Update History for IBM iDoctor for IBM i - Job Watcher V5R4

April 23, 2010 – C00783

---> C00783 - Collection Services - Bug Fix Request
In CSI, the IFS graphs numbers in the by job and higher groupings are higher than they should be because of a problem with the SQL.
It was not handling the fact that IFS statistics are reported the same for all threads of a job.

---> C00783 - Job Watcher - Bug Fix Request
When displaying the collections within a monitor it was possible for the gui to incorrectly show collections starting with the same characters as the monitor name but not part of the monitor.

---> C00782 - Job Watcher - Usability
When stepping through the Start Monitor Wizard, it was possible to hit next multiple times on the basic options screen and have the same ptf check/warning window popup multiple times.

---> C00782 - Job Watcher - Usability
Starting or restarting a monitor in the GUI with the list of monitors shown, should no longer require the user to manually refresh the list of monitors in order to see the new one.

---> C00782 - Job Watcher - Bug Fix Request
Modified the scheduling support so that it no longer needs to retrieve the job scheduled entry number from the job log in most cases (unless the scheduled command string is > 500 characters). This should resolve a situation where the job scheduled entry number could not be retrieved properly.

---> C00782 - General - Usability
If the window title contains the <Collection system> parameter (the system the collection was created on) it will be shown as blank if the <system> parameter is the same system.

---> C00782 - Job Watcher - Bug Fix Request
In JW, opening some of the serverside output files would show an SQL error due to the incorrect order by being used.

---> C00782 - Collection Services - Bug Fix Request
In CSI, in the thread wait time signature graphs, the submitter job name was incorrectly shown as blank in the flyover window.

---> C00782 - Installation - Enhancement
Add option to the install on the summary panel to allow a user to create the setup log files in the desired directory. This can be required if doing a server-side only install and the default of C:\program files\ibm\idoctor for iseries\setuplogs will not work due to lack of ADMIN rights for the user.

---> C00782 - Collection Services - Enhancement
In CSI, from the CS objects (MGTCOL objects) folder, the create collection window now contains all the parameters available on the CRTPFRDTA command (categories, interval, time range).

---> C00781 - Job Watcher - Bug Fix Request
Fixed several issues with the function to start a monitor over multiple partitions. It will also be available now with the latest 5.4 JW server builds.

The graphing support (from JW and the initiating system) is working properly now. The flyover window will now contain the system name so you can tell which partition you are looking at in each view.

---> C00781 - General - Enhancement
Graph flyover window now contains the system name and the system the data was collected on (if different from the system name where data resides now).

---> C00781 - General - Bug Fix Request
Time range selections were broken with the previous client.

---> C00781 - General - Bug Fix Request
Fix a crash with the report generator function if opening a tree table view (CSI disk configuration or PEX TPROF reports).

---> C00781 - Collection Services - Bug Fix Request
In CSI, fix divide by 0 error on queries that use SYIFTA (occurs if system has no interactive capacity)

---> C00780 - Disk Watcher;Job Watcher - Usability
In Job Watcher and Disk Watcher monitors added a check to make sure the JOBQ is not held as the monitor runs. If the JOBQ becomes held during collection, the montior will end to avoid ending up with lots of monitor submission jobs sitting on the JOBQ.

This applies to all releases for both monitor commands except JW 5.4.

---> C00780 - General - Bug Fix Request
When writing your own SQL, if you specify the same table name in multiple libraries the data would not be queried correctly because iDoctor's use of OVRDBF. With this change, when the GUI encounters multiple files in the same query but with different libraries, OVRDBF will not be used for these files. If the files contain multiple members then a warning will be shown explaining that only results from the 1st member in each file can be shown and if this is not desired SQL aliases will need to be written to point to the correct members.

---> C00780 - Installation - Bug Fix Request
Renamed the install program from setup.exe to InstalliDoctor.exe to avoid the possibility at install time that a different setup program will be launched (from the user's path environment variable)

Additionally, registry entries at install time will no longer be written to HKEY_CLASSES_ROOT and will instead be placed under HKEY_CURRENT_USER to resolve an issue on Windows 7.

---> C00780 - Collection Services - Bug Fix Request
Fixed broken drill downs in the CSI interval summary interface from the pages that show lists of jobs and disks

---> C00780 - General - Usability
Add cancel button to the Choose database members window that will cancel running the current query.

---> C00780 - Collection Services - Bug Fix Request
Add Disk other waits to the default set of wait types graphed for the CSI wait graphs at 5.3 and 5.4 (16 bucket mapping). This was not shown for the "Collection overview time signature" graph but was at the rankings and single job over time graph types.

---> C00780 - Collection Services - Bug Fix Request
When making bar selections in time range graphs, just clicking the first bar and the last bar of the desired time range will cause the time range between them to be used as the drill down. Previously this required shift to be held down. Also pressing the control key while clicking can be used to do the same thing

although is not required. If > 2 bars are selected on a time range graph then the 1st and last bars are used to determine the time range.

---> C00779 - Collection Services - Bug Fix Request
In CSI, when listing collections in a library, the stored procedure that determines information about these collections will no longer rely on a file that requires *ALLOBJ authority.

---> C00779 - Collection Services - Bug Fix Request
 In CSI, the rankings filtered by option did not work with any graph types except wait graphs but is working now for the rest of the job based graphs

---> C00779 - Job Watcher - Enhancement
In 7.1+ JW, under the CPU graphs folder added a new graph:

Scaled CPU ratio - This graph shows ratio of scaled CPU to unscaled CPU for all jobs reported in the QAPYJWTDE file.

---> C00779 - Job Watcher - Usability
For 7.1 JW, in the server side files folder, when displaying the QAPYJWTDE file, the SQL now includes the 4 new fields added at 7.1 at the end of the file (scaled CPU times and virtual processor delay time)

---> C00779 - General - Bug Fix Request
In the previous client, the open report into new data viewer and open into existing data viewer options were not working correctly.

---> C00779 - Collection Services - Bug Fix Request
In CSI 5.4 in the interval summary property pages, the wait buckets tab had an SQL error. Also at that release the wait bucket totals tab has been removed (because it is derived from a 6.1 file qapmisum)

---> C00779 - General - Bug Fix Request
Switching the graph type to step or a line graph causes the colors defined in the legend to not be used in the lines and instead all black lines are shown.

---> C00779 - Collection Services - Usability
Change the way rates and utilizations are calculated in CSI when ranking jobs so they are based on the time period selected or the entire collection rather than based on the job's run duration.

---> C00778 - Job Watcher - Bug Fix Request
In JW for summarized collections, the ability to drill down from a Classic JVM overview to the list of JVMs for the desired interval was not showing up.

---> C00778 - Collection Services - Bug Fix Request
The duplicate as table, duplicate as graph options from a graph view was not working properly with the latest client.

---> C00778 - Collection Services - Bug Fix Request
In CSI 5.4 for the wait graph - Job wait time signature for job xyz , corrected a problem where all wait buckets would be initially shown instead of only the interesting waits.

---> C00778 - Job Watcher - Usability
Change the way rates are calculated in JW when ranking jobs so they are based on the time period selected or the entire collection rather than based on each individual job's run duration. This means the order of the jobs when graphing rates vs totals will be identical since the time duration used to compute the rate will be the same for all jobs in the graph.

Also in the I/O graphs -> by thread folder under the collection added I/O totals graphs for those graphs that only provided a rates graph.

---> C00778 - Job Watcher - Enhancement
In JW, for summarized collections, the ability to drill down on a list of jobs for the selected user, generic user, pool, etc was not showing up.
For example, this allows the user to drill down from pool rankings to see the list of jobs with the highest disk page faults for the desired pool.

---> C00777 - Collection Services - Enhancement
In CSI for 7.1 in the external storage rank graphs made the following changes:

1) The analysis step will create a new field in the rank statistics by time interval file called RANK_IOCALC which is really the true Disk I/Os calcuated from the Rank I/Os. This field will be the disk I/Os for the rank in the interval with a different calculation depending on the raid type. Raid 10's calculation is IOs = Reads + (2 * Writes).
All other raid types cannot be calculated accurately and the value will be NULL.

2) When drilling down from a rank "rankings" graph, you can now select multiple ranks and graph them together over time (as a step graph) or you can select a single rank and graph it over time as a bar graph. Use the "Selected Ranks" menu for the step graphs and the "Selected Rank ID" menu for bar graphs. (note you can also change the graph type now by right-clicking the legend)

3) The Selected Ranks drilldown menu includes an option called RAID 10 disk I/O rates , which shows rates over time for each rank selected using the new RANK_IOCALC field.

---> C00777 - Collection Services - Enhancement
In CSI for 7.1 in the external storage link graphs made the following changes:

1) There are now 2 sets of graphs, 1 for SCSI links and one for PPRC links.

2) When drilling down from a link rankings graph, you can now select multiple links and graph them together over time (as a step graph) or you can select a single link and graph it over time as a bar graph. Use the "Selected Links" menu for the step graphs and the "Selected Link ID" menu for bar graphs. (note you can also change the graph type now by right-clicking the legend)

3) SCSI link data is now shown as total I/Os only instead of showing reads/writes separately.

4) The ECKD data is no longer graphed as it is for system z and not i.

---> C00777 - General - Enhancement
Users can now right-click the legend to change the graph type using the "set graph type (primary Y)" menu.

---> C00777 - Collection Services - Enhancement
In CSI, from any job rankings graph, you can now drill down to another job rankings graph of any grouping using the current selection as a filter.
This lets you do things like graph a list of threads for the desired pool for example.

Use the menu called Rankings filtered by selected (thread, job, pool, etc)

---> C00777 - Collection Services - Bug Fix Request
The calculations for write cache overrun and avg response time done by the CSI situational analysis function were off for external disk. Now the numbers will match those given in the disk unit rankings graphs.

---> C00776 - Collection Services - Enhancement
Added the following new external storage link and rank statistics graphs

Link graphs -> by time interval
which provides drill downs to rank the links by their interface ID for the selected time range
Link graphs -> by interface ID
which provides drill downs to graph the selected link over time

Rank graphs -> by time interval
which provides drill downs to rank the ranks by rank ID for the selected time range
Rank graphs -> by rank ID
which provides drill downs to graph the selected rank over time.

---> C00776 - General - Usability
The iDoctor main window titlebar will now contain the iDoctor client build number.

---> C00776 - Collection Services;Job Watcher - Usability
In CSI and JW, in the Interval Summary interface, wait bucket totals tab, the wait buckets are now shown in a vertical list with 1 wait bucket per row instead of all wait buckets in one row as done previously. Sorting by any field also maintains the sort direction when scrolling through intervals.

---> C00776 - Job Watcher - Bug Fix Request
In JW 5.4, with the previous internal client only, the SQL that builds the General section of the Interval Summary pages had a syntax error.

---> C00776 - Job Watcher - Bug Fix Request
In JW and CSI, when displaying the create job summary reports in a library, and the job filter contained an underscore, the filter string would indicate both the job and subsystem filter was used when only the job filter was used.

---> C00776 - Job Watcher - Usability
When performing a JW search, the title bar of the report now contains the system name and start and end times of the collection (if this type of title has been selected for use within preferences.

---> C00776 - Collection Services - Enhancement
Added Bad Current Waits and Disk units tabs to CSI Interval Summary property pages.

The disk units tab allows you to filter by the desired ASP and select a unit and drill into the graphs over time.

---> C00776 - Job Watcher - Bug Fix Request
In JW, fixed some problems with the search function where you could type a value in to search on and instead of searching on that value, it would search from a value from the search history that contained similar characters.

---> C00775 - General - Bug Fix Request
On the main window the copy to clipboard function (ctrl+C) from the list side of the component views and the clock icon was not working after the changes that added user-defined queries and graphs to the collection menu. This was due to a conflict in the command IDs used.

---> C00775 - Installation - Bug Fix Request

If an ODBC error occurs during installation when trying to make the connection, the connection string used will now be returned in the error message.

---> C00775 - Collection Services - Enhancement
In CSI, add new interval summary property pages if you double-click a bar in one of the overview graphs. The new property pages includes tabs for:

Wait buckets - Shows the top jobs experiencing the wait bucket selected. The list features drill down capability into the selected job.
Physical I/Os - Physical disk I/O totals
Logical I/Os - Logical I/Os, SQL cursor stats
Wait bucket totals - This page shows the total time, counts, and contributing jobs in each wait bucket
Situations - Shows each situation found in the interval (if the situational analysis option has been ran)

All of the above includes a general section showing various CPU utilizations from QAPMSYSTEM as well as other high level disk and memory pool stats.

Additional property pages for interval details - (for a selected thread and interval) are not yet defined in CSI.

---> C00774 - Collection Services - Bug Fix Request
In CSI, fixed a crash when opening a create job summary report due to changes made in previous client.

---> C00774 - Job Watcher - Usability
Scrolling through intervals in the JW interval summary pages is faster now.
Also changed the interval arrow buttons to horizontal to look like the interval detail pages.

---> C00774 - General - Bug Fix Request
When saving a user-defined query or graph, Data Explorer was missing as a possible component.

---> C00774 - Job Watcher - General
For code cleanup reasons, most of the code that handled V5R2 and V5R1 Job Watcher on the GUI side has been removed. If you need to analyze collections for old/unsupported releases, use the iDoctor client included with the server builds (at those releases) on our external or internal websites.

Also added a check at startup if connected to a V5R2 or earlier system, Job Watcher will no longer be launchable.

---> C00774 - Job Watcher - Bug Fix Request
Update the checking for summarized JW colletions to check that the gap file is up to date to the latest version. If it isn't the collection summarized status will indicate that the gap file must be recreated. This will avoid SQL errors when graphing a single job/thread over time.

---> C00774 - Job Watcher - Enhancement
In JW 5.4 and higher, for summarized collections, in the wait ranking graphs by thread/job/job user/etc, added the current user profile as a flyover field to the graph. If multiples are found within the current job grouping then the value will contain "Multiples found".

---> C00774 - Collection Services - Bug Fix Request
In CSI disk graphs the read cache hit %, write cache overrun %, write cache efficiency % and device I/O rates should be more accurate when dealing with configurations having both external and internal storage in the same ASP.

---> C00774 - General - Usability

If the SQL statement contains comments, and the user tries to use query definition interface, prompt them asking if they really want to because the comments will be lost.

Also add a preference on the Confirm tab called "Confirm usage of query definition if the SQL contains comments (which will be lost)" which controls whether the confirmation is shown.

---> C00774 - General - Usability
When using the Duplicate as temporary SQL view option add a check to see if the query contains physical files with multiple members. This option cannot work in that instance so a message will be shown blocking the attempt. SQL views cannot be created over aliases or multi member physical files, only over SQL tables or physical files with only 1 member.

---> C00774 - General - Enhancement
In the iDoctor Report Generator add an option that has allows for N pages of data to be captured per graph/table. In tables the pages captured are vertical pages (rows) not horizontal ones (columns).

---> C00773 - Disk Watcher - Bug Fix Request
In DW 6.1+ when connecting to a system where no Disk Watcher definitions have been created yet, and using the Start Collection option , the GUI would display an error about the qapydwdfn file not existing in QUSRSYS. WIth this fix, the IBM-supplied definitions will be automatically loaded and the file created.

---> C00773 - General - Bug Fix Request
If the system contained a very large number of disk units and the user tried to use the Disk units folder or the disk unit selection option in the Add Disk Watcher Definition interface, the GUI would crash due to the receiver variable being set too small.

---> C00773 - General - Usability
On the manage and import query and graph definition screens, removed the VRM filter option since the VRM is no longer used when deciding which user-defined queries/graphs to show.

Also removed the Copy and Edit options on the Manage query/graph definitions window since they were only needed to move query/graph definitions to a new release.

On the manage query defintions and import query defintions windows, also added the SQL statement within the query definition to the list of fields. You can place the mouse over it for a more readable flyover and/or copy/paste the SQL statement cell contents to the clipboard.

---> C00773 - Job Watcher - Bug Fix Request
In the STRJWMON command changed the check to see if the definition exists. The method used of checking for the command string (being non blank) inside the qapyjwdfn file was not reliable.

---> C00773 - Job Watcher - Usability
Stepping through most tabs in the JW interval details interface with the outer buttons (find next available or find previous available call stack) is much faster now.
Stepping through call stacks is 2 to 3 times faster.

---> C00773 - Job Watcher - Usability
Stepping through the SQL statements in JW with the outer buttons (find next available or find previous available SQL statement) is faster now. If you uncheck the display host variables checkbox and then step through the available SQL statements the response time should be near 0..

---> C00773 - General - Usability
Display the x-axis label when using variable width mode so that the time range size (collected interval size, 1 min intervals, etc) is still shown on the graph.

---> C00773 - Collection Services - Enhancement
Added Situational Analysis support to CSI at 5.4+.
To initiate situational analysis right-click the collection and pick the Analyses -> Run Situational Analysis
option.   You may also right-click the collection and pick Analyses -> Analyze Collection.... and then from
there press the Preferences button to configure the thresohlds associated with the Situations detected by
CSI.

All the time interval based Wait graphs will include different background colors each identifying a situation.
We could add these to some of the other key graph types (CPU/IO/Disk) if desired, let me know if this is
necessary.

The current list of situations and the default minimum thresholds are shown below:

1. Interactive feature utilization high - 100%
2. Write cache overruns - 20%
3. High disk response times - 15 ms
4. High faulting in the machine pool - 10 faults per sec
5. High disk storage consumed - 90%
6. Jobs ineligible to run - 3 instances for a job per interval
7. Large change in pool sizes - 50% change from one interval to the next.

Any changes to thresholds are saved on the user's PC and will be used on future analysis runs.

A file identifying the situations found is also available under the server-side output files folder.

---> C00773 - General - Enhancement
In Preferences, removed the PEX Analyzer and Job Watcher title format preferences and replaced this with
a Window title format preference at the top of the Display tab that covers all components.

Also added new parameter markers for the title format including
<collection system> = The system the collection was created on
<collection start> = Start time of the collection in MM/DD HH:MM AM/PM format
<collection end> = End time of the collection in MM/DD HH:MM AM/PM format

Note: In some IBM internal components like Plan Cache Analyzer not all of these fields are available.

---> C00773 - Collection Services;Pex Analyzer - Bug Fix Request
When working with a generic job ranking graph in CSI and PEX, using the Change SQL Parameters option
to change the generic job name length worked fine, but drilling down did not work because it didn't use the
changed length value for the next graph.   This issue has been corrected.

---> C00773 - Collection Services;Job Watcher - Usability
Add a button on the Data Viewer toolbar that is used to show all the idle waits (include all buckets) for wait
bucket "job" and larger grouping graphs in CSI and Job Watcher.
This is a toggle button that allows you to quickly see the idle waits and then click again to see the
"interesting" waits.
Previously the idle waits were not shown for "job" or higher groupings.

---> C00773 - Pex Analyzer - Enhancement
In PEX, the active collections folder now contains the start and resume times for the collections listed.
Also the properties interface from collections in this folder will show the status and start time now.

---> C00773 - General - Bug Fix Request

When saving a user-defined graph containing multiple x axis fields the 2nd and 3rd x axis fields will now be saved properly and shown when the graph is reopened.

---> C00772 - General - Bug Fix Request
Added a fix if using kerberos authentication, in the previous client it would still prompt you for a password because the cwbCO_SetUserIDEx API was still being called.

---> C00772 - Collection Services - Bug Fix Request
Fixed a bug with CSI in the previous client (771) with the execution of stored procedure QIDRCSCRTA that was preventing CSI graphs from being opened properly.

---> C00772 - Collection Services - Bug Fix Request
In CSI, the machine pool sizes and rates graphs had an incorrect label for the machine page faults per second field (was listed as machine pool size)

---> C00772 - General - Enhancement
When defining drilldowns inside of an SQL statement's comments and you are filtering on multiple selections, you can define a field called <<MAXXYZ>> where the MAX part of the field name indicates that the last record of the selection should be used when substtituing the parameter marker. If the parameter marker does not start with MAX then the 1st record of the selection is used to retrieve the replacement value.

You can also define filter using a parameter marker like FIELD <<IN_FIELD>> where the replacement output will create an IN clause taking the value for FIELD from every row selected. FIELD IN(1, 2, 3, 4, 5)

---> C00772 - Job Watcher - General
In the start monitor commands (STRJWMON, STRPAMON, STRDWMON) removed the check to see if the JW subsystem was active since it was not needed.

---> C00771 - Installation - Bug Fix Request
In the install wizard, when checking to make sure the JW or PEX subsystem description exists and it already exists, ensure that the subsytem pool number is set properly to the value specified in the wizard by the user by issuing a CHGSBSD command on it.

---> C00771 - General - Enhancement
Added a new preference to the misc tab (default off) called "Add column table names to the Record Quick View window"

Also on the Record Quick View window the column's table name will be show next to each column's name if it is known (not derived from an expression).

---> C00771 - Job Watcher;Pex Analyzer - Bug Fix Request
Fixed a possible stored procedure conflict between PEX analyses and JW summaries.
Both used a procedure in QIDRGUI called QIDRCRTFM but with different implementation. The Pex procedure uses the same name now, but the JW summary uses a procedure by a different name QIDRCRTF.

---> C00771 - General - Enhancement
The dynamic drill down support (added from sql comments) provided in the previous client has changed so that all the drill downs created contain all the comments from initial query at the top before the sql statement to run.

---> C00771 - General - Bug Fix Request

Allow selections to be made in graphs that show negative values.

---> C00771 - Job Watcher - Usability
In Job Watcher, for summarized collections at 5.4+, the selected thread/job/user wait buckets graph will allow the SSD Improvements Estimator to be used now without needing to modify the SQL statement. Previously required lines were commented out because they require the summary files to be recreated. Older collections that have not done this yet will need to be resummarized (gap file step).

---> C00771 - General - Bug Fix Request
The dynamic drill downs embedded within the SQL statement added in the previous client should work now in Job Watcher, CSI, DW.

---> C00771 - Collection Services - Enhancement
WLE import has changed to coincide with recent changes to the WLE website so that the OS Type passed to WLE is now "IBM i" at 5.4 and higher instead of "OS/400".

---> C00771 - Collection Services;Job Watcher - Usability
The SSD Improvements Estimator menu will no longer show up on the graph popup menu if it cannot be used (required fields don't exist).

---> C00771 - General - Bug Fix Request
Fixed a possible crash in the Component views related to the improper clean up of property sheets attached to a folder in the tree.

---> C00771 - Collection Services - Enhancement
In CSI 5.4+, make the the import data to WLE option accessible from a CSI graph if a time range has been selected. (the numbers provided to WLE will be based on the selected time period).

---> C00771 - General - Enhancement
Add support for the following new graph types:

Step
Area
Vertical overlapping bar
Horizontal overlapping bar

These graph types can be set on the general tab of the graph definition properties.

---> C00771 - General - Enhancement
In the record quick view, show the table alias name before the column name if it is known.

---> C00771 - General - Enhancement
When signing on to a system, iDoctor will now make use of the configured signon setting defined in System i Navigator (Properties -> Conneciton tab for a system). You can use the options like "Use Windows username and password" to avoid needing to signon through iDoctor if your Windows password matches the username and password of the System i you are connecting to.   Also makes use of the Client Access password cache to avoiding prompting for a password unless needed.

If you still wish to be prompted for a password everytime you start iDoctor, just make sure you have the "Prompt every time" option set within System i Navigator.   (Properties->Connection tab).

The 1st time you transfer a file to another system you will typically be prompted for a password to the remote system now since it was not provided at startup.

---> C00771 - General - Usability
If the connection is dropped reestablish connections automatically. This should be a lot better than before although you most likely will need to close and reopen any graphs.
If you still have failures or crashes with this client, let me know what you were doing in iDoctor when you tried to reestablish connections at idoctor@us.ibm.com.

---> C00770 - Collection Services - Bug Fix Request
In CSI the wait graphs -> disk page faults rankings graphs were not working correctly if accessed via the collection menu.

---> C00770 - General - Bug Fix Request
Fixed a memory leak with the cache for collection popup menus.

---> C00770 - General - Bug Fix Request
Fixed possible crash when using the Generate Reports option with the previous client.

---> C00770 - General - Enhancement
Added support for drill downs to be dynamically added into iDoctor using comments in the SQL statement. This works for both user-defined and iDoctor supplied table views (should we desire to add some this way in the future).

Here is an example of the syntax required to add a drill-down menu called "QAYPEPROCI queries":

 -- Examples of "Dynamic drill downs"
-- based on data (row selections) in the current query
-- take advantage of current substitution parameters << >> in iDoctor SQL editor
-- Filter on any field in the next drill down like this WHERE FIELD1 = '<<FIELD1>>'
--00--HEADER QAYPEPROCI queries
--00--END
--01--HEADER List <<LIBNAME>>/QAYPEPROCI rows by Names
--01--
SELECT * FROM <<LIBNAME>>/QAYPEPROCI ORDER BY QPRPGN, QPRMNM, QPRPNM
--01--END
--02--HEADER Sum <<LIBNAME>>/QAYPEPROCI by Header Type
--02--SELECT COUNT(*) AS CNT, QPRHDR FROM <<LIBNAME>>/QAYPEPROCI GROUP BY QPRHDR ORDER BY CNT DESC
--02--END
--03--HEADER List all <<LIBNAME>>/QAYPEPROCI Type MI rows
--03--SELECT * FROM <<LIBNAME>>/QAYPEPROCI WHERE QPRHDR = 'MI' ORDER BY QPRPGN, QPRMNM, QPRPNM
--03--END
--04--HEADER Sum <<LIBNAME>>/QAYPEPROCI Header Type MI by name
--04--SELECT COUNT(*) AS CNT, QPRPGN FROM <<LIBNAME>>/QAYPEPROCI WHERE QPRHDR = 'MI' GROUP BY QPRPGN ORDER BY CNT DESC, QPRPGN
--04--END
--05--HEADER List all <<LIBNAME>>/QAYPEPROCI Header Type MI rows with non-Q names
--05--SELECT * FROM <<LIBNAME>>/QAYPEPROCI WHERE QPRHDR = 'MI' and SUBSTR(QPRPGN,1,1) <> 'Q' ORDER BY QPRPGN, QPRKEY
--05--END
--06--HEADER Sum <<LIBNAME>>/QAYPEPROCI Header Type MI by non-Q names
--06--SELECT COUNT(*) AS CNT, QPRPGN FROM <<LIBNAME>>/QAYPEPROCI WHERE QPRHDR = 'MI' and SUBSTR(QPRPGN,1,1) <> 'Q' GROUP BY QPRPGN ORDER BY CNT DESC, QPRPGN
--06--END
--07--HEADER List <<QPRKEY>> occurrences in <<LIBNAME>>/QAYPEIAD

--07--SELECT * FROM <<LIBNAME>>/QAYPEIAD WHERE x'<<QPRKEY>>' in (QIAKEY1,QIAKEY2,QIAKEY3,QIAKEY4,QIAKEY5,QIAKEY6,QIAKEY7,QIAKEY8,QIAKEY9,QIAKEY10,QIAKEY11,QIAKEY12,QIAKEY13,QIAKEY14,QIAKEY15,QIAKEY16) ORDER BY QRECN
--07--END

---> C00770 - General - Bug Fix Request
After opening property sheets (record quick view/query properties) from a table view or graph a crash would occur if the user closed the table or graph view while leaving the property sheet open and then tried to navigate to other tabs in the property sheet.

The fix for this problem is to close the child property sheets of a table or graph view when the user tries to close the table or graph. When this happens the user will be given a message indicating that N dependent property sheet views are open and must also be closed. At this point the user can continue and close the table or graph and all dependent property sheets or cancel the operation.

ThIs confirmation only happens when working with tables and graphs in the Data Viewer and not from the iDoctor main window.

Also add a new preference on the Confirm tab that controls whether or not a confirmation message is shown.

---> C00769 - General - Usability
Right-clicking a collection is faster now. Added caching of collection menus so that if the same type of menu is opened twice the response time will be greatly reduced.

---> C00769 - Collection Services - Enhancement
In CSI, at 5.4 and higher, added a function that allows the user to take the selected collection's information and import it into Workload Estimator.

Use the new menu on a collection called "Launch Workload Estimator".
A webpage will be shown with the average CPU and disk statistics for the collection. Upon pressing Submit the data will be sent to Workload Estimator for further analysis.

This is very similar to the function in PDI called "Size next upgrade" except for it works at 5.4/5.3 without converting the data to 6.1 first.

---> C00768 - General - Enhancement
Add an option to the iDoctor Report Generator window that allows the user to filter the list of available graphs by the graph folder (top level) such as "wait graphs, cpu graphs, user-defined graphs, etc"

Because of this change also needed to modify how items were selected, instead of using checkboxes, you now need to select and press the Add Selected buttons to add reports to the included list.

---> C00768 - Dir Dump - Enhancement
In the DirDump component the server-side output files, user-defined queries and user-defined graphs folders are now available under the collection. You can also define user-defined reports now.
Also removed the Start DirDump menu option since it was not implemented and there are no plans to add this function.

---> C00768 - Collection Services;Disk Watcher;Job Watcher;Pex Analyzer;External Storage;Dir Dump;Plan Cache Analyzer;Memory Watcher;NMON Investigator - Enhancement
When right-clicking a collection in most iDoctor components, user-defined graphs and user-defined queries menus will be available to allow you to quickly access any user-defined queries and graphs you have defined.

---> C00768 - General - Usability
Removed some unnecessary code called (needed for the server-side output files folder but not for user-defined) when building the list of user-defined queries and graphs.

---> C00768 - General - Usability
When building a list of objects in a component view made a change so that any CHKOBJs done are never done twice for the same object and/or file/member. The main spot where you may notice a big performance improvement is when listing the user-defined queries and user-defined graphs folder contents.

---> C00768 - General - Enhancement
Add user-defined reports to the list of reports available from the iDoctor Report Generator

---> C00768 - Job Watcher - Enhancement
In JW, added new situational analysis reporting options similar to the call stack reports based on the selected situation.
If coming from the Detail reports -> Situation details report or the Situations tab on the Interval Summary property pages, right-clicking a situation will now provide a new menu called "Selected situation" with the following reports available:

Total occurrences: all intervals
Total occurrences: this interval
Total occurrences: by job/task
Total occurrences for this job/task: all intervals
Jobs with this situation: all intervals
Occurrences by interval
Occurrences by interval for this job/task

From these you can double-click a job/thread to get into the call stack or perform further drilldowns.

---> C00768 - Collection Services - Enhancement
Added pool size and pool faulting rate graphs to CSI under the new folder "Memory pool graphs".
The following set of graphs were added:

Memory pool consumption (by percentage)
Memory pool sizes (by percentage)
Memory pool sizes
Memory pool faulting rates
Memory pool activity levels
Transitions to ineligible state
Machine pool sizes and rates

Support was also added to allow multiple collections to be graphed at the same time using these graphs (although adding many is slow)

---> C00767 - General - Enhancement
In the connection list view the menu shown when right-clicking a system now includes an option called Delete Obsolete Analysis Files with 4 suboptions (JW/CSI/PEX/All). This allows you to remove any iDoctor analysis files that are no longer used by the GUI or in the case of PEX Analyzer have been replaced by SQL-based analyses.

---> C00767 - General - Usability
Make the Filter libraries values saved specific for each system rather than use the same filter for all systems.

---> C00767 - Collection Services - Enhancement
Made a change to the CSI summarization process so that the summary files now allow null data to be added to them. This was due to a rare situation where Collection Services did not collect data for 2 intervals rendering all data after that point unavailable to the user.

---> C00767 - General - Enhancement
Add a new command called QIDRGUI/DLTOLDSUM at V5R4 and higher that will remove all iDoctor obsolete summary and analysis files from the current system. The command includes a collection type parameter (default value *JW) with possible values of:

*JW - Job Watcher (CRTWCHSUM output files)
*CS - Collection Services Investigator (CRTCSSUM output files)
*PA - PEX Analyzer classic analyses (except taskswitch)
*ALL

A menu option will be added in the GUI on the Connection list view to run this command.

---> C00767 - Collection Services - Enhancement
Add support for log page 33 and 34 for DS6K/DS8K external storage analysis to CSI at 7.1. This includes link and rank statistics. If file qapmxstgd is found, an analysis will be available by right-clicking the collection to process the data. After the analysis runs, the data will be available under folder External storage link and rank statistics.
Currently this folder contains reports that show the raw data parsed out of the file as well as files that show the delta values between intervals. A subfolder "By time interval" is also available to graph the link data over time.

The latest server builds (QIDRGUI library) must be installed.

---> C00766 - General - Enhancement
On the Connections View added a new menu called "Load iDoctor Stored Procedures" that will load all iDoctor stored procedures on the desired system that are needed for the new QIDRGUI/STRIDRSUM or QIDRGUI/RSTIDRDTA commands.

---> C00765 - General - Enhancement
Created a new iDoctor command at 5.4/6.1 in library QIDRGUI called RSTIDRDTA. (Restore iDoctor Data)

This command restores the iDoctor data found in the given save file to the desired library and afterwards will optionally run the STRIDRSUM command.

---> C00765 - General - Enhancement
Created a new iDoctor command at 5.4 and up in library QIDRGUI called STRIDRSUM (Start iDoctor Summaries).

This command will start all iDoctor summary and analysis processes (with the default parameters used by the GUI) for any collections found in the library. It will also expand any *MGTCOL objects found in the library into the same library if this has not yet been done.

This command requires that the stored procedures created by the GUI on the system exist in QIDRGUI. If some do not exist, then some analyses will not be performed. Typically these get created by summarizing/analyzing something once on a system using the GUI.

In the future a "load all stored procedures" option will be added to the GUI from the Connections view to make it easy to load all the possible stored procedures used by iDoctor onto the desired system..

---> C00765 - Job Watcher - Usability
In JW 5.4+, if a collection is summarized while it is still active the summary data will contain less data than the regular JW data. Change the Summarized flag value of 'Yes' to 'Yes (intervals missing, must resummarize) to indicate that data is missing and a resummarization is required to see it all.

---> C00765 - Job Watcher - Usability
For Jw 5.4+, modified the check for summary files to include the newer QAIDRJWANL_DTL_<<MBRNAME>>. This file is now required for the collection overview time signature graph.

---> C00765 - Collection Services - Bug Fix Request
In CSI 7.1, the wait graphs were not showing the wait bucket names and colors correctly in the Y axis.

---> C00765 - Job Watcher - Usability
In JW 5.4+, if a collection is summarized while it is still active the summary data will contain less data than the regular JW data. Change the Summarized flag value of 'Yes' to 'Yes (intervals missing, must resummarize) to indicate that data is missing and a resummarization is required to see it all.

---> C00765 - General - Enhancement
Add a 4 hour time grouping to the clock icon.

---> C00764 - General - Bug Fix Request
Changed the authorities on the iDoctor report databases shipped so that Windows users with restricted access user accounts no longer get several database access authority errors at startup.

---> C00764 - Collection Services - Enhancement
In Collection Services Investigator (CSI) at 5.4 and higher, added 2 new graphs to the Communication graphs folder:
Average IOP utilizations, and
Maximum IOP utilizations

These graphs show IOP utilizations broken down by disk and comm IOPs over time.
Let me know if you have interest in seeing drilldowns to rank the IOPs for a time period so you can see which IOPs were busy vs which were not..

---> C00764 - Job Watcher - Bug Fix Request
Fix added for the 5.4 JW situational analysis overview graph's SQL. It was only showing a max of 1 situation per type.

---> C00763 - Job Watcher - Enhancement
For 5.4+ JW, added a new graph under the wait graphs folder called
Situational analysis overview time signature

This graphs show the number of situations that occurred for each type over time.

Collections that were summarized with older builds may need to be resummarized (just the create situations file part) so that the required table QAIDRJWANL_DTL_<<MBRNAME>> is created.

---> C00763 - Collection Services - Bug Fix Request
In CSI fixed a potential problem where the graphs wouldn't work if the aliases that were previously created were now invalid.

---> C00762 - Collection Services - Enhancement

Graphing multiple collections in CSI at 5.4 is available now if the collection has been summarized. However this does not work for all types of graphs such as the wait graph rankings by job/thread. Previously this only was available at 6.1+.

---> C00761 - General - Enhancement
All graphs now have new options available in the popup menu under a submenu called "Duplicate".

"Duplicate -> using a Temporary SQL View" will reopen the graph using an SQL View defined in QTEMP over the existing report. This can be used to simplify the SQL statement, allowing the Query Definition interface to work against nearly all graphs defined in iDoctor. The only exception are graphs defined over multiple collections, since aliases are used in those SQL statements and a view can't be created over an alias. The Query Definition interface doesn't work against some complex queries so this action will allow it to work properly.

"Duplicate -> as Table" replaces the previous Duplicate as table option that creates a table view over the current graphs data into a new Data Viewer or an existing one.

"Duplicate -> as Graph" allows you to reopen the same graph in a new Data Viewer or an existing one. .

---> C00761 - Collection Services - Enhancement
In CSI 5.4+, the collection overview wait graphs now show batch and interactive CPU utilization on the second Y axis.

---> C00761 - General - Bug Fix Request
When listing the threads for a job within the Active jobs folder, the run priority field was shown twice.

---> C00761 - Job Watcher - Usability
When saving (or transfering) a library via the GUI, high data compression will now be used. Previously the parameter was not specified on the savlib command.

---> C00761 - Job Watcher - Enhancement
When starting a JW monitor from the GUI, added a check to make sure the monitor library's collections (if any exist) are of the same OS release as the partition.

---> C00760 - General - Bug Fix Request
Identified a rare but annoying crash when viewing/scrolling graphs from within the graph control code. Had to do with the mechanism used for the caching of pen objects used for drawing.

---> C00760 - Collection Services - Enhancement
The disk graphs in CSI 5.3 or 5.4 where qapmdisk file does not contain the response time bucket fields has an additional graph available: I/O rates totals with cache statistics.

---> C00759 - Disk Watcher - Usability
In Disk Watcher, the path ranking graphs are no longer segregated by ASP.

---> C00759 - Collection Services - Usability
In CSI when working with disk graphs, made a change so that ASP segregation only occurs with the by time interval, and by disk unit graphs. All other groupings no longer filter by ASP (default is ASP = 1).

---> C00759 - Collection Services - Bug Fix Request
In CSI 5.4 for the disk graphs folder, if QAPMDISK does not contain the 6 response time buckets the 5.3 CSI disk graphs are shown instead. This is a reduced set of graphs but at least will allow some disk information to be shown.

Also changed the collection status to no longer mention the qapmdisk file is incompatible with the disk graphs folder.

---> C00759 - Collection Services - Bug Fix Request
In CSI 5.3 disk graphs, fixed some drill downs that didn't work right. (selected disk unit over time, or disk path over time).

---> C00759 - General - Enhancement
The math mode icon in the Data Viewer now works for the table view beneath each graph.

---> C00759 - Job Watcher - Bug Fix Request
The refresh interval icon in the JW interval detail pages was messed up with the previous client.

---> C00759 - Installation - Bug Fix Request
Fixed some issues if the user used the Enter key on the keyboard to move through the install. Holding it down on the Wizard's summary screen would cause the finish button to be hit multiple times causing the install to fail. Pressing the enter key on the install status window would also abort the install prematurely without warning.

---> C00758 - Collection Services - Bug Fix Request
Made another change to the CSI disk graphs read cache hit % for ds8000s. In one case the numbers exceeded 100% but hopefully this is correct now for all environments.

---> C00758 - General - Enhancement
Added a new button on the toolbar in the Data Viewer (far right) that allows math to be performed over cells from the selected rows (in a flyover) as the user moves their mouse pointer over the data. The options are None (normal flyover behavior), Sum, Average, Min and Max, Percent of, Delta (current - prior) and Delta (prior - current).

Note: This feature works off of data loaded into the table view's memory cache. If you try to select rows outside of the cache and they aren't loaded in you will get a warning in the flyover message. Right now a maximum of 2000 rows can be loaded. You load this data in by scrolling through the desired rows.

---> C00757 - Job Watcher - Bug Fix Request
The Create Job Summary option in JW was broken with clients 755 and 756. An SQL error on field QCOUNT05 would be given. This was due to the new fields added to the qaidrjwgap_<<MBRNAME>> file.

---> C00757 - Collection Services - Bug Fix Request
In the Ext Stg cache statistics graphs for a single unit/asp over time, the flyover fields disk drive type and disk type descirption were showing up blank previously.

---> C00757 - Collection Services - Bug Fix Request
Fix an SQL error when trying to run the CSI disk graphs drilldowns for a single unit/path/ioa/iop or asp over time at 5.4. and 6.1

---> C00757 - Collection Services - Bug Fix Request
In CSI, made a slight name change to the external storage cache statistics graphs. Graph names beginning with "I/O" have changed to "IO" in order to not conflict with the disk graphs when performing drilldowns. The default drilldown selections are made based on the name of the graph and too often it was possible to drilldown from the regular disk graphs to the external storage disk graphs and vice versa.

---> C00757 - Collection Services - Bug Fix Request

Fix a bug with the CSI disk graphs, the read cache hit % is half what it should be when dealing with multipath DS LUN data. Also changed the color of the read cache hit % on the corresponding external storage graphs to match the color used on the regular disk graphs.

For the new CSI external storage cache statistics graphs, several fields like I/O rates, counts were double what they should have been when dealing with multipath DS LUN data. Also removed the External storage cache statistics IOA and IOP ranking graphs due to accuracy issues.
The empty Disk type graphs sub folder has also been removed.

The analysis data must be regenerated in order to fix the doubled rates and counts in the external storage cache statistics graphs.

---> C00757 - Collection Services - Enhancement
In CSI, under the wait graphs folder, added a set of graphs "Disk page faults rankings" that allow the user to rank threads, jobs, generic jobs, etc by ones having the highest amount of disk page fault time.

---> C00757 - Job Watcher - Enhancement
Create a situational analysis detail file that contains all the job hits for every interval.

The file is named QAIDRJWANL_DTL_<<MBRNAME>> and can be viewed within the JW server-side output files folder,

Double-clicking an interval and viewing the Situations tab will now show all jobs experiencing the situation within the interval instead of only 2 max. If the updated file has not been created yet, then the panel will look the same as previously.

However the Detail reports -> Situation details report has changed to use the new file. This means existing collections will need to be resummaried to recreate the situations file.

---> C00757 - Collection Services - Enhancement
In CSI 5.4 and up, added a disk configuration folder under the collection which contains information about the ASPs, IOPs, IOAs and units.
Information about the IOAs including the read/write cache sizes is included. Currently two reports are provided showing the same disk configuration data. One is a flat table, the other is a tree. The tree provides counts and percentages of the units/IOAs/IOPS/ASPs within each prior level grouping.

A new file qidrgui/qaidrccins is required for these reports, so either the latest server builds need to be put on or from the GUI the user can right-click the CSI collection and choose "Analyses -> Run Create disk configuration QAIDRCCINS table". This only needs to be done once per system.

---> C00757 - General - Enhancement
On the Import Graph Definitions and Import Query Definitions Windows, added a button to delete user-defined queries/graphs stored on the server (files QAIDRGPH09 and QIDRSQL05) or from the selected local database.

---> C00756 - Job Watcher - Bug Fix Request
When ending or holding a JW monitor, occasionally CPF4102 would be issued by QIDRENDJW1 and not handled properly.

---> C00756 - Job Watcher - Bug Fix Request
When starting a monitor, the program that checks if the subsystem is active could sometimes fail, indicating the subsystem was not active when in fact it was.

---> C00756 - Job Watcher - Bug Fix Request
Fix a bug that would occur if more than 256 JW collections existed in a library or monitor's library: the list of collections would not be viewable.

---> C00755 - Job Watcher - Bug Fix Request
In Job Watcher, when drilling down on a summarized collections, several of the drilldown submenus (job user, current user, generic job, etc), would list the wrong menu description as the default option within each group (1 level previous in the list).

---> C00755 - Job Watcher - Bug Fix Request
With the previous client, some Job Watcher drilldowns did not work due to changes made for the SSD Improvements Estimator. This was due to the Qcount05, and QCOUNT06 fields not existing in the qaidrjwgap_<<MBRNAME>> summary file and the attempted changes to get them from the Qapyjwtde caused some drilldowns to fail (like selected generic job over time). These fields have been added to the summary file but rather than add them to the SQL and break all summarized collection drilldowns, I've added a comment and a required line which must be uncommented in order to utilize the SSD Improvements Estimator in summarized JW reports for a single job/thread grouping over time.   Of course the collection must also be resummarized (build gap file step) in order for this to work.

In summary: if the SSD Improvements Estimator fails (which is normally due to COUNT05, COUNT06 not existing in the query), check the SQL for any comments and further action to make it work.

---> C00755 - Job Watcher - Usability
Temporarily disabled the ability to collect monitors over multiple partitions at once in preparation for upcoming external builds. The function is not yet ready for customer use.

---> C00755 - Job Watcher - Bug Fix Request
The first time starting a monitor on a system, ensure that the IBM-supplied definitions get loaded and shown in the lists on the Basic Options page of the Start Monitor Wizard.

---> C00754 - Collection Services;Job Watcher - Enhancement
Made improvements to the SSD Improvements Estimator. Run time improvements are more apparent now if looking at just a single job as wait times are moved to the left.

---> C00754 - Installation - Usability
Changed the error handling in the install when FTPing the save files to the system fails. More information should be returned in the log now.

---> C00754 - Job Watcher - Bug Fix Request
In JW on the interval details SQL Statement tab - improve the parsing to handle host variable parameter markers where ":H :H" is equal to 1 parameter.

Also host variable values that look like strings will now be enclosed in single quotes when parsed back into the sql statement.

---> C00753 - Collection Services;Job Watcher - Enhancement
Added a function available from JW and CSI graphs called "SSDs Improvement Estimator" which modifies the current wait bucket graph based on the input provided showing how the run/wait signature could look if I/Os were handled by SSDs. This function is only accessible if the IBM internal components enabler is applied to the PC.

Scrolling the graph or refreshing the graph will cause the estimation to be lost until it is applied again. The changes only effect the graph data loaded into memory, not the underlining table data or SQL statement.

---> C00753 - Job Watcher - Bug Fix Request
Fix a bug with the Job Watcher summary, the thread list file wasn't being created after the change in the previous client.

---> C00753 - Collection Services - Enhancement
Some graphs in the old csi graphs folder in CSI, do not have their colors set properly. This causes the legend to not work right.

---> C00753 - General - Bug Fix Request
Fixed an issue with the transfer to testcase.boulder.ibm.com function for a customer where extended passive mode needed to be turned off using the sendepsv 0 command before transfering the data. A checkbox was added to the GUI which disables extended passive mode. This can also be used for other servers where this is needed.

---> C00752 - Job Watcher - Bug Fix Request
Customer reported a situation when running a monitor, collections numbered 900-999 would not be deleted if the max collections was 99. This issue has been fixed for V5R4 and up for all 3 monitor commands (STRJWMON/STRPAMON/STRDWMON), the latest server builds are required..

---> C00752 - Collection Services - Bug Fix Request
For the CSI disk graphs, the BY IOA grouping SQL had a problem which caused not all IOAs to be shown and some of the IOAs would be double what they should have been.

---> C00752 - Job Watcher - Usability
In the 540+ JW summary, the create thread list option has been embedded within the create interval summary file option, as it is always required there was no need to have it be a separate option.

---> C00752 - Disk Watcher;Job Watcher;Pex Analyzer - Usability
Ending or holding a monitor from the GUI will now run the command in the remote command status view in order to not tie up the GUI.

Also with the latest server builds at 540 and up, the ENDPAMON will submit a job to do it's processing. This command runs ENDPEX so on large collections will take some time.

---> C00751 - Collection Services - Enhancement
Add a series of graphs for the CSI external storage data (file QAPMXSTGV). This is generated for DS8K/DS6K external storage boxes.
To access these graphs, first run the option under the Analyses menu when right-clicking the collection.
Once the required files are built, the graphs will appear under the "External storage cache statistics" folder.

If you previously ran this analysis, you will need to rerun it in order to pick up new files/fields created for the graphs.

---> C00750 - Job Watcher - Bug Fix Request
The help text for the STR monitor commands (STRJWMON/STRPAMON/STRDWMON), incorrectly stated that the monitor name can be 8 characters long when in fact it now can only be a max of 7 characters long. This was fixed for 5.4 and up and the latest server builds are required.

---> C00750 - Job Watcher - Bug Fix Request
Drilldown in Job Watcher on a bucket that is under the "Advanced" menu under Thread rankings -> Wait Graphs would not show the default drill down option at the top of the menu if the Preferences -> Job Watcher -> Display advanced reporting options checkbox was not checked

The default menu drilldown will now work and the advanced menu option buckets will be shown now even if this preference is unchecked.

The preference now only controls the display of the "Detail reports" menu option and some Wait graphs options like wait bucket counts and wait bucket contributing jobs graphs.

---> C00750 - Job Watcher - Enhancement
In V5R4+ JW the SQL provided for the "Detail reports" menu options have been modified to include comments and examples for how you might modify them to suite your own needs.

Also for collections that have been summarized, the detail reports menu options have been added to the Job rankings graphs. This allows you to filter these reports by the selected job (groups) from the previous chart.

---> C00750 - Collection Services;Job Watcher - Enhancement
In the Search window for CSI/JW, added a Current wait search option that allows the user to search for job records that experienced the desired wait type (or in JW also allows searching on the wait object name found at the end of each interval).

---> C00750 - General - Bug Fix Request
Made fixes to the QIDRGUI/ADDIDRUSR and QIDRGUI/RMVIDRUSR commands at 5.4
They will no longer fail if the QAPEXFTR file doesn't exist yet in QUSRSYS.
It will also update the required files/commands needed to run super collections.

---> C00750 - Job Watcher - Usability
Graphs in Job Watcher that show reads and writes were using inconsitent colors. Now reads are consistently a red color and writes are consistently a green color. This matches the colors used in the other components.

---> C00750 - Collection Services - Enhancement
In CSI, added disk graphs groupings by the disk type field (DSTYPE) in file QAPMDISK. Look for the by disk type folder under Disk graphs to access this.

Also added filtering support via the Change SQL Parameters interface. You can modify the list of disk types (all initially included) in order to filter the data or to perfom comparisons of one type of disk with another.

Also added a parameter shown on the Change SQL Parameters interface to set the last character of the disk model field (DSMDLN).

In addition added a disk type description flyover field to the disk graphs to offer extra information about what the type of disk the dstype number value applies to.

---> C00750 - Job Watcher - Enhancement
In JW 5.4+, add a new situation that identifies jobs having INELible to run thread status.

---> C00750 - Collection Services;Job Watcher - Enhancement
Added a drop down list with prior search history in the Job Watcher and CSI Search and Create Job Summary windows.

---> C00749 - Disk Watcher;Job Watcher - Enhancement
By customer request, I've added an option on the final page of the following wizards that allows you to specify the USER parameter on the SBMJOB command when starting monitors and some types of iDoctor collections:

Start iDoctor Monitor Wizard
Start Disk Watcher Collection Wizard
Start Job Watcher Collection Wizard
Start Super Collection Wizard.

Press the Submit job options button to access this capability.

---> C00749 - Collection Services - Bug Fix Request
Disable variable width mode for CSI disk ranking graphs.   The option was available but didn't work (not possible with graphs that aren't time-based).

---> C00749 - General - Bug Fix Request
When running a "WITH" query in the SQL editor, the file/member overrides in the SQL Statement were not being set properly. This caused some unexpected behavior.

---> C00749 - Job Watcher - Bug Fix Request
Fix a bug with the super collection support, the collection job would sometimes not ever get set in the qaidrsuper control file.

---> C00749 - Job Watcher - Bug Fix Request
The change user function on the components window did not disconnect the previous user profile and end the jobs it was using like it should have.

---> C00749 - Job Watcher - Enhancement
Enable the capability to do a Collection Search in JW over multiple collections at once.

---> C00749 - Collection Services - Bug Fix Request
The ability to graph multiple CS collections at once is now working correctly again in CSI with latest builds.

---> C00749 - General - Enhancement
Make the find window non-modal for easier usability. Added find history, to keep track of previous searches (with a remove button to remove the desired history entry from the list).
Also updated the list of columns to include an "All columns" option..

---> C00749 - Collection Services;Job Watcher - Enhancement
When searching over multiple collections in JW/CSI, the browse options will also include data from all selected collections. Also updated the window to display the selected collections (up to a max of 5 shown) , the total intervals and the total time range of the desired collections.

Note: In JW this only works (for multiple collections) if the collection summary option to create the aliases has been performed.

---> C00748 - General - Enhancement
Added a new option to the list of connections called "Remove iDoctor stored procedures". This function will remove all iDoctor stored procedures from library QIDRGUI, which is sometimes helpful if one or more of the stored procedures are not registered properly. The GUI will recreate them as needed.

---> C00748 - General - Enhancement
Added an Edit menu to the list of connections that allows you to modify the connection settings.

---> C00748 - General - Bug Fix Request
Some customers were reporting stored procedures shipped in library QIDRGUI were not being registered properly and as a result they would fail when trying to run Job Watcher summaries. The simple solution to this is to no longer ship the stored procedures in QIDRGUI (which I've done) with the latest server builds.

This means at startup the 1st time iDoctor is ran from the GUI there will be anywhere between 3-20 second delay while stored procedures are created on the system via the GUI.

This applies to QIDRGUI library at 5.4+.

---> C00748 - Job Watcher - Bug Fix Request
Starting a super collection would not work properly if the Collection Services option was not checked. The super collection job would not be started due to the call to QIDRSSMON supplying the wrong number of parameters.

---> C00748 - Collection Services - Bug Fix Request
In CSI, when opening a graph and then going to the properties -> query tab the filenames listed might show <<UNIONEND>> as part of the filename.

---> C00747 - Installation - Bug Fix Request
The GUI was not recognizing the install directory if the default was not used.

---> C00747 - Disk Watcher;Job Watcher;Pex Analyzer - Usability
The 2 messages shown when PTFs are not found, have been combined into 1 window, which allows you to continue anyway or cancel. The window also now includes the system name where the PTFs were not found.

---> C00746 - Collection Services - Enhancement
Add support for collection of CPU / disk / memory information for all partitions collected by IBM Director to CSI at V5R4 and higher to the System graphs (Director) folder. This includes the following graphs that each allow you to select 1 or 2 partitions in order to see the desired partitions over time:

System CPU utilization
System I/O counts
System I/O rates
System disk space
System memory

---> C00746 - Collection Services - Bug Fix Request
In CSI the clock icon was not working on the Thread wait time signature graph due to an SQL error. The graph only worked at the "Collected interval size".

---> C00745 - Job Watcher - Bug Fix Request
The split collection option from a Job Watcher graph (after making a time range selection) was not showing up..

---> C00745 - Job Watcher - Bug Fix Request
Occassionally when viewing a table and scrolling to the bottom the last row was not showing up.

---> C00745 - Job Watcher - Bug Fix Request
In Job Watcher when drilling down to the Thread wait time signature for XYZ graph, the graph title was showing up blank.

---> C00745 - General - Bug Fix Request
Fix a bug where the act of canceling a long running graph or table was not behaving properly and the client could end up hanging for a long time or crash.

---> C00745 - Job Watcher - Usability
On the JW Search window, the Include system tasks checkbox will now be checked by default.

---> C00744 - General - Bug Fix Request
The alternate views menu in the legend was not always showing the currently selected graph (in bold) in the popup menu.

---> C00744 - Job Watcher - Bug Fix Request
In Job Watcher picking an alternate view graph for a job rankings graph will no longer cause the description on the X axis to revert back to the parameter name <<OBJTYPEDESC>>.

---> C00744 - Job Watcher - Bug Fix Request
Some default drilldowns were missing from Job Watcher with client 743.

---> C00742 - General - Bug Fix Request
The Generate Reports function is working now if Firefox is the default web browser.

---> C00742 - Job Watcher - Bug Fix Request
In Job Watcher, on the create job summary's Thread totals and Job Totals reports, drilling down will no longer prompt the user for the desired member. The Thread totals and job totals output files now contains the 1st (mininum) and last member (maximum) found within the summary. The member used on the drilldown is the 1st one.

---> C00742 - Job Watcher - Bug Fix Request
In Job Watcher, splitting a collection will now invoke a stored procedure to perform the split. This will call the CPYWCHCMD to split the files, then run some SQL to properly set the collection size, last interval collected and start and end times in the qapyjwruni and/or qaidrjwri files.

---> C00742 - General - Bug Fix Request
The save collection option (from the tree side of a component view) was incorrectly showing the transfer collection window.

---> C00742 - Job Watcher - Enhancement
When saving a graph as a JPEG (or using the Generate Reports function), and a failure occurs during the save of an image, added error handling that will display the windows error message and return code for the reason of the failure (out of memory, bad filename, etc).

---> C00741 - Job Watcher - Bug Fix Request
Deleting a user-defined query or graph in Job Watcher could fail from the user-defined graphs/reports folder but would work from the manage interface.

---> C00741 - General - Bug Fix Request
If users set the bars per page in a graph > 100, it could the client to fail to draw the bars properly when scrolling and crash the client.

---> C00741 - General - Enhancement
You can now click the Save button on the toolbar of the Data Viewer to save the current graph AND legend to a JPG image.

---> C00741 - General - Enhancement
Added a preference to always show the date on time range graphs for the Collected interval size, 30 secs and 1 min time interval groupings. This option is called "Always show dates on time range graph X-axis" on the Miscellaneous tab of the Preferences interface.

---> C00741 - General - Enhancement

Add 15 min time interval grouping option to iDoctor.

---> C00741 - Job Watcher - Enhancement
In rare cases, users could be presented with Job Watcher graphs that would show a higher than expected amount of CPU queueing. This effected summarized and unsummarized collections.

The Job Watcher summary's gap file (QAIDRJWGAP_<MBRNAME>) has been changed as follows to handle a possible CPU queueing overreporting issue in the Job Watcher data:

1. When filling the wait times from Idle waits in file QAPYJWSTS into the gap file, the ENUM field will be used instead of the wait bucket field when determining which bucket the job was in for the interval's duration.
2. A new field BKTMIS has been added to the end of ther GAP file. It will contain a value of 1 if there was a enum/wait bucket mismatch in the QAPYJWSTS file for the current record generated in the GAP file. The only known possible mismatch involves the CPU queueing bucket and an enum for the last known wait before the job entered CPU queueing.

Sometimes the QAPYJWSTS file would report the job was in CPU Queuing during the interval which it was, but in most cases it wasn't for the entire interval's duration and usually CPU queuing was an (unknown much) smaller portion when compared to the wait duration. Therefore we are using the enum to determine the bucket to report the wait duration as.

The SQL used to graph unsummarized collections has also been changed as a result of this change.

August 12, 2009 – C00740

---> C00739 - General - Enhancement
If the Tip of the Day videos were not found on the PC when trying to display them, the GUI will ask the user if they wish to download them automatically.

---> C00738 - General - Enhancement
Any table view column that contains a number (0-100) followed by a % sign will now draw a shaded box (PDV-style) filling the cell's space by the appropriate percentage.

This feature currently just applies to the TPROF table views that show percent hits, but would also work for user-defined tables.

---> C00738 - General - Usability
Removed the menu option Query definition -> member selection as it was too confusing to go to the member selection window instead of the query definition window.

Changing the member will require opening a different collection, or using the Choose database members option from the SQL editor.

Also clicking query definition button on the toolbar will now show the query definition record selection interface instead of prompting the user for a member.

---> C00738 - Job Watcher - Usability
Made a change to the Job Watcher summary. The slowest part that does an update to the gap file to fill in user profiles, priorities is now turned off by default.   This part is not required and only increases accuracy of the job priority ranking graphs and the current user profile ranking graphs.

---> C00738 - General - Usability
In the legend, for the list of "Available fields", changed the color used for fields that have no non-zero values to grey. Previously in some windows schemes it looked like a selection.

---> C00738 - Job Watcher - Bug Fix Request
Drill down from a JW call stack report for a specific job was not working.

---> C00736 - Job Watcher - Bug Fix Request
Job Watcher user-defined tables and graphs that were created from a new SQL editor were not showing up in the user-defined graphs/tables folder.

---> C00736 - Job Watcher - Bug Fix Request
If 2 data viewers are open, and working with the call stack in the 2nd… performing a call stack report drill down would incorrectly open the report in the 1st Data Viewer.

---> C00736 - General - Enhancement
The hierarchical tree reports in the TPROF analysis are now cancelable (by closing the window) and provide progress information in the Data Viewer's status bar.

---> C00735 - Job Watcher - Bug Fix Request
In client 734, the save as option for a table was incorrectly showing the panel for the save graph definition instead.

---> C00734 - General - Usability

On the Save Graph Definition and Save Query Definition windows, IBM internal components are no longer listed unless the user is authorized to use them.

---> C00734 - Job Watcher - Usability
In the JW collection wizard, specifying a collection or library name and then loading a definition into the wizard will no longer cause the library or collection name change to be lost.

---> C00734 - Job Watcher - Usability
In JW, when loading a definition into the Start Collection Wizard, if the definition did not include the ending options parameter then the default end options will now be 1 GB or 1 hour (whichever happens first)

---> C00734 - Collection Services - Bug Fix Request
In CSI, the disk graphs write cache efficiency calculation was incorrect. Also if the current graph contains external storage disks then this percentage will be shown as 0 (it cannot be calculated).   It will only work for internal disks.

---> C00734 - General - Enhancement
User defined queries/graphs that are usable against V7R1 can be defined now .

Also user-defined queries/graphs are no longer VRM dependent. You no longer need to define your queries and graphs for each release you wish to run them on. When saved the GUI will keep track of the release the user-defined query/graph was defined for and display this information to you when listing the user-defined queries/graphs.

In addition, the version field on the Save Graph Definition and Save Query Definition windows has been removed.

---> C00733 - Disk Watcher - General
Updated the list of required DW PTFs on the website and in server builds V5R4 and higher.

---> C00733 - General - Bug Fix Request
Fixed a potential crash/hang when using the query definition interface on a table view and a long running query. Placing the mouse cursor over a column that would normally display a cell flyover while the query is running would cause the crash/hang.

---> C00733 - General - Usability
The member selection page of the query defintion has been removed.
Using the menu option Query definition -> Member selection will now reset the members used for the current report, and prompt you for the desired member(s) using the "Choose Database Members for Query" window.

---> C00733 - Website (external) - General
Updated text in section D) Payment Terms, in our Price Quotes document (SOW) sent to US Customers.

---> C00733 - Website (external) - General
The IBM Service Agreements included with the price quotes (for US Customers) has been updated.

---> C00731 - Job Watcher - Bug Fix Request
For V5R3 JW, when drilling down sometimes the graph would not show the job name correctly in the graph title.

---> C00731 - Job Watcher - Bug Fix Request
For V5R3 JW, the Wait Buckets tab within the interval details window was giving an SQL error.

---> C00730 - Disk Watcher - Enhancement
Enabled Disk Watcher for release V7R1.   Performed testing of graphs and drill downs, creating collections, monitors, and the trace summary function,

---> C00730 - Collection Services - Enhancement
Enabled Collection Services Investigator for release V7R1.   Performed testing of graphs and drill downs,

---> C00730 - Job Watcher - Enhancement
Created V7R1 Job Watcher builds.
All major functions tested (graphs, monitors, super collections, definitions, collection management) and they appear to be working. These were tested using driver 150.

Note: The classic JVM graphs have been removed at V7R1 since the classic JVM is no longer available at V7R1.

---> C00730 - Job Watcher - Enhancement
In V5R4 or higher JW, when summarizing collections, the same set of stored procedures will now be used for all releases, instead of using a different set of procedures for each release.

---> C00730 - Job Watcher - Bug Fix Request
In V5R4 or higher JW, for summarized collections fixed an SQL error with the State transitions rankings graph.

---> C00730 - Job Watcher - General
In V5R4 or higher JW, when building the list of collections in a monitor, if the stored procedure used to build the list fails for some reason, the client will build the list the old (slower) way instead.

---> C00729 - Job Watcher - General
In V5R4 JW or higher, when building the list of collections in a library, if the stored procedure used to build the list fails for some reason, the client will build the list the old (slower) way instead.

---> C00728 - General - Bug Fix Request
Fixed an issue where users may not be able to rerun a query successfully if the query contained an SQL created table and the query to that table failed, and then was later corrected and reran.   The data viewer would show no records were found when data should have been available.

---> C00728 - Job Watcher - Bug Fix Request
Fix a bug in the JW summary for 5.4 and higher. The situational analysis part would fail if the call stack was not collected.

---> C00728 - Installation - Bug Fix Request
Fix bug in the uninstaller for the client side, the desktop icon was not being removed.

---> C00727 - Installation - Bug Fix Request
Fix bug in installer that was causing the creation of an unnecessary (empty) folder for iDoctor on the desktop.

---> C00726 - General - Bug Fix Request
In the query definition interface, on the record selection tab, the list filter option would add an extra set of commas between values in the SQL statement.

---> C00725 - Disk Watcher - Bug Fix Request
When drilling down in Disk Watcher, the user could be prompted with a signon box unexpectedly.

---> C00724 - Job Watcher - Bug Fix Request
In the Job Watcher interval details panels replaced the label next to the Job or task name with one of the following depending on the TDE type.
Primary thread:
Secondary thread:
System task:

Previously the label always showed "Job information:".

Also placing the mouse over the Job or task button will show the taskcount value in the flyover.

---> C00724 - Job Watcher - Bug Fix Request
In the Job Watcher interval details panels, the LDIO, IFS, and transaction panels could show the data for the wrong interval when working with a job with multiple threads.

---> C00724 - General - Bug Fix Request
If the ODBC connection was using a QZDASSINIT job name instead of QZDASOINIT several parts of the GUI would not be usable.
Updated server builds (QIDRGUI library) and new client is required to fix.

---> C00723 - Disk Watcher - Enhancement
For V5R3 and up, deleting a disk watcher collection will now use the QIDRWCH/DLTDWCOL command instead of removing each Disk Watcher file member separately.

---> C00723 - Job Watcher - Usability
For V6R1 JW, increased the default value for the STRJWMON/STRDWMON command parameter MAXSIZE to 4 GB (4096 MBs).

---> C00723 - General - Enhancement
iDoctor will now properly retrieve and show the field descriptions from SQL created tables that had column names greater than 10 characters.

---> C00722 - Job Watcher - Enhancement
The JW interval details - wait buckets panel will now show information for intervals where CPU was not used from the QAPYJWSTS file.

---> C00722 - Disk Watcher - Bug Fix Request
Fixed a bug in the Start Disk Watcher Collection wizard that would occur if the target library doesn't already contain file QAPYDWRUNI.

---> C00722 - General - Bug Fix Request
Fix a possibility in Job Watcher that a user would be prompted to select a member for the query when displaying the call stack for a holder job.

---> C00722 - General - Bug Fix Request
If the monitor commands were ran from the green screen, using the *SAME parameter values for the collection size or collection duration parameters the collections submitted would be sumitted with an incorrect command string.

The latest server builds are required for this fix.

---> C00721 - Job Watcher - Usability
When displaying a holder call stack, fixed a possiblity that the holder jobname would not be shown if the holder did not have a QAPYJWTDE record.

---> C00721 - General - Bug Fix Request
If an IBM internal user, and the internal website's security certificate is not trusted, added a message at startup to indicate the client update checking failed for this reason. It must be installed via IE or Firefox in order for the automatic update checking to work for internal users.

---> C00720 - Collection Services - Enhancement
In CSI when working with data on a V6R1 system that has been converted from V5R3/V5R4 , the wait bucket graphs are now available.

---> C00720 - Collection Services;Pex Analyzer - Bug Fix Request
The V5R4 SAVPFRCOL command was not saving CS or PEX mgtcol objects correctly. If only 1 collection was specified but multiple CS or PEX MGTCOL objects existed in the library, all of them would be saved.

This also effects saving or transferring CS or PEX mgtcol objects to another system.

---> C00720 - Job Watcher;Pex Analyzer - Bug Fix Request
When starting a monitor on the Basic options screen, the maximum collection size parameter will be greyed out if creating a PEX monitor or a JW monitor at V5R4 or earlier since the parameter doesn't apply to those instances (the definition is used instead to define the maximum size of each collection).

Also fixed a remote command error when starting a PEX monitor if the default maximum collection size value of 1024 on the Basic options screen was changed.

---> C00720 - Job Watcher - Enhancement
Added a new JW situation that (attempts to) determine if concurrent write support is enabled or not.
Added a new JW situation that (attempts to) determine if journal caching is not being properly utilized.

The SQL for these situations may need tweaking still depending on how they perform with customer data.

---> C00720 - Job Watcher - Enhancement
When summarizing a V5R4/V6R1 JW collection, the summarization prompter window now has a new button called Preferences. Pressing this button allows the user to customize which situational analyses should be executed as well as the limits used for some of the situations.

---> C00719 - General - Bug Fix Request
With client 718, the Save -> View As... function would not correctly save all records (the records in the cache from the view would not be included).

---> C00719 - Job Watcher - Bug Fix Request
On V5R4M0 and V5R3M0 systems, the PTF checking would incorrectly check the PTFs against the list for V5R4M5 and V5R3M5.

New external builds will be released soon to correct this. (Only the client needs to be updated).

---> C00719 - General - Bug Fix Request
Defining a query in Job Watcher with wait bucket fields like this would cause the wrong wait bucket descriptions to be shown for wait buckets 10 and higher.

SUM(QTIME11), SUM(QTIME12), ...

The above previously would show bucket 1s and bucket 2s descriptions intead of 11 and 12.

May 18, 2009 – C00718

---> C00718 - Job Watcher - Bug Fix Request
In V5R3 JW, if the QAIDRJWRI file doesn't exist, the GUI will rebuild it from QAPYJWRUNI. This is needed in order for the GUI to display V5R3 collections that have been summarized. At V5R4 and higher, the GUI just relies on the QAPYJWRUNI file.

---> C00718 - General - Usability
When using the Save -> View As... (to CSV) file function the query will no longer be reran upon completion..

---> C00718 - General - Bug Fix Request
Column headings messed up when saving reports to CSV files that contain commas in the headings.

---> C00716 - Job Watcher - Bug Fix Request

On the iDoctor member selection window, using the Update all button will now change the collection parameter (internally). This will avoid SQL errors problems and future prompting of the member when viewing things like call stacks in the interval details.

One example of this is when drilling down from a job within the Thread totals report in the Create Job Summary folder. You will be asked which member this drill down should apply to. The selection will "stick" when working with further drill downs.

---> C00716 - Job Watcher - Bug Fix Request
Drilling down from a graph showing a holder ^ (caret indicator) for V5R4 and higher JW into the Holder-> Wait graphs menu would incorrectly show "Subsystem wait time signature" reporting options instead of "Thread wait time signature"

---> C00716 - Collection Services;Disk Watcher;Job Watcher - Enhancement
For all time-based graphs, the interval number will now always be included in the x-axis value on the flyover.

---> C00715 - Job Watcher - Bug Fix Request
In JW, right-clicking a job in a rankings chart to show the popup menu then right-clicking a different job in the same graph right away would cause the popup menu to show the wrong job name in the default drill down.

---> C00715 - Job Watcher - Bug Fix Request
Drilling down using the Display call stack menu in the Synchronized table view would show the call stack for the wrong job.

---> C00715 - Job Watcher - Bug Fix Request
Drilling down from a JW rankings chart in the Synchronized table view for a graph would show the incorrect job name as the menu for the default drill down.

---> C00715 - Job Watcher - Bug Fix Request
In V5R4 JW, when listing collections in a library, added a check to see if SYSPSTAT in QSYS2 exists before trying to run the stored procedure that lists the collections in a library (which the procedure requires). If the file is not on the system, the other/slower method to build the list will be used instead.

---> C00715 - Job Watcher - Enhancement
In the JW Create Job Summary function made the following enhancements/fixes:

- Several values in the Job totals by collection report were too high (QAPYJWPRC fields)
- Selecting multiple collections from a list and using the Create Job Summary menu would preselect those same collections on the Create Job Summary window.
- The same set of stored procedures are used for 5.4 and 6.1 now.
- The job name filter can now be the full 26 characters of the tdejobname field (previously only 10).
- When creating the summary over just 1 collection, the last 2 files that gives totals across all collections are no longer built.
- The thread totals (for all collections) report did not include thread ID.

---> C00715 - Job Watcher - Bug Fix Request
In Job Watcher drill down to a graph for a specific job/thread from an interval and a taskcount in some cases would not position the user to the interval given in the previous report.

---> C00715 - General - Bug Fix Request
On the main window, the View -> Remote Command Status View menu wasn't working.

---> C00714 - General - Usability
You can now scroll through graph legend vertically with the mouse wheel.

---> C00714 - General - Bug Fix Request
Scrolling the synchronized table view with the mouse wheel will not properly scroll the graph as well.

---> C00714 - General - Enhancement

Added a new menu called "Choose Database Members" to the SQL editor that clears the current member selections and brings up the member selection window.

---> C00714 - General - Usability
When defining a new query in the SQL editor you will no longer be prompted for the desired member every time you run the query (just once!).

---> C00714 - General - Usability
Updated the label shown on the transfer collection window when sending data to IBM to include the full file name format.

Also added .idr. to the filename as this will indicate the data is iDoctor data that should be restored and summarized/ready for analysis on the Ecurep debug system.

---> C00714 - General - Bug Fix Request
The Check PTFs function will fail more gracefully now if the system does not have the required level of the PTF checker program.

---> C00714 - General - Enhancement
Increased the max amount of data read by the FTP "Quick View" function on text files to 512KB from 10KB.

---> C00714 - General - Bug Fix Request
Fixed a bug in table views if the data contained more than 100 million records, no records would be shown. The maximum records that can be shown is 100 million. You can change the sort or filter to see the rest of the data.

---> C00714 - Collection Services;Disk Watcher;Job Watcher;Pex Analyzer - Usability
Running a summary in JW, CS, DW or an analysis in PEX Analyzer will now cause an automatic refresh of the window after the summary completes rather than requiring the user to manually refresh the window.

---> C00714 - Job Watcher - Usability
Trying to summarize a V5R3 JW collection on a V6R1 system would fail with program not found errors. Added an appropriate error message.

---> C00714 - Collection Services - Usability
You can now select multiple collections in CSI and use the Summarize option to do them all at once. Previously this would only perform the summary for the 1st collection.

---> C00714 - Collection Services - Usability
Trying to graph multiple collections at once in V5r4/v5R3 CSI would show an error message. Now it will just open the graph once for each collection as this is only available at V6R1.

---> C00714 - General - Enhancement
In the FTP component, if the connection is lost (times out) , it will be reestablished automatically without needing to retype the password.

---> C00714 - General - General
Since the call stacks under the Active jobs folder didn't work at V5R3, I've disabled it. You can still view statistics for jobs/threads however.

---> C00713 - General - Enhancement
All major collection types (PEX, CS, DW, JW) can now be saved using the Save... menu for the collection. You can save multiple collections or monitors as well.

After using this option a the new folder "Saved collections" will show a record identifying the save file which you can use to restore the data, or send the data to another system or ftp server directly from the desired save file.

---> C00713 - Job Watcher - Enhancement
In Job Watcher on the interval details tab, when stepping through intervals, the inner buttons will work how the outer buttons did previously (stepping to the next/previous interval), while the outer buttons will step to the next/previous CPU using interval.

In the case of the call stack or SQL tab the outer buttons will jump to the next/prev collected call stack/SQL statement.

---> C00713 - General - Bug Fix Request
Removed an SQL error that would be shown when going to the Scheduled jobs folder if the idoctor job scheduler file hasn't been created yet on the system.

---> C00713 - General - Enhancement
In V5R3 and higher, in library QIDRGUI the FTPFILE command has been updated:

Local file/remote file parms increased to 100 chars.
You can specify a *GET or *PUT of the file as the optional FUNCTION parameter.
Most parameters are mixed case now which fixes problems with mixed case passwords or filenames.

User profile has been increased to 50 bytes.

The script/log files size has been increased to support longer path/file names.

Also the FTPBATCH command has been removed since it's no longer needed.

---> C00713 - General - Enhancement
In library QIDRGUI the program that performs PTF checking has been updated (added another parameter). The latest server/client builds are required in order for PTF checking to work correctly. This was done in order to allow PTF checking for Disk Watcher.

---> C00713 - General - Enhancement
At V5R4, library QIDRGUI now includes a command SAVPFRCOL which is used to save performance data collections in the desired library. It supports saving of PEX, CS, JW, DW collections as well as PEX/CS MGTCOL object collections. A command of the same name that performs similar function exists in library QSYS at V6R1 and higher.

---> C00713 - Job Watcher - Bug Fix Request
In V5R3 Job Watcher, the thread rankings drill downs from the overview graphs were missing with the latest client.

---> C00713 - General - Enhancement
Added PTF checking when starting a super collection that includes, Disk Watcher or PEX data for the required Disk Watcher or PEX PTFs.

---> C00713 - General - Enhancement
Added PTF checking when starting a Job Watcher, Disk Watcher or PEX monitor for the collection type(s) selected within the monitor wizard on the basic options screen.

---> C00713 - General - Enhancement
Added PTF checking when starting a Disk Watcher collection.

---> C00713 - General - Enhancement
Based on customer request, added an option to the list of connections that allows the user to perform a PTF check for all components that provide PTF checking against the selected system(s). The results of the check and any missing PTFs are shown in the Check PTF Results window.

---> C00712 - Installation - General
Updated the uninstall CL program at 530 and higher to handle removal of the new iDoctor control files QAIDRJWM2, QAIDRDWM2, QAIDRPAM2

---> C00712 - Collection Services - Bug Fix Request
Drilldowns in CSI were showing the wrong graphs in some cases using client 711.

---> C00712 - Job Watcher;Pex Analyzer - Enhancement
Updated the monitors wizard on the basic options screen to include the ASP limit option. The value shown is taken from data area QMONASPLMT. If the used disk space is greater than this value then the user will not be able to start the monitor unless the limit is increased up to a max of 99%. Changing the value from this panel will update the data area which is checked by the STRJWMON/STRDWMON/STRPAMON commands.

---> C00712 - Collection Services - Bug Fix Request
Fixed crash in CSI when opening job summary reports with client 711

---> C00712 - Job Watcher;Pex Analyzer - Enhancement
Added several enhancements to the Monitors support at V5R3 and higher (latest server builds required)

PEX Monitors are now available at V5R3 and higher.
Add the Monitors folder to the PEX Analyzer GUI.

Added a data area QIDRWCH/QMONASPLMT which contains the maximum allowed asp limit. If the ASP for the monitor library's available disk space has exceeded the value specified in the data area then the monitor will end. The default is 90%.

The maximum length of a monitor name has decreased from 8 to 7.
The maximum number of collections in a monitor has increased from 99 to 999

The RSTJWMON/RSTDWMON/RSTPAMON commands have been removed. You can now restart a monitor using the start commands: STRJWMON/STRDWMON/STRPAMON

The monitor control files have changed as follows due to parameter changes on the STR monitor commands:
QAIDRJWM1 -> QAIDRJWM2
QAIDRPAM1 -> QAIDRPAM2
QAIDRDWM1 -> QAIDRDWM2
The first time a monitor is started on a system, the existing monitor information will be copied over to the new file.

In STRPAMON added a parameter called ENDPEX data option with these values/meanings:
*LIB = The monitor will create the PEX database files for each collection.
*MGTCOL = The monitor will create PEX MGTCOL objects for each collection.
*SUSPEND = The monitor will issue ENDPEX with the *SUSPEND option instead of dumping any data.

---> C00711 - General - Usability
Changing the bars per page (via the graph defintion->set bars per page menu/not via the preferences interface) for a graph, and then changing to a different alternate view will now preserve the bars per page setting when loading the alternate view graph.

---> C00711 - General - Usability
Using the preferences window from a graph view will cause the graph to be refreshed upon closing the preferences window.

---> C00711 - General - Usability
Using the menu within a graph

Graph definition -> Set bars per page -> Set value... will now take the user to the general page of the graph definition properties to set the bars per page override there instead of going to preferences.

---> C00711 - Pex Analyzer - Bug Fix Request
The generic job name length preference did not work in the PEX Analyzer graphs that provided generic job name groupings and drill downs.

---> C00711 - Pex Analyzer - Bug Fix Request
In PA, for the sql-based analysis, collection overview -> CPU consumed rankings, the drill downs were showing duplicate menu options.

---> C00711 - Disk Watcher - Enhancement
When working with user-defined queries/graphs in Disk Watcher or the QAPYDW* files that include fields STUNITNUM or TCUNITNUM you will now be able to right-click and graph the selected disk unit, path or ASP from the current record.   Note however that this will only be available if the collection has trace data and has been summarized.

---> C00711 - Job Watcher;Pex Analyzer - Bug Fix Request
JW FTP Transfer function was failing if an IP address was used instead of the remote system name. This was due to a bug in the client side when issuing the command to FTPJWCOL.

---> C00711 - Collection Services - Enhancement
When working with user-defined queries/graphs in CSI or the QAPM* files that include fields JBTDE or JWTDE you will now be able to right-click and graph the selected thread from the current record. In order to display the job name /user /nbr correctly in the drill down graph, fields JBNAME, JBUSER, JBNBR must also be included in the query.

Also if the query contains field DSARM you will be able to graph the selected disk unit over time from the current record.

Additionally the "Collection overview" graphs from CSI are also available now from these table views.

---> C00711 - Job Watcher - Bug Fix Request
When working with user-defined queries/graphs in JW that include the taskcount and interval fields:, drilling down into the call stack would cause the user to be prompted for a member name unnecessarily. You can also right-click and drill into the thread using the selected thread menu now..

---> C00711 - Job Watcher - Bug Fix Request
On the call stack window pressing the Show Holder button would not show the holder's name correctly at the top of the window if the holder did not have any TDE file records.

---> C00711 - Job Watcher - Bug Fix Request
Viewing the call stack for a job/interval that did not use any CPU would not show the full job name properly on the interval details panel if the option was invoked from some areas in the GUI (like the Thread wait time signature for XYZ: 001 graph)

---> C00711 - Job Watcher - Bug Fix Request
For V5R4/V6R1 JW when working with unsummarized collections, the thread signatures wait bucket rankings chart would show no data if the bucket and interval(s) selected contained no wait bucket times > 0 in the TDE file. This would only happen if the selected wait was only present in the STS file.

---> C00710 - Plan Cache Analyzer - Bug Fix Request
When displaying a plan cache analyzer collection, most of the fields were being shown as blank incorrectly with the latest builds.

---> C00710 - General - Enhancement
You can now drill down from the Syncronized table under the graph, just like you were drilling down from the bars in the graph with a couple of exceptions:

1) You can't select the desired wait bucket to drill down by default from the table.
2) If you are working with rows in the table that aren't in the graph (because the number of rows currently shown in the table is greater then the number of bars currently shown in the graph) then the drill downs won't be available.

---> C00710 - General - Usability
Changed how selection and "quick view" works in graphs to be consistent with how they work in a table.

A DOUBLE-CLICK ON A BAR IN A GRAPH WILL NOW OPEN THE QUICK VIEW WINDOW INSTEAD OF A SINGLE CLICK! (ahem in case you weren't listening)

Essentially, a single click will select a bar (without needing to hold down ctrl/shift) You still need to hold down shift to make a range selection however (just like a table).

I did this to help avoid the annoyance of opening the quick view window when not desired.

---> C00710 - General - Usability
The preference for new users on the Miscellaneous tab "Maximize reports opened into Data Viewers" is now on/checked by default.

---> C00710 - General - Usability
Updated the minimum level of client access needed on the PC in order to run iDoctor to V5R3 or higher.

---> C00710 - Job Watcher;Pex Analyzer - Enhancement
Add support to run PEX Analyzer monitors at V5R4. New commands STRPAMON, RSTPAMON, etc have been added to the QIDRWCH library and the GUI now provides the option to start a PEX monitor at V5R4. The latest server builds are required.

---> C00710 - General - Bug Fix Request
Duplicate as table menu option not working from the Synchronized table view.

---> C00710 - General - Bug Fix Request
The selected row of the syncrhonized table was not always matching up with the graph.

---> C00709 - Job Watcher - Enhancement
When using the graph compare function in JW , selecting the bottom graph and clicking the alternate view button to show the alternate view window, the situation filters were not shown in the list. This will effectively allow you to compare where different situations are occurring for the same graph.

---> C00709 - Job Watcher - Enhancement
In 540 and higher JW, updated the the detail reports -> Waits -> Longest "bad" current wait details and added a new one in the same location: Longest "bad" current wait details summarized

---> C00709 - Job Watcher - Bug Fix Request
From the call stack reports and object reports , the drill down menu into the "selected thread" was missing.

---> C00709 - Job Watcher - Bug Fix Request
Drilling down further into interval details from the object reports (launched from interval detail panel) would cause some field values to be shown incorrectly because the field type was being incorrect set for those fields to HEX.

---> C00709 - Collection Services - Bug Fix Request
In CSI for the CS (mgtcol) objects, on the properties window the wait buckets tab was being displayed when it should not have been and giving an SQL error.

---> C00709 - Collection Services - Bug Fix Request
When right clicking on a CS object in CSI -> Create collection, if you type in a library that doesn't exist, the remote command for crtpfrdta was being submitted immediately without checking to see if the library exists.

---> C00709 - Job Watcher - Enhancement
In V5R3 and higher JW, on the call stack panel, and looking at a job that is a holder, added a label that indicates how many jobs were waiting on the holder and of those how many different holder call stacks are in the data..

---> C00708 - General - Bug Fix Request
When using the report generator function modified the implemtentation that creates the report directory so that any error messages and the GetLastError rc from Windows will now be shown to the user.

---> C00708 - Collection Services - Enhancement
For V5R4+ CSI, added the IP address family and formatted IP address to the end of the job search report

---> C00708 - Job Watcher - Bug Fix Request
Moving through the call stacks using the arrows could sometimes cause the interval to go downwards incorrectly, if certain intervals contained multiple call stacks for the same taskcount (happens if holder call stacks were captured).

---> C00708 - Collection Services - Bug Fix Request
For V5R4+ CSI, reports that include the IP address where showing invalid values for the IP address due to the GETIPADDR function using the DETERMINISTIC setting.

---> C00708 - Collection Services;Disk Watcher - Bug Fix Request
When working with the disk reports in CSI or DW, if asp 1 was not collected, and only a single asp was collected , then no data would be displayed in the graph. The GUI will set the SQL to graph the ASP collected.

---> C00708 - Job Watcher - Bug Fix Request
Fix call stack reports so they exclude counts for holder call stacks.

---> C00708 - Collection Services - Enhancement
In V5R3/V5R4 CSI for the 16 bucket graphs, added gate time to the collection overview.

---> C00708 - Job Watcher - Bug Fix Request
In JW, starting a collection from an existing definition, would not correctly pull the interval duration value from the definition. This could cause the collection to be started with an interval duration of 10 instead of the desired value from the definition.

---> C00707 - Job Watcher - Usability
Graphing multiple JW collections is not supported at V5R3 but the check wasn't working and was allowing a user to try to do so anyway which resulted in no report being generated.

---> C00707 - Job Watcher - Enhancement
The situation analysis background and the dashed lines shown on graphs when graphing multiple collections at once are now displayed when variable width graphing mode is enabled.

---> C00707 - Installation - Enhancement
Updated the Install program on the FTP parameters page to include a checkbox that lets the user choose to perform or skip the ftp connection check (to QTEMP library).

---> C00707 - Pex Analyzer - Enhancement
Added the job name and changed the sort order of the IFS analysis report in PEX Analyzer to job and then time since collection started..

---> C00707 - Collection Services - Enhancement
When deleting collections in CSI, the GUI will now call the OS command DLTPFRDTA instead of using its own implementation.
In addition, the GUI will now delete any summarization or analysis files if they exist.

---> C00707 - Collection Services - Enhancement
Updated the CSI interval summary program at V5R3 and higher so the fields in the interval summary file can contain null values. In one case, certain intervals of the JOBMI file were not available, causing errors on the SQL insert.

---> C00707 - General - Usability
When deleting a monitor the GUI attempts to delete the job scheduled entries associated with the monitor if any exist. If none have ever been created an SQL error would be shown about a file not existing. Removed this message.

---> C00707 - Job Watcher - Usability
After hitting the apply button on the alternate views window, focus is moved to the graph automatically so that flyovers will work without first needing to click on the graph.

---> C00707 - Job Watcher - Enhancement
The situational analysis filter will now set the color to the situation color instead of leaving it red for intervals where multiple situations occurred.

---> C00707 - Disk Watcher - Enhancement
The V5R4 DW monitor commands now allow you to specify the maximum size of each collection.
Previously this was only supported at V6R1. Latest server builds are required.

Also updated the V5R4 monitors GUI to support this max size option (which applies only to DW since JW/PEX max size is specified within the definition at that release).

---> C00707 - Collection Services - Enhancement
Added wait bucket counts to the CSI graphs (overview and graphs for a single thread/job over time) at V5R4 and V6R1

---> C00706 - Job Watcher - Bug Fix Request
In JW, the holder indicator (>) was not showing up when there was a holder on some drill down graphs.

---> C00706 - Job Watcher - Enhancement
In the alternate views window, the graph background drop down now contains a list of all situations (for JW/DW). Setting it to a specific situation will act as a filter to only show the desired situation that occurred within the background for the graph.

---> C00706 - Disk Watcher - Enhancement
For DW stats, it's fairly common to miss I/Os if the interval size is set too large because the buffer can become full. Added a situational analysis option to DW Stats to look for this and displays it on the graph background (in red) if it occurs.

---> C00705 - Pex Analyzer - Bug Fix Request
Added fixes for the PEX Analyzer SQL-based IFS analysis.

---> C00705 - Job Watcher - Bug Fix Request
In JW when working with the TDE file from the server-side output files folder, the drill down into the selected thread graphs were not showing up.

---> C00705 - Installation - Bug Fix Request
Fixed a bug in the installer when it tries to test the ftp connection. On a failure the FTP error message and rc were not being displayed in the message box.

---> C00704 - General - Bug Fix Request
Fix a bug in the table and graph views where if you change an SQL statement and you have a syntax error, you would no longer be able to run queries from that view.

---> C00704 - General - Usability
Improved performance in the JW interval detail property pages. Information on these pages is retrieved much faster now when changing the current interval especially when working on a slower internet connection.

---> C00704 - General - Usability
Improved performance in the parsing of SQL statements by reducing the number of CHKOBJs that occur. Also cached these checks when performing drill downs so they are not made again.

---> C00704 - General - Usability
Improved performance in several places in the GUI that are used to build a list of items (definitions, monitors, etc).

---> C00704 - Collection Services - Bug Fix Request
In CSI, fixed a bug with the SQL in the drill down graph for a single generic job "Generic job wait time signature for XYZ" graph if the generic job name length preference was changed from the default.

---> C00704 - Job Watcher - Bug Fix Request
Fixed a couple of broken drill downs in JW 5.4, 6.1 for summarized collections.
- Thread wait time signature for XYZ wasn't filtering on a taskcount (was showing all taskcounts)
- Job signatures ranked by XYZ for <<OBJTYPE>> <<OBJDESC>> wasn't filtering on the object type.

---> C00704 - Job Watcher - Bug Fix Request
In V5R3 JW, drill downs from the search function did not utilize the summary files when they existed.

---> C00703 - General - Usability

SQL tables and SQL aliases are no longer shown in the member selection window.
Also overrides are no longer made against SQL tables (unnecessarily).

---> C00703 - General - Enhancement
The graph compare function is now available in all components except for the old style ("flattened") graphs
in PEX Analyzer. Previously this was only available in JW, CSI and DW.

---> C00702 - Collection Services - Enhancement
At V5R4 and higher, all CSI graphs for a single job (grouping) over time now provide these additional fields

  Job current user profile
  Subsystem name
  Job pool
  Submitter job user
  Submitter job
  Job accounting code
  Server Type
  Remote IP address family
  Remote IP address
  Remote IP port
Pages allocated since job start
Pages deallocated since job start
Net pages allocated since job start
Accumulated total job CPU time (seconds)

For string fields: if multiple (different) values are found for the interval , the field will be blank.

---> C00702 - General - Enhancement
When working with the bottom graph of a comparison graph view, you can now change the time interval size
(when applicable) and the Change SQL Parameters window is also available.

---> C00702 - Collection Services;Disk Watcher - Enhancement
When opening a disk graph in CSI or DW that is designed to only show you one ASP at a time, a new
window will be shown that lets you pick which ASP(s) you wish to graph.
If 2 ASPs are selected, then a comparison graph is shown allowing you to more easily compare the 2 ASPs.

If you wish to not get this prompt and always default your graphs to ASP 1, then you can check the box
shown at the bottom of the window.
To reenable this option in the future if it's been turned off, visit the Confirm tab of the Preferences interface.

Note: The lists of ASPs will also contain an option where all ASPs are listed together. Choosing that option
will average all the ASPs together into the graph (although typically this is not recommended).

---> C00702 - General - Enhancement
The libraries folder in the component views, now contains the library filter (within the folder name) if one is
set.

---> C00702 - Job Watcher - Bug Fix Request
Changed the sort on the Physical I/O activity totals graphs for summarized collections in V5R4/V6R1 JW to
TOTIO instead of all 8 IO fields added together..

---> C00702 - Collection Services;Job Watcher - Enhancement
Added a preference on the Miscellaeous tab that lets you indicate how many characters should make up the
generic job names on the generic job grouping graphs in CSI and JW. The possible values are 2 - 9 with a
default value of 6 (the previous default).

---> C00702 - General - Bug Fix Request
The Edit -> Preferences menu when asked from the tree side of a component view didn't work.

---> C00702 - General - Enhancement

The SQL log now contains the QZDA job number and (PC time) timestamp.
Also removed a blank line between the statement and the line indicating how long it took to run.
Access this feature from the View -> SQL Message Log menu within the Data Viewer.

---> C00702 - Job Watcher - Usability
Displaying the list of JW collections in a library should be a little faster now
(the GUI was calling a stored procedure to check if the summary files were valid twice unnecessarily).

---> C00702 - General - Enhancement
Added an SQL tab to the Preferences interface that contains the following options:

Log messages to log file (default is now OFF)
Maximum data viewer connections (Number of QZDASOINIT jobs to utilize in the data viewer)

---> C00702 - General - Enhancement
Added a Confirm page to Preferences where all the confirmation options from the Miscellaneous
preferences tab will now be located.

---> C00702 - General - Enhancement
Made a change in the iDoctor database code, so it will no longer remove overrides in situations where none
need to be removed.

---> C00702 - General - Usability
Updated the error message given when an access code was generated for the wrong OS VRM (the server
code and OS VRM match but the access code doesn't)

The old message was
Access code is invalid (OS release mismatch?)
the new message is
Access code is invalid. It appears to apply to a different OS VRM. Contact support at idoctor@us.ibm.com.

---> C00702 - General - Usability
Updated the error message given when a user upgrades their system but forgets to upgrade the iDoctor
server code.

The old message was
Access code is invalid (OS release mismatch?)
the new message is
Access code is invalid. A previous release of the iDoctor server code is installed and must be updated. Visit
our website to install the correct version of the server code for this system. http://www-
912.ibm.com/i_dir/idoctor.nsf

---> C00702 - General - Usability
In the component views, refreshing a large list of items (like a list of active jobs), should look cleaner now.

---> C00702 - Collection Services - Bug Fix Request
Fixed some additional problems with the CSI disk graphs for V5R4 and higher for mirrored units.

---> C00702 - Collection Services - Bug Fix Request
Fixed a crash that would occur if you try to right-click a CSI graph as it opens.

---> C00702 - General - Usability
Add a history for the library filter dialog so that it remembers the last generic names entered and displays
them in the drop down list. Also added a remove button to clear the selected entry from the list.

---> C00701 - Collection Services - Bug Fix Request
Fixed some problems with the CSI disk graphs for V5R4 and higher:

Disk units did not include the arm letter if there was one.
Read cache hit%, write cache efficiency%, write cache overrun % were off for mulitpath disk units.
The disk path rankings graphs did not show all the paths.

The IOP rankings graphs was missing some of the IOPs.

---> C00701 - Job Watcher - Usability
The create aliases step for the JW summary at R5.4, R6.1 is now enabled by default.
These aliases are required for performing some functions like graphing multiple collections at once.

---> C00700 - Job Watcher;Pex Analyzer - Enhancement
Updated the Job Watcher monitors support for V5R4 to provide the capability to automatically run PEX stats
flat collections for the top CPU QZDASOINIT jobs.

Added a new JW definition Q5MINQZDAS. Use the Reload IBM-supplied definitions option from the
Monitors window (basic options- actions button) to make this definition available.

When definition Q5MINQZDAS is used, the monitor will first collect all QZDASOINIT jobs for 5 minutes at 5
second intervals.
Every 5 minutes from then on the monitor will look for the top cpu burning QZDASOINIT jobs and start a
PEX stats flat collection against them (up to 5 max)
The PEX stats collections are named the same as the monitor + 01-99. After 99 PEX collections it will wrap.
You will end up with many more PEX collections than JW collections.

There is also a new JW command in QIDRWCH library called RUNPXSTATS that you can manually run if
desired against a JW collection to collect PEX Stats flat for the top N CPU burning jobs in the collection over
the last 5 minutes (the collection wouldn't necessarily need to be a monitor or even QZDASOINIT jobs).

---> C00700 - Job Watcher - Enhancement
Updated the required PTFs for JW at 530/540 to include the fixes for SQL collection support.

---> C00700 - Collection Services - Enhancement
In CSI for all releases, updated the job runtime shown in ranking graphs. It was inaccurate in some cases,
especially when the job duration was less than the duration of 1 interval.

---> C00698 - General - Enhancement
The threads within a job within the Active jobs folder now provides a drill down into the current call stack for
the selected thread.

You can configure the fields shown for the call stack using the select fields... menu found when right-clicking
the thread ID from the tree side.

---> C00698 - General - Enhancement
Most iDoctor components now contain an "Active Jobs" folder which provides WRKACTJOB like
functionality from the iDoctor client.
You can sort by a statistic and refresh to keep tabs on the top CPU burners, etc.
There is also a filter option to filter the list by name, user, number, current user or minimum CPU %.
Use the Select fields... menu option when right-clicking the folder to rearrange fields or add additional fields
(from the Open list of Jobs API)

Expanding a job shows the threads and the thread statistics available for each.

You can start Job Watcher or PEX Analyzer collections or add jw /pex definitions using the selected jobs
within the Active jobs folder. You can also end the selected jobs or view job logs.

---> C00698 - Job Watcher - Bug Fix Request
Updated the job descriptions shipped in the QIDRWCH (Job Watcher) library so that the initial library list
value is now *SYSVAL. Previously the value of QIDRWCH, QGPL, QTEMP was causing problems on some
systems.

---> C00698 - General - Bug Fix Request
When selecting a folder in the tree side of component views, the status bar was not identifying the
description of the selected folder as it should (like when selecting items in the list view).

---> C00698 - General - Enhancement

From the new ASPs and Disk Units folders you can now select one or more ASPs or disk units, and right click to create either a PEX filter or Disk Watcher definition from the desired selections.

---> C00698 - General - Enhancement
Most iDoctor components now contain two new folders showing the ASPs and Disk Units configured on the current system.
The ASPs folder allows drill down to view the Disk Units within an ASP.
The Disk Units folder provides WRKDSKSTS type of statistics with updates provided with each refresh.
(also includes total I/O, total sizes). Use the right-click menu option Reset Statistics on the Disk Units folder to restart the collection of disk statistics.

Use the Select fields... menu option when right-clicking the folder to rearrange fields or add additional fields (from the Open list of ASPs API)

The status bar of the Main window shows the times for first disk statistics snapshot, and the last one.

---> C00698 - Disk Watcher - Enhancement
When starting a Disk Watcher collection at V6R1, the default definition selected will now be QTRC.
Previously one was not preselected.,

---> C00698 - Disk Watcher - Enhancement
When starting a Disk Watcher collection, added a check in the GUI to ensure that collection files in the target library are of the same release as the current one.
This will prevent the collection from starting up and immediately failing due to record format problems.

---> C00698 - Disk Watcher - Bug Fix Request
When starting a collection in Disk Watcher by first selecting a definition from the Definitions folder, the selected definition was not being visibly selected in the list of definitions in the Definition Selected panel of the Start Disk Watcher Wizard.

Also made a change on this panel, so that the list of definitions is only retrieved once the 1st time the interface is loaded, and not everytime the user changes the definition type from IBM to user-defined or vice versa.

---> C00698 - Disk Watcher - Usability
In DW V6R1, removed an unnecessary delay when opening the basic options screen of the Disk Watcher collection/definition wizard.

---> C00698 - General - Bug Fix Request
Fixed a bug in several components, the timestamps in the alternate view graph descriptions for thread rankings charts would sometimes not be shown in the correct format or at all.

---> C00698 - General - Enhancement
When creating a CSV file with iDoctor, the default setting is now to include the field headings as a header row. You can modify this on the File tab of Preferences.

---> C00697 - Pex Analyzer - Enhancement
Added several additional enhancements to the PEX save/restore analysis at V5R4/V6R1.
You can now drill down from any of the table views, into SR event details for the desired thread.
Added a SR thread rankings folder as another possible starting point.

---> C00696 - Pex Analyzer - Enhancement
Added an SQL-based Taskswitch analysis to PEX Analyzer (V5R4/V6R1). The files available so far are very low level and not "done" yet. (taskswitch object index file).

---> C00696 - General - Bug Fix Request
The Edit -> Open in New Data Viewer option not working right when data viewer 1 is open.

---> C00694 - Pex Analyzer - Enhancement
Updated several reports in the PEX Analyzer save/restore analysis, and added support for the analysis at V5R4.

---> C00694 - Job Watcher - Bug Fix Request
In JW 540+, for summarized collections, the drill down graphs Physical I/O activity totals by thread and
Logical I/O activity totals by thread are now properly sorted by total I/Os instead of total rates.

---> C00693 - Pex Analyzer - Enhancement
For PA V5R4+, for the TPROF report drill down "Call stack hits for procedure 'XYZ'"
added call level 0 (basically the procedure you clicked on when doing the drill down) to the call stacks
shown.

Also added these fields:
Procedure code start address
Procedure offset (instruction address - procedure code start address)
Module timestamp

This report is also now available from the Classic TPROF analyses at V5R4 and higher.

---> C00693 - Job Watcher - Usability
If summarizing multiple JW collections using the batch option, only 1 informational message will be shown
now instead of one message per collection.

---> C00693 - Job Watcher - Bug Fix Request
If deleting Job Watcher collections from the green screen or by a monitor overwriting one, the summary files
were not being deleted (via the DLTJWCOL command).
The latest server builds must be installed to resolve this.

---> C00692 - Pex Analyzer - Enhancement
For PA V5R4+, added a save/restore event analysis the parses out the data in the QAYPEMIUSR file into
several reports. This is a work in progress and all the reports are not yet complete.

---> C00692 - Job Watcher - Enhancement
In JW 540+, added the holder name and wait object information fields to the drill down:
Detail reports -> Waits -> Longest current waits for bucket XYZ

---> C00692 - Job Watcher - Bug Fix Request
Removed a message that could appear when starting or restarting a monitor regarding file QAPEXDFN in
QUSRSYS not found, if no PEX definitions had ever been created on the system.

---> C00692 - General - Bug Fix Request
Fixed a possible crash when opening reports in the Data Viewer, and then trying to close the Data Viewer as
they open.

Jan 22, 2009 – C00691

Job Watcher changes

---> C00690 - Job Watcher - Bug Fix Request
In JW, the drill down options from the call stack or wait object reports when previously coming from the
Thread wait time signature for XYZ graph were incorrect.

---> C00690 - Job Watcher - Bug Fix Request
In the 1st JW Create Job Summary report, the elapsed seconds field was mislabeled as elapsed
microseconds.

---> C00689 - Job Watcher - Enhancement
Added missing field descriptions for some graphs in Job Watcher.

---> C00689 - Job Watcher - Enhancement
In the Job Watcher J9 reports, that show the java thread state field (JTSTATE), you can now perform a
flyover over that field and get a text description of what the numeric value means.
If a value isn't shown correctly for certain states, let me know.

---> C00688 - Job Watcher - Bug Fix Request
The Create Job Summary option found in Job Watcher graphs after selecting a time range, was showing up incorrectly for V5R3 JW graphs, or higher release collections that were not yet summarized. This option requires V5R4 and higher and the collection to be summarized.

---> C00688 - Job Watcher - Bug Fix Request
In Job Watcher ranking graphs, when multiple bars were selected, the Create Job Summary and Split Collection options were showing up in the popup menu when they should not have been.

---> C00688 - Job Watcher - Bug Fix Request
Fixed a few missing drilldowns for Selected thread information from some of the table reports in Job Watcher.

---> C00686 - Disk Watcher;Job Watcher - Enhancement
When displaying the list of monitors, the list now contains the collection overlap parameter in seconds.

---> C00686 - Job Watcher - Usability
Drilling into a monitor that has a name that starts the same as other monitors in the same library could cause more collections to show up under the monitor than it should have.

---> C00686 - Job Watcher - Usability
Updated the set of IBM-supplied Job Watcher definitions at V5R3/V5R4, so that the default run duration is 2 hours 10 minutes instead of 60 minutes. This will ensure that the default monitor collection duration of 60 minutes will be less than the most likely definitions used. In cases where the definitions have already been loaded, the user will either need to increase the definitions duration value "*NBRSEC" appropriately , or use the Reload IBM-supplied definition menu option found when clicking the Actions button on the Basic Options screen of the Monitor Wizard.

---> C00686 - Job Watcher - Enhancement
When starting a Job Watcher monitor at V5R3 and V5R4, added checking to ensure that the collection duration + overlap time is less than or equal to the definition's collection duration.

This doesn't apply to Disk Watcher monitors or V6R1 Job Watcher because the run time isn't part of the definition in those cases.

---> C00685 - Job Watcher - Bug Fix Request
Drill down after using the Job Watcher search function wasn't working properly if the collection wasn't summarized or collected at V5R3.

---> C00684 - Job Watcher - Bug Fix Request
Attempting to start a monitor that already exists in the GUI would fail, without any error message given.

---> C00684 - Disk Watcher;Job Watcher - Bug Fix Request
The collection overlap parameter for the JW and DW monitors wasn't working properly causing possible short gaps in the monitor data on busy systems.

Latest server builds are required to fix.

---> C00684 - Job Watcher - Enhancement
In V6R1 JW, added a new piece to the JW summary that looks for long transactions that occurred in the collection. Currently 2 types are detected. Type 1 = Long running SQL, Type 9 = Generic transaction (means no wait types that are typically idle occurred during the time period).

The long transactions file is QAIDRJWTXNSUM_<<COLNAME>> and is available in the server-side output files folder after being created.

The JW server builds must be installed for this to work.

---> C00684 - Job Watcher - Enhancement
In V5R4 JW+, updated the call stack summary , renaming field MININT to INTERVAL.
Also when opening the call stack summary file QAIDRJWSTKSUM_<<COLNAME>> from the server-side output files folder, you can now double-click a record and view the call stack for any records listed in the file.

---> C00682 - Job Watcher - Enhancement
In JW, if on an overview graph and a time range is selected, the popup menu now contains an option Create Job Summary which will open the Create Job Summary window letting you perform the summary for the desired time period.

---> C00682 - Disk Watcher;Job Watcher - Enhancement
Based on customer request, when deleting a monitor, the job scheduled entries associated with the monitor will now be automatically deleted as well.

---> C00682 - Job Watcher - Enhancement
In JW V5R4 and higher, (for summarized collections only) on the interval summary wait buckets tab, jobs that didn't use CPU but are in a wait for the entire interval are now visible.

---> C00682 - Job Watcher - Enhancement
For JW V5R4 and higher (for summarized collections only), added a Bad Current Waits tab to the Interval Summary pages that shows the worst current waits in an interval with drill down support to see the graphs for any of those jobs.

The wait buckets used for this at V5R4 were:
        5, 7, 11, 13, 14, 15, 16, 17, 19, 32
and for V6R1 they were:
5, 7, 8, 11, 14, 15, 16, 17, 18, 19, 31, 32

---> C00681 - Job Watcher - Bug Fix Request
Fixed a bug when displaying properties for a Job Watcher definition at V6R1 that contained the *ACTGRPDTL collection option. This caused the window to indicate that activation group detail was not part of the definition when in fact it was.

---> C00681 - Job Watcher - Enhancement
For V5R4 and higher JW, if the collection is summarized, several additional interval summary property pages are now available.
Situations (lists the situations found in the interval)

These all list totals for the interval: Physical I/Os, Logical I/Os, IFS, Transactions, Other statistics , Wait bucket totals

---> C00681 - Job Watcher - Enhancement
The call stack reports are now available from the jobs listed in these interval summary property pages: wait objects, holders and wait buckets.

---> C00681 - Job Watcher - Enhancement
Added CPU utilization % to the JW interval summary pages.

---> C00681 - Job Watcher - Bug Fix Request
In some graphs in JW, the flyover would not show the Y2 axis description and value.

---> C00681 - Job Watcher - Enhancement
The reports shown on JW interval summary property pages (Wait objects, holder, wait buckets) are a bit faster now.   Also fixed a bug where data was being reretrieved unnecessarily when switching between pages.

---> C00680 - Job Watcher - Enhancement
For V5R4 and higher JW, on wait bucket rankings and JVM rankings graphs, added a new drill down menu "Call stacks" which provides reports that show (up to 50 level deep) call stacks for all threads in the the selected job in a single report for either the current interval or all intervals in the time range the rankings graph applies to.

---> C00680 - Job Watcher - Enhancement
For V5R3 and higher JW, the code that builds the drill downs was restructed Let me know if anything appears out of place.

---> C00675 - Job Watcher;Pex Analyzer - Bug Fix Request
Fixed a crash with the transfer collection function, if it was used from the tree side instead of the list side of the component view.

---> C00674 - Job Watcher - Bug Fix Request
In JW 540 and higher, the Job rankings graph would always show 1 and only 1 task but the bar actually represented All System Tasks. The task name shown was just 1 task within the set. Changed the name of this bar to 'All System Tasks'.

---> C00673 - Job Watcher - Bug Fix Request
Fixed a crash in the JW collection wizard, on the basic options panel, if the library specified contained collections of a prior release.

---> C00669 - Job Watcher - Enhancement
For V5R4 and higher JW, added a new situation that identifies intervals where the avg sync write response time is > 3 ms.

---> C00669 - Job Watcher - Enhancement
For V5R4 and higher JW, for summarized collections, added a report showing details about the situations occurring for the selected time range.
This option is available under the menu: Detail reports -> Situation details.

From this report you can view the call stack or view one of the detail graphs for the desired situation job/thread.
You can also perform a flyover of the situation ID field to get the situation name.

---> C00669 - Job Watcher - Enhancement
For V5R4 and higher JW, for summarized collections, added graphs showing totals as alternate views of graphs showing rates. The following graphs are new:

Disk pages allocated/deallocated totals
Disk reads and writes totals
Physical I/O activity totals
Logical I/O activity totals

IFS lookup cache totals
IFS opens totals
IFS reads totals
IFS creates/deletes totals

Transaction totals

---> C00669 - Job Watcher - Enhancement
In JW V5R4 and higher, for summarized collections, combine the synchronous read and writes response time graphs into a single graph.

---> C00667 - Disk Watcher;Job Watcher - Bug Fix Request
Fixed a bug when restarting a monitor, changing the definition listed in the wizard did not update the definition used when the monitor was started.

---> C00667 - Job Watcher - Usability
Added a new IBM -supplied definition to Job Watcher called Q3MINQZDAS which includes only QZDASOINIT jobs with 3 minute intervals that is intended to be used by the FNDJOB function to dump information about long running SQL statements.

---> C00667 - Job Watcher - Usability
In JW 530/540, when loading a definition into the start collection wizard, the conflict and bad waits call stack options will no longer be turned on if they were not specified on the command string for the definition loaded.

Collection Services Investigator changes
---> C00689 - Collection Services - Bug Fix Request

In CSI, the ASP capacity was listed in megabytes when it was really gigabytes on the Collection properties System tab.

---> C00689 - Collection Services - Bug Fix Request
Drive capacity was being calculated incorrectly in the CSI disk rankings graphs.

---> C00689 - Collection Services - Enhancement
Added missing field descriptions for the Disk Graphs and Communication graphs in CSI.

---> C00689 - Collection Services - Bug Fix Request
Fixed a bug where the V5R3 CSI Disk Graphs by time interval didn't allow the time interval size to be changed.

---> C00686 - Collection Services - Enhancement
Added missing field descriptions (for additional fields listed in the graph legend) for several graphs in CSI.

---> C00686 - Collection Services;Pex Analyzer - Enhancement
In PEX and CSI added a Wait Buckets tab to collection properties that shows the wait buckets/enum mapping at V5R3 and higher

---> C00686 - Collection Services - Enhancement
Modified all job-based CSI rankings charts at V5R3 or higher so that they include a flyover showing the total number of contributing threads/tasks.
This is helpful if you need to know how many threads are included within a generic job or other grouping.

Also when performing the next drill down showing a single job grouping over time, these graphs now show the total contributing threads/tasks value as the Y2. In the case of the CPU graphs, 2 new graphs were added that
allow you to either display the Y2 as the total contributing threads/tasks or showing a Y2 value as before (the systems CPU time/utilization).

---> C00682 - Collection Services - Enhancement
Added support to do a Create Job Summary in CSI, based on a selected time range from an overview graph.

Also, when displaying reports in the Job Summaries Folder in CSI and JW, changed the filter column so that the timestamps are formatted more cleanly.

---> C00679 - Collection Services - Usability
The popup menu for CSI graphs and loading the next report is faster now.

---> C00678 - Collection Services - Usability
In CSI, for the wait bucket thread ranking graphs made a change so that instead of showing all buckets, only the "most interesting" set of wait buckets will be shown,. The set of "interesting" wait buckets is configurable using the Edit -> Wait Bucket Preferences menu on the Main Window.

I made this change because in many cases when dealing with CSI data, the idle waits are too overwhelming/uninteresting.

---> C00678 - Collection Services - Enhancement
In CSI, the SQL was changed to use aliases on the QAPM collection file members instead of database overrides. A stored procedure will create the aliases before the SQL statements are ran for the desired collection(s). This was done in order to support graphing of multiple collections simultaneously.

---> C00678 - Collection Services - Enhancement
For V6R1 CSI, added the capability to graph over multiple collections at once in the same library. This only works for V6R1 collections. Select the desired collections from the CSI component view's list side and then right-click and choose the graph of interest. You will be asked if the graph should be created for all collections selected, click yes. From then on the drill downs for rankings and the single object over time graphs should apply to this same set of colllections.

---> C00678 - Collection Services - Enhancement
Enable the capability to do a Collection Search in CSI over multiple collections at once by selecting the desired collections in the CSI component view's list side and then using the Search popup menu.

After the search results are shown , the drill downs are over the same set of collections provided for the desired job/thread.

Note: At V5R3/V5R4, the drill downs are only for the single selected collection member for the job/thread (due to SQL complexity issues).
At V6R1, the drill downs are for all selected collections and the job/thread.

---> C00676 - Collection Services - Bug Fix Request
Fixed a bug in CSI, where the System Collected on VRM was listed incorrectly if the collection had been converted from V5R3 or V5R4 to V6R1..

---> C00676 - Collection Services - Enhancement
CSI Create Job Summary didn't work on collections converted from V5R3 or V5R4 to V6R1.

---> C00676 - Collection Services;Pex Analyzer - Enhancement
The SQL-based Analysis Window in PEX Analyzer and CSI, now supports the option to submit the desired analyses to a batch job.

---> C00676 - Collection Services - Enhancement
In CSI for V6R1, added support to analyze the new Collection Services file QAPMXSTGV. It contains external storage DS8K (for X32) data and will be produced in the future once the required PTFs are made available.

If the collection contains this file, a new collection menu option 'Analyze' is available which contains a submenu Run External Storage (cache statistics). After using this option, a couple of reports are available under the External storage cache statistics folder.

For an example of the statistics available, see iDoc610, library EXTSTG. The generated SQL tables are named QAIDRCSEXT_MBRNAME and QAIDRCSEXT_TOT_MBRNAME.

---> C00674 - Collection Services - Bug Fix Request
Fixed a bug in CSI, in the Old CSI graphs folder. There were two versions of the same graphs shown and one set did not work (giving an SQL error).

---> C00670 - Collection Services - Enhancement
In CSI, in the disk graphs, for the selected "disk object type" graphs, added a drill down into the rankings graphs for the selected time range.

Also corrected a bug in the SQL statement for this type of graph where the ASP number was incorrectly forced to 1.

---> C00670 - Collection Services - Enhancement
In CSI, in the disk graphs , the current ASP number will now be listed in the graphs title.

You can change the current ASP number if desired using the Change SQL Parameters menu and this will also update the graph title.

Disk Watcher changes

---> C00690 - Disk Watcher - Enhancement
In DW, added flyover descriptions for the values shown for fields TCIOFUNCTN and TCSMIOTYPE in the QAPYDWTRC file and the Top 25 longest I/Os report.

---> C00690 - Disk Watcher - Enhancement
In DW, on the Detail reports -> Top 25 longest I/Os report, added the Storage management I/O type field.

---> C00689 - Disk Watcher - Bug Fix Request
When displaying Disk Watcher collections in a library or monitor, the description field was not being retrieved.

Also there is a problem with the Disk Watcher description not being set correctly in the STRDW command at V5R3/V5R4 which means trying to set a description at those releases won't work (unless a PTF is released to fix).

---> C00688 - Disk Watcher - Bug Fix Request
Made a change to the IBM-supplied Disk Watcher definitions, QTRC, QTRC1MIN, QFULL, QFULL1MIN to ensure that they include COLRNG (*NONE) COLCOND(*NONE)

At V5R3/V5R4, on some systems, not including these parameters will cause STRDW to fail (see APAR SE34876). Use the Reload IBM-supplied definitions action from the Monitors or Super Collections interfaces to reload the IBM Disk Watcher definitions on your system.

---> C00688 - Disk Watcher - Usability
Add the Start New Monitor menu from definitions listed in the Disk Watcher Definitions view.

---> C00688 - Disk Watcher;Job Watcher;Pex Analyzer - Bug Fix Request
Fix a bug when listing monitors on a system. No Disk Watcher or PEX monitors would be listed if no Job Watcher monitors had been created first.

---> C00688 - Disk Watcher;Job Watcher - Enhancement
When starting a monitor, the collection type selected by default will be based on the component you are using at the time. Previously it was always set to Job Watcher by default.

---> C00686 - Disk Watcher;Job Watcher - Enhancement
For V6R1 JW and DW monitors , added a parameter to the basic options screen that lets you set the maximum collection size (in megabytes) for each collection running in the monitor.

Latest server builds must be installed in order for monitors to work with this client build.

---> C00686 - Disk Watcher;Job Watcher - Enhancement
When displaying the list of monitors, the list now contains the collection overlap parameter in seconds.

---> C00678 - Disk Watcher - Enhancement
Added the capability to schedule Disk Watcher collections.

---> C00674 - Disk Watcher - Enhancement
Add field descriptions to the DW Trace interval summary files.

---> C00670 - Disk Watcher - Enhancement
In DW, in the stats rankings graph, changed the sorting to be based on avg response time instead of the unit number, path or asp number..

---> C00670 - Disk Watcher - Enhancement
In DW, the current ASP number will now be listed in the graphs title and all disk graphs except for ASP rankings will work against only 1 asp at a time.
You can change the current ASP number if desired using the Change SQL Parameters menu and this will also update the graph title.

General iDoctor GUI changes

---> C00685 - Installation - Enhancement
Added the capability to remove all iDoctor objects/libraries from the desired system by right-clicking the system name from the connections view and choosing "Uninstall iDoctor".

This option can also be used by executing the following commands manually from the green screen:

```
CRTDUPOBJ OBJ(QIDRUNINST) FROMLIB(QIDRGUI) OBJTYPE(*PGM) TOLIB(QTEMP)
CALL PGM(QTEMP/QIDRUNINST)
```

Because this was implemented as a CL program in library QIDRGUI, the latest server builds are required in order for the uninstall option to work.

---> C00691 - General - General
Due to their size, I moved the tip of the day videos out of the normal builds and moved them to a separate download on the main downloads page on the website.

---> C00690 - General - Bug Fix Request
Fixed an error on Vista where at startup there would be registry errors shown due to being unable to write certain keys.

---> C00689 - General - Enhancement
Moving the mouse over a field in the legend will now show additional information about that field in the status bar of the Data Viewer including the field short name and the SQL expression (if applicable/available) from the SQL statement for the field.

---> C00689 - General - Usability
Modified scrolling in the graph legend, so it doesn't scroll such huge amounts per click anymore.

---> C00689 - General - Bug Fix Request
In Graph definition properties for the Y1 and Y2 axis pages. removing the axis description and then clicking on the list would readd the description

---> C00689 - General - Usability
Made a change to the graph legend for the primary Y axis fields. Any fields (besides the 1st one) that are all 0 in the graph data, are no longer shown in the list of primary Y axis fields for easier readability of the available data. These fields are basically hidden, even though they still exist in the Primary Y axis fields and if you change the query or scroll (if non zero data is encountered) they will become visible again in that section of the legend.

Primary Y axis fields that are all zeros are now listed in the Available fields section but are greyed out. Trying to add them to the primary Y axis will do nothing because it's already part of the Primary Y axis internally.

If there is interest I could add a preference to have this part of the legend work either the old way or the new way. Let me know if you are interested.

---> C00689 - General - Bug Fix Request
Writing an SQL query in the SQL editor against an SQL alias should work properly now . Previously the GUI would prompt you for a member and fail.

---> C00687 - General - Enhancement
Added a new preference on the Misc tab called "Confirm system name when opening new SQL editor", that determines whether or not the list of available systems are shown when opening a brand new SQL editor.

---> C00684 - General - Bug Fix Request
Fixed a problem for users on Citrix servers when the iDoctor GUI tries to determine if it has been installed correctly and fails. The registry settings that identify the client installation were incorrectly placed in HKEY_CURRENT_USER instead of HKEY_CLASSES_ROOT

---> C00683 - General - Bug Fix Request
The right-click option to open a graph or table into a new data viewer, wasn't working when the preference "always open a report into an existing data viewer" was checked.

---> C00683 - General - Bug Fix Request
Fixed a problem with the creation of user defined queries and graphs over the SQL table summary files. They would not be usable on any collection other than the current one.

---> C00682 - General - Usability
When changing the time interval size in a graph, the graph will now be respositioned to the 1st bar or start of the graph. Previously the position was the same bar # as before, which didn't match where you were in the graph.

---> C00682 - General - Enhancement
The Windows menu on the Main Window, now contains the list of open Data Viewers, which allows you to open any of them from that menu.

---> C00681 - Installation - Bug Fix Request
Fixed a problem in the installer, where the abort install option wouldn't cancel the install right away.

Also made a change so that the error from windows is shown in a popup when unable to create the setup log file.

---> C00679 - General - Bug Fix Request
Fixed a bug where the Filter libraries menu on the Libraries folder didn't work if used from the list side instead of tree side of a component view.

---> C00678 - General - Bug Fix Request
The SQL editor doesn't seem to get refreshed with the correct sql statement after the query definition interface is used from a table view.

---> C00674 - General - Bug Fix Request
Library descriptions were not being displayed if the user , scrolled the list of libraries using the mouse wheel instead of the scrollbar.

---> C00669 - General - Bug Fix Request
Fixed a possible long delay /client freeze when doing a flyover when the data in the flyover contained a large amount of text. It will now be truncated to a more reasonable size.

---> C00667 - General - Enhancement
Added a new component tentatively called the Plan Cache Analyzer.
It offers the ability to run a collection containing plan cache and SQL query data from the system based on several options from the QQQOOOCACH API.

On the server side a new library QPLANCACHE will be restored when installing the server portion of iDoctor. This library contains a command STRPCCOL that is used to start a collection.

On the client side, a wizard is available to start the collection and if the analyze option is taken a set of graphs ranking the plan cache statistics is available.

---> C00667 - Installation - Bug Fix Request
The MS Visual Studio redistributable should install silently now.

Sep 17, 2008 – C00664

Job Watcher changes
---> C00662 - Job Watcher - Bug Fix Request
In JW 540/610, fixed a rare overflow problem on the CPU utilization fields in the interval summary file.

Collection Services Investigator changes
---> C00663 - Collection Services - Enhancement
In CSI, added the disk graphs to V5R4 collections that were collected using PTF SI23390 or higher. If the QAPMDISK file does not include the response time buckets then these graphs will not work. The IO Adapter drill down /rankings graphs aren't available though at that release since the field doesn't exist.

In CSI for V5R3 added a reduced set of only 5 disk graphs with drill downs.

---> C00663 - Collection Services - Enhancement
In CSI for all releases, added a Collection search function similar to the one provided in Job Watcher. This allows a user to generate a report showing job records based on the desired job name, user, number, subsystem, pool id, or current user profile. From these reports users can drill down into the graphs for the desired job over time.

Disk Watcher changes
---> C00662 - Disk Watcher - Bug Fix Request
In DW, corrected a couple of minor issues with the trace report drill down "Detail reports -> Top 25 longest I/Os"

---> C00661 - Disk Watcher - Bug Fix Request
In Disk Watcher when creating definitions (or starting collections at V5R3/V5R4), the COLRNG and COLCOND parameters will be set to *NONE by default. Previously they were not set at all which could cause a failure when running the STRDW command.

General iDoctor GUI changes
---> C00664 - General - Bug Fix Request
Fixed a crash that would occur if the user tried to rerun an SQL statement using the Execute menu from the SQL editor while the query was still running.

---> C00663 - General - Bug Fix Request
Fixed some problems with the alternate views functions from rankings graphs.
1. Duplicates entries would sometimes show up in the list.
2. The wait bucket description listed wouldn't change as it should based on the last bucket selected.
3. The current time range selection could sometimes be lost when choosing an alternate view graph.

---> C00662 - General - Enhancement
For all graphs in all components, added a function available by right-clicking on the graph that let's you change any parameters defined in the SQL statement such as the taskcount, asp number or time range. This is available under the "Change SQL Parameters" menu.

---> C00662 - General - Bug Fix Request
When looking at a vertical bar graph that has a secondary axis consisting of only a single bar, the 2nd axis was not visible. With this build, a point is drawn in the middle of the single bar at the height appropriate for the axis.

August 22, 2008 – C00659

Job Watcher changes

---> C00659 - Job Watcher - Bug Fix Request
Changed usage of file QSYS/QADBXREF which is public *EXCLUDE to QSYS/QADBXLFN so that users without *ALLOBJ authority can run Job Watcher without authority errors. This affected the program that builds the list of collections in a library and the Create Job Summary view.

---> C00659 - Job Watcher - Bug Fix Request
In JW 540 and 610, fixed the following issues in the create job summary function.
  -- In the thread totals and job totals files it was possible to have incorrect metrics shown because the same taskcount number was used in different collections but for completely different jobs.
  -- Removed the Contributing collections field from the Job totals file since it wasn't accurate and there wasn't a good way to correct it.
  -- Runtimes were not accurate in the Thread totals and Job totals files. The times were the run times of all threads in all collections added together.

---> C00659 - Job Watcher - Enhancement
In JW 540 and 610, when viewing the Create Job Summary reports, "Thread totals by collection" and "Job totals by collection", right-clicking a job or thread in the report will provide the graphs/reports for the selected collection/thread over time as well as the collection overview graphs.

---> C00659 - Job Watcher - Bug Fix Request
In JW 540 and 610, when viewing the Create Job Summary reports, double clicking a record will no longer produce an SQL error and incorrectly try to show the interval detail property pages.

---> C00659 - Job Watcher - Bug Fix Request
In JW 540 and 610 for summarized collections only, on the interval summary interface, on the holding threads/tasks tab, the Holder menu options for a record in the list was not showing up.

---> C00659 - Job Watcher - Bug Fix Request
In JW 540 and 610, the CPYWCHCMD (split function) was not setting the TO Text description parameter correctly. This caused the collection size to not be correct for split collections at that release.

---> C00659 - Job Watcher - Bug Fix Request
For V6R1 JW,the collection menu Wait graphs -> by thread -> Advanced -> Thread signatures ranked by Disk op-start contention was listed twice.

---> C00658 - Job Watcher - Bug Fix Request
Modified the SQL statement behind the V5R4/V6R1 JW, summarized collection overview time signature so that it does a left outer join instead of an inner join.

---> C00658 - Job Watcher - Bug Fix Request
The JW command DLTJWCOL will now delete all QJW_* user spaces in the collection library. These unnecessary userspaces are generated by Job Watcher and were not being cleaned up.

---> C00658 - Disk Watcher;Job Watcher - Enhancement
For DW and JW monitors, all releases, added some extra error handling to the start and restart monitor commands. If an error occurs when trying to read the control file to determine the monitor's status, the monitor will send a message to the QSYSOPR message queue, wait 10 minutes then try to read the control file again. This will allow the monitor to keep running should the error resolve itself (if it's a locking issue) or through system operator intervention.

The latest server builds are needed to enable this change.

---> C00658 - Disk Watcher;Job Watcher - Usability
Deleting a monitor now will happen "in batch" through the remote command status view, to not tie up the GUI thread.

---> C00656 - Job Watcher - Usability
The menu options for a super collection should be working properly now. They are all tied to the Job Watcher portion of the super collection now and previously some were not.

---> C00656 - Job Watcher - Usability
When working with super collections, if the Job Watcher collection under the super collection has been deleted, all the menu options are now disabled except for Delete.

---> C00656 - Disk Watcher;Job Watcher - Bug Fix Request
At V5R3 and V5R4 JW, the monitor would fail to collect data using the correct collection name if a 10 character library name was used along with a 8 character monitor name.

The 1st collection (and all others) would be named ABCDEFGH0 instead of ABCDEFGH01. This caused only the most recent collection of the monitor to be available as the same collection name was repeatedly reused.

---> C00656 - Job Watcher - Bug Fix Request
In JW 530+ , the collection size listed for split collections is inaccurate.

---> C00655 - Job Watcher - Bug Fix Request
In JW, when displaying collections in a library and the library unexpectedly contains a member with no RUNI record, the stored procedure would fail and no collections would be displayed.

---> C00652 - Job Watcher - General
When deleting a monitor the command will no longer try to end the monitor if the status file indicates that the job is not active.

---> C00652 - Job Watcher - Usability
In JW call stacks , the menu Display Full Procedure Name will always be enabled now. Previously it was only enabled if the procedure was very long.

---> C00652 - Job Watcher - Bug Fix Request
Fixed a bug in the JW V6R1 server code that was causing a collection to fail to be deleted.

---> C00651 - Job Watcher - Bug Fix Request
In V6R1 JW, in the add definition wizard, specifying specific jobs would always also include all tasks by default. I've changed this so that tasks are not included by default to be consistent with the behavior of the interface at the previous release.

---> C00651 - Job Watcher - Enhancement
For V5R3/V5R4 JW, on the Add Jobs window of the Collection Wizard. Specifying a Current user profile along with a job filter and pressing the Add button will allow you to include jobs that match the job filter AND the current user profile. Previously the current user profile field only modified the list of jobs shown when pressing the Refresh button.

It's not possible to select jobs matching the job filter AND the current user profile at V6R1 due to a change in the commands at that release. You can select the current user profile separately from the job name at that release but this will give you selections matching either condition, not both.

---> C00649 - Job Watcher - Bug Fix Request
In JW 530+, in the seizes and locks time signature graph, the default drill down when clicking on an object lock or db record lock bucket would incorrectly display the seize contention bucket.

---> C00649 - Job Watcher - Usability
In JW when running an SQL statement and one of the files doesn't exist, the error given should indicate which file in which library doesn't exist, instead of the cryptic message "*FIRST in *N type *MEM not found.".

---> C00649 - Job Watcher - Bug Fix Request
In 530+ JW, right clicking on a wait bucket ranking graph, would set the selected thread/taskcount in memory, and then when performing an alternate view for the same graph and the query is reran, the selection would incorrectly act as a filter, excluded all the jobs/threads except the previously selected one.

---> C00649 - Job Watcher - Bug Fix Request
For V6R1 JW, the graph, Thread signatures ranked by bucket type, the bucket "other waits" was not shown when it should have been. Same problem also occurred in its drill down graph.

---> C00648 - Disk Watcher;Job Watcher;Pex Analyzer - Enhancement
At V5R3+, modified the super collection wizard on the basic options screen so that the user can pick between IBM-supplied or user-defined definitions.

Also added code that will apply IBM-supplied definitions for JW/PA/DW.
If no IBM-supplied definitions are found they will be automatically loaded upon entering this screen, but if some are missing, they can be reloaded using the action button's reload IBM-supplied definitions option for each collection type.

If you have a preference for additional IBM-supplied definitions that should be added (especially for PEX), let me know.

---> C00643 - Job Watcher - Bug Fix Request
In JW 540 and higher, fixed a bug where a single task would be erronesouly listed in the create job summary reports Job totals by collection and Job totals with bad data. Tasks are only shown in the thread totals by collection or thread totals reports.

---> C00642 - Job Watcher - Bug Fix Request
The summarize in batch option for V6R1 JW, wasn't working. The user would get an error file QAIDRJWSQL not found.

---> C00640 - Job Watcher - Bug Fix Request
The JW transfer function was broken at V5R4 and V6R1 if the target release was a V5R4 or V6R1 system.

---> C00640 - Job Watcher - Bug Fix Request
Changed the stored procedure that returns the list of JW collections in a library because it was failing to return anything on 1 V6R1 system.

---> C00640 - Job Watcher - Bug Fix Request
Fixed a bug that would occur when trying to view a V5R4/V6R1 JW collection on a V5R3 system. An error stating that stored procedure QIDRJWCHKS could not be found would be shown.

---> C00638 - Disk Watcher;Job Watcher - Enhancement
Added a new function that generates an HTML report consisting of the desired set of summary graphs. Use the "Generate Reports..." menu for a DW, PEX, CSI or JW collection.
The default web browser will be opened to show the HTML report after the reports have been captured to JPG files.

If you wish to send the report to another user, you can either print to PDF if you have a PDF writer installed , or send the HTML file + JPG files from the directory specified on the report generator window.

Note: As reports are running you can switch to other windows, but right before screenshots are taken the data viewer must move to the front of all windows in order to get the screenshot. This will happen automatically but may be concerning to some users.

---> C00637 - Disk Watcher;Job Watcher - Bug Fix Request
The display job log menu option for a monitor that has ended is working again.

---> C00637 - Disk Watcher;Job Watcher - Enhancement
Added a Start Monitor... menu to the compoent icon's for Job Watcher and Disk Watcher.

---> C00637 - Job Watcher - Usability
On the iDoctor components window, if the Job Watcher server code is not installed to the latest required level, the Disk Watcher and CSI components will be avaiable (though the server code should still be installed to avoid problems).. Previously CSI and DW components would only be shown if Job Watcher was listed as "Available".

---> C00636 - Job Watcher - Bug Fix Request
For V6R1 JW, the copy collection function didn't work.
Latest server builds are required in order to fix this.

---> C00635 - Job Watcher - Bug Fix Request

In JW or CSI, if the graph shows CPU utilization, and the actual value is greater than 100%, switch to use automatic scaling so that the values above 100% are visible.

---> C00635 - Job Watcher - Enhancement
In V5R3/V5R4 JW, in the Start Collection Wizard on the options page, setting the definition to *NONE, will reset the parameter values in the wizard back to the defaults.

---> C00634 - Job Watcher - Enhancement
In JW, 530 and higher, from the collection overview graphs, two new reports have been added and are usable if the JW collection is a super collection containing PEX data with events: BASEVT *PRCCRT, BASEVT *PRCDLT

The reports show missing CPU and I/Os not reported by JW in the Collection Overview graph. Time periods where this has happened can be identified by comparing the average collection CPU utilization vs the average partition CPU utilization. When the partition CPU utilization is higher these reports could yield useful information.

The drill down menus are:

Detail reports -> Missing contributors -> Ungraphed detail detected by PEX collection
Detail reports -> Missing contributors -> Ungraphed totals detected by PEX collection

The CPU and I/O values shown are estimates.

---> C00634 - Job Watcher - Bug Fix Request
Corrected an SQL error with the V5R4 and higher JW report
Details reports -> SQL -> SQL Statements

---> C00634 - Job Watcher - Bug Fix Request
In JW, server side output files view, the delete function for output files was enabled when it should not have been.

---> C00634 - Job Watcher - Bug Fix Request
For V5R3 JW and higher, on the interval details, JVM tab, the JVM handle was not being shown correctly for a classic JVM.

---> C00634 - Job Watcher - Enhancement
For V6R1 JW, added a set of J9 JVM graphs, JVM statistics on the Java Virtual Machine interval details tab and J9 call stacks on the call stack tab. The J9 Java entries are embedded within the regular JW call stacks. J9 java call stack entries aren't usable with the call stack reports.

Note: from the J9 JVM rankings graphs, to see the worker threads for the primary JVM thread, use the menu Selected thread -> J9 JVM Graphs -> J9 JVM threads. From there you can see the call stacks for the desired thread.

---> C00633 - Job Watcher - Enhancement
Updated several icons in the GUI: the Job Watcher icon, the JW/DW/CS collection icons and the super collection icons.

---> C00633 - Job Watcher - Bug Fix Request
Fixed a couple of SQL parsing problems that were causing the wrong data/member to be shown in two places in Job Watcher if the library contains multiple collections.

The start time and initilization time on the Collection Properties - General tab.

The JVM overview graphs.

For the JVM graphs if you are on an older client the workaround is the following:

Open the SQL editor
Right-click the editor, Find....
Search on QAPYJWINTI) INTI

Add a new line or two between QAPYJWINTI and the)
Hit F4 or F5 to rerun the query and the correct member for this file will be picked up.

---> C00632 - Job Watcher - Enhancement
Added support to define Job Watcher super collections. These collections include Job Watcher data, but also allows for collection of Disk Watcher, PEX Analyzer and/or CSI data at the same time. All collections will have the same name, and reside in the same library.

Right-click on the Job Watcher icon and choose Start Super Collection... to access this function.

Super collections are stored in the new Super Collections folder under the Job Watcher icon.

Expanding a super collection reveals a folder for each collection type available with each containing the reports available from the applicable iDoctor component.

The super collection is monitored by job QIDRSSMON. Once the Job Watcher collection ends, then all other collections will be end as well (CS data will be copied, PEX collection will end, Disk Watcher will end).

File QUSRSYS/QAIDRSUPER keeps track of the super collections on a system.

The latest server builds must be installed in order to use this function successfully (works at V5R3 and higher).

---> C00631 - Disk Watcher;Job Watcher;Pex Analyzer - Enhancement
Changed the interface that displays a definition's properties in PEX Analyzer to match the same look as the one used in Job Watcher. Also created an interface to display a Disk Watcher's definition properties.

---> C00631 - Job Watcher - Enhancement
Modified the Job Watcher component view so that the Job Watcher icon now contains several views each containing a different logical element of Job Watcher from the current system. The views are:

Libraries - Filterable list of libraries on the system containing Job Watcher data.
Monitors - Allows you to work with the monitors on the system. You can also double-click on a monitor now to see the collections and reports within.
Super Collections - Allows you to work with the super collections on a system. These contain a folder for each collection type collected within the super collection.
Definitions - List of JW Definitions on the system.
Scheduled Jobs - Allows you to work with any iDoctor scheduled jobs on the system. This interface replaces the use of the iDoctor Scheduled Jobs View.

---> C00630 - Collection Services;Job Watcher - Bug Fix Request
Fixed a bug in CSI and JW where drilling down into rankings graphs would occassionally include the interval before the selection.

---> C00630 - Job Watcher - Testing
Fixed a JW 540 and higher bug where users would get an SQL error on the Classic JVM graphs if the aliases from the summary process have not been created.

Collection Services Investigator changes

---> C00656 - Collection Services - Enhancement
In CSI, from the CS objects folder, added an option to create a collection from a MGTCOL object. A window is shown allowing the user to specify the destination library and collection name along with the description. This will then run a CRTPFRDTA command with the desired parameters.

---> C00639 - Collection Services - Enhancement
CSI collections now offer a report menu containing the complete set of summary graphs similiar to the menu provided in JW/DW.

---> C00639 - Collection Services - Enhancement
Added a new function that generates an HTML report consisting of the desired set of summary graphs. Use the "Generate Reports..." menu for a DW, PEX, CSI or JW collection.

The default web browser will be opened to show the HTML report after the reports have been captured to JPG files.

If you wish to send the report to another user, you can either print to PDF if you have a PDF writer installed , or send the HTML file + JPG files from the directory specified on the report generator window.

Note: As reports are running you can switch to other windows, but right before screenshots are taken the data viewer must move to the front of all windows in order to get the screenshot. This will happen automatically but may be concerning to some users.

---> C00635 - Collection Services - Enhancement
In JW or CSI, if the graph shows CPU utilization, and the actual value is greater than 100%, switch to use automatic scaling so that the values above 100% are visible.

---> C00633 - Collection Services - Enhancement
The folders available under CSI have signficantly changed. Instead of showing libraries containing CS data, the following folders are available:

Libraries - Libraries containing CS database file collections (filterable).

CS objects - A list of all CS management collection objects on the system

Super collections - Work with the super collections on the current system

Scheduled jobs - iDoctor jobs scheduled to run

---> C00632 - Job Watcher - Enhancement
Added support to define Job Watcher super collections. These collections include Job Watcher data, but also allows for collection of Disk Watcher, PEX Analyzer and/or CSI data at the same time. All collections will have the same name, and reside in the same library.

Right-click on the Job Watcher icon and choose Start Super Collection... to access this function.

Super collections are stored in the new Super Collections folder under the Job Watcher icon.

Expanding a super collection reveals a folder for each collection type available with each containing the reports available from the applicable iDoctor component.

The super collection is monitored by job QIDRSSMON. Once the Job Watcher collection ends, then all other collections will be end as well (CS data will be copied, PEX collection will end, Disk Watcher will end).

File QUSRSYS/QAIDRSUPER keeps track of the super collections on a system.

The latest server builds must be installed in order to use this function successfully (works at V5R3 and higher).

---> C00630 - Collection Services;Job Watcher - Bug Fix Request
Fixed a bug in CSI and JW where drilling down into rankings graphs would occassionally include the interval before the selection.

Disk Watcher changes
---> C00659 - Disk Watcher - Bug Fix Request
In DW, using the select fields menu against a library would show a blank field at the bottom.

---> C00659 - Disk Watcher - Enhancement
In DW, added 3 new graphs to Disk Watcher stats mode -> by interval.

- Page count rates categorized totals
- Page count rates categorized writes
- Page count rates categorized reads

---> C00658 - Disk Watcher - Bug Fix Request
Fixed a problem with the DW stored procedure that lists collections in a library. An SQL error would occur if a collection name was an SQL reserved word.

---> C00658 - Disk Watcher - Bug Fix Request
Fixed an SQL error that could occur with stored procedure QIDRDWGRI, if the collections in a Disk Watcher monitor were viewed under the Monitors library in Job Watcher, before the Disk Watcher component was launched.

---> C00658 - Disk Watcher;Job Watcher - Enhancement
For DW and JW monitors, all releases, added some extra error handling to the start and restart monitor commands. If an error occurs when trying to read the control file to determine the monitor's status, the monitor will send a message to the QSYSOPR message queue, wait 10 minutes then try to read the control file again. This will allow the monitor to keep running should the error resolve itself (if it's a locking issue) or through system operator intervention.

The latest server builds are needed to enable this change.

---> C00658 - Disk Watcher;Job Watcher - Usability
Deleting a monitor now will happen "in batch" through the remote command status view, to not tie up the GUI thread.

---> C00658 - Disk Watcher - Bug Fix Request
Fixed a chance that the Disk Watcher collections under a monitor would not be visible (if a different collection had the same name prefix as the monitor).

---> C00656 - Disk Watcher;Job Watcher - Bug Fix Request
At V5R3 and V5R4 JW, the monitor would fail to collect data using the correct collection name if a 10 character library name was used along with a 8 character monitor name.

The 1st collection (and all others) would be named ABCDEFGH0 instead of ABCDEFGH01. This caused only the most recent collection of the monitor to be available as the same collection name was repeatedly reused.

---> C00655 - Disk Watcher - Bug Fix Request
In DW, when displaying collections in a library and the library unexpectedly contains a member with no RUNI record, the stored procedure would fail and no collections would be displayed.

---> C00648 - Disk Watcher;Job Watcher;Pex Analyzer - Enhancement
At V5R3+, modified the super collection wizard on the basic options screen so that the user can pick between IBM-supplied or user-defined definitions.

Also added code that will apply IBM-supplied definitions for JW/PA/DW.
If no IBM-supplied definitions are found they will be automatically loaded upon entering this screen, but if some are missing, they can be reloaded using the action button's reload IBM-supplied definitions option for each collection type.

If you have a preference for additional IBM-supplied definitions that should be added (especially for PEX), let me know.

---> C00639 - Disk Watcher - Bug Fix Request
When creating a disk watcher collection with a range condition, the final screen would list the range type incorrectly if the limit type selected on the previous screen didn't equal the range type.

---> C00638 - Disk Watcher;Job Watcher - Enhancement
For DW and JW collections at V5R3 and higher, added a new function that generates an HTML report consisting of the desired set of summary graphs. Use the "Generate Reports..." menu for a DW or JW collection. Support for this function in CSI, will come in a future update.

The default web browser will be opened to show the HTML report after the reports have been captured to JPG files.

If you wish to send the report to another user, you can either print to PDF if you have a PDF writer installed , or send the HTML file + JPG files from the directory specified on the report generator window.

Note: As reports are running you can switch to other windows, but right before screenshots are taken the data viewer must move to the front of all windows in order to get the screenshot. This will happen automatically but may be concerning to some users.

---> C00637 - Disk Watcher;Job Watcher - Bug Fix Request
The display job log menu option for a monitor that has ended is working again.

---> C00637 - Disk Watcher - Enhancement
Updated the graphs in DW trace mode so that they include the following:

I/O counts categorized totals
I/O counts categorized writes
I/O counts categorized reads

I/O time categorized totals
I/O time categorized writes
I/O time categorized reads

Also removed the graphs
I/O rates categorized others
Deferred queue time ratio categorized others

---> C00637 - Disk Watcher;Job Watcher - Enhancement
Added a Start Monitor... menu to the compoent icon's for Job Watcher and Disk Watcher.

---> C00637 - Disk Watcher - Enhancement
Scheduling a DW monitor should be working now.
Also updated the scheduled jobs view to recognize scheduled Disk Watcher monitors jobs.

---> C00637 - Disk Watcher - Enhancement
The DW monitor support is now available for release V6R1. Latest JW server builds required.

---> C00637 - Disk Watcher - Enhancement
Add new trace DW drill down showing top 25 IOs

---> C00636 - Disk Watcher - Usability
In the DW Stats graphs, changed the graph titles so they matched the naming convention used by the trace graphs. Also changed usage of 'pool' to 'disk pool', 'unit' to 'disk unit' and 'path' to 'disk path' in the graph titles throughout DW.

---> C00636 - Disk Watcher - Bug Fix Request
Made a change to DW when it lists the collections in the library at V5R4 and higher. If the stored procedure that is normally used to build the list fails, the SQL method through the client will be attempted instead.   This allows DW collections to be viewable on idoc540 despite the sql system error/problem there.

---> C00636 - Disk Watcher - Enhancement
Added Disk Watcher monitor server commands to releases V5R3 and V5R4. Commands for V6R1 will come in a future update.

These commands are very similar to the Job Watcher monitor commands and include:

STRDWMON, RSTDWMON, HLDDWMON, RLSDWMON, ENDDWMON, DLTDWMON

Also added support in the GUI to work with or start monitors in Disk Watcher for either Job Watcher or Disk Watcher. The same Monitors folder is also available in Job Watcher which lets you work with Disk Watcher monitors from Job Watcher component if you wish.

The interface for starting/restarting a monitor has been redesigned into a 4 page wizard which lets you start a monitor for both types at once if you desire.

---> C00636 - Disk Watcher - Enhancement
Added Disk Watcher monitor server commands to releases V5R3 and V5R4. Commands for V6R1 will
come in a future update.

These commands are very similar to the Job Watcher monitor commands and include:

STRDWMON, RSTDWMON, HLDDWMON, RLSDWMON, ENDDWMON, DLTDWMON

Also added support in the GUI to work with or start monitors in Disk Watcher for either Job Watcher or Disk
Watcher. The same Monitors folder is also available in Job Watcher which lets you work with Disk Watcher
monitors from Job Watcher component if you wish.

The interface for starting/restarting a monitor has been redesigned into a 4 page wizard which lets you start
a monitor for both types at once if you desire.

---> C00636 - Disk Watcher - Bug Fix Request
Made a change to DW when it lists the collections in the library at V5R4 and higher. If the stored procedure
that is normally used to build the list fails, the SQL method through the client will be attempted instead.   This
allows DW collections to be viewable on idoc540 despite the sql system error/problem there.

---> C00636 - Disk Watcher - Usability
In the DW Stats graphs, changed the graph titles so they matched the naming convention used by the trace
graphs. Also changed usage of 'pool' to 'disk pool', 'unit' to 'disk unit' and 'path' to 'disk path' in the graph
titles throughout DW.

---> C00636 - Disk Watcher - Bug Fix Request
Corrected an overflow problem in the DW trace summary function which could occur due to some numeric
fields being too small which caused records not to be written to the summary file. When this occurred,
graphs and drill downs reported fewer I/Os, smaller response times than expected.

---> C00634 - Disk Watcher - Enhancement
In DW, added the hardware resource file QAPYDWHDWR to the list of files available in the "Collection files"
menu for stats collections

---> C00634 - Disk Watcher - Bug Fix Request
When right-clicking on DW collection then selecting 'Collection files' you see the hardware file listed as
QAPYDWHDW... it should be QAPYDWHDWR... The selection of the file will fail because that file does not
exist.

---> C00633 - Disk Watcher - Bug Fix Request
In Disk Watcher V5r3/V5R4, the command string shown in the Creation Settings page would be incorrect if
the COLCOND, COLRNG parameters were used. The *RSPTIME and *SRVTIME parameters were listed
as *RESPTIME and *SERVTIME.

Note: There is still a known problem where the COLRNG parameter will combine the low and high range
values together (missing a space). Fixing that would require a PTF.

---> C00631 - Disk Watcher;Job Watcher;Pex Analyzer - Enhancement
Changed the interface that displays a definition's properties in PEX Analyzer to match the same look as the
one used in Job Watcher. Also created an interface to display a Disk Watcher's definition properties.

---> C00631 - Disk Watcher - Enhancement
Modified the Disk Watcher component view so that the Disk Watcher icon now contains several views each
containing a different logical element of Disk Watcher from the current system. The views are:

Libraries - Filterable libraries containing Disk Watcher data as before.
Definitions - List of Disk Watcher Definitions on the system.
Scheduled Jobs - Allows you to work with any iDoctor scheduled jobs on the system. This interface
replaces the use of the iDoctor Scheduled Jobs View.

---> C00631 - Disk Watcher - Enhancement

Added the ability to define and work with (delete/change) definitions in DW for V5R3 and V5R4.

---> C00631 - Disk Watcher - Bug Fix Request
Using the DW Add Definition Wizard's Conditional Collection options wouldn't work if the response time or service time options were used at V6R1 because the command string wasn't being generated correctly.

General iDoctor GUI changes

---> C00658 - General - Enhancement
Added a preference to the misc tab, called "Confirm when stopping the Report Generator function". When checked Closing the Data Viewer or Component View performing the report generator function will show a confirmation prompt before the report generator is stopped.

---> C00658 - General - Bug Fix Request
Fixed a bug in the report generator. If a user started the report generator, and while it's running closed down either the Data Viewer or the component view it was started from, the iDoctor client would crash.

---> C00658 - General - Bug Fix Request
Fix an issue with the transfer library function where it would not prevent a user from trying to transfer their library to the current system.

---> C00658 - General - Usability
When a graph is running or returns no data, the graph area will contain the graph title as well as the status of the graph. Made this change, so it would be captured and available for users of the Generate Reports function.

---> C00658 - General - Usability
If the Preference on the Misc tab - Maximize reports opened into Data Viewers is checked,
then all views including property sheets (call stacks) and job logs, opened in the Data Viewer will also be maximized. Made this change because otherwise when opening a call stack or record quick view, all graphs would no longer be maximized.

---> C00658 - General - Enhancement
The iDoctor Update History window, now offers buttons to find the next or previous occurrence of the desired search string.

---> C00656 - General - Bug Fix Request
Scrolling through a table or graph should be slightly faster now. 2 fetches were being done per page instead of just 1.

---> C00652 - General - Bug Fix Request
Made a change to the ODBC driver/connection properties so that extended dynamic package support is no longer being used.

I did this because of the somewhat random SQL0440 errors on some stored procedures called by iDoctor because of a possible bug with SQL packages at V6R1.

---> C00652 - General - Usability
If the required server code isn't found, updated the message so it's phrased as a recommendation to install the latest server builds rather than a requirement and lists the functions changed/added by the recommended server code.

---> C00649 - General - Bug Fix Request
When saving user-defined queries/graphs that utilize summary files (which contain SQL long tables that contain the collection name as part of the name) they will be usable against other collections in other libraries or the same library .

The collection name will be resolved when the query is opened. Previously these queries would only be usable if accessed against the same collection.

---> C00648 - General - Enhancement

On the report generator window, added options to save or load the current list of reports to generate based on a user-defined name.

Also added a button to toggle the graphs available from checked/unchecked.

---> C00646 - General - Bug Fix Request
Fixed a possible crash that could occur when opening a graph if the code that displays the graph flyover is invoked before the thread that retrieves the data completes.

---> C00645 - General - Bug Fix Request
When scrolling through large tables, sometimes data would incorrectly appear out of sequence because the fetched buffer from ODBC could get into a state where it would contain the wrong data every other fetch. The solution of reissuing the same fetch again to get the correct data works, but does cause extra network delay.   Need to investigate if the actual cause is iDoctor code, ODBC driver or system DB PTF level.

---> C00640 - General - Enhancement
Added a preference to the Misc tab called "Maximize reports opened into Data Viewers" which will auto maximize all reports opened into Data Viewers. The default is off/unchecked.

---> C00640 - General - Enhancement
Moved the preference "Always open a new report into an existing Data Viewer" from the Job Watcher tab to the miscellaneous tab. It also should be working properly from collection right-click context menus for all components.

---> C00640 - General - Enhancement
Removed the disable caching preferences on the PEX Analyzer and Job Watcher preferences tabs and created a single preference on the miscellaneous tab to replace it that applies to all components.

---> C00639 - General - Enhancement
On the main window, the clock icon to set the preferred time range interval size is now settable from any component. Made this change since now PEX Analyzer will also have time range graphs.

---> C00637 - General - Usability
Since a lot of new users are familiar with the common practice of left clicking a column twice to sort in descending sequence, this will now work in iDoctor. Previously left clicking would always sort ascending.

Note: You can right-click the column to sort descending, instead of left clicking twice.

---> C00637 - General - Bug Fix Request
When viewing a spool file (job log) from the server in the GUI, corrected some formatting problems.

---> C00636 - General - Usability
Made a change to the scheduling window so that the checkbox at the top is checked by default the 1st time the window is shown in order to show the calendar and time fields.

---> C00636 - General - Usability
Made a change to the scheduling window so that the checkbox at the top is checked by default the 1st time the window is shown in order to show the calendar and time fields.

---> C00635 - General - Enhancement
If the graph legend is not currently shown (via toolbar button) or set to 0 width in preferences, expanding the graph or data viewer should not display the legend.

---> C00634 - General - General
Moved all iDoctor stored procedures into library QIDRGUI instead of library QUSRSYS.

---> C00633 - General - Bug Fix Request
The schronized table view wasn't scrolling down correctly, it would snap back up after scrolling down. Scrolling via the graph scroll bar would scroll the table correctly though.

Apr 30, 2008 – C00618

Job Watcher changes

When creating or restarting a monitor, the check on the definition to determine if the definition is smaller than the collection size was incorrect, causing a message to be incorrectly displayed stating that the definition size was only for 1 minute duration.

Fixed a bug where the "flyover fields" section of the graph quick view property page would be incorrect.

When using the JW work with monitors view on a system that has never created a monitor before removed an error message stating that the monitor control file QAIDRJWM1 in QGPL could not be found.

In JW V5R4 and higher, added extra checks to see if a summary is still in progress or not.
  An SQL table QAIDRJWSUM_COLNAME_INPROGRESS containing the job running the summary will exist while the summary is in progress.
  This will allow you to tell if a summary is running in the GUI, that another user started.
  The summarized status field within the list of collections will show "In progress" if the summary is currently running.

Collection Services Investigator changes

Fixed a bug in CSI, in the Create job Summary function at V5R4. If the collection contained 32 wait buckets in file QAPMJOBWT, the CPU and IO numbers would be double what they should have been.

Disk Watcher changes

When bringing up Disk Watcher for the 1st time on a system, the Start Disk Watcher Wizard is shown, after submitting a collection for the 1st time if it failed, or a refresh is done before the collection files are created, don't show the wizard a 2nd time.

When deleting a DW collection, file QAPYDWHDWR was missing from the list of files to delete the collection members from.

For V5R3 and higher Disk Watcher, added an option in the Start Disk Watcher Wizard to collect the hardware resource file.

This option is on the advanced page at V5R3 and V5R4, and on the basic options page at V6R1 (since there isn't an advanced option page at V6R1)

General iDoctor GUI changes

Fixed a bug where it was possible to get graph tooltips to show up through the main window or other application windows.

Synchronized table disappears when SQL Editor is opened/closed

When opening or resizing a graph view while the synchronized table view is minimized, the status bar would incorrectly display a negative number when showing the number of bars/records visible of the total.

Updated the iDoctor Components window so when it is resized larger, the list of components resizes bigger and also the window cannot be resized smaller than its initial size.

Made a bug fix to query parsing where the libraries weren't being filled in from a user-defined query. The parsing failed because there was a new line and not a space before the GROUP BY.

In graph views, only show flyovers if the graph has the current focus. In other words if a graph is behind another window, don't show flyovers for the graph.

We are making this change to help alleviate problems we sometimes have with flyovers showing up when we don't want them to.

Automatically add single quotes for string fields entered into the Record Selection page of the Query definition interface.

When resizing graph views that have a synchronized table at the bottom, need to make sure the synchronized table doesn't become visible when resizing the graph larger unless the button on the toolbar has been pressed to show the table. It tended to appear at undesirable times.

Can't get graph or table tooltips to show up if the Data Viewer is moved to the 2nd monitor in a multi-display environment.

Fixed a crash when using the Query definitions Manage interface, changing the local database, and then using the Manage definitions interface again.

Changed the component menu User-defined reports -> Query(graph) definitions -> Manage... to Manage/Export... to make it more obvious this is where the export function resides.

In the remote SQL Statement Status View, using the display job log option would often not work, returning the error, The job number for the ODBC job could not be retrieved.

Add the Default time range interval size setting to Preferences on the Display tab.
This preference can also be set using the clock icon on the Main Window's toolbar. Also when working with graphs and using the time range interval size button, added a menu at the top "Set current as default" that will make the current setting the default for future graphs opened.

Apr 21, 2008 – C00612

Job Watcher changes
When working with the definitions on a system for the 1st time when no definitions have been created yet, removed the error message file QAPYJDWDFN in QUSRSYS not found.

Added a create job summary function to Job Watcher that produces job totals for the desired jobs based on a generic job name or subsystem name. Use the Create Job Summary menu for a library or V5R4 or higher collection. The collection must be summarized in order to use this function.

For V5R4 JW and higher, changed the order that the summary program stored procedures are created in to fix a possible failure when updating stored procedures to the current levels. The situational analysis stored procedure must be created before the interval summary file procedure.

For JW V5R4 and higher, I've added support to submit the selected summary functions to batch to a QIDRJWSUM job. This is done by creating a file/member in QUSRSYS/QAIDRJWSQL containing the calls to the stored procedures that need to be executed.

A checkbox on the summary prompter window allows you to indicate if this should be done instead of running them in the Remote SQL statement status view (a separate GUI QZDASOINIT job).

Some advantages to running in batch are:
1) You can run start up many summaries simultaneously
2) You can start the summary then kill your GUI session and it will keep on running
3) You can start up multiple summaries on multiple systems without waiting on the remote sql statement status view to execute them in order.

A disadvantage is it's not as easy to tell which step you are on (although this can be deduced by looking at the status field and see which files aren't created yet).

For V5R4 JW and higher, added a new situation which will show up as a background color on the overview charts if found.:

 -- Problem: Contention on DB in use table possibly due to high number of opens and closes
 -- Recommendation: Reduce number of opens.

For V5R4 JW and higher, added a new situation which will show up as a background color on the overview charts if found.:

 -- Problem: high number of creates and/or deletes by multiple jobs,
 -- where all of the objects are owned by the same user profile.
 -- Symptom: contention on the user profile
 -- Recommendation: reduce the number of creates and deletes.
 -- If a temporary workfile is being created, clear the file and re-use rather than
 -- creating and deleting. If the creates and deletes cannot be reduced consider using
 -- multiple user profiles.
 -- Callstack: qdbcrtme, qdbdltfi, quscrtus or other create or delete programs
 -- Wait obj type: user profile
 -- Wait type : (100/Rex) Seize: exclusive

For V5R4 JW and higher, added a new situation which will show up as a background color on the overview charts if found.:

 -- Problem: Fixed allocated length setting on a varchar or lob type column is defaulted to 0 or is set too small.
 -- When the setting is too small, the dbio tasks are used to update the aux portion of the file.
 -- Updates are synchronous to keep the aux portion in synch with changes in the fixed portion of the record.
 -- Recommendation: Choose a larger setting for the fixed allocated length, a rule of thumb is to choose a length
 -- that 90% of the accesses will fit within.
 -- Symptom: many dbio tasks are in synch write wait for several intervals
 -- Task: DBIOnn tasks
 -- Wait: (167/SWt) Mainstore/logical-dasd-io: dasd write
 -- Callstack: dbxoio

Collection Services Investigator changes
No changes made

Disk Watcher changes

In DW 530 and higher for collections containing both stats and trace data, the stats graphs were not available if the collection had not yet been summarized.

General iDoctor GUI changes

Opening a server-side output file that had key fields which happened to be SQL reserved words would cause an SQL error.

Fixed some crashing problems that would occur if the connection to the server is lost when either shutting down the client or when attempting to reconnect. As long as the user/password entered previously is still valid, there should be no noticeable effect when automatically reconnecting besides a slight delay.

When performing collection or monitor scheduling on a daily or weekly basis, it appears that subsequent calls after the 1st would fail because the command to be executed was removed from the control file that keeps track of the commands to execute. In order to correct this, I've changed the command attached to the scheduled job so it now contains the actual SMBJOB command to execute rather than a call to program QIDRSTRSCD. Program QIDRSTRSCD is only used now if the command string scheduled exceeds the job scheduler limit of 512 bytes. In addition command strings are only added to file QGPL/QAIDRJBSCD if they exceed 512 bytes and are not removed by calling program QIDRSTRSCD.

Canceling a long running query and then attempting to run another query from the same connection would give the error "Prepared statement S0 in use".

The query parsing is a bit smarter now when looking for table names inside a select statement. Changes made to eliminate redundancy.

When launching a component and building the list of libraries, made a change so that the library descriptions are not retrieved until you need to see them. This means that I don't retrieve the library descriptions (for libraries not visible) until you try to scroll down. This offers a big startup performance improvement on some systems.

When closing a Data Viewer that contains actively running queries or graphs, it should no longer take two tries. Previously the 1st click on the X would close the views within the Data Viewer and the 2nd would close the Data Viewer itself.

Mar 27, 2008 – C00609

Job Watcher specific changes

In JW V5R3 and higher, on the interval details transaction tab, the labels were wrong for the values shown.

The tip of the day interface now displays the Tip number before the text script for each video.

Added a new tip of the day that describes how to use call stack reporting options from the procedure information file.

Increased volume on most of the Tip of the day videos.

When the install checks the PTFs after installing Job Watcher V5R4 and logs the results in the setup log file, the checking would incorrectly fail on PTF SI27652 due to a buffer being too small.

Collection Services Investigator changes

In CSI, V5R3 and higher, added a check to ensure that the 7 required CS files needed to use the reports exist. If any are missing, their names will be listed in the Status field for the collection and no graphs will be available for that collection.

General iDoctor GUI changes

On the iDoctor components window, if the expiration is set to 12/31/39 (effectively *NOEXP when generated the code) it will be shown as "Never" instead now.

On some PCs, users would get the message when starting iDoctor, "the application configuration is incorrect". This was due to iDoctor not being able to find all required DLLs. They have been moved to a different location and this problem should no longer occur.

 I'm enabling the query definition interface again, even if the statement contains comments or new line characters which can tend to break the SQL parsing which tells the interface how the query is structured. Several of the queries i tried worked fine though using the interface to change a query will mean any comments in the SQL editor will be lost (not permanently if saved, just for the current view). PLEASE NOTE: On many of the rankings charts the job name field you need to filter on is now called OBJNAME as part of the query rewrite for CSI and JW. If in doubt, double-click the X axis field in the legend to see the field name you need to perform your record selection on.

Mar 21, 2008 – C00606

Job Watcher specific changes

The V5R4 and higher JW summary program now builds the files with indexes. It also should run a bit faster now.

In V5R4 JW or higher, it's now possible to analyze multiple JW collections into the same graph. This has only been enabled for summarized collections at V5R4 in some select reports. To do this you need to select the desired collections from the main window (from the list side) right-click and choose the wait graphs -> by time interval -> collection overview time signature. This will ask you if you wish to combine the collections into a single graph. Indicate yes to perform this option. If the new JW summary option to create the aliases for the collection files has not been executed this option won't work. The resulting graph is time based so the data will be ordered as such.

When working with multiple collections in V5R4 and higher JW corrected problems with the wrong member being used for the following: the detail reports from the collection overview graph, the internal detail property pages and the interval summary property pages. Unlike the graphs that provide data for all selected collections, each of these functions will be based off of just 1 collection for the time period selected or clicked on in the preceding graph. Note: The detail reports drill downs will not provide information spanning multiple collections even if a time selection is made that exceeds a single collection..

In V5R4 and higher JW, for summarized collections, the Collection overview time signature graph now contains information about possible performance situations found during the summarization process. If your collection was already summarized with an older client, you must run the summary program and choose the option "Create performance situations file" in order for the SQL for this graph to continue to work. You can do this via the summary prompter and avoid rerunning the rest of the steps in the summary program. Any situations found are indicated by colored backgrounds behind the time periods of interest. Different colors are used for each. No drill downs into these situations are available yet, but the SQL behind the graph contains the fields SITSTRTIME (start time of situation), SITENDTIME (end time of situation), SITID (id number as shown above or X if multiple situations occurred), and SITTASKCOUNT (the taskcount of the job/thread that experienced this situation, or X if multiple situations occurred). Use the legend and add these fields as flyover fields or use the Duplicate as table menu to show the table view for the graph

In V5R4 and higher JW, for summarized collections, the following situations have been defined and might be shown in the collection overview time signature graph as different colored backgrounds for the time periods where the situation occurred.
1. DARK BLUE - Job(s) are waiting for faulting on the seize/lock table. This can indicate that the seize/lock table has become large due to some job having a high number of locks.
2. DARK GREEN - Job(s) appear to be constantly starting and ending commitment control. They should NOT be doing this.
3. YELLOW - Job(s) may be executing poorly implemented SQL statements.
4. LIGHT BLUE High percentage of jobs were NOT caught by Job Watcher because they started and ended within the collection interval size.
5. PURPLE - Force-end-of-data used to force data to disk. This can cause seize contention on the file in a multi-threaded application.

For V5R4 and higher JW, change the JOBINIT and JOBTERM fields in the QAIDRJWSUM_<MBRNAME> table so that they allow null values to avoid a warning message returned by the summary program in some cases.

Time gaps when analyzing multiple collections are now represented by a dashed vertical line. Currently this only applies to V5R4 and higher JW.

For JW V5R4 and higher, fixed a bug in the summary program that would cause the wrong data to be shown if multiple collections were summarized having the same name, residing in different libraries.

In V5R4 and higher JW, for summarized collections, the IFS rankings graphs didn't work due to an error in the SQL statement

In JW V5R4 and higher, added the creation of the performance issues file (a.k.a "Situational Analysis") QAIDRJWANL_<MBRNAME> to the summary program. Also added this file to the server side output files list.

In JW V5R4 and higher, fixed a bug that was causing the drill downs from the summary program (SQL tables) server-side output files folder from working because the SQL was trying to point to the wrong member.

In JW V5R4 and higher, the summary programs gap file QAIDRJWGAP_COLNAME, contains the following new fields at the end of the file: JOBSTART = Job started in previous interval JOBEND = Job ended in this interval. Note: Job Watcher cannot report any statistics in the partial intervals where the job started and ended.

For V5R4 JW and higher, added a new stored procedure to the summary program QIDRJWCA that builds aliases over the JW collection files using the naming convention filename_collectionname. (example: QAPYJWTDE_RUN1) These are needed by the GUI when graphing multiple collections at once.

For V5R4 JW and higher, when using the Summarize option, you will be prompted to indicate which functions of the summary program should execute. This will allow you to run the create alias step for several collections quickly without having to rerun the entire summary. The choices made are saved to the registry on the PC for the next time a summary is made.

In Job Watcher, when listing collections in a library for V5R4 and higher, updated the checking to see if the required summary files exist to be much more robust. 4 files are required now: QAIDRJWSUM_<MBRNAME>, QAIDRJWGAP_<MBRNAME>, QAIDRJWTL_<MBRNAME> and QAIDRJWANL_<MBRNAME>. The summarized field will provide a message indicating the files missing if only some of them are.

Corrected an SQL error when working with the thread wait time signature drill down from a wait bucket rankings chart that would occur if the Summary program, create aliases step hasn't been done yet (even when analyzing just 1 collection).

For V5R4 and higher JW, added field descriptions to the situation fields shown in the Collection overview time signature graph's table view.

In V5R4 and higher JW, added the job name/thread id for both taskcounts listed in the performance issues summary file QAIDRJWANL_MBRNAME. The file must be rebuilt in order for the Collection Overview time signature graph to work.

In V5R4 and higher JW, for summarized collections, the collection overview time signature graph now provides fly over information about any performance situations shown including a problem description of the situation and the job/thread associated with it.

In V5R4 and higher JW, for summarized collections, add the contributing jobs/threads value as a Y2 for the run/wait time signature graphs for a single object over time.

In V5R3 and V5R4 JW, when creating a collection using the auto generate collection name option, the collection properties creation settings tab would not be available because the member saved into QAIDRJWDFN was slightly different than the collection member.

Added 1-hour time intervals as a time range graphing option. I did this so it could be used in Collection Services Investigator but the option will also show up in JW.

When starting a collection in iDoctor, from the component icon, the following default libraries will be used unless changed by the user. Disk Watcher - QDWDATA PEX Analyzer - QPADATA Job Watcher - QJWDATA

The JW Tip of the Day will now open maximized by default so that the larger videos will be fully visible without having to resize the window.

Fixed a couple of instances where total writes was computed using TOTWRT field in the TDE file instead of the sum of 4 sync/async write fields. TOTWRT field should be ignored is the advice of the performance experts.

In V5R3 JW and higher, drilling into a single time interval will now correctly show the starting and ending interval time of the time range. Previously the start and end time would be listed as the end time of the interval.

In V5R3 JW and higher, made the To library default value when performing a collection split, the From library value instead of blank.

In V5R3 JW and higher, from the collection overview graphs, if a time range has been selected (by holding down shift + clicking), a Spilt Collection menu will be shown, allowing you to create a new collection from the one being analyzed based on the time period selected.

In V5R3 JW and higher, made significant improvements to the table views available under the "Detail Reports" menu from the Collection overview graphs. Several new reports have also been added including TPROF (top call stack detected while CPU in use), activation group, sockets and SQL information. Drilling down into these reports now applies to the entire time period selected instead of just the 1st interval of the time period selected. The following reports are now available (if the appropriate optional files have been generated in some cases): Interval -> Basic interval information (INTI file records) Interval -> CPU statistics (QAIDRJWCPU file records) Job -> Threads using CPU Job -> Tasks using CPU Job -> Threads using CPU when call stack taken Process -> Primary threads using CPU Process -> Primary threads using CPU (including TDE file fields) Process -> Activation group details Classic JVM -> Java virtual machine statistics Classic JVM -> Java threads waiting TPROF -> All procedures TPROF -> Classic JVM (Java) procedures TPROF -> All procedures by thread TPROF -> Classic JVM (Java) procedures by thread SQL -> SQL statements SQL -> Open cursor lists SQL -> Prepared statement area Sockets -> Socket details Sockets -> TCP endpoint details Sockets -> Socket and TCP endpoint details Waits -> Longest current waits for <<BUCKETDESC>> (V5R4+ only) Waits -> Longest "bad" current waits (V5R4+ only)

In V5R3 JW and higher, the Details Reports -> Process -> Activation group details report has two drill down reports available under the "Selected activation group" menu if you right click on a single record (activation group) in the table. This allows you to see the program information details for the activation group.

The JW command DLTJWCOL should now be more likely to be able to delete collections that experience iDoctor GUI QZDASOINIT job open cursors/files from them not being closed by the system.

 In JW V5R3 and up, for unsummarized collections, added the graph Thread wait time signature with dispatch CPU breakdown for <<JOBDESC>> ..

In JW V5R3 or higher, changing the alternate view for a thread signatures ranked by graph, would result in the number of buckets shown being incorrectly reduced. The X axis label would also incorrectly be shown as <<OBJTYPEDESC>>.

In JW V5R3 and higher, when displaying file QAPYJWTDE from the server side output files folder, i've removed the SEIZE and TOTWRT fields from the default list of fields displayed. These fields do NOT contain valid values and should never be used.

From the Job Watcher component view, when selecting multiple collections and using the summarize menu, all selected collections will be summarized instead of only the 1st one.

In the iDoctor application property pages for Job Watcher - Server configuration tab, modified labels in the Default job run priorities section so that collections are referred to as "collections" instead of "job watches"

Modified the reports under the menu Detail reports -> Job and Detail reports -> Process so that only records that had CPU > 0 are returned.   Also modified the Details reports -> Classic JVM -> Java virtual machine statistics so the data is sorted only by GC heap size rather than Interval first and then GC heap size.

JW auto refresh not scrolling to end of graph when looking at collection overview time signature
Made a change to the JW auto refresh support such that the views will refresh even if they don't have the current focus. Because of the multithreading improvements, this allows for many views to be open (tiled) in the same data viewer and they can be all refreshing periodically if the graphs shown are from an active JW collection.   It would be wise to be mindful of how many views you have open and refreshing, as the more views open, the more queries will be executed on the server.   You can change how often these views refresh on the Job Watcher preferences tab.

Fixed several problems with JW monitors at V5R3 and higher: Latest server builds are required to fix these issues. Deleting a collection with files open hid the collection from view in GUI, but didn't actually delete it. If this happens (such as when a monitor is running and trying to delete older collections), the collection will be flagged as unable to be deleted and the collection status will change to show this in the GUI. Using the hold option doesn't end the collection currently running. The collection kept running until the parameter specified on the definition's termination was met. To be safe, when holding a monitor, I'm attempting to end all 99 possible collections for the monitor. Ending a monitor only stopped the collection with the highest collection number rather than all 99 possible ones. So if two were running because of a transition, the older one would keep running. The start/restart monitor commands wasn't very good about checking if the monitor was held or released. Now it checks the status every 5 seconds to see if it changed instead of the collection duration that could've been an hour.

In the Job Watcher monitors view, replaced the field title "Collection Job" with "Monitor Job".

Don't show the JW tip of the day when displaying the collections in a JW monitor.

If a collection can't be deleted in JW at V5R3 and higher, the status will change now in the GUI to "Delete request failed". This requires latest server builds.

Displaying the list of collections in a Job Watcher library should be considerably faster now, especially on slower connections.

General iDoctor GUI changes

All graphs in iDoctor now contain an equivalent table view of the data behind the graph under the graph view. The table view under the graph is always initially hidden, but if you wish to compare the table data with the graph, you can open it up to see the data.   Rather than relying on a separate query/connection to produce the table, the table view is based on the exact data from the graph that is

already cached in memory.   This allows the scrolling to be much faster than it would be otherwise. The table and graph are synchronized in their scrolling and selections. Scrolling the table scrolls the graph, and vice versa. Note that sorting the table will also sort the graph (which may not be a desirable thing to do for time based graphs.) This can have some usefulness on rankings graphs to quickly change the ranking based on another statistic. The SQL editor above the graph is tied to both table and graph, and making changes there refreshes both the table and the graph. I've also added a toolbar button to show/hide the table portion for a graph.

This client adds a button to the main windows' toolbar that allows the user to set the desired views built when opening a graph between 1 of 2 modes. The default way is to build a SQL editor, graph and synchronized table. If this button is pressed, then this changes to 2 graphs views where the bottom graph allows you to specify a different alternate view using the Data Viewer's alternate view toolbar button.   Either graph can be changed using the alternate views or time range selection option to make comparisons. Scrolling is synchronized as long as the number of bars in each graph is the same. This works for CSI, DW and JW.

Fixed a crash that could occur if an iDoctor user queries a file containing a DBCLOB type field.

When scheduling a collection or monitor it's a two-part process, running the ADDJOBSCDE command and then adding a record to file QAIDRJBSCD that contains the command string tied to the scheduled job. Added a check to remove the scheduled entry from the system if the 2nd part failed.

The summary programs for JW, DW, HA and CSI are now run able for collections that were created on a different OS VRM. This means you can do a create job summary for V5R3 CS collections on a V5R4 system, or create a V6R1 JW summary on a collection on a V5R4 system. As part of this change the stored procedures generated in QUSRSYS are now named QIDRJW54*, QIDRCS54*, QIDRJW61*, etc

In the iDoctor application property pages for some components on the Server configuration tab, changed the section that displayed the required PTFs installed, to display the PTFs NOT installed if any are not currently installed.   The PTF check is performed when showing this screen in order to accomplish this. Note: Another reason I'm doing this change is because the PTFs listed previously were not correct if the system was running V5R3M5 or V5R4M5.

When starting the GUI, the client will automatically download the latest PTF information from the external iDoctor website. The PTF lists are saved to the window's registry and when a user attempts to start a collection; the PTF list will be checked against those on the system by updating the appropriate data areas on the server to contain the latest PTF levels. Previously the automatic PTF checking was dependent upon the server builds being reinstalled to update the PTF data areas on the system. Note: This function requires that the iDoctor GUI have access to the Internet, and if this is blocked by a firewall or not possible, it will not work. The PTFs shipped in the data area on the server will be checked for instead which may not be the latest ones.

Fix a bug where gaps could sometimes appear in the graphs, when drag scrolling through them.
If the connection behind the remote SQL statement status view is lost, another one will be automatically started in its place. Previously resolving this required restarting the client.

I've created a stored procedure that is used to speed up most of the SQL statements executed by iDoctor (including graphs/tables in the Data Viewer), by requiring less network traffic. This effected many functions in all components. The performance gains are greatest on slow connections.

When starting the client, a screen is shown to display the most recent updates made to iDoctor in a list. You can check the box to not show this screen again, and can also bring this back from the Miscellaneous tab of the Preferences interface.

 In the iDoctor application property pages - General tab, removed the iDoctor Server version as it is redundant with the OS version now (as of V5R3)

There should be slight speed improvements when opening or scrolling through graphs and tables in the Data Viewer (and elsewhere). I optimized the part of the code that looks up field indexes by field name to use a keyed hash table rather than loop through an array.

When installing the client, added a check to see if any instances of the iDoctor client are open, and if they are give the user a message to close them and abort the install.

The change environment option from the Connections view didn't work.

For the Query properties tab for a table or graph, added an apply button next to the Report description field. This will change the title of the report or graph the query applies to. Also made the SQL statement field read only since the SQL editor should be used instead now.

The save library window now lets you specify a target release earlier than the current release when connected to V5R4 or V6R1 systems.

When pressing help buttons when working on a V6R1 system, you should now go to the V6R1 documentation. Also when pressing help buttons, the documentation will now be looked for in the iDoctor install directory, and if not there it will be opened from the external website (slower). You can download the documentation to the iDoctor install directory using the documentation installer download from the page http://www-912.ibm.com/i_dir/idoctor.nsf/documentation.html

Added a button on the toolbar of the Data Viewer to show/hide the table behind the current graph view. Also fixed the SQL editor button to behave appropriately with the new graph/table splitter view.

Corrected a performance problem with the way bar selection was implemented. The more total records in the query behind the graph, the slower the graph would be when making bar selections. This began to become noticeable > 100,000 records or so.

For all components, fixed a problem that would occur when changing the library names in the SQL editor for the files queried behind a graph view. The overrides were not being updated to use the new files, and the previous library was being used instead. This would cause misleading results.

The SQL parsing code should be a bit smarter now when dealing with queries that contain both single member SQL tables and JW physical files. Previously problems could arise when working with the SQL editor in some situations.

Fix a bug defining an SQL statement containing an alias name longer than 10 characters would cause the querying parsing to try to incorrectly issue an override against the alias.

Jan 30, 2008 – C00588

The summary program for V5R4 Job Watcher should be working correctly with this client. In the previous client only (C00586), one of the required stored procedures was missing causing a failure when trying to summarize collections.

Corrected a problem where a graph would not be updated after modifying the SQL statement and changing the libraries listed in the query to a different library.

Jan 16, 2008 – C00586

The summary program for V5R4 Job Watcher is much faster now.

Dec 12, 2007 – C00585

When working with Time interval graphs using the Collected Interval Size mode, the drill down to rankings charts would sometimes include/exclude data outside of the interval selected. How much the time range was off would vary depending on various factors.

Nov 29, 2007 – C00584

Added a Tip of the day interface to the startup of Job Watcher that offers several short videos. If you choose not show this option again, it is accessible in the future using the Help -> Job Watcher Tip of the

Day menu from the iDoctor main window. These videos are installed on the PC under the C:\Program Files\IBM|iDoctor for iSeries\videos directory.

Made a significant change to the time range/time interval graphs. There is now a toolbar button on the main window that allows you to set the desired time range between 1 of 6 options with the default being 1 minute. This simplifies greatly the time range graph folders. Once into a time range graph in the Data Viewer, there is a new "time" icon on the toolbar you can click to cycle through any of the desired. This also applies to the JVM time range graphs (both summarized and detail level for 1 job).

I've combined the summary graph folders, wait graphs -> by time ranges and wait graphs -> by interval into a single folder called wait graphs -> by time interval and replaced both folders with this one. The new folder contains some of the more useful reports from the wait graphs by interval folder like the Current wait duration time signature but now allows it to be graphed with one of the desired time range graph interval sizes

Combined the summary graph folders, CPU graphs -> by time ranges and CPU graphs -> by interval into a single folder called CPU graphs -> by time interval and replaced both folders with this one.

Removed the following folders for summarized collections: IFS graphs -> by interval, JVM Graphs by interval and Other graphs by interval. These folders are replaced by the use of time interval graphs that allow the interval size to be set to the collected interval size

I've made the x-axis of the time range graphs vary based on the interval size. The x-axis contains lower level information like interval and timestamp as the interval size decreases and contains more information like the date but no seconds as the interval size increases.

Moved the interval size label in time range graphs, to the X-axis label instead of on the graph title.

Several new job groupings graphs are now available for summarized collections from the collection overview time signature graph.

Produced a new JW summary program. The intent here is to provide a faster mechanism to produce the summary and wait bucket gap data.   Use the Summarize option under a JW collection to use this option.   The primary files are single member sql tables with the names QAIDRJWSUM_<member name> and QAIDRJWGAP_<member name>. These files are currently available under the server side output files folder.

Added a prompt to JW V5R4 when running a summary that will give the user the option to delete the old version summary files with the default being set to "Yes".

From a rankings charts in summarized collections, if multiple bars are selected the user will be given additional drill down options. Selected Threads (separate) - produces a different graph for each thread selected Selected Threads (combine) - produces a single chart combining the values of the selected threads together Another menu will also only be shown for wait bucket rankings charts (currently doesn't work) that will show a comparison of a single wait bucket over time for the selected objects: Selected Threads compare Dispatched CPU Note that Threads, could be any of the job grouping types: Pool, Job name, etc.

The by interval graphs in Job Watcher for unsummarized collections are now by time interval graphs which support the various time range graphing options. This effects V5R3 and higher. Also made the set of wait graphs -> by time interval shown nearly the same as those shown for summarized collections for consistency.

Updated report folder structures for Job Watcher to hopefully simplify the experience and make it easier to get at the reports of interest

Some JW drill down options have changed such that the default /recommended choice will be a level higher with sub level menus beneath them in the same list. Previously it was not possible with the way I was building the menus from the database to have graphs and subfolders at the same level but it's possible now. By making this change the recommended choices will be easier to get at.

Enabled selection of time ranges and elements (such as a thread) in bar graphs. This has allowed for simplification of drill down choices, by putting time range selection in the graph itself rather than in the popup menu options. To use this feature, hold down the shift key and click a bar on a vertical bar graph in JW that indicates the start of the time range. Then click another bar that indicates the end of the desired time range. The start/end of your selection can be on different pages of the graph. To cancel your selection, click anywhere on the graph except for on a bar or line. When working with rankings charts, you can use the ctrl key to make individual selections.

Updated JW graphs to utilize the bar graph selection option. If a selection has been enabled then the default menu when you right click will be over the time range of the selection, rather than the single interval you clicked on.

Horizontal bar graphs (rankings charts) now support bar selection. Use SHIFT to select a range of entries, and use CTRL to select/deselect a single bar.

Added an option from the Holder menu called Holder chase which produces a report of up to 5 levels deep for the current interval showing the holder of the holder of the holder, etc along with other useful information like the wait object and wait information.

Added a button on the toolbar in the Data Viewer to redisplay the current graph with variable width bars such that wider bars represent longer intervals. This option is only available for stacked vertical bar graphs in some Job Watcher and Disk Watcher graphs. When displaying a graph in this manner the x-axis will change to date/time scaling.

Added several summary JVM graphs. These graphs can be drilled down to a rankings graph showing jvm jobs for an interval. From there a specific JVM job can be picked and statistics graphed with the same interval options available from the summary graphs.

On the interval details pages, on the SQL tab, added a button to launch the SQL statement collected by Job Watcher into the Run SQL Scripts window. Note that the system name passed to run SQL scripts is not the system the collection currently is located on, but is the system name where the collection was created.

On the interval details pages, the SQL tab should no longer give SQL errors when pressing the find next/find previous buttons if the collection doesn't have a QAPYJWSQL member.

On the interval detail property pages, all 3 priority values are now listed. Current XPF, Current LIC and Original LIC

Fixed a bug in interval detail property pages where the subsystem and job function were not being displayed when they should have been.

On the interval detail property pages, if a thread has ended during an interval as indicated by the QAPYJWSTS file, the current state will be listed as "terminated".

On the interval detail property pages, if looking at an interval for a job where CPU was not used, many fields shown are now retrieved from the last active CPU record (like current user profile, pool number, priority, etc). Previously these fields would be blank.

On the interval details pages, the JVM tab should be working correctly again. Previously was giving an SQL error. Also in prior versions of the client, the initial heap size and maximum heap size values were incorrect because the data in the files is stored in KB and not bytes as previously assumed. This has been corrected.

Add key fields like LDIO, IFS, and Transaction counters from the QAPYJWPRC file to the search function output. This affects the job/task name, subsystem, pool, and current user profile and taskcount search option

Removed the JW graph "Job wait time signature graph with sliced intervals" since support for variable width bars has been added which makes this option obsolete.

Changed the graphs for a single object over time, time range graphs so that the user can effectively zoom in/out using the time range graphing option

Removed all normalized graphs defined in Job Watcher. Now any of the time-based graphs in Job Watcher can be normalized by the user when desired by clicking the button in the toolbar of the Data Viewer.

The alternate views option is now available for 530 or higher Job Watcher I/O graphs for a single job.

Removed the JW graph "Job wait interval pct signature" since support for normalizable bar heights has been added which makes this option obsolete.

Added a graph called Current wait duration time signature under the wait graphs -> by interval folder. This graph displays only the "bad" current waits that were detected while snapshots were being taken. This has the effect of showing where these 'bad' waits occurred even if they are relatively small. The wait buckets that are included in this graph are 5, 7, 11, 13, 14, 15, 16, 17, 19, 32 for 530/540.

Added a graph for a single job called "Job wait time signature with current wait contributions" to display the job wait time signature graph such that each interval that used CPU contains 2 components, the wait bucket times before the current wait occurred, and the wait time for the current wait. Intervals that did not use CPU just have 1 component, which is the idle wait time. This graph should be viewed with the variable width bar option turned on. Using the normalize option as well will give the bars relatively fixed heights

Updated the parameters shown on the transfer collection window when using the testcase server ftp option so they match the current ones. Added a field for the PMR number which is used as part of the file name. Also updated the server command QIDRWCH/FTPJWCOL so it handles file names of up to 20 characters now instead of 10. This change requires both server and client builds at the latest levels in order to work.

Renamed the Job Watcher menu, "Selected holding job/thread" back to "Holder" and moved its location back to the 1st spot in the popup menu. Also added the "display call stack" option under the Holder menu again.

If a JW collection is running, and the refresh button is clicked on the LPAR CPU tab in the collection properties, the intervals collected field will now be correctly updated in addition to the list.

The Start Job Watcher Wizard has been renamed to Start Job Watcher Collection Wizard.

In the Start Job Watcher Collection Wizard, I've replaced the ending options screen with the 3 termination options found on the WCHJOB command. As part of this change, the max data value on the startup options panel has been effectively moved to the termination options panel. The default termination settings are 1 GB max size, or 100 intervals (whichever comes first).

In the Start Job Watcher Collection Wizard, renamed the startup options panel to basic options. Also added the ability to automatically generate the collection name when using this interface from the current date/time. The library will also default to QJWDATA instead of blank.

Fixed a divide by zero error that was sometime causing data format errors in the QZDASOINIT job log for the wait graphs by time interval, Collection overview time signature

Fixed a problem where the scheduled job's file was not reusing deleted records. Also fixed a bug where the file was not having its records removed after a job that was scheduled executes. Also added code to delete the record when scheduled jobs are deleted from the GUI. Reinstalling the latest QIDRGUI library is required for this fix.

Added a new report called "Longest current waits for <selected BUCKET Description> for interval X" accessible from the Detail reports menu from one of the "wait graphs by interval' Also added another report here called "Longest "bad" current waits" which shows the current waits found for an interval using the same
buckets shown on the Current wait duration time signature graph.

The "Job signatures ranked by XYZ" graphs have been renamed to "Thread signatures ranked by XYZ" to make room for the new Job signatures ranked by graph which is a grouping by initial taskcount (all threads) instead of taskcount   This also effects the "Job wait time signature for XYZ" graph which has been renamed to "Thread wait time signature for XYZ" graph

The wait graphs -> by job folder has been renamed to wait graphs -> by thread because of the creation of the new job rankings graphs

Multi selection from table views then drilling down can often behave badly because it will produce 1 graph per selection. Look into changing these so the selections are grouped together in ways that make sense to produce a single graph. This could involve include each taskcount in the resulting query, or using the time fields available in the selected records.

If a drill down action from rankings chart, or a table will result in more than 5 graphs being produced, added a warning message to make sure the user really wanted to do this

Changed the ship default preferences for Job Watcher summary graphs showing threads. Previously the default was only to show interval 1 to 10 or the last 10 intervals. I've changed this to 1 to 999999 and last 999999 intervals since most people look at the entire collection anyway

Added the I/O pending page faults per second field to the Thread rankings - > Page faults graph. Also added this field to the collection level I/O graphs by thread, page faults graph.

In the Other graphs -> State transitions graph, the field active to inelig should be yellow (255, 255, 0) to match the other graphs of this type

The unsummarized collection overview time signature graph will now exclude Job name = SMLOGICALDASDASM records in the TDE file

Updated the iDoctor GUI to compile using the new compiler Visual Studio 2005.

From the Run SQL View, added a menu "Launch SQL in Run SQL Scripts..." that will take the current SQL statement, place it in a temporary file and launch it in iSeries Access for Windows Run SQL Scripts for the purpose of either running it within that interface or using visual explain.   Note that iDoctor uses System naming convention (library/file) and iSeries Navigator uses SQL naming convention (library.file). You will have to change the default connection properties (Connection->JDBC Setup... Format tab, Naming Convention) in order for the queries created within iDoctor to be executable.

Added a button on the toolbar in the Data Viewer to normalize time interval based graphs. This will divide each value of the graph by the interval duration, which will make bars more consistent in height. This button is available on several Job Watcher graphs.

Added 2 new Preferences to the Graph views section of the Display tab.   "Use normalize option" -when checked will cause the normalize option to be used when a graph is first opened (if it's available.) "Use variable width bar option" - when checked will cause the variable width bar option to be used when a graph is first opened (if it's available.)

Property sheets that are not modal (meaning you can go back to other screens while they are open) are now displayed inside of a view.   This provides the following benefits/changes that users have been asking for: 1) These property sheets are no longer "always on top". You can click another view and it will move in front of the property sheet view. 2) Property sheet views can be resized to any size. 3) "Drill-down" actions from a property sheet view will now open on top of the property sheet view, like any other drill down in the Data Viewer.

Added a preferences on the misc tab called "Show non-modal property pages in a view". If you don't like the new property sheet views, you can uncheck this and go back to the way it was before.

Property sheet views that don't have a need for OK, Cancel buttons will no longer show these buttons (for the purpose of saving screen space). The user can press escape or the X to close the window just like any other view

Added a menu option called "Filter libraries..." from the component icon to filter the list of libraries shown. This setting applies to all components that display collection libraries. This library filter is shown next to the component's name in the tree when active. This can be helpful if there are many libraries on the system or if the connection in use is slow.

Added a scheduling preferences page that allows the user to change the default scheduled date and time, from 1 week in the future, to whatever the user prefers (in +days/hours from current date/time).

When opening a report from the list portion of a component tree/list window using the popup menu, added an option called "Edit" below the normal "Open table" or "Open graph" menu. This option allows a user to quickly open a user-defined or iDoctor-supplied table or graph such that the SQL editor is initially open without actually attempting to run the query. If the query will take a long time run, this can be used to open and modify the SQL statement first. This technique can also be used in graphs to modify the graph settings before running the query behind the graph

Added a new view, which processes database SQL statements in a separate thread. These requests can be queued up and are processed one at a time. This view is called the Remote SQL Statement Status View and is currently used by the V5R4/V6R1 JW summary program as well as the "Execute in batch" option from the SQL editor

The popup menu in the SQL editor has a new option, Execute in batch. This will run the SQL statement in a separate thread and show the status of execution in the remote SQL statement status view. If an SQL select statement is provided, you will be optionally prompted for an SQL table/library name in order to create a table using the SQL statement given

Added new icons to the remote command status view such that the icons change based on the status of the request (waiting to run, completed successfully, error). These same icons are used on the remote SQL statement status view.

iDoctor can now handle SQL statements that contain SQL table names

Opening a graph/table in the Data Viewer should no longer tie up the GUI and you should be able to cancel this action by closing the graph/table view being opened if desired. While a report is being opened you should be able to do other things in the GUI like view already open reports/graphs or go back to the main window and browse other collections. If you are currently opening a report, you can go to another view and drill down to another view while the 1st one is running if you wish. A maximum of 3 data viewer connections are used to accomplish this which means you can have a maximum of 3 queries running simultaneously from the same iDoctor client. If you exceed this the GUI thread may appear to hang and execution will no longer occur simultaneously for these extra requests. If you try to scroll a graph/table while another one is being loaded, you may experience delays in the GUI thread (if the same connection was used for the report being scrolled and the new one being opened) but it should work once the previous query finishes. As part of this change I've removed the preference "Show query cancel window when filtering or resorting a report" on the Misc tab as this option is no longer needed. The multithreaded behavior should now act the same whenever opening/requerying any graph/table in the Data Viewer.

Show the report/graph title on the view being opened before the SQL statement finishes.

Add/Remove programs in the Windows control panel now properly recognize the iDoctor GUI. You can use this option to uninstall the client.

Made a change how data viewer connections are handled, so that instead of only using a maximum of 3 connections if 3 simultaneous queries are executed, 5 connections are available to the Data Viewer's views where each view that is created uses a different connection by cycling through them. By doing this it will allow the user to open view A and view B, and scroll through view A while view B is running a query. Previously view A and view B would likely be using the same connection which meant trying to scroll view A while view B is still running a query would hang the client until view B finished it's query! The maximum value of 5 is a registry setting if you wish to change this:

[HKEY_CLASSES_ROOT\IBM.AS400.Network\3RD PARTY EXTENSIONS\IBM.PEX\DataViewer]
"QZDASOINITConnections"="5"

The toolbars in the main window and in the Data Viewer are now dockable (can move them to wherever you want

When refreshing a table, the hourglass on the cursor is often lost. To resolve this, show a message on the table when SQL statements are running in the background similar to what is shown for graphs "The query behind this report is running, please wait..."

Since graph tool tips can sometimes get in the way of what you are doing, I've added the ability to turn them on/off by pressing CTRL+T while in a graph. The status bar in the data viewer now includes a message after the graph control memory used message that indicates if tool tips are currently enabled or disabled.

Fix a bug on the query definition -> member selection page where users could not change the member their query is running against.

Fixed a problem with graphs having some portion of their bars missing every 2000 or so intervals of continuous scrolling if the number of bars per page is set to > 100.

If an error is detected in a graph definition, such as a required field does not exist in the SQL statement, the error message will be drawn right onto the graph view now instead of shown in a popup window.

Fixed some problems in the manage query definitions interface such as getting SQL errors when copying many query definitions at once.

Fixed a bug in the query parsing that was causing too many CHKOBJS to be performed when checking for the existence of files. This would slow down the opening of reports and on slower connections could greatly decrease response times.

The hide/show borders option in the graph legend wasn't working if the graph definition was defined such that the field had the same border color as the fill color. If this is the case the color will switch to black and then back to the original color.

If an SQL error occurs when launching any graph, the sql editor/error window should be visible above the graph right away rather than requiring the user to open the editor and reexecute the query.

The option to save a graph as a JPEG (Save As... menu when right-clicking on a graph) has been removed from the GUI. It didn't work since the June builds when the graph control was recompiled to work with the latest MS compiler and I can't easily fix it. The preferred method to achieve this is use Alt-Print Screen which copies the current window to the clipboard. Then you can paste this image into Windows Paint or something similar and save the image to the desired format. Also removed the Copy menu (which copied the graph view only to the clipboard) because you can just use Alt-Print Screen to do this which will include the legend and the copy function did not

Add check in the install upfront for system value QALWOBJRST.

The Graph views section on the Preferences -> Clipboard Tab has been removed since the copy to clipboard of the graph image option was removed from the GUI in favor of telling users to use Alt-Print Screen instead

May 11, 2007 – C00542

Flyover text for long text strings in table views (like SQL statements) should no longer be cutting off text on the right side of the window.

Fix a bug with Job Watcher monitors involving file QAIDRJWRI which caused the monitor to fail to start.

After first use, the show holder/show waiter button text was flipped around from what it should have been in the call stack window.

Corrected a situation that could happen with the thread list summary file QAIDRJWTL where it could become out of date, if ran for an active collections and the interval detail property pages would show missing information if the job was created after QAIDRJWTL was initially built

The interval detail property pages should refresh a bit faster now

Apr 18, 2007 – C00539

In interval details property pages: If showing the call stack for a waiter, and the holder record for the same interval doesn't exist because it wasn't collected, then clicking the show holder button will show the holder's name from the previous screen instead of blank.

Fix to the summary program CRTWCHSUM: it can now handle collections with more than 100,000 intervals.

The collection overview time signature graph incorrectly labeled the secondary Y-axis as 'Total active threads/tasks' when it should have been 'CPU Utilization'.

Fix a seemingly random SQL error that would show up on drill downs stating that the SQL statement is not a SELECT statement.

Apr 3, 2007 – C00535

Updated PTF checking for V5R3 and V5R4 Job Watcher to include the latest PTFs.

Because of the way some users have been taught to use Job Watcher I have moved the pie charts in the time range graphs to appear as the 4th chart per type instead of the 1st chart.

Produced a fix for Job Watcher monitors related to being unable to restart a monitor.

Fix bug in the client when starting or restarting job watcher monitors from the GUI where the job name is incorrectly set to QSTRJWMON when restarting a monitor and QRSTJWMON when starting.

The menu options to work with the holder's call stack and reports should now be working correctly.

On the Display tab of the Preferences window added an option "Display patterns". When checked the graphs will display line hatchings and patterns instead of solid colors.

Added options to select the desired fill patterns from the legend and graph definition primary Y-axis fields panel. These selections are not savable in a user-defined graph and are only used for the current graph.

The record quick view window no longer has a checkbox to "allow multiple rows" to be compared. This is in effect always on now. If the user selects multiple rows in a table (up to 30) and uses the record quick view menu, they can be viewed vertically.

The window that prompts a user for a member has been revamped. Instead of the user needing to specify the member for every single file found in the query, a list of all tables found in the query is built and then sent to a single dialog. This window allows the user to pick a member from a list and then indicate that it should be used for all files found in the query.

Query parsing should be a bit faster. Removed redundant checking for the existence of file/members in the query.

Fixed some flickering problems in the list of fields on the field selection panel of the query definition interface.

If the query behind the graph is changed/reran and it returns no data, the graph will now display the message "The query behind this graph returned no data" rather showing the graph data from the local cache.

If the graph definition window is resized, the coordinates are remembered for next use.

For readability purposes, the field column in the list for the graph definition -> primary Y-axis fields window now only displays the short field name instead of the long description and the field name.

In the open file window, pressing the enter key will no longer close the window, but will instead trigger the browse button to refresh the list of files

Added a menu option to the legend to hide/show a bar's black border. This can offer the effect of emphasizing a field in the graph

The legend can now be hidden/shown using the toolbar when the legend has focus. Previously this button only was enabled when the graph had focus.

The legend will no longer force lower case characters in field descriptions.

The creation date shown for user-defined queries and graphs should now show the creation time as well as the date

Fixed some multi-threaded timing issues that would cause canceling a long running query to crash the client. As a side effect of this the amount of time the user has to cancel the query is reduced, once the query has executed and it is to the point of building the columns and fetching rows; that part cannot be canceled.

Fields removed from the legend (or dropped into the 'available fields' section) will now retain their properties, such as the field description, and color. These fields are temporary and are listed first among the 'available fields'. If added back to the graph, the field description, color/pattern will become visible again instead of being randomly set.

Feb 12, 2007 – C00520

Added a new category of high level graphs: Wait graphs by time ranges. These graphs show wait bucket summaries by larger interval ranges allowing a user to see a higher level collection overview graph for the entire collection with less scrolling through it required. The interval sizes are fixed to either the entire collection (in a pie chart), or 1 minute, 10 minute, or 1 hour intervals. These graphs are only displayed if the collection has been summarized and are accessible by right clicking on a collection or from the wait graphs by time ranges folder under a collection.

Added job wait bucket drill down graphs from the time range graphs. These allow a user to see the jobs for the desired time range. By default the time range will be the same range as for the interval selected (either 1 min, 10 min, or 1 hr)

Changed the interval details quick view property page so that it no longer has the general graph information section above the listing of the selected point's fields. This was necessary to alleviate problems building that window for some types of graphs.

Trying to open a summary graph over a collection that is active and currently running a summary, will result in an annoying error message because of a lock. This fix requires an update to the server builds.

The check for incompatible files has been updated to handle a special case where the check would fail.

Updated the server program that performs a check for summary file record format compatibility to utilize a more generic approach for file format checking.

There was a rare situation when the JW monitor control file would contain multiple records for a single monitor. This caused a problem restarting a monitor. Eliminated this possibility.

Updated the install so all commands issued are library qualified with QSYS. Though unlikely this could cause a potential problem depending on what was in the library list.

Deleting a Job Watcher collection using the DLTJWCOL command will now delete the userspace associated with it as well.

Some of the I/O graphs by interval were not using the summary files like they should have been if the collection had been summarized.

The I/0 graphs by interval now share the same SQL statement (except for LDIOs). This means the alternate view option is now enabled for them.

In the I/O graphs by interval, the Page faults and Physical I/O activity graphs now have a 2nd line on the Y2 axis to show IO pending faults per second

For summarized collections added new graphs: IO/Transaction/State transition time range graphs and drill down options.

Some drill down menu options have been renamed and restructured to offer better organiziation and flow.

Added a call stack drill down report created by Larry if a single stack entry is selected: Total occurrences: by job. This shows the number of times the call stack entry was detected per job/thread in the collection.

Created an advanced folder under Wait graphs -> by interval named normalized times. This contains the collection overview normalized time signature graphs. Also added a new graph to this folder that shows the (estimated) break down of disk page faults, into IOP times and disk page fault times.

In the JW collection wizard, the last executed sql statement options on the sql panel has been reworded to indicate either the 'active or last executed sql statement' applies.

In JW renamed the preference, "display advanced wait graph options" to "Display advanced reporting options (when available)" since it no longer applies just to wait graphs.

Renamed the wait buckets 'others' menu options that referred to wait buckets normally not used for drill down (such as the wait bucket 'idle/waiting for work") to "advanced".

Made some pretty big changes to how the summary report folders and drill down menu options are built. The majority of these folder names, descriptions and drill down names are now controlled through the local queries database iDocQueries.mdb. Some of the names, and hierarchy of reports have been moved. Also some reports are considered advanced and only show up if the Job Watcher advanced reporting option has been checked under the Preferences window. This will allow a user to remove some level of complexity in the number of reports by unchecking this option

The Data Viewer now has a new button on the toolbar which provides alternate views for the graph with current focus. Look for the button that looks like a line graph next to the legend on/off button. Clicking this brings up a list of any available alternate views for the current graph. This mostly applies to Job Watcher graphs.

For all components/releases, when throwing an error message after a remote program call in iDoctor, one of the CPF messages was not being returned in the error window.

For all releases, in the client portion of the install program, updated a few remote commands that were issued that did not qualify the library as QSYS, such as the ALCOBJ command. Though unlikely this could cause a potential problem depending on what was in the library list.

Added a preference on the display panel of the preferences interface for the default legend size (as a percentage of the graph window). This preference allows you to show more or less of the legend (from 0 to 50%) every time a graph is opened.

Removed the "reuse these settings" checkbox from the field selection page of the query definition interface. This function has become obsolete since users can now create their own user-defined queries.

When showing a graph of a query consisting of only a single record, the graph would often not be displayed until the screen was refreshed or clicked on by the user.

Fix a problem with calculating record counts on some types of advanced queries.

Fix a SQL parsing error that would happen if a select statement started with an expression having more than 1 set of parentheses. No field name would be displayed in the column heading.

When writing your own queries and defining expressions in the field list of the select statement, you can now leave off the AS <field> after the expression. The GUI will display nnnnn as the short name and the expression as the field description. This fixes the inability to define expressions without an AS clause in iDoctor.

When writing your own queries in iDoctor you can now define field descriptions right into the SQL statement (surrounded by double quotes) and have them be displayed as multiline headings in the table view. Example: FIELD AS "My description"

Updated the query definition's parsing code to ensure that only the outer most where, order by, group by and having clauses are shown in the query definition interface. Previously the parsing would become messed up if inner select statement's where, group by, or order by clauses were shown in the query definition interface and then modified. Modifying the inner selects requires the user to change the sql statement manually (not supported from the graphical interface).

Fix a bug on the group by query definition panel. If fields are built from expressions and used in the group by and then later updated through the interface, the expression's alias name would be passed into the sql statement instead of the group by's expression itself which is incorrect SQL syntax.

The query definition parsing code should now be able to handle expressions in a where clause that have been defined with redundant sets of parentheses.

When reconstructing a query after working with the record selection window, an extra space added between the from clause and the where clause has been removed.

Support added to create user-defined pie charts. The Pie chart is now a selectable option on the general tab of the graph definition interface

Made a change in the query definition parsing such that if the parsing cannot be done successfully, the query definition interface menu options are greyed out. This interface is not intended to be used for all possible types of SQL syntax. This should help prevent users from trying to use the interface after a complex query has been ran and encoutering problems because the original SQL statement was not parsed correctly.

Add the create alternate view menu that is shown on all graph view popup menus to all table view popup menus. This will allow a user to either make a copy of their current table view in the same data viewer (perhaps for comparison purposes), or to take their current table view and begin working with the raw sql statement more easily using the table view with SQL editor.

Flyover field descriptions defined in user-defined graphs within the flyover tab of the graph definition interface are now saved with the graph definition.

When defining flyover fields in the graph definition interface, added a check to ensure no more than 10 can be added to the list of flyover fields.

Nov 6, 2006 – C00508

Some of the job detail reports can become stale (missing newly created jobs) when analyzing a collection in the GUI that is still running. A file is being built under the covers named QAIDRJWTL by the client the first time the collection is analyzed by the GUI. This file contains 1 record for every job that exists in the collection.   The fix included in this client is to rebuild the file every time a report is launched from the report folders underneath the collection if the collection is still running.

The call stack report options for a specific job are no longer displayed if no taskcount exists in the preceding report (typically a call stack search output, or by looking at the QAPYJWPROC file). Trying to run these reports when no taskcount was available would result in an SQL error.

The SAVJWCOL and FTPJWCOL commands now have parameters to specify the target release and data compression parameters to apply to the save file. These parameters will allow a user to transfer a JW collection to a system of a different release or to compress their data before sending it over FTP. The transfer window in the GUI has also been updated to provide these two new parameters.

Fix a rare but serious SQL problem in the collection overview time signature graphs in V5R3 and V5R4 Job Watcher. If the CPU utilization % was > 1000% the query would fail and the graph would not be produced. This happened due to a large amount of transferred CPU being reported in a CPU intensive collection.

The FTP collection transfer functions in JW and PA should now work against a unix FTP server. The FTP commands like 'put' were being issued in uppercase and this did not work on that type of FTP server.

Updated the PTF list for Job Watcher at release 530, 535, and 540.

When restarting a JW monitor from the GUI, the user will now be prompted with the Start JW Monitor dialog, allowing them to change some of their parameters before restarting the monitor. This also gives the user the option to setup automatic holding/releasing of their monitor on the restart request. If the user wishes to change their definition they may do so using the View button and then editing the command string that makes up the definition before restarting the monitor.

Fix a tab order problem in the transfer collection window.

Oct 5, 2006 – C00504

When starting a Job Watcher monitor, specifying the maximum collection value of 99 would cause the monitor to not run properly.

Sept 19, 2006 – C00502

Added a menu "Alternative views" in the legend which if available provides the option to display other graphs available using the current query. The graphs shown all utilize an identical SQL statement which makes this feature possible.

Trailing spaces in the Add jobs window in the Job Watcher collection window, causes the selection to be incorrectly stated as being invalid.

The access code expiration date will be shown on the components dialog when connecting to a system if it is known.   An expiration date can only be determined and displayed if server builds dated after Sept 19th, 2006 are installed. For a green screen solution to determine

expiration dates, call program QYPBASE/QIDRCHKEXP (for releases V5R1/V5R2) or QIDRGUI/QIDRCHKEXP (V5R3+) The expiration dates will be returned in your job log if any access codes are installed on the system.

When checking access codes in the client, added more informational messages on the reason the access was denied including: access code not found, access code is invalid (OS release mismatch?) and access code has expired.

Fix a bug where monitors could not be scheduled; there was an extra set of single quotes added to the command string on the OVRLAP and TEXT parameters.

Sept 1, 2006 – C00500

Updated the icons for several of the iDoctor components. Also made it so reports /graphs opened in the Data Viewer will match the icon for the component being opened.   Added icons to the iDoctor components window and made it wider to make it easier to see any status error messages.

In an instance where all objects but one were deleted from library QIDRWCH on a system, corrected an incorrect message saying the required access code was not found.

August 4, 2006 – C00497

On some systems users complained of having scheduled jobs failing to start. The command string was occasionally getting extra random characters added to it, causing it to fail.

By customer request, add QAPGMDESC file to the list of server side output files shown in a collection in both Job Watcher and PEX Analyzer at release V5R3 and higher.

The journal serialization bucket had a misspelling in the word serialization.

Due to a query parsing error, the object waited on graph for a selected job in fails with an SQL error.

June 19, 2006 – C00496

The flyovers for long cells in table views should look a little cleaner now. It won't wrap the flyover text unless the cell length is at least 100 bytes.

The GUI will now set the QPTFOVR data area to '0' if the user answers in the GUI to NOT override the PTF checking. Previously the GUI would set the data area to a '1' the first time the user overrides PTF checking and did not provide a way to set it back to a '0'. QPTFOVR will also be set back to 0 by the GUI if the GUI detects the PTFs are installed so the user does not get a warning in the collection job log.

When submitting scheduled jobs in PEX Analyzer or Job Watcher they will now be submitted to run on the default PEX Analyzer or Job Watcher job queue. In PEX Analyzer this is QCTL in QSYS and in Job Watcher this is configured at install time and stored in data area QIDRJOBQ in the job watcher library.

Fix some problems with creating rule definitions from the GUI.

Fix a bug with the Start JW Monitor interface: it would not honor the hold/release/end scheduling options if the job watch monitor itself was scheduled.

Added a Java virtual machine interval details panel for a single job/interval. This panel is only displayed if the job/thread has a JVM or is a secondary thread of a JVM. Some of the fields shown on the panel are Java thread name, last waited on object name, GC heap size, and the GC reference counters. This panel is only shown for V5R3/V5R4 Job Watcher.

In the start collection wizard, on the call stack tab, changed the last checkbox label to indicate that the bad waits collected are those without a holder condition.

Fix an SQL syntax error with the Physical I/O activity rates by job graph.

May 8, 2006 – C00492

Updated the interface used for copying/transferring a collection so that it has a consistent look for all components that copy or transfer collections.

The menu User-defined reports ->Select local database.... has been renamed to User-defined reports->Set local database...

Added a new menu option under the user-defined reports menu for all components that allows a user to manage their query or graph definitions. The menu is User-defined reports ->Query definitions->Manage... OR User-defined reports->Graph definitions->Manage...

Added a new manage definitions interface, that lets a user work with their current local query/graph definitions database. This interface allows a user delete or to make copies of their definitions and set them to another release or component. The interface also allows a user to export definitions to the server. These definitions exported to the server can then be imported by other users, by using the menu User-defined reports->Query definitions->Import from server... OR User-defined reports->Graph definitions->Import from server...

When opening a data viewer, the component the data viewer was opened from will be remembered. When opening a table from the open file window or creating an SQL editor/table split view, and when a query is defined and saved, the default component selected on the save query or save graph window will match the component the data viewer was created from. The default version will be the version of the system, rather than blank if it is not known

The SQL statement within the Query tab in a report's properties is now editable. After making the change the user must refresh the report manually (F5) to pick up the changes in the SQL statement. This applies to both table views and graph views.

In the record selection page of the query definition interface, if a field is defined as an expression but is also a regular field name in the file, the user should now be able to filter on this field when the field was not listed previously in the list of fields.

Fix a bug with the legend where fields were not dragable if certain small fonts were used.

Add an interface to work with iDoctor scheduled jobs (for PEX or JW). This view is accessible from the View ->Scheduled Jobs menu or the Work with Scheduled Jobs menu found when right-clicking the Job Watcher icon. This interface supports the ability to submit the job immediately, hold/release or delete the scheduled job.

Within Job Watcher or PEX Analyzer add menu options to the View menu of the main window. These menus allow the user to views applicable to the current system/component view being looked at. Some of the options are Definitions, Monitors and Scheduled Jobs.

Fix a build problem with the Java PTF group checking program. It was requiring ALLOBJ authority in order to produce a correct check.

Add the ability to schedule a job watcher monitor.

Add ability to schedule a date/time for ending, holding and releasing job watcher monitors in both the GUI and the green screen STRJWMON command. The green screen RSTJWMON (restart) command also allows the user to schedule the end, hold, release dates but not via the GUI. Updated command help text for STRJWMON, RSTJWMON. The hold/release options can be used to hold on a monitor on a Friday and startup again on a Monday.

The STRJWMON and RSTJWMON commands now have a check to ensure that the subsystem specified in data area QIDRJOBQ is active before starting a monitor.

When a monitor is started from the GUI, added a check in the GUI to ensure that the definition is defined so the collections will run for at least as long as the maximum collection duration parameter on the monitor. If not a warning is shown informing the user that gaps will occur in the monitor because each collection will end before the next is started. The client will advise the user to modify their definition to ensure that there are no gaps.

If a monitor fails to start this will now be properly flagged in the GUI as an error with the monitor. Typically if a monitor fails to start this means that the jobd or jobq parameters on the definition are incorrect. (if definitions are very old this could happen). Also added a menu option on the monitors view to display the job log for the monitor job if it is available on the system.

Fix some problems with JW monitors that could occur if the parameters specified on the definition for the job description or job queue are no longer valid. If this happens the monitor job will end rather than experiencing an active monitor job with no collector jobs running.

When defining a rule definition that contains a function (Avg/Rate/Pct) the function is not saved/restored from/to the rule definition properly.

Fix a bug with the call stack reports where it was possible to see duplicate records in the "Jobs with this occurrence" reports.

Added a preference on the Job Watcher tab, that controls whether or not additional wait graphs by interval and wait graphs by job options should be shown. These options allow the user to display wait graphs in the desired way depending on the files available. Wait graphs by interval has the following options   Pre-summarized = Displays the QAIDRJWIS2 file contents in the graph if the collection has been summarized. Actives only - Summarizes the collection by interval using file QAPYJWTDE Actives and idles (For V5R4 only)- Summaries the collection by interval using files QAPYJWTDE and QAPYJWSTS (includes idle waits).   Wait graphs by job has the following options: Actives only - Summarizes the collection by job using file QAPYJWTDE. The idle time of jobs is not shown on the graph but this graph will run faster (sometimes much faster) than the actives and idles option. Actives and idles (If collection summarized or V5R4). If collection summarized this option will use the IG2 file , otherwise if it's V5R4 , then this graph is produced using the QAPYJWTDE and QAPYJWSTS files.

All wait graphs by interval will now show CPU utilization as the Y2 line and total active and total idle threads as flyover fields. Also added the Collection overview waits per interval sec time signature and Collection overview with dispatch CPU time signature graphs to all views available under the Wait graphs by interval folder.

Add a new "Advanced" page to the collection options window of the Job Watcher collection wizard. This page allows the user to specify the following parameters on the WCHJOB command: ASP threshold override, DB2 availability (*ITVEND/*END), Min CPU stats gather time, Use WCHJOB block mappings (*YES/*NO). Also added these parameters to those saved/restored when creating/restoring a definition.

Add an interface to list , delete , edit job watcher definitions at releases V5R3 and higher. Use the Work with Definitions menu from the Job Watcher icon, or the View -> Definitions menu from the Job Watcher component view.

Increase the JW V5R3M5 PTF checking to include the latest MF PTF, MF37676.

Fix a build problem with the Java PTF group checking program at releases V5R2 and up. It was requiring ALLOBJ authority in order to produce a correct check.

Allow multiple job watches to be collected within the same library at release V5R3M5.

Add a command to delete a job watcher definition, RMVJWDFN.

Allow the wait bucket property page for a JW collection to be shown even if the collection is still running.

If a job or task was idle throughout the entire collection and the job or task name is not known, display "task count nnn" as the label instead of / / : 0000000. Also display "task count nnn" in the job information field of the interval detail property pages for these idle jobs or tasks.

February 28, 2006 – C00479

In the Collection Wizard, add a button to save the current definition (JW command string) on the final page of the Wizard.

Add a position indicator in the toolbar of the data viewer. This allows a user to enter the desired record/bar position and press enter or press the button next to the edit box to have the graph or table scrolled to the desired point.

Add popup menu options to build table views from any graph either with the SQL edit split view or without. The menu options are Create alternate view -> Table Create alternate view -> Table with SQL editor

When first opening a table view the widths of the columns should be resized better than before. Sometimes field data longer than the header was not completely visible.

When resizing columns in a table view, they should now be remembered again after performing a refresh.

Rename adjust scale popup menu for graphs to adjust primary scale. Add adjust secondary scale menu option for graphs. This will work just like the adjust scale menu for the primary axis but instead works on the secondary axis. This allows a quick way for user to see above 100% utilization in the CPU utilization graphs (as an example).

February 14, 2006 – C00476

Job Watcher V5R4 released.