

iDoctor What's New Oct – Mar 2011



Ron McCargar
iDoctor Development

iDoctor resources

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

iDoctor update history: embedded into the GUI. *Tip: Search the update history on the area of interest to see if your question is answered there.*

YouTube Channel (20+ videos): <http://www.youtube.com/user/IBMiDoctor?feature=mhum>

IBM i 7.1 Technical Overview – Covers all updates from 6.1 GA to 7.1 GA: <http://www.redbooks.ibm.com/redbooks.nsf/RedpieceAbstracts/sg247858.html?Open>

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

Nov–Mar 2011 – New Builds

Released new builds on Feb 3rd.

Generally try to do 4 major updates a year.

7.2 internal builds not yet available, probably April or May.

Nov–Mar 2011 - Overview

Current Focus Areas

Side-by-side comparisons

GUI Performance Improvements

IASP support

Monitors

CSI/JW/PEX updates

Top 10 reasons to update to latest builds

Important things to know

Nov–Mar 2011 – Current Focus Areas

Making users more productive (always our #1 goal!)

Comparing data in iDoctor more effectively:

Side-by-side comparison mode

Monitors

Several updates based on customer requests

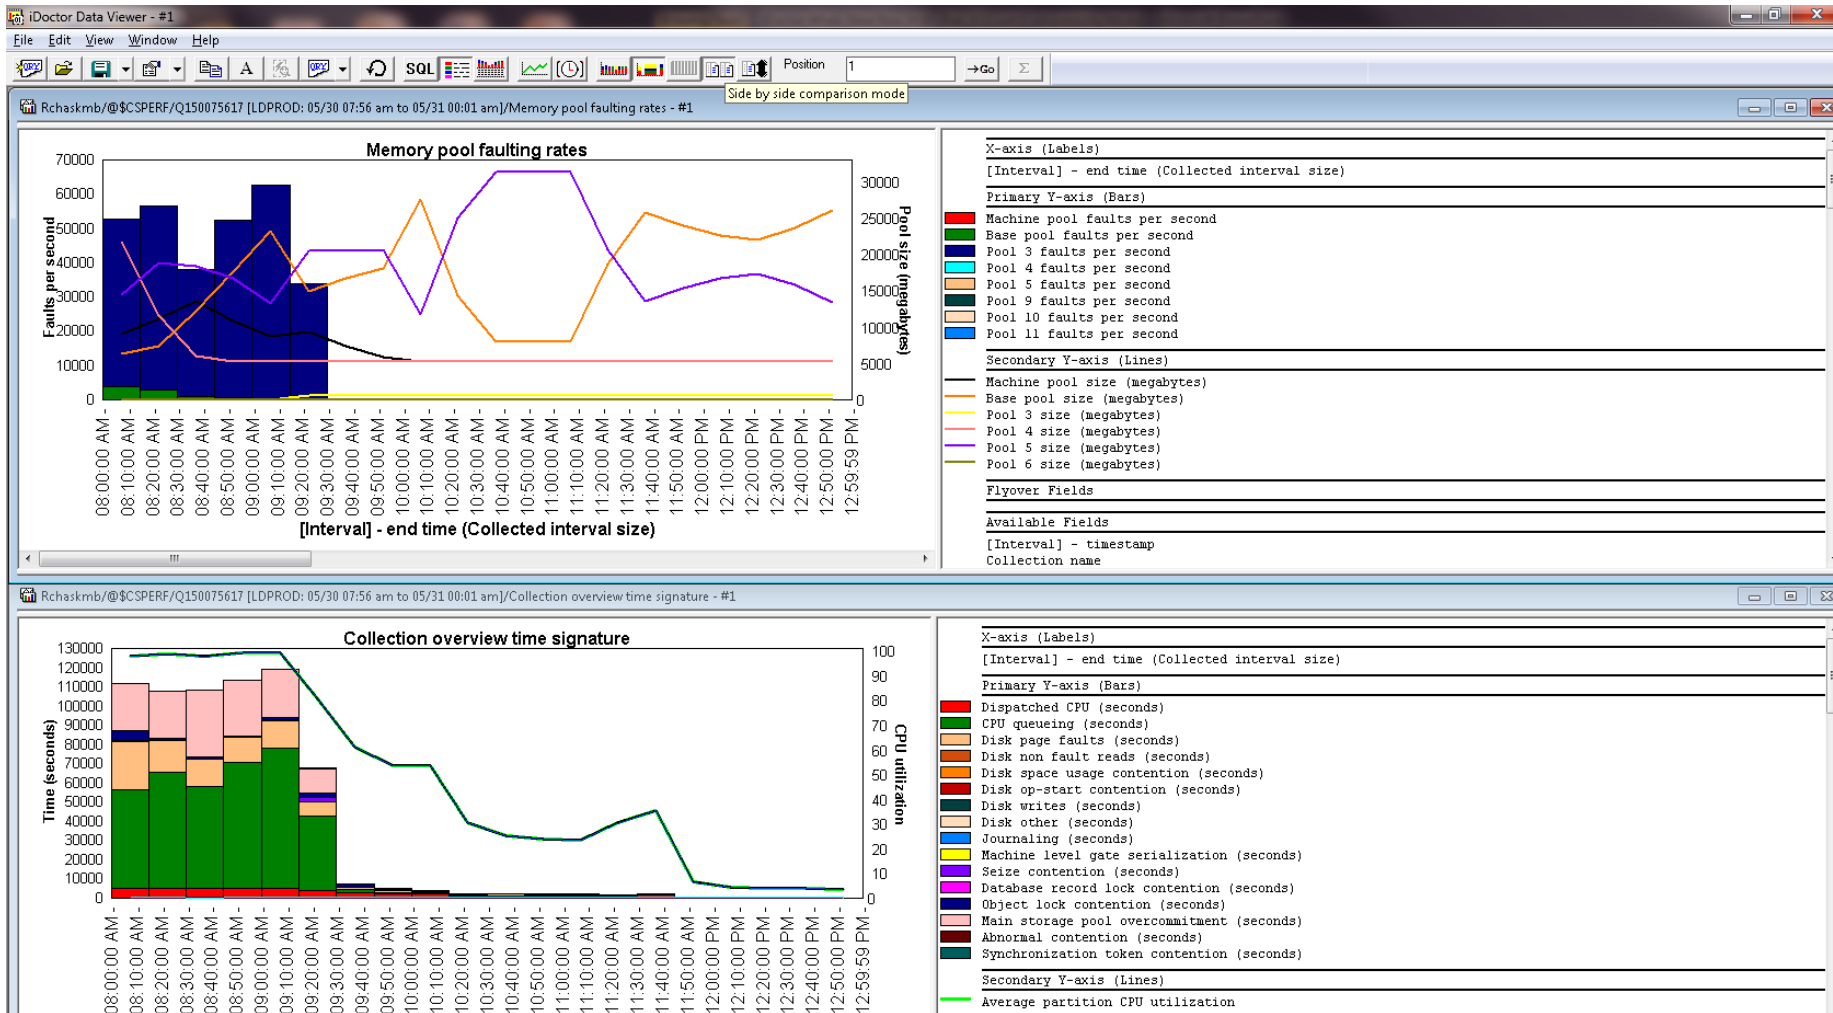
External storage

Server-side VIOS data collection enablement still in progress.
GUI development to start when server-side is done (May?).

Surface new metrics/features added to OS

Nov-Mar 2011 – Side-by-side comparisons

Side-by-side comparisons allow you to sync up the scrolling and Y axis scaling of any 2 graphs (or tables) in a Data Viewer.



Nov–Mar 2011 – Side-by-side comparisons

Example usage:

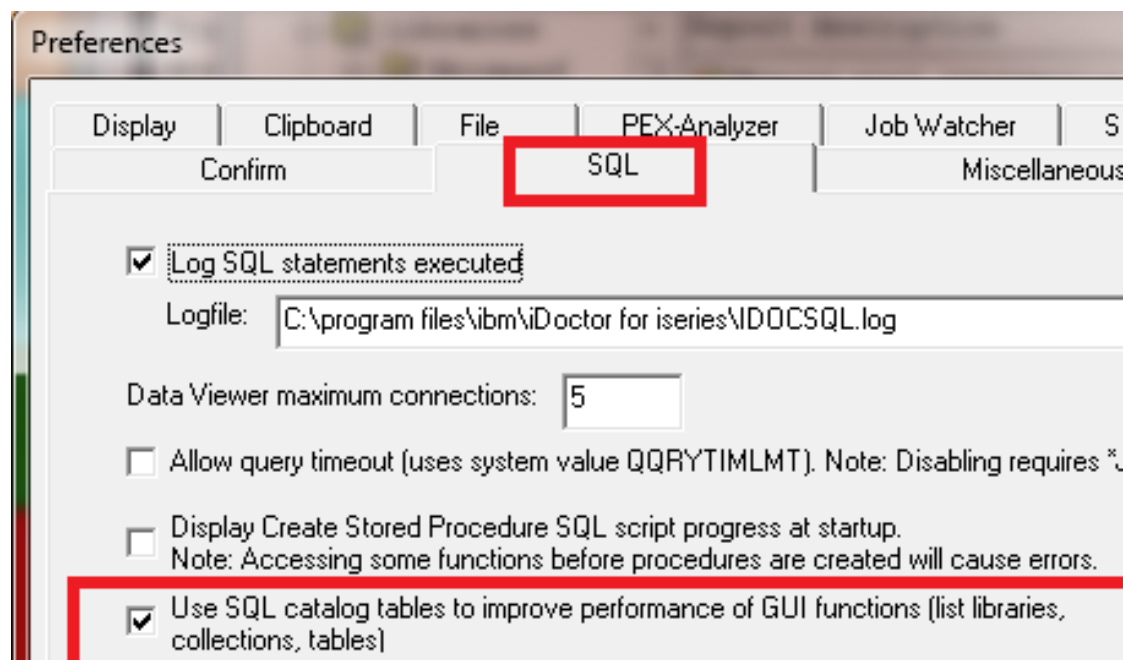
http://www.youtube.com/watch?v=0qkg_cmglAo&feature=channel_video_title

Future plans:

1. Make SSD Estimator invoke a side-by-side comparison
2. Possibly replace “graph compare mode” with side-by-side comparisons. The advantage of being able to compare any graph or table makes it a better choice in most cases.

Nov–Mar 2011 – GUI Performance Improvements

New preference on SQL tab controls whether SQL methods are used to replace older ones. This has performance advantages in most cases but doesn't always work on some systems (due to various issues).



Nov–Mar 2011 – GUI Performance Improvements

The following features are faster if the previously mentioned option is enabled:

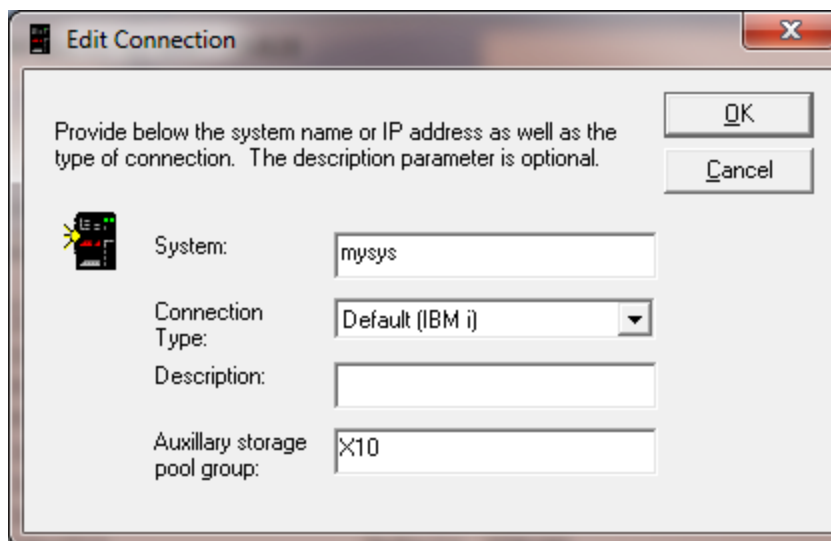
1. Listing libraries that contain performance data.
2. Listing collections in libraries for 5.4+ JW, CSI
3. Listing collections in libraries for 5.4+ PEX is only faster if the 'Hide-classic analysis options' preference is enabled.
4. Listing the contents of the server-side output files folder (in all components)

Nov–Mar 2011 – IASP Support

Added support to view collections stored in libraries created in Independent ASPs:

http://www.youtube.com/watch?v=43OauwCub3k&feature=channel_video_title

1. Controlled using the ASP group parameter on the Add/Edit Connection Window.



Provide below the system name or IP address as well as the type of connection. The description parameter is optional.

System: mysys

Connection Type: Default (IBM i)

Description:

Auxillary storage pool group: X10

OK Cancel

2. ASP group value used must be the IASP database name shown in System i Navigator and the ASP group name shown in WRKDEVD *ASP

Nov–Mar 2011 – Monitors Enhancements

These enhancements apply to 5.4+ STRJWMON, STRPAMON, STRDWMON.

1. The Collection Overlap (OVLAP) parameter is obsolete. Monitor will detect that new collection has started before ending prior.
2. The Collection Duration (STRGAP) parameter can now be specified in minutes with a decimal point (i.e. 59.5 minutes)
3. When restarting a monitor if the max historical collections has been reduced added checks to delete the extra ones.
4. Deleting collections (or other time consuming tasks) is done in submitted jobs to not increase each collection's duration beyond the specified time limit.

Nov–Mar 2011 – JW Monitors Enhancements

These apply to 5.4+ STRJWMON.

1. Added Resubmit Collections (RESUBMIT) parameter to submit new collections if a collection fails to start or quits early.
2. Added Max Consecutive Resubmits (MAXTRIES) parameter to indicate the number of times collections will be resubmitted if the RESUBMIT parameter is set to *YES and the current collection has ended prematurely.

Nov–Mar 2011 – Collection Services Investigator (CSI)

Added Interval Detail pages (wait buckets, I/Os, etc) similar to the ones in Job Watcher to CSI.

Added J9 JVM graphs showing heap consumption and more.

<http://www.youtube.com/watch?v=4zPOshEwzE&feature=related>

Added Virtual I/O graphs

http://www.youtube.com/watch?v=zrpG0b0af78&feature=channel_video_title

Added a new report to the disk configuration folder called "Capacity (in GBs) by ASP with paths"

In the disk graphs, added a new graph (for each disk graph type) called Total queue elements and average length.

Nov–Mar 2011 – Job Watcher (JW)

Job Watcher Installation video

http://www.youtube.com/watch?v=pURHegltOTQ&feature=channel_video_title

Job Waits and IBM i White Paper for 6.1/7.1

[https://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/\\$FILE/Job_Waits_White_Paper_61_71.pdf](https://www-912.ibm.com/i_dir/idoctor.nsf/3B3C112F7FBE774C86256F4000757A8F/$FILE/Job_Waits_White_Paper_61_71.pdf)

Discussion Video <http://www.youtube.com/watch?v=7t5qoPIUJlw&feature=related>

In the Interval Details interface, added a button to go to primary thread from a secondary thread.

http://www.youtube.com/watch?v=A_PZFK9ldl8&feature=related

Added a new Call Stack Summary analysis to identify the call stacks, waits and objects associated with the most frequently occurring call stacks found in the collection.

http://www.youtube.com/watch?v=NyycbX6eCL4&feature=more_related

Added new reports under the Detail reports -> Opens menu that shows the top programs causing opens (as detected by the Job Watcher call stacks captured).

Nov–Mar 2011 – SQL Tables

Added New Options to the SQL Tables Interface

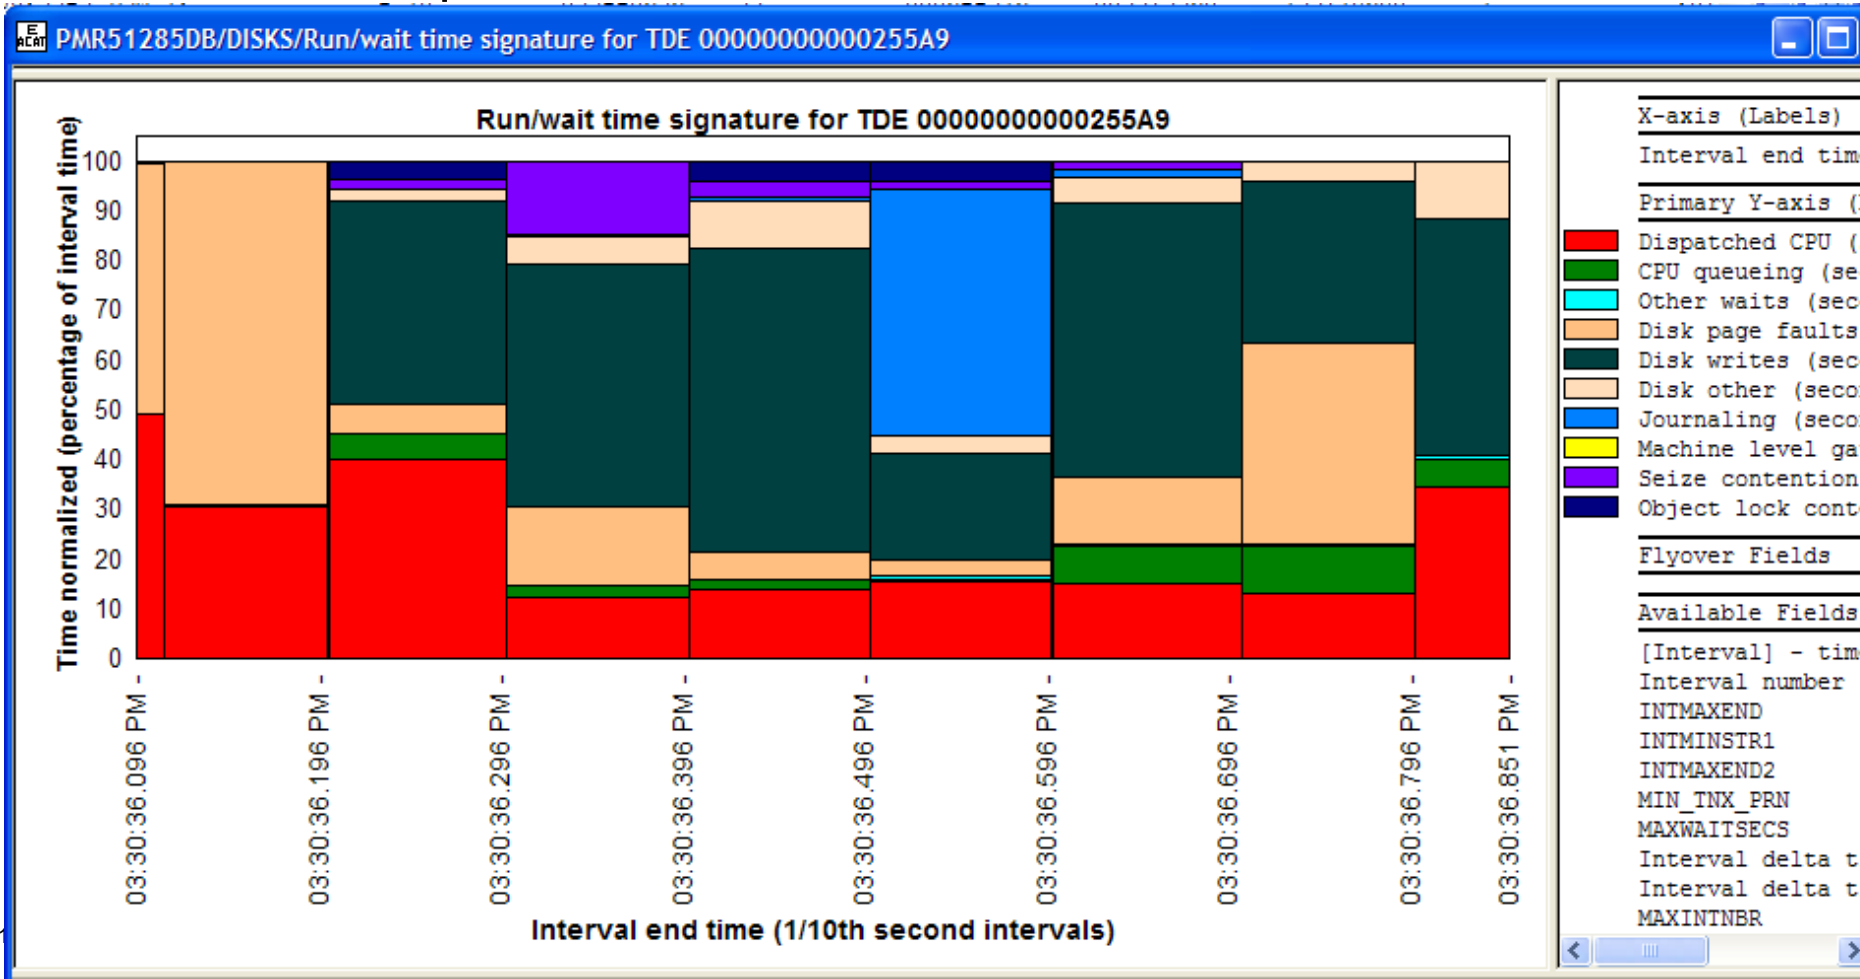
1. Create Merged Table – Create a new SQL table from all selected.
2. Open Merged Table – Runs a query to ‘union all’ selected tables together.
3. Launch Compare Wizard – Build customized reports over the desired tables to summarize and compare data.

Added special comparison options based on this for PEX Stats

Nov-Mar 2011 – PEX Taskswitch


In PEX Taskswitch, added new graphs to show wait bucket times for a desired job/task (TDE). Also added several additional drill downs from this graph in the Taskswitch analysis.

New overview reports also added.



Nov-Mar 2011 – PEX Taskswitch

This graph shows each taskswitch “transaction” because the X-axis indicates the time range is “Collected interval size”. Some transactions can be very long and some are very small.



Nov–Mar 2011 – PEX Call Stacks

1. When displaying call stacks in PEX Analyzer, added call level 0 from qaypebase (if the call stack is a base event and it's available)
2. On the PEX Call Stack Window, added a text box with the current QRECN and a refresh button that allows you to display a different call stack for a different user-entered QRECN within the current collection

Nov–Mar 2011 – Objects owned by User

Utility function to manage/cleanup iDoctor data or other data stored on the system.

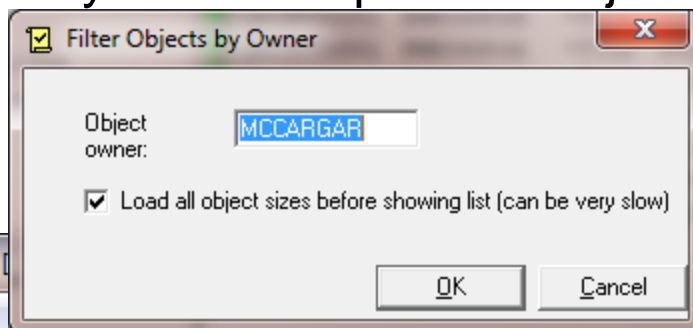
The screenshot shows the IBM iDoctor interface. The left pane displays a tree view of folders under 'Job Watcher'. The right pane shows a table with 'Folder Name' and 'Description' columns.

Folder Name	Description
All objects	Work with the objects on the partition owned by MCCARGAR
All files	Work with the physical files, logical files, aliases and save files on the
iDoctor files	Work with the collection files (tables) on the partition owned by MCCARGAR
iDoctor analysis files	Work with the analysis files (tables) on the partition owned by MCCARGAR

The tree view on the left includes folders such as Libraries, Monitors, SQL tables, Super collections, Definitions, Saved collections, Scheduled jobs, Active jobs, ASPs, Disk units, and Objects owned by MCCARGAR. The 'Objects owned by MCCARGAR' folder is expanded to show 'All objects', 'All files', 'iDoctor files', and 'iDoctor analysis files'.

Nov–Mar 2011 – Objects owned by User

- Object sizes initially not all loaded, right-click and choose Filter Objects by Owner to preload object sizes owned by the user (often slow)



IBM iDoctor for IBM i C00832 - [

File Edit View Window



Job Watcher

- + Libraries
- + Monitors
- + SQL tables
- + Super collections
- Definitions
- + Saved collections
- Scheduled jobs
- + Active jobs
- + ASPs
- + Disk units
- Objects owned by MCCARGAR
 - + All objects
 - + All files
 - + iDoctor files
 - + iDoctor analysis files

Object	Library	Type	Attribute	Description	ASP device	Size (KBs)
Qaidr00007	PMR06984B	*FILE	Physical file	Job totals by collection	*SYSBAS	480
Qaidr00001	PMR06984B	*FILE	Physical file	CSI - Interval Summary	*SYSBAS	344
Qaidr00002	PMR06984B	*FILE	Physical file	CSI - Aggregated Interval Summary	*SYSBAS	272
Qaidrjwcpu	MCCARGAR54	*FILE	Physical file	JOB WATCH - CPU Utilization statistics	*SYSBAS	84
G_tpr00008	IBM2105_1	*FILE	Physical file		*SYSBAS	84
Qaidrpcr1	MCCARGAR1	*FILE	Physical file		*SYSBAS	52
Qaidr00001	OSC0722	*FILE	Physical file		*SYSBAS	40
Qaidr00004	TESTING	*FILE	Physical file	Situational analysis detail file	*SYSBAS	40
Qaidr00005	PMR06984B	*FILE	Physical file		*SYSBAS	40
Qaidr00001	PMR06984A	*FILE	Physical file		*SYSBAS	40

haskmb: Job Watcher\Objects owned by MCCARGAR\iDoctor analysis files

Nov–Mar 2011 – Top 10 reasons to update to latest builds

1. In CSI SQL aliases were moved to QTEMP to resolve problems with them residing in collection library.
2. Graphing data with 'variable-width' mode now support microseconds accuracy instead of just seconds.
3. Fixed a hang if starting a JW collection on the system that did not have any definitions created yet. (Workaround was to do ADDJWDFN once first).
4. In JW 6.1 when listing collections in a library, if QAIDRJWRI file was present in the library but had no records it would cause the GUI to not show the collections in the library.
5. Fix for a lack of hits in PEX TPROF "Hits by component" report

Nov–Mar 2011 – Top 10 reasons to update to latest builds

6. In JW, in the unsummarized collection overview graphs, the CPU utilization was calculated as the maximum CPU within each time period, instead of the average.
7. STRPACOL (or the Start PEX collection Wizard) would fail to start a collection if the IFS additional event types (*IFSEVTS) had been included.
8. The rates in CSI disk path rankings graphs were double what they should have been if graphing multiple collections.
9. In the ASPs folder, expanding an asp with no disk units configured within would cause the GUI to hang.
10. QIDRPA/STRPACOL infinite loop and possible system crash if library QIDRGUI does not exist!!

Nov–Mar 2011 – Important things to know

1. The JW Collection Summary analysis no longer tries to hide wait times in the Collection overview graph for system tasks that report bad waits that normally really aren't (like SMLOGICALDASDADM.)
2. Data stored on libraries that contain a dot/period cannot be analyzed with iDoctor.
3. Added a Preference on the PEX Analyzer tab called "Hide classic analysis options at 5.4 and 6.1" which is checked by default. With this option checked, only the SQL-based analysis options will be shown, and the analyses will be shown directly under the collection rather than under the SQL-based analyses folder.