

Update History for iDoctor for iSeries – Job Watcher for OS/400 V5R1

Feb 12, 2007 – C00520

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

Fix a tab order problem in the transfer collection window.

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.

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.

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.

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.

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 the ability to schedule a job-specific or system-wide job watch at release V5R2 or V5R1.

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.

February 28, 2006 – C00479

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

Fixed a problem with collection properties. The collected tasks field is blank on the Tasks page.

Fixed a problem with job-specific mode interval details property pages. The job information field is blank when viewing the details for a system task.

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.

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 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

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 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.

Enable the collection search function developed in V5R3 JW, for use over V5R2 and V5R1 system-wide job watches.

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.

When looking at the interval summary property pages for a system-wide job watch, the display call stack menu is incorrectly shown for jobs in the list. The call stack is not available for V5R1 system-wide job watches.

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.

July 15th, 2005 – C00423

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

New commands ADDIDRUSR and RMVIDRUSR at release V5R2 have been added to library QYPBASE. 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)
*USE to ADDPEXDFN/RMVPEXDFN cmds in QSYS (for the GUI create pex definition interface)
These commands require SECADM/ALLOBJ

Authority checking for Job Watcher added to the GUI just before starting a collection. At V5R1, this check will be for *SERVICE This check also added to the WCHJOB and WCHSYS commands.

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

Fixed a bug with the SQL statement behind the Holding threads/tasks system -wide interval properties at releases V5R2 and V5R1.

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\Doctor 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 – C00411

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

May 16th, 2005 – S00130-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.

Added support to the GUI 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.

Add a check to make sure the call stack file exists before showing the call stack tab on the interval details page within Job Watcher.

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

Improve performance when building the list of job watches in a library. It takes too long if there are many collections in the library.

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

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

The WCHSYS command default for parameter WCHOUT is now *FILE2 instead of *FILE1.

For system-wide job watches, if the file QPYRTSWLG2 is not in the collection, disable graphing support since most graphs use this file in the graphs.

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 QPYRTSWLG2 should now correctly display the enum description.

Updated the job descriptions shipped with Job Watcher to reduce the amount of logging occurring by default when running a job watch. This will prevent job watches ending when the message queues become full from repetitive messages occurring each interval when running the WCHSYS command. Also updated the code behind WCHSYS command so that the CHGJOB it issues does not change with the logging level currently in effect for the job. That way the JOBD logging level will be used instead.

March 28th, 2005 – S00117-C00391

Fixed an issue where the buckets 28-32 don't display in the time signature graphs. This was due to a problem dealing with the unused buckets 14-18.

In job-specific mode, going into the properties for a detail graph (run/wait signature for example) will cause the job information to be displayed incorrectly as the X-axis timestamp value.

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 QAIDRSQLO4 (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 an install issue where program ADDWCHMSG would fail in some situations.

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.

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.

August 16th, 2004 - C00342

Added graph categories. New folders and popup menus under a job watch make it easier to find the graph of interest. Many new graphs have also been added.

Improved the rule definition interfaces.

All job watches now collect data at priority 1.

Fixed some issues with the query definition interfaces.

Changed the default for the WCHSYS command so that the files produced will be useable by the GUI. This is the *FILE2 option on the WCHOUT parameter.

Added property pages for many views within system-wide summary graphs and detail graphs.

The number of bars per page to show on a graph can now be defined depending on if the graph is a vertical or horizontal bar graph. This may be done through the General page of the Preferences window.

To improve usability added the ability to display the detail graphs from any table view or summarized graph view that has the required fields available to complete the query: TASKCOUNT, NAME, THREADID.

Added code to keep track of the amount of graph data open at one time on the PC and prevent users from opening another graph if getting close to running out of memory.

February 4th, 2003 - C00288

Added a new "components window" at startup which provides a list of all components installed on the system, so a user can see what components are installed prior to attempting to start one. You can also apply an access code, and change the user profile signed onto a system from this window.

Enhanced the interface for defining rule definitions in the Job watch Wizard.

When a Job Watch rule definition is used in the creation of a Job Watch, a copy of it will be made into a new file QPYRTJWRDB (job-specific) and QPYRTJWSDB (system-wide) depending on the type of Job Watch. The member name will match the Job Watch name.

Added a property page to display information about the state of the system at the time the Job Watch was started.

When a Job Watch is started, the command string used to start the Job Watch is saved into a new file QPYRTJWDFN (job-specific) and QPYRTSWDFN (system-wide). This is used to display the job watch creation settings panel in the property pages for a Job Watch.

Added a new report accessible from the system-wide graph views (Detailed and Summarized). The menu is called "Display Interval Detail Record" when right-clicking on a bar in a detailed graph and "Detail Records" when right-clicking on a bar in a summarized graph. The report will show all the records from the file QPYRTSWLG2 matching the specified job/interval.

Removed the Print button on the Graph Coordinate Properties window because it did not work properly.

If the connection behind the remote command status view is lost, it will now be automatically reconnected when the next request is made.

Removed the Job-specific 'Current activity' reports.

Fixed a bug where the client would potentially hang if any server program it calls 'Function checks'. Now the client will get control back and provide message to the user.

Fixed a bug with the call stack parsing done on the client.

A new green screen command ANZJWSTK is now available in the Job Watcher library QPYRTJW.

December 17th, 2003 - C00281

In system-wide mode fixed a problem where the wrong command was being issued to reset the wait buckets.

Added the ability to read in an existing system-wide rule into the Job Watch Wizard for the purpose of reusing it.

November 19th, 2003 - C00272

Added support for new parameters on the WCHSYS command. This is the command that starts a system-wide Job Watch. These new options are available on the data configuration window of the Start Job Watch Wizard when in "system-wide" mode.

Added an estimated collection running time to the options page of the Job Watch wizard based on the maximum intervals and interval size values.

Fixed a problem when signing on a system where the focus was not on the user id field if the client did not already know one.

Removed duplicate error messages that would sometimes be displayed to the user.

Most of the summarized Job Watcher reports and graphs now summarize on the TASKCOUNT field instead of the NAME, THREADID fields. This was necessary to fix a problem where multiple task counts could be used within the same job name, thread ID.

Added a warning message to the install if the system value QALWOBJRST is not set to *ALL or *ALWYSYSTT. The install requires this value because we ship objects that have system state and adopt authority attributes set.

We have removed the support for launching the iDoctor application from iSeries Navigator.

The bottom of the summary page on the Start Job Watch Wizard will now show the CL command string that the client will issue.

Added entries in the list of output files to display OBJTYPE and SEGTYPE in the folder "Server-side output files" for both job-specific and system-wide Job Watches.

Fixed some problems with the positioning of the flyover windows in the graphs.

Enhanced the properties for a Job Watch to include totals information and other basic information such as when the Job Watch started and ended and how long it ran.

The "Version Management Wizard" is no longer launched when the client is first started after an install has occurred. This wizard is still available by right clicking on the Job Watcher icon.

Fixed a problem that would occur if the connection was terminated. Previously it could not be reestablished without restarting a client.

In system-wide Job Watcher, right clicking on an interval in a detailed run/wait signature graph can now create a summarized run/wait signature graph over a single interval.

Added Job watcher definitions. These allow the user to save/reuse their settings in the Job Watcher Wizard for future use. The definitions are stored in file QPYRTJWDFN in library QGPL.

Removed the zoom/zoom out feature from the graph views. It was determined to be not very useful.

Improved the speed of the client when scrolling horizontally through data in table view.

May 20th, 2003 - C00225

In the Data Viewer fixed a bug with the save as function. There was a problem saving data with very long records (such as 10000+ bytes per record for the SQL statement file in Job Watcher).

Changed the 'copy' menu on a tool tip in the data viewer to 'copy popup text to clipboard'

May 9th, 2003 - C00222

Fixed a problem with the call stack window.

May 8, 2003 - C00219

Removed several dependencies on ODBC in favor of the Client Access DB APIs. The API implementation uses blocking so in some cases bringing up a list of things (user-defined graphs/queries) should be noticeably faster.

Added more details to the Job Watcher call stack window including the current wait object, holder, current wait duration, lic wait type and current user profile. Also added flyover text to show the program descriptions if available when putting the mouse over a program name in the call stack. Another new feature is the job information drop down list that allows a user to easily switch to a different job's call stack within the same interface.

Fixed a hard crash in the client within the Job Watcher graphs that would happen if the call stack or interval details dialogs were open and then a bar was clicked to get the bar details.

Removed code that was making the graph title in user-defined graphs views lower case after the first character.

Fixed a problem with the menu for user-defined job-specific Job Watcher summary graphs.

When running a "user-defined system-wide Job Watch intervalized over a specific job" type graph the client will now prompt for a job name/thread id to substitute into the query definition.

Fixed a minor spacing problem with the flyover windows in the system-wide Job Watcher time interval graphs.

Fixed a bug in the job-specific job watch rule definition window where the job watch would not run if the 'action' was to call a program and the first parameter on the pgm call was left blank.

Added the 5 Top N parameters to the Start Job Watch Wizard options page. These will show when creating a System-wide Job watch and are available by clicking the Top N options button on the options page.

Added the ability to specify the number of intervals to run a job-specific Job Watch launched from a system-wide Job Watch rule.

Fixed a problem with all uppercase text in the graph view 'flyover' windows.

Added the ability to modify the having clause of a query within the query definition interface on the 'group by' page.

Made several enhancements to the query definition interface and syntax checking on SQL that iDoctor makes to feed queries into the interface.

Fixed a bug where the data viewer window would restore itself to its previous size if it was maximized and a new report was opened into the existing data viewer.

The graph control supports a limited number of points displayable at one time. If this limit were exceeded (somewhere around 15000 points) iDoctor would crash. Instead of doing this I have modified the graph control code to display an out of memory error message instead and allows the user to continue using the client. The most recently opened graphs will show as empty windows however and will need to be closed down.

March 11, 2003 - C00196

Improved the client/server version compatibility checking. This will prevent users from using the tool when the appropriate client or server level is NOT installed. The warning provided to users in the client was often ignored.

March 6, 2003 - C00195

In the install, cleaned up the error message that would come back if user lacks authority to install the tools. Added text on the server parameters page of the install wizard to clearly indicate the required special authorities. Also changed a label on the access code page to clarify that the access code is component specific.

Added a check at startup of the client for the required Job Watcher PTFs on the server.

Added a warning message that is displayed when selecting the 'full speed' option in the Start Job Watch wizard.

Fixed a bug where the detailed graphs didn't work from within the summary graph when the summary graph was launched by right clicking on the job watch folder instead of from the summary graphs folder.

For all graphs, added the preference on the "preferences" window called "Enable automatic scaling". If this option is turned on then the graph will behave the same as they do today, and as you scroll through data in the graph, the scaling adjusts based on the maximum height of all bars. If you turn this option off, then the scaling is fixed based on the maximum height of the first page of the graph displayed. Scrolling through other pages will not change the scaling.

Reorganized the graph definition interfaces and added new "graph type" and "show legend" options. The graph type option lets you pick between vertical stacked bar graph, vertical bar

(side by side bars per x point), horizontal stacked bar graph and horizontal bar (side by side bars per x point).

Fixed a caching problem that could sometimes cause user defined graphs to not be openable if one was deleted and then recreated over different fields within the same iDoctor session.

Changed the behavior of the popup windows that come up when the text of a cell isn't completely visible. There will be no popup window anymore for the first column of cells in a list view. This is to prevent problems right-clicking a row in the list and having the popup window get in the way.

A single Data Viewer window can now open reports from multiple systems. Previously all reports in a Data Viewer had to be for the same system. As part of this change, you can now select the system you want to use on the 'Open file dialog' and the 'SQL query view' within the Data Viewer.

February 12, 2003 - C00189

Fixed a problem with the Start Job Watch Wizard: system-wide rule selection page (define action popup dialog). The user was not able to select an action type "Start Job Watch" or "Start Job Watch (over holding job)". **Note:** Doing the rule initiated Job Watch for 10 intervals is the only option currently available until changed on the server side.

Cleaned up the finish page text on the last page of the Start Job Watch wizard when creating a rule file for a System-wide Job Watch.

Fixed bug where it was not possible to delete old query definitions that were converted to the new file format after client 185.

Fixed bug for user-defined 'group by' queries not being rerunnable (meaning to open into the query definition interface and then rebuilding the statement by closing the interface) within the client if they contained expressions in the group by clause.

February 10, 2003 - C00187

Migrated the iDoctor-supplied SQL into the new file QIDOCSQL04 from QAYPESQL03. This will allow us to support multiple versions of queries if Job Watcher files change from iDoctor release to release. This same file is also used to store user-defined queries. At startup the existing user-defined queries will be converted to the new format.

Migrated the supplied-supplied graph definitions into the new file QIDOCGPH06 from QAYPEGPH05. This will allow us to support multiple versions of queries if Job Watcher files change from iDoctor release to release. This same file is also used to store user-defined graphs. At startup the existing user-defined graphs will be converted to the new format.

You can now create user-defined queries and graphs just like in PEX Analyzer for any job-specific Job Watch. These may be scoped to the current instance (watch), library or system. When scoped to the system level it means you can define the query or graph once and reuse it for any job watch.

Updated the version management wizard code for Job Watcher to fix some problems with new files not being checked for.

Added a "High" option for the SQL statement detail parameter in the Start Job Watch Wizard. This option provides even more detail such as the open cursors and prepared statements found within the jobs being watched.

Added a check to prevent users from creating a Job Watch of either type in QTEMP, QSYS, QGPL, QPYRTJW, and QPYBASE

Added the capability to define rules and actions over a system-wide Job Watch. This may be used to start job-specific Job Watches over poorly running jobs identified by the rules you define.

Fixed a problem with the way windows get stacked on top of each other when opening multiple views in the same data viewer.

Renamed "Clear Query Definition" menu to "Reset". Reset will reset the query to its original state. This means if you open an existing query, change the query, and then use the Reset menu, the query will revert back to what it was when it was first opened.

Added a warning message in the Query Definition Record Selection page when you have made changes without adding them to the list. It was a common error among users to fill out the page and not click the "Add Filter" button before clicking okay.

Removed the "work with queries" and "work with graphs" interfaces from Job Watcher. Added user-defined queries and user-defined graphs folders underneath a job-specific Job Watch instead.

November 30, 2002 - C00172

Within Query Definition enabled support to specify user-defined fields in the group by part of a query.

Within Query Definition made a change so that the list of fields will be automatically/immediately updated on the sort, record selection and group by pages, after any new fields are created on the field selection page. Previously you had to run the query and then reopen query definition to see the new fields defined on the other pages.

List columns that have been manually resized by the user will no longer get resized automatically to show all the data when scrolling horizontally.

Fixed a die-hard bug with the AND/OR logic part of the interface on the query definition/record selection window. The And/OR column in the filter list has been moved from the first column to the last column as it makes more sense that way (and simplifies my code).

Fixed a bug where lower case strings in an Sql statement were being converted to upper case erroneously by the query definition processing code.

October 25, 2002 - C00163

Added an iDoctor client jobs panel on the application properties pages accessible by right clicking on the PEX Analyzer or Job Watcher icon. This lets you control the time slice and priority given to the iDoctor client jobs QZDASOINIT and QZRCSRVS.

When scrolling through data in a table view horizontally I will now automatically resize the currently visible columns so all data is visible (only if the column has not been resized smaller than 15 pixels in case someone intentionally has hidden columns). This will change the interface to be a bit more useable when the column headers are much shorter than the data content.

When saving a user-defined query or graph definition, the default description is now the report name instead of the title bar of the table view.

Fixed some problems with the Boolean AND/OR options on the record selection page of query definition. The add selected button would add the exact opposite Boolean value of what it was supposed to after the 2nd entry.

Disabled the 'clear query definition' option for group by queries as it would usually break the query. The original intent on this option is to remove the filter (where clause) and sort order (order by clause).

October 17, 2002 - C00161

Fixed a problem where record quick view didn't work from the list portion of the tree/list view.

Fixed the Job Watcher Current activity IO activity report. Some of the fields were missing.

Removed all references in reports in the GUI to the terraspace stack size fields because they are currently disabled and not being collected by Job Watcher. These fields are: TSUSRSTK and TSSYSSTK in file QPYRTJWA

Opening Job Watcher graphs is now a bit faster than before. The graph and query definitions are now cached on the client after being retrieved from the server the first time.

Added some new IFS summary and detail graphs for Job-specific Job Watches. These graphs are: IFS reads, IFS opens, IFS creates/deletes, and IFS lookup cache statistics

Fixed some bugs in the Job-specific summary graphs that would cause queries to fail when N is a very large value. The summary graphs affected were: Page Faults, I/O activity, DASD waits, DASD reads and writes

October 14, 2002 - C00160

Added the ability to run System-wide Job Watches. There are now an entirely new set of reports and graphs over the system-wide Job Watches. The Start Job Watch Wizard has been modified to accommodate creating either system-wide or job-specific Job Watches. Let us know what you think about these reports/graphs.

Added a new DLTWCH command to delete a job watch. This is called via the GUI and can also be used by green screen users. It supports deleting either system-wide or job-specific job watches.

On the Job Watcher Server configuration page (found by right-clicking on the Job watcher icon) added a line displaying the Job Watcher PTF level installed on the system.

Added new graphs in Job Watcher for Java heap analysis.

Fixed some bugs and improved performance within the Start Job Watch Wizard.

Renamed the concept of defining "triggers" in Job Watcher to "rules" to coincide with database people's concerns.

Added several new fields as possible rules-file fields in a job-specific Job Watch.

Added categories to the list of fields on the page used to define the rules-file conditions to make finding things from the list easier. Any suggestions on these categories or the names of the categories themselves let me know.

Modified the fonts used in the Job Watcher graphs to be slightly smaller. Changed the display of a job/thread in a graph to only show 8-byte thread ID instead of 16 byte.

If the iDoctor client needs a subsystem and it is not started, added a prompt asking the user if it should be started. Previously you had to start via the green screen. The most noticeable place this affects is prior to starting the Start Job Watch wizard but this same check is done several other places.

In the install for PEX Analyzer and Job Watcher added the option to create the subsystem if the user confirms that should be done. Previously it had to be created via green screen.

When displaying a time interval graph in Job Watcher the last 3 digits (were always 0 cause they were usecs) are no longer shown.

Added time signature graphs, these show the total bucket times (in seconds) instead of as percentages.

I've added the Job Watch name and the Job Watch description given when it was created to all hover help/fly over window graphs in Job Watcher.

Fixed a scrolling problem in horizontal bar graphs.

July 31, 2002 - C00140

Removed the option to view detail and summary graphs BY JOB. This isn't possible at the moment based on the way those graphs have been designed.

Fixed a problem where the end of line character would sometimes not be written out when copying a set of rows to the clipboard from a table view.

Fixed a problem where the Copy button was broken on the work with queries window.

When working with tool tips in a Table View, they will no longer overlap other windows.

The open file dialog now allows you to open multiple file/members at once. Also double-clicking on a member from this window will open the member now.

Fixed a problem in the SQL last interval report where the default fields were not being set at install time.

Fixed problem where if the connection were lost, the client would not prompt the user for a password if the password changed.

June 17, 2002 - C00126

The install will now add a shortcut to iDoctor for iSeries on the desktop.

The install will add several shortcuts under the \Program Files\ iDoctor for iSeries\ folder in the Start Menu. One of these shortcuts is to a new uninstall program. The uninstall program removes the iDoctor for iSeries client from the PC. Other shortcuts under the start menu are useful website links for iDoctor for iSeries and a shortcut to launch iDoctor exe. As part of the changes to add shortcuts to launch iDoctor for iSeries, I have removed all dependencies on

Operations Navigator. Ops Nav does not need to be installed to run iDoctor for iSeries, and launching iDoctor from within Ops Nav is still supported but is optional.

Fixed a bug in the JW Wizard that made it impossible to select a task first and then a job within the same wizard instance.

June 13, 2002 - C00125

Within the "System watcher" query in the Add Jobs window of the Start Job Watch wizard all 32 buckets (if defined in the bucket description (QPYRTJWAT2) file) are now displayed. Before only some were shown and the descriptions were hard coded.

Added support for loading existing trigger definitions into the Job Watch wizard. On the options page specify an existing trigger definition by selecting from the drop-down list in order to do this. (The values from the trigger file will be automatically filled into the trigger options and trigger selection pages of the wizard. Any changes made overwrite the existing trigger file member). If you provide a name of a trigger that doesn't exist, it will be created based on what you enter on the trigger options and trigger conditions pages of the wizard.

Added bucket descriptions for the QTIME01-QTIME32 and QCOUNT04-QCOUNT32 in the field selection drop-down list within the trigger selection page of the JW Wizard.

Added several LDIO and IFS fields to the trigger selection page of the JW Wizard. This lets you define a JW trigger over these new fields from the QPYRTJWAPR file.

Fixed a bug where I would delete the JW status user space and flag the JW as unusable if the library name read from the user space did not match the current library. This was making restoring a job watch to a different lib not possible.

June 10, 2002 - C00124

Added a 'Watched Threads' folder under a Job Watch folder to allow faster access to the list of threads being watched.

Fixed bug when displaying list of job watches. If the duration value of "as fast as possible" was selected the value "1001" was incorrectly being displayed instead of "**FULLSPEED**".

May 29, 2002 - C00122

Fixed a bug in the Job Watch properties window for building the list of jobs in the watch.

May 28, 2002 - C00121

Implemented dynamically built multiline column headings in all list views throughout the application! I will always use the least number of lines as possible if multilines are not necessary, but will use up to 3 lines for long descriptions. This obviously is a big usability improvement when you are looking at many fields with long descriptions but narrow data content in terms of screen space. Also changed the default behavior when displaying a column where some of the data is not visible to use a > character instead of 3 dots ... This was done to conserve screen space and improve usability.

Added several (about 8) new intervalized graphs. These graphs show the same type of data that the summary graphs currently show, except it is for a specific job/thread and the data is shown as rates or counts for each interval in the Job Watch. All the intervalized graphs also

have menu options to look at the call stack or the interval details window (shows bucket count, times, wait object, etc for current interval).

Added the ability to find the next or previous occurrence of a specific string (part of an offset, program name, proc name, or module name) out of a job's call stack, within the Call Stack Window.

I changed the title of the iDoctor-supplied Job Watcher graphs so the first word is either "Summarized" or "Intervalized" as appropriate. Summary graphs summarize each thread in the job watch over the last N intervals.

I enabled support to modify the query and graph definition behind the intervalized graphs. This may lead to some interesting possibilities like hiding certain fields (bars) in the graph or changing the sort or filter for any of the intervalized graphs. You could also do things like produce graphs only when a certain wait object occurred or whatever is desired. You can't save these modified graphs yet however. That is something I'll be working on for the future, so you can apply your custom graphs to any Job Watch on your system.

Fixed some issues with auto refresh of Job Watcher graphs to improve performance of the client. Fixed a bug where the auto refresh timer would not factor in the time it actually took to run the query (sometimes 3 to 5 seconds) which could cause it to refresh too quickly and cause more sluggish response to the user. Auto refresh will also now only happen for the active Job watch graph or table view that has the current focus.

Added averages times to the detail reports 'Run/wait signature counts' and 'Run/wait signature times'. Added the proper bucket descriptions to these reports when viewed via the Interval Details window.

Fixed some problems with selecting jobs off a job queue in the Start Job Watch Wizard.

May 1, 2002 - C00110

On the server side minor fixes were made and added the ability to reduce/eliminate invocation stack harvesting. This allows a higher sampling rate if desired. The client has not yet added the option for this feature in the Create Job Watch Wizard but it will be available soon.

Fixed a problem where if a user lacks authority to run Job Watcher, they would not be prompted to try to sign on with a different ID.

April 19, 2002 - C00108

Added some basic options for library folders so they are available in all components.

Improved the general page of the library properties to provide additional information like create object authority. Also improved the speed of the calculate button that compute the object size. Now it comes out of an API instead of running a command to an outfile and then doing a query.

Fixed a bug on systems without feature code 2924 (US English) installed. Previously when starting a database connection the error "The database server returned a general warning" would occur, and no queries would work.

Fixed a couple of small setup wizard issues.

March 27, 2002 - C00105

The Job Watcher PTFs are now automatically installed by the setup wizard.

Added an interface in the Create Job Watch Wizard, to view performance statistics for all jobs on the system. This can be helpful for seeing what types of waits are happening for all jobs on the system before even starting a Job Watch. With this interface you can either 'restart statistics' to begin a new delta computation between two complete system snapshots -or- use 'refresh list' to rebuild the statistics based on when the last 'restart statistics' was done. Note: For unknown (to be determined) reasons this query will cause a "low-level SQL system error" after being executed approximately 45 times within the same session. If this happens shut down iDoctor and restart. We are working with database to get a DB PTF available to fix this problem.

The auto refresh of graphs and table views in the Data Viewer was not working properly. This has been fixed.

Fixed a first time sign on error that would happen when connecting to a system via an iDoctor.exe shortcut instead of through Ops Nav.

Added a call to *AFTERIPL option in the WCHJOB command at startup of Job Watcher. This ensures that the 32-wait buckets information in LIC are properly initialized and ready for use.

March 6, 2002 - C00101

Added an interface and code for changing the user's password via iDoctor if it has expired.

Fixed a bug with the call stack window that could cause a crash at a later time, especially if you had multiple call stack windows open. The call stack window is also much faster now (removed some unnecessary code).

March 5, 2002 - C00100

Added an 'Interval Details' menu available in the job run/wait signature graph. This brings up a window that lets you see all the details from the QPYRTJWA file for a particular interval/bar in the graph (like holders, objects waited on, basic job statistics, etc).

March 4, 2002 - C00099

Fixed some bugs in the trigger support. A temporary copy of the trigger file was being left in the user's library (file name started with T1). When trying to reuse an existing trigger it would not work (was replacing existing one with an invalid trigger).

March 1, 2002 - C00098

Fixed bug where the job/run wait signature graph would not show new intervals on a refresh for an active job watch job. (I was not recomputing total number of bars)

Fixed a bug in the Data Viewer status bar that shows the total number of bars per graph where it could sometimes get erroneously set to "Bars 1 - 1 of 1". This would happen when the automatic refresh for a job watcher graph was turned on and multiple graphs were open and one in the background was doing a refresh.

To decrease excessive load on the client and fix the problem above, auto refreshes in the data viewer of Job Watcher graphs/reports will now only happen for the view that has the current focus.

Added a detailed set of menus for a 'job watch' folder that let you open ANY of the summary reports/graphs without drilling down through the various report folders.

In query definition, group by fields will now properly parse from an SQL statement into the interface.

To prevent syntax errors, SQL field expressions in query definition are no longer present as fields to select from on the record selection and group by pages

Added a Job Watcher version management wizard to help get rid of old files and prevent errors. It runs after a new installation (from the current PC) or using the 'Manage incompatible Job Watches' menu from the 'Job Watcher' icon. This wizard identifies all libraries on the system that contain Job watches incompatible with the current version installed. These should be deleted or at the very least do not create new Job Watches within these libraries as syntax errors will occur. These deletes are generic file deletes (QPYRTJW*) using the Remote command window in a separate thread. The benefit of this approach is you can do other things in the GUI while the deletes are taking place, the draw back is in order to tell if any of the files didn't get deleted, you need to look at this window and check for any errors.

Added an option in the Job Watcher preferences page (Edit -> Preferences, Job Watcher tab) that lets you determine if new reports should be opened in an existing data viewer (default) or in a brand new data viewer.

February 15, 2002 - C00096

Added support in the database processing code to handle CCSID(13488) fields (a.k.a. CWBDB_PCGRAPHIC fields).

Fixed a bug in the connection list view where an extra set of unused column headers was being added to the list.

February 13, 2002 - C00095

Fixed a bug in the code that keeps track of the user spaces containing field definitions in QTEMP. This code was not properly qualifying per system, so if you ran a query on 2 different systems over files in identical libraries/file names/record format names users would get a non-fatal error like "user space T141414401 in QTEMP not found".

February 1, 2002 - C00093

The query definition member selection and field selection pages have changed significantly. Member selection page shows all active overrides for the query, and you can click on the member in the list to bring up a drop down list to change it.

On field selection page I am trying to make it more useable by allowing the freedom to enter any field expression, and field names desired, by directly editing the field expression, field name or field description right into the list. If a user changes field descriptions they will reflect in the query being viewed, but do not get saved (as of right now). I got rid of old interface at the top because it wasn't very useable to advanced users. The field selection page shows all fields from the all files in the query that are in the top level SELECT field list from the SQL statement.

January 21, 2002 - C00088

Added vertical scrolling to horizontal stacked bar graphs. Fixed some presentation problems where the data was displayed in reverse sort order per page.

Fixed menu for summary graphs so 'open table(s)' is no longer an option.

Fixed bug in the summary tables where the bucket descriptions were not shown in the list. However, when viewing the select fields window the bucket descriptions are still not there yet.

January 10, 2002 - C00086

Added the graph for detailed queuing bucket statistics (32 wait buckets). Added support for viewing the graphs by job or by job/thread.

Fixed a crash in the graph properties for the job watcher summary graphs. Changed the look of graph properties slightly to make it more readable.

Made some fixes to handle older watches. If a watch has a prior version user space status, then the Job Watch will show a status of 'unusable' and the only thing you could do is delete it or view its basic properties/job log. A watch is also unusable if no status user space is available, and the watch is not running and not on a jobq.

Cleaned up the job watch properties 'general' page. Modifying a job watch description is no longer possible. Modifying interval duration is no longer possible.

Some of the summary graphs now default to being shown horizontally. This is not yet configurable (between vert/horz) but will be in the future.

For the summary reports (tables) made the following changes: Added rates for socket fields, Added rates for numeric exception fields, Added rates for state transition fields, Added rates for DASD pages fields, Combined DASD waits with I/O activity report

Added queuing bucket descriptions to current activity queuing bucket reports.

Improved graph properties on the job watcher summary graphs so you don't see SQL expressions for field descriptions (more work to do here though)

Reworked the organization of the sections in a job watch such that we now have the option to view the same reports grouped by job instead of job/thread. Within the reports by job there is also a report for activation groups (last interval) only. Expanding the act grp report shows the pgms within the act grp for that interval. You can also view the detailed object properties for the program object (code from Object Explorer).

November 21, 2001 - C00076

There is a new Job Watcher wizard to replace the Create Job Watch dialog, but it does not have trigger support yet.

Fixed a bug where job watches started from a job on a job queue would not be visible in the Job Watcher component unless the job started at interval 1 (highly unlikely).

Fixed a crashing bug in the Job /Run wait signature graph that would happen when doing a refresh of the graph.

Added the ability to browse files/members within the data viewer via the **Open File Dialog**. You can specify a generic file name like QPY* or whatever and see the results in a list.

Improved performance of all API calls by setting the auto extend flag of the temporary user space using chg user space attributes API instead of pre allocating the user space size. Silly me.

All table views should open about 25% faster now based on misc improvements.

Started adding tool tips to various controls in the application. Added tool tips to input controls in the setup wizard.

When creating output files from the data viewer to the PC, the end of line characters were not being properly written out to the file.

Fixed a problem when it was not possible to create an output file on the PC greater than 32765 records.

After many request I have finally changed the behavior of the main application window and data viewer windows so that the main window is no longer the "parent" window and whenever you click on any of these window it becomes the top most visible window and all other windows are moved behind it.

Fixed a problem in the install wizard where some user settings were being overwritten after doing an update.

Put in a special check to see if a logical file is queriable via SQL prior to trying to open it so the error message is better. The two cases that would cause an error when trying to query in SQL are if the logical file is an SQL index, or if the logical file contains more than one record format.

October 4, 2001 - C00062 (Client Update)

Added a "show job log" button on the Job Watch Properties window. This will show you the job log of an *ACTIVE or *JOBQ batch mode job watch, or it will show the spool file if it exists for a batch mode Job Watch in *OUTQ status. If the spool file or job log is not available then this button will not be shown.

Added the display of Job Watches that have been started but aren't running because they are sitting on the job queue.

Disabled menus and options as appropriate for the new batch mode Job Watch status types (for job status of *OUTQ, *JOBQ, *ACTIVE).

Added job queue name, library and subsystem name to the list of jobs sitting on job queues on the create job watch window.

Changed internal use of list jobs API to return the routing data instead of making a separate call to the retrieve job information API for each job.

Fixed a refresh problem that would only occur when stopping a Job Watch that is the first job watch in the list.

Fixed the "write interval now" menu to only be displayed for job watches that are executing from the client (not in batch mode).

Cleaned up several problems with the Job Watch Properties window.

October 2, 2001 - C00061 (Client Update)

When installing Job Watcher added a page for defining the Job Queue and Subsystem description needed for running the Job Watch batch jobs.

Job Watches can now be executed as a batch job on the server. This provided several advantages.

Redesigned the Create Job Watch window.

Fixed problem with Job Run/Wait Signature graph where the percentage values could display over 100% because DELTAUS was being used as the total instead of QTIME01+QTIME02+QTIME03.

Switched to displaying the end of interval time stamp fields instead of the start of interval timestamp fields.

Added several additional checks and appropriate messages when creating a Job Watch. SBMJOB failed check, duplicate watch name check, job watcher environment data area not found.

Fixed problems with the auto refresh and always display new data options (user preferences window and Data viewer view menu options).

Fixed install problem where the initial library list contained an invalid library.

September 25, 2001 - C00059 (Client Update)

Added Job Watcher help links to all interfaces and all reports (via new help menu on report popup menu).

Added a Job Watcher preferences page to the Preferences Window.

Fixed a problem when opening a job watch report if the list of reports is sorted in an order other than the default. The wrong report was being opened in this case.

Fixed a PEX Analyzer dependency bug because QYPECLLNC1 was being called in SMTRACE to get a library description. I am now using the retrieve object description API to get the library description.

September 21, 2001 - C00058 (Client Update)

Fixed a problem with the ending conflict information report prompting for a member when it should not have been.

Fixed a problem with the ending conflict information report's option to create a job watch over the holding tde.

August 23, 2001 - C00050 (Client Update)

Changed all instances of objects and libraries named QYP* to QPY*

Fixed a problem with the Stack View sometimes containing bad data towards the bottom of the stack.

Disabled sorting of a stack view by clicking column headings. This was meaningless.

July 30, 2001 - C00047 (Client Only Update)

Added a check in the setup wizard when installing job watcher for the special ptf. This is done by attempting to call wchjob to build the TDELIST as discussed with Jay.

Added a refresh button on the Create Job Watch dialog.

Fixed a problem where a freeze would happen if more than one client tried to open the list of jobs at a time. This is because the TDE list file was being incorrectly built in library QYPRTJW instead of library QTEMP in the user's job.

July 25, 2001 - C00046 (Client Update)

This update enhances how the queries are saved and loaded into the client.

July 23, 2001 - C00045 (Client Update)

Fixed a couple of problems in the custom graph definition support. (Labels not being updated on the graph and hover help labels also not being updated).

July 19, 2001 - C00043 (Client Update)

Added a group by panel on query definition. Right now the user must know that you need to use column functions over fields in the field list in order to make a group by query. Haven't figured out yet how to make it easier for an SQL new user.

Graph definitions are now tied to the user profile. This means that when viewing a table you can only work with graph definitions previously created with the same user profile you are currently connected to the system with.

Improved the performance of the member selection page of query definition.

Added the ability to start a job watch over a holding TDE within the End of interval snapshot view. Right click on the table and choose the analyze menu to get to this feature.

Fixed a crashing bug that would happen if you try to stop more than one watch at once. I have changed this so only one watch can be stopped at a time (by menu invocation).

Fixed a bug where if you right click on a library in the list and not the tree and choose the create job watch menu, the create job watch dialog would not properly create a job watch.

Fixed a bug where the delete of a job watch would not refresh the list or tree view. Also modified the delete so it takes place in another thread instead of tying up the GUI thread.

July 11, 2001 - V5R1 S00025-C00039, V4R5 S00022-C00039, V4R4 S00022-C00039

Added the ability to define custom graph definitions, and save/restore them from the server. Use the new Graph Definition menu over a table view to access this feature.

Added a window to delete graph templates or open a graph template. This window is accessible by a new Graph Definition -> Manage List... menu for a table view.

Added the ability to change a graph definition, without closing down and reopening the graph.

Field names are now retrieved via OS/400 APIs rather than through the catalog file in QSYS/QADBILLB via SQL ODBC.

The OS/400 object APIs are now used to build the list of libraries instead of a combination of DSPOBJD to an output file and then reading it with ODBC. The performance improvement is considerable.

Made some user interface changes to the Output Preferences Dialog. A checkbox has been added which lets the user either turn on/off the setting to include column headings in copy to clipboard data.

Improved the performance slightly when opening up a new tree/list component window in the iDoctor for iSeries main window.

The title of a Data Viewer window now contains the system it is associated with. A data viewer can contain table or graph views from multiple systems but the title will match the system the data viewer is tied to. This system name determines what system to connect to when doing features like Open Custom File, Interactive SQL, or the Query List features.

Added the ability to view multiple records in the Record Quick View dialog. You can also select up to 30 records in a list, right click and choose the 'Record Quick View' menu to view the selected records vertically instead of horizontally.

Added decorative animations to the following screens: Login, Application Properties, and the Analysis Wizard

Implemented a disable-caching panel in the Application Properties page. Right click on the component icon and choose properties to get to this. By checking the disable caching option caching will not be used and every time a selection changes in the tree, data will be requested from the server.

Added the ability to type in an SQL where clause expression on the record selection page of query definition. I've done this the quick and easy way with a toggle from "value" to "expression". If the state is "value" then all of the existing syntax checking for a single numeric or string value takes place. If the state is "expression" then no syntax checking takes place and it's totally up to the user to enter a correct SQL where clause expression.

When the user actually clicks OK on the query definition page would the query be attempted to be executed as is and they would get the standard SQL syntax error from the server.

Previously if you typed a query such as **select annbr from smtrace/qa440anlt** from the interactive query view, and then went into query definition, you could only work with the fields in the field list (in this case annbr). I have made a change so that you can interactively work with all other fields in the table(s) add/hide, filter, sort, not just the fields in the select statements field list.

Fixed a problem that would occur when viewing graphs on a PC with Client Access V5R1. Graphs containing data where every point value was 0 would sometimes crash the application.

Fixed a couple problems when saving the entire contents of a table view to a file.

Fixed the support for minimum VRM checking on the server to include a logical check for the build VRM AND the build number. The build VRM was not being checked.

June 25, 2001 - C00026 (Client Only Update)

Building a list of job watches in a library is now faster because of the use of OS/400 APIs to build the lists.

Fixed the nasty crash when stopping the 2nd job watch.

I removed the size field for a job watch. This field was not accurate because a job watch consists of multiple files, and this value only reflected the member size within file QYPRTJWA.

June 15, 2001 - C00024 (Client Only Update)

Implemented authorization checking for job watcher on the client side.

June 6, 2001 - C00018 (Client Only Update)

Added a check for V5R1 OS/400 when attempting to start a job watcher session.

Reworked the stack view to only contain one column for each interval stack instead of 2.

Added a menu so you can open the stack for a bar in the job watch graph view.

June 6, 2001 - C00017 (Client Only Update)

Renamed the table Most Recent Wait Info (as of the end of each interval) to End of interval snapshot.

Added the fields CURRUP, RECCNFLCT to the end of interval snapshot query.

Added a green pulse-like animation on the create job watch dialog

Fixed a crash that would happen if trying to open an empty graph.

Added a friendly error displayed to the user if they try to create a watch without specifying a name.

May 29, 2001 - V1R1M0 C00015 (Client Update only)

Improved notification messages in the status bar while list of things are being created.

Added dialog icons on the delete collection and delete analysis dialogs.

Fixed a problem with the tree/list view that was sometimes causing the tree to behave strangely (adding duplicate items) after an item in the list was expanded.

Made a change/(usability fix) to the tree/list to always highlight the current selection in the tree even when the tree does not have focus. Also changed the tree to not expand every time a tree item is selected for better usability with potentially large numbers of libraries, collections, etc.

May 22, 2001 - V1R1M0 C00013 (Client Update only)

Support for basic Job Watcher functionality added. This includes the ability to select up to 10 jobs or tasks to monitor and then look at various reports over the real time data collected from these jobs.