

Update History for iDoctor for iSeries - PEX Analyzer for OS/400 V5R1

Feb 12, 2007 – C00520

Removed the Total processors in LPAR field from the System tab of collection properties. This field had the same value of 'max processors during collection' so was not needed.

Removed an annoying member selection popup message that would come up when attempting to view the collection properties of an active PEX collection.

Added a new property page for completed collections showing the list of files in the collection and the number of records they contain. This also displays any PEX files that do not exist in the collection. This is most useful for releases prior to V5R4 when members were created for every PEX file in order to determine if the collection is missing any required files for the desired analysis.

In the Add PEX Definition Wizard, on the job selection panel, fixed a bug where selecting a job of *ALL/NAME/*ALL was not allowed for Stats Hier. This has been changed to only disallowing the selection of *ALL/*ALL/*ALL for Stats Hier.

The context menu for active collections no longer has a greyed out suspend option. The suspend function of ENDPEX is not supported by PEX Analyzer.

The PEX Analyzer popup menu for a table view, should now look a lot more like the Job Watcher popup menu for a table view. Some of the menus were in different positions, such as the set font and preferences menus.

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

The client checks the number of events to determine if the status "contains no events" should be shown. Updated this check so the status is not shown if the number of events reported in the PEX DB file QAYPERUNI s < 0. Also update the # of events field displayed if the events reported in QAYPERUNI is < 0. This happens if the total events collected by PEX is between 2147483647 and 4294967296. This change to the client allows the client to display the correct number of events in the PEX collection up to 4294967296. This field will wrap if more than 4294967296 are collected.

Fix a tab order problem in the transfer collection window.

Sept 19, 2006 – C00502

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.

PEX collections that are listed as incomplete or unusable can now be deleted from the GUI (via delete menu).

Add the ability to schedule a PEX collection.

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

Fix a crash in the PEX Analyzer install that occurred when PT1 was not installed on the server.

Add the ability to graph the Event counts by job file in the Events analysis in PEX Analyzer.

Disable the query definition toolbar button in the Data Viewer for PEX Analyzer 'flattened' graphs to prevent problems if the user changes the query.

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.

Replace references in PEX Analyzer call stacks in the GUI to "PASE" with "LIC"

November 8, 2005 – C00459

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

PEX Analyzer graph views that show dynamic numbers of colors are showing all black colors after the 1st 14 colors. Generate the colors randomly to ensure a better distribution of colors.

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

Add PTF checking to the STRIDOCOL/STRPACOL commands at V5R1 and up, and update the GUI to perform this same checking before running a PEX collection.

The PTFs currently checked for at release V5R1 are: MF33253 MF31783 MF30692 MF33253 MF34625 MF34455 TC03147

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.

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.

July 15th, 2005 – C00423

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

New commands ADDIDRUSR and RMVIDRUSR at release V5R1 have been added to library QYPBASE. The commands will 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 PEX Analyzer added to the GUI before starting a collection. At V5R3 this check will be for *SERVICE OR access to function group QIBM_SERVICE_TRACE. At V5R2, V5R1, this check will be for *SERVICE This check has also been added to the STRIDOCOL command.

In the PEX Analyzer analysis wizard:

Added a check to see if a job has been selected on the subset job/task page. If a job or task has not been selected, the selection window will be shown after clicking Next.

When selecting a job from the Browse jobs window, the thread id is not being properly displayed on the subset job/task page of the wizard.

When using the menu option for a library "Delele all PEX analyses" the file QAYPEANAL (V5R1 analysis tracking file) should be deleted. When using the menu option for a library "Delele all PEX collections and analyses" the file PF150 should also be deleted.

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

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

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

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

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

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

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

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

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

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

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

May 20th, 2005 – C00411

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

May 16th, 2005 – S00076-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.

March 28th, 2005 – S00071-C00391

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

The PEX Analyzer version management wizard was not working because of the progress bar not being initialized.

Opening a table and a graph at the same time (using the open table and graph menu) and then going into the table and viewing the property pages will cause a crash.

When displaying a list of collections in a library, the TYPE column does not display "User-defined-xxx" for collections that were user-defined. The client is incorrectly displaying "PEX Analyzer supplied-xxx" for the TYPE column

After deleting one or more user-defined graphs, a refresh of the list of graphs does not occur. Also the delete of a user-defined query has a prompt that is not the standard for deletes.

The netsize analysis graphs should now always be graphed with side-by-side vertical bars

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

January 13th, 2005 – C00367

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

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

December 20th, 2004 – C00366

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

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

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

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

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

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

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

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

September 13th, 2004 - C00345

Creating a PEX collection from the GUI using a user-defined PEX definition caused the job selection within the PEX definition to be changed incorrectly to the job that submitted the collection.

August 16th, 2004 - C00342

Combined the interfaces for Record quick view and properties into one set of property pages. Additional panels such as the Call stack panel and Procedure information panel will display when available for the analysis data being viewed.

System-scoped user defined queries and graphs are now stored in library QUSRSYS rather than QYPBASE.

Removed the analysis categories from the PEX Analyzer Analyze data "fast path" menu and now just list the names of the analyses available instead of having another level of submenus showing each analysis category. This should simplify the process of running default analyses since there are usually only 5 or 6 choices possible for a given collection.

Many programs and objects renamed to be more consistent with object naming standards.

New support allowing generic task names to be entered in the PEX collection and definition wizards. This also affects the STRIDOCOL command.

Selecting the New field button on the Query definition -> Field selection window will now add the new field after the first selected field in the table (if any).

Improved the display of 'unreadable hex fields' in the GUI.

Fixed a bug with copy to clipboard from the Record quick view window.

Added a check to prevent users from trying to open multiple PEX Analyzer output file reports as graphs if some of them are not graphable. The Open Graphs menu should only be enabled if every selected report in the list is graphable.

Added a check on the graph definition - primary Y-axis page to prevent a user from adding the same field to the graph more than once

February 5th, 2004 - C00288

Added a new "iDoctor 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.

Fixed a couple of job/task selection issues in the Analysis Wizard.

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

On the summary page of the PEX collection wizard, changed the label "Collection system information: " to "Collect system information:"

Fixed a problem when deleting an analysis. The hourglass goes away too soon (before the delete actually completes).

Fixed the behavior of collections that have ended in error. Added support for a new icon that indicates that the collection is in error vs. a normal collection. These collections can be deleted and properties are available (in order to see the job log from the general page).

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

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 when changing a PEX Stats or Profile definition that contains 1 or more tasks. On the job/task selection page the task selection option will be incorrectly defaulted to *NONE.

Enhanced usability of the Data Area/data queue reports. Added flyover help for some of the fields with values like RDAL, CDAL to spell out what the acronyms mean.

Fixed a bug with single object PEX collections not being visible until they have completed collecting.

December 17th, 2003 - C00281

Added the ability to stop a running PEX collection from the client. This is the *STOP option on the ENDDOCCOL command.

Fixed a bug where a phantom iDoctor for iSeries icon was showing up in iSeries Navigator. Launching iDoctor from iSeries Navigator is no longer available.

Fixed some problems with collections not showing the proper status if the job running the collection ended in error.

Added the option to select tasks in the PEX Collection Wizard and PEX Definition Wizard when creating a Stats type collection.

Added the 'duration of trace(us)' field to the general page of collection properties. This is the total time of collection in microseconds. Unlike PRTPEXPRT this value does not include the time it took to end the collection and produce the output files.

All static fields in collection properties are now selectable for copy and paste purposes.

Fixed a problem with system-scoped user defined queries created from the server-side output files folder within a PEX analysis.

November 19th, 2003 - C00272

Improved validation checking in the PEX Collection Wizard and PEX Definition Wizard.

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

Added a warning message to the install if the system value QALWOBJRST is not set to *ALL or *ALWSYSSTT. 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.

Improved the usability of the program bracketing events options for creating PEX stats definitions in the PEX Definition Wizard.

The bottom of the summary page on the PEX Collection, PEX Analysis Wizards will now show the CL command string that the client will issue.

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

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 PEX Analyzer icon.

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

Fixed a bug with the delete collection function.

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.

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

May 20th, 2003 - C00225

Fixed several issues with CPYIDOCOL, SAVIDOCOL, RSTIDOCOL, and CVTIDOCOL

Fixed a problem on the advanced options window of the Start Collection Wizard with parameter validation of the filter name parameter.

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'

Fixed a bug in PEX Analyzer where the creation settings properties panel did not show up for single object collection.

CPYPEXDTA - includes fix for bug where if source and target library are the same but from and to member names are different

May 8th, 2003 - C00219

Added new PEX analyzer preferences to indicate the default job queue name and library to use when starting a job watch. The selected values are saved on the client and will be preloaded into Start PEX Collection Wizard, every time it is used.

Added a new folder 'server-side output files' within each analysis in PEX Analyzer. This is essentially showing all files for the collection (the QAYPE* and SMTR* files).

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 for example) should be noticeably faster.

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

Fixed a bug in the PEX Analyzer properties window where certain panels were visible for active running analyses that should not have been visible.

The icon for a user-defined graph is now a graph icon instead of a query icon.

Made a change so that output files are not listed if the appropriate file/member no longer exists. This will prevent users trying to open files that don't exist and getting an error.

Fixed a bug in PEX Analyzer where the creation settings panel in collection properties only showed up if the job that ran the collection was still present on the system.

Added a new 'size' panel on the collection properties and analysis properties windows. This provides information on the disk space used for collections and analyses.

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.

Fixed a problem with the QYPINT/CPYPEXDTA command whereby an error occurs when the from and to collection member names are different and the from and to library names are different. An unhandled CPF9815 escape message appeared as a result of this error

April 9, 2003 - S00009 - C00196

Fixes the 'decimal data error' problem that occurs under GTPROF when ran in batch.

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.

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 bug where the report description was not being shown correctly for the 'query selected' queries in Task Switch.

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.

When viewing the Call stack for a PDIO report, added text to the list that says, "Call stack not found" if the call stack could not be located.

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 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.

Fixed bug where the child analysis member did not display in the list of child analyses.

Fixed a caching problem with the 'query selected' popup menu where it only remembered the first type of selection made (1 row selected vs. 2 or more rows selected), since the menu will display differently (in some cases) depending on the number of rows selected.

Fixed a caching problem with the 'analyze' popup menu where it only remembered the first type of selection made (1 row selected vs. 2 or more rows selected), since the menu will display differently (in some cases) depending on the number of rows selected.

February 10, 2003 - C00187

Improved the user-defined queries and user-defined graphs support.

Several new "query selected" menu options are available within the Task Switch analysis. These let you run queries based on the data you select to drill down into the reports. Also added the ability to select a range of records as input to the query selected option. This option is used in some of the reports to indicate a start and end time range to apply to the query.

Fixed bug in the PEX definition wizard at V4R5. Stats hier mode does not support the Program Bracketing events parameter.

Fixed bug in the PEX definition wizard at V4R5. Stats hier mode does not support *ALL/*ALL for the Jobs parameter. Added an error message on the job selection page.

Fixed a hang that would happen when trying to open a PEX definition that no longer exists (was deleted after the client was last refreshed).

Fixed a bug on the TPROF options page with the Filter percentage input validation checking. If invalid input was provided the error message popped up more than once.

Fixed a bug in collection properties on the collection jobs and collection tasks pages that made it impossible to sort the list of jobs or tasks by TDE ID.

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

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.

Fixed some usability problems with the time subsetting control in the Analysis Wizard.

November 30, 2002 - C00174

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 in the PEX definition wizard that happened when the user selected no trace types events and some individual events. CALLRTN trace type events were being included when the definition was created erroneously.

Suppressed errors by the QRYPEX command that would show in the GUI when it function checks. This happens when the state of the PEX collection is in transition and the command doesn't handle it.

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

Added support for child analyses over SWAFD records (field OPR) within file SMTRMODL9 (part of the task switch analysis).

The timestamp parameters are now editable on the child analysis prompter window when running a task switch child analysis of type 'TDE list' or 'TDE summary'.

Fixed some problems with the PDIO call stack support. Some offset used in the client were bad causing the data to show out of alignment.

The Stats call level indent GUI option works properly now.

October 30, 2002 - C00166

Fixed a bug when creating a PEX collection via the GUI or STRIDOCOL where program GETSTATS 1 on the server side would fail with CPF0819 (Variable or substring of variable too small to hold result) when the number of threads to be used on ENDPEX time was not start *CALC.

Made a small usability fix in the PEX definition wizard on the trace options page.

Fixed a bug where STRIDOCOL was failing when using an existing PEX definition.

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 timeslice 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.

Fixed a bug where PEXcollections created in the GUI could not have a space in the description or you'd get an error when trying to create it.

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.

Now caching the information for the Analyze and Query popup menus for task switch and stats hier analyses to get rid of a slight delay when right-clicking on a table view. The delay will still be there the first time a report is right-clicked because the client needs to retrieve the information.

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).

Fixed a bug where the queries behind user-defined graphs in PEX Analyzer were not being saved properly.

October 17, 2002 - C00161

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

Fixed bug with the user-defined queries in PEX Analyzer where they sometimes could not be opened.

October 8, 2002 - C00158

Added a new PEX Analyzer Trace PEX Events analysis. This analysis provides a nice breakdown of the number of each type of PEX event that occurred in the collection per thread.

On the Application properties - server configuration page added a details button next to the job queue name that lets you see all the job queue properties - (same info as for a JOBQ object in Object Explorer). Also added a field that shows whether or not the subsystem is active or not.

Added to documentation information on how to collect Websphere PEX events so PEX Analyzer can analyze them.

Added a popup menu item "Display call stack" when within a Physical Disk I/O analysis output file that contains the STCKIDX field and a single record is selected.

Added the ability to FTP PEX collections (single object format or PEX DB files format) from one server to another via the client.

Added several new green screen commands to simplify working with PEX collections. These commands are used via the GUI triggered by the appropriate menu option. The commands are: DLTIDOCOL (delete collection), CVTIDOCOL (convert single object file to PEX DB files), SAVIDOCOL (save a collection to a save file), RSTIDOCOL (restore a collection saved via SAVIDOCOL from a save file). All of these commands work with either the single object or pex db files type of collection.

July 31, 2002 - C00140

Made a change so PEX Assistant will work if the system is authorized for PEX Analyzer.

Enabled the 'Copy' and 'FTP' functions for PEX collections into PEX Assistant.

Fixed a bug in the PEX definition wizard, where certain types of PEX profile definitions were not being retrieved properly. Also fixed a bug in the PEX definition and PEX collection wizards where job selection was not available on type PROFILE definitions.

Greatly improved the PEX definition wizard and PEX Collection Wizard interfaces for defining the programs to profile in a PEX Profile definition. You can now choose to either browse for pgms within libraries, modules within pgm, and procedures within modules, or simply just type in the desired values. You can also leave the add window open and select any existing value from the program selection list and the selected values will be plugged into the add pigmy window making it easier to make changes.

On any analysis that allows job/task subsetting you may now select multiple jobs or tasks which will submit a separate analysis for each job/task selected. The following conditions apply:

- For Trace analyses you cannot select jobs and tasks to be submitted within the same wizard session. You can either select as many jobs as you want or as many tasks as you want. When selecting multiple jobs the thread type setting (to include all threads within primary thread or just a specific thread) will apply to all threads selected. You cannot have a mix of different thread type values within the same wizard session.
- The rest of the subset criteria will be identical for each analysis submitted.
- You can no longer filter the browse jobs or browse tasks window by generic job name. Sort and Find functionality are built into these views. It complicated the interface and my code too much to keep browse in during the case when multiples were already selected.

This update adds many enhancements/changes to the child analysis support. There are new child analysis interfaces for Stats Hierarchical and Task Switch Trace analyses.

- Child analyses are now located in a special folder underneath the (parent) analysis called "Child analysis reports." This new folder allows the user to view the status of any batch analyses in progress, or view, delete any finished analyses. Child analyses allow you to drill-down deeper into a particular piece of data (such as in stats get a report only for a specific program). A child analysis consists of just one output file/member in the users library.

- Child analyses are created via the GUI using the Analyze popup menu found by right clicking on a table view from an analysis that supports child analyses. If the 'Analyze' menu is not shown then the particular report you are looking at does not provide any child analyses.

- Note: The Stats child analysis wizard has been removed in favor of the new interface.

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.

Added the 'total events' field to the events property page of collection properties.

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.

Enabled the Query Selected menu (pre-defined queries) within the Data Viewer for Task Switch analyses at V4R5 and higher.

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.

June 10, 2002 - C00124

Fixed a bug where the correct number of lines for a multiline header were sometimes not being shown.

When displaying a list of libraries it was possible to see a library that did not exist if the following scenario occurred.

- Create a PEX collection via the GUI - this creates an IDOCCOL job with routing data identifying the collection and library name.
- Delete the library the collection was created in.
- Assuming the IDOCCOL job still exists on the system, the library was shown at the end of the list which is caused by reading the routing data of all IDOCCOL jobs on the system to determine the appropriate additional libraries to display.

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.

Changes to the SQL Query View

- Disabled the execution of SQL statements other than SQL SELECT statements via this view.
- Fixed a bug where the user was not being prompted for a different member if the SQL (library/file) changed and the previous member supplied was no longer valid.
- Improved speed of the member selection window that is displayed via the 'SQL Query View'. Decreased the font of the member selection dialog to be consistent with rest of app. Added icons to the dialog.

Changes to the Graph details window

- Improved the look/formatting such that it is consistent regardless if you are looking at a user-defined or an iDoctor-supplied graph's details.
- When viewing the graph details for a user-defined graph all fields in the associated DB record are now shown as "other" fields. Previously only the fields for the X, Y1, Y2 axes were shown.

Fixed a bug where a library could not be renamed.

On the server side fixed a corrupt file G_DBCFP04.

In the PEX Definition and PEX Collection wizards changed the Add Jobs window to initially not bring up a list of all jobs on the system. You can set up generic name to configure which jobs are displayed or if desired type in your selections and use the 'Add' button to add them to the job selection page in the wizard without needing to take the time to build the list of all jobs on the system.

In the PEX Collection Wizard, added verification of the collection name and library name parameters on the Options page. Also added checks to make sure an attempt is not being made to create a PEX collection into a library already containing PEX collections of a prior release.

Fixed a bug the Create PEX Collection and Create PEX Definition wizards, where any add jobs/add tasks/add programs popup window could sometimes not be reopened after the initial time it is opened in the same wizard session. Also fixed a related bug with the open file dialog in the Data Viewer.

Changes to the Query Definition - Record Selection window.

Added the ability to define record selection filter criteria over SQL field expressions you've already created.

All operator types are now enabled when the selected field is a custom created field (an SQL expression), i.e. It's up to the user to make sure he/she doesn't try to use the wrong operator based on the type of expression the field is for.

May 1, 2002 - C00110

Added the ability to FTP PEX collections to other systems. Use the Transfer to... menu after selecting one or more PEX collections to access this feature.

Added an interface for copying one or more PEX Collections. Use the Copy... menu after selecting one or more PEX collections.

In the Analysis Wizard, time subsetting is enabled again for Trace-based analyses.

Job selection in the analysis wizard now shows an appropriate error message if you try to do a generic job user or generic job number search via the Browse button. The error msg will suggest doing a generic job name search and then clicking the browse button, which is the only generic search method I support. You can specify a specific job user or job number if desired. Of course you can still sort by job user and do a 'find' within the Browse Jobs window.

Fixed a bug in the query definition where the case sensitivity of a mixed case string within an SQL statement would be lost on a sort.

Made some slight cosmetic changes to the browse jobs and browse tasks dialog. This was mostly just alignment to make it look better. Fixed bug in browse tasks dialog where it was saying "All jobs" instead of "All tasks" in the label at the top.

Fixed a memory leak that would occur if you opened the browse jobs or browse tasks dialogs in the Analysis Wizard more than once before closing down the analysis wizard.

Added some code to work around an ADDPEXDFN bug in the PEX Definition Wizard. If you try to specify a *Profile definition with the Cpu interval sample less than 1 ms the ADDPEXDFN command will fail. You can specify less than 1 ms for Trace without a problem though. This applies to 510 and higher. A fix was also added on the server side to STRIDOCOL to handle this.

Fixed a crashing bug that would happen if you were executing remote commands in the remote command view and closed the view before all the remote commands finished. The crash would actually occur once all the remote commands submitted had finished. You can hide and show the remote command view at anytime now and it will correctly display which command it is currently working on and its status.

A problem was reported where the user was getting a "not able to start a remote program call connection" right after launching iDoctor from Ops Nav. This was been fixed.

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

On the server side fixed a problem where a residual lock may be left on the SMTRMOD physical file after 'Delete Analysis'

On the server side fixed a bug in QYPINT/CPYPEXDTA command

Added quite a bit of PEX Analyzer documentation for many new interfaces such as the PEX Definition wizard, the PEX Definitions view and properties, and the new dialogs for copying and FTPing collections.

April 19, 2002 - C00108

Fixed a bug in the STRIDOCOL command and in the PEX Collection Wizard when trying to specify a *STATSFLATN collection an error "you cannot specify *ALL jobs" would result even though an individual job was specified.

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

- Copy... Copy a library into a new one.
- Save... Lets you save the library into a save file.
- Rename - lets you rename a library.
- Delete - some components didn't have this option.
- Clear - Delete all the objects in a library.

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.

April 13, 2002

New server-side support released for PEX Analyzer V1R1M1.

New Physical DASD IO Analysis (GPDIO)

This command replaces the old GIO and GIOTIME analysis commands. It produces a set of summary files that can be joined to answer the majority of questions relating to physical DASD operations.

Whenever one of these files stores the low-level DASD operation type of

FT - fault

PO - page out

RA - read asynchronous

RS - read synchronous

WA - write asynchronous

WS - write synchronous

then the minimum, maximum and average IO size and minimum, maximum and average IO time provided.

Program stack by IO type - Only available at OS/400 V5R1

Using the GUI, the user can click on the new STCKINDX (Stack Index) field to show the last 50 MI program/method/procedure names in the call stack at the time the PDIO operation occurred. The program/method/procedure names are only available *IF* PEX TRACE program events of type MI Entry and MI Exit were collected.

The STCKINDX field is available in the following 6 GPDIO summary files:

Output File Name	Output File Description
G_IOTHOZS	IO Total by Job-Thread/Object/IO-Type/Stack-Index
G_IOTHZS	IO Total by Job-Thread/IO-Type/Stack-Index
G_IOTOHZS	IO Total by Object/Job-Thread/IO-Type/Stack-Index
G_IOTOZS	IO Total by Object/IO-Type/Stack-Index
G_IOTVOZS	IO Total by Interval/Object/IO-Type/Stack-Index
G_IOTYZS	IO Total by ObjectSegment-Desc/IO-Type/Stack-Index

This new data is not available in PEX Analyzer 111 at OS/400 V4R5.

New WebSphere Analysis (GWS)

iDoctor PEX Analyzer 111 includes support for WebSphere version 3.5.3. With this analysis you can view all events for the application server and evaluate the overall workflow in regards to events processed. There is also support for database pool events to track the pool usage as connections are allocated and freed. In addition to this there is detailed servlet support as well as servlet summary information that can be obtained for individual calls as well as specific servlets used during this collection. This information can be very useful for identifying bottlenecks and/or potential problems with the application.

The minimum WebSphere requirements are:

OS/400 V5R1:

WebSphere Application Server Version 3.5.3 to 3.5.x
Java Toolbox V5R1M0.

OS/400 V4R5:

WebSphere Application Server Version 3.5.3 to 3.5.x
Java Toolbox V4R5M0 or V5R1M0.

New Copy PEX Data command (CPYPEXDTA)

Copies PEX Data from the collection QAYPE* files from one member to another member. This enables users to easily isolate interesting PEX collection data.

This command will prevent you from copying PEX data into a library already containing PEX data collected at a different OS/400 version/release level.

PEX Analyzer work files and analysis output files are **not** copied.

New verify PEX Collection command (VFYPEX)

You provide an OS/400 version/release and the PEX collection name and library.

When issued from the green screen, this command:

- checks whether the PEX data is at the OS/400 version/release specified
- provides total event count and total missed event count
- provides a record count for all PEX QAYPE* files with >= 1 record

This command can be useful when trying to determine whether a PEX collection is both complete and usable at a particular OS/400 version/release level.

March 27, 2002 - C00105

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

Fixed a problem in the PEX definition wizard when changing a PEX definition that contains individual jobs. The job selection checkbox was not being checked which would lead someone to believe the definition did not contain any specified jobs.

March 13, 2002 - C00102

Whenever an analysis is about to be run or checking for libraries containing obsolete analyses, the file SMTRWEBS is also checked (in addition to SMTRMOD) for incompatible record format IDs with SMTRACE.

Added support for viewing the PEX event call stack from the PDIO analysis files (or any PEX report that contain the STCKINDX field). To view this feature double-click on the STCKINDX field value within the appropriate table view.

March 6, 2002 - C00101

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

March 1, 2002 - C00098

Added several options for managing PEX collections and analyses:

Added a PEX Analyzer version management wizard which lets a user choose to remove old PEX collections and/or obsolete PEX Analyzer analysis files. This wizard runs after an installation or at anytime using the Manage incompatible PEX collections and analyses menu for the PEX Analyzer icon.

Added a menu option to delete all PEX analyses for a library.

Added a menu option to delete all PEX collections and analyses for a library.

All of these deletes above are accomplished using generic file deletes (QAYPE*, SMTR*, G_*) 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 to be deleted didn't get deleted, you need to look at this window and check for any errors.

Added support for copying PEX collections on both the client and server sides. Use the CPYPEXDTA command in QYPINT to move a PEX collection from one library to another -or- use the Copy menu for one or more collections, and then use the Paste menu for a library folder in order to copy the collections to the library.

Added support for starting PEX collections in suspended mode with the Create PEX Collection Wizard. Use the standby mode checkbox on the options page in order to do this. This can be very useful if you want to set up a PEX collection for a test, and then to 'resume'

it right after your test begins. Starting a PEX collection normally can take several seconds. Also added a resume menu for a suspended collection that is only enabled if it was started in standby mode using STRIDOCOL or the GUI Create PEX collection wizard.

Added support for the new STRIDOCOL, iDoctor-supplied PEX definition types in the Create PEX collection wizard. These are the *HEAPSTG, *WEBSPHERE, and *USR_TRNSACT definition types. Note the websphere analysis is not yet available but should be within a few more weeks.

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 were being added to the list.

February 13, 2002 - C00095

Added check when starting PEX Analyzer that the user has the following special authorities: *ALLOBJ, *SPLCTL, *JOBCTL

When viewing an analysis in the client that has iDoctor-supplied queries (queries we ship) there will be two 'report folders' One folder is for the iDoctor-supplied queries, the other is for the analysis output files. If no iDoctor-supplied queries exist for an analysis then this extra layer is not present.

Added support for retrieving descriptions for SQL field expressions in the iDoctor-supplied queries. This is done by defining the field expression name and description in file QAnnnRPTC for each query in file QAnnnRPT.

Added support for the *HEAPSTG definition type in the PEX collection wizard.

Fixed bug for the following scenario: If an analysis is running on system in library x and an analysis has been submitted and is "waiting on the job queue" to be created in library y the analysis waiting on job queue would not show up in library y in the GUI.

Fixed a bug in the code that keeps track of the userspaces 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 "userspace T141414401 in QTEMP not found".

February 1, 2002 - C00093

This client has support for running the new Physical Disk I/O analysis and viewing the SQL based reports.

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, 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, 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.

Field selection page shows all fields from the all files in the query that are in the top-level SELECT field list.

January 21, 2002 - C00088

Changed analysis start time stamps to display in ISO format (to be consistent with elsewhere in the application).

Fixed bug where 'Stats flat merge Yes' collections were shown in the GUI as type 'stats flat merge no'

When deleting a collection the collection job was not being deleted. Since the existence of a collection job will show up in the GUI as a 'no data found' collection (possibly one that failed) I am now deleting the collection job after the delete collection program successfully deletes the associated collection.

After a collection was deleted the refresh was not working right. Cleaned that up.

November 21, 2001 - C00076

A wizard for creating PEX collections has been created. This is an interface over the STRIDOCOL command. There are two modes of operation in the wizard. Taking "Basic" mode lets the user follow a series of questions to determine the "definition type" parameter (problem type in STRIDOCOL). "Advanced" mode skips these questions. The wizard features job, task and program browse and selection capabilities via the GUI interface as well as selecting the additional events via by checking them off in a list.

Improved status of collections in progress by showing the current number of events collected as well as the type of collection (Trace, Stats, Profile). It will also indicate if the collection is stopped or suspended.

Added menu options for collections in progress to end the collection in various ways based on the options on the ENDDOCOL command: *CRTDBFILES, *CRTSVCFILE, *DELETE, *RESTART

Added some additional information to the collection properties general page: Total time, total events, missed events, total jobs, total threads, and total tasks.

For all PEX Analyzer graphs I added code to support showing 2 additional fields per graph in the popup tool tip window that you get by putting your mouse over the bar, (not graph coordinate properties that you get by clicking on a bar and it shows everything). These additional fields will also show in the status bar (as space permits). For example the TPROF graph will show the procedure name in the tool tip window and the status bar.

Added the ability to browse PEX definitions on the system. Choose the View | PEX definitions menu when a PEX Analyzer view is active to get to this feature. The properties menu will show you the contents of a PEX definition in a format nearly identical to

PRTPEXDFN spool file output. I still need to add event detailed descriptions to these properties.

Removed a seemingly random crashing bug caused by a control in the time subsetting page of the analysis wizard. This page is temporarily disabled while I write replacement time selection code.

TDE ID in the Collection properties collection jobs and collection tasks pages is now shown in 8-byte hex instead of decimal. Thanks to Kent Milligan for help on this one.

Made improvements to the PEX Collection wizard by disabling certain options as appropriate for the selected iDoctor-supplied PEX definition.

Fixed a bug at V5R1 in the collection events page where the counter number was showing for trace collections.

September 20, 2001 - C00057

Added the ability to customize the PEX Analyzer report title default in the Preferences window by adding your own constants. The dropdown list is now editable and your changes will of course be saved for future use.

Fixed a bug in the Analysis wizard on the job/task subsetting page that would happen when selecting a task name with mixed-case (upper and lower case) from the Browse Tasks Window. A query to validate if the task exists was failing due to a problem handling the mixed-case task name.

Fixed a bug in the Analysis wizard on the job/task subsetting page. Selecting a job and using the All thread option would generate an analysis containing all TDEs in the collection instead of all threads for the particular TDE.

September 19, 2001 - C00056

Added collection and analysis filtering to the "Work with Queries" and "Work with Graphs" interfaces. Added 'Work with Queries' and 'Work with Graphs' menus for collections and analyses so you can see the queries or graphs only for a particular collection or a particular analysis.

Fixed a problem on the collection properties general page. The suspended time value would overflow in some cases.

August 23, 2001 - C00050 (Client Update)

Fixed a problem on the server side that was occurring because a SAVLIB was happening over a PEX db library and held a lock on the 'list analyses for a collection file, QAYPEANAL' in that library that prevented the status of 'in progress' analyses from being updated.

Fixed a problem where it was possible to click the delete icon from the toolbar at times when not appropriate (like when a library is shown). You can't delete a library in PEX Analyzer. The tool would display the delete analysis dialog inappropriately.

Made some changes with regards to analysis and collections status to work with changes on the server side listed above.

August 9, 2001 - C00048 (Client Update)

In the data viewer I changed the title of table and graph views to display the collection name immediately following the report title.

Improved the display of several numeric fields in collection properties by adding decimal places. Formatted the memory size on the collection properties system page to a more readable format (usually GB if memory is over 1 GB instead of x,xxx,xxx,xxx KB)

Fixed in collection properties (events page) where the PCMO interval was listed as milliseconds when it should have been microseconds. This PMCO Interval field is milliseconds at V4R4 and V4R5 however.

July 30, 2001 - C00047 (Client Update)

Fixed a bug in the analysis fast path menu options, where invalid analysis types were sometimes being shown. thanks to Deb Hatt for finding this one.

Made most property windows modeless instead of modal (except query definition and graph definition). This means that you can open multiple property windows (like collection or analysis properties) at the same time. The way this works though is that if the underlining object for a property window in the tree/list view is refreshed or deleted, the property sheet you are looking at will be destroyed as well.

You can now open up as many record quick view dialogs at one time as you want. However if you close the list the quick view dialog is dependent on, the quick view dialog will also close.

Modified the title of the record quick view dialog to indicate, the view the data is coming from.

In the record quick view dialog the record column heading has been changed to "record nn" where nn is the record number based on the current sort sequence.

Fixed a problem where the program to get the record count of a member on the server was being called from library SMTRACE instead of library QYPBASE. This program is called by all components and is therefore a big problem when trying to run a component when PEX Analyzer is not already installed on the system! Thanks to Jay Kurtz for letting me know about this one.

July 25, 2001 - C00046 (Client Update)

This update enhances how 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 to use.

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.

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.

Added graph coordinate properties for the custom graphs.

July 7, 2001 - V5R1 S00025-C00038, V4R5 S00022-C00038, V4R4 S00022-C00038

STATS hierarchical analyses have been changed to more accurately represent the true chronological call flow. You will be able to work with any existing STATS HIERARCHICAL analyses, however we **recommend** that you delete and then recreate any STATS HIER analyses created over **PEX data at the same OS/400 version/release/modification-level as your iSeries current OS/400 version/release/modification-level.**

Made a change so if a stats hier field Pgm name or MI Cpx Name is double-clicked and a description does not exist, the text "No description available" will be displayed, rather than showing record quick view. Record quick view also happens on a double click but only on any other field being double-clicked.

All fields in the PEX Analyzer tree/list view are now sortable.

Fixed a bug on the collection events property page when running Client Access V5R1. Some events were not being displayed.

Fixed a problem at OS/400 V5R1 where V4R4 collections were incorrectly showing a status of "Collection not supported" in the GUI.

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 6, 2001 - V1R1M0 V5R1 S00021-C00016, V4R5 S00017-C00016, V4R4 S00018-C00016

The IOP information gathered for DASD operations had the IOP Type field incorrectly truncated to 3 chars. It has been changed to 4.

STATS hierarchical analysis had a problem when 0 detail records existed for an individual job/thread (partition) - this is now fixed and thanks to the Rochester Benchmarking Center for reporting this error.

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)

Now displaying the field short name in the record quick view dialog.

If an analysis output file gets manually deleted and is at a later time attempted to be read by the GUI, it will no longer display more than one error message back to the user.

Fixed a problem where fields defined as CCSID 500 could not be read in properly.

Fixed a problem where Null data would be displayed as garbage characters within the data viewer.

Fixed a problem where it was possible to go into the query definition for a file after the file was deleted and this was causing a subsequent crash.

MAJOR performance improvement when using query definition of a very large file like task switch trace in the data viewer.

May 18, 2001 - V1R1M0 C00012 (Client Update only)

Hover help has been added to graphs (V5R1 only). The popup help includes the labels (field descriptions) next to each value.

Fixed a problem where negative zoned values (effected netsize analysis) were not always being displayed properly in the Data Viewer.

Previously if the active log file was deleted while logging was activated, the logging events would be lost because the file would not get recreated.

May 10, 2001 - V1R1M0 S00016-C00007(for V4R4), S00015-C00007(for V4R5), and S00019-C00007(for V5R1)

GUI generated analyses have a unique member identifier assigned. This object being used to store this identifier was being deleted by the QYPINT/STRIDOCOL command. The result is that the relationship between the original PEX data input to an analysis and the subsequent iDoctor-created analysis file(s) was not being preserved.

Fixed a small graph display problem at V5R1.

May 5, 2001 - V1R1M0 S00010-C00006(for V4R4), S00011-C00006(for V4R5), and S00014-C00006(for V5R1)

Switched to a different self-extracting .exe tool (from Winzip) to improve the install process. Previously the iDocSetup.exe would not always get launched and there wasn't any kind of status bar or visual indication what was happening in files were unzipped.

Added a warning message when trying to sort a very large number of records (400,000 or more).

Analyses run in batch now use the job description SMTRACE/IDPAJOB1 instead of using the default job description for the QZRCSRVS server jobs.

May 5, 2001 - V1R1M0 S0009-C0005(for V4R4), S0010-C0005(for V4R5), and S0012-C0005(for V5R1)

As beta test period is now ending completion the **STRIDOCOL** command in library QYPINT now checks for authorization to the iDoctor for iSeries PEX Analyzer tool.

April 26, 2001 - V1R1M0 S0008-C0005(for V4R4), S0009-C0005(for V4R5), and S0011-C0005(for V5R1)

New DASD fields in SMTRMOD and associated SMTRMODLn files in library SMTRACE. We've added additional information about physical disk operations such as the IOP.

We've made changes to our authority checking to be specific to a version/release of iDoctor for iSeries. This requires all existing customers running V4R5 or higher to have new access codes generated. We have sent out new access codes to all of our registered customers, but if we missed you, please contact us and we'll get you a new access code.

Major performance improvements in the client code!

Save As, Save Selection As and Find within a table view are 3 to 4 TIMES FASTER! Scrolling or opening graph or table views is also much faster than before.

Improved SQL Query View.

Improved Query Definition Interface

Windows can be tiled vertically now, instead of just horizontally.

Double-clicking on a row of data in a table will bring up the 'Record Quick View Dialog'

Added a View | Field Descriptions (toggle) menu in the data viewer. This allows you to toggle between short and long field names with a single click. You can still set this value through the output preferences dialog.

Graph coordinate details are more flexible. Now you can choose to click on multiple points and compare them within one window if you like. The properties window can also be left open as you click on different points on the graph and the properties window will get updated instead of needing to close and reopen the window each time.

The ODBC connection timeout property page has been removed.

When listing the reports within an analysis 2 new fields have been added. Number of records, and an indicator to show if the report is graphable.

F4 or the Execute menu can be used to execute an SQL query in the SQL query view. This has replaced the use of the enter key to execute a query.

Setting the font within a tree/list view window will now update the font in both parts (tree and list) of the window.

The help link to the Analysis intervals property page is fixed.