

What's New with iDoctor Oct 2013-Mar 2015 (iDoctor GUI Builds 1040-1140)

Ron McCargar
IBM i Global Support Center



Power is performance redefined

Deliver services faster, with higher quality
and superior economics



Agenda

- **Overview**

- Goals, components, unique features and website.

- **What's new with iDoctor (October 2013 – March 2015):**

- **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 8 external and several additional 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) <u>iDoctor – VIOS Investigator</u> IBM i Disk Watcher (PT1) / STRDW
Plan cache analysis	<u>iDoctor – Plan Cache Analyzer</u>
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](https://www-912.ibm.com/i_dir/idoctor.nsf/F204DE4F34767E0686256F4000757A90/$FILE/iDoctorV7R1.pdf)



What's New with iDoctor: Overview

- **General GUI Enhancements**
- **Collection Services Investigator**
- **Job Watcher**
- **PEX Analyzer**



General – Support for tabbed style MDI

The latest builds now feature a tabbed-style user interface with 3 options to choose from:

Grouped – Tabs can be grouped together into 1 or more groups in order to make comparisons.

Note: Views cannot be tiled or cascaded when using this style.

Standard – This style provides tabs for each view and still allows for cascading and tiling of views.

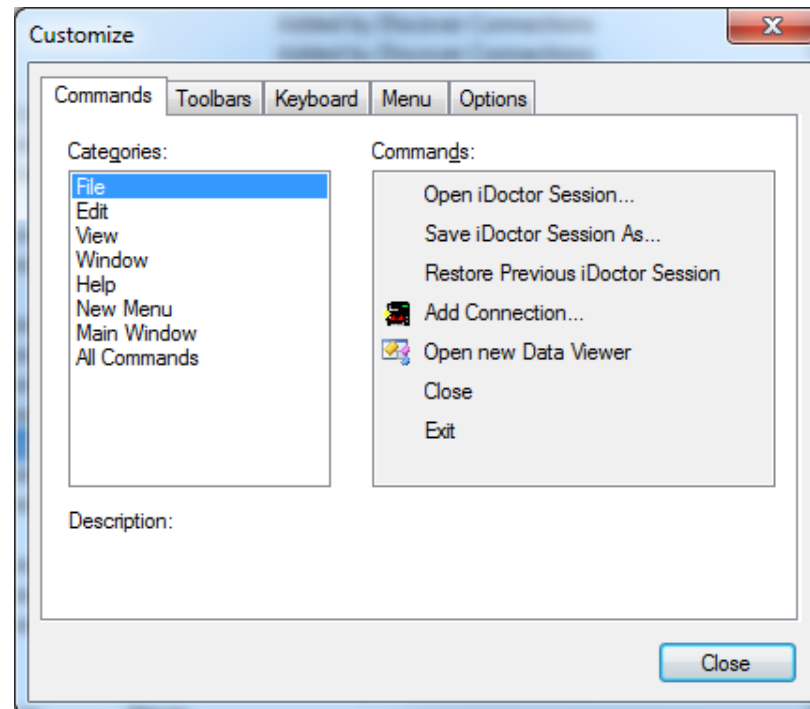
None – This is the same as the original iDoctor MDI GUI design where no tabs are used.

See the [View -> MDI Tabbed Style](#) menu to modify the current style used.

General – Menu and toolbar customizations

Users can now customize the menus and toolbars shown on the Main Window and Data Viewer.

Use the View -> Customize menu or for toolbars you can also click the drop down arrow at the far right corner of the toolbar to add/remove buttons from the toolbar.





General – Full screen mode

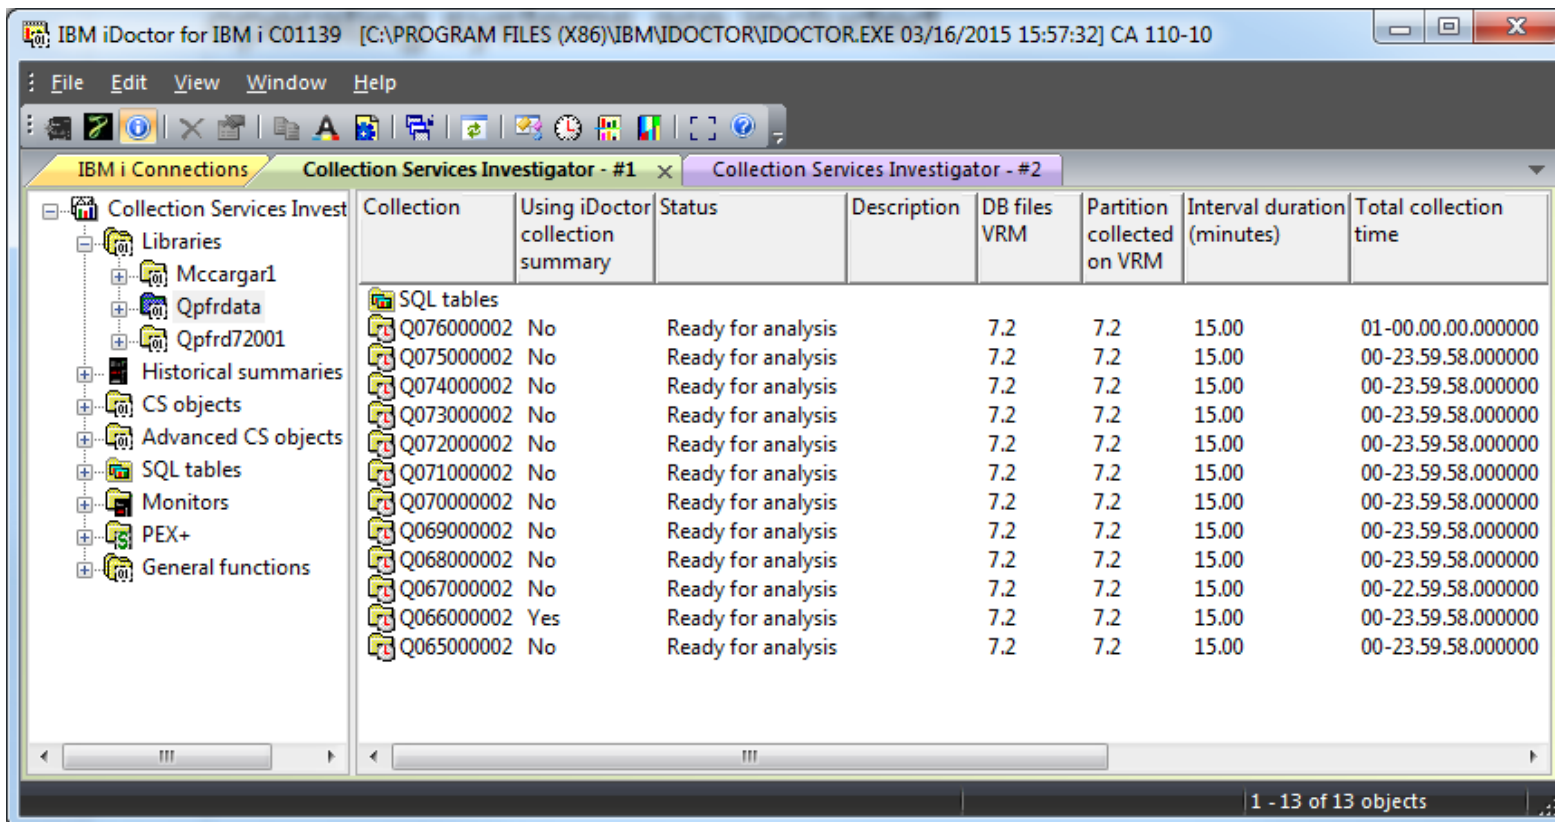
This mode lets you view the current Main Window or Data Viewer contents full screen with no toolbars and menus.

Use the View -> Full Screen menu to activate this.

General – Application Look

This option allows you to quickly change the look of the entire application by using the View -> Application Look menu.

Several Windows color themes and styles from past Windows operating systems are included.



The screenshot shows the IBM iDoctor application window for IBM i C01139. The main window is titled "Collection Services Investigator - #2" and displays a table of collection data. The table has the following columns: Collection, Using iDoctor collection summary, Status, Description, DB files VRM, Partition collected on VRM, Interval duration (minutes), and Total collection time. The table lists 13 objects, all of which are "Ready for analysis" and have an interval duration of 15.00 minutes. The total collection time for each object is 00-23.59.58.000000, except for the first object which is 01-00.00.00.000000.

Collection	Using iDoctor collection summary	Status	Description	DB files VRM	Partition collected on VRM	Interval duration (minutes)	Total collection time
SQL tables							
Q076000002	No	Ready for analysis		7.2	7.2	15.00	01-00.00.00.000000
Q075000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q074000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q073000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q072000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q071000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q070000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q069000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q068000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q067000002	No	Ready for analysis		7.2	7.2	15.00	00-22.59.58.000000
Q066000002	Yes	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000
Q065000002	No	Ready for analysis		7.2	7.2	15.00	00-23.59.58.000000



General – iDoctor no longer requires admin level authority

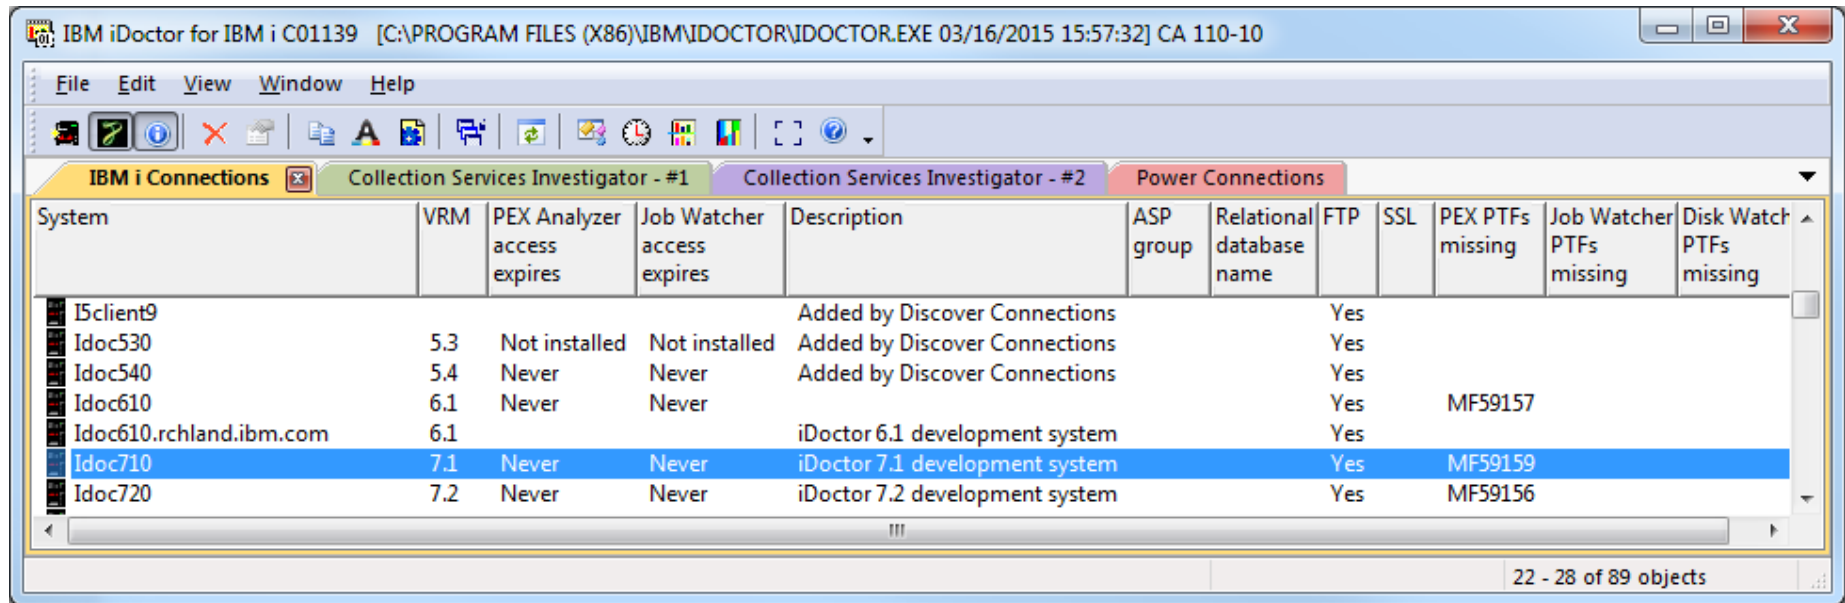
Changes have been made so that iDoctor no longer requires admin level authority in order to be used.

These changes should allow users to run iDoctor on a citrix server without complications.

Note: The check for new server builds button on the iDoctor components screen still requires it because it will run the installation in the background which requires it.

General – “My Connections” is now IBM i Connections

The My Connections list has been divided into an IBM i Connections interface as well as a Power Connections interface for all of your non IBM i connections.



IBM iDoctor for IBM i C01139 [C:\PROGRAM FILES (X86)\IBM\IDOCTOR\IDOCTOR.EXE 03/16/2015 15:57:32] CA 110-10

File Edit View Window Help

IBM i Connections Collection Services Investigator - #1 Collection Services Investigator - #2 Power Connections

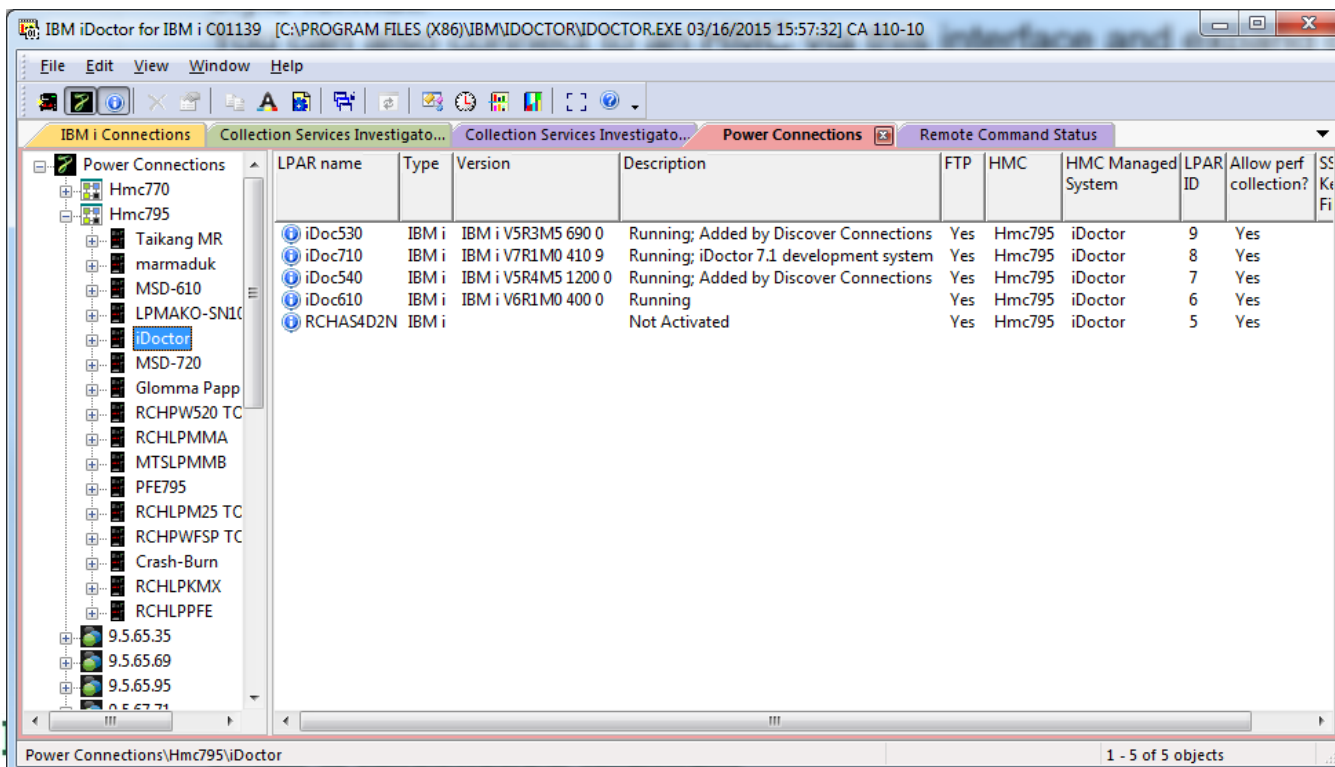
System	VRM	PEX Analyzer access expires	Job Watcher access expires	Description	ASP group	Relational database name	FTP	SSL	PEX PTFs missing	Job Watcher PTFs missing	Disk Watch PTFs missing
I5client9				Added by Discover Connections			Yes				
Idoc530	5.3	Not installed	Not installed	Added by Discover Connections			Yes				
Idoc540	5.4	Never	Never	Added by Discover Connections			Yes				
Idoc610	6.1	Never	Never				Yes		MF59157		
Idoc610.rchland.ibm.com	6.1			iDoctor 6.1 development system			Yes				
Idoc710	7.1	Never	Never	iDoctor 7.1 development system			Yes		MF59159		
Idoc720	7.2	Never	Never	iDoctor 7.2 development system			Yes		MF59156		

22 - 28 of 89 objects

General – Power Connections

The Power Connections interface allows users to work with their non IBM i connections (HMCs, VIOS, AIX, etc) with an expandable tree/list style format.

You can also connect to an HMC via this interface and expand down to the managed system and LPAR levels. In cases where the LPAR has iDoctor installed/available, additional functionality is easily available to allow you to work with performance data on desired LPARs.



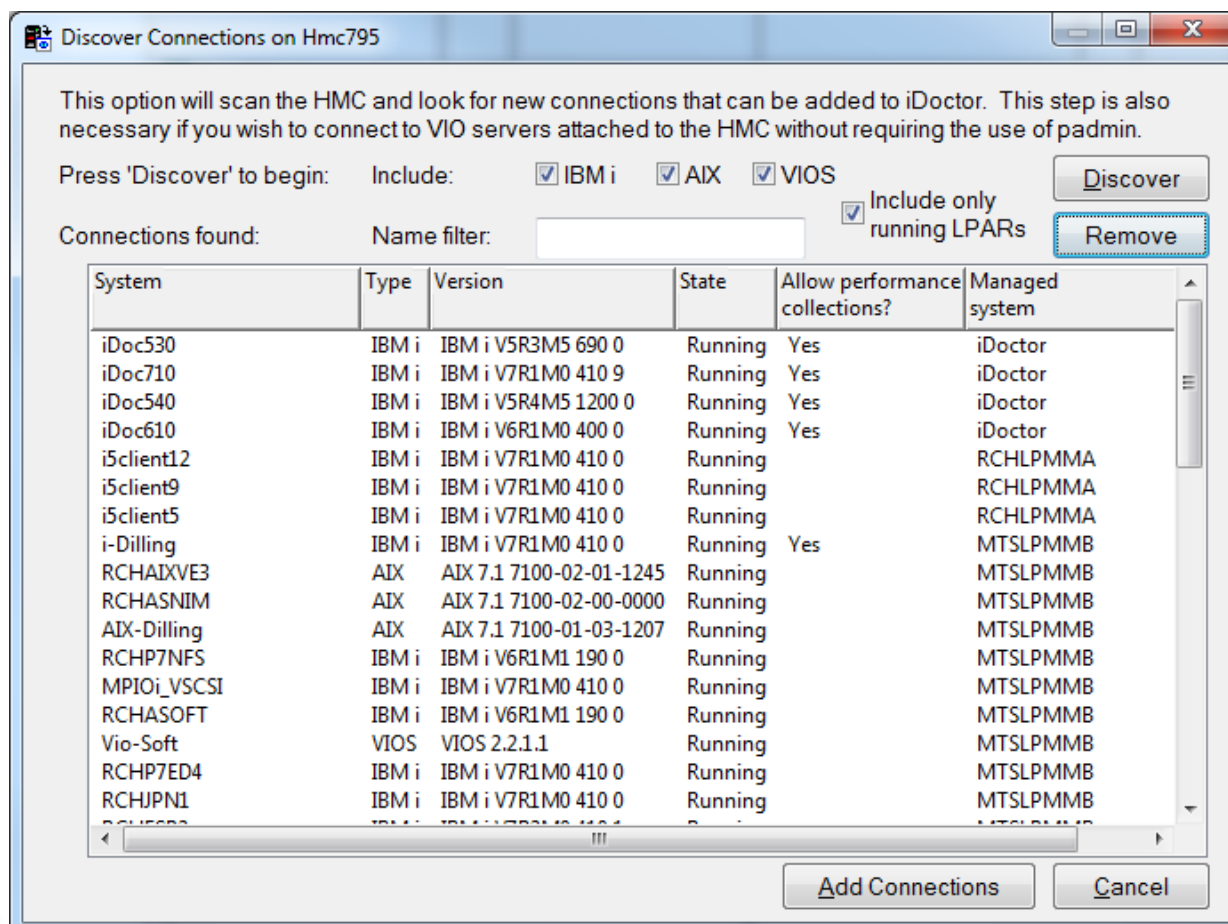
The screenshot shows the IBM iDoctor for IBM i software interface. The window title is "IBM iDoctor for IBM i C01139 [C:\PROGRAM FILES (X86)\IBM\IDOCTOR\IDOCTOR.EXE 03/16/2015 15:57:32] CA 110-10". The interface has a menu bar (File, Edit, View, Window, Help) and a toolbar. The main area is divided into a tree view on the left and a table on the right. The tree view shows a hierarchy of connections, including "Power Connections", "Hmc770", "Hmc795", and various LPARs like "Taikang MR", "marmaduk", "MSD-610", "LPMMAKO-SN10", "iDoctor", "MSD-720", "Glomma Papp", "RCHPW520 TC", "RCHLPMMA", "MTSLPMMB", "PFE795", "RCHLPM25 TC", "RCHPW520 TC", "Crash-Burn", "RCHLPKMX", "RCHLPPFE", and several versions of "9.5.65". The table on the right lists the LPARs with columns for LPAR name, Type, Version, Description, FTP, HMC, HMC Managed System, LPAR ID, Allow perf collection?, and SS. The table data is as follows:

LPAR name	Type	Version	Description	FTP	HMC	HMC Managed System	LPAR ID	Allow perf collection?	SS
iDoc530	IBM i	IBM i V5R3M5 690 0	Running; Added by Discover Connections	Yes	Hmc795	iDoctor	9	Yes	
iDoc710	IBM i	IBM i V7R1M0 410 9	Running; iDoctor 7.1 development system	Yes	Hmc795	iDoctor	8	Yes	
iDoc540	IBM i	IBM i V5R4M5 1200 0	Running; Added by Discover Connections	Yes	Hmc795	iDoctor	7	Yes	
iDoc610	IBM i	IBM i V6R1M0 400 0	Running	Yes	Hmc795	iDoctor	6	Yes	
RCHAS4D2N	IBM i		Not Activated	Yes	Hmc795	iDoctor	5	Yes	

The status bar at the bottom of the window shows "Power Connections\Hmc795\iDoctor" and "1 - 5 of 5 objects".

General – Power Connections – HMC Discover Connections

When connected to an HMC you can use the Discover Connections option to quickly find and add connections to your Power Connections and/or IBM i connections interfaces in iDoctor.





General – Sessions

With the latest builds the GUI will keep track of each graph and table opened or closed and save to a 'history' folder the state of your current session. This includes most interfaces used on the Main Window and all your graphs/tables on all Data Viewers opened.

From the File menu you can use the following new options to work with sessions:

File -> Open iDoctor Session...

File -> Save iDoctor Session As...

Note: This option lets you pick the file name and where to save it. Your current iDoctor session is now ALWAYS saved to the most recent file in the history folder.

File -> Restore Previous iDoctor Session

This option will open the 2nd most recent session from the history folder. Use this if you have crashed, etc and need to restore your previous session.



General – Shortcut support removed

The ability to create shortcuts for individual reports/graphs in the Data Viewer has been removed.

Instead users can save their current session (main window and all data viewers to an .idr file.)



General – Keep connections alive preference

On the Preferences -> Misc tab added the following new option to keep iDoctor connections alive on networks that drop inactive connections:

- Keep connections alive by sending a request every N minutes when GUI is idle (must restart GUI after changing)

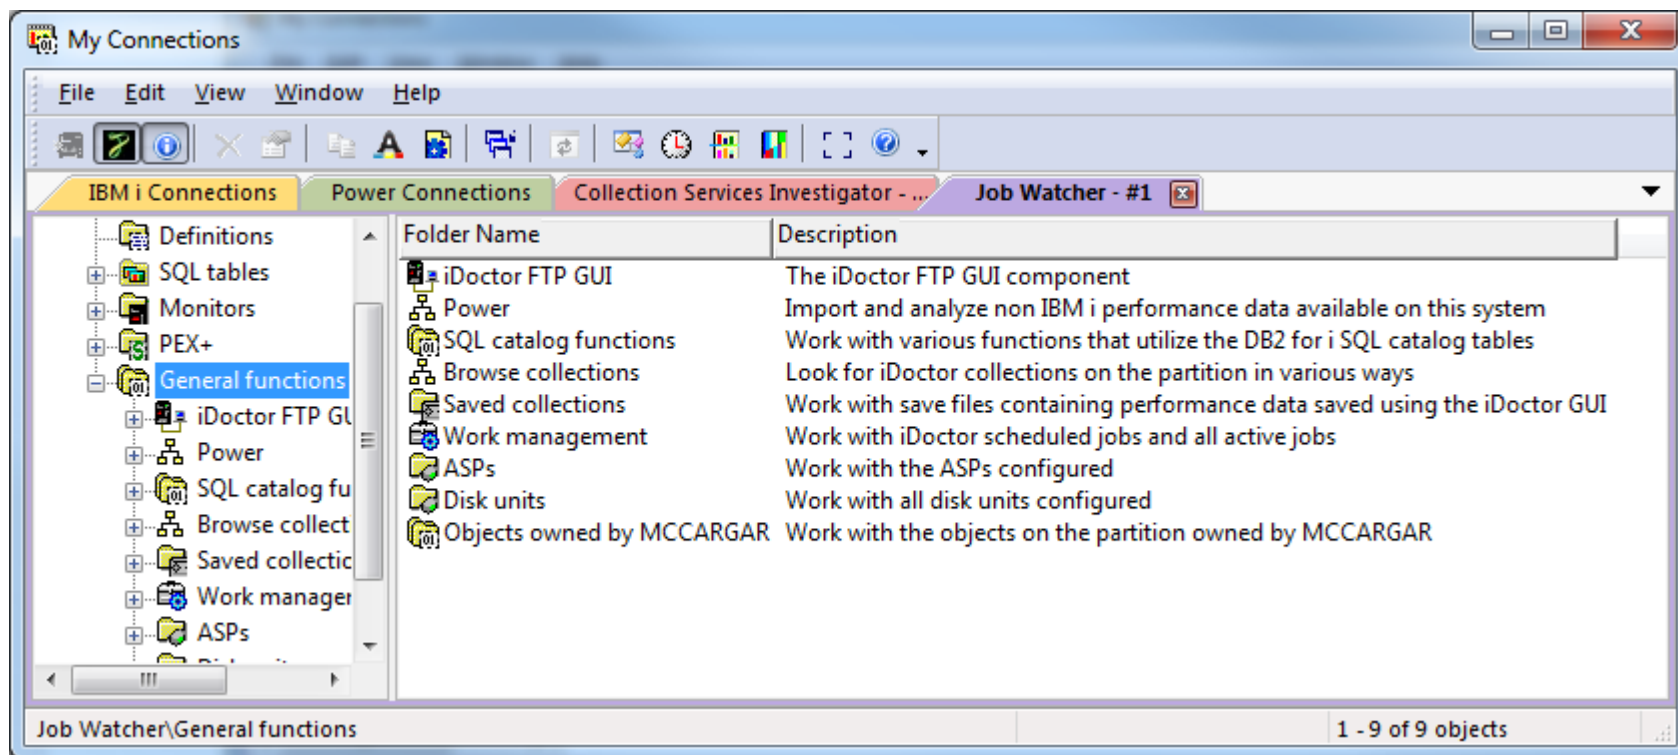
N is specified in the text box next to this option.

This option is unchecked by default so if you want to use this you must check the box, confirm or change the value for N then restart the GUI.

General – New General Functions folder

Several miscellaneous functions provided by iDoctor have been grouped together into a new General Functions folder found under the component you are working with. These functions are general use functions that apply to any of the IBM i iDoctor components.

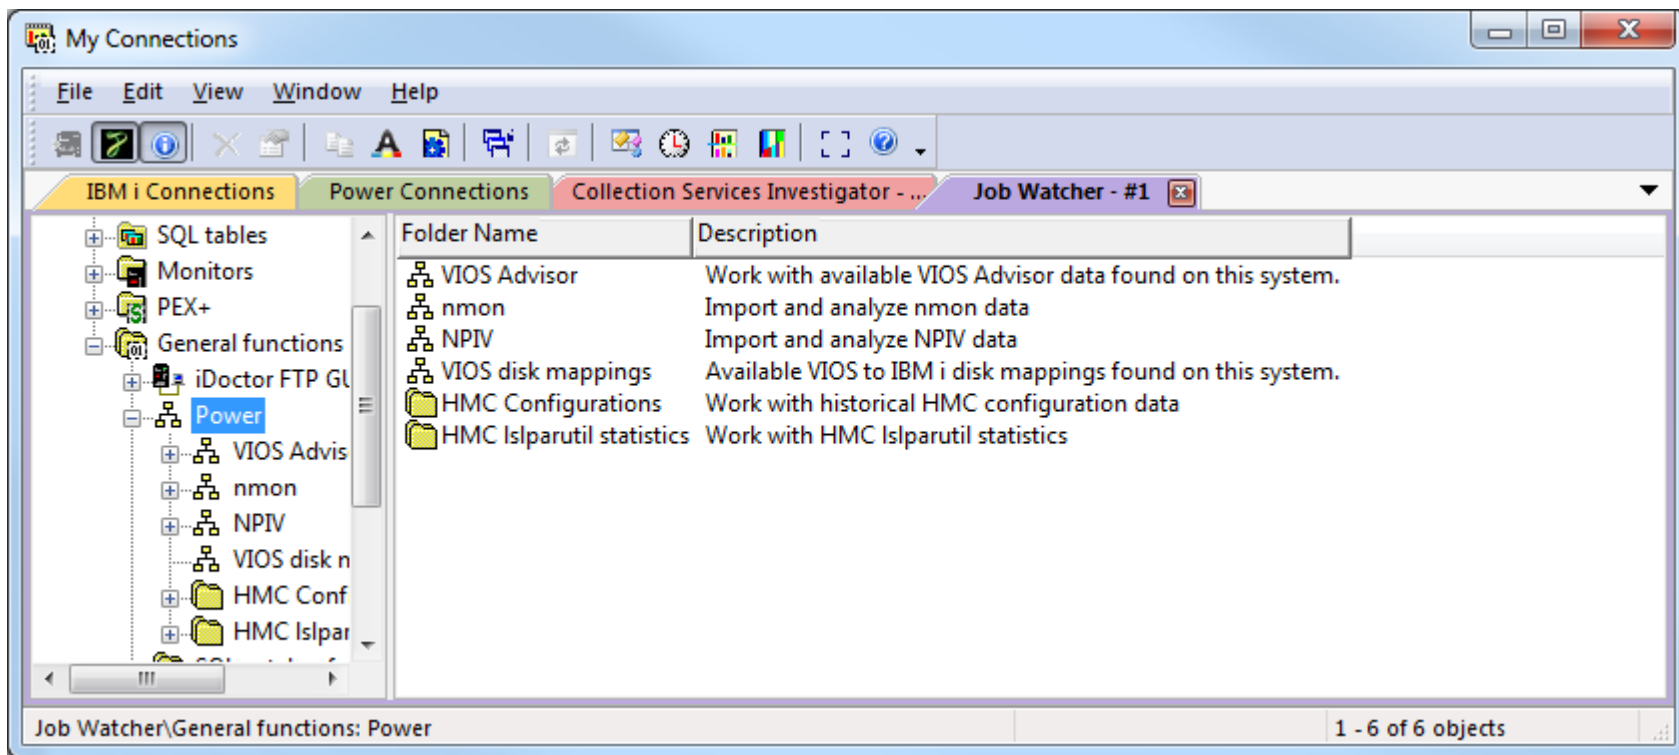
This folder currently contains the following:



Folder Name	Description
iDoctor FTP GUI	The iDoctor FTP GUI component
Power	Import and analyze non IBM i performance data available on this system
SQL catalog functions	Work with various functions that utilize the DB2 for i SQL catalog tables
Browse collections	Look for iDoctor collections on the partition in various ways
Saved collections	Work with save files containing performance data saved using the iDoctor GUI
Work management	Work with iDoctor scheduled jobs and all active jobs
ASPs	Work with the ASPs configured
Disk units	Work with all disk units configured
Objects owned by MCCARGAR	Work with the objects on the partition owned by MCCARGAR

General – Power folder

The Power folder allows users to work with nmon, npiv and other types of non IBM i performance data either by importing it from the IFS of the IBM i you are working with or by uploading files from your PC to the IBM i for analysis.



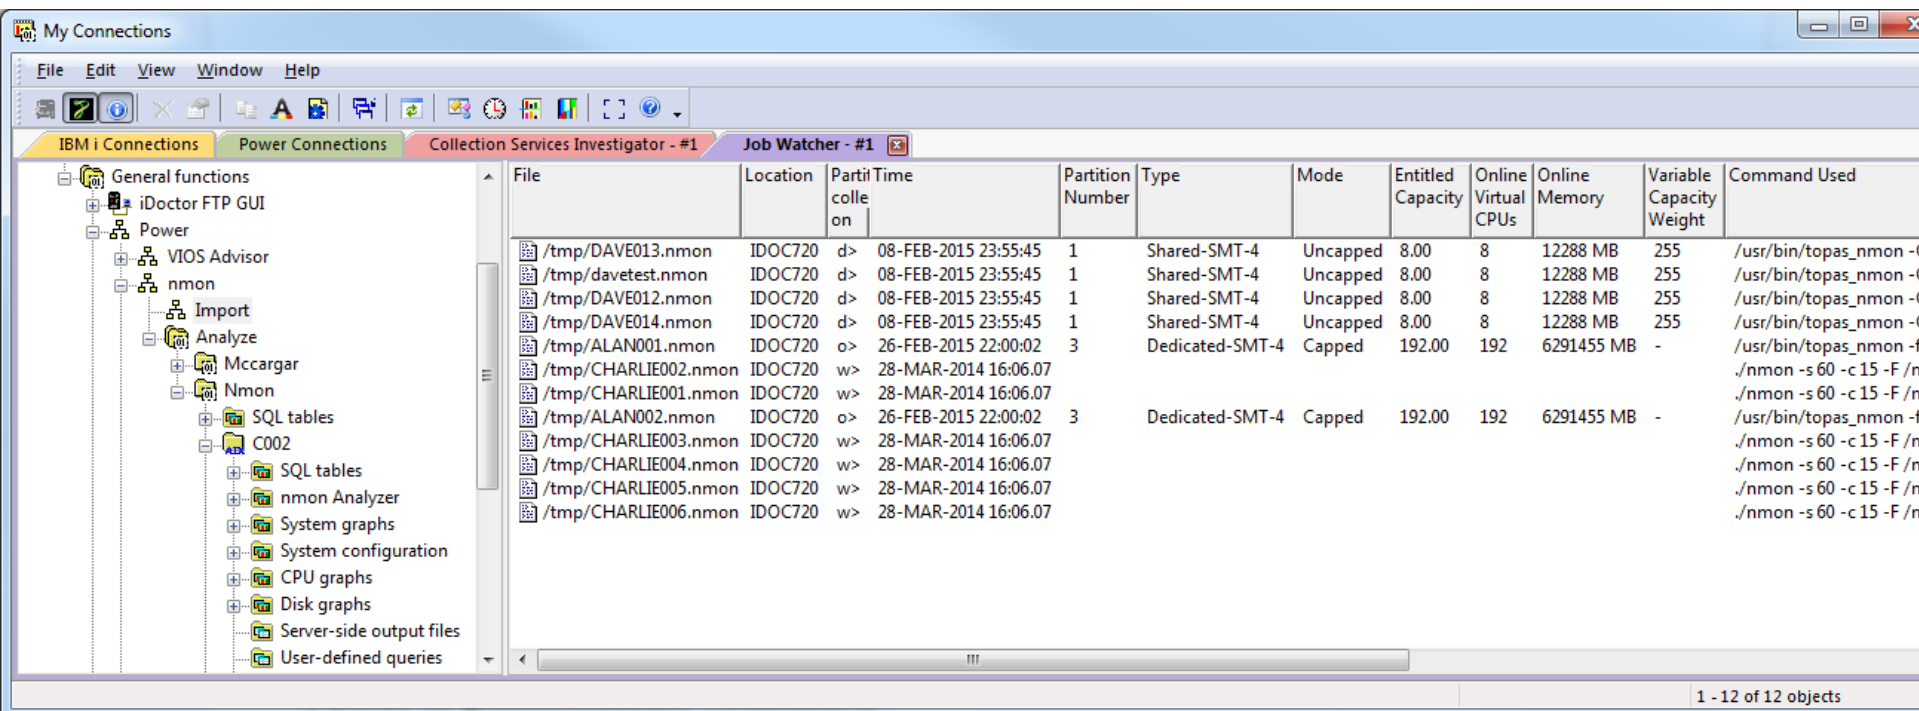
The screenshot shows the IBM i Connections interface. The left pane displays a tree view of folders under 'General functions', with the 'Power' folder selected. The right pane shows a table of folders and their descriptions.

Folder Name	Description
VIOS Advisor	Work with available VIOS Advisor data found on this system.
nmon	Import and analyze nmon data
NPiV	Import and analyze NPiV data
VIOS disk mappings	Available VIOS to IBM i disk mappings found on this system.
HMC Configurations	Work with historical HMC configuration data
HMC Isparutil statistics	Work with HMC Isparutil statistics

Job Watcher/General functions: Power 1 - 6 of 6 objects

General – Power – nmon import

The nmon import/analysis mechanism has been enhanced to handle much larger number of characters per line and rows and much greater numbers of columns and disks.



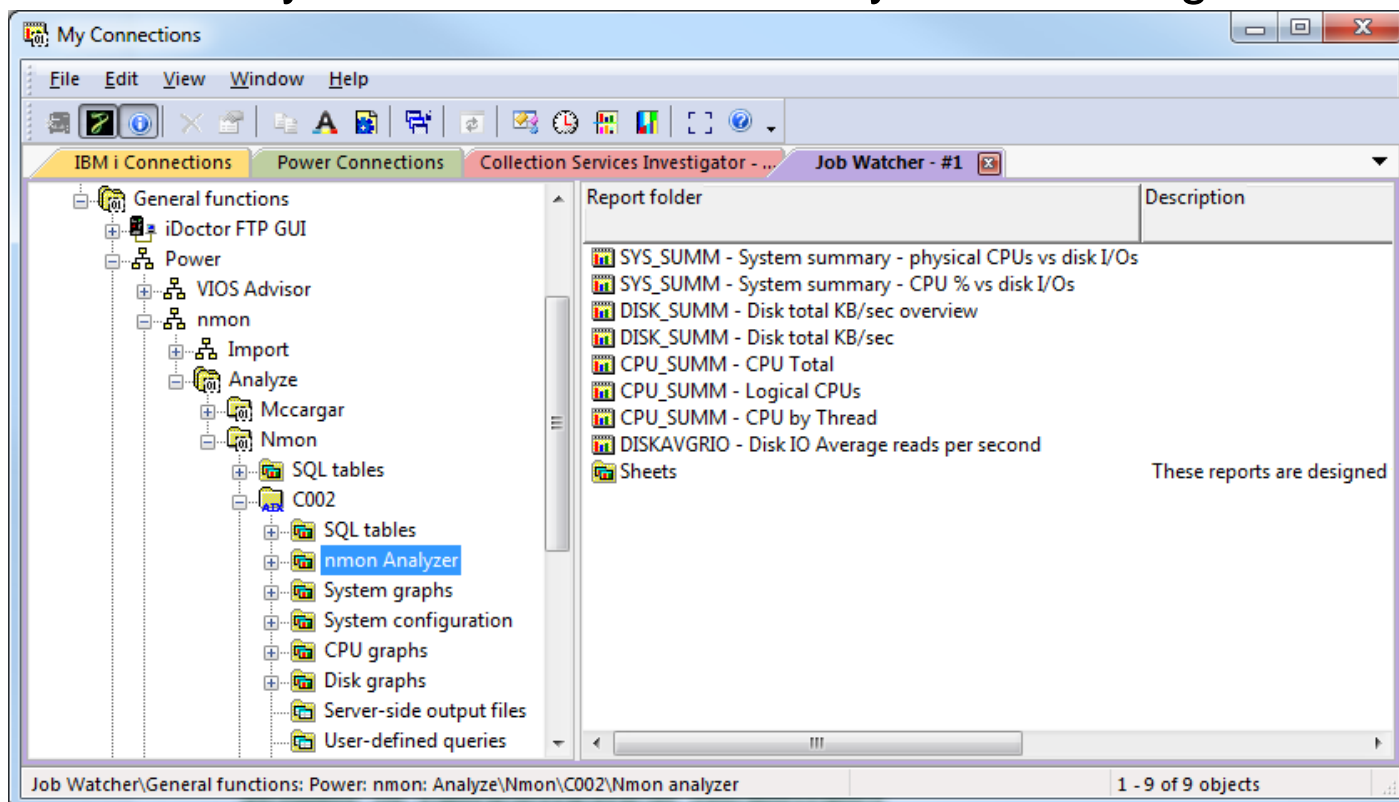
The screenshot shows the IBM i Connections Job Watcher interface. The left sidebar displays a tree view with 'Power' > 'nmon' > 'Import' selected. The main pane shows a table of job entries with the following columns: File, Location, Partition, Time collected, Partition Number, Type, Mode, Entitled Capacity, Online Virtual CPUs, Online Memory, Variable Capacity Weight, and Command Used.

File	Location	Partition	Time collected	Partition Number	Type	Mode	Entitled Capacity	Online Virtual CPUs	Online Memory	Variable Capacity Weight	Command Used
/tmp/DAVE013.nmon	IDOC720	d>	08-FEB-2015 23:55:45	1	Shared-SMT-4	Uncapped	8.00	8	12288 MB	255	/usr/bin/topas_nmon -f
/tmp/davetest.nmon	IDOC720	d>	08-FEB-2015 23:55:45	1	Shared-SMT-4	Uncapped	8.00	8	12288 MB	255	/usr/bin/topas_nmon -f
/tmp/DAVE012.nmon	IDOC720	d>	08-FEB-2015 23:55:45	1	Shared-SMT-4	Uncapped	8.00	8	12288 MB	255	/usr/bin/topas_nmon -f
/tmp/DAVE014.nmon	IDOC720	d>	08-FEB-2015 23:55:45	1	Shared-SMT-4	Uncapped	8.00	8	12288 MB	255	/usr/bin/topas_nmon -f
/tmp/ALAN001.nmon	IDOC720	o>	26-FEB-2015 22:00:02	3	Dedicated-SMT-4	Capped	192.00	192	6291455 MB	-	/usr/bin/topas_nmon -f
/tmp/CHARLIE002.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n
/tmp/CHARLIE001.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n
/tmp/ALAN002.nmon	IDOC720	o>	26-FEB-2015 22:00:02	3	Dedicated-SMT-4	Capped	192.00	192	6291455 MB	-	/usr/bin/topas_nmon -f
/tmp/CHARLIE003.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n
/tmp/CHARLIE004.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n
/tmp/CHARLIE005.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n
/tmp/CHARLIE006.nmon	IDOC720	w>	28-MAR-2014 16:06:07								./nmon -s 60 -c 15 -F /n

General – Power – nmon – Analyze folder

This folder lets you work with nmon data that is ready to be analyzed. Several additional graphs have been added; many are designed to look like the graphs found in the nmon Analyzer Excel spreadsheet for consistency reasons.

Note: nmon Analyzer is now maintained by Ron McCargar.





General – Graphing Preferences changes

The default X-axis and Y-axis graph font size is now 9 instead of 11.

The default max x-axis characters for a vertical bar graph is now 20 instead of 30.

Also made improvements to the graph auto resize fonts code.

General – GDI limit registry files

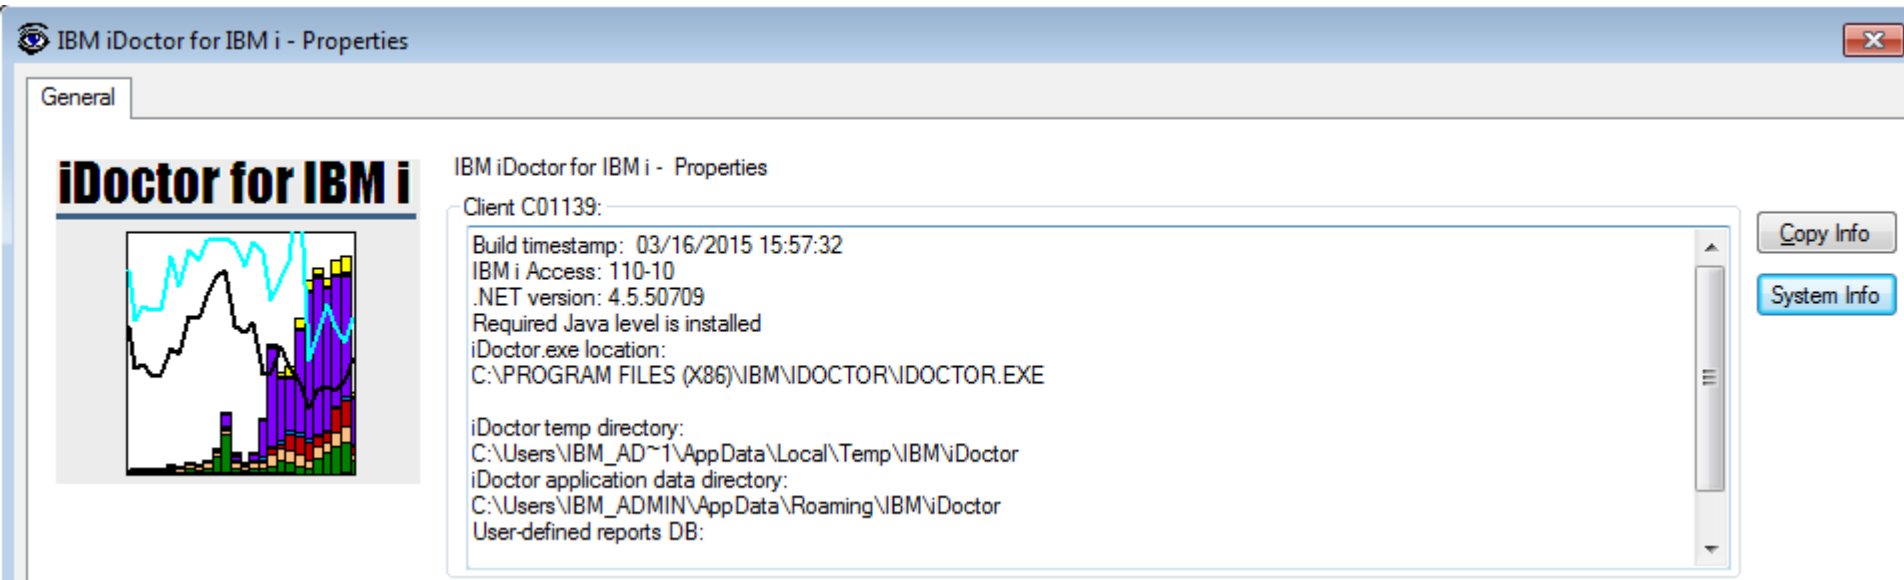
After installation in the iDoctor install directory made the following changes to address a limitation with the number of graphs you can open in iDoctor:

1. Added file SetGDILimitTo40K.reg which will set the GDIProcessHandleQuota setting in the Windows Registry under HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Windows to 40000. This will allow 4 times more graphs to be opened in iDoctor.
2. Added file SetGDILimitTo10K.reg which will set the GDIProcessHandleQuota back to the Windows 7 default value.

General – Application properties buttons

The Help -> About window (or Application Properties -> General tab) now has 2 new buttons:

1. The "Copy Info" button will copy the text on the screen to the clipboard as text (the Copy button at the bottom of the screen copies the data in RTF format which causes issues with the path names not being correct).
2. The "System Info" button will launch Window's System Information utility.



The screenshot shows the "IBM iDoctor for IBM i - Properties" dialog box with the "General" tab selected. On the left, there is a logo for "iDoctor for IBM i" and a line graph with a bar chart overlay. The main area displays system information for Client C01139:

```
Client C01139:
Build timestamp: 03/16/2015 15:57:32
IBM i Access: 110-10
.NET version: 4.5.50709
Required Java level is installed
iDoctor.exe location:
C:\PROGRAM FILES (X86)\IBM\IDOCTOR\IDOCTOR.EXE

iDoctor temp directory:
C:\Users\IBM_AD~1\AppData\Local\Temp\IBM\iDoctor
iDoctor application data directory:
C:\Users\IBM_ADMIN\AppData\Roaming\IBM\iDoctor
User-defined reports DB:
```

On the right side of the dialog, there are two buttons: "Copy Info" and "System Info".



General – Support for SSL connections

When editing your IBM i connection, there is now a checkbox Use SSL you can check if you need to make an SSL connection.

This currently will change your remote command/program call and ODBC connections to use SSL.

Note: Unfortunately support for secure FTP is not yet available in all areas of iDoctor so some GUI functions won't work yet on those environments.



General – Find window changes

In the Data Viewer the interface for doing a Find within table views has been simplified. The direction function and column drop down lists have been removed.

You can also now use F3 to find the next occurrence or Shift+F3 to find the previous occurrence after clicking back on the table. Note: These same options are also available in the Data Viewer's Edit menu.

You can also now use the Find window on lists within the Main Window. Use the Edit -> Find menu to bring up this window.



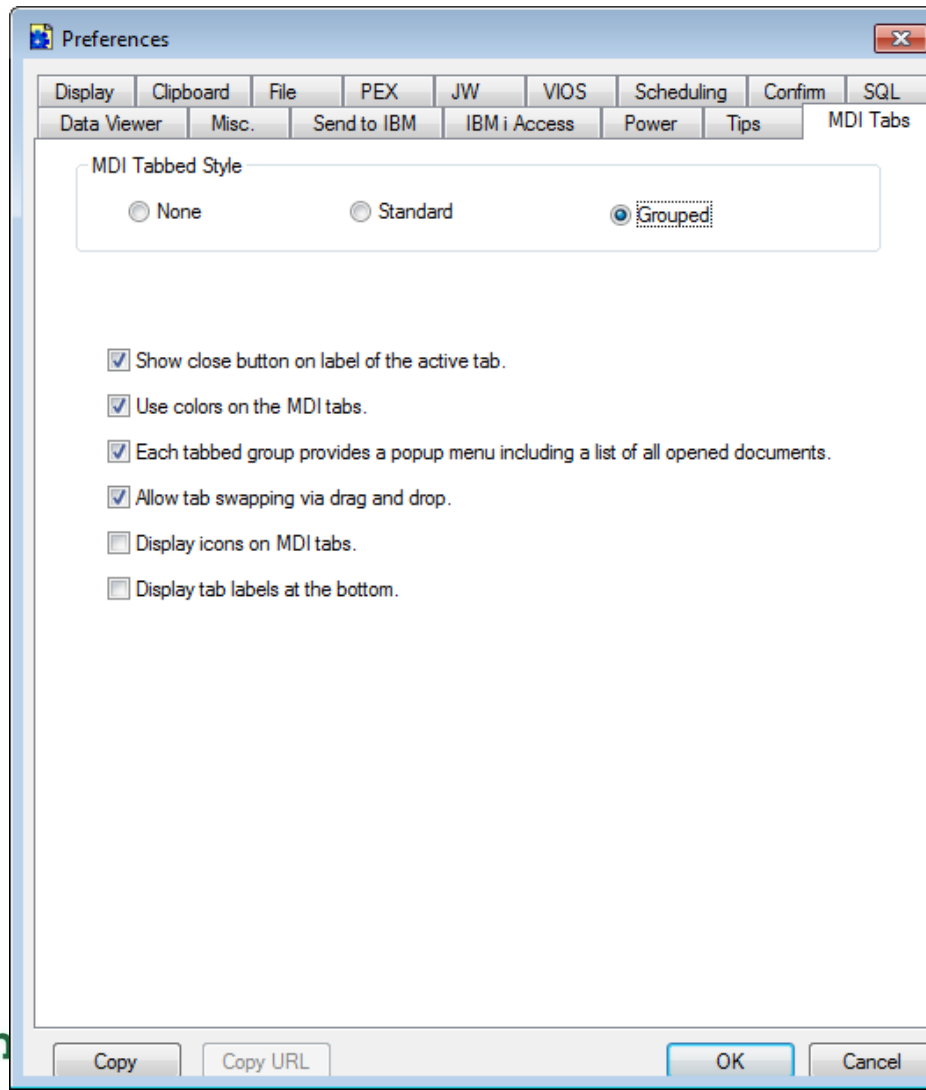
General – SQL Catalog functions folder

This folder contains a Tables folder for allowing you to collect metrics to determine which tables would be good candidates to move to SSDs. You can also use this folder to view real time statistics for all tables on the system.

This folder also contains a subfolder called “IBM i Services” that provides reports over available data provided with recent technology refreshes made to IBM i and DB2 for i in the SQL catalog tables. These reports include system values, PTFs and user and groups information.

General – Preferences – MDI Tabs

Added a set of preferences to control behaviors of the MDI tabbed style interface.





General – Preferences – Data Viewer

In the Preferences -> Data Viewer tab added a new option called "Cache popup menus in memory" which will speed up the display of drill down popup menus shown in the Data Viewer considerably after they are initially loaded.

This option is checked by default.



General – Job queues, subsystems and job description changes

1. Only one job queue and subsystem will now be used for all of iDoctor, and is installed with the base support (QIDRGUI library.) The default value for the JOBQ is QGPL/QIDRJW and the default subsystem description value is QSYS/QIDRJW. The values used by the user at install time are stored in data area QIDRGUI/QIDRJOBQ.
2. For PEX Analyzer at 6.1 only , the job queue and subsystem used for classic analyses is now hard coded to QIDRPA/QIDRPA for both the JOBQ and SBSDB. This separate JOBQ/SBDB is required because classic analyses can only support 1 job running at a time.
3. Added a new command at install time IDOCSBA/INSTSBSD that is used to install the job queue and subsystem for iDoctor after running IDOCSBA/INSTIDOCBA.
4. All of the INSTIDOCJW command parameters used at install time are now deprecated at 6.1 and higher.
5. All of the INSTIDOCCHA command parameters used at install time are now deprecated at 6.1.



General – Job queues, subsystems and job description changes (continued)

6. All of the INSTIDOCPA command parameters used at install time are now deprecated at 6.1 and higher.
7. The install program has been updated to only prompt the user once for a job queue and subsystem (previously 3 different screens were used at 6.1 for JW/HA/PA.)
8. The install program check that restricted the number of GUI install threads to 1 if the same job queue and subsystem description was used has been removed.
9. The QIDRBCH and QIDRLOWBCH *CLS and *JOB objects have been moved to QIDRGUI. Previously they existed in QIDRPA, QIDRWCH and QIDRHAJ libraries.
10. Starting a collection in PEX Analyzer will now default to using the new iDoctor subsystem and job queue (previously the one defined by the Job Watcher installation) instead of QCTL.



General – Green screen program to check access code exp dates

Added a new program to QIDRGUI called CHKEXPDATE.
If you call this program like this:

```
CALL PGM(QIDRGUI/CHKEXPDATE)
```

The PEX Analyzer expiration date will be found in
QIDRGUI/CHKCODE20 data area.

The Job Watcher expiration date will be found in
QIDRGUI/CHKCODE21 data area.



General – Convert 1 record bar chart to pie

Added a new menu option for graphs:
Graph Definition-> Set bars per page -> 1

Use this if you want to convert a single job in a rankings graph to a pie chart by also using the menu Graph Definition -> Set Graph type -> Pie



General – PEX+ super collections (QMGTOOLS/STRPEX_)

In the GUI, replaced the super collections support (combined JW, PEX and CS collection) with the one provided with Must Gather Tools (STRPEX_ command.) In the iDoctor GUI this is called PEX+ and a PEX+ folder will appear in most components as long as the system is licensed to both PEX Analyzer and Job Watcher.

The PEX+ folder displays the libraries found on the system that contain STRPEX_ created data. Each library will show the Active pex sessions, PEX Database files (if any), PEX mgt col objects, JW data (if any), and CS data (if any) associated with the super collection.

Normal iDoctor functions can be applied to any of these for viewing/working the data.

In order to create a new super collection, you can right-click the PEX+ folder and use the Start Collection menu. This will launch the Start PEX Collection Wizard which has been modified with additional options for creating a PEX+ collection.



General – IBM i Access Client Solutions support

IBM i Access Client Solutions with the Windows Application Package installed is now supported as a possible replacement for the IBM i Access for Windows prerequisite that iDoctor requires in order to connect to an IBM i.



General – Connections export function

In the IBM i Connections or Power Connections views you can now use the Export Connections menu to save your desired connections to a Windows registry file.

One reason for doing this is if you need to uninstall and reinstall a different version of client access or move your connections to another PC.



General – QIDRGUI/IDRINSTALL command

Added a new command QIDRGUI/IDRINSTALL which can be used to install iDoctor and Must Gather Tools from the green screen.

This command also supports installing the save files from a local library instead of using the iDoctor FTP site. This can be used on LPARs that are unable to reach the iDoctor FTP site.



General – Install the GUI silently

Added options to the Install that allow the GUI to be installed silently (no prompting.)

To do this, you will need to unzip the iDoctorClientUpdate.exe (using WinZip or other compatible software) to a directory of your choice. Then execute InstalliDoctor.exe with these parameters:

```
InstalliDoctor.exe -hidegui -nostart
```

-hidegui will cause the GUI to not be shown at all
-nostart will not start the GUI when done.



General – Change sensitive user data analysis

CSI, Job Watcher and Disk Watcher all now provide an analysis called Change sensitive user data which will modify any customer sensitive information within to random values.

Use the Restore sensitive user data analysis to return the data that has been modified by this analysis to its original state.



General – Sort from graph legend option

In the graph legend added a new menu called Sort with options for Ascending and Descending that allow a user to resort the graph by the selected column.

This action only works in normal (not "flattened") graphs that have a table below the graph.

General – Mouse wheel scrolling in graphs

In graphs users can now use the mouse wheel to scroll the data up or down (or left or right).

By default the amount scrolled per mouse wheel click is 5% of the total number of bars in the report (or an entire page if 5% is > than the number of bars per page).

Also added a preference on the Display tab called "Mouse wheel scroll percentage of total bars:" which lets you control/change this value of 5% to something else if you desire.

Using the mouse wheel in graphs while holding down the CTRL key will cause the total number of bars per page to increase or decrease by this mouse wheel scroll percentage. The graph will be redrawn with the new desired total bars per page with each movement of the mouse wheel while the CTRL key is held down.



General – Graph zoom

Added the capability to zoom into a desired selection within a graph. This is done by left clicking on the background of the graph and holding the mouse down and moving it left or right to select several bars. The background will temporarily change to black during this process. Release the mouse and the selected bars will be redrawn as the entire graph.

To zoom back out, just click on the whitespace again in the graph. You can also zoom in more than once. For example: from 300 bars down to 100 then again down to 50. Left clicking each time will zoom back out at each level.

Also added a new preference on the Display tab called "Minimum bars needed to perform zoom:" with a default value of 5. This means you must select at least N bars in order for the zoom function to work. This is provided to avoid users accidentally zooming in when they didn't intend to. Decrease or increase this value as desired.



General – Graph threshold line

In a graph to make the threshold line move (or create one), just click on the graph to give it focus, then start pressing the following keys on the keyboard's keypad

- + (to go up 1%)
- (to go down 1%)
- * (goes up 10%)
- / (goes down 10%)

You can also hold any of these keys then let up after a bit to move the position faster.

You can also add this line using the Graph Definition -> Set threshold menu.



General – Override client codepage setting

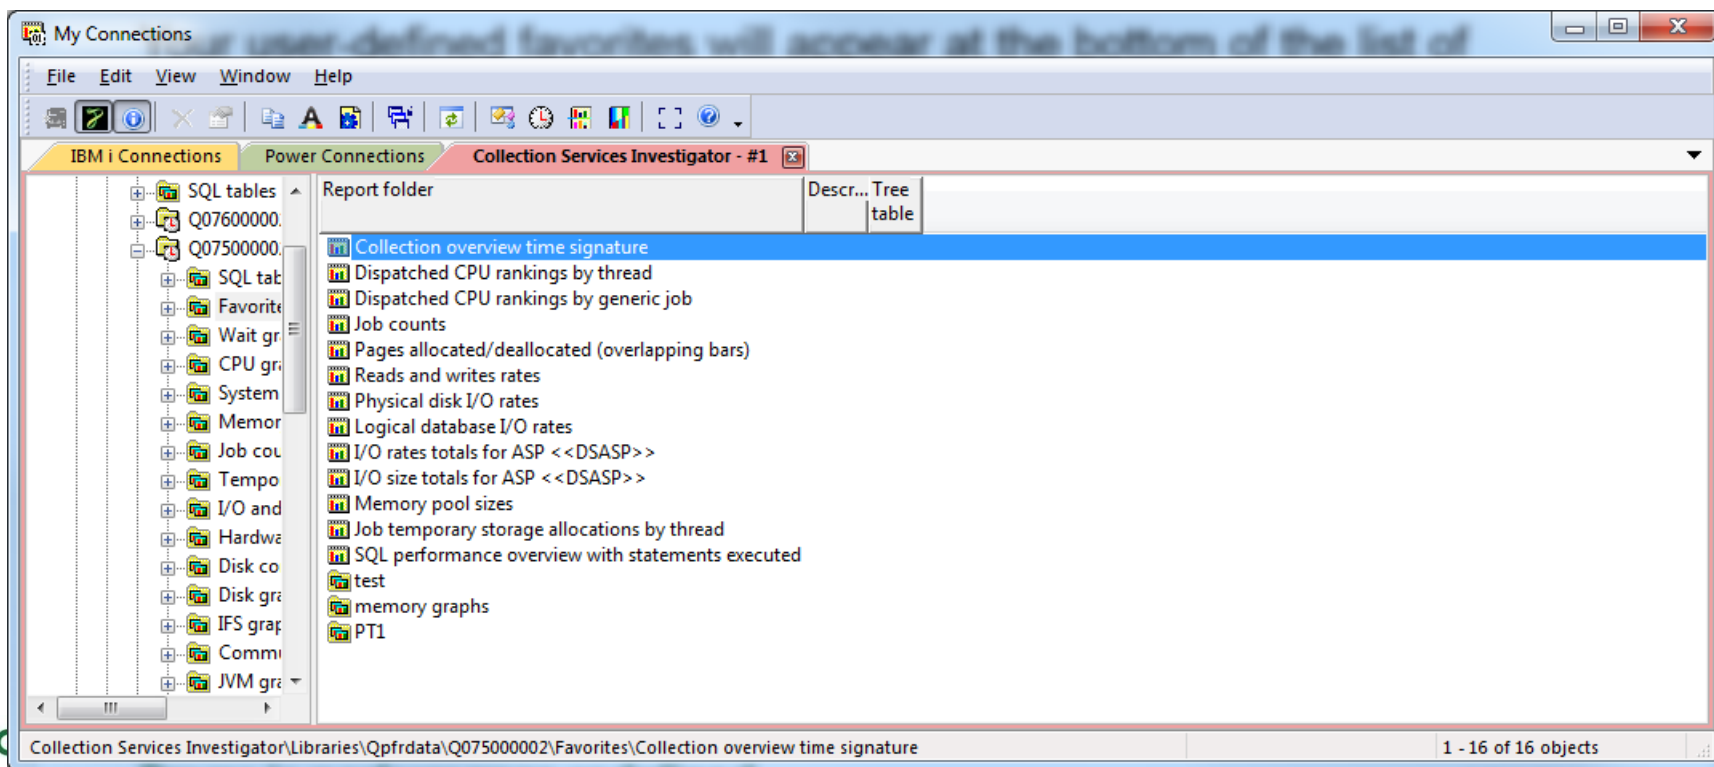
In the Preferences -> SQL tab added a new option called "Override client codepage setting with " that includes a textbox to provide the new value. The default is unchecked. In some DBCS environments (Chinese System i Access for Windows installed on a Chinese Windows PC) it may be necessary to change this override value to 1252 (Windows English code page).

This effects the CCSID used in the ODBC driver settings used by iDoctor for each IBM i connection. A restart of the GUI is required for the change to take effect.

Collection Services Investigator – Favorites folder

The new Favorites folder in CSI provides easy access to the most widely used graphs. You can also define your own lists of favorite graphs and reports using the save list option within the iDoctor Report Generator interface.

Your user-defined favorites will appear at the bottom of the list of graphs within the Favorites folder.



The screenshot displays the Collection Services Investigator (CSI) application window. The title bar reads "My Connections" and "Collection Services Investigator - #1". The interface includes a menu bar (File, Edit, View, Window, Help) and a toolbar. The main area is divided into a left-hand tree view and a right-hand list view. The tree view shows a hierarchy of folders: "IBM i Connections", "Power Connections", and "Collection Services Investigator - #1". Under "Collection Services Investigator - #1", there is a "Report folder" which contains a "Favorites" folder. The "Favorites" folder is expanded, showing a list of graphs and reports. The list view shows the following items:

- Collection overview time signature (highlighted)
- Dispatched CPU rankings by thread
- Dispatched CPU rankings by generic job
- Job counts
- Pages allocated/deallocated (overlapping bars)
- Reads and writes rates
- Physical disk I/O rates
- Logical database I/O rates
- I/O rates totals for ASP <<DSASP>>
- I/O size totals for ASP <<DSASP>>
- Memory pool sizes
- Job temporary storage allocations by thread
- SQL performance overview with statements executed
- test
- memory graphs
- PT1

The status bar at the bottom of the window shows the path: "Collection Services Investigator\Libraries\Qpfrdata\Q07500002\Favorites\Collection overview time signature" and the text "1 - 16 of 16 objects".



Collection Services Investigator – Historical Summary analysis

Added the following new graphs to the CSI historical summary analysis reports under the SQL folder:

SQL statements executed (7.2+)

SQL logical database I/O totals (7.2+)

SQL logical database I/O rates (7.2+)

SQL compressions (7.2+)

SQL cursor and native DB total opens (6.1+)

SQL cursor and native DB opens rates (6.1+)

Collection Services Investigator – SVC/Storwize graphs

Added a new set of graphs to analyze SVC/Storwize performance data under the SVC/Storwize graphs folder:

Virtual disk statistics - Summarized and ranked vdisk statistics

Managed disk statistics - Summarized and ranked managed disk statistics

Node statistics - Summarized and ranked node-level statistics

Flattened by virtual disk name - Flattened virtual disk statistics (one color per vdisk)

Flattened by managed disk name - Flattened managed disk statistics (one color per mdisk)

Flattened by node name - Flattened node-level statistics (one color per node)

Flattened by drive ID - Flattened node-level statistics (one color per drive ID)

Flattened by generic virtual disk name which will show one color per generic virtual disk name.

Use the QMGTOOLS/SVCIMPORT command to create the necessary database files.



Collection Services Investigator – Hardware configuration reports

Under Hardware -> Resource configuration added 4 new reports:

Hardware summary

Sockets summary

Memory summary

Planar summary

These give high level overviews of the cores, planars and memory available on the physical system.



Collection Services Investigator – Collection properties - System

1. Processor cores (Total / Installed / Available) is now shown as

On demand processor cores active (was called processor cores installed previously)

On demand processor cores maximum (was called processor cores total previously)

On demand processor cores available for purchase

2. Added a new value called 'Installed processor count on physical system'

3. Changed the 'On demand memory' label to 'On demand memory maximum'

4. Changed the 'On demand memory available' label to 'On demand memory available for purchase'



Collection Services Investigator – Management collection properties

Using the properties menu on a *MGTCOL object now provides a new tab called Management collection that lets you view the categories of data included within the object and the collection intervals and time periods included within.

This allows you to determine if the data included is what you expect.

Note: In order for data to be usable by CSI you will need at least *SYSLVL,, *JOBMI, *JOBOS at a minimum. Also ideally the time interval sizes will be somewhat consistent at least between *SYSLVL and the other categories or graphing issues will occur.

Collection Services Investigator – CPU graphs

In CSI added the following graphs under the CPU Graphs folder that show CPU cores (Virtual processor units) over time.

CPU utilization breakdown by core vs total CPU utilization
CPU utilization breakdown by core

At 7.1 and earlier the graphs will show up to 64 cores. At 7.2 the graphs will show up to 96 cores.

If all 64 cores are shown, then only about 200 bars per page can be used.

If all 96 cores are shown then only about 100 bars per page can be shown and you will need to adjust the graph bars per page setting accordingly in order to view these.



Collection Services Investigator – CPU SQL graphs

In CSI at 7.2, under the CPU graphs folder added a new graph called SQL CPU utilization that compares the SQL CPU utilization (unscaled) to the scaled CPU utilization along with the total CPU utilization on the 2nd Y axis.



Collection Services Investigator – 7.2 CPU graphs

At 7.2 the graph CPU graphs -> Shared pool CPU utilization now contains the "Scaled physical shared pool CPU utilization" on the Y2. This is how the physical shared pool CPU utilization would look if it were running nominally (or without changes to CPU speed by active energy management.)

Under the "CPU graphs -> Power 8 or higher" folder added 2 new graphs:

Instruction counts

Run-state processor virtual time



Collection Services Investigator – Temporary storage graphs

In CSI at 7.2, added a series of temporary storage graphs under the collection (under the temporary storage folder.)

All 3 types are provided (overview, rankings, selected job/job grouping over time.)

Collection Services Investigator – SQL graphs

At 7.2 only added the following graphs to CSI:

SQL graphs:

- SQL statements executed
- SQL logical database I/O totals
- SQL logical database I/O rates
- SQL compressions

- SQL statements executed rankings (contains 6 rankings graphs)
- SQL logical database I/O totals rankings (contains 6 rankings graphs)
- SQL logical database I/O rates rankings (contains 6 rankings graphs)
- SQL compressions rankings (contains 6 rankings graphs)

- SQL statements executed for <<job grouping XYZ>>
- SQL logical database I/O totals for <<job grouping XYZ>>
- SQL logical database I/O rates for <<job grouping XYZ>>
- SQL compressions for <<job grouping XYZ>>



Collection Services Investigator – SQL – Plan cache graphs

In CSI at 7.2 under the SQL folder added a folder called Plan cache which includes several graphs over the new collection services file QAPMSQLPC. These are high level SQL plan cache statistics.

Collection Services Investigator – System graphs (HMC)

Added a new folder under System graphs (HMC) called "LPAR graphs by OS type" containing the following graphs:

LPAR CPU time by OS type

LPAR cycles per instruction and CPU time by OS type

LPAR instructions per second and CPU time by OS type

LPAR entitled CPU time by OS type

LPAR memory allocated by OS type

Also added these same 5 graphs under subfolders called:

LPAR graphs by shared processor status

LPAR graphs by virtual shared pool ID



Collection Services Investigator – Hardware – resource configuration

Added a new set of reports over the hardware resource config file (QAPMHDWR) which provides outputs similar to the DSPHDWRSC command. Look under the Hardware -> Resource configuration folder to find the following reports:

- Hardware resources with totals
- Hardware resources
- Storage resources with totals
- Storage resources
- Comm resources

- Coupled resources
- Processor resources
- Local workstation resources
- Tape units by IOP/IOA
- Disk units by IOP/IOA
- Optical storage units by IOP/IOA
- Tape units and libraries
- Disk units
- Optical storage units
- Comm ports



Collection Services Investigator – Save/restore wait bucket graphs

In CSI at 7.1+ updated the hardware graphs -> save/restore folder to include a new graph showing all the wait buckets for the jobs doing save/restore operations called "Save/restore wait bucket times".

You can also drill down from this graph to get a job rankings chart for only the jobs doing saves/restores.

The Hardware graphs -> Save/restore folder also contains the following 4 different types of graphs with different colors depending on how the data is grouped together:

- All jobs, operations together over time.
- Per operation type (over time)
- Per job (over time)
- Per job and operation type (over time)



Collection Services Investigator – Hardware graphs - Tape

The following graphs are available that show each tape over time together on the same graph:

Tape wait time overview

Advanced tape wait time overview

Tape time spent waiting on drive

Tape percent time spent waiting on drive

Tape total I/O size rates

Tape read size rates

Tape write size rates

A subfolder called "All drives" shows totals for all tape drives together. A subfolder called "Selected tape drive only" shows stats for just 1 tape drive. This one is a possible drill down from the overview charts.

You will also see these 2 graphs that include disk, tape, and ethernet statistics all in one:

Disk, tape and ethernet I/O size rates

Disk, tape and ethernet I/O size rates with total MB/sec



Collection Services Investigator – Communication graphs

Added graphs for Domino server jobs under the Domino subfolder.

Added graphs for HTTP server jobs under the HTTP subfolder.

Added graphs for TCP/IP statistics under the TCP/IP subfolder.

Added a new folder called TCP/IP interfaces showing statistics per TCP/IP interface or line.

This folder also now contains graphs for Ethernet, Sockets and SSL statistics.

Collection Services Investigator – PT1 reports

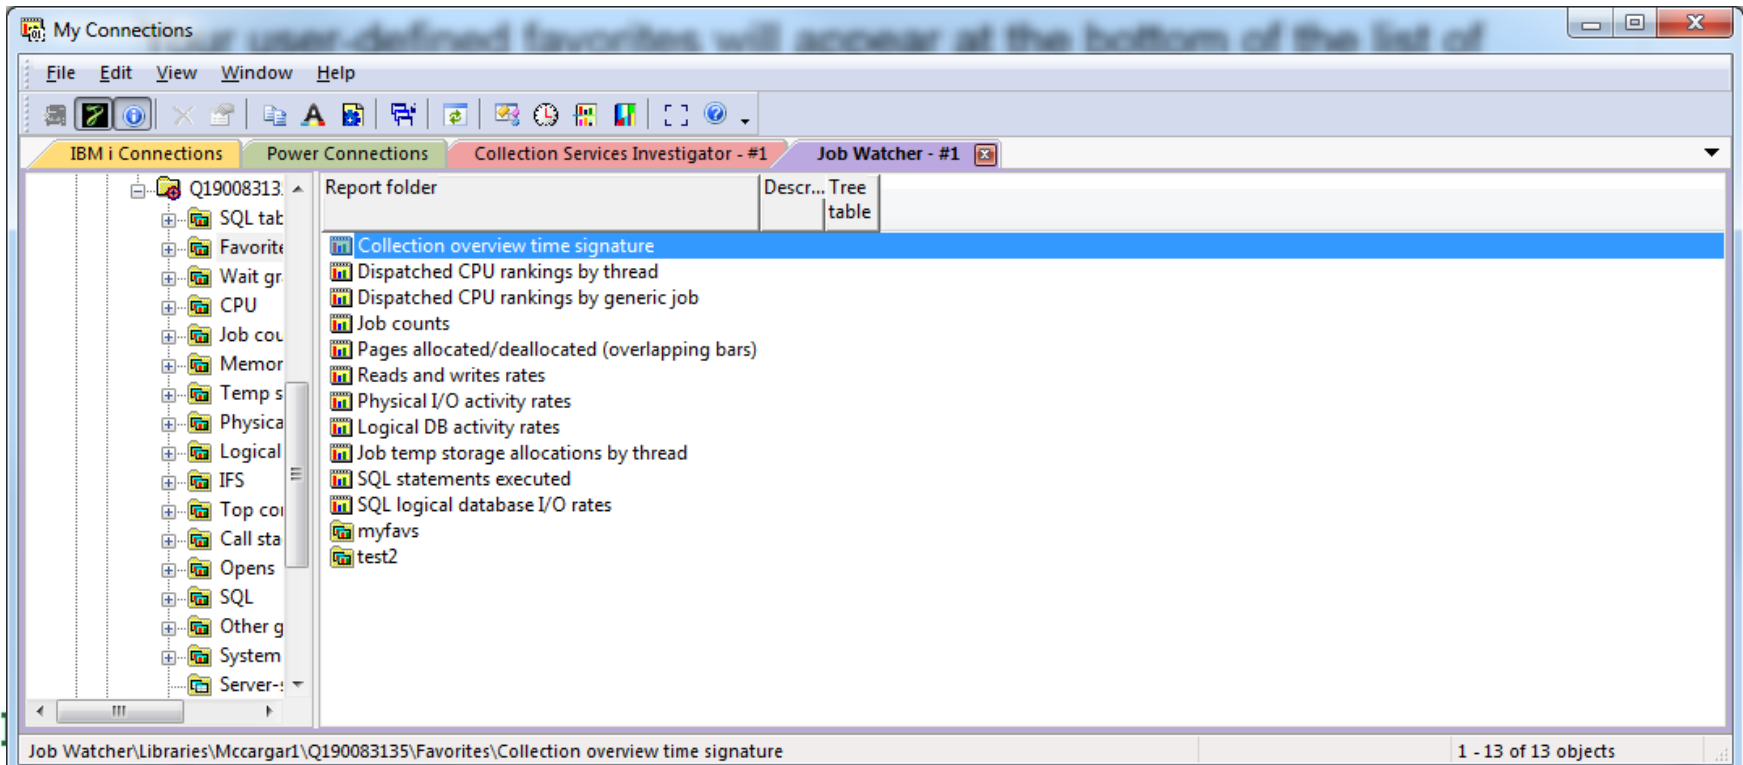
Added a new report folder called "PT1 reports" that contains the following set of reports similar to what you'll find in the green screen Performance Tools reports:

- System report - Communications Summary (Ethernet only)
- Component report - TCP/IP Activity
- Component report - TCP/IP Activity Totals
- Component report - TCP/IP Activity by line
- Component report - HTTP Server Activity
- Component report - HTTP Server Activity Totals
- Component report - Domino Server Activity
- Component report - Domino Server Activity Totals
- Component report - Extended Database Journaling Summary
- Component report - Database Journaling Summary
- Component report - Exceptions Per Second
- Component report - Exception Totals

Job Watcher – Favorites folder

The new Favorites folder in Job Watcher provides easy access to the most widely used graphs. You can also define your own lists of favorite graphs and reports using the save list option within the iDoctor Report Generator interface.

Your user-defined favorites will appear at the bottom of the list of graphs within the Favorites folder.



The screenshot displays the Job Watcher application window. The interface includes a menu bar (File, Edit, View, Window, Help), a toolbar, and a tabbed interface with tabs for IBM i Connections, Power Connections, Collection Services Investigator - #1, and Job Watcher - #1. The main area is divided into a left-hand tree view and a right-hand list view. The tree view shows a hierarchy of folders under 'Q19008313', including 'SQL tab', 'Favorite', 'Wait gr', 'CPU', 'Job col', 'Memor', 'Temp s', 'Physica', 'Logical', 'IFS', 'Top coi', 'Call sta', 'Opens', 'SQL', 'Other g', 'System', and 'Server-'. The 'Favorite' folder is expanded, showing a list of report folders. The list view displays a table with columns 'Descr...' and 'Tree table'. The table contains the following entries:

Descr...	Tree table
Collection overview time signature	
Dispatched CPU rankings by thread	
Dispatched CPU rankings by generic job	
Job counts	
Pages allocated/deallocated (overlapping bars)	
Reads and writes rates	
Physical I/O activity rates	
Logical DB activity rates	
Job temp storage allocations by thread	
SQL statements executed	
SQL logical database I/O rates	
myfavs	
test2	

The status bar at the bottom of the window shows the path 'Job Watcher\Libraries\Mccgar1\Q190083135\Favorites\Collection overview time signature' and the text '1 - 13 of 13 objects'.



Job Watcher – STRJWMON jobs ending abnormally

When a Job Watcher monitor (or Disk Watcher monitor) ends a collection, it will no longer trigger the job to end abnormally if the collection does not exist at the time it is ended.

That could happen if the job ended on its own due to the definition settings, before the monitor was able to end it.

The latest server builds are required. (QIDRWCH library)



Job Watcher – STRJWMON command OVLAP parameter

The help text for this parameter has changed as follows:

Collection overlap (OVLAP)

OVLAP

This value (in seconds) will be added to the JW definition's collection duration, to ensure that data is still being collected while the next collection is being started. This value should be set to at least how long it takes for a collection to start in your environment.



Job Watcher – Additional green screen command changes

1. The following commands now have a new parameter called "Run default analyses (RUNSUM)" *YES/*NO with default *NO:

ENDJWCOL, ENDDWCOL
ENDJWMON, ENDDWMON, ENDPAMON

This allows the default "summarization" analyses to be ran when collections or monitors end based on the command used.

2. Commands DSZJWCOL and SPLJWCOL existed only at 6.1 and have been removed at that release. These commands have been replaced by functions in the GUI.

3. Removed the ADDPRDACS command (use the same command in library QIDRGUI instead.)



Job Watcher – new 50 level call stack summary reports

Under detail reports -> 50 level call stack summary-> added the following 2 new options that let you filter the call stacks returned by a desired program name or procedure name :

1. 50 level call stacks containing user-specified program name
2. 50 level call stacks containing user-specified procedure name



Job Watcher – Memory graphs

Added a new folder called "Memory" that contains the following graphs:

Pages marked easy to steal (7.2)

Pages marked easy to steal by pool (7.2)

Memory page demand

Memory page demand by pool

Net page frames requested

Net page frames requested by pool

Pages marked easy to steal rankings (7.2 with 10 groupings)

Memory page demand rankings (with 10 groupings)

Net page frames rankings (with 10 groupings)



Job Watcher – Job counts graphs

Under the Job counts folder added the following new graphs:

Job counts rankings (with 7 job groupings)

Short-lived rankings (with 7 job groupings)



Job Watcher – CPU graphs

Renamed the "CPU graphs" folder to "CPU" and also made several changes:

Replaced the CPU graphs -> by thread folder with CPU -> Dispatched CPU rankings (with 10 groupings,)

Added the following new graphs:

CPU consumption

CPU utilization

Active processors

CPU consumption rankings (with 10 groupings)

CPU utilization rankings (with 10 groupings)



Job Watcher – Wait graphs

Under the wait graphs folder, replaced by the Wait graphs -> by thread folder with Wait graphs -> dispatched CPU rankings with 10 job groupings within if the collection is summarized.

Note: If the collection is not summarized then only a "by thread" job grouping is available.



Job Watcher – Rankings graphs changes

Made several changes to the graphs related to how drill downs are performed to be more consistent with the drill downs found in Collection Services Investigator. Also added many new job rankings graphs (with 10 groupings) for most metrics where possible.

The menu to drill down is now called "Rankings" with options within that to select the type of metric then the job grouping.

Job Watcher – SQL graphs

Added the following new graphs under the "SQL" folder at 7.2 ONLY:

- SQL statements executed
- SQL logical database I/O totals
- SQL logical database I/O rates

Also added the following set of ranking graphs each with these 10 job groupings available (thread, job, job user, generic job, current user, pool, priority, subsystem, job type, job function):

- SQL statements executed rankings
- SQL logical database I/O totals rankings
- SQL logical database I/O rates rankings

Note: The other new SQL metrics (CPU, PDIO, clock time, etc) at 7.2 contain BAD data and are not yet graphable. A PTF will be required. Once the data is fixed, additional graphs can be added.

Job Watcher – Temporary storage graphs

Added the following new graphs if the collection has been summarized:

- Job temporary storage allocations
- Job temporary storage allocations rankings
- Job temporary storage allocations by thread
- Job temporary storage allocations by job
- Job temporary storage allocations by job user
- Job temporary storage allocations by generic job
- Job temporary storage allocations by current user
- Job temporary storage allocations by pool
- Job temporary storage allocations by priority
- Job temporary storage allocations by subsystem

The 1st graph will be more accurate if the collection was created with the "force 1st interval" option set on the definition to collect all jobs in the 1st interval. Summarizing the collection is required in order to determine the last active record in the QAPYJWTDE file for each job.



Job Watcher – Wait object library

For 7.2 only, added the wait object library to several drill down reports where needed and the interval detail property pages.



Job Watcher – Collection summary by TDE type analysis

Added a new analysis to JW called Collection Summary by TDE type. You must run the collection summary analysis before running this one.

This analysis produces an SQL table adding up totals by type of TDE (meaning primary job, task , and secondary thread).



Disk Watcher – STRDWMON command OVLAP parameter

The help text for this parameter has changed as follows:

Collection overlap (OVLAP)

OVLAP

This value (in seconds) will be added to the DW definition's collection duration, to ensure that data is still being collected while the next collection is being started. This value should be set to at least how long it takes for a collection to start in your environment.



PEX Analyzer – Maximum events to collect option

In the PEX Collection wizard – Options screen users can now specify an optional value for the maximum number of events to collect.

If a value is provided then the collection will end after this number of events has been collected.

Note: This program/check will wake up every second to see if the requested number of events has been reached. Because of this the number of events actually collected will usually exceed the limit provided, but typically not by much.

The latest server builds must be installed in order to use this option.



PEX Analyzer – PMCO interval size mismatch check

In the Start PEX Collection wizard, it will now check to see if an active PEX collection is already running with a different CPU interval size than the one being proposed to start. PEX does not allow 2 different PEX collections to be active on the system at the same time with different CPU interval sizes. If a mismatch exists, the user will be warned and asked if they would like to change their interval size to match the existing one and continue (or just quit.)

Note: This option will not correctly detect interval sizes that use decimal values like .1 or 1.5 due to a limitation with the IBM i API used to perform this check. In those cases the checking will be incorrect and may allow a collection to start (and then fail).



PEX Analyzer – Call stacks analysis

The call stacks analysis has additional options:

1. You can now choose to include/exclude statement numbers in the output.
2. You can now choose to generate call stacks by generic job/task.



PEX Analyzer – Activation groups analysis

Added an Activation groups analysis to PEX Analyzer. Currently this analysis just provides a single table based report that summarizes the events collected.



PEX Analyzer – TPROF (with PMCO format 4)

At 7.1 and higher, after running the TPROF analysis on a collection that contains PMCO format 4 (MCLI) events, an additional folder will be listed under the TPROF folder called "Cacheline breakouts"

This folder contains a new report "Cacheline breakout with threshold count" and from that report several additional reports are available for more detailed MCLI analysis on a per cache line basis. All of these reports apply only to Power 7 systems and are:

Hits by procedure for specific cache line

Hits by job or task for specific cache line

Seize/lock object resolution for specific cache line

Call stacks for specific cache line

Call stacks with threshold met for specific cache line



PEX Analyzer – Green screen command changes

1. Removed the ADDPRDACS command (use the same command in library QIDRGUI instead.)



PEX Analyzer – Start PEX Collection Wizard changes

1. You can now kick off a PEX+ collection from this interface.
2. The STRPACOL command problem types are now no longer used. Instead the GUI provides a new set of IBM-supplied types some using new names that are more consistent with those used by QMGTOOLS (i.e. *DISK instead of *PDIO_TIME.)
3. ADDPEXDFN is used to create the IBM-supplied types with your changes.
4. Added the new problem types on the additional events screen.
5. Removed the option to specify the format for the taskswitch outq event. (SWOQ)
6. Added a checkbox on the options screen that lets you indicate if PMCO events shall be included. This can be used to modify any of the IBM-defined definitions at run-time.
7. In trace mode all jobs and all tasks will now be included by default. Previously all jobs and no tasks was the default.
8. On the options screen you can now specify the maximum number of events to collect.



PEX Analyzer – PEX objects – Create Collection menu

In PEX Analyzer, the PEX objects -> Create collection menu option now supports multiple selected collections at once. If this is done, then all collections will be created into the same library and collection name as each existing *MGTCOL object.



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, SPECcapc, SPECchpc, 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
ABAQUS	Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results. http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seq=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mssoftware.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