

Update History for IBM iDoctor for IBM i – PEX-Analyzer V5R3

May 18, 2009 – C00718

---> C00718 - Pex Analyzer - Bug Fix Request

The STRPACOL command at V5R4 and higher, would not collect taskswitch events if the following command was used:

```
STRPACOL COLLECTION(SIMPLE)
  PROBLEM(*TPROF)
  DURATION(5)
  ADDEVTS(*TASKSWITCH)
```

---> C00718 - General - Usability

When using the Save -> View As... (to CSV) file function the query will no longer be rerun upon completion..

---> C00718 - General - Bug Fix Request

Column headings messed up when saving reports to CSV files that contain commas in the headings.

---> C00716 - Collection Services;Disk Watcher;Job Watcher - Enhancement

For all time-based graphs, the interval number will now always be included in the x-axis value on the flyover.

---> C00715 - General - Bug Fix Request

On the main window, the View -> Remote Command Status View menu wasn't working.

---> C00714 - General - Usability

You can now scroll through graph legend vertically with the mouse wheel.

---> C00714 - General - Bug Fix Request

Scrolling the synchronized table view with the mouse wheel will not properly scroll the graph as well.

---> C00714 - General - Enhancement

Added a new menu called "Choose Database Members" to the SQL editor that clears the current member selections and brings up the member selection window.

---> C00714 - General - Usability

When defining a new query in the SQL editor you will no longer be prompted for the desired member every time you run the query (just once!).

---> C00714 - General - Usability

Updated the label shown on the transfer collection window when sending data to IBM to include the full file name format.

Also added .idr. to the filename as this will indicate the data is iDoctor data that should be restored and summarized/ready for analysis on the Ecurep debug system.

---> C00714 - General - Bug Fix Request

The Check PTFs function will fail more gracefully now if the system does not have the required level of the PTF checker program.

---> C00714 - General - Enhancement

Increased the max amount of data read by the FTP "Quick View" function on text files to 512KB from 10KB.

---> C00714 - General - Bug Fix Request

Fixed a bug in table views if the data contained more than 100 million records, no records would be shown. The maximum records that can be shown is 100 million. You can change the sort or filter to see the rest of the data.

---> C00714 - Collection Services;Disk Watcher;Job Watcher;Pex Analyzer - Usability

Running a summary in JW, CS, DW or an analysis in PEX Analyzer will now cause an automatic refresh of the window after the summary completes rather than requiring the user to manually refresh the window.

---> C00714 - General - Enhancement

In the FTP component, if the connection is lost (times out) , it will be reestablished automatically without needing to retype the password.

---> C00714 - General - General

Since the call stacks under the Active jobs folder didn't work at V5R3, I've disabled it. You can still view statistics for jobs/threads however.

---> C00713 - General - Enhancement

All major collection types (PEX, CS, DW, JW) can now be saved using the Save... menu for the collection. You can save multiple collections or monitors as well.

After using this option a the new folder "Saved collections" will show a record identifying the save file which you can use to restore the data, or send the data to another system or ftp server directly from the desired save file.

---> C00713 - General - Bug Fix Request

Removed an SQL error that would be shown when going to the Scheduled jobs folder if the idocor job scheduler file hasn't been created yet on the system.

---> C00713 - General - Enhancement

In V5R3 and higher, in library QIDRGUI the FTPFILE command has been updated:

Local file/remote file parms increased to 100 chars.

You can specify a *GET or *PUT of the file as the optional FUNCTION parameter.

Most parameters are mixed case now which fixes problems with mixed case passwords or filenames.

User profile has been increased to 50 bytes.

The script/log files size has been increased to support longer path/file names.

Also the FTPBATCH command has been removed since it's no longer needed.

---> C00713 - General - Enhancement

In library QIDRGUI the program that performs PTF checking has been updated (added another parameter). The latest server/client builds are required in order for PTF checking to work correctly. This was done in order to allow PTF checking for Disk Watcher.

---> C00713 - General - Enhancement

At V5R4, library QIDRGUI now includes a command SAVPFRCOL which is used to save performance data collections in the desired library. It supports saving of PEX, CS, JW, DW collections as well as PEX/CS MGTCOL object collections. A command of the same name that performs similar function exists in library QSYS at V6R1 and higher.

---> C00713 - General - Enhancement

Added PTF checking when starting a super collection that includes, Disk Watcher or PEX data for the required Disk Watcher or PEX PTFs.

---> C00713 - General - Enhancement

Added PTF checking when starting a Job Watcher, Disk Watcher or PEX monitor for the collection type(s) selected within the monitor wizard on the basic options screen.

---> C00713 - General - Enhancement

Added PTF checking when starting a Disk Watcher collection.

---> C00713 - General - Enhancement

Based on customer request, added an option to the list of connections that allows the user to perform a PTF check for all components that provide PTF checking against the selected system(s). The results of the check and any missing PTFs are shown in the Check PTF Results window.

---> C00712 - Installation - General

Updated the uninstall CL program at 530 and higher to handle removal of the new iDoctor control files QAIDRJWM2, QAIDRDWM2, QAIDRPAM2

---> C00712 - Job Watcher;Pex Analyzer - Enhancement

Updated the monitors wizard on the basic options screen to include the ASP limit option. The value shown is taken from data area QMONASPLMT. If the used disk space is greater than this value then the user will not be able to start the monitor unless the limit is increased up to a max of 99%. Changing the value from this panel will update the data area which is checked by the STRJWMON/STRDWMON/STRPAMON commands.

---> C00712 - Job Watcher;Pex Analyzer - Enhancement

Added several enhancements to the Monitors support at V5R3 and higher (latest server builds required)

PEX Monitors are now available at V5R3 and higher.

Add the Monitors folder to the PEX Analyzer GUI.

Added a data area QIDRWCH/QMONASPLMT which contains the maximum allowed asp limit. If the ASP for the monitor library's available disk space has exceeded the value specified in the data area then the monitor will end. The default is 90%.

The maximum length of a monitor name has decreased from 8 to 7.

The maximum number of collections in a monitor has increased from 99 to 999

The RSTJWMON/RSTDWMON/RSTPAMON commands have been removed. You can now restart a monitor using the start commands: STRJWMON/STRDWMON/STRPAMON

The monitor control files have changed as follows due to parameter changes on the STR monitor commands:

QAIDRJWM1 -> QAIDRJWM2

QAIDRPAM1 -> QAIDRPAM2

QAIDRDWM1 -> QAIDRDWM2

The first time a monitor is started on a system, the existing monitor information will be copied over to the new file.

In STRPAMON added a parameter called ENDPEX data option with these values/meanings:

*LIB = The monitor will create the PEX database files for each collection.

*MGTCOL = The monitor will create PEX MGTCOL objects for each collection.

*SUSPEND = The monitor will issue ENDPEX with the *SUSPEND option instead of dumping any data.

---> C00711 - General - Usability

Changing the bars per page (via the graph definition->set bars per page menu/not via the preferences interface) for a graph, and then changing to a different alternate view will now preserve the bars per page setting when loading the alternate view graph.

---> C00711 - General - Usability

Using the preferences window from a graph view will cause the graph to be refreshed upon closing the preferences window.

---> C00711 - General - Usability

Using the menu within a graph

Graph definition -> Set bars per page -> Set value... will now take the user to the general page of the graph definition properties to set the bars per page override there instead of going to preferences.

---> C00711 - Pex Analyzer - Bug Fix Request

The generic job name length preference did not work in the PEX Analyzer graphs that provided generic job name groupings and drill downs.

---> C00711 - Pex Analyzer - Bug Fix Request

In PA, for the sql-based analysis, collection overview -> CPU consumed rankings, the drill downs were showing duplicate menu options.

---> C00711 - Job Watcher;Pex Analyzer - Bug Fix Request

JW FTP Transfer function was failing if an IP address was used instead of the remote system name. This was due to a bug in the client side when issuing the command to FTPJWCOL.

---> C00710 - General - Enhancement

You can now drill down from the Synchronized table under the graph, just like you were drilling down from the bars in the graph with a couple of exceptions:

- 1) You can't select the desired wait bucket to drill down by default from the table.
- 2) If you are working with rows in the table that aren't in the graph (because the number of rows currently shown in the table is greater than the number of bars currently shown in the graph) then the drill downs won't be available.

---> C00710 - General - Usability

Changed how selection and "quick view" works in graphs to be consistent with how they work in a table.

A DOUBLE-CLICK ON A BAR IN A GRAPH WILL NOW OPEN THE QUICK VIEW WINDOW INSTEAD OF A SINGLE CLICK! (ahem in case you weren't listening)

Essentially, a single click will select a bar (without needing to hold down ctrl/shift) You still need to hold down shift to make a range selection however (just like a table).

I did this to help avoid the annoyance of opening the quick view window when not desired.

---> C00710 - General - Usability

The preference for new users on the Miscellaneous tab "Maximize reports opened into Data Viewers" is now on/checked by default.

---> C00710 - General - Usability

Updated the minimum level of client access needed on the PC in order to run iDoctor to V5R3 or higher.

---> C00710 - Job Watcher;Pex Analyzer - Enhancement

Add support to run PEX Analyzer monitors at V5R4. New commands STRPAMON, RSTPAMON, etc have been added to the QIDRWCH library and the GUI now provides the option to start a PEX monitor at V5R4. The latest server builds are required.

---> C00710 - General - Bug Fix Request

Duplicate as table menu option not working from the Synchronized table view.

---> C00710 - General - Bug Fix Request

The selected row of the synchronized table was not always matching up with the graph.

---> C00708 - General - Bug Fix Request

When using the report generator function modified the implementation that creates the report directory so that any error messages and the GetLastError rc from Windows will now be shown to the user.

---> C00707 - Installation - Enhancement

Updated the Install program on the FTP parameters page to include a checkbox that lets the user choose to perform or skip the ftp connection check (to QTEMP library).

---> C00707 - Pex Analyzer - Enhancement

Added the job name and changed the sort order of the IFS analysis report in PEX Analyzer to job and then time since collection started..

---> C00707 - General - Usability

When deleting a monitor the GUI attempts to delete the job scheduled entries associated with the monitor if any exist. If none have ever been created an SQL error would be shown about a file not existing. Removed this message.

---> C00705 - Pex Analyzer - Bug Fix Request

Added fixes for the PEX Analyzer SQL-based IFS analysis.

---> C00705 - Installation - Bug Fix Request

Fixed a bug in the installer when it tries to test the ftp connection. On a failure the FTP error message and rc were not being displayed in the message box.

---> C00704 - General - Bug Fix Request

Fix a bug in the table and graph views where if you change an SQL statement and you have a syntax error, you would no longer be able to run queries from that view.

---> C00704 - General - Usability

Improved performance in the JW interval detail property pages. Information on these pages is retrieved much faster now when changing the current interval especially when working on a slower internet connection.

---> C00704 - General - Usability

Improved performance in the parsing of SQL statements by reducing the number of CHKOBJS that occur. Also cached these checks when performing drill downs so they are not made again.

---> C00704 - General - Usability

Improved performance in several places in the GUI that are used to build a list of items (definitions, monitors, etc).

---> C00703 - General - Usability

SQL tables and SQL aliases are no longer shown in the member selection window.

Also overrides are no longer made against SQL tables (unnecessarily).

---> C00703 - General - Enhancement

The graph compare function is now available in all components except for the old style ("flattened") graphs in PEX Analyzer. Previously this was only available in JW, CSI and DW.

---> C00702 - General - Enhancement

When working with the bottom graph of a comparison graph view, you can now change the time interval size (when applicable) and the Change SQL Parameters window is also available.

---> C00702 - General - Enhancement

The libraries folder in the component views, now contains the library filter (within the folder name) if one is set.

---> C00702 - Collection Services;Job Watcher - Enhancement

Added a preference on the Miscellaneous tab that lets you indicate how many characters should make up the generic job names on the generic job grouping graphs in CSI and JW. The possible values are 2 - 9 with a default value of 6 (the previous default).

---> C00702 - General - Bug Fix Request

The Edit -> Preferences menu when asked from the tree side of a component view didn't work.

---> C00702 - General - Enhancement

The SQL log now contains the QZDA job number and (PC time) timestamp.
Also removed a blank line between the statement and the line indicating how long it took to run.
Access this feature from the View -> SQL Message Log menu within the Data Viewer.

---> C00702 - General - Enhancement

Added an SQL tab to the Preferences interface that contains the following options:

Log messages to log file (default is now OFF)

Maximum data viewer connections (Number of QZDASOINIT jobs to utilize in the data viewer)

---> C00702 - General - Enhancement

Added a Confirm page to Preferences where all the confirmation options from the Miscellaneous preferences tab will now be located.

---> C00702 - General - Enhancement

Made a change in the iDoctor database code, so it will no longer remove overrides in situations where none need to be removed.

---> C00702 - General - Usability

Updated the error message given when an access code was generated for the wrong OS VRM (the server code and OS VRM match but the access code doesn't)

The old message was

Access code is invalid (OS release mismatch?)

the new message is

Access code is invalid. It appears to apply to a different OS VRM. Contact support at idoctor@us.ibm.com.

---> C00702 - General - Usability

Updated the error message given when a user upgrades their system but forgets to upgrade the iDoctor server code.

The old message was

Access code is invalid (OS release mismatch?)

the new message is

Access code is invalid. A previous release of the iDoctor server code is installed and must be updated. Visit our website to install the correct version of the server code for this system. http://www-912.ibm.com/i_dir/idoctor.nsf

---> C00702 - General - Usability

In the component views, refreshing a large list of items (like a list of active jobs), should look cleaner now.

---> C00702 - General - Usability

Add a history for the library filter dialog so that it remembers the last generic names entered and displays them in the drop down list. Also added a remove button to clear the selected entry from the list.

---> C00700 - Job Watcher;Pex Analyzer - Enhancement

Updated the Job Watcher monitors support for V5R4 to provide the capability to automatically run PEX stats flat collections for the top CPU QZDASOINIT jobs.

Added a new JW definition Q5MINQZDAS. Use the Reload IBM-supplied definitions option from the Monitors window (basic options - actions button) to make this definition available.

When definition Q5MINQZDAS is used, the monitor will first collect all QZDASOINIT jobs for 5 minutes at 5 second intervals.

Every 5 minutes from then on the monitor will look for the top cpu burning QZDASOINIT jobs and start a PEX stats flat collection against them (up to 5 max)

The PEX stats collections are named the same as the monitor + 01-99. After 99 PEX collections it will wrap. You will end up with many more PEX collections than JW collections.

There is also a new JW command in QIDRWCH library called RUNPXSTATS that you can manually run if desired against a JW collection to collect PEX Stats flat for the top N CPU burning jobs in the collection over the last 5 minutes (the collection wouldn't necessarily need to be a monitor or even QZDASOINIT jobs).

---> C00698 - General - Enhancement

The threads within a job within the Active jobs folder now provides a drill down into the current call stack for the selected thread.

You can configure the fields shown for the call stack using the select fields... menu found when right-clicking the thread ID from the tree side.

---> C00698 - General - Enhancement

Most iDoctor components now contain an "Active Jobs" folder which provides WRKACTJOB like functionality from the iDoctor client.

You can also sort by a statistic and refresh to keep tabs on the top CPU burners, etc.

There is also a filter option to filter the list by name, user, number, current user or minimum CPU %.

Use the Select fields... menu option when right-clicking the folder to rearrange fields or add additional fields (from the Open list of Jobs API)

Expanding a job shows the threads and the thread statistics available for each.

You can start Job Watcher or PEX Analyzer collections or add jw /pex definitions using the selected jobs within the Active jobs folder. You can also end the selected jobs or view job logs .

---> C00698 - General - Bug Fix Request

When selecting a folder in the tree side of component views, the status bar was not identifying the description of the selected folder as it should (like when selecting items in the list view).

---> C00698 - General - Enhancement

From the new ASPs and Disk Units folders you can now select one or more ASPs or disk units, and right click to create either a PEX filter or Disk Watcher definition from the desired selections.

---> C00698 - General - Enhancement

Most iDoctor components now contain two new folders showing the ASPs and Disk Units configured on the current system.

The ASPs folder allows drill down to view the Disk Units within an ASP.

The Disk Units folder provides WRKDSKSTS type of statistics with updates provided with each refresh. (also includes total I/O, total sizes). Use the right-click menu option Reset Statistics on the Disk Units folder to restart the collection of disk statistics.

Use the Select fields... menu option when right-clicking the folder to rearrange fields or add additional fields (from the Open list of ASPs API)

The status bar of the Main window shows the times for first disk statistics snapshot, and the last one.

---> C00698 - General - Bug Fix Request

Fixed a bug in several components, the timestamps in the alternate view graph descriptions for thread rankings charts would sometimes not be shown in the correct format or at all.

---> C00698 - General - Enhancement

When creating a CSV file with iDoctor, the default setting is now to include the field headings as a header row. You can modify this on the File tab of Preferences.

---> C00697 - Pex Analyzer - Enhancement

Added several additional enhancements to the PEX save/restore analysis at V5R4/V6R1.

You can now drill down from any of the table views, into SR event details for the desired thread.

Added a SR thread rankings folder as another possible starting point.

---> C00696 - Pex Analyzer - Enhancement

Added an SQL-based Taskswitch analysis to PEX Analyzer (V5R4/V6R1). The files available so far are very low level and not "done" yet. (taskswitch object index file).

---> C00696 - General - Bug Fix Request

The Edit -> Open in New Data Viewer option not working right when data viewer 1 is open.

---> C00694 - Pex Analyzer - Enhancement

Updated several reports in the PEX Analyzer save/restore analysis, and added support for the analysis at V5R4.

---> C00693 - Pex Analyzer - Enhancement

For PA V5R4+, for the TPROF report drill down "Call stack hits for procedure 'XYZ'"

added call level 0 (basically the procedure you clicked on when doing the drill down) to the call stacks shown.

Also added these fields:

Procedure code start address

Procedure offset (instruction address - procedure code start address)

Module timestamp

This report is also now available from the Classic TPROF analyses at V5R4 and higher.

C00692 - Pex Analyzer - Enhancement

For PA V5R4+, added a save/restore event analysis that parses out the data in the QAYPEMIUSR file into several reports. This is a work in progress and all the reports are not yet complete.

---> C00692 - General - Bug Fix Request

Fixed a possible crash when opening reports in the Data Viewer, and then trying to close the Data Viewer as they open.

Jan 22, 2009 – C00691

PEX-Analyzer changes

---> C00691 - Pex Analyzer - Bug Fix Request

In PEX Analyzer, when working with the Classic Analyses that contain an IBM-supplied queries folder, added a check to ensure that if the required file for the query isn't available (typically SMTRMOD), then the reports in that folder won't be shown.

---> C00690 - Pex Analyzer - Enhancement

In the PEX Analyzer stats hier report, added a drill down that shows you the call flows for the selected record's REFID.

---> C00688 - Pex Analyzer - Bug Fix Request

When creating a stats flat classic analysis, and subsetting by job, the GUI was showing that each job/thread was filtered for ALL threads when in fact it wasn't.

---> C00686 - Pex Analyzer - Usability

In PEX Analyzer, for the SQL-based Collection overview reports over the QAYPETASKI file: made a change so that the I/O reports shown no longer have different colors to distinguish between jobs in the collection and not in the collection.

Also it was determined that I/O values in QAYPETASKI are not deltas for the life of the PEX collection but are rather cumulative totals for each job since it started. Therefore these reports will only show jobs/tasks that were started after the PEX collection begins collecting data.

---> C00686 - Collection Services;Pex Analyzer - Enhancement

In PEX and CSI added a Wait Buckets tab to collection properties that shows the wait buckets/enum mapping at V5R3 and higher

---> C00680 - Pex Analyzer - Bug Fix Request

Fixed in a bug in PA 540, 610, in the MI user event analyses where the program returned could be a LIC program when the program should have been returning the first non-Q MI program. Updated PEX Analyzer server builds are required for this fix.

---> C00679 - Pex Analyzer - Enhancement

In PEX Analyzer, migrated the MI user event analyses (opens/closes, Idios, etc) back to V5R4. Latest server builds are required.

---> C00679 - Pex Analyzer - Enhancement

Changed the SQL-based analyses user program resolution program so that instead of showing "UNKNOWN" when no Q named program was found in the call stack, the current level program name will be shown regardless of the name.

Although the program name is not a 'user' program, it may be better than showing 'UNKNOWN' as far as users are concerned.

---> C00679 - Pex Analyzer - Bug Fix Request

In PA, some of the new SQL-based analyses were sorting their reports by QTITSP which has been changed to the more granular QTITIMN.

---> C00679 - Pex Analyzer - Enhancement

In PA (V5R3 and higher), in the Start Collection Wizard when using one of the iDoctor problem types, the default event format value for PMCO and Taskswitch is now format 2 (previously format 1).

---> C00679 - Pex Analyzer - Enhancement

In V5R4/V6R1 PA, the STRPACOL command (and the Start Collection Wizard) will now include format 2 events by default for all the MI user events problem types (*DB_OPEN, *DB_LDIO, etc).

Also if collecting storage events (at V6R1 only) using the *NET_OBJSIZE problem type (events *CRTSEG, *DLTSEG, *EXDSEG, *TRUNCSEG) format 2 events will be collected.

Not collecting with Format 2 will now require creating your own PEX definition.

---> C00679 - Pex Analyzer - Bug Fix Request

In PEX Analyzer, on the definition tab for a collection, the format type was not correct if the event type was not one of the base events and the event type was > 1.

---> C00678 - Pex Analyzer - General

For V5R3 and V5R4 PA, removed several SQL-based analyses, except for TPROF and (FMT2) call stacks, events, and the collection overview.

In the future enhancements to the SQL-based analyses will be done in V6R1 and higher only.

---> C00678 - Pex Analyzer - Bug Fix Request

For V5R3 and higher PA, the call stacks analysis is now only available if FMT2 events were collected.

---> C00675 - Job Watcher;Pex Analyzer - Bug Fix Request

Fixed a crash with the transfer collection function, if it was used from the tree side instead of the list side of the component view.

---> C00674 - Pex Analyzer - Enhancement

In PA, SQL-based analyses can now be submitted from the SQL-Based Analyses folder's popup menu.

---> C00674 - Pex Analyzer - Enhancement

In PA, classic analyses can now be submitted from the Classic Analysis folder's popup menu.

---> C00674 - Pex Analyzer - Enhancement

The graph compare function is now available in PEX Analyzer.

---> C00673 - Pex Analyzer - Bug Fix Request

Changed how SQL-based analyses are created in the GUI so that the user interaction/flow is more similar to the one for the Classic Analyses. A new menu "SQL-based Analyses" is available on the collection popup menu. This contains submenus for each SQL-based analyses that is available based on the events in the collection. Once an analysis has been ran a folder representing the analysis and all of it's reports should then be available under the collection or in the collection's popup menu.

---> C00673 - Pex Analyzer - Bug Fix Request

Corrected a bug in the Add PEX definition wizard where specifying program events in the definition for V5R4 and higher wouldn't work because an incorrect command string was being generated.

---> C00672 - Pex Analyzer - Enhancement

In the TPROF SQL-based analysis, reports that show the STRADR (instruction address) field allow a drill down that displays all call stack hits with the most frequent stacks listed first.

---> C00671 - Pex Analyzer - Enhancement

Added an LDIO SQL-based analysis to PEX Analyzer under the menu Database I/O analysis. Some reports similar to what is provided to the classic analysis are available.

---> C00669 - Pex Analyzer - Enhancement

Added a Data Area analysis menu found when right-clicking the collection in PEX Analyzer (if data area events have been collected) which provides an SQL based report similar to the SMTRDTAA file.

General iDoctor GUI changes

---> C00685 - Installation - Enhancement

Added the capability to remove all iDoctor objects/libraries from the desired system by right-clicking the system name from the connections view and choosing "Uninstall iDoctor".

This option can also be used by executing the following commands manually from the green screen:

```
CRTDUPOBJ OBJ(QIDRUNINST) FROMLIB(QIDRGUI) OBJTYPE(*PGM) TOLIB(QTEMP)
CALL PGM(QTEMP/QIDRUNINST)
```

Because this was implemented as a CL program in library QIDRGUI, the latest server builds are required in order for the uninstall option to work.

---> C00691 - General - General

Due to their size, I moved the tip of the day videos out of the normal builds and moved them to a separate download on the main downloads page on the website.

---> C00690 - General - Bug Fix Request

Fixed an error on Vista where at startup there would be registry errors shown due to being unable to write certain keys.

---> C00689 - General - Enhancement

Moving the mouse over a field in the legend will now show additional information about that field in the status bar of the Data Viewer including the field short name and the SQL expression (if applicable/available) from the SQL statement for the field.

---> C00689 - General - Usability

Modified scrolling in the graph legend, so it doesn't scroll such huge amounts per click anymore.

---> C00689 - General - Bug Fix Request

In Graph definition properties for the Y1 and Y2 axis pages. removing the axis description and then clicking on the list would readd the description

---> C00689 - General - Usability

Made a change to the graph legend for the primary Y axis fields. Any fields (besides the 1st one) that are all 0 in the graph data, are no longer shown in the list of primary Y axis fields for easier readability of the available data. These fields are basically hidden, even though they still exist in the Primary Y axis fields and if you change the query or scroll (if non zero data is encountered) they will become visible again in that section of the legend.

Primary Y axis fields that are all zeros are now listed in the Available fields section but are greyed out. Trying to add them to the primary Y axis will do nothing because it's already part of the Primary Y axis internally.

If there is interest I could add a preference to have this part of the legend work either the old way or the new way. Let me know if you are interested.

---> C00689 - General - Bug Fix Request

Writing an SQL query in the SQL editor against an SQL alias should work properly now . Previously the GUI would prompt you for a member and fail.

---> C00687 - General - Enhancement

Added a new preference on the Misc tab called "Confirm system name when opening new SQL editor", that determines whether or not the list of available systems are shown when opening a brand new SQL editor.

---> C00684 - General - Bug Fix Request

Fixed a problem for users on Citrix servers when the iDoctor GUI tries to determine if it has been installed correctly and fails. The registry settings that identify the client installation were incorrectly placed in HKEY_CURRENT_USER instead of HKEY_CLASSES_ROOT

---> C00683 - General - Bug Fix Request

The right-click option to open a graph or table into a new data viewer, wasn't working when the preference "always open a report into an existing data viewer" was checked.

---> C00683 - General - Bug Fix Request

Fixed a problem with the creation of user defined queries and graphs over the SQL table summary files. They would not be usable on any collection other than the current one.

---> C00682 - General - Usability

When changing the time interval size in a graph, the graph will now be repositioned to the 1st bar or start of the graph. Previously the position was the same bar # as before, which didn't match where you were in the graph.

---> C00682 - General - Enhancement

The Windows menu on the Main Window, now contains the list of open Data Viewers, which allows you to open any of them from that menu.

---> C00681 - Installation - Bug Fix Request

Fixed a problem in the installer, where the abort install option wouldn't cancel the install right away.

Also made a change so that the error from windows is shown in a popup when unable to create the setup log file.

---> C00679 - General - Bug Fix Request

Fixed a bug where the Filter libraries menu on the Libraries folder didn't work if used from the list side instead of tree side of a component view.

---> C00678 - General - Bug Fix Request

The SQL editor doesn't seem to get refreshed with the correct sql statement after the query definition interface is used from a table view.

---> C00674 - General - Bug Fix Request

Library descriptions were not being displayed if the user , scrolled the list of libraries using the mouse wheel instead of the scrollbar.

---> C00669 - General - Bug Fix Request

Fixed a possible long delay /client freeze when doing a flyover when the data in the flyover contained a large amount of text. It will now be truncated to a more reasonable size.

---> C00667 - General - Enhancement

Added a new component tentatively called the Plan Cache Analyzer.

It offers the ability to run a collection containing plan cache and SQL query data from the system based on several options from the QQQOOOCACH API.

On the server side a new library QPLANCACHE will be restored when installing the server portion of iDoctor. This library contains a command STRPCCOL that is used to start a collection.

On the client side, a wizard is available to start the collection and if the analyze option is taken a set of graphs ranking the plan cache statistics is available.

---> C00667 - Installation - Bug Fix Request

The MS Visual Studio redistributable should install silently now.

Sep 17, 2008 – C00664

PEX-Analyzer changes

None

General iDoctor GUI changes

---> C00664 - General - Bug Fix Request

Fixed a crash that would occur if the user tried to rerun an SQL statement using the Execute menu from the SQL editor while the query was still running.

---> C00663 - General - Bug Fix Request

Fixed some problems with the alternate views functions from rankings graphs.

1. Duplicates entries would sometimes show up in the list.
2. The wait bucket description listed wouldn't change as it should based on the last bucket selected.
3. The current time range selection could sometimes be lost when choosing an alternate view graph.

---> C00662 - General - Enhancement

For all graphs in all components, added a function available by right-clicking on the graph that let's you change any parameters defined in the SQL statement such as the taskcount, asp number or time range. This is available under the "Change SQL Parameters" menu.

---> C00662 - General - Bug Fix Request

When looking at a vertical bar graph that has a secondary axis consisting of only a single bar, the 2nd axis was not visible. With this build, a point is drawn in the middle of the single bar at the height appropriate for the axis.

August 22, 2008 – C00659

PEX-Analyzer changes

---> C00657 - Pex Analyzer - Bug Fix Request

Fixed problems with job/task selection in the PEX Analyzer analysis wizard. This was causing some reports like stats hier to show no data.

---> C00656 - Pex Analyzer - Enhancement

Added the Generate Reports function to PEX Analyzer. It works against all of the SQL based summary reports.

---> C00656 - Pex Analyzer - Bug Fix Request

When summarizing a PEX collection that has call stack format 2 events, made a change to the stored procedure QIDRPACSTK that runs against the call stacks, so that it no longer returns an error when performing the insert, for event types where no data is found.

---> C00656 - Pex Analyzer - Bug Fix Request

Made a change to the STRPACOL command at 530 and higher, so that the library name the collection is to be created in is saved onto the PEX session's userspace description. This was needed in order for PEX Analyzer's Active collections view, to work appropriately with STRPACOL initiated collections.

Previously trying to end collections created using the GUI, from the active collections view, could cause problems. The ENDPEX command should not be used to end collections started with STRPACOL. ENDPACOL command should be used instead.

---> C00654 - Pex Analyzer - Bug Fix Request

Fixed a new bug where the Generate reports option for the SQL-based TPROF in PEX Analyzer wouldn't run.

---> C00654 - Pex Analyzer - Usability

Fix a usability issue in PEX Analyzer where the Call stack tab could be shown, even though no call stacks were collected.

---> C00654 - Pex Analyzer - Enhancement

Expanding a PEX collection will now display a folder called "SQL-based Analyses" which contains the SQL-based reports that were previously only accessible from the collection popup menu.

Also renamed the Analyses folder under a collection to Classic Analyses.

Also the popup menu option for a Collection, Analyze Data, has been renamed to Classic Analyses to coincide with folder it relates to.

---> C00653 - Pex Analyzer - Enhancement

Added the PEX Speed Team's PDIO SQL based analysis to the iDoctor client under the PDIO analysis menu for a collection. Just the starting point tables are provided currently.

Note: This analysis only works currently at V6R1.

---> C00649 - Pex Analyzer - Usability

In PA 530+, in the events analysis reports added checking to ensure that event category menu options are only shown if the collection contains them.

---> C00648 - Pex Analyzer - Enhancement

The event analysis graph, taskswitch events for thread XYZ now contains the raw cpu used for the thread as the Y2.

The other graphs of this type include the field in the SQL statement (and legend) but does not show it, because depending on the events collected it will give an incomplete picture of CPU consumption.

---> C00648 - Pex Analyzer - Enhancement

In PA 530+, changed the code that retrieves a PEX 16 level call stack to be SQL based, replacing the use of an RPG program.

Also from the new SQL based TPROF reports, made a fix so that the call stacks are accessible for reports that show the (starting) instruction address.

From the old tprof analysis, on the tprof graph, fixed a bug where the correct stack for the program name clicked on was not being shown.

---> C00647 - Pex Analyzer - Enhancement

In PA 530+, added a menu called "TPROF analysis" for trace collections that contain PMCO events that allows the user to build or view the SQL based TPROF reports that were created by the PEX speed team.

---> C00646 - Pex Analyzer - Enhancement

In PA 530+, in the events analysis reports, added drill down reports to show the raw pex data records based on the current context. (selected time range, job/thread and/or events).

---> C00646 - Pex Analyzer - Enhancement

In PA 530+, in the events analysis reports, added event rankings drill down graphs that show the individual events occurring for any of the event types.

---> C00646 - Pex Analyzer - Enhancement

In PA 530+, for the new SQL based reporting options, added checking to only show the menus that apply to the current PEX collection. (example: only show taskswitch menu if collection contains taskswitch events).

---> C00646 - Pex Analyzer - Bug Fix Request

In PA 530/540, the report Collection overview -> Collection thread summary is working now.

---> C00645 - Pex Analyzer - Enhancement

For V5R3+ PA, added a set of reports/graphs under the Heap Analysis menu for a collection which can be used if the collection contains heap events.

Let me know if you want changes.

---> C00645 - Pex Analyzer - Enhancement

In PA 530 or higher, added a step to the summary program that will produce data for heap events equivalent to that provided by the HEAPTOOL/RTVHEAPSTK command.

The 3 files generated may be viewed using the collection menu: Collection files -> Summary -> QAIDRPA_HSTACKS, QAIDRPA_HSTACKS2, and QAIDRPA_HSTACKS3,

The 16 level fmt 2 call stacks are also produced by the summary step and can also be viewed under the same menu.

---> C00644 - Pex Analyzer - Enhancement

In several places in the new PEX Analyzer reports, when an OBJVALUE field is available (either a taskcount, job name, generic job name, etc). The following drill down reports are available.

All events for <<OBJTYPE>> <<OBJDESC>>
Taskswitch events for <<OBJTYPE>> <<OBJDESC>>
PMCO events for <<OBJTYPE>> <<OBJDESC>>

Let me know if other similar graphs for different sets of events are desired.

---> C00644 - Pex Analyzer - Enhancement

In PA 530+, under the collection events analysis folder, the high level reports available now includes:

Collection events summary
Collection events detail by thread
By time interval
 Collection events overview
 Etc
Event totals
 Event totals by <<OBJTYPE>>

---> C00644 - Pex Analyzer - Enhancement

In PA 530+, added additional fields from file QAYPETASKI to the Collection thread summary report.

Task Database CPU
Storage pages allocated
Storage pages deallocated
Database and non-db page faults
Times added to collection
Java Thread Name
Home node
Affinity level
Affinity identifier
Preferred affinity resource identifier
Scaled running (CPU) microseconds at last start
Scaled running (CPU) microseconds at last stop
Scaled running (CPU) time
Subsystem name

---> C00643 - Pex Analyzer - Enhancement

In PA 530 and higher, when summarizing a collection, any Format 2 events that provide 16 level call stacks will be summarized in a new file QAIDRPA_STACK_<<COLNAME>>.

This indicates the number of call stacks that occurred for each type of call stack.
To view the PEX summary files, use the Collection menu -> PEX collection files -> summary.

To view only call stacks for a certain event type, set the EVTTYPE and EVTSUBTYPE fields in the call stack summary file. (example evttype = 3 and evtsubtype = 8 for PMCO call stacks).

---> C00643 - Pex Analyzer - Enhancement

In PA 530 and higher, updated the pex collection property pages that collection jobs and tasks so that the data is correctly sorted by CPU time descending.

Also removed the % fields as the new taski graphs can be used instead to see this type of information.

---> C00643 - Pex Analyzer - Enhancement

In PA 530 and higher, added additional drill down options from the event analysis queries to show the events by job grouping for a specific event type. These reports are still a work in progress.

---> C00642 - Pex Analyzer - Enhancement

In PA 530 and higher, added a new set of reports/graphs over the QAYPETASKI file data. This is accessible from the Collection overview menu. The graphs available show CPU, PDIO, Reads and writes, disk page allocated and page faults for the collection by either thread, job, job user, generic job, subsystem or pool.

---> C00642 - Pex Analyzer - Enhancement

I had to modify the PEX Analyzer events summary file in order to better handle timestamp selection drill downs. The PA summary step must be redone in order for some of the PA Events Analysis queries to work like the events overview graphs.

---> C00642 - Pex Analyzer - Enhancement

Variable width bar mode is working now for the PEX events analysis graphs. Also added drill downs into event by thread graphs for the selected time range from these graphs.

---> C00641 - Pex Analyzer - Bug Fix Request

Sorting a query that contains a tab before the "order by" instead of space, would cause a syntax error.

---> C00641 - Pex Analyzer - Bug Fix Request

In PA 530/540, the events analysis drill downs weren't working right. Also fixed SQL error with the events analysis by time interval graphs at those releases.

---> C00640 - Pex Analyzer - Enhancement

In PA V5R3 and higher, added additional reporting options under the Events analysis folder which replaces the existing events analysis and also includes a report showing the events property page into the Data Viewer. This is called the Collection events summary report.

This report offers two additional drill downs for the selected event type showing either the event counts over time, or the event counts by job/thread.

---> C00640 - Pex Analyzer - Enhancement

In PA V5R3 and higher, renamed the collection menu PEX events summary to Events analysis. Also added a menu Taskswitch analysis that contains the Collection Summary report built using SQL against the raw PEX files.

---> C00640 - Pex Analyzer - Bug Fix Request

In PA, the copy collection function didn't work at release V6R1.

---> C00639 - Pex Analyzer - Enhancement

Added a summary function to PEX Analyzer for V5R3 and higher trace collections. This function currently only produces an event interval summary file containing a record per second of the collection, with event counts that occurred for each event category and every event type. The file generated is QAIDRPAINTI_<Collection name>

Use the Summarize menu for a collection to access this feature.

When displaying collections in a library, added a Summarized field to indicate if the summary has been ran not or yet. You may need to use the Select fields... menu in order to add this field to the list of fields shown, if you have changed the fields shown in PEX Analyzer within a library to something other than the shipped default field list.

Once summarized a collection will have a new menu PEX event summaries containing a Collection events overview graph showing the event counts by category over time as well as alternate view graphs that break down each individual event within category over time. These graphs support changing the time interval size and will be used for some future drill downs into the PE X trace events using time range selection

---> C00634 - Pex Analyzer - Enhancement

When right-clicking a PEX collection added a set of menu options that will allow you to quickly open one of the PEX database files.

The menu options are

PEX core files - for the most commonly used files

PEX definition files - files that store some of the values from the PEX definition

PEX advanced files - the rest of the files
PEX summary files - displays files created by the summarize option

---> C00634 - Pex Analyzer - Enhancement
In Pex Analyzer, changed the contents of a PEX collection so it now contains 4 folders

Analyses - the analyses within the collection
Server-side output files - the pex collection and SMTR* files
User-defined queries
User-defined graphs.

This allows you to access the PEX files without generating one of the analyses.

---> C00633 - Pex Analyzer - Enhancement
The folders available under PEX Analyzer has significantly changed. Instead of showing libraries containing PEX data, the following folders are available:

Libraries - Libraries containing PEX collections (filterable)

Active collections - Work with the active PEX sessions on the current system. This is like an ENDPEX equivalent view. Since super collections use STRPEX/ENDPEX it was necessary to have this view to monitor the status of the PEX portion of a super collection.

PEX objects - A list of all PEX management collection objects on the system

Super collections - Work with the super collections on the current system

Definitions - Available definitions found on the current system

Filters - Available filters found on the current system

Scheduled jobs - iDoctor jobs scheduled to run

General iDoctor GUI changes

---> C00658 - General - Enhancement
Added a preference to the misc tab, called "Confirm when stopping the Report Generator function". When checked Closing the Data Viewer or Component View performing the report generator function will show a confirmation prompt before the report generator is stopped.

---> C00658 - General - Bug Fix Request
Fixed a bug in the report generator. If a user started the report generator, and while it's running closed down either the Data Viewer or the component view it was started from, the iDoctor client would crash.

---> C00658 - General - Bug Fix Request
Fix an issue with the transfer library function where it would not prevent a user from trying to transfer their library to the current system.

---> C00658 - General - Usability
When a graph is running or returns no data, the graph area will contain the graph title as well as the status of the graph. Made this change, so it would be captured and available for users of the Generate Reports function.

---> C00658 - General - Usability
If the Preference on the Misc tab - Maximize reports opened into Data Viewers is checked, then all views including property sheets (call stacks) and job logs, opened in the Data Viewer will also be maximized. Made this change because otherwise when opening a call stack or record quick view, all graphs would no longer be maximized.

---> C00658 - General - Enhancement
The iDoctor Update History window, now offers buttons to find the next or previous occurrence of the desired search string.

---> C00656 - General - Bug Fix Request

Scrolling through a table or graph should be slightly faster now. 2 fetches were being done per page instead of just 1.

---> C00652 - General - Bug Fix Request

Made a change to the ODBC driver/connection properties so that extended dynamic package support is no longer being used.

I did this because of the somewhat random SQL0440 errors on some stored procedures called by iDoctor because of a possible bug with SQL packages at V6R1.

---> C00652 - General - Usability

If the required server code isn't found, updated the message so it's phrased as a recommendation to install the latest server builds rather than a requirement and lists the functions changed/added by the recommended server code.

---> C00649 - General - Bug Fix Request

When saving user-defined queries/graphs that utilize summary files (which contain SQL long tables that contain the collection name as part of the name) they will be usable against other collections in other libraries or the same library .

The collection name will be resolved when the query is opened. Previously these queries would only be usable if accessed against the same collection.

---> C00648 - General - Enhancement

On the report generator window, added options to save or load the current list of reports to generate based on a user-defined name.

Also added a button to toggle the graphs available from checked/unchecked.

---> C00646 - General - Bug Fix Request

Fixed a possible crash that could occur when opening a graph if the code that displays the graph flyover is invoked before the thread that retrieves the data completes.

---> C00645 - General - Bug Fix Request

When scrolling through large tables, sometimes data would incorrectly appear out of sequence because the fetched buffer from ODBC could get into a state where it would contain the wrong data every other fetch. The solution of reissuing the same fetch again to get the correct data works, but does cause extra network delay. Need to investigate if the actual cause is iDoctor code, ODBC driver or system DB PTF level.

---> C00640 - General - Enhancement

Added a preference to the Misc tab called "Maximize reports opened into Data Viewers" which will auto maximize all reports opened into Data Viewers. The default is off/unchecked.

---> C00640 - General - Enhancement

Moved the preference "Always open a new report into an existing Data Viewer" from the Job Watcher tab to the miscellaneous tab. It also should be working properly from collection right-click context menus for all components.

---> C00640 - General - Enhancement

Removed the disable caching preferences on the PEX Analyzer and Job Watcher preferences tabs and created a single preference on the miscellaneous tab to replace it that applies to all components.

---> C00639 - General - Enhancement

On the main window, the clock icon to set the preferred time range interval size is now settable from any component. Made this change since now PEX Analyzer will also have time range graphs.

---> C00637 - General - Usability

Since a lot of new users are familiar with the common practice of left clicking a column twice to sort in descending sequence, this will now work in iDoctor. Previously left clicking would always sort ascending.

Note: You can right-click the column to sort descending, instead of left clicking twice.

---> C00637 - General - Bug Fix Request

When viewing a spool file (job log) from the server in the GUI, corrected some formatting problems.

---> C00636 - General - Usability

Made a change to the scheduling window so that the checkbox at the top is checked by default the 1st time the window is shown in order to show the calendar and time fields.

---> C00636 - General - Usability

Made a change to the scheduling window so that the checkbox at the top is checked by default the 1st time the window is shown in order to show the calendar and time fields.

---> C00635 - General - Enhancement

If the graph legend is not currently shown (via toolbar button) or set to 0 width in preferences, expanding the graph or data viewer should not display the legend.

---> C00634 - General - General

Moved all iDoctor stored procedures into library QIDRGUI instead of library QUSRSYS.

---> C00633 - General - Bug Fix Request

The synchronized table view wasn't scrolling down correctly, it would snap back up after scrolling down. Scrolling via the graph scroll bar would scroll the table correctly though.

Apr 30, 2008 – C00618

PEX-Analyzer changes

No changes

General iDoctor GUI changes

Fixed a bug where it was possible to get graph tooltips to show up through the main window or other application windows.

Synchronized table disappears when SQL Editor is opened/closed

When opening or resizing a graph view while the synchronized table view is minimized, the status bar would incorrectly display a negative number when showing the number of bars/records visible of the total.

Updated the iDoctor Components window so when it is resized larger, the list of components resizes bigger and also the window cannot be resized smaller than its initial size.

Made a bug fix to query parsing where the libraries weren't being filled in from a user-defined query. The parsing failed because there was a new line and not a space before the GROUP BY.

In graph views, only show flyovers if the graph has the current focus. In other words if a graph is behind another window, don't show flyovers for the graph.

We are making this change to help alleviate problems we sometimes have with flyovers showing up when we don't want them to.

Automatically add single quotes for string fields entered into the Record Selection page of the Query definition interface.

When resizing graph views that have a synchronized table at the bottom, need to make sure the synchronized table doesn't become visible when resizing the graph larger unless the

button on the toolbar has been pressed to show the table. It tended to appear at undesirable times.

Can't get graph or table tooltips to show up if the Data Viewer is moved to the 2nd monitor in a multi-display environment.

Fixed a crash when using the Query definitions Manage interface, changing the local database, and then using the Manage definitions interface again.

Changed the component menu User-defined reports -> Query(graph) definitions -> Manage... to Manage/Export... to make it more obvious this is where the export function resides.

In the remote SQL Statement Status View, using the display job log option would often not work, returning the error, The job number for the ODBC job could not be retrieved.

Add the Default time range interval size setting to Preferences on the Display tab. This preference can also be set using the clock icon on the Main Window's toolbar. Also when working with graphs and using the time range interval size button, added a menu at the top "Set current as default" that will make the current setting the default for future graphs opened.

Apr 21, 2008 – C00612

PEX-Analyzer changes

In PA, if a collection's VRM cannot be determined removed an error message that would be show when trying to call the QvrmCOLLC1 program. These unusable collections will also no longer be visible in the GUI.

General iDoctor GUI changes

Opening a server-side output file that had key fields which happened to be SQL reserved words would cause an SQL error.

Fixed some crashing problems that would occur if the connection to the server is lost when either shutting down the client or when attempting to reconnect. As long as the user/password entered previously is still valid, there should be no noticeable effect when automatically reconnecting besides a slight delay.

When performing collection or monitor scheduling on a daily or weekly basis, it appears that subsequent calls after the 1st would fail because the command to be executed was removed from the control file that keeps track of the commands to execute. In order to correct this, I've changed the command attached to the scheduled job so it now contains the actual SMBJOB command to execute rather than a call to program QIDRSTRSCD. Program QIDRSTRSCD is only used now if the command string scheduled exceeds the job scheduler limit of 512 bytes. In addition command strings are only added to file QGPL/QAIDRJBSCD if they exceed 512 bytes and are not removed by calling program QIDRSTRSCD.

Canceling a long running query and then attempting to run another query from the same connection would give the error "Prepared statement S0 in use".

The query parsing is a bit smarter now when looking for table names inside a select statement. Changes made to eliminate redundancy.

When launching a component and building the list of libraries, made a change so that the library descriptions are not retrieved until you need to see them. This means that I don't

retrieve the library descriptions (for libraries not visible) until you try to scroll down. This offers a big startup performance improvement on some systems.

When closing a Data Viewer that contains actively running queries or graphs, it should no longer take two tries. Previously the 1st click on the X would close the views within the Data Viewer and the 2nd would close the Data Viewer itself.

Mar 27, 2008 – C00609

General iDoctor GUI changes

On the iDoctor components window, if the expiration is set to 12/31/39 (effectively *NOEXP when generated the code) it will be shown as "Never" instead now.

On some PCs, users would get the message when starting iDoctor, "the application configuration is incorrect". This was due to iDoctor not being able to find all required DLLs. They have been moved to a different location and this problem should no longer occur.

I'm enabling the query definition interface again, even if the statement contains comments or new line characters which can tend to break the SQL parsing which tells the interface how the query is structured. Several of the queries i tried worked fine though using the interface to change a query will mean any comments in the SQL editor will be lost (not permanently if saved, just for the current view). PLEASE NOTE: On many of the rankings charts the job name field you need to filter on is now called OBJNAME as part of the query rewrite for CSI and JW. If in doubt, double-click the X axis field in the legend to see the field name you need to perform your record selection on.

Mar 21, 2008 – C00606

PEX-Analyzer changes

iDoctor for iSeries PEX Analyzer has been renamed to IBM iDoctor for i5/OS PEX-Analyzer.

When starting a collection in iDoctor, from the component icon, the following default libraries will be used unless changed by the user. Disk Watcher - QDWDATA PEX Analyzer - QPADATA Job Watcher - QJWDATA

Updated the server code for 2 programs that used the CHGSYSLIBL command without handling a possible not authorized to use the CHGSYSLIBL command problem. This change applies to the QIDRPA (PEX Analyzer) and QIDRGUI libraries at V5R3 and higher.

When creating a PEX profile collection from the collection wizard, ensure that the program(s) specified exist before adding to the list of programs.

Fix a bug where call stacks could not be shown in PEX Analyzer if file QAYPEBASE was empty even if a valid QRECN was provided in the current record.

In PEX Analyzer, in the add pex definition interface on the event selection screen, the format type can now be specified for any event type at V5R4+. Previously it only worked for base events.

In PA, on the PEX Definition Wizard, on the Add Program Selection window, you can now specify *LIBL as the library or *ALL as the program.

On the Not in Collection tasks property page for a collection in PEX Analyzer, the total number listed was 0 when it should not have been.

In PEX Analyzer, V5R3 and higher, updated the server program that submits an analysis to batch so that it now specifies the OUTQ parm as *CURRENT and MSGQ parm as *NONE. Previously the SBMJOB command would fail if the user profile specified a MSGQ or OUTQ that didn't exist.

General iDoctor GUI changes

All graphs in iDoctor now contain an equivalent table view of the data behind the graph under the graph view. The table view under the graph is always initially hidden, but if you wish to compare the table data with the graph, you can open it up to see the data. Rather than relying on a separate query/connection to produce the table, the table view is based on the exact data from the graph that is already cached in memory. This allows the scrolling to be much faster than it would be otherwise. The table and graph are synchronized in their scrolling and selections. Scrolling the table scrolls the graph, and vice versa. Note that sorting the table will also sort the graph (which may not be a desirable thing to do for time based graphs.) This can have some usefulness on rankings graphs to quickly change the ranking based on another statistic. The SQL editor above the graph is tied to both table and graph, and making changes there refreshes both the table and the graph. I've also added a toolbar button to show/hide the table portion for a graph.

This client adds a button to the main windows' toolbar that allows the user to set the desired views built when opening a graph between 1 of 2 modes. The default way is to build a SQL editor, graph and synchronized table. If this button is pressed, then this changes to 2 graphs views where the bottom graph allows you to specify a different alternate view using the Data Viewer's alternate view toolbar button. Either graph can be changed using the alternate views or time range selection option to make comparisons. Scrolling is synchronized as long as the number of bars in each graph is the same. This works for CSI, DW and JW.

Fixed a crash that could occur if an iDoctor user queries a file containing a DBCLOB type field.

When scheduling a collection or monitor it's a two-part process, running the ADDJOBSCDE command and then adding a record to file QAIDRJBSCD that contains the command string tied to the scheduled job. Added a check to remove the scheduled entry from the system if the 2nd part failed.

The summary programs for JW, DW, HA and CSI are now run able for collections that were created on a different OS VRM. This means you can do a create job summary for V5R3 CS collections on a V5R4 system, or create a V6R1 JW summary on a collection on a V5R4 system. As part of this change the stored procedures generated in QUSRSYS are now named QIDRJW54*, QIDRCS54*, QIDRJW61*, etc

In the iDoctor application property pages for some components on the Server configuration tab, changed the section that displayed the required PTFs installed, to display the PTFs NOT installed if any are not currently installed. The PTF check is performed when showing this screen in order to accomplish this. Note: Another reason I'm doing this change is because the PTFs listed previously were not correct if the system was running V5R3M5 or V5R4M5.

When starting the GUI, the client will automatically download the latest PTF information from the external iDoctor website. The PTF lists are saved to the window's registry and when a user attempts to start a collection; the PTF list will be checked against those on the system by updating the appropriate data areas on the server to contain the latest PTF levels. Previously the automatic PTF checking was dependent upon the server builds being reinstalled to update the PTF data areas on the system. Note: This function requires that the iDoctor GUI have access to the Internet, and if this is blocked by a firewall or not possible, it will not work. The PTFs shipped in the data area on the server will be checked for instead which may not be the latest ones.

Fix a bug where gaps could sometimes appear in the graphs, when drag scrolling through them. If the connection behind the remote SQL statement status view is lost, another one will be automatically started in its place. Previously resolving this required restarting the client.

I've created a stored procedure that is used to speed up most of the SQL statements executed by iDoctor (including graphs/tables in the Data Viewer), by requiring less network traffic. This effected many functions in all components. The performance gains are greatest on slow connections.

When starting the client, a screen is shown to display the most recent updates made to iDoctor in a list. You can check the box to not show this screen again, and can also bring this back from the Miscellaneous tab of the Preferences interface.

In the iDoctor application property pages - General tab, removed the iDoctor Server version as it is redundant with the OS version now (as of V5R3)

There should be slight speed improvements when opening or scrolling through graphs and tables in the Data Viewer (and elsewhere). I optimized the part of the code that looks up field indexes by field name to use a keyed hash table rather than loop through an array.

When installing the client, added a check to see if any instances of the iDoctor client are open, and if they are give the user a message to close them and abort the install.

The change environment option from the Connections view didn't work.

For the Query properties tab for a table or graph, added an apply button next to the Report description field. This will change the title of the report or graph the query applies to. Also made the SQL statement field read only since the SQL editor should be used instead now.

The save library window now lets you specify a target release earlier than the current release when connected to V5R4 or V6R1 systems.

When pressing help buttons when working on a V6R1 system, you should now go to the V6R1 documentation. Also when pressing help buttons, the documentation will now be looked for in the iDoctor install directory, and if not there it will be opened from the external website (slower). You can download the documentation to the iDoctor install directory using the documentation installer download from the page http://www-912.ibm.com/i_dir/idoctor.nsf/documentation.html

Added a button on the toolbar of the Data Viewer to show/hide the table behind the current graph view. Also fixed the SQL editor button to behave appropriately with the new graph/table splitter view.

Corrected a performance problem with the way bar selection was implemented. The more total records in the query behind the graph, the slower the graph would be when making bar selections. This began to become noticeable > 100,000 records or so.

For all components, fixed a problem that would occur when changing the library names in the SQL editor for the files queried behind a graph view. The overrides were not being updated to use the new files, and the previous library was being used instead. This would cause misleading results.

The SQL parsing code should be a bit smarter now when dealing with queries that contain both single member SQL tables and JW physical files. Previously problems could arise when working with the SQL editor in some situations.

Fix a bug defining an SQL statement containing an alias name longer than 10 characters would cause the querying parsing to try to incorrectly issue an override against the alias.

Jan 30, 2008 – C00588

The PEX Analyzer server builds have been updated to correct a problem that would occur if the user was not authorized to run the CHGSYSLIBL command.

Corrected a problem where a graph would not be updated after modifying the SQL statement and changing the libraries listed in the query to a different library.

Nov 29, 2007 – C00584

The PEX Analyzer stats flat analysis for all jobs/threads is now available from PEX Stats Hier collections at V5R3 or higher. Also added the default menu for this analysis to V5R3+ stats hier collections. The latest PEX Analyzer server builds are required for this fix to work.

When working with PEX Analyzer graphs and you open a graph that is too complex and get an out of memory error, the error will be shown inside the graph view itself rather than in a popup box

Turning on graph patterns/hatchings in PEX Analyzer graphs, will now display the hatchings used in the legend. Previously the patterns were only shown on the graph itself and not the legend.

Corrected some problems with the flyovers for PEX Analyzer graphs, where some labels were not being shown correctly.

Fixed a problem where the PDIO call stack option was sometimes not available when it should have been

Some iDoctor supplied reports were not having their queries parsed correctly because they contained analysis files and PEX files in the same query which requires different member names to be used. An example is the ASM Event Details query for the Net Size Analysis

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.

Creating PEX collections from MGTCOL objects in the GUI will now delete the userspace in QUSRSYS created by the CRTPEXDTA command. Previously the existence of this userspace would cause the initial refresh to give an incomplete collection status which required a 2nd refresh.

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

A check has been added to the Add PEX Definition Wizard on the stats bracketing events option page. Specifying both *JVA and *PRC is prevented since this causes an error in the PEX collection when started. The options are mutually exclusive and cannot both be specified.

In the Add PEX Definition Wizard, removed the merge job data checkbox when creating a stats definition. This option has no meaning at 530 or higher since PDC forces stats to merge NO.

In the stats flat analysis we have removed file G_STATSF because it is redundant with the other relatively new files produced in that analysis that show job and program level stats breakdown.

To alleviate a customer problem, most major PEX Analyzer programs (the collector, and analysis commands) have been updated so that any other libraries above QSYS in the library list are removed from the current job's library list before attempting to run. Custom written versions of common commands like ALCOBJ, CPYF were causing these commands to fail because they were above QSYS in the library list and being executed instead of the expected ones.

In collection properties on the definition panel, the stats program bracketing events were previously not being displayed correctly.

When opening PEX Analyzer analysis output files , the order by of the select statement was not being properly set.

The merge yes criteria is no longer displayed on the PEX definition properties window. The MRGJOB parameter on the ADDPEXDFN string is also no longer shown in the GUI.

When determining which user-defined graphs/queries to display, the decision will now be based on the file name of the query instead of the analysis number associated with the query/graph definition. This change will allow a user to define a query outside the analysis reports, save it and still be able to associate it with the desired analysis based on the filename used in the query.

Removed command CVTPACOL and its CPP QIDRCVTCOL from the server builds since it is no longer used.

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 QAYPERUNIS < 0. Also update the # of events field displayed if the events reported in QAYPERUNIS < 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.

The client should now detect and display and allow actions against PEX MGTCOL objects that were created outside of the GUI using STRPEX, ENDPEX instead of STRPACOL.

Allow the ability to create the PEX collection files from a MGTCOL object using the same collection name as the MGTCOL object.

If CRTPEXDTA is issued outside of the GUI, the client should now be able to display those collections created in this manner. A second refresh of the list of collections may be required for these collections however.

The SAVPACOL and FTPPACOL commands now have parameters to specify the target release and data compression parameters to apply to the save file. These parameters will allow a user to transfer a PEX collection to a system of a different release or to compress their data before sending it over FTP. The transfer window in the GUI has also been updated to provide these two new parameters.

The FTP collection transfer functions in JW and PA should now work against a unix ftp server. The ftp commands like 'put' were being issued in uppercase and this did not work on that type of ftp server.

Updated the PTF list for PEX Analyzer at release 530, 535 and 540. We removed a PTF that should not have been listed as required.

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.

By customer request, add QAPGMDESC file to the list of server side output files shown in a collection in both Job Watcher and PEX Analyzer at release V5R3 and higher.

July 13, 2006 – S00126-C00496

The TAPE_IO definition type will now contain the events *TAPDONE, *TAPSNDCMD and *SAVRST.

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

Fix to server programs that compute CPU time and lpar partial processor delay time in taskswitch analysis.

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

In the collection wizard on the advanced options page, provide the ability to wrap the collection if the max collection size has been reached.

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

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

Updated the required V5R3M0 and V5R3M5 PTFs.

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"

Fixed a bug in the GUI with showing the status of collections having greater than 20 million events. Sometimes a message would be displayed with a "receiver variable too small to hold result" error.

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

Help text added for the CPYPACOL, DLTPACOL, FTTPACOL, SAVPACOL, RSTPACOL green screen commands.

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 19, 2005 – S00112-C00454

Updated the required V5R3 PEX Analyzer PTFs to MF33172 MF36743 MF36702
MF33681 MF36689 MF34457 MF34994 MF37038

October 17, 2005 – S00111-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.

If overwriting a collection that already exists, any existing analyses will be first deleted before the new collection is created.

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 V5R3 are: MF33172 MF33255 MF33747 MF34889 MF33292 MF35005 MF34457 MF35302 SI16573

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

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

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

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

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

These commands require SECADM/ALLOBJ

Authority checking for PEX Analyzer added to the GUI before starting a collection. At V5R3 this check will be for *SERVICE OR access to function group QIBM_SERVICE_TRACE. This check has also been added to the STRPACOL 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.

Made a change to how the client determines the QRECN for the V5R3 PEX Analyzer call stack properties window from the G_TPROF output file. This should fix a problem where the call stack was not always visible.

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 cwbntbl 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.

Added support to view V5R2/V5R1 PEX collections and analyses on a V5R3 system. The following are part of this change:

- Copying, saving, transferring collections of an earlier release than the current system release is now possible through the GUI.
- Call stacks for V5R2, V5R1 PDIO analyses can now be viewed on a V5R3 system.
- Transferring collections is no longer restricted to systems of the same release, but the system cannot be of an earlier release than the collection being transferred (can't transfer a V5R2 collection to a V5R1 system).
- Copying or restoring a collection will no longer give an error if the PEXAnalyzer control file doesn't exist.
- To prevent problems from occurring with single object collections, they will only be shown in the GUI if they were created at the same release as the current system.

Updated the TPROF analysis to contain 3 new reports. Hit Count Totals by Job-thread/task
Hit Count Totals by Pgm within Job-thread/task and Hit Count Totals by Pgm

Fixed a bug with the TPROF analysis files G_TPROF1 and G_TPROF2. These files were not created correctly if tasks existed in the collection data greater than 10 bytes long.

Added new iDoctor supplied queries for the Netsize analysis and the TPROF analysis.

March 29th, 2005 – S00075-C00391

The Taskswitch analysis should now run much faster than before.

March 28th, 2005 – S00074-C00391

Fixed the following issues with the server-side IO analyses: (1) problems with independent ASP unit numbers (2) problems with I/Os to unconfigured units (3) incorrect WAITSLEE times, rarely, if CPSW events occur during switched out times.

Fixed an overflow problem with Stats flat analysis. The CPU fields that are defined as 9 byte Zoned in files G_STATSFPG and G_STATSFJB are too small for large analyses. The old versions of these files that may exist on the system should be deleted before running this analysis again in the same library. A special check for this record format difference is not yet in place.

The RSTPACOL command refers to a program that does not exist when trying to copy the control file record. This is the final step of the process to FTP a collection to another system

Removed *STATSFLATN from the GUI collection wizard as an option and replace *STATSFLATY with *STATSFLAT. The new stats flat analysis provides merged and unmerged reports in the same analysis so no distinction for merge jobs yes/no is necessary at releases above V5R2.. Also updated the basic mode of the collection wizard to not ask the question about whether the user wanted merge yes/no.

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

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

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

January 13th, 2005 – C00367

Fixed an issue with the IOP names/types listed for some SAN drives in the SMTRMOD file.

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.

Fixed an issue in the events property page for a PEX collection where the formats are not listed for anything other than category BASEEVT.

Rebuilt QIDRPAS save file in the library QIDRPA to fix some broken commands due to the library name change from SMTRACE to QIDRPA.

Fixed a bug with the total samples value on the CPU totals page for PEX Profile collections at release V5R3.

Bug with calculating collection overhead removed on the CPU totals page for a PEX Stats collection at release V5R3. The number incorrectly includes jobs/tasks not in the collection.

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.

PEX Definition wizard incorrectly referred to command STRIDOCOL instead of STRPACOL.

The missed events field on the general page of collection properties is now visible.

August 16th, 2004 - C00342

New version released!

Call stack display and collection options for the TPROF and Taskswitch analyses.

Support added to the PEX Definition Wizard for all of the new ADDPEXDFN parameters (ADDTHDOPT, INTERVAL randomize, and event formats)

Support added for viewing/changing V5R3 PEX definitions.

New reports added for the statistical flat analysis.

New option allowing generic task names to be entered in the PEX collection and definition wizards.

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

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.

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.

Naming changes

New library QIDRPA replaces SMTRACE/QYPINT.

New library QIDRGUI replaces QYPBASE.

Commands renamed from *IDOCOL to *PACOL.

