IBM iDoctor for IBM i
Data Viewer

IBM iDoctor for IBM i Development Team

8 February 2022

Abstract

Provides in-depth coverage of all major GUI functions for all components at 7.2 and higher. This covers the Data Viewer interface in the iDoctor GUI.

Changes

8 Feb 2022 – Created new document to separate the documentation into different documents for ease of maintenance. This is section 6 from the previous version of the documentation.

Table of Contents

1 Overview

The Data Viewer is a frame window within IBM iDoctor for IBM i used for displaying tables and graphs over data on the system. You can have as many Data Viewer windows open at one time as you want. The data behind the views within a Data Viewer may come from any number of systems desired which allows for easier comparisons.

These views are manipulated in the same way as within the Main Window using groups of tabs (or they optionally can be tiled, cascaded, etc.). Tables and graphs are typically opened from the component views in the Main Window. There are also drill down options to open tables and graphs directly from the Data Viewer window instead of going back to the Main Window.

The data behind a table or graph in the Data Viewer is produced using an SQL query. These SQL statements usually are IBM-defined but the capability to define and open user-defined graphs and queries in the Data Viewer is also supported. The SQL statement behind tables and graphs can either be manipulated in the query definition interface or using the SQL editor found above every table or graph by clicking the SQL button on the toolbar of the Data Viewer.

In addition to running SQL statements to produce table and graph views the Data Viewer also has support for viewing spool files containing job logs.

An example of a Data Viewer showing a graph from Job Watcher is shown below:

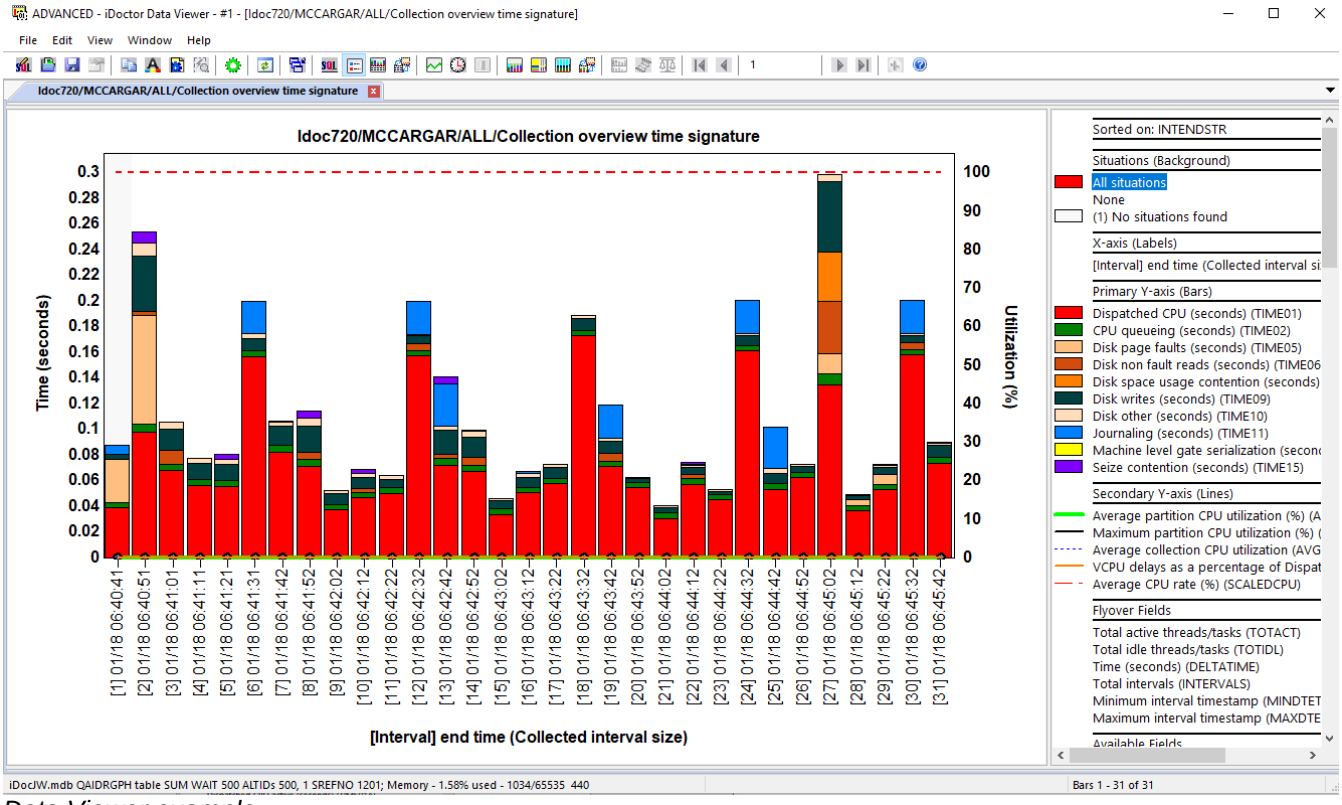

Data Viewer example

2 Toolbar

This section covers the toolbar buttons shown on the Data Viewer.

Note: Some of these options are only enabled depending on the current focus of the mouse pointer.

The Data Viewer Toolbar

The following describes the different toolbar buttons available in the Data Viewer and their purpose:

Button	Description
	**New SQL Query** Opens a new instance of an SQL Query View. The SQL Query View is used to create a query using Structured Query Language (SQL). The top portion of the view is an area where you can enter an SQL statement (also known as the SQL Editor) and the bottom portion is the result or output from the statement above.
	**Open File** This option allows you to open any library/file/member on the system using the Open File/SQL Table Window. This window lets you browse for the file or SQL table you wish to open.
	**Save As** This option allows you to either save the contents of a table view to a file or if viewing a graph to a .jpg image. When using this option on a table the entire contents of the table are saved. You can choose between rich text, comma separated, and tab separated text formats. If you wish to include/exclude the header in the saved file, see the Preferences -> File tab.
	**Properties** Use this option to view the properties for the current selection in the graph or table view.
	**Copy (Ctrl+C)** Copies the current selection from the current table, graph, or selection made within the SQL Editor to the clipboard. If a table view or SQL editor has current focus, this is only enabled when something has been selected. For graph views this will copy an image of the current graph (without the legend) to the clipboard. If you wish to include the legend in your image, then use Shift+Windows+S to draw a selection rectangle over the desired area of the window and copy it to the clipboard.
	**Set Font (Ctrl+D)** This button displays the Set Font window allowing you to set the font for the current view. This only works for tables, the SQL editor or when viewing spool file output. **Note:** This does not change graph fonts.
	**Edit Preferences** Displays the iDoctor Preferences interface.
	**Find (Ctrl+F)** This button displays the Find Window, allowing you to locate the specified text in the active view. **Note:** This is not applicable to graphs.
	**WRKACTJOB** Displays or hides the WRKACTJOB Pane. Unlike the command this view lets you work with active or inactive jobs on the system (to view spool files of completed jobs.) You can also graph some job statistics in limited ways.
	**Refresh (F5)** Refresh the currently active table or graph view.
	**Window Manager** This button will display the Window Manager which lets you work with a list of all views that are opened. This lets you find and activate the desired view/window or close one or more views quickly.
	**SQL Editor** Shows or hides the SQL Editor containing the SQL statement(s) behind the current graph or table view.
	**Hide/Show Legend** Shows or hides the graph's legend. This button will display the legend when it is pressed in and not show it otherwise.
	**Hide/Show Table** Shows or hides the Synchronized Table View under most graphs. This provides the raw data behind the SQL statement and graph currently shown. This table view is synchronized with the graph in all scrolling and selections made.

	**Toggle Table/Graph** When working with a graph, this button lets the user switch quickly to see either the entire table or back to the graph.
	**Alternate Views** Displays a list of available alternate views available for the current graph and/or table. Alternate views (in most cases) allow you to quickly redisplay the view in another way using the data already retrieved. In some cases, the SQL may need to be reran to produce the view.
	**Set Time Grouping** Modifies the time grouping used for the current time-based graph. The SQL statement behind the graph will be reran using a different time grouping, based on the choice made.
	**Time Filter** This option lets the user filter the graph by day of week and/or the time of day. This can be useful if you want to exclude weekends or certain hours of the day. **Note:** It is only enabled on some types of graphs (CSI Historical Summaries).
	**Normalize** Flattens the bars in a time-based graph where intervals are of varying durations. This option can be useful if there are wide variations in the time taken to produce intervals shown on the graph. For example, if Job Watcher was slow to initialize and the 1st 2 intervals took 10 times longer than the rest of the intervals, the rest of the intervals may be barely visible unless this option is turned on. When normalizing a graph, each value on the primary Y axis, is divided by the interval's duration applicable to each Y-axis value. This can create a "flattening" effect to bar heights by drawing time values based on relative contributions.
	**Variable-Width Bars** Toggles the use of variable width bar mode. Intervals that took longer to collect are drawn with wider bars than intervals taking less time. This allows you to see if the collection did not collect intervals at a consistent rate. Typically, this can happen if the system is overburdened and the collection itself cannot be performed optimally.
	**Show/Hide Idle Waits** Shows (or hides) the idle waits in wait bucket graphs in CSI and JW. This toggle will show all wait buckets including ones that are typically not of interest or revert to the original set. **Note:** You can configure which wait buckets are considered "idle" and which are not in the Wait Bucket Preferences.
	**Toggle Graph Format** Changes the graph format for the current graph. Typically, this is used to toggle a rankings graph between vertical bars to horizontal bars. It is less often used to toggle time-based graphs between lines and bars.
	**Side-by-side Comparison** This enables or disables Side-by-Side Comparison Mode. This option is only enabled if 2 or more graphs or tables exist in the current Data Viewer.
	**Synchronize Scrolling** Use this option to synchronize the scrolling of data while in Side-by-Side Comparison Mode. This means when both views will scroll together when one of them is scrolled.
	**Synchronize Y1-Scaling** Use this option to synchronize the Primary Y-axis scaling while in Side-by-Side Comparison Mode. This means when both graphs will use the same min/max values on the primary Y-axis.
	**First Row Set** This button will take the user to the 1st row of data in the current report. **Note:** iDoctor defines a block of X rows as a "row set.". By default, this is **100** rows but is configurable in **Preferences -> SQL -> Max rowset size**.

Previous Row Set
This button will move the current report to the previous row set of data. Typically, this moves the current record position 100 rows backward.

1

Record Position
This text field on the toolbar indicates the current scroll position (or row position) within the report being viewed. You can type a new value and press Enter into this field and the graph or table will move to that position.

Next Row Set
This button will move the current report to the next row set of data. Typically, this moves the current record position 100 rows forward.

Last Row Set
This button will take the user to the last row set of data in the current report.

Math Mode
Use this option to perform math functions over the selected rows in the current table or table beneath a graph. Clicking this button provides a menu where you can pick which function to perform (or none). After selecting an option placing the mouse over cells in the selected rows will show the results.

The possible math functions are

None – No math function is performed; this is the default.
Sum – Adds up all values for the current column's selected rows.
Average – Average of all values for the current column's selected rows.
Min and Max – Displays the min and maximum values from the current column's selected rows.
Percent of – Displays the percentage of the current cell's value of the total from the current column's selected rows.
Delta – Displays the difference between the rows indicated.

The following is an example of using the math function to add up the values in a column for the selected rows. The total is shown in the flyover.

[Interval] - timestamp (TIMEINT)	Interval end timestamp (INTENDSTR)	Interval number (INTERVAL)	Dispatched CPU (seconds) (TIME01)	CPU queueing (seconds) (TIME02)	Disk page faults (seconds) (TIME05)	Disk non fault reads (seconds) (TIME06)
[1] 01/18 ...	2022-01-18...	1	.0389	.0045	.0332	0
[2] 01/18 ...	2022-01-18...	2	.0980	.0064	.0846	.0032
[3] 01/18 ...	2022-01-18...	3	.0678	.0050	0	.0107
[4] 01/18 ...	2022-01-18...	4	.0563	.0048	0	0
[5] 01/18 ...	2022-01-18...	5	.0554	.0048	0	.0003
[6] 01/18 ...	2022-01-18...	6	.1571	.0043	0	0
[7] 01/18 ...	2022-01-18...	7	.0822	.0053	0	0
[8] 01/18 ...	2022-01-18...	8	.0716	.0050	0	.0059

25

About
This option displays the properties for iDoctor. This button performs the same action as the Help -> About menu.

3 Status Bar

The Data Viewer status bar provides additional details about the current view you are working with in the grey area at the very bottom of the window.

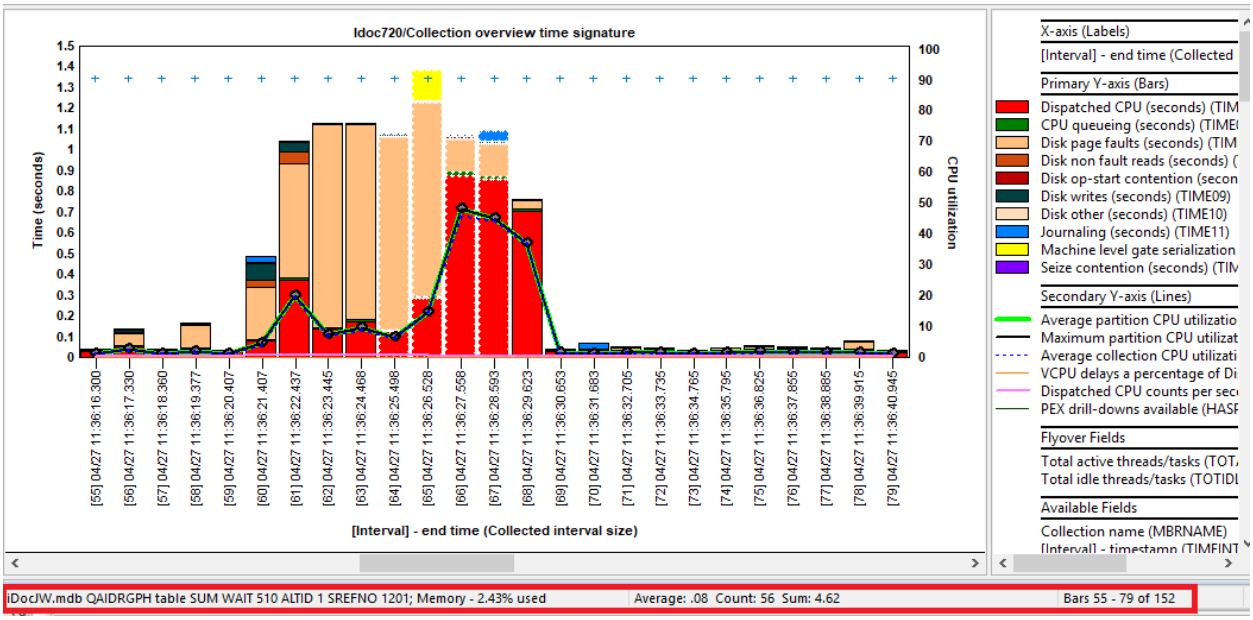

Data Viewer Status Bar (red box around it) and Graph

The status bar has 3 parts; some of which will only be filled based on the current view and status:

1) General Status Pane
2) Math Pane
3) Position Indicator Pane

3.1 General Status Pane

This portion of the status bar identifies different things depending on the current selection.

1) If an iDoctor graph is shown and the mouse is not over any data (over white space), then this information identifies the .mdb database (MS Access) and where in the database the report is located. The available memory (based on GDI objects) is also available here. For example:

   iDocJW.mdb QAIDRGPH table SUM WAIT 510 ALTID 1 SREFNO 1201; Memory - 4.72% used

   a. iDocJW.mdb identifies the database as iDocJW.mdb in the iDoctor install directory (typically C:\program files (x86)\ibm\idoctor)

   b. QAIDRGPH table SUM WAIT 510 ALTID 1 SREFNO 1201 means table QAIDRGPH where SQRYCAT = SUM and SQRYCATSUB = WAIT, GRAPHID is 510 and SREFNO = 1201 in QAIDRGPH and in table QAIDRSQL.

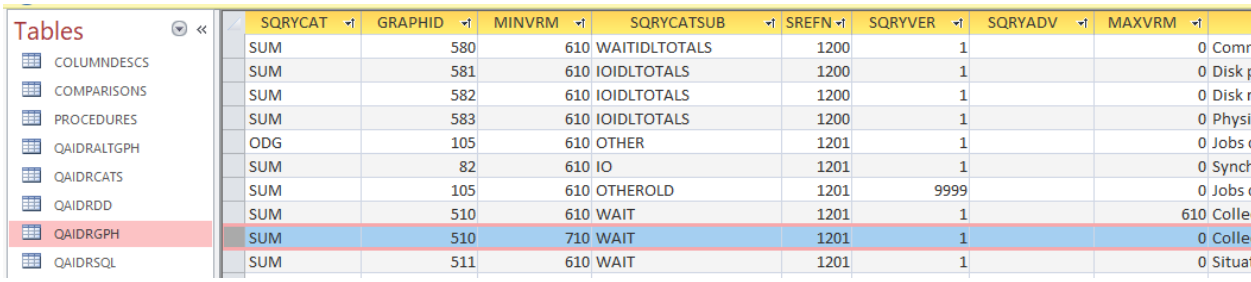

iDocJW.mdb QAIDRGPH table GRAPHID = 510, SQRYCAT = 'SUM' and SQRYCATSUB = 'WAIT'

c. The memory used is 4.72%. This is based on GDI objects and not based on normal RAM. You can increase this if you have access to change your Windows registry settings, by using the menu Edit -> Increase Windows GDI limit menu from the Main Window. This is a one-time step and increases the maximum GDI limit in all windows application from 10000 to 64000 objects. Because the graphing mechanism consumes large numbers of these, this may be desirable depending on how iDoctor is used.

2) If an iDoctor graph is shown and the mouse is over a selection, the general status bar pane contains the data relative to that point.

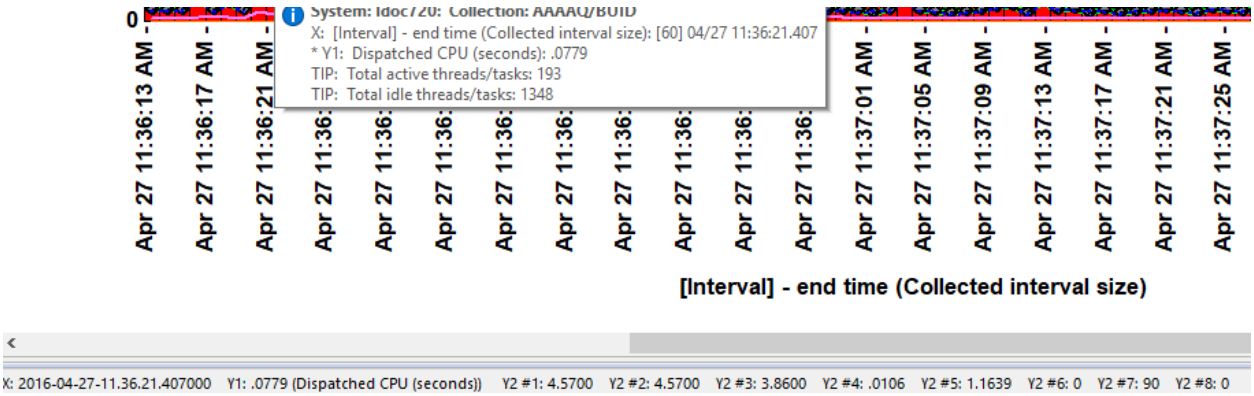

3) If an iDoctor report is shown, then this information identifies the .mdb database (MS Access) and where in the database the report is located.

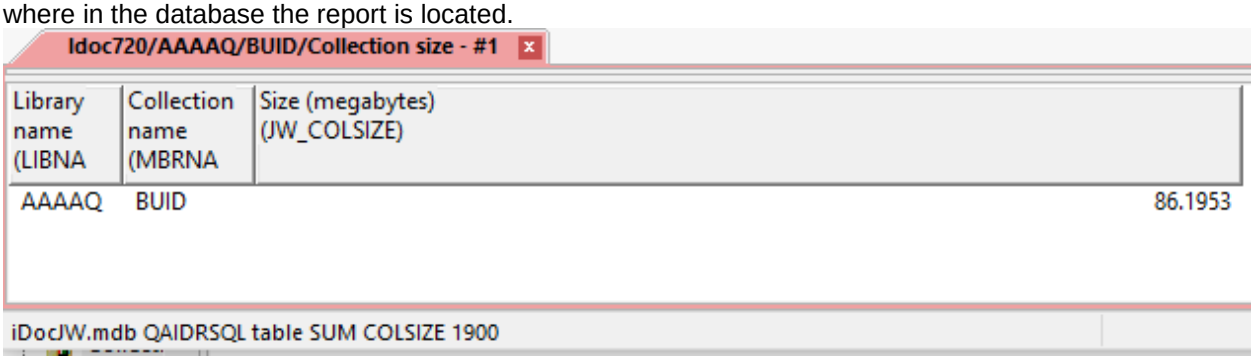

a. iDocJW.mdb identifies the database as iDocJW.mdb in the iDoctor install directory (typically C:\program files (x86)\ibm\idoctor)

b. QAIDRSQL table SUM COLSIZE 1900 means table QAIDRSQL where SQRYCAT = SUM and SQRYCATSUB = COLSIZE and SREFNO = 1900.

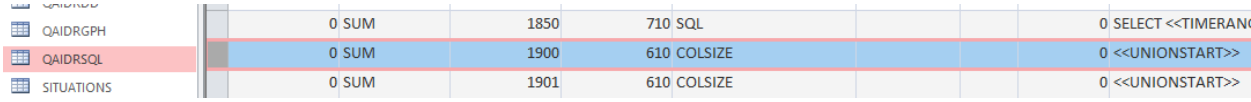

iDocJW.mdb QAIDRSQL table SREFNO = 1900, SQRYCAT = 'SUM' and SQRYCATSUB = 'COLSZE'

3.2 Math Pane

The Math Pane will only be filled in some situations where a selection has been made in a table or graph.

A typical usage scenario would be if the user needs to know the sum of several values in a table. By selecting the desired cells, the average, count, and sum of the values selected will be given.

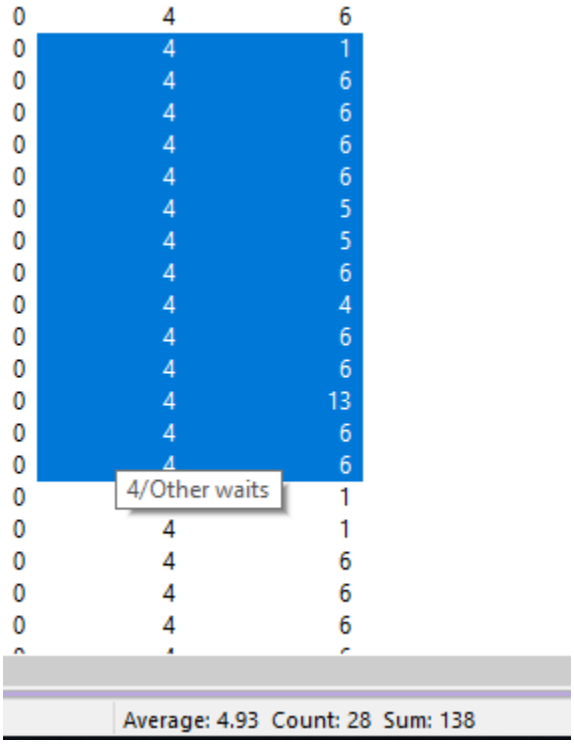

Math Pane Example

Selecting the entire rows of a table can also be used by the Math Pane and all non-numeric fields will be discarded from the calculations.

Note: A maximum limit of 150 rows and 50 columns applies. Exceeding this will cause no data to be shown in this pane.

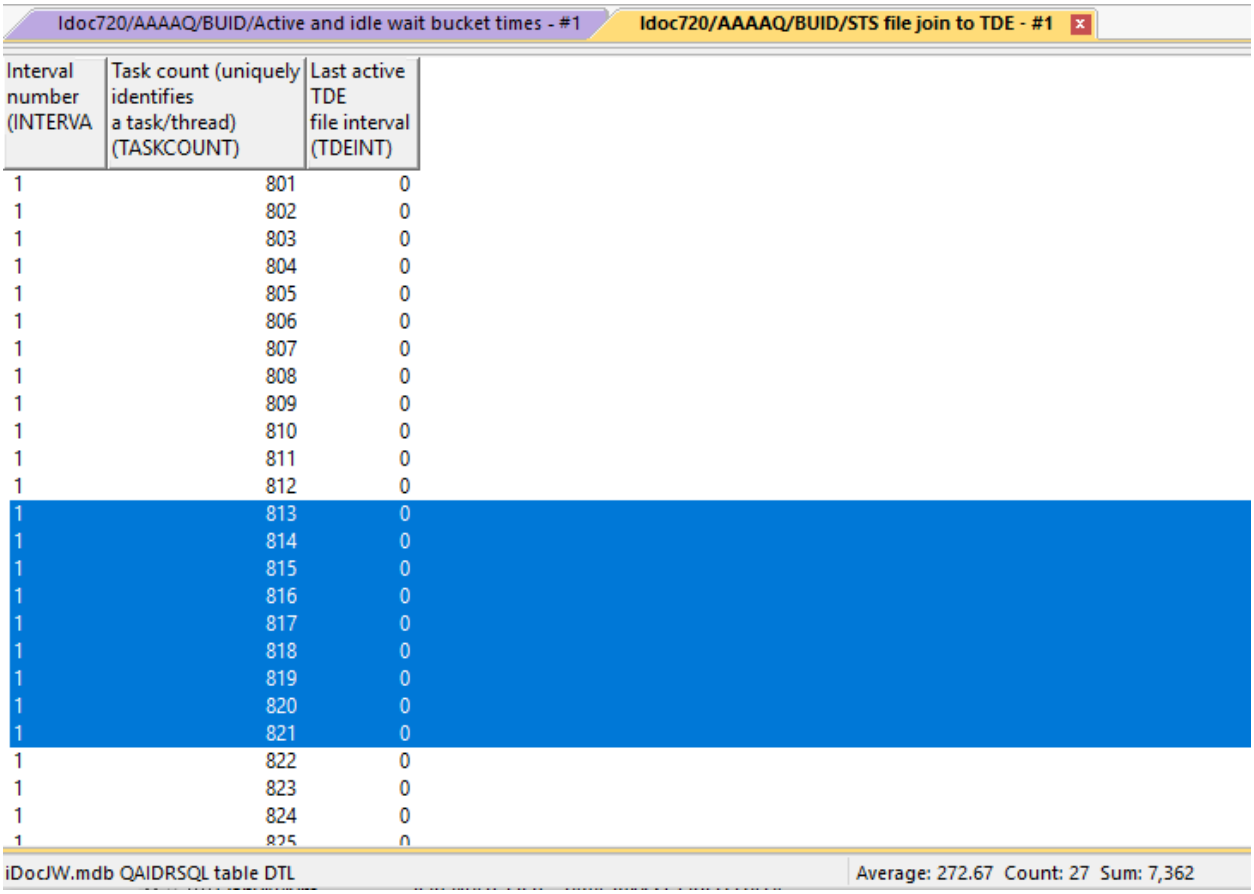

Math Pane Row Selection Example

This feature also applies to graphs when making selections. The numbers on the Y1-axis will be added up and shown in the Math Pane.

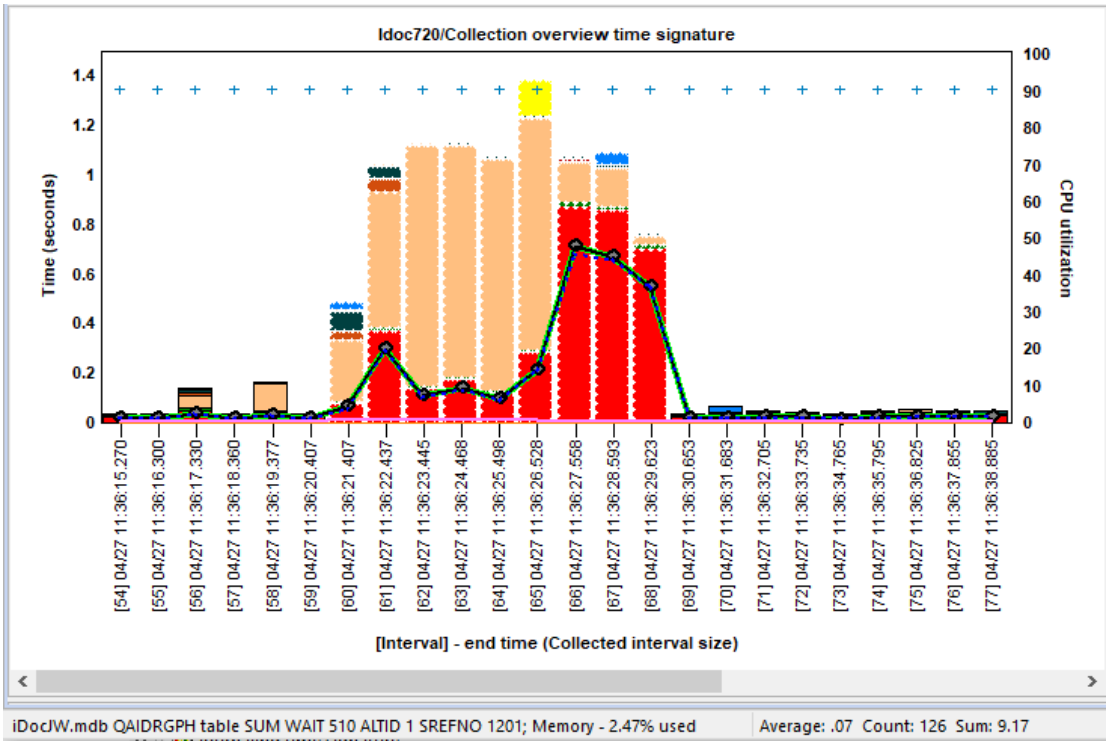

Graph Math Pane Example

3.3 Position Indicator Pane

This pane displays for the current view, the total records/bars and the scrolled location within that total.

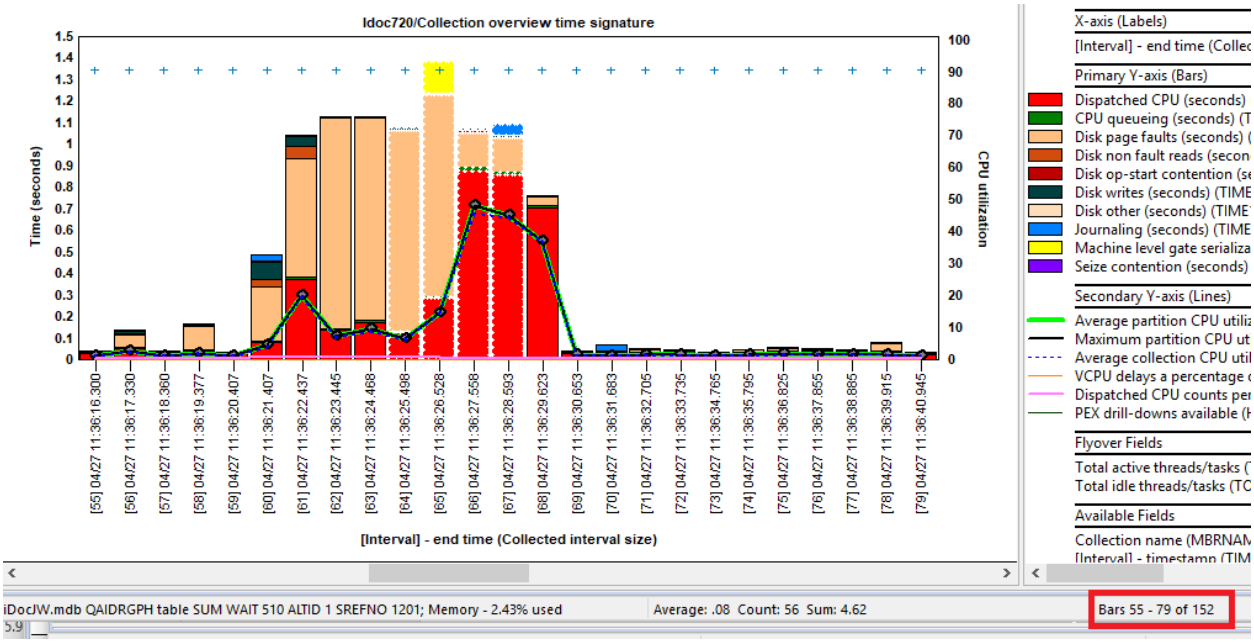

Position Indicator Pane Example

4 Menu Options

This section discusses the menu options available within the iDoctor Data Viewer. This only covers the menus available at the top of the Data Viewer window and does not cover the popup menus available within views displayed inside the Data Viewer.

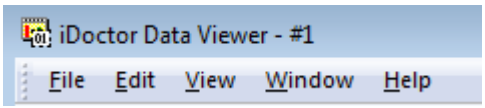

The Data Viewer Menus

The table below outlines the different types of menu operations that may be performed within the Data Viewer.

File Menu	Description
New SQL Query	Opens a new instance of an SQL Query View. The SQL Query View is used to create a query using Structured Query Language (SQL). The top portion of the view is an area where you can enter an SQL statement (also known as the SQL Editor) and the bottom portion is the result or output from the statement above.
Open File/Member	This option allows you to open any library/file/member on the system using the Open File/SQL Table Window. This window lets you browse for the file or SQL table you wish to open.
Save -> View As...	This option allows you to save the contents of a table view to a file. When using this option, the entire contents of the table are saved. You can choose between rich text, comma separated, and tab separated text formats.
Save -> Selection As...	The option allows you to save the **selected** contents of a table to a file. When using this option only the selected records or block of cells are written to the file. When using this option, you can choose between rich text, comma separated, and tab separated text formats. This option is not available for graph views.
Save -> Query Definition...	This option allows you to save the current table's query definition to the user-defined reports database.
Save -> Graph Definition...	Allows the graph definition behind the current graph view to be saved to the user-defined reports database.
Close	This menu will close the active view in the Data Viewer.
Print	This menu allows you to print the active graph view or contents of the SQL Editor. Before using this option, set focus into the desired SQL Editor or graph view by clicking inside the view. **Note:** This option is not currently available for table views.
Close Data Viewer	Use the menu to close the Data Viewer and everything in it. **Tip:** If queries are currently executing when performing this action, you may have to request this option a 2nd time since the queries will be canceled on the 1st attempt to close the Data Viewer.

Edit Menu	Description
Undo (Ctrl+Z)	Undo changes made in the SQL Editor with current focus.
Zoom Out (Alt+Z)	If working in a graph view and you've zoomed in already, this option will allow you to zoom back out 1 level.
Zoom Out All	If working in a graph view and you've zoomed in 1 or more times, this option will restore the graph to its original state.
Cut (Ctrl+X)	Cut the current selection from the SQL Editor to the clipboard.
Copy (Ctrl+C)	Copies the current selection from the current table, or SQL Editor to the clipboard. If a table view or SQL Editor has current focus, this is only enabled when something has been selected. For graph views this will copy an image of the current graph (without the legend) to the clipboard. If you wish to include the legend in your image, then use Alt-Print Screen to copy the Data Viewer and everything in it to the clipboard. You may want to maximize the current graph view within the Data Viewer before doing this.
Paste (Ctrl+V)	Paste the current text selection on the clipboard into the SQL Editor.
Find... (Ctrl+F)	This option opens the find window for the SQL Editor or table view **with current focus**. Click on a table or SQL editor first to give the desired view focus if necessary. When used in an SQL editor it allows you to quickly search for the next occurrence of a value in the SQL statement. This option allows you to reposition the current record position in a table view, based on some input you supply. The Find Dialog will be displayed, and you can use it to search for a specific value.
Find Next (F3)	Find the next occurrence of a value within the SQL Editor or table view. This option doesn't apply to graph views.
Find Previous (Shift+F3)	Find the previous occurrence of a value within the SQL Editor or table view.
Replace (Ctrl+H)	Displays a window allowing you to perform text replacement in the SQL editor.
Set Font (Ctrl+D)	This menu displays a window allowing you to set the font used for the table views in the IBM iDoctor for IBM i application. This option does not apply to the graph views. The font sizes used in the graph views are controlled in the Preferences window.
Preferences…	This menu displays a window allowing you to set user Preferences for the application.
Select All	Selects the entire contents of the SQL Editor or all rows in a table.

View Menu	Description
MDI Tabbed Style	Use this option to change the current MDI tabbed style being used. There are 3 styles of MDI tabs available in iDoctor: 1) None – this is a classic Windows MDI without tabs 2) Standard – allows users to tile and cascade but you cannot create groups of MDI tabs to compare with other tabs. 3) Grouped – Tabs cannot be tiled or cascaded but you can create groups of MDI tabs in order to make comparisons. 1) A Preferences option is also available which will open up the Preferences for the MDI tabs.
Report Visibility	This option is used to control the level of detail in the list of reports in 3 ways: Basic, Detailed, Advanced. Basic mode will show the fewest reports and Advanced will show all reports. For new users Intermediate is recommended. **Note:** You can tell which mode is in use by looking at the first word in the title bar of the Main Window or Data Viewer. 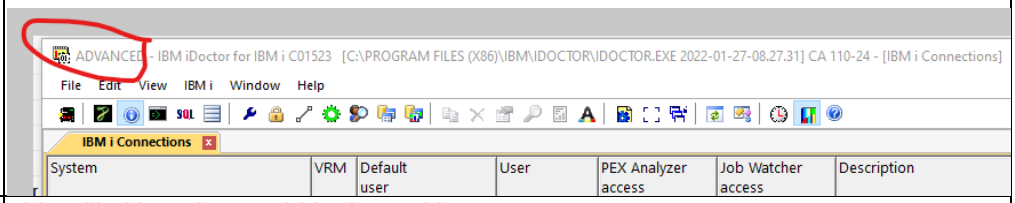
Toolbar	This will either show or hide the tool bar.
Status Bar	This menu will either show or hide the status bar.
Refresh (F5)	This menu will refresh the currently active view.
Resize Column Widths (F8)	This option will resize the columns shown in the currently show list view. F8 can also be used to do this action.
Field Names	This option will toggle the display of column headings so only the SQL given field name is shown. This applies only to table views.
Field Descriptions	This option will toggle the display of column headings to long field descriptions (if known). This applies only to table views.
Field Names and Descriptions	This option will toggle the display of column headings to both short field names and descriptions. This applies only to table views. Collection name (MBRNAME) | [Interval] - timestamp (TIMEINT) | Interval number (INTERVAL) | Interval end timestamp (INTENDSTR) | Minimum interval timestamp (MINDTETIM) BUID [1] 04/27 11:35:19.745 1 2016-04-27-11.35.19.745000 2016-04-27-11.35.17.779720 BUID [2] 04/27 11:35:21.805 2 2016-04-27-11.35.21.805000 2016-04-27-11.35.19.745001

Window Menu	Description
Cascade	Use this menu to rearrange all views in the Data Viewer in an overlapping sequence starting in the upper left corner of the window. **Note:** This option is not shown when the MDI tabbed style is set to Grouped.
Tile Horizontally	Use this menu to rearrange all views in the Data Viewer such that each view will have an equal distribution of the available height in the Data Viewer. The views will not overlap each other. **Note:** This option is not shown when the MDI tabbed style is set to Grouped.
Tile Vertically	Use this menu to rearrange all views in the Data Viewer such that each view will have an equal distribution of the available width in the Data Viewer. The views will not overlap each other. **Note:** This option is not shown when the MDI tabbed style is set to Grouped.
Close All	This option can be used to close all open tabs/views.

Help Menu	Description
IBM iDoctor for IBM i website	Launches your web browser and takes you to the iDoctor website.
IBM iDoctor for IBM i downloads	Launches your web browser and takes you to the iDoctor downloads page.
IBM iDoctor for IBM i documentation	Launches your web browser and takes you to the iDoctor documentation page.
About	This will display version information for the IBM iDoctor for IBM i client.

5 SQL Query View

The SQL Query View refers to the view you get when you press the SQL Query View button  on the Data Viewer toolbar. Pressing this button will initially provide you with an option to select the system name you wish to write an SQL statement against like the following:

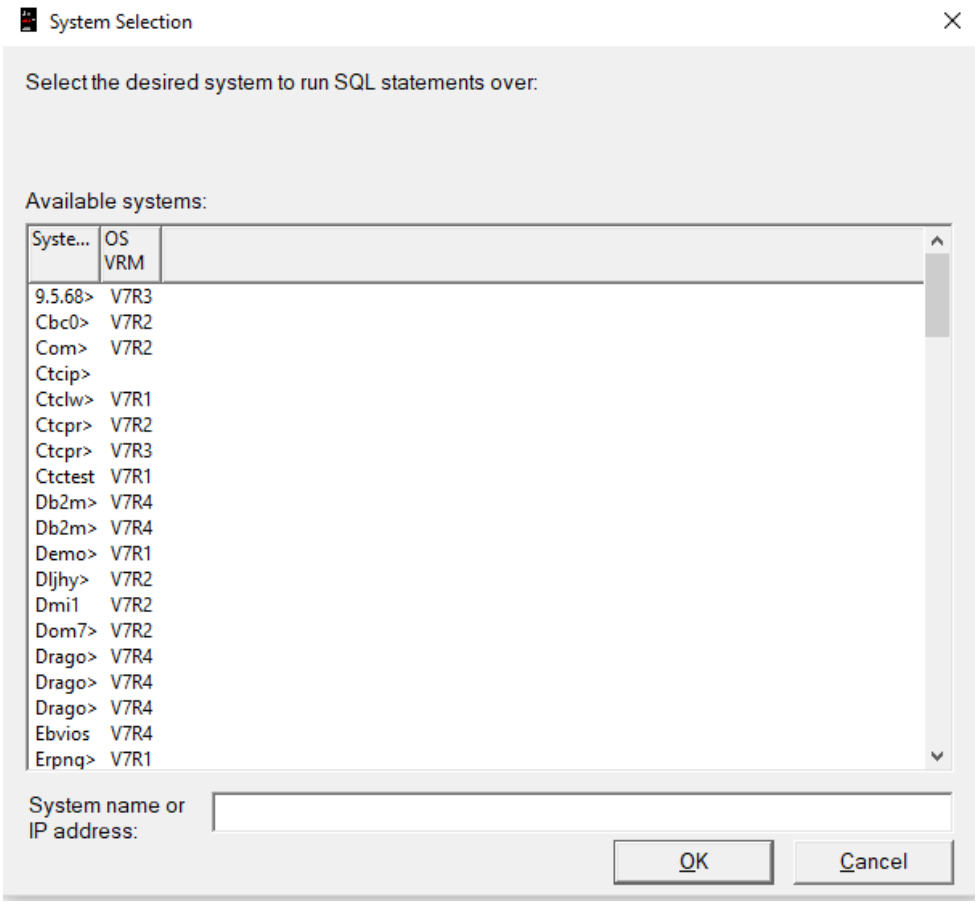

System Selection Window

Tip: If you do not want to see the System Selection Window and just have iDoctor use the current system you are working with, then uncheck the “Confirm system name when opening new SQL editor” option in Preferences -> Confirm.

After selecting a system, you will be presented with an empty SQL query view where you can write any SQL statement(s) you wish against the system name provided.

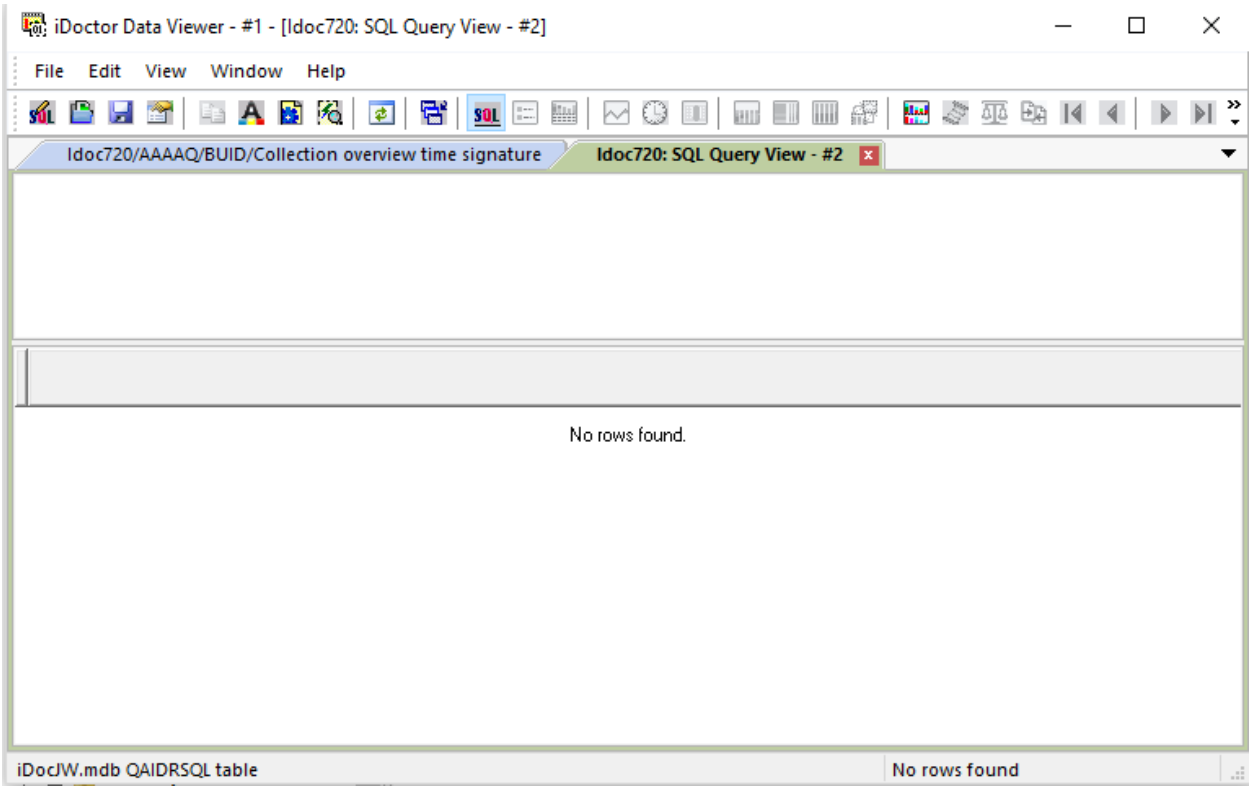

SQL Query View

At the top of this view is an SQL Editor and the bottom the results of the execution of that statement will be displayed. If a SQL select statement or WITH statement are used the results are shown as a table. See the next section for more information about the SQL Editor in iDoctor

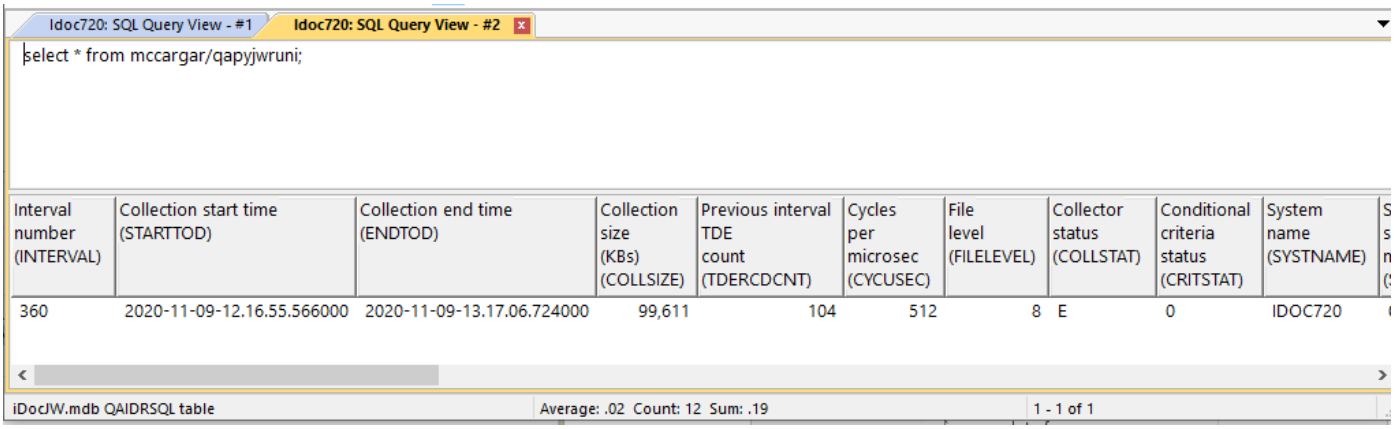

SQL Query View with data

6 SQL Editor

The SQL editor lets you execute and display the results of one or more SQL statements.
Note: If multiple SQL statements are to be ran, then each must end with a semicolon.

The top portion of the view is an area where you can enter SQL statement(s). This bottom view is either a table view or graph view or in the case of errors, the results and error messages are shown instead of a table or graph.

Note: All iDoctor SQL statements (and RUNSQL on green screen) by default use System naming convention and not SQL naming convention where “library/file” syntax is used and not “library.file”. Other tools such as Run SQL Scripts uses SQL naming convention by default. Use the convert menu options if necessary to copy and paste SQL statements between iDoctor and Run SQL Scripts.

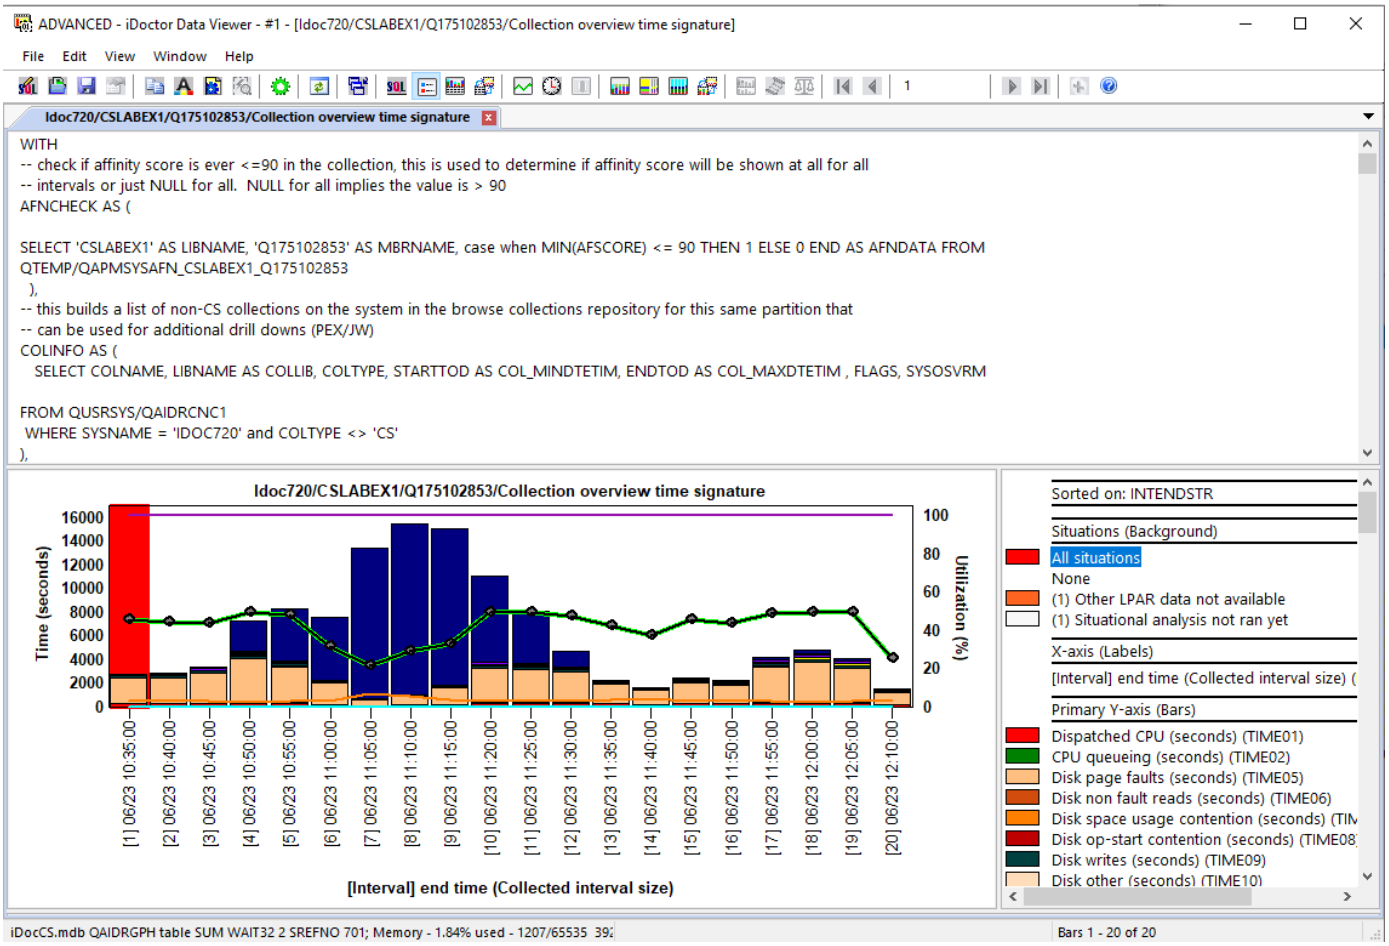

SQL Editor with attached graph

The queries you create with this view may be saved and restored for later use and their definitions can be viewed and manipulated using the query definition interface.

The SQL Editor is used to execute one or more SQL statements in the order that they appear. Each SQL statement must end with a semicolon. If you want the results to be viewable under the editor, the last statement must be an SQL WITH or SQL SELECT statement. Typically statements before the SELECT or WITH could be calls to SQL stored procedures or commands to create aliases or drop tables, etc.

To execute your SQL statement, right-click on the SQL Editor and choose the Execute menu or press the F4 key. This will cause the SQL statement(s) to be executed and the table or graph will be redisplayed.

Tip: If you wish to run a specific statement (out of more than 1), then click somewhere with that statement and right-click and use the Execute -> Selected menu option.

You can also add any comments to the SQL Editor using two dashes like in this example:

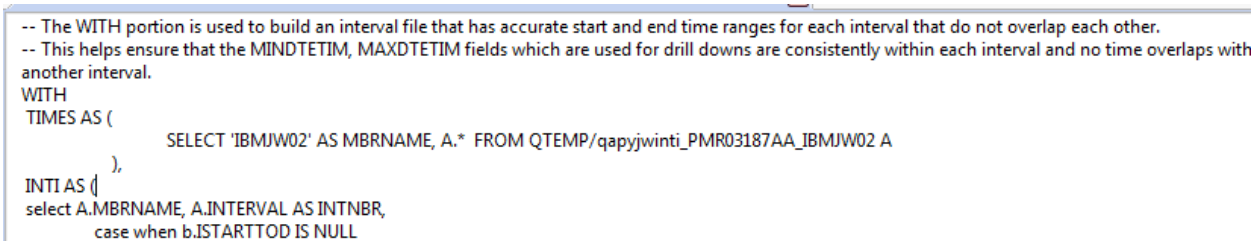

SQL Editor showing comments in 1st two lines

Most table and graph views in iDoctor's Data Viewer provide an SQL Editor in a hidden view above them. Some comparison modes do not provide this. You can show or hide the SQL Editor using the SQL button on the toolbar. You can also open a new SQL Editor with an attached table view called the SQL Query View, using the 1st icon on the toolbar of the Data Viewer.

Note: The SQL Editor can be used to run any SQL statements desired within the current QZDASOINIT job. This allows advanced users to define their own stored procedures, drop tables, and create tables or indexes using the SQL editor in iDoctor. For a history of the SQL statements issued and results, use the SQL message log (View -> SQL Message Log menu.) But to enable this logging you will need to use Preferences -> SQL tab and configure this there.

Note: With the January 2022 builds, the **OVRDBF** command is no longer used by iDoctor. Aliases must be created to point to the correct collection in multiple member physical files. By default, iDoctor uses QTEMP aliases having a naming convention of QTEMP/<LIBRARY>_<FILE>_<MEMBER>.

The following options are available in the SQL Editor's popup menu:

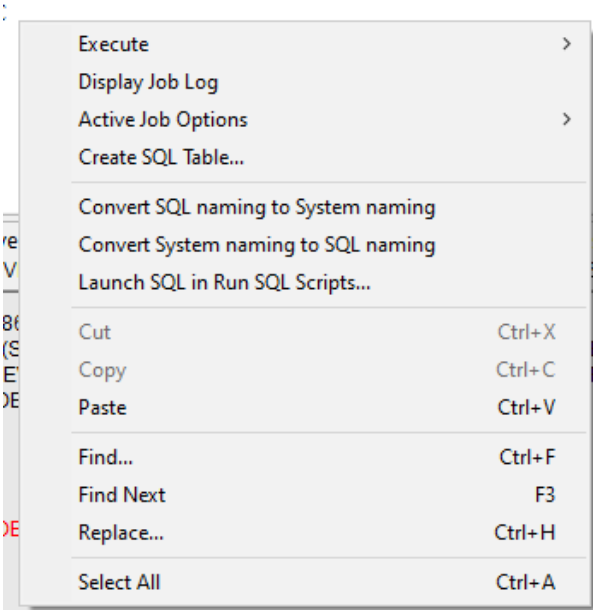

SQL Editor Popup-Menu

Popup Menu	Description
Execute -> All (F4)	Execute the SQL statement(s) within the SQL Editor. If multiple statements are used, then all statements are executed from top to bottom and each must end with a semicolon. If necessary, comment out any statements you do not wish to execute.
Execute -> Selected	Executes only the selected SQL statement(s) or where the mouse pointer has current focus if none are selected.
Execute -> All in batch	This attempts to execute all SQL statements in the Remote SQL statement status view. **Tip:** You cannot use this option to run SELECT or WITH statements there. Also, any QTEMP aliases may need to be recreated.
Execute -> Selected in batch	This attempts to run the **selected** SQL statements in the Remote SQL statement status view.
Display Job Log	Displays the job log for the QZDASOINIT job that is responsible for running the SQL statements shown in the current table or graph.
Active Job Options	This contains a set of options applicable to QZDASOINIT job processing the requests within this interface such as viewing the call stack, open files, or searching the job log. For more information, see the Main Window documentation.
Create SQL table	This option will create an SQL table on the IBM i based on the SELECT or WITH statement in the SQL editor. This can be useful if you wish to save the data provided by iDoctor to your own file for later viewing.
Convert SQL naming to System naming	This option can be used to modify the SQL statement so that all tables using SQL naming convention (library.table) will change to System naming convention instead (library/table).
Convert System naming to SQL naming	This option can be used to modify the SQL statement so that all tables using System naming convention (library/table) will change to SQL naming convention instead (library.table).
Launch SQL in Run SQL Scripts	This option will copy the contents of the SQL Editor into a temporary file and open it in the Run SQL Scripts tool provided in IBM i Access Client Solutions.
Cut	Cut the current selection from the SQL Editor to the clipboard.
Copy	Copies the current selection from the selection made within the SQL Editor to the clipboard.
Paste	Paste the current text selection on the clipboard into the SQL Editor at the current position.
Find	This option opens the find window for the SQL Editor allowing you to quickly search for the next occurrence of a value in the SQL statement.
Find Next	Find the next occurrence of a value within the SQL Editor.
Replace	Displays a window allowing you to perform text replacement in the SQL Editor.
Select All	Selects the entire contents of the SQL Editor.

If executing the SQL statement within the SQL Editor results in an error, the SQL Editor will show the error instead of the expected graph or table view results. Correct the error and rerun the statement to display the graph or table attached to the SQL Editor.

Example of an error shown in the SQL Editor

6.1 Create SQL Table

This option will create an SQL table on the IBM i based on the SELECT or WITH statement in the SQL editor. This can be useful if you wish to save the data provided by iDoctor to your own file for later viewing.

Note: This only works if the SQL editor contains a single SELECT or WITH SQL statement.

This will temporarily modify the SQL editor to contain an appropriate CREATE TABLE SQL statement then change it back to the original statement after execution if the Show CREATE TABLE checkbox is unchecked.

By using this option, you will be prompted for the library and table name where you wish the data to go:

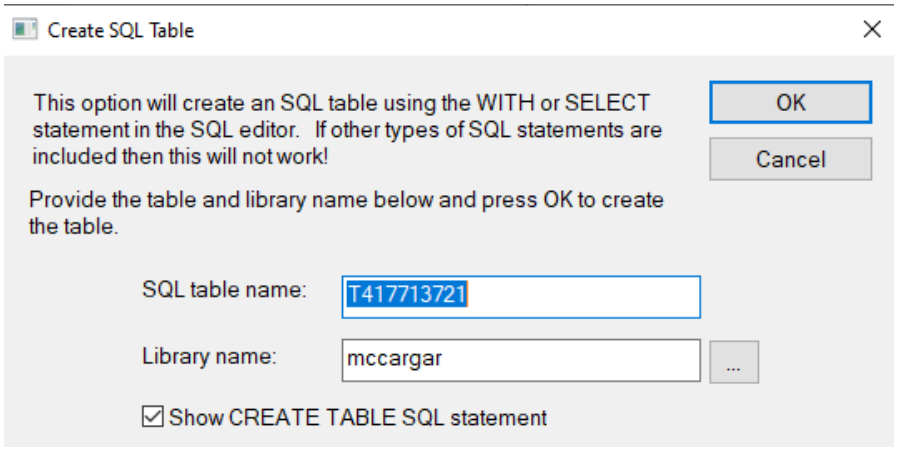

Create SQL Table Window

Afer successful use the graph or table view will show something like this to let you know that the table was created. **Note:** This feature requires client 1525 or higher.

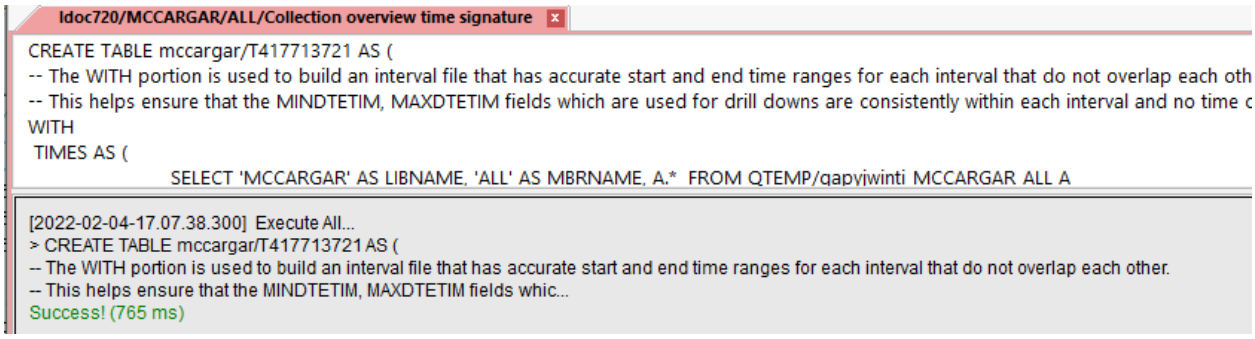

Create SQL Table Success

Tip: If the “Show CREATE TABLE SQL statement” checkbox is unchecked, then reexecute the SQL statement to restore the previously visible graph or table. If that option is checked, then rerunning the SQL statement will fail unless a new table name is used.

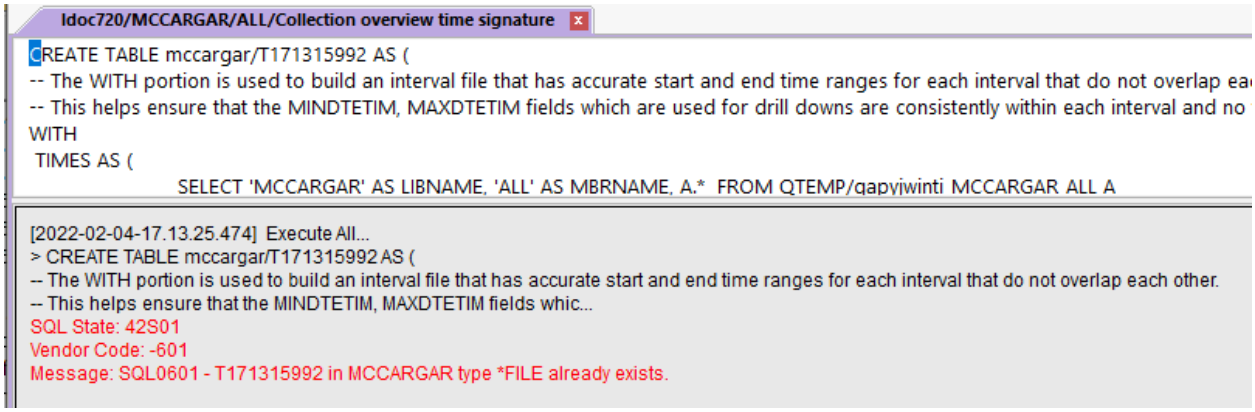

Create SQL Table Failure

6.2 Launch SQL in Run SQL Scripts

This option will copy the contents of the SQL Editor into a temporary file and open it in the Run SQL Scripts tool provided in IBM i Access Client Solutions.

Tip: If this option does not work and you have IBM i Access Client Solutions installed, then you have not yet installed 32-bit ACS on your Windows machine using the javascript provided in the install download. After unzipping the files, you will need to expand the Windows_Application folder and run **install_acs_32_allusers.js**.

Note: The 32-bit option is needed because iDoctor is a 32-bit application and it will be launching ACS Run SQL Scripts.

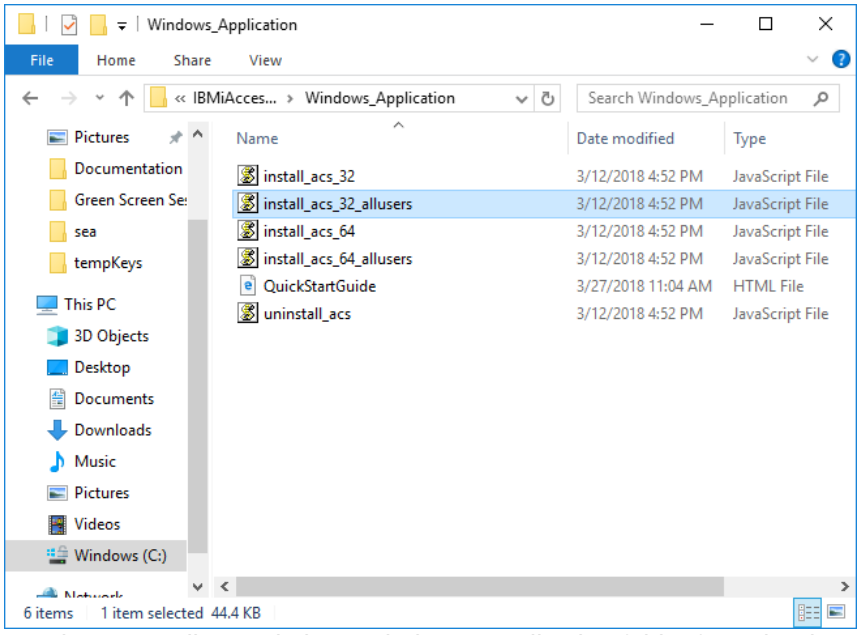

IBM i Access Client Solutions Windows_Application folder from the download image

Note: Run SQL Scripts uses SQL naming convention by default and iDoctor uses System naming convention. Use the Convert System naming to SQL naming convention option before using this menu option.

7 Open File/SQL Table Window

This option allows the user to open any library/file/member on the system. A window is displayed where the user can browse for the desired type of file and member to open.

The following 5 types of objects can be browsed and opened with this window:

1) SQL tables
2) Physical files
3) Logical files
4) Aliases
5) Views

Of these 5 types only physical files and logical files will potentially fill the list with members to choose from. SQL tables, aliases and views can be opened by double-clicking them without needing to select a member.

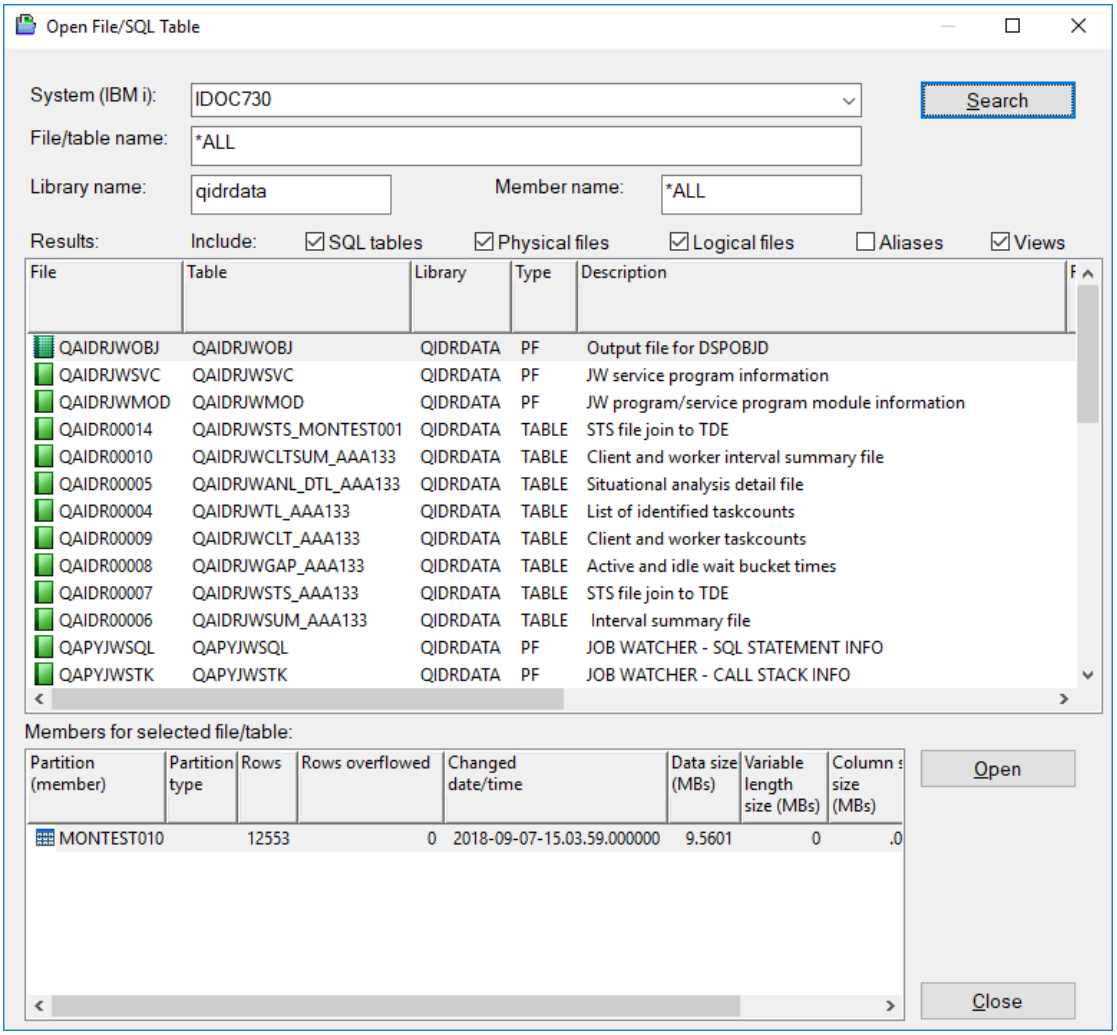

Open File/SQL Table Window

The following table describes the interface elements within this window.

Option	Description
System (IBM i)	The IBM i system to connect to and look for files to open. **Note:** This drop-down list of system names is the same as those found in the IBM i Connections View. Add a system there for it to appear automatically in this list.
File/table name	The name of the physical file, SQL table, alias, view or logical file to search for. This can be *ALL or blank to include all file names or it can be a generic file name like QAYPE*.
Library name	The name of the library name to search for tables. This can be *ALL or blank to include all libraries or it can be a generic library name like MC*.
Member name	The member name to search for within the Members for selected file/table list.
Include options	This list of checkboxes indicates which types of objects should be included in the list of results.
Results	This is the list of objects (files, aliases, views, etc) that match the filtering options above after pressing the Search button.
Search button	The search button will query the IBM i for the objects of interest based on the filters provided.
Members for selected file/table list	This is the list of members for the selected file in the results list (file list) and that also match the member name filter. **Note:** This list is not applicable if the currently selected type of object in the results list is VIEW or ALIAS. This list will always be empty for those types.
Open button	Opens the selected library/file/member or SQL table, alias or view in the Data Viewer.

8 Table Views

A table view shows data from database files via SQL statements executed on the system. The user can display many records in a table view and use the scroll bar to quickly move to the records desired via relative positioning.

The SQL behind the table view can be modified at any time using the SQL Editor. If you don't wish to modify the SQL either the Query Definition interface or Filters may be used instead. Filters may be added by right-clicking the desired column and using the Add Filter menu.

The data may also be sorted by clicking the desired column to sort by. Left click will sort a column in ascending sequence. Click again to resort in descending sequence. If you hold the down the SHIFT key the next column clicked will be added to the existing sort sequence. You can also change the sort order by modifying the ORDER BY clause of the SELECT statement in the SQL Editor. A final option is to use Sort menu found by right-clicking a column in the graph's legend.

Data in a table view may be selected for copy and paste to a file or to the clipboard. A set of records -or- a block of cells may be selected at any one time. Click the left mouse button and drag across the cells desired to make a block selection. Once a selection is made, use the Edit -> Copy (Ctrl+C) menu to copy the current selection to the clipboard. Use the File -> Save Selection As... menu to write the selection to a file.

Tip: Making cell/row selections will also fill the Math Pane in the status bar with the sum of all numeric values in the selection.

The record Position Indicator Pane in the status bar will show which records are currently being viewed out of the total possible in the active view.

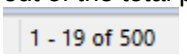

Position Indicator Pane

An example of a table view is the following:

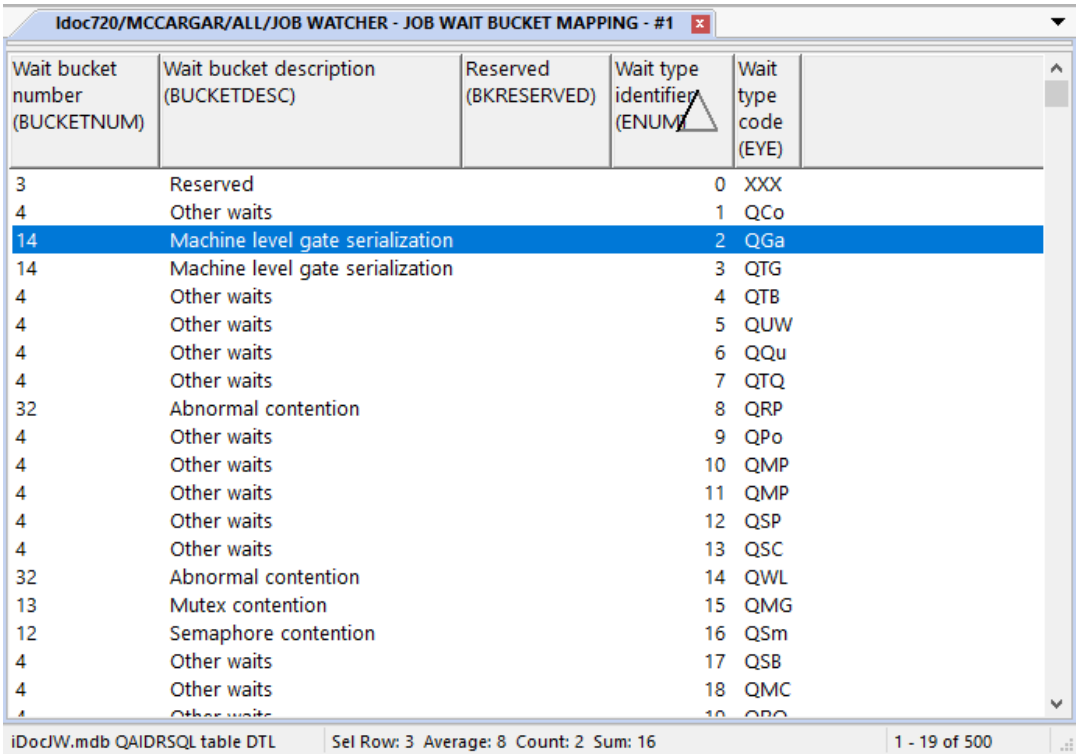

Table View Example

8.1 Row Menu Options

A popup menu is available by right-clicking on any row within the table. The following options are available:

Menu	Description
Record Quick View	This option will display a vertical view of the current row(s) selected. If multiples are selected this option can be used to show a comparison between two rows in a side-by-side view.
Copy	Copies the current text selection to the clipboard. This may consist of rows or block of cells. Make a block selection by holding down the left mouse button and draw a box.
Find...	This menu allows the user to reposition the current record in a table view, based on input supplied if matching information is found.
Save -> View As...	This option allows you to save the contents of a table view to a file. When using this option, the entire contents of the table are saved. You can choose between rich text, comma separated and tab separated text formats.
Save -> Selection As...	The option allows you to save the **selected** contents of a table to a file. When using this option only the selected records or block of cells are written to the file. When using this option, you can choose between rich text, comma separated and tab separated text formats. This option is not available for graph views.
Save -> Query Definition...	This option allows you to save the current table's query definition to the local user-defined reports database on the PC. Query definitions are saved into the user-defined queries folder under collections and can be reused. A window will be displayed asking for a description of the query and which component it applies to (if this is not already known).

Set Font	This menu displays a window allowing customization of the font used for all table views.
Preferences	This menu displays a window allowing the user to set customized settings for the IBM iDoctor for IBM i application.
Graph Definition -> Define New	This contains an option to create a new user-defined graph from the current report.
Query Definition -> Field Selection	This allows the user to modify the list of columns, rearrange them or change which ones are visible.
Query Definition -> Filters	This allows the user to define filters and apply them to the current report.
Query Definition -> Reset	This option allows the user to discard changes made to the SQL statement using the above query definition menu options.
Query Definition -> Save As…	Use this option to save the current SQL statement as a user-defined report.

Duplicate as table view	This option creates a copy of the current report as a new table view in the Data Viewer.
Properties	Displays the properties for the current report. The information shown in the property pages varies based on the type of report being viewed.
Search Google for 'X'	This option will open the default web browser and do a search on the contents of the table cell or column header you right-clicked on.

Depending on the type of report shown in a table view, other menus applicable to that report type will be shown. These are mentioned in the documentation for each of the components.

8.2 Column Menu Options

Right-clicking on a column provides additional options for filtering, sorting or hiding the column from view.

The following options are available:

Menu	Description
Sort descending	Changes the sort order of the SQL statement, removing existing fields in the sort and adding the current field in descending order to the sort sequence.
Sort ascending	Changes the sort order of the SQL statement, removing existing fields in the sort and adding the current field in ascending order to the sort sequence.
Edit…	This option lets you change the column description shown for the selected column.
Add filter...	Displays the Filter interface for tables and provides options for defining filtering based on the current selection (column, row) in the table.
Remove selected filter	This option removes the filter defined for the current column. When filters are defined, the column header text is shown in a red color.
Remove all filters	This option removes all filters defined in the SQL statement that were added by the Table Filtering Interface.
Hide	Removes the selected column from view.
Unhide all columns	Redisplays all columns that were previously removed using the Hide menu option.
Search Google for 'X'	This option will open the default web browser and do a search on the contents of the table cell or column header you right-clicked on.

8.3 Making Row Selections

Row selections are made by selecting (clicking on) the desired row.
If you desire to select a continuous range of rows, then hold down the shift key and click on the 1st row and then click again on another row. All rows between the 1st and last selection will be selected.

After the selection is made you can copy it to the clipboard by pressing the Ctrl+C keys or using the Copy  button on the toolbar. You can also export the desired rows to an Excel (CSV format) file using the File -> Save Selection As menu. In addition, your numeric fields will be added up in some situations and shown in the Math Pane.

8.4 Making Cell Selections

Cells in table views are selected in iDoctor by performing the following action:

1) Left-click and hold the button down on the desired (cell). The point clicked on should be the upper left position of the set of cell(s) desired to be selected.
2) With the button held down , move the mouse down and to the right. During this process a box is drawn.

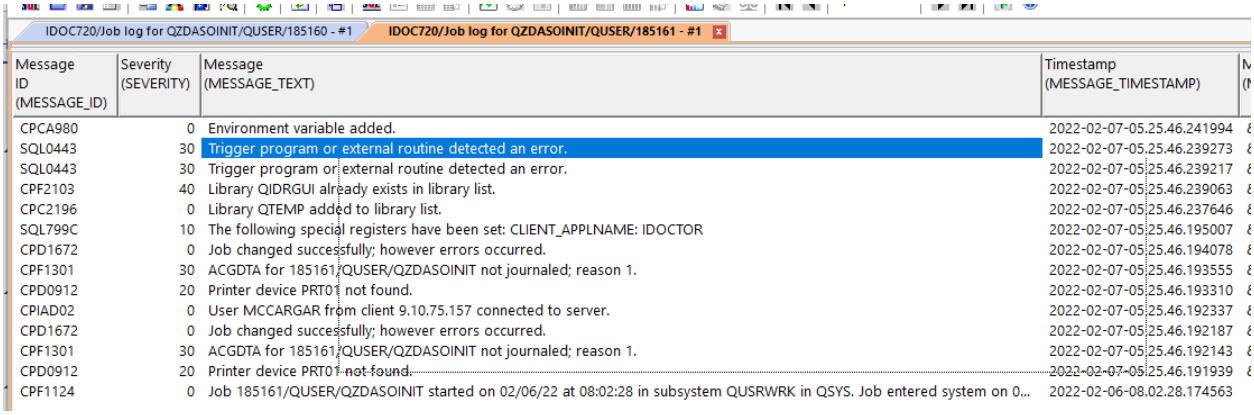

Message ID (MESSAGE_ID)	Severity (SEVERITY)	Message (MESSAGE_TEXT)	Timestamp (MESSAGE_TIMESTAMP)
CPCA980	0	Environment variable added.	2022-02-07-05.25.46.241994
SQL0443	30	Trigger program or external routine detected an error.	2022-02-07-05.25.46.239273
SQL0443	30	Trigger program or external routine detected an error.	2022-02-07-05.25.46.239217
CPF2103	40	Library QIDRGUI already exists in library list.	2022-02-07-05.25.46.239063
CPC2196	0	Library QTEMP added to library list.	2022-02-07-05.25.46.237646
SQL799C	10	The following special registers have been set: CLIENT_APPLNAME: IDOCTOR	2022-02-07-05.25.46.195007
CPD1672	0	Job changed successfully; however errors occurred.	2022-02-07-05.25.46.194078
CPF1301	30	ACGDTA for 185161/QUSER/QZDASOINIT not journaled; reason 1.	2022-02-07-05.25.46.193555
CPD0912	20	Printer device PRT01 not found.	2022-02-07-05.25.46.193310
CPIAD02	0	User MCCARGAR from client 9.10.75.157 connected to server.	2022-02-07-05.25.46.192337
CPD1672	0	Job changed successfully; however errors occurred.	2022-02-07-05.25.46.192187
CPF1301	30	ACGDTA for 185161/QUSER/QZDASOINIT not journaled; reason 1.	2022-02-07-05.25.46.192143
CPD0912	20	Printer device PRT01 not found.	2022-02-07-05.25.46.191939
CPF1124	0	Job 185161/QUSER/QZDASOINIT started on 02/06/22 at 08:02:28 in subsystem QUSRWRK in QSYS. Job entered system on 0...	2022-02-06-08.02.28.174563

Making a cell selection rectangle in a table

3) Release the mouse button at the desired point, all cells within the box drawn will be selected.

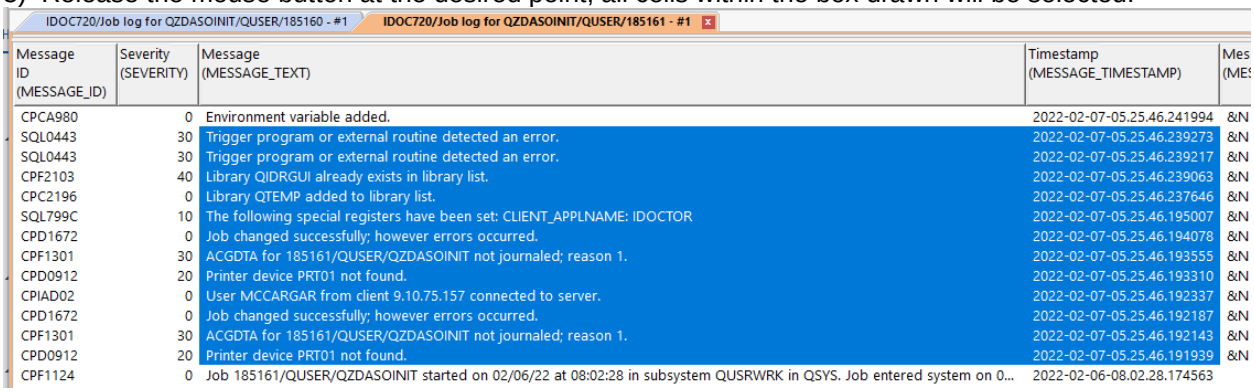

Message ID (MESSAGE_ID)	Severity (SEVERITY)	Message (MESSAGE_TEXT)	Timestamp (MESSAGE_TIMESTAMP)	Mes (MES
CPCA980	0	Environment variable added.	2022-02-07-05.25.46.241994	&N
SQL0443	30	Trigger program or external routine detected an error.	2022-02-07-05.25.46.239273	&N
SQL0443	30	Trigger program or external routine detected an error.	2022-02-07-05.25.46.239217	&N
CPF2103	40	Library QIDRGUI already exists in library list.	2022-02-07-05.25.46.239063	&N
CPC2196	0	Library QTEMP added to library list.	2022-02-07-05.25.46.237646	&N
SQL799C	10	The following special registers have been set: CLIENT_APPLNAME: IDOCTOR	2022-02-07-05.25.46.195007	&N
CPD1672	0	Job changed successfully; however errors occurred.	2022-02-07-05.25.46.194078	&N
CPF1301	30	ACGDTA for 185161/QUSER/QZDASOINIT not journaled; reason 1.	2022-02-07-05.25.46.193555	&N
CPD0912	20	Printer device PRT01 not found.	2022-02-07-05.25.46.193310	&N
CPIAD02	0	User MCCARGAR from client 9.10.75.157 connected to server.	2022-02-07-05.25.46.192337	&N
CPD1672	0	Job changed successfully; however errors occurred.	2022-02-07-05.25.46.192187	&N
CPF1301	30	ACGDTA for 185161/QUSER/QZDASOINIT not journaled; reason 1.	2022-02-07-05.25.46.192143	&N
CPD0912	20	Printer device PRT01 not found.	2022-02-07-05.25.46.191939	&N
CPF1124	0	Job 185161/QUSER/QZDASOINIT started on 02/06/22 at 08:02:28 in subsystem QUSRWRK in QSYS. Job entered system on 0...	2022-02-06-08.02.28.174563	

Completed cell selection

After the selection is made you can copy it to the clipboard by pressing the Ctrl+C keys or using the Copy  button on the toolbar.

8.5 Filter

The Filter window is accessed by right-clicking a column and using the Add Filter menu.
Filters can be defined on one or more columns and each filter will modify the SQL statement (within the where clause) to perform the desired filtering. The column headers that have filters applied are drawn with a red color.

This following is an example of the Filter window:

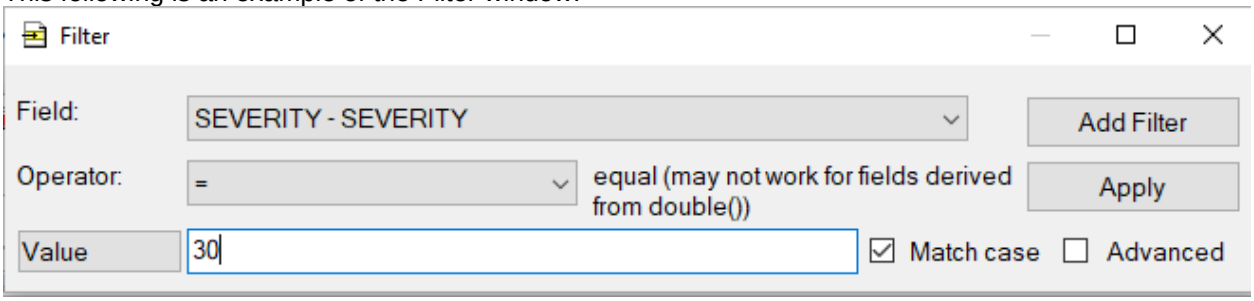

Filter Window

The following options are available:

Element	Description
Field	This is the desired field to filter on. By default, this is the same field that was right-clicked.
Operator	The type of operation to use for this filter. The possible values are: 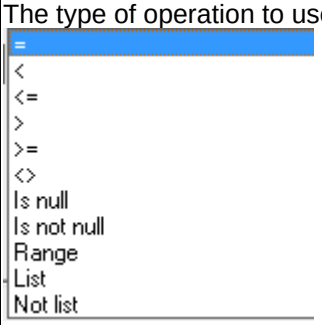 **Note:** the = (equal) operator may not work for fields derived from the double function.
Value	This is the value to apply to the filter. Generally, the filter is something like FIELD OP VALUE where OP is the operator, FIELD is the field to filter on and VALUE is a constant numeric or text string.
Add/Update Filter	This button will add or update the desired filter in the table. The change does not take effect until the Apply button is pressed or the table is refreshed.
Apply	This button will rerun the query behind the table and apply any changes made to filters defined.
Advanced checkbox	This checkbox will hide/display the additional options on this window. They are typically not needed unless you wish to define more advanced options. See the section on the Query Definition ->Filters interface for more information (these options behave the same as in that interface)

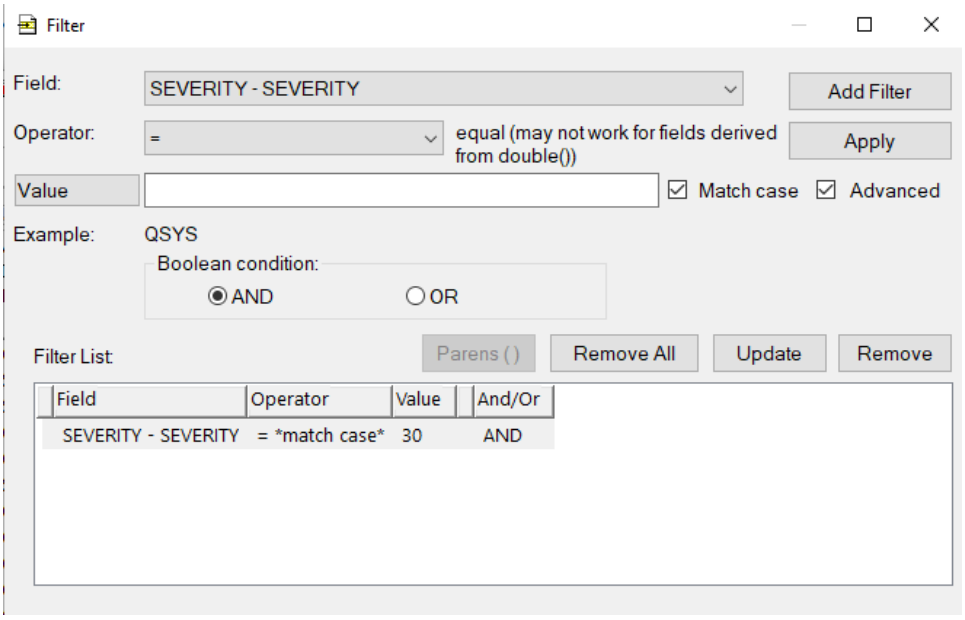

Filter Window with Advanced options displayed

8.5.1 Adding a Filter

The following shows an example of adding a Filter to a table:

1) Right-click desired column, then use the Add Filter… menu.

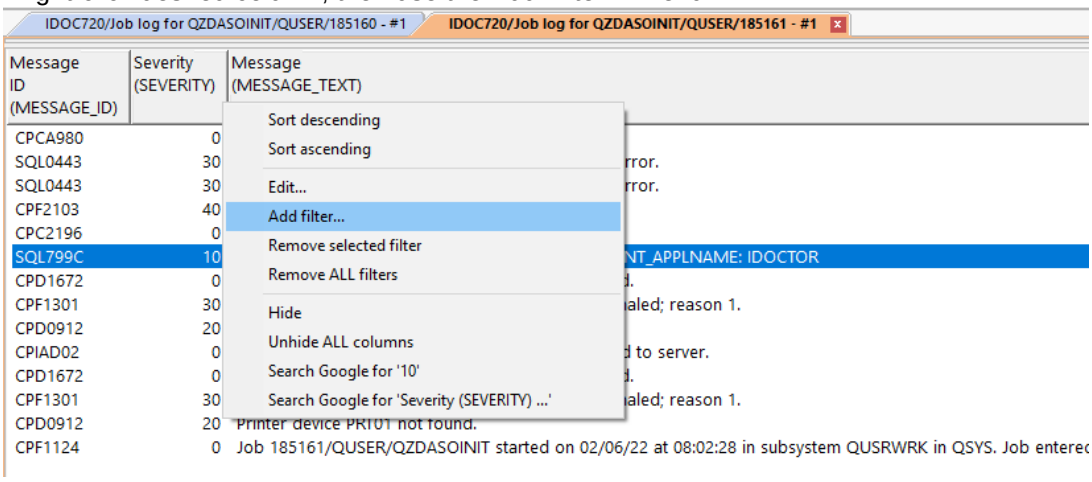

Column menu – Add Filter

2) Type desired value and select Operator value and click Add Filter.

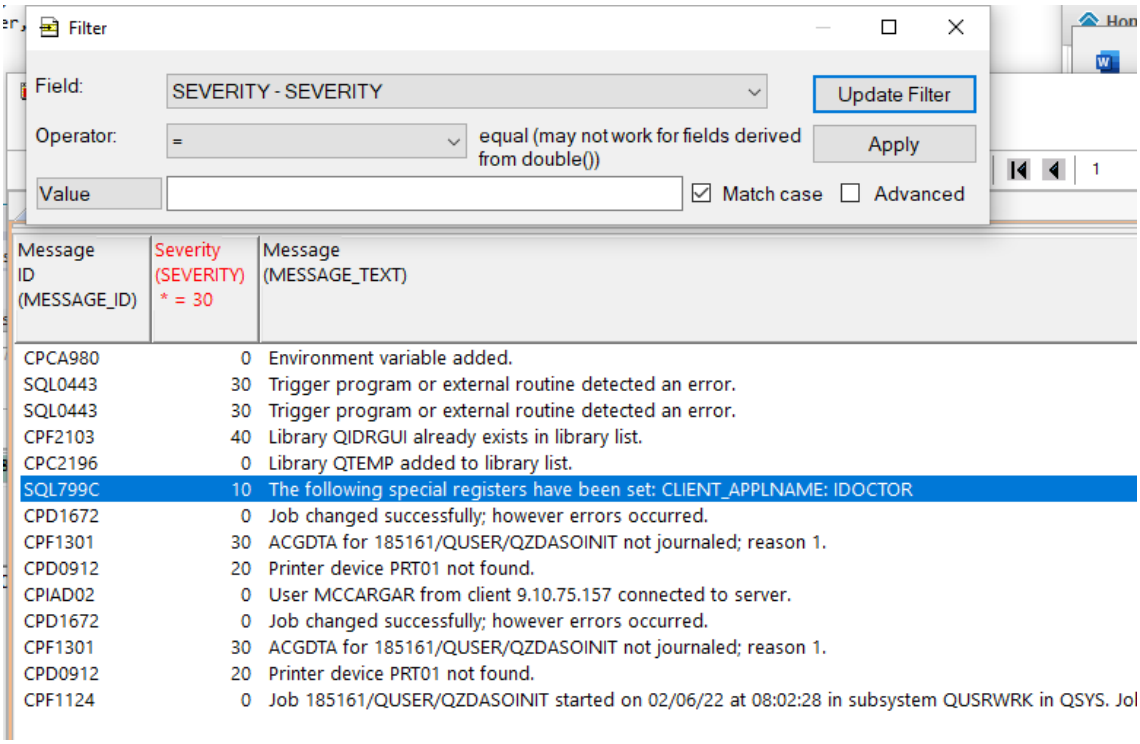

Table with Filter window, Apply button not yet pressed

3) Pressing the Apply button reruns the query using the desired filter and updates the results.

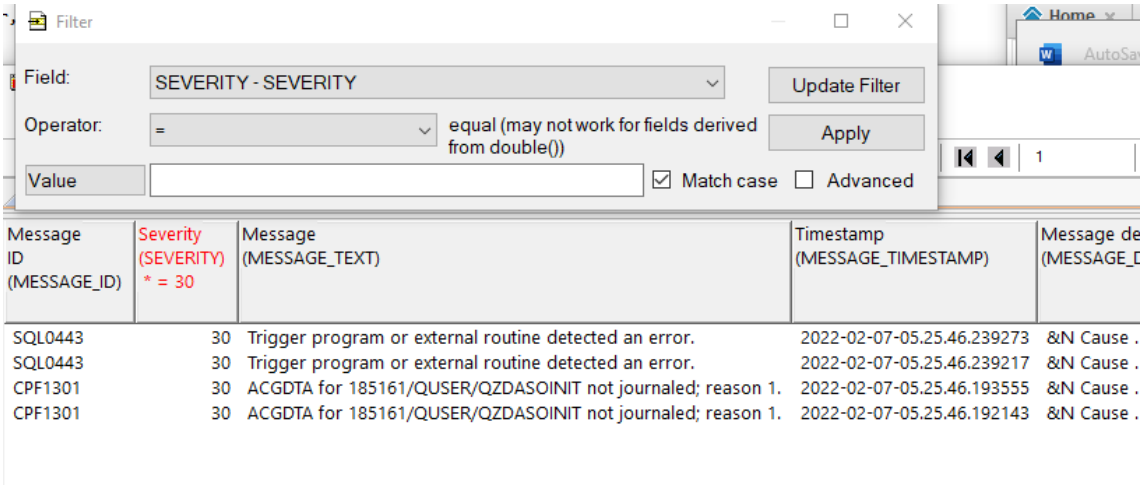

Table with Filter window, Apply button has been pressed

8.5.2 Removing a filter

The process to remove a filter from a table looks like this:

1) Right-click the table's column showing the filter.

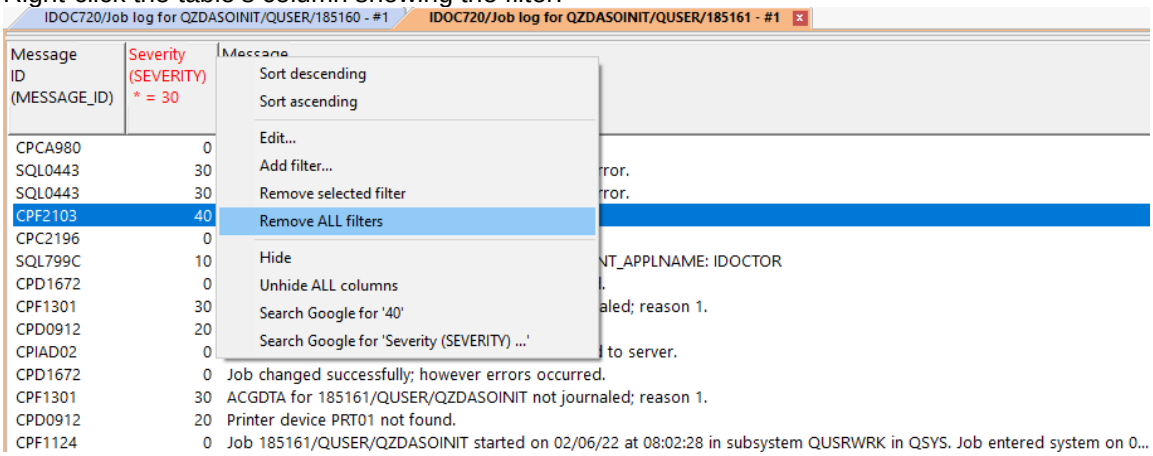

2) Select the Remove ALL Filters menu or Remove selected filter and the filter will be removed.

8.5.3 SQL Statement Changes

When filters or the Query Definition interface options are used, the SQL statement is modified so that the existing SQL statement becomes a subselect of a new statement. For example, the SQL behind the table shown in the previous section is:

```
SELECT * FROM (SELECT MESSAGE_ID, SEVERITY, MESSAGE_TEXT, MESSAGE_TIMESTAMP,
MESSAGE_SECOND_LEVEL_TEXT  AS MESSAGE_DETAILS, ORDINAL_POSITION AS POS
FROM TABLE(QSYS2/JOBLOG_INFO('185161/QUSER/QZDASOINIT')) A
) GUI_FILTER WHERE UPPER(SEVERITY) = 30 ORDER BY POS DESC
```

The identifier "GUI_FILTER" is used by the iDoctor GUI to indicate that this special type of filtering has been defined. Removing or changing this value in the SQL Editor will cause the filters to be no longer usable via the GUI.

8.6 Find Window

The Find Window allows a user to perform a search over a Table View. Use the Edit -> Find... menu (Ctrl+F) or right-click on a Table View and choose the Find... menu to use this option. Find allows the user to search for a text string within a specific column.

An example of the find window is shown below:

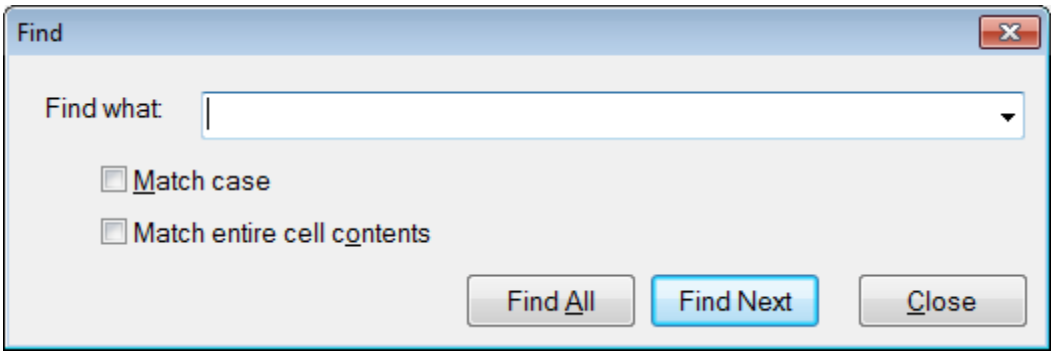

Find Window

After providing a search term to look for you can close the window and use the Edit -> Find Next (F3) or Edit -> Find Previous (Shift+F3) menu to look for the next/previous occurrences without needing to have this window visible.

8.7 Query Definitions

Tables and graphs are created via an underlying query definition or SQL Statement. The query definition defines exactly how data is to be retrieved and from what files(s), SQL tables, views or aliases. The Query Definition Interface is an interface over (the outermost portion of) an SQL statement. Most table and graph views in iDoctor provide a query definition menu that lets the user work with the SQL Statement behind the report.

The Query Definition Interface allows a user to customize the query for the active table or graph within the Data Viewer. Right-click on the view and use the Query Definition menu to open the Query Definition Interface for the desired tab.

The tabs within the interface are:

Field selection - indicates the order of the fields and the ones to display or hide
Filters - used to filter out or only include records that meet certain characteristics

8.7.1 Field Selection

The field selection panel allows you to hide or reorder the fields in the associated table view. You can also use the Reset button to discard all changes.

Instructions for performing each of these types of operations follows:

8.7.1.1 Working with field visibility

Visible fields are indicated by a checkmark in the Show? column within the Field List. If a field is not checked, then it will not be shown.

You may use the Toggle Selected button to check/uncheck the checkbox for the selected fields. This can be handy when you want to hide or show many fields at once.

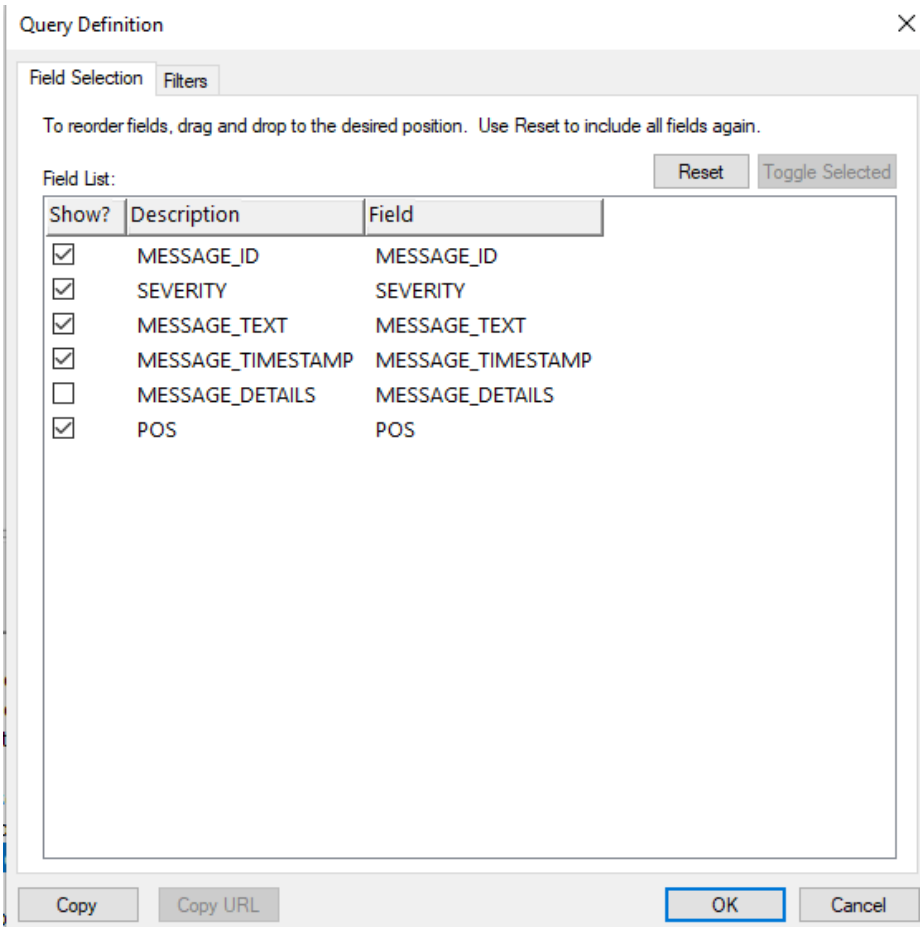

Query Definition -> Field Selection Window

Tip: Removing fields that are currently in the ORDER BY clause of the SQL statement will cause the field(s) to be removed from the order by.

8.7.1.2 Reordering fields

The order that the fields are displayed in the Field List is the same as they will be shown in the table view.

To reorder fields:

1. Select the fields you wish to reorder using the mouse and ctrl/shift keys.
2. Press the left mouse button over one of the selected fields and hold it down.
3. Drag the selection to your desired position in the list. You can scroll through to the bottom of the list if desired.
4. Release the left mouse button.

8.7.2 Filters

The Filters tab allows a user to limit the number of records returned in the active table or graph view.

An example of this interface:

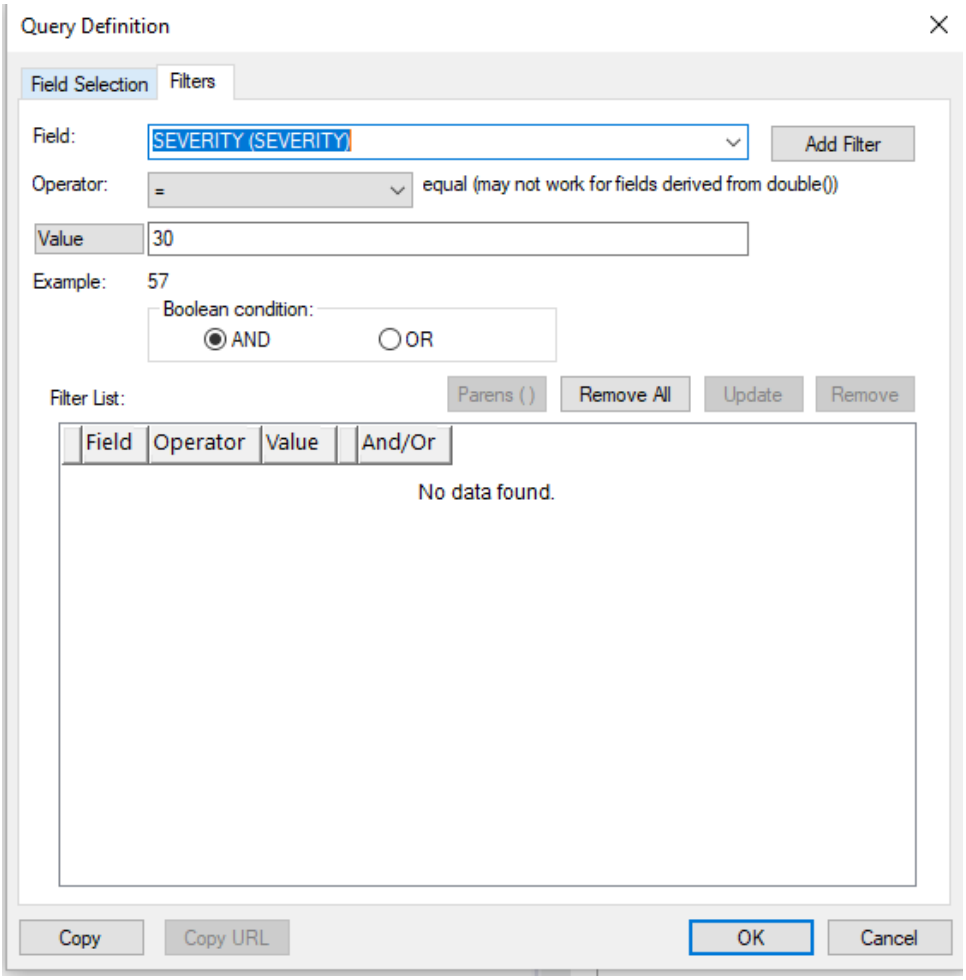

Query Definition -> Filters Tab

As the selection changes in the list, the interface objects above the list will change based on the current selection. This allows the user to quickly change values in the filter list by selecting any item in the list, changing any values from the fields above the list, and clicking the 'Update' button. The 'Update' button will update the selected row in the filter list.

A description of the interface follows:

Option	Description
Field drop-down list	This is a list of every field in the current report. Select a field to filter by before clicking the 'Add Filter' or 'Update' buttons. The short name of a field may also be entered.
Operator list	This is a list of every operator available for the currently selected field. A text field has a different set of available operators than does a numeric field. The set of operators is also different for a timestamp field. The operators 'Field contains', 'Field starts with', 'Field ends with', 'Field xxx', etc are not valid for numeric and timestamp fields. The following operators are supported on this page: Equal Less than Less than or equal to Greater than Greater than or equal to Not equal Is null Is Not null Range List Not List Field contains Field starts with Field ends with Field does not contain Field does not start with Field does not end with
Value text box	Use this textbox to enter the value to apply to the current field using the selected operator. The value should match the format presented by the 'Example' label directly beneath the text box. Text fields should have their values enclosed in 'single quotes' and if the operator is 'Range', 'List' or 'Not list' then more than one values each separated by a space is expected. Whenever entering a value, follow the example provided.
Add Filter button	This button creates a new filter and adds the filter to the Record Selection Filter List.
Value/Expression button	This button allows the user to enter a valid SQL expression instead of a single value. This provides greater flexibility but requires that you know SQL syntax. Any errors in the SQL statement will prevent the query from running and will cause an SQL error message.
AND/OR options	Use this to indicate whether two filters should be ANDed together or OR'd together.
Parens ( ) button	The 'Parens ( )' button allows grouping of multiple filters in the Record Selection Filter List into a single logical expression by placing parentheses around the set of filters. If parentheses already exist for the starting and ending record in the selected range, the parenthese will be removed by pressing this button.
Remove All	This button will clear the list of filters.
Update button	This provides the ability to change the selected filter from the Record Selection Filter List.
Remove button	This button allows the user to remove one or more records from the Record Selection Filter List.
Filter List	This is a list of all of the active filters to be applied to the report. Use the 'Add Filter' button to add a filter to the list. Press the OK button on the bottom of the Query Definition dialog to close the dialog and display the report using the filters from the list.

8.7.2.1 Adding a Filter

1) First select the field from the Field drop down list.
2) Select the desired operator from the operator list. **Tip:** Character fields have more operators available.
3) Type in the value that the operator should test for.
4) Press the Add Filter button to add the filter to the list.
5) Press the OK button to close this interface and run the query using the new filter.

8.7.2.2 Grouping Filters

By selecting more than one entries in the list and pressing the 'Parens ( )' button the user can add or remove a set of parentheses. To remove parentheses around multiple filters, select the range of records that contain the starting and ending parentheses and click the 'Parens ( )' button. Parentheses are necessary in order to make complex evaluations in the where clause of an SQL statement such as: CPUTIME >10 OR (IO > 1000 AND JOBNAME LIKE 'QZ%')

8.7.3 Save Query Definition (Save As...)

Query Definitions are saved using the Query Definition -> Save As... menu for an active table view. The query definition behind a graph view is saved using the Graph Definition -> Save As... menu for a graph view. All Query Definitions are saved into the user-defined reports database. You can specify the folder name to save the query definition into within this database. These are accessed later from the User-defined reports folder for the collection type you are working with. Typically, these queries can be reused on any collection data of the same type (i.e. Job Watcher.)

An example of the Save Query Definition interface is shown below:

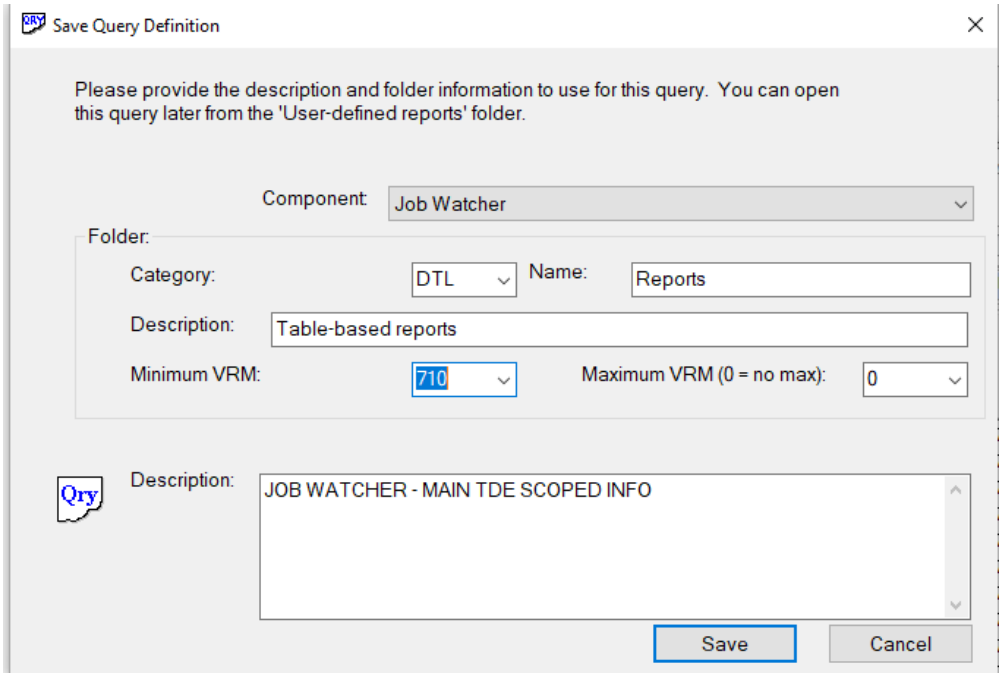

The interface elements within this window are described in more detail below:

Option	Description
Component	The name of the component this query should be visible in.
Category	3-character identifier for the folder the query should be saved into. If you wish to define new folder give it a name not already in the drop down list like 'DX1'
(Folder) Name	The name of the folder to store the query into
(Folder) Description	A long description to give the user-defined folder. (optional)
Minimum VRM	The minimum IBM i VRM that this query should be visible too in nnn format (i.e. 610, 710, 730, etc.) If the collection was created on a system older than this value, then this query will not appear.
Maximum VRM	The maximum IBM i VRM that this query should be visible too. Use the value of 0 for no maximum.
Description	The user-defined description for the query. This description can be up to 250 characters long.

8.8 Properties

This section describes the properties available for all table views in iDoctor. Additional tabs are available in some components for certain types of reports.

8.8.1 Record Quick View

This tab is part of the property pages for a table view. This interface is covered in the Main Window PDF under The Main Window -> Record Quick View.

8.8.2 SQL

This tab is part of the property pages for a table view. This interface is covered in the Main Window PDF under The Main Window -> Object Properties -> SQL.

8.8.3 Columns

This tab is part of the property pages for a table view. This interface is covered in the Main Window PDF under The Main Window -> Object Properties -> Columns.

9 Graph Views

The graph views in iDoctor display line, bar, area and pie charts built using SQL statements executed against data on the system. There are several different types of graphs supported: vertical stacked bar, vertical bar (side-by-side), horizontal stacked bar, horizontal bar (side-by-side), pie chart, area, vertical overlapping bar and horizontal overlapping bar.

In most cases, each color in the graph represents a different field from the query and each stacked bar represents a single row in the query results returned. In a few cases in iDoctor, a 'flattening' technique is used where a single stacked bar may be built from many rows and different values of a single field make up the various colors. (See 12x loop advanced graphs in CSI)

The graphs provide an optional attached legend identifying the fields in the SQL statement by color and where used within the graph. This button  on the toolbar may be used to hide or show the legend.

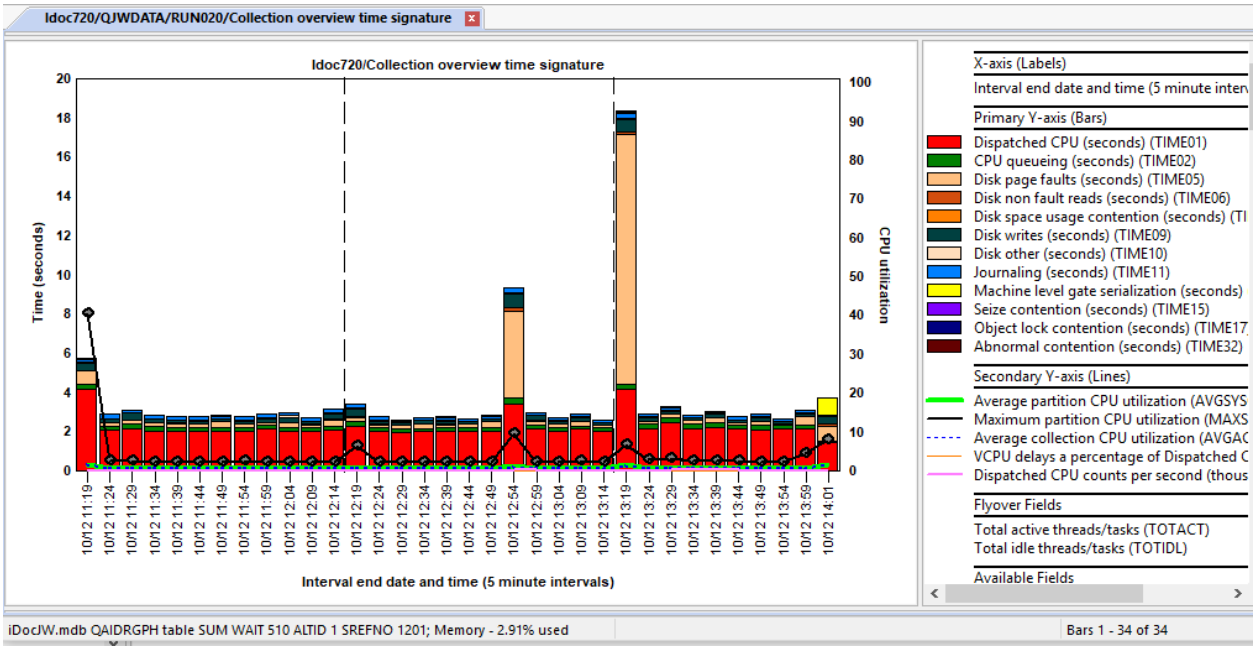

Job Watcher graph built from multiple collections

If necessary, use the scroll bars to navigate through the data shown in the graphs. Due to the potential to view vast amounts of data at one time, the graph data is shown a page at a time. The number of bars shown per page is configurable through the Preferences interface. When scrolling through the data the scale of the axes can be set to adjust automatically. This is another option on the Preferences interface.

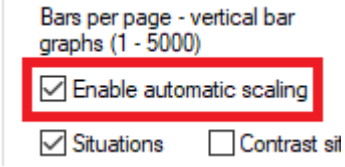

Preferences -> Display -> Enable automatic scaling

If automatic scaling is disabled, then the graph scale will be set to the maximum/minimum values of the first page shown in the graph.

Additional information about each piece of data in the graph is available by moving the mouse over the bar of interest. A flyover help window will appear in yellow providing this information. Some of this information is also displayed in the status bar as the mouse moves from bar to bar. The user can also

click on any bar to get a complete look at all the information for that particular piece of the graph and any other applicable data that goes with it (interval, job, etc)..

The Position Indicator Pane in the status bar indicates exactly which bars are being viewed out of the total possible.

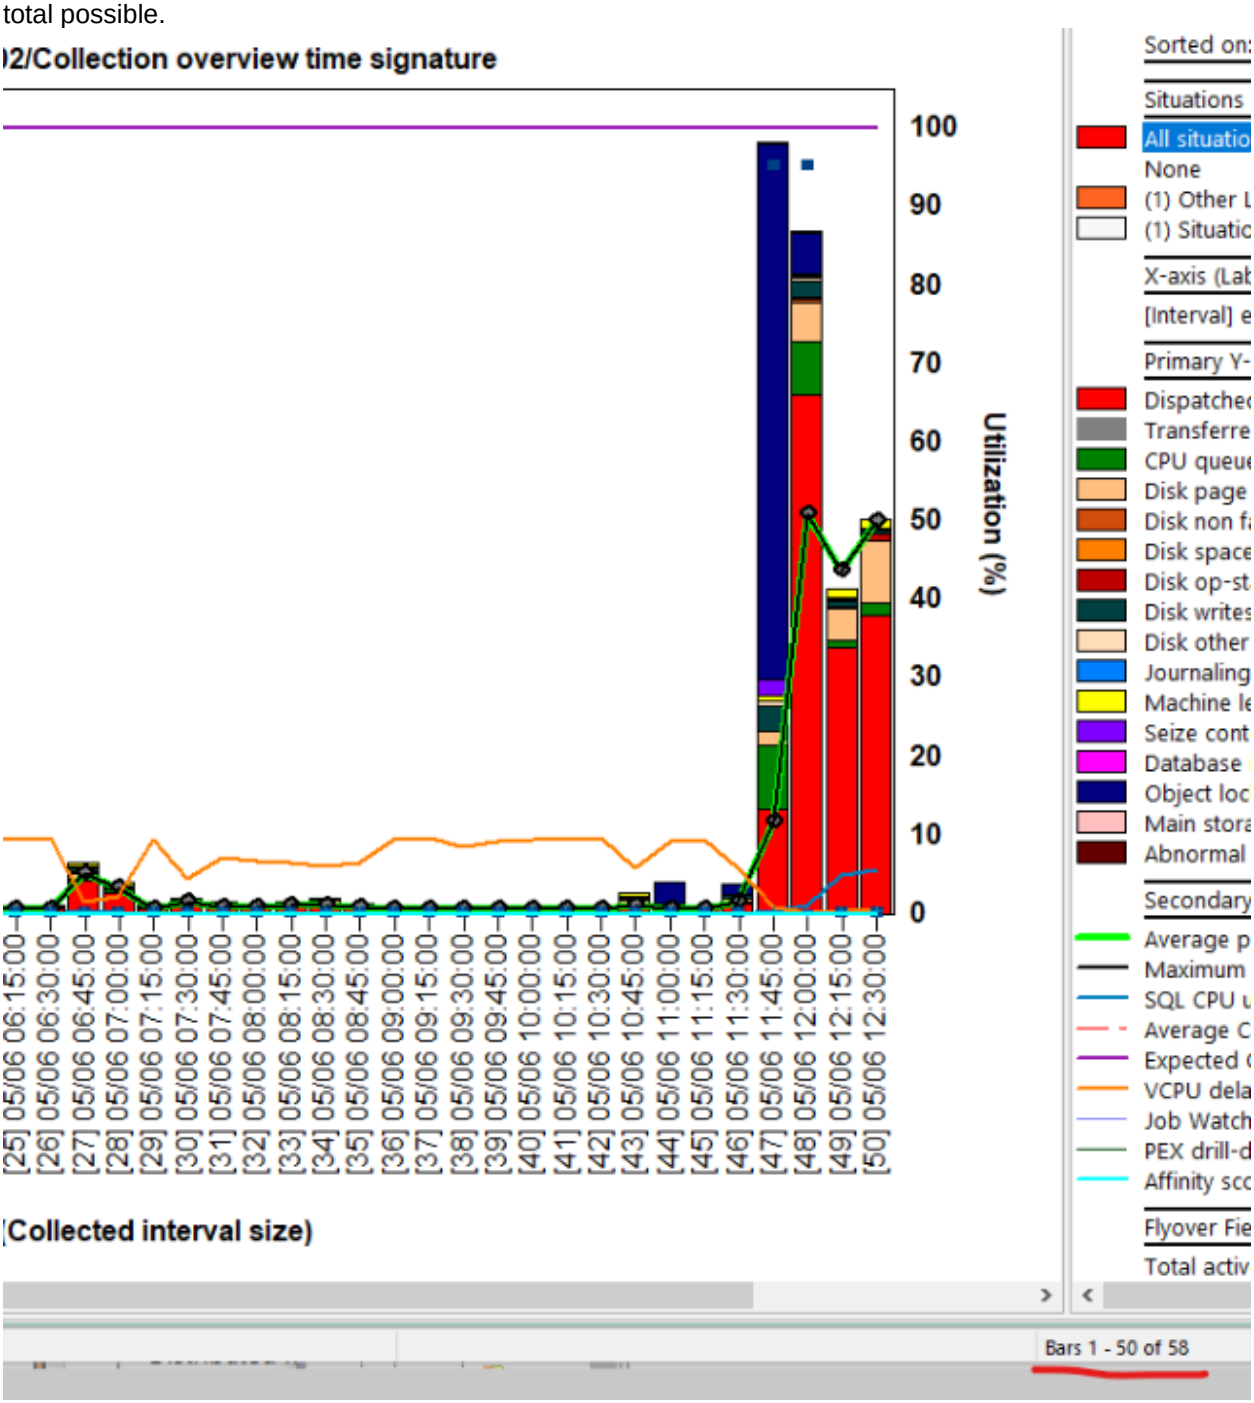

Graph position indicator pane in status bar (bottom right)

9.1 iDoctor-supplied graphs

iDoctor-supplied graphs are graphs shipped by IBM within iDoctor. Typically, iDoctor-supplied graphs will have additional drill-down options that are not accessible from the user-defined graphs.

9.2 User-defined graphs

User-defined graphs are created by the user and saved into a graph definition within the iDoctor User Defined Reports Database.

A user-defined graph can be initially created either from a table view or by modifying and saving an iDoctor-supplied graph. Creating a graph from a table view is done using the **Graph Definition -> Define New...** popup menu of a table view.

9.3 Vertical vs Horizontal

All iDoctor graphs fall into two categories: vertical and horizontal.

9.3.1 Vertical graphs

Vertical graphs have an X-axis (labels) at the bottom, the primary Y-axis is on the left and the optional secondary Y-axis is on the right (always lines.) The primary Y-axis is what defines the graph type (see the next section.)

The list of graph types that are vertical are:

1. Area
2. Lines
3. Pie
4. Vertical bar

9.3.2 Horizontal graphs

Horizontal bar graphs have an X-axis (labels) on the left, the primary Y-axis is on the bottom and no secondary Y-axis.

The list of graph types that are horizontal are:

1. Gantt
2. Horizontal bar

9.4 Graph Types

iDoctor supports many different types of graphs. Most iDoctor graphs are stacked bar graphs with either horizontal or vertical bars. The vertical bar graphs in iDoctor will often have lines on the Y2-axis.

Use the graph popup menu Graph Definition -> Set Graph type to change the graph type.

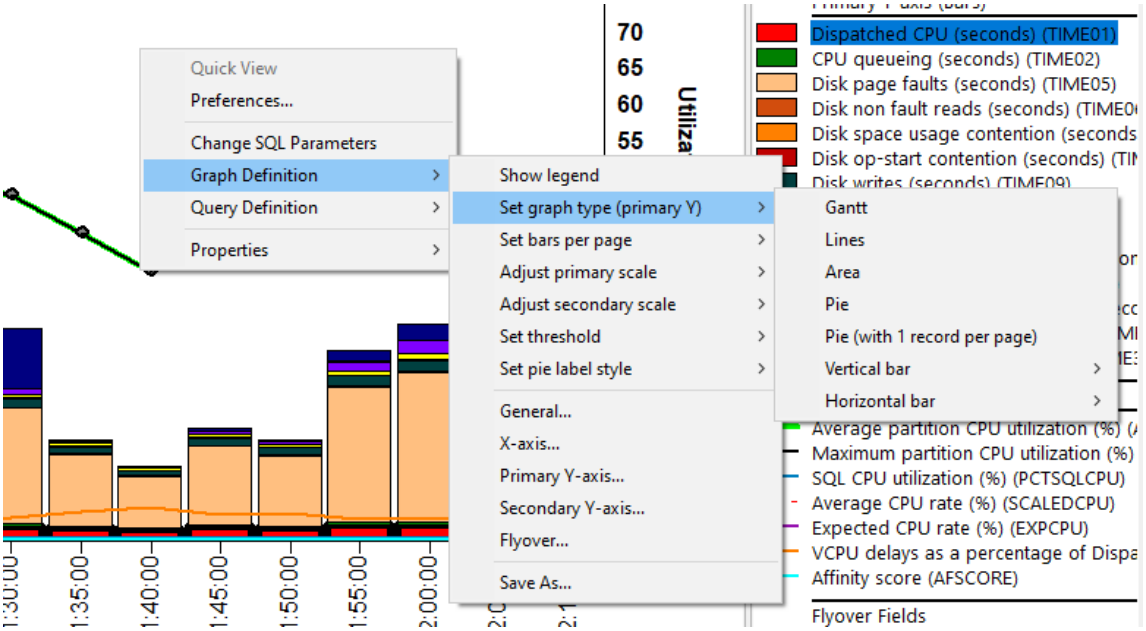

Graph Definition -> Set Graph type menu

Note: The graph type refers **only to the primary Y-axis**. If supported a secondary Y-axis showing lines may also exist.

Examples for each type of graph available are listed below:

9.4.1 Gantt

Gantt charts are used in iDoctor to display where a job executed over time in comparison to other jobs.

The time of day is listed on the bottom and each job starts and ends throughout that period as shown in the graph.

Note: Keep in mind the duration of the job could be longer than what is visible here depending on which wait buckets (or other metrics) are graphed. Only if ALL wait buckets are included would the duration of the jobs be accurate in this style of graph.

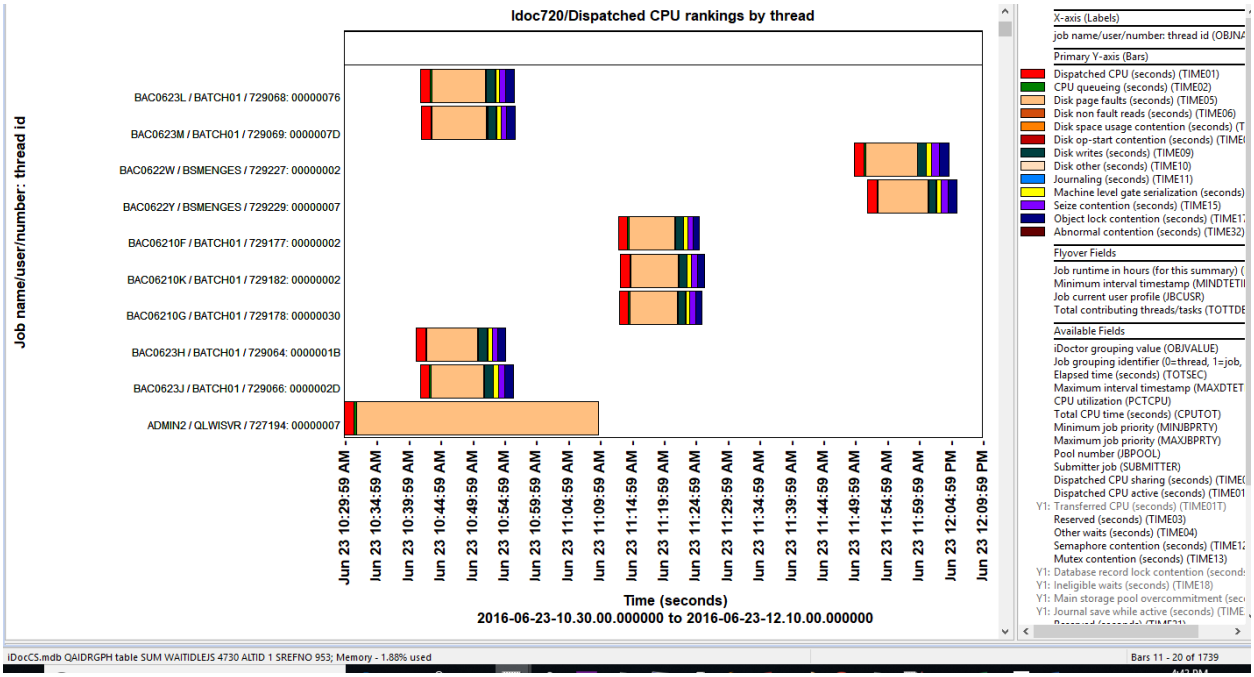

Gantt chart in Collection Services Investigator (excluding idle waits)

To include all wait buckets on this chart, press this button  on the Data Viewer toolbar. Pressing it again would toggle these idle wait off.

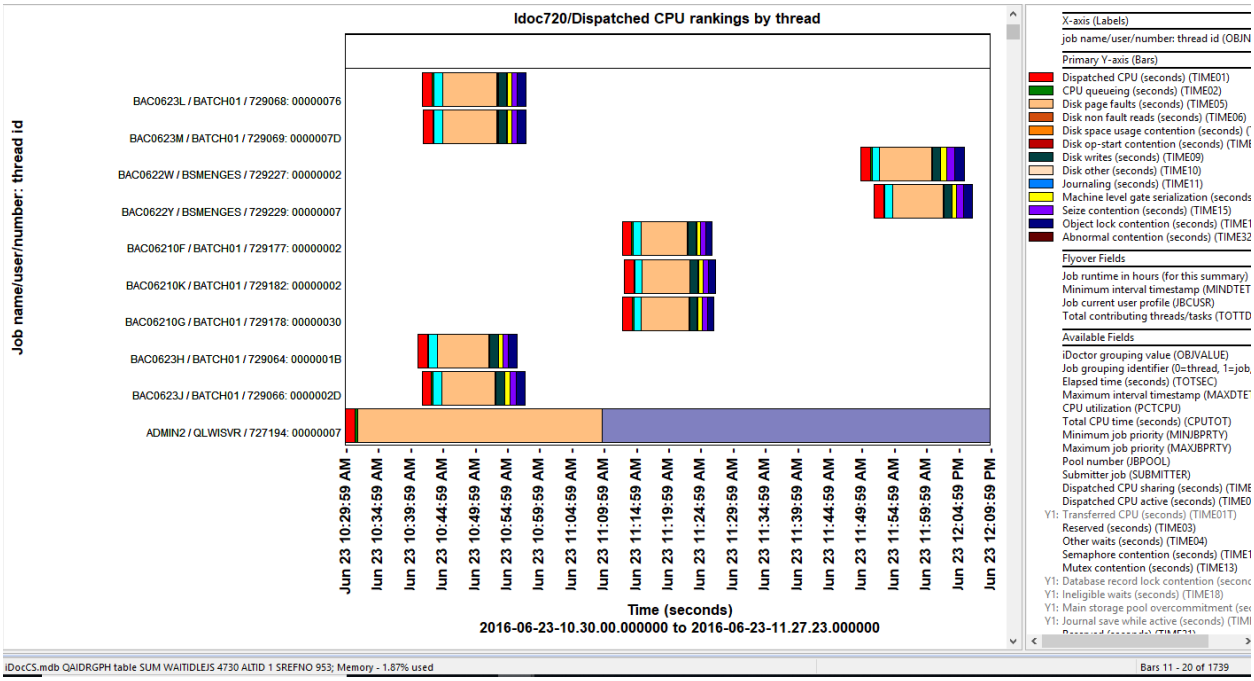

Gantt chart including idle waits

9.4.2 Lines

A lines graph in iDoctor can be used in iDoctor if desired by the user. It changes the metrics on the primary Y-axis to each be a line.

Note: The secondary Y-axis line options for widgets are not currently available on the primary Y-axis for this type of graph.

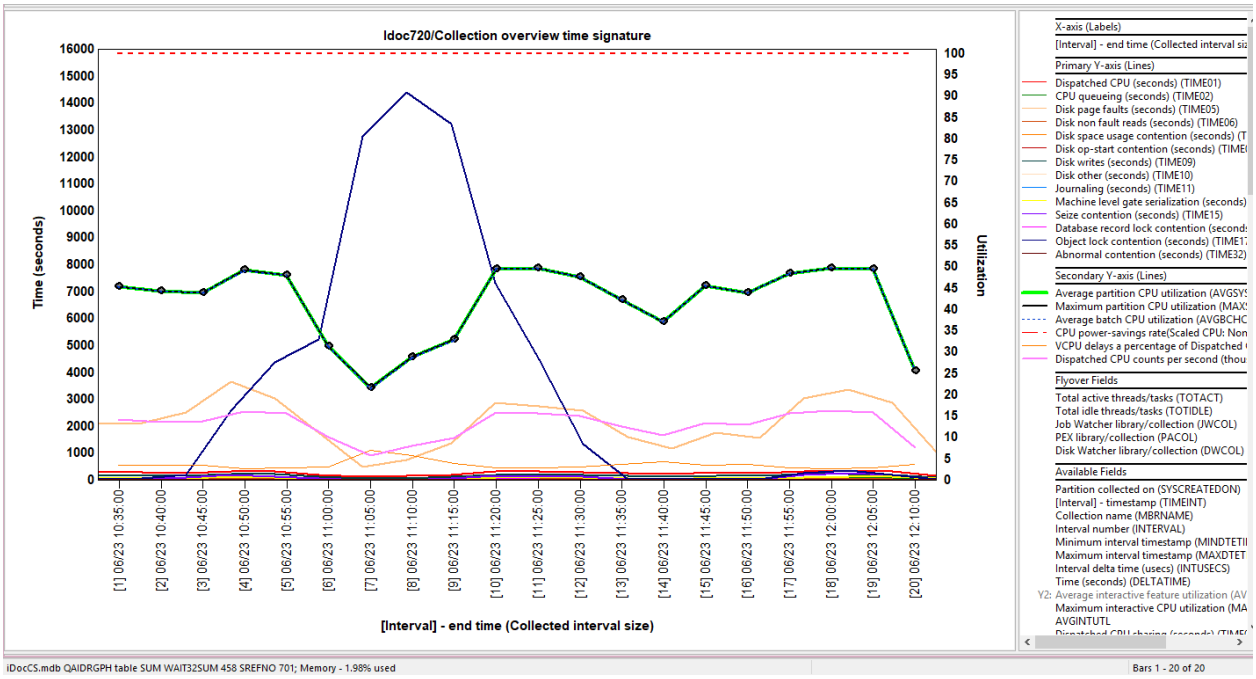

Line graph example

9.4.3 Pie

Pie charts display the metrics on the Y1-axis as percentages of the total of all values.

Tip: This type of chart works best if the bars shown per page is set to 1. You can do this using the Graph Definition -> Set bars per page -> 1 menu option or simply use the Graph Definition -> Set graph type -> Pie (with 1 record per page) option to do this in 1-step.

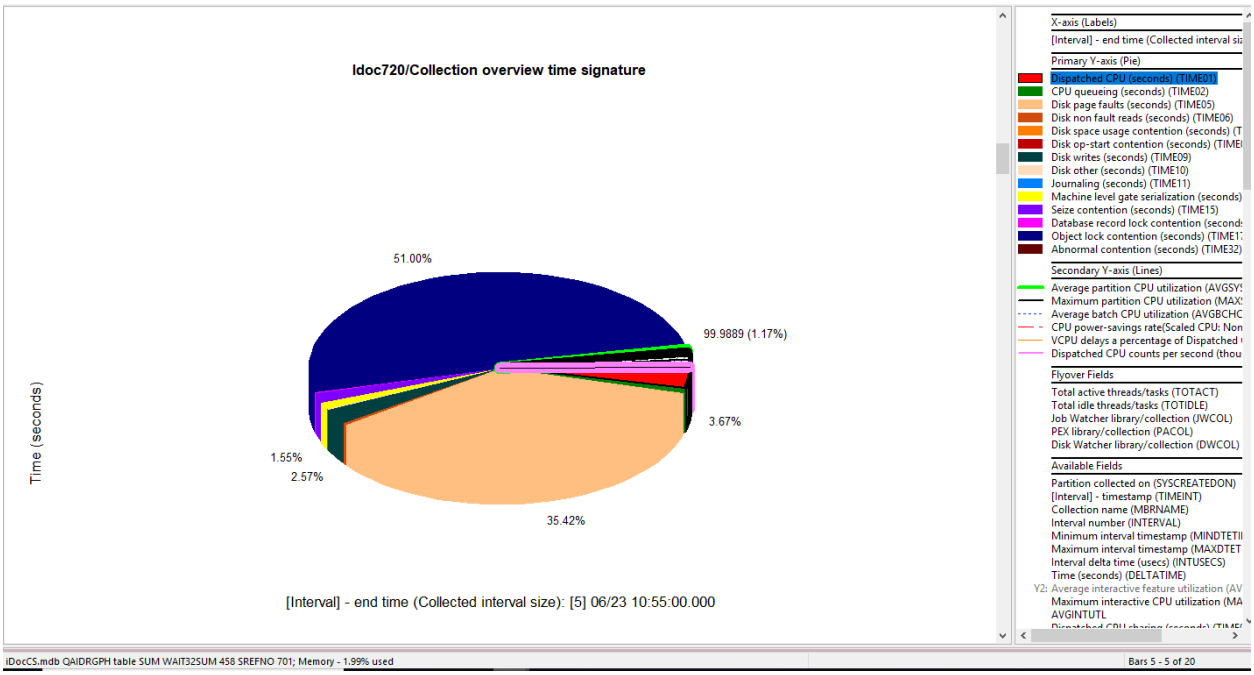

Pie chart example

9.4.4 Vertical bars stacked

This type of graph is the most common type used in iDoctor. It allows both a Y1 and Y2 axis, giving multiple types of data on the graph simultaneously. The Y1-axis contains stacked bar colors.

The Y2-axis in all iDoctor graphs are lines. These can be configured from the legend by right-clicking on them to add widgets (shapes) to the line points if desired. Various styles of lines can also be used (dashed, dots, etc) and the colors and widths may also be adjusted.

Tip: Graphs of this type allow the use of the Toggle Graph Format button on the toolbar.

1. Pressing this button for a time-based graph will convert the graph to Lines on the Y1 axis instead.
2. Pressing this button for a rankings graph will convert the graph to Horizontal stacked bars and the Y2 axis will not be available.

If you need to see the total value of all colors for a record or time interval, then simply click on it to select it and the Math Pane in the status bar will give you the total SUM of those values.

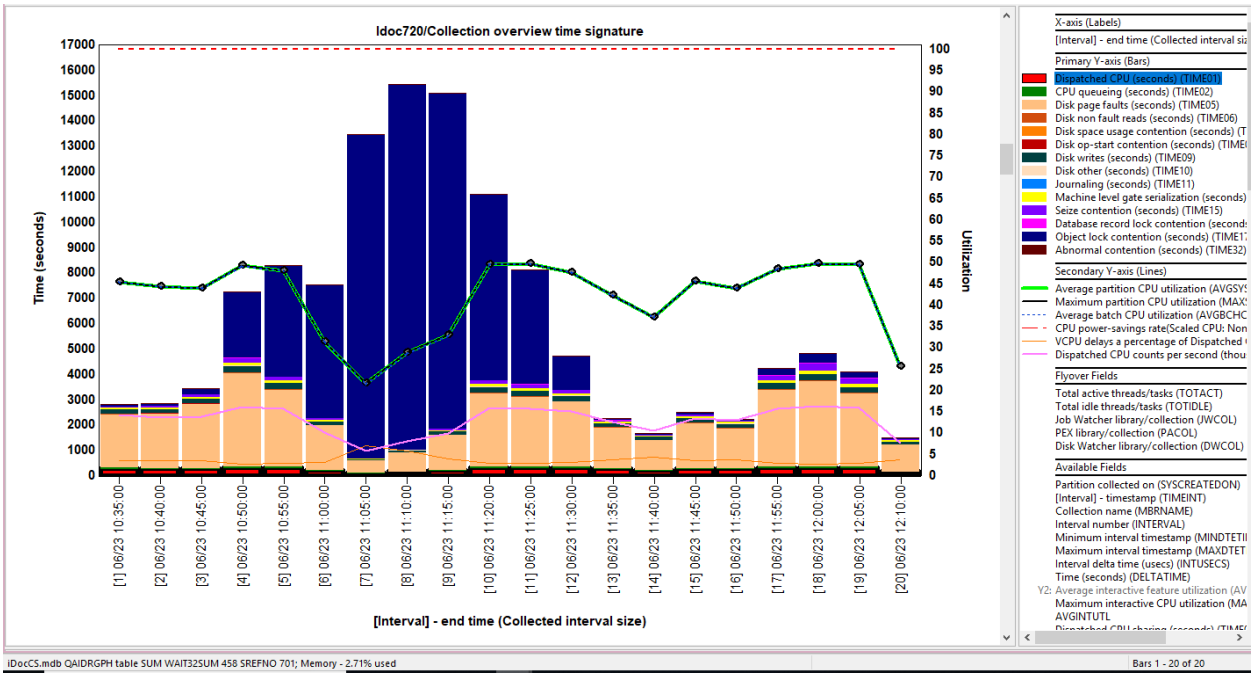

Vertical stacked bar graph

9.4.5 Vertical bars side-by-side

Vertical side-by-side graphs will show vertical bars, but each metric per X-axis point is shown side-by-side instead of stacked. This type of graph is best used with graphs only having 2 to 4 metrics on the Y1-axis. More than that is allowed but are difficult to read and seldom used for that reason.

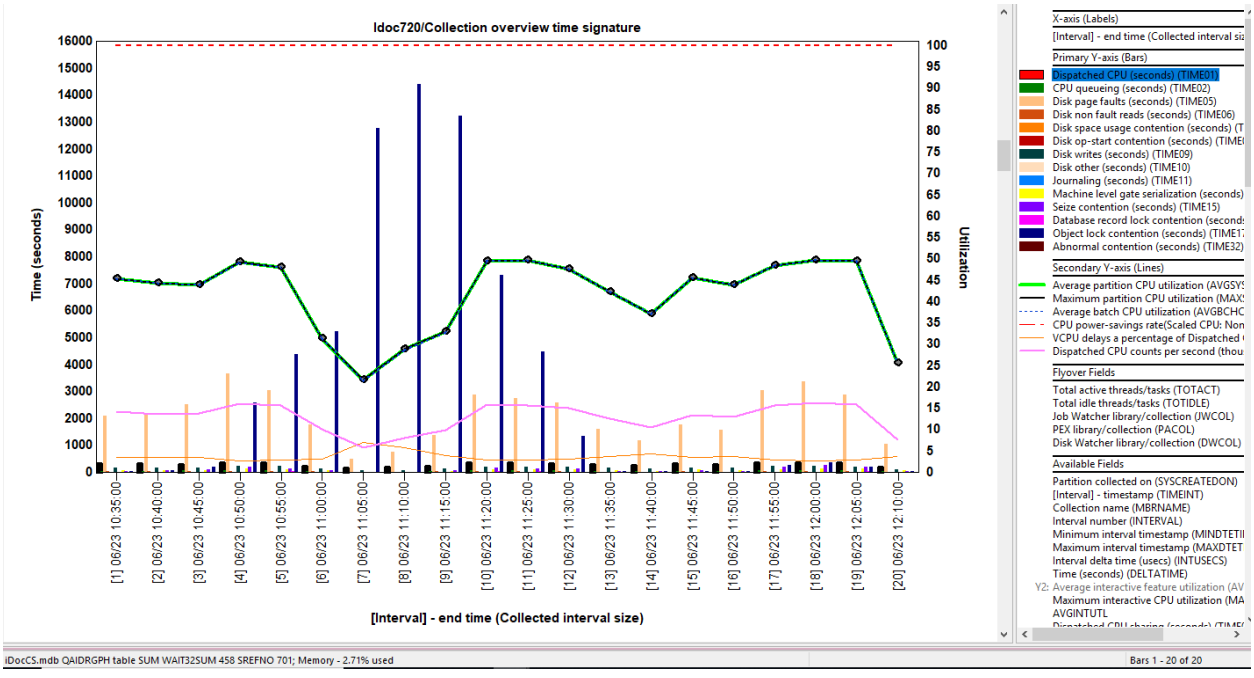

Vertical bar side-by-side graph

9.4.6 Vertical bars overlapping

Vertical overlapping bar graphs will show you the maximum metric per interval at the top of each bar. The metrics overlap each other and only the highest values will appear at the top of each bar/interval.

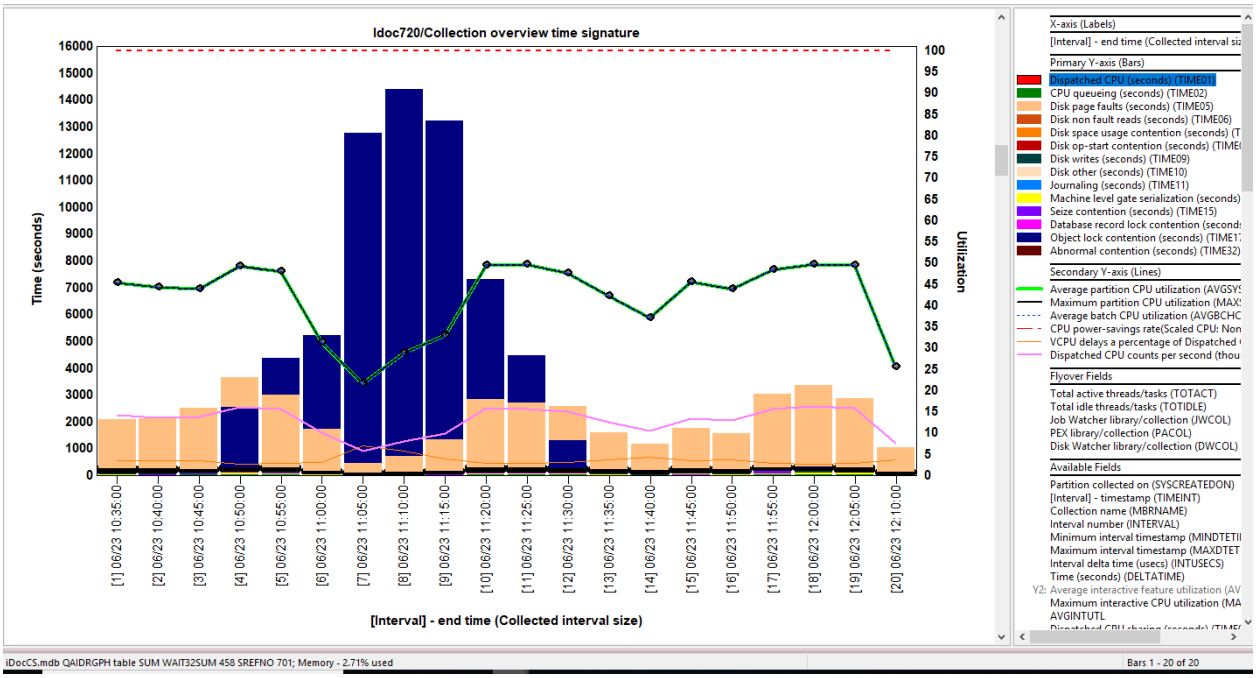

Vertical overlapping graph

9.4.7 Horizontal bars stacked

This type of graph is typically used in iDoctor for Ranking graphs. This graph type does not include a Y2-axis but since the labels are drawn horizontally are usually easier to read when listing job names or object names.

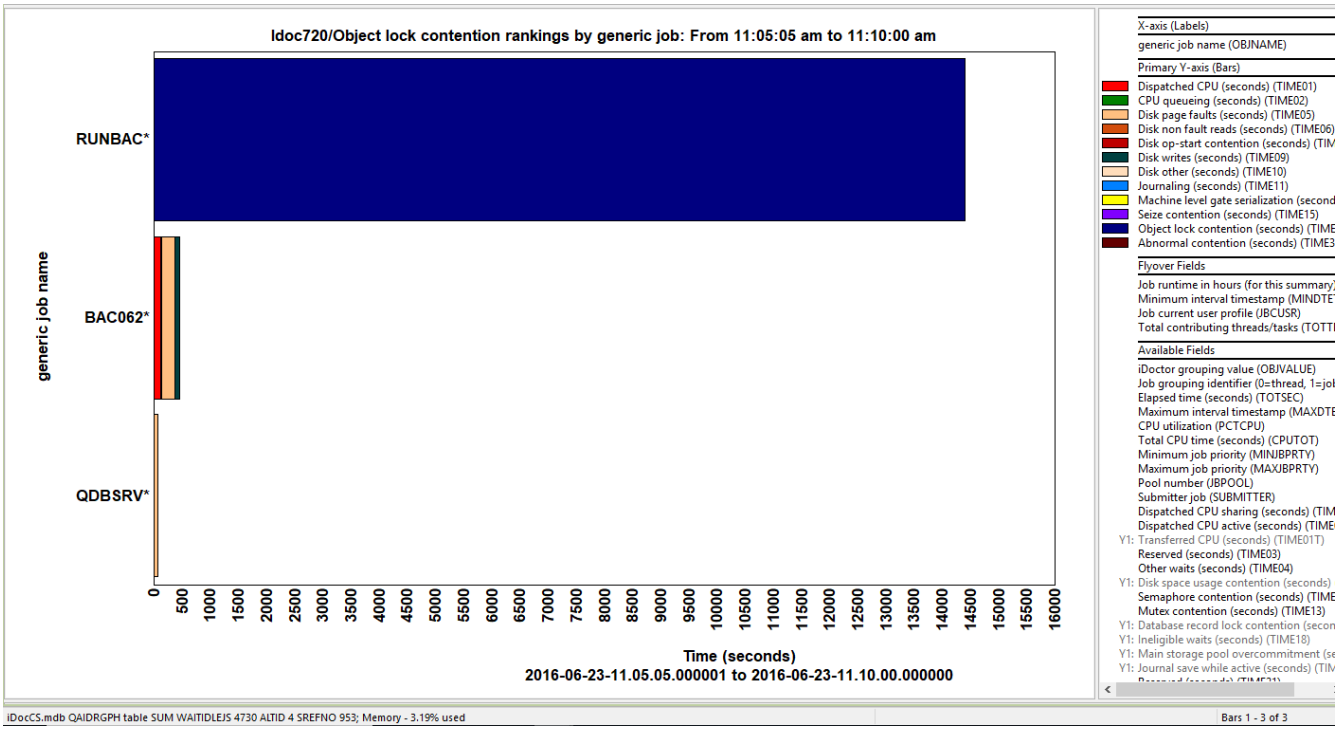

Horizontal bars stacked graphs

Tip: Graphs of this type allow the use of the Toggle Graph Format  button on the toolbar. Pressing this button will convert the graph to Vertical stacked bars instead.

9.4.8 Horizontal bars side-by-side

Horizontal side-by-side graphs will show horizontal bars, but each metric per point is shown side-by-side instead of stacked. This type of graph is best used with graphs only having 2 to 4 metrics on the Y1-axis. More than that is allowed but are difficult to read and seldom used for that reason.

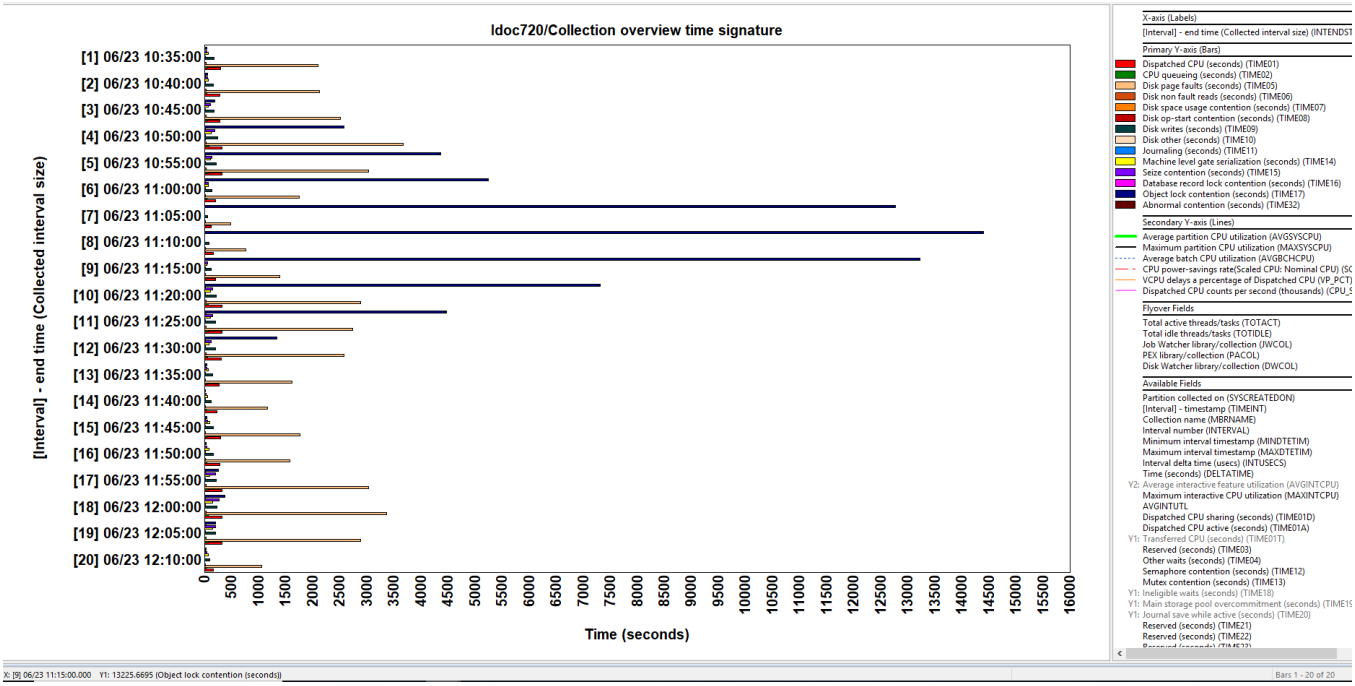

Horizontal bars side-by-side

9.4.9 Horizontal bars overlapping

Horizontal overlapping bar graphs will show you the maximum metric per label at the end of each bar. The metrics overlap each other and only the highest values will appear at the end of each bar on the right side.

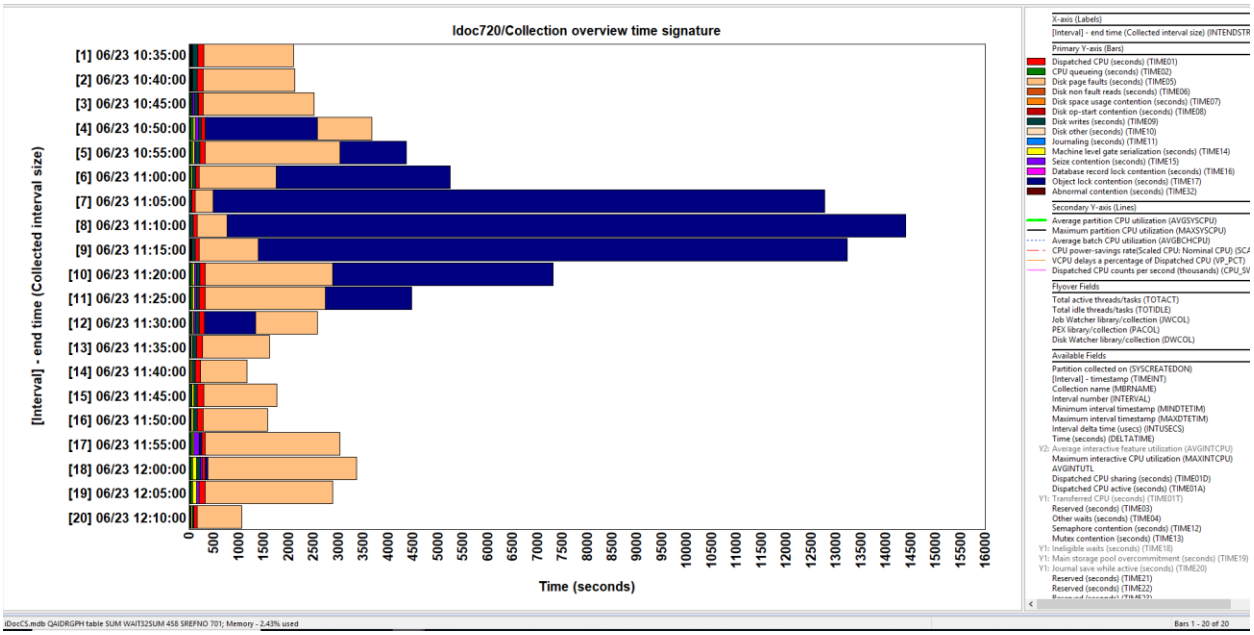

Horizontal bars overlapping

9.5 Graph Popup Menu

iDoctor graphs typically offer at the top the most likely drill downs you may wish to use next. Depending on the component and the graph being used, you will see options like Rankings, Detail reports and more.

From an Overview (time-based) chart typically you will select an area of interest and drill down into a rankings chart next. This will rank jobs, or disk units, or whatever type of data you are graphing into logical grouping and sorts them.

An example graph popup menu is:

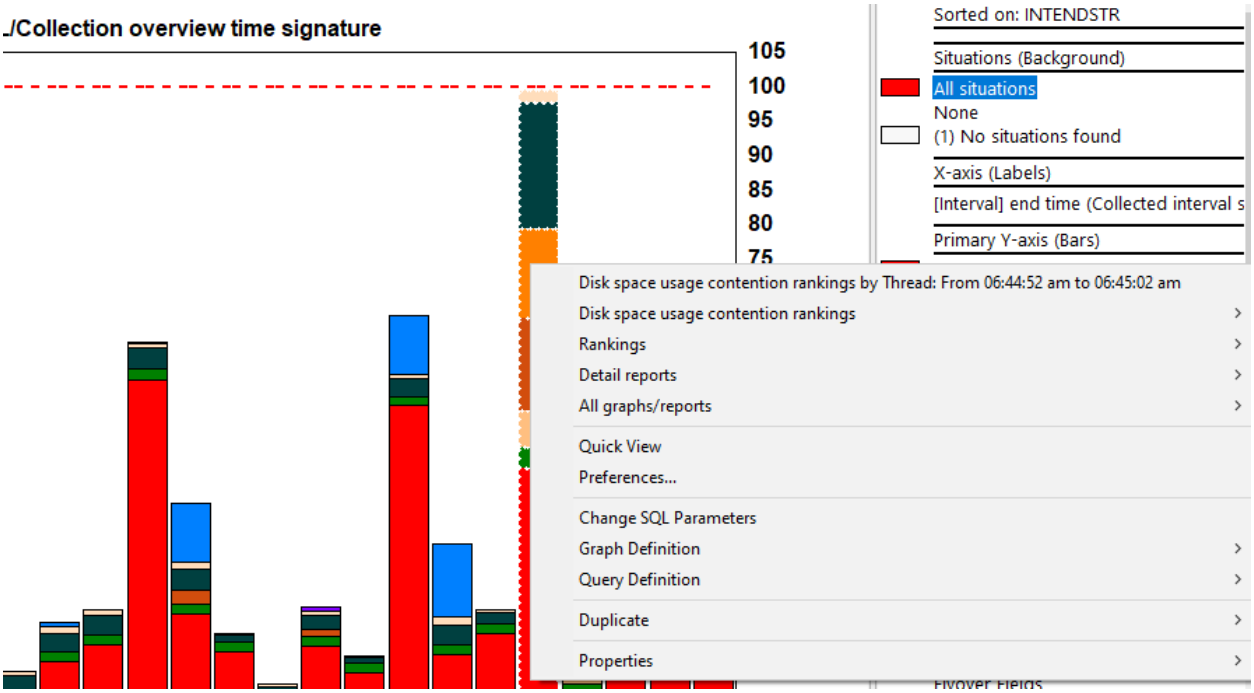

Drilling down from a Collection overview time signature graph

The graph and report menu options are discussed in more detail in the PDFs applicable for each component.

The following table describes the menu options available:

Menu	Description
Quick View	Quick View works similarly to Record Quick View for tables. It provides an interface to view the metrics behind the selected bar in a vertical list.
Preferences	Displays the Preferences window. With this interface the user can change the number of bars shown per page and customize font and other graph settings.
Change SQL Parameters	Displays the Change SQL Parameters interface allowing you to modify settings in the SQL statement behind the graph. In several graphs in iDoctor parameters are provided to allow users to filter the data more easily.
Graph Definition -> Show legend	Displays the legend if it is not already visible.
Graph Definition -> Set graph type	Allows you to modify the graph type of the primary Y-axis. See the previous section for more information.
Graph Definition -> Set bars per page	Allows the user to quickly set the bars shown per page, or to reset this value back to the default.
Graph Definition -> Adjust primary scale	Allows the user to quickly adjust the primary axis maximum value based on a percentage of the current maximum scale value shown. Use the Reset option to set the scale based on the largest value in the graph.
Graph Definition -> Adjust secondary scale	Allows the user to quickly adjust the secondary axis maximum value based on a percentage of the current maximum scale value shown. Use the Reset option to set the scale based on the largest value in the graph
Graph Definition -> Set threshold	This option will draw a threshold line straight across a vertical bar graph at the desired percentage of the Y2-axis. For example, this graph shows a threshold line added at 90%. 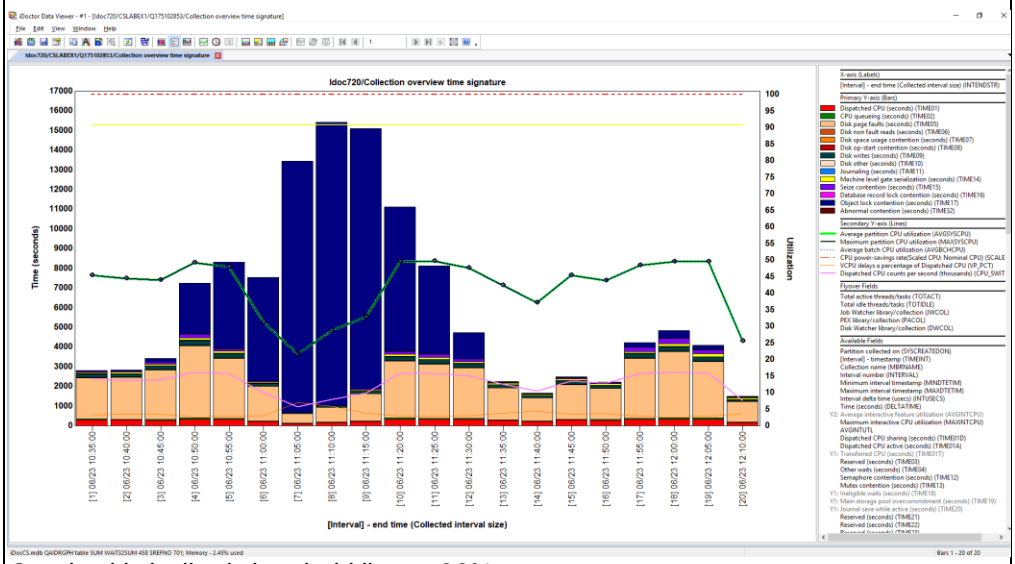 *Graph with (yellow) threshold line at 90%*
Graph Definition -> Set pie label style	This applies only to Pie charts and allows customization of how the labels appear on the graph. 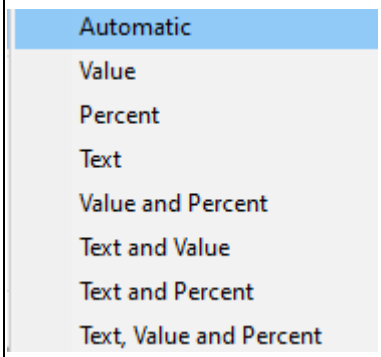 *Set pie label style options*

Graph Definition -> General	Displays the graph definition interface's General tab.
Graph Definition -> X-axis	Displays the graph definition interface's X-axis tab.
Graph Definition -> Primary Y-axis	Displays the graph definition interface's Primary Y-axis tab.
Graph Definition -> Secondary Y-axis	Displays the graph definition interface's Secondary Y-axis tab.
Graph Definition -> Flyover	Displays the graph definition interface's Flyover tab.
Graph Definition -> Save As	Saves the graph definition.
Duplicate as Table	Produces a new table view based on the SQL Statement used to produce the graph.
Duplicate as Graph	Produces a duplicate copy of the graph into the desired Data Viewer. This lets you move a graph from one Data Viewer to another
Properties	Displays the properties for the currently selected point on the graph as well as other information such as the SQL statement behind the graph view.

Other popup menu items are shown depending on the type of data/analysis being viewed. These additional options are covered under the documentation for the appropriate collection type.

9.6 Legend

All graphs provide an optional legend. The legend shows all the fields displayed on the graph and the fields defined in the SQL Statement that are available.

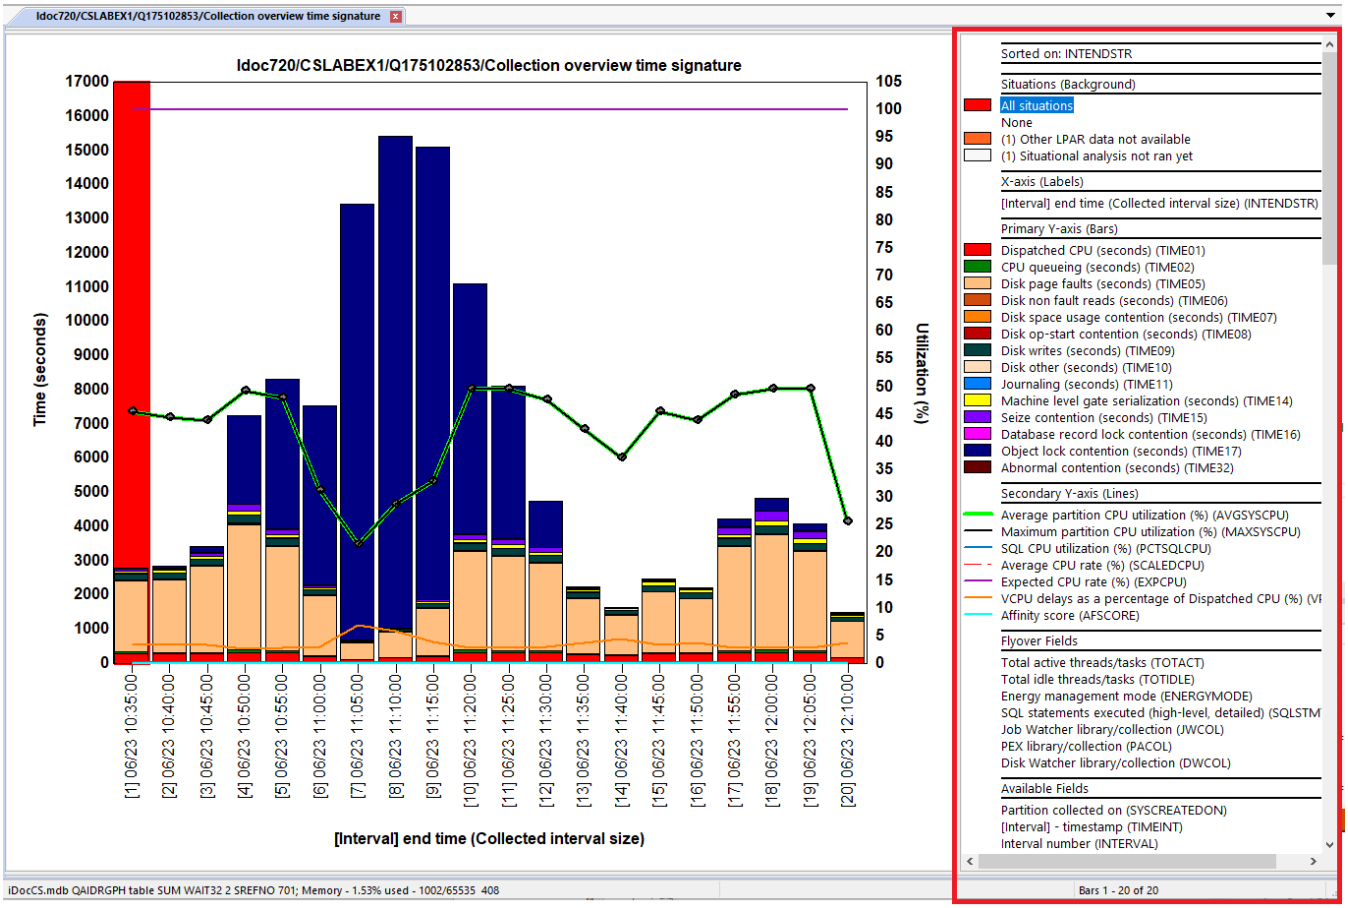

Overview graph with legend

The Legend contains the following sections, each representing an aspect of the graph:

Interface Element	Description
Sorted On	This section is just a label that identifies how the graph is sorted.
Situations (Background)	Contains a list of situations found in the collection and how many hits for each. Situations are shown in background colors in some of the vertical bar graphs in CSI, Disk Watcher and Job Watcher. **Tip:** Click on one of the options like All situations, None or a specific situation to change the graph background appropriately.
X-Axis (Labels)	Contains the fields shown that make up the labels along the X-axis. Up to 5 fields may be used to construct the X-axis labels.
Primary Y-Axis	Identifies the list of fields and colors (and patterns) that make up the colors on the graph's Primary Y-axis. You can drag and drop fields in this section to rearrange their order shown on the graph.
Secondary Y-Axis (Lines)	Displays the list of fields and colors used for the secondary Y axis lines shown on the graph. These fields are only displayable for vertical graphs.
Flyover Fields	Displays the list of fields to be displayed as optional flyovers when placing the mouse over an area on the graph.
Available Fields	This section lists all fields that are not defined in any of the previous sections in the legend.

Tip: Drag and drop is supported for all sections of the legend except the Sorted on and Situations sections.
This allows fields to be moved from one section of the graph to another.

You can also manipulate the legend and graph using the Legend's popup menu:

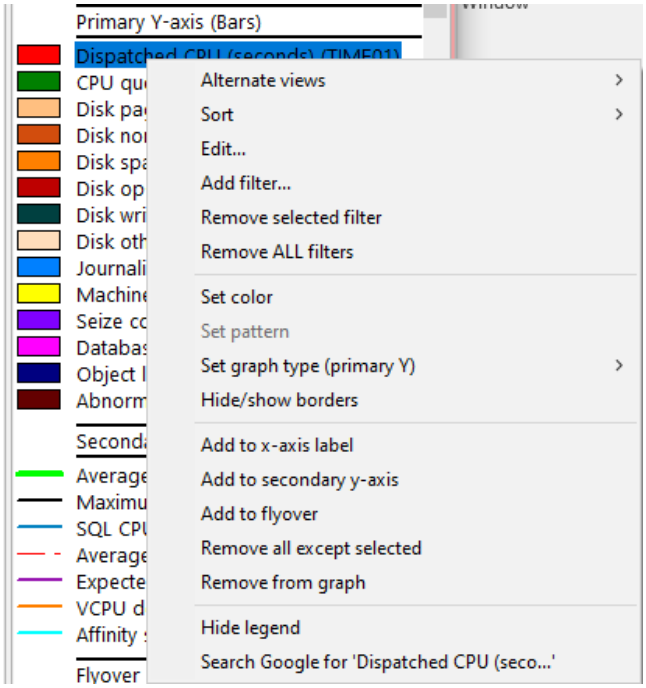

Legend Popup Menu

Popup Menu	Description
Alternate Views	If available, allows the user to quickly switch the graph to another IBM-supplied graph built from the current SQL statement.
Sort	Select 1 or more fields and use this option to rerun the SQL statement, changing the order of the data shown on the graph. **Tip:** If you select multiple fields then the values are added together before sorting. For example, this graph was sorted by ALL Primary Y-axis fields so the highest total for all buckets is shown first: 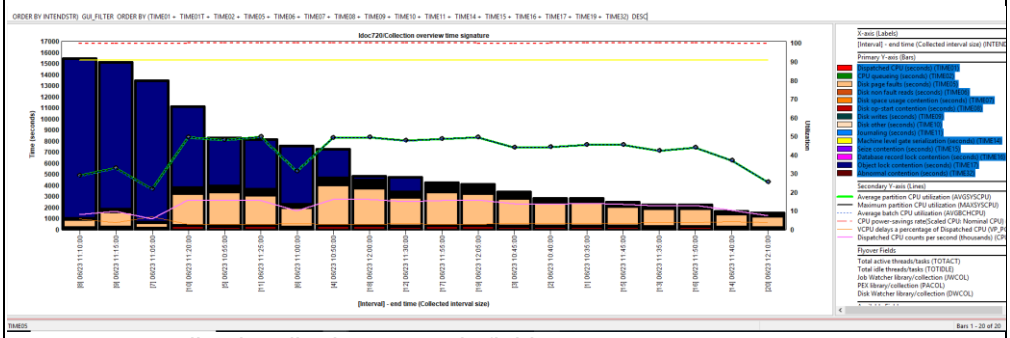 *Sort ->Descending by all primary Y-axis fields*
Edit…	This option shows the Edit Column interface which lets you modify a column's settings such as description and color. In iDoctor, a field's short name identifies it uniquely across all components.
Add Filter…	This lets you define a filter to add to the graph to reduce the data shown. **Tip:** If a filter is applied the graph will be updated and the filter value will be shown in the graph legend with a red font. 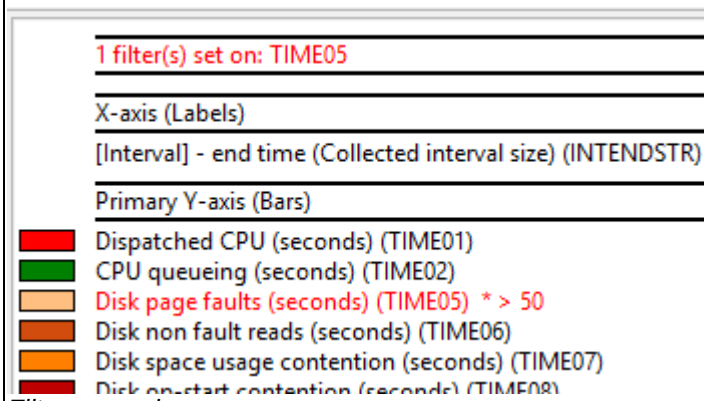 *Filter example*
Remove selected filter	If the currently selected field in the legend has a filter this option will remove it.
Remove ALL filters	This will remove all filters from the graph.
Set color	Allows the user to change the selected field's color using the Color Window.
Set pattern	Allows the user to change the pattern for the selected bar. These patterns are only visible if the Display Patterns option in the Preferences interface is enabled.
Set graph type	Allows the user to modify the graph type of the primary Y-axis.
Hide/show borders	This option allows you to remove or redisplay the thin border around every bar in the graph. Removing the border around a field in the bar graph can cause it to become more noticeable.
Add to X-axis label	Adds the selected field to the X-axis
Add to primary Y-axis	Adds the selected field to the Primary Y-axis

Add to secondary Y-axis	Adds the selected field to the Secondary Y-axis
Add to flyover	Adds the selected field to the Flyover.
Remove all except selected	Removes all fields from the current section of the graph you are working with except the selected one.
Remove from graph	Removes the selected field from the graph and adds it to the list of Available fields.
Hide legend	Hides the legend.
Search google for XYZ	This will open the default web browser and perform a Google search for the selected field description/name.

Set color	Allows the user to change the selected field's color using the Color Window.
Set pattern	Allows the user to change the pattern for the selected bar. These patterns are only visible if the Display Patterns option in the Preferences interface is enabled.
Set graph type	Allows the user to modify the graph type of the primary Y-axis.
Hide/show borders	This option allows you to remove or redisplay the thin border around every bar in the graph. Removing the border around a field in the bar graph can cause it to become more noticeable.
Add to X-axis label	Adds the selected field to the X-axis
Add to primary Y-axis	Adds the selected field to the Primary Y-axis
Add to secondary Y-axis	Adds the selected field to the Secondary Y-axis
Add to flyover	Adds the selected field to the Flyover.
Remove all except selected	Removes all fields from the current section of the graph you are working with except the selected one.
Remove from graph	Removes the selected field from the graph and adds it to the list of Available fields.
Hide legend	Hides the legend.
Search google for XYZ	This will open the default web browser and perform a Google search for the selected field description/name.

9.7 Filter

The Filter interface is accessed by right-clicking a field in the legend and using the Add Filter menu. Filters can be defined on one or more columns and each filter will modify the SQL statement (within the where clause) to perform the desired filtering. The text for fields that have filters applied are drawn with a red color.

The following shows an example of a graph with the Filter window open (but the filter is not yet applied):

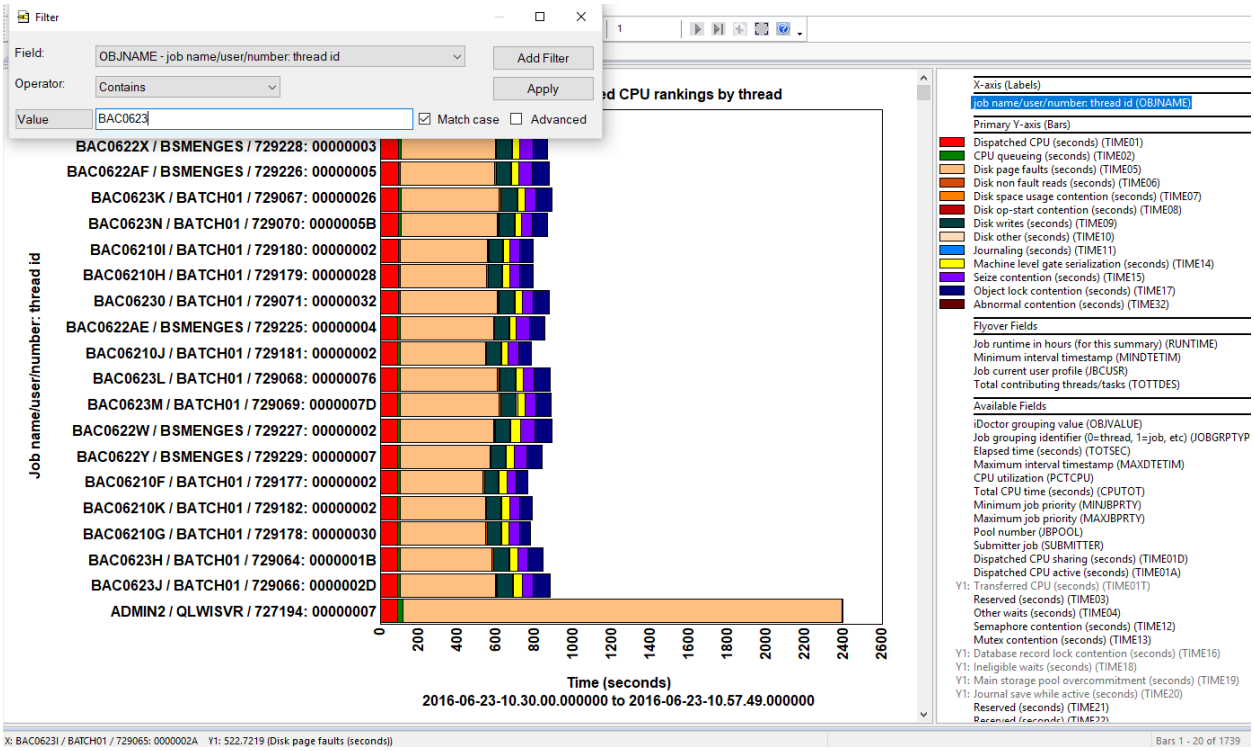

Graph with Filter window, Apply button not yet pressed

Pressing the Apply button reruns the query using the desired filter and updates the results.

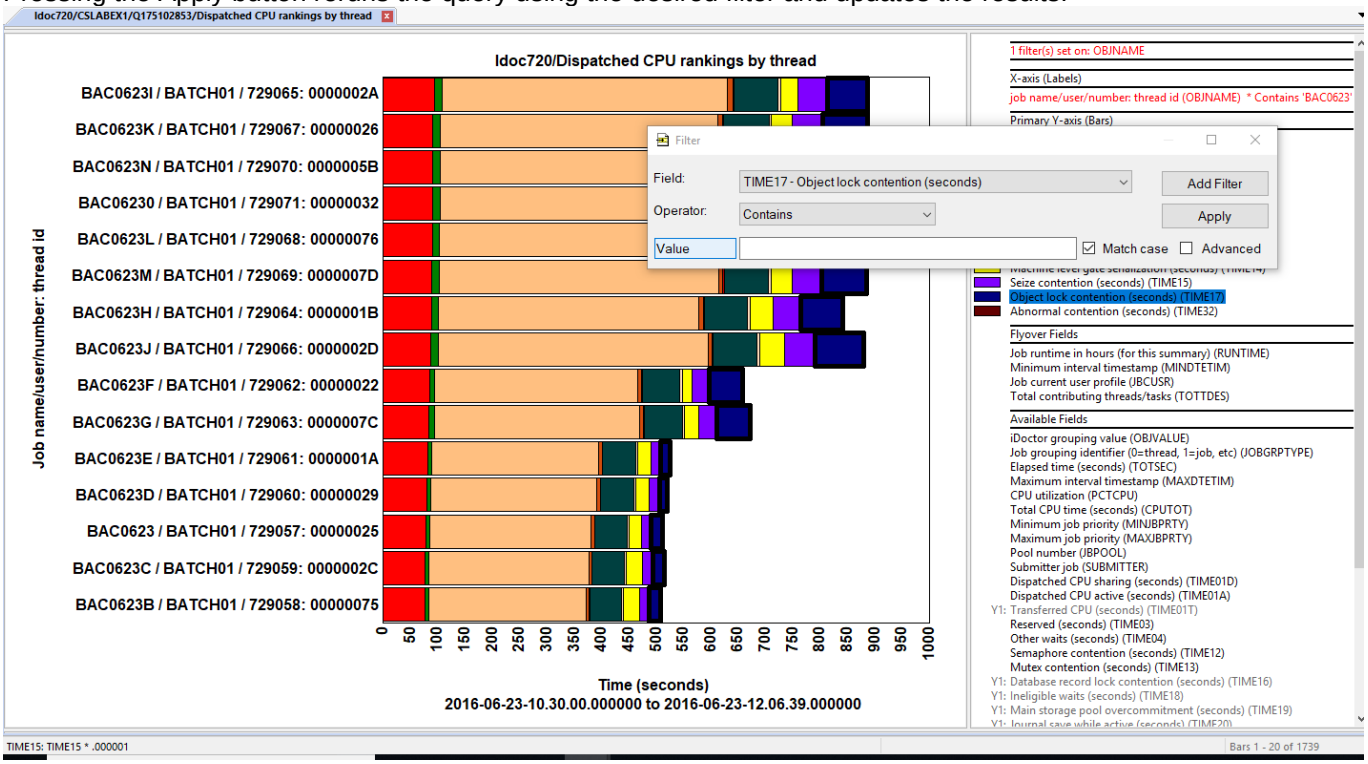

Graph with Filter window, Apply button has been pressed

The following options are available:

Option	Description
Field	The field to filter the graph data on. By default, this is the same field that was right-clicked (or last clicked on from the legend.)
Operator	The type of operation to use for this filter. The possible operators available depends on the type of the field selected. **Numeric field operators:** 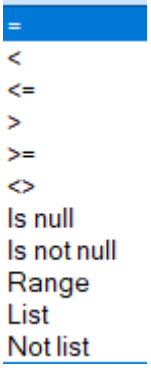 **Text field operators:** 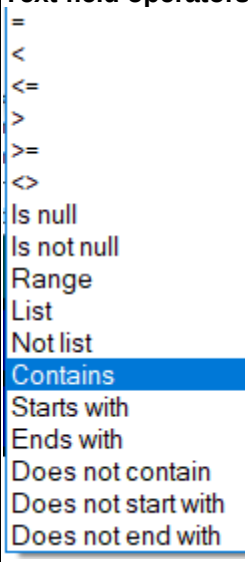 Note: the = (equal) operator may not work for fields derived from the double function.
Value	This is the value to apply to the filter. Generally, the filter is something like FIELD OP VALUE where OP is the operator, FIELD is the field to filter on and VALUE is a constant numeric or text string.
Match case	This checkbox indicates whether a case-sensitive search will be performed.
Add/Update Filter	This button will add or update the desired filter in the graph. The change does not take effect until the Apply button is pressed or the graph is refreshed.
Apply	This button will rerun the query behind the graph and apply any changes made to filters defined.
Advanced checkbox	This checkbox will hide/display the additional options on this window. Generally, they are not needed unless you wish to define more advanced options.

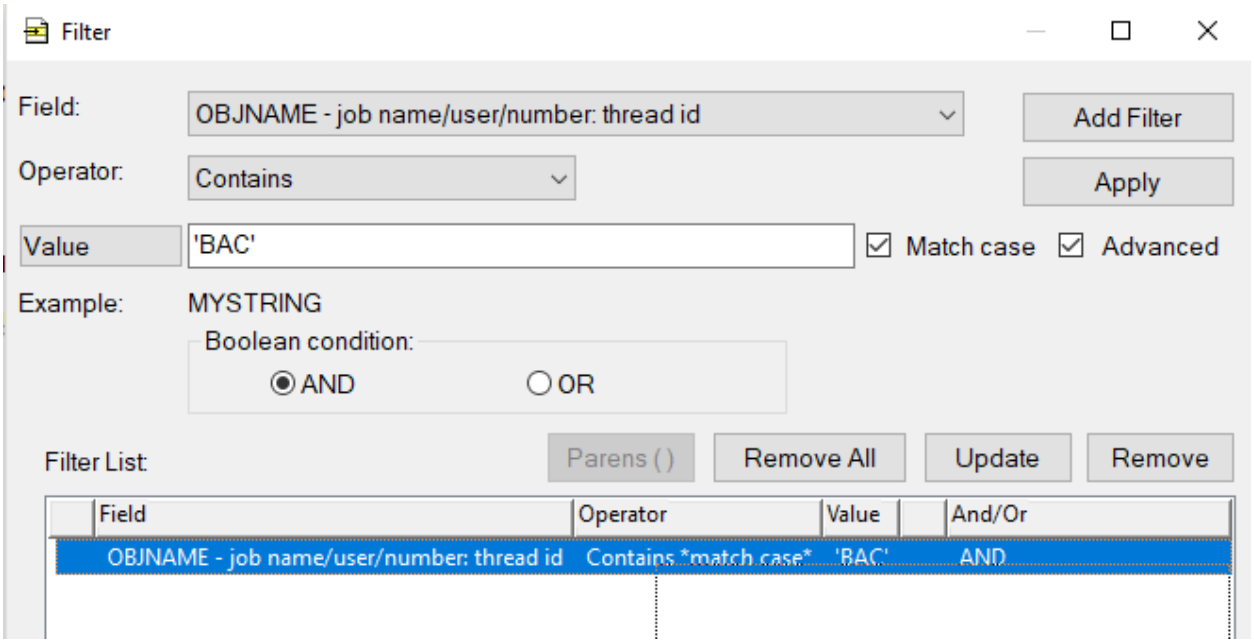

Filter Window with Advanced options displayed

Option	Description
Boolean condition	When using multiple filters this allows you to define if you want ALL conditions true (use AND) or EITHER one (use OR). By default, all filters must be true and use the AND Boolean condition.
Parens ()	This let you place parentheses around the multiple filters selected. This allows you to control how the logic is performed and the order of operations to have special filtering like: (X = 1 AND Y = 2) OR Z = 3 This button will only be enabled if 2 or more fields in the list are selected.
Remove All	Removes all filters from the list. Hit the Apply button to rerun the SQL statement and update the graph
Update	This will update the selected filter in the list based on any changes in the fields above.
Remove	Removes the selected filter(s) from the list.
Filter list	The list of filters to apply to the current SQL statement.

9.7.1 SQL Statement Changes

When filters or the Query Definition interface options are used, the SQL statement is modified so that the existing SQL statement becomes a subselect of a new statement. For example:

```
SELECT * FROM (SELECT MESSAGE_ID, SEVERITY, MESSAGE_TEXT, MESSAGE_TIMESTAMP,
MESSAGE_SECOND_LEVEL_TEXT  AS MESSAGE_DETAILS, ORDINAL_POSITION AS POS
FROM TABLE(QSYS2/JOBLOG_INFO('185161/QUSER/QZDASOINIT')) A
) GUI_FILTER WHERE UPPER(SEVERITY) = 30 ORDER BY POS DESC
```

The identifier "GUI_FILTER" is used by the iDoctor GUI to indicate that this special type of filtering has been defined. Removing or changing this value in the SQL Editor will cause the filters to be no longer usable via the GUI.

9.8 Properties

This section describes the Properties interface for iDoctor graphs. Additional tabs are available in some components and are covered in the documentation for those components.

9.8.1 Quick View

The Quick View page contains all the information about a bar in the graph from the data retrieved by the SQL Statement. Access this screen by double-clicking on any bar in a graph view or use the Quick View menu in the graph's popup-menu.

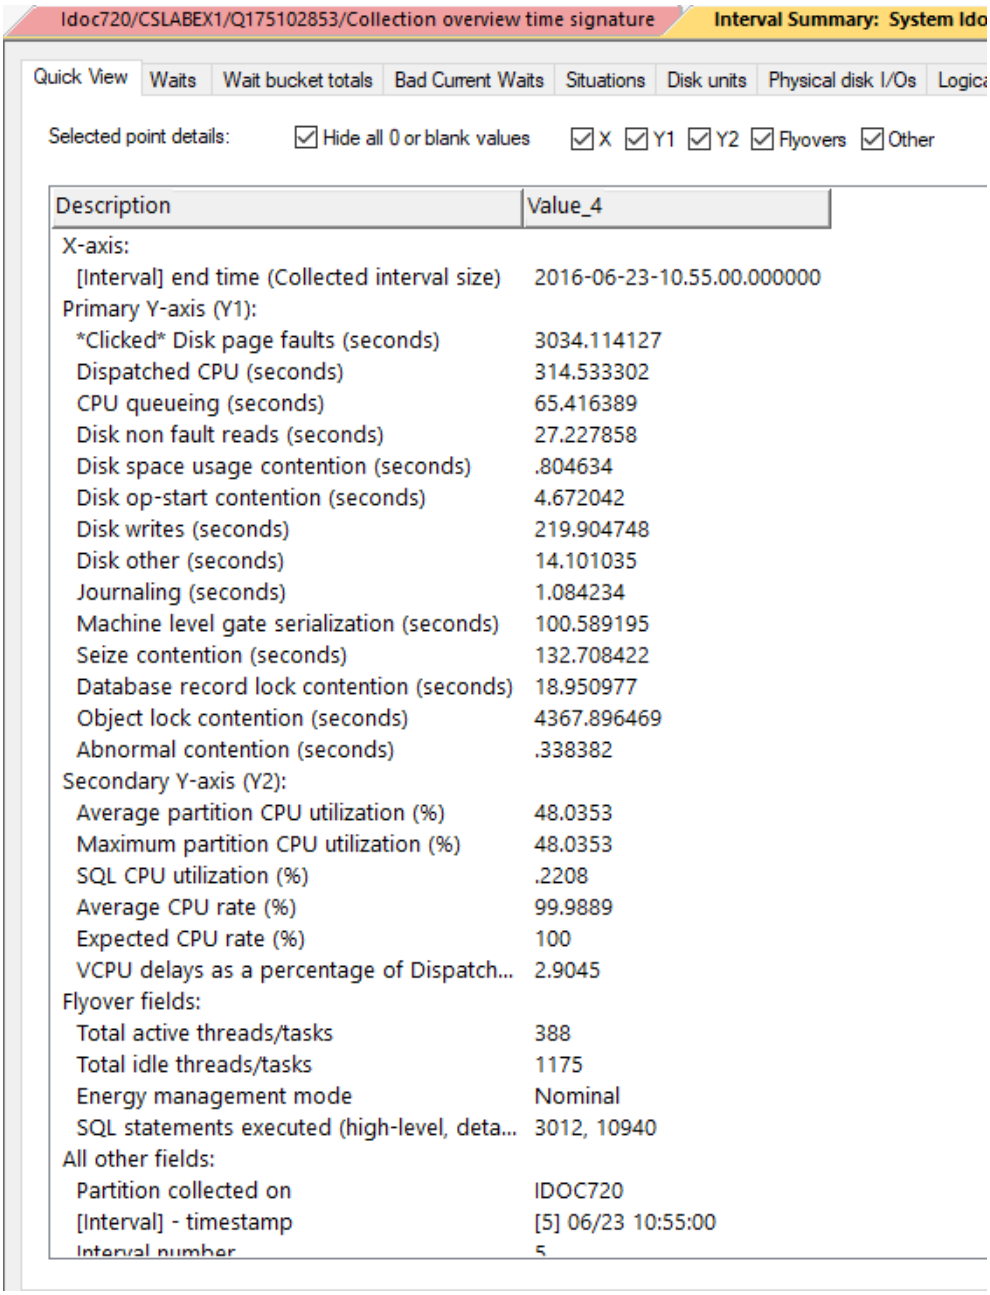

Description	Value_4
X-axis:	
[Interval] end time (Collected interval size)	2016-06-23-10.55.00.000000
Primary Y-axis (Y1):	
Clicked Disk page faults (seconds)	3034.114127
Dispatched CPU (seconds)	314.533302
CPU queueing (seconds)	65.416389
Disk non fault reads (seconds)	27.227858
Disk space usage contention (seconds)	.804634
Disk op-start contention (seconds)	4.672042
Disk writes (seconds)	219.904748
Disk other (seconds)	14.101035
Journaling (seconds)	1.084234
Machine level gate serialization (seconds)	100.589195
Seize contention (seconds)	132.708422
Database record lock contention (seconds)	18.950977
Object lock contention (seconds)	4367.896469
Abnormal contention (seconds)	.338382
Secondary Y-axis (Y2):	
Average partition CPU utilization (%)	48.0353
Maximum partition CPU utilization (%)	48.0353
SQL CPU utilization (%)	.2208
Average CPU rate (%)	99.9889
Expected CPU rate (%)	100
VCPU delays as a percentage of Dispatch...	2.9045
Flyover fields:	
Total active threads/tasks	388
Total idle threads/tasks	1175
Energy management mode	Nominal
SQL statements executed (high-level, deta...	3012, 10940
All other fields:	
Partition collected on	IDOC720
[Interval] - timestamp	[5] 06/23 10:55:00
Interval number	5

Graph Properties – Quick View Example

Tip: The check boxes provide options for hiding or showing different parts of the graph's data in the interface.

9.8.2 SQL

This tab is part of the property pages for a graph view. This interface is covered in the Main Window PDF under The Main Window -> Object Properties -> SQL.

9.8.3 Columns

This tab is part of the property pages for a graph view. This interface is covered in the Main Window PDF under The Main Window -> Object Properties -> Columns.

9.9 Graph Definitions

In iDoctor, users can define graphs over data generated by any SQL statement desired. Graphs are defined using a graph definition in iDoctor. Graph definitions supply the unique information that builds a user-defined or iDoctor-supplied graph. Like query definitions, graph definitions are stored in a report database.

A graph definition defines everything needed to display the graph including the query definition (SQL statement.) Whenever a graph definition is saved, the query definition is also saved.

The menu to create a new graph is the **Graph Definition -> Define New...** option within a table view. This action will display the graph definition interface. You can also open the graph definition interface for a graph using the graph definition menu or by double-clicking the legend in a graph.

There are several panels in the interface that make up the graph definition. These pages are discussed in greater detail in the next sections. A summary of the pages that make up a graph definition is shown below:

Page Name	Description
General	Defines the general features of the graph, like the type of graph and the graph's title.
X-axis	Defines the field(s) to show on the X-axis as well as the text to separate them if desired.
Primary Y-axis	This page defines the fields, colors, patterns and descriptions to use for the bars in the graph. Up to 32 different fields/colors may be defined in the graph definition.
Secondary Y-axis	This page identifies the secondary Y-axis. This axis consists of multiple lines of the desired color and width. This axis can only be shown on vertical bar graphs.
Flyover	Lists the additional fields to show when the mouse is placed over a bar in the graph.
SQL	This tab contains the parameterized SQL statement for the graph. iDoctor uses many parameters which are replaced at run-time and these parameter markers are revealed with this view.

9.9.1 General

The general page lets the user define the graph description, the type of graph to display as well as the number of bars to show on the graph if this should differ from the value shown on the preferences interface.

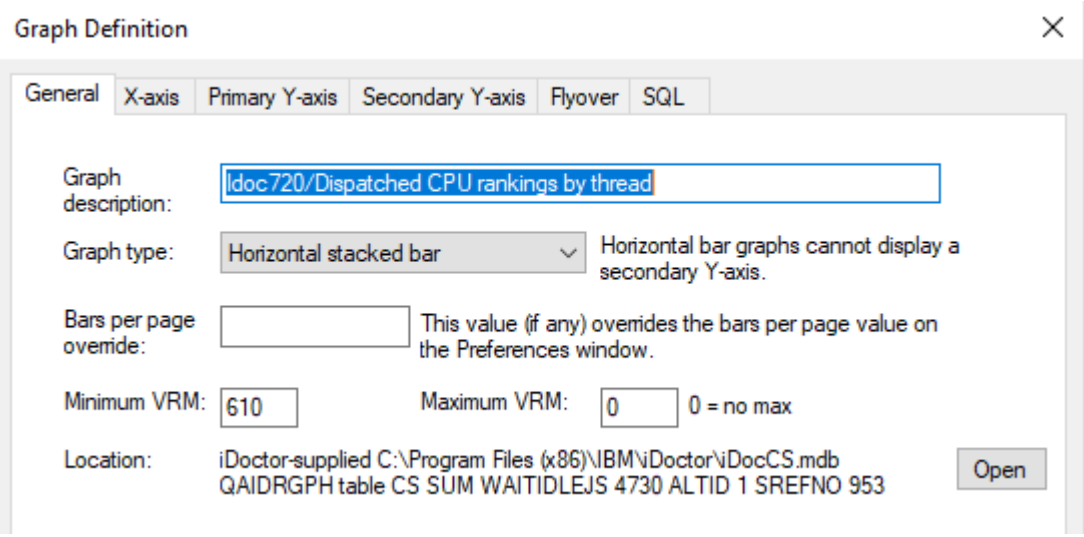

Graph Definition General Tab

Option	Description
Graph Description	A description of the graph that is displayed as the graph's title.
Graph type	This identifies the type of graph and refers to the look of the Primary Y-axis specifically.
Bars per page override	This value can be used to optionally specify the number of bars to show per page on this graph. If a value is not specified on this page, then the applicable bars per page value on the Preferences window will be used instead.
Minimum VRM	This is the minimum VRM of the collection data that this graph will appear for. This is checked against the IBM i OS VRM for the system that created the data collection. **Note:** For iDoctor-supplied graphs, this value is read-only.
Maximum VRM	This is the maximum VRM of the collection data that this graph will appear for. This is checked against the IBM i OS VRM for the system that created the data collection. **Note:** For iDoctor-supplied graphs, this value is read-only.
Location	This identifies where the graph definition is located within the indicated report database.
Open button	This will open the report database where this graph definition is located. If the DB is located on the PC then this will only work if MS Access is installed.

9.9.2 X-axis

The X-axis page lets you define the fields to display as the X-axis label. Up to 3 fields may be used to make up the label and the text to separate each field may be specified here as well.

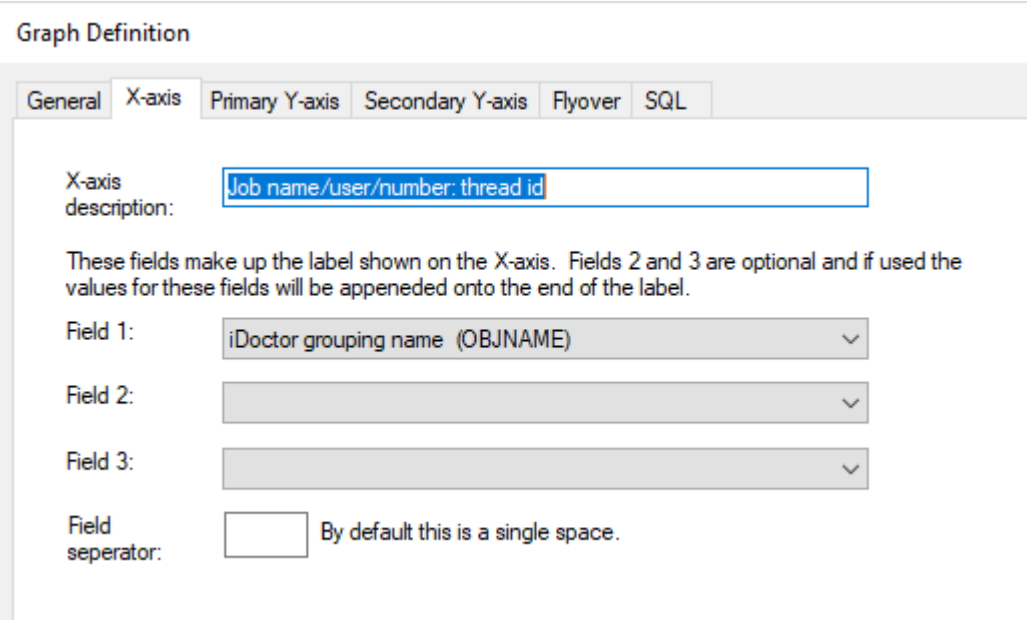

Graph Definition X-axis

Options	Description
X-Axis Description	The description to display under the X-Axis on the graph.
Field 1	The field to use as the 1st X-axis field on the graph.
Field 2	The optional field to use as the 2nd X-axis field on the graph.
Field 3	The optional field to use as the 3rd X-axis field on the graph.
Field separator	The text to separate the fields in the X-axis label if multiple fields are used. By default, this is a single space.

9.9.3 Primary Y-axis

Use the Primary Y-axis page to define the fields that should be displayed on Y1-axis for the graph. Each field has several options to customize how it appears on the graph.

Tip: In iDoctor if the graph type is a horizontal (bar) graph then the Y1-axis appears at the bottom of the graph and not the left.

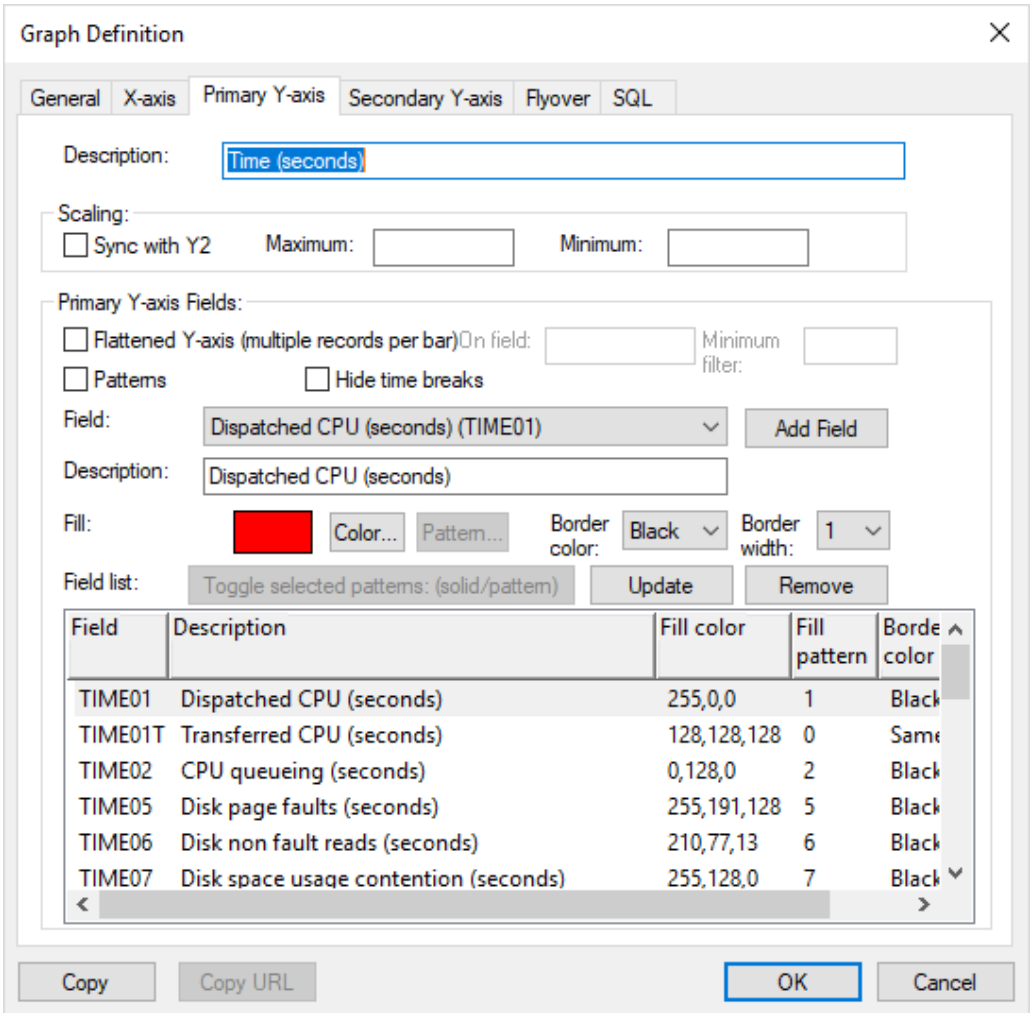

Graph Definition Primary Y-axis

A table describing the options on this screen follows:

Option	Description
Primary Y-axis Description	A description of the primary Y-axis.
Scaling - Sync with Y2	This value indicates if the scaling (min/max values) of the primary Y-axis will be in sync with the secondary Y-axis or not.
Scaling - Maximum	This value (if not blank) will be the maximum value shown on the Y1 axis. **Note:** If values occur in the data beyond this value, then the graph data will be truncated.
Scaling - Minimum	This value (if not blank) will be the minimum value shown on the Y1 axis. **Note:** If values occur in the data beyond this value, then the graph data will be truncated.
Flattened Y-axis	If checked, the graph becomes a "flattened" graph. Flattened graphs are built such that each stacked bar is defined from multiple rows in the data. **Note:** Special SQL syntax is required for this to work properly. For an example, see the Memory pool graphs -> Flattened style in CSI. The SQL statement must contain a ROW_NUM field.
(Flattened-)On field	This value indicates the field the flattened graph is using to define the colors/values shown. The flattened-on field is used to determine the colors on the graph. For example, if this is a memory pool graph and the flattened-on field is POOL (the pool number), then pool 1 might be red, pool 2 green, etc.
(Flattened) Minimum Filter	This is used typically to filter the data in the graph by a maximum number of occurrences/colors. For example, if each color is a disk unit and the SQL is setup properly to rank the disk units, this could define the max disk units to include in the results. The SQL parameter used is <<STACKEDFILTER>>. If this parameter does not exist in the SQL statement, then modifying this value will have no effect. An example graph in iDoctor that use this setting is the CSI – System graphs - LPAR CPU time.
Patterns	This indicates if patterns (hatchings) will be shown on the graph. It overrides the same setting in **Preferences –> Display -> Patterns**.
Hide time breaks	This setting allows the user to turn off the vertical dashed lines used to indicate breaks between different collections when graphing multiple collections.
Field	Allows selection of a field to add to (or modify in) in the field list. Changing the field will update the description to match the field description for the selected field.
Description	The description that identifies the data in the graph for the current field. This description will be displayed in the graph's legend.
Fill Color Button	Displays the Color window so the color to use for the field selected (in the list) can be modified. If no color is selected a color will be randomly assigned.
Fill Color Pattern Button	Shows the Fill Pattern Selection window that allows a user to modify the hatching (pattern) to use for the current field. These patterns will only be shown if the Display Patterns checkbox is checked or the Display Patterns option within the Preferences interface is checked.

Border color	If the graph is a bar graph, then this changes the color for the border. The default border width is 1 pixel. The possible values are: 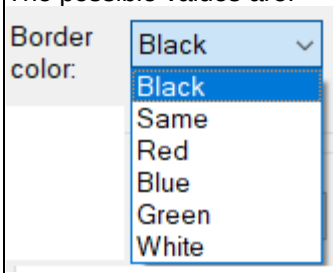 A value of “Same” means the color will match the fill color for the field (which in effect means no border will be visible. **Tip:** If the number of bars per page on the graph exceeds 250 then the border is automatically removed from the graph to avoid the graph appearance changing too much and looking more and more black with many bars. 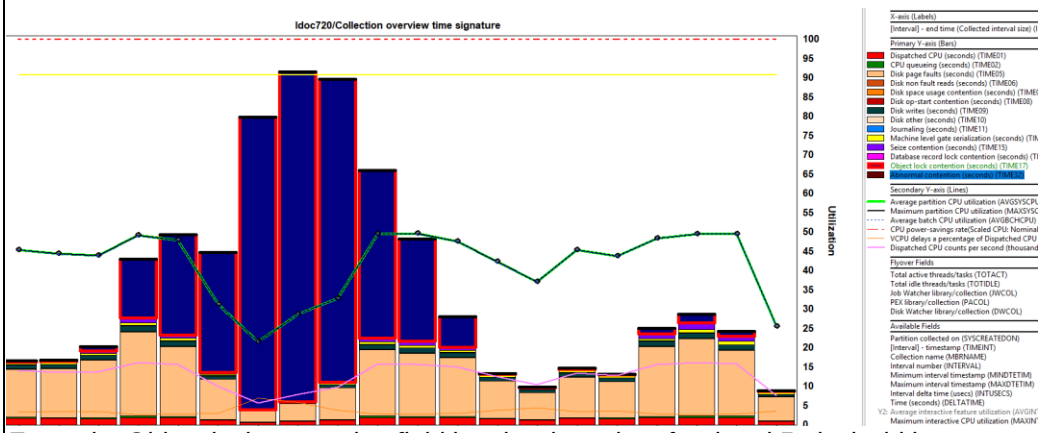 *Example: Object lock contention field has border color of red and 5 pixel width*
Border width	If the graph is a bar graph, then this changes the width for the border. The default border width is 1 pixel.
Toggle selected patterns	This option will toggle the selected fields' pattern setting to solid fill or back to a pattern. Use this button to more easily construct a graph where only 1 or 2 fields use a pattern and the rest show a solid fill pattern. **Tip:** This is only enabled if the Patterns checkbox is checked.
Update	The update button is used to modify the selected field in the field list.
Remove	This option will remove the selected fields from the list.
Field list	Displays the field names, descriptions and colors and more to use for the fields on the Y1-axis of the graph. By clicking OK on this window, any changes to this list will be applied to the graph.

9.9.3.1 To Add a Field

1. Select the field you wish to use for the new field from the Field drop-down list.
2. If desired, modify the description of the field from the field description.
3. If desired, define a color and pattern for this Y-Axis field. If this is not done, a color will be automatically assigned. The pattern only applies if the **Display -> Preferences -> Patterns** checkbox is checked.
4. Click the Add Field button to add the field to the list of fields.

9.9.4 Secondary Y-axis

This page allows the user to define a secondary Y-axis consisting of 1 or more lines on a graph. This axis is not visible for horizontal graphs and only will appear for vertical graphs. It will always appear on the right-side of the graph.

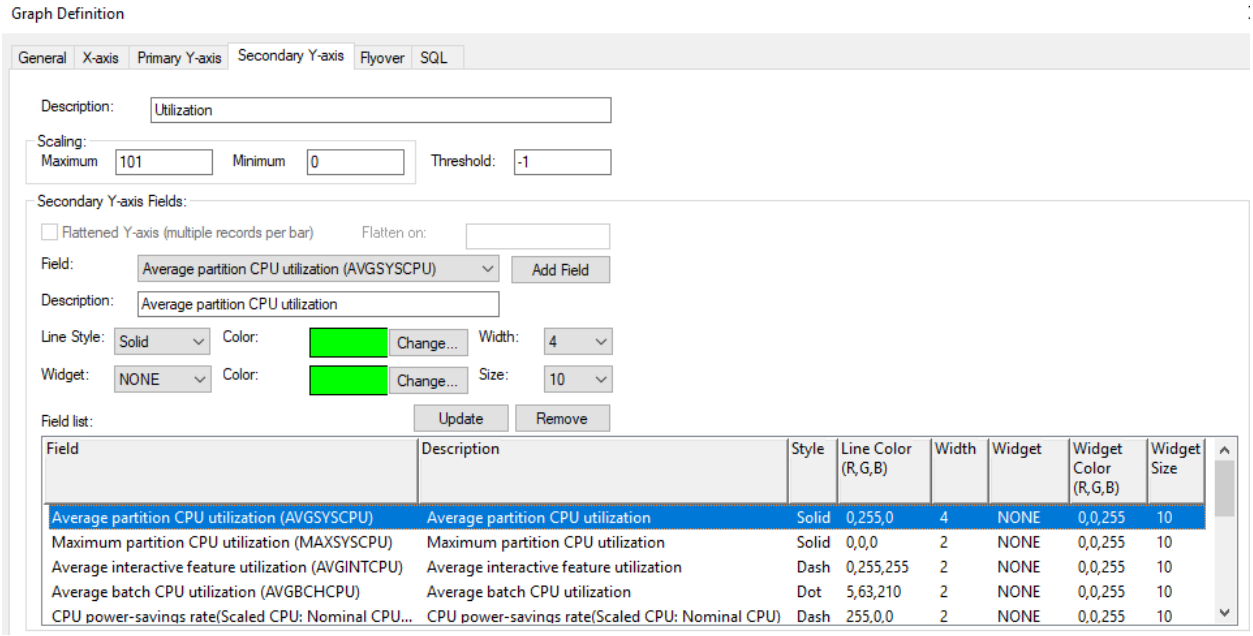

Graph Definition Secondary Y-axis

A table describing the options on this screen follows:

Option	Description
Description	The title to give the secondary Y-axis.
Scaling - Maximum	This value (if not blank) will be the maximum value shown on the axis.
Scaling - Minimum	This value (if not blank) will be the minimum value shown on the axis.
Scaling – Threshold	Specifying a value here (other than -1) will draw a yellow threshold line at the value given.
Flattened Y-axis	This value is always read on but included here as it does affect the Y2-axis. To change it you must do so from the Primary Y-axis tab. If checked, the graph becomes a "flattened" graph. Flattened graphs are built such that each stacked bar is defined from multiple rows in the data. **Note:** Special SQL syntax is required for this to work properly. For an example, see the Memory pool graphs -> Flattened style in CSI. The SQL statement must contain a ROW_NUM field.
Flatten on	This value is always read on but included here as it does affect the Y2-axis. To change it you must do so from the Primary Y-axis tab. This value indicates the field the flattened graph is using to define the colors/values shown. The flattened-on field is used to determine the colors on the graph. For example, if this is a memory pool graph and the flattened-on field is POOL (the pool number), then pool 1 might be red, pool 2 green, etc.
Field	The list of fields available to add to the secondary Y-axis.
Description	The description for the field to show on the legend to add to the secondary Y-axis.
Line Style	This changes the style for the line. The possible values are: 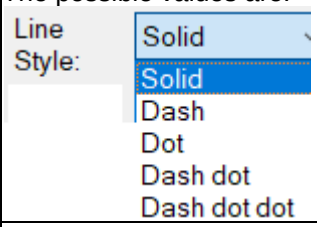
Line Color	The color to use for the current line.
Line width	The number of pixels wide to draw the current line.
Widget (name)	This allows each point of the line to contain an optional shape called a widget. These shapes have different names and the choices are: 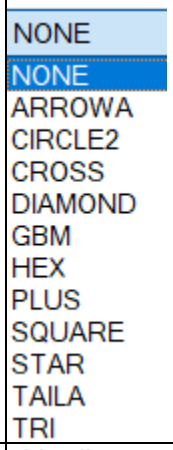
Widget color	This allows each point of the line to contain an optional shape called a widget. This is the color to draw the widget with. Press the Change button to change the color.

Widget size	This allows each point of the line to contain an optional shape called a widget. This is the size to draw the widget. The bigger the number, the bigger the widget will appear. 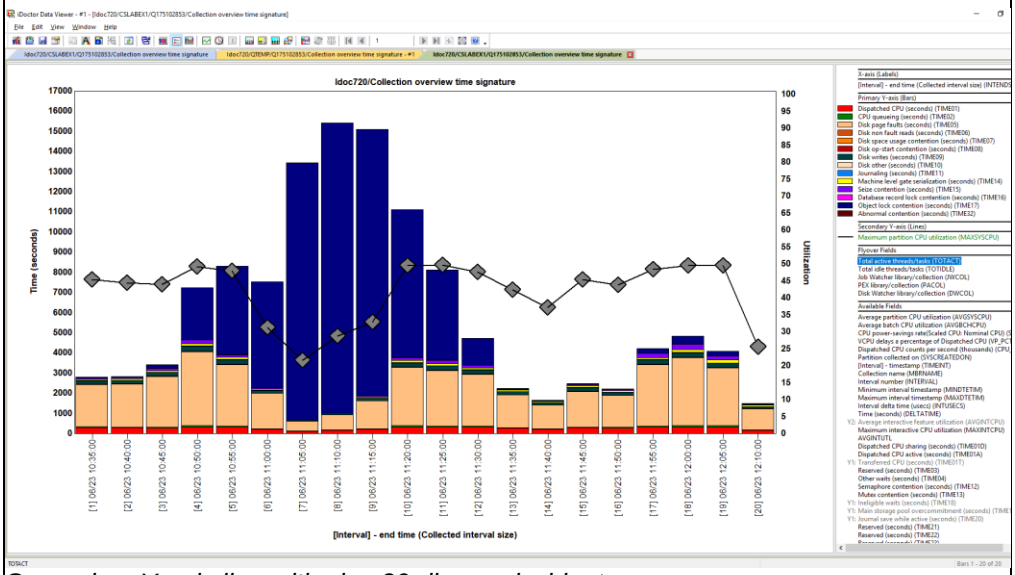 *Secondary Y-axis line with size 30 diamond widget*
Update	The update button is used to modify the selected field in the field list.
Remove	This option will remove the selected fields from the list.
Field list	Displays the field names, descriptions, colors and more to use for the fields on the Y2-axis of the graph.

9.9.5 Flyover

This page allows the user to define up to additional fields to show on the flyover window for the current graph. Flyovers are shown when the mouse is placed over a point of interest on the graph (like a bar or line).

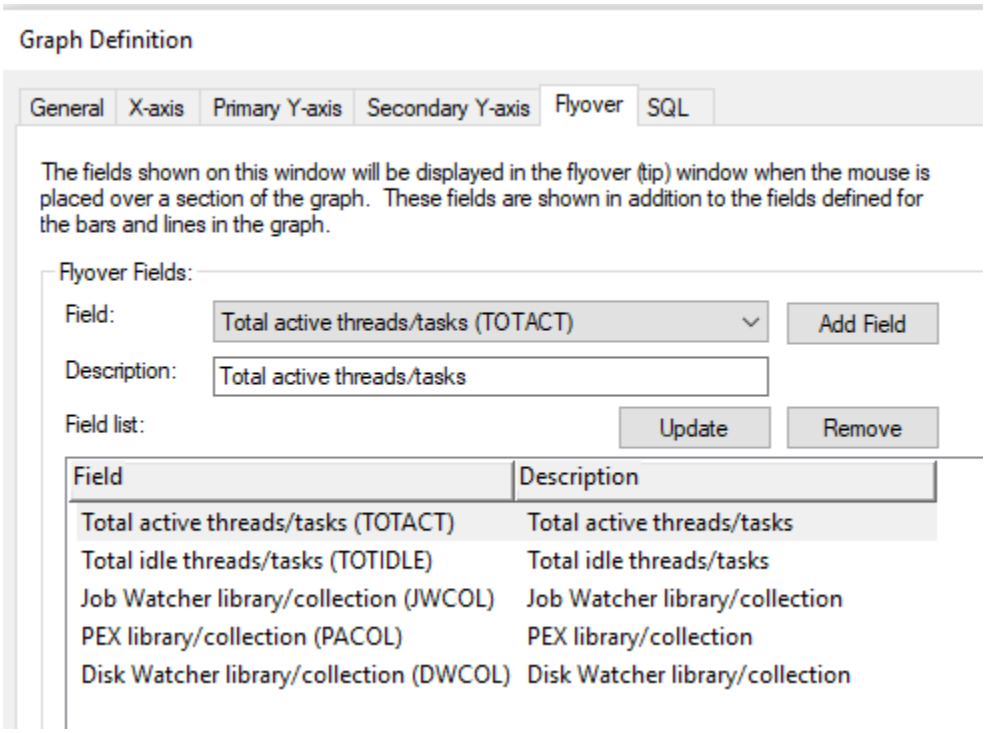

Graph Definition - Flyover

9.9.6 SQL

This tab displays the parameterized version of the SQL statement behind the graph. iDoctor uses many parameters for iDoctor-supplied SQL statements and this interface will reveal them.

The window also features a text field where the desired input can be searched for within the statement.

Note: The SQL statement cannot be edited via this interface unless you are viewing the graph definition via the User-Defined Reports Database. It cannot be edited for IBM defined graphs.

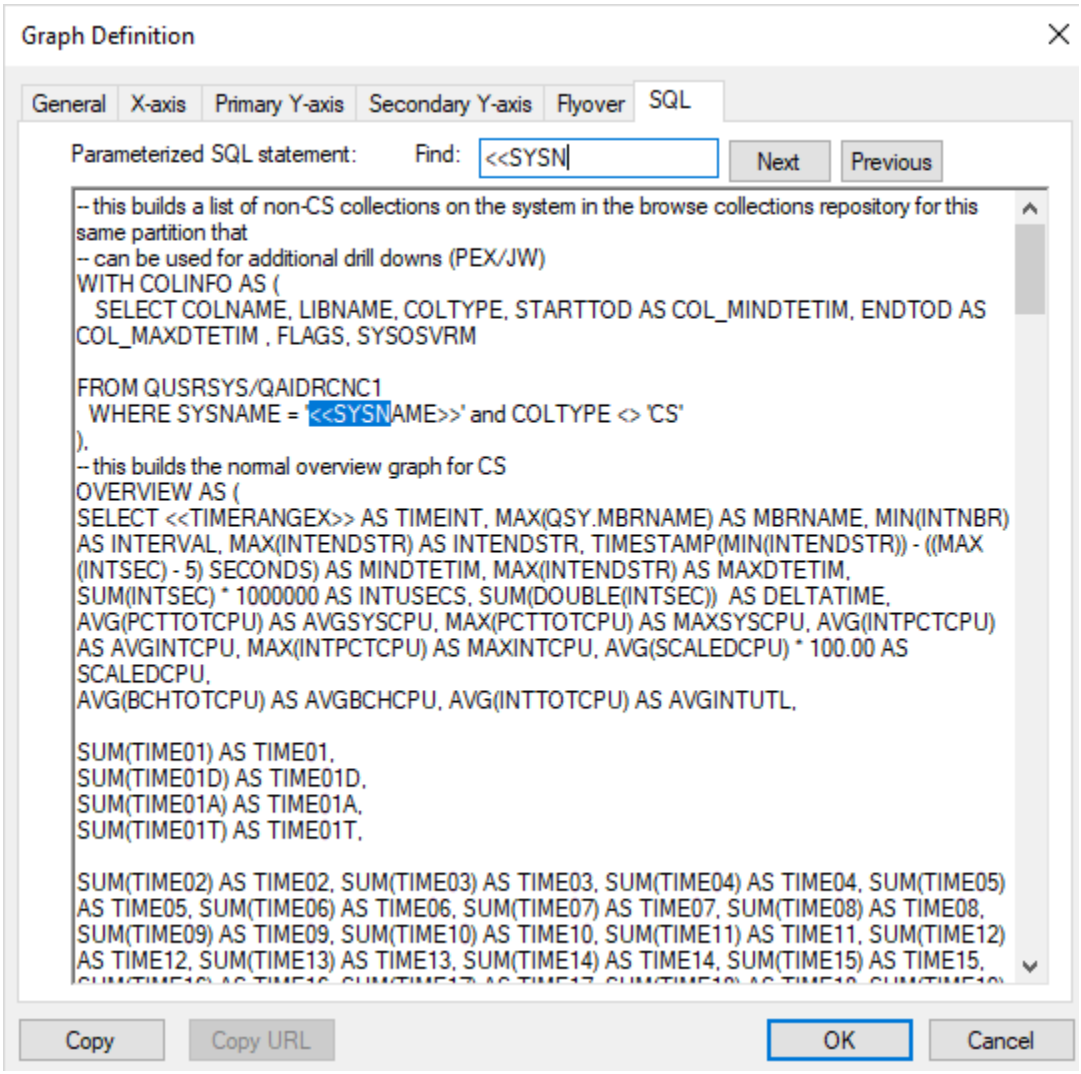

Graph Definition -> SQL

9.9.7 Save Graph Definition (Save As...)

Graph Definitions are saved using the Graph Definition -> Save As... menu for a graph view. All Graph Definitions are saved into the user-defined reports database. You can specify the folder name to save the query definition into within this database. These are accessed later from the User-defined reports folder for the collection type you are working with. Typically, these queries can be reused on any collection data of the same type (i.e. Job Watcher.)

An example of the Save Graph Definition interface is shown below:

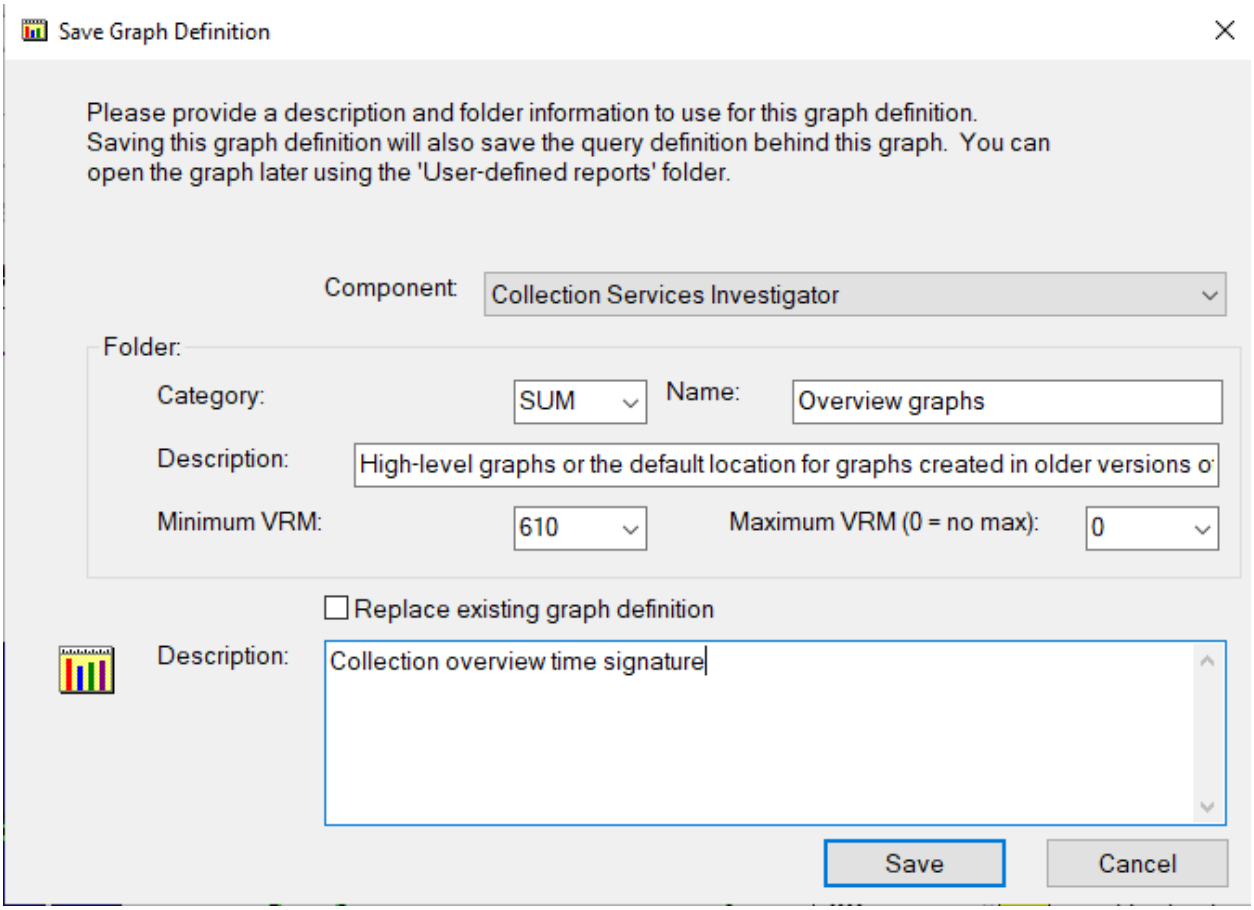

The interface elements within this window are described in more detail below:

GUI element	**Description**
Component	The name of the component this graph should be visible in.
Category	3-character identifier for the folder the graph should be saved into. If you wish to define new folder, give it a name not already in the drop down list like 'DX1'
(Folder) Name	The name of the folder to store the query into
(Folder) Description	A long description to give the user-defined folder. (optional)
Minimum VRM	The minimum IBM i VRM that this query should be visible too in nnn format (i.e. 610, 710, 730, etc.) If the collection was created on a system older than this value, then this query will not appear.

10 Synchronized Table View

Most graphs in iDoctor will have available beneath it a table view of data. This is the raw data behind the SQL statement used to build the graph. By default, this data is hidden from view unless you need it.

Press the  button on the Data Viewer toolbar while focus is on a graph view to open the table behind the graph. **Note:** This option is not available when using the ASP comparison functions.

This view works just like a regular table view except it is synchronized with the graph. Any scrolling and selections made are also done in the table and vice versa.

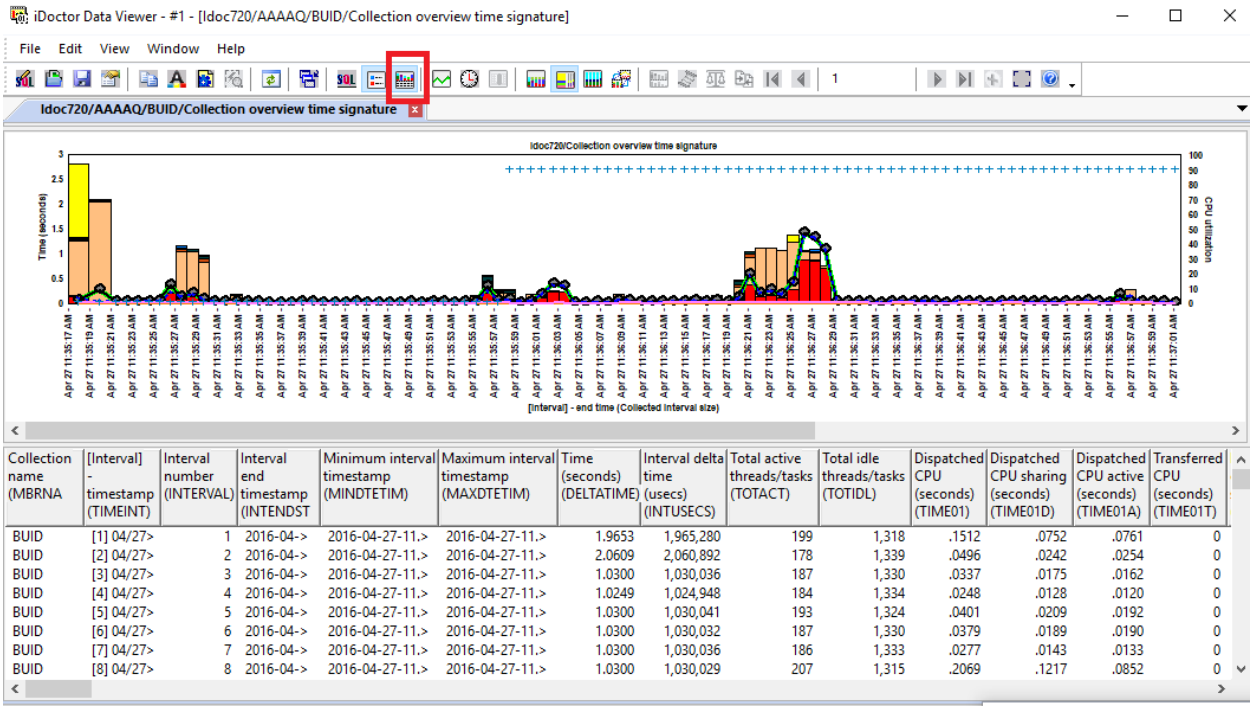

Collection name (MBRNA	[Interval] - timestamp (TIMEINT)	Interval number (INTERVAL)	Interval end timestamp (INTENDST	Minimum interval timestamp (MINDTETIM)	Maximum interval timestamp (MAXDTETIM)	Time (seconds) (DELTATIME)	Interval delta time (usecs) (INTUSECS)	Total active threads/tasks (TOTACT)	Total idle threads/tasks (TOTIDL)	Dispatched CPU (seconds) (TIME01)	Dispatched CPU sharing (seconds) (TIME01D)	Dispatched CPU active (seconds) (TIME01A)	Transferred CPU (seconds) (TIME01T)
BUID	[1] 04/27>	1	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.9653	1,965,280	199	1,318	.1512	.0752	.0761	0
BUID	[2] 04/27>	2	2016-04->	2016-04-27-11.>	2016-04-27-11.>	2.0609	2,060,892	178	1,339	.0496	.0242	.0254	0
BUID	[3] 04/27>	3	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0300	1,030,036	187	1,330	.0337	.0175	.0162	0
BUID	[4] 04/27>	4	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0249	1,024,948	184	1,334	.0248	.0128	.0120	0
BUID	[5] 04/27>	5	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0300	1,030,041	193	1,324	.0401	.0209	.0192	0
BUID	[6] 04/27>	6	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0300	1,030,032	187	1,330	.0379	.0189	.0190	0
BUID	[7] 04/27>	7	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0300	1,030,036	186	1,333	.0277	.0143	.0133	0
BUID	[8] 04/27>	8	2016-04->	2016-04-27-11.>	2016-04-27-11.>	1.0300	1,030,029	207	1,315	.2069	.1217	.0852	0

Synchronized Table View and Graph

11 Alternate Views

This feature allows a user to quickly change a graph and show it in a different way using data already loaded.

Tip: New in 2021 is the ability to modify the look of some tables. Enhancements were added when working with Job Summary reports in CSI and the QAPMSYSTEM table.

Press the  button to activate this pane when looking at a graph or table.

Be aware that some alternate view graphs will require the SQL statement to be reran and some do not. If the SQL statement contains parameters that are different for the alternate view then the SQL query must be reran. "Flattened-style" graphs also must be ran again.

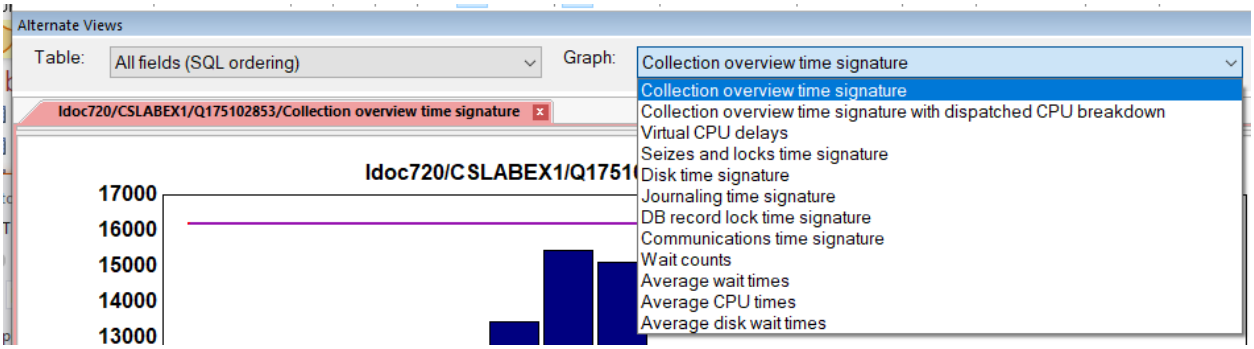

Alternate Views Pane

Option	Description
Table	This option lets you modify the table below the graph in various ways. All graphs supply these options: 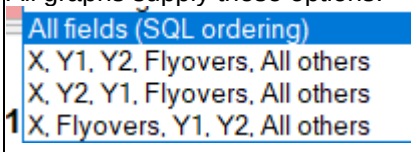 Some graphs have additional options like the CSI Collection overview graph also includes these: 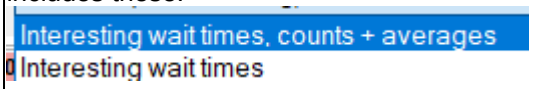 For some tables it allows you to modify the columns shown if special content has been added to iDoctor. The places so far this support has been added include: - CSI Job Summary SQL tables - CSI QAPMSYSTEM file
Graph	This is the available list of alternate view graphs. Selecting one from the list will change the graph.

11.1 Table alternate view example

If a user wants to see the Y2 fields near the front of the table, they can choose the "X, Y2, Y1, Flyovers, All Others" option from the Table drop down list.

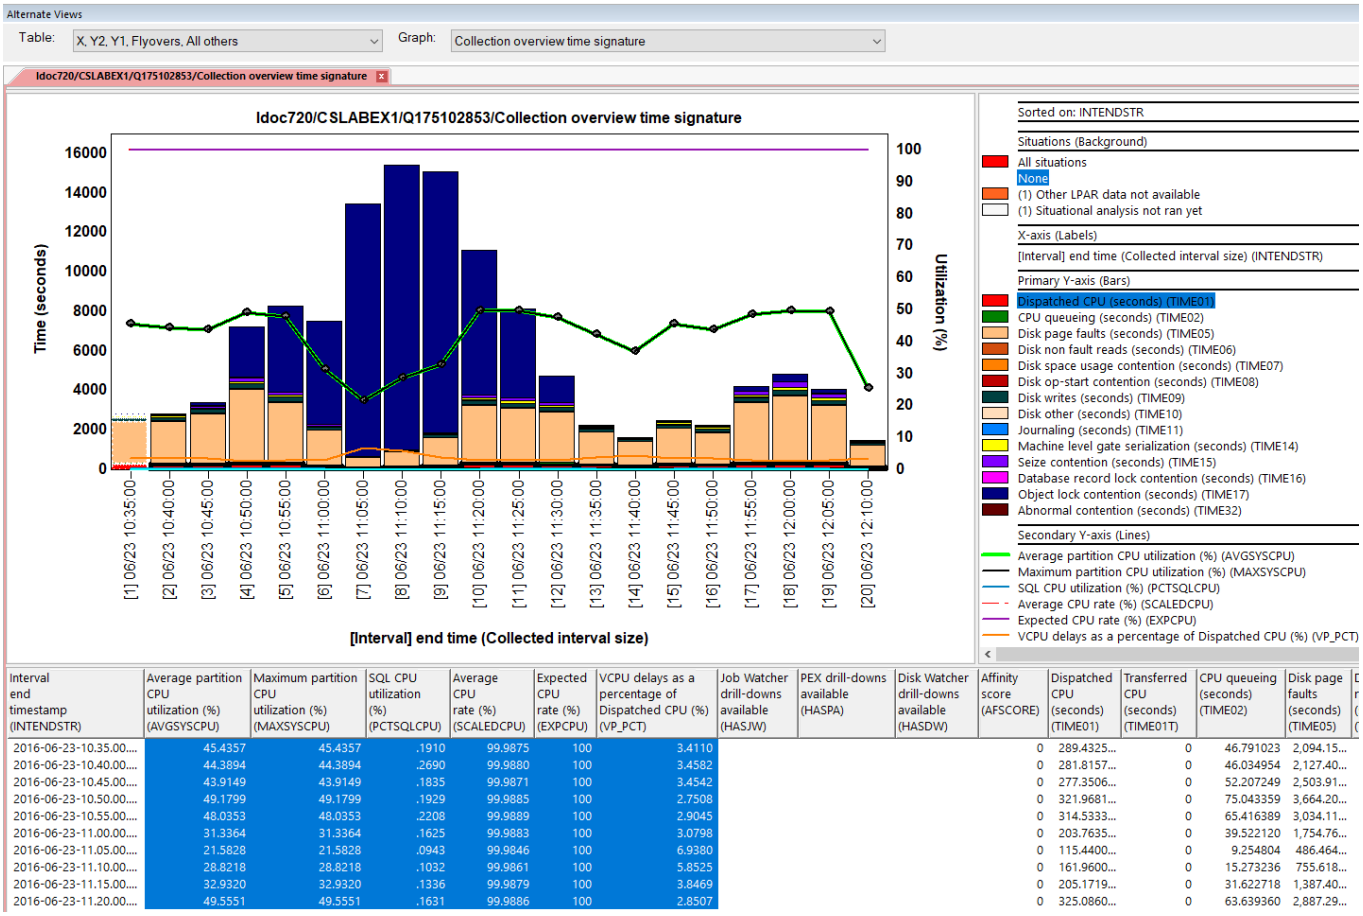

Interval end timestamp (INTENDSTR)	Average partition CPU utilization (%) (AVGSYSCPU)	Maximum partition CPU utilization (%) (MAXSYSCPU)	SQL CPU utilization (%) (PCTSQLCPU)	Average CPU rate (%) (SCALEDCPU)	Expected CPU rate (%) (EXPCPU)	VCPU delays as a percentage of Dispatched CPU (%) (VP_PCT)	Job Watcher drill-downs available (HASJW)	PEX drill-downs available (HASPA)	Disk Watcher drill-downs available (HASDW)	Affinity score (AFSCORE)	Dispatched CPU (seconds) (TIME01)	Transferred CPU (seconds) (TIME01T)	CPU queueing (seconds) (TIME02)	Disk page faults (seconds) (TIME05)
2016-06-23-10.35.00....	45.4357	45.4357	.1910	99.9875	100	3.4110				0	289.4325...	0	46.791023	2,094.15...
2016-06-23-10.40.00....	44.3894	44.3894	.2690	99.9880	100	3.4582				0	281.8157...	0	46.034954	2,127.40...
2016-06-23-10.45.00....	43.9149	43.9149	.1835	99.9871	100	3.4542				0	277.3506...	0	52.207249	2,503.91...
2016-06-23-10.50.00....	49.1799	49.1799	.1929	99.9885	100	2.7508				0	321.9681...	0	75.043359	3,664.20...
2016-06-23-10.55.00....	48.0353	48.0353	.2208	99.9889	100	2.9045				0	314.5333...	0	65.416389	3,034.11...
2016-06-23-11.00.00....	31.3364	31.3364	.1625	99.9883	100	3.0798				0	203.7635...	0	39.522120	1,754.76...
2016-06-23-11.05.00....	21.5828	21.5828	.0943	99.9846	100	6.9380				0	115.4400...	0	9.254804	486.464...
2016-06-23-11.10.00....	28.8218	28.8218	.1032	99.9861	100	5.8525				0	161.9600...	0	15.273236	755.618...
2016-06-23-11.15.00....	32.9320	32.9320	.1336	99.9879	100	3.8469				0	205.1719...	0	31.622718	1,387.40...
2016-06-23-11.20.00....	49.5551	49.5551	.1631	99.9886	100	2.8507				0	325.0860...	0	63.639360	2,887.29...

Table Alternate View: X, Y2, Y1, Flyovers, All others

11.2 Interesting waits times, counts + averages example

This option is used to reorder fields and put all information about each wait bucket together in the list.

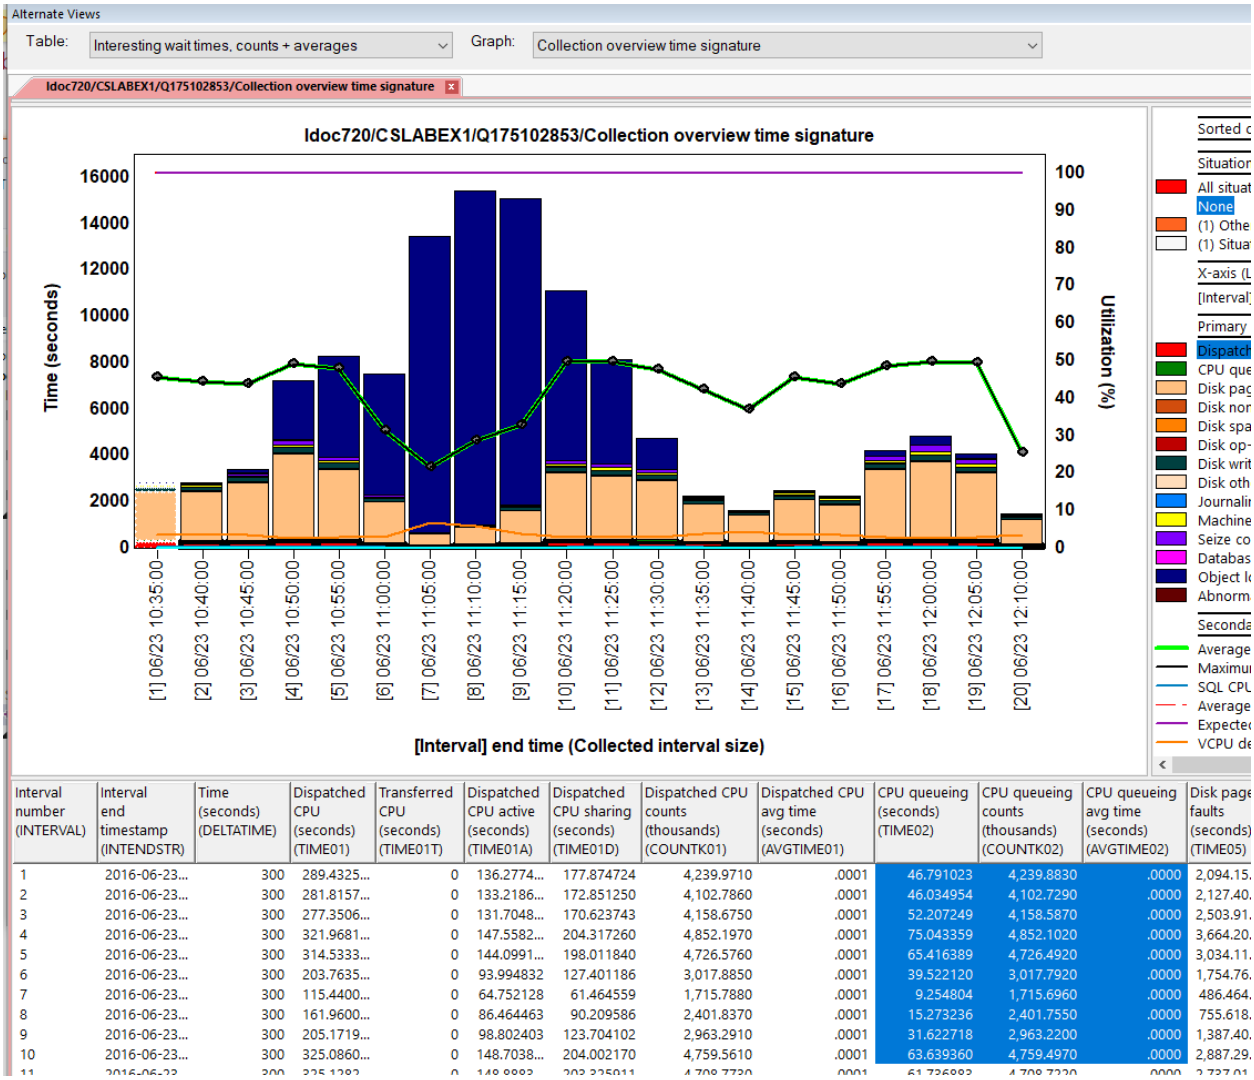

Interval number (INTERVAL)	Interval end timestamp (INTENDSTR)	Time (seconds) (DELTATIME)	Dispatched CPU (seconds) (TIME01)	Transferred CPU (seconds) (TIME01T)	Dispatched CPU active (seconds) (TIME01A)	Dispatched CPU sharing (seconds) (TIME01D)	Dispatched CPU counts (thousands) (COUNTK01)	Dispatched CPU avg time (seconds) (AVGTIME01)	CPU queueing (seconds) (TIME02)	CPU queueing counts (thousands) (COUNTK02)	CPU queueing avg time (seconds) (AVGTIME02)	Disk page faults (seconds) (TIME05)
1	2016-06-23...	300	289.4325...	0	136.2774...	177.874724	4,239.9710	.0001	46.791023	4,239.8830	.0000	2,094.15.
2	2016-06-23...	300	281.8157...	0	133.2186...	172.851250	4,102.7860	.0001	46.034954	4,102.7290	.0000	2,127.40.
3	2016-06-23...	300	277.3506...	0	131.7048...	170.623743	4,158.6750	.0001	52.207249	4,158.5870	.0000	2,503.91.
4	2016-06-23...	300	321.9681...	0	147.5582...	204.317260	4,852.1970	.0001	75.043359	4,852.1020	.0000	3,664.20.
5	2016-06-23...	300	314.5333...	0	144.0991...	198.011840	4,726.5760	.0001	65.416389	4,726.4920	.0000	3,034.11.
6	2016-06-23...	300	203.7635...	0	93.994832	127.401186	3,017.8850	.0001	39.522120	3,017.7920	.0000	1,754.76.
7	2016-06-23...	300	115.4400...	0	64.752128	61.464559	1,715.7880	.0001	9.254804	1,715.6960	.0000	486.464.
8	2016-06-23...	300	161.9600...	0	86.464463	90.209586	2,401.8370	.0001	15.273236	2,401.7550	.0000	755.618.
9	2016-06-23...	300	205.1719...	0	98.802403	123.704102	2,963.2910	.0001	31.622718	2,963.2200	.0000	1,387.40.
10	2016-06-23...	300	325.0860...	0	148.7038...	204.002170	4,759.5610	.0001	63.639360	4,759.4970	.0000	2,887.29.

CSI Collection Overview with Interesting wait times, counts + averages table alternate view option

11.3 QAPMSYSTEM example

To reorder columns in this table in different ways, open the table under Server-side output files, then use the Alternate Views pane.

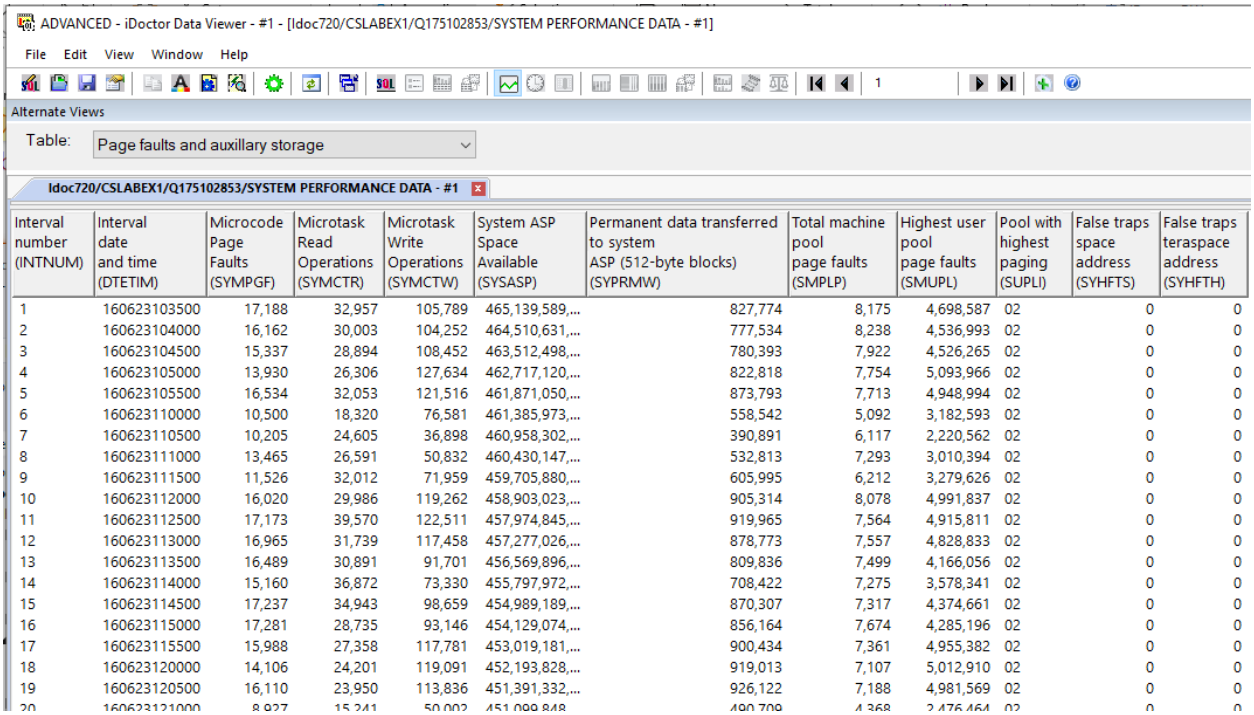

Interval number (INTNUM)	Interval date and time (DTETIM)	Microcode Page Faults (SYMPGF)	Microtask Read Operations (SYMCTR)	Microtask Write Operations (SYMCTW)	System ASP Space Available (SYSASP)	Permanent data transferred to system ASP (512-byte blocks) (SYPRMW)	Total machine pool page faults (SMPLP)	Highest user pool page faults (SMUPL)	Pool with highest paging (SUPLI)	False traps space address (SYHFTS)	False traps teraspace address (SYHFTH)
1	160623103500	17,188	32,957	105,789	465,139,589,...	827,774	8,175	4,698,587	02	0	0
2	160623104000	16,162	30,003	104,252	464,510,631,...	777,534	8,238	4,536,993	02	0	0
3	160623104500	15,337	28,894	108,452	463,512,498,...	780,393	7,922	4,526,265	02	0	0
4	160623105000	13,930	26,306	127,634	462,717,120,...	822,818	7,754	5,093,966	02	0	0
5	160623105500	16,534	32,053	121,516	461,871,050,...	873,793	7,713	4,948,994	02	0	0
6	160623110000	10,500	18,320	76,581	461,385,973,...	558,542	5,092	3,182,593	02	0	0
7	160623110500	10,205	24,605	36,898	460,958,302,...	390,891	6,117	2,220,562	02	0	0
8	160623111000	13,465	26,591	50,832	460,430,147,...	532,813	7,293	3,010,394	02	0	0
9	160623111500	11,526	32,012	71,959	459,705,880,...	605,995	6,212	3,279,626	02	0	0
10	160623112000	16,020	29,986	119,262	458,903,023,...	905,314	8,078	4,991,837	02	0	0
11	160623112500	17,173	39,570	122,511	457,974,845,...	919,965	7,564	4,915,811	02	0	0
12	160623113000	16,965	31,739	117,458	457,277,026,...	878,773	7,557	4,828,833	02	0	0
13	160623113500	16,489	30,891	91,701	456,569,896,...	809,836	7,499	4,166,056	02	0	0
14	160623114000	15,160	36,872	73,330	455,797,972,...	708,422	7,275	3,578,341	02	0	0
15	160623114500	17,237	34,943	98,659	454,989,189,...	870,307	7,317	4,374,661	02	0	0
16	160623115000	17,281	28,735	93,146	454,129,074,...	856,164	7,674	4,285,196	02	0	0
17	160623115500	15,988	27,358	117,781	453,019,181,...	900,434	7,361	4,955,382	02	0	0
18	160623120000	14,106	24,201	119,091	452,193,828,...	919,013	7,107	5,012,910	02	0	0
19	160623120500	16,110	23,950	113,836	451,391,332,...	926,122	7,188	4,981,569	02	0	0
20	160623121000	8,927	15,241	50,002	451,099,848	490,709	4,368	2,476,464	02	0	0

QAPMSYSTEM with Page faults and auxiliary storage perspective enabled

The possible choices for this file are:

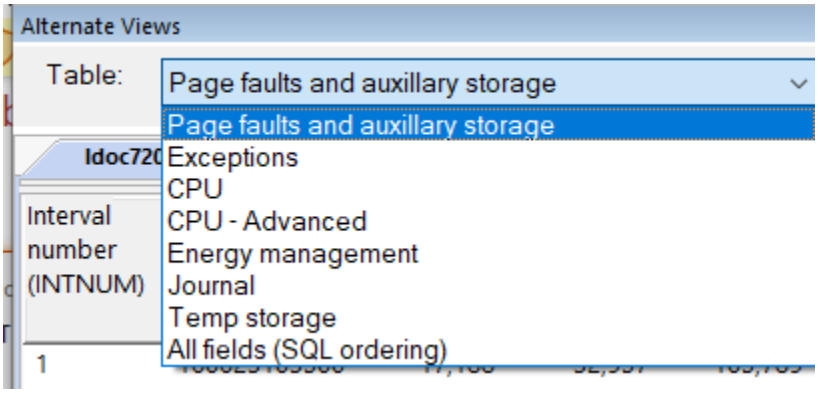

Alternate Views Table options for QAPMSYSTEM

12 Set time grouping (clock icon)

The clock icon on the Data Viewer toolbar allows a user to modify how a time-based graph is grouped. Many groupings are available and make it easier to get the big picture view of thousands of records of data more quickly rather than scrolling through it or trying to graph thousands of points on a single screen (often not feasible!)

To enable this option, press this  button and a list of options will be shown. Picking the desired grouping will rerun the SQL statement using the desired time grouping.

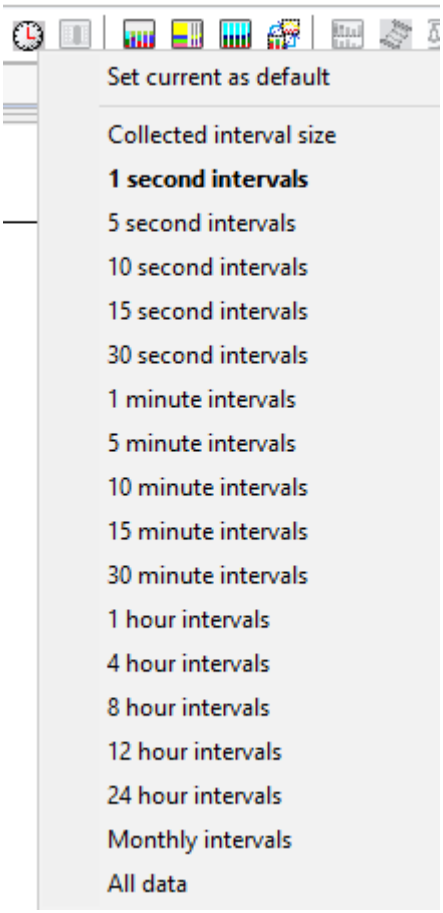

Clock icon menu

Note: The number of options here will vary a bit based on the time grouping that was used to create the current collection. For example, if using Collection Services and the collection was created with 15 min time intervals, then no options between 1 second and 15 minutes would be shown.

Some examples of how these settings will change the graph X-axis labels are described below:

Option	Description
Set current as default	This will use the currently shown grouping on the graph and use this is a default preference for all iDoctor graphs
Collected interval size	This will graph the data at whatever grouping that was used to create the original data. This option will include the interval number in the X-axis label (unless variable width bar mode is enabled.) 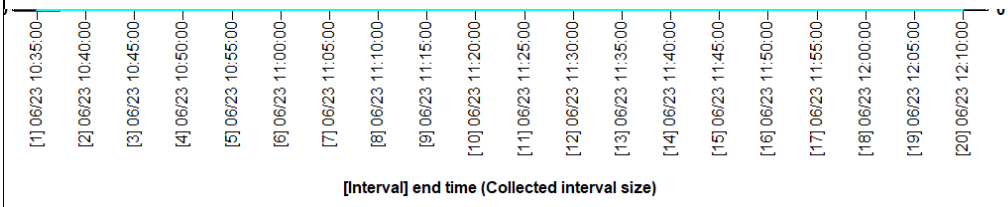
1 second intervals	Groups the data at 1 second intervals 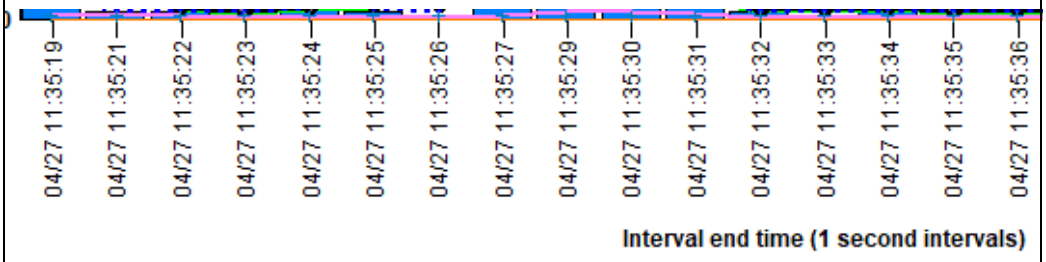
5 second intervals	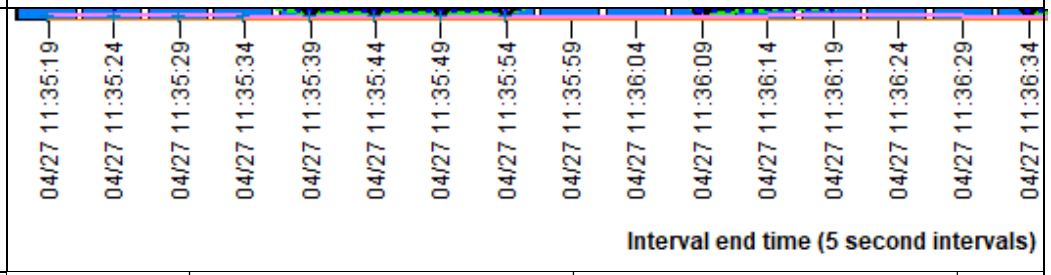
1-hour intervals	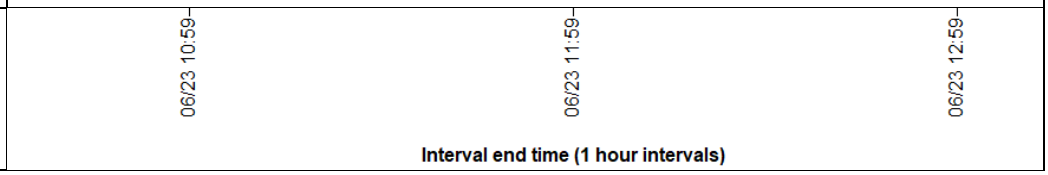

13 Normalize option

This option  can be useful if there are variations in the time taken to produce intervals shown on the graph. For example, if Job Watcher was slow to initialize and the 1st 2 intervals took 10 times longer than the rest of the intervals, the rest of the intervals may be barely visible unless this option is turned on.

When normalizing a graph, each value on the primary Y axis, is divided by the interval's duration applicable to each Y axis value. This can create a "flattening" effect to bar heights by drawing time values based on relative contributions. Keep in mind that the graph was changed to be percentage based on the amount of time each interval took and what you are seeing are not the true values collected.

The following example shows the difference between a graph with the normalize option NOT used and then with it applied. In this case the Job Watcher collection had a longer than normal 2nd interval because the option to collect all threads active or idle on interval 1 was enabled.

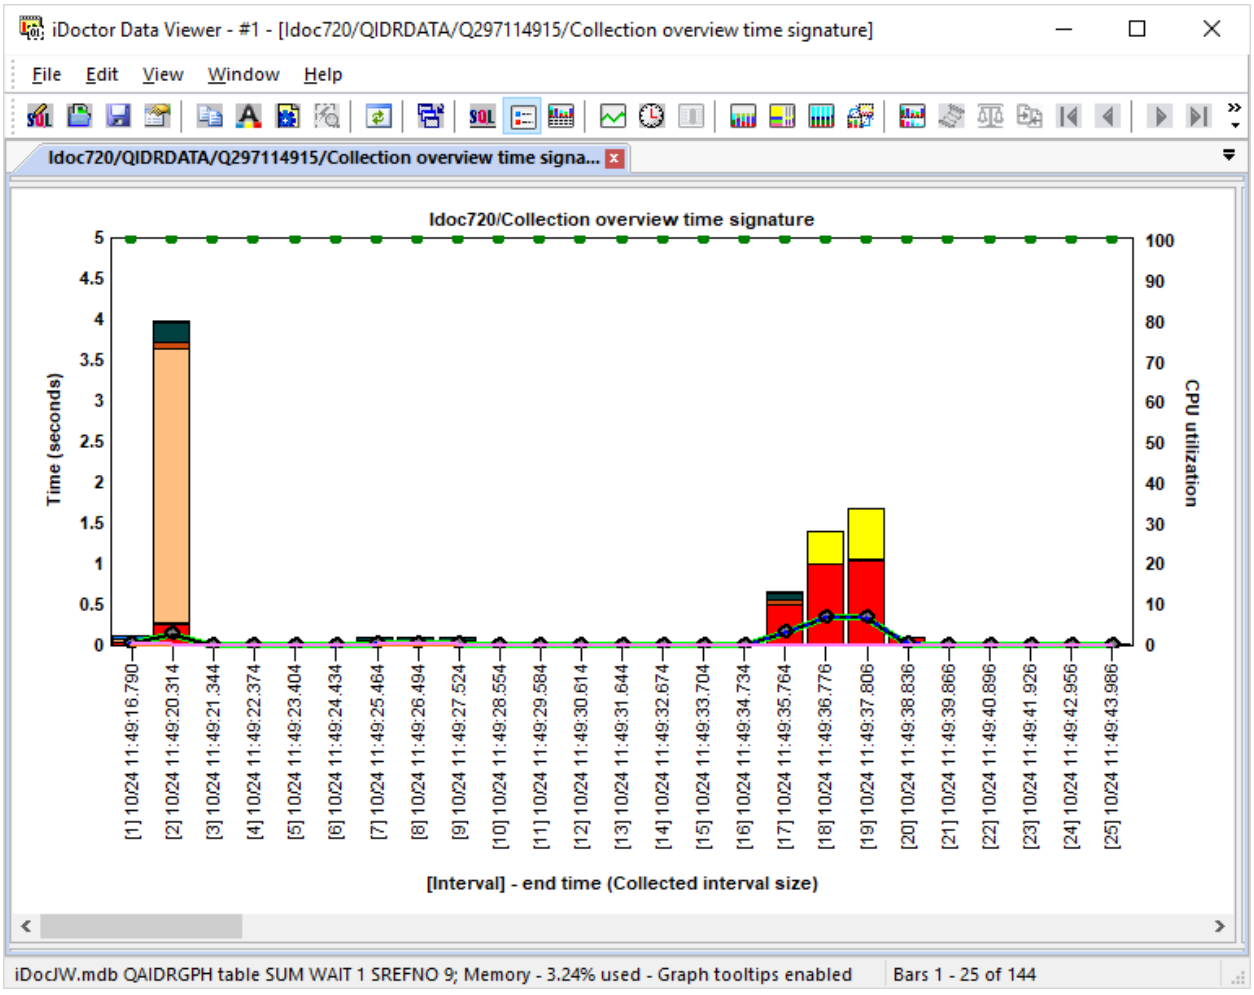

Collection overtime graph

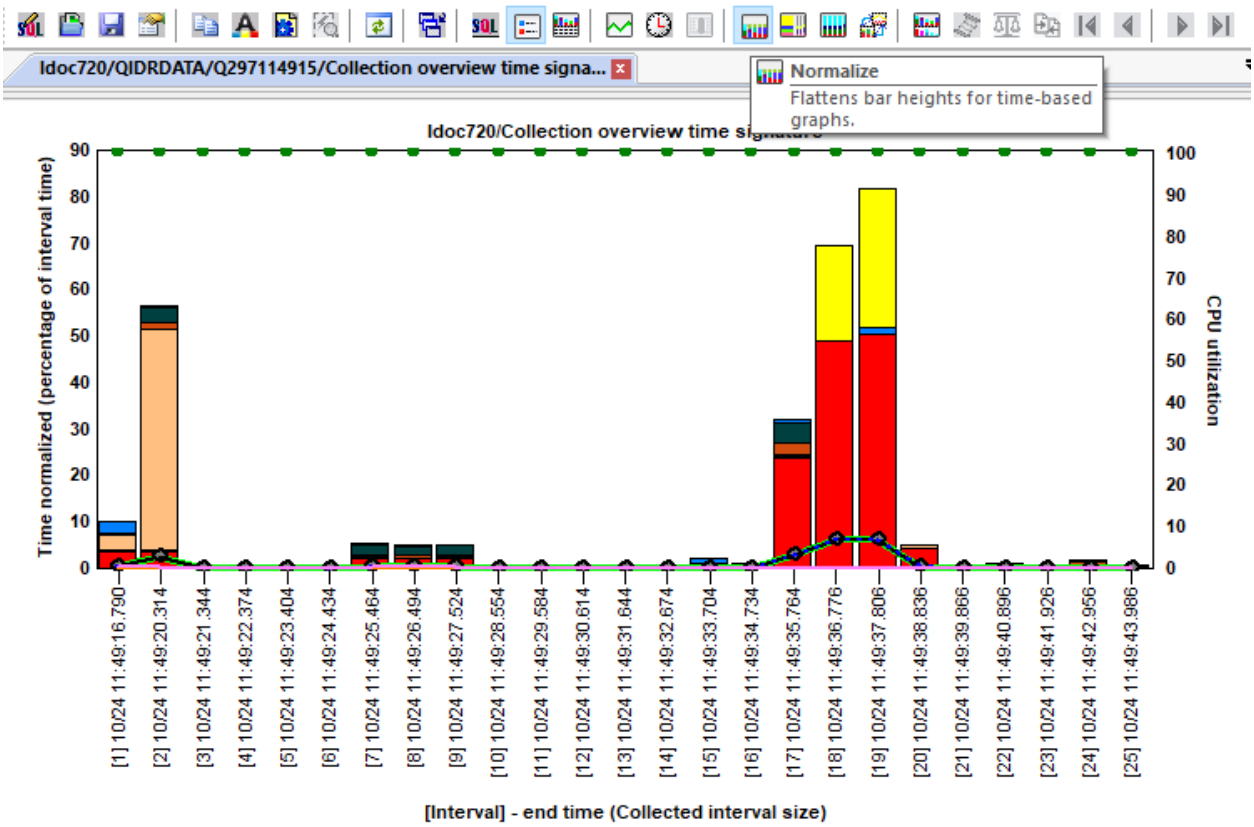

Collection overview graph with Normalize option

14 Variable-width bar mode option

This option  can also be useful if there are variations in the time taken to produce intervals shown on the graph. For example, if Job Watcher was slow to initialize and the 1st 2 intervals took 10 times longer than the rest of the intervals, then this fact may not be obvious unless this option is turned on.

This graphing technique will change the graph so the amount of time each interval took is visually graphed when normally they are not. Longer intervals will have wider bars and short intervals will have thinner bars. This mode also changes the X-axis labeling so that each bar is not labeled. Instead, the true time periods are graphed where they occurred.

In the following example the 2nd interval was 4 times longer than the others and can be easily seen with variable-width bar mode enabled.

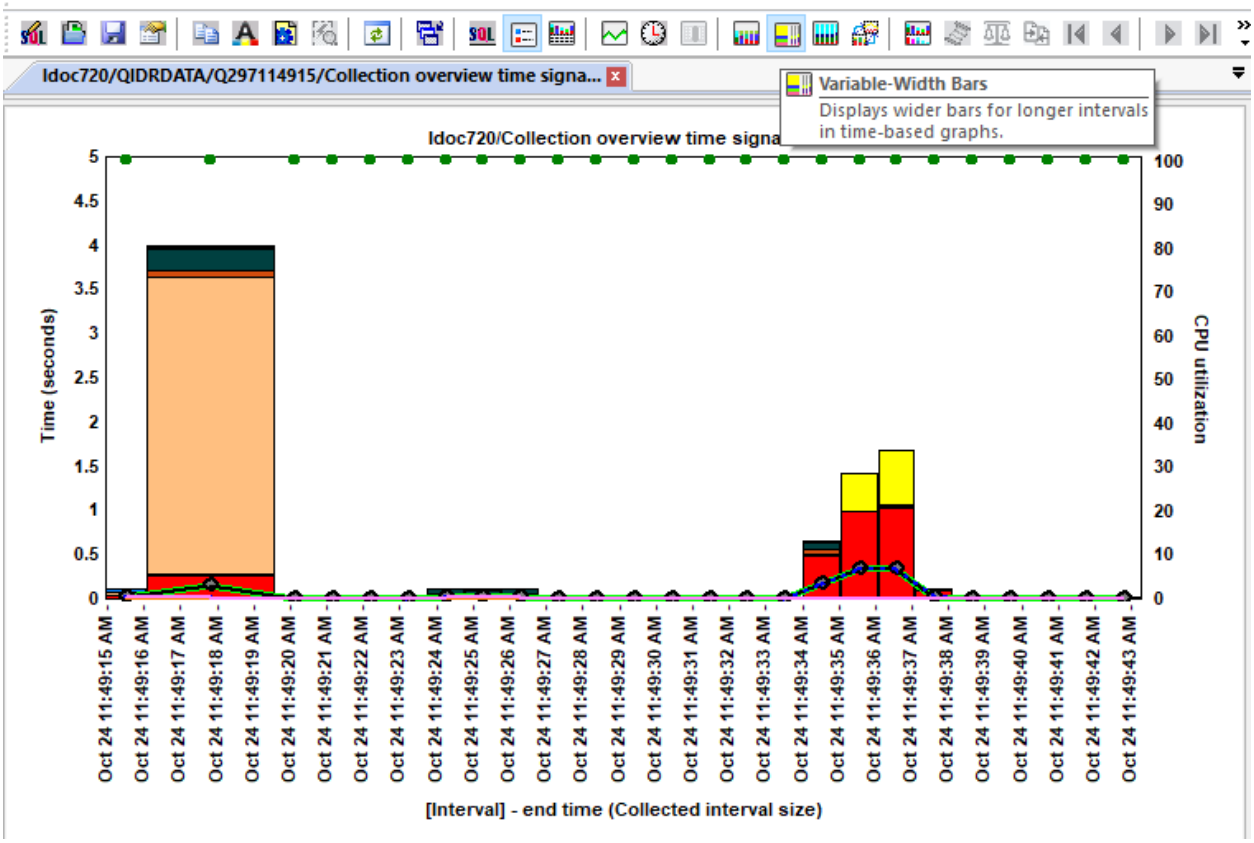

Collection overview graph with variable-width bar mode enabled

15 Toggle graph format

This button on the toolbar allows a graph to be quickly modified and presented in a different way.

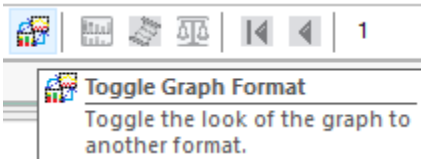

There are two types of graphs that have this option available:

1) Rankings graphs – toggles from horizontal to vertical bars.
2) Time-based overview graphs – toggles from stacked vertical bars to lines

Please note that using this option is like setting a preference. iDoctor will remember this the next time this graph or any other graph in the same folder is opened.

In this example from Job Watcher, a rankings graph is toggled from vertical to horizontal bars.

Note: This removes the Y2-axis and the avg response time metric but makes the label easier to read.

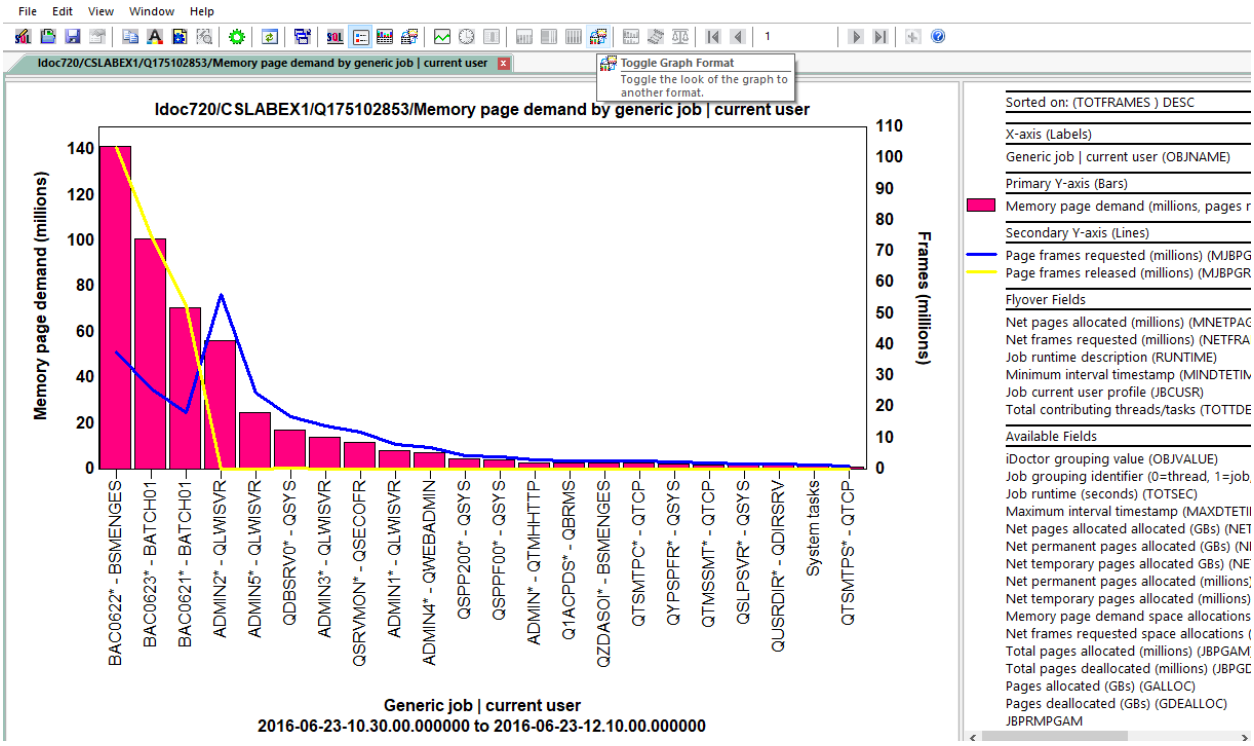

JW –> Page allocations –> Memory page demand by generic job | current user

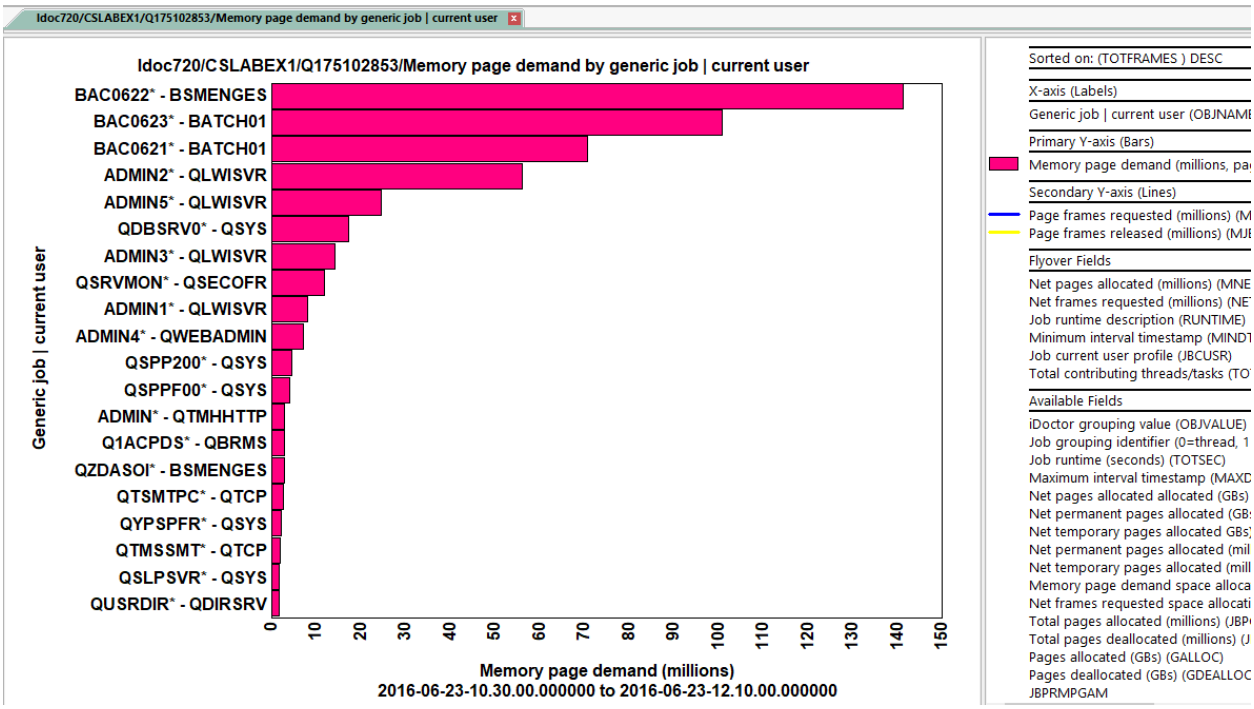

JW –> Page allocations –> Memory page demand by generic job | current user (horizontal bars)

16 Side-by-side comparison mode

This option is enabled by pressing this button  when 2 or more graphs or reports exist in a Data Viewer. If only 2 graphs exist in the Data Viewer, then the view will be split so 1 graph is on the top and 1 is on the bottom so they can be compared.

If more than 2 graphs exist in the Data Viewer when pressing this button, then you will be prompted for the other graph you wish to compare the current one with.

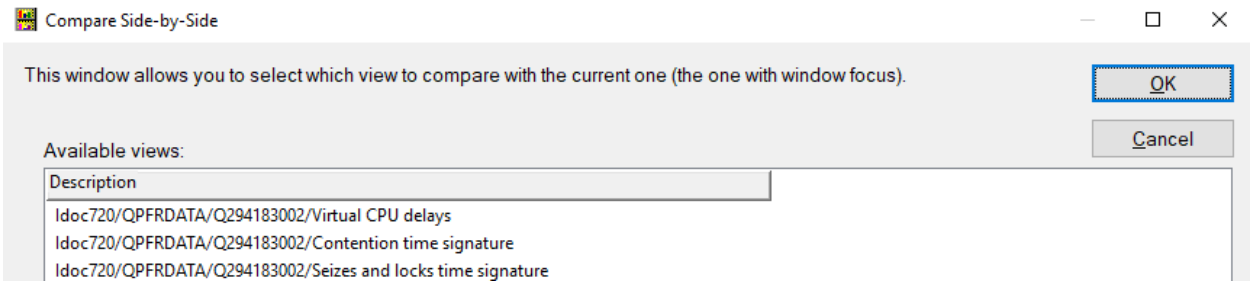

Compare Side-by-Side window

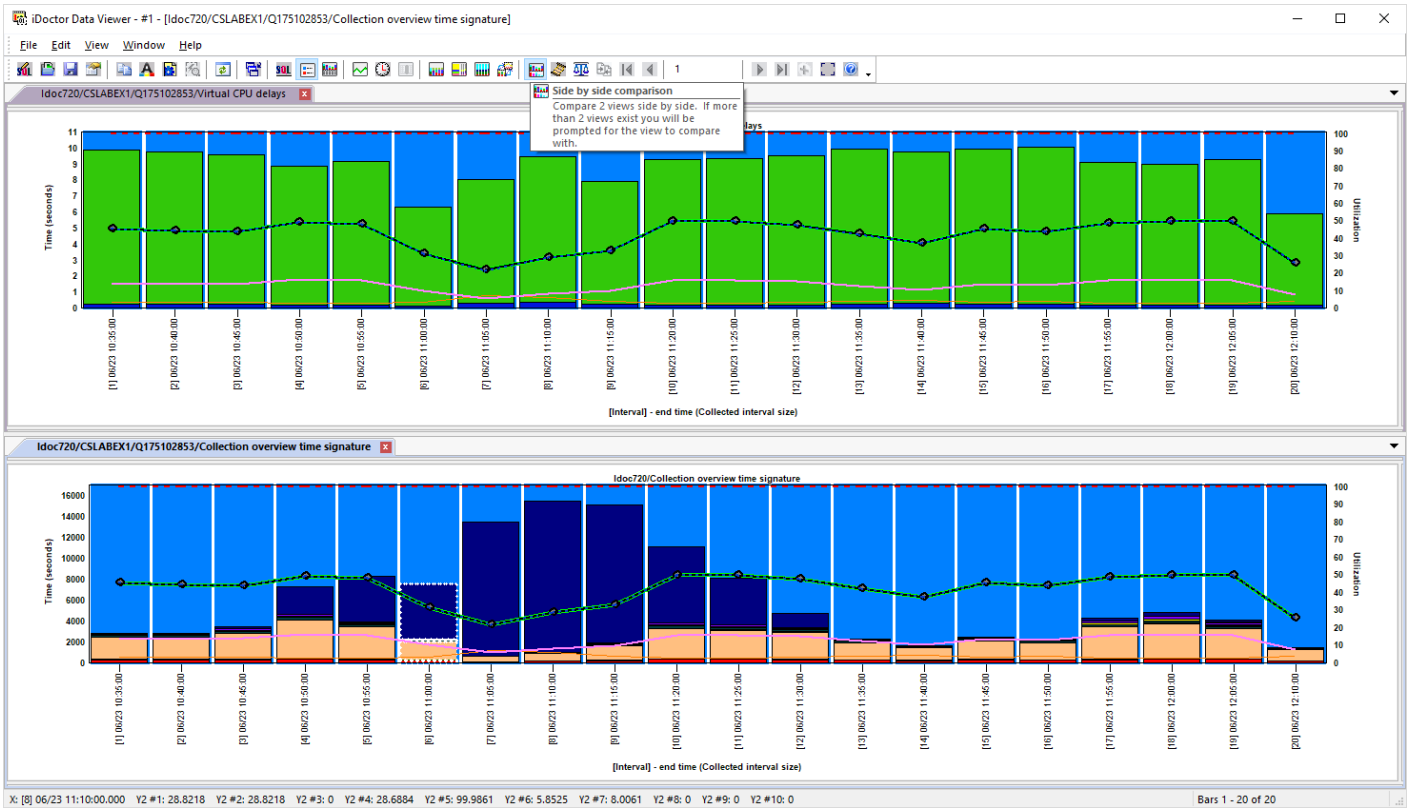

Side-by-side comparison mode

After enabling this graph mode additional toolbar buttons will be enabled on the Data Viewer toolbar providing additional options for comparing the graphs:

Button	Description
	Use this option to synchronize the scrolling of data while in Side-by-Side Comparison Mode. If pressed, then both graphs will scroll together when one of them is scrolled.
	Use this option to synchronize the Primary Y-axis scaling while in Side-by-Side Comparison Mode. If pressed, then both graphs will use the same min/max values on the primary Y-axis.

17 Spool File Views

The Data Viewer can be used to display the contents of spool files on the server. Whenever a job log for a collection is viewed that has already ended the job log is displayed in this view. This is also used when viewing spool files on an output queue in the IBM i Explorer component.

The spool file viewer will read in the entire contents of the spool file into the viewer. Although this will cause delays when reading large files this allows the user to more quickly perform a text search using the Find feature on the toolbar after the data is loaded into the client.

Tip: For large files it is best to download them as PDFs and view using Acrobat or your web browser using those options in the IBM i Explorer component.

An example of a Spool File View is shown below:

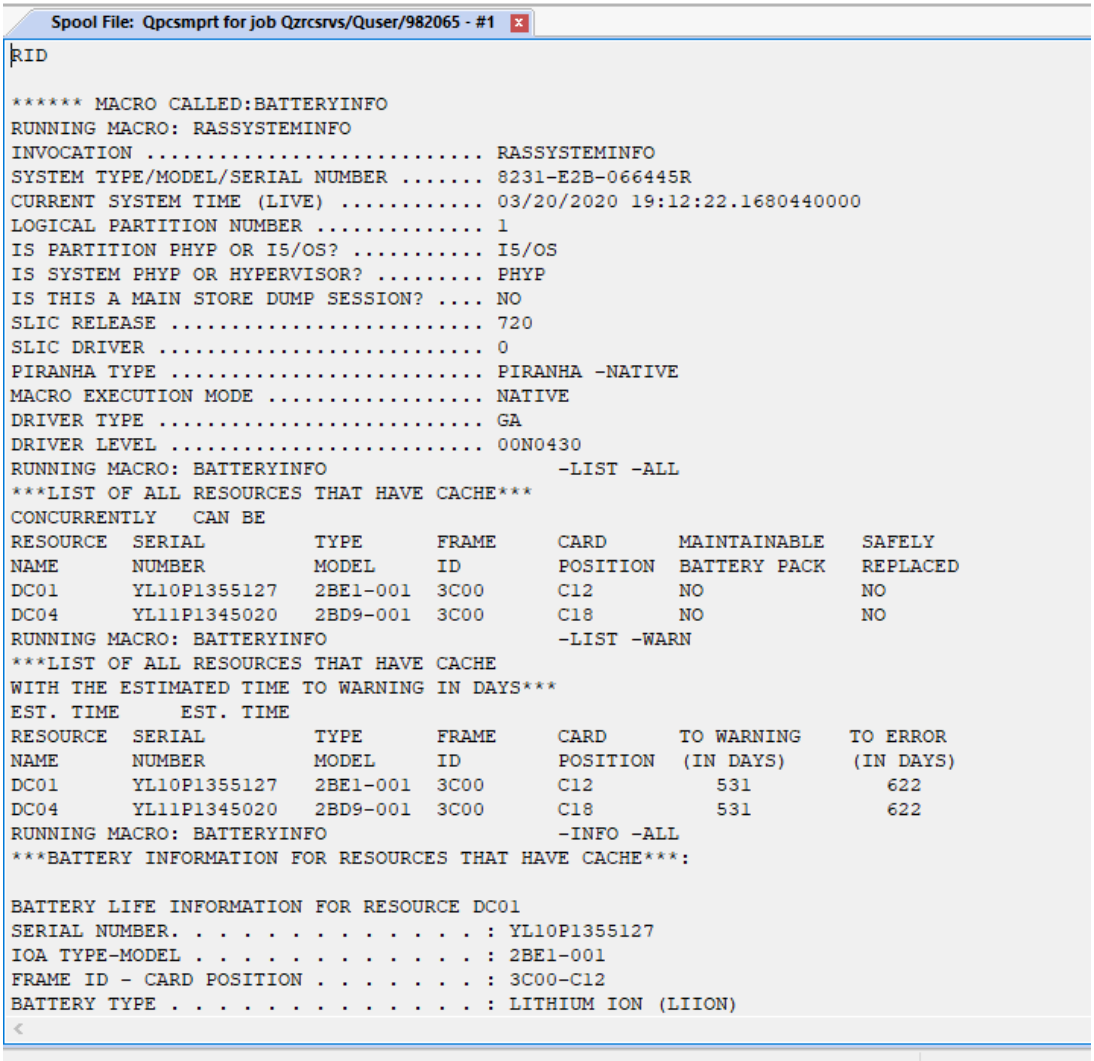

Spool File View from a completed job

18 Change SQL Parameters

This interface allows the user to modify the iDoctor-defined parameters included in an SQL statement. In some cases, the interface will appear when first opening a report if a parameter value is unknown and needs to be filled in by the user. Right-click a graph and use the Change SQL Parameters… menu to open this window.

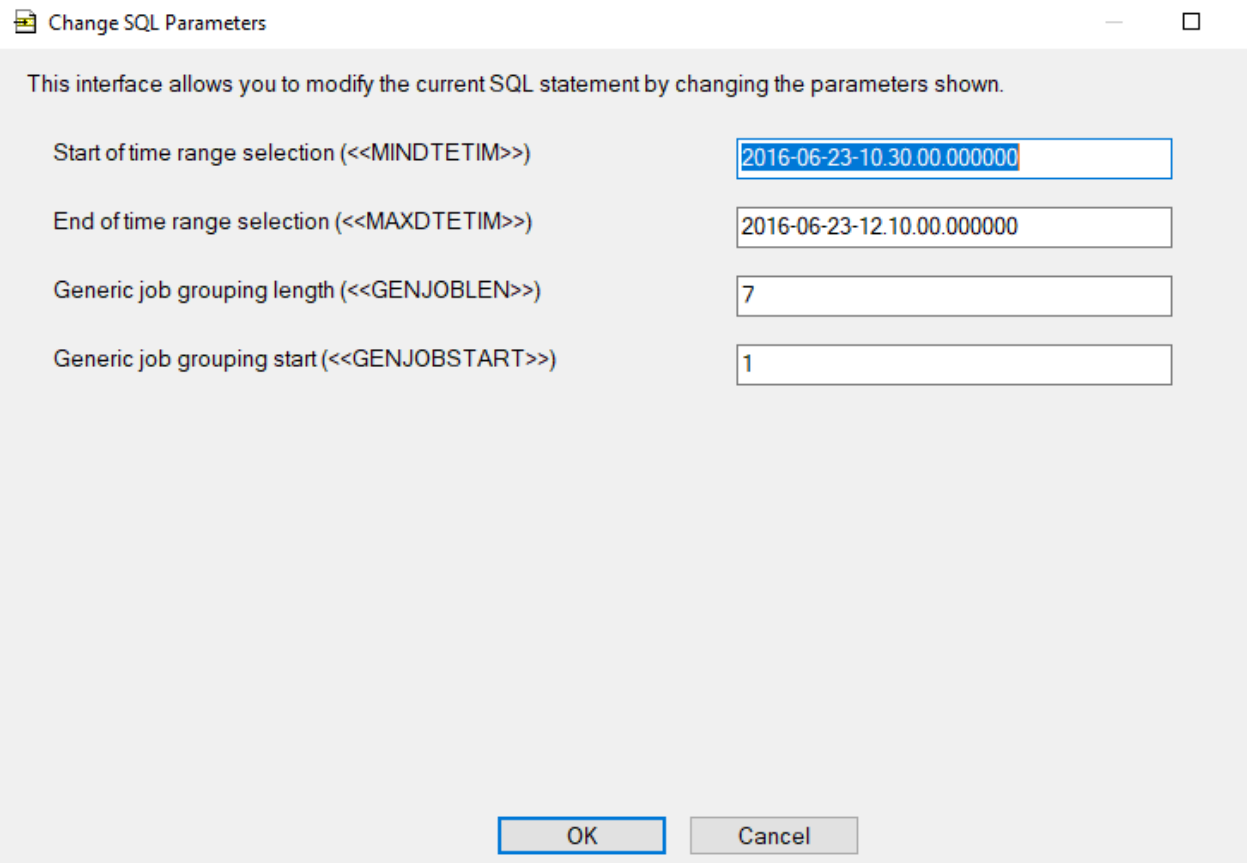

Change SQL Parameters Window

Each parameter (that users can modify) in the SQL statement will be listed on this screen. Additional parameters exist but are not changeable therefore not discussed here.

18.1 Change generic job grouping length

This shows how a user could change the generic job grouping length from 7 to 5 in CSI or Job Watcher.

1) Use Change SQL Parameters… menu

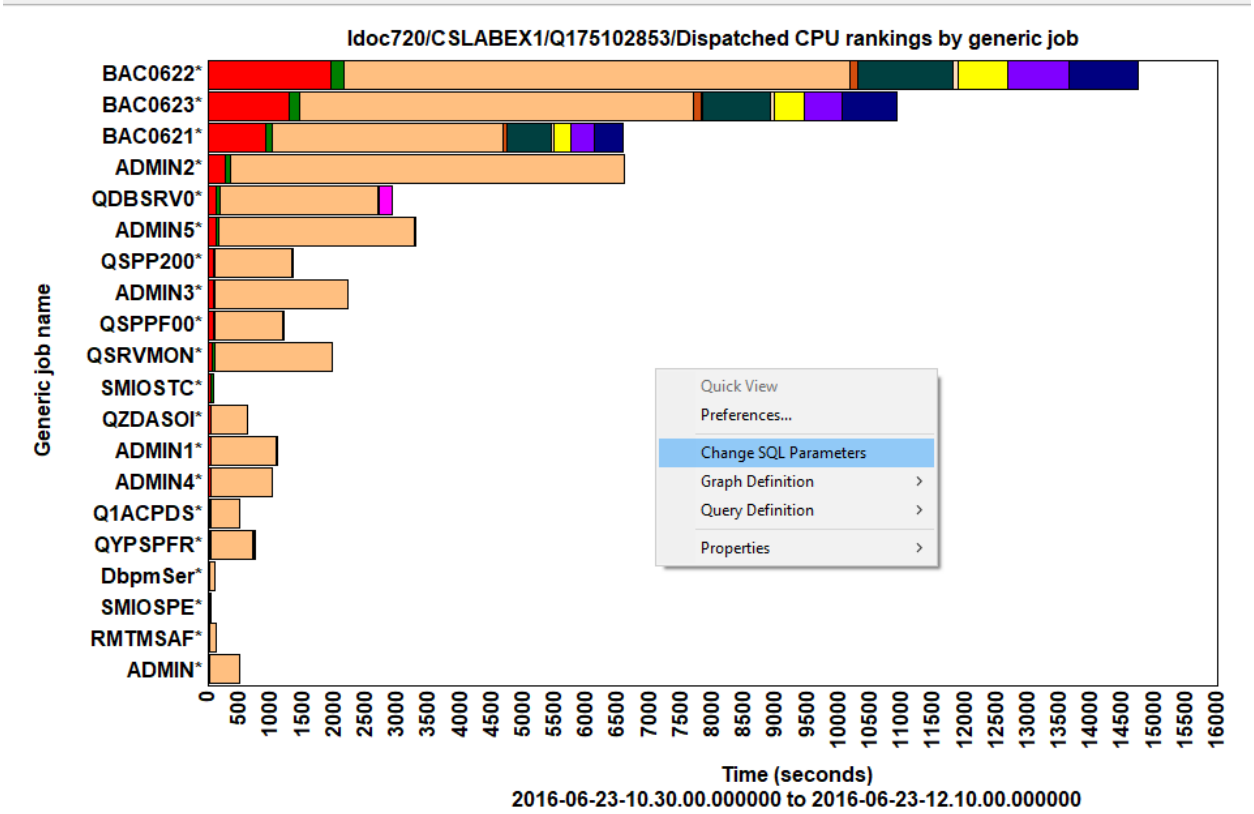

Generic job rankings with a name grouping length of 7

2) Fill in the value 5 from the generic job grouping length and hit OK.

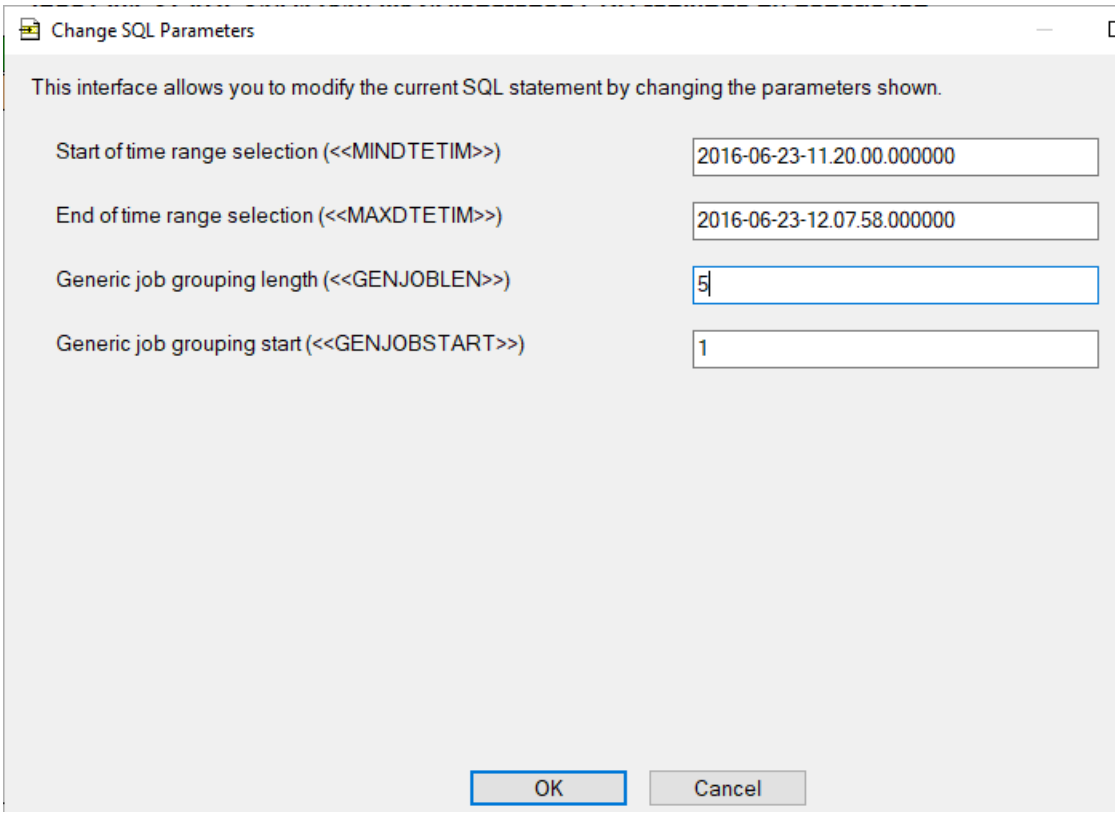

Change SQL Parameters – Generic job grouping length to 5

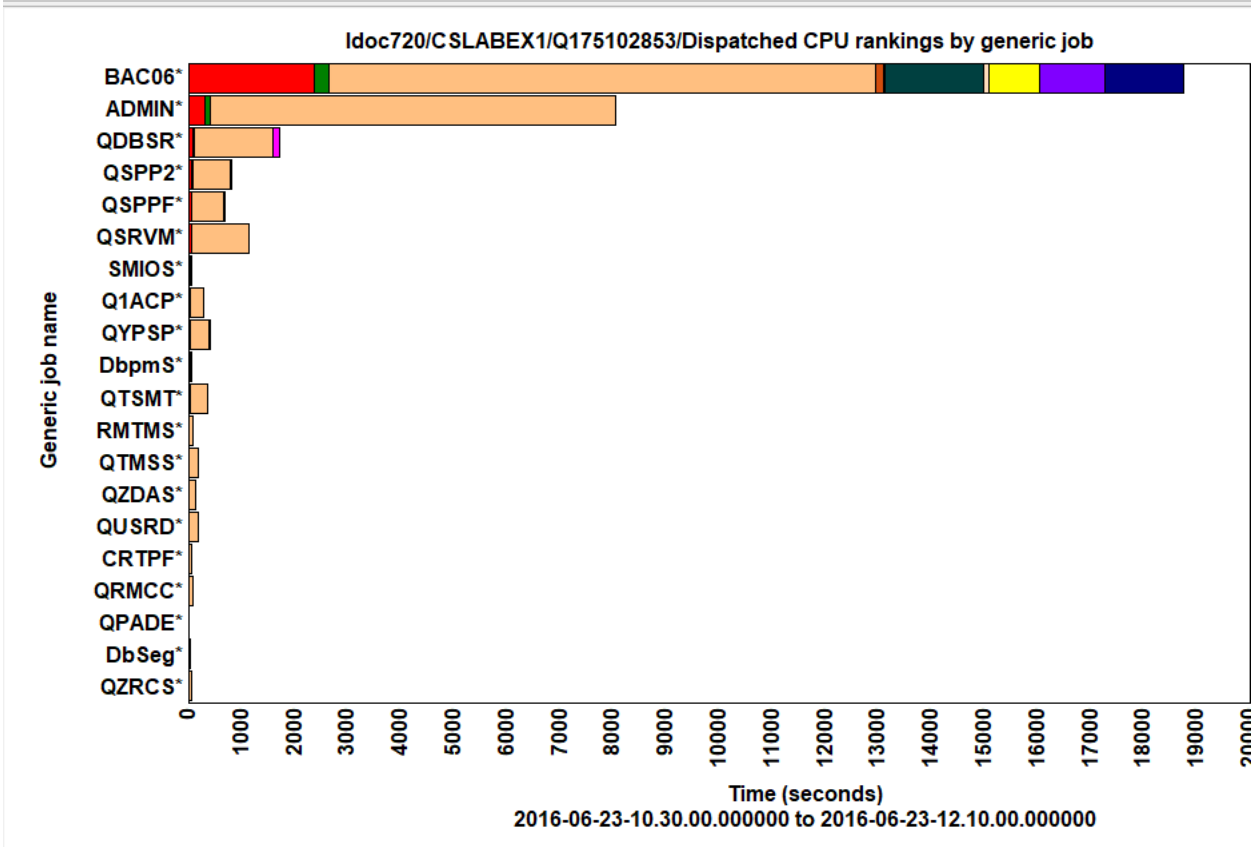

Generic job rankings with a name grouping length of 5

19 Edit Column

The Edit Column window is used to modify the attributes for a field/column in an iDoctor table or graph.

From tables this option is accessed by right-clicking on a column heading and using the Edit… menu. For graphs this option is accessed by right-clicking on a column in the graph legend and using the Edit… menu. Once the edit column window is opened from the graph legend you can click on other columns in the legend, and it will load their settings.

Note: Changing a column's attributes will make those changes permanent and visible throughout iDoctor. The column's changes are saved in the User-Defined Reports Database. You can remove those changes later by going into database's Column settings overrides folder and deleting the column.

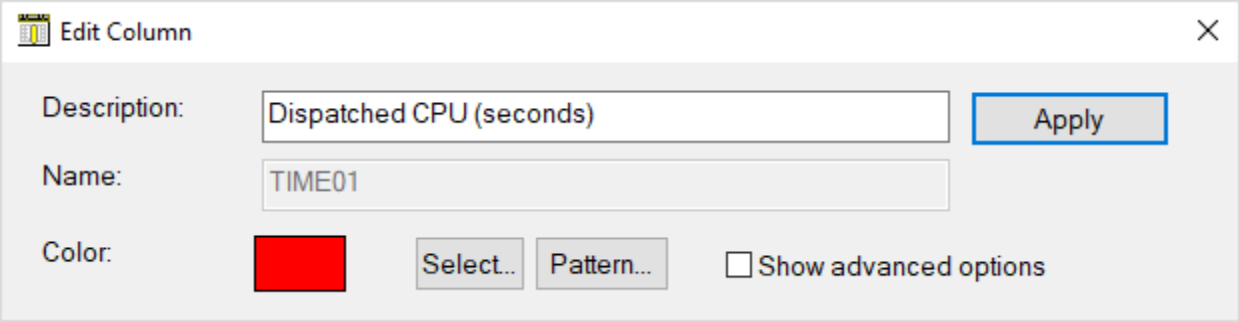

Edit Column

Option	Description
Description	The description to give the field name.
Apply button	This will apply the changes to the graph / report you are working with. If this is an iDoctor-defined field then these changes are saved and reused until removed from the User-Defined Reports Database -> Column settings overrides folder. **Tip:** Columns that have been changed from the iDoctor's defaults will have a green font. 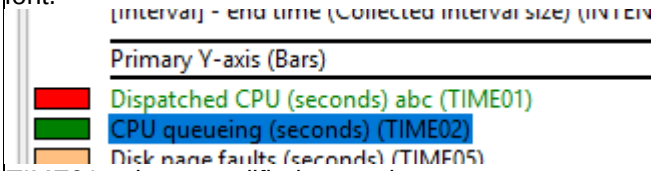 *TIME01 column modified example*
Name	This is the name of the field or column in the graph/report.
Color	The color to give the field when on a bar graph or line. Applies only when the field is shown on a graph.
Select button	Allows the user to modify the color used by showing the Color Window.
Pattern	Allows the modification of the hatching/pattern used on graphs for this field, but only if the **Preferences -> Display -> Patterns** checkbox is checked.
Show advanced options	Displays additional settings for graphing.

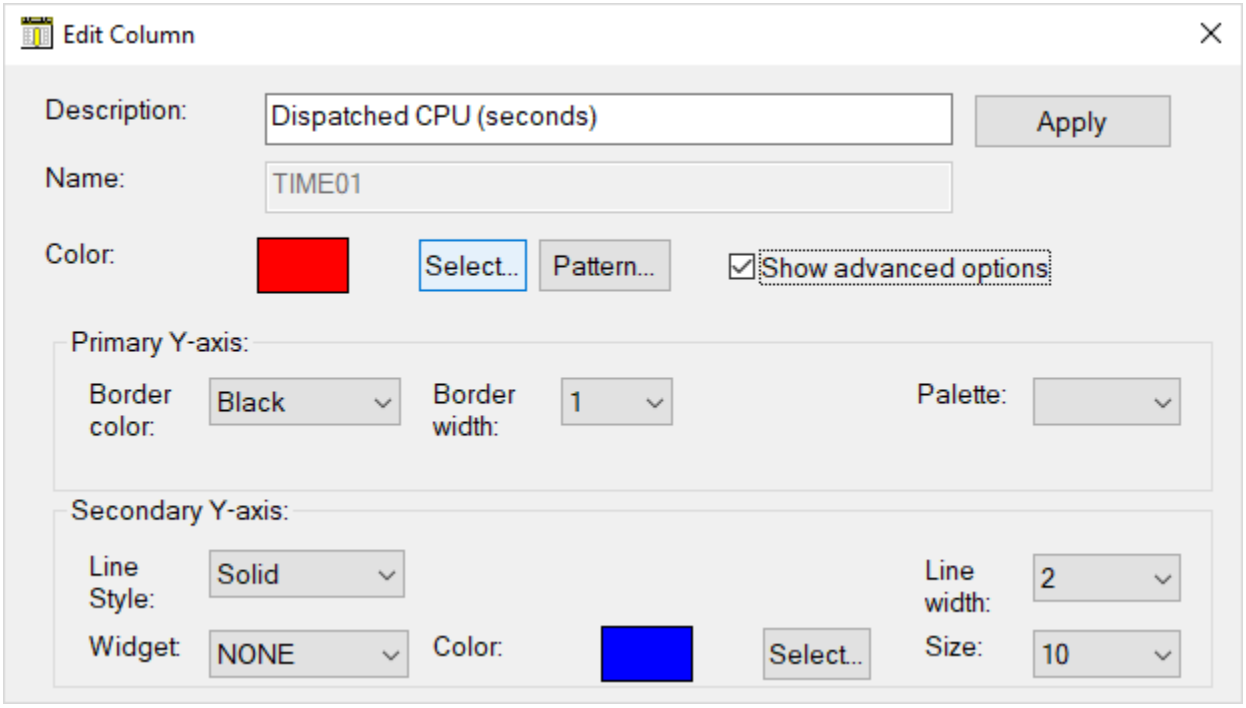

Edit Column -> Show advanced options checked

Option	Description
Primary Y-axis border color	If the field is used on a graph's primary Y-axis and the graph is a bar graph, then this changes the color for the border. The default border width is 1 pixel. The possible values are: 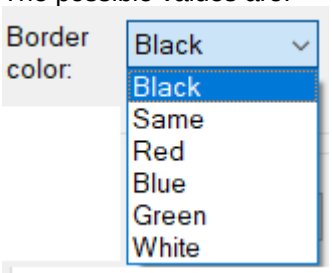 A value of "Same" means the color will match the fill color for the field (which in effect means no border will be visible. **Tip:** If the number of bars per page on the graph exceeds 250 then the border is automatically removed from the graph to avoid the graph appearance changing too much and looking more and more black with many bars. 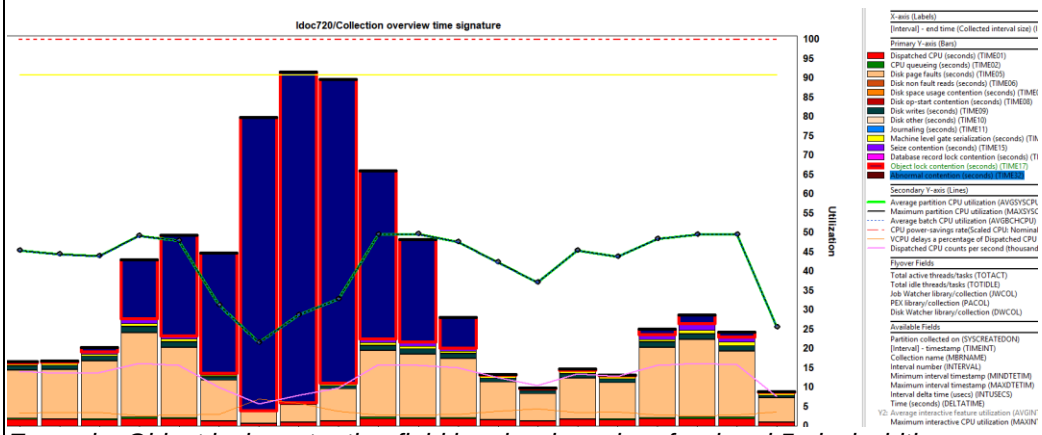 *Example: Object lock contention field has border color of red and 5 pixel width*
Primary Y-axis border width	If the field is used on a graph's primary Y-axis and the graph is a bar graph, then this changes the width for the border. The default border width is 1 pixel.

Primary Y-axis Palette	For flattened style graphs only, the colors used are dynamic because there are a variable number of values shown on the graph. This value lets you define a palette (a series of many colors predefined and shipped with iDoctor) for the field. **Note:** When drawing flattened style graphs, each color will be used from the Palette one at a time and when all colors are used the colors will wrap and be used again. The possible Palettes are: 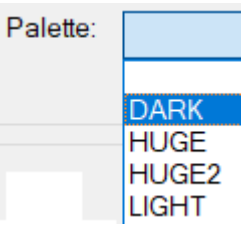 Using the 1st option (blank) will cause a random color to be assigned each time. This is the default setting. DARK is a series of darker colors and LIGHT are lighter colors. The HUGE palettes cover all colors are predefined, so the same colors will be used each time the same graph is reopened for the same data. Otherwise if no Palette is defined the colors will be randomly defined every time and they may not look good depending on chance.
Secondary Y-axis line style	If the field is used on a graph's secondary Y-axis, this changes the style for the line. The possible values are: 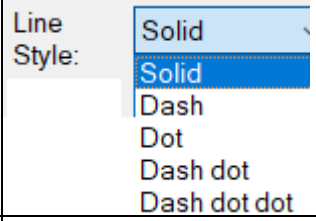
Secondary Y-axis line width	If the field is used on a graph's secondary Y-axis, this changes the width for the line. The possible values are 1-5 pixels.
Secondary Y-axis Widget (name)	If the field is used on a graph's secondary Y-axis, this allows each point of the line to contain an optional shape called a widget. These shapes have different names and the choices are: 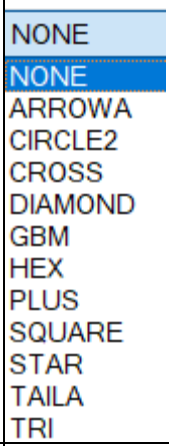
Secondary Y-axis Widget color	If the field is used on a graph's secondary Y-axis, this allows each point of the line to contain an optional shape called a widget. This is the color to draw the widget with. Press the select button to change the color.

Secondary Y-axis Widget size	If the field is used on a graph's secondary Y-axis, this allows each point of the line to contain an optional shape called a widget. This is the size to draw the widget. The bigger the number, the bigger the widget will appear. 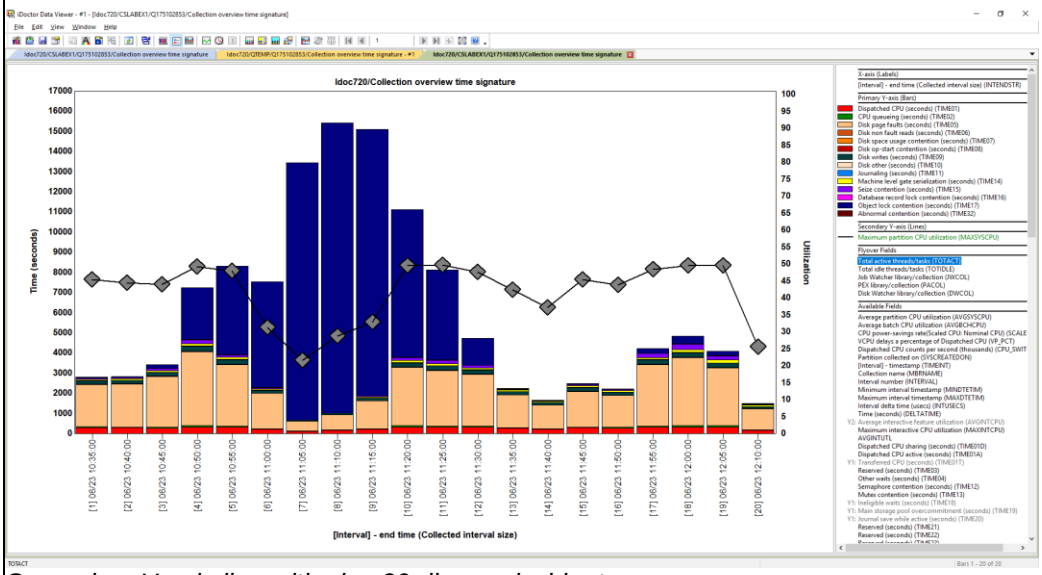 *Secondary Y-axis line with size 30 diamond widget*

20 Color Window

When modifying the color of a field in iDoctor, this window is used. Colors for iDoctor fields are defined using RGB values (Red/Green/Blue) where each color is a number between 0 and 255.

Examples:

Red = 255,0,0

Green = 0,128,0 or 0,255,0

Blue = 0,0,255

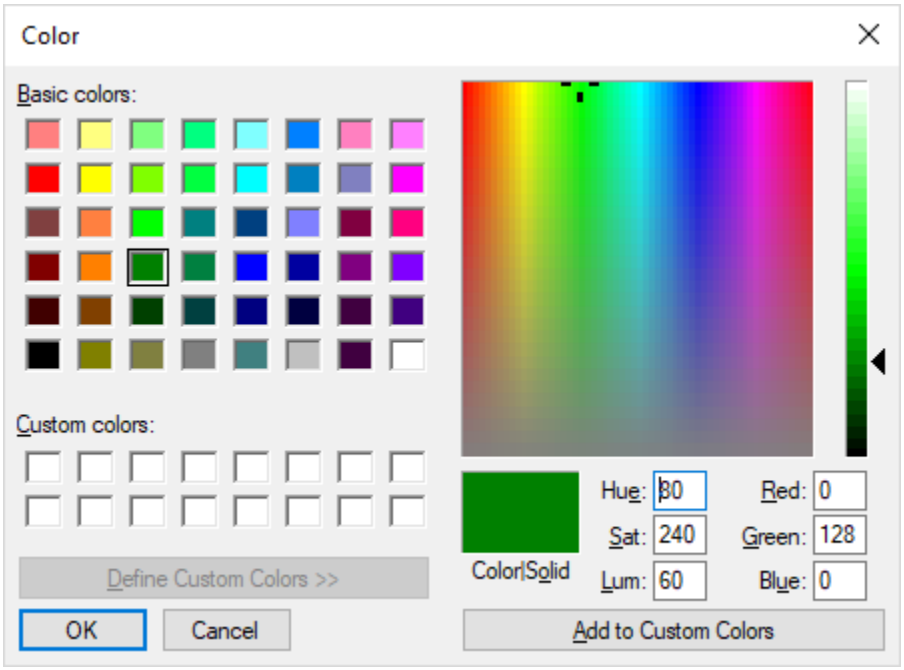

Color Window

The RGB values are shown in the boxes in the bottom right-corner of this interface. You can create your own colors here and save / reuse here if desired.

21 Fill Pattern Selection

This window allows a user to select the type of pattern they would like for the current field you are working with. This option is available from the Edit Column window or from the Graph Definition Primary Y-axis page.

Note: Making changes here has no effect on the graphs unless the **Preferences -> Display -> Patterns** is checked -or- the Patterns checkbox on the **Graph Definition -> Primary Y-axis** page is checked.

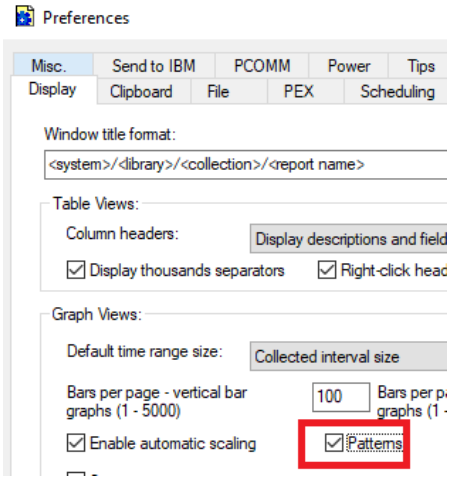

Preferences -> Display -> Patterns

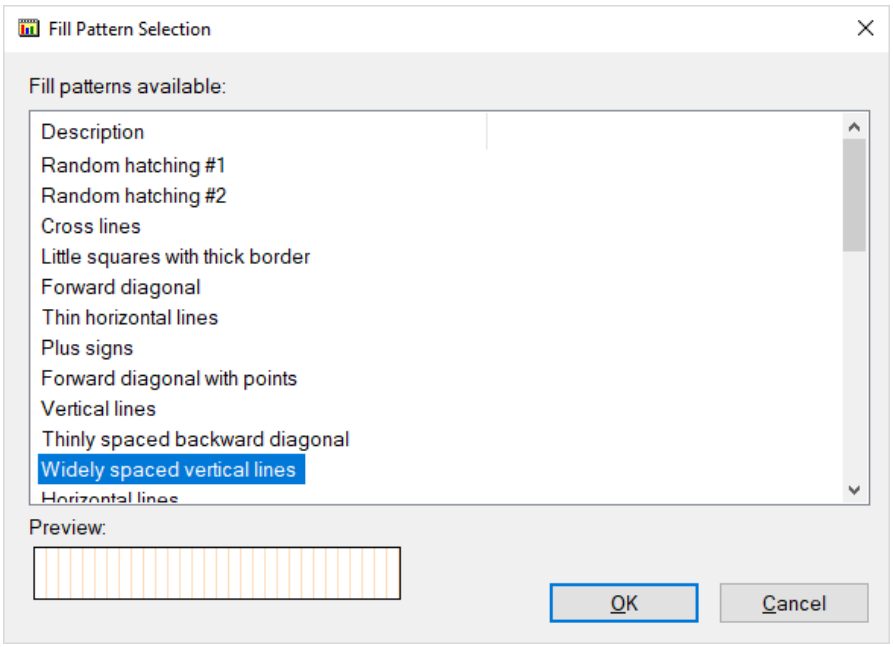

Fill Pattern Selection