

Update History for IBM iDoctor for IBM i – Job Watcher V5R3

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 idocor 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 definition->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 Synchronized 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 than 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 synchronized 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 implementation 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 Miscellaneous 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 multipath 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 also 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 re-retrieved 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 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 restructured. 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 collections.

---> 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 repositioned 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

None

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 erroneously 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 component 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 available (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 initialization 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 occasionally 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 similar 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 significantly 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 occasionally 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 component 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 QAPYDWHDR 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 QAPYDWHDR... it should be QAPYDWHDR... 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 synchronized 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.

Collection Services Investigator changes

No changes made

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 QAPYDWDWR 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.

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.

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

In V5R3 JW, for the graphs for a single job/thread over time, drilling up for a single bar would only include the 1st interval of the bar instead of the entire time range if the job was idle in the 1st interval shown on the graph.

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

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.

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.

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

In the Split Collection window added a checkbox named "Include summary files" to include/exclude summary files from the collections created by this option. For this change to work the latest JW server builds are needed.

For summarized collections, the graph "Thread signatures ranked by XYZ" did not allow you to filter the data by Current user profile (field CURRUP) using the Qry Dfn interface since that field didn't exist in this query. I've added it to the query, which allows you to add a filter using the query definition -> record selection menu and filling in CURRUP as the field with the desired value. For all other releases and V5R3 unsummarized collections, this capability is already supported.

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/O 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 organization 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 encountering 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 OVLAP 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.

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 draggable 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.

Fix a build problem with the Java PTF group checking program. It was requiring ALLOBJ authority in order to produce a correct check.

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.

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) - Summarizes 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

Add V5R3 JW PTF MF38866 to the required PTF checking

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

Renamed some of the call stack reports (menu names) to try to alleviate some confusion about the scope of the data shown in the reports.

Added a message to the STRJWMON help text indicating that the job should be submitted to batch.

When using the find next/find prev buttons in the JW interval details call stack or SQL page, the search will take place over the call stack or SQL file instead of the TDE file. This will allow the user to find the next/previous collected call stack or SQL statement which does not always match whether or not a record was collected in the TDE file. Also fixed a small memory leak with the find next/prev buttons that occurred when the search returned no results.

Added several new drill down options on the job watcher interval details page if a wait object is present for release V5R3 and up.

On the interval detail property pages on the holding thread page remove the holding thread TDE ID field. This field does not exist at release V5R3 and higher.

On the interval details property pages rename the field Holding thread or task to Holding job or task since the thread ID for the holder is not known at this release. Also rename the Title of the holding thread/task page to holding job/task

Fixed an issue with the Synchronous read and write response time graphs by job. They were incorrectly showing Occurrences (counts) instead of response time in the primary Y-axis.

When starting or restarting a Job Watcher monitor, the definition (command string used to create the collection) will now be saved for each collection created within file QAIDRJWDFN in the user's collection library. This will allow the creation settings property page for the collection to be visible for collections created by JW monitors.

In the summarized collection overview graph's query, add parens around the fields in the where clause that are OR'd together. This will fix a problem users were having when adding filters to the end of that query and getting unexpected results because of the OR in the preexisting query.

Modified the transfer collection window in PEX Analyzer and Job Watcher to contain options to transfer the collection to an FTP site or to the boulder Testcase FTP site. Increased parameter limits to make this feasible for sending data to an FTP directory, not just a library on an iSeries. Also modified the FTPJWCOL and FTTPACOL command to handle these parameter changes.

In JW added drill down menu options from the collection overview graphs to drill into jobs by an interval range consisting of the current range viewed on the page, or by a user-specified range. This affects the Run/wait graphs by job and the IO graphs by job popup menus. Merged the CPU graphs by job menu options with the wait graph by job menu.

Added a properties menu when working with Job Watcher reports in a Data Viewer to view the collection's properties.

Renamed the "System-wide interval details" property pages in JW (which show a summary of data for a single interval across the collection) to "Interval summary"

Right-clicking on a line in a JW collection overview graph will no longer show machine level gate serialization graph by job as the 1st menu item of the popup menu.

The description for the buckets in the interval summary - wait buckets property page incorrectly shows the units as microseconds instead of seconds.

Long running queries can now be canceled in all components.

Reduced refresh/screen display delay when dropping a selection into a list after dragging it.

Double-clicking on a tooltip in a table view will now act the same as if the user double-clicked on the table itself.

Renamed the 'confirmations' preferences tab to 'miscellaneous'. Added the enable animations checkbox from the display tab to the miscellaneous tab. Added options to this tab to hide system names from window titles and to allow the user the option of always having the cancel query window shown even for queries that are estimated to run very quickly.

Fix an issue with the automatic scaling of the Y2 axis not working if auto scaling was turned off and then back on again.

In all components/releases enhance the graph legend so that available fields in the graph that are not yet shown in the graph are draggable into the desired parts of the graph legend (X-axis label, Y-axis, Y2-axis). Once this is done the graph will instantly update to show the changes. Fields can be dragged into, out of and within the various areas of the graph. You can reorder the fields in a stacked bar graph if desired. Colors can also be set for the bars and lines directly from the legend. The only restriction is PEX Analyzer graphs that have been flattened (multiple records that are used to construct a single bar) cannot have fields dragged into or out of the X axis and the primary Y axis.

Redesigned some parts of the graph definition interface specifically the x-axis page and the secondary Y-axis pages.

Add a flyover tab to the graph definition interface, where the user can define the fields to be shown in the graph tip window (in addition to the fields for the bars and lines in the graph). Up to 10 flyover fields can be defined in the graph. Also add a submenu under the graph definition popup menu entry to get into this page of the interface.

Double-clicking the legend in certain areas (x-axis, bar, lines, flyovers) will show the graph definition property pages and open to the corresponding page

Fixed a potential crash that could happen if all bar values on the graph were 0.

Made a large change to the user-defined query and graph support. User defined queries and graphs are now stored in a local database of the user's choice. This offers a number of improvements, most noticeable are the performance improvements, but this also provides the ability to write a query or graph over data on one system and apply it to any system without having to move files around from one server to the other. In order to avoid people having to redo all their query and graph definitions I've created an import interface that allows someone to import their definitions from either the server of their choice or someone else's iDoctor user-defined definitions database. The default name of the user defined definitions database is C:\program files\ibm\iDoctor for iSeries\UserQueries.mdb. Use the menu "User-defined reports" when right-clicking on the component icon for any component in order to access the import interface

The filter selection page of query definition needs to handle setting up where clauses for Hex defined fields that are not binary. Only apply Char(hex(field)) if the value provided is double the length of the field and the field is defined as Hex. An example of where this was a problem was the FULLADDR field in file SMTRMOD.

December 2, 2005 – C00463

Removed some background flickering (desktop icons) when refreshing table views and graph views.

Fix some problems with missing order bys in some Job Watcher graphs. In some cases this could cause the data to be out of normal sequence.

Added a command and menu option in the GUI to restart an existing Job Watcher monitor. The command is RSTJWMON. The definition used for the monitor must still exist in the definition file or this option will fail.

The Job watcher install will now end the Job Watcher subsystem specified. This is to prevent the install from failing if the subsystem is already active.

Add parameters on the application property server configuration page to set the collection job's run priority and the priority of the collection summary (CRTWCHSUM job). These priorities are handled through job description and class QIDRBCH and QIDRLOWBCH.

Added a check in the collection wizard for the number of job selection entries. This must be 100 entries or less.

When a summarization is currently running, show a percent complete indicator in the status field on the screen that displays the list of job watches in a library.

Fixed an issue where the wait bucket property page for a job watch was not always being displayed.

November 8, 2005 – C00459

Fix an SQL error with rate calculation in the Job Watcher reports if a jobs total run time was less than 1 second.

Add support for continuous collections via Job Watcher Monitors. New commands added: STRJWMON, DLTJWMON, ENDJWMON, ENDJWCOL, HLDJWMON, RLSJWMON. Add an interface in the GUI to work with Job Watcher monitors. This is accessible by right-clicking on the Job Watcher icon and choosing the Work with Monitors menu.

Add command ADDJWDFN. This command provides a simple way for a green screen user to create a Job Watcher definition using a reduced/simplified set of parameters from the WCHJOB command. The

parameters are: collection duration, interval duration, max collection size (MB), include call stacks, include sql statements, include comm data, force 1st interval. This command retrieves the job queue for the SBMJOB from the data area QIDRJOBQ in QIDRWCH. Also add an interface for this command from the Start JW Monitor window.

Add a button to view a definition on the Job Watch Wizard starting options page.

Add a menu "Filter..." found when right-clicking on the user-defined queries or graphs folder. This menu will bring up a window where the user can define if they want to filter the list of user-defined queries and graphs by user name or generic file name.

The default field ordering for the user-defined graphs folder does not match the user-defined queries folder like it should.

Fix an occasional client crash that could occur if opening a graph that contains no data. This also changes slightly the look of the graph view with no data. Now a text message will be drawn on the window saying that the graph contains no data.

Add checks to see if the user is not authorized to write to the registry when creating ODBC connections and setting preferences, and provide a message with information on how to fix the problem.

Add a button to remove system names from the list of systems shown in the installation. You can also remove systems by right-clicking on the system list and using the remove menu.

Add a property page for user defined queries and graphs that displays the SQL statement for the report.

October 24, 2005 – C00455

Updated the PTF requirements to check for MF37430 and Java group PTF level 5.

Some SQL host variables are unreadable in the GUI. Some are stored in ASCII, some are stored in EBCDIC, check the byte ranges to handle displaying both.

When displaying host variables within an SQL statement, an ORDER BY was not properly added to the query that builds the list of host variables, so they were out of sequence in some cases when placed back into the SQL statement.

The program that does PTF checking is now stored in QIDRGUI/QYPBASE instead of the component library. This requires the latest client and server code to be in sync.

When checking for required PTFs show all PTFs not on the system instead of just the first one.

If a field is a HEX field and in record selection of query definition the user tries to use that field, add CHAR(HEX()) automatically around the field in the where clause so that the expression will be evaluated correctly in SQL.

When opening a table in the Data Viewer, if there are key field(s) defined add them to the default order by of the SQL statement when the table is first opened. The sort by page of query definition will also revert back to this sort sequence if the 'default' button is pressed.

October 17, 2005 – C00454

Add iDoctor-supplied definition choices to the collection wizard on the Options page.

If the default summarization is currently running the Summarized field in the list of collections will show "In progress" to indicate this.

Fixed in the summary function the calculation of active and dispatched CPU time. The summarization will need to be rerun on existing collections to pick up this fix.

Add a warning in the V5R3 collection wizard if the user has selected last executed sql statements with all jobs, activation groups with all jobs, comm data with all jobs, or force 1st interval with all jobs. This can be a resource intensive operation.

Add Find Next/Find Previous/Refresh buttons on the interval details property pages to find the Next active record collected in the TDE or LG2 file. This does not apply to the job-specific property pages for V5R1/V5R2.

In Job Watcher, in the search window add an option to specify the data Viewer to open the search results in

Removed duplicate 'display call stack' menus that showed up on some popup menus.

Enhance the install program to allow multiple partitions and components to be installed in one batch process.

Reestablish the ODBC connection automatically if the connection is lost.

When retrieving the last saved position of a window if the coordinates are out of the display's visible range (due to changing monitors or resolution) center the window instead of using the saved position.

Migrate all iDoctor queries and graphs to the PC in a local database to improve performance when loading the query and graph definitions.

Add a graph view display preference to 'always show the graph legend'. If this is checked the graph legend is always shown and cannot be hidden without unchecking this box. The default is to have this option turned off. This change was added based on feedback from the last iDoctor class. This preference overrides the show legend checkbox within the graph definitions general page.

Modify all graph views so that they become a split view with the legend optionally shown in a fully scrollable view to the right of the graph. This should resolve issues with being unable to see the full legend in the graph views in some instances. Also add a hide menu when right clicking on the legend that will remove the legend.

Changed max Y2 axis value used in graphs of 100 to 101, so that the line is visible if the value is 100 across the entire graph.

If the minimum client or server level could not be found, ask the user if they wish to continue anyway.

Add a new graph definition submenu to set the bars per page shown in the graphs

Add a new graph definition submenu to set the maximum Y-axis scale shown in the graphs

Add an icon on the Data Viewer toolbar to easily hide/show the legend.

Change the way quick sorting a column works. Left click = sort by column ascending. Right click = sort by column descending. Shift + left click = add column ascending sort to end of existing sort (if any) Shift + right click = add column descending sort to end of existing sort (if any) If the column is already in the

existing sort and shift is held down it is modified based on whether the column was left or right clicked. Note that multi column sort is not available if the list was not created via a SQL query. This applies to the component views (lists of libraries and collections)

In the Data Viewer's File -> Save As menu, add a menu to save the query and graph definitions. Also add these menus to the Save menu on the toolbar

Add a new Confirmation Preferences Window tab containing two checkboxes. These control whether or not a confirmation window is shown when closing the iDoctor GUI or when closing Data Viewers.

Add a confirmation window when closing a Data Viewer or the iDoctor client. This window includes a checkbox saying 'do not show this window again' in case the user doesn't want the confirmation windows

On the find window the find's progress window had a tendency to disappear behind other windows, and pressing the cancel button on the find window would not cancel a find already in progress. Disable the find window while a find is in progress.

Disable from view the code that handles migrating old version of the query and graph definition files to the current version. This includes the 'migrate query and graph definition menu' found by right-clicking on the component icon in the tree.

August 18, 2005 – C00436

Move the PTF checking to just before starting a collection rather than when signing on to a system.

Update required PTFs to: MF36310 SI17681 JAVA GROUP 4 SF99269

Delete any existing V5R2/V5R1 iDoctor libraries on the server during installation.

August 9, 2005 – C00431

Add a new menu under the graph definition menu called "Adjust scale" containing several percentages which can be used to quickly adjust the maximum scale of the Y axis based on a percentage of the current maximum. This also must contain a reset option to go back to the default scale for the graph.

The option to disable the automatic scaling in graphs should now work for the secondary Y axis

When retrieving a spool file in the GUI (normally when looking at a job log), the option to cancel the retrieval should now be working properly and letting the user see partial results of what has been retrieved so far.

Add file QAIDRJWTL to the list of files processed within commands DLTJWCOL and CPYJWCOL.

Add a new CPU utilization graph added to the CPU graphs by job folder for collections that have been summarized.

The option to collect idle threads during the 1st interval is now disabled if a rule definition is defined in the collection wizard.

Fixed a bug in the client startup, where on certain systems it would incorrectly check for the V5R3M5 PTFs instead of the V5R3M0 PTFs.

August 1, 2005 – C00428

On the server signon screen of the install, check for the required authorities (*SECADM, *ALLOBJ) before moving on with the rest of the install.

Add a check at GUI startup to see if SP1-SP4 of Client Access V5R2 is installed on the PC. If it is prevent use of the GUI as a bug with the ODBC driver in these service packs causes missing bars in the graphs which can be very misleading.

Fix a bug with the Help->About menu prompting for a signon when it should not have been.

Customer requested that routing data in Job Watcher and Heap Analyzer be removed from the SBMJOB command issued by the GUI and instead have the routing data stored in the job description. As part of this change the job descriptions used by the GUI were renamed to QIDRBCH.

Changing the interval field manually in the interval summary or detail property pages will require the enter key to be pressed for the change to take effect. Pressing the up down arrows will behave the same as it does now (retrieve the next/previous interval without requiring enter key to be pressed).

Added a 'Display call stack' menu for Job Watcher table views that contain a taskcount and an interval. Also added this menu as a submenu under the holder to show the call stack for a holder. On the interval summary property pages (objects waited on and holder) these menus are also available to drill into the call stacks for the waiter or the holder.

Added drill down support from the interval detail property pages on either the job or the holder. Left click on the job info or holder name to access the menu providing the detail reports for either of these jobs. Note however in some cases the holder has not been collected by job watcher (if the holder did not use CPU) which would result in an empty bar or lack of call stack /job data.

Added drill down support from the interval summary property pages on any of the jobs listed.

The Cpu graphs by job -> Cpu queueing menu is not working from the collection overview graph.

The collection size sort order should be numeric when viewing the fields in a list of job watches.

If a rule is defined, a new panel will be shown to describe the rule settings in the property pages for the collection.

Added new V5R3 JW call stack reports: Total occurrences for this job/task for all intervals, Occurrences by interval for this job/task

Corrected a problem with job/thread selection by main storage pool within the WCHJOB command.

Modified the colors for bucket 14 (machine level gate serialization) and bucket 23 (socket accepts). New colors are yellow and light blue respectively.

In the collection wizard, the estimated collection end time was incorrect if the duration value provided is < 1 or contains a decimal point.

In the collection overview graph and contention overview graph, add bucket 14 to the graph but exclude any contributions from task SMLOGICALDASDADM. Accumulation of this bucket for all other jobs/tasks is generally considered bad and of interest. Also when drilling into the jobs within an interval exclude SMLOGICALDASDADM as well. Changes also have been made to the summary function to exclude these numbers from the total but already summarized data will show times from this job in

bucket 14 within their graphs. If drilling into this bucket you see a yellow bar with no name then it is probably this task contributing to bucket 14 and can be ignored.

Created a new file built by the GUI the 1st time its needed to contain a unique list of taskcounts, job /task names, and thread ids. This file will show up in the server side output files folder. It is created under the covers by the GUI whenever a taskcount is known but the thread id and job name for the taskcount is not known (such as when opening a graph from a table view where the job name field is not included in the list of fields).

July 15th, 2005 – C00423

iDoctor for iSeries is now available for use on DBCS systems.

New commands ADDIDRUSR and RMVIDRUSR at release V5R3 have been added to library QIDRGUI. These commands will be used to grant/remove authority to the function groups QIBM_SERVICE_TRACE and QIBM_SERVICE_JOB_WATCHER. The commands will also add/remove individual authorities to the following system objects:

- *ALL to QAPEXDFN file in QUSRSYS (for the GUI create pex definition interface)

- *ALL to QAPEXFTR file in QUSRSYS (for the GUI create pex filter interface)

- *USE to ADDPEXDFN/RMVPEXDFN cmds in QSYS (for the GUI create pex definition interface)

- *USE to ADDPEXFTR/RMVPEXFTR cmds in QSYS (for the GUI create pex filter interface)

These commands require SECADM/ALLOBJ

Authority checking for Job Watcher added to the GUI just before starting a collection. At release V5R3 this check will be for *SERVICE OR access to function group QIBM_SERVICE_JOB_WATCHER. This check also added to the WCHJOB command.

Added the ability to see the collection overview graphs at any time from most of the job watcher graphs. Also added this feature from table views having a valid interval and taskcount fields in the selected record.

On the Quick View tab of the properties for some of the Job Watcher graphs, all of the fields are not being shown in the "All other fields:" section of this panel

Added the starting interval and ending interval to the V5R3 collection properties general page. Also added checks to see if the collection has been split and updated the fields on this panel accordingly.

Added a new collection search function for V5R3 Job Watcher collections. The possible search types are job or task name, subsystem, pool, current user profile, taskcount, sql statement and call stack pgm/mod/procedure. The search will display a table view in the data viewer which can be used to graph and drill down further into the desired job records. Right-click on a job watch for the new "search..." menu.

Included the license agreements (in various languages for all external components) as part of the install image. These are installed to the directory C:\Program files\IBM\iDoctor for iSeries\license

Added some checking to the min/max scale values on the graph definition->primary Y axis page. A minimum scale value cannot be specified without also specifying a maximum scale value.

Fixed a problem with the control automatic resizing code for dialog windows (open file window, browse jobs window in pex analyzer, select fields window). If the window is resized larger than normal and reopened, the controls do not properly fill the larger window area and will shrink smaller than they are supposed to if the window is resized smaller.

Added a new menu "save" which lets users save the data in a table view, or list view into a csv excel file, or tab delimited file. You can either save the entire list or the selected cell block or selected records. This is now available on places such as the job watcher call stack.

Fixed a bug when writing column descriptions to a text file. New line characters were sometimes incorrectly added to the file as part of the description.

The help buttons in the GUI and in the install now work if the default web browser is Internet Explorer, Netscape or Firefox. Other browsers have not been tested.

The license agreements for either the trial or purchase agreement (based on the component being installed) are now viewable from the license agreement page of the installation. Non english agreements will be shown within the default web browser in order to handle DBCS character display.

If a code page conversion error occurs (CWBNL0101) and the required code page table is not found it will be now be attempted to be retrieved automatically from the managing system using the cwbnltbl command. The user must restart iDoctor before trying the request again however

Added a check to the install wizard on the FTP options page to validate that the FTP connection to the server is good before continuing with the install.

Fixed a bug with establishing an ODBC connection. A call to program QIDRGETJBN (which is used to retrieve the ODBC jobs job name) is being called more often than it needs to be and is adding extra trigger exception error messages to the QZDASOINIT job log.

The iDoctor property page "iDoctor client jobs" is now only displayed if the user has *JOBCTL special authority. Without *JOBCTL the user cannot modify the run priority or time slice values.

May 20th, 2005 – S00087-C00411

Updated the required PTFs to MF35297 and SI17681

Enhanced the interval summary property pages.

Fixed a problem where copying a job watch would not work if the collection had been summarized.

Fixed some issues with the client not handling NULL data correctly when running queries.

May 16th, 2005 – S00082-C00408

Replaced use of the iSeries Access for Windows DB APIs in the GUI with ODBC API calls.

Added support allowing iDoctor to work on DBCS (double-byte character set) systems. The client and server has ran successfully on Japanese and Korean systems. Testing on other language systems is in progress.

Added a function to transfer any library shown in the GUI to another system using FTP.

The last used user id should no longer be saved into the same place used by iSeries Navigator. Instead the last used user id will be saved into an area of the Windows registry reserved only for the iDoctor client. This will mean that the user id used for iDoctor and the user id used for iSeries Navigator may be different and will be remembered independently of each other.

Fixed an issue with system names greater than 30 characters causing errors when establishing the ODBC connection.

When opening a user-defined query or graph definition, added a check to ensure that the primary file in the report exists before attempting to go any further. This will reduce the number of errors messages the user gets.

The check for the required fields in the interactive queries/graphs has been moved to before the menu is opened. This will prevent the menu from showing up if it cannot be used.

Fixed an overflow issue with the scrollbar when dragging through a graph having with a large number of bars.

Updated the required PTFs to: MF34840, SI17681

Fixed an error MCH1202 (decimal data error) that occurred while running Job Watcher

Added a warning message when opening a summary graph if the collection has not yet been summarized. Summarization is highly recommended to capture information about threads that did not use CPU each interval.

Added support to work with V5R2/V5R1 collections on a V5R3 system. This support includes the ability to copy, delete, ftp and view all the graphs and reports available to the collection at each release.

Modified the interface for the call stack collection options so that the dynamic collection options (call stacks for threads in a bad wait) is a separate option from the call stack option for active threads (threads using CPU). That way both options can be used at the same time. Using both options is also now the default selection in the GUI.

When viewing a call stack, added a button to display the holder for the current stack if one exists. While viewing the holder, the button can be pressed again to go back to the waiter job's call stack. In some instances the holder's call stack may not have been collected if the holder was idle during the interval and the advanced call stack options were not being used. If this is the case this button will still exist (no extra check is done) and it will show an empty stack. Clicking on the button again will go back to the waiter stack

Added support to split a job watch into smaller pieces. The ability to split by time or an interval range is provided as well as the option to split the collection into multiple sections (by minutes or by interval size). This option is available for all types of job watches except V5R1/V5R2 job-specific mode collections.

Added a check to make sure the call stack file exists before showing the call stack tab on the interval details page within Job Watcher. Also use this check when showing the "Display Call Stack" menu so that it is only shown when the stack data exists.

The following fields are not being displayed properly on the SQL tab of the interval details window: job subsystem, thread status, current state, job function and pool.

Improved performance when building the list of job watches in a library.

Added new commands to copy, ftp, save and restore collections and options to copy and ftp these collections from the GUI.

Add commands to copy, ftp, save and restore job watches.

Added options in the GUI to copy and ftp job watches to another system.

Selecting multiple collections within a library, then right-clicking and using one of the graph/report options will now open the same report for each collection selected. This should make it a little easier to do comparisons between two collections.

Placing the mouse over field BLOCKENUM in Job Watcher within file QAPYJWTDE should now correctly display the enum description.

Tasks in the job watcher data sometimes incorrectly showed a job status value of "held" within file QAPYJWTDE. This field should always be blank.

Fix a GUI memory leak with the Job Watcher CPU property page

March 28th, 2005 – S00065-C00391

Added a new graph Collection overview with CPU utilization time signature (for collections that have been summarized.) The graph will show CPU utilization as the Y2 axis instead of total threads active/idle.

Add a new detailed wait graph menu shown from within the Job Watcher output files that contain a TASKCOUNT field, and for graphs that show lists of jobs. This new graph will be the run/wait time signature with sliced intervals. This is produced by running the command CRTWCHSLC under the covers which fills data into file QAIDRJWIG3 for a specific taskcount and slice size. It provides a different perspective on the time signature graphs showing the percentage of bucket times within each user-specified time slice.

The Wait graphs by job (if collection has been summarized) should have the following changes: Jobs that started later in the collection will be offset appropriately based on the interval start time. The flyover should display the starting interval and ending interval of the summary for each job. The flyover should display the percentage of each buckets time against the total time.

Modified the CPU/CPUq by high/low priority jobs graph so the Y2 axis contains a line for the high priority jobs and a line for the low priority jobs. Also created another version of this graph with CPU utilization as the Y2 axis (only available if collection has been summarized).

In the Job Watch collection wizard, allow *ALL to be entered for the taskname parameter

Before running a summarization (CRTWCHSUM command), check to see if the Job Watcher subsystem is running. If it's not, ask the user if it should be started. If they say no, then abort the summarization

Add the wait bucket description property page for job watches.

In the collection wizard, within the advanced collection options, set focus to the page and field after the error message is shown if one of the fields has not been entered correctly.

Update the queries behind the Job Watcher graphs to resolve issues with time stamp confusion (interval end vs interval start) and to provide visual indicators in the detail graphs to highlight certain bars of interest. This version will provide an arrow below the intervals for any records that have a holding job. This will make it easier to spot intervals where holders occurred, and to tell the user where the holder menu options will be available. Also when looking at detail graphs for a job that has not been summarized (idle intervals not filled), no time stamp will be provided just the interval number.

When sending a job watch to another system, provide a drop down list of possible remote system names to choose from. The list should only contain systems from the connections view matching the release of the current system

When creating the default summarization files from the GUI, first check that the record formats of the files match those in QIDRWCH. If they do not they should be deleted before the CRTWCHSUM command is executed. Also if the record formats do not match, do not use the queries over the summary files because additional fields were added. The default non-summarized queries over the TDE file should be used instead. The summarized field for a list of job watches should show "No, must recreate" if the record formats don't match. Before issuing the summary from the GUI, attempt to allocate all files to ensure that another job does not have a lock on the files which would cause the CRTWCHSUM command (issued by the client) to fail.

Fixed an issue with the AVGTIME field calculation in the interval summary sync response time graph.

Include the RTVSTKDATA created output file QAIDRJWSTK in the list of server-side output files. This file will only show up if the member name of the collection is a member within file QAIDRJWSTK.

Added support for defining a rule definition when creating a job watch.

When a job watch is created that uses a rule definition, a copy of the rule definition will be created in file QAIDRJWRDB, with the member name matching the collection name. This file will also be displayed in the list of server-side output files.

Change the field description for the ENUM field in file QAIDRJWENM to "Specific block ID number" from "Enum description". The field is actually the ENUM ID, not the enum description.

In graphs and reports containing the wait bucket fields (QTIME01-QTIME32), the bucket descriptions from file QAIDRJWBKT will now be used if they are available.

Not all files were being copied properly by the QIDRCOPYGEN PGM used by the CPYJWCOL command. If an error occurred on one of the files (such as the file doesn't exist), the command would quit rather than continue on to the next file.

Add the following files to the list of V5R3 JW, server-side output files: QAIDRJWLOG - for new interface to V5R3 PEX Flight Recorder approach QAIDRJWIX1 - for a new JW Very Long Transaction details file, QAIDRPXIS1 - for some PEX work that will be linked to JW collections and QAIDRJWIG3

Added call stack reports. These reports allow the user to see the total number of instances where a procedure/module/program was used in the collection by which jobs and for which intervals. These reports are available by right-clicking on an entry in the V5R3 Job Watcher call stack property window, or when right-clicking on a report that contains the TBTADDR field from the QAPYJWPROC file

Added the ability to graph the holder for a job/thread within an interval if one exists. These menus should be available within details graphs or within any table view that contains the HTASKCNT (holder's taskcount) and INTERVAL fields where HTASKCNT > 0

If viewing file QAIDRJWIX1 and opening a detail graph, position to the first interval of the transaction of the user's selected record in the file

When viewing the properties for a single job, the label 'current wait' should read 'current or last wait'. Also change the 'current wait duration' label to read 'wait duration'. Below the thread status list the current run/wait state of the job from the CURRSTATE field in the TDE file. Also display the POOL id of the job next to the priority field

When displaying the property pages, retrieve data over the QAIDRJWIG2 file if it exists. This will provide more information in the property pages for idle/missing TDE records.

The scheduling option in the collection wizard doesn't pick the correct date selected from the calendar in most cases

Opening a server-side output file, then opening a graph view, will cause the report being opened to lose the flags that keep track of whether or not the job watch has been summarized. This means that the unsummarized detail reports (without idles filled in) will be incorrectly used instead of the summarized reports

The menu Display Call Stack, should now show up for any report or graph opened in Job Watcher that contains a TASKCOUNT field > 0 and an INTERVAL field > 0. For user-defined reports/graphs, until the query or graph has been saved as a Job Watcher query/graph and reopened, this menu will not be available

Fixed a problem with the query definition -> reset option failing for reports that were spawned from other graph views/table views.

When viewing a call stack within the property pages, added a menu option and window to display the full procedure name. The procedure name can be up to 1000 bytes long which is too many to display in the procedure name column of the call stack view.

If it cannot be determined whether or not call stack and SQL data has been collected, the call stack and SQL property pages will be shown by default

When installing Job Watcher, if the PTFs are not on the system and the user is prompted to get the PTFs, the install can crash after this point if the user decides to hit the finish button again while the install is completing. This triggers a 2nd instance of the install to run which crashes the install application.

If the TASKCOUNT and INTERVAL is known and opening a detail graph, position to the interval that the user came from.

When using the popup menus for the wait bucket graphs while in a user-defined graphs they should now should more than just bucket 1. The bucket description array was not being loaded properly in the instance of user-defined graphs over job watcher data

In JW, the detailed property pages for a single job do not currently show up when clicking a job in the run/wait signature job summary graphs when the range of intervals selected is > 1. If the range of intervals is > 1, show the first interval of the range in the detailed property pages

When dragging the scrollbar for a graph view, an indication of the current scroll position is now provided in the status bar.

Property pages that are non-modal (meaning you can do other things while these windows are active) can now be minimized. An example of this type of property page is the interval details properties for a job within Job Watcher.

Improved performance when loading the list of queries/graphs in a folder by not loading the column descriptions for each query or graph until necessary (just before opening).

Modified the quick view page of a bar in a graph's property pages so that each piece of the graphs data is listed separately. X-axis, primary y-axis, secondary y-axis, other primary y-axis fields, and special fields should each be listed in separate sections at the top of the list of fields for easier readability.

Enhance the code that builds the graphs so that if all data for a bar (field) in the graph is zero, the bar will not be added to the graph control's data set. This will save memory on the graph control, improve performance while scrolling and make the legend more readable by only showing colors for data that exists in the current page of each graph.

When working with the fields for the Graph definition -> Primary Y axis panel, the list of fields can now be reordered using drag and drop.

For the Graph definition -> Primary Y axis panel allow the user to determine the border color and the border thickness (1 to 5 pixels). A color for the field must already be defined for either of these options to work.

For the Graph definition -> Secondary Y axis panel allow the user to enter two Y2 axis lines and to specify the color and legend description for each line

In the legend for graph views the secondary Y axis fields (if any exist) will be shown as thin lines instead of bars

Allow the user to specify a width for the Y2 axis lines (between 1 and 5 pixels) If a width has not been selected the default line width will be 2

The minimum value for the Y axis scale was not being loaded properly from a saved graph definition.

Fixed a problem with hard to read graph flyovers. Depending on where the graph/mouse pointer is, sometimes the entire text of the flyover is cut off and not readable.

Added a splash screen with trademark/copyright statements to the install and iDoctor GUI. Clicking on the window will make it go away.

Within the install program, cleaned up the server side install parameters page: removed the "create if necessary" checkbox, the user will get prompted with a Yes/No question instead. Removed the checkboxes on the final page of the install if installing just the client since they only apply to a server installation. If unable to connect to the system, improved the error messages given.

Fixed a problem when loading a user-defined graph that has user-created field expressions in the underlining query. The descriptions defined for these field expressions were not being loaded properly into the query definition->field selection window.

Fixed an issue with hidden fields being no longer shown in the field selection page of query definition if a field expression has been defined over a hidden (not checked) field.

Fixed an issue when right-clicking on a flyover window in a graph. The event did not get properly sent to the parent window, causing the popup menu for the graph to not be displayed.

The function to create a text file on the PC from a table views data did not work correctly if the record length was relatively small (< 100 bytes)

Make the parsing of SQL statements work regardless of the number and placement of carriage returns and line feeds within the SQL statement. There were several instances where problems with query definition would occur depending on where the new line characters were within the statement. Note: if a user of the SQL interactive view modifies the query definition through the query definition interface, any carriage returns and line feeds in the SQL statement part of this view will be lost. These carriage returns/line feeds are also lost when the query definition is saved

Improved the window that pops up if an SQL statement syntax error occurs. The window now contains the message id of the error, the 2nd level error text/recovery and the complete SQL statement in a copy-enabled window.

When retrieving the field SQLSTMT from file QAIDRSQL04 (iDoctor user-defined query file) for any component/release do not strip off the trailing blanks if any on the SQL statement. Normally all fields are stripped of trailing blanks when read into the client.

Fix a SQL parsing bug that would occur if the numeric operators were used in a WHERE clause without spaces before and after the operator, (=, <=, etc)

Bringing up a window within a window will often cause the focus to be incorrectly reverting to a window in the background (not in the iDoctor application) once both windows are closed.

January 13th, 2005 – C00367

Fixed a problem when changing the query under a graph that would cause the scroll position to become invalid because of the number of records in the query changed to a smaller number. Now the scroll position will be moved to the first page if this happens.

When editing an expression in the field selection window of query definition, the popup edit window does not word wrap long expressions within the textbox. Also resizing the window should resize the text box so more is visible at one time making it easier to edit.

December 20th, 2004 - C00366

The limit of 32,000 bars in graphs has been removed. The new limit is 2,147,483,647

Improved the display of the legend in graph views by automatically resizing the fonts shown in the legend proportional to the size of the graph.

User-defined queries and graphs are now listed alphabetically by name (by default) rather than by date.

When a user resizes columns in table views or list boxes in the application, the sizes will be saved into memory and when the view is refreshed the column widths will be retained rather than being lost. These widths are only temporarily saved and are not saved as part of a query definition. They are also lost if the query definition changes (change field ordering for example). In the component list views, the column widths are only saved until the current tree position changes.

Fixed issues with the interactive SQL statement view: 1) Incomplete prompting of members after initial run. 2) SQL column headings not multiline on initial run, etc. 3) SQL statements containing embedded carriage return/line feeds failed to run in some cases.

From the Server-side Output Files folder (and also or in the Analysis Reports folder in PEX Analyzer) provide a menu option to display the properties for one of the files in the list. This will provide the file's object description including detailed field information for the file.

Fixed issues with the query definition interface in parsing 3 or 4+ way joins in SQL statements.

Pressing the Enter key while a property page is visible will no longer close the window. The escape key may still be used to do this.

Added new types of Job Watcher graphs showing IO statistics ranked by job.

On the interval details property pages, add labels on the Transactions, IFS, and Logical IOs property pages that indicate that the numbers are coming from the job level and not the thread level. These statistics are only counted across all threads within a job.

For graphs in the Wait Graphs by Interval folder, replace the Y2 axis description "Total active threads/tasks" with "Total collected threads/tasks". "Collected" is a much more accurate description due to the new force 1st interval write option at V5R3, and the fact that jobs in a bad wait are collected at V5R2/V5R1 even if they didn't use CPU.

The wait bucket colors used in the graphs are now more consistent across all releases.

Added the collection overview time signature graph. Also updated all queries behind the wait graphs by interval so they use the same query each containing all 32 wait buckets. These values are displayable by clicking on any bar in the graph.

Added a CPU utilization graph under the CPU graphs folder if the idle thread graph files have been created. The CPU times shown should include CPU, CPU01-CPU (wait on chip time) and CPU queueing. The Y2 axis will be the percentage of CPU used.

The install and client startup is now checking (and requiring) PTFs MF34003 and SI15800

Fixed a bug with RSTJWCOL (restores a Job Watch after an FTP operation). The allow object differences parameter should not be set to *ALL because this will cause the collection to not be restored properly in many cases during an FTP operation if the file level identifiers (or object owner) of the source and target files do not match.

Fixed a bug with CPYJWCOL (copies a Job Watch). Copying a collection in a library with other collections having a similar name will cause extra records to be added to the control file in the target library. This was due to a bug in the program that does the copy of the control file record. Also the files produced by CRTWCHSUM will also be copied by the CPYJWCOL command. The control file record from file QAIDRJWIS0 for the default analysis should also be copied along with files QAIDRJWIS2 and QAIDRJWIG2

In the collection wizard, provide an option to schedule the date and time to run the job watch. This will utilize the scheduled date and time parameters on the SBMJOB command. The new panel should appear after the ending options page of the Wizard.

When submitting the default summarization analysis for a job watch, provide the command string in the window so that it may be copy and pasted to another window.

In the collection wizard when specifying job selection always OR job selection criteria together. Previously different categories of job selection were logically ANDed together.

Added the option "transition to none" on the ending options page of the collection wizard. This will only be visible on this page if the user has selected specific jobs/tasks to watch.

Added the option to "Force 1st interval write" on the job/task option page of the collection wizard.

For Job Watcher reports, the properties menu for a table view should always be enabled, even if the query produced no data. In this case only the query tab showing the current sql statement should be shown.

If a job run/wait signature graph (for a single job) is opened, scroll to the 1st interval that contains data within the graph. The user will still be able to scroll backwards into empty intervals if desired if they are trying to make a comparison with another graph at that point in time.

From the holder interval summary page, new menu options are available for entries in the list of holder to allow the user to see the interval details for a specific holder job or waiter job in the list. Note: it is possible that the holder was not collected and in that case not much information is provided beyond what can be derived from the QAPYJWSTS file.

In the collection wizard on the ending option page, enhanced the options for interval ending to show estimated completion time and for time ending to show a drop down list allowing the user to pick how to enter the time value (in secs, mins, hours or days)

At release V5R3, added support to start a job watch over a job on a job queue.

On the interval summary holder property page, the field ordering has been switched around to show the holding job/task first and waiter job/task second. Also the default sort sequence is now 'wait time descending'

Added support for the CRTWCHSUM command. This command creates additional files that fill in and total 'idle waits'. Because by default Job Watcher only collects records for jobs that are using CPU each interval, the idle waits are important in order to get a more accurate view of a job's run/wait time signature. A new menu option under a job watch allows the user to create the "default" graph files containing idle records. These graph files will now show up under the server side output files folder. If the graph files exist for a collection, the reports will be over these files instead of the QAPYJWTDE file. A new field exists in the list of job watches in a library after the collection size indicating if the extra graph files have been created

In a JW job summary graph over a specific interval, when opening up a detail graph for a specific job, scroll to the interval number indicated in the summary graph within the detail graph by default.

On the add jobs page of the collection wizard, the list of jobs will no longer go invisible then become visible again when hitting the refresh button.

On the JW interval details property pages on the object waited on tab, if a physical file is the wait object and the wait type is DB record lock, the record number is listed after the file name.

In the Job Watcher interval summary graphs, added new popup menus allowing drill down into the jobs occurring for the selected interval and bucket. These menus show job run/wait time signatures for a single interval ranked by a wait/bucket type.

On the interval summary property pages on the objects waited on tab if there is a wait object in record lock conflict, the record number is now listed after the file name. The object name also now has a slash between library and file for readability.

On the interval summary property pages on the objects waited on tab and holders tab modified the description above the list boxes on the window to clarify that the lists are built from data collected at the end of the interval, and do not reflect objects waited on or holders that occurred throughout the entire interval.

On the interval summary property pages if they are invoked from a interval summary time signature graph by clicking on a specific bucket, that bucket will be selected and show by default on the Wait buckets page when the wait buckets page is visited.

For V5R3 JW, there's a problem with opening the list of active jobs on the system from the collection wizard on systems that do not have library QGY. At V5R2/V5R1 this was a work management library and contained an API to build the list of jobs. This API is now in library QSYS at V5R3. Didn't catch this because most systems seem to have library QGY at V5R3 as well so the API exists in both libraries.

When the wait graphs by job folder is expanded, show the worst types of waits first, and include an "other graphs" folder at the end of the list that contain the rest of the (less interesting) wait buckets. This will mirror the way the popup menu for these graphs behave.

Make the server-side output files folder in all components have a consistent set of fields available within it. Allow the option to select which fields should be visible. Fix the labels in this folder in Job Watcher (the column headings are reversed) See detailed description for a list of fields available in this interface: The default fields shown and ordering is 'Output file, Description, Records'

For JW V5R3 in the collection wizard on the add jobs page, there is a button to add a filter and a button to add a specific job name/user/number from the list after refreshing the list (based on the filter). Filling in a specific job name/user/number into the filter section and hitting the add filter button will place *ALL for the job number instead of the job number entered.

For V5R3 JW, the collection wizard asks for taskcounts to be entered in HEX format when the command expects taskcounts in decimal format. Updated the GUI interface to ask for the taskcounts in decimal format and perform parameter validation to prevent users from entering the value in hex format. Also added parameter validation for the "current user profile" option.

When running Job Watcher at release V5R3 remove the And/Or column on the job selection page of the job watcher wizard.

When running the GUI at release V5R3 and looking at a file that is not an iDoctor-supplied report remove the option to save a user-defined query definition or graph definition as type "job watcher-job-specific". Replace the "job watcher system-wide" option with "Job Watcher"

September 13th, 2004 - C00345

On the Interval Details Property Pages added the holder's task count (in parentheses) to all property pages that shows the holding job/thread/task information. Also fixed a display problem with the holding information field if the holder was a task.

Fixed some problems with the creation settings property page for a Job Watch.

If the user specified a collection with dynamic call stacks, the call stacks were not visible through the GUI.

August 16th, 2004 - C00342

New version released!

Collection engine

Completely rewritten engine providing faster collection with more options.

Dramatically less data collected per time interval and a corresponding dramatic reduction in the resources consumed by the collecting job. V5R3 Job Watcher adopts a Collection Services like design, in that only a tiny amount of info is put into DB2 files for jobs/threads/tasks that have been completely idle for a snapshot interval.

"Only collect data on active threads/tasks" change, and a change that reduces the number of DB2 bytes consumed for Call Stack information, system-wide Call Stack harvesting is now feasible. Job Watcher now opens new possibilities for system-wide collection. In the V5R2 version, it was not even possible to collect Call Stacks for every thread/task, only for those that were "running" or in certain types of predetermined waits. It would have been impractical to allow all threads/tasks to have their Call Stacks harvested, because the engine did not differentiate between active and idle threads/tasks.

New ways to select which jobs/tasks to collect: System-wide, unique or generic job or task name, pool ID, subsystem, task count, current user profile.

Option to write out data in huge blocks rather than every interval, providing significant collection time performance improvement at the cost of delays in seeing the data.

Multiple ways to end a collection: DB size limit, number of seconds, number of intervals, selected jobs/tasks become inactive.

ASP threshold option to end collection if total disk space goes over a specified limit.

GUI

New graphs (examples: Transactions, CPU by priority, System-wide seize contention per interval).

New graph categories to make it easier to locate the data of interest.

New property pages at the summarized (over multiple jobs) and detailed levels (a single job) to consolidate information and make analysis easier.

New options to manipulate Job Watcher data providing copy, FTP, save, restore support using the GUI or green screen commands.

Easier to use Collection Wizard.

Automatic font resizing in graph views.

Enhancements to the query definition interface.

Improved call stack viewer, with information about why the call stack was collected and an option to view the call stack for the Holding job if appropriate.

Database sizes are now listed for each Job Watch.

Unreadable Hex data in database files now automatically readable without needing to write additional queries.

Option to view the host variables parsed back into SQL statements collected

Naming Changes

New library QIDRWCH replaces QPYRTJW.

New library QIDRGUI replaces QYPBASE.

Database files named QAPYJW* instead of QPYRT* and shipped with OS in library QSYS.