

Update History for IBM iDoctor for IBM i – Heap Analyzer Heap Analysis Tools V5R3

Sep 17, 2008 – C00664

No Heap Analyzer-specific updates occurred this build for this release, but see the Job Watcher update history for general GUI enhancements

Nov 29, 2007 – C00584

Updated the compare function. There is now a stored procedure that gets called to replace the use of the HCOMPARE command in the GUI. You must select two object table dumps for the Create Comparison Report menu to be enabled when right-clicking on the selected collections. This will create a new file visible in the server side output files folder for the 1st collection selected of the 2.

Added a summary report for all object table dump in a library. Right click on a library and choose the "Create Object Table Summary" menu to perform this option. A new file QAIDRHASUM will be created in the library.

Added the QAIDRHASUM file to the server-side output files folder if it exists for any object table dump in a library.

The 'root finder' function of Heap Analyzer has been removed from the GUI. The collection wizard no longer provides this option and these types of existing collections are no longer visible.

Fixed a bug in Heap Analyzer install where it could not be installed if Job Watcher wasn't already installed on the system

Updated the iDoctor GUI to compile using the new compiler Visual Studio 2005.

From the Run SQL View, added a menu "Launch SQL in Run SQL Scripts..." that will take the current SQL statement, place it in a temporary file and launch it in iSeries Access for Windows Run SQL Scripts for the purpose of either running it within that interface or using visual explain. Note that iDoctor uses System naming convention (library/file) and iSeries Navigator uses SQL naming convention (library.file). You will have to change the default connection properties (Connection->JDBC Setup... Format tab, Naming Convention) in order for the queries created within iDoctor to be executable.

Added a button on the toolbar in the Data Viewer to normalize time interval based graphs. This will divide each value of the graph by the interval duration, which will make bars more consistent in height. This button is available on several Job Watcher graphs.

Added 2 new Preferences to the Graph views section of the Display tab. "Use normalize option" -when checked will cause the normalize option to be used when a graph is first opened (if it's available.) "Use variable width bar option" - when checked will cause the variable width bar option to be used when a graph is first opened (if it's available.)

Property sheets that are not modal (meaning you can go back to other screens while they are open) are now displayed inside of a view. This provides the following benefits/changes that users have been asking for: 1) These property sheets are no longer "always on top". You can click another view and it will move in front of the property sheet view. 2) Property sheet views can be resized to any size. 3) "Drill-down" actions from a property sheet view will now open on top of the property sheet view, like any other drill down in the Data Viewer.

Added a preferences on the misc tab called "Show non-modal property pages in a view". If you don't like the new property sheet views, you can uncheck this and go back to the way it was before.

Property sheet views that don't have a need for OK, Cancel buttons will no longer show these buttons (for the purpose of saving screen space). The user can press escape or the X to close the window just like any other view

Added a menu option called "Filter libraries..." from the component icon to filter the list of libraries shown. This setting applies to all components that display collection libraries. This library filter is shown next to the component's name in the tree when active. This can be helpful if there are many libraries on the system or if the connection in use is slow.

Added a scheduling preferences page that allows the user to change the default scheduled date and time, from 1 week in the future, to whatever the user prefers (in +days/hours from current date/time).

When opening a report from the list portion of a component tree/list window using the popup menu, added an option called "Edit" below the normal "Open table" or "Open graph" menu. This option allows a user to quickly open a user-defined or iDoctor-supplied table or graph such that the SQL editor is initially open without actually attempting to run the query. If the query will take a long time run, this can be used to open and modify the SQL statement first. This technique can also be used in graphs to modify the graph settings before running the query behind the graph

Added a new view, which processes database SQL statements in a separate thread. These requests can be queued up and are processed one at a time. This view is called the Remote SQL Statement Status View and is currently used by the V5R4/V6R1 JW summary program as well as the "Execute in batch" option from the SQL editor

The popup menu in the SQL editor has a new option, Execute in batch. This will run the SQL statement in a separate thread and show the status of execution in the remote SQL statement status view. If an SQL select statement is provided, you will be optionally prompted for an SQL table/library name in order to create a table using the SQL statement given

Added new icons to the remote command status view such that the icons change based on the status of the request (waiting to run, completed successfully, error). These same icons are used on the remote SQL statement status view.

iDoctor can now handle SQL statements that contain SQL table names

Opening a graph/table in the Data Viewer should no longer tie up the GUI and you should be able to cancel this action by closing the graph/table view being opened if desired. While a report is being opened you should be able to do other things in the GUI like view already open reports/graphs or go back to the main window and browse other collections. If you are currently opening a report, you can go to another view and drill down to another view while the 1st one is running if you wish. A maximum of 3 data viewer connections are used to accomplish this which means you can have a maximum of 3 queries running simultaneously from the same iDoctor client. If you exceed this the GUI thread may appear to hang and execution will no longer occur simultaneously for these extra requests. If you try to scroll a graph/table while another one is being loaded, you may experience delays in the GUI thread (if the same connection was used for the report being scrolled and the new one being opened) but it should work once the previous query finishes. As part of this change I've removed the preference "Show query cancel window when filtering or resorting a report" on the Misc tab as this option is no longer needed. The multithreaded behavior should now act the same whenever opening/requerying any graph/table in the Data Viewer.

Show the report/graph title on the view being opened before the SQL statement finishes.

Add/Remove programs in the Windows control panel now properly recognize the iDoctor GUI. You can use this option to uninstall the client.

Made a change how data viewer connections are handled, so that instead of only using a maximum of 3 connections if 3 simultaneous queries are executed, 5 connections are available to the Data Viewer's views where each view that is created uses a different connection by cycling through them. By doing this it will allow the user to open view A and view B, and scroll through view A while view B is running a query. Previously view A and view B would likely be using the same connection which meant trying to scroll view A while view B is still running a query would hang the client until view B finished it's query! The maximum value of 5 is a registry setting if you wish to change this:
[HKEY_CLASSES_ROOT\IBM.AS400.Network\3RD PARTY EXTENSIONS\IBM.PEX\DataViewer]
"QZDASOINITConnections"="5"

The toolbars in the main window and in the Data Viewer are now dockable (can move them to wherever you want)

When refreshing a table, the hourglass on the cursor is often lost. To resolve this, show a message on the table when SQL statements are running in the background similar to what is shown for graphs "The query behind this report is running, please wait..."

Since graph tool tips can sometimes get in the way of what you are doing, I've added the ability to turn them on/off by pressing CTRL+T while in a graph. The status bar in the data viewer now includes a message after the graph control memory used message that indicates if tool tips are currently enabled or disabled.

Fix a bug on the query definition -> member selection page where users could not change the member their query is running against.

Fixed a problem with graphs having some portion of their bars missing every 2000 or so intervals of continuous scrolling if the number of bars per page is set to > 100.

If an error is detected in a graph definition, such as a required field does not exist in the SQL statement, the error message will be drawn right onto the graph view now instead of shown in a popup window.

Fixed some problems in the manage query definitions interface such as getting SQL errors when copying many query definitions at once.

Fixed a bug in the query parsing that was causing too many CHKOBJs to be performed when checking for the existence of files. This would slow down the opening of reports and on slower connections could greatly decrease response times.

The hide/show borders option in the graph legend wasn't working if the graph definition was defined such that the field had the same border color as the fill color. If this is the case the color will switch to black and then back to the original color.

If an SQL error occurs when launching any graph, the sql editor/error window should be visible above the graph right away rather than requiring the user to open the editor and reexecute the query.

The option to save a graph as a JPEG (Save As... menu when right-clicking on a graph) has been removed from the GUI. It didn't work since the June builds when the graph control was recompiled to work with the latest MS compiler and I can't easily fix it. The preferred method to achieve this is use Alt-Print Screen which copies the current window to the clipboard. Then you can paste this image into Windows Paint or something similar and save the image to the desired format. Also removed the Copy menu (which copied the graph view only to the clipboard) because you can just use Alt-Print Screen to do this which will include the legend and the copy function did not

Add check in the install upfront for system value QALWOBJRST.

The Graph views section on the Preferences -> Clipboard Tab has been removed since the copy to clipboard of the graph image option was removed from the GUI in favor of telling users to use Alt-Print Screen instead

May 11, 2007 – C00542

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

Apr 3, 2007 – C00535

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

Heap Analyzer users should now be able to open a Data Viewer, and then open files and write custom queries from the Heap Analyzer component view.

Updated the install so all commands issued are library qualified with QSYS. Though unlikely this could cause a potential problem depending on what was in the library list.

Add the ability to delete an object root finder in the GUI and using the DLTHEAPWCH command.

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.

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.

Made a change so that opening a file or using the SQL query window within Heap Analyzer does not require a valid JW or PA access code. Heap Analyzer now has full functionality in the Data Viewer.

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.

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.

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 20, 2006 – C00476

Updated the V5R3 Heap Analyzer PTF checking to include PTF SI22708. This PTF fixes a problem where an infinite loop was possible in the WCHJVA API, if there are no threads in a JVM. It can be immediately applied. If any collections are active, they should be stopped first.

February 14, 2006 – C00476

A data area QPTFOVR will now be shipped in library QIDRHAJ. If this data area contains a '1' the PTF checking will not take place when running the WCHJVA command. When running via the GUI, the user will be asked if they wish to override the PTF checking and continue anyway, if they say yes a '1' will be placed in this data area in order for the collection command to work without the required PTFs on the system

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.

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

Added new graphs and reports (and drill downs) in Heap Analyzer for the object table dump and object create profile modes. The object create profile features a graph showing the object create heap sizes per thread with right-click drill down options to view the object classes within the thread and then the stacks for a specific class name. The counts and object sizes are available in each level of these reports.

The start root finder menu should only be enabled in the root finder tree and from the object table dump/profile views if the Java job the collection is over is still running

Fixed an issue with being unable to view user-defined queries and graphs in Heap Analyzer.

Modified the save query and save graph window so that the component type selected can be heap analyzer - object table dump, heap analyzer - object create profile, or heap analyzer - object root finder so the user can create user defined queries and graphs and have them show up only under the appropriate type of collection within Heap Analyzer

The toolbar shown in the Data Viewer for a Heap Analyzer graph/report is now the same as the one shown for Job Watcher graph/report. Previously some functions were disabled.

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

Reestablish the ODBC connection automatically if the connection is lost.

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

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

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

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

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

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

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

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

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

Change the way quick sorting a column works. Left click = sort by column ascending. Right click = sort by column descending. Shift + left click = add column ascending sort to end of existing sort (if any) Shift + right click = add column descending sort to end of existing sort (if any) If the column is already in the existing sort and shift is held down it is modified based on whether the column was left or right clicked. Note that multi column sort is not available if the list was not created via a SQL query. This applies to the component views (lists of libraries and collections)

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

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

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

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

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

August 18, 2005 – C00436

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

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

August 9, 2005 – C00431

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

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

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

August 1, 2005 – C00428

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

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

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

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

July 15th, 2005 – C00423

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

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

- *ALL to QAPEXDFN file in QUSRSYS (for the GUI create pex definition interface)
- *ALL to QAPEXFTR file in QUSRSYS (for the GUI create pex filter interface)
- *USE to ADDPEXDFN/RMVPEXDFN cmds in QSYS (for the GUI create pex definition interface)
- *USE to ADDPEXFTR/RMVPEXFTR cmds in QSYS (for the GUI create pex filter interface)

These commands require SECADM/ALLOBJ

Authority checking added for Heap Analyzer just before starting a collection in the GUI. 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 also added to the WCHJVA command at all releases

Fix a bug with the object table dump compare option. The first report shown for the compare would always contain 0 records.

Allow heap analyzer collections to be expandable to view partial data for collections that ended in error.

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 \IBMiDoctor 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 – S00082-C00408

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

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

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

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

user id used for iSeries Navigator may be different and will be remembered independently of each other.

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

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

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

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

March 28th, 2005 – S00028-C00391

Updated the required PTFs. These PTFs listed on the download page for Heap Analyzer are required in order for the root finder collection option to work.

The root finder collection option, and analysis tree is now available.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

When retrieving the field SQLSTMT from file QAIDRSQLO4 (iDoctor user-defined query file) for any component/release do not strip off the trailing blanks if any on the SQL statement. Normally all fields are stripped of trailing blanks when read into the client.

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

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

January 13th, 2005 – C00367

Fixed 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

Fixed an issue with the iDoctor-supplied queries folder. The descriptions of the queries in this folder were not displayed correctly.

September 13th, 2004 – C00345

This build enabled support for Heap Analyzer at release V5R3.

Added new property pages for Heap Analyzer collections including a General page and Creation Settings page.

Added slashes to the job names listed in the list of collections in a library within Heap Analyzer.