Power Systems

IBM

iDoctor update

Ron McCargar – idoctor@us.ibm.com

IBM

LUG

Agenda

- Overview
  - Goals, components, unique features and website.
- Latest major updates to iDoctor
  - Covering the last 2 years.
- Questions?
  - Contact idoctor@us.ibm.com or mccargar@us.ibm.com

Overview

- iDoctor is a suite of dynamic performance tools
  - Developed by the IBM i Global Support Center.
  - Sold via Lab Services

- We cover all areas of performance
  - Started in V4R5
  - 7 IBM i components under “IBM i Connections”(2 free)
  - 1 non-IBM i components under “Power Connections” (free)

iDoctor goals

- Broaden the user base for Performance Investigation
  - Operators, programmers, consultants, etc.
- Simplify and automate processes
- Provide quick access to data
- Provide more analysis options
- Reduce the dependency on PEX traces

PRODUCTIVITY

IBM i Component	Description
Collection Services Investigator (CSI)	High-level system/job monitoring Specializes in job waits analysis
Job Watcher (JW)	Medium-level job analysis Specializes in job waits analysis
PEX Analyzer (PA)	Low-level system/job/disk tracing stats, profiling
Disk Watcher (DW)	Disk stats and traces See also CSI – disk graphs and/or PEX Analyzer – Physical Disk I/O
Plan Cache Analyzer	Analyses the system plan cache.
iDoctor FTP GUI (free)	Work with the IFS and libraries/objects.
Must Gather Tools (free)	GUI for QMGTOOLs library **Note: GUI** mostly untouched in last 2 years so much is likely broken.

Power Component (non IBM i)	Description
Power Connections (free)	Manage your HMC, VIOS, AIX or Linux connections and drill down for more details. Includes support to graph nmon data on the PC (with SQLite) or on an IBM i. On VIOS: Collects or analyzes nmon, npiv or VIOS Advisor data. On VIOS: Includes support to automatically download/install and run the latest PerfPMR.

iDoctor GUI

- Built on Windows
  - Superior flexibility and functionality
  - Support to run on Citrix servers
  - Runs on Linux KVMs, VirtualBox, Parallels, etc.

- All components offer a similar user experience

- Requires:
  - System i Access for Windows or IBM i Access Client Solutions with the Windows add-on package.
    - Note: only if using the IBM i components.
  - .NET 4.0 or higher
  - Visual Studio 2012 redistributable package

CSI – System graphs

- New graph 7.2+
  - LPAR VCPU delays with HMC CPU utilization by partition
  - LPAR CPU time with entitlement consumed from HMC perspective
  - LPAR entitled CPU time

  - LPARs with capacity changes
  - LPARs with changes to memory allocated
  - LPARs with changes to current virtual processors
  - LPARs with changes to current processing capacity

.

CSI – System graphs -> TLBIEs

- TLBIEs subfolder – 26 graphs

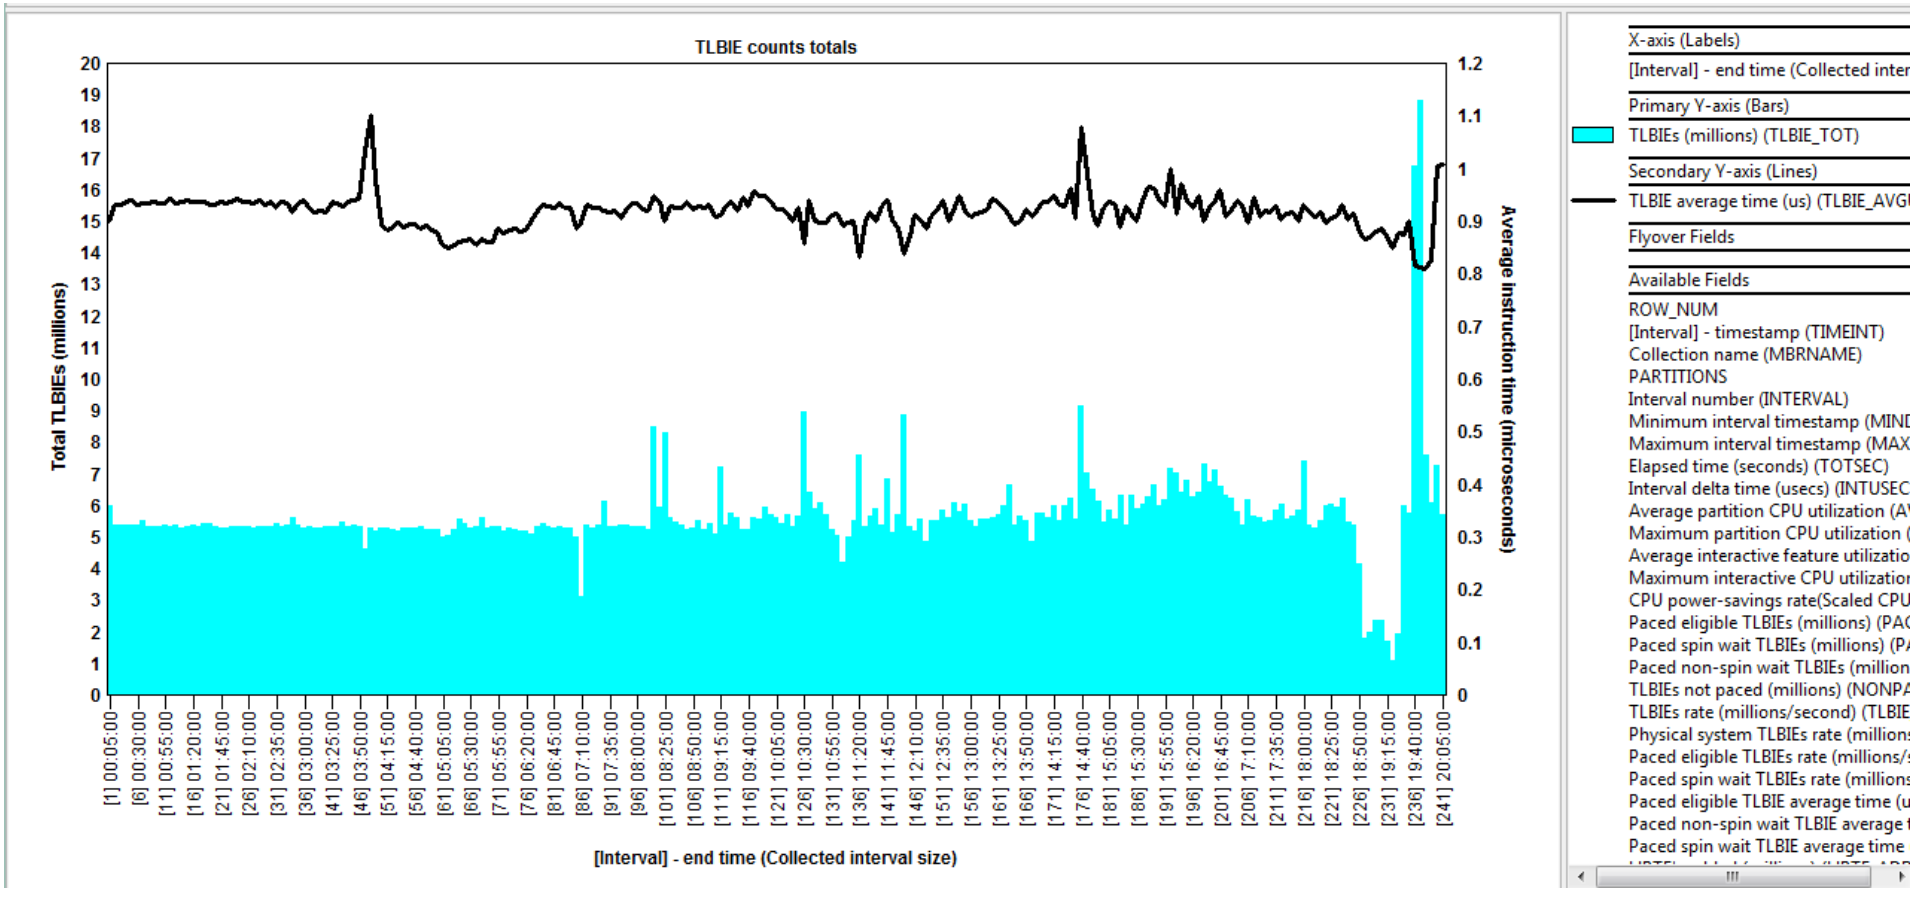

CSI – System graphs -> Interrupts

- 7.2+ with PTFs – 7 graphs

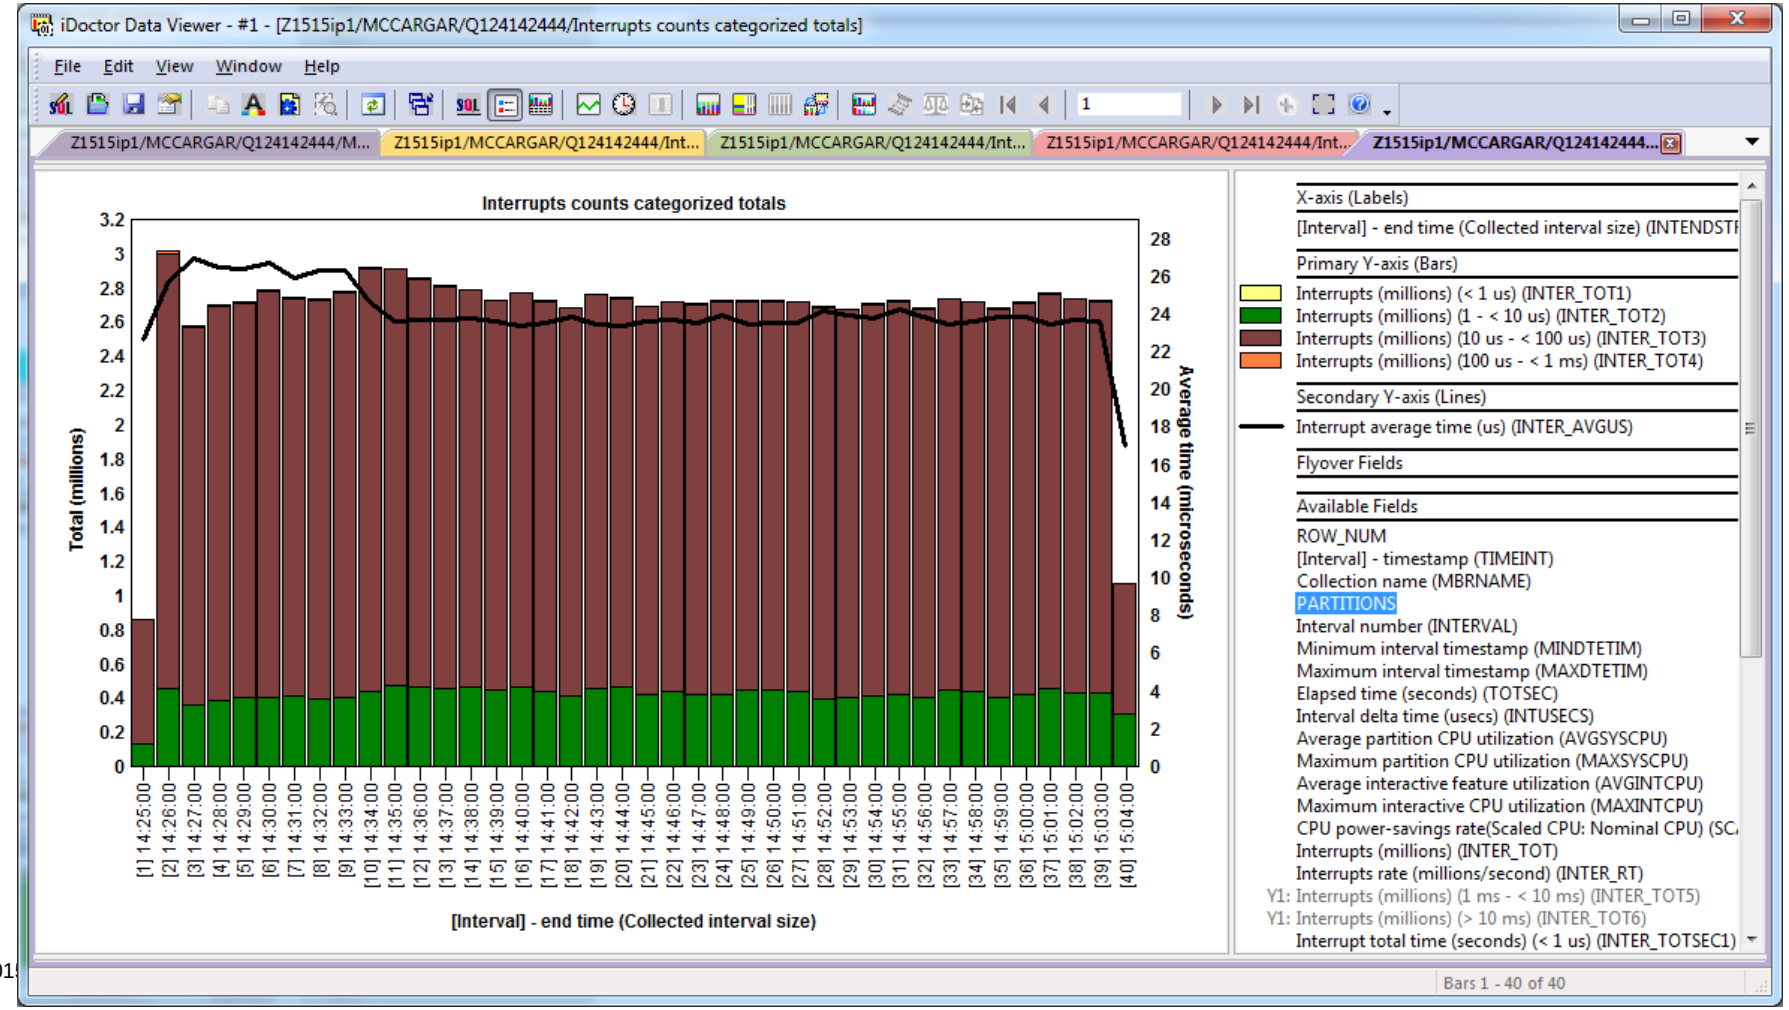

CSI – Disk graphs -> Free space map

- 7.2+ with PTFs

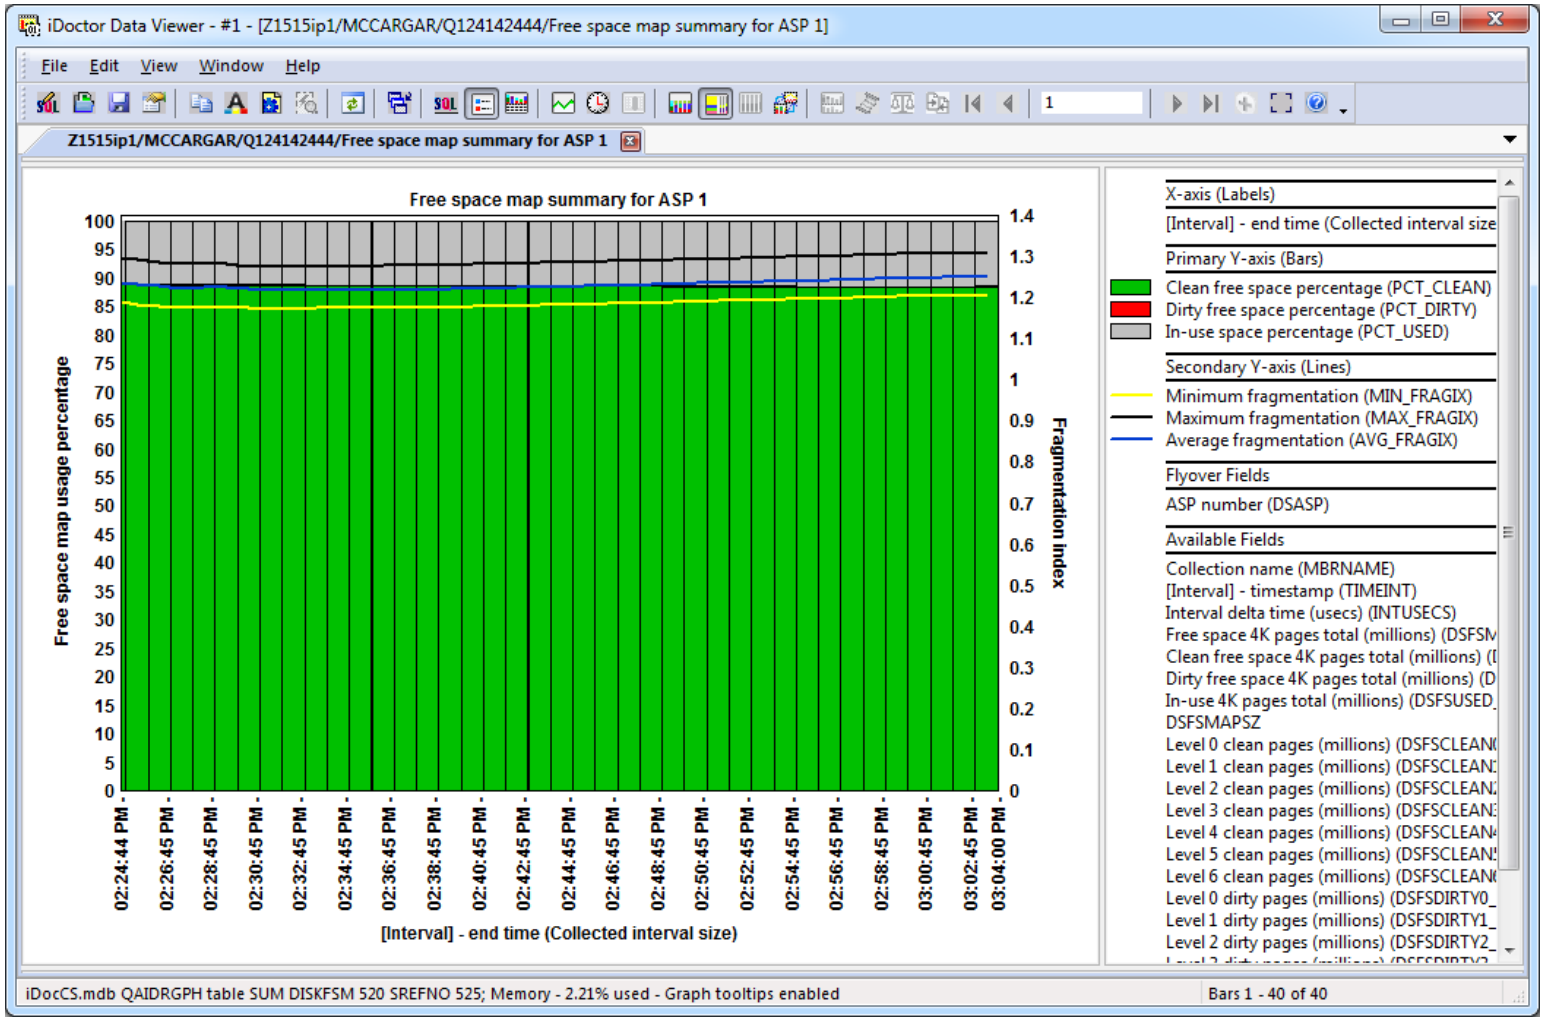

CSI – Disk graphs -> Free space map

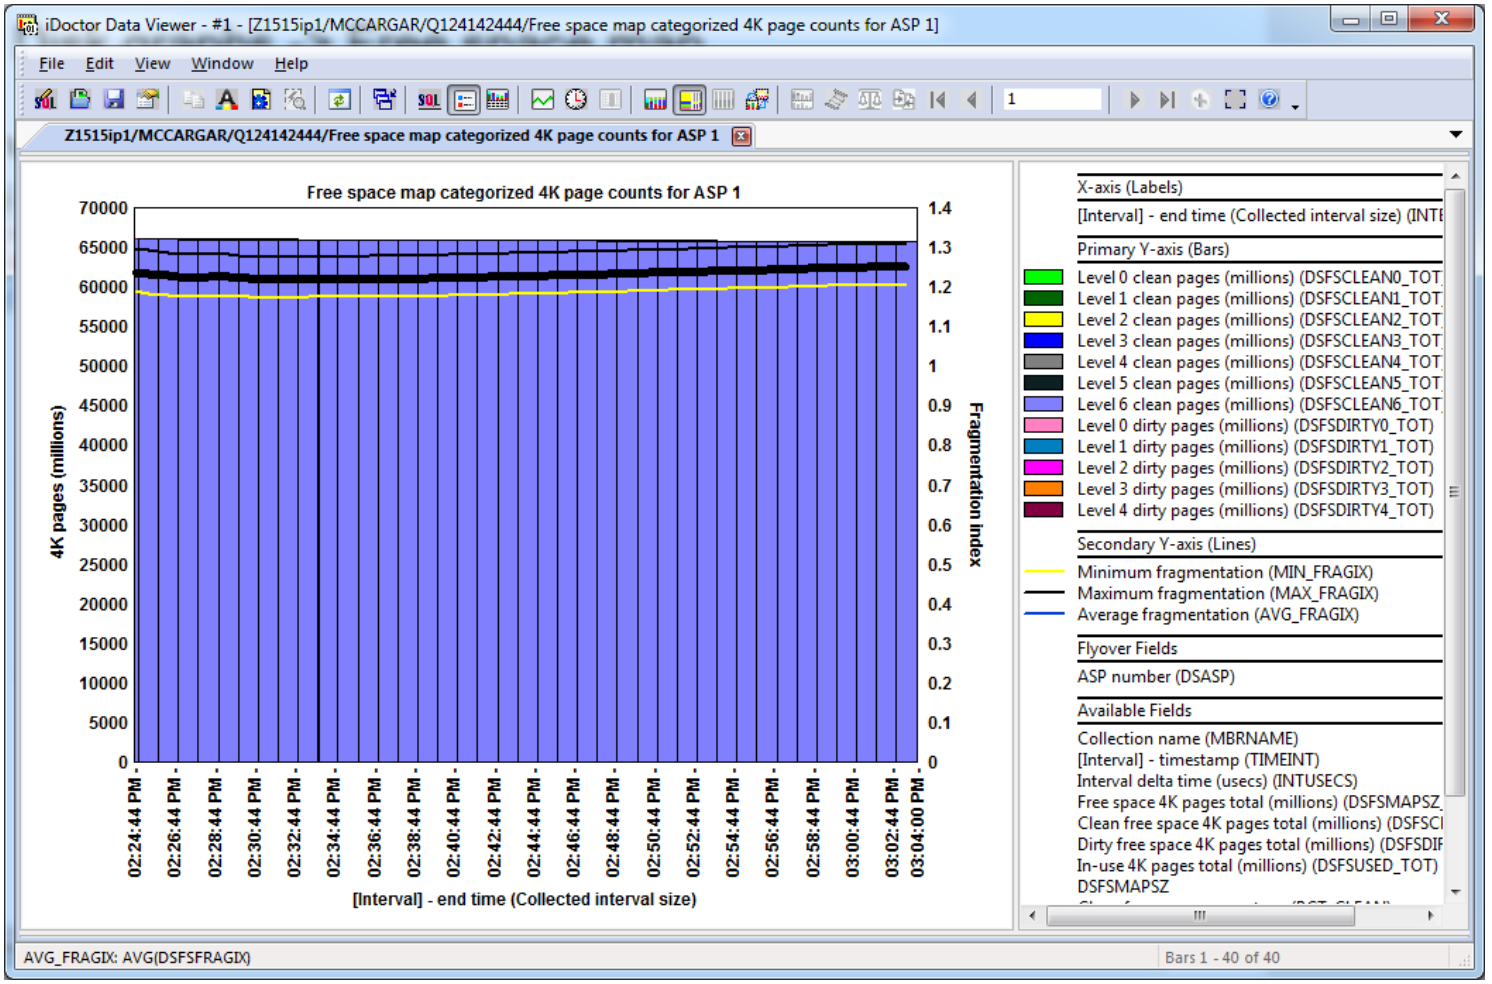

CSI – Disk graphs -> Free space map

- Most fragmented units listed first in graph legend. Top 20 shown by default.

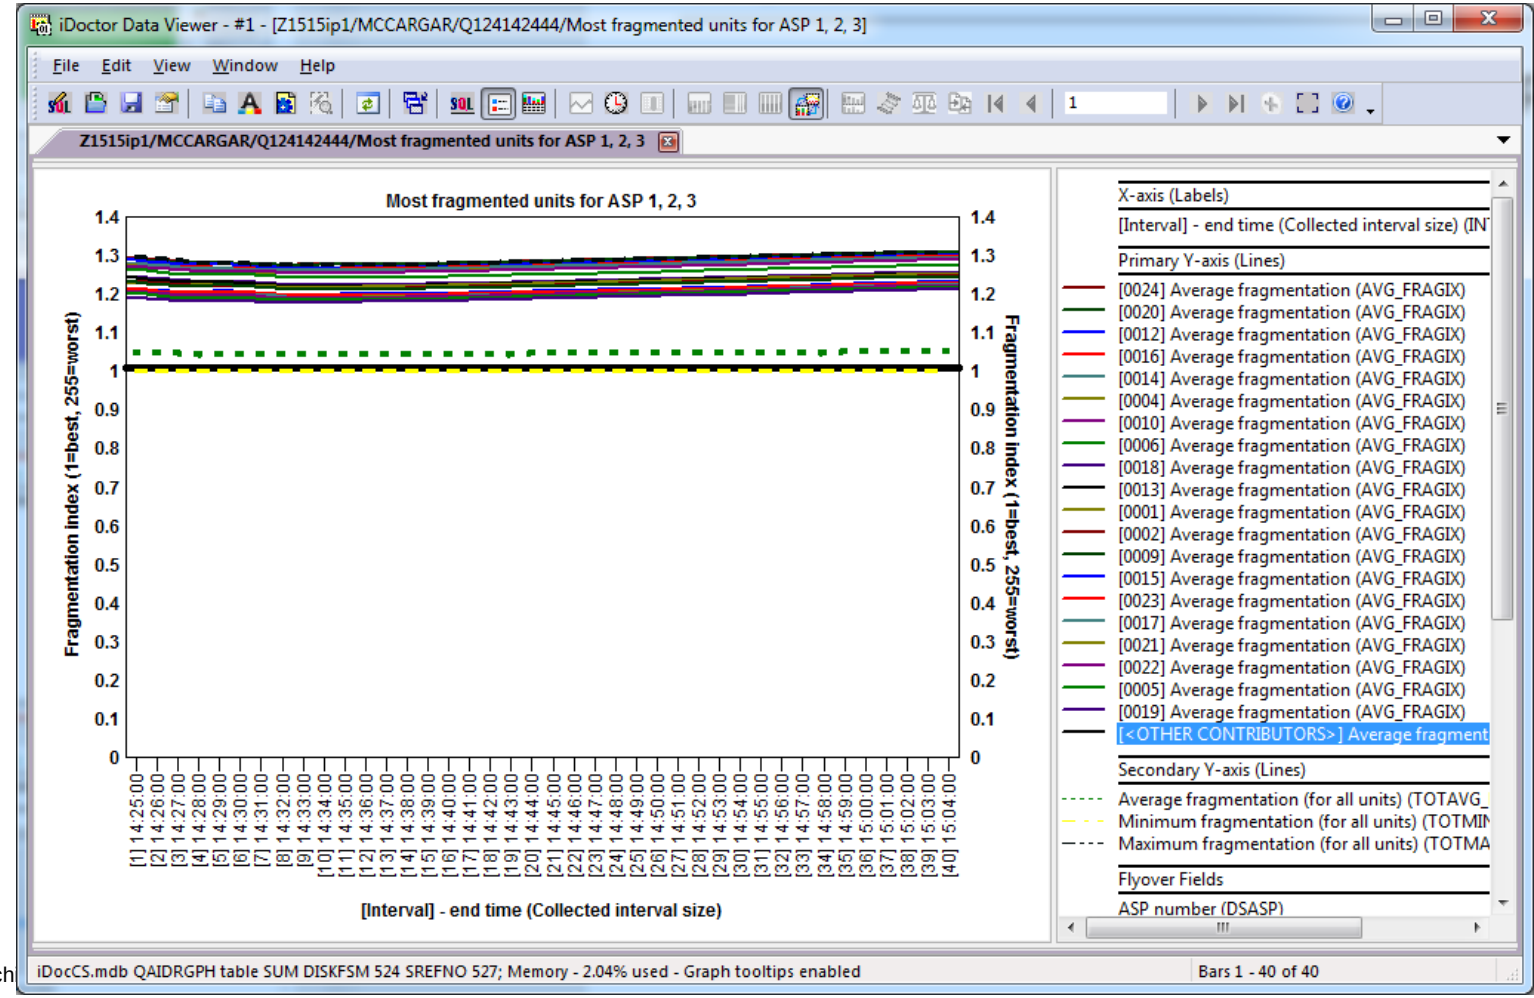

CSI – Memory pool graphs

- New graph
  - Machine pool consumption breakdown

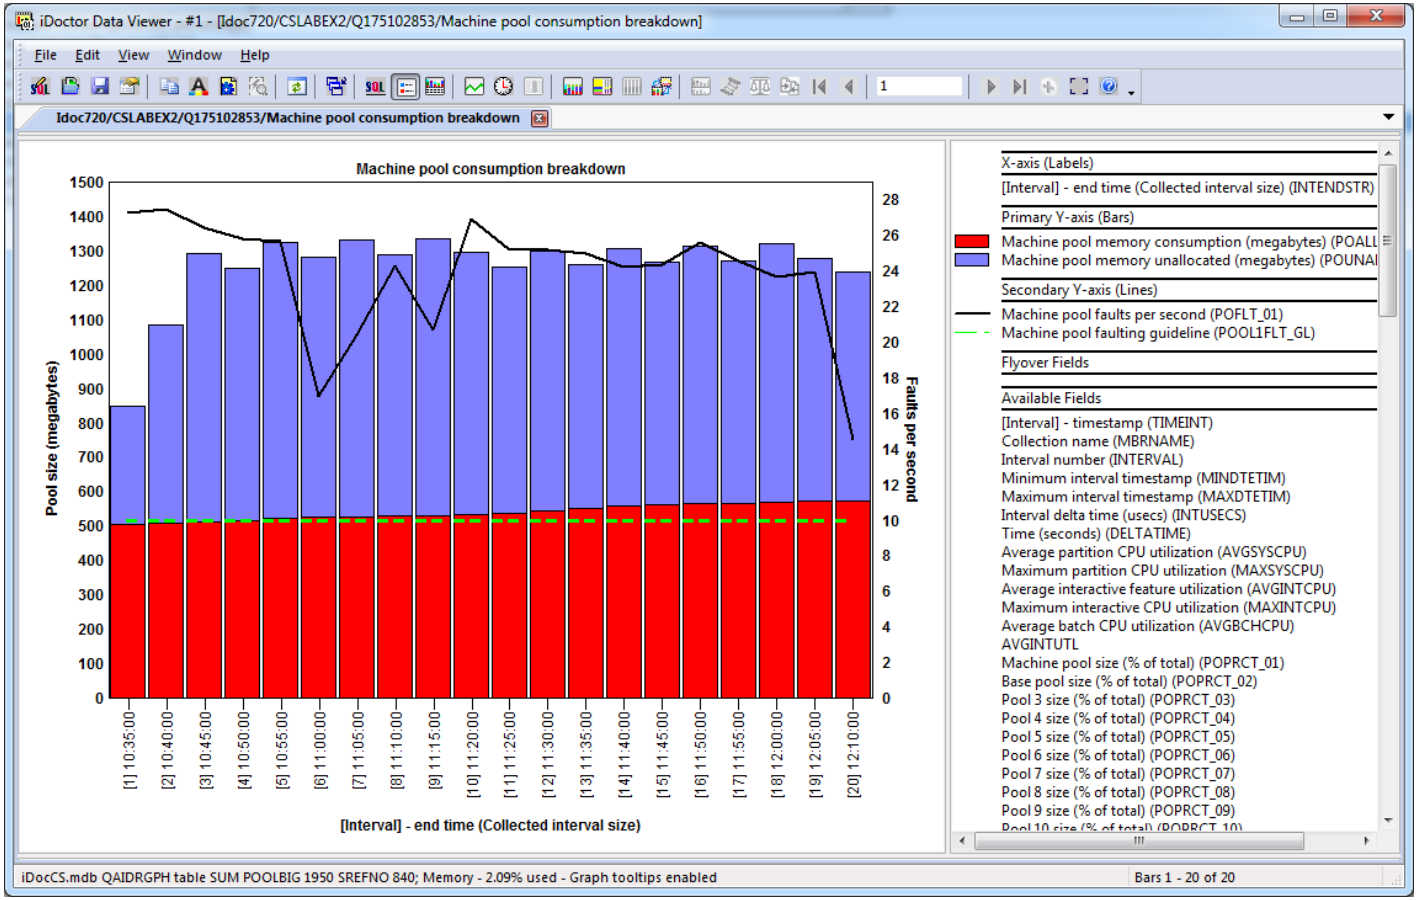

CSI – External storage links and ranks graphs

- New Links graph types
  - Flattened by serial number
  - By DS serial, interface ID

- New Ranks graph types
  - Flattened by DS serial
  - By DS serial, rank ID

CSI – External storage cache statistics

- Graph types:
  - Flattened on ASP
  - Flattened on LSS ID
  - By disk unit
  - By disk path
  - By ASP
  - By disk serial
  - By LUN ID
  - By LUN ID and path
  - By LSS ID
  - By ASP, extent pool
  - By LUN ID, disk unit, extent pool
  - By LSS ID, extent pool

CSI – System graphs -> LPAR configuration

- 3 reports showing CPU / memory capacity

Collection name (MBRNAME)	Partition (PART...	Partition ID (PARTITIONI...	Virtual shared pool ID (HPPOOLI...	OS	Processor sharing (PROCSHARE)	Current virtual processors (HPVPRC)	Current processing capacity (HPPRCCC)	Memory allocated (GBs) (HPMEMC_GBS)
Q004000002	AIX->	0025	0	Other OS	Capped sharing	1	.2000	1
Q004000002	AIX>	0015	0	Other OS	Uncapped sharing	0	0	.2500
Q004000002	Clien>	002C	0	IBM i	Capped sharing	0	0	.2500
Q004000002	DO>	0007	0	IBM i	Capped sharing	1	.1000	50
Q004000002	IGSC>	0023	0	Other OS	Uncapped sharing	0	0	.2500
Q004000002	LPD>	0003	0	IBM i	Uncapped sharing	20	10	50
Q004000002	MOT>	0028	255	IBM i	Not shared	0	0	0
Q004000002	MPI>	001A	0	Other OS	Capped sharing	1	1	4.2500
Q004000002	MPI>	001B	0	Other OS	Capped sharing	1	1	3.2500
Q004000002	MPI>	0022	0	IBM i	Capped sharing	1	.2000	3.2500
Q004000002	MTS>	0012	0	Other OS	Uncapped sharing	4	4	8
Q004000002	Mak>	0009	0	Other OS	Uncapped sharing	1	1	4
Q004000002	Mak>	0017	0	Other OS	Capped sharing	1	1	4.2500
Q004000002	Mak>	0008	0	Other OS	Uncapped sharing	1	1	4
Q004000002	Mak>	000C	0	Other OS	Uncapped sharing	1	1	4
Q004000002	Mak>	0011	0	Other OS	Uncapped sharing	1	1	4.2500
Q004000002	Mak>	0013	0	Other OS	Uncapped sharing	1	1	4.2500
Q004000002	Mak>	0004	0	Other OS	Uncapped sharing	1	1	4.2500
Q004000002	Mak>	0018	0	Other OS	Uncapped sharing	1	1	4.2500
Q004000002	Mak>	0019	0	Other OS	Uncapped sharing	1	1	4.2500
Q004000002	RCH>	0001	255	IBM i	Not shared	1	1	16
Q004000002	RCH>	0002	0	IBM i	Uncapped sharing	4	4	32

CSI/JW – new job groupings

- All rankings graphs now have two new types of groupings:
  - By generic job, current user
  - By thread, current user

- You may also drill down from these on the selection made.

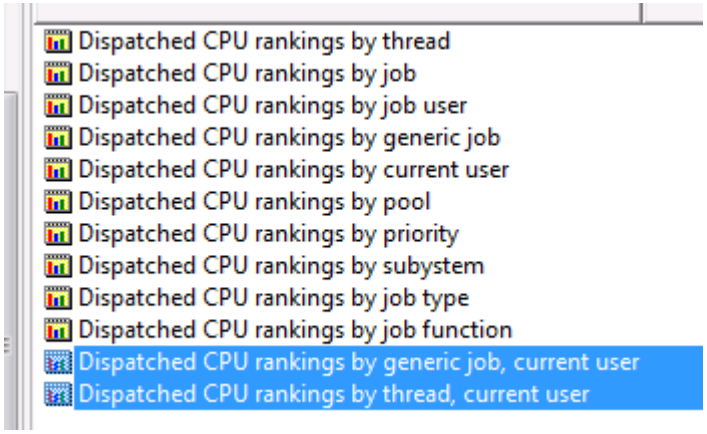

CSI/JW – Graph Jobs

- New option when you right-click a collection called "Graph Job(s)."
- Graph the thread wait time signature graph for any job /thread/task in any collection (for the component you are currently working with) on any system!
- Add a second collection/job and do a comparison (optionally.)

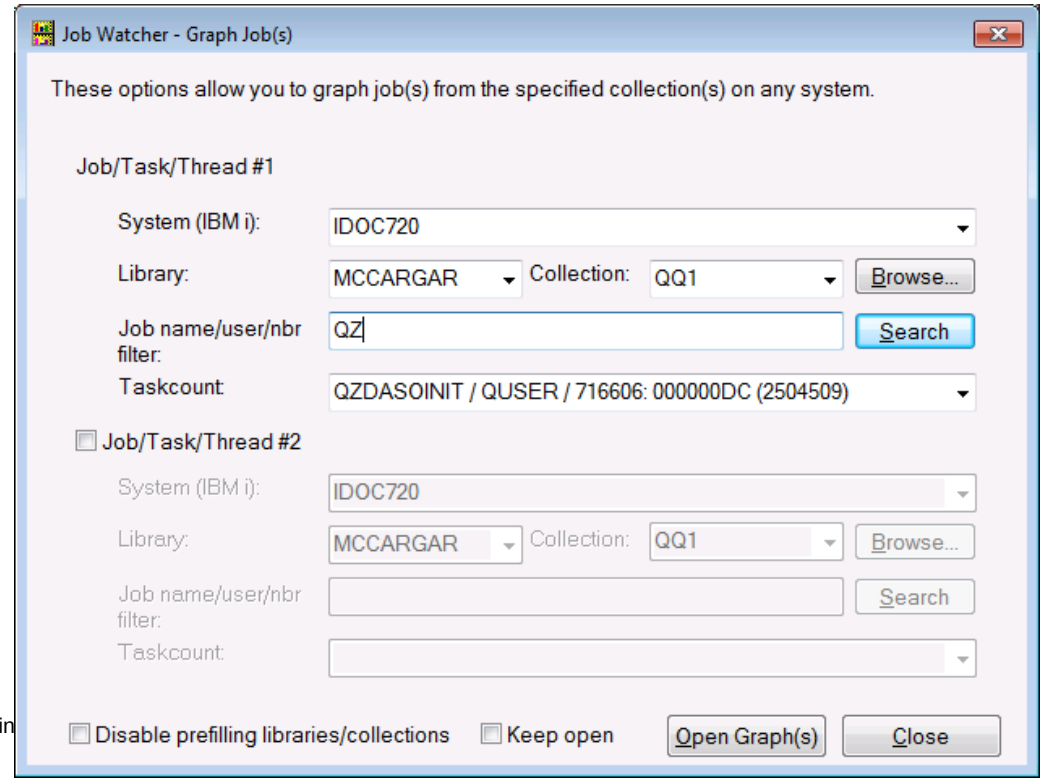

Transfer function

- In all of the IBM i components updated the **Transfer to** function for all types of performance collections to now include the following options:

  PC
  Another IBM i
  FTP Server
  IBM – Testcase
- IBM – ECUREP
- IBM - Blue Diamond Lab

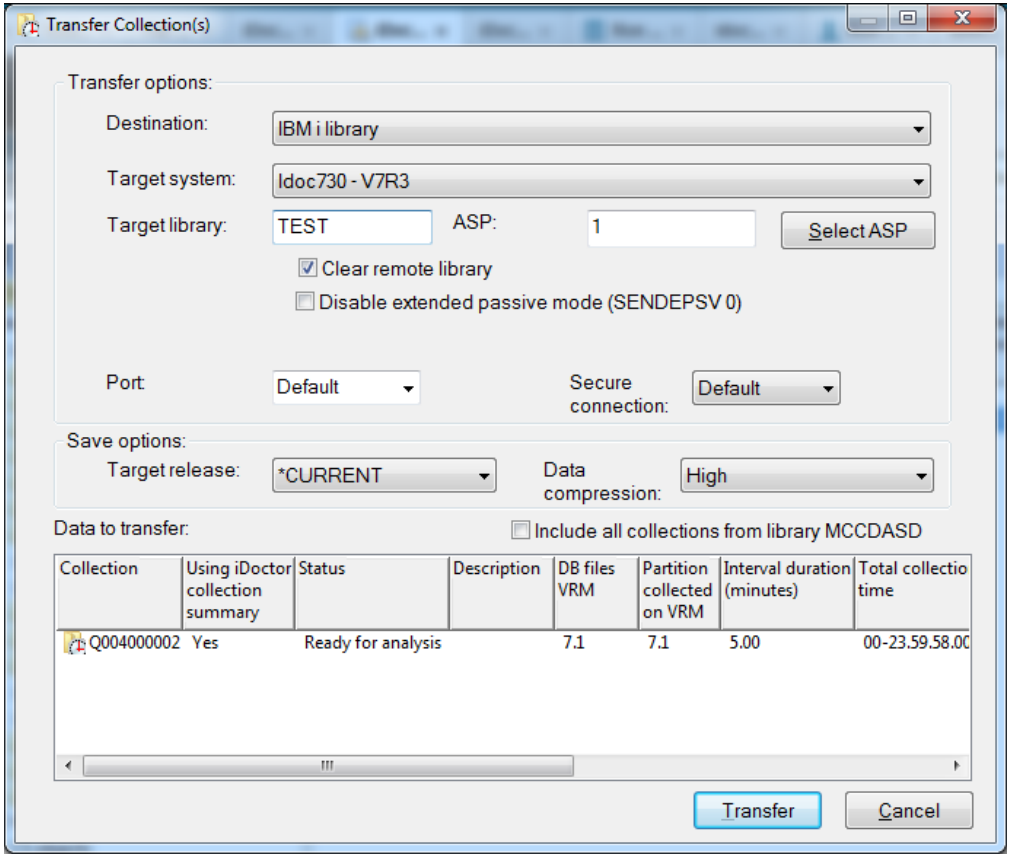

Data Viewer toolbar changes

- Added 4 new buttons to control movement within the current result set. A row set is by default 1000 rows but can be controlled via Preferences -> Data Viewer.
- **Note:** These were created mainly for use with SQLite but also has some benefit when using IBM i as the DB to move through data quickly. When using SQLite, the row set size defines how much data is loaded into memory at one time.

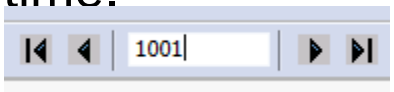

  - First Row Set – jump to the beginning
  - Previous Row Set – go back N rows
  - Next Row Set – move forward N rows
  - Last Row Set – jump to the final row set (or N rows from the end)

Data Viewer toolbar changes

- Added a button to toggle a graph from vertical stacked bars to lines. Note: The option is only enabled if the graph is a vertical stacked bar or a line graph.

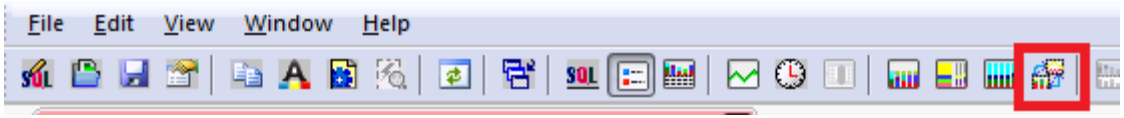

Data Viewer graph changes

- In the legend the user can now press or hold down the up or down arrows on the keyboard to step through the fields. After each selection is made the graph will update.

- Left clicking on the graph no longer zooms out. Right-click and use the Zoom out menu instead. You can also zoom in multiple times now.

- Graph flyovers:
  - 1. Y2 fields that have a value of 0 are no longer shown.
  - 2. Flyover fields with a blank value are no longer shown.
  - 3. An * is now displayed next to the Y1 or Y2 field the mouse was over in the flyover.
  - 4. Each type of field will now have an identifier before it like X:, FT: (flattened field), Y1:, Y2: and TIP: (flyovers.)
  - Also added a preference on the Tips tab called "Include all non-zero Y2 fields in graph flyovers" with a default value of checked.

Data Viewer graph legend changes

- Remove all except selected menu option
  - Quickly remove everything from the graph (in that section) except the selection.

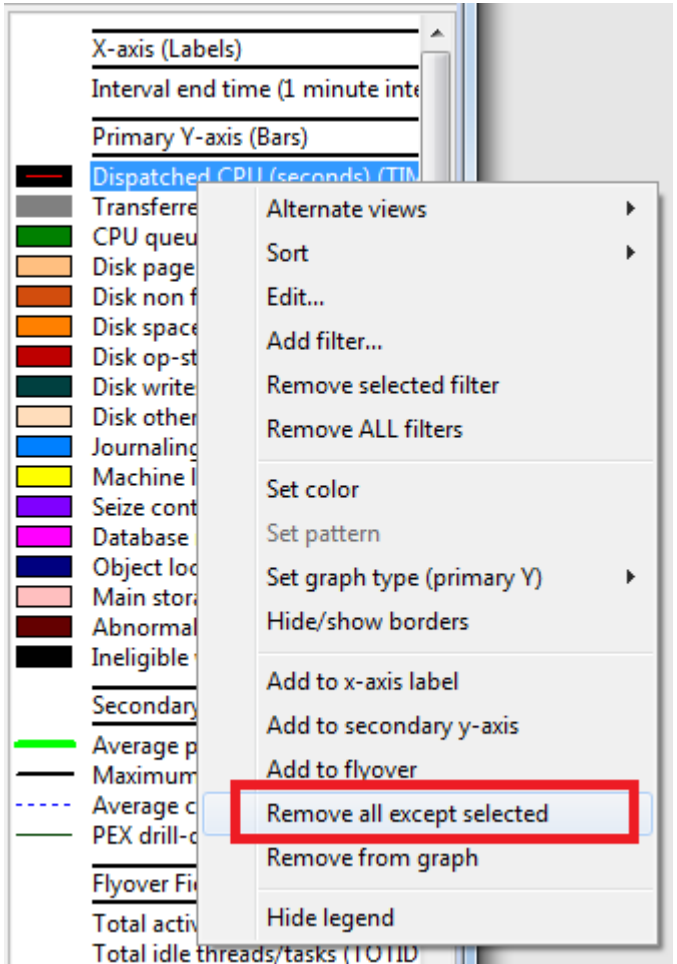

Power Connections: VIOS Monitor Support

- Added the ADDMONDIR, RMVMONDIR commands for continual data collection on VIOS and sent to IBM i for analysis (additional setup required.)

- General functions -> Power -> Monitors folder

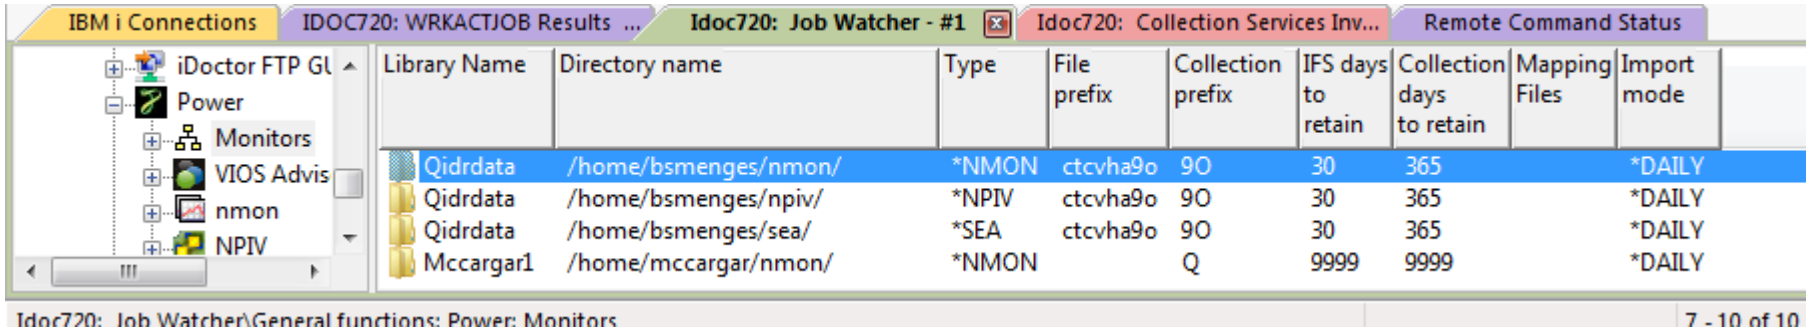

Power Connections

- Added support to collect and analyze SEA (Shared Ethernet Adapter) data on VIOS.
- Added support to merge multiple nmon, npiv or sea data files into a single iDoctor collection.
- Added these functions for nmon, npiv or SEA collection types:
  - Copy
  - Rename
  - Save
  - Transfer
  - Analysis -> Change sensitive user data

Power Connections

- Added support to analyze (graph) NPIV or SEA data captured by NMON.
- Added several additional NPIV graphs and drill downs from physical fiber channels to virtual.
- Added a (RESET) flag within the NPIV data to indicate if the adapter was reset during data collection.
- When graphing nmon data, only folders and graphs will show up now for data that actually exists.
- Added support to graph FCSTAT within the nmon support.

Power Connections – HMC lslparutil support removed

- It was broken for a long time and no one reported it.
- It is somewhat obsolete, the support in CSI: “System graphs” has better graphing options as long as the Allow performance data collection flag is enabled.

PEX

- In PEX TPROF added a new drill down from any report that contains the QPRKEY field (Traceback table address): Link register resolution -> Advanced link register resolution for procedure XYZ.
- Also added TPROF -> Cacheline breakouts -> Call stack summary with cache line counts
- PEX event analysis graphs updated to handle the latest PEX event changes.

PEX Definition Wizard

- Now handles the following new events :
  - 1. At 7.2+ the *ACTDTA base event.
  - 2. At 7.1+ the *CHGSEGATR storage event
  - 3. At 7.2+ the *LOGSNSEND disk event
  - 4. At 7.2+ the 7 new database events

- Added support in the PEX Analyzer definition wizard for the subsystem and thread ID options for the JOB parameter.

- Added support for the TRCTYPE (trace type) parameter.

- Added support to load the PEX Definition from the ADDPEXDFN command string now stored within the PEX definitions file.

Job Watcher graphs

- SQL folder – 4 new graphs
- Communication graphs folder – 10+ graphs (if data captured)

Job Watcher (or PEX) single call stack modes

- Basic
  - Call level, pgm, module, procedure
- Detailed
  - Basic fields + offset, LIC statement number
- Advanced
  - Detailed fields + instr address, proc start/end address, TBT address.

Plan Cache Analyzer

- For 7.1+ plan cache snapshots:
  - Under Plans added new subfolders
    - With indexes advised
    - Without indexes advised
      - Note: These graphs are the same as the normal plan graphs except they show only plans where indexes were advised or not and includes the key columns advised in the flyover for the "with indexes advised" graphs.
  - Added Index Advice (tables) folder

Generate Reports function

- General collection information added to the report
- Added Preferences: bars per page, # of pages to capture, etc.
- Buttons added to rearrange the reports in the list.
- Can schedule reports to be generated at a later time.

Preferences

- Added a preference on the Misc tab called Maximize the Main Window that is unchecked by default.
- Added a preference on the Data Viewers tab called Maximize Data Viewers that is checked by default.
- Preferences -> Power all of the File transfer preferences have been removed. These are now defined on each connection.
- Moved the Preference on the Misc tab -- "Check for new builds when starting the iDoctor GUI" to the Confirm tab.

Monitors

- Added a GUI option for the Run default analyses parameter on the STRJWMON/STRDWMON/STRPAMON commands to the GUI.
- On the Basic Options screen of the Start iDoctor Monitor Wizard and is called "Run analyses automatically"
- Note: In some situations this could be resource intensive.

Collection Repository Graphs

- General Functions -> Browse collections folder provides 8 graphs showing the iDoctor collection repository in various ways.

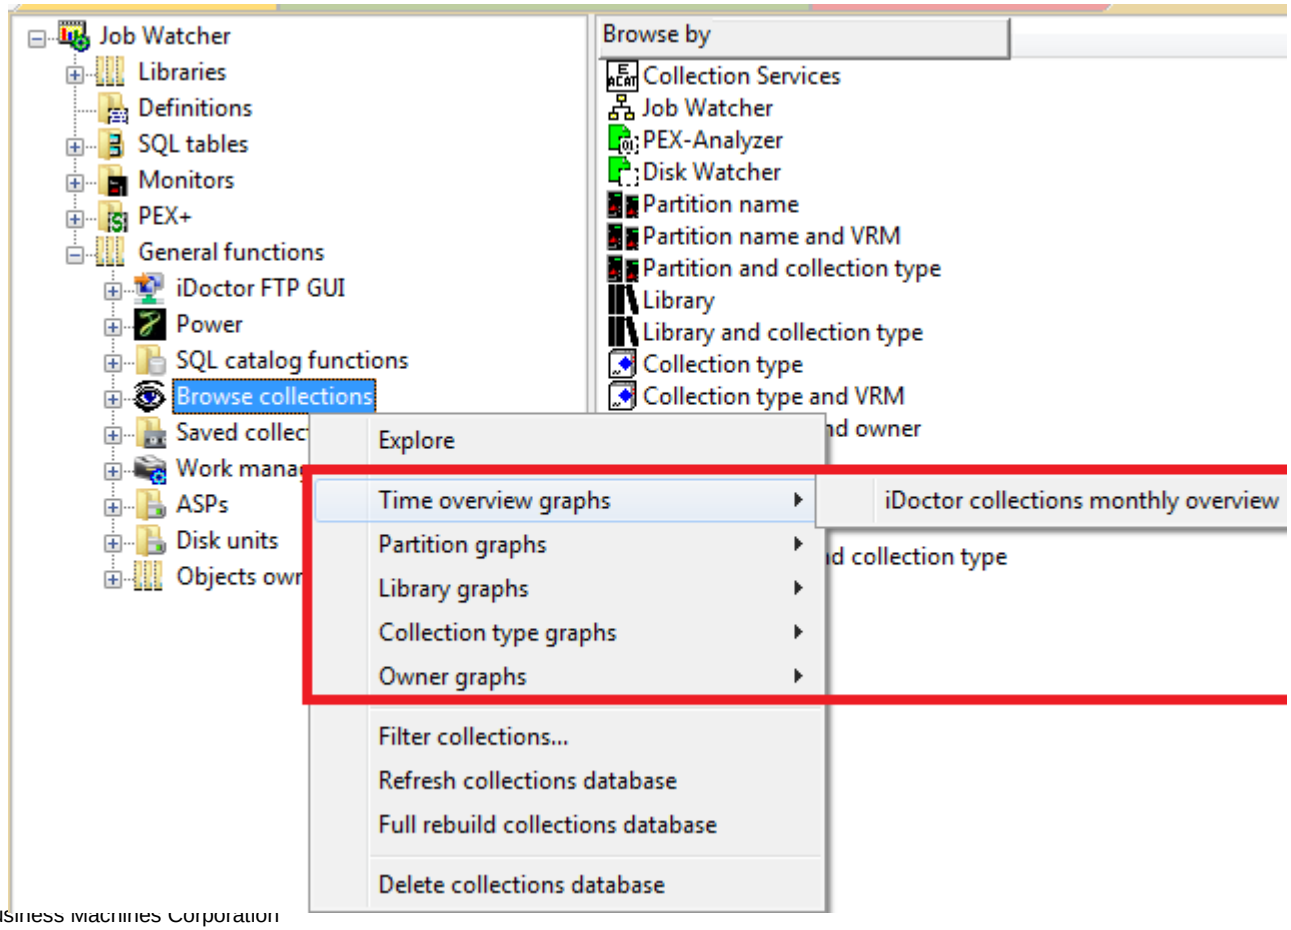

Collection Repository required for Collection Overview

- In CSI, JW the collection repository must exist in order for the Collection Overview graph to work.

- This provides drill downs into different components/collections where data is available for the same LPAR and time period you are viewing.

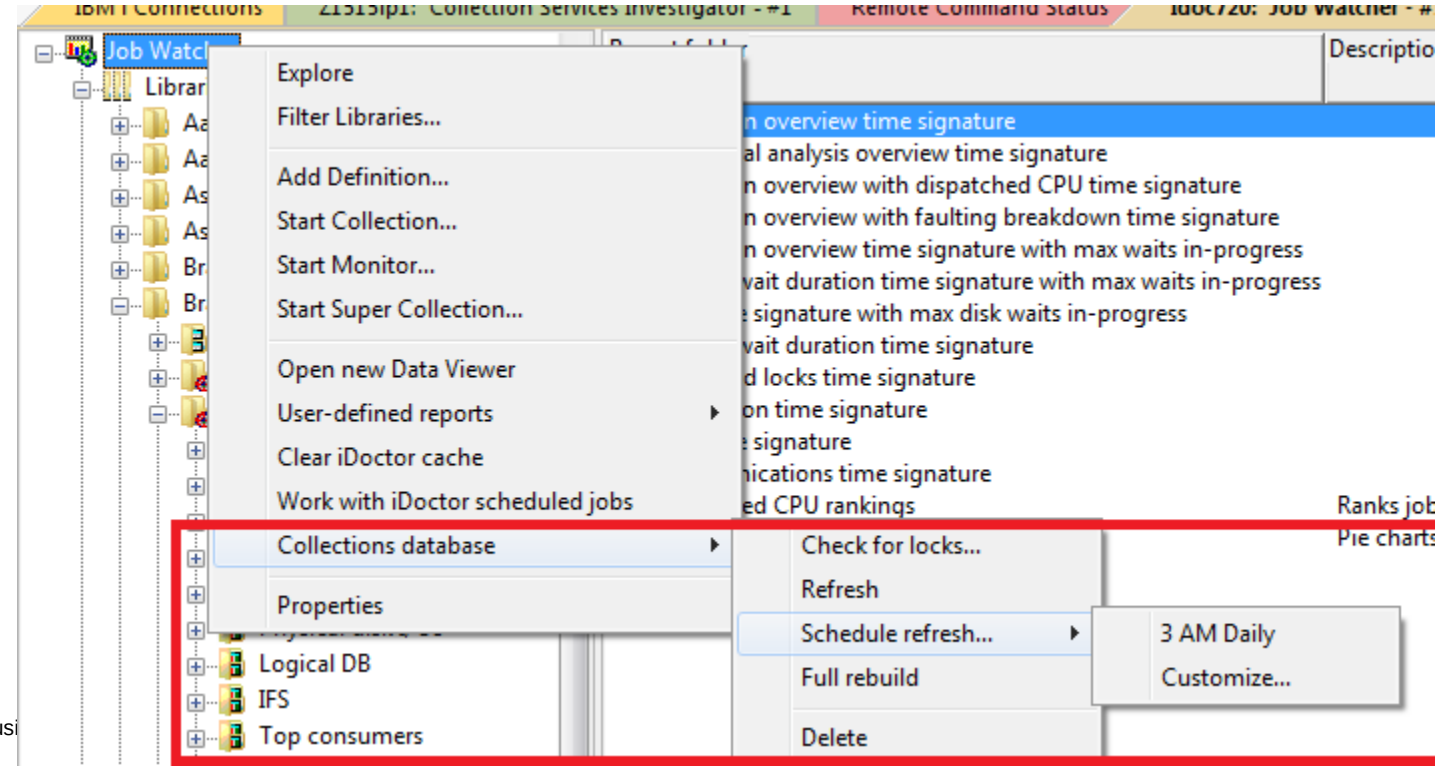

Collection Repository required for Collection Overview

- Example drill down from JW -> PEX:

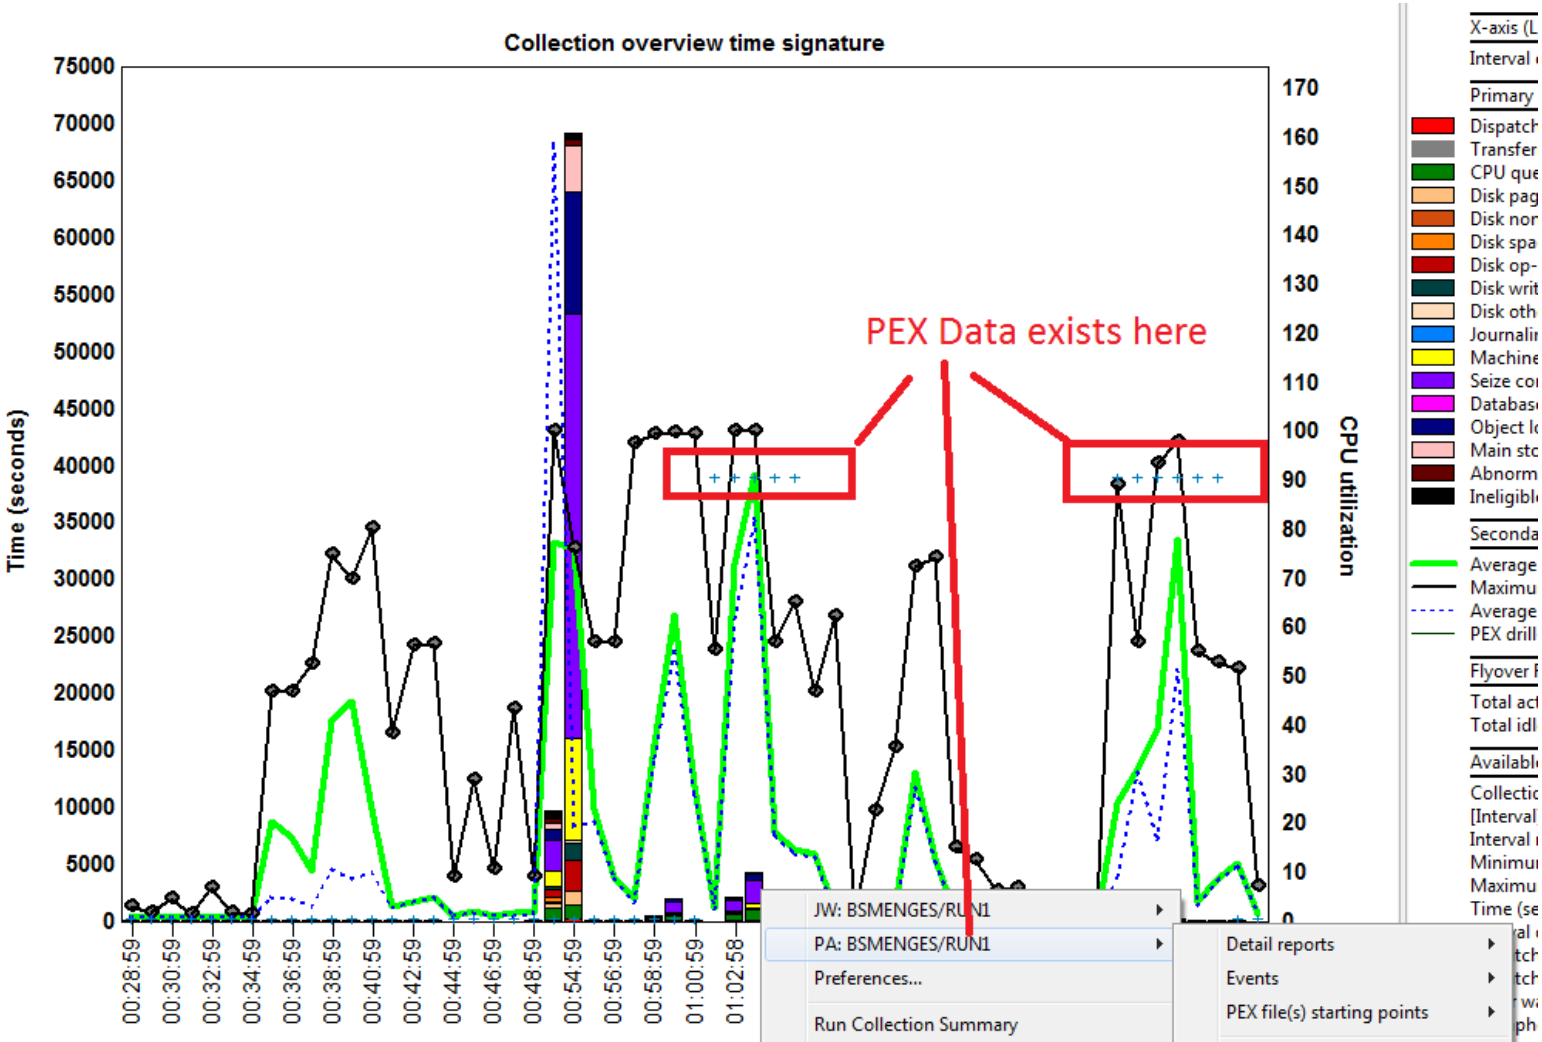

Collection Repository Graphs

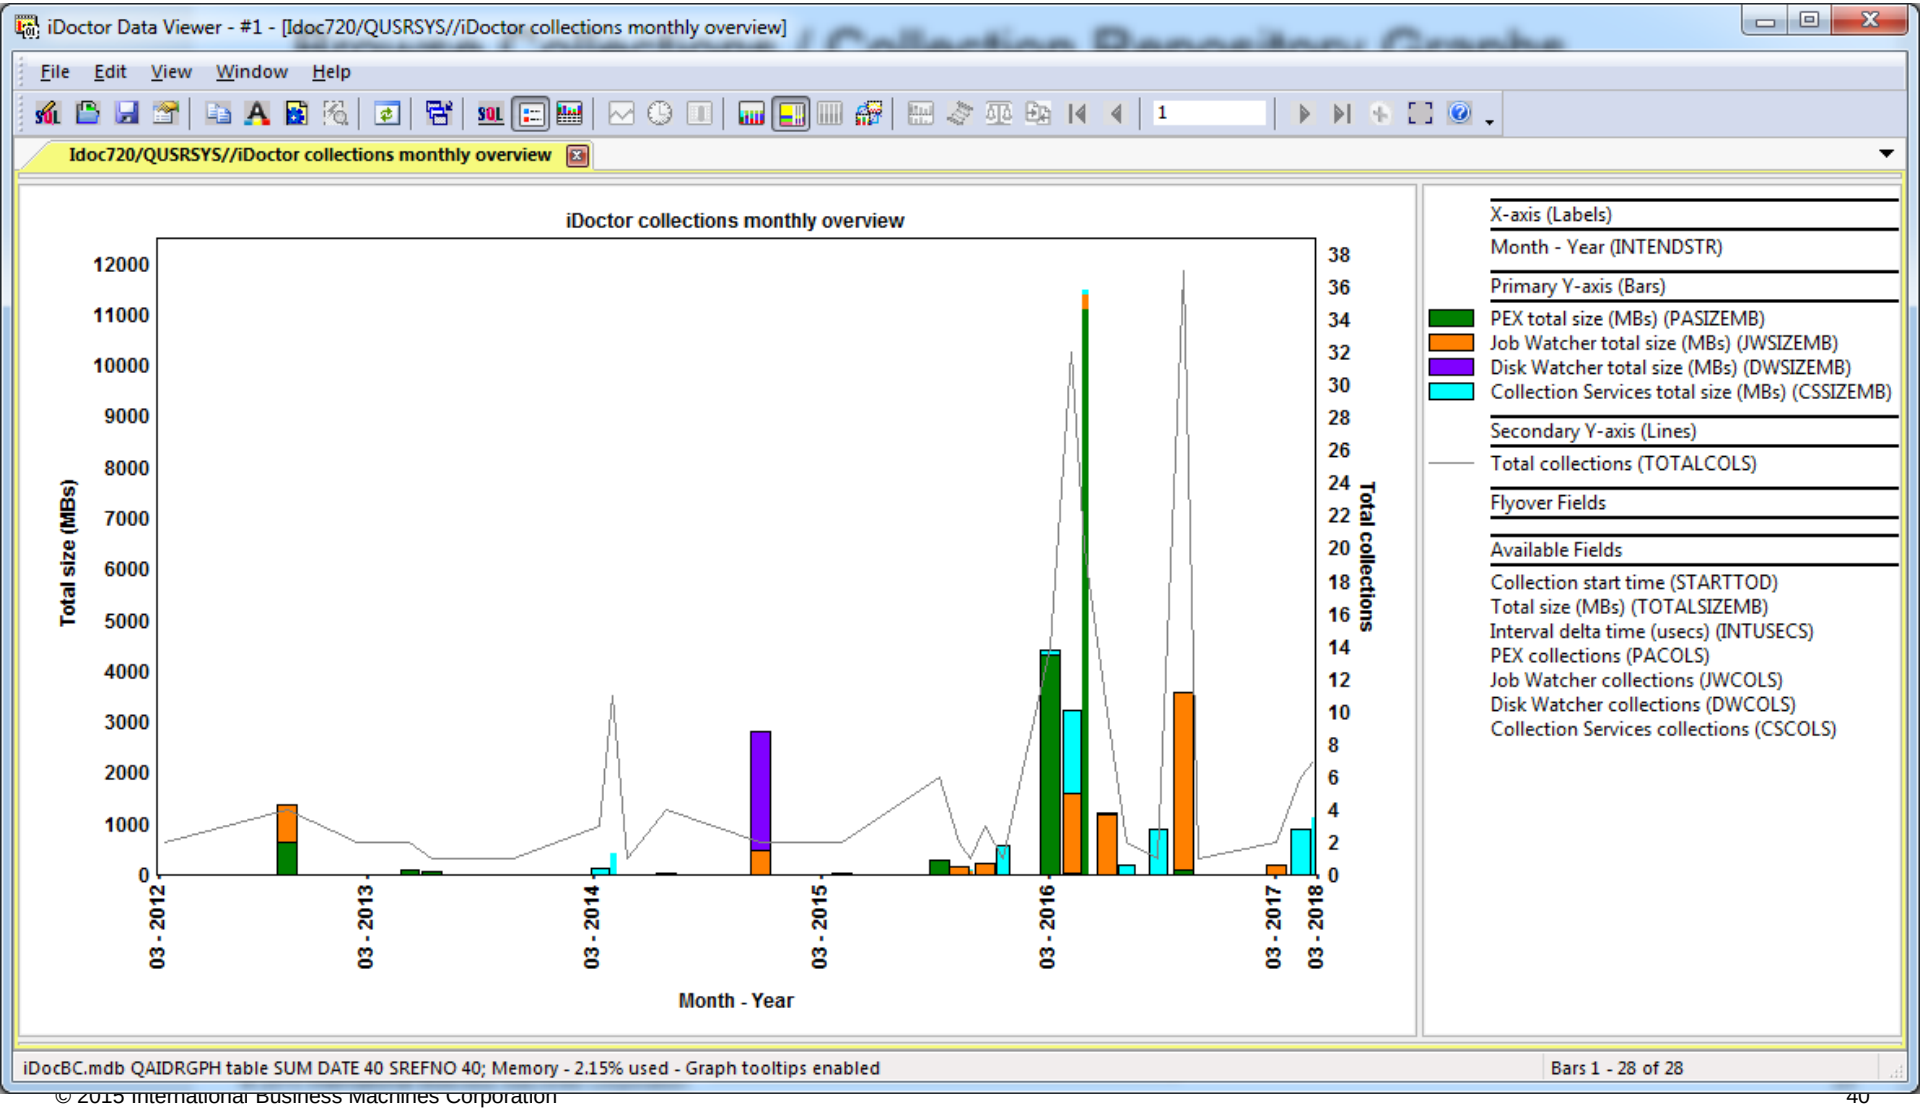

Collection Repository – Collecting Browsing Option

- Added a new preference on the Miscellaneous tab called "Use the collections database to see which collections exist" (default unchecked) If you check this option and the collection database has been created then you will experience much faster response times when listing collections in Job Watcher, PEX, Disk Watcher or CSI. Keep in mind however when using this option the collections repository DB must be recreated first (full rebuild option) , the server builds must be reinstalled to the latest levels and also the data is cached which in some cases will not reflect changes made to the data if they were done outside of the iDoctor GUI. Actions such as copying/deleting/renaming collections or running analyses will be updated automatically in the collections DB. However actions like copying entire libraries or individual file/members will not be immediately recognized and the collections repository must be refreshed in order to see changes like this.

Widgets

- Special identifying figures that can be added to any line in a graph.

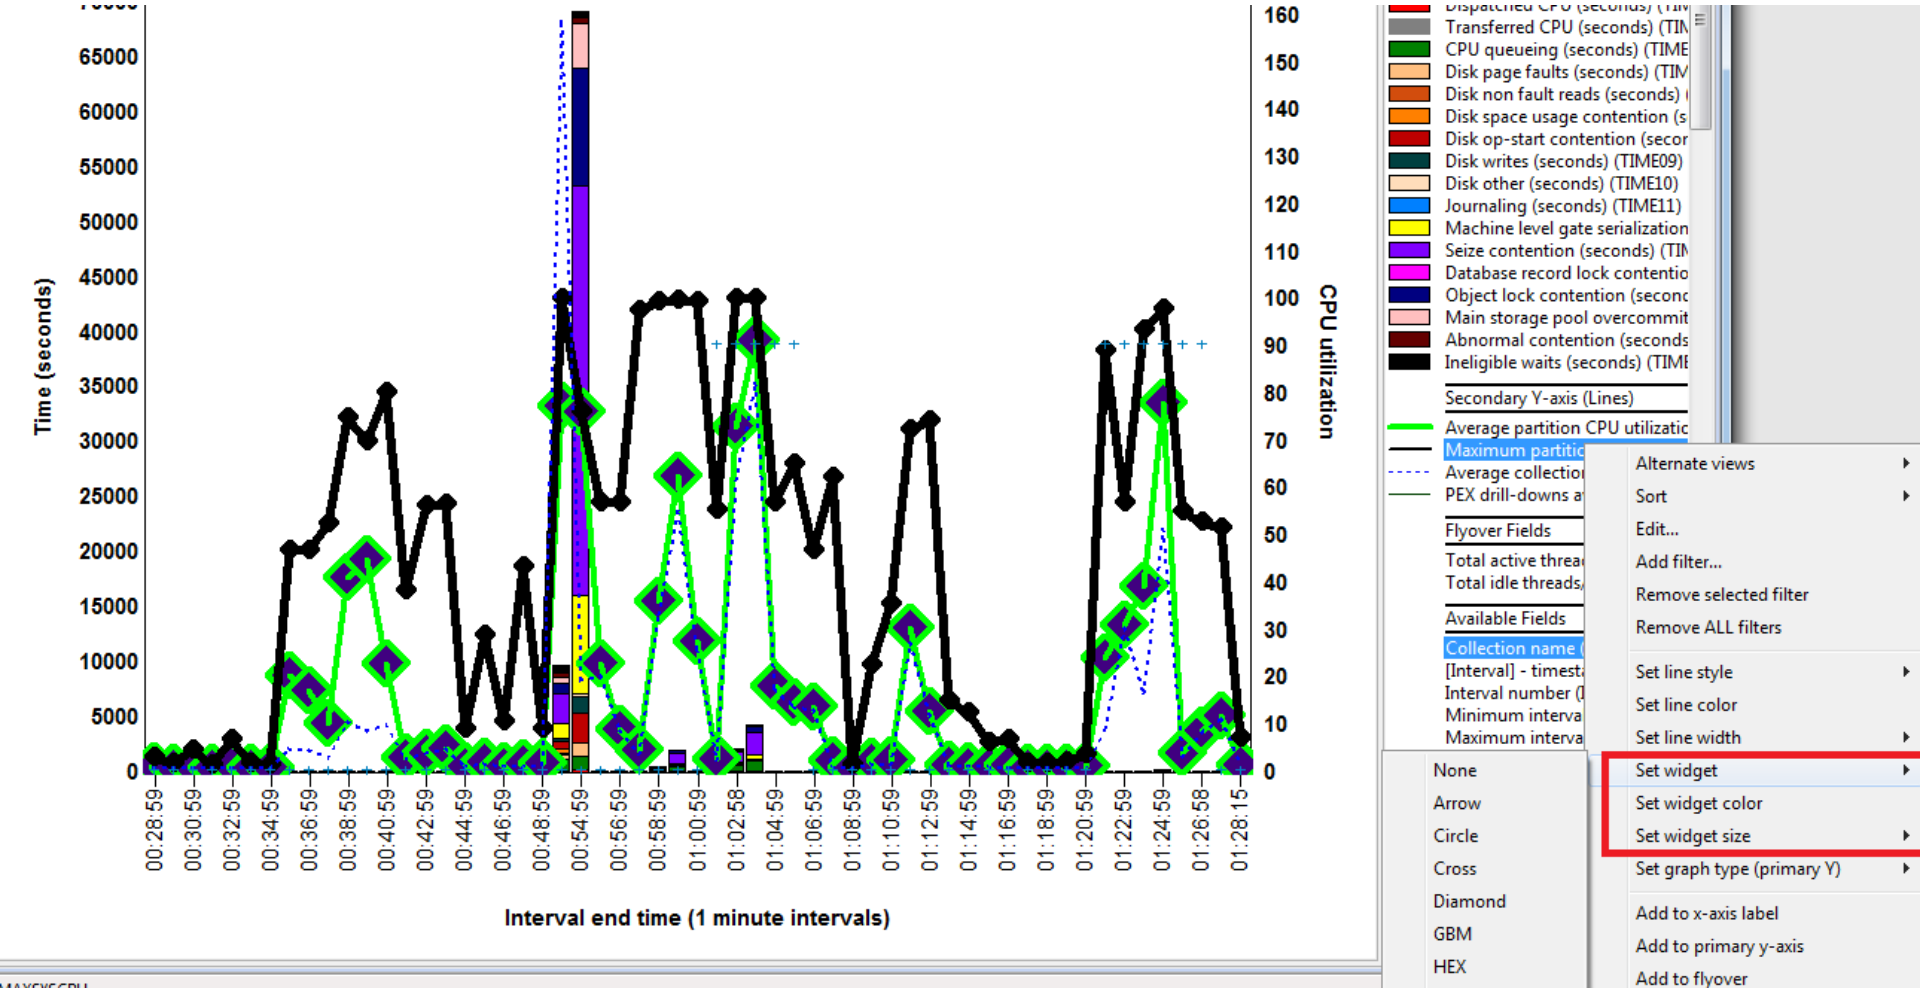

IBM i command interfaces

- Added free functions for WRKOBJ, WRKOBJLCK, WRKLNK, WRKACTJOB, WRKUSRPRF and a table search function

- No access code required!

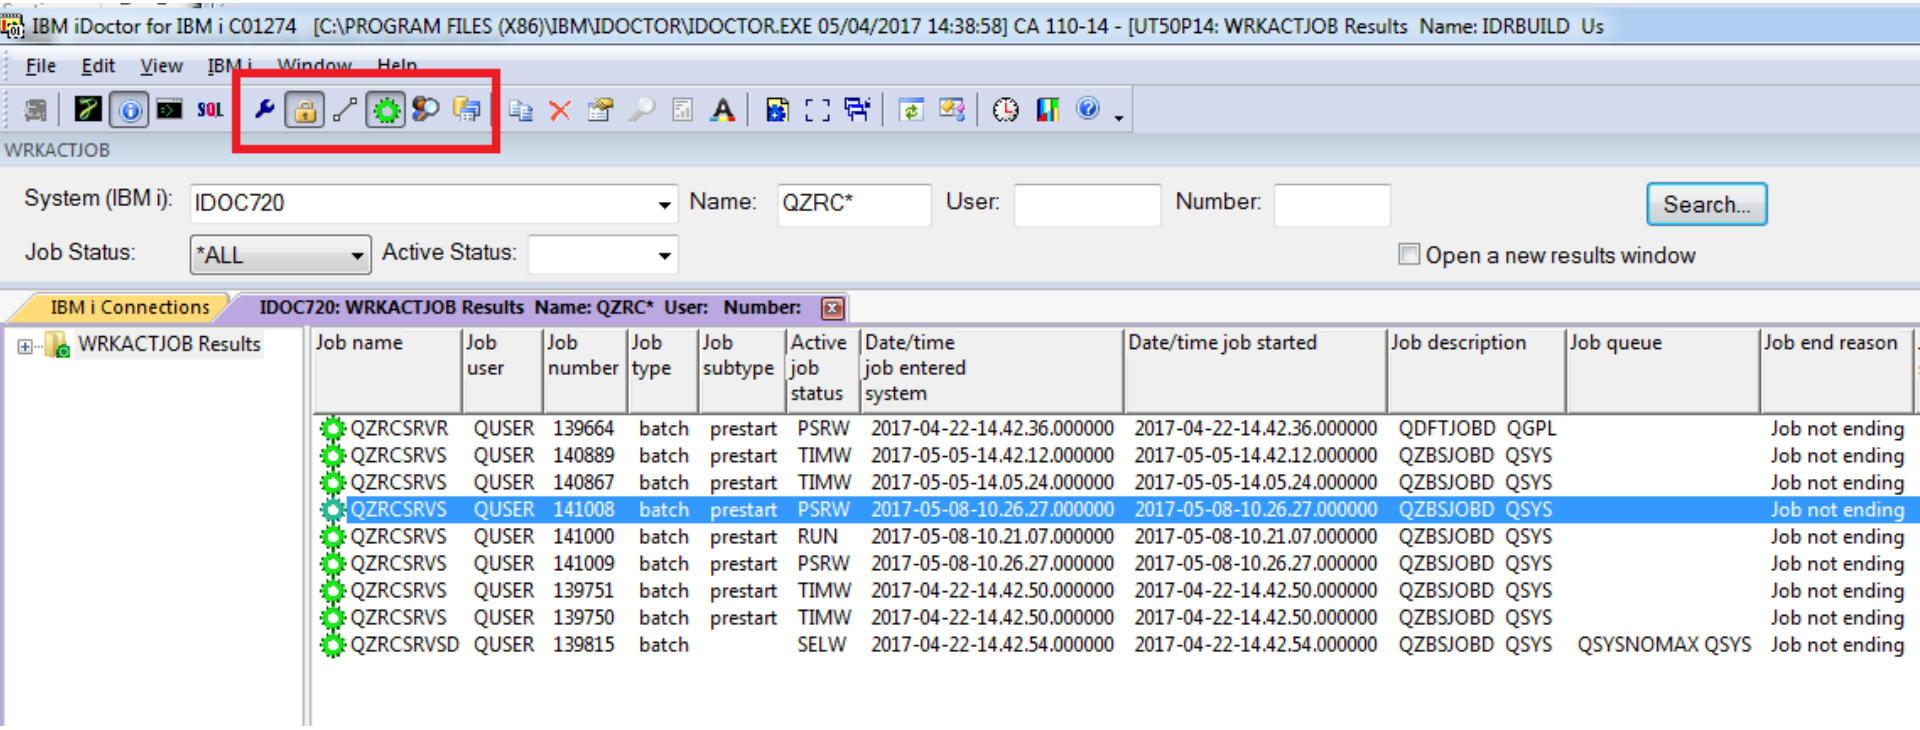

Problem reporting interface

- Use Help -> About -> Report a problem button to report an iDoctor problem to us. Thanks

- Note: A history of installed iDoctor build versions on the PC is also show on the Help -> About window within the text box (scroll down.)

iDoctor Resources

- iDoctor e-mail list: usage tips, build updates, PTF info, etc.
  - Send requests to mccargar@us.ibm.com
- iDoctor Website:
  - https://www-03.ibm.com/systems/services/labservices/iDoctor/
- Documentation:
  - http://public.dhe.ibm.com/services/us/igsc/idoctor/iDoctorV73.pdf
- YouTube Channel is currently unavailable.