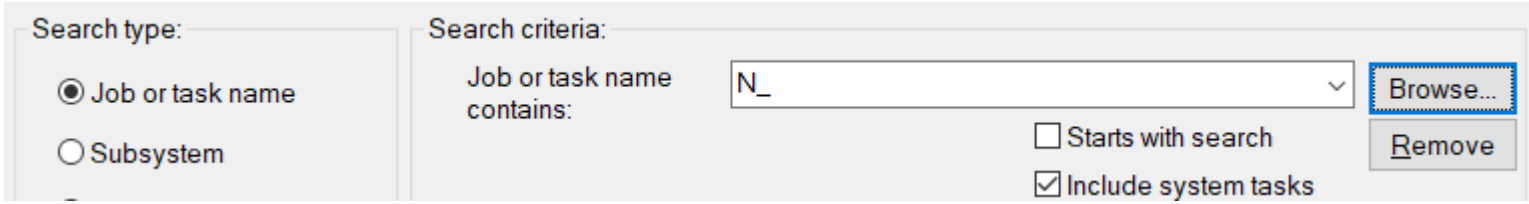

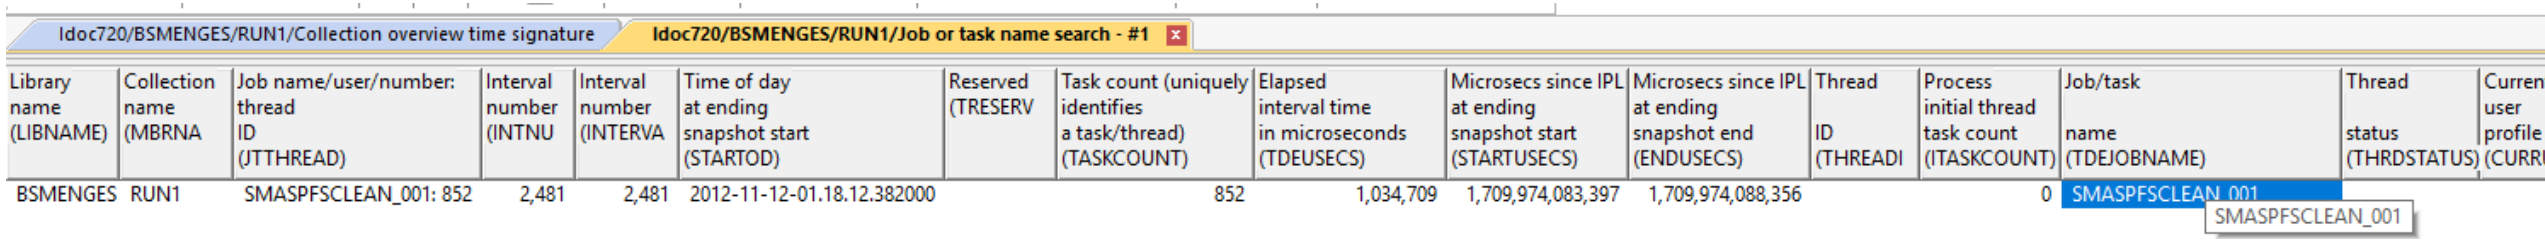

August 2018 (1327) – CSI/JW – Collection Search “Starts with” option

- Added a "Starts with search" checkbox (default is checked) that indicates if the search value should apply to the beginning of data being searched or contained anywhere.

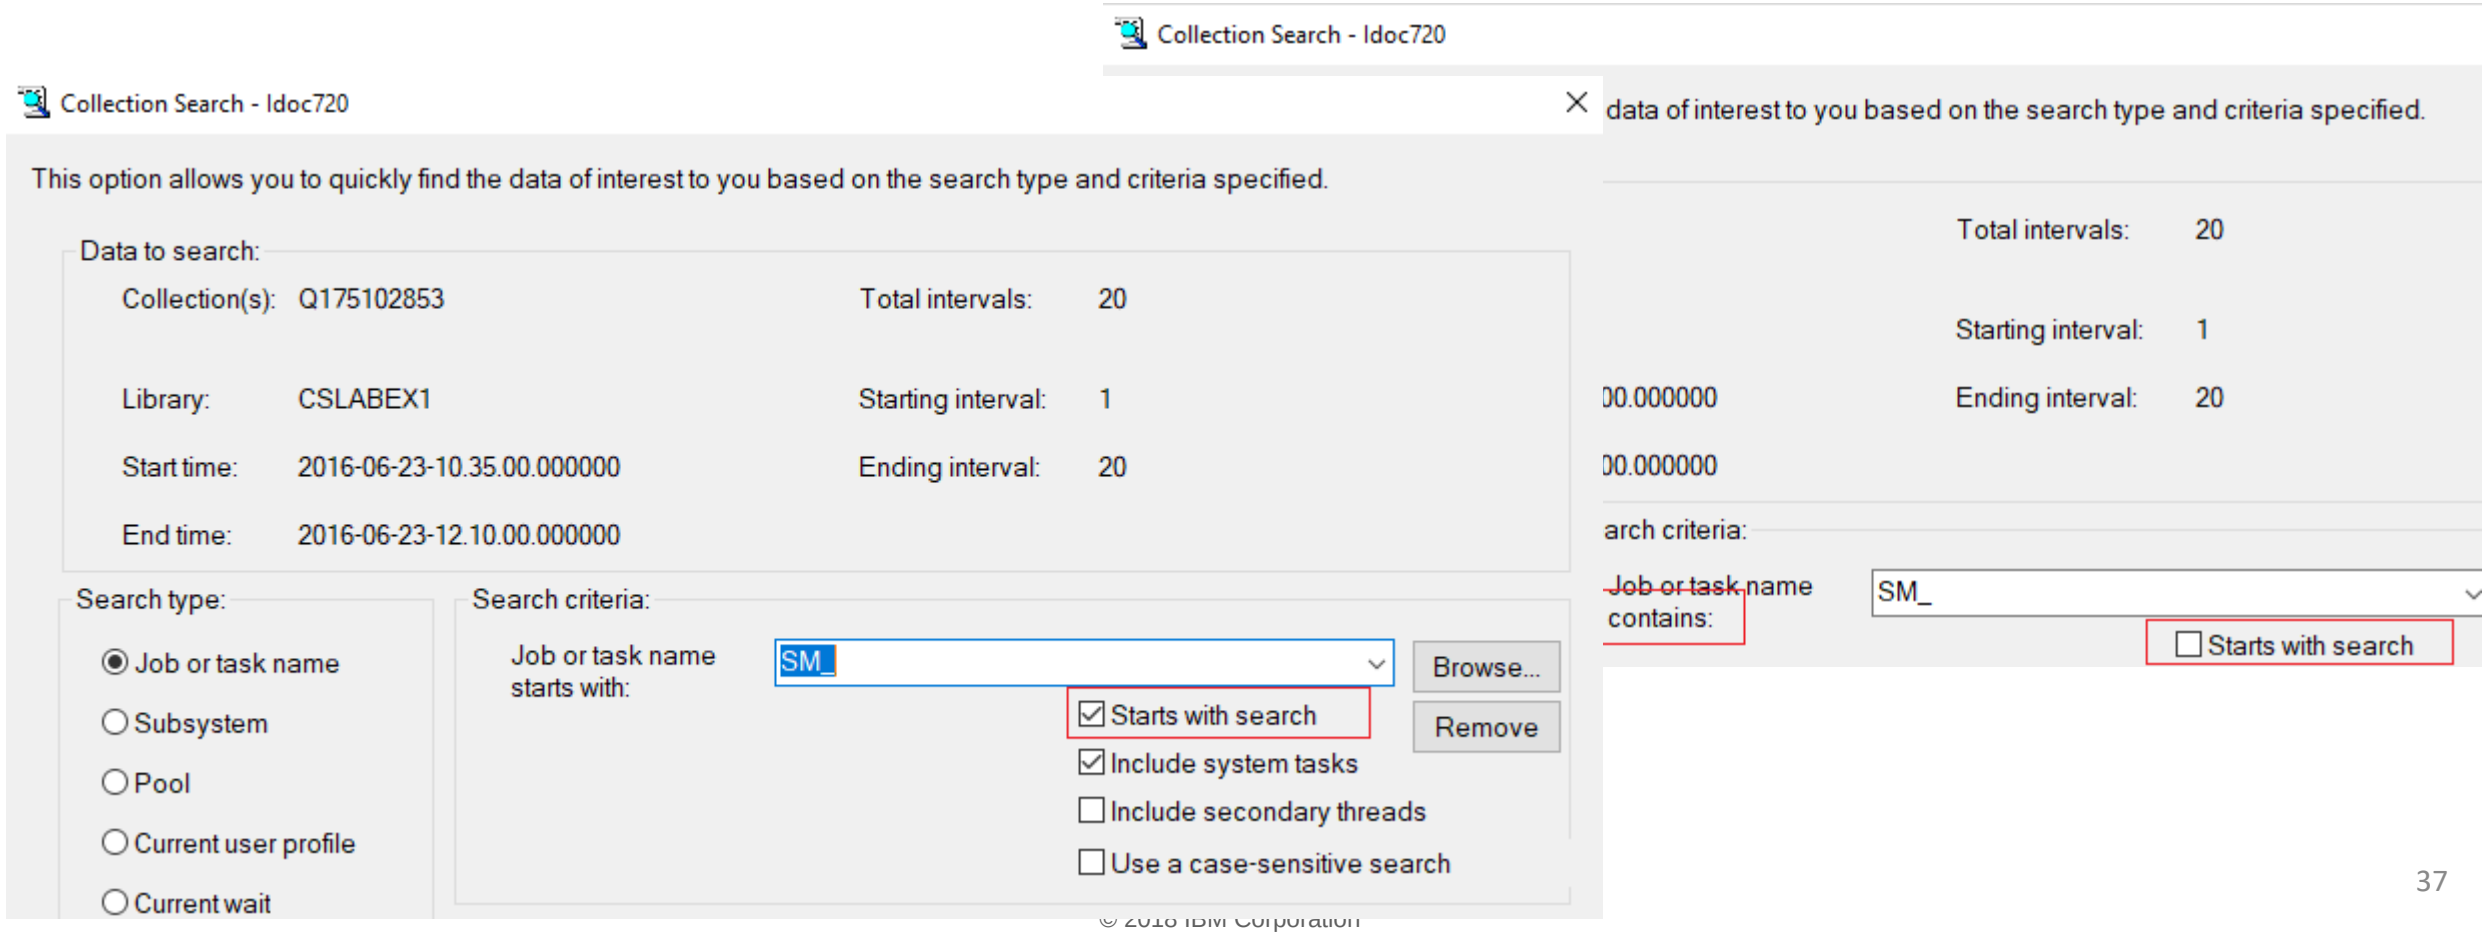

August 2018 (1327) – CSI/JW – Collection Search Browse update

- The browse function will now use the value entered in the textbox (if any) next to it as a filter to reduce the results shown. The new "Starts with" (or contains) option also now applies to the browse window.

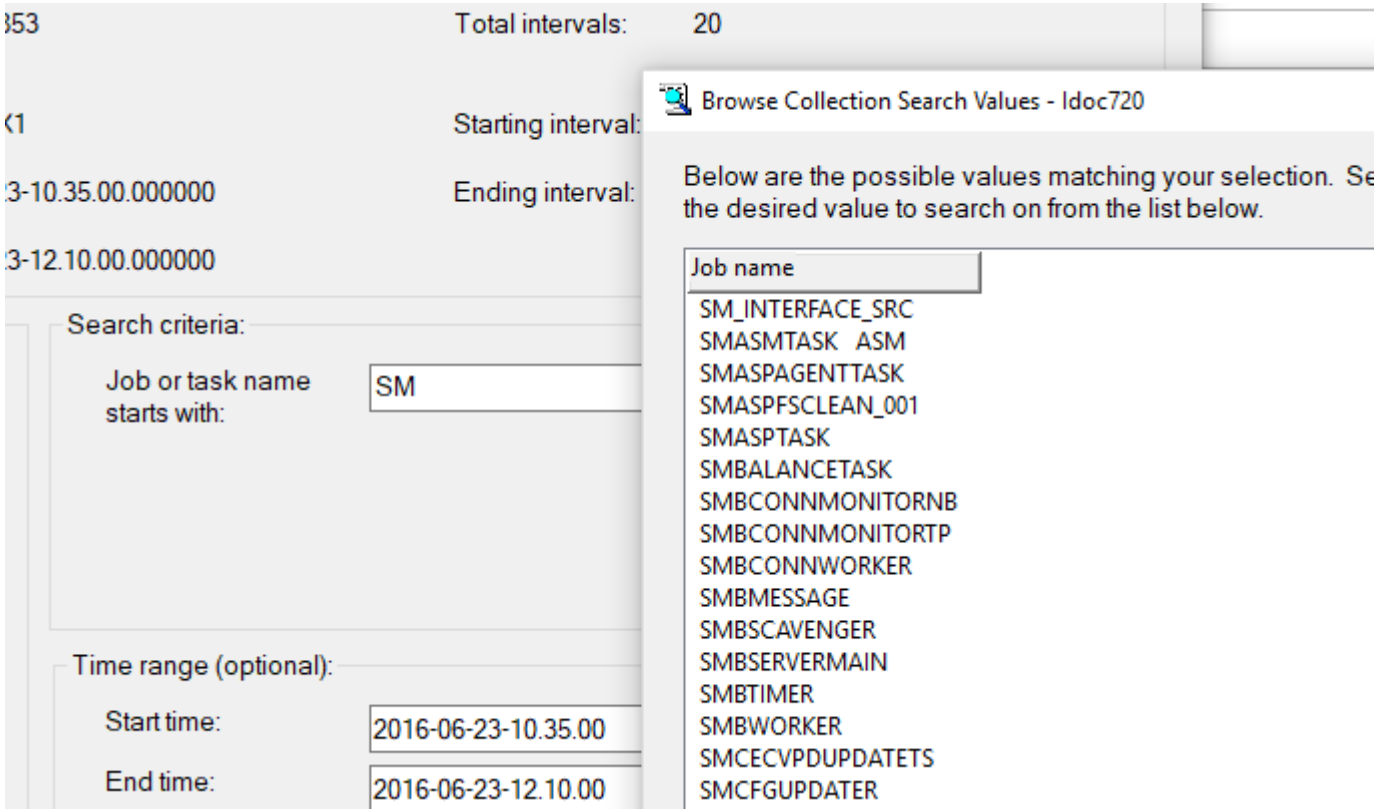

August 2018 (1327) – CSI/JW – Collection Search updates

- The option to browse using search type "Current wait" (i.e. enum) did not work at all.

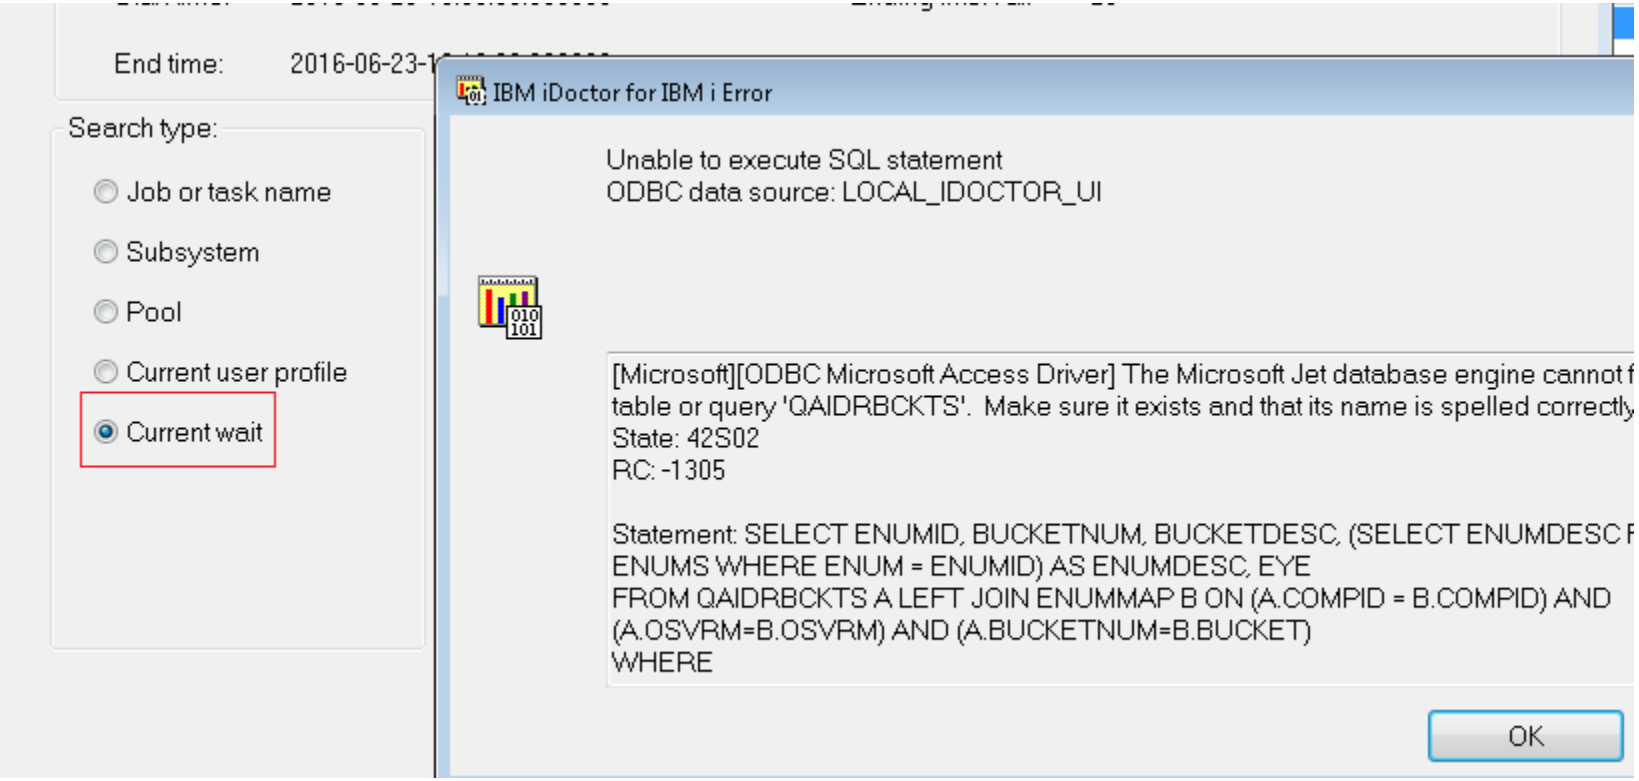

August 2018 (1327) – JW – Collection Search SQL statement option

- Added an option to do a case-sensitive search (or not) on the SQL statement.

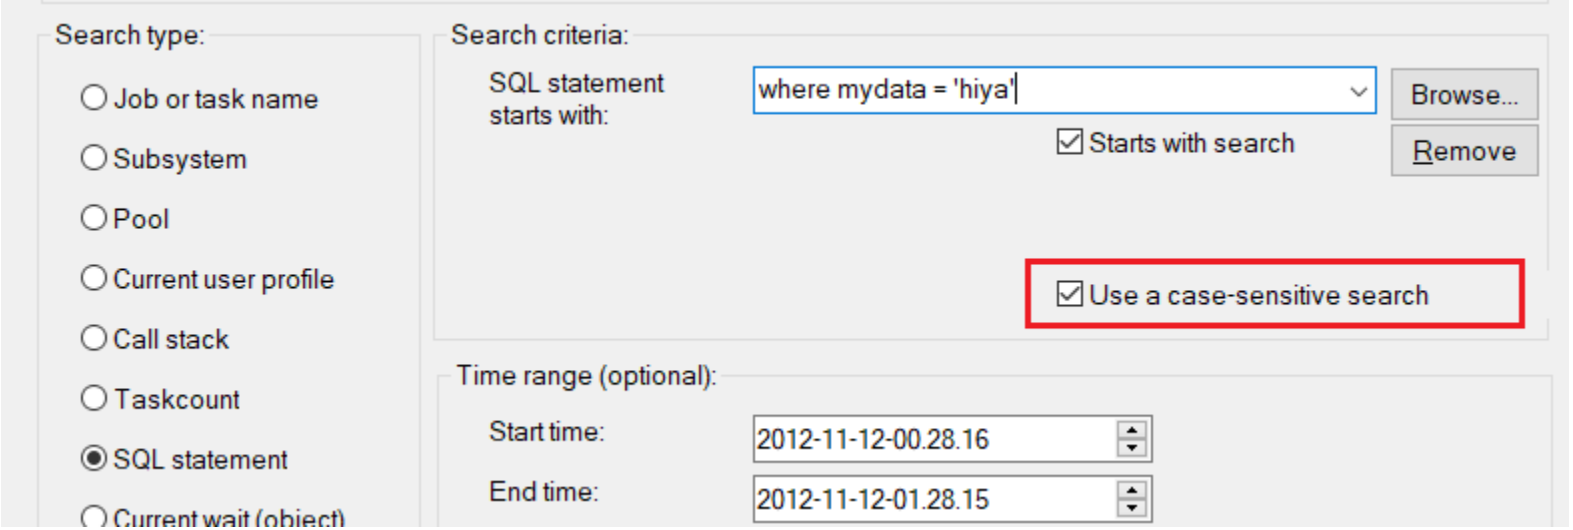

August 2018 (1327) – Display job log error

- Improved the error handling when using the Display job log option and it fails to find the required program lstjbl.exe. It will now clearly state the reason for the error.

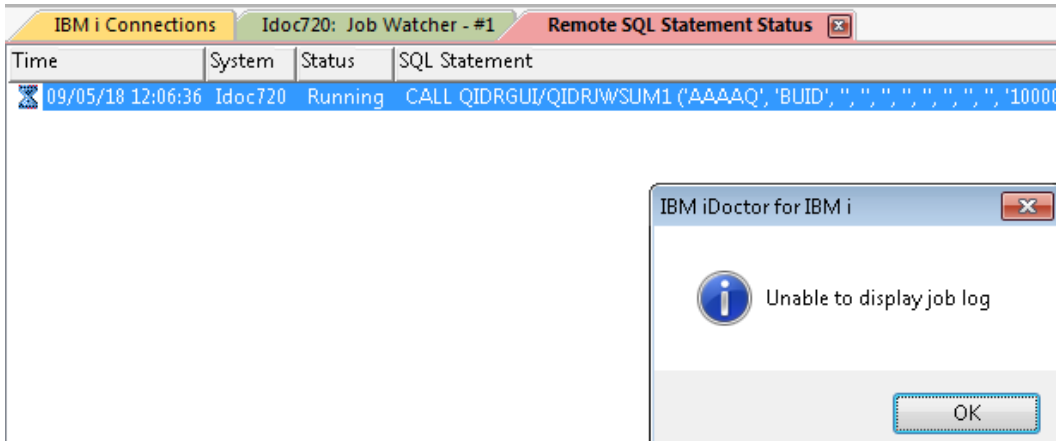

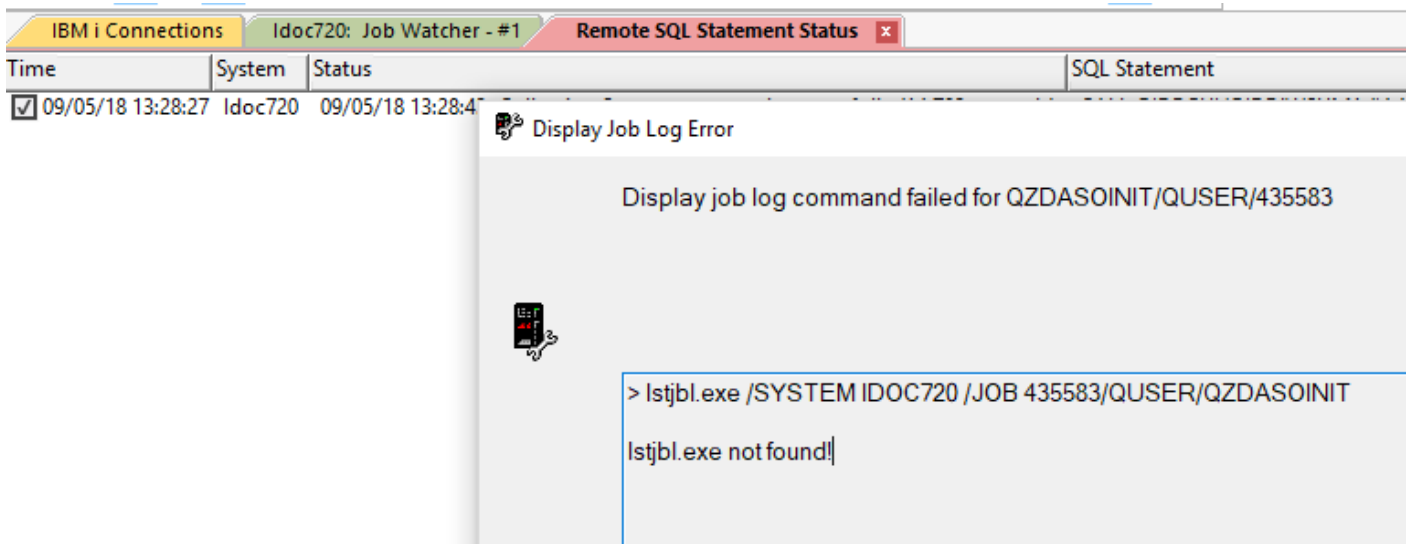

August 2018 (1327) – Installation secure FTP support broken

- In the installer, the support to FTP securely (“Use SSL” option) was broken and has been fixed. Also fixed a potential issue when doing the FTP connection test, it may try to use the wrong directory and cause a failure.

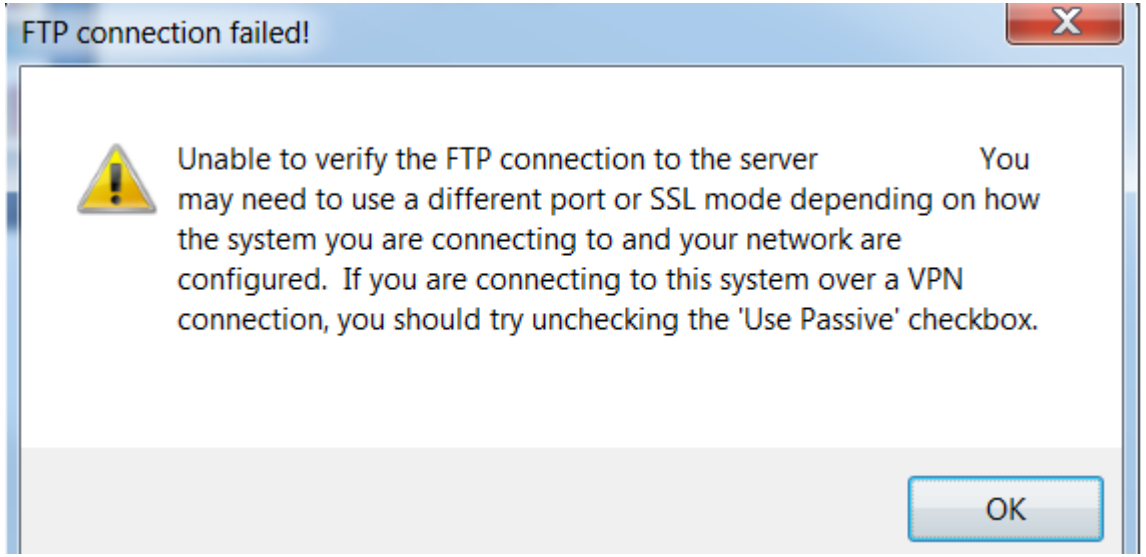

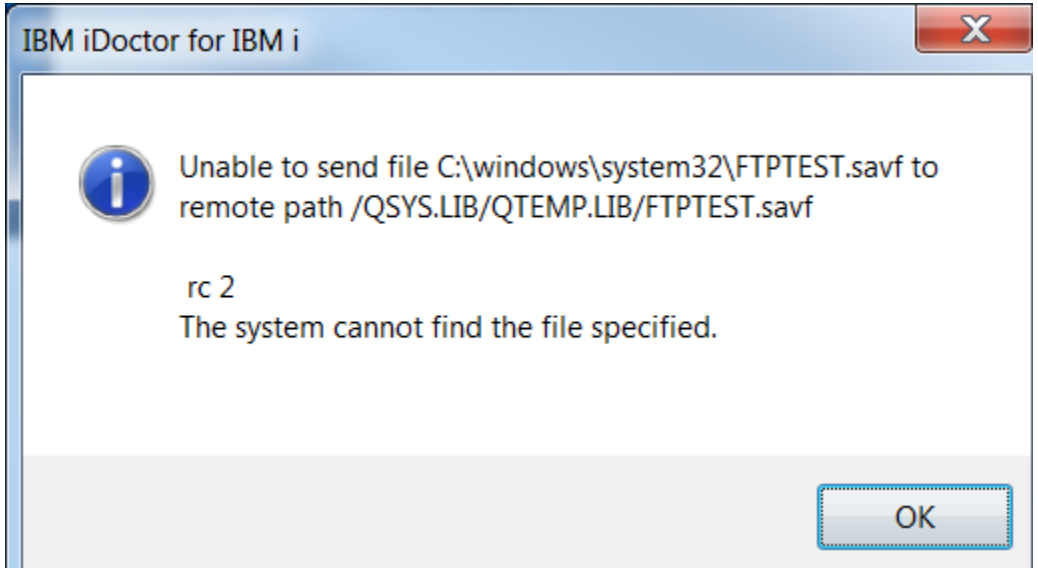

August 2018 (1327) – CSI - CPU Graphs -> P8 or higher update

- Most of the instruction counts are now shown in billions instead of "in units of 1048576 instructions" to improve usability.

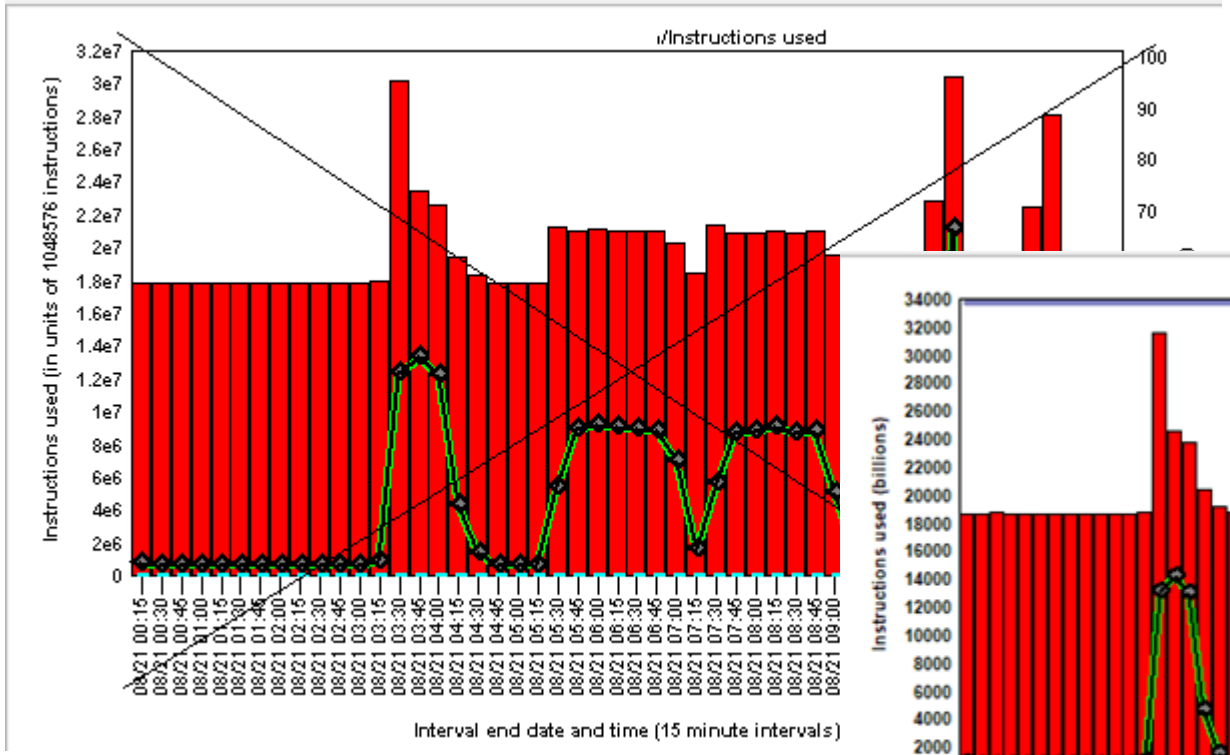

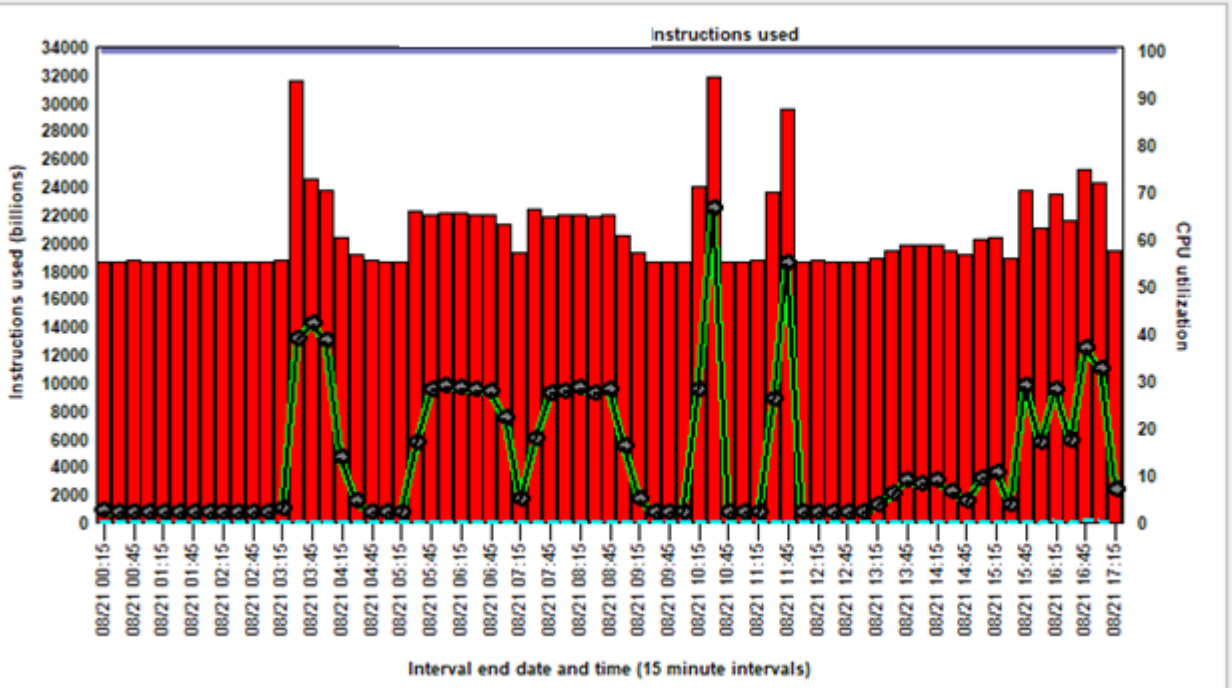

August 2018 (1327) – CSI - CPU Graphs -> P8 or higher update

- The instruction rates are now shown in millions per second instead of "in units of 1048576 instructions per second" to improve usability.

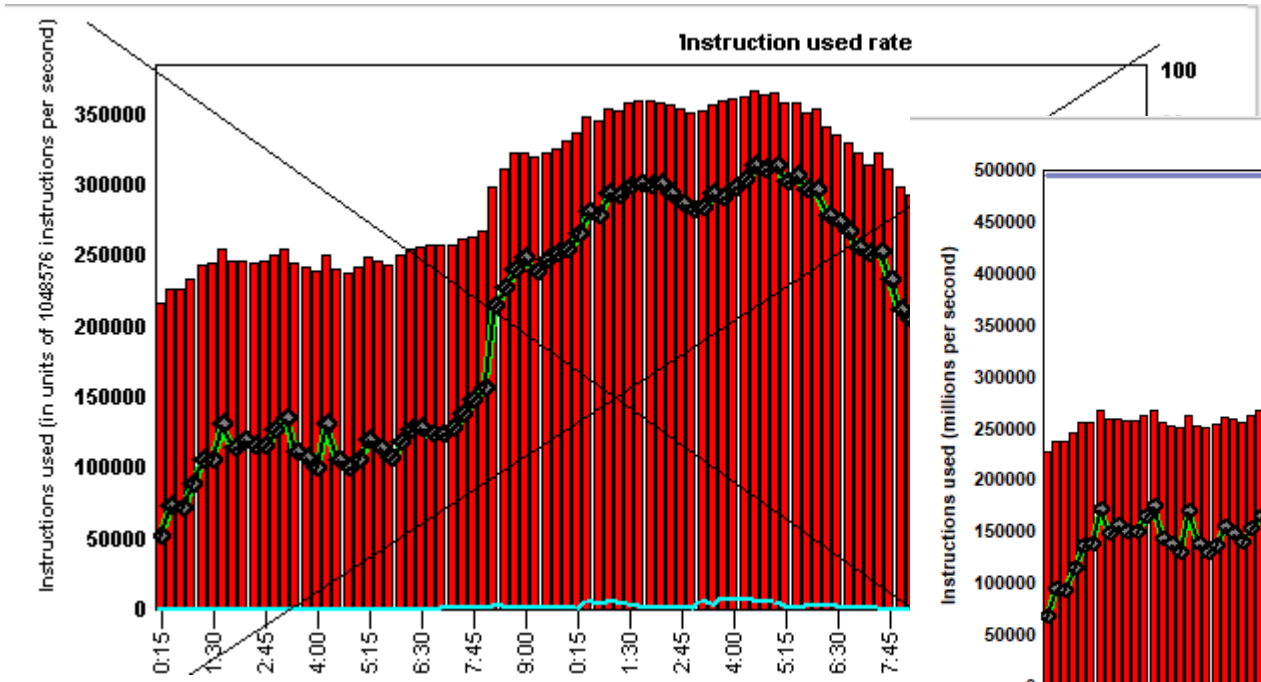

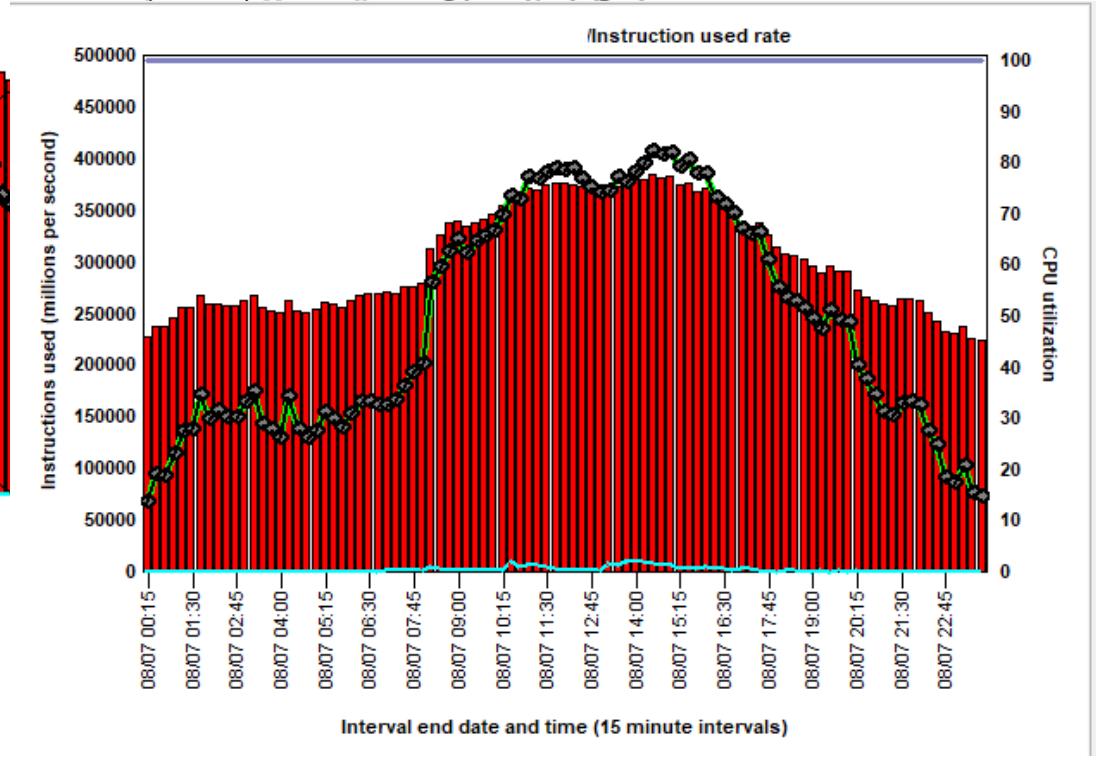

August 2018 (1327) – CPU Graphs -> P8 or higher used vs charged ratio

- These graphs on the Y2 now include a percent of instructions used vs charged (same value for both is at 100%):
  Instructions used
  Instructions charged
  Instruction used rate
  Instruction charged rate

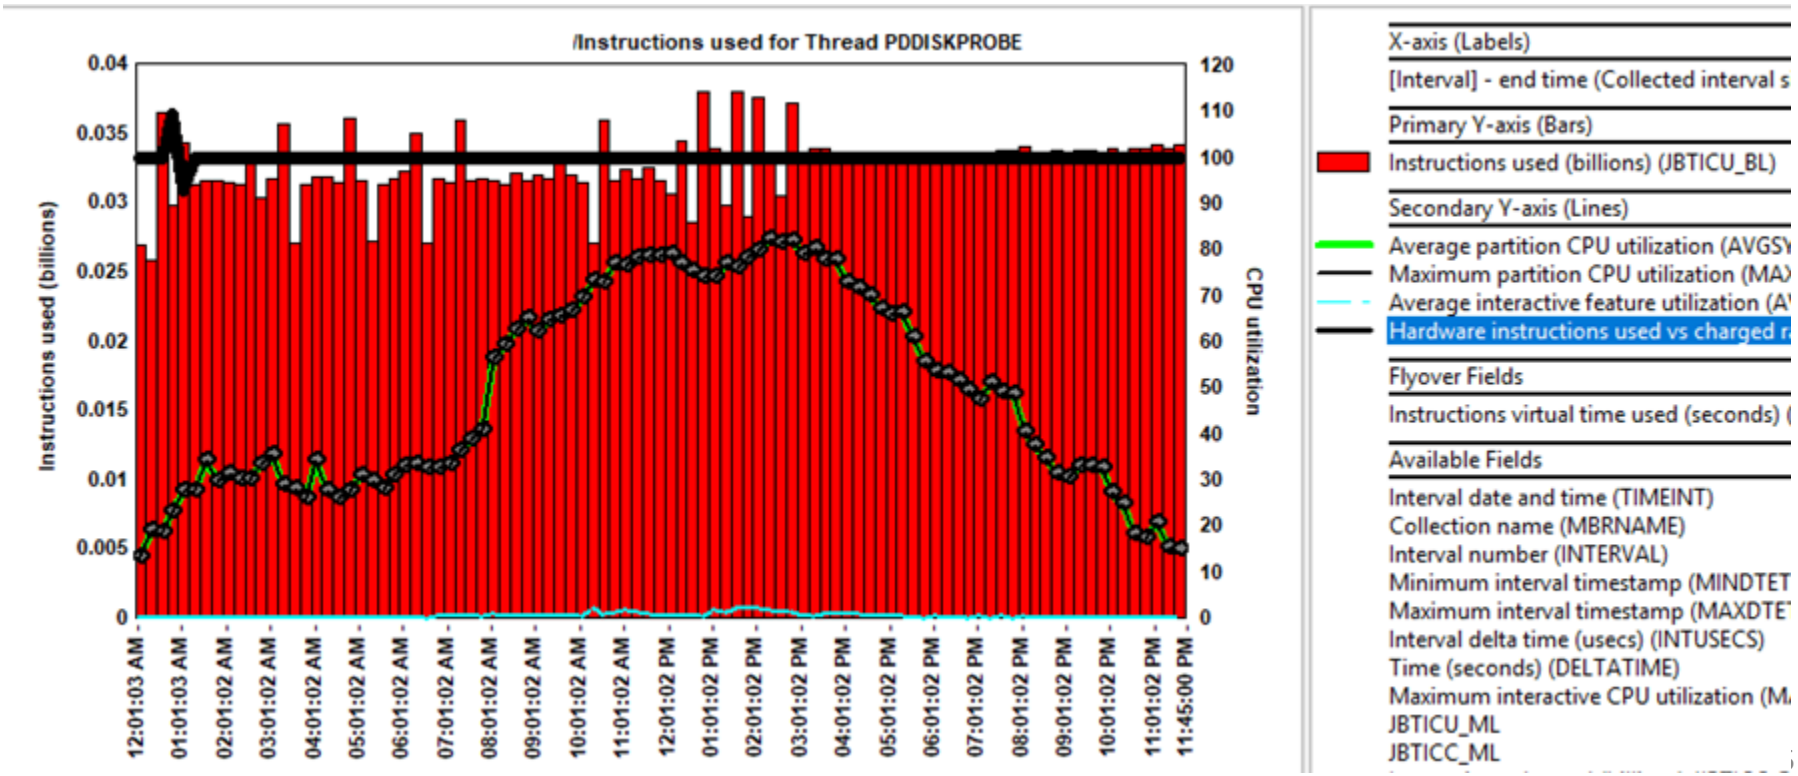

August 2018 (1327) – CSI - CPU Graphs drill down removed

- In CSI the drill down graph Selected Threads/Job/etc -> CPU graphs -> Hardware instructions per second has been removed as it was a duplicate graph with Instructions used rates.

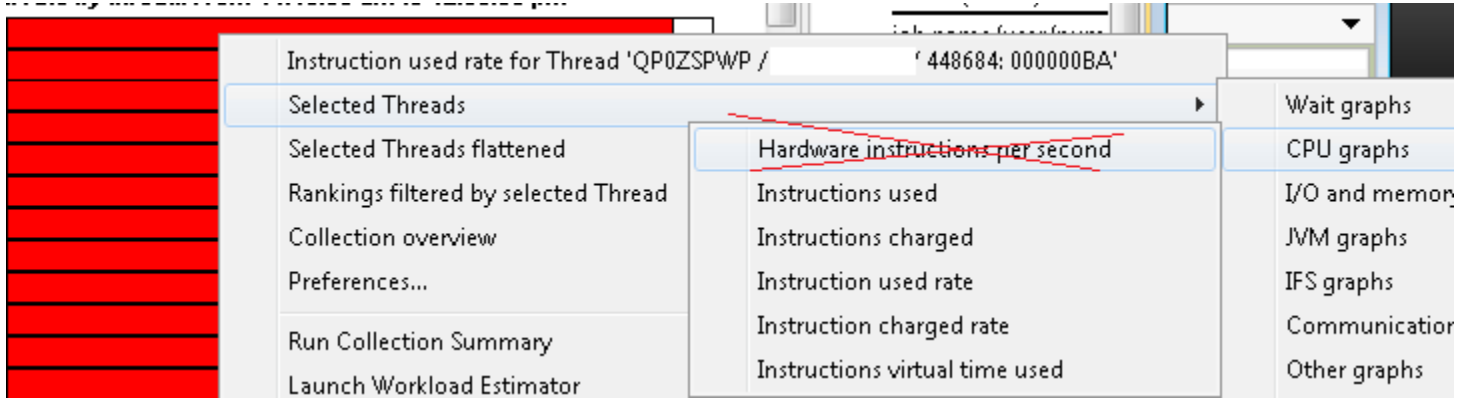

August 2018 (1327) – CSI - Incorrect idle wait calculations

- In CSI, for unsummarized collections, the <OBJTYPE> wait time signature for <XYZ> wait bucket graph calculated the idle waits incorrectly in the SQL statement **for all groupings above thread**. (job, generic job, etc)

- The SQL was incorrectly trying to compare a generic job name to the TDE ID as a filter…

```
FROM QTEMP/QAPMSYSTEM QSY LEFT OUTER JOIN QTEMP/QAPMJOBWTG QWT ON QSY.INTNUM =
QWT.INTNUM
WHERE CHAR(HEX(JWTDE)) IN ('DEFAUL*')
```

August 2018 (1327) – JW -> Waits -> Objects waited on drill downs

- These graphs will now show the "interesting" wait bucket times associated with each wait object instead of showing the number of intervals each wait object was found in the data.

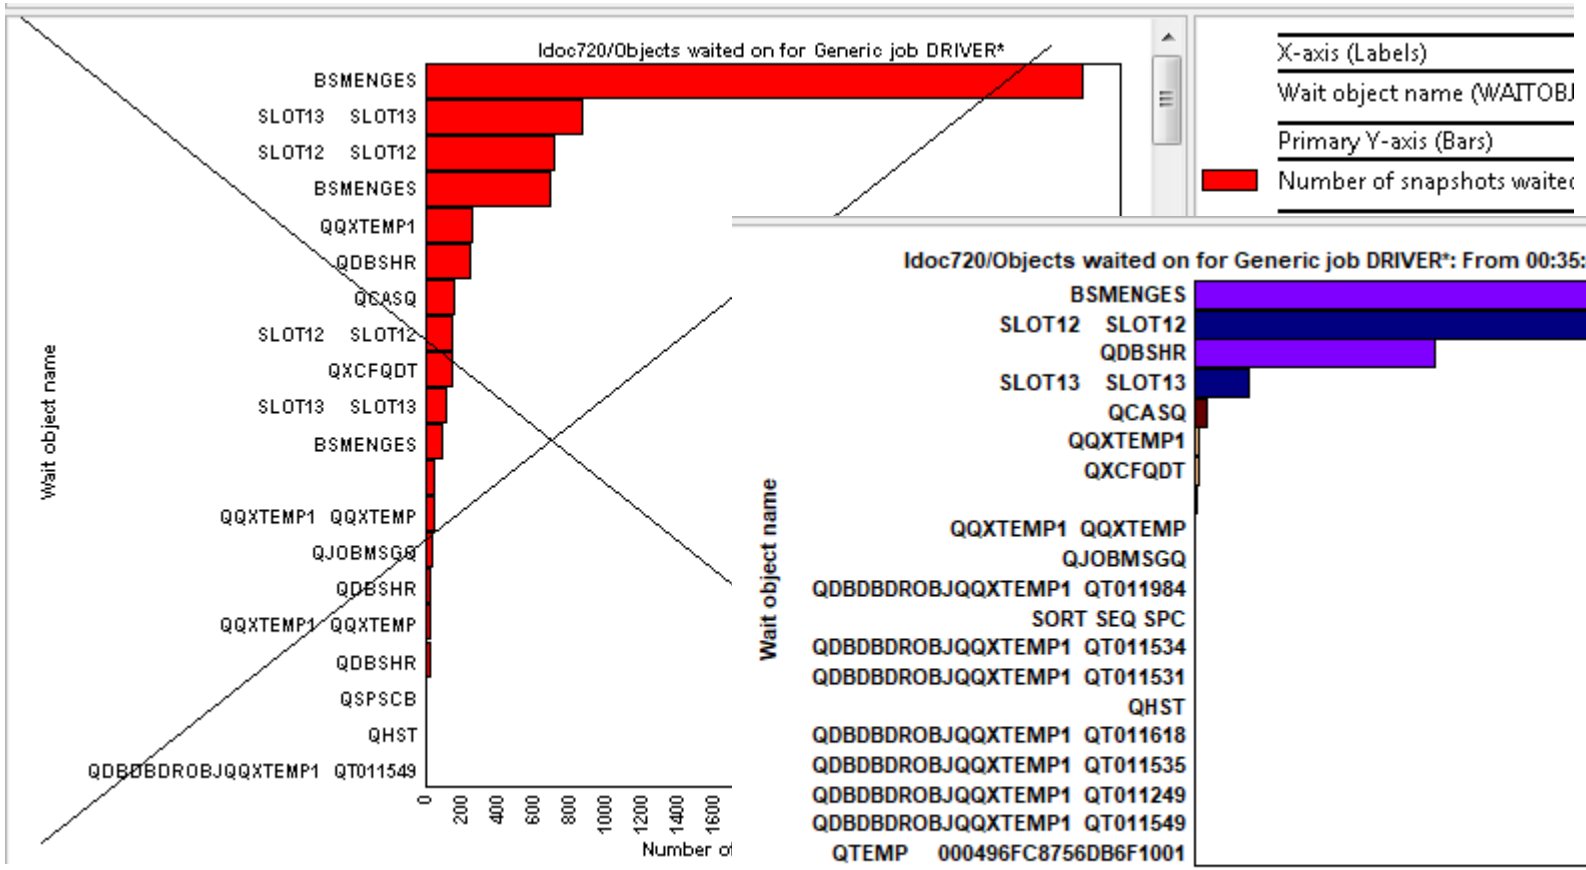

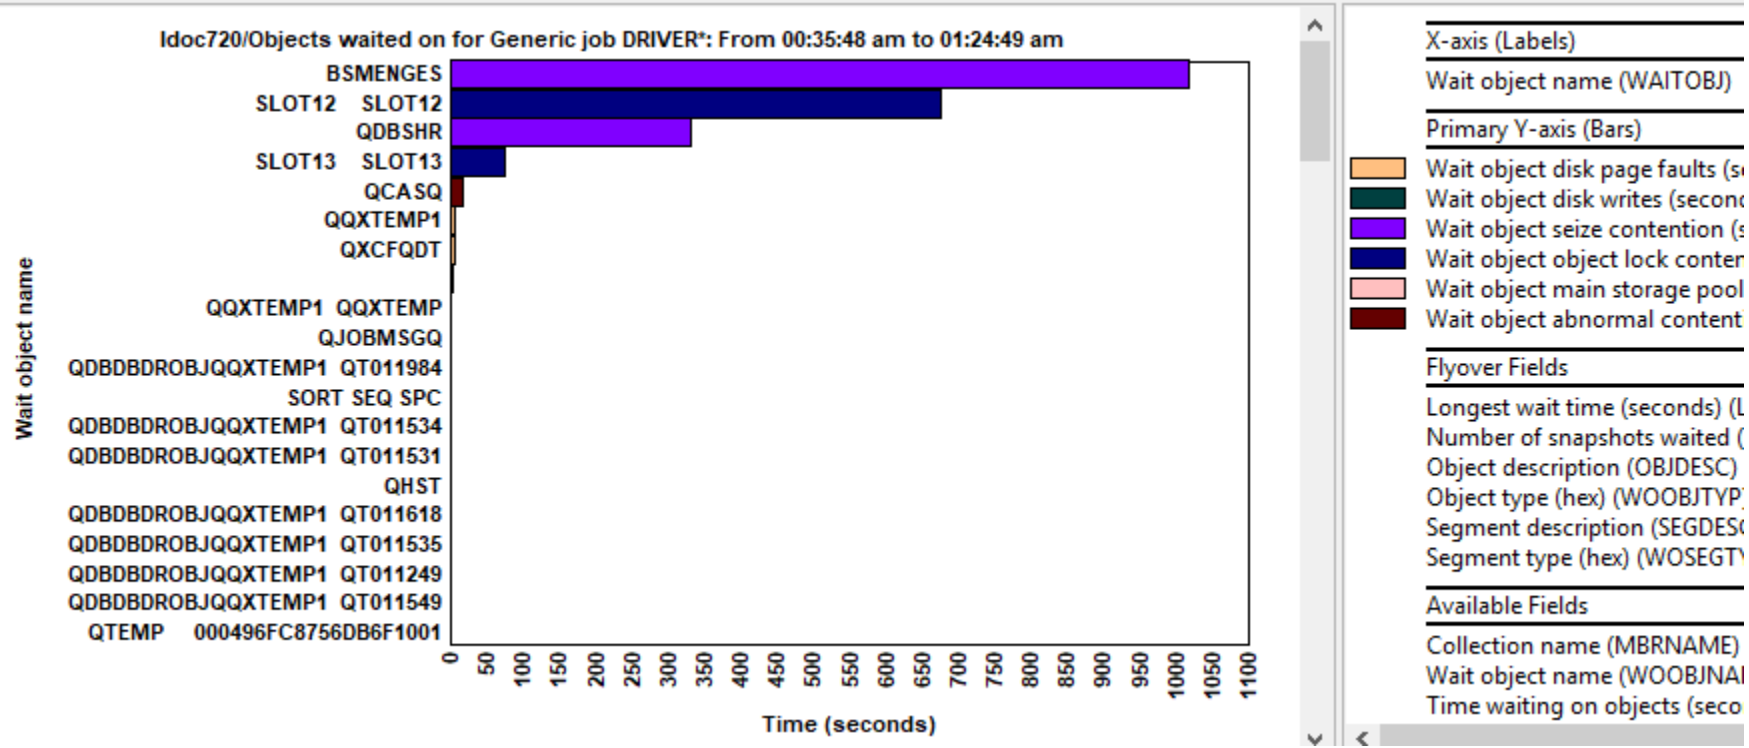

August 2018 (1327) – JW -> Waits -> Objects waited on drill downs

- These graphs previously did not support multiple collections but do now.

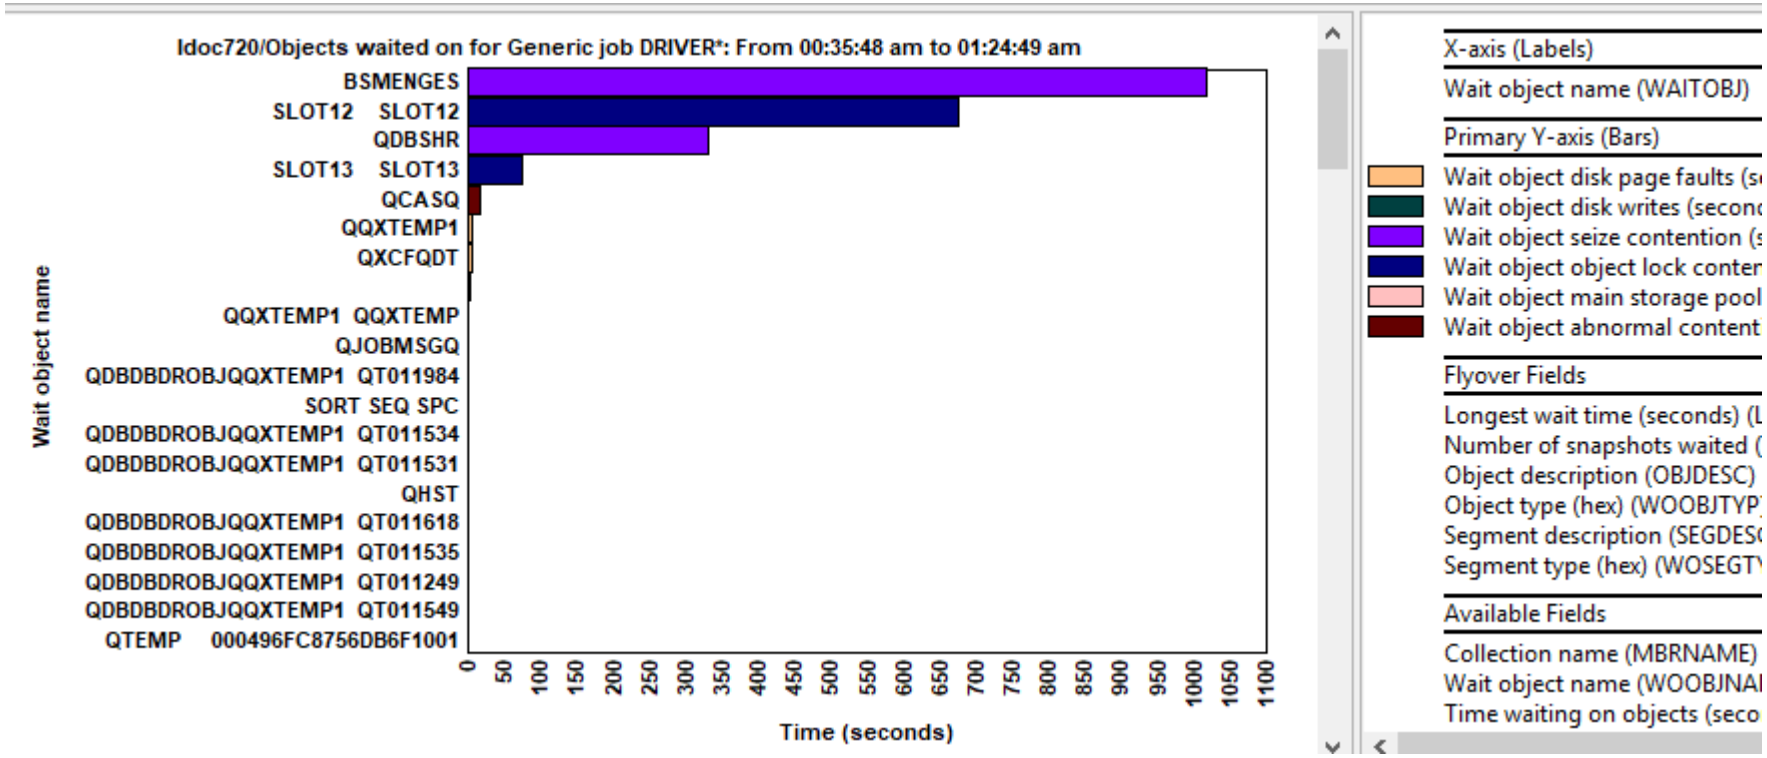

August 2018 (1327) – JW -> Time waiting on objects

- These graphs add up “interesting” current wait durations while a wait object was found and provides analysis capability of the waits associated with the wait objects.

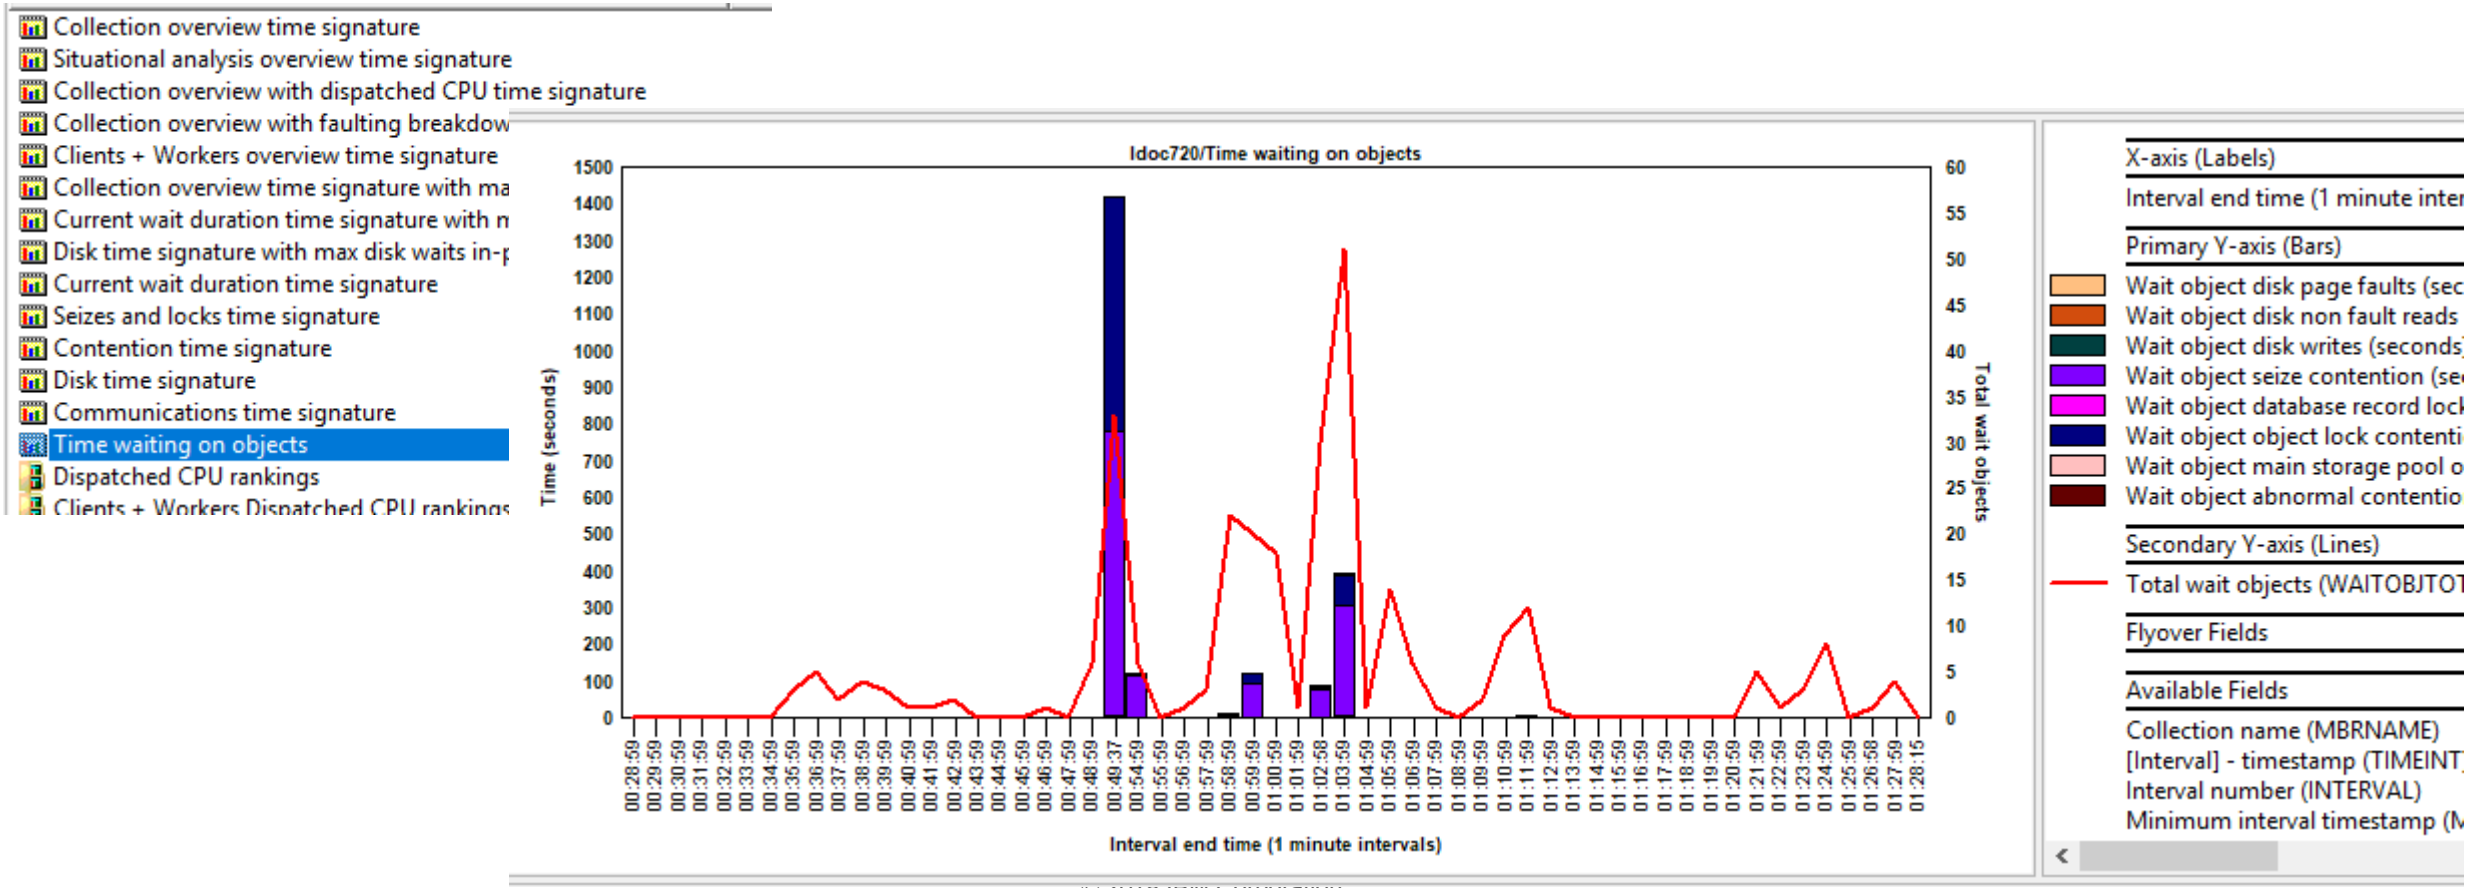

August 2018 (1327) – JW -> Time waiting on objects

- Generic job - wait object job grouping example.

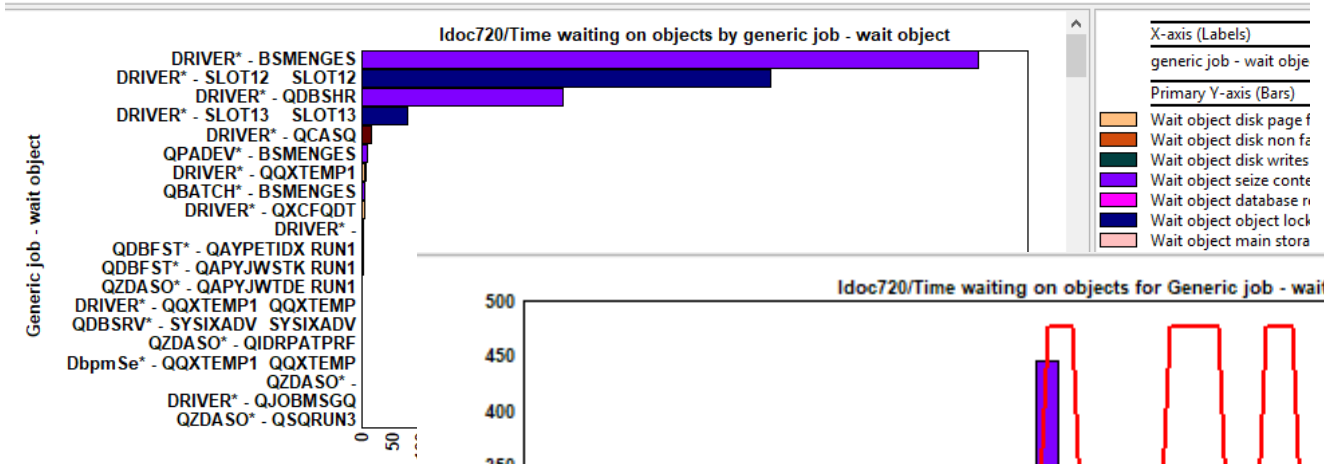

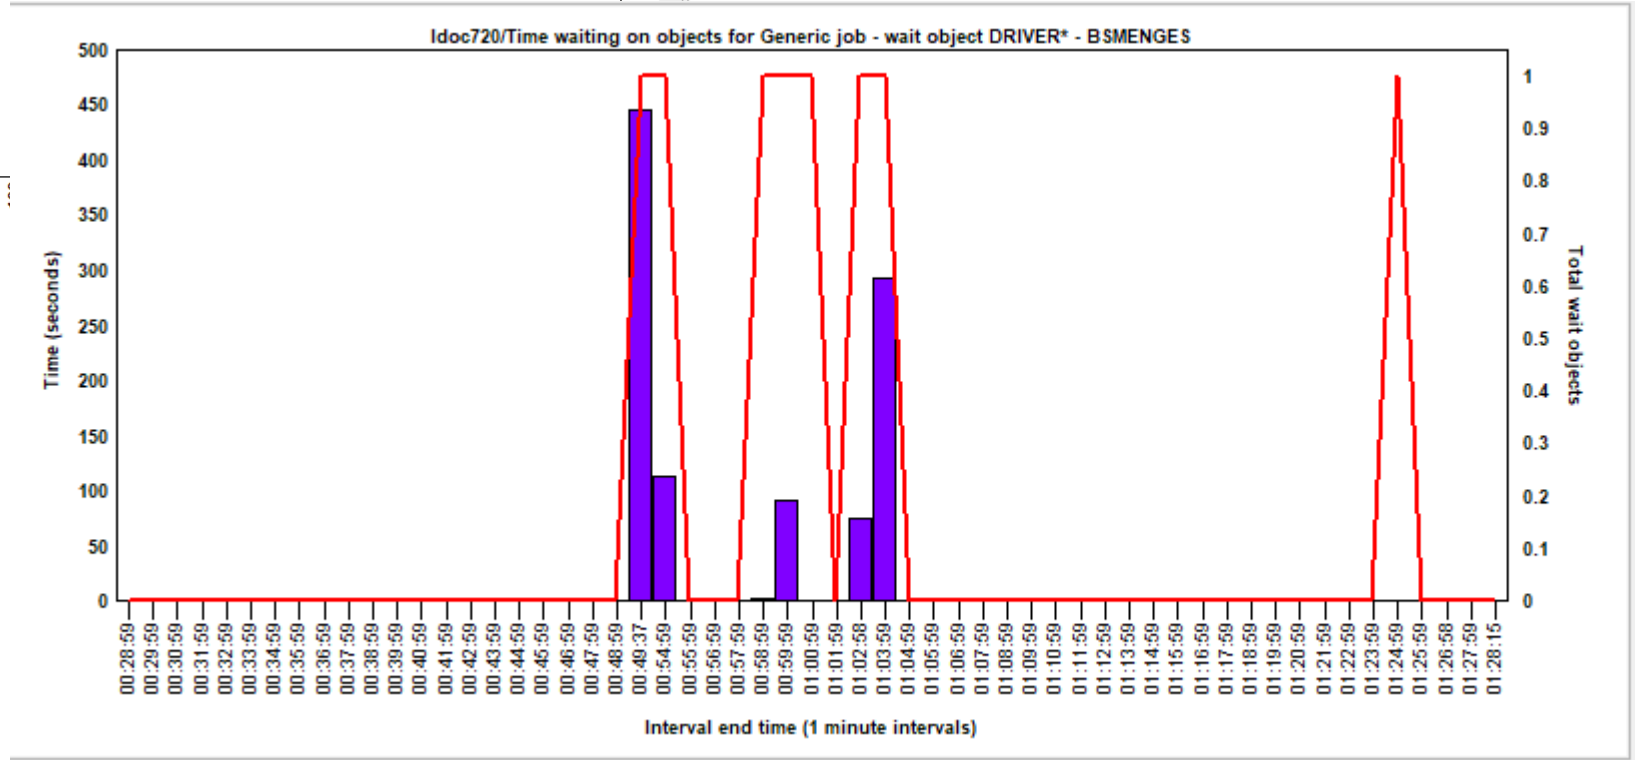

August 2018 (1327) – Fixed bad flyover descriptions in flattened graphs

- In some flattened graphs, the Y1 flyover description was incorrect and has been fixed.

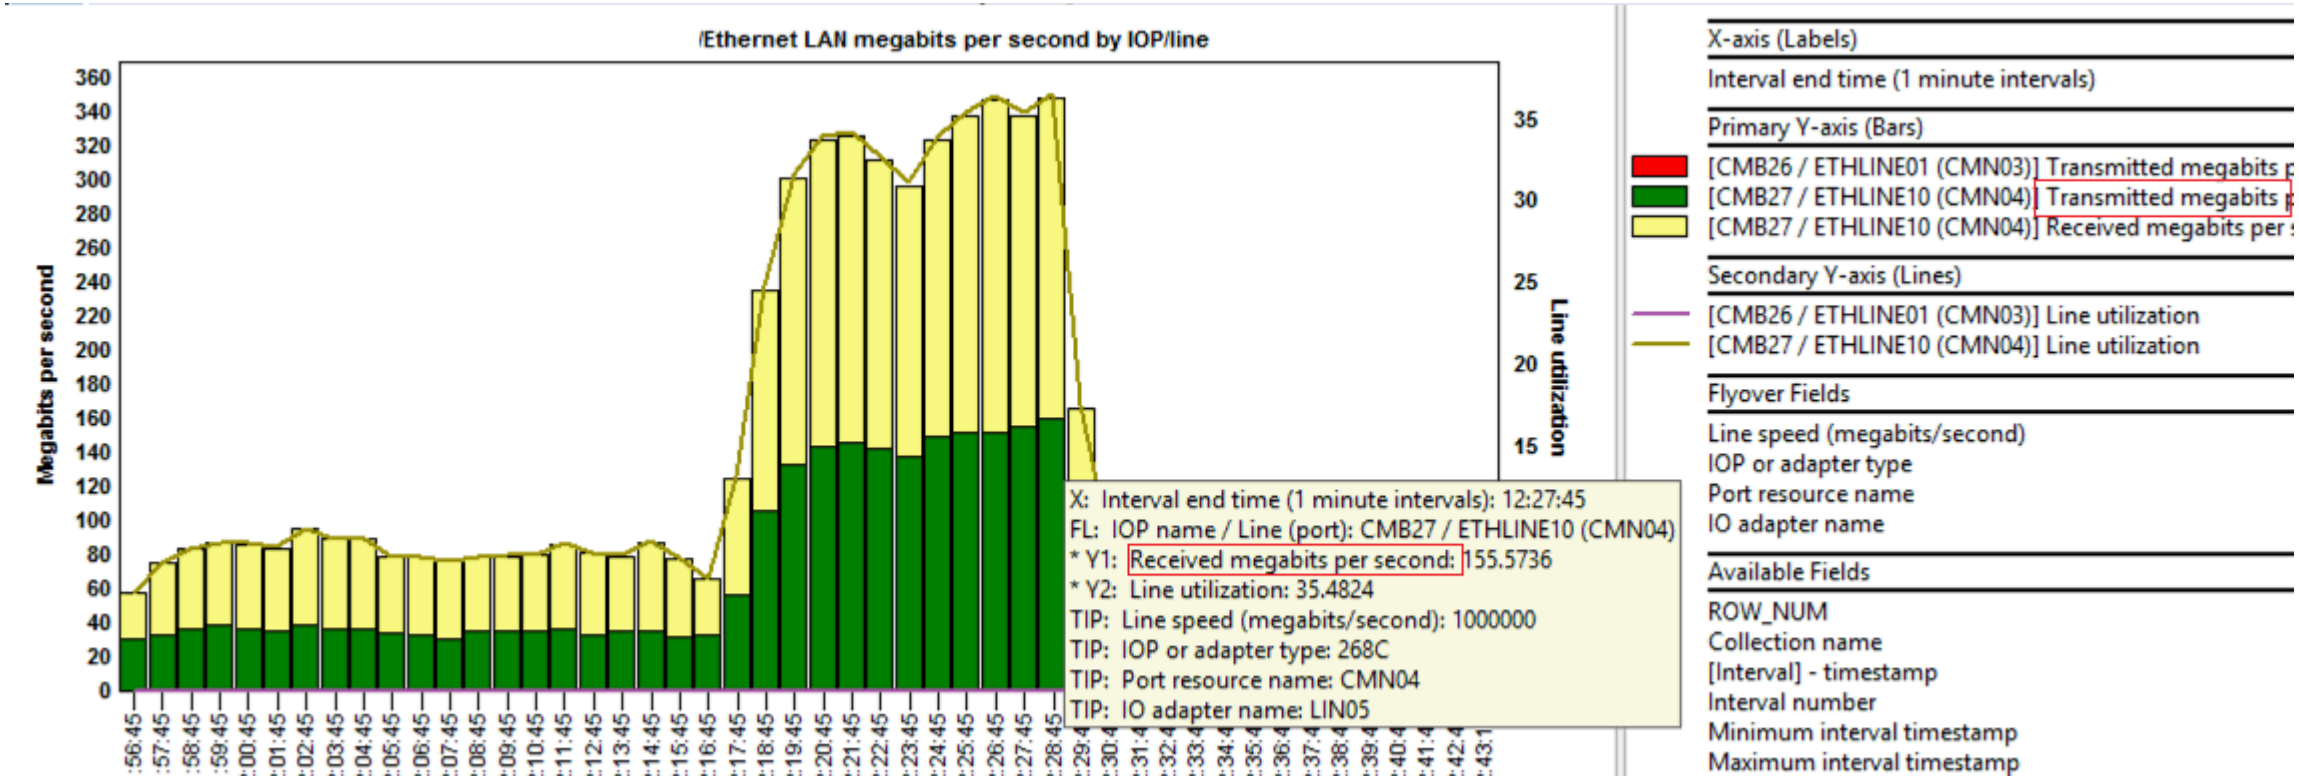

August 2018 (1327) – CSI – Favorites folder updates

- The latest CSI Favorites folder contents are shown below.
  - New graphs = green star, changed graphs = pink star.
  - Removed the following graphs:
    - Dispatched CPU rankings by thread
    - Job temporary storage allocations by thread

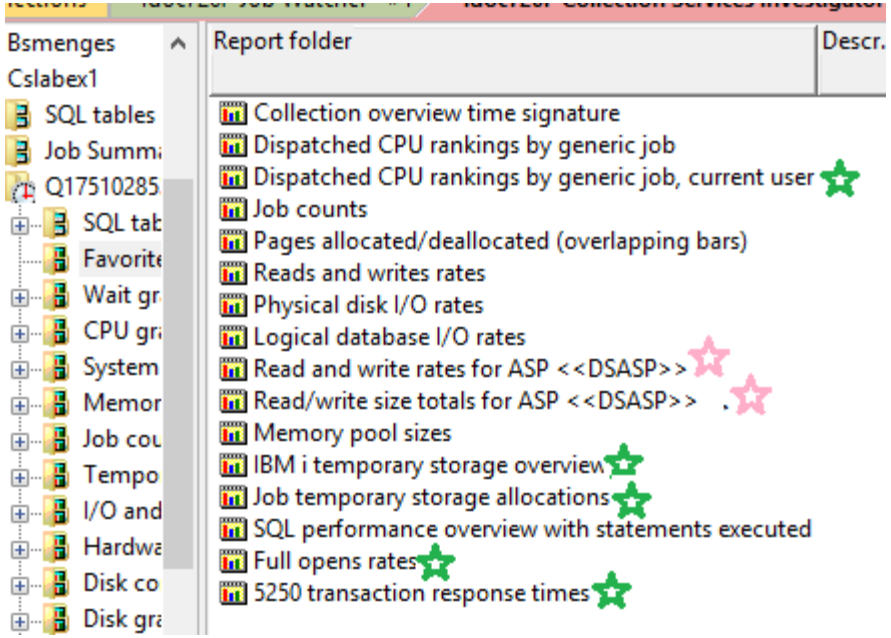

August 2018 (1327) – JW – Favorites folder updates

- The latest JW Favorites folder contents are shown below.
  - New graphs = green star!

  - Removed the following graphs:
    - Dispatched CPU rankings by thread
    - Job temporary storage allocations by thread

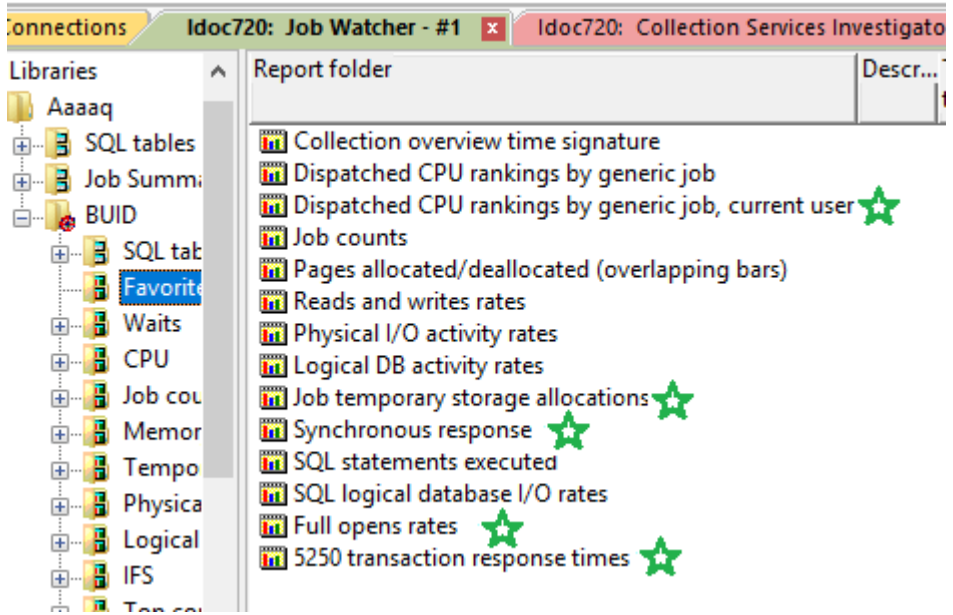

August 2018 (1327) – CSI – Active virtual processors graph

- The values shown can now exceed 32 virtual processors.
  - Data now comes from QAPMSYSVP instead of QAPMSYSCPU which allows this.
  - Simplified the field descriptions shown to be:
    Total virtual processors (TOTVPS)
    Active virtual processors (ACTVPS)

August 2018 (1327) – JW – Generic job for the selected thread drill down

- Renamed the drill down to "Generic job for selected <<OBJTYPE>>" i.e. thread/user/etc since this option is available from more types of rankings than just thread rankings.

- The rates graphs showed inaccurate rate values if the interval selection was greater than the collected interval size.

- The Generic job X* memory page demand graph did not work.

August 2018 (1326) – Power Connections issues fixed

- Fixed bugs in many spots related to the connection settings for Power systems (VIOS, HMC, etc) getting confused / saved in the location for IBM i connections instead.
  - FTP transfers may fail, because the wrong connection settings were being used.
  - This also effected functions that show data (open file), transfer files, or the edit option.

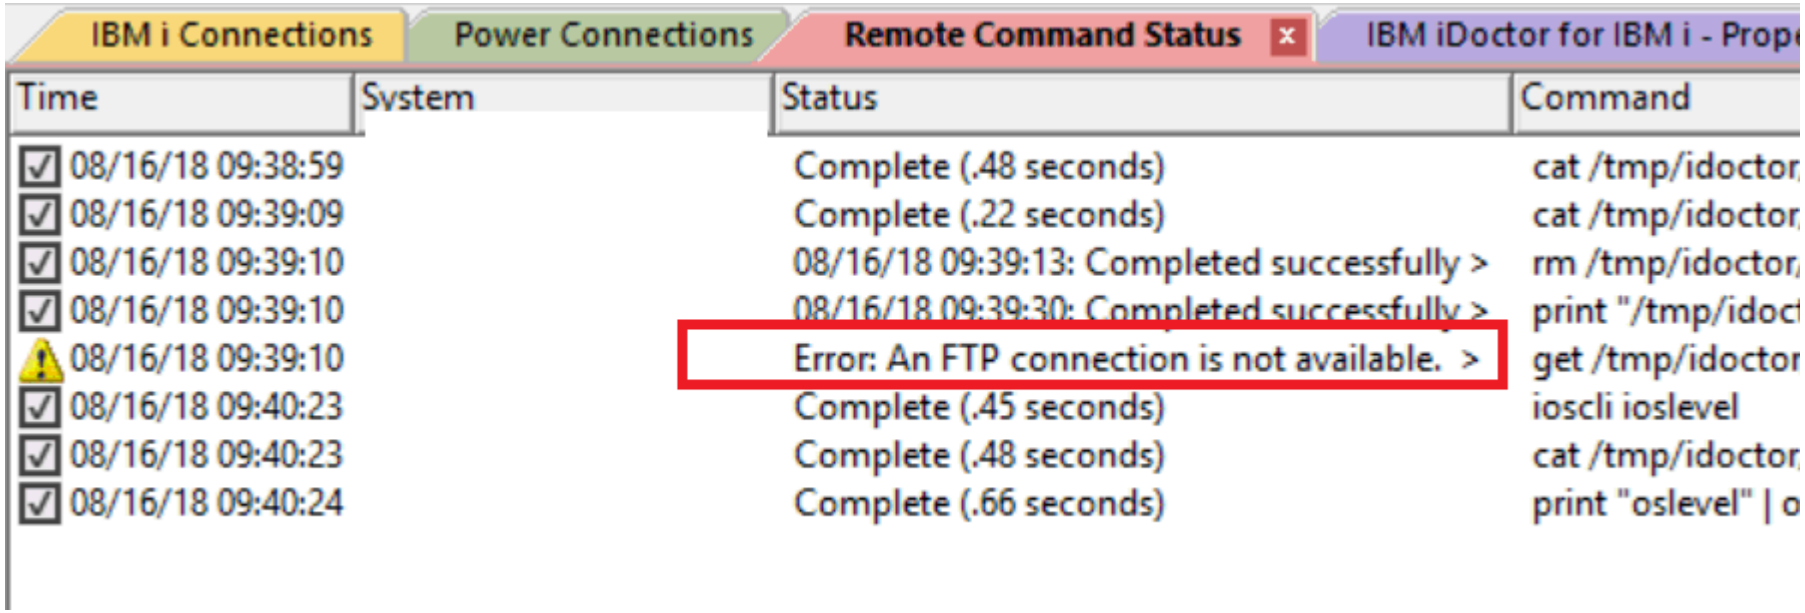

August 2018 (1326) – Power Connections find window

- In the find window, the preference for the search directories can now be saved for each type of data (nmon, npiv, etc.) Previously all types of data shared the same value.
- Also the text on the window that refers to multiple system(s) when doing a find has been removed as this support no longer exists.

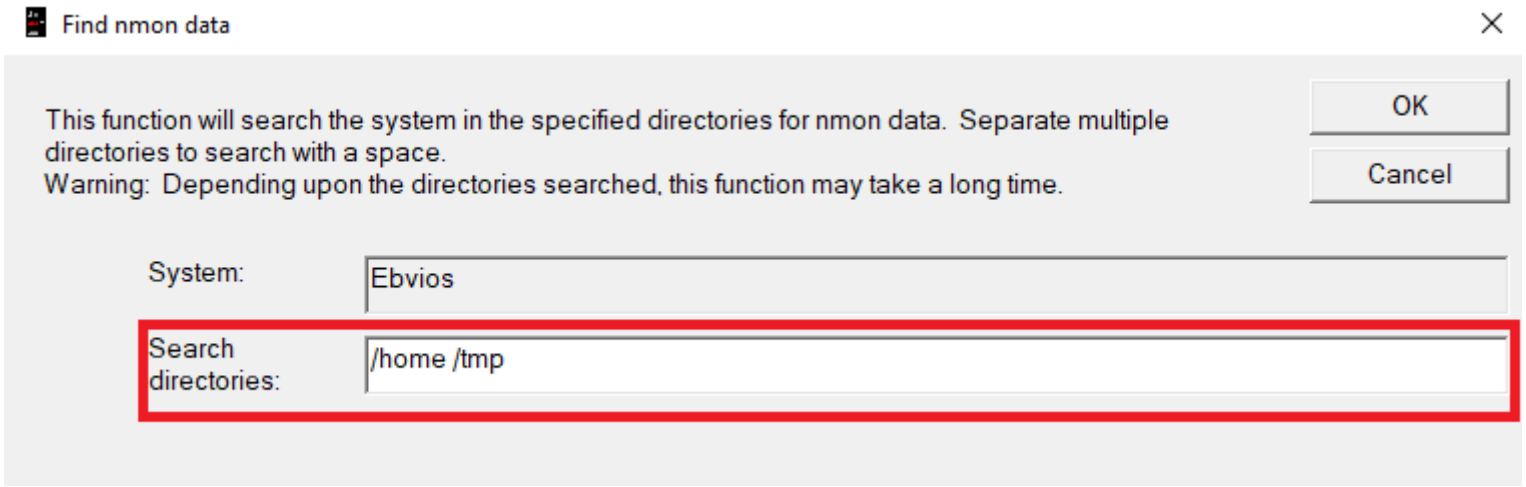

August 2018 (1326) – JW temp storage graphs inaccurate

- In Job Watcher, the Temporary storage - job temporary storage ranking graphs and selection over time were inaccurate and have been fixed. The values given were considerably smaller than they should have been in some situations.

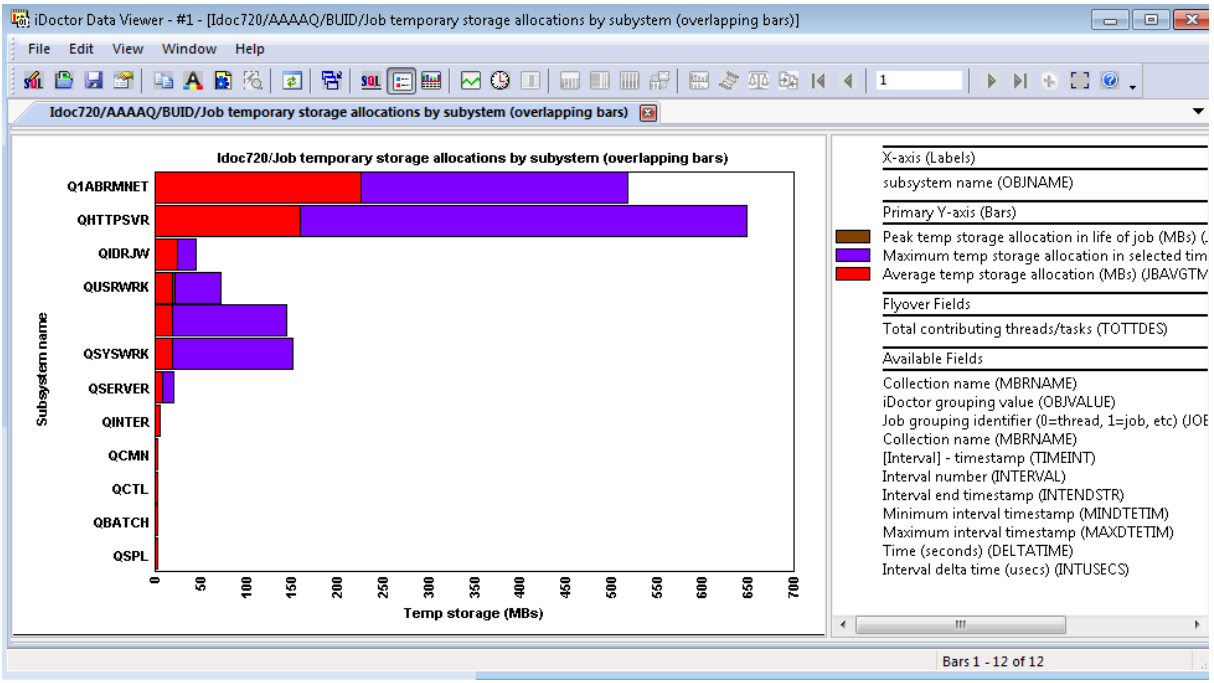

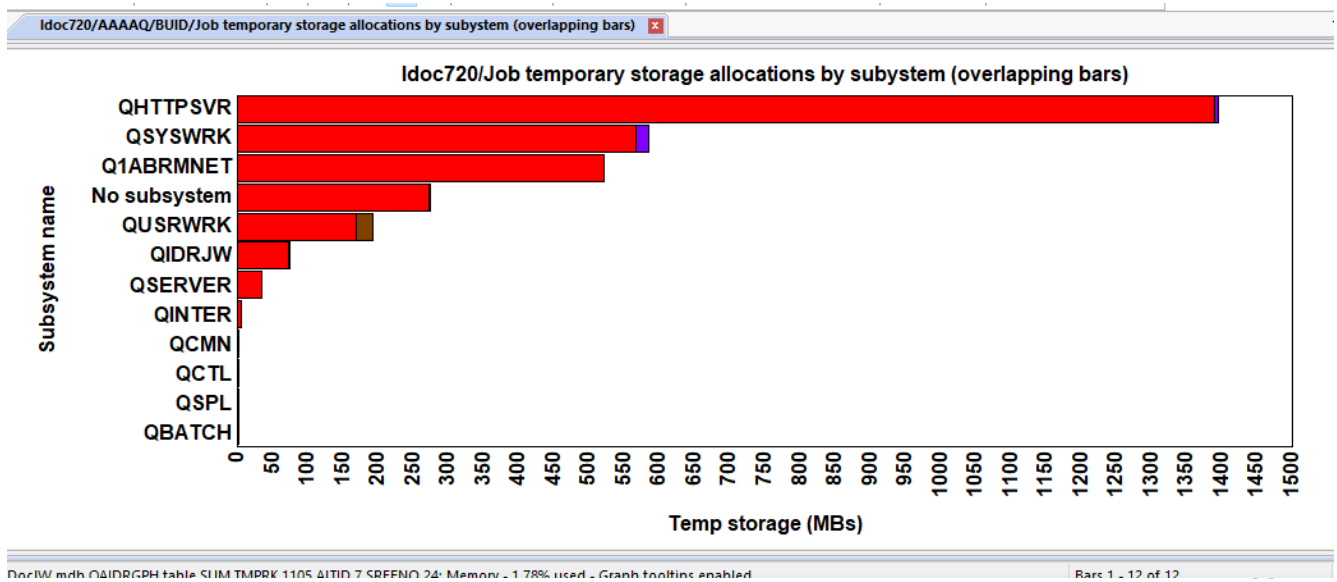

August 2018 (1326) – JW Job counts missing “always idle work”

- The Job Watcher job counts graphs were inaccurate as they did not include jobs/threads that never used CPU during the collection. Because they never used CPU we do not know if they are jobs, threads or tasks.

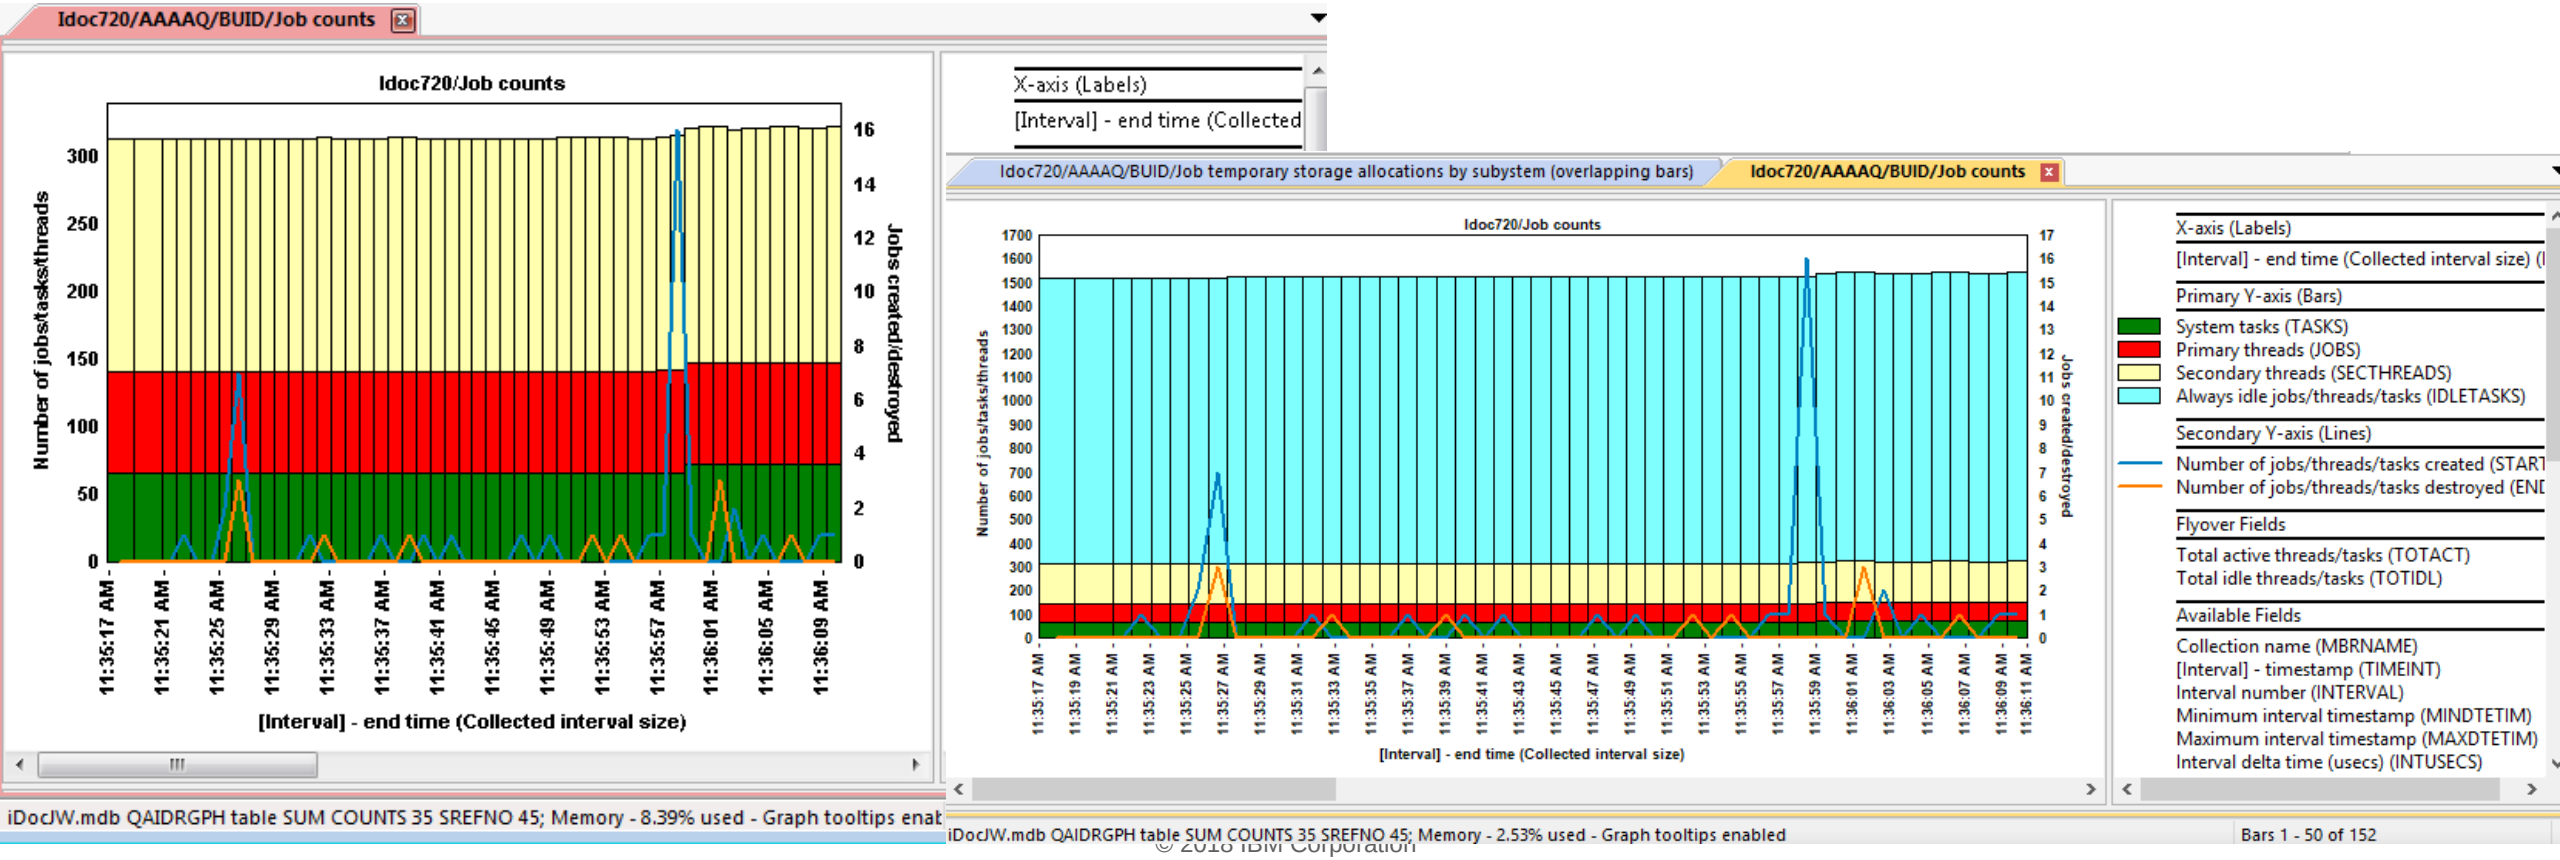

August 2018 (1326) – CSI ranking graphs – “No subsystem”

- In CSI fixed a problem in some of the ranking graphs where you may see "All system tasks" appear in the results when ranking by subsystem. This was really just data (not necessarily system tasks) where no subsystem was assigned to the job. This has been renamed to "No subsystem”.

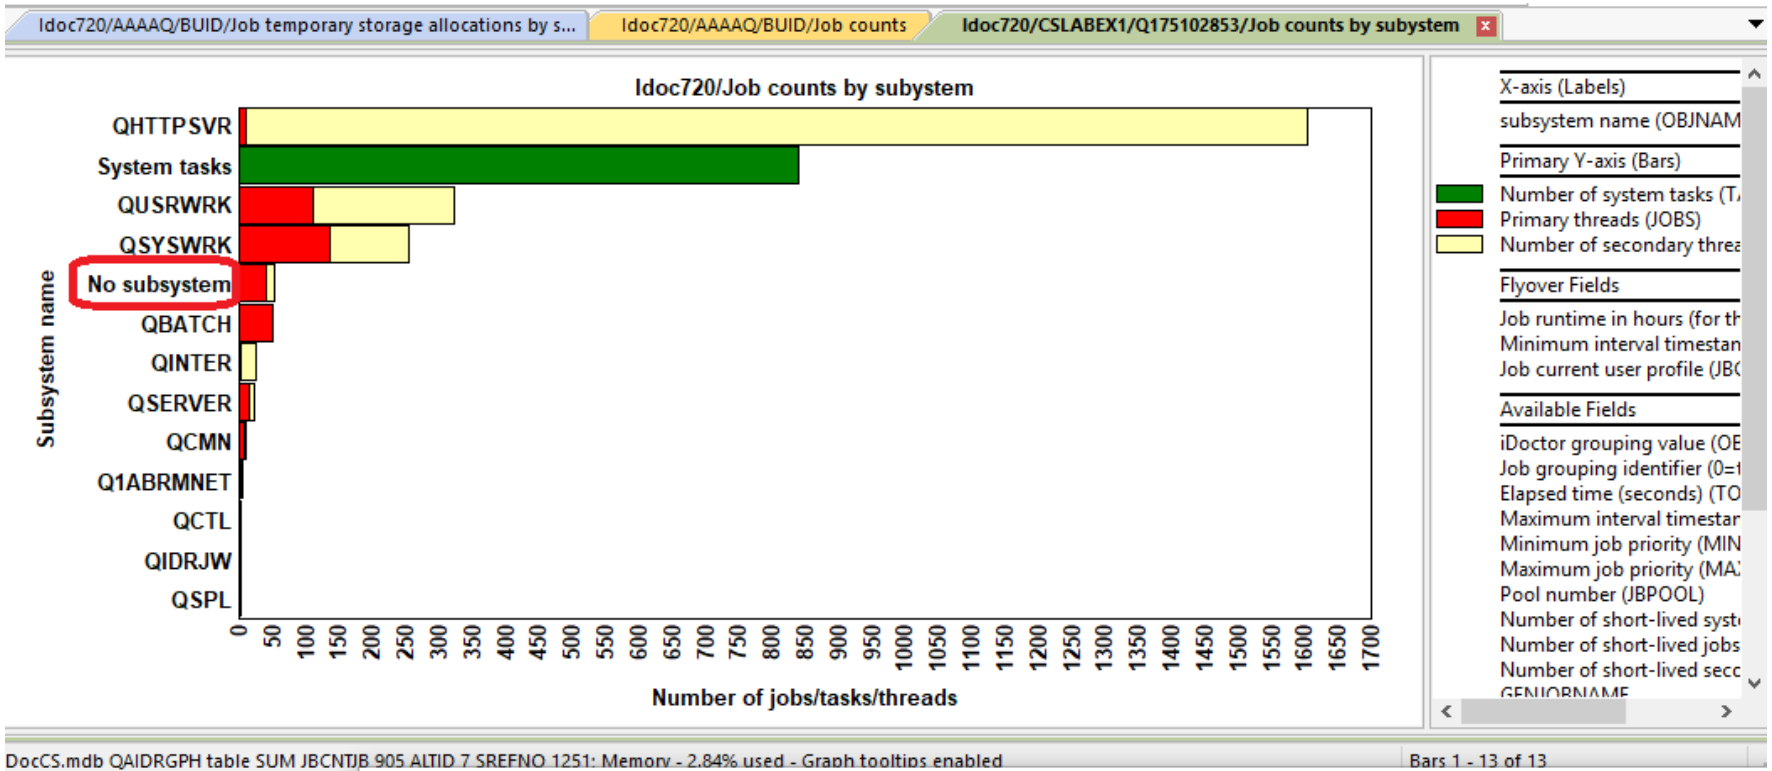

August 2018 (1326) – JW ranking graphs – “No subsystem”, “No user”…

- In JW in ranking graphs if no value is found applicable for the current grouping (i.e. subsystem, current user, etc) then instead of seeing blank you will now see text such as “No subsystem”. This provides the ability to drill down further on this entry where previously this was not possible.
  - For example you could drill down to see the jobs within the “No subsystem” grouping.