

Update History for iDoctor for iSeries - PEX Analyzer for OS/400 V5R2

May 11, 2007 – C00542

Flyover text for long text strings in table views (like SQL statements) should no longer be cutting off text on the right side of the window.

Apr 3, 2007 – C00535

Creating a PEX collection from the GUI and specifying that it should be built as a PEX MGTCOL object will no longer cause the status of this collection while it's active to not be visible.

Collection properties for PEX MGTCOL objects are now behaving correctly. Previously it was showing blank fields that did not apply.

On the Display tab of the Preferences window added an option "Display patterns". When checked the graphs will display line hatchings and patterns instead of solid colors.

Added options to select the desired fill patterns from the legend and graph definition primary Y-axis fields panel. These selections are not savable in a user-defined graph and are only used for the current graph.

The record quick view window no longer has a checkbox to "allow multiple rows" to be compared. This is in effect always on now. If the user selects multiple rows in a table (up to 30) and uses the record quick view menu, they can be viewed vertically.

The window that prompts a user for a member has been revamped. Instead of the user needing to specify the member for every single file found in the query, a list of all tables found in the query is built and then sent to a single dialog. This window allows the user to pick a member from a list and then indicate that it should be used for all files found in the query.

Query parsing should be a bit faster. Removed redundant checking for the existence of file/members in the query.

Fixed some flickering problems in the list of fields on the field selection panel of the query definition interface.

If the query behind the graph is changed/reran and it returns no data, the graph will now display the message "The query behind this graph returned no data" rather showing the graph data from the local cache.

If the graph definition window is resized, the coordinates are remembered for next use.

For readability purposes, the field column in the list for the graph definition -> primary Y-axis fields window now only displays the short field name instead of the long description and the field name.

In the open file window, pressing the enter key will no longer close the window, but will instead trigger the browse button to refresh the list of files

Added a menu option to the legend to hide/show a bar's black border. This can offer the effect of emphasizing a field in the graph

The legend can now be hidden/shown using the toolbar when the legend has focus. Previously this button only was enabled when the graph had focus.

The legend will no longer force lower case characters in field descriptions.

The creation date shown for user-defined queries and graphs should now show the creation time as well as the date

Fixed some multi-threaded timing issues that would cause canceling a long running query to crash the client. As a side effect of this the amount of time the user has to cancel the query is reduced, once the query has executed and it is to the point of building the columns and fetching rows; that part cannot be canceled.

Fields removed from the legend (or dropped into the 'available fields' section) will now retain their properties, such as the field description, and color. These fields are temporary and are listed first among the 'available fields'. If added back to the graph, the field description, color/pattern will become visible again instead of being randomly set.

Feb 12, 2007 – C00520

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 the 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 V5R2 are: MF32368 MF33254 MF31568 MF31041 MF35300 MF31375 MF33171 MF31083 MF33749 MF34590 MF34456 TC03252

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 V5R2 have been added to library QYPBASE. The commands will also add/remove individual authorities to the following system objects:

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

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

These commands require SECADM/ALLOBJ

Authority checking for PEX Analyzer added to the GUI before starting a collection. 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.

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

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

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

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

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

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

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

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

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

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

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

May 20th, 2005 – C00411

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

May 16th, 2005 – S00087-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.

The Taskswitch analysis should now run much faster than before.

March 28th, 2005 – S00081-C00391

Made changes to server side analyses to ensure proper handling of independent ASP dasd units in physical I/O traces.

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.

Removed support for including collection overhead in the analysis for Stats Hierarchical collections. This option is only supported at release V5R1.

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.

Added the following events to PEX Analysis collection type *TASKSWITCH: CPUSWTIN *CPUSWTOUTP.

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 ENDIDOCOL 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.

There is a new PEX Definition property page within collection properties. This page contains the PEX definition used to create the collection and replaces the definition jobs and definition tasks pages of collection properties.

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.

Added support for the new 'collect system information' parameter.

Fixed a bug with the delete collection function.

Added help text for the CPYIDOCOL, DLTIDOCOL, CVTIDOCOL, FTPIDOCOL, SAVIDOCOL, RSTIDOCOL commands.

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

Some data reporting problems with the data queue and data area analyses has been fixed.

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 bug where the child analyses (if any) did not get deleted when the parent analysis is deleted.

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.

Fixed a bug in the build where CPYIDOCOL command was missing its program in library QYPINT. This was preventing users from copying PEX collections via the GUI.

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

Finished the support in the GUI for PEX Filters. You can view/change/delete and create PEX filters from the client. To work with PEX filters, right-click on the PEX Analyzer icon and use the "work with filters" menu.

Fixed a bug when creating a PEX collection via the GUI or STRIDOCOL where program GETSTATS1 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 ENDPLEX 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 PEX collections 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

Shipped V5R2 PEX Analyzer. Many enhancements/changes were implemented to support the changes to the PEX database files at V5R2.

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 options panel of the PEX Collection Wizard added an 'advanced' button that provides new parameters: 1. Output format (to produce PEX DB files or a single object for easy transport) 2. ENDPEX priority, 3. Number of threads during ENDPEX, 4. PEX Filter

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 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.