What's New with iDoctor July 2012-Oct 2013
(iDoctor GUI Builds 950-1039)

Ron McCargar
IBM i Global Support Center

Agenda

- **Overview**
  - Goals, components, unique features and website.
- **What's new with iDoctor (July 2012 – October 2013):**
- **Questions?**
  - Contact idoctor@us.ibm.com or mccargar@us.ibm.com

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 5 external and 5+ IBM internal components.**

Goals of iDoctor

- **Broaden the user base for Performance Investigation**
  - enable Operators, Programmers, IS Management
  - as well as Performance Specialists, Consultants
- **Simplify and automate processes**
- **Provide quick, immediate access to collected data**
- **Provide more analysis options**
- **Reduce the dependency on PEX traces**

iDoctor Functional Areas

Functional Area	Performance Tool
High-level system/job monitoring	**iDoctor – Collection Services Investigator** **iDoctor - HMC Walker (beta testing)** IBM i Performance Tools (PT1) Management Central Monitors
Medium-level system/job monitoring	**iDoctor – Job Watcher** WRKSYSACT IBM i Job Watcher (PT1) / STRJW
Low-level system/job tracing, stats, profiling	**iDoctor – PEX Analyzer** IBM i Performance EXplorer / PRTPEXRPT
Disk stats/tracing	**iDoctor – Disk Watcher** **iDoctor – Collection Services Investigator** **iDoctor – PEX Analyzer (PDIO analysis)** **iDoctor – VIOS Investigator** IBM i Disk Watcher (PT1) / STRDW
Plan cache analysis	**iDoctor – Plan Cache Analyzer**
Determine if SSDs could help performance	**iDoctor – Collection Services Investigator** SSD Analyzer Tool for IBM i

iDoctor GUI

- **It's a Windows client offering superior flexibility and functionality**
- **All components offer a similar user experience**
- **The latest GUI builds provide access for iDoctor components installed on servers running IBM i V5R4 or higher.**
  - Note: Latest GUI builds won't work with V5R3 or earlier systems.
- **Requirements:**
  - System i Access for Windows (not needed if only using HMC Walker)
  - .NET 4.0 or higher
  - Trial or License keys for Job Watcher (includes DW, CSI) and PEX Analyzer component usage.

iDoctor Resources

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

- **iDoctor Website:**
  http://www-912.ibm.com/i_dir/idoctor.nsf/

- **Presentations (What's New, etc):**
  http://www-912.ibm.com/i_dir/idoctor.nsf/downloadsDemos.html

- **YouTube Channel (20+ videos):**
  - https://www.youtube.com/user/IBMiDoctorForIBMi
    - These videos are also available on IBM.COM if your company blocks YouTube.
    - Just click the appropriate links titled "Video name on IBM.COM" from the Video Library pages on our website:
      - https://www-912.ibm.com/i_dir/idoctor.nsf/videos.html

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

- **Documentation:**
  https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf

iDoctor Team

- **Larry Cravens**
  - Design, Performance Expert, Taskswitch, Wait Bucket Analysis
- **Ron McCargar**
  - Design, GUI, Website, Monitors, Builds, Documentation, Videos
- **Brad Menges**
  - Design, Performance Analyst, Education (on-site visits), Sales
- **Nguyen Nguyen**
  - Design, QMGTOOLS, Performance Analyst
- **Paul Stimets**
  - Sales
- **Shane Smith**
  - Design, GUI and server side development
- **Chris Pilcher (Business Manager)**
- **Lab Services (Sales Contacts)**
  - Karen Anderson, Mark Even (USA)
  - Jin-Ming Liu (AP)
  - Virginie Cohen, Jean-Francois Soulard (EMEA)

What's New with iDoctor: Overview

- **General GUI Enhancements**
- **Installation**
- **Must Gather Tools**
- **HMC Walker**
- **Job Watcher**
- **Collection Services Investigator**
- **PEX Analyzer**

General – New Components Added Oct 2013

Must Gather Tools (QMGTOOLS) is now available externally. This component will assist support with data collection of needed metrics to help solve performance problems.

HMC Walker is now in beta testing. It primarily captures configuration data from the HMC and lslparutil data (CPU utilization, memory, etc) across all LPARs attached to the HMC from all physical systems.

Please send comments to idoctor@us.ibm.com

General – Requirements

- **.NET 4.0 or higher**
  - –http://www.microsoft.com/en-us/download/details.aspx?id=30653
- **Visual C++ Redistributable for Visual Studio 2012 Update 1 or higher**
  - –http://www.microsoft.com/en-us/download/details.aspx?id=30679
  - –(install the 32-bit version)
- **System i Access for Windows (optional)**
  - –Note: It's required for everything except HMC Walker.
  - –http://www-03.ibm.com/systems/power/software/i/access/windows_sp.html

  If on Windows 7 or higher System i Access for Windows will also need:

  http://www.microsoft.com/en-us/download/details.aspx?id=26347

  Install both 32-bit and 64-bit versions of the above if you have 64-bit Windows installed.
- **Oracle Express edition (if desired to analyze HMC/AIX performance data on the PC)**
  - –http://www.oracle.com/technetwork/products/express-edition/downloads/index.html

General – Transfer function updates

- **Sending data to IBM will now use IBMSDDUU (secure method.)**

- **Added options to download/upload data using SSL FTP.**

General – Google search options

- **Added Google search options embedded within the GUI in several places:**
  - Update history window
  - Selected text in a table view
  - Check PTFs window (select the PTF, then right-click to search for it)
  - Error windows

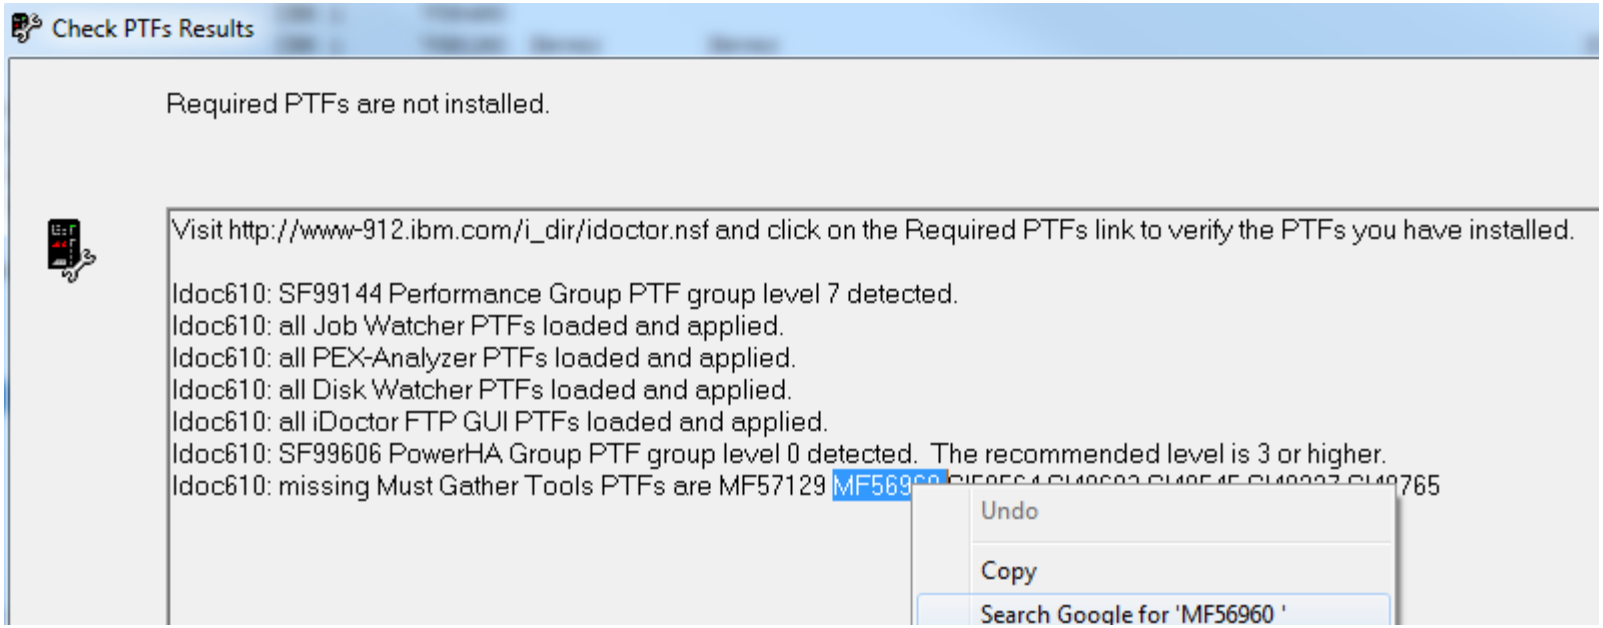

General – New Preferences

- On the Display tab you can now specify a character limit for the x axis label on horizontal bar graphs.

- On the Display tab, added an option called “Force legend width percent on resize”. For example if the legend width percent is 20%, any time the graph is resized, the legend will always remain at 20% of the total width of the graph window if the option is checked.

- Added a tab called IBMSDDUU to configure options related to sending data to IBM.

- On the SQL tab added preferences for the estimated timeout limit, query temp storage limit estimate and the QAQQINI options file library.

General – Changed Preferences

- The “Display advanced reporting options” preference on the Data Viewer tab has been moved to the Miscellaneous tab. It was also renamed to “Display advanced options”.

- The Super collections and Browse Collections folders in iDoctor are now only visible if the “Display advanced options” preference on the miscellaneous tab is checked (unchecked by default.) This option also controls whether or not the “Detail reports” drill down menu in Job Watcher is shown.

General – iDoctor client job preference

- Added a preference called "Remove libraries above QSYS in the library list (requires *ALLOBJ.)" If used, this effects all QZRCSRVS and QZDASOINIT jobs started by iDoctor. The default is unchecked. You can use this if the customer environment you are on, has non standard libraries/commands above QSYS in the library list causing issues.

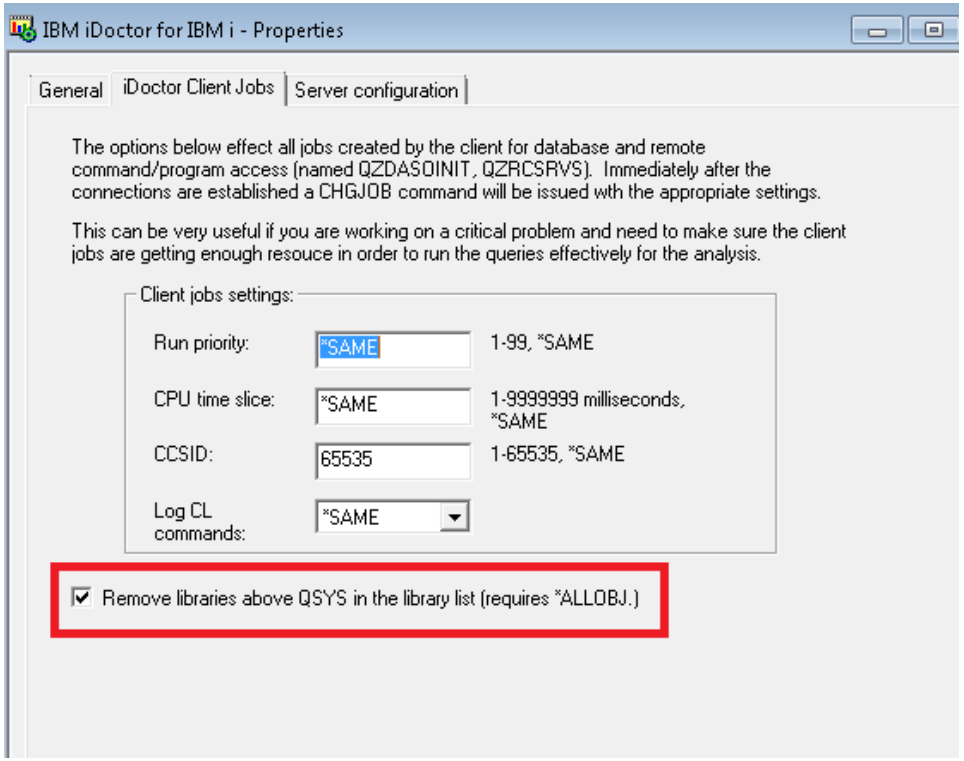

General – iDoctor temp directory

- The application properties General tab will now list the iDoctor temp directory. This folder contains several log files which may be needed for debug purposes.

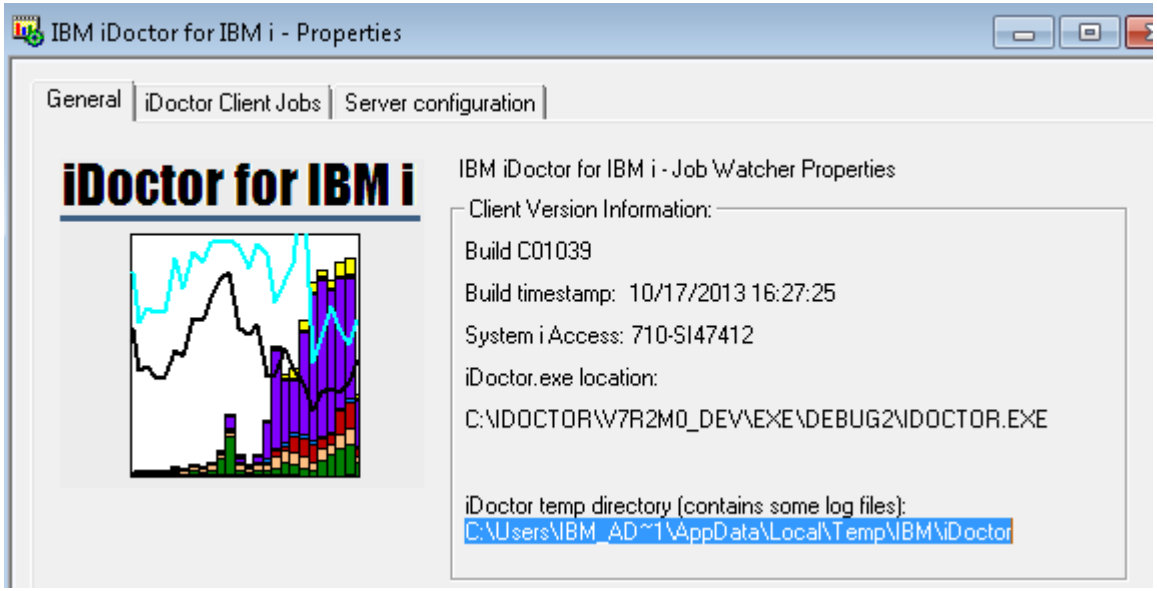

General – Debug report location listed in status bar of Data Viewer

- When sending a bug report, please include the status bar of the Data Viewer in your screenshots which indicates where in the repository of graphs and tables your report resides. This will make it easier to find and correct the issue.

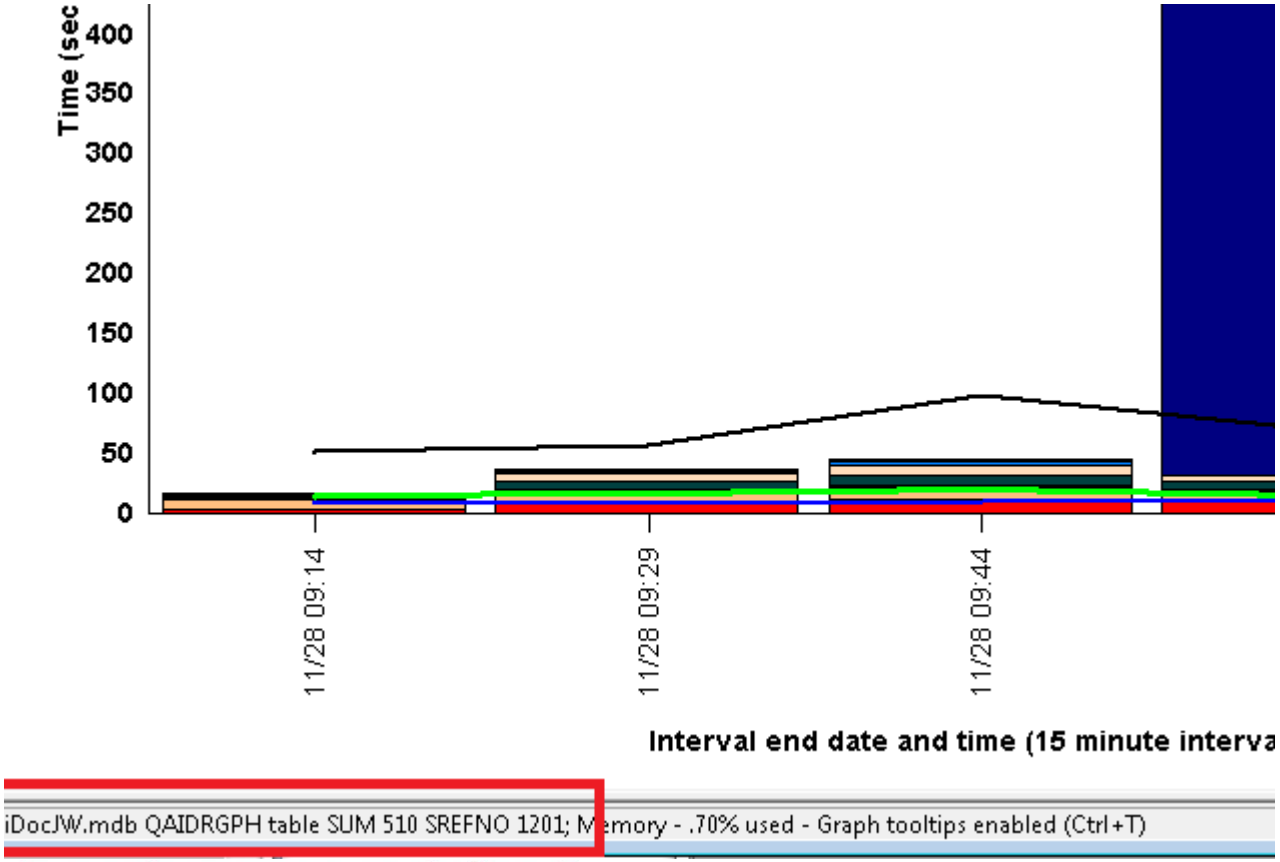

General – Debug messages in QIDRGUI/ADDPRDACS command

- When the command is ran via the green screen it will now list messages regarding the success/failure of any access code applied.

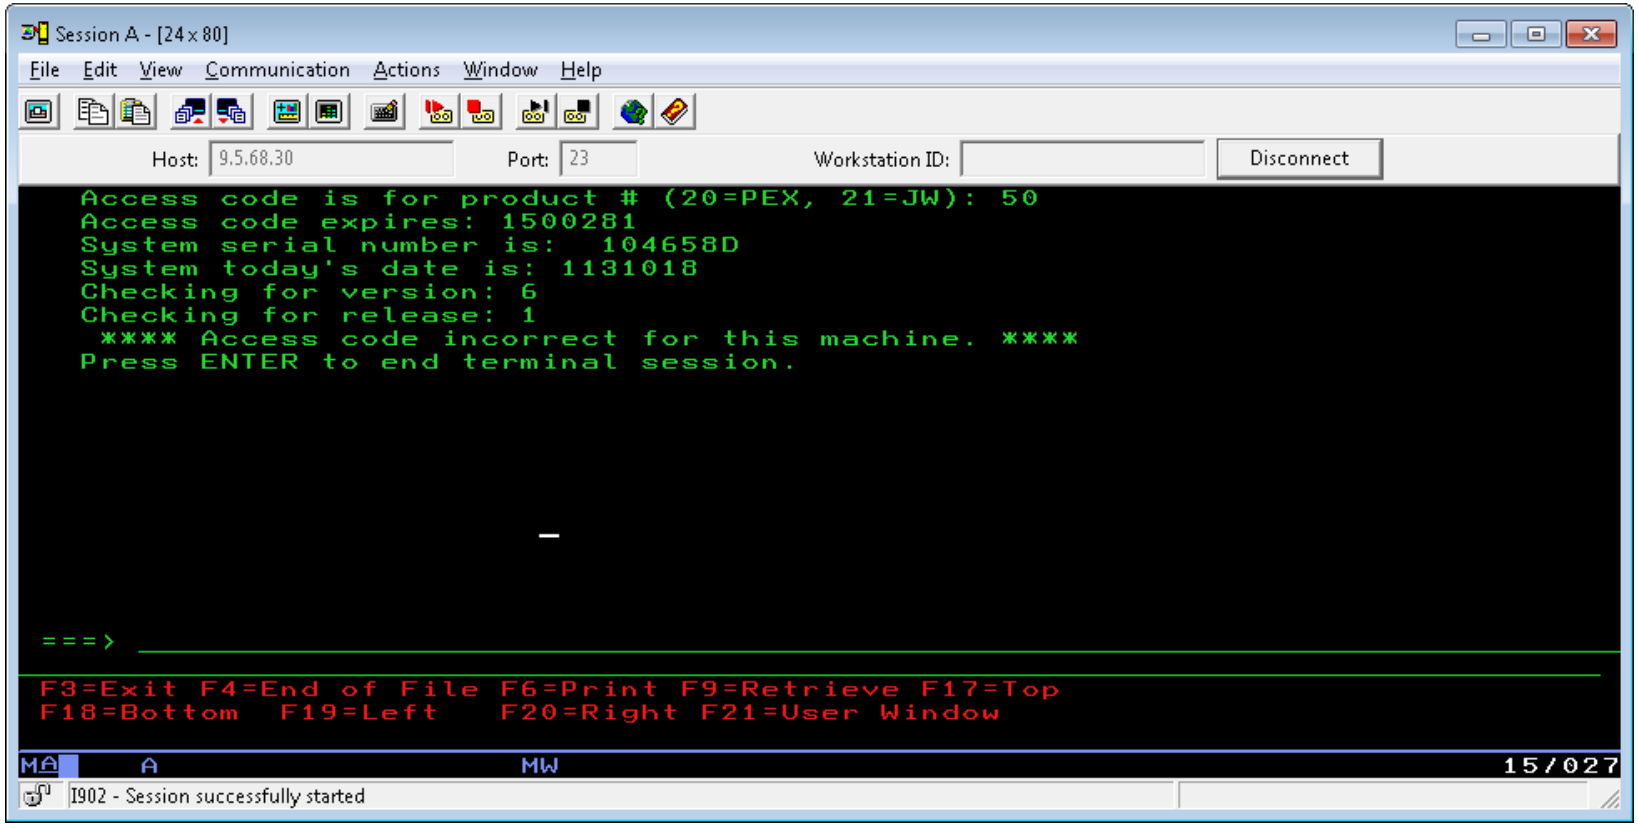

General – Multi selection in graph legend

- You can now select multiple items in the graph legend and perform actions against them.

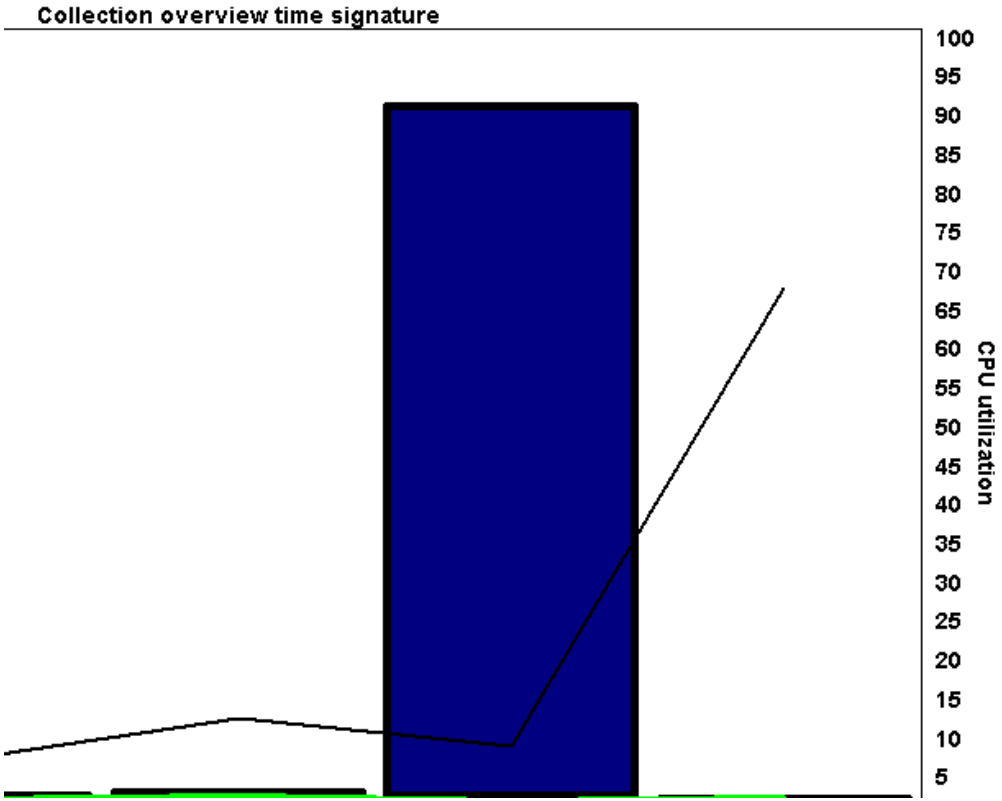

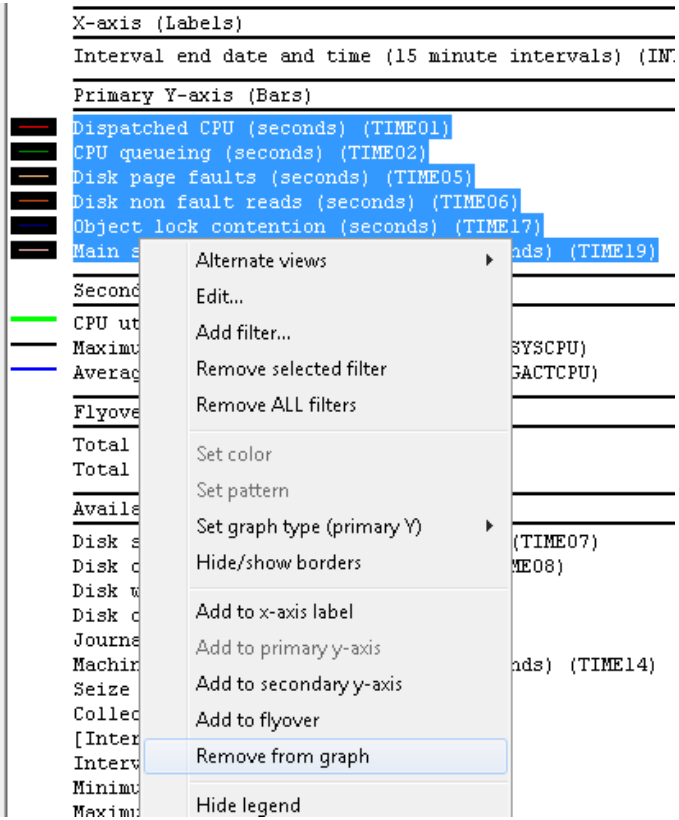

Installation – Check new server builds

- On the components window, the build date column tells you how old your server builds are.

- Use the Check for new server builds button to install the latest server builds. This will suspend the GUI, auto install the latest server builds with default options, then resume the GUI.

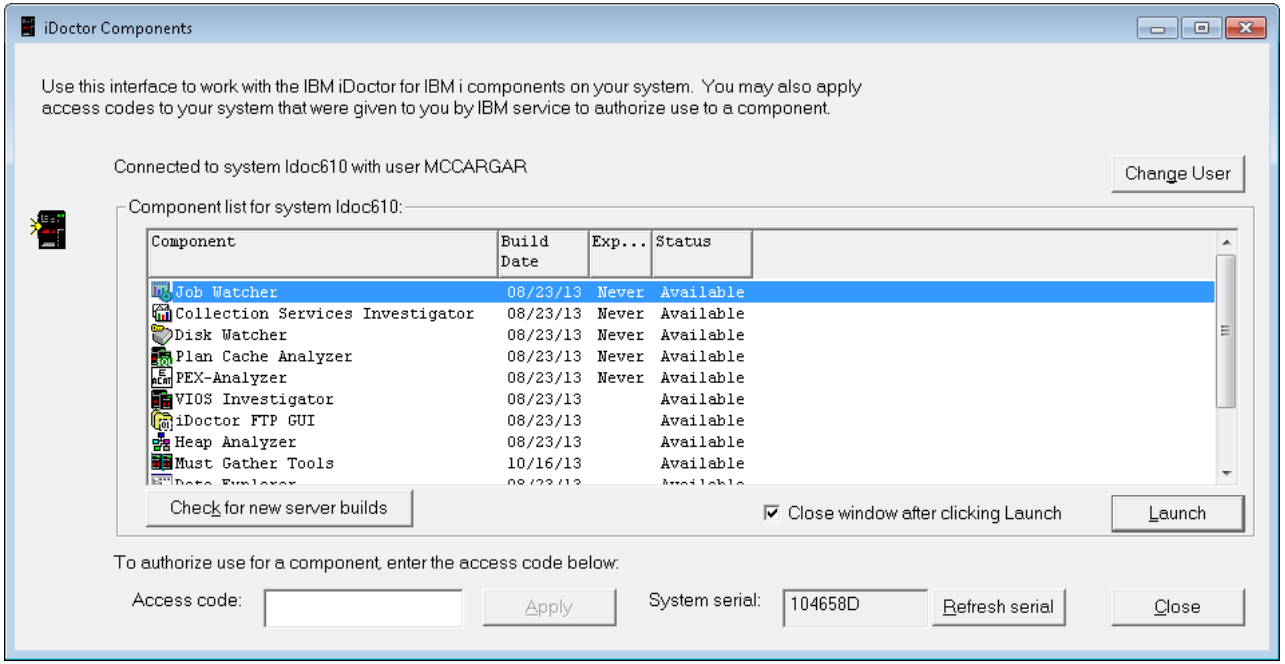

Installation – Updates

- The server builds are now downloaded when the install runs at run time from the iDoctor FTP site. There are settings to use FTP or HTTP. Settings to specify your proxy server with user/pwd are also provided if this is required for your environment.

- Updated the license agreement page to show all of them that now apply.

- Added a checkbox on the component selection screen that indicates if the job queue and subsystem screens should be shown (otherwise the default settings are used.)

- Updated the FTP connection screen to include an option to use SSL FTP when uploading the save files to an IBM i.

Must Gather Tools

- A free offering, but it's also a work in progress. Some functions shown may not work yet.

- Available at 6.1 or higher only.

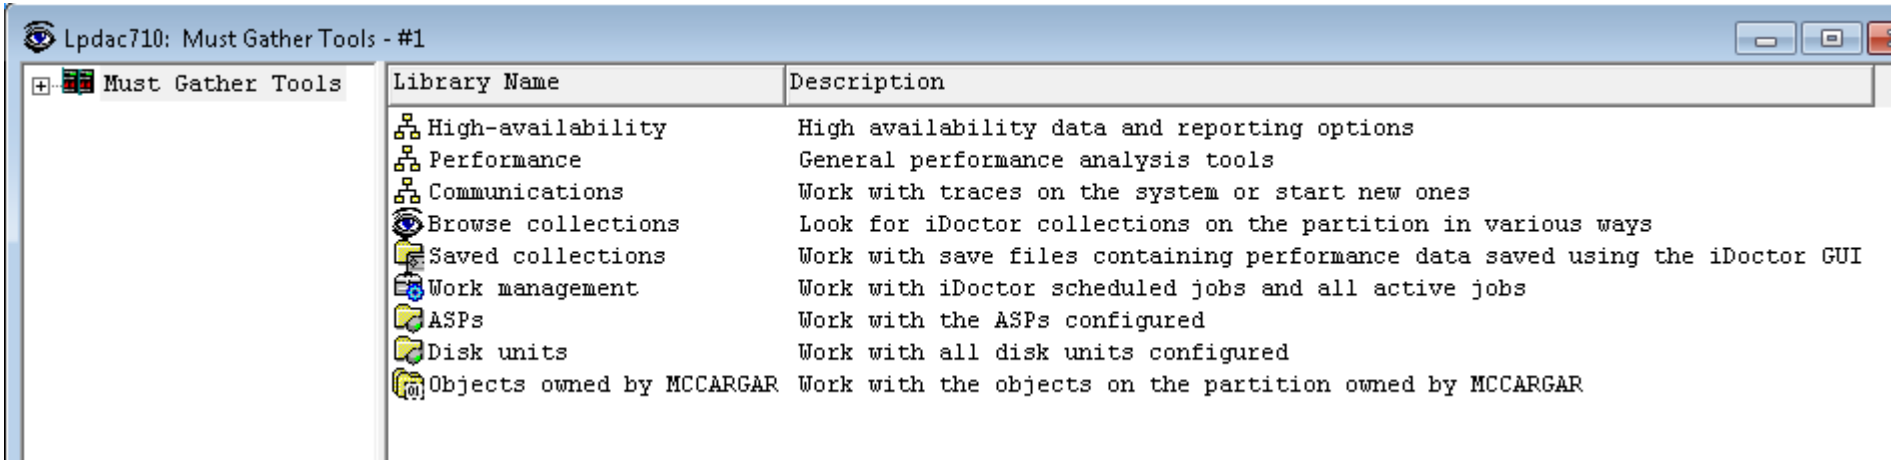

Must Gather Tools - Menu

- The menus found on the Must Gather Tools folder, and the subfolder below it are designed to match the GO MG menu found on the green screen.

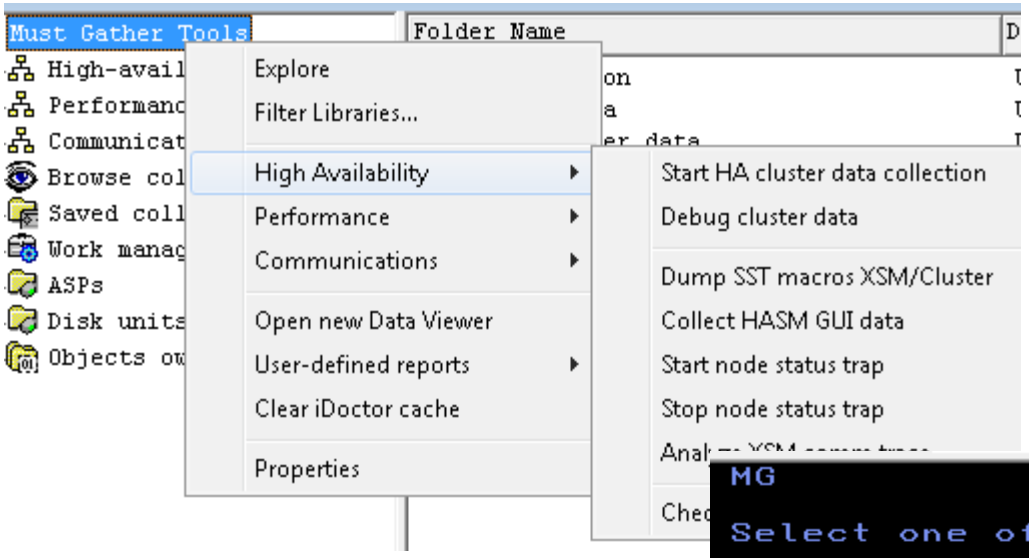

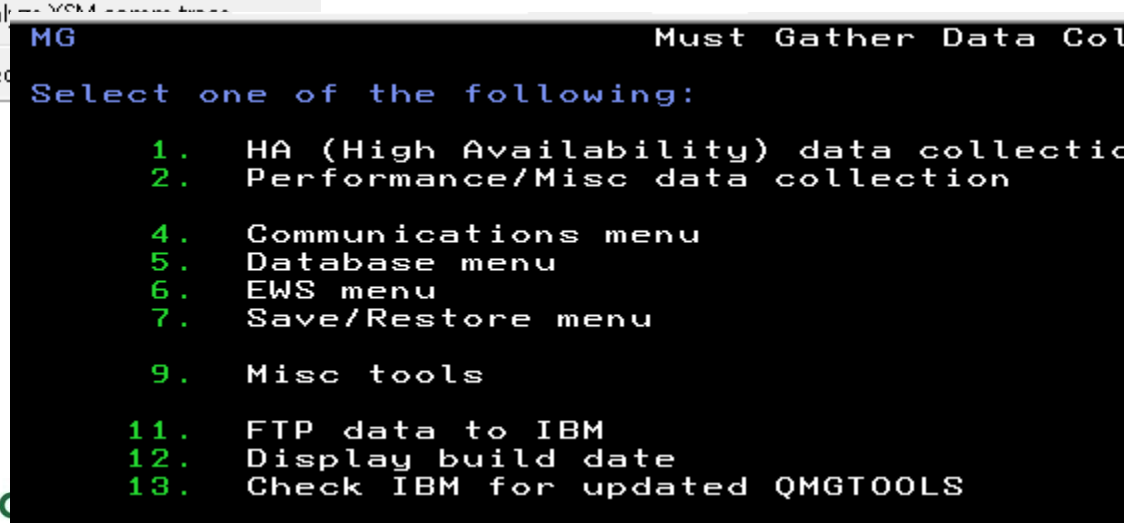

Must Gather Tools – High availability Menu

- This menu is designed to match the GO QHASTOOLS menu found on the green screen and provides all the HA options in Must Gather Tools.

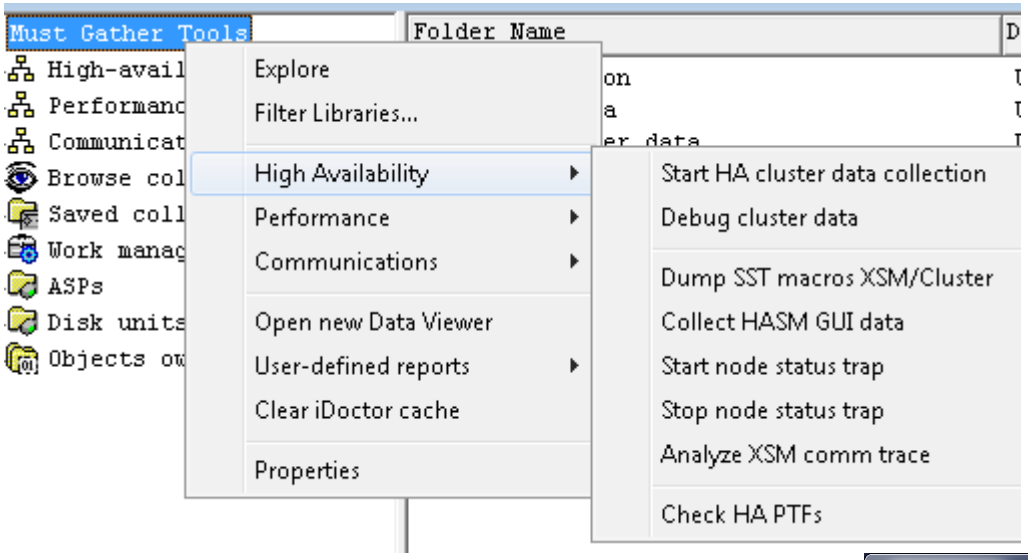

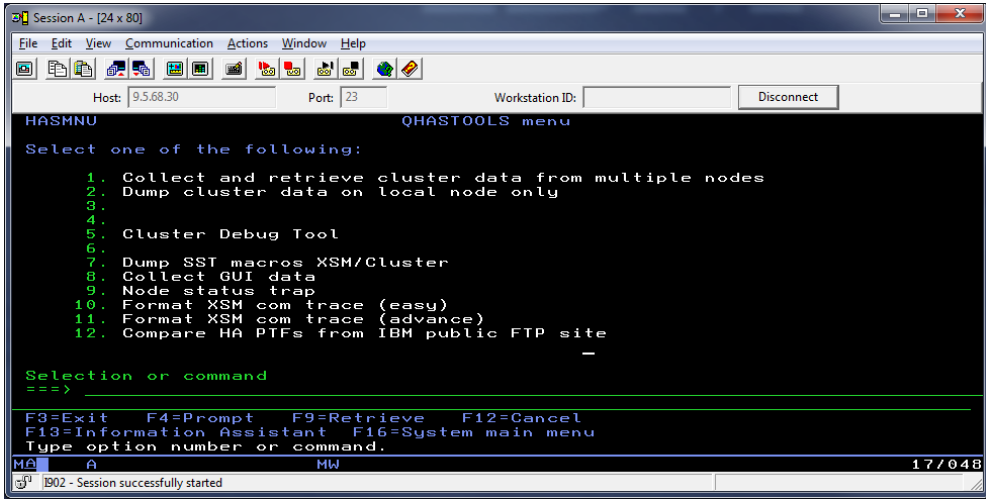

Must Gather Tools – Cluster data capture

- High availability -> Start HA cluster data collection option

- By default it will connect to all nodes and collect cluster data. GUI then downloads it all back to the PC and zips it up.

- Next data is sent to IBM.

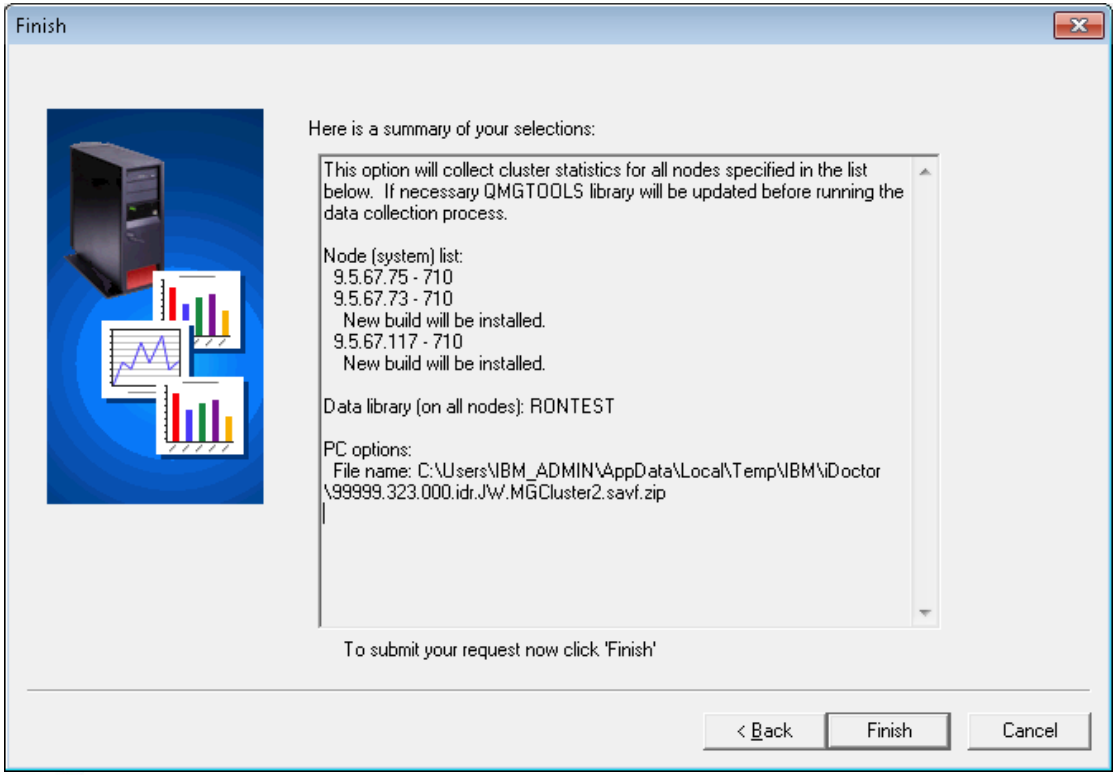

Must Gather Tools – Work with cluster data

- High availability -> Cluster data folder to work with captured data.

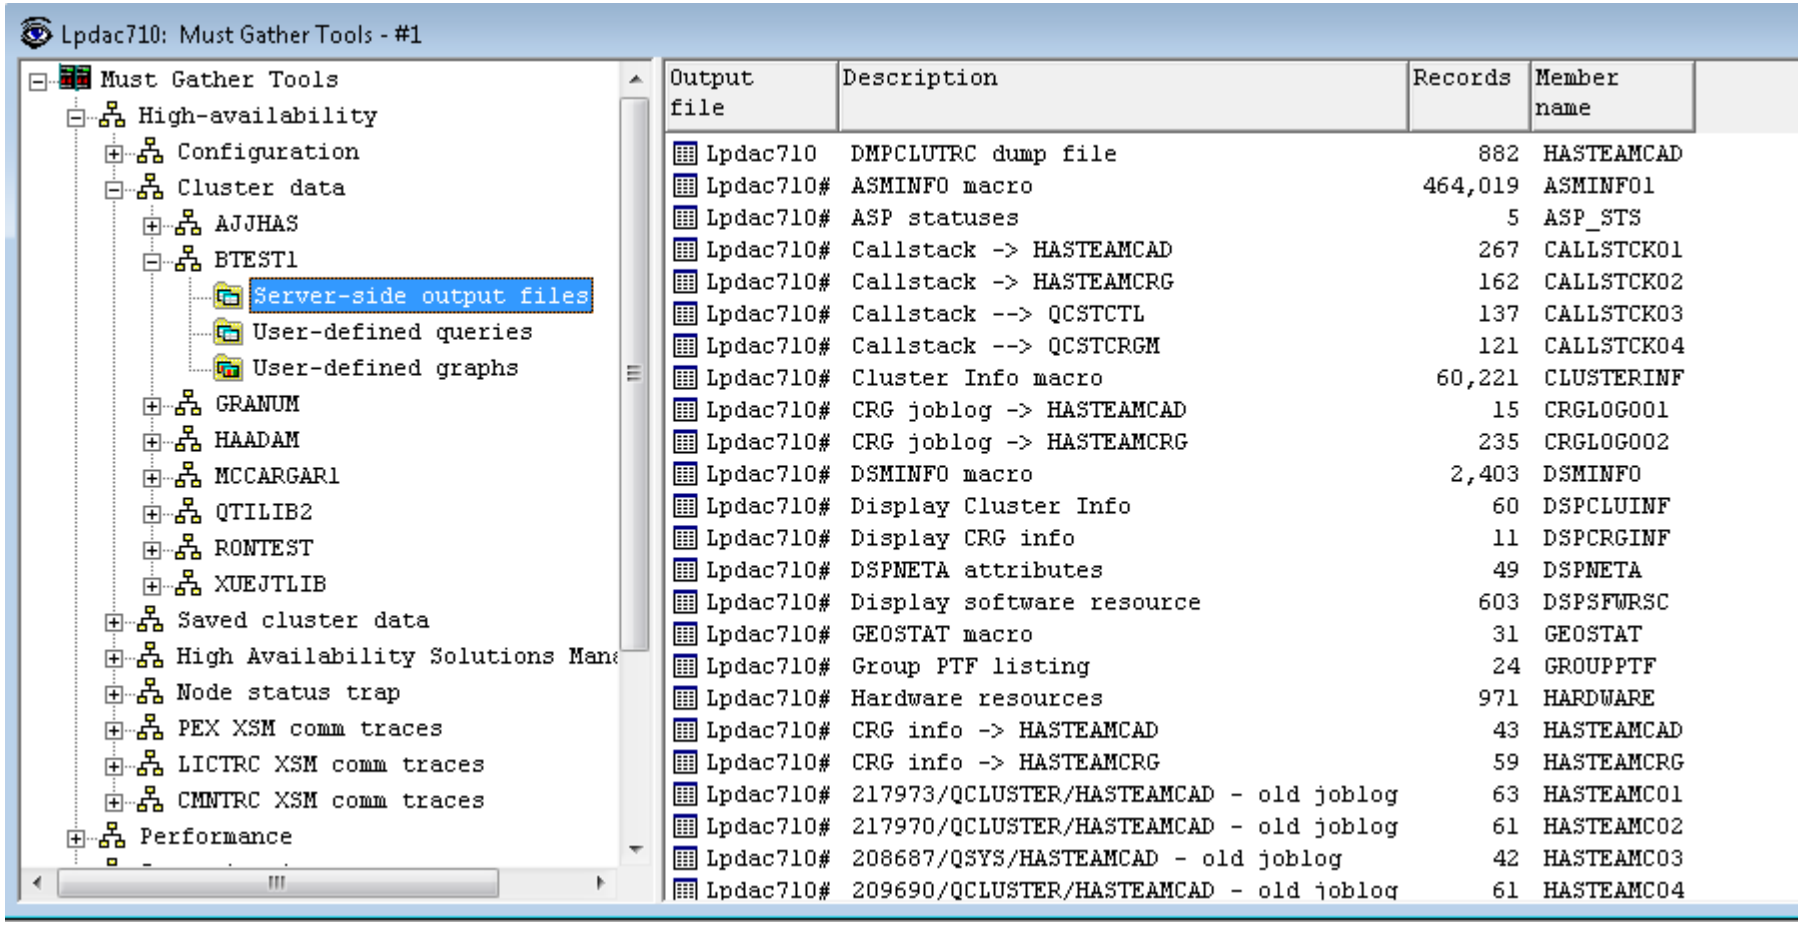

Output file	Description	Records	Member name
Lpdac710	DMPCLUTRC dump file	882	HASTEAMCAD
Lpdac710#	ASMINFO macro	464,019	ASMINFO1
Lpdac710#	ASP statuses	5	ASP_STS
Lpdac710#	Callstack -> HASTEAMCAD	267	CALLSTCK01
Lpdac710#	Callstack -> HASTEAMCRG	162	CALLSTCK02
Lpdac710#	Callstack --> QCSTCTL	137	CALLSTCK03
Lpdac710#	Callstack --> QCSTCRGM	121	CALLSTCK04
Lpdac710#	Cluster Info macro	60,221	CLUSTERINF
Lpdac710#	CRG joblog -> HASTEAMCAD	15	CRGLOG001
Lpdac710#	CRG joblog -> HASTEAMCRG	235	CRGLOG002
Lpdac710#	DSMINFO macro	2,403	DSMINFO
Lpdac710#	Display Cluster Info	60	DSPCLUINF
Lpdac710#	Display CRG info	11	DSPCRGINF
Lpdac710#	DSPNETA attributes	49	DSPNETA
Lpdac710#	Display software resource	603	DSPSFWRSC
Lpdac710#	GEOSTAT macro	31	GEOSTAT
Lpdac710#	Group PTF listing	24	GROUPPTF
Lpdac710#	Hardware resources	971	HARDWARE
Lpdac710#	CRG info -> HASTEAMCAD	43	HASTEAMCAD
Lpdac710#	CRG info -> HASTEAMCRG	59	HASTEAMCRG
Lpdac710#	217973/QCLUSTER/HASTEAMCAD - old joblog	63	HASTEAMC01
Lpdac710#	217970/QCLUSTER/HASTEAMCAD - old joblog	61	HASTEAMC02
Lpdac710#	208687/QSYS/HASTEAMCAD - old joblog	42	HASTEAMC03
Lpdac710#	209690/QCLUSTER/HASTEAMCAD - old joblog	61	HASTEAMC04

Must Gather Tools – Debug cluster data

- High availability -> Debug cluster data menu will produce a report/analysis of the captured data in the specified library.

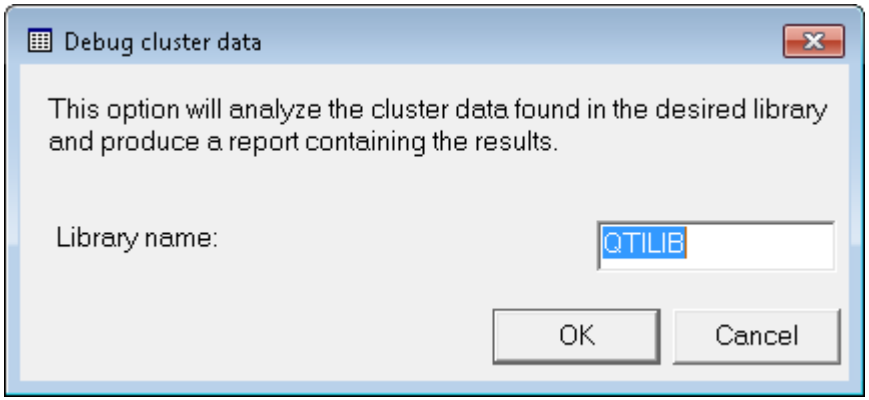

Must Gather Tools – Saved cluster data

- High availability -> Saved cluster data folder can be used to view the save files of cluster data already captured. You can restore/delete/transfer to IBM this data.

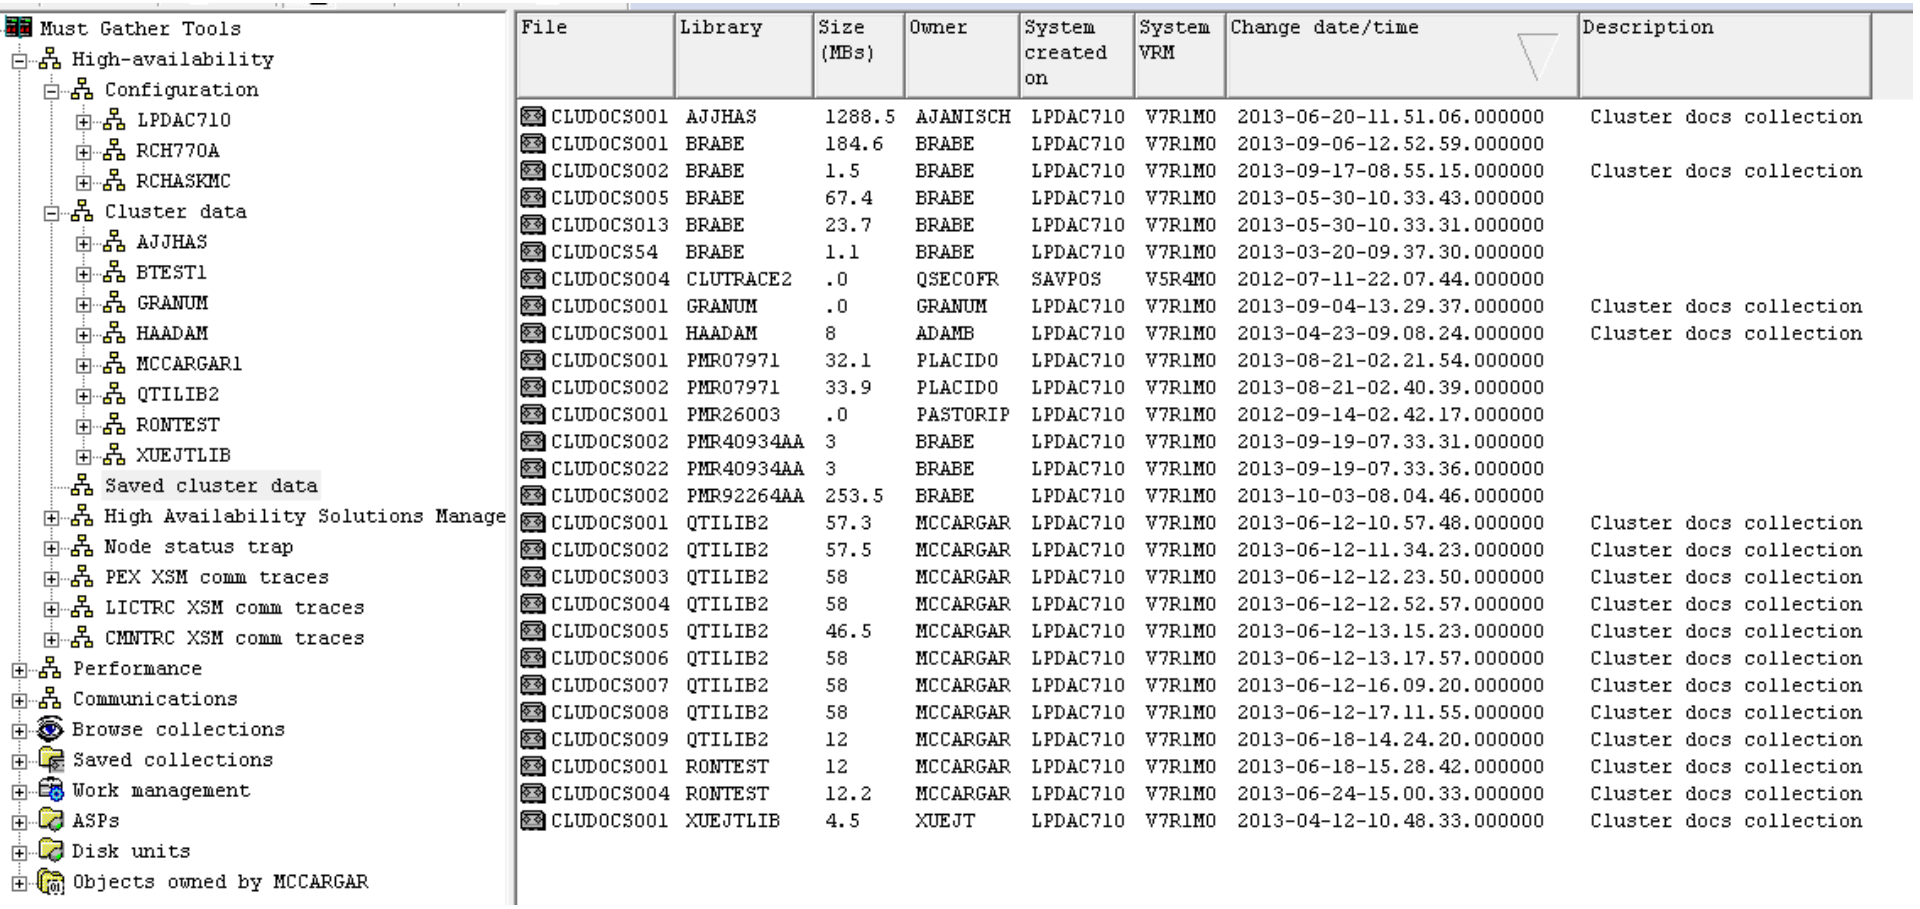

File	Library	Size (MBs)	Owner	System created on	System VRM	Change date/time	Description
CLUDOCS001	AJJHAS	1288.5	AJANISCH	LPDAC710	V7R1M0	2013-06-20-11.51.06.000000	Cluster docs collection
CLUDOCS001	BRABE	184.6	BRABE	LPDAC710	V7R1M0	2013-09-06-12.52.59.000000	
CLUDOCS002	BRABE	1.5	BRABE	LPDAC710	V7R1M0	2013-09-17-08.55.15.000000	Cluster docs collection
CLUDOCS005	BRABE	67.4	BRABE	LPDAC710	V7R1M0	2013-05-30-10.33.43.000000	
CLUDOCS013	BRABE	23.7	BRABE	LPDAC710	V7R1M0	2013-05-30-10.33.31.000000	
CLUDOCS54	BRABE	1.1	BRABE	LPDAC710	V7R1M0	2013-03-20-09.37.30.000000	
CLUDOCS004	CLUTRACE2	.0	QSECOFR	SAVPOS	V5R4M0	2012-07-11-22.07.44.000000	
CLUDOCS001	GRANUM	.0	GRANUM	LPDAC710	V7R1M0	2013-09-04-13.29.37.000000	Cluster docs collection
CLUDOCS001	HAADAM	8	ADAMB	LPDAC710	V7R1M0	2013-04-23-09.08.24.000000	Cluster docs collection
CLUDOCS001	PMR07971	32.1	PLACIDO	LPDAC710	V7R1M0	2013-08-21-02.21.54.000000	
CLUDOCS002	PMR07971	33.9	PLACIDO	LPDAC710	V7R1M0	2013-08-21-02.40.39.000000	
CLUDOCS001	PMR26003	.0	PASTORIP	LPDAC710	V7R1M0	2012-09-14-02.42.17.000000	
CLUDOCS002	PMR40934AA	3	BRABE	LPDAC710	V7R1M0	2013-09-19-07.33.31.000000	
CLUDOCS022	PMR40934AA	3	BRABE	LPDAC710	V7R1M0	2013-09-19-07.33.36.000000	
CLUDOCS002	PMR92264AA	253.5	BRABE	LPDAC710	V7R1M0	2013-10-03-08.04.46.000000	
CLUDOCS001	QTILIB2	57.3	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-10.57.48.000000	Cluster docs collection
CLUDOCS002	QTILIB2	57.5	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-11.34.23.000000	Cluster docs collection
CLUDOCS003	QTILIB2	58	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-12.23.50.000000	Cluster docs collection
CLUDOCS004	QTILIB2	58	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-12.52.57.000000	Cluster docs collection
CLUDOCS005	QTILIB2	46.5	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-13.15.23.000000	Cluster docs collection
CLUDOCS006	QTILIB2	58	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-13.17.57.000000	Cluster docs collection
CLUDOCS007	QTILIB2	58	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-16.09.20.000000	Cluster docs collection
CLUDOCS008	QTILIB2	58	MCCARGAR	LPDAC710	V7R1M0	2013-06-12-17.11.55.000000	Cluster docs collection
CLUDOCS009	QTILIB2	12	MCCARGAR	LPDAC710	V7R1M0	2013-06-18-14.24.20.000000	Cluster docs collection
CLUDOCS001	RONTEST	12	MCCARGAR	LPDAC710	V7R1M0	2013-06-18-15.28.42.000000	Cluster docs collection
CLUDOCS004	RONTEST	12.2	MCCARGAR	LPDAC710	V7R1M0	2013-06-24-15.00.33.000000	Cluster docs collection
CLUDOCS001	XUEJTLIB	4.5	XUEJT	LPDAC710	V7R1M0	2013-04-12-10.48.33.000000	Cluster docs collection

Must Gather Tools – High Availability Solutions Manager GUI

- High availability -> High Availability Solutions Manager GUI folder is used to work with the zip files in the IFS that contain debug data for the HASM GUI.

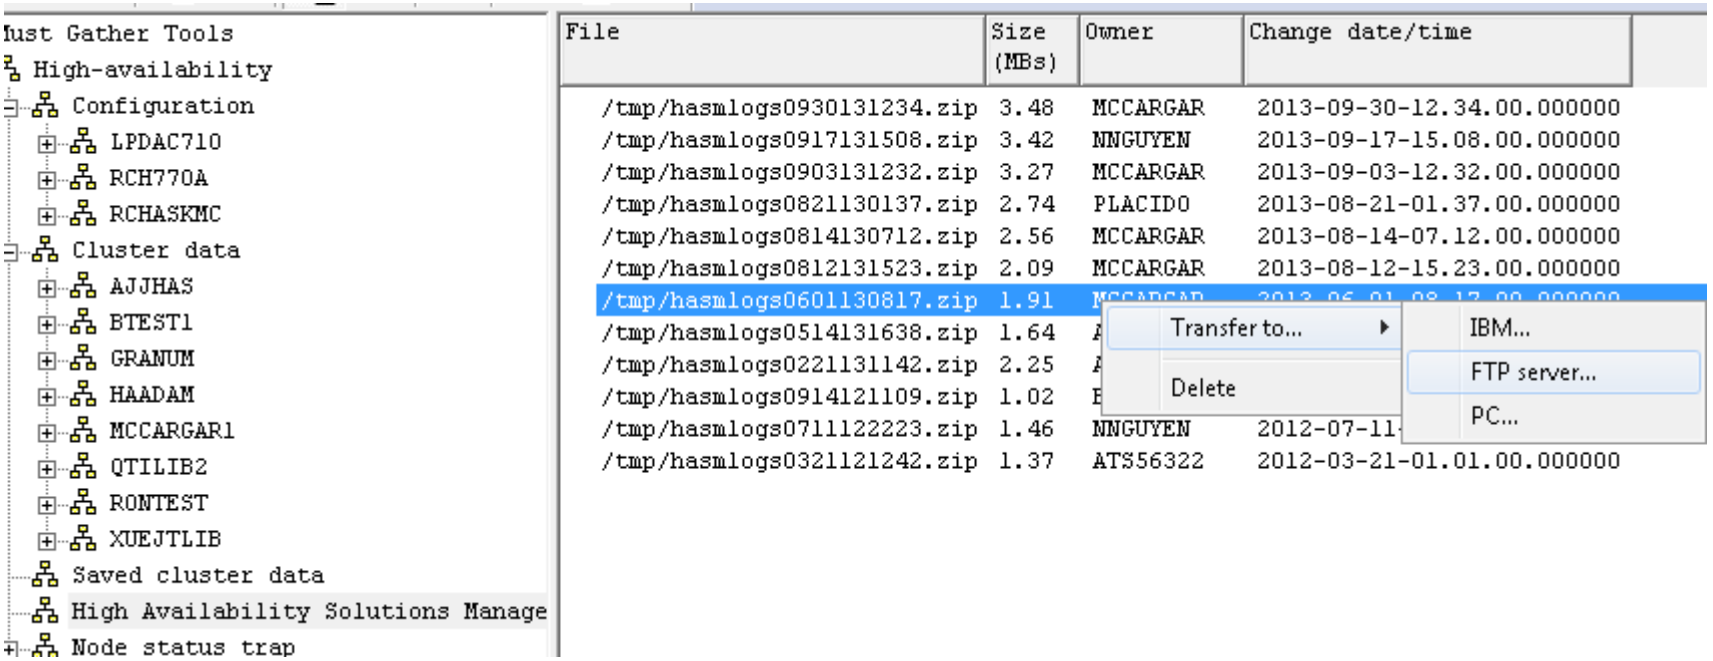

Must Gather Tools – Performance menu

- The menu provides access to similar options available via the GO QPERF command

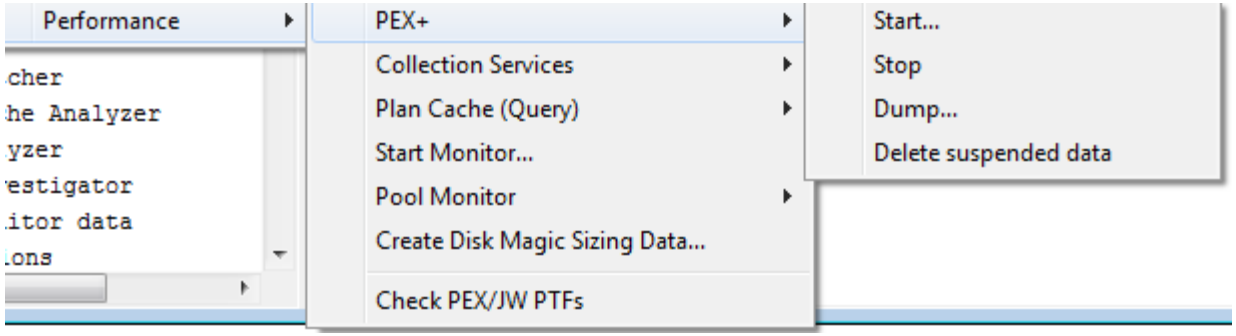


```
Performance Menu                    QPERF Menu

Select one of the following:

      1.  Combined PEX/JobWatcher/Collection Service trace   PEX+
      2.
      3.  Collections Services                                 _
      4.  Query Performance
      5.  Monitors
      6.  PEX clean up
      7.
      8.  PEX status (PEX started by QMGTOOLS)
      9.
     10.  Compare PEX/JW PTFs from IBM public FTP site
     11.  POOL Info Monitor
     12.  Gather Disk Magic sizing data
     13.  Job Watcher Status
     14.  Remote Command Exit Program
Selection or command
```

Must Gather Tools – Performance folder

- The options shown under the Performance folder will vary depending upon the license keys applied to the system. The folders such as Collection Services Investigator and Job Watcher provide ALL GUI functions for the applicable component (inside of the Must Gather Tools GUI.)
  - PEX+ (not yet implemented)
  - Monitors (requires JW and PEX license)
  - Collection Services Investigator (JW license)
  - Job Watcher (JW license)
  - Disk Watcher (JW license)
  - Plan Cache Analyzer (JW license)
  - PEX-Analyzer (PEX license)
  - VIOS Investigator
  - Pool monitor data

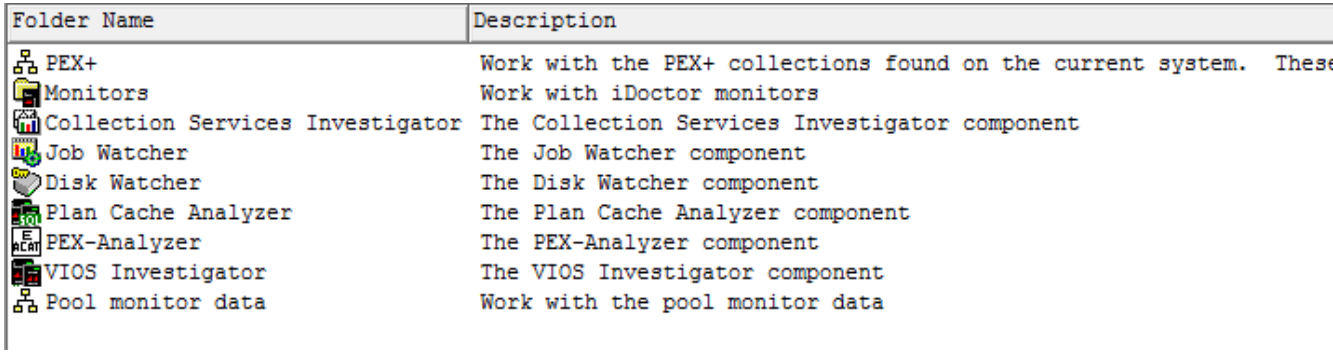

Folder Name	Description
PEX+	Work with the PEX+ collections found on the current system. These
Monitors	Work with iDoctor monitors
Collection Services Investigator	The Collection Services Investigator component
Job Watcher	The Job Watcher component
Disk Watcher	The Disk Watcher component
Plan Cache Analyzer	The Plan Cache Analyzer component
PEX-Analyzer	The PEX-Analyzer component
VIOS Investigator	The VIOS Investigator component
Pool monitor data	Work with the pool monitor data

Must Gather Tools – Performance folder

- Performance – Monitors folder

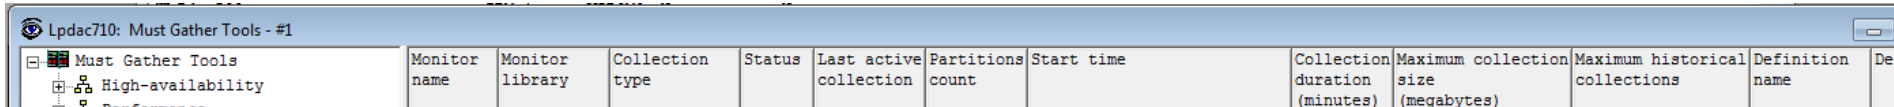

Monitor name	Monitor library	Collection type	Status	Last active collection	Partitions count	Start time	Collection duration (minutes)	Maximum collection size (megabytes)	Maximum historical collections	Definition name
JWMON	PMR16073MN	Job Watcher	Ended	JWMON095		2013-09-30-11.55.02.211450	60	8192	5	OPTMON
KSE	KEDWARDS	Job Watcher	Ended	KSE001		2013-09-11-09.32.15.362173	60	2096	3	Q10SEC
T	QIDRDATA	Job Watcher	Ended	T147		2013-09-04-18.36.29.119445	60	4096	5	Q5SECSQL
AAA	MCCARGAR1	Job Watcher	Ended	AAA001		2013-03-11-15.07.23.815504	60	4096	5	ABC
PAMON	IBMPEX	PEX-Analyzer	Ended	PAMON003		2012-05-20-22.42.07.751516	5		40	TPROFFMT2
PEXMON	PEXMONIBM	PEX-Analyzer	Ended	PEXMON515		2012-01-03-16.12.26.448961	2		15	FRONTMON

- Collection Services Investigator

Libraries	Libraries containing Collection Services Investigator collections (filterable)
Historical summaries	All data generated by the Historical Summary analysis (or STRCSMON command) on the system
CS objects	A list of all Collection Services management collection objects on the system
SQL tables	Work with the SQL-based tables generated by iDoctor analysis processes (library filterable)

- Job Watcher

Libraries	Libraries containing Job Watcher collections (filterable)
Definitions	Work with definitions used for creating collections
SQL tables	Work with the SQL-based tables generated by iDoctor analysis processes (library filterable)

- Plan Cache Analyzer

Plan cache snapshots	A list of all Plan Cache Analyzer snapshots on the system
Plan cache dumps	A list of all Plan Cache Analyzer dumps on the system
SQL performance monitors	Contains all DBMON data found on the current system
SQL tables	Work with the SQL-based tables generated by iDoctor analysis processes (library filterable)

Must Gather Tools – Performance – Pool Monitor - Start

- Use this menu to start a pool monitor collection.

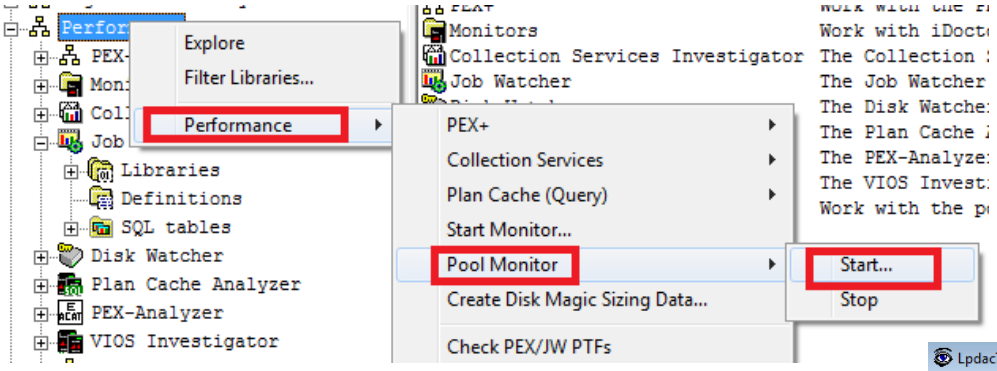

- Which shows this interface:

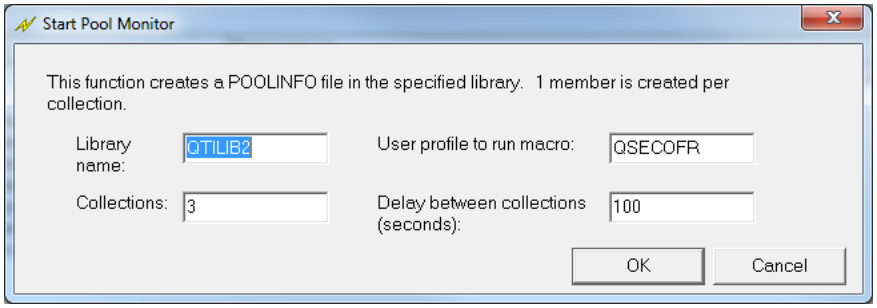

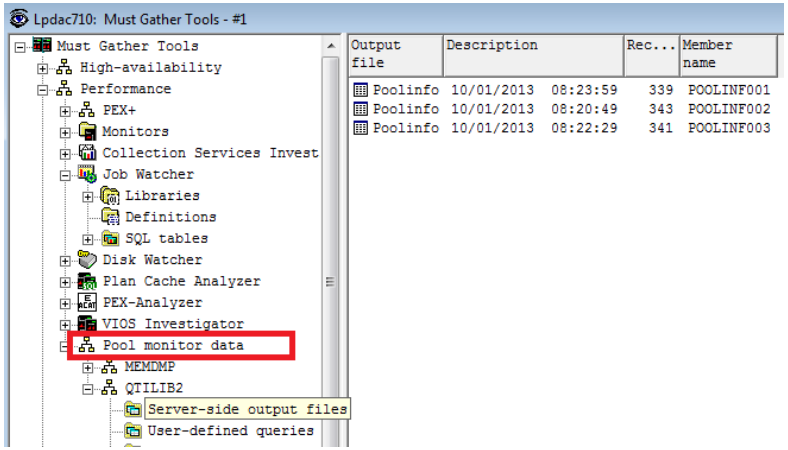

- And you can view the results under Performance – Pool Monitor data

Must Gather Tools – Performance – Pool Monitor report

- An example of the data returned by the Pool Monitor

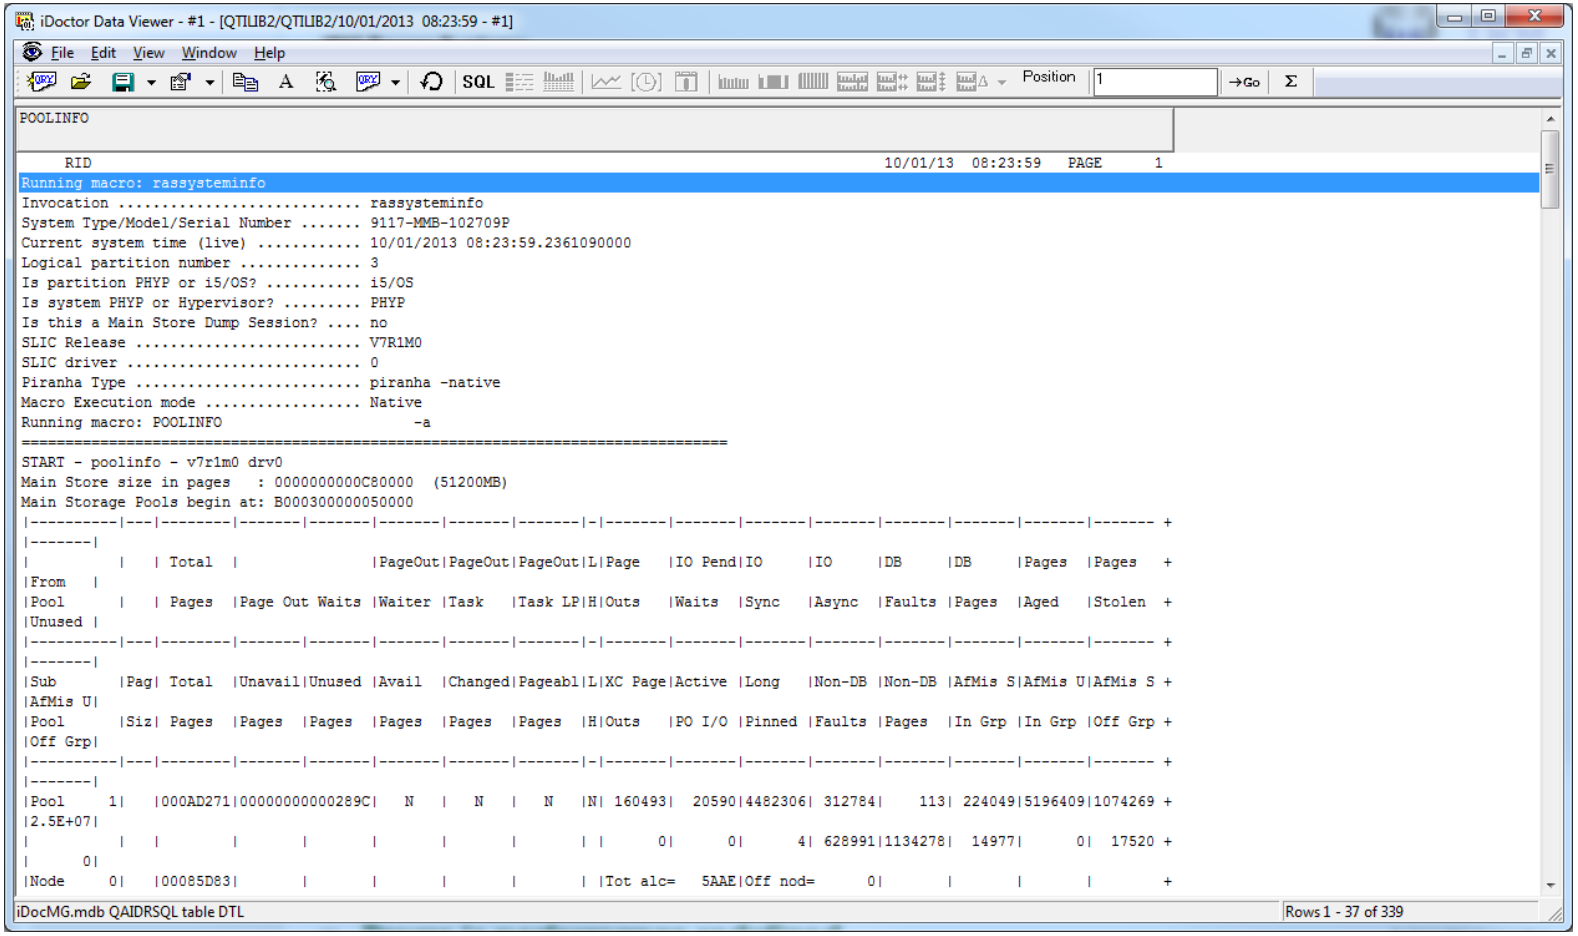

Power is performance redefined

Must Gather Tools – Performance – Create Disk Magic Sizing Data

- Use this menu to create Disk Magic sizing data

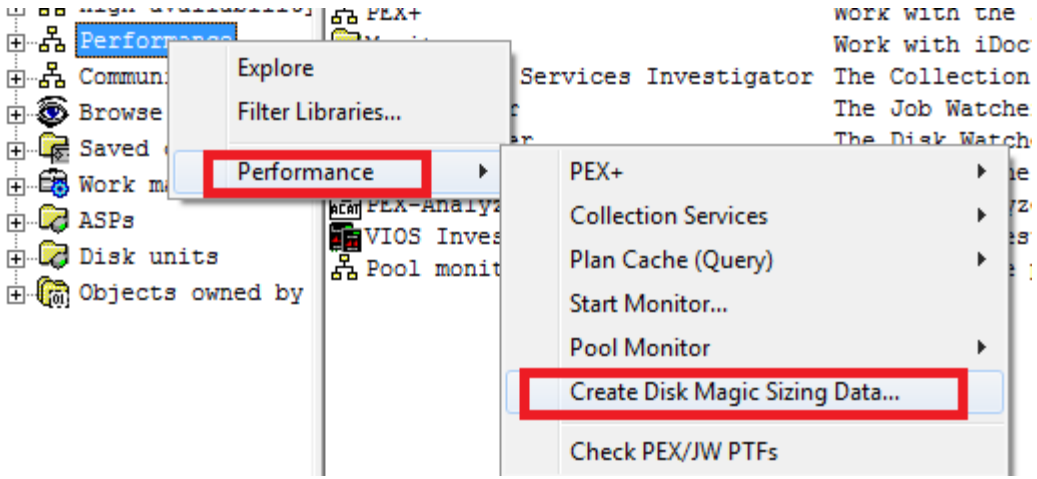

- Which shows this interface: (you must provide the library and collection name to process)

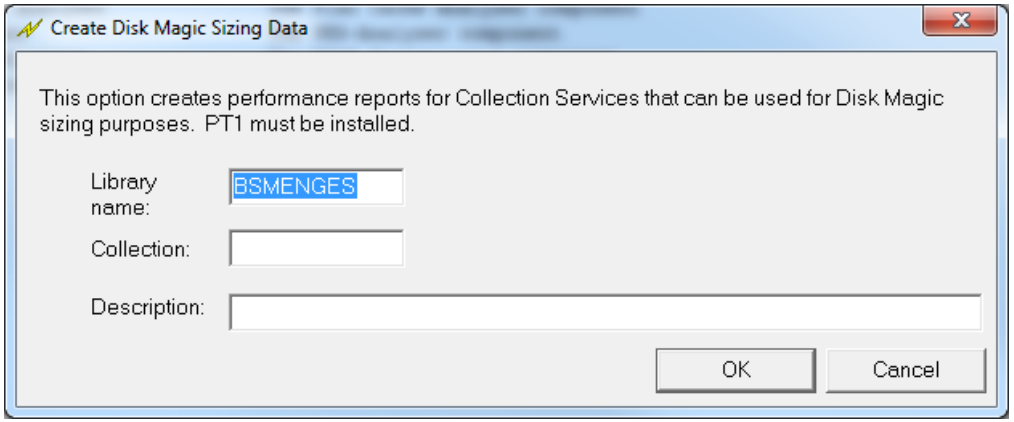

Must Gather Tools – Communications menu

- The menu provides access to the same options available via the GO COM command plus options to use the STRCMNTRC command.

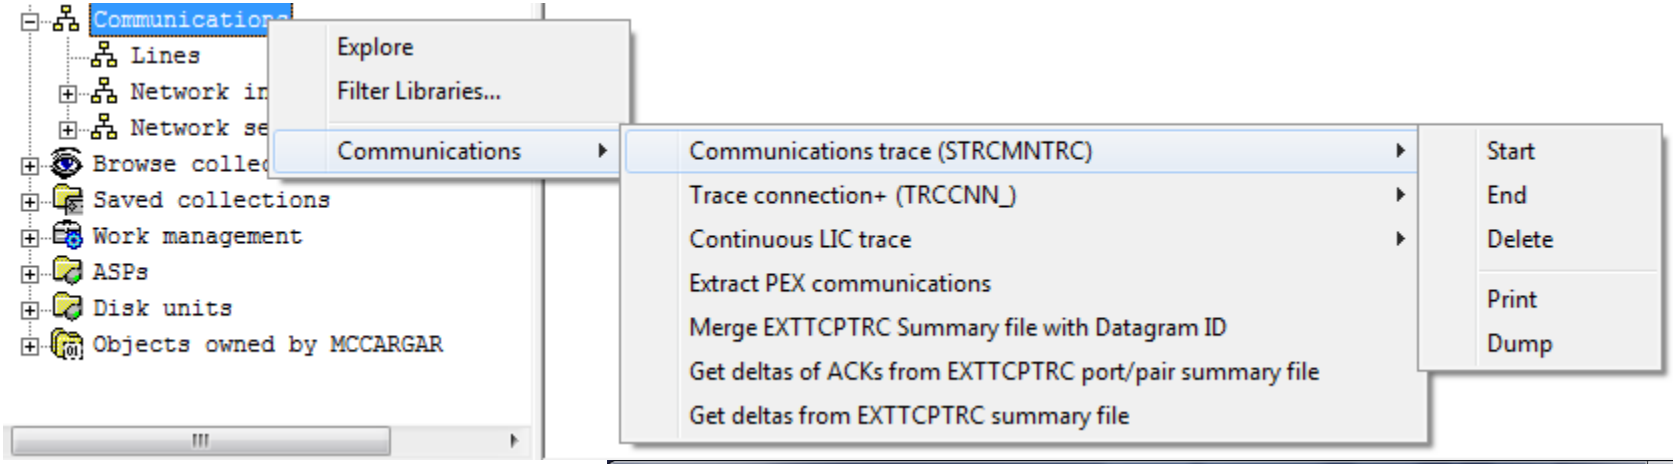

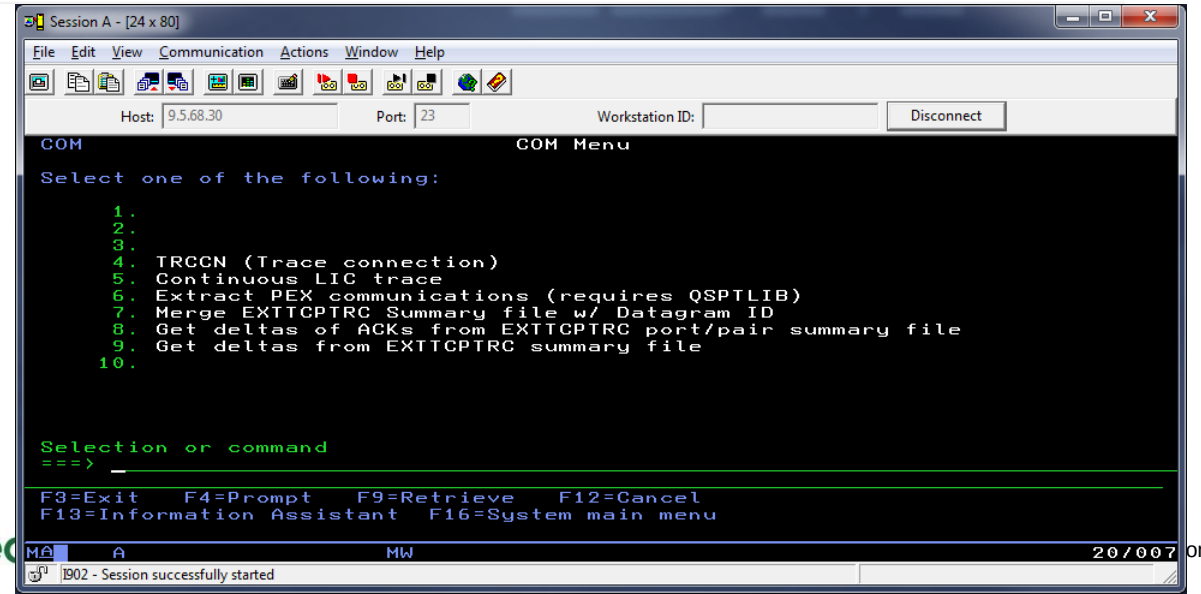

Must Gather Tools – Communications folder

- Communications folder provides options to work with the lines, network interfaces and network servers on the system. Comm traces can be started, ended, deleted, printed or dumped from the GUI.

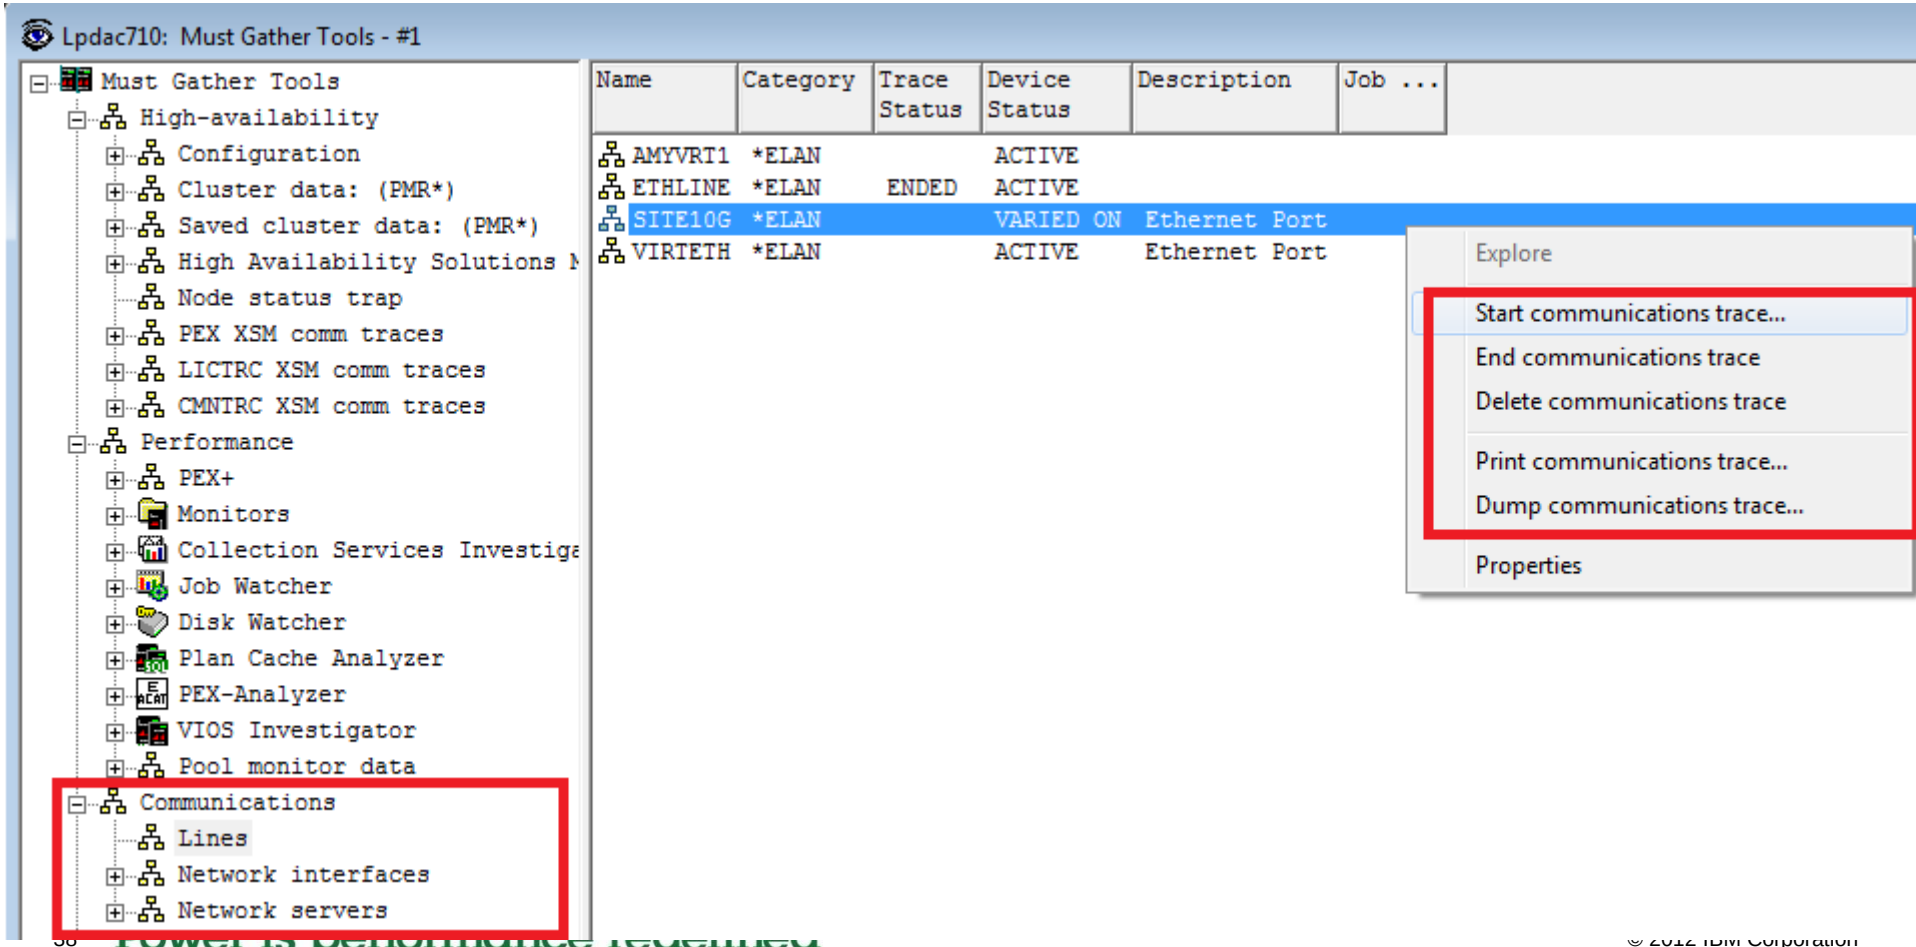

HMC Walker Introduction

HMC Walker is an HMC GUI that provides configuration details and performance metrics across all LPARs attached to the HMC.

Provides CPU and memory statistics across 1 or more ‘physical systems.’

Performance data for any type of LPAR can now be graphed with iDoctor using this offering.

HMC Walker - Managed (Physical) system CPU graph example

Shows CPU time and max LPAR CPU utilization over the last 60 days.

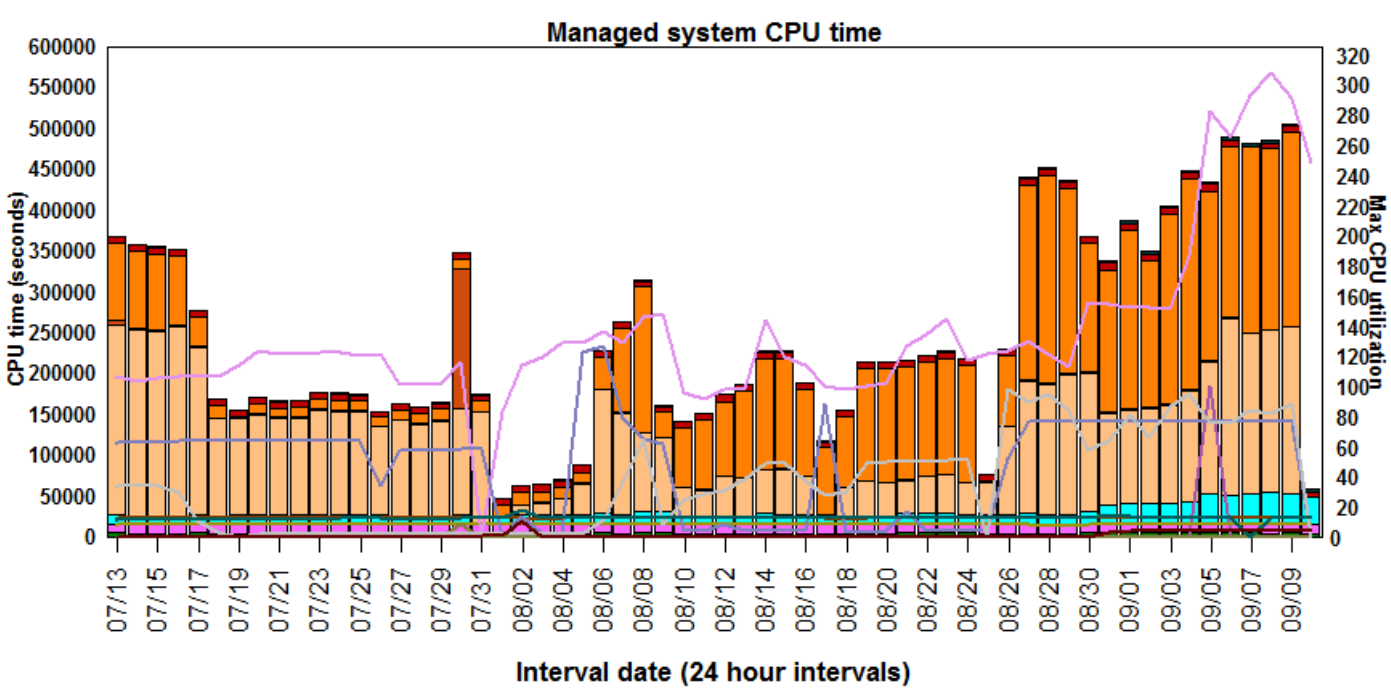

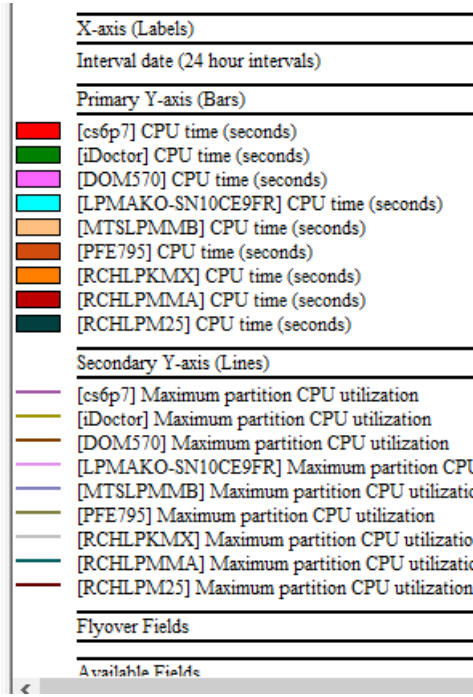

HMC Walker - LPAR CPU time graph example (ALL systems)

Same as previous except showing the LPARs instead.

This graph is showing AIX, VIOS and IBM i.

Note: 30 sec CPU filter, means LPARs that used < 30 seconds of CPU per day are bundled together.

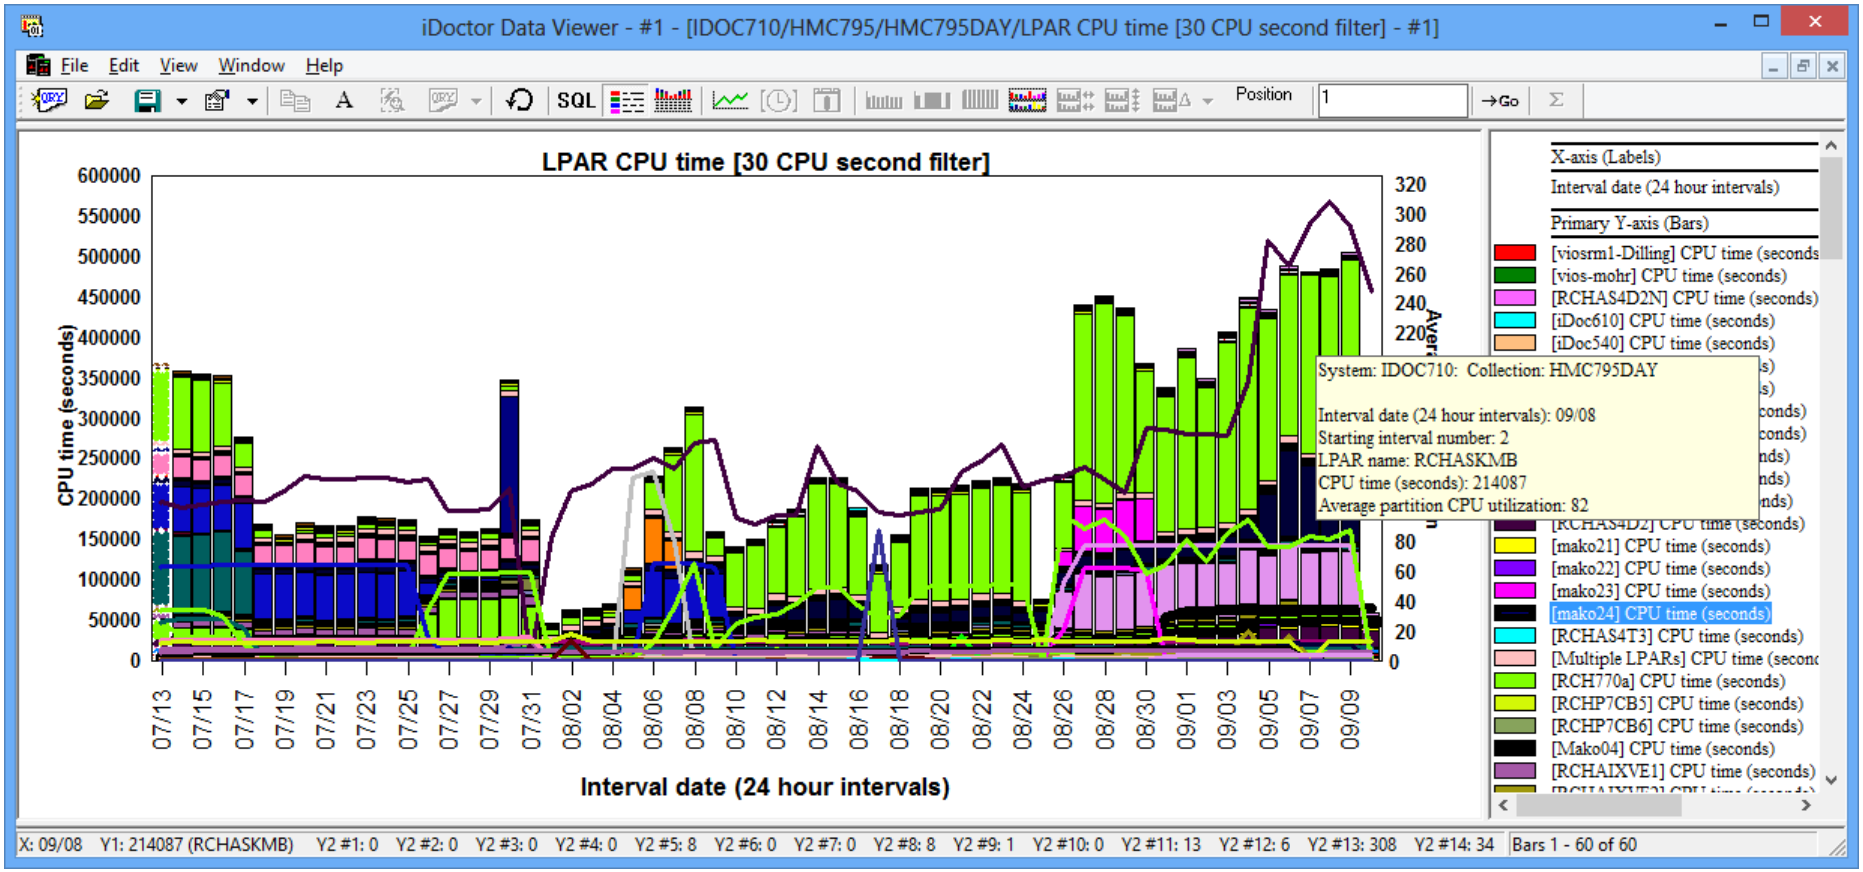

HMC Walker - Create an HMC connection

After starting the iDoctor GUI. Right-click the connections list and use the Add Connection menu.

Set the connection type to HMC and fill in the HMC name or IP address.

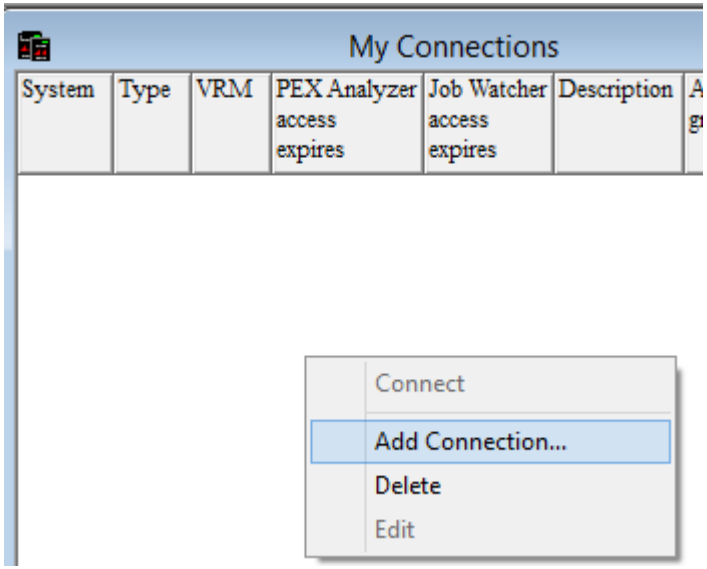

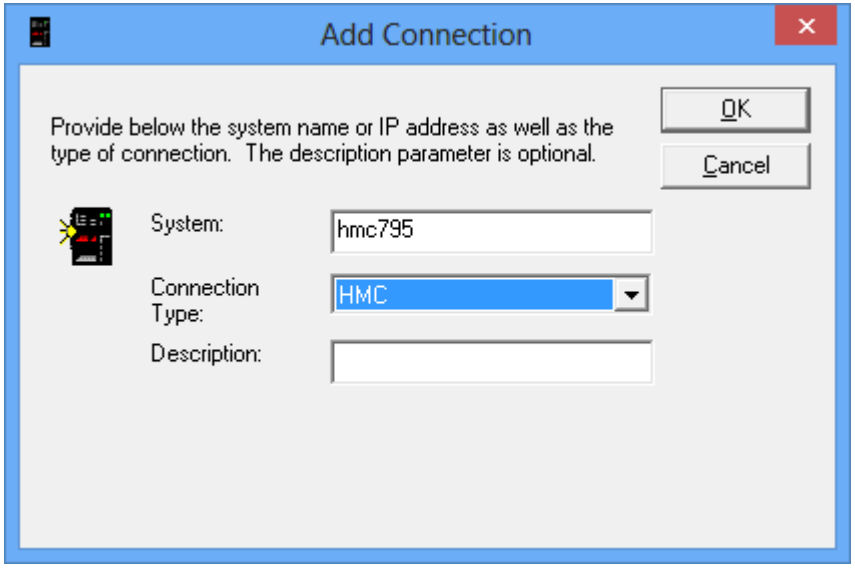

HMC Walker - Open (double-click) the connection

Launch the connection using the connect menu or double-click it.

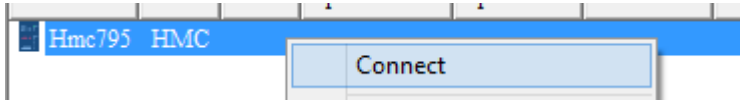

If you get this message, then you must install the IBM internal components enabler (Window's registry file):

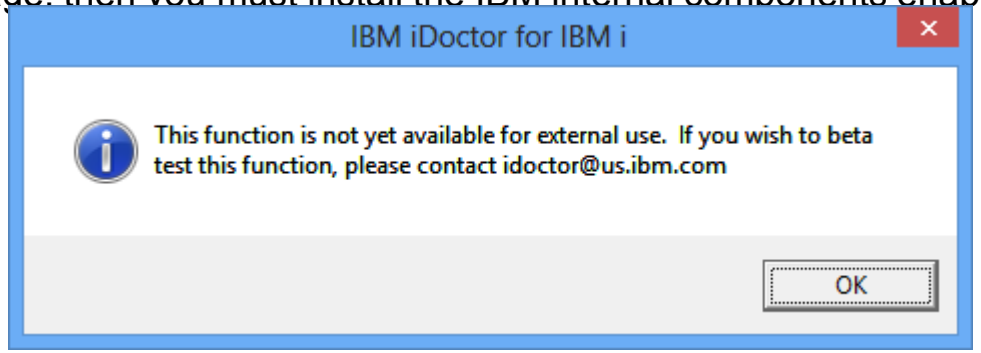

Then sign on:

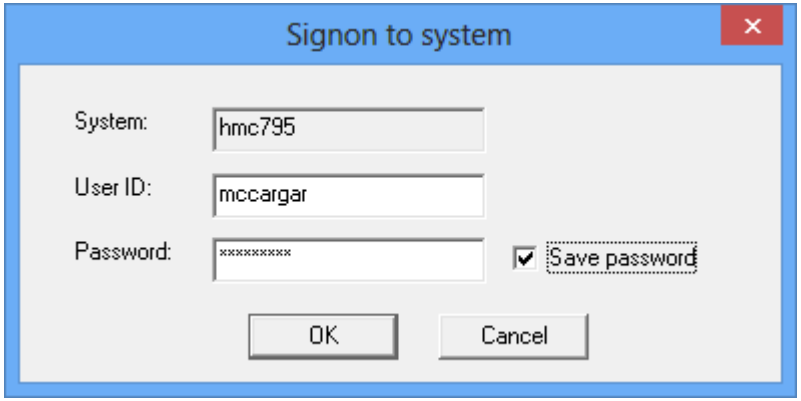

HMC Walker Component View

Initially the view will only contain options to manage/work with your HMC.

Once performance data has been captured additional views will be shown here.

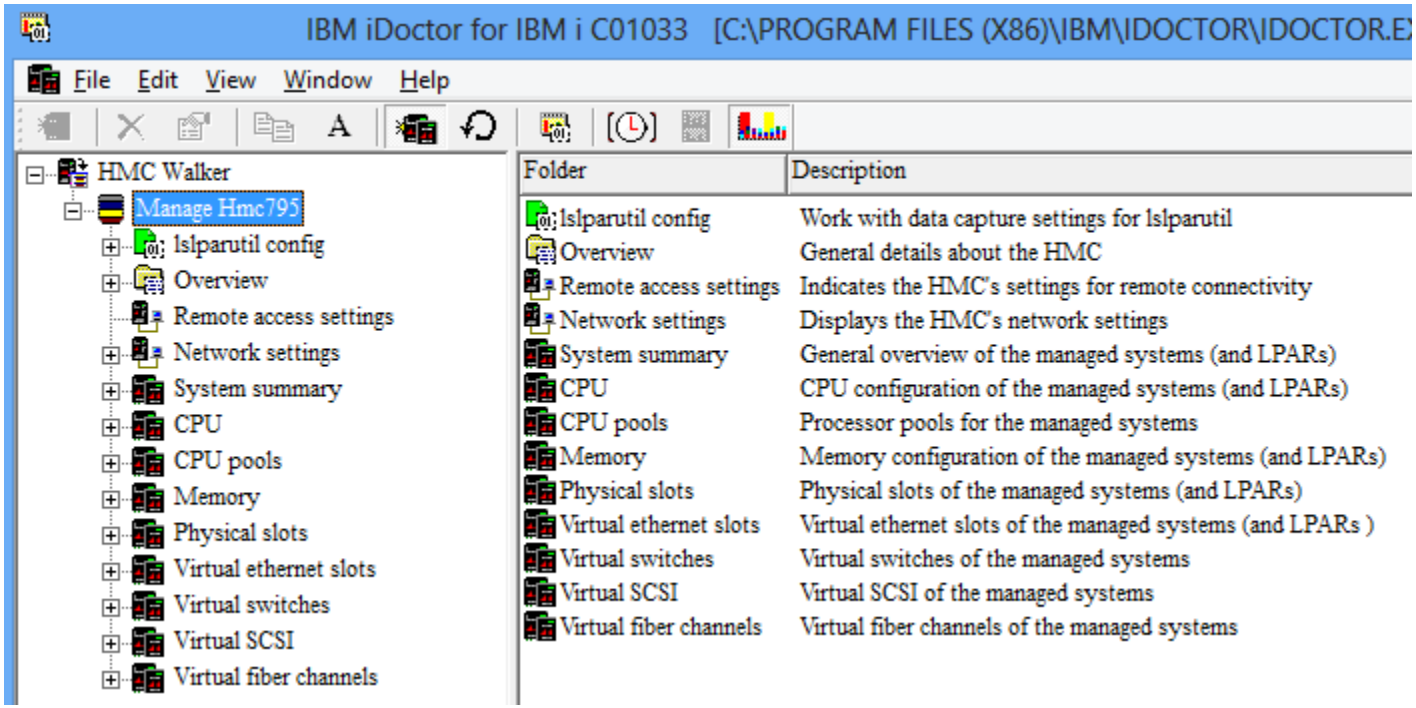

Important: Right-click the HMC Walker icon and choose the “Set analysis database” menu to select where data collected by HMC Walker should be stored for analysis. (If you already have IBM i systems in your connections list you will be automatically prompted to pick one.)

HMC Walker - Set analysis database

DB2 on IBM i is the default. (fill in your desired IBM i system name here or select a different option.)

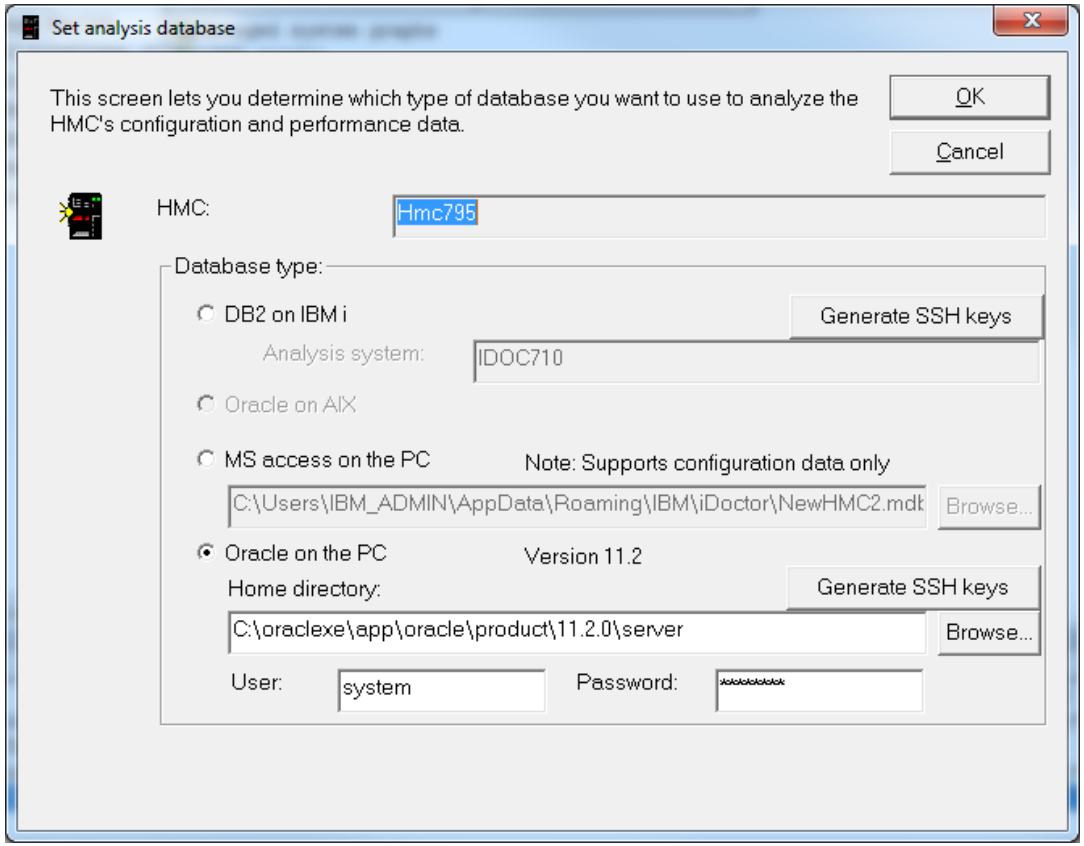

If you want to use Oracle on the PC, then specify where you installed it to and the user and password used when you installed Oracle.

HMC Walker - Generate SSH keys (if using IBM i as the DB)

The following information is shown to help you set up a secure SSH connection between the HMC and your IBM i analysis system. This is a required step if you want to use an IBM i for analysis.

In order to setup a secure SSH connection between IBM i and the HMC, please do the following steps:

1. Open a green screen session to the IBM i and sign on to the system.

2. From the CL command line run the following command:

> QSH

3. From QSH run the following commands:

$ cd /QIBM/ProdData/iDoctor/scripts

$ hmcKeyGen.sh <your hmc name> <your hmc user name>

(example: hmcKeyGen.sh hmc795 mccargar)

(when prompted with "Password:", please enter your hmc user's password)

4. Now to confirm that the SSH key generation is successful, issue the following command and you should not be prompted for a password:

$ ssh <your hmc user name>@<your hmc name>

(example: ssh mccargar@hmc795)

For more information please visit:

http://www-01.ibm.com/support/docview.wss?uid=nas1315c113cf5dd9ea0862570de0062e1ce

http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=%2Fp7ha1%2Fsettingupsecurescriptexecution.htm

HMC Walker - Generate SSH keys (if using PC/Oracle as the DB) – page 1

In order to setup a secure SSH connection between the PC and the HMC, please do the following steps:

1. **Download the windows installer package for Putty from this page:**
   http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html
   This is the link under "A Windows installer for everything except PuTTYtel" in bold.

2. **Go to the directory you installed Putty to Using Windows Explorer and open puttygen.exe**

3. **Click the Generate button and move the mouse around within the Key area until the key has been generated.**

4. **Click the Save public key button to a file of your choice. You will need to edit this file later.**

5. **Click the Save private key button and call it something like hmc1_prvkey.ppk and save it to the directory Putty is installed in.**

6. **Open your public key file in wordpad (not notepad). We have to convert the key format to OpenSSH format.**

Your key will look something like this:

```
---- BEGIN SSH2 PUBLIC KEY ----
Comment: "rsa-key-20131003"
AAAAB3NzaC1yc2EAAAABJQAAAIEAs6oe2BJwYnYNysrsNvwn+SHCePnm2QcPfVbq
MGp4QMOIojERQz+Jw9Iz+7lpgxhRnc/GF7z0hFAPgXx5/gTA7qtEXpsAEGMk3ts0
opt0eUPBY+fUVC0mbU8P6pJW/XoEe1zme/C+HVaoe569go1D9NXyvhpzujpOyXG+
jtahrFs=
---- END SSH2 PUBLIC KEY ----
```

HMC Walker - Generate SSH keys (if using PC/Oracle as the DB) – page 2

7. Remove the 1st 2 lines and the last line. Add "ssh-rsa " at the beginning and remove all new line characters so the entire string is on one line.

8. At the end add username@hmcname. The changed public key file should look something like this:

ssh-rsa AAAAB3NzaC1yc2EAAAABJQAAAIEAs6oe2BJwYnYNysrsNvwn+SHCePnm2QcPfVbqMGp4QMOlojERQz+Jw9Iz+7lpgxhRnc/GF7z0hFAPgXx5/gTA7qtEXpsAEGMk3ts0opt0eUPBY+fUVC0mbU8P6pJW/XoEe1zme/C+HVaoe569go1D9NXyvhpzujpOyXG+jtahrFs= mccargar@hmc795

9. Next open an SSH connection to the HMC using Putty.exe.

Copy and paste the following command (replacing your key with mine) to apply your key to the HMC. It's very important that all new lines are removed or this won't work!

mkauthkeys --add 'ssh-rsa AAAAB3NzaC1yc2EAAAABJQAAAIEAs6oe2BJwYnYNysrsNvwn+SHCePnm2QcPfVbqMGp4QMOlojERQz+Jw9Iz+7lpgxhRnc/GF7z0hFAPgXx5/gTA7qtEXpsAEGMk3ts0opt0eUPBY+fUVC0mbU8P6pJW/XoEe1zme/C+HVaoe569go1D9NXyvhpzujpOyXG+jtahrFs= mccargar@hmc795'

10. To confirm this is working open putty.exe again to your HMC. The key should be used when signing on instead of a password. You will see something like this:

login as: mccargar

Authenticating with public key "rsa-key-20131003"

Last login: Fri Oct 4 07:43:32 2013 from rmccargar.rchland.ibm.com

HMC Walker Component View (with an IBM i DB)

If an IBM i DB is used, then 3 options are provided:

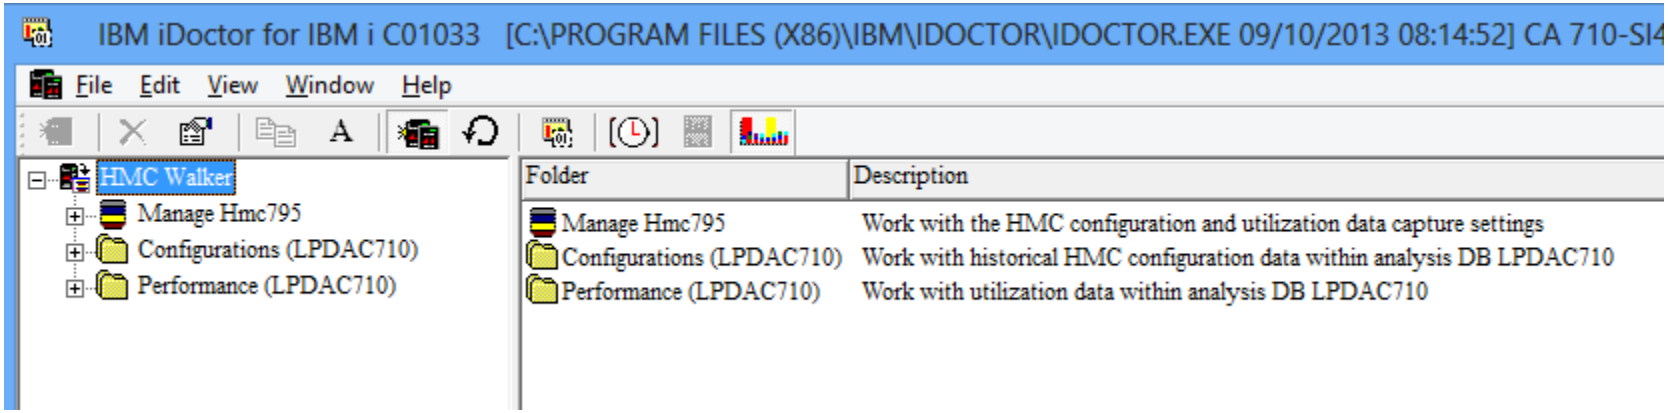

Configurations folder is historical HMC configuration data.

The Performance folder is lslparutil HMC performance data previously (or currently being) captured and stored on the IBM i.

HMC Walker Component View (with Oracle as DB)

If the local Oracle DB is used, then these options are provided:

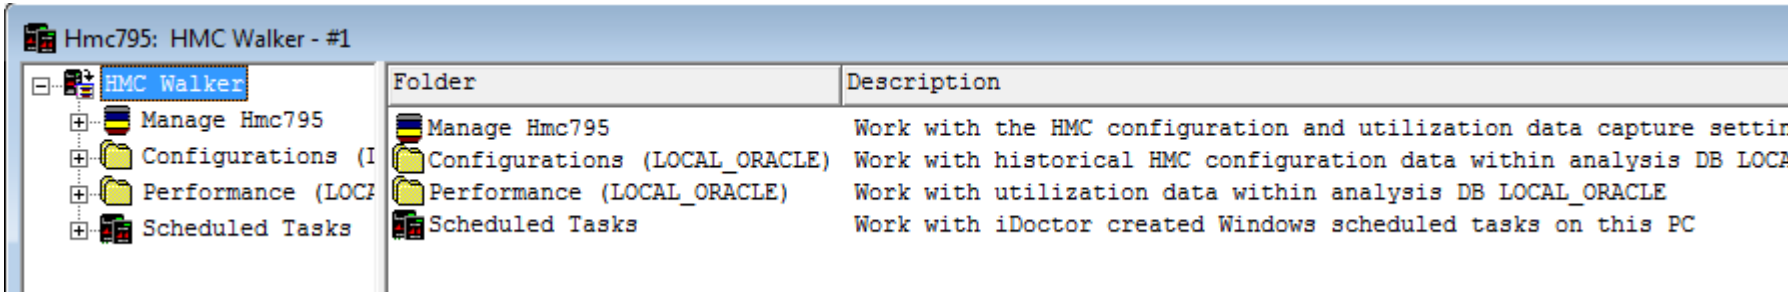

This view also contains a scheduled tasks folder that lists the iDoctor HMC Walker tasks that have been created for you in the Windows Task Scheduler on the current PC.

HMC Walker - Manage HMC -> lslparutil config folder

Be sure that lslparutil is configured to collect data under the Manage HMC -> lslparutil config folder. It probably is not collecting any data.

Once turned on data is automatically collected 24x7 for the desired physical systems for all LPARs on each. It's best to set the sample rate to be the same for all physical systems if you want to graph them at the same time.

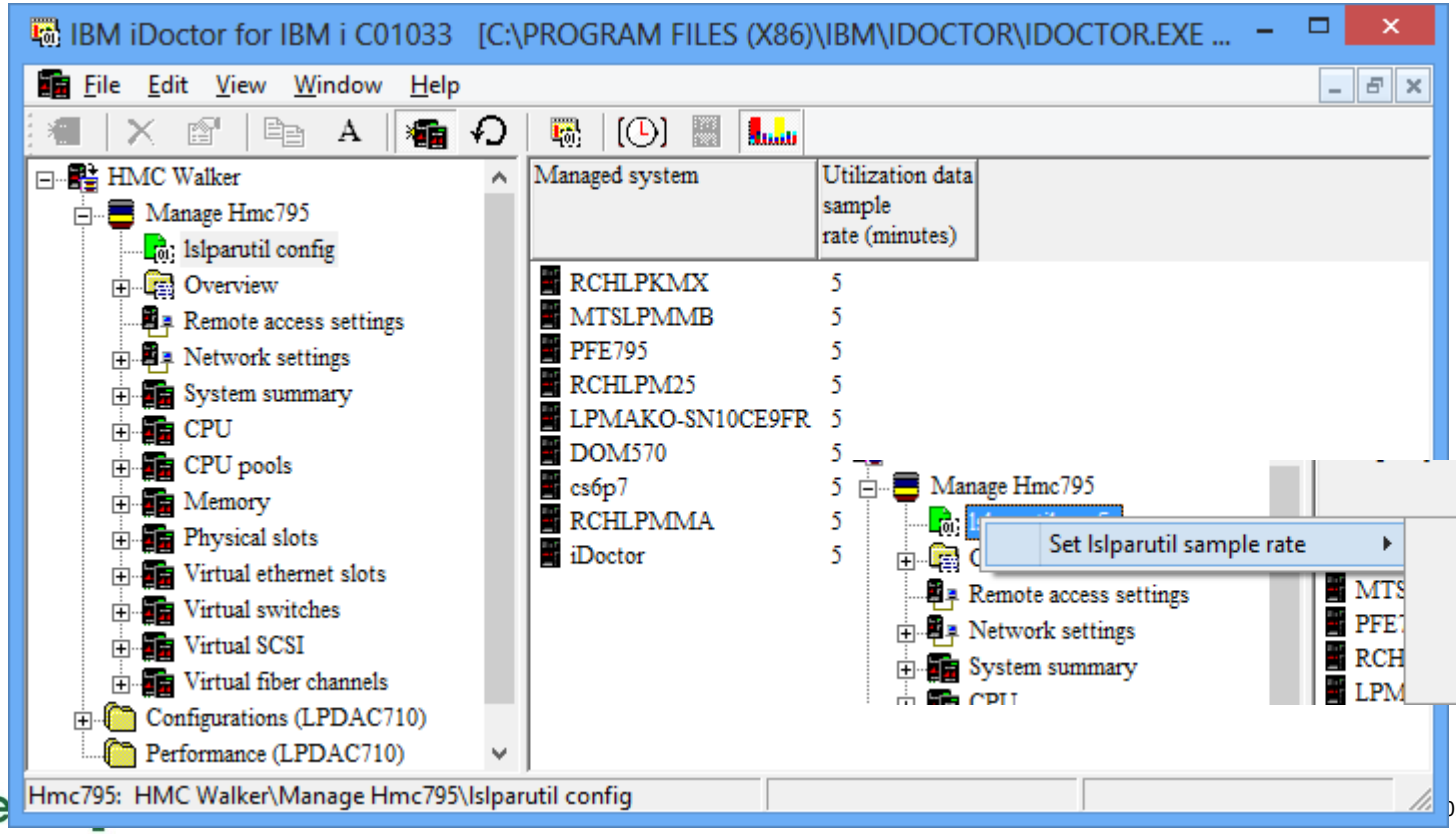

HMC Walker - Collecting lslparutil data (explained)

The oldest data will be automatically removed from the HMC once the size becomes too large. Hourly events saved 2 months, daily samples saved 2 years, monthly events saved 10 years.

Capturing data works by specifying how many minutes, hours, days prior to the current time you want to capture. Options exist to also allow you to specify the desired start and end time of the collection instead.

Right-click HMC Walker and use the Build performance collection menu. You can filter on one or more managed systems if desired.

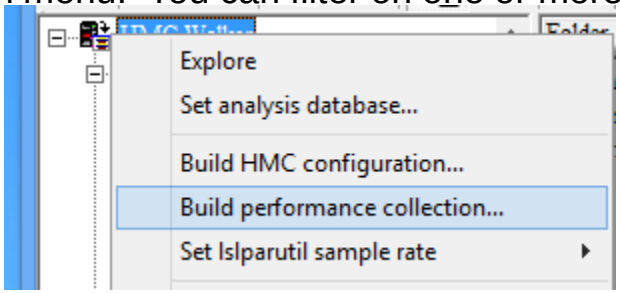

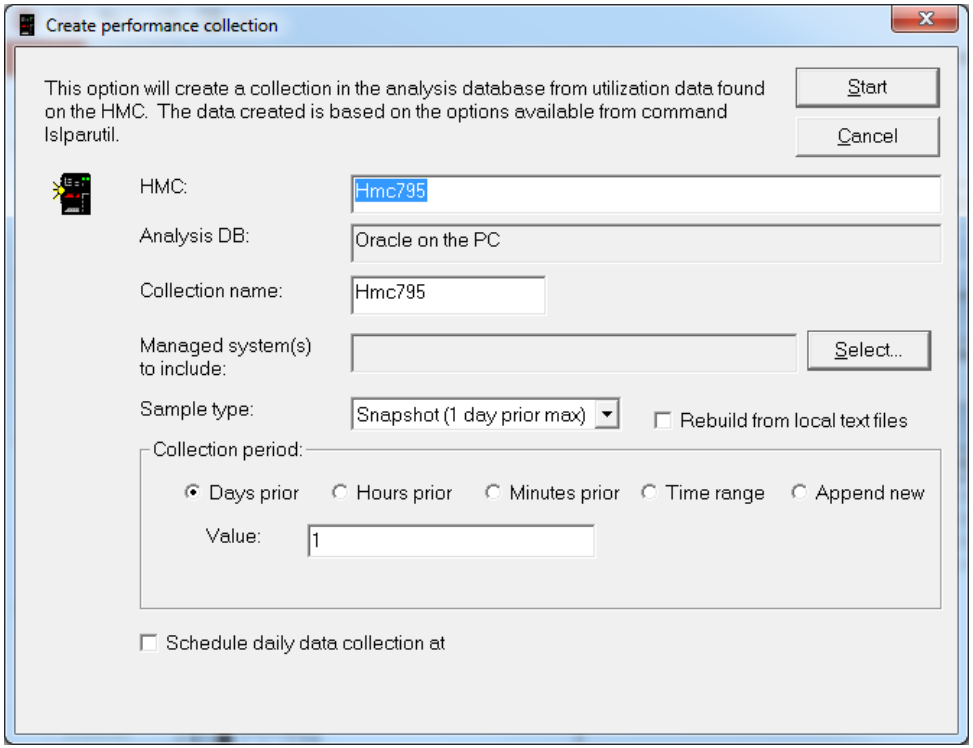

HMC Walker - Build performance collection window

This window allows you to specify the schema (library if using IBM i), collection name and managed systems to include.

Press the select button to view and select specific managed systems to include in the collection (or leave blank to include all of them.)

The sample type can be snapshot (whatever the sample rate is set to), hourly, daily or monthly. This lets you create graphs over longer periods of time if desired. **Note:** Of course if you just started collecting you will have to wait until the data exists.

Click the schedule daily collection at button to create a Windows Task Scheduler entry to collect the data off of the HMC every day at the desired time. The action drop down let you choose if you want the data appended to your collection (recommended) or create a new collection for each day.

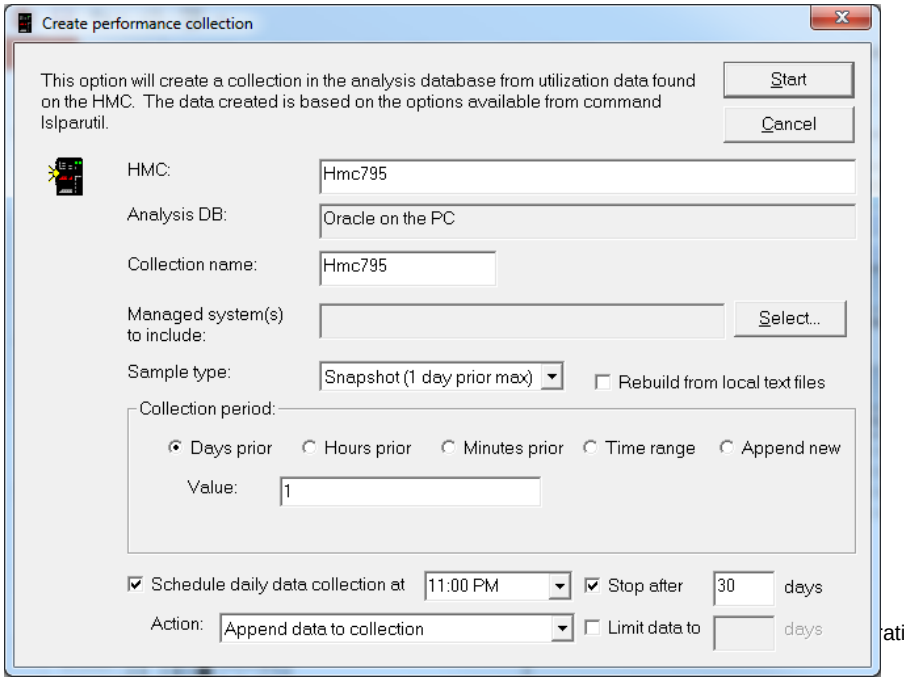

HMC Walker – Scheduled Tasks folder

This folder lists the Windows Scheduled Tasks related to HMC Walker that have been created on this PC.

From this view you can run the task immediately, delete it or view properties.

Note: The layout of columns will change in the future to be more usable and the properties are not yet implemented.

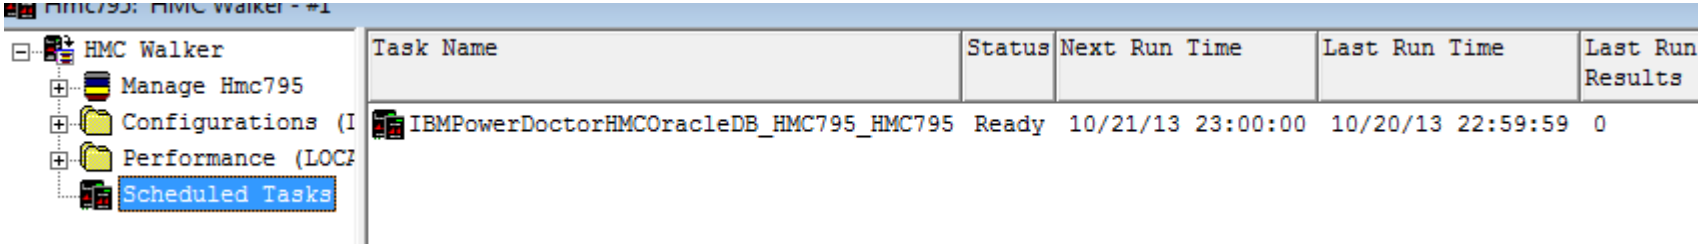

HMC Walker - Viewing Performance collections

After the collection has been captured to the database, the Performance folder will display it.

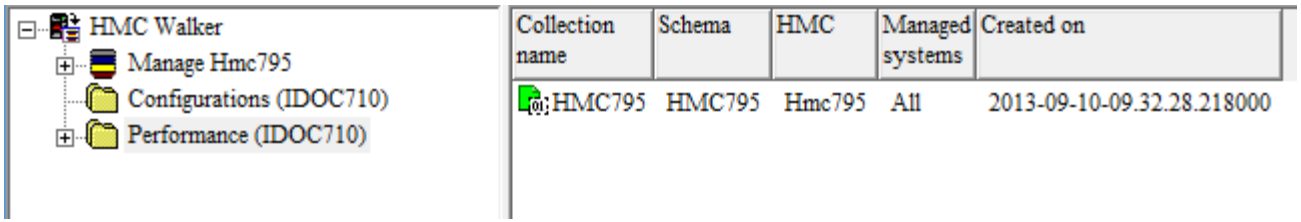

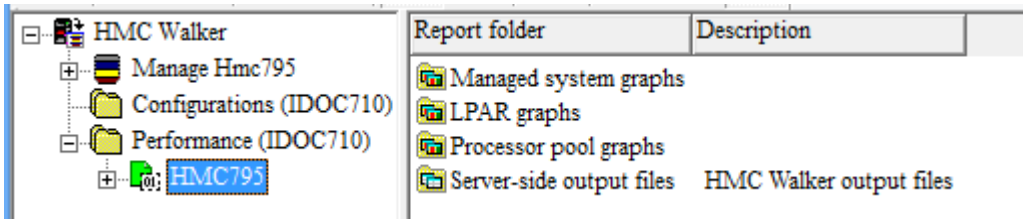

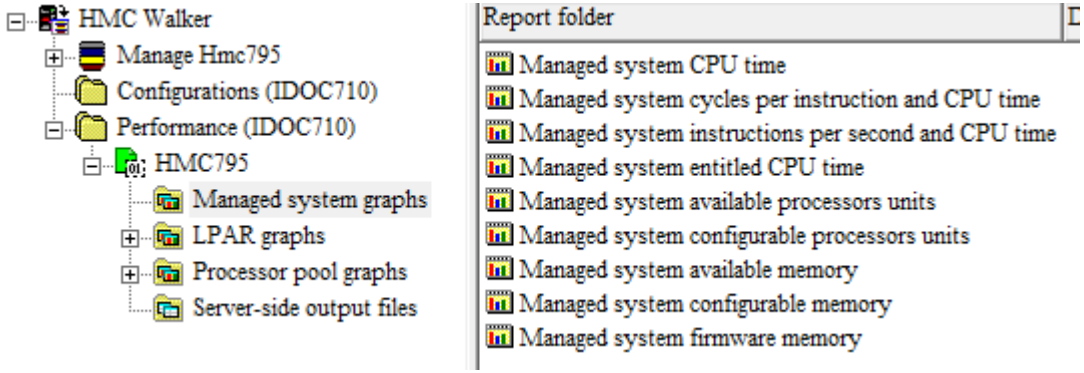

HMC Walker - 60 day graph example (investigate KMB)

The light green in this chart represents CPU time used by LPAR rchaskmb. Because this is an IBM i LPAR we can use Collection Services Investigator (in the default CS lib) to see which jobs are burning CPU. The high CPU burn on KMB has been happening for several days.

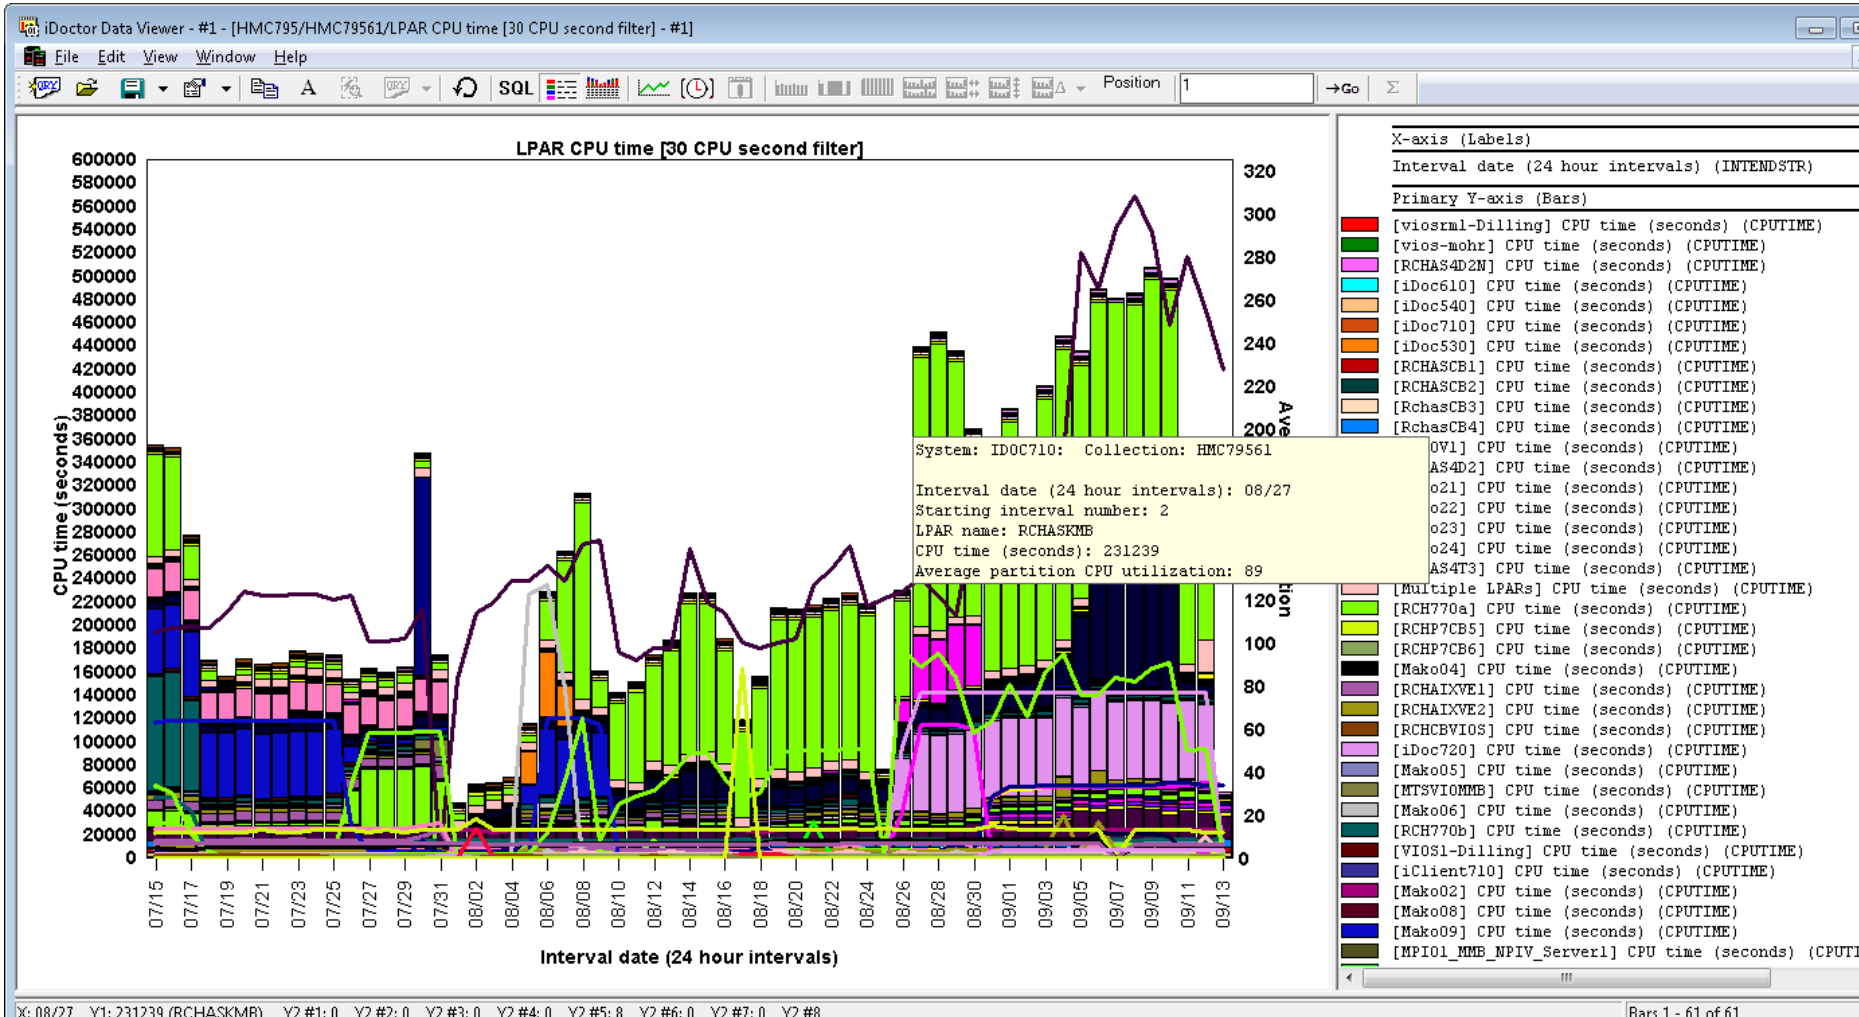

HMC Walker - Investigating CPU burn on Sept 12th

This graph on KMB within CSI for Sept 12th shows which jobs used the CPU that day and the other waits experienced.

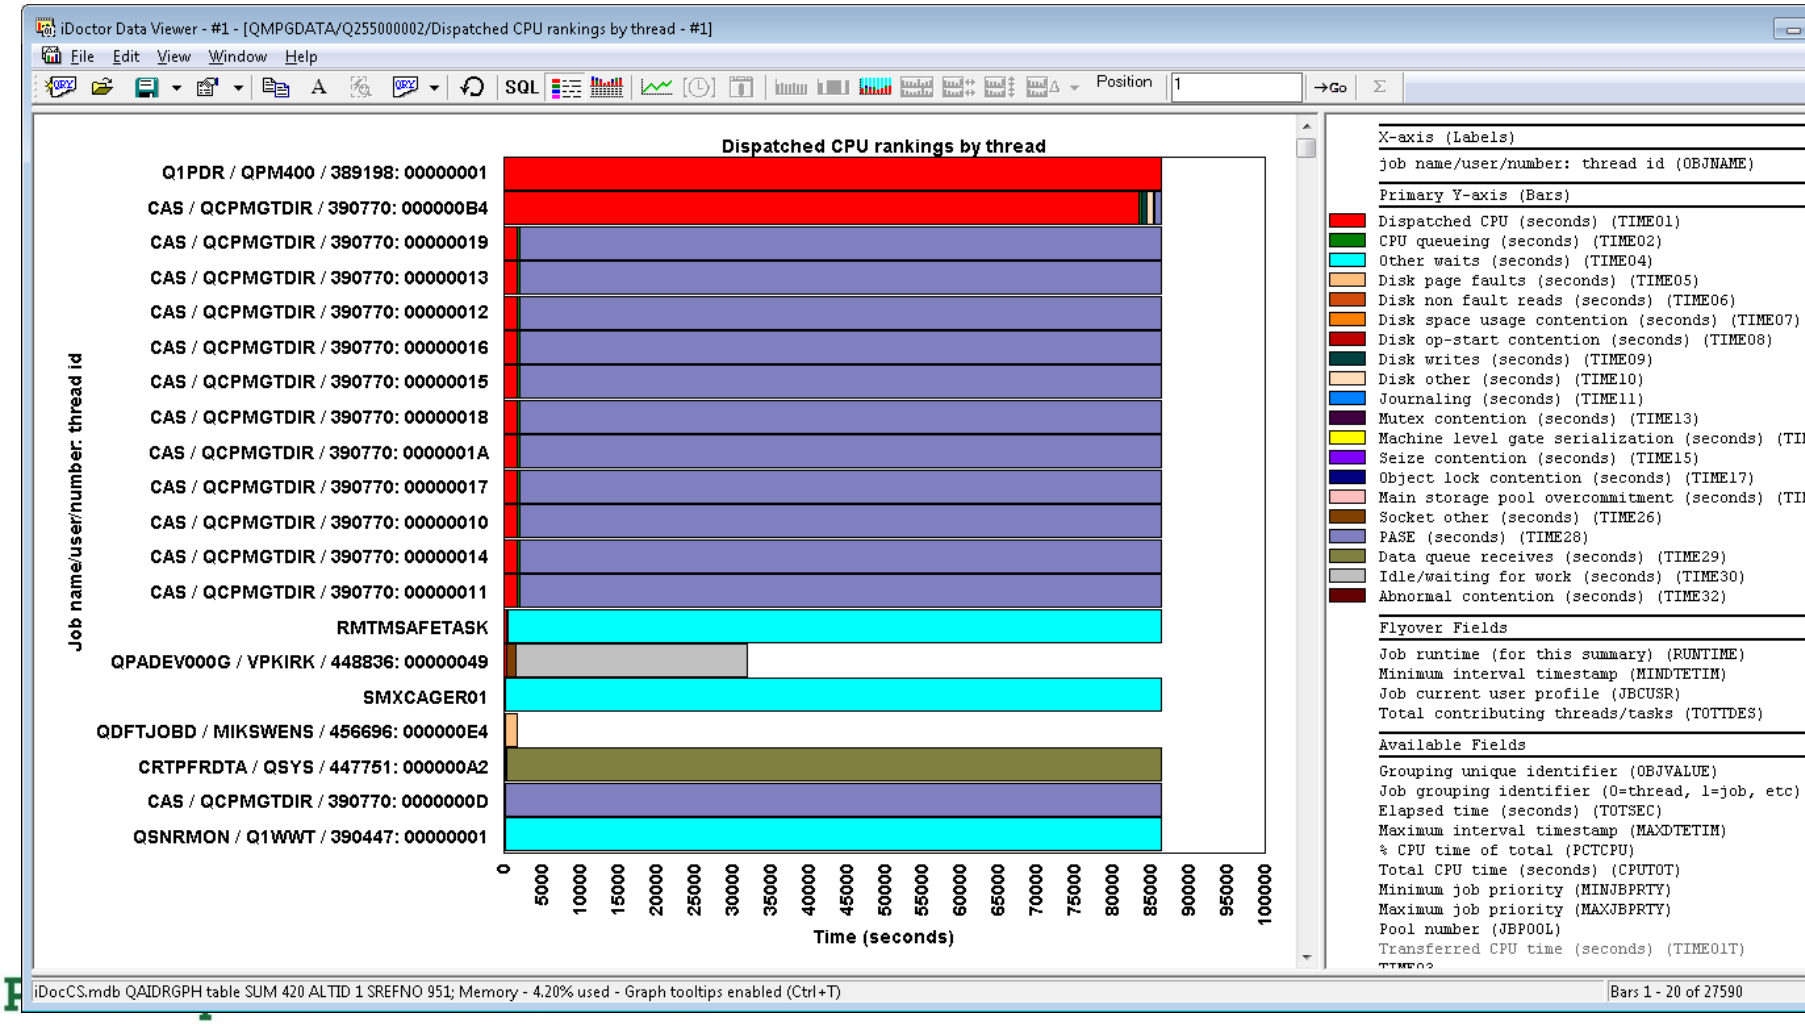

HMC Walker - Build configuration

Use the menu Build HMC configura

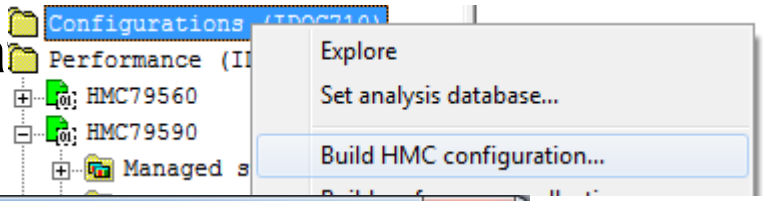

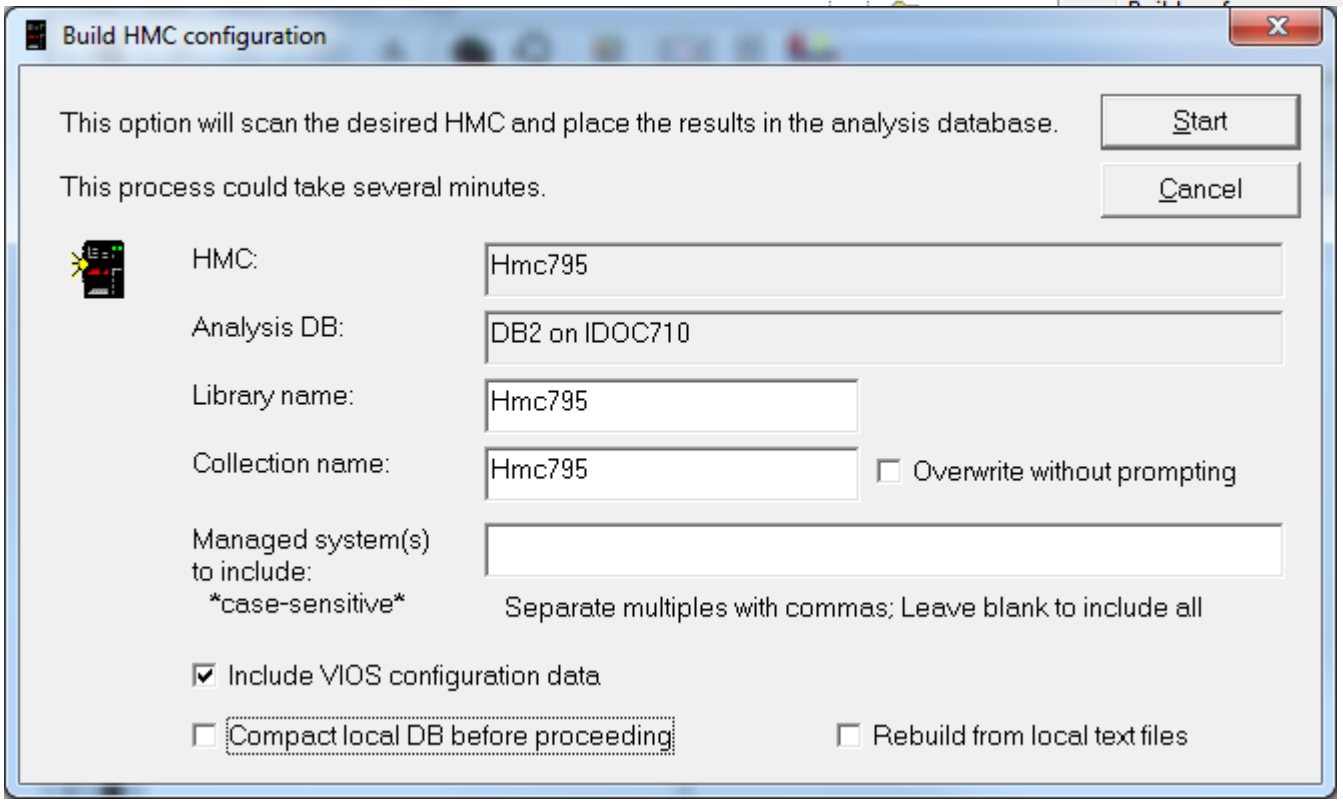

HMC Walker - Configurations folder

Contains the list of configurations that have been captured and stored in the analysis DB (in this case IBM i iDoc710.)

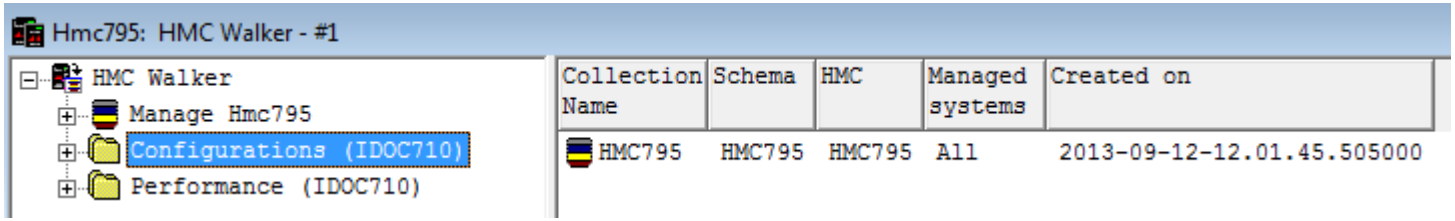

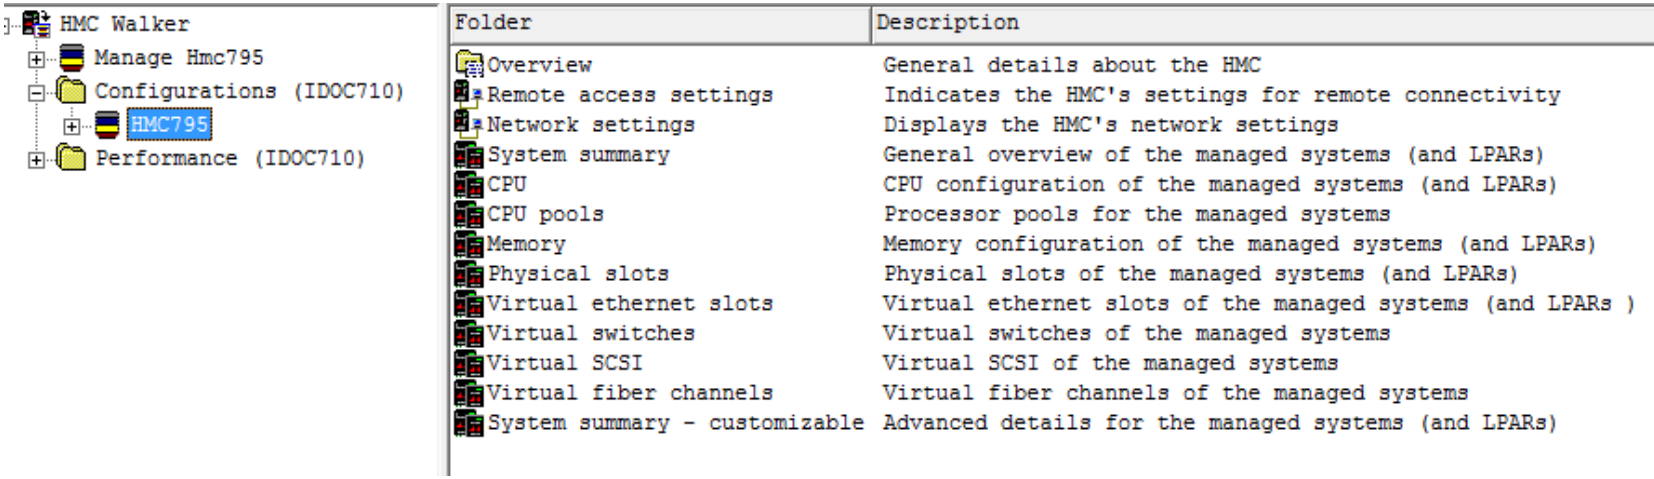

HMC Walker - Overview example

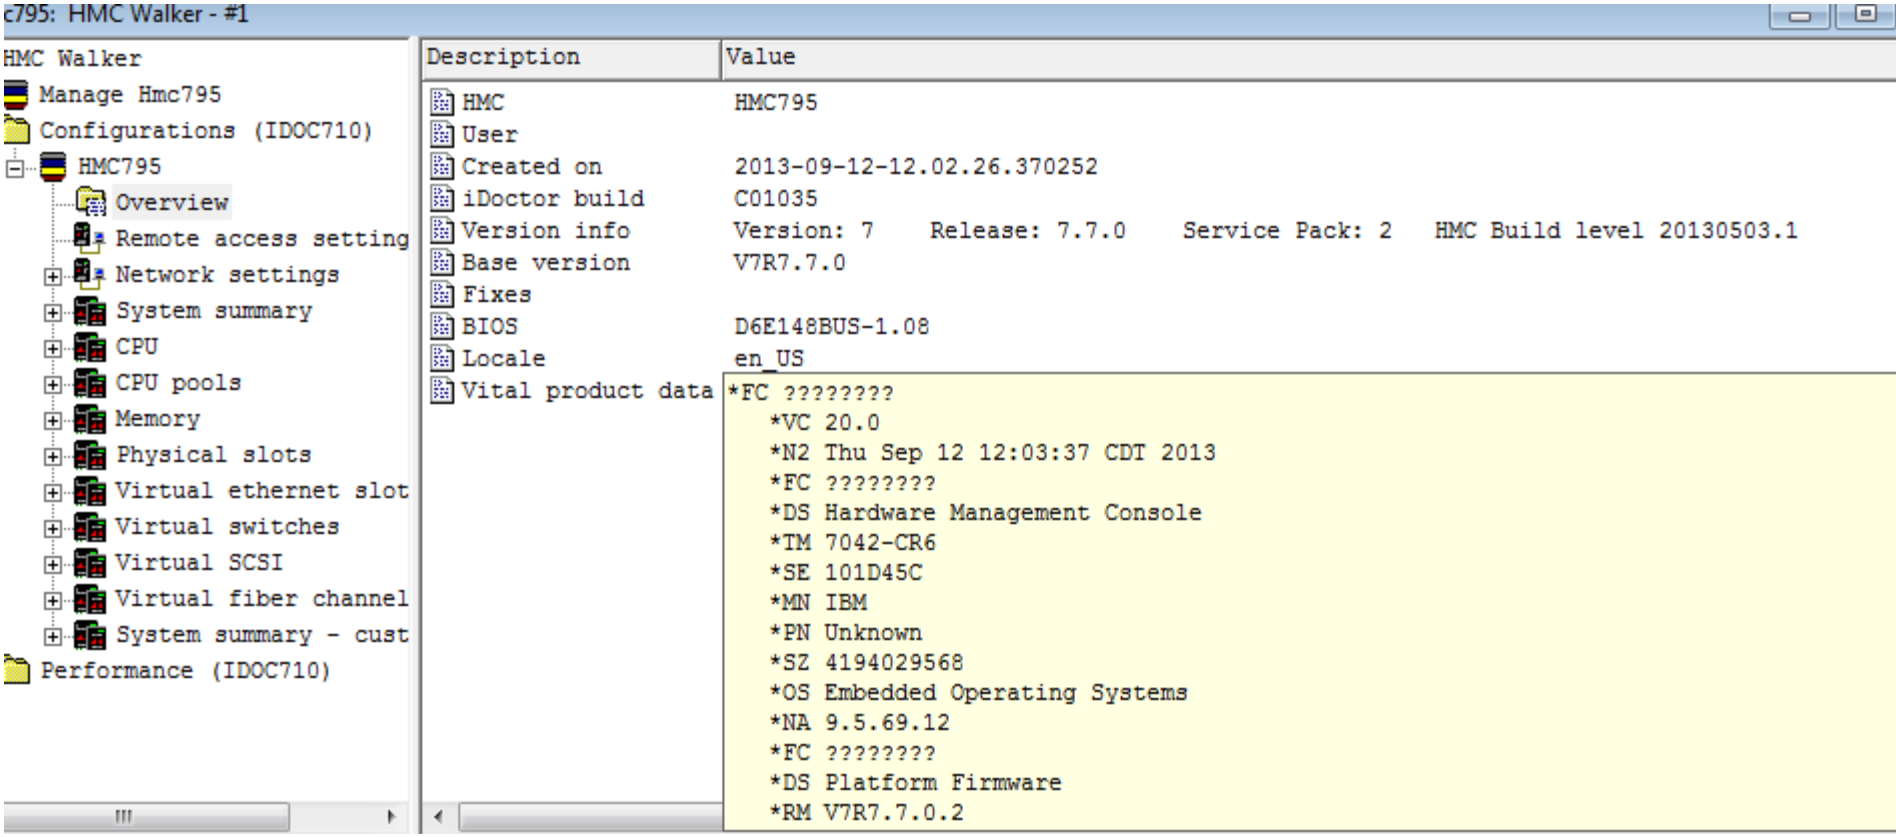

Job Watcher – TDETYPE added to rankings graphs

- TDETYPE field indicates T (task), P (primary thread), or S (secondary thread)

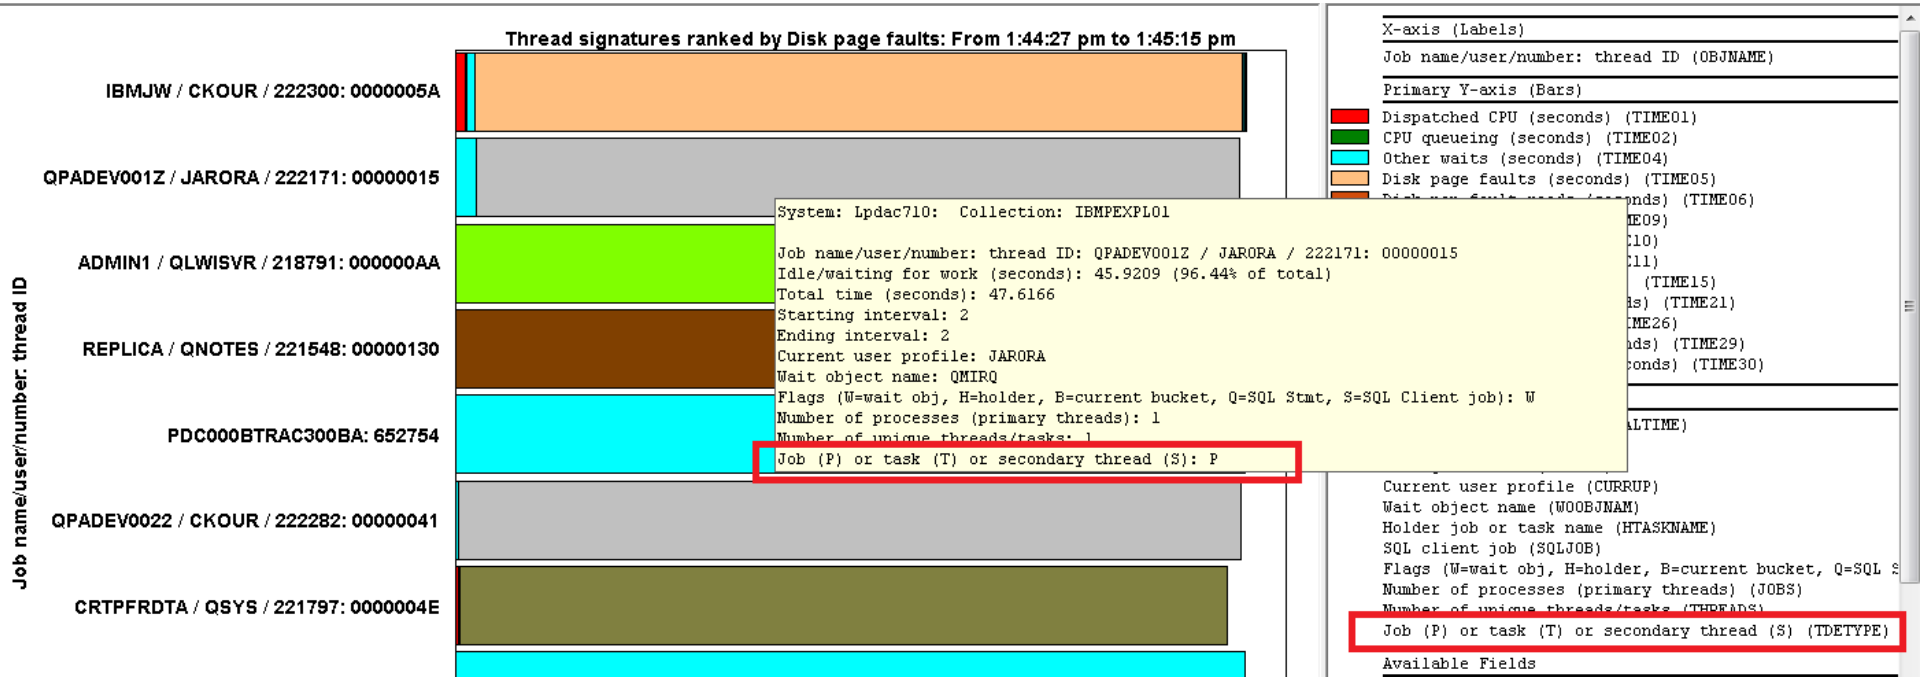

Job Watcher – Search now shows total call stacks

- Call stack search now includes the total number of call stacks each found entry applies to.

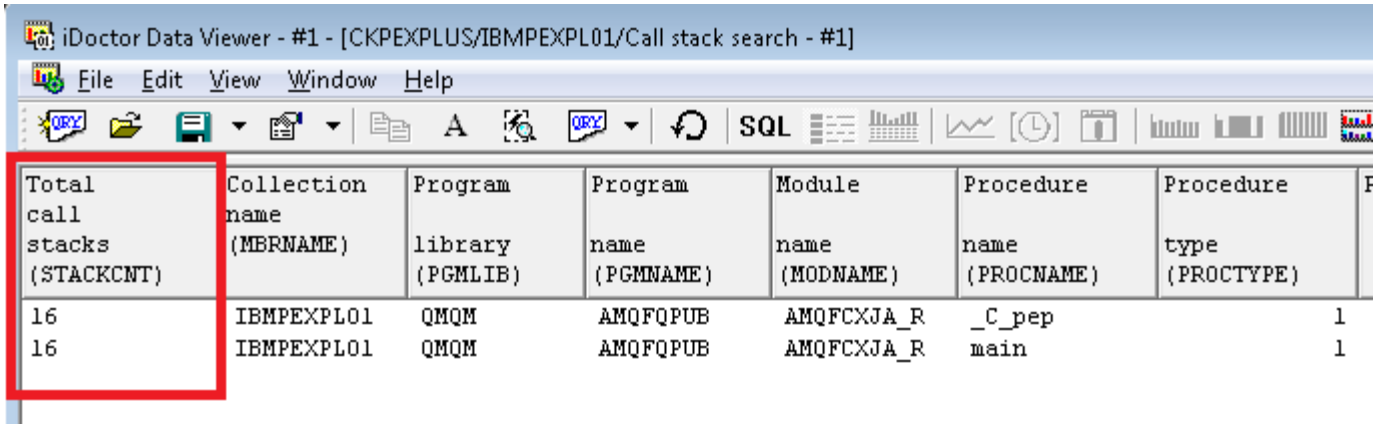

Total call stacks (STACKCNT)	Collection name (MBRNAME)	Program library (PGMLIB)	Program name (PGMNAME)	Module name (MODNAME)	Procedure name (PROCNAME)	Procedure type (PROCTYPE)	
16	IBMPEXPL01	QMQM	AMQFQPUB	AMQFCXJA_R	_C_pep		1
16	IBMPEXPL01	QMQM	AMQFQPUB	AMQFCXJA_R	main		1

Job Watcher – Search now shows total call stacks

- Added a new call stacks report /drill down called Jobs and programs (14 levels) calling the selected pgm/procedure: all intervals.

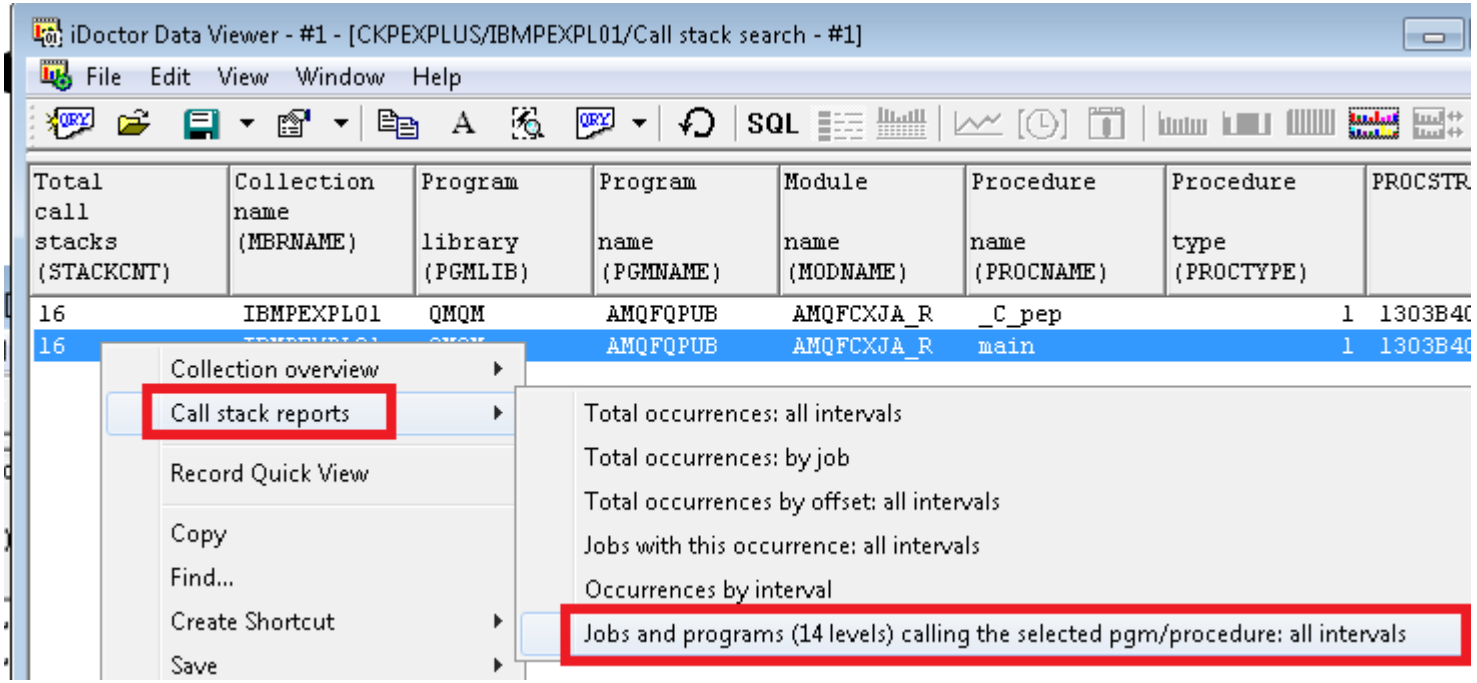

Job Watcher – New J9 JVM graph updates

- J9 JVM graphs – J9 JVM collection wait buckets (wait buckets graph but only includes the J9 JVM jobs)

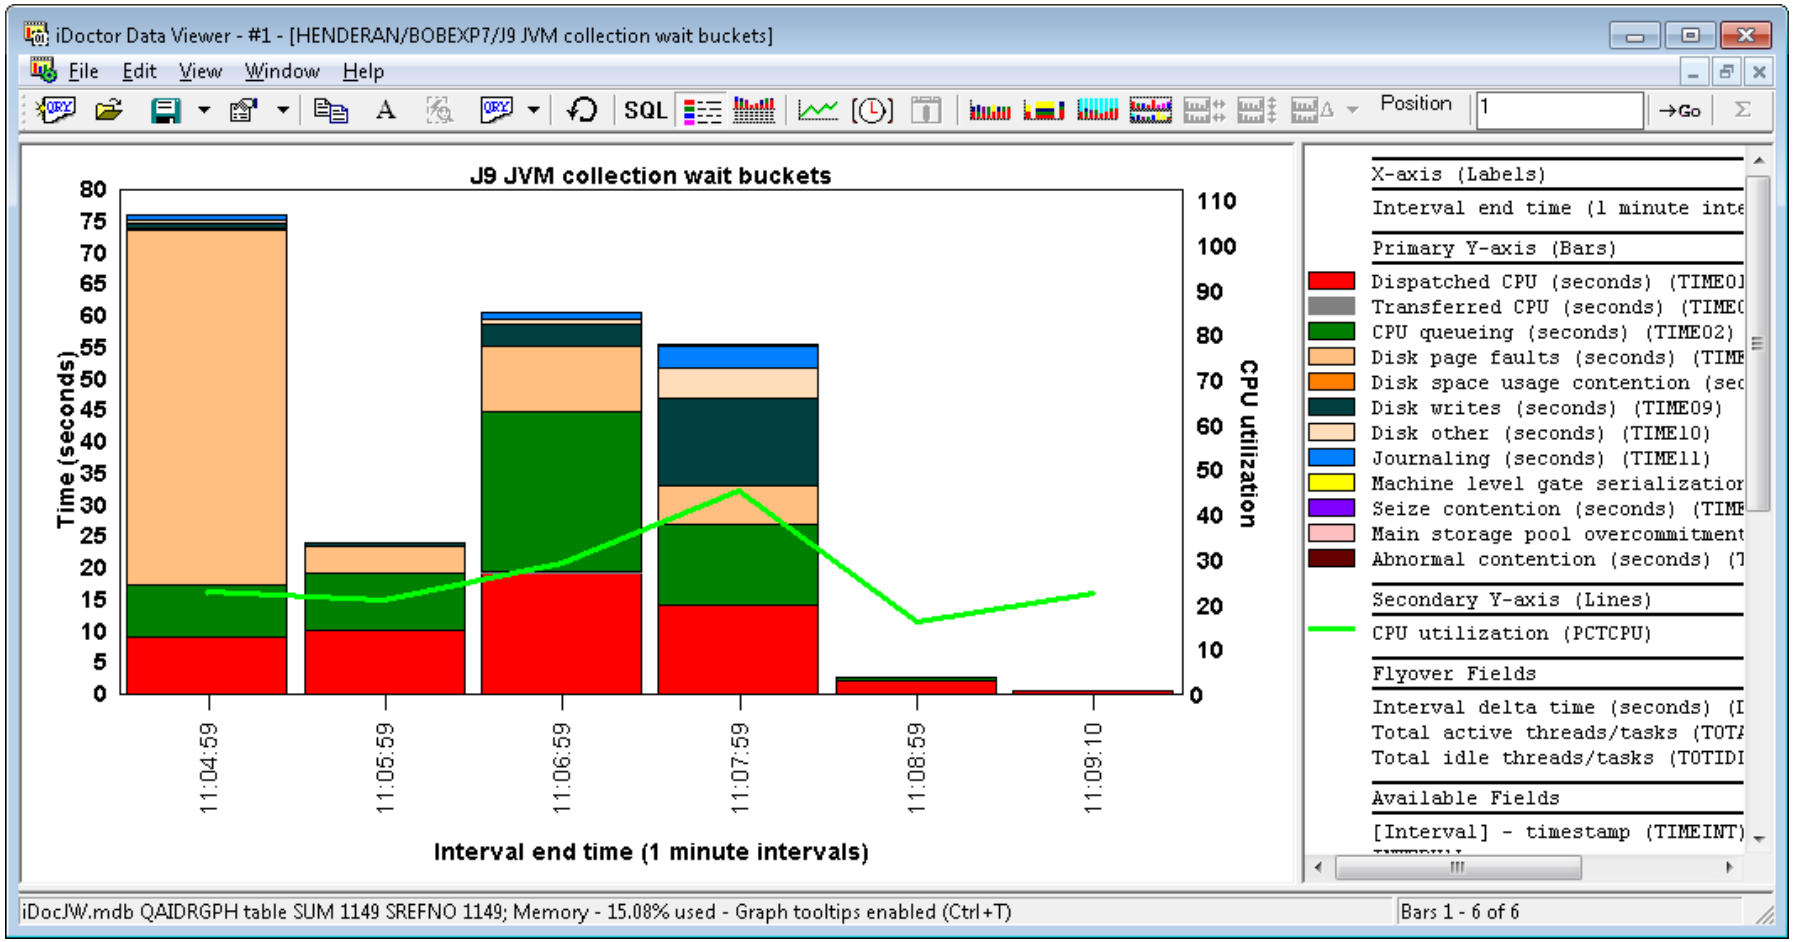

Job Watcher – New J9 JVM graph updates

- J9 JVM graphs – J9 JVM collection wait buckets (thread) rankings

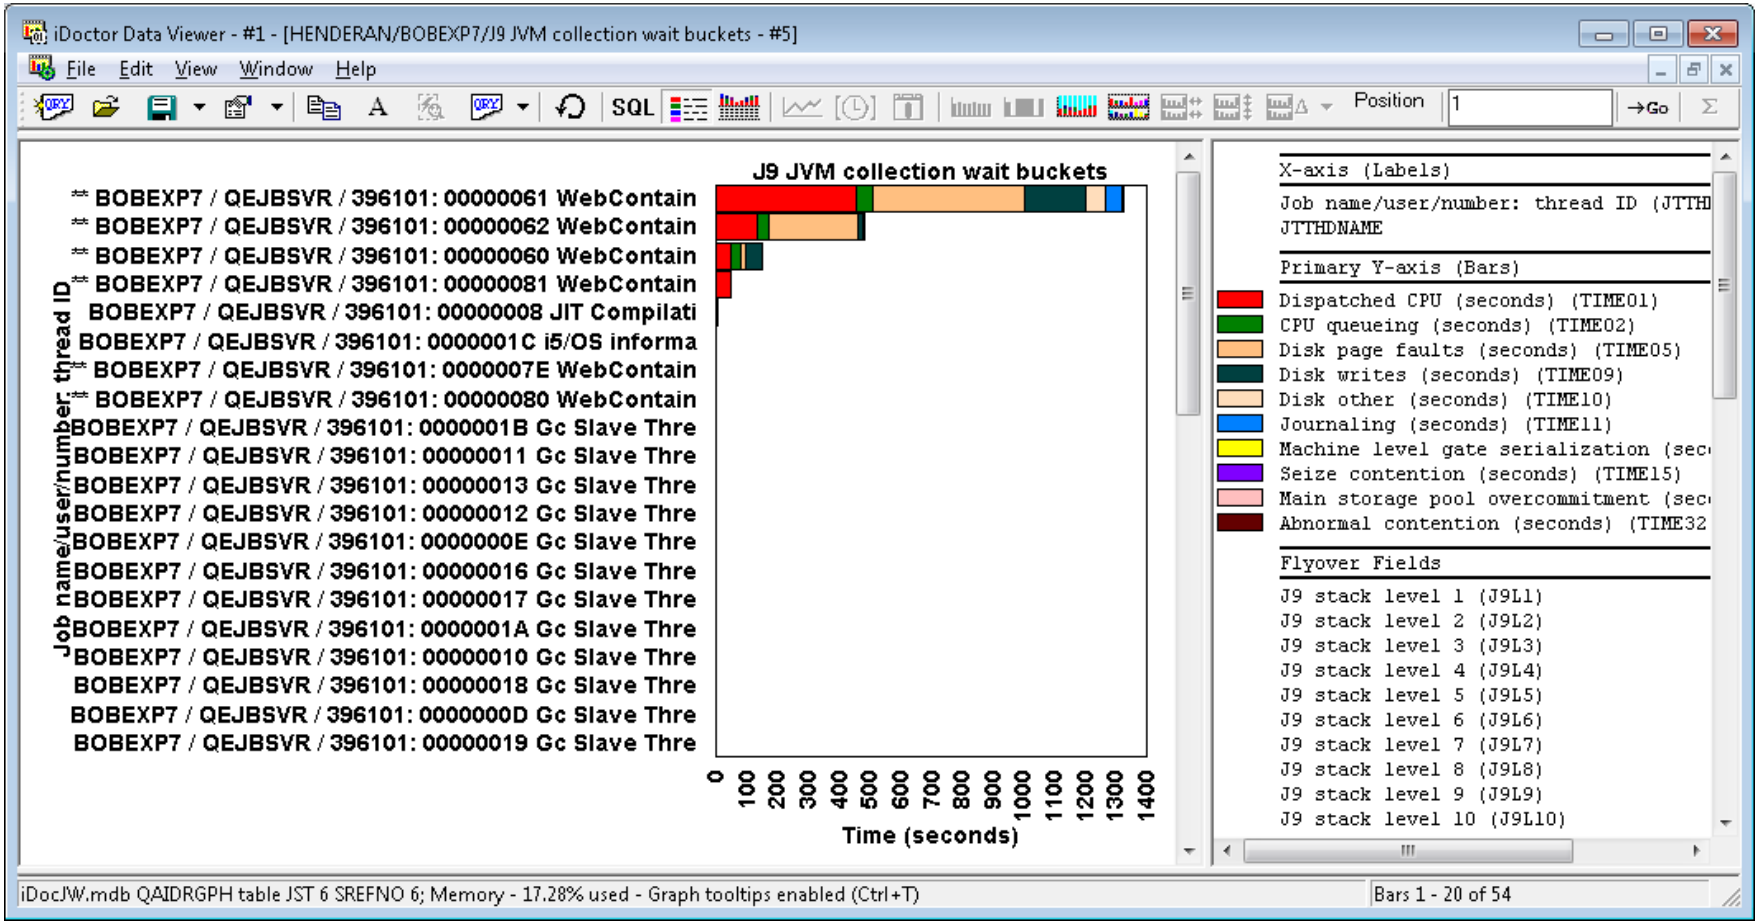

Job Watcher – Synchronous response graph updates

- I/O and memory page graphs – Synchronous response graph now includes “in progress” reads and write response times for I/Os that have not yet completed.

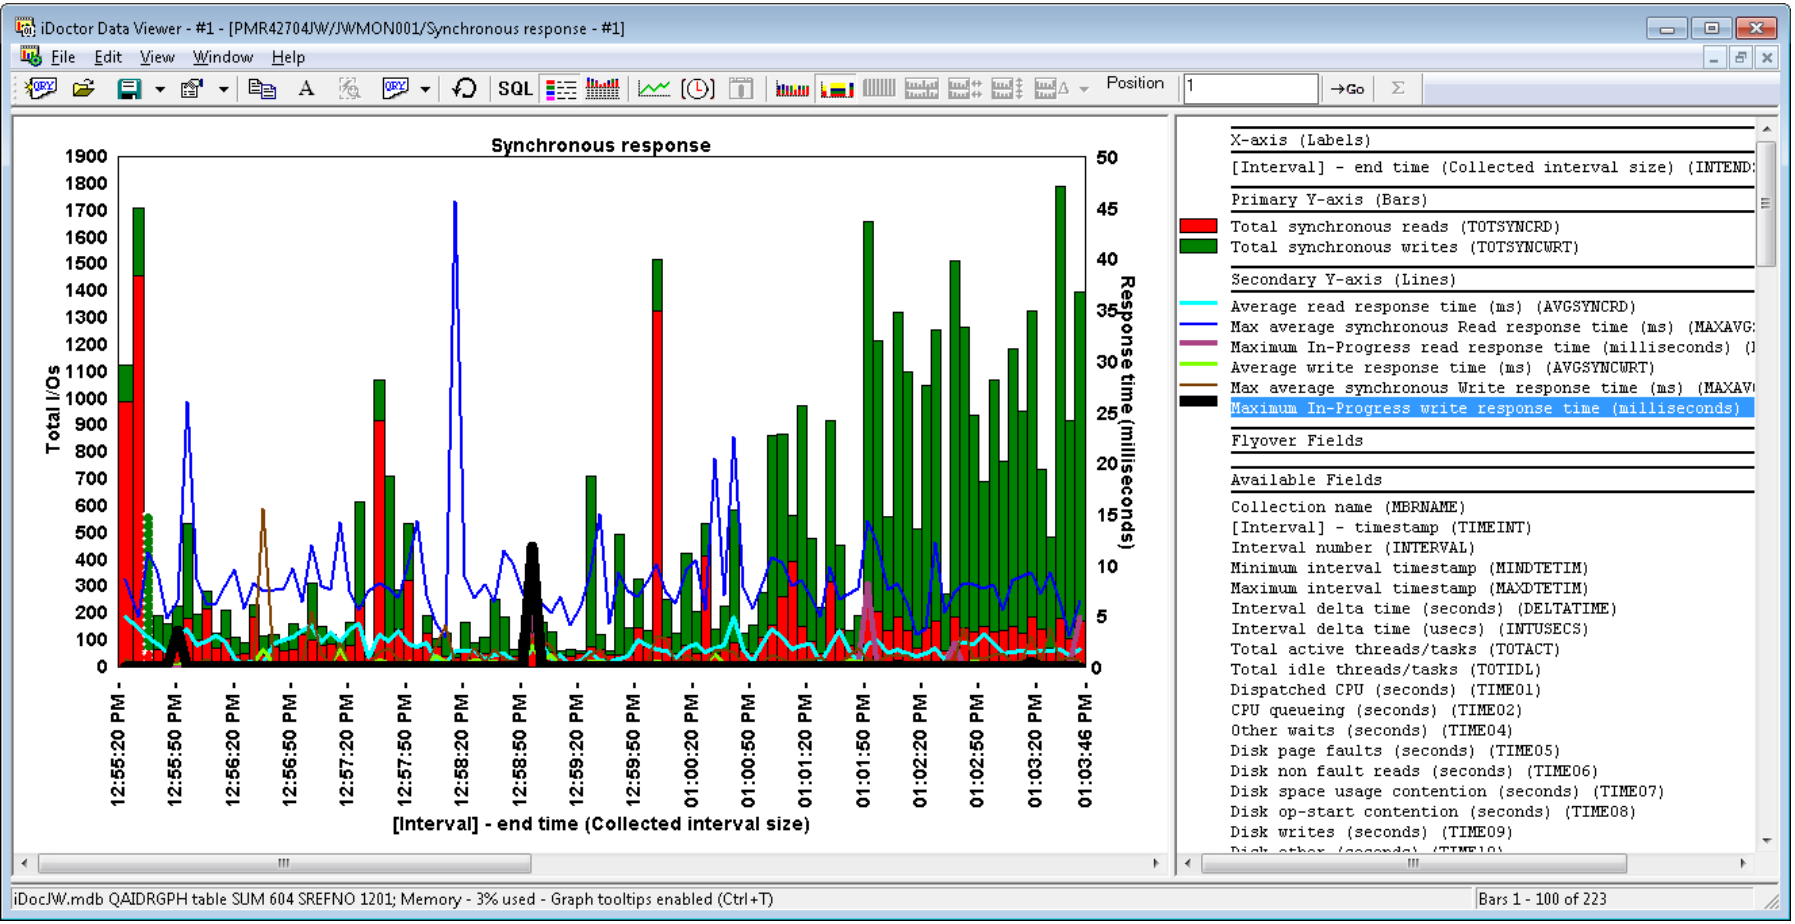

Job Watcher – Synchronous response graph drill down for “in progress I/O”

- Use Detail reports - Waits – Current wait details for synchronous reads/writes

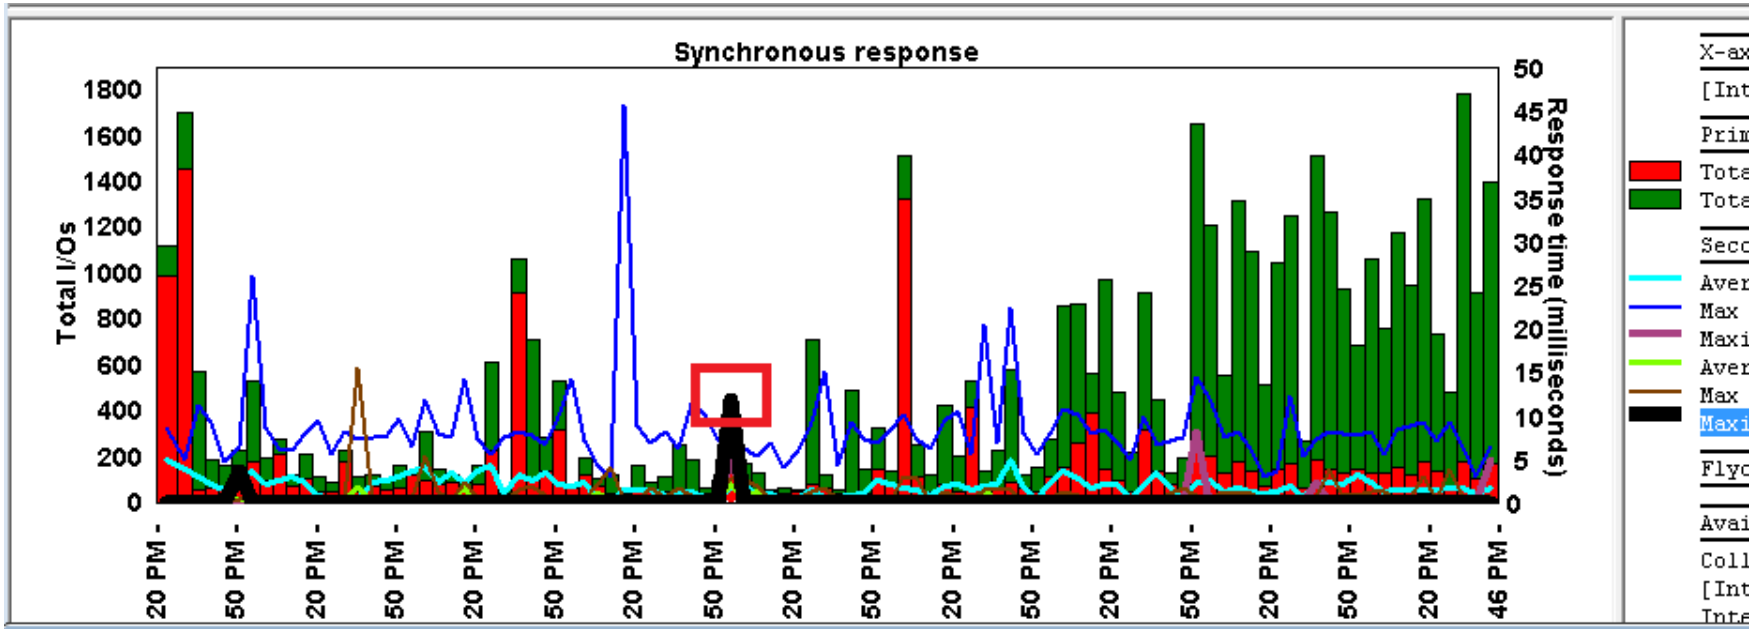

PMR42704JW/JWMON001/Current wait details for synchronous reads/writes - #1

Interval number (INTERVAL)	Job name/user/number: thread ID (JTTHREAD)	Current wait duration (usecs) (CURRWTDUR)	BLOCKBCKT	ENUM	LICWO	Wait object name (WOOBJNAM)
43	LOGASIO / QNOTES / 051552: 00000004	11,855	9/Disk writes	174	GTA	
43	SERVER / QNOTES / 051491: 00000032	4,993	5	161	SFt	TitleTReque:982A6AD7 000000.

Job Watcher – Synchronous response graph updates

- **I/O and memory page graphs – New Rankings graphs:**
  - Synchronous response (sorted by avg read)
  - Synchronous response (sorted by avg write)
  - Synchronous response (sorted by MAX avg read)
  - Synchronous response (sorted by MAX avg write)

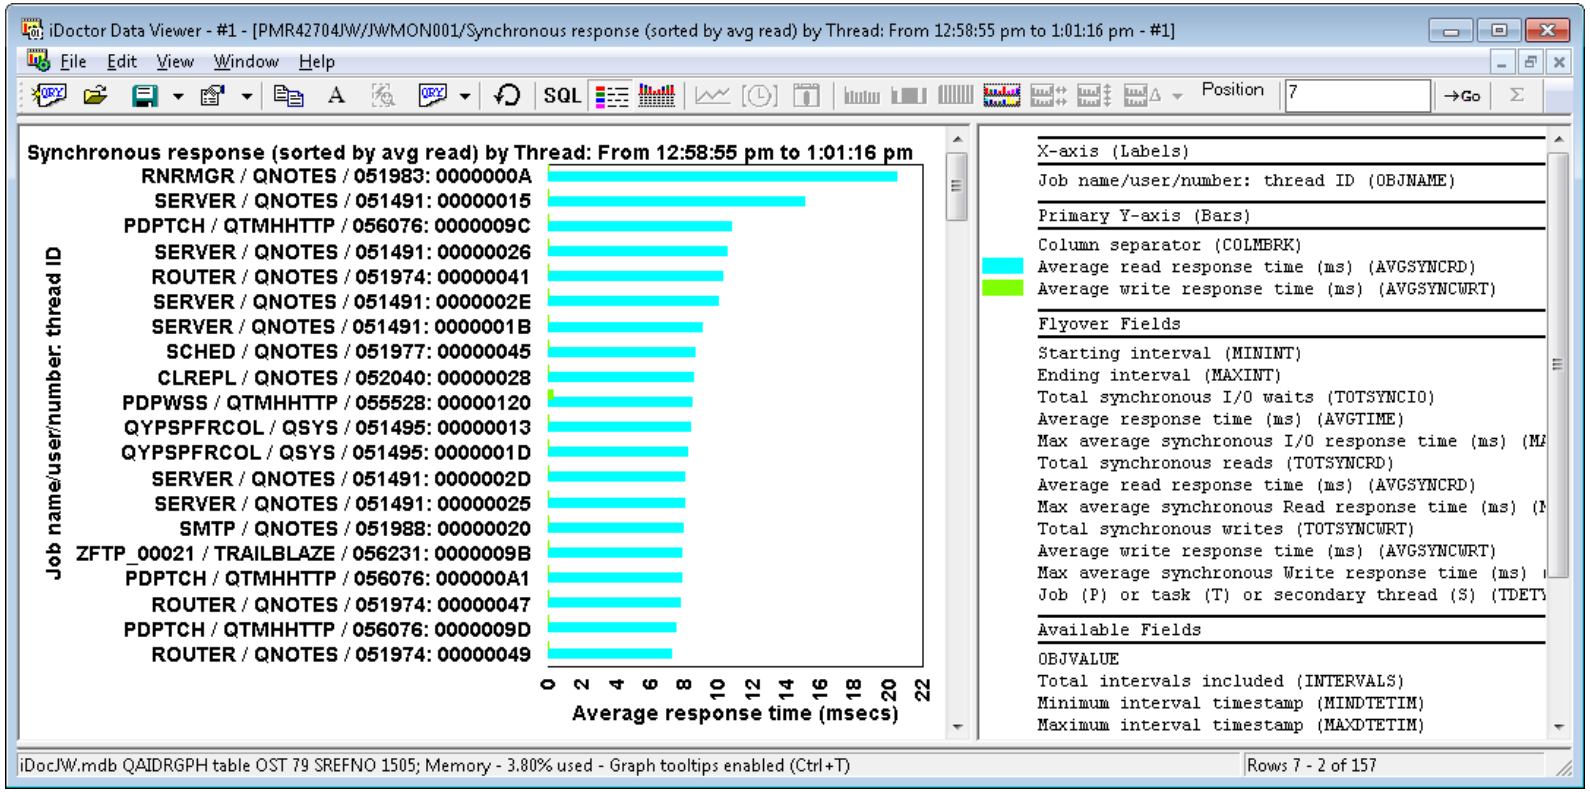

Job Watcher – Job counts

- If the collection summary analysis has been done, added a job counts graphs folder under the collection containing the following graphs:
  - Job counts
  - Net jobs created
  - Short lived job counts [excludes jobs that lived and died in a single collected interval]
  - Jobs created/destroyed

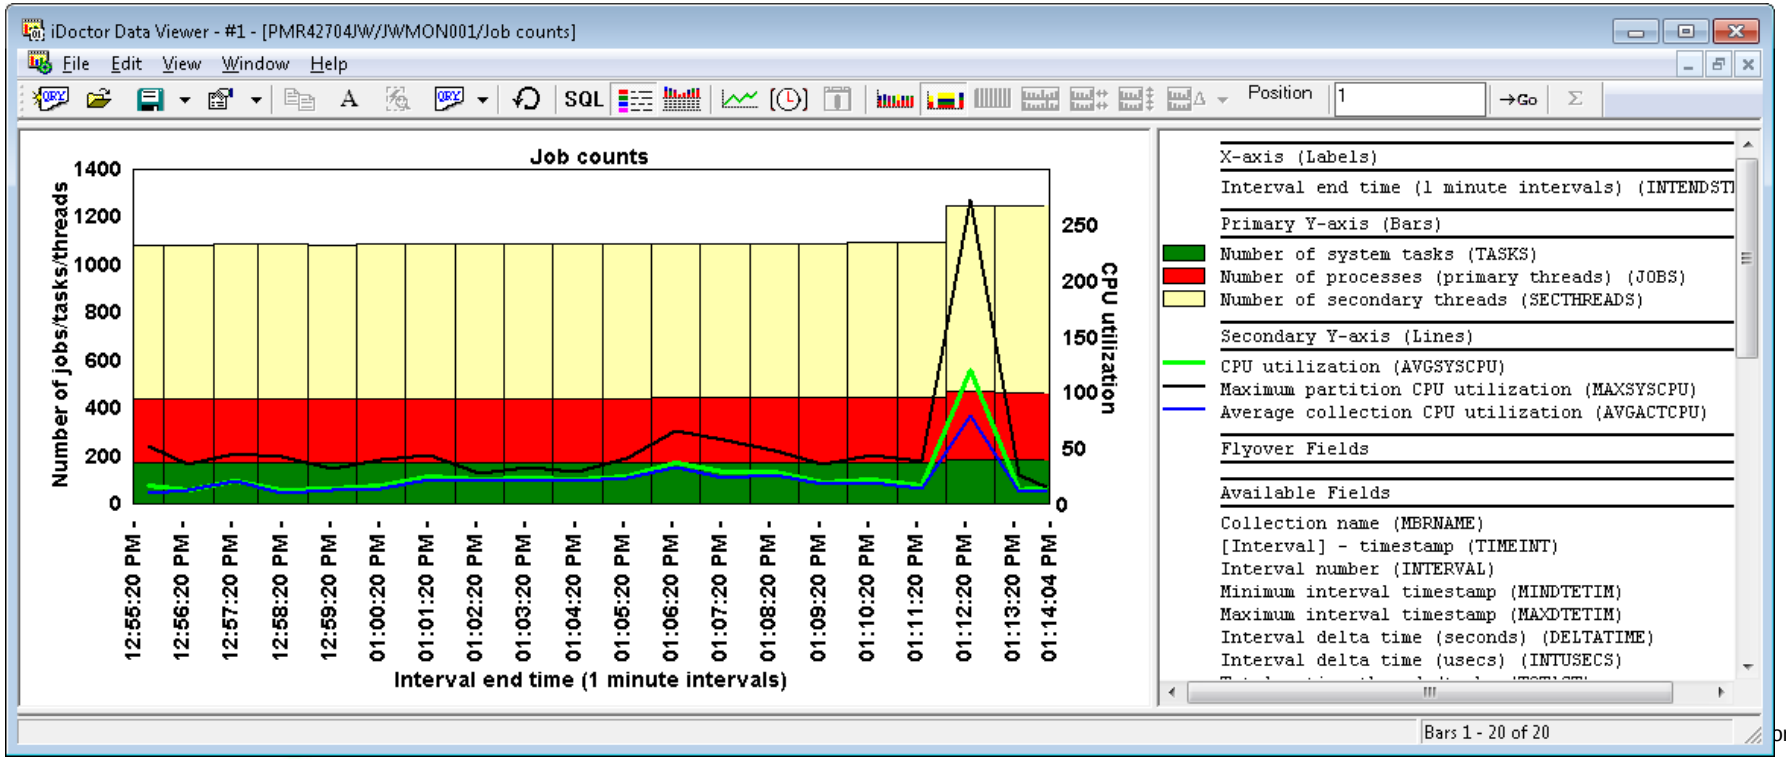

Job Watcher – Job counts – Net jobs created

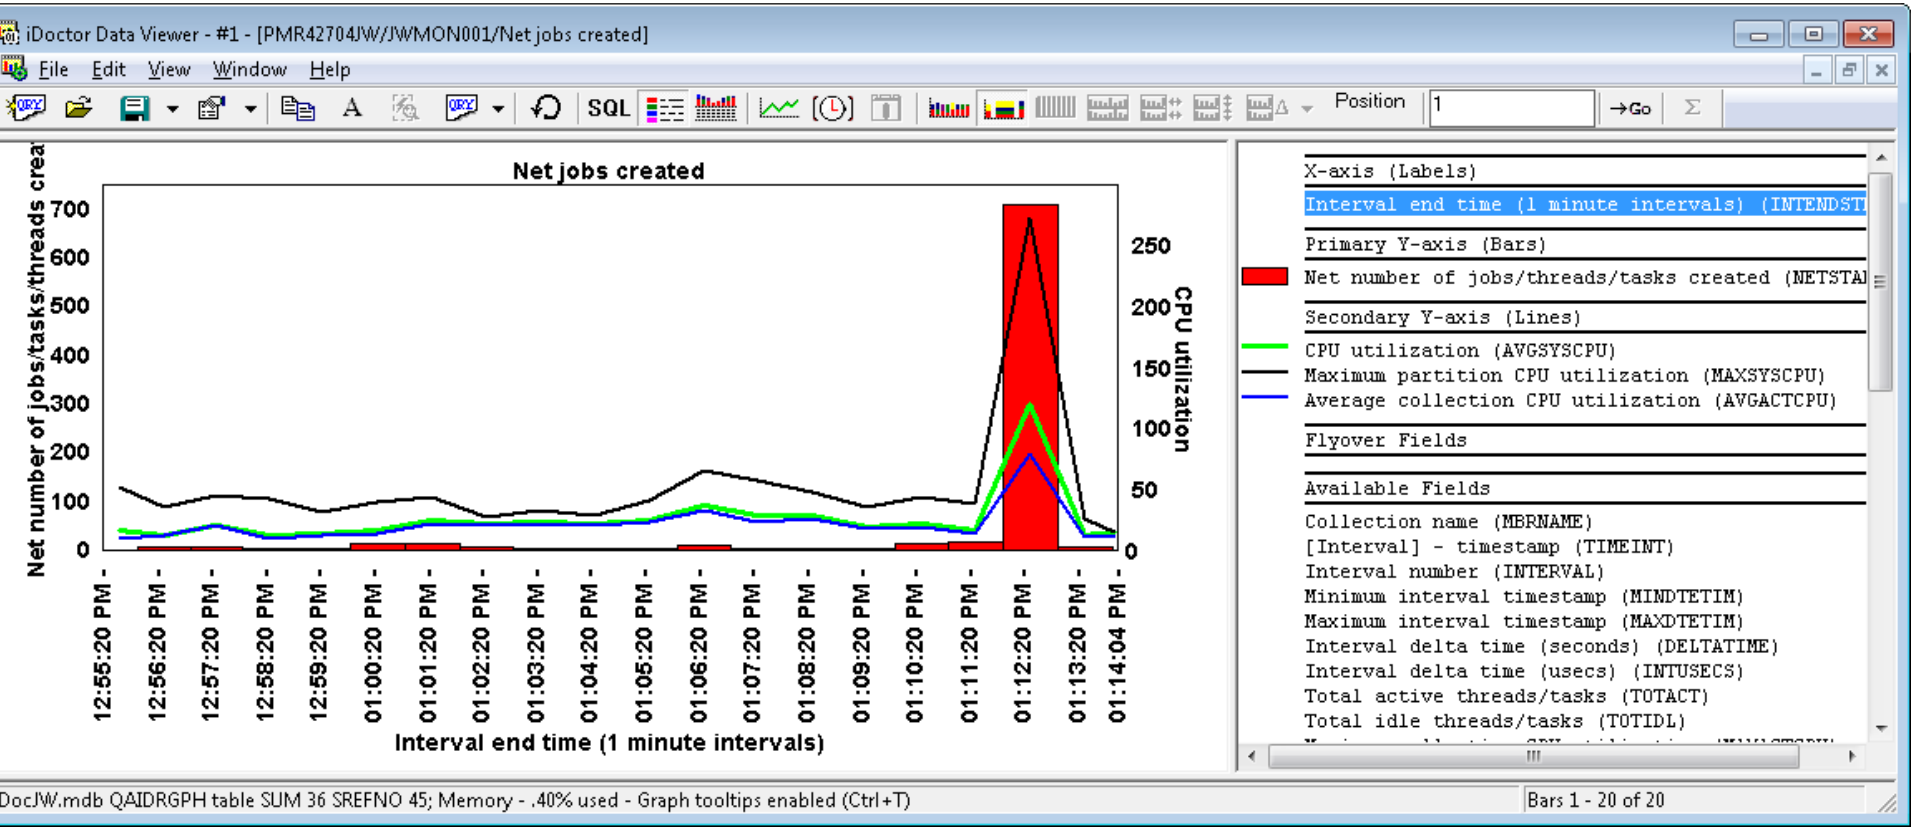

Job Watcher – Job counts drill downs

- Added drill downs called job counts by job, job counts by generic job, etc.
- Also added Short lived job counts by thread, Short lived job counts by generic job, etc

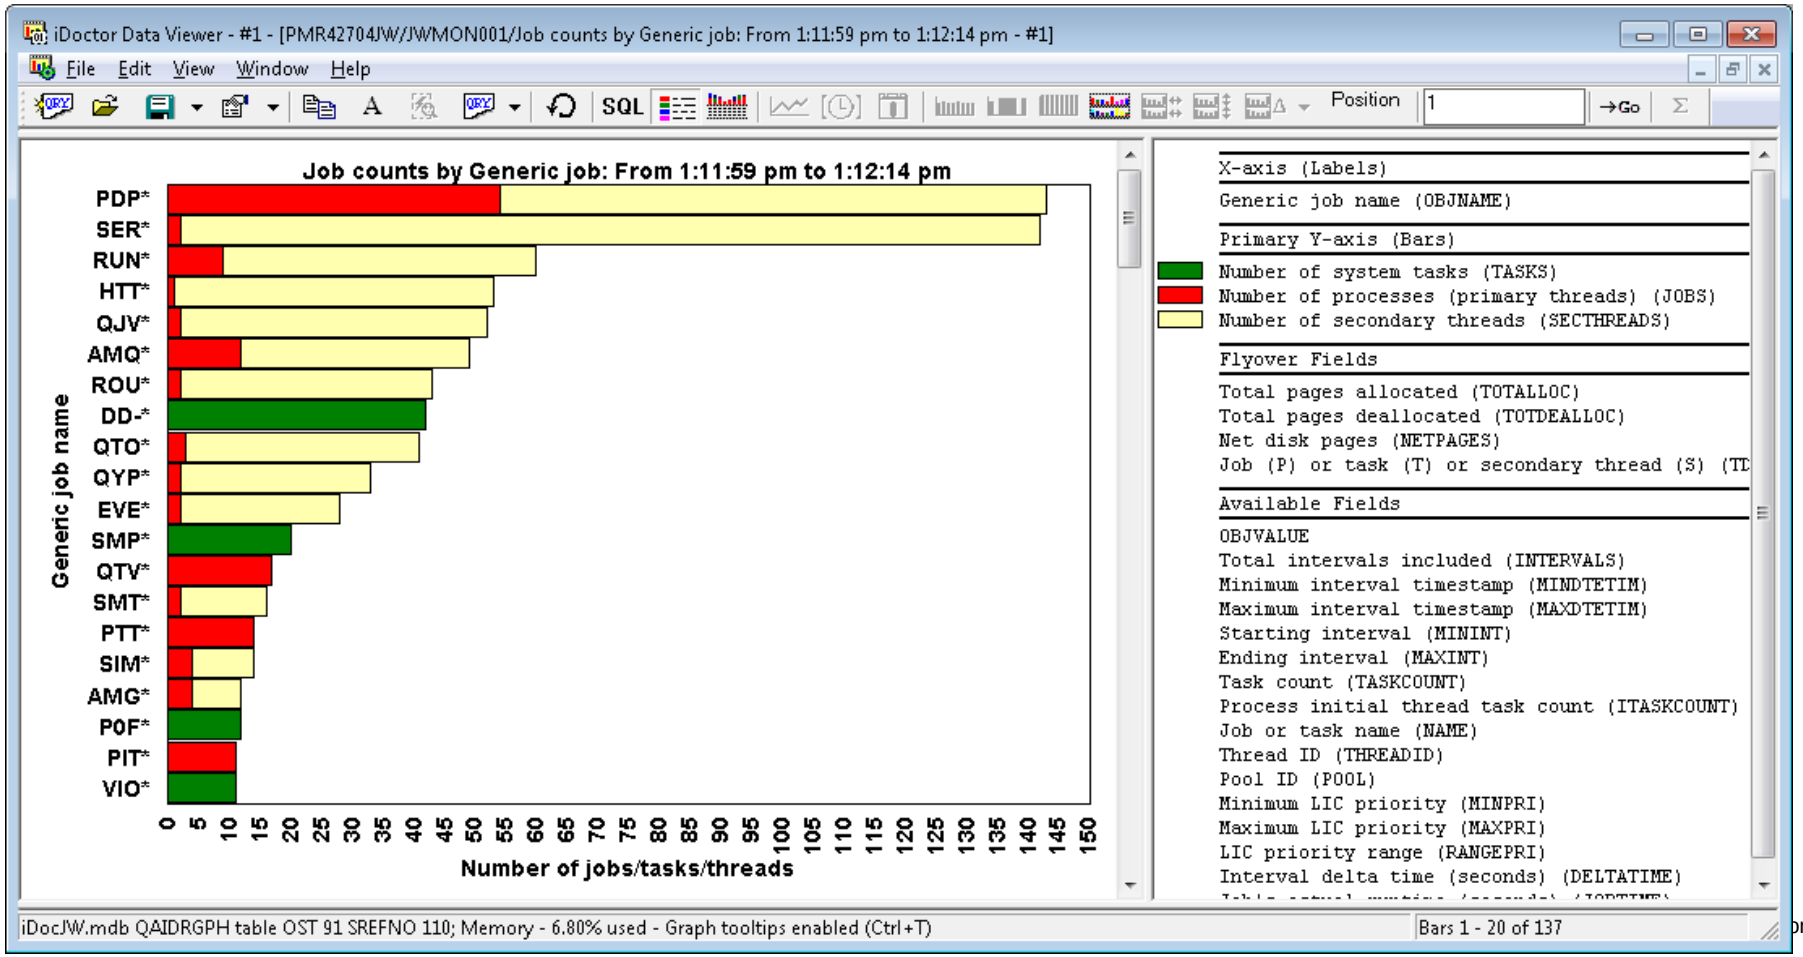

Job Watcher – Create job summary analysis changes

- The JW create job summary analysis now provides the same enhancements recently added to the CSI create job summary. This applies to 5.4 and up Job Watcher. The following filters may be used:
  - Up to 10 job name filters
  - Up to 10 current user profiles
  - Up to 10 subsystem names
  - Start time, end time
  - Minimum run time (hours) and
  - Minimum CPU used (seconds)

- For drill downs from the table views you can now drill down into the selected thread(s) over time

Job Watcher – Create job summary analysis report changes

- The report options have been simplified with the graphing options reduced to "wait graphs rankings" and "other graphs rankings". Each of these menus has grouping options such as "by job", "by collection, thread", "by collection, job".
- Note: These changes also apply to Collection Services Investigator.

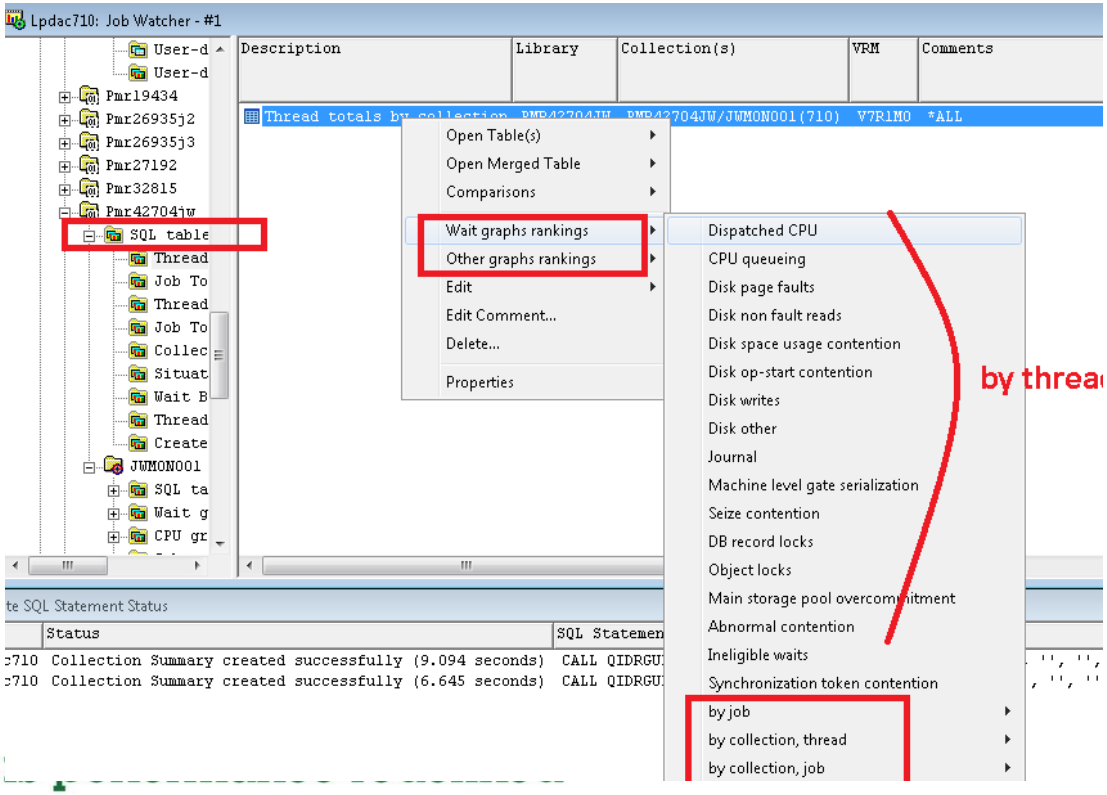

Job Watcher – Create job summary analysis report changes

- When drilling down from tables in the Data Viewer a menu called "Filter by" appears which lets you control whether all jobs, selected jobs, or whether you will be prompted for a generic job filter when determining what to include on the drill down graphs called "Wait graphs rankings" or "Other graphs rankings".
- Note: These changes also apply to Collection Services Investigator.

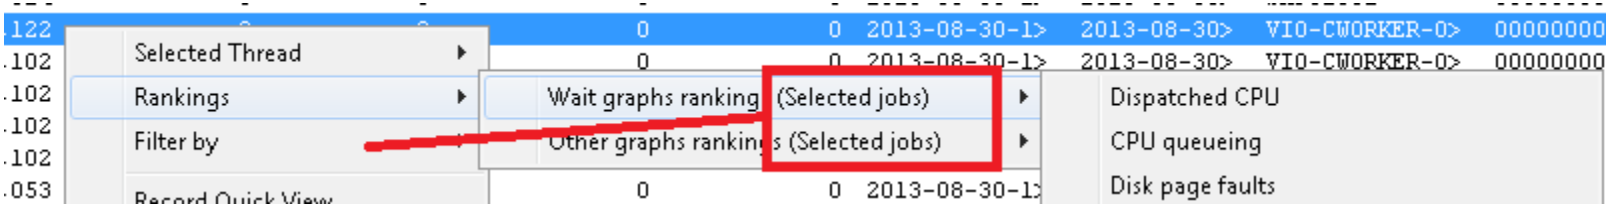

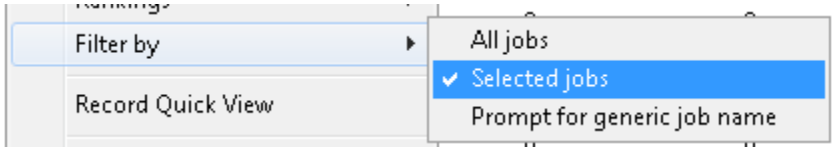

Job Watcher – Call stack summary analysis changes

- You can now specify either 16 or 50 call levels when right-clicking a collection.

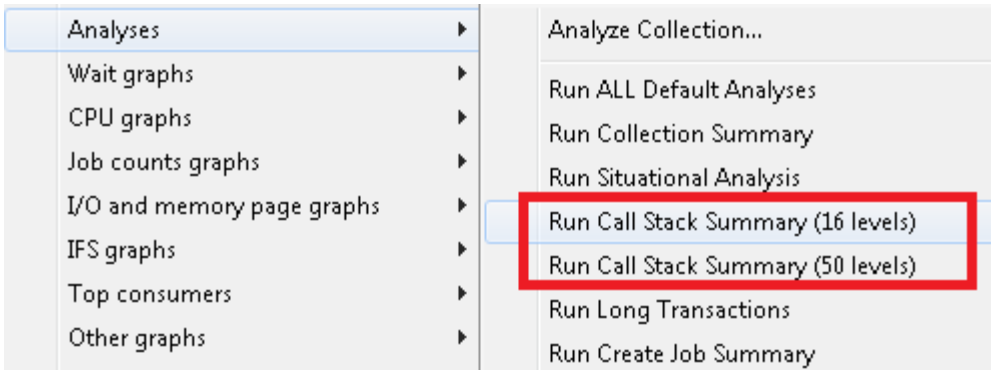

- Also if you want more levels than 50 (or a different value) you can run the stored procedure manually using the following examples in any SQL editor:
  - CALL QIDRGUI/QIDRJWCSS ('LIB', 'COL', 10)
  - CALL QIDRGUI/QIDRJWCSS ('LIB', 'COL', 35)
  - CALL QIDRGUI/QIDRJWCSS ('LIB', 'COL', 99)

Job Watcher – New wait graphs

- Added the following new graphs under the wait graphs folder:
  - Collection overview time signature with max waits in-progress
  - Current wait duration time signature with max waits in-progress
  - Disk time signature with max disk waits in-progress

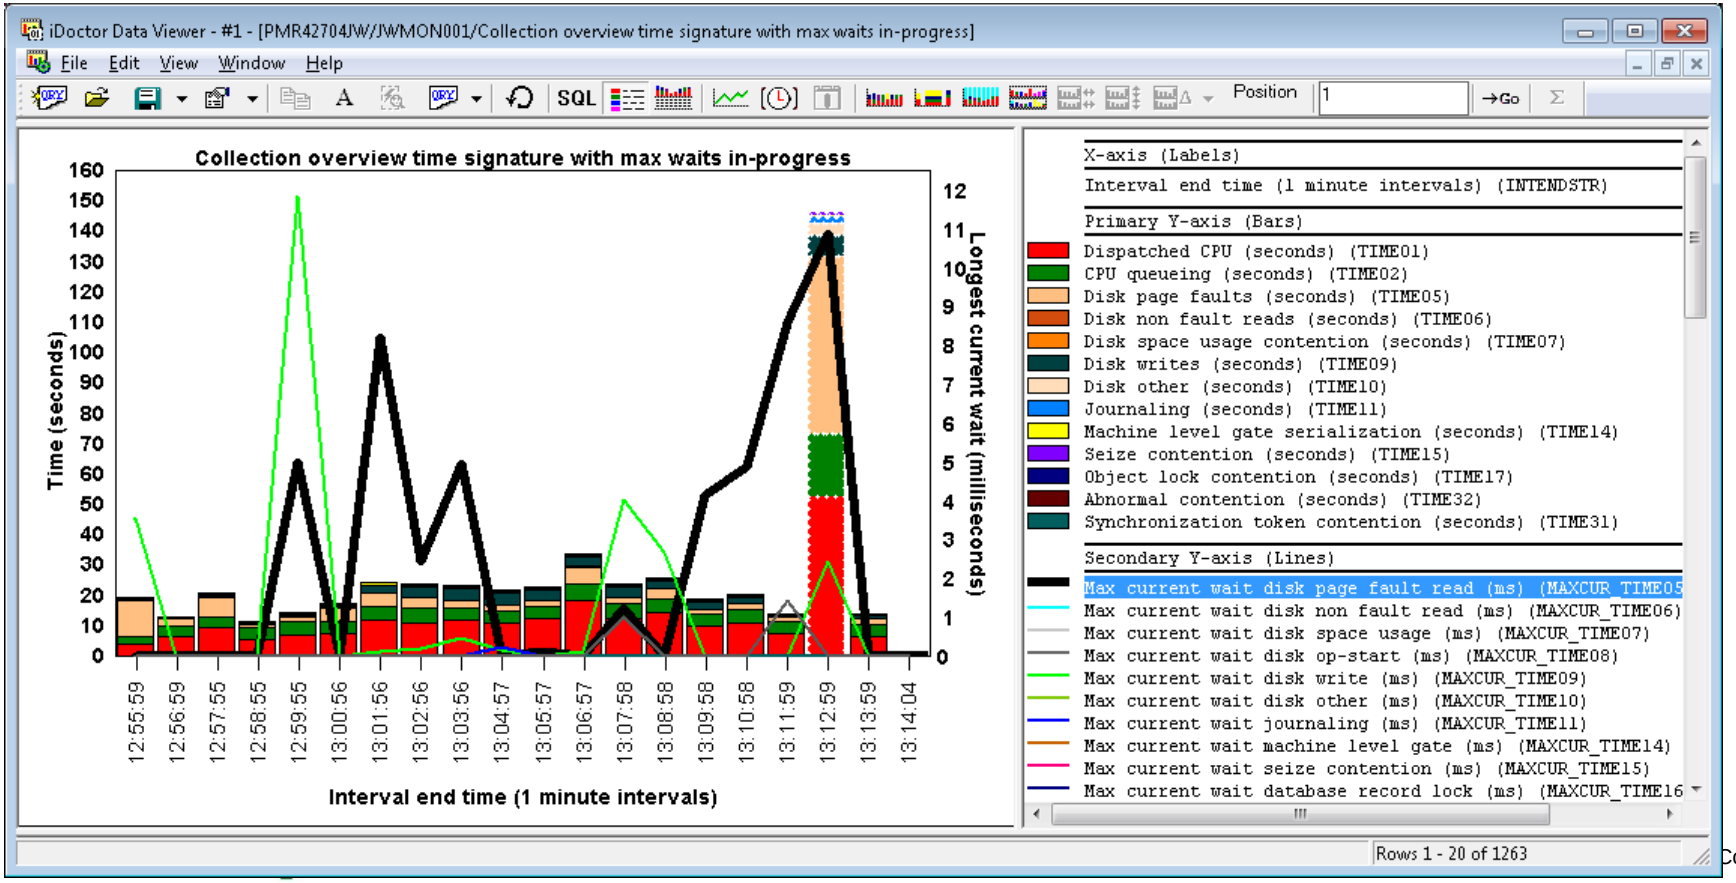

Job Watcher – New wait graphs

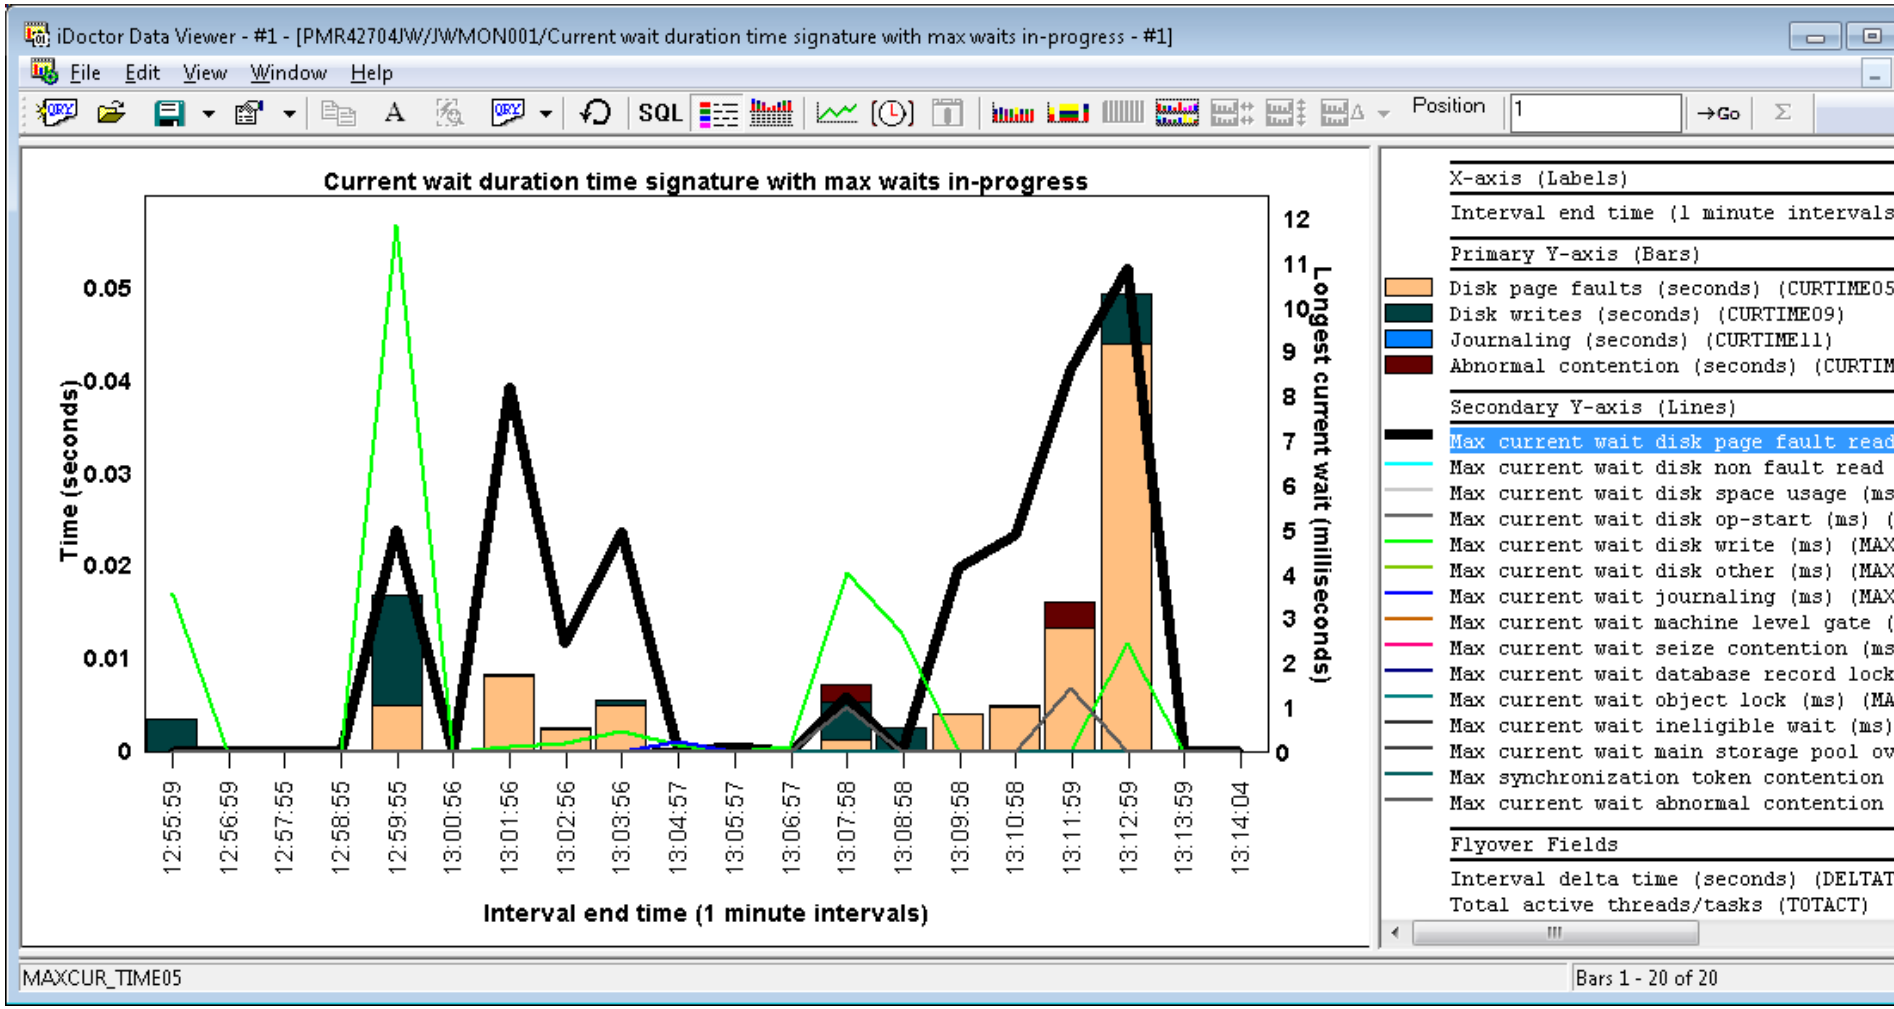

Job Watcher – New wait graphs

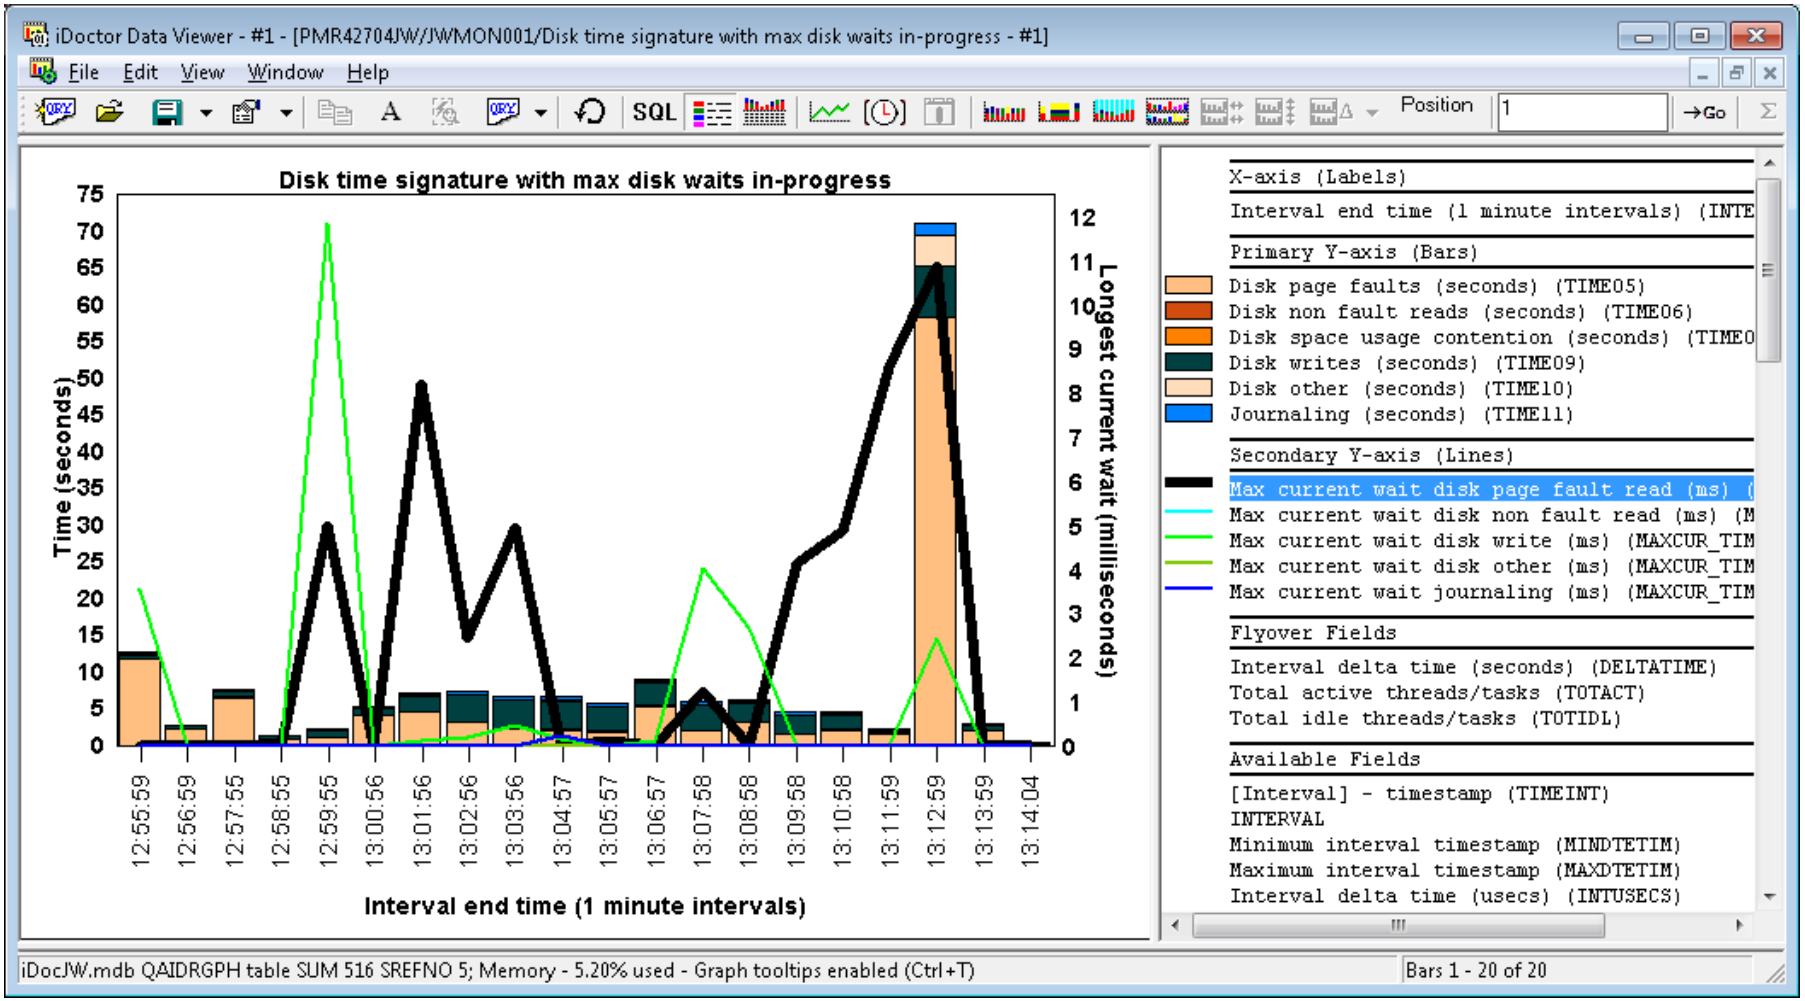

Collection Services Investigator – Ethernet graphs

- Added the following Ethernet LAN usage graphs to iDoctor under the communication graphs folder:
  - Ethernet LAN megabits per second by IOP/line
  - Ethernet LAN frames by IOP/line
  - Ethernet LAN frame rates by IOP/line
  - Ethernet LAN congestion by IOP/line
  - Ethernet LAN megabits per second by IOP
  - Ethernet LAN frames by IOP
  - Ethernet LAN frame rates by IOP
  - Ethernet LAN congestion by IOP

  - Note: All of these have the utilization rates on the 2nd Y-axis. (Note: if you want to see the IOA / LINXx value, see the flyover.)

Collection Services Investigator – Ethernet graphs

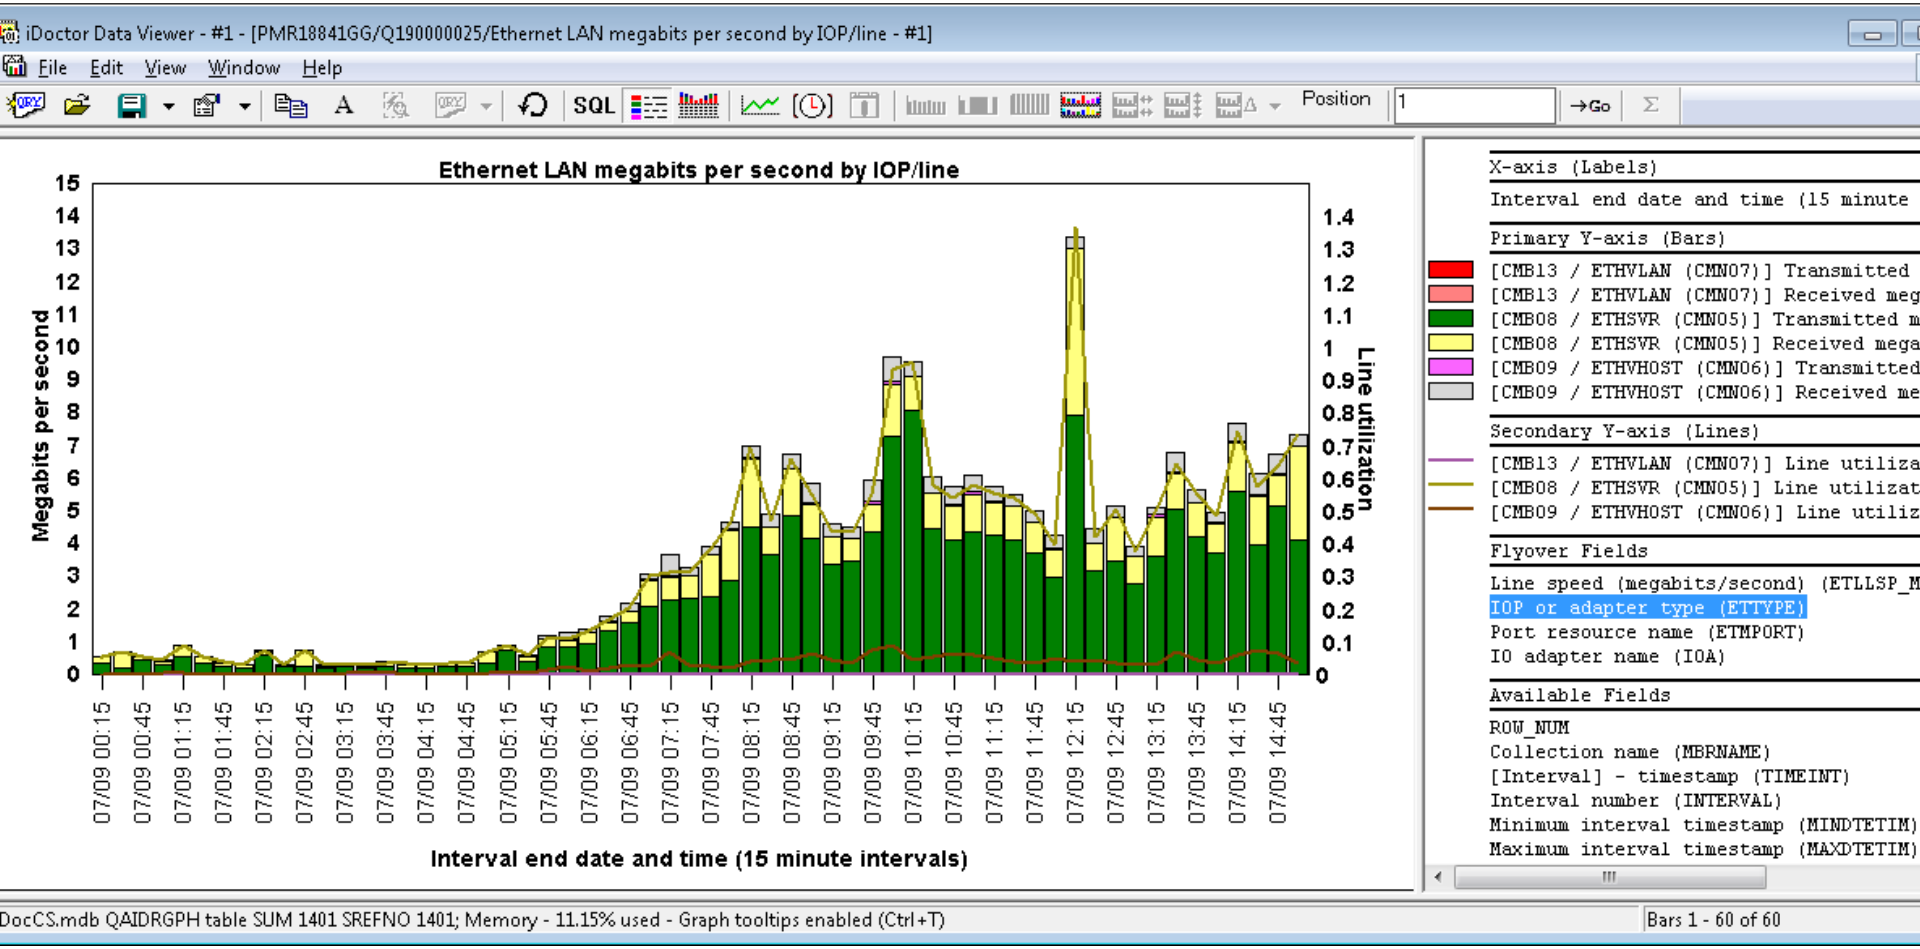

Collection Services Investigator – TLBIE graphs

- **New at 7.1 only if the file QAPMSYSINT exists with the required data.**
- **You can copy CS collections from other partitions into the same library to view the data from multiple LPARs in the same graph.**

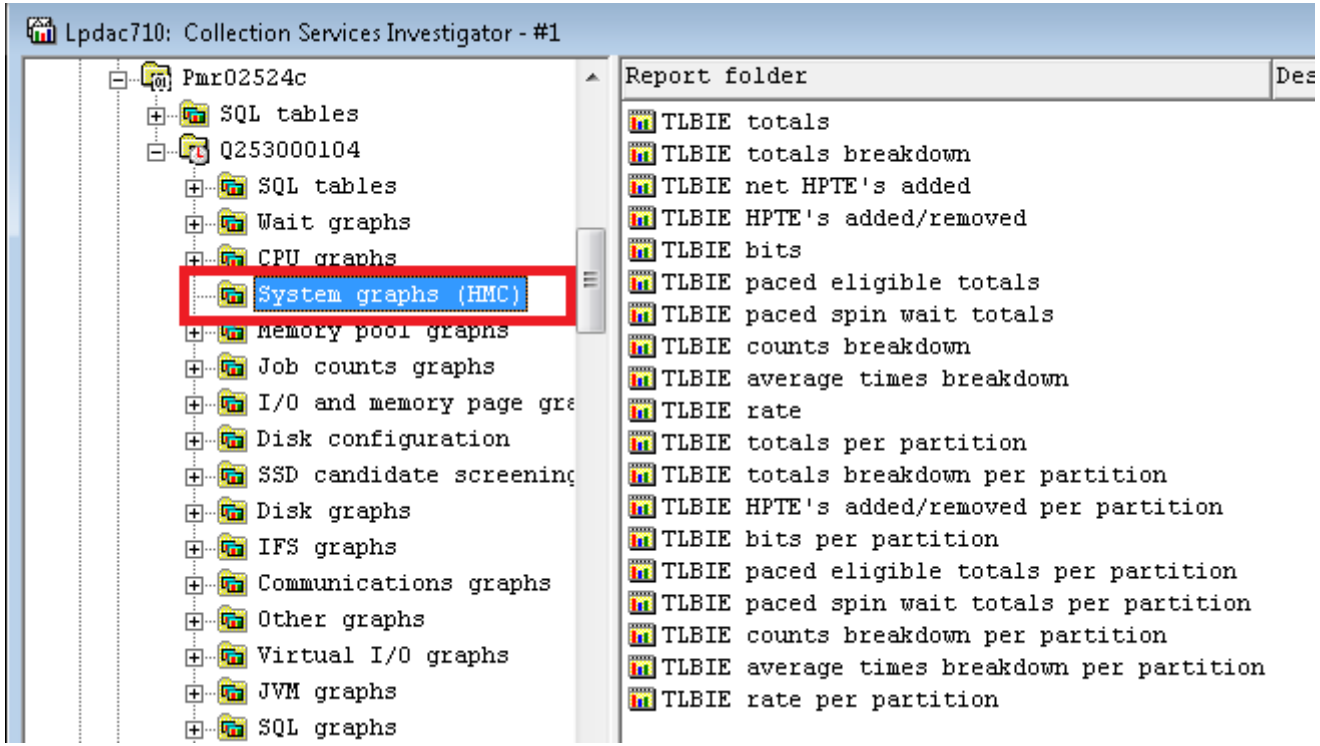

Collection Services Investigator – System graphs - TLBIE totals

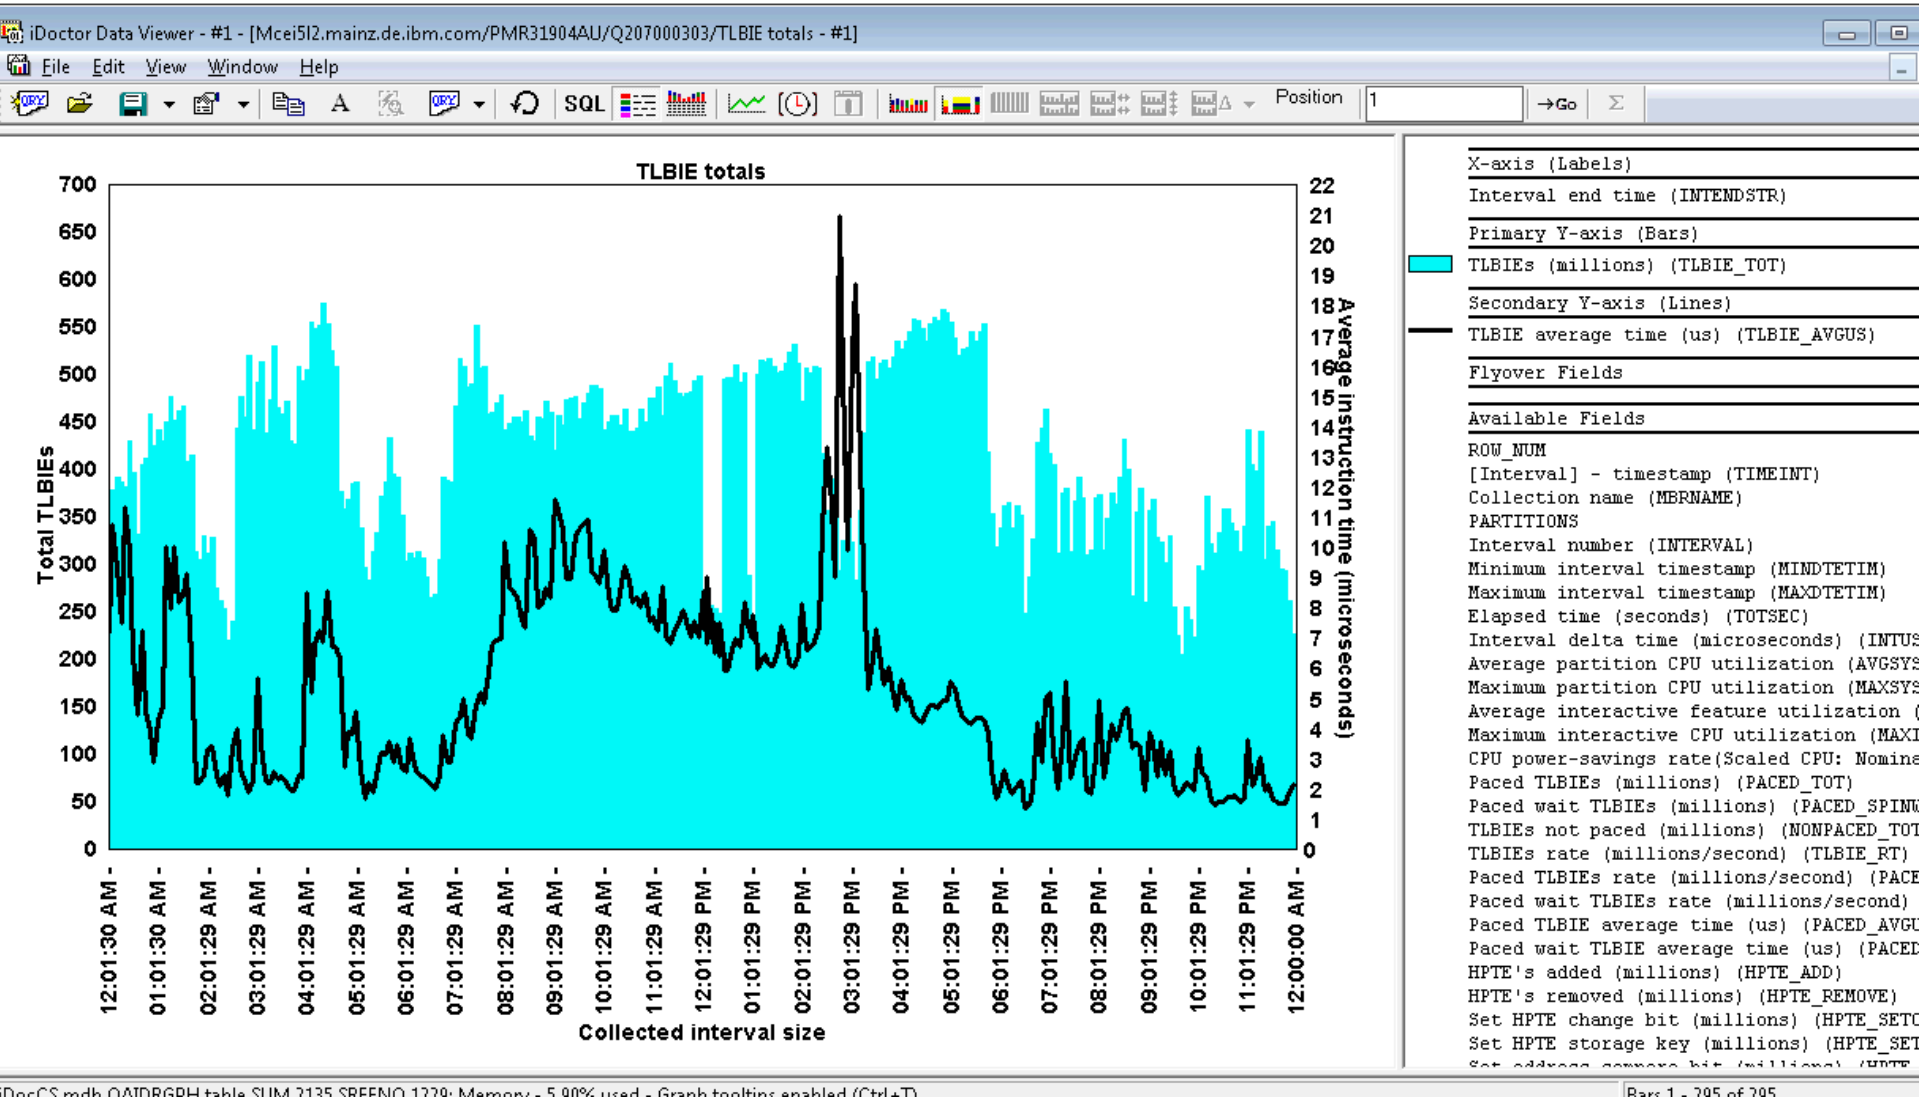

Collection Services Investigator – System graphs - TLBIE totals per partition

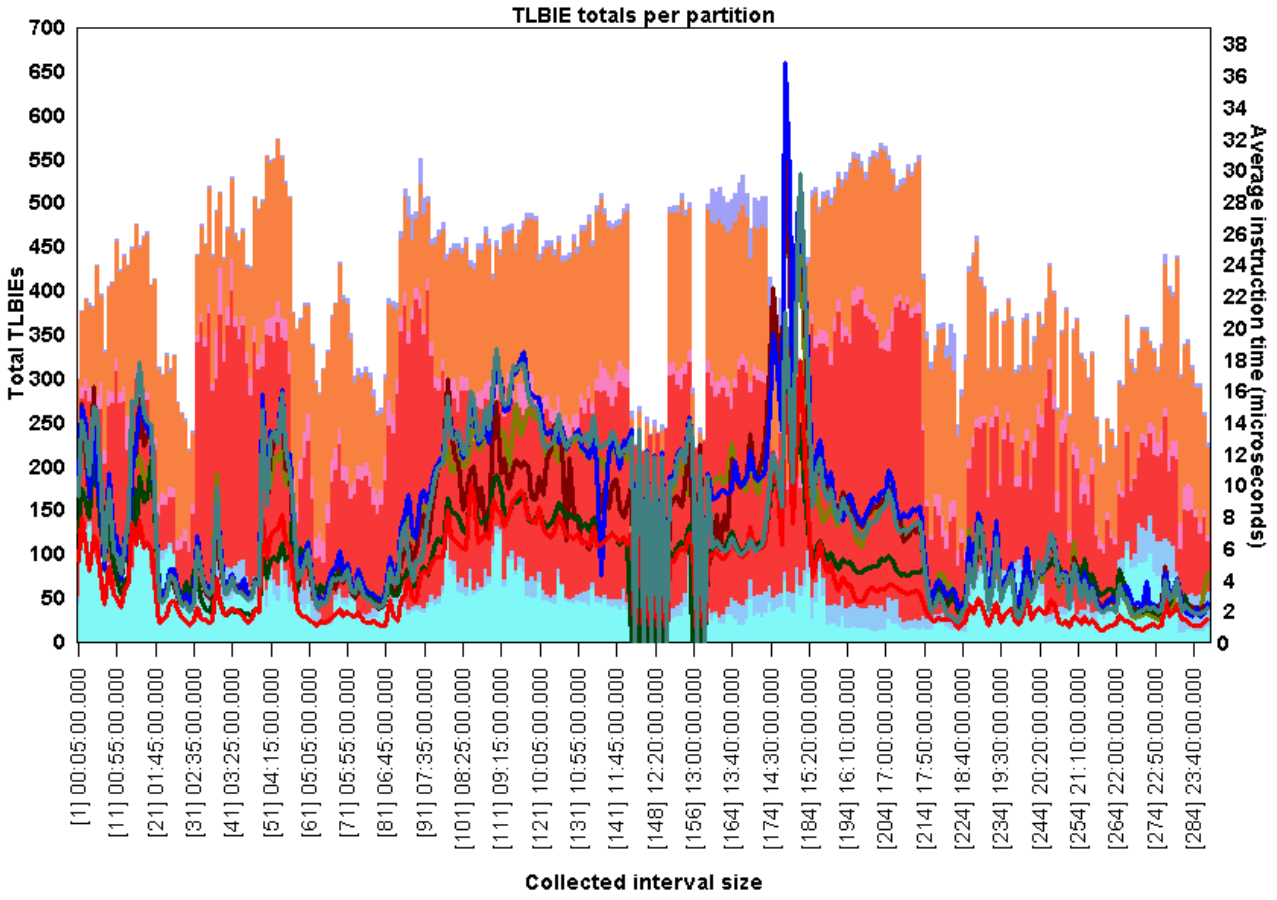

X-axis (Labels)

Interval end time (INTENDSTR)

Primary Y-axis (Bars)

- [D005D1D6] TLBIEs (millions) (TLBIE_TOT)
- [APSI0042] TLBIEs (millions) (TLBIE_TOT)
- [E005D1D6] TLBIEs (millions) (TLBIE_TOT)
- [APSI0040] TLBIEs (millions) (TLBIE_TOT)
- [B005D1D6] TLBIEs (millions) (TLBIE_TOT)
- [APSI0041] TLBIEs (millions) (TLBIE_TOT)

Secondary Y-axis (Lines)

- [D005D1D6] TLBIE average time (us) (TLBIE_AVGUS)
- [APSI0042] TLBIE average time (us) (TLBIE_AVGUS)
- [E005D1D6] TLBIE average time (us) (TLBIE_AVGUS)
- [APSI0040] TLBIE average time (us) (TLBIE_AVGUS)
- [B005D1D6] TLBIE average time (us) (TLBIE_AVGUS)
- [APSI0041] TLBIE average time (us) (TLBIE_AVGUS)

Flyover Fields

Available Fields

ROW_NUM
[Interval] - timestamp (TIMEINT)
Collection name (MBRNAME)
Interval number (INTERVAL)
Minimum interval timestamp (MINDTETIM)
Maximum interval timestamp (MAXDTETIM)
Elapsed time (seconds) (TOTSEC)
Interval delta time (microseconds) (INTUSECS)
Average partition CPU utilization (AVGSYSCPU)
Maximum partition CPU utilization (MAXSYSCPU)
Average interactive feature utilization (AVGINTCPU
Maximum interactive CPU utilization (MAXINTCPU)
CPU power-savings rate(Scaled CPU: Nominal CPU) (S
TLBIEs (millions) (TLBIE_TOT)
Paced TLBIEs (millions) (PACED_TOT)
Paced wait TLBIEs (millions) (PACED_SPINWAITTOT)
TLBIEs not paced (millions) (NONPACED_TOT)
TLBIEs rate (millions/second) (TLBIE_RT)

Collection Services Investigator – System graphs - TLBIE totals breakdown

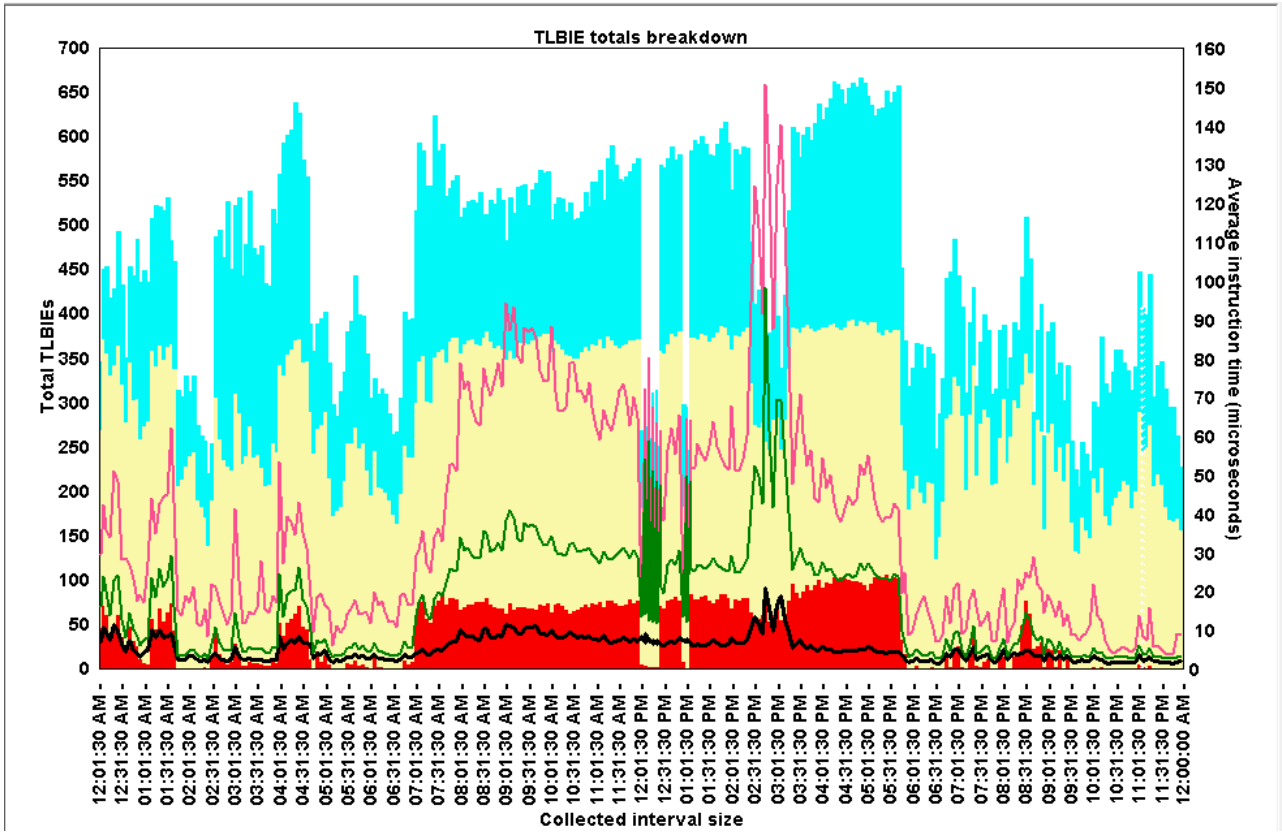

X-axis (Labels)

Interval end time (INTENDSTR)

Primary Y-axis (Bars)

- Paced wait TLBIEs (millions) (PACED_SPINWAITTOT)
- Paced TLBIEs (millions) (PACED_TOT)
- TLBIEs not paced (millions) (NONPACED_TOT)

Secondary Y-axis (Lines)

- TLBIE average time (us) (TLBIE_AVGUS)
- Paced wait TLBIE average time (us) (PACED_AVGSPINWAI
- Paced TLBIE average time (us) (PACED_AVGUS)

Flyover Fields

Available Fields

ROW_NUM
[Interval] - timestamp (TIMEINT)
Collection name (MBRNAME)
PARTITIONS
Interval number (INTERVAL)
Minimum interval timestamp (MINDTETIM)
Maximum interval timestamp (MAXDTETIM)
Elapsed time (seconds) (TOTSEC)
Interval delta time (microseconds) (INTUSECS)
Average partition CPU utilization (AVGSYSCPU)
Maximum partition CPU utilization (MAXSYSCPU)
Average interactive feature utilization (AVGINTCPU)
Maximum interactive CPU utilization (MAXINTCPU)
CPU power-savings rate(Scaled CPU: Nominal CPU) (SCA
TLBIEs (millions) (TLBIE_TOT)
TLBIEs rate (millions/second) (TLBIE_RT)
Paced TLBIEs rate (millions/second) (PACED_RT)
Paced wait TLBIEs rate (millions/second) (PACED_SPIN
HPTE's added (millions) (HPTE_ADD)
HPTE's removed (millions) (HPTE_REMOVE)
Set HPTE change bit (millions) (HPTE_SETCHGBIT)
Set HPTE storage key (millions) (HPTE_SETSTGBIT)
Set address compare bit (millions) (HPTE_SETADDRCMPB
Set HPTE page protect bits (millions) (HPTE_SETPGPRB
Get HPTE status bits (millions) (HPTE_GETSTS)

Collection Services Investigator- System graphs - TLBIE totals breakdown per partition

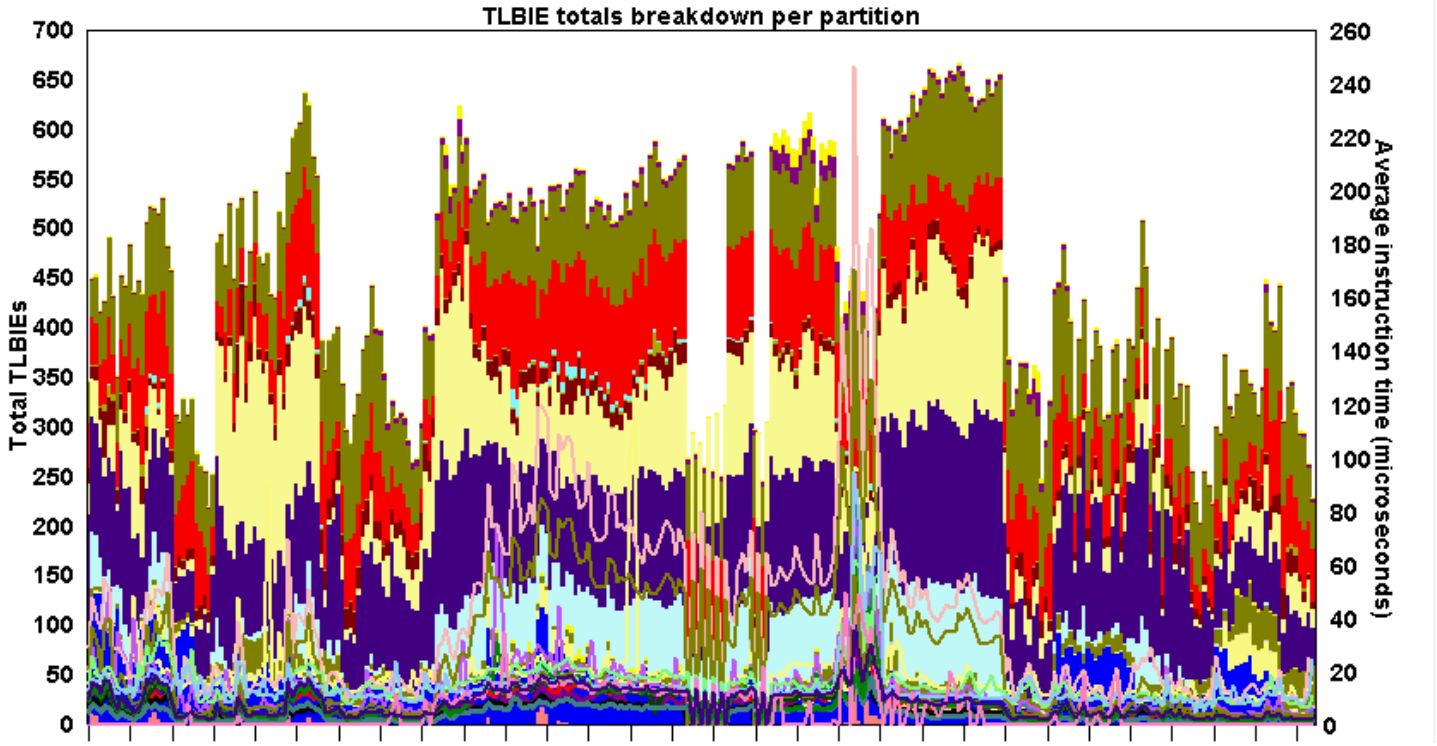

Collection Services Investigator – New disk graphs

- **Added new types of graphs that show a different color per ASP, disk type, disk group or IOA type.**

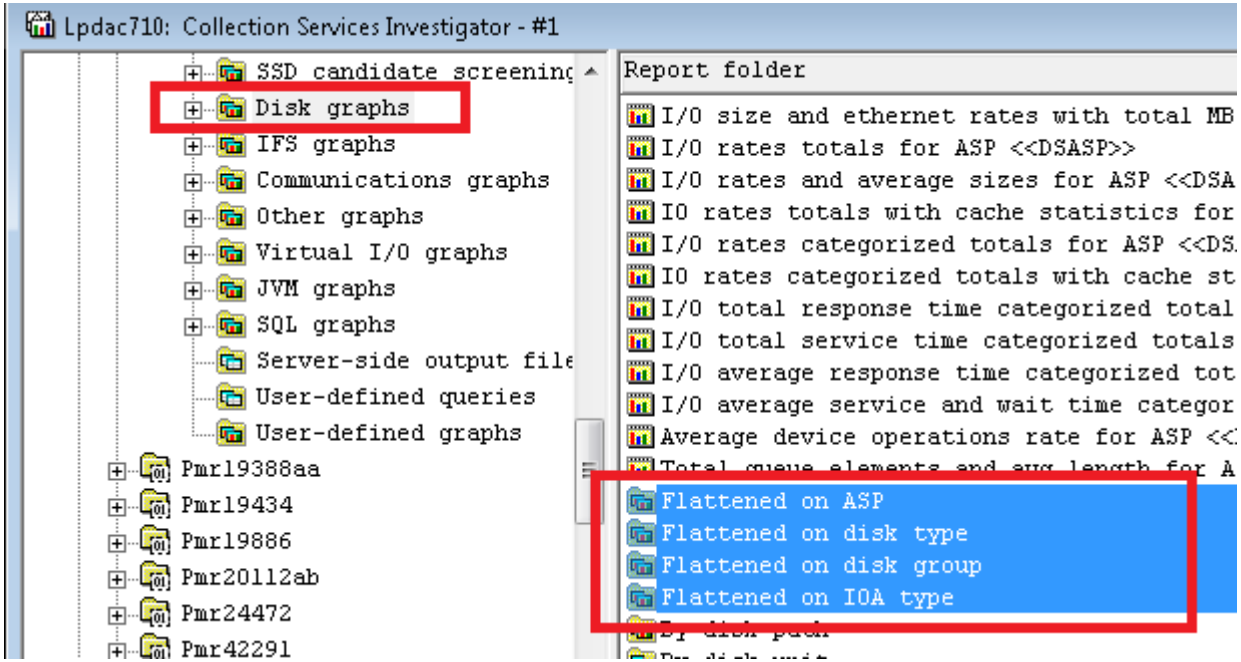

Collection Services Investigator – Disk graphs – flattened by disk type example

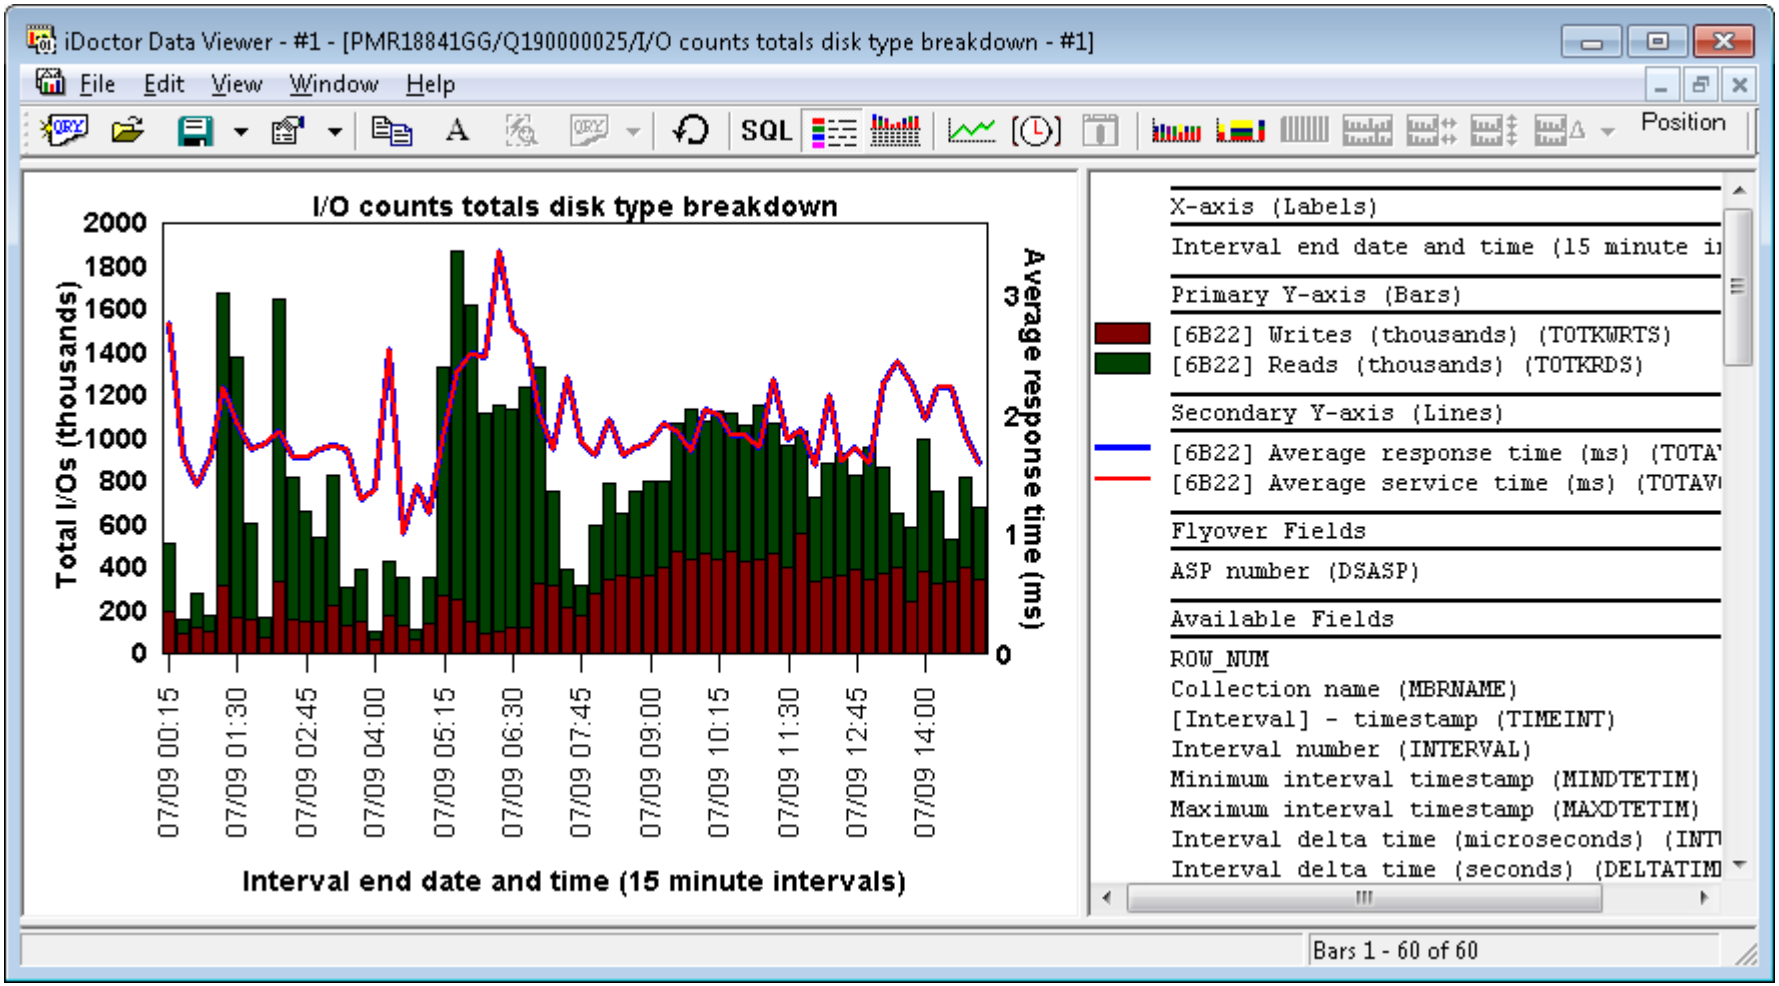

Collection Services Investigator – Memory pool graphs

- Added new graphs:
  - Memory pool consumption
  - Memory pool sizes
  - Flattened type - Memory pool consumption
  - Flattened type - Memory pool unallocated space available

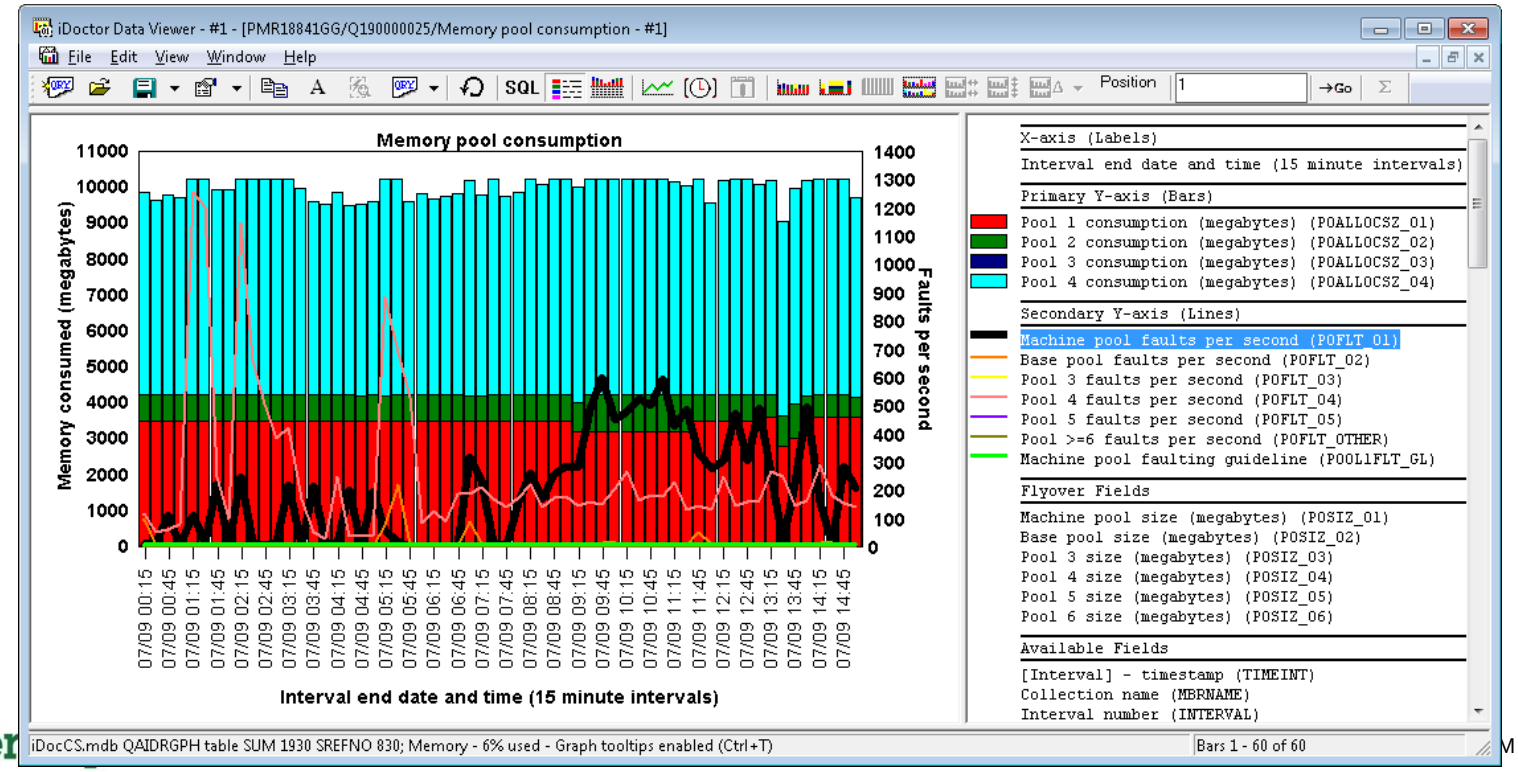

Collection Services Investigator – Memory pool graphs

– Memory pool sizes

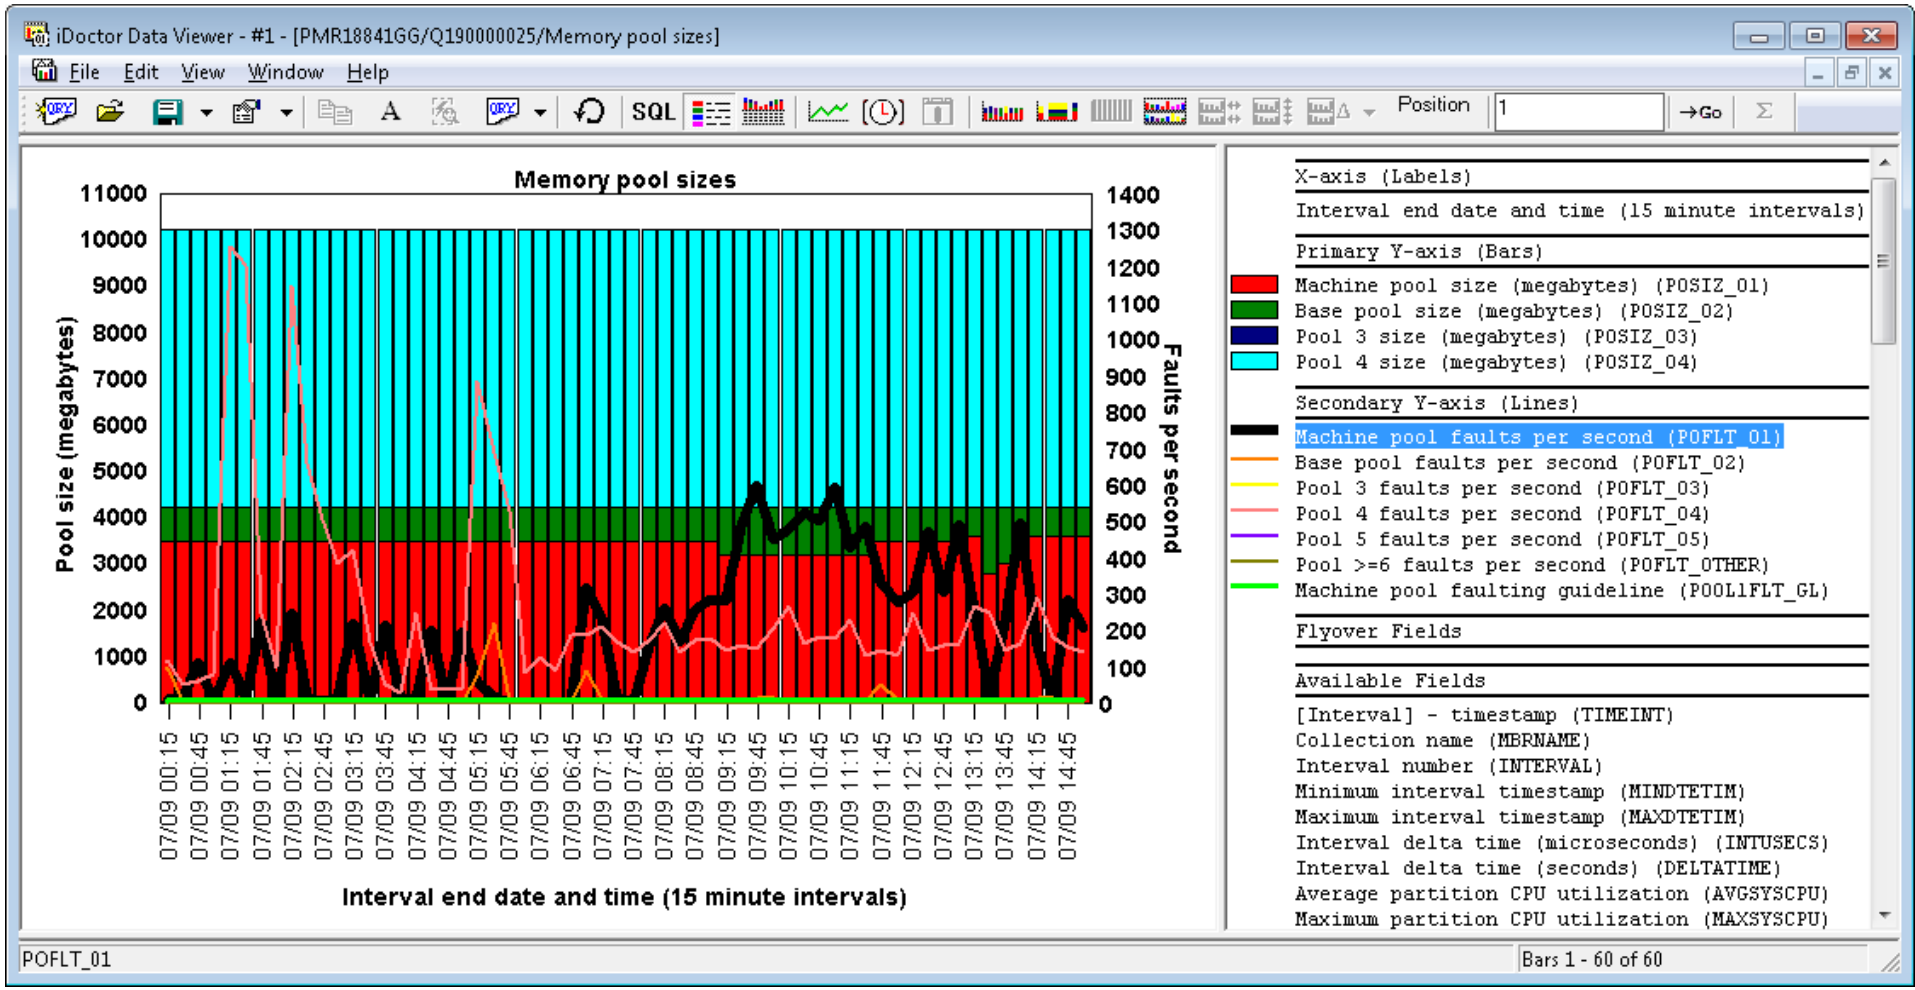

Collection Services Investigator – Memory pool graphs

- Added new graphs:
  - Flattened type - Memory pool consumption

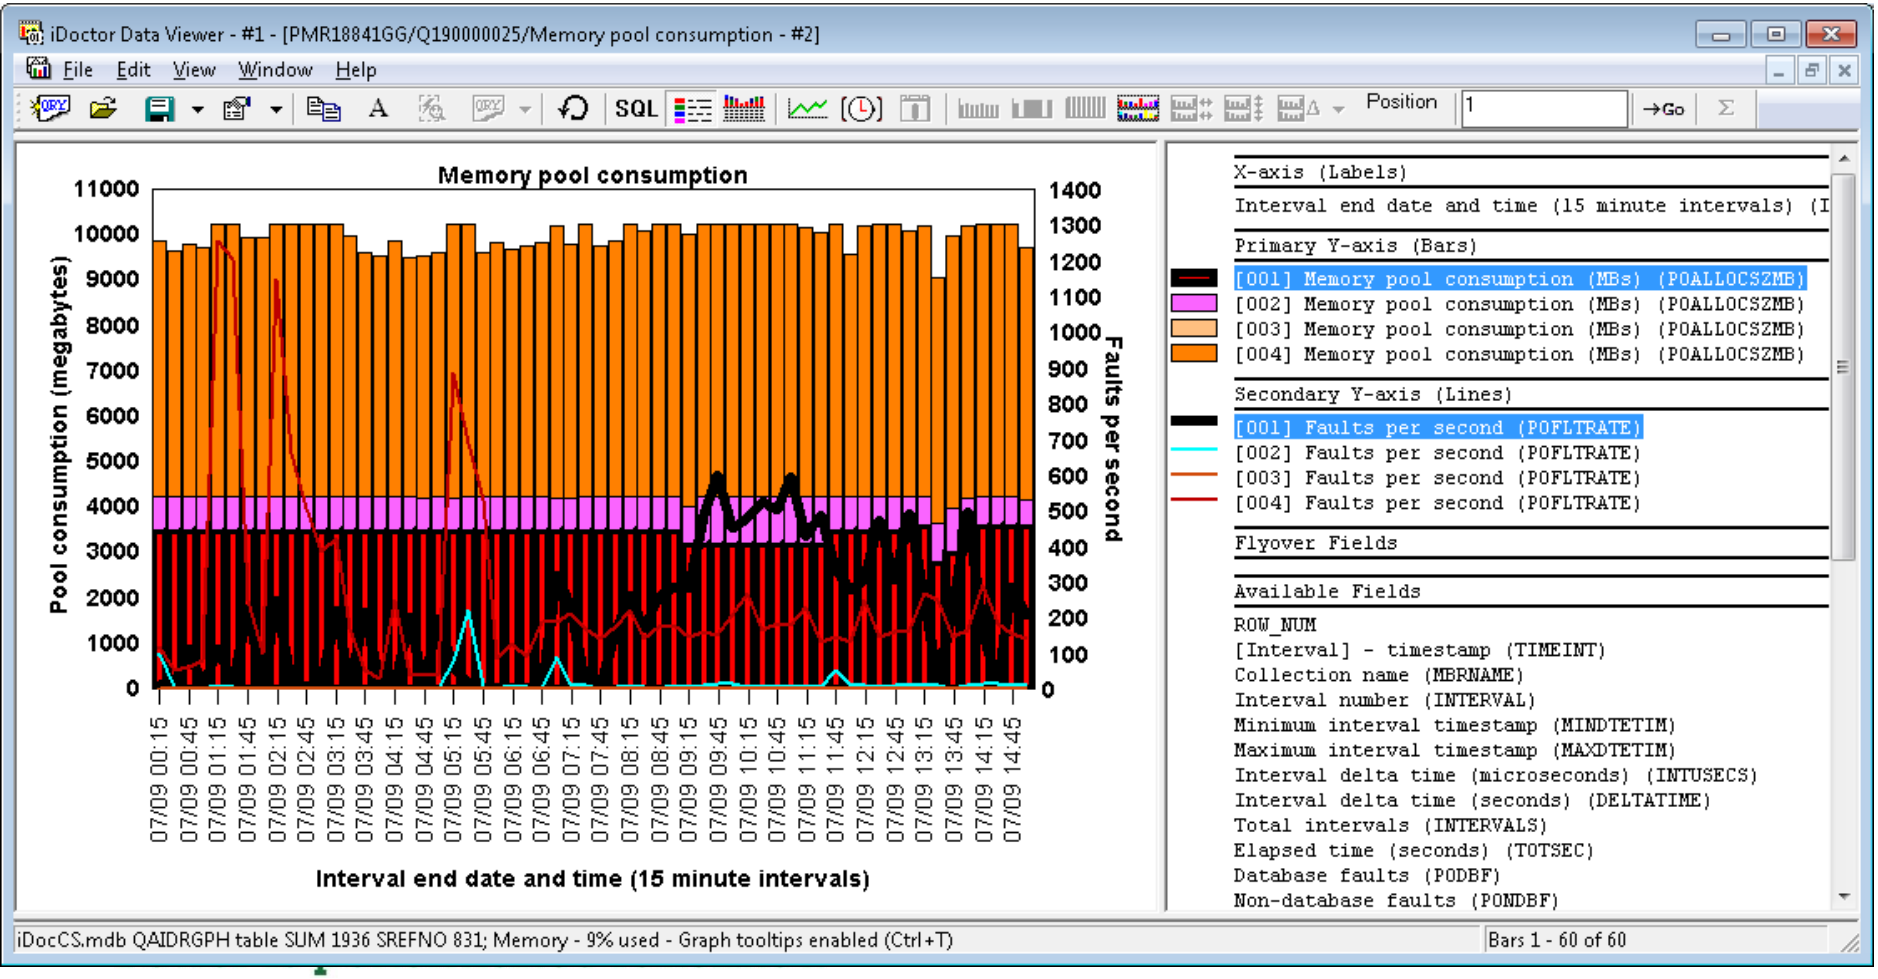

Collection Services Investigator – Memory pool graphs

- Added new graphs:
  - Flattened type - Memory pool unallocated space available

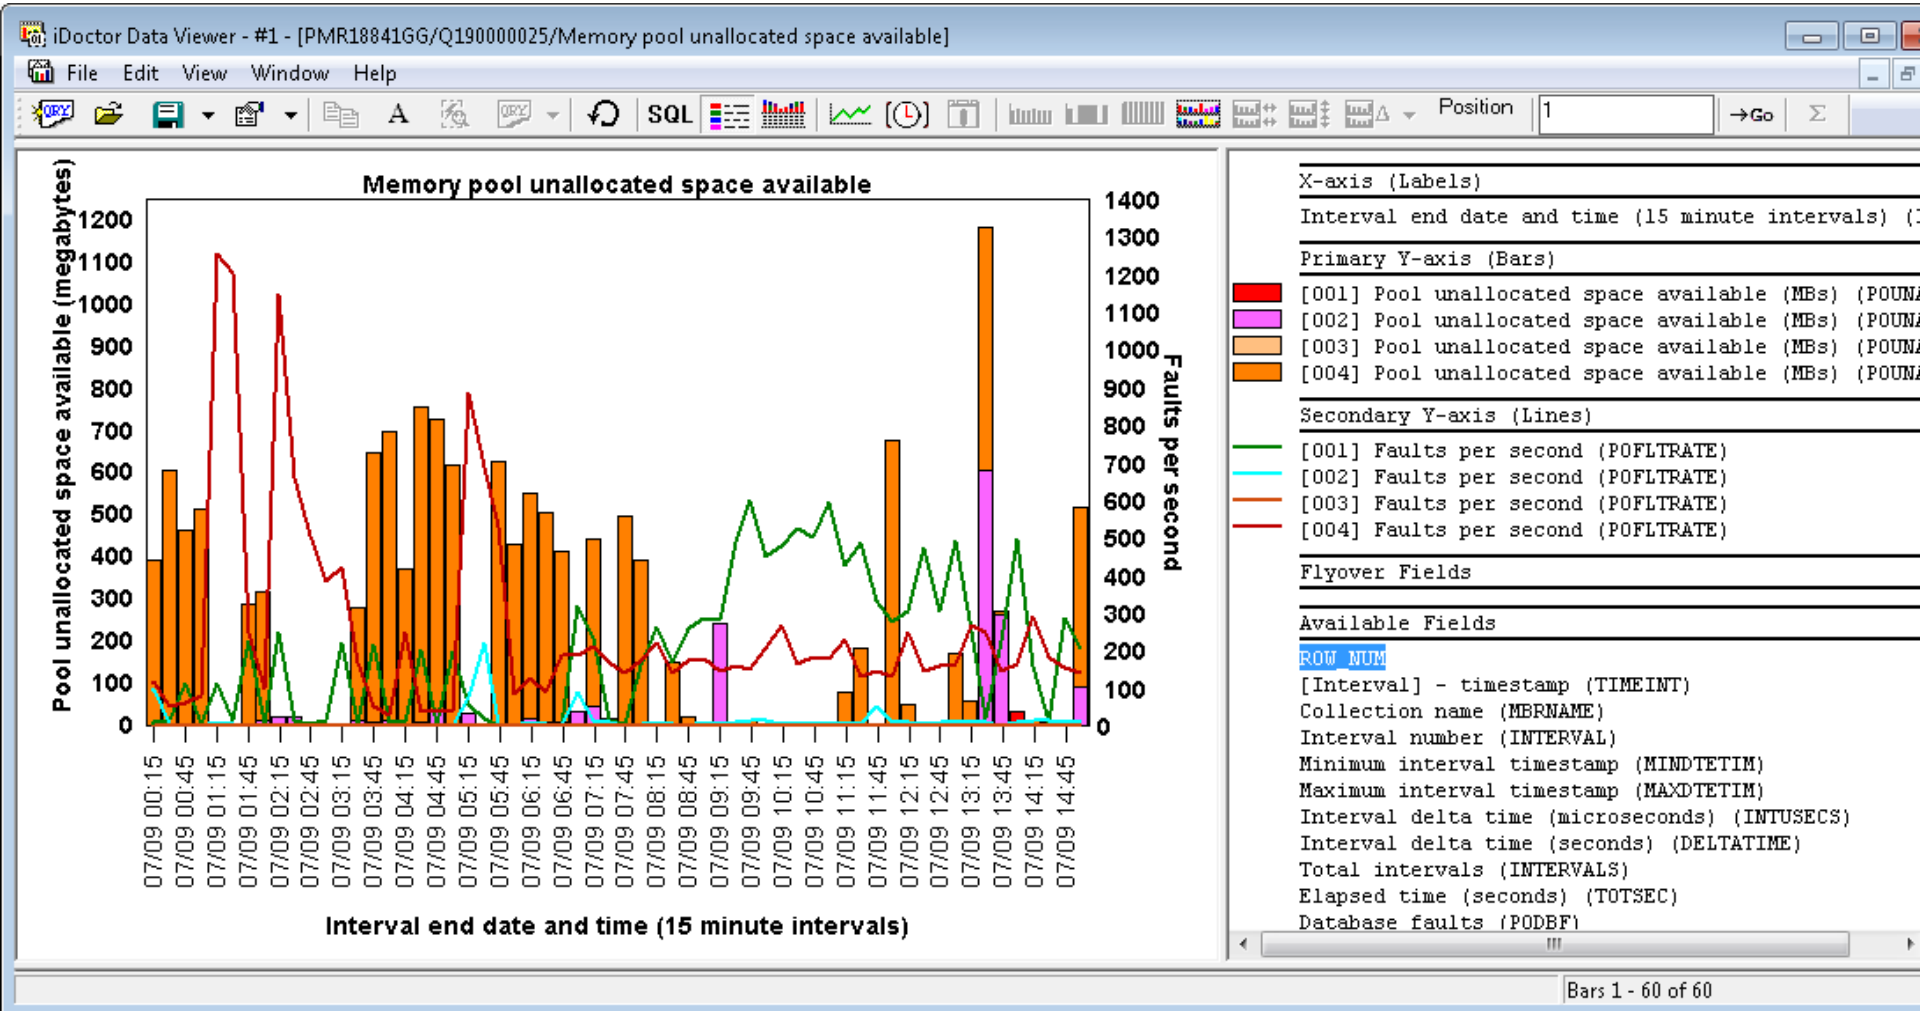

Collection Services Investigator – New job counts graphs

- Added a graph called Short-lived job counts that shows all jobs/tasks/threads that were created and destroyed within a single Collection Services interval over time.

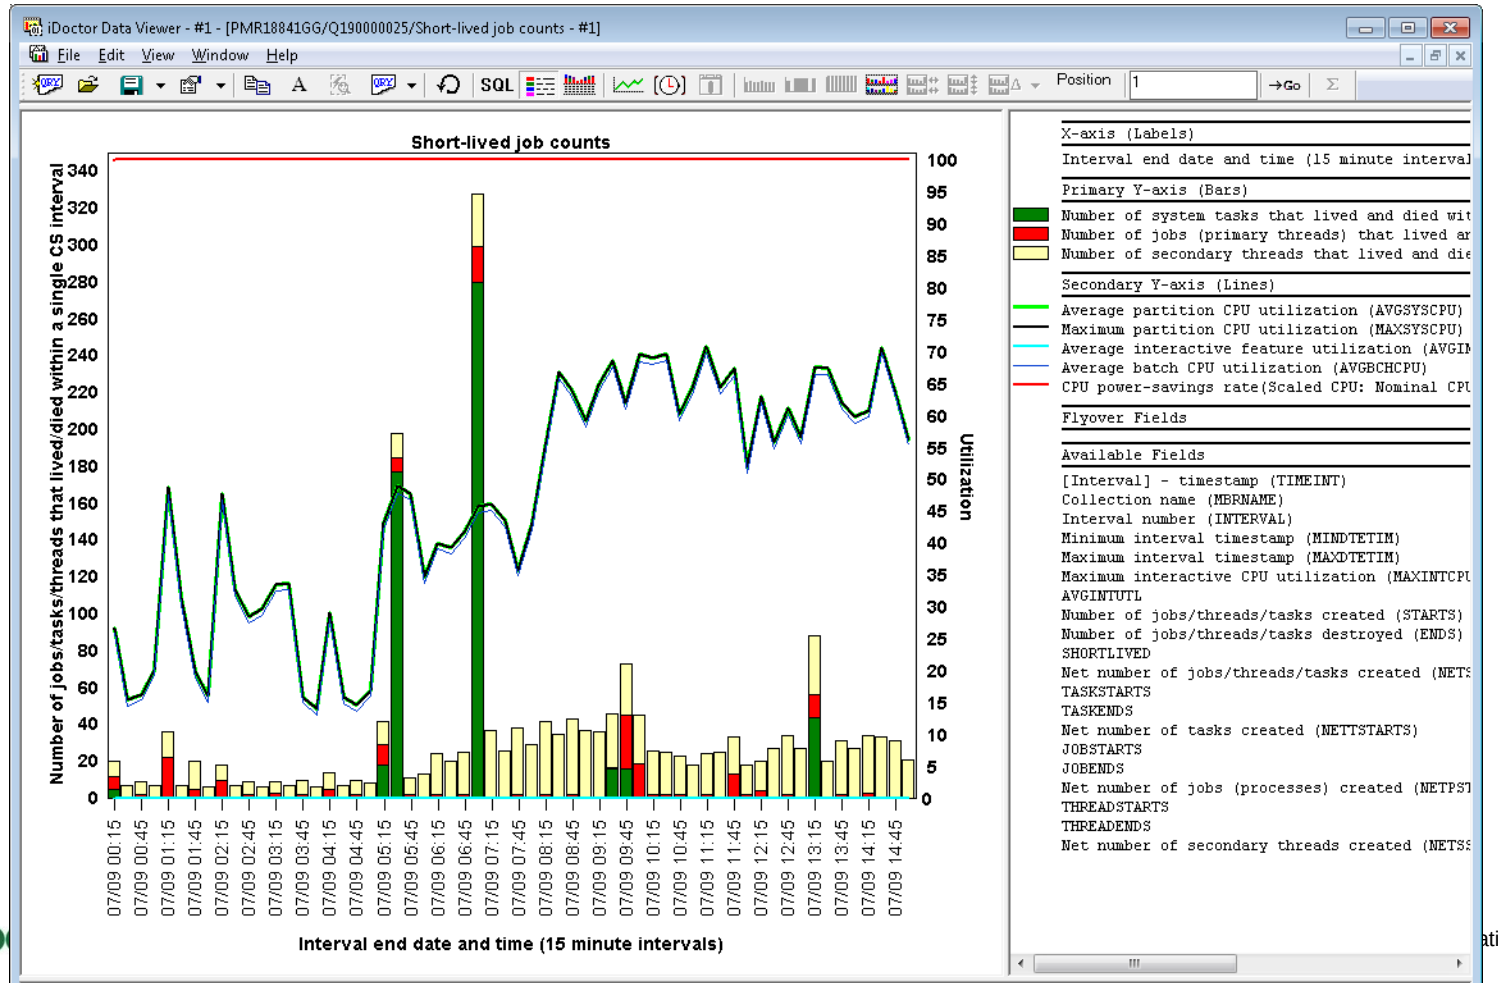

Collection Services Investigator – New job counts graphs

- Added a new folder called "Short-lived job counts rankings" (which also can be used as a drill down)
- At 7.1+, the job counts graph and jobs created/destroyed graph have been updated to include the new short-lifespan counts provided by Collection Services. The job counts rankings graphs have also been updated at 7.1 to include these counts.

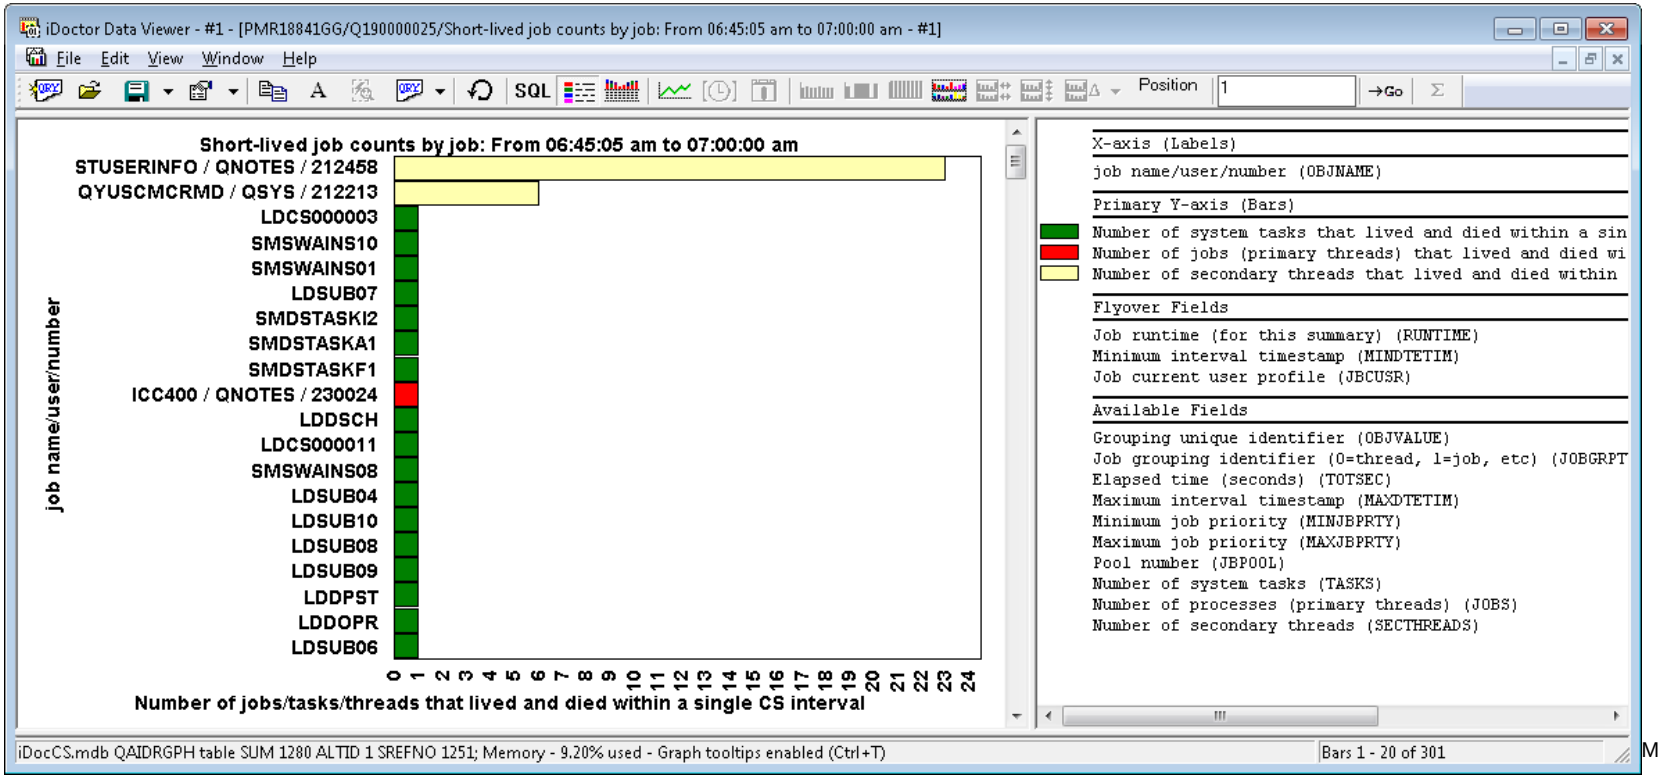

Collection Services Investigator – Wait bucket rankings can now become gantt charts

- Right click the legend – Set graph type – Gantt

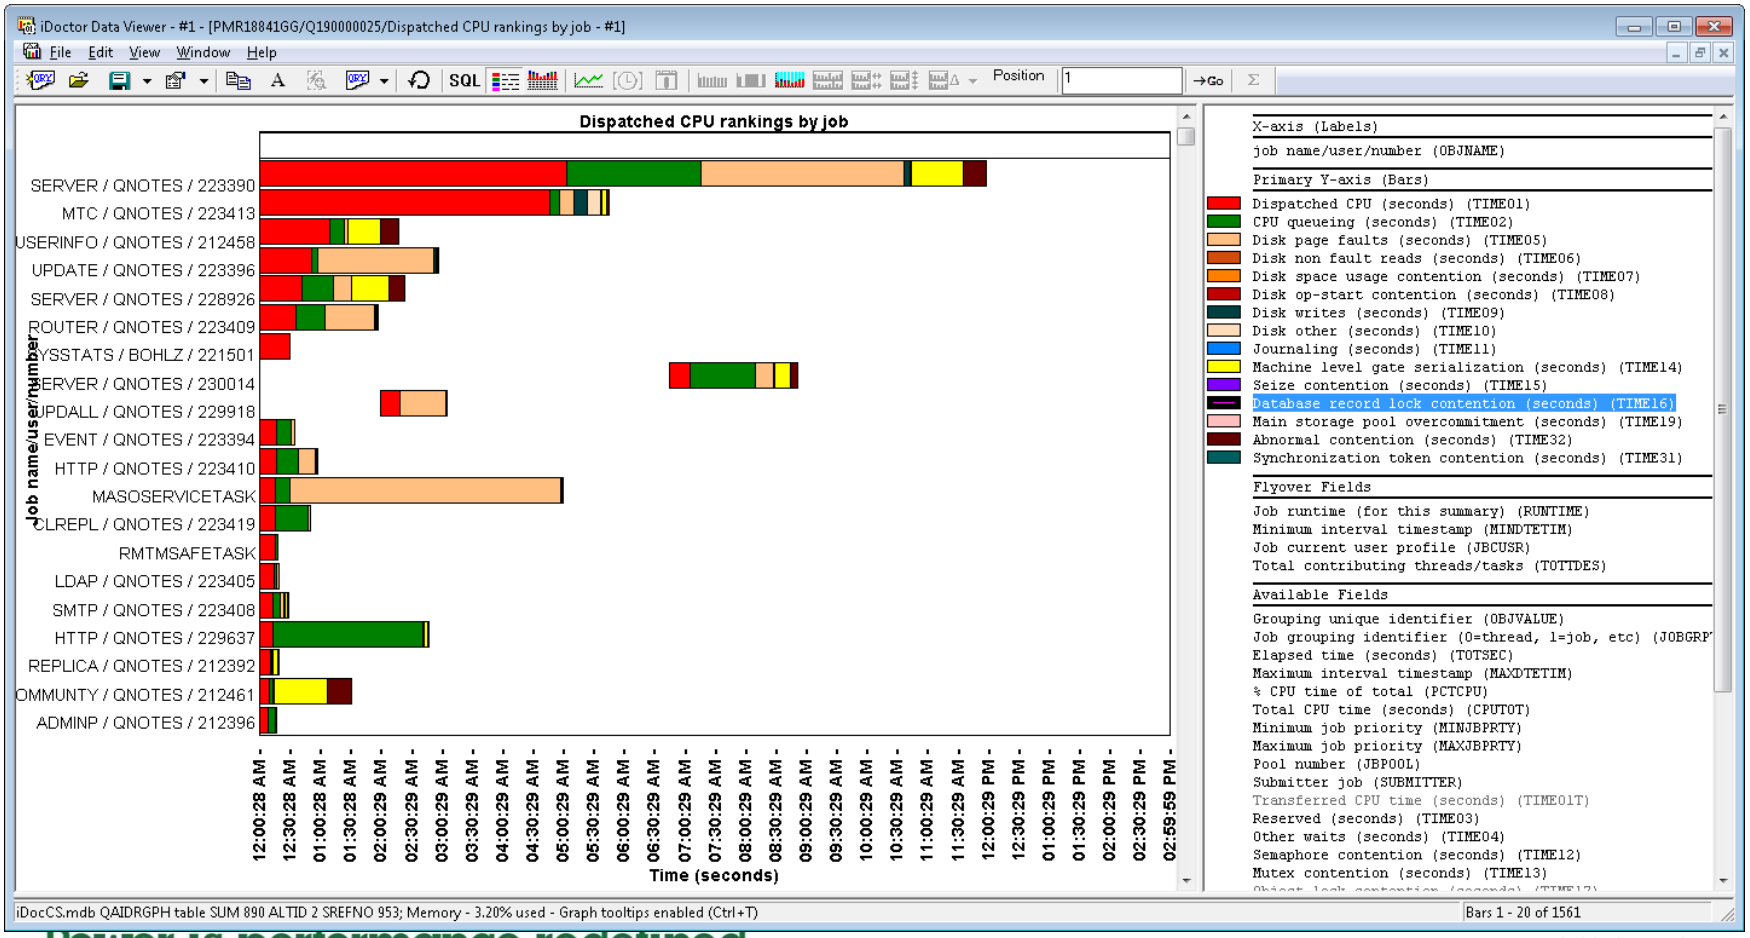

Collection Services Investigator – Workload capping graphs

- The following graphs appear ONLY if file QAPMSYSWLC exists in the library:
  - **Collection overview with workload capping time signature**
  - Workload capping delay
  - Workload capping processors assigned
  - Workload capping delay by [thread/generic job/etc]

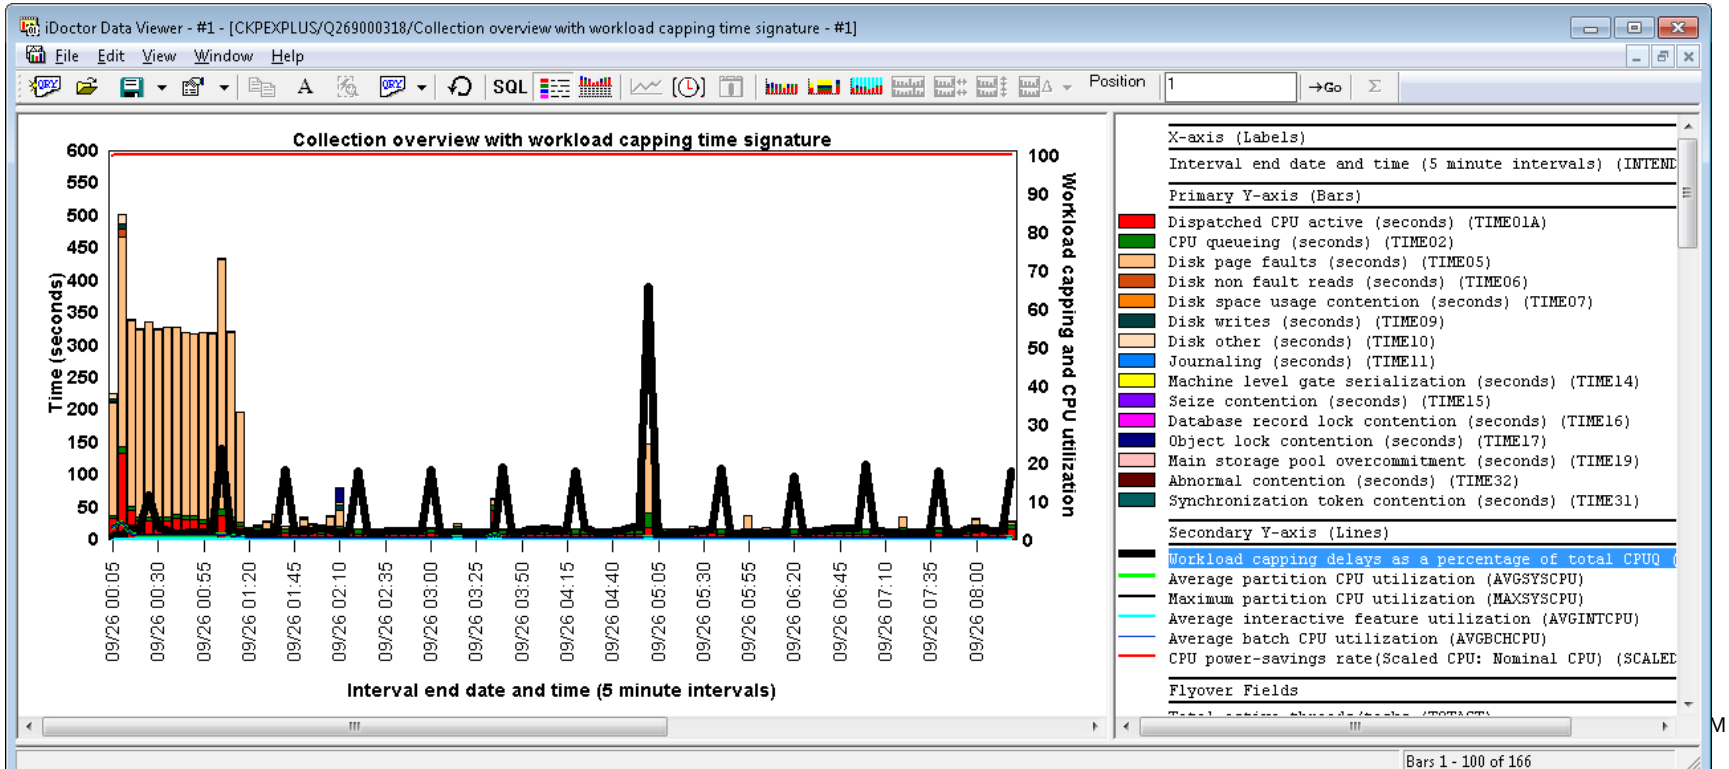

Collection Services Investigator – Workload capping graphs

- The following graphs appear ONLY if file QAPMSYSWLC exists in the library:
  - Collection overview with workload capping time signature
  - **Workload capping delay**
  - Workload capping processors assigned
  - Workload capping delay by [thread/generic job/etc]

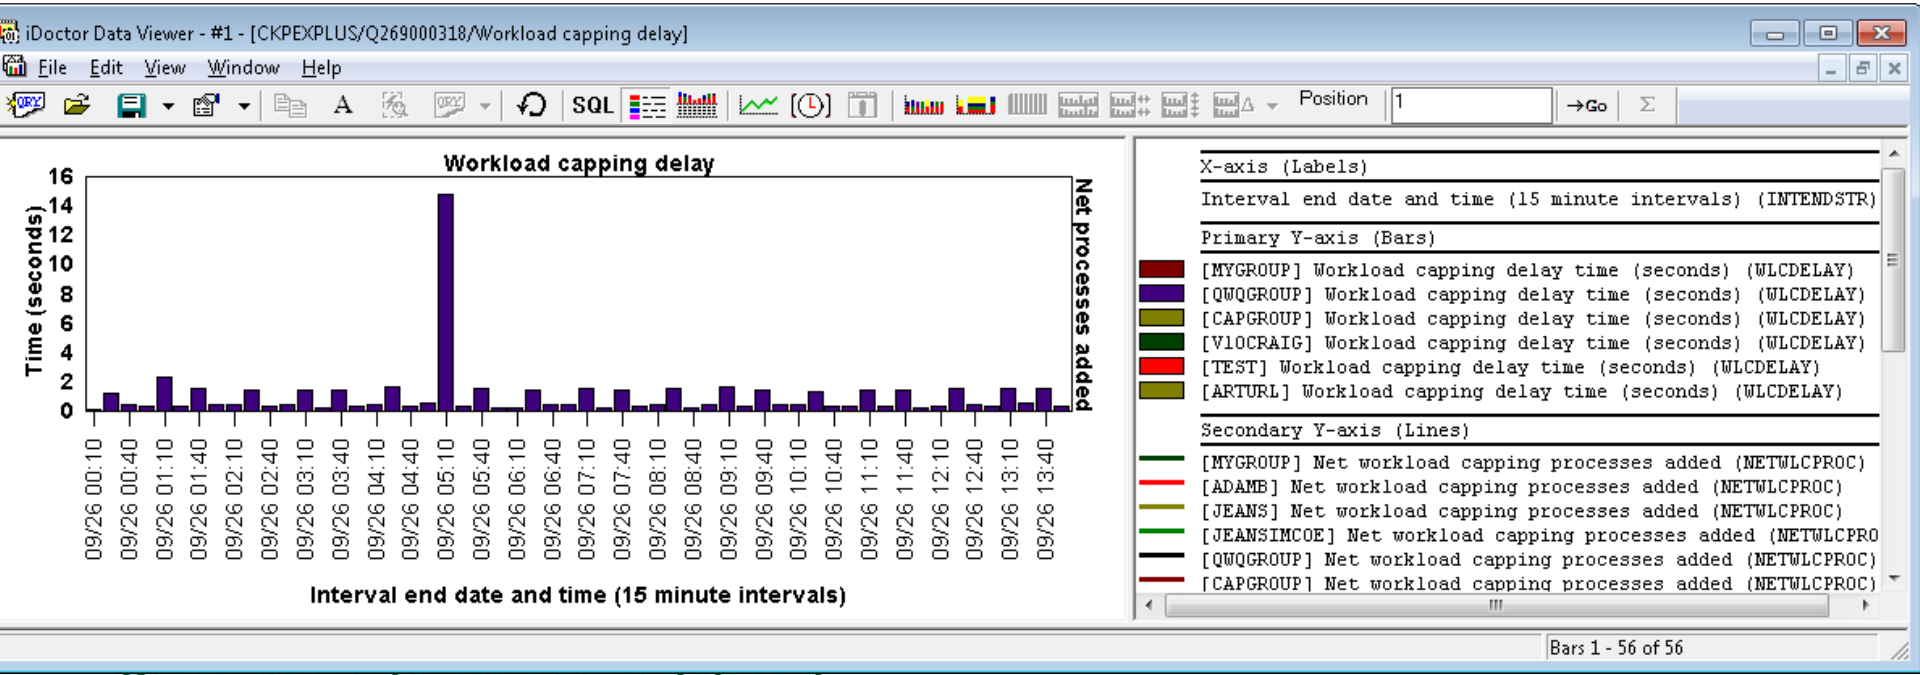

Collection Services Investigator – Workload capping graphs

- The following graphs appear ONLY if file QAPMSYSWLC exists in the library:
  - Workload capping processors assigned

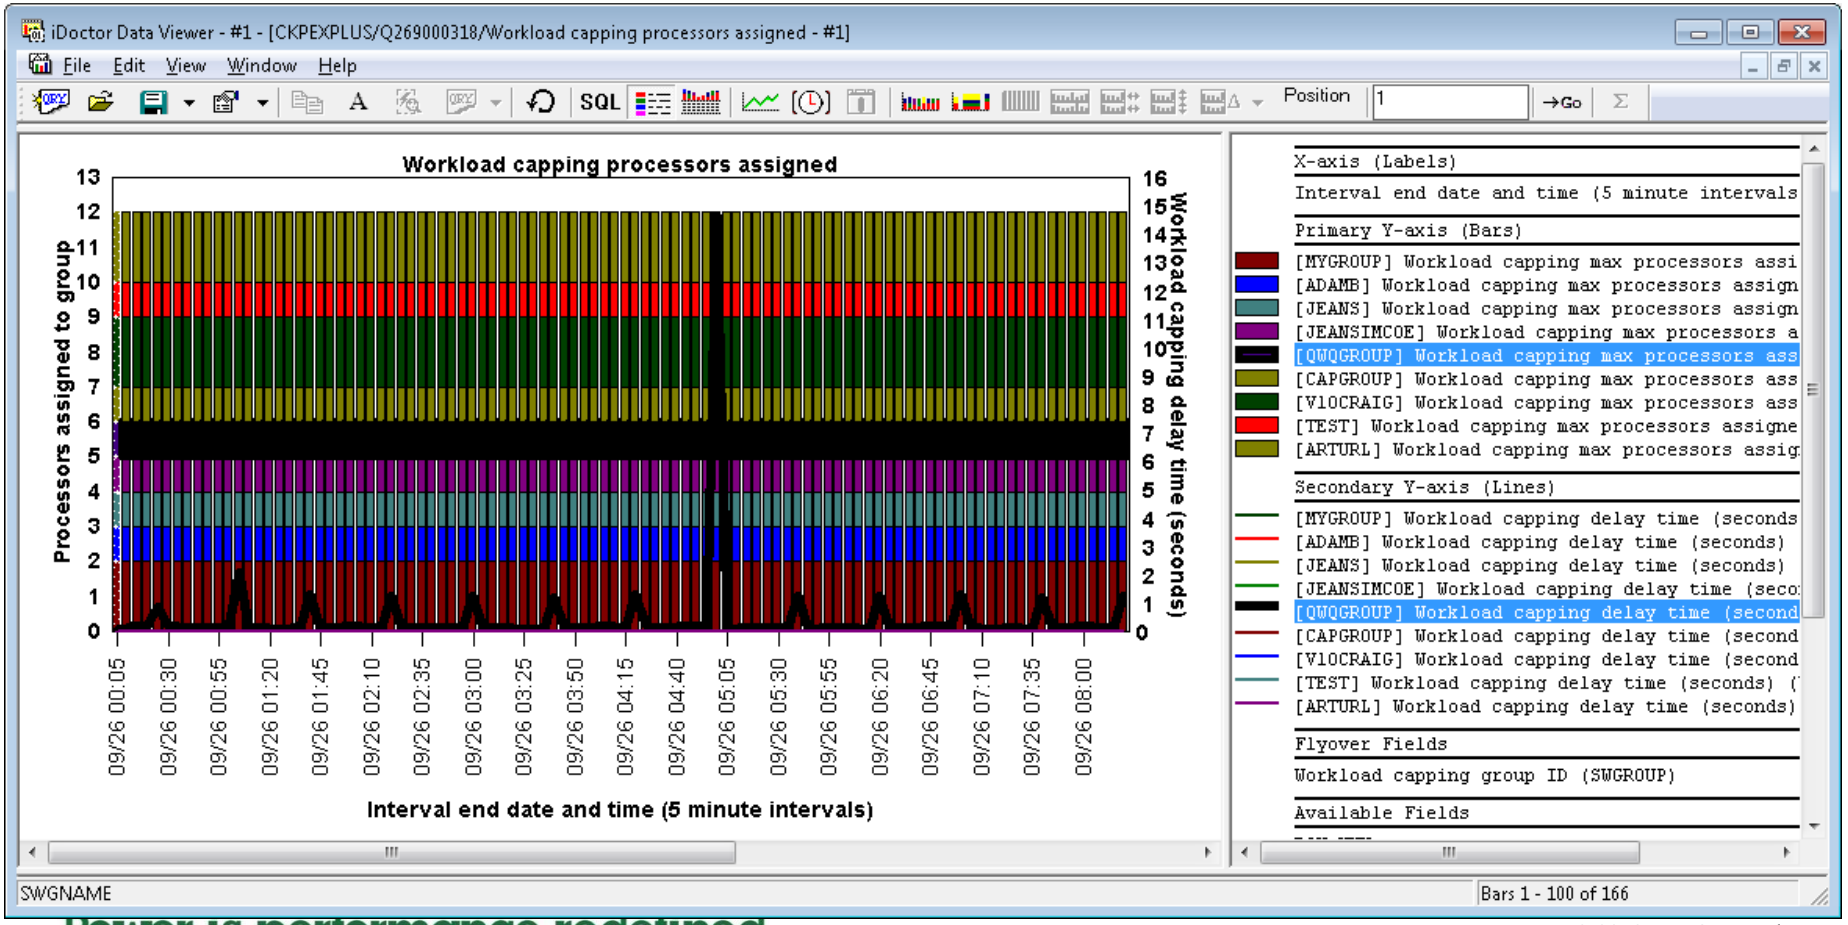

Collection Services Investigator – Collection Services menu

- Added options to view the Collection Services settings or change them, as well as cycle, end or start Collection Services. Right-click the Collection Services Investigator icon for these options.

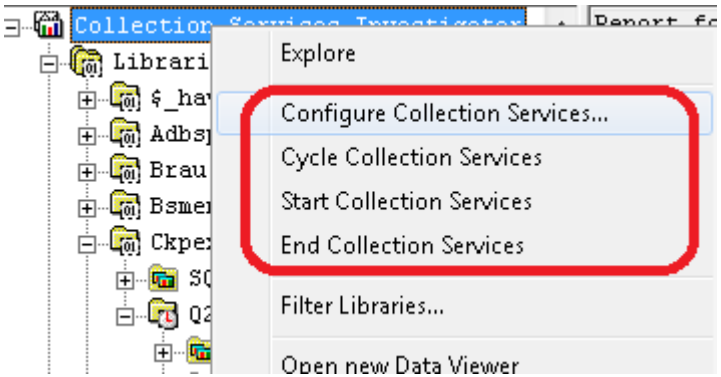

Collection Services Investigator – Configure Collection Services

- Added a window to configure collection services default settings.

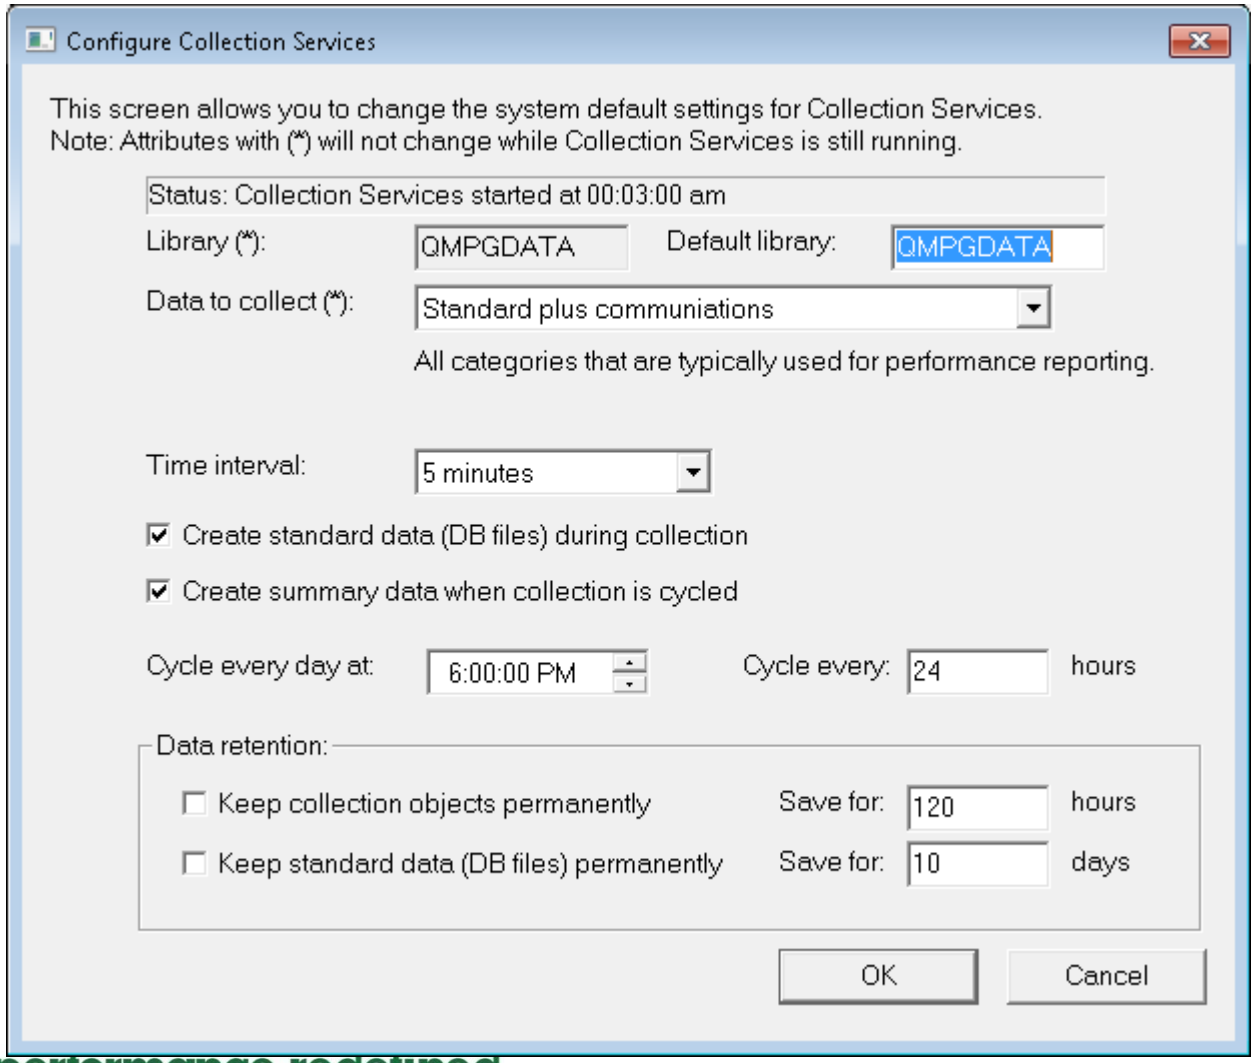

Collection Services Investigator – Disk configuration report updates

- Updated the report "Capacity (in GBs) by ASP with paths" so it now shows Min and max drive size as well as the resource (drive) status (operational/inoperative/etc)
- Added new reports:
  - Capacity (in GBs) by ASP/IOP with paths
  - Capacity (in GBs) by ASP/IOP/IOA with paths
  - Disk configuration (non-operational disks only)
- The disk configuration report also now shows the capacity for each drive as well as the disk resource status (operational, not detected, inoperative.)

Collection Services Investigator – IO and memory page demand graphs

- **Added new graph called Memory page demand**

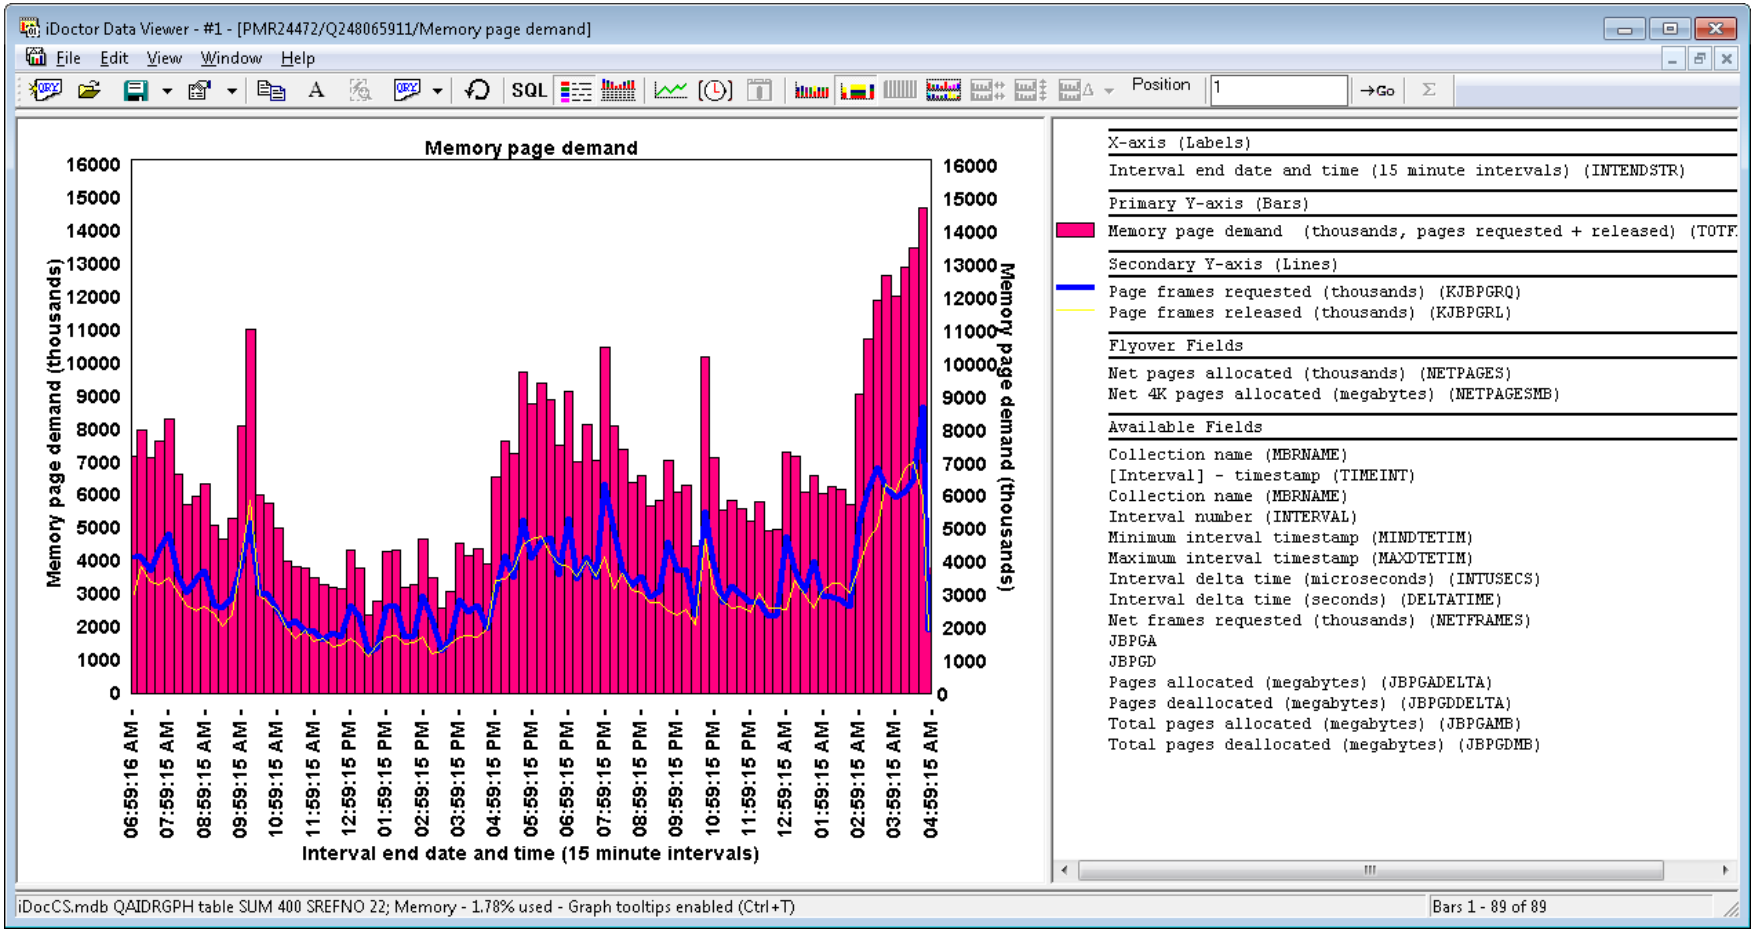

Collection Services Investigator – IO and memory page demand graphs

- **Rankings - Memory page demand**

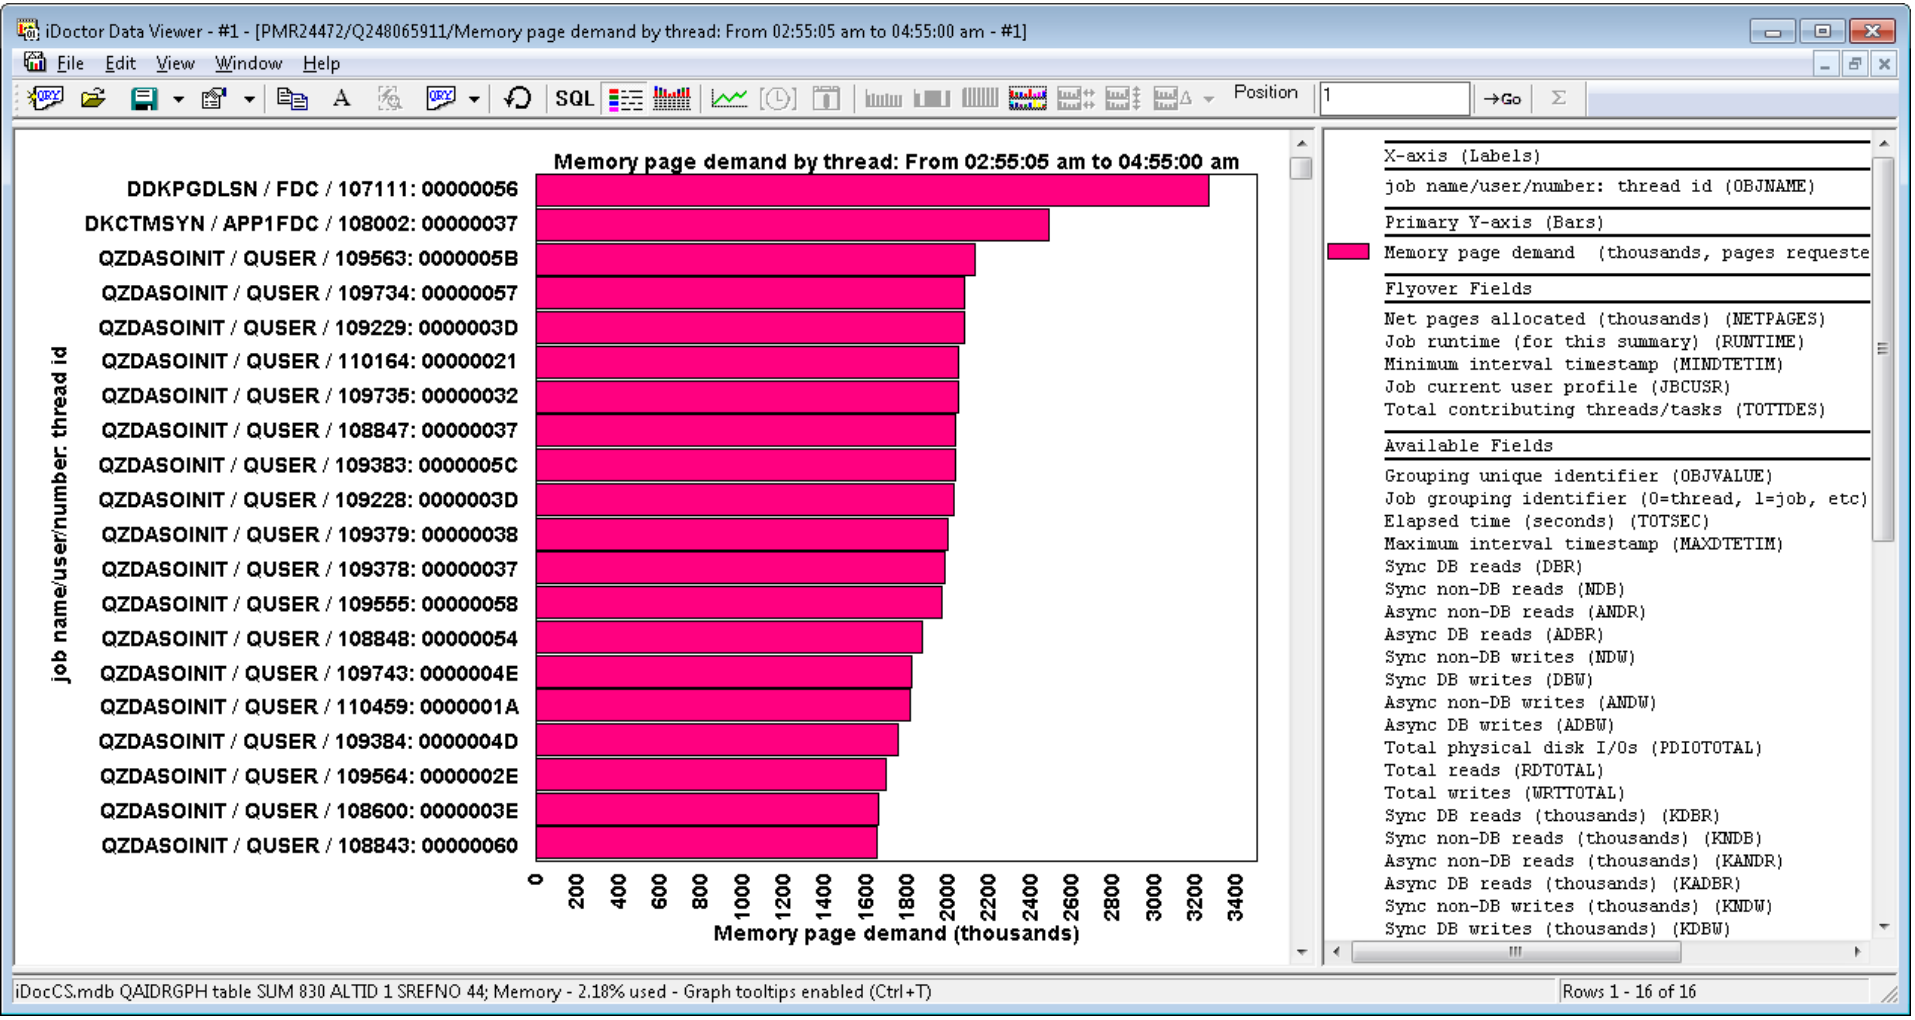

Collection Services Investigator – SQL graphs

- At 7.1, added SQL rankings and SQL graphs for a selected job grouping over time for the new SQL performance statistics. The data is provided with IBM i 7.1 TR5.

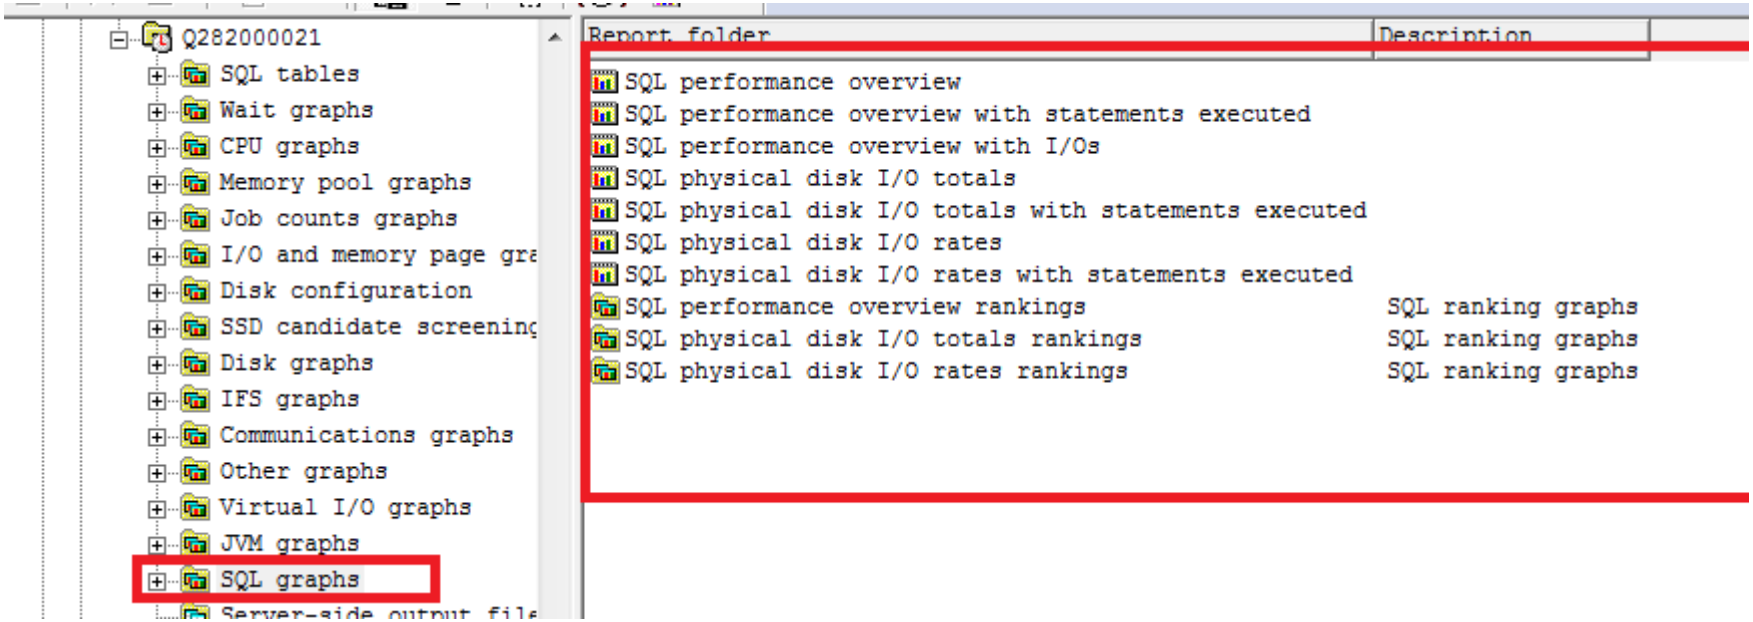

Collection Services Investigator – SQL graphs

- SQL performance overview graph

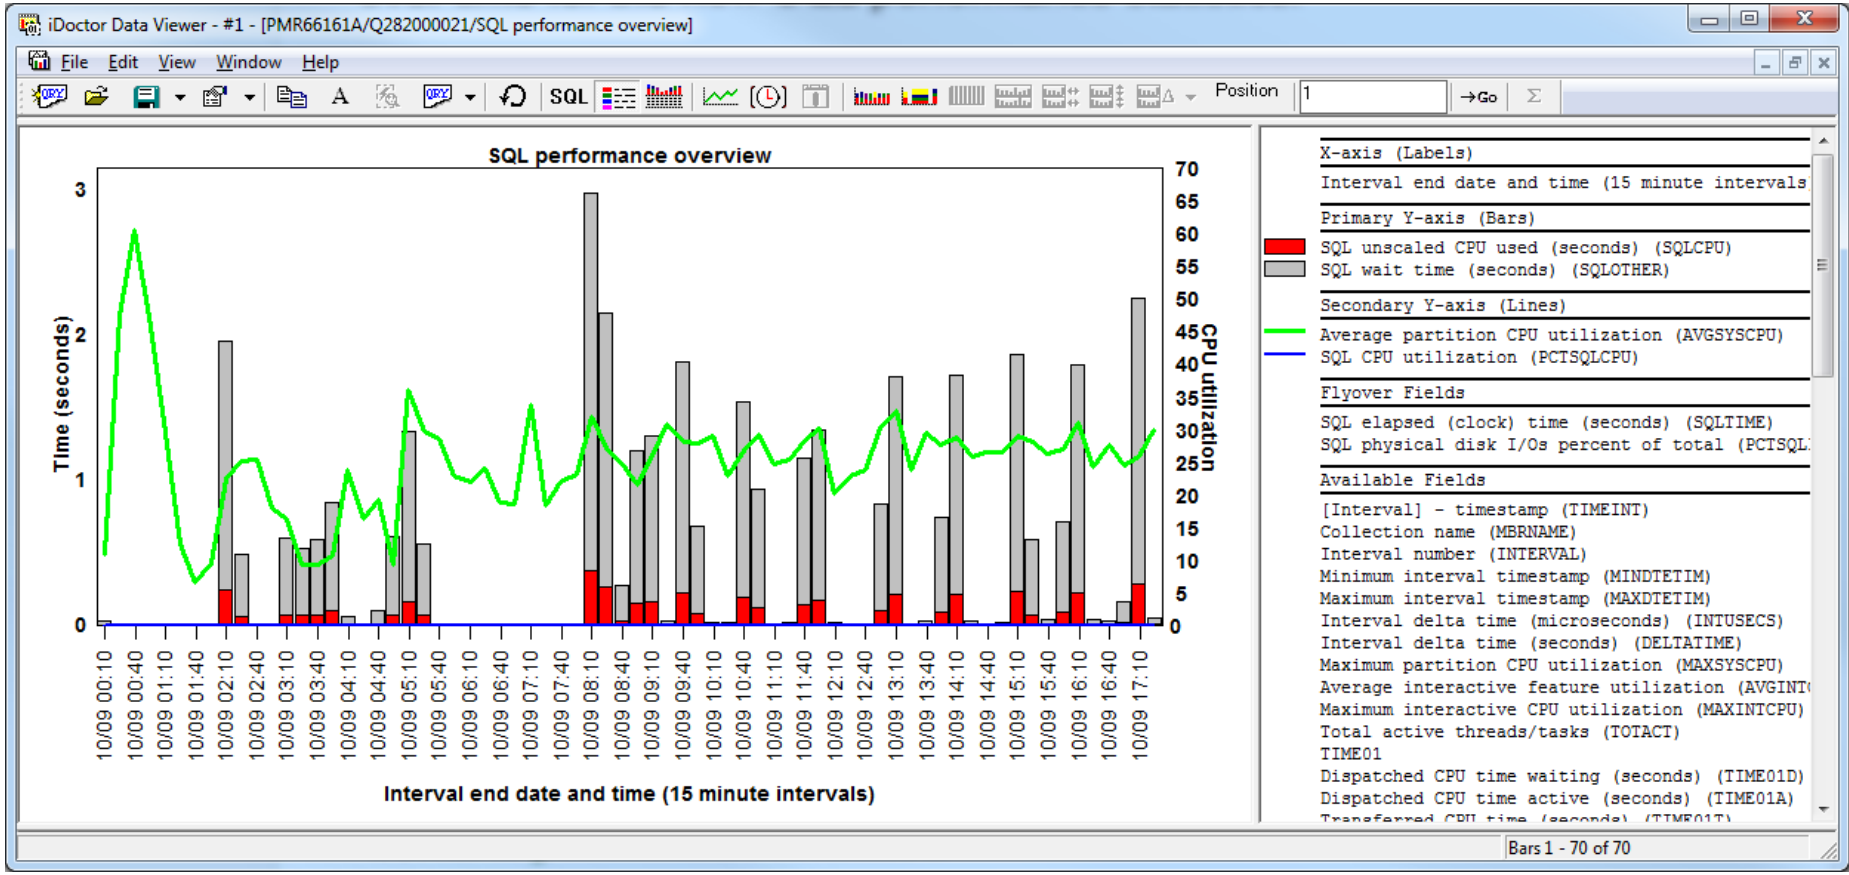

Collection Services Investigator – SQL graphs - Rankings

- SQL performance overview by current user example

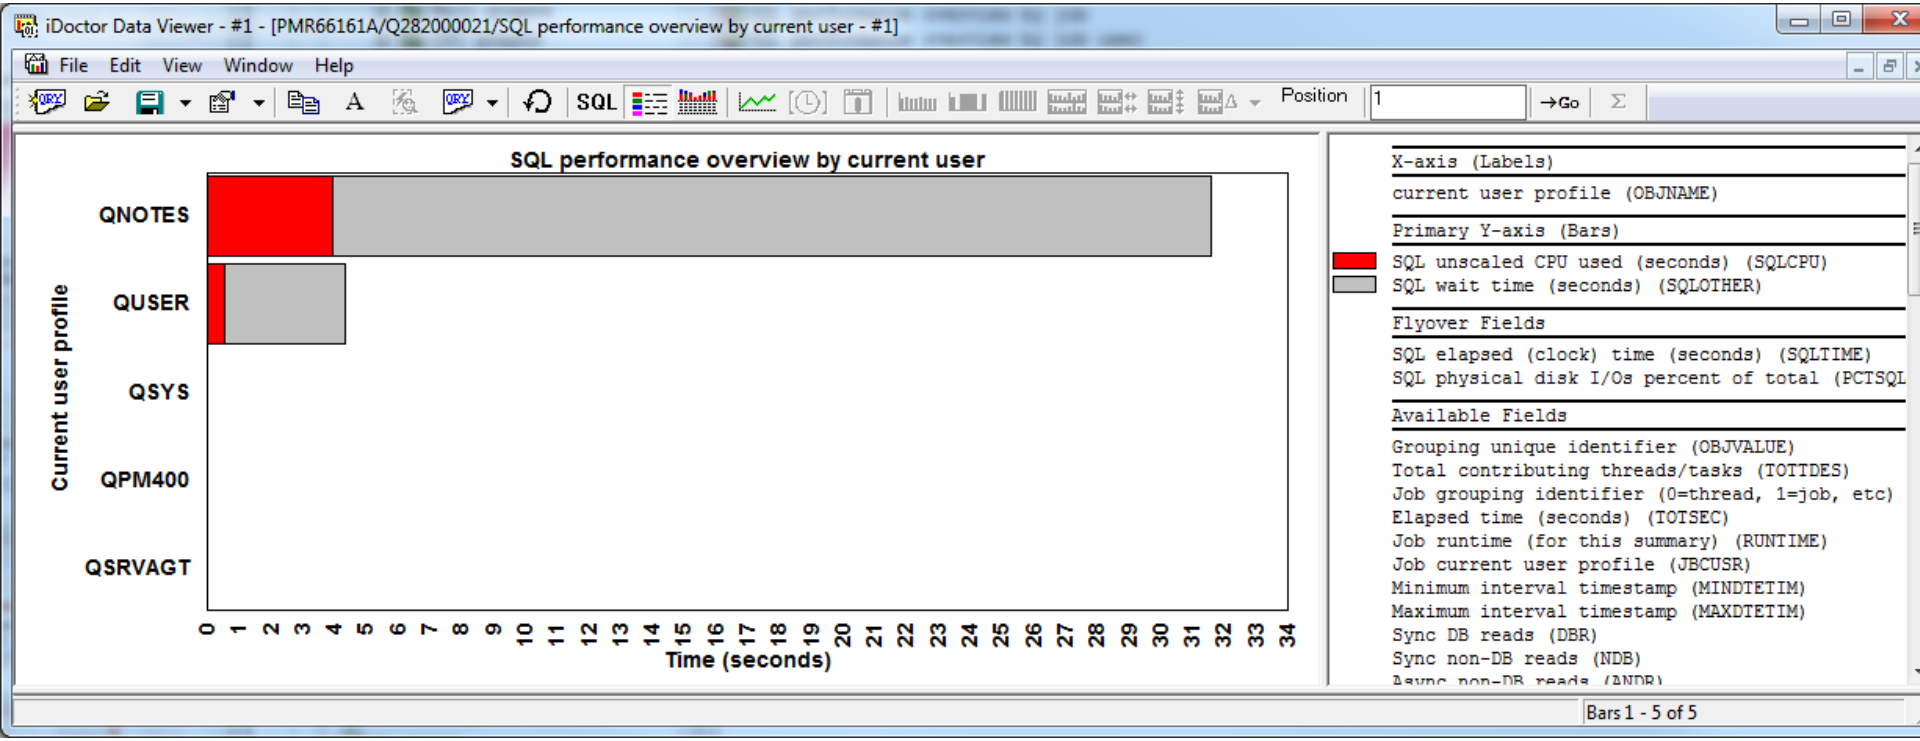

Collection Services Investigator – SQL graphs - Rankings

- SQL physical disk I/O totals by thread

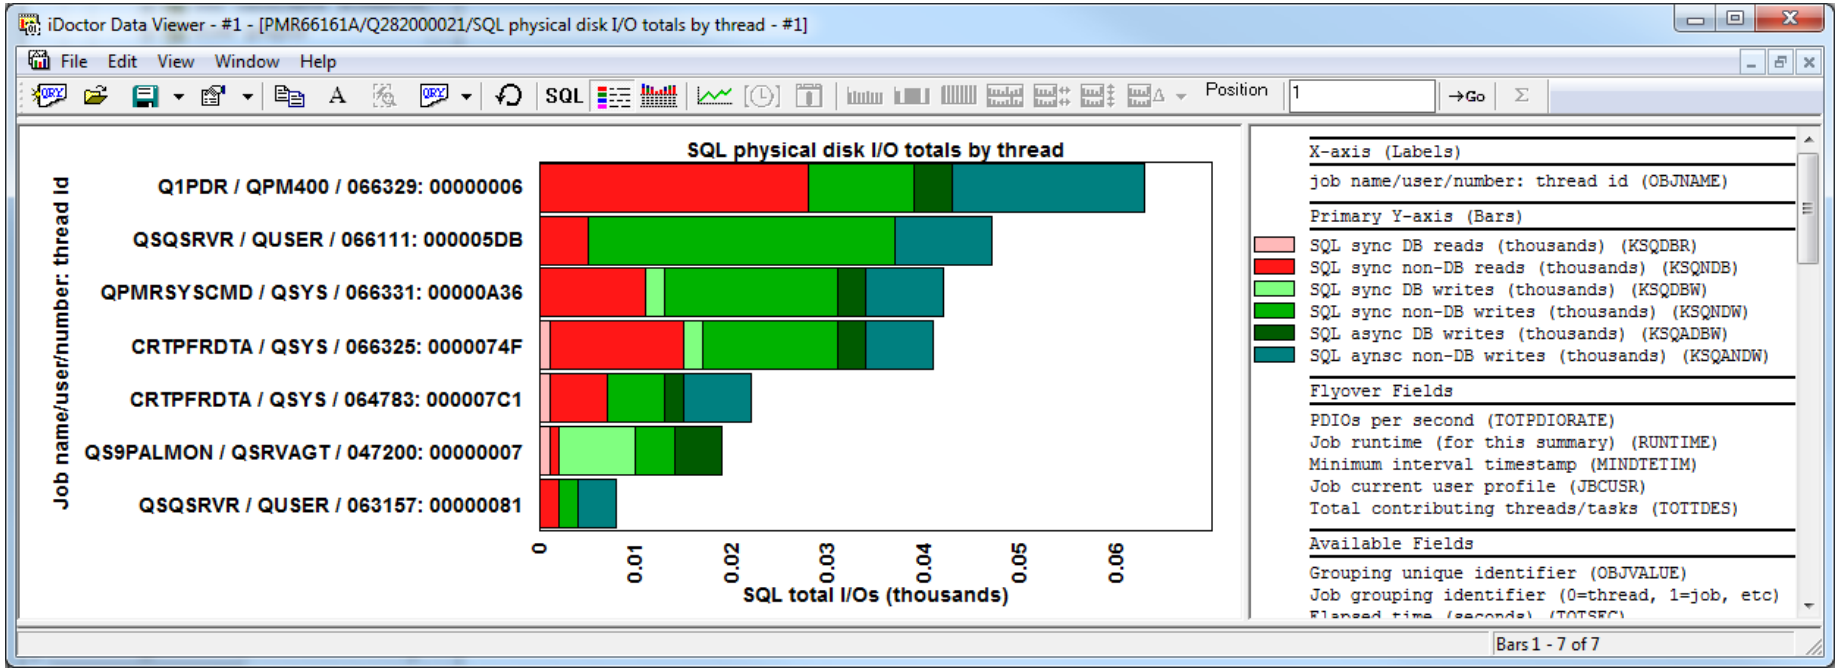

Collection Services Investigator – CS objects folder

- **You can now filter the folder by generic library name or owner.**

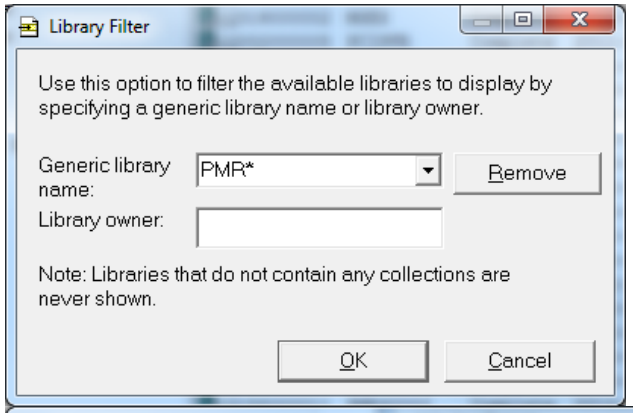

IBM iDoctor for IBM i C01039 [C:\PROGRAM FILES (X86)\IBM\IDOCTOR\IDOCTOR.EXE 10/19/2013 09:31:38] CA 710-SI47412 - [Lpdac710: Collection Services Investiga]

File Edit View Window Help

- Collection Services Investigator
  - Libraries
  - Historical summaries
  - CS objects: (PMR*)
  - SQL tables
  - Monitors
  - Browse collections
  - Super collections
  - Saved collections
  - Work management
  - ASPs
  - Disk units
  - Objects owned by MCCARGAR

Collection name	Collection library	Status	Start time	Last update time	Collection interval	Size (MB)	Retention period	Owner	Partition collected on	Partitic collecte on VRM
Q253000104	PMR02524C	Complete	2013-09-10-00.01.04	2013-09-10-13.21.11	5 minutes	291.4	Permanent	QCOLSRV	C104C43P	V7R1M0
Q013000102	PMR19388AA	Complete	2013-01-13-00.01.02	2013-01-14-11.10.29	15 minutes	177.2	24 hours	QCOLSRV	ECHO	V7R1M0
Q248065911	PMR24472	Complete	2013-09-05-06.59.12	2013-09-06-05.00.09	5 minutes	1,304>	Permanent	QDFTOWN	A2029P21	V6R1M0
Q233000004	PMR46297	Complete	2013-08-21-00.00.05	2013-08-22-00.00.04	5 minutes	286.9	432 hours	QCOLSRV	AS01	V7R1M0
Q234005027	PMR46346	Complete	2012-08-21-00.50.27	2012-08-21-23.14.44	15 minutes	1,930>	720 hours	QSECOFR	SYS38A	V6R1M0
Q236012630	PMR46346	Complete	2012-08-23-01.26.30	2012-08-23-23.14.25	15 minutes	2,033>	720 hours	QSECOFR	SYS38A	V6R1M0
Q363000108	PMR48973P7	Complete	2012-12-28-00.01.08	2012-12-29-00.01.08	15 minutes	156.9	720 hours	QCOLSRV	PROTEC7	V7R1M0
Q270000104	PMR55232	Complete	2013-09-27-00.01.04	2013-09-28-00.01.04	15 minutes	8,263>	72 hours	QCOLSRV	PC04030	V7R1M0
Q309010017	PMR59215AA	Complete	2012-11-04-01.00.17	2012-11-04-23.01.04	15 minutes	73.6	120 hours	QCOLSRV	OCC01XX4	V7R1M0
Q155000027	PMR92002	Complete	2012-06-03-00.00.27	2012-06-03-17.12.21	5 minutes	1,127>	192 hours	QCOLSRV	VF09	V7R1M0
Q155175312	PMR92002	Complete	2012-06-03-17.53.12	2012-06-04-00.00.26	5 minutes	398.8	192 hours	QCOLSRV	VF09	V7R1M0
Q156000027	PMR92002	Complete	2012-06-04-00.00.27	2012-06-05-00.00.23	5 minutes	2,092>	192 hours	QCOLSRV	VF09	V7R1M0
Q157000024	PMR92002	Complete	2012-06-05-00.00.25	2012-06-06-00.00.10	5 minutes	2,130>	192 hours	QCOLSRV	VF09	V7R1M0
Q158000011	PMR92002	Complete	2012-06-06-00.00.11	2012-06-07-00.01.04	5 minutes	2,144>	192 hours	QCOLSRV	VF09	V7R1M0
Q159000108	PMR92002	Complete	2012-06-07-00.01.08	2012-06-07-17.16.52	5 minutes	1,430>	192 hours	QCOLSRV	VF09	V7R1M0
Q159171719	PMR92002	Complete	2012-06-07-17.17.20	2012-06-08-00.00.15	5 minutes	575.3	192 hours	QCOLSRV	VF09	V7R1M0

1 - 35 of 35 objects

Collection Services Investigator – Other graphs

- **Added the following graphs to the Other graphs folder:**
  - **SQL cursor and native DB total opens**
  - SQL cursor and native DB opens rates
  - SQL cursor and native DB total opens rankings (all 8 types)
  - SQL cursor and native DB opens rates rankings (all 8 types)
  - **Note:** The following PTFs are needed in order for the data behind the graphs to be correct: At 5.4 PTF SI44181 At 6.4 PTF SI44182 At 7.1 PTF SI44183

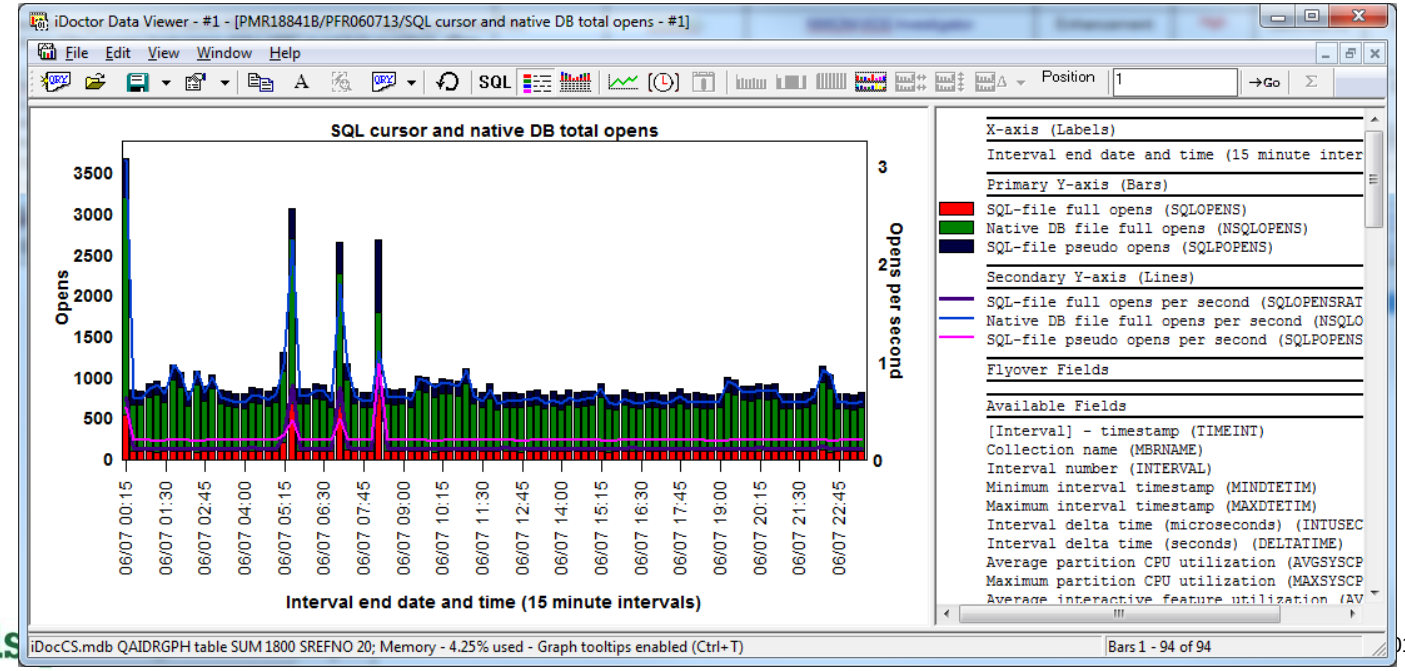

Collection Services Investigator – Other graphs

– **SQL cursor and native DB opens rates**

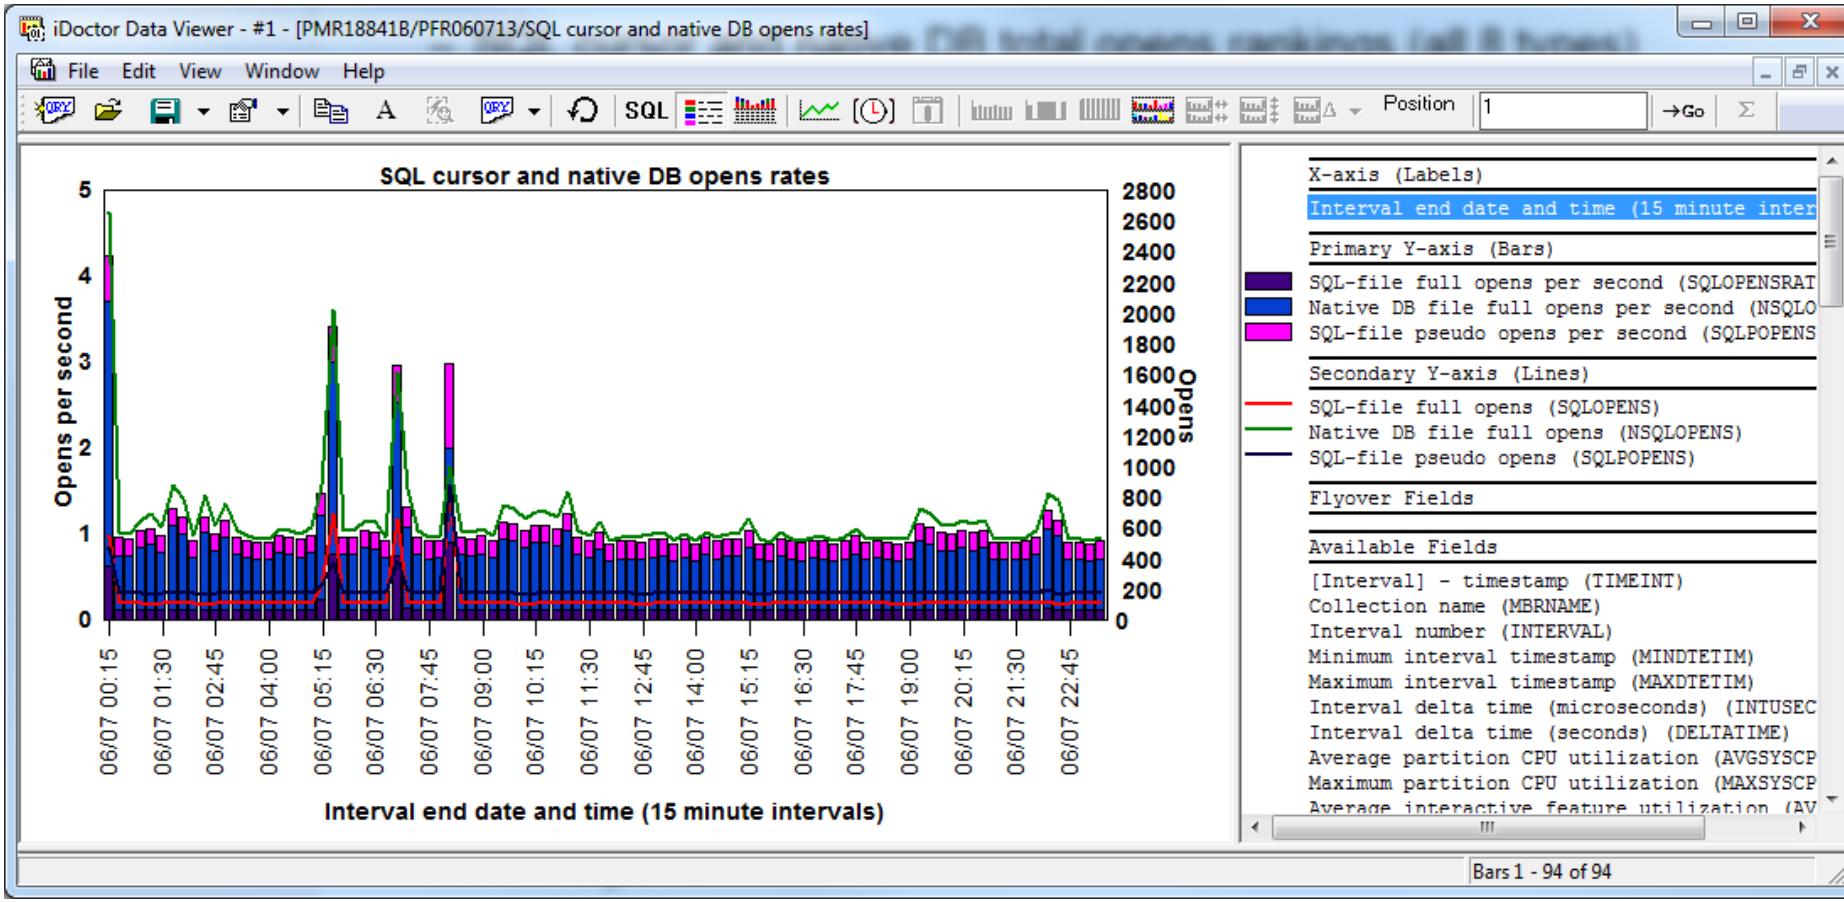

Collection Services Investigator – Other graphs

– SQL cursor and native DB total opens by thread example

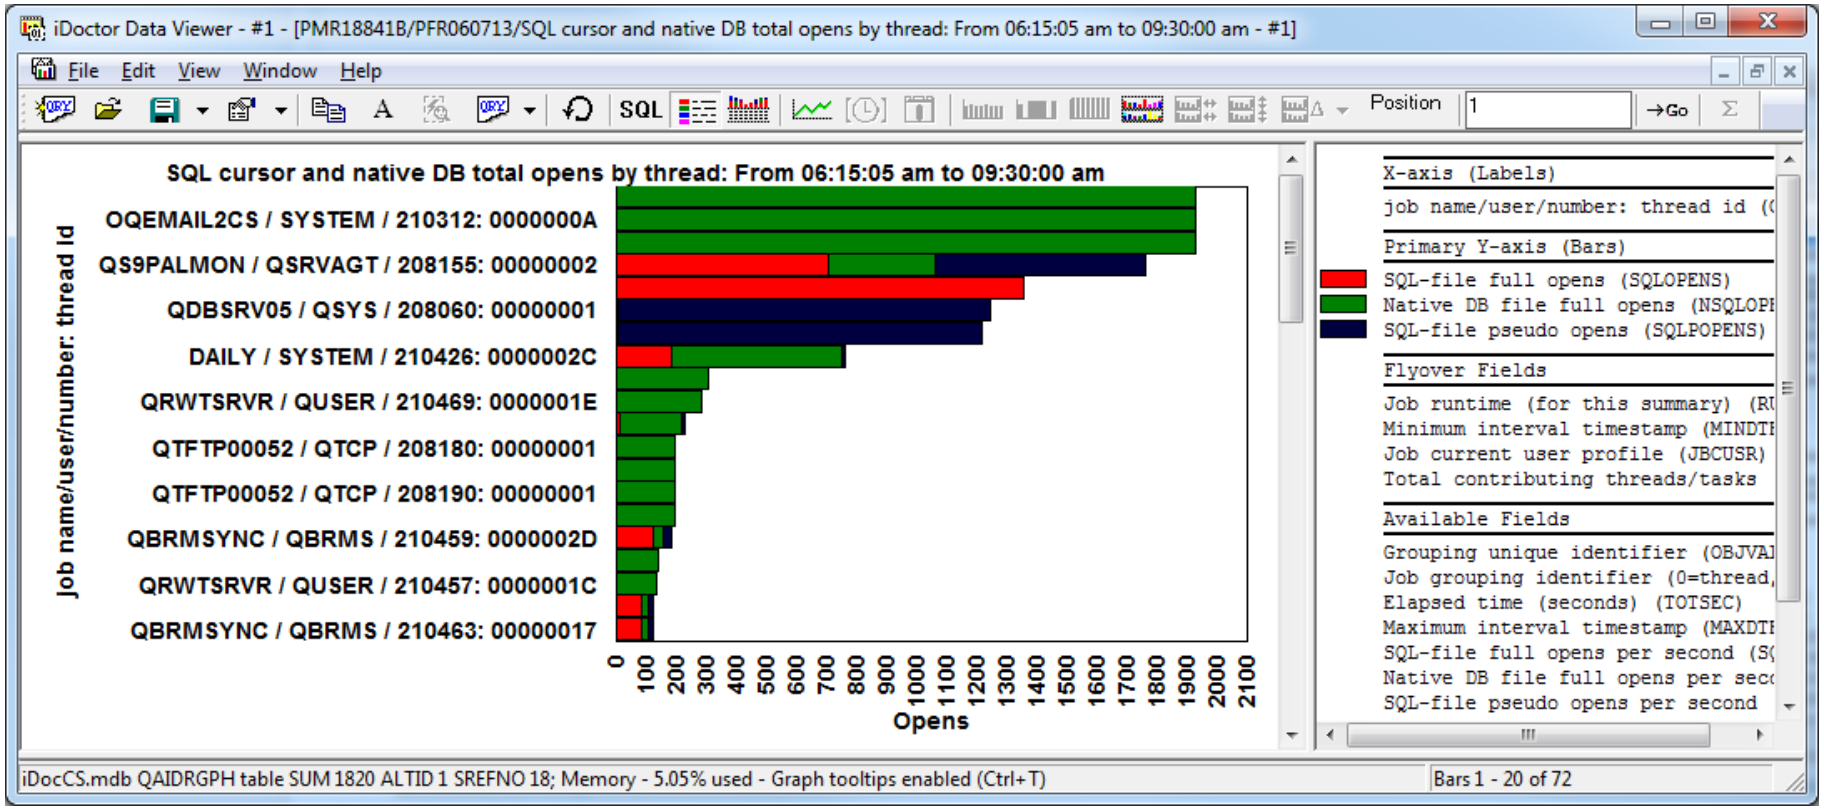

Collection Services Investigator – Other graphs

– SQL cursor and native DB opens rates by generic job

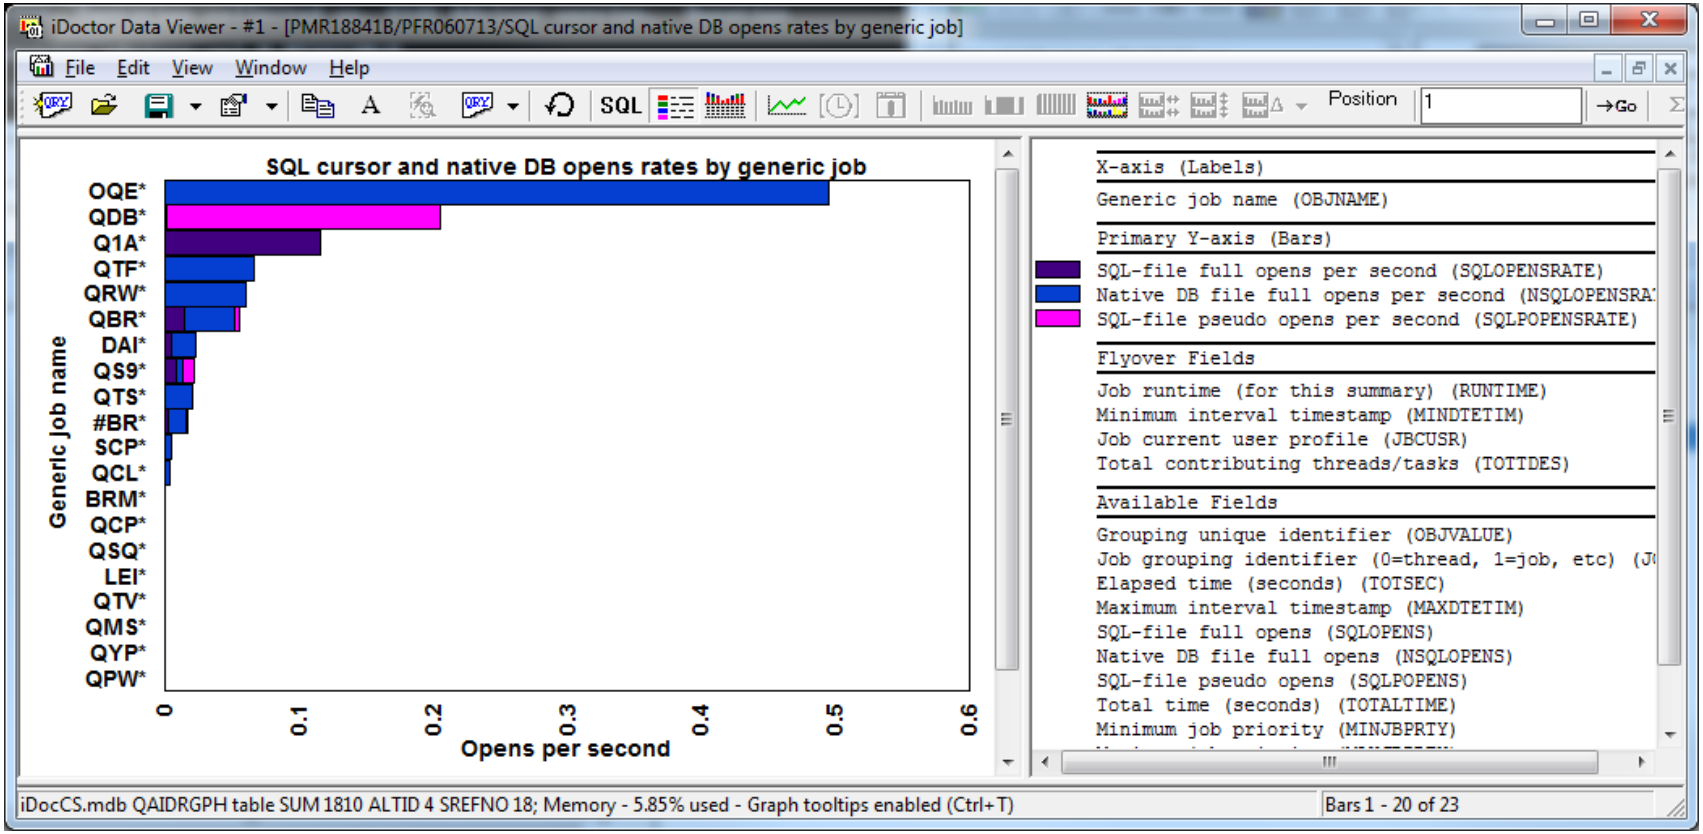

Collection Services Investigator – Memory pool graphs for > 1 TB

- **Added a set of memory pool graphs that can be used if the pool sizes exceed 1 TB.**
  - At 6.1 PTF SI47480 is required.

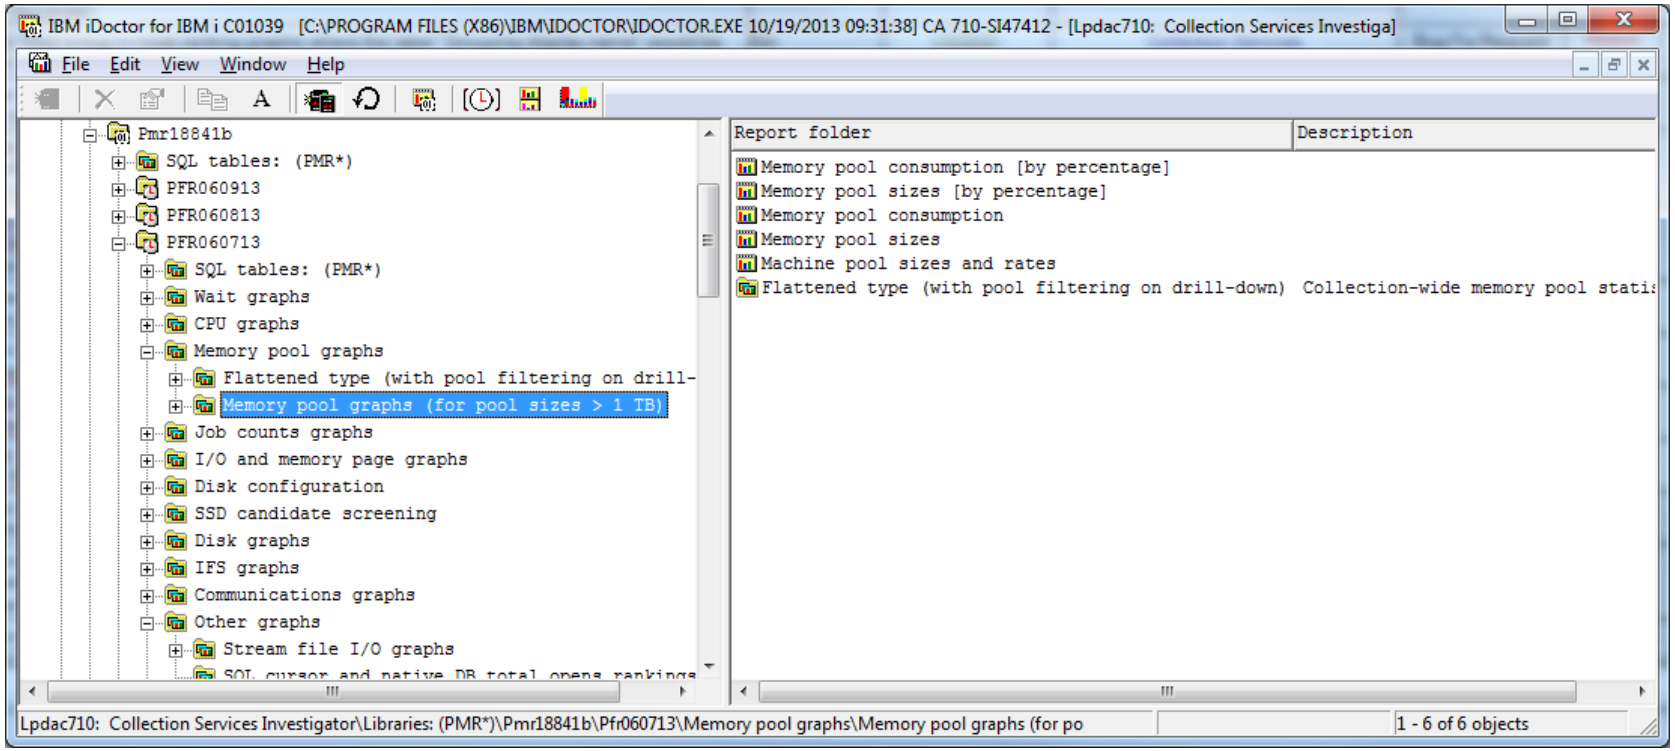

PEX Analyzer – Classic analysis programs/commands removed

- We removed the classic analysis commands and programs at 7.1.
- They were the commands named G*.
- They no longer work at 7.1 because of changes to the PEX database files.

PEX Analyzer – Cycles per instruction reports added

- Added reports under PEX file(s) starting points - Job statistics:
  - Cycles per instruction
  - Cycles per instruction rankings

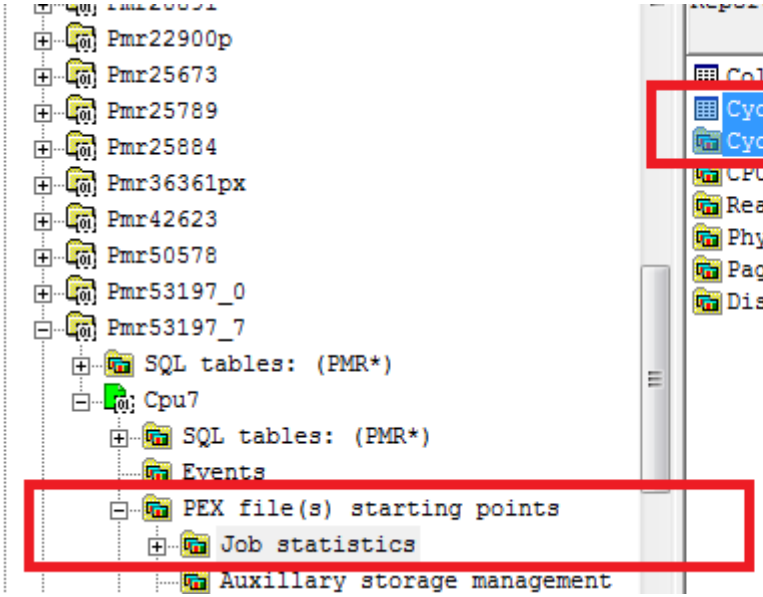

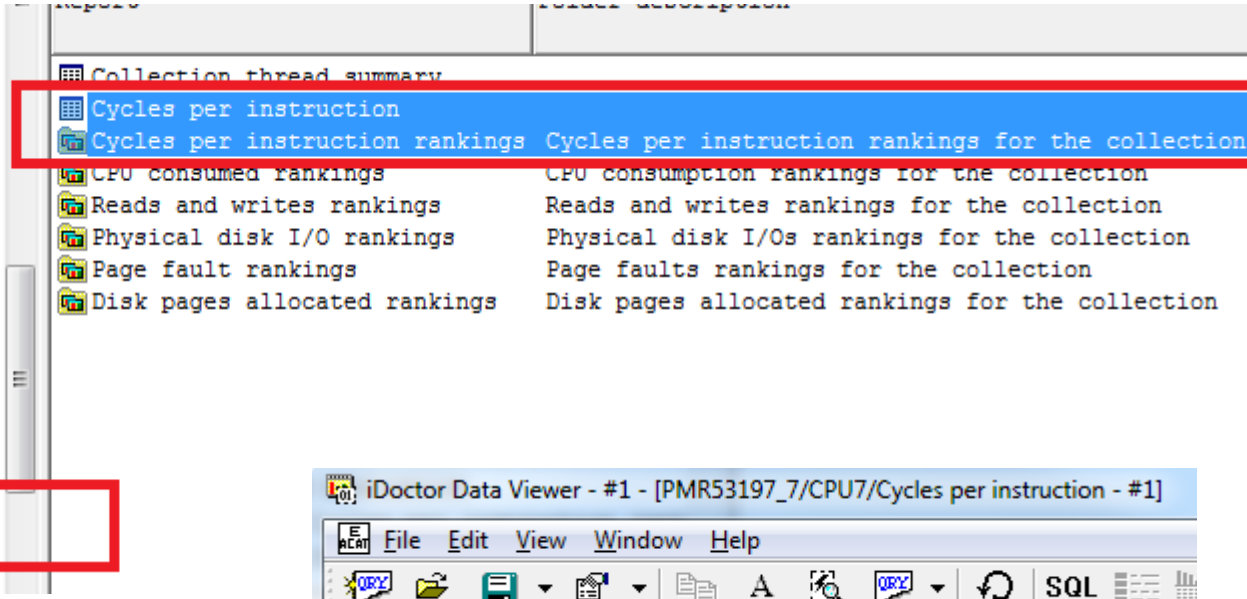

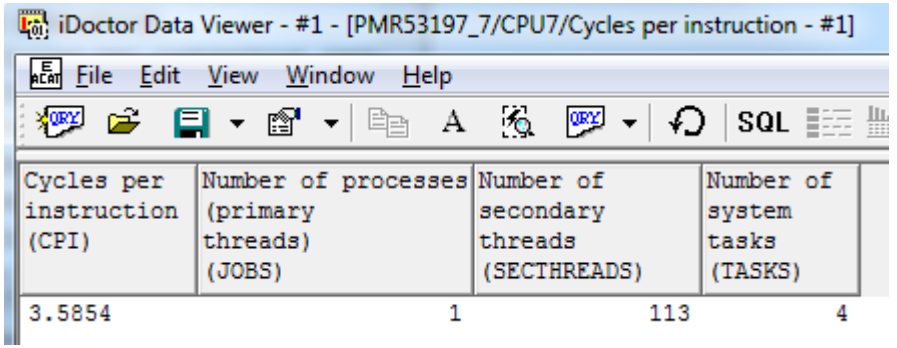

PEX Analyzer – Cycles per instruction reports added

- Cycles per instruction by thread ordered by total CPU cycles

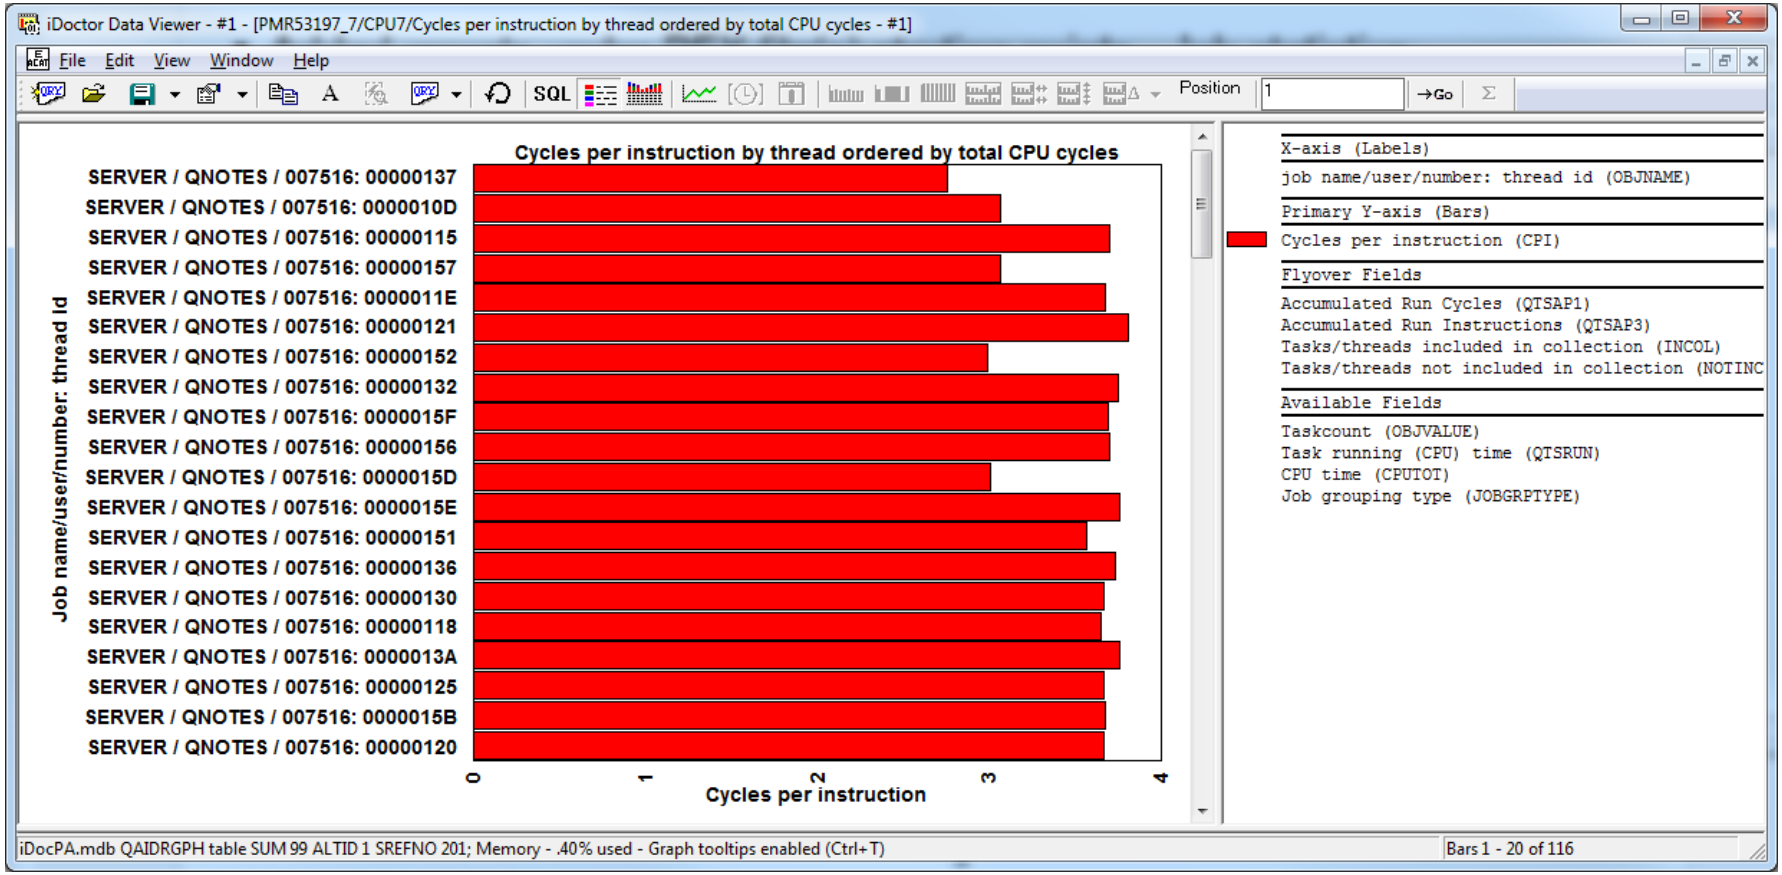

PEX Analyzer – Stats hier for N levels reports

- Added a new analysis and set of reports called Stats hier for N levels

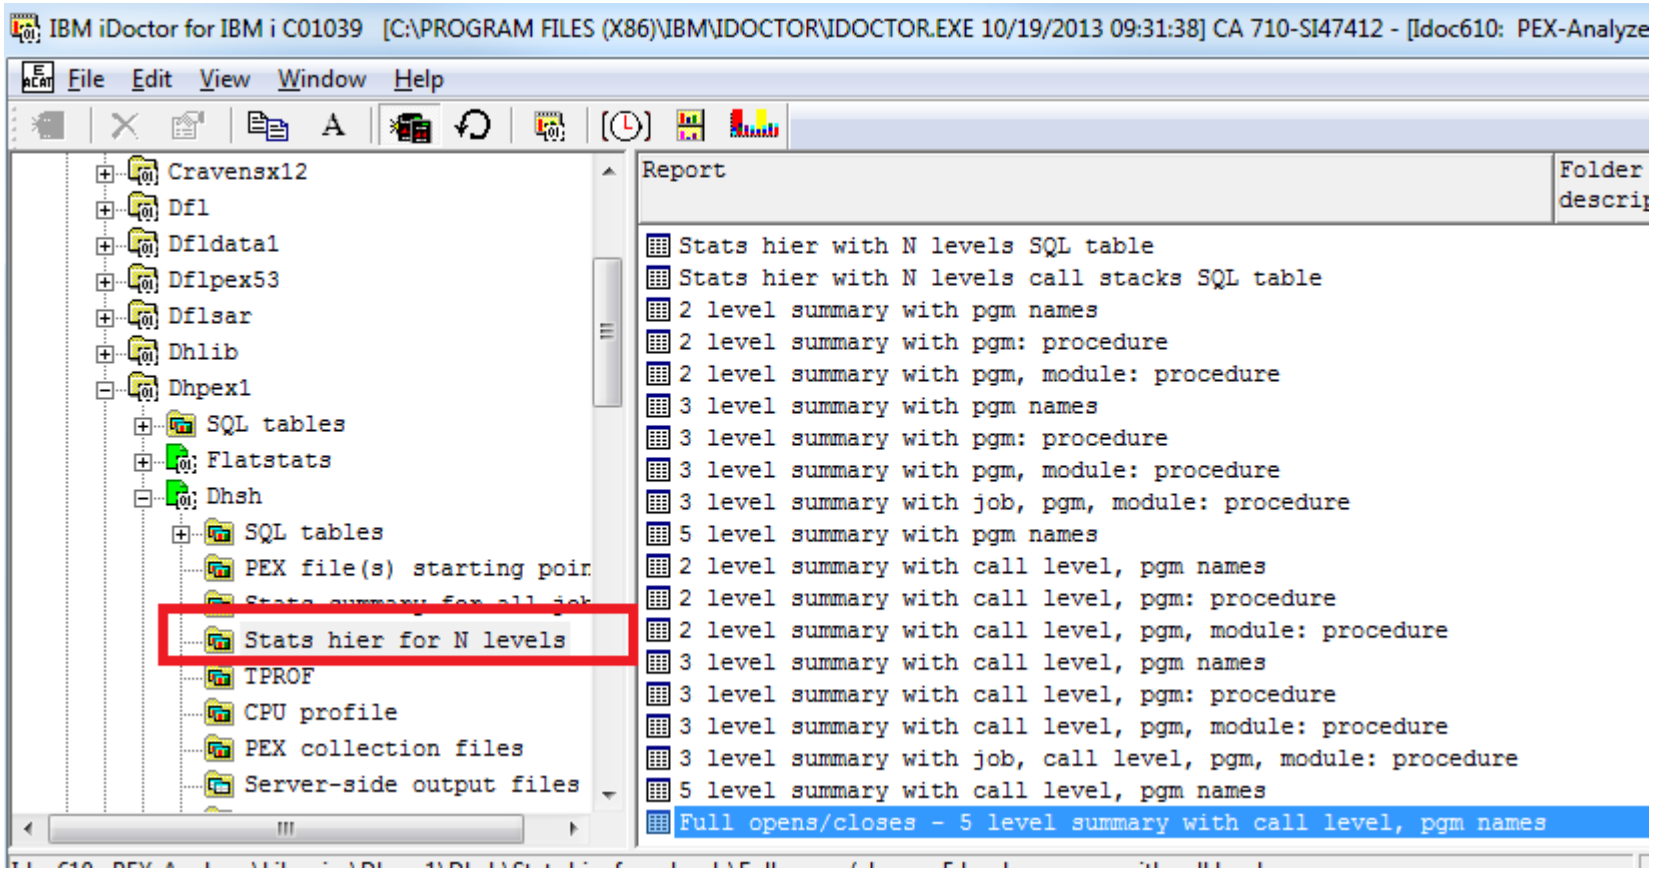

PEX Analyzer – Stats hier for N levels reports

- Full opens/closes – 5 level summary with call level, pgm names

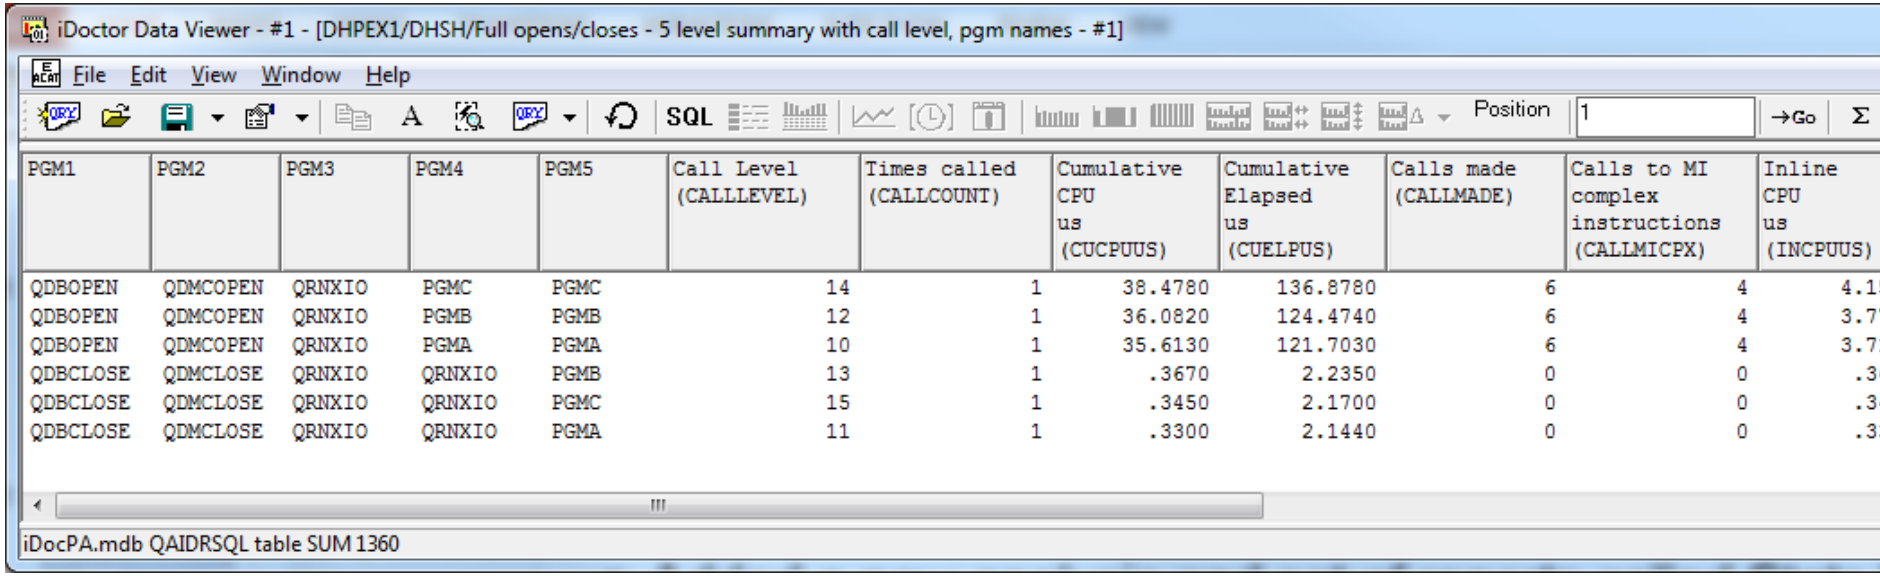

PGM1	PGM2	PGM3	PGM4	PGM5	Call Level (CALLLEVEL)	Times called (CALLCOUNT)	Cumulative CPU us (CUCPUUS)	Cumulative Elapsed us (CUELPUS)	Calls made (CALLMADE)	Calls to MI complex instructions (CALLMICPX)	Inline CPU us (INCPUUS)
QDBOPEN	QDMCOPEN	QRNXIO	PGMC	PGMC	14	1	38.4780	136.8780	6	4	4.1
QDBOPEN	QDMCOPEN	QRNXIO	PGMB	PGMB	12	1	36.0820	124.4740	6	4	3.7
QDBOPEN	QDMCOPEN	QRNXIO	PGMA	PGMA	10	1	35.6130	121.7030	6	4	3.7
QDBCLOSE	QDMCLOSE	QRNXIO	QRNXIO	PGMB	13	1	.3670	2.2350	0	0	.3
QDBCLOSE	QDMCLOSE	QRNXIO	QRNXIO	PGMC	15	1	.3450	2.1700	0	0	.3
QDBCLOSE	QDMCLOSE	QRNXIO	QRNXIO	PGMA	11	1	.3300	2.1440	0	0	.3

PEX Analyzer – Taskswitch report updates

- Added several column descriptions.
- Added JTTHREAD (job thread/task name) to several reports near the beginning.
- Added WAKER_JTTHREAD (waker's job thread name) where needed.
- Added EVENTDESC field to show the event info based on the new preference.
- Also added new types of reports in 2 folders at the bottom containing reports like:
  - Disk waits summary by job/thread/enum
  - Disk waits summary by job/thread/enum/object
  - Bad waits summary by job/thread/enum
  - Bad waits summary by job/thread/bucket
  - Bad waits summary by job/thread/object/bucket
  - Disk waits summary by object/enum
  - Disk waits summary by object/bucket
  - Bad waits summary by object/enum
  - Bad waits summary by object/bucket

PEX Analyzer – Page fault reports

- Added new reports for page fault events under the **PEX file(s) starting points – Page faults** folder:
  - Page faults by generic object name
  - Page faults by object
  - Page faults by object key

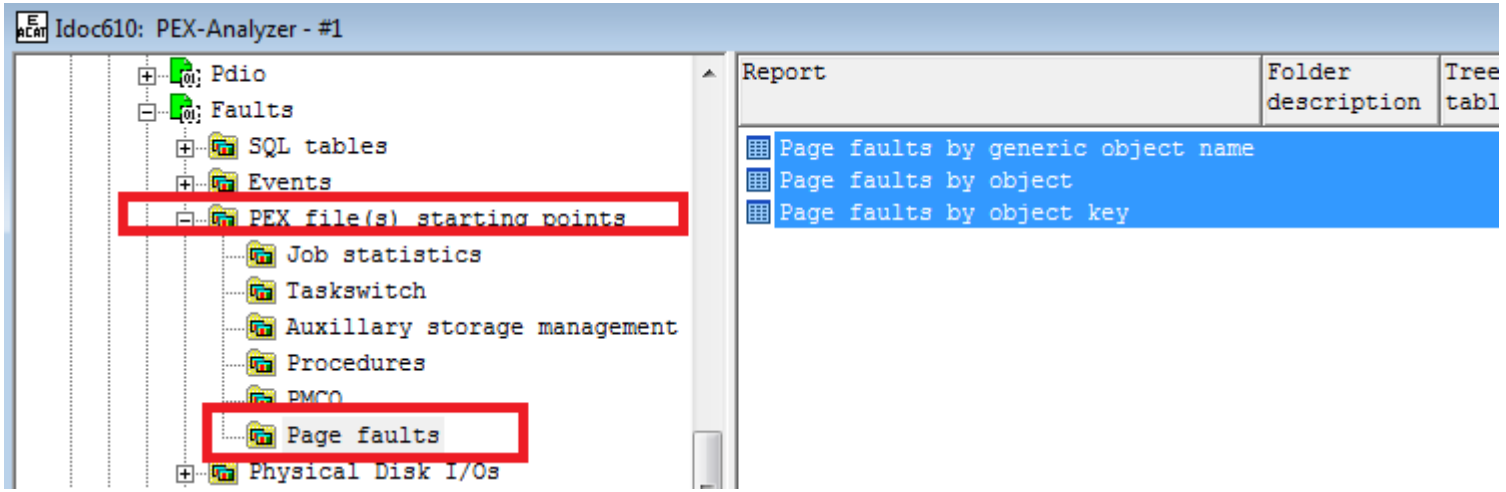

PEX Analyzer – Page fault reports

– Page faults by generic object name

iDoctor Data Viewer - #1 - [MCCARGAR/FAULTS/Page faults by generic object name - #1]

File Edit View Window Help

SQL Position 1 →Go Σ

Generic object name (GENOBJNAME)	Page Fault (EVT_6)	Page Fault Start (EVT_6_2)	Page Fault End OK (EVT_6_3)	Page Fault End with Error (EVT_6_4)
QTEM*	426	213	213	0
SQLS*	246	123	123	0
SQL.*	238	238	0	0
TERA*	192	148	44	0
QQSP*	135	135	0	0
QAYP*	57	29	28	0
MWS*	54	27	27	0
*	39	39	0	0
QWMS*	30	15	15	0
DB E*	24	12	12	0
MCCA*	24	12	12	0
QYPE*	22	11	11	0
QAPM*	18	9	9	0
QSQS*	17	17	0	0
PERM*	14	7	7	0
L/L*	8	4	4	0
Q040*	8	4	4	0
QIDR*	6	3	3	0
HOUN*	6	3	3	0
QSQR*	4	2	2	0
SMTR*	4	2	2	0
SYSI*	4	2	2	0
QRPL*	4	2	2	0
ENDP*	4	2	2	0
QWCC*	4	2	2	0
T172*	3	2	1	0
QJOB*	2	1	1	0

iDocPA.mdb QAIDRSQL table SUM 800 Rows 1 - 27 of 31

PEX Analyzer – Page fault reports

– Page faults by object

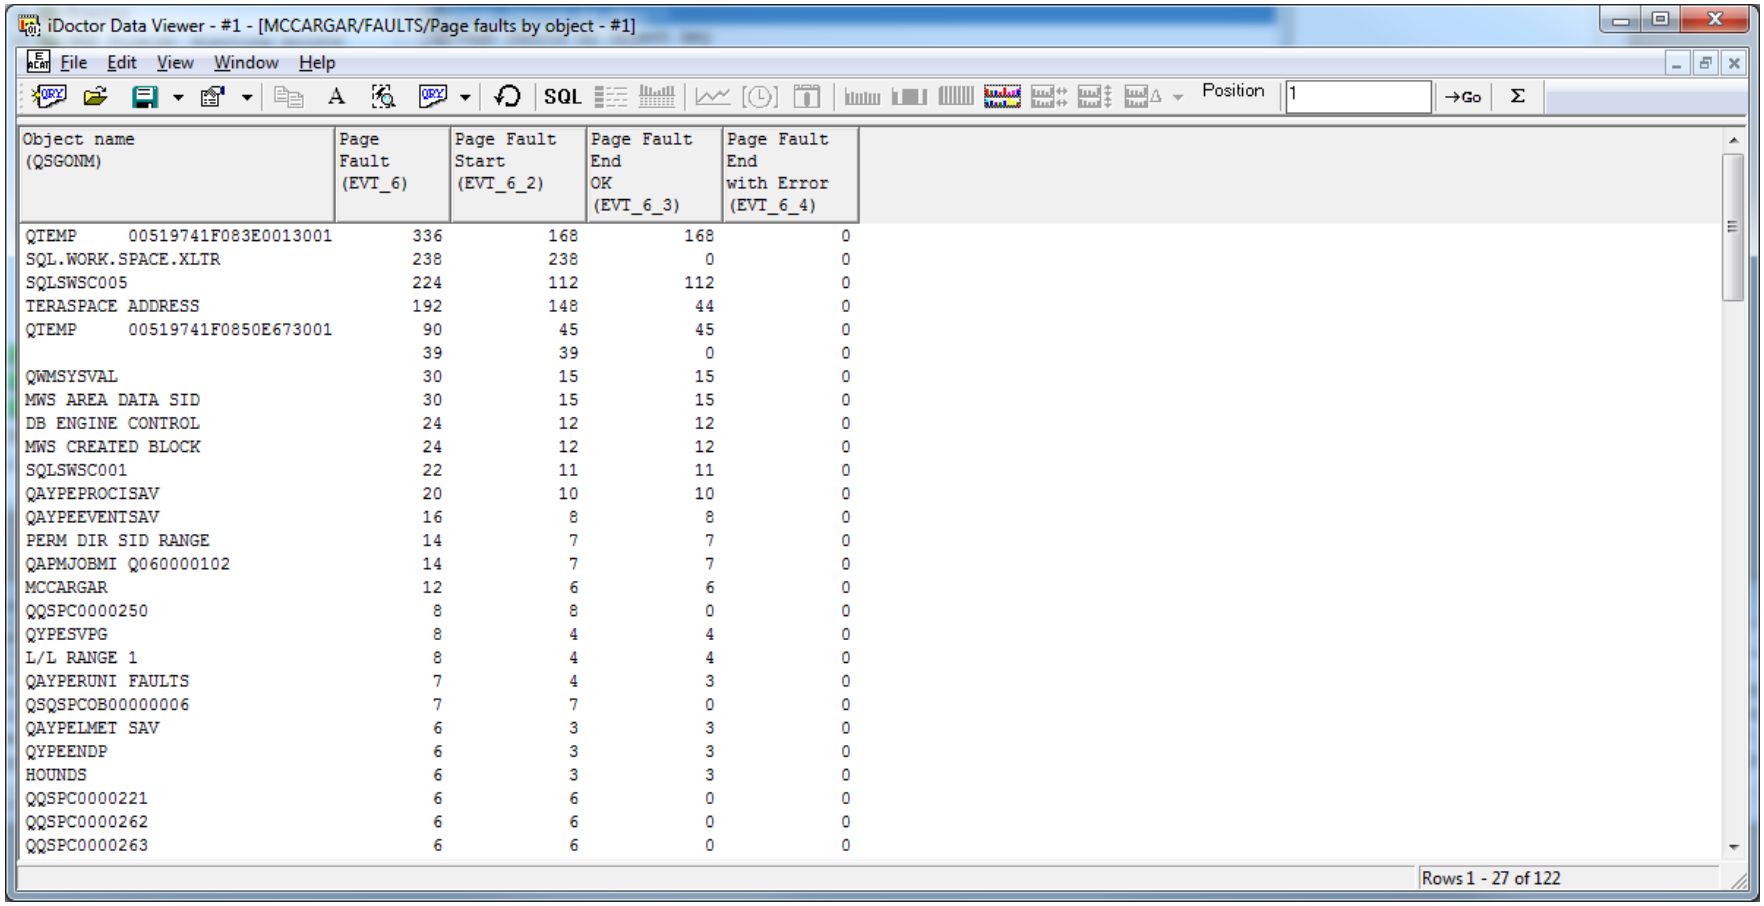

Object name (QSGONM)	Page Fault (EVT_6)	Page Fault Start (EVT_6_2)	Page Fault End OK (EVT_6_3)	Page Fault End with Error (EVT_6_4)
QTEMP 00519741F083E0013001	336	168	168	0
SQL.WORK.SPACE.XLTR	238	238	0	0
SQLSWSC005	224	112	112	0
TERASPACE ADDRESS	192	148	44	0
QTEMP 00519741F0850E673001	90	45	45	0
	39	39	0	0
QWMSYSVAL	30	15	15	0
MWS AREA DATA SID	30	15	15	0
DB ENGINE CONTROL	24	12	12	0
MWS CREATED BLOCK	24	12	12	0
SQLSWSC001	22	11	11	0
QAYPEPROCISAV	20	10	10	0
QAYPEEVENTSAV	16	8	8	0
PERM DIR SID RANGE	14	7	7	0
QAPMJOBMI Q060000102	14	7	7	0
MCCARGAR	12	6	6	0
QQSPC0000250	8	8	0	0
QYPESVPG	8	4	4	0
L/L RANGE 1	8	4	4	0
QAYPERUNI FAULTS	7	4	3	0
QSQSPCOB00000006	7	7	0	0
QAYPELMET SAV	6	3	3	0
QYPEENDP	6	3	3	0
HOUNDS	6	3	3	0
QQSPC0000221	6	6	0	0
QQSPC0000262	6	6	0	0
QQSPC0000263	6	6	0	0

PEX Analyzer – Convert HEX values preference

- Added a preference to avoid converting hex values to decimal which if used will improve performance.
- If the option is used to show the values as HEX the column descriptions will indicate this.
- Placing your mouse over any of these hex values will show the decimal value in the flyover.

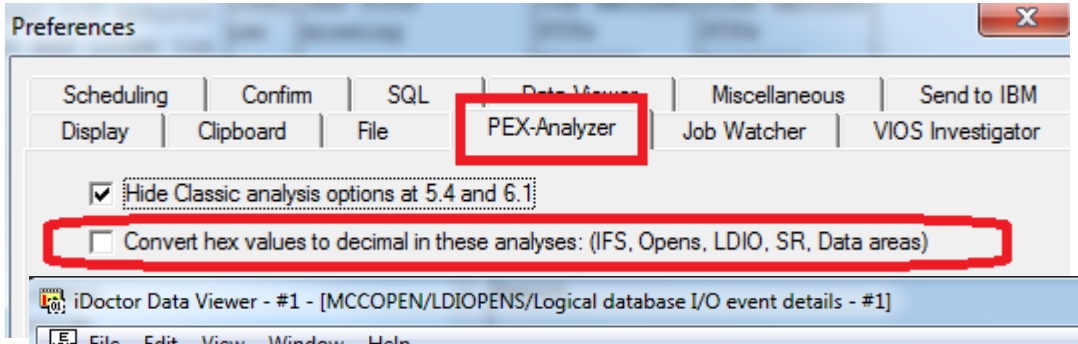

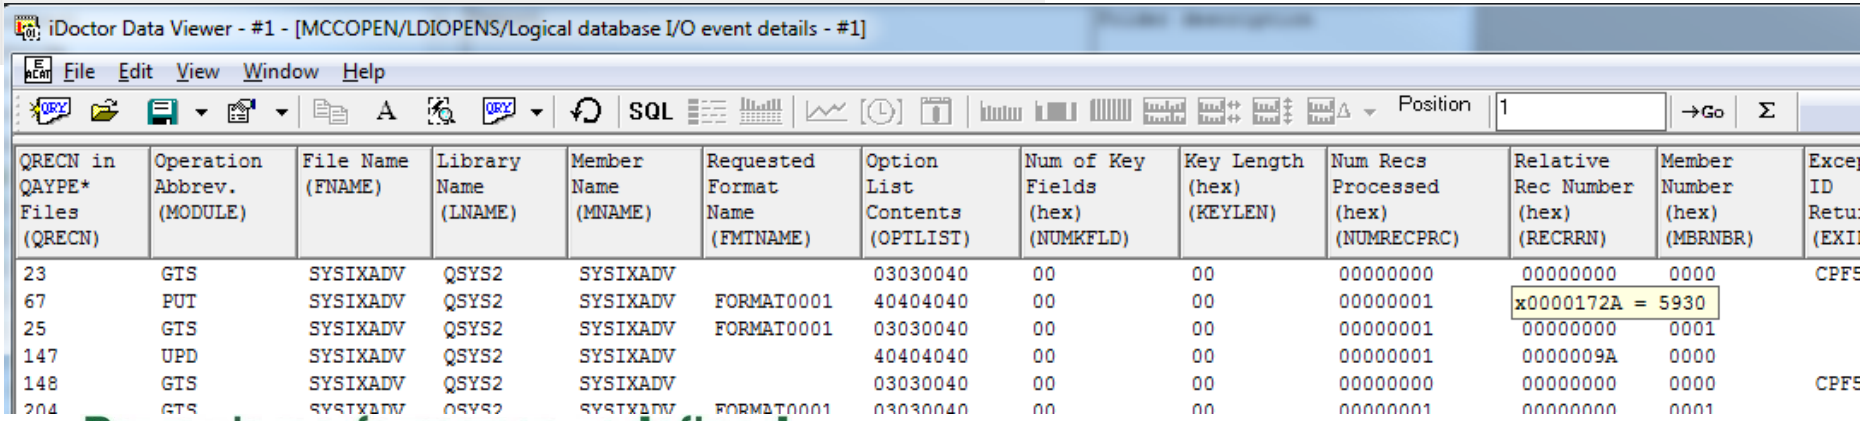

PEX Analyzer – TPROF analysis – MCLI options

- Added MCLI options within the TPROF analysis reports under the following folder:
  - Hits by SDAR
  - Hits by SDAR/seize object

Note: If the folders are not shown or do not contain anything then you do not have the required PEX definition needed to produce the required data.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: http://www.**ibm.com**/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
AltiVec is a trademark of Freescale Semiconductor, Inc.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.

Revised December 2, 2010

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
	Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seg=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- **rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.**

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

==========

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007