IBM iDoctor for IBM i
Disk Watcher

IBM iDoctor for IBM i Development Team

11 February 2022

Abstract

Provides in-depth coverage of all major GUI functions for all components at 7.1, 7.2, 7.3 and 7.4. This document covers the Disk Watcher component.

Changes

8 Feb 2022 – updated for 2022

Table of Contents

1 Overview

Disk Watcher provides the user with the ability to collect either a statistical summary of disk performance data or a trace of all disk I/O events that occur on a system. The trace mode is recommended as it provides more options for analyzing the data and determining potential disk problems.

The Disk Watcher GUI provides many graphs with drill downs for each mode of collection (statistical or trace). Using Disk Watcher, the user can take a trace and summarize the trace data into an interval size desired for the purpose of easily graphing the statistics at either or broad or detailed levels.

Data is collected in Disk Watcher using commands that are included with IBM i which are:

- ADDDWDFN – Adds a Disk Watcher definition to the system
- STRDW – Starts a Disk Watcher collection
- ENDDW – Ends an active Disk Watcher collection

2 Starting Disk Watcher

Disk Watcher is a component of the iDoctor suite of tools. After launching iDoctor, the Disk Watcher component is started from the IBM i Connections List View by double-clicking on the desired system.

A list of available components will appear on the next window. Double-click on the Disk Watcher component or select Disk Watcher and click the Launch button to continue

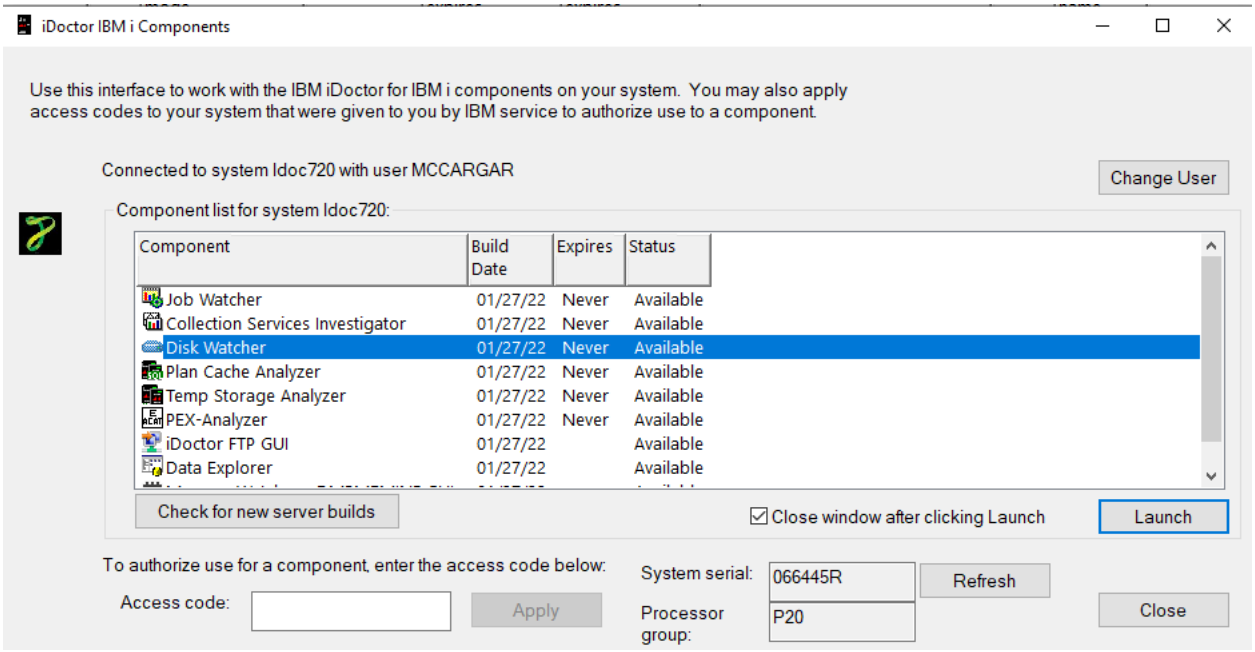

iDoctor Components Window

3 Disk Watcher Component View

The Disk Watcher folder contains a list of folders, each providing different features available. Collections can be displayed in various ways, either under the Libraries folder on a per library basis, or under the Monitors folders to show Disk Watcher collections under a monitor.

This also provides options for working with the Disk Watcher Definitions that exist on the system. These are used for defining the aspects for what data is collected. Several IBM-supplied definitions exist, or the user can make their own.

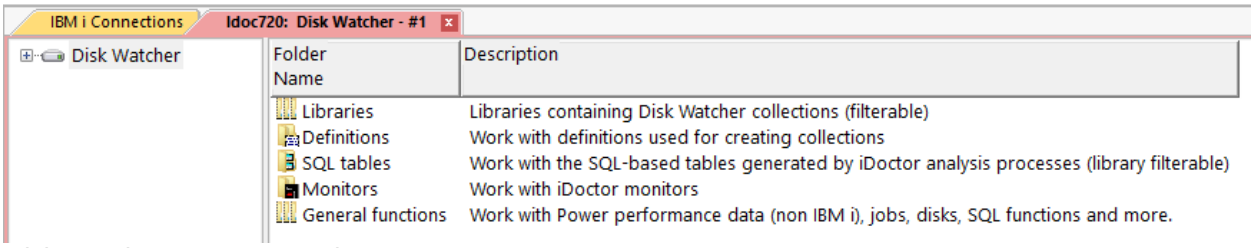

Disk Watcher Component View

3.1 Root Folder Menu Options

The following Disk Watcher specific menu options are available by right clicking on the 'Disk Watcher' folder:

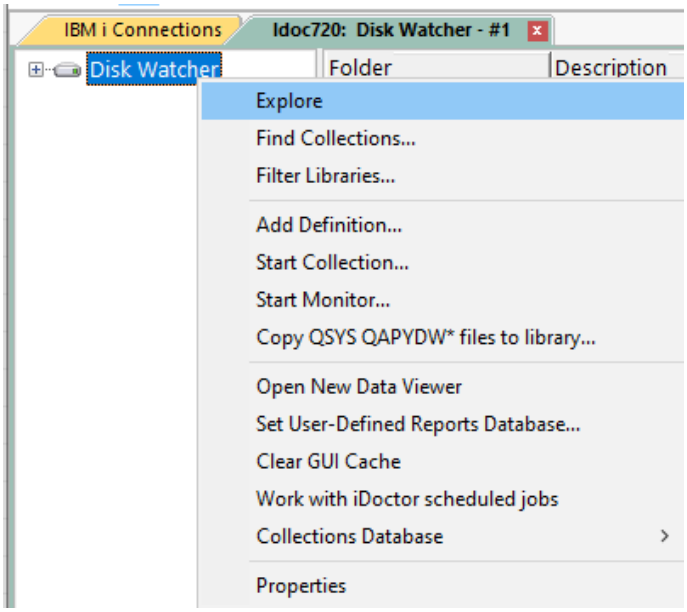

Disk Watcher Popup-Menu

Menu	Description
Find Collections…	This option displays the Find Collections interface which provides the ability to look for collections matching user-defined characteristics. Example SQL statements are provided. The results of these queries are available under the General Functions -> Find collections results folder.
Filter Libraries…	This option allows you to filter the libraries shown in the Libraries and SQL Tables folders by a generic library name or library owner. This is useful for speeding up the display of the list if the system contains many libraries containing collections (and/or SQL tables).
Add Definition	This option displays the Add Disk Watcher Definition Wizard. The definition defines characteristics about the collection such as which data options to collect. **Note:** A definition is required when starting a collection.
Start Collection	This menu will open the Start Disk Watcher Wizard where the user can define and run a collection.
Start Monitor	This menu will open the Start Monitor Wizard for iDoctor where the user can start a Job Watcher, PEX Analyzer or Disk Watcher monitor. Monitors are designed to provide 24x7 collection of performance data.
Copy QSYS QAPYDW* files to library…	This option can be used to copy all QAPYDW* files from QSYS to the desired library. This is an optional step that prepares a library for collecting performance data. It also allows the library to appear under the Libraries folder in Disk Watcher. 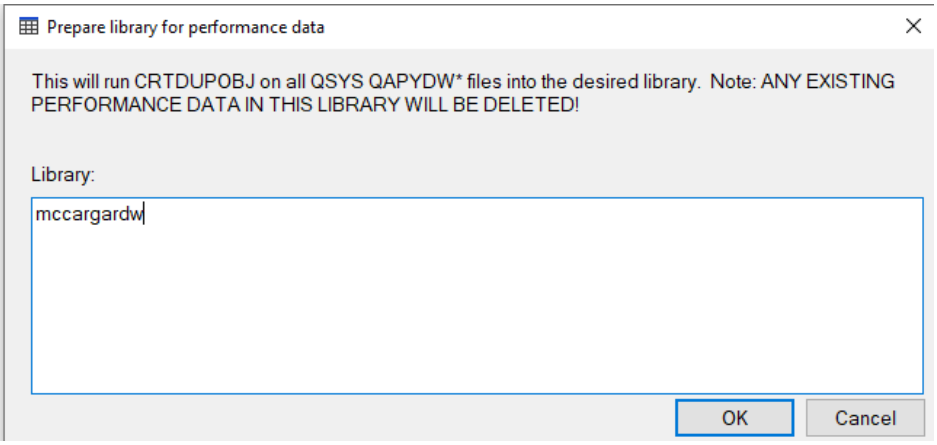 *Prepare library for performance data window* **Note:** The reason this exists, is in some situations where previous release DB files are used to collect performance data on the current release, then various problems will occur.
Open New Data Viewer	Opens a new Data Viewer window. This window is used to display tables and graphs on the system. You can open iDoctor-defined reports into this window or you can also open any database file or SQL table and display the results in graph or table form.
Set User-Defined Reports Database	This option allows the user to view/modify the currently used user-defined reports database. The database can either be an MS Access file or a library on an IBM i. The database stores the information needed to build the user-defined tables and graphs shown in iDoctor.
Clear GUI cache	This option clears everything loaded in the GUI's cache (like menus, graph definitions, query definitions, stored procedure versions installed, etc)
Work with iDoctor scheduled jobs	This option is a shortcut to the General functions –> Work management -> Scheduled jobs folder. It shows all the iDoctor created scheduled jobs that exist on the current IBM i system.

Collections database	The iDoctor collection database identifies all collections on the system and can be used to facilitate the drill down from one component to another in some situations. It also is used to improve performance when browsing collections on the system.
Properties	Use this menu to display version information for the current component. The build level of the GUI is also displayed here.

4 Definitions

A Definitions folder is provided in Disk Watcher to allow the user to work with the Disk Watcher definitions that exist on the current system. An example of this interface is:

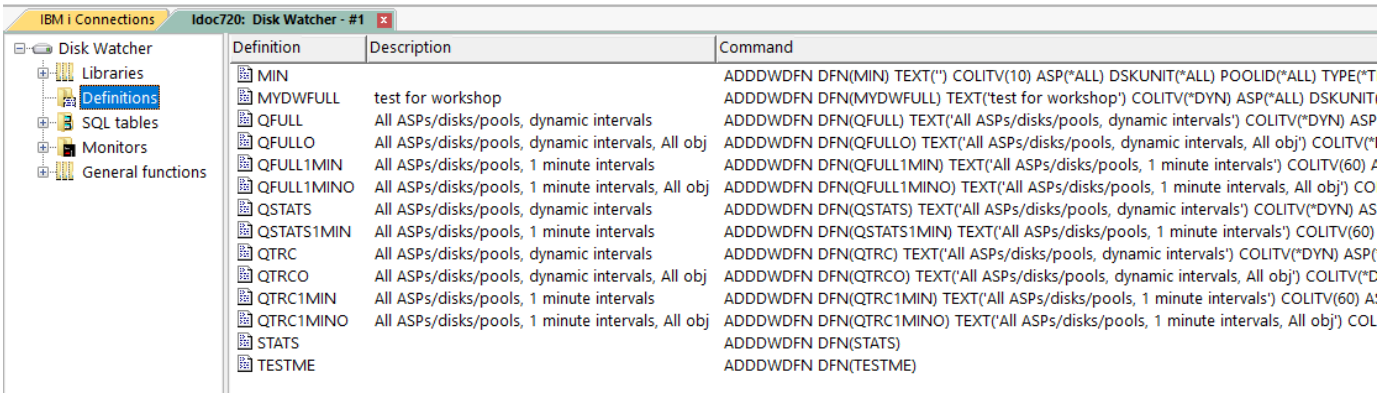

Definitions Folder

Note: Right-clicking the Definitions folder provides an option to rebuild the IBM-supplied definitions which is sometimes needed if none of the Q* definition names appear.

The following options are available when right clicking on one or more definitions in the list:

Field	Description
Change Definition	Opens the Add Disk Watcher Definition Wizard and loads the selected definition into it so it can be changed. **Note:** This is only applicable to definitions that aren't named Q*.
Add Definition	Opens the Add Disk Watcher Definition Wizard in order to create a new definition.
Start Collection	Opens the Start Disk Watcher Collection Wizard using the selected definition.
Start New Monitor	Opens the Start iDoctor Monitor Wizard using the selected definition.
Delete	Removes the selected definitions from the system.
Properties	Displays the properties for the selected Disk Watcher definition.

4.1 Add Disk Watcher Definition Wizard

Use this interface to create or change a user-defined Disk Watcher definition. This option is accessible by right clicking the Disk Watcher root folder and using the Add Definition… menu.

It can also be accessed using the Add Definition menu for a definition.

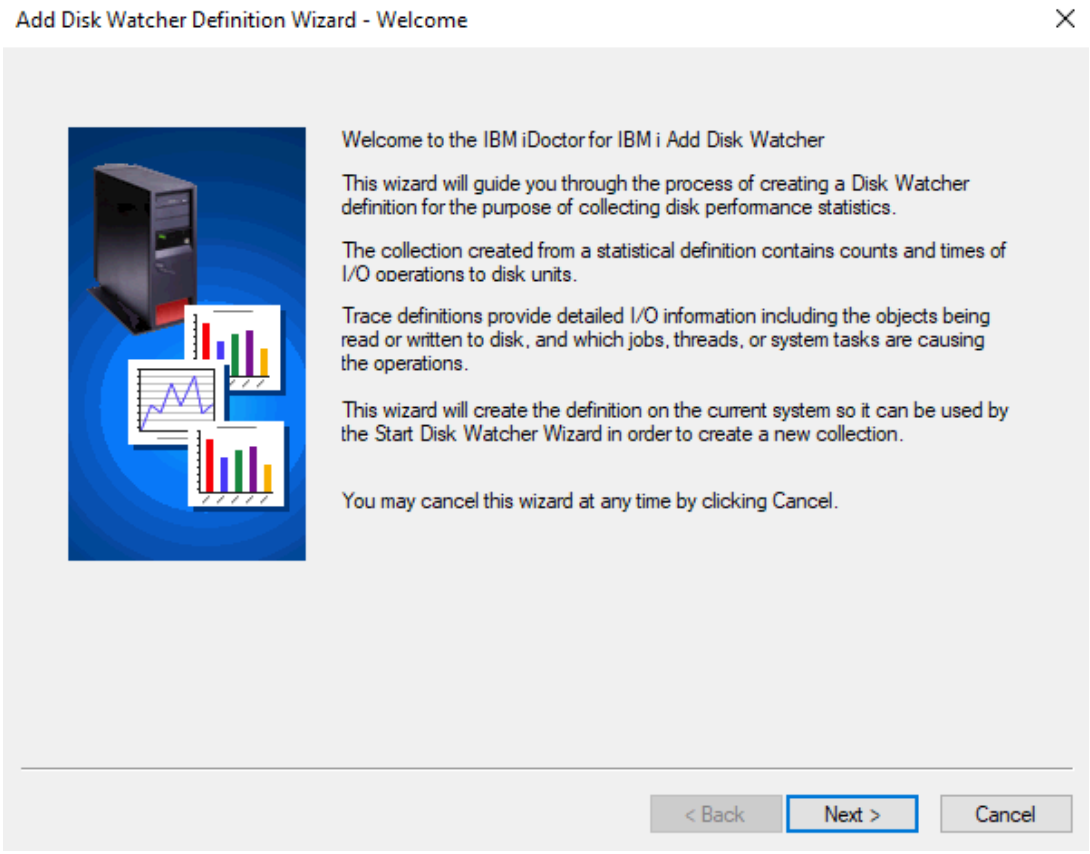

Add Disk Watcher Definition Wizard – Welcome

4.1.1 Basic Options

This interface allows the user to provide general information about the definition.

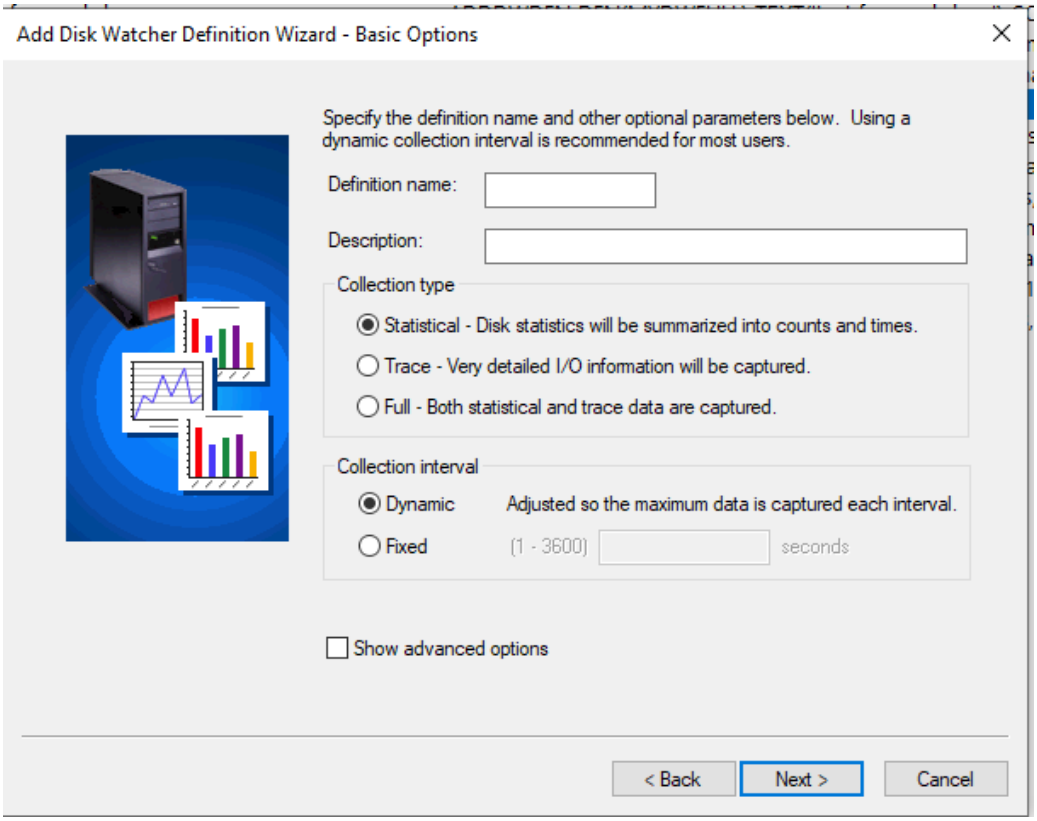

Add Disk Watcher Definition Wizard – Basic Options

The following table describes some of the less obvious options on this screen:

Option	Description
Definition name	The definition name will be added as a new member name in file QAPYDWDFN in QUSRSYS. This cannot start with Q when using this interface.
Collection type	Statistical, Trace or Full (Both)
Collection interval	Dynamic or Fixed. For easier graphing it is often wise to use Fixed at an interval duration like 1 second or 5 second intervals. Dynamic can also be used, but in some cases may require use of variable-width bar mode to better understand the data.

4.1.2 Advanced Options

These settings are rarely needed and only visible if the “Show advanced options” checkbox on the previous window is checked.

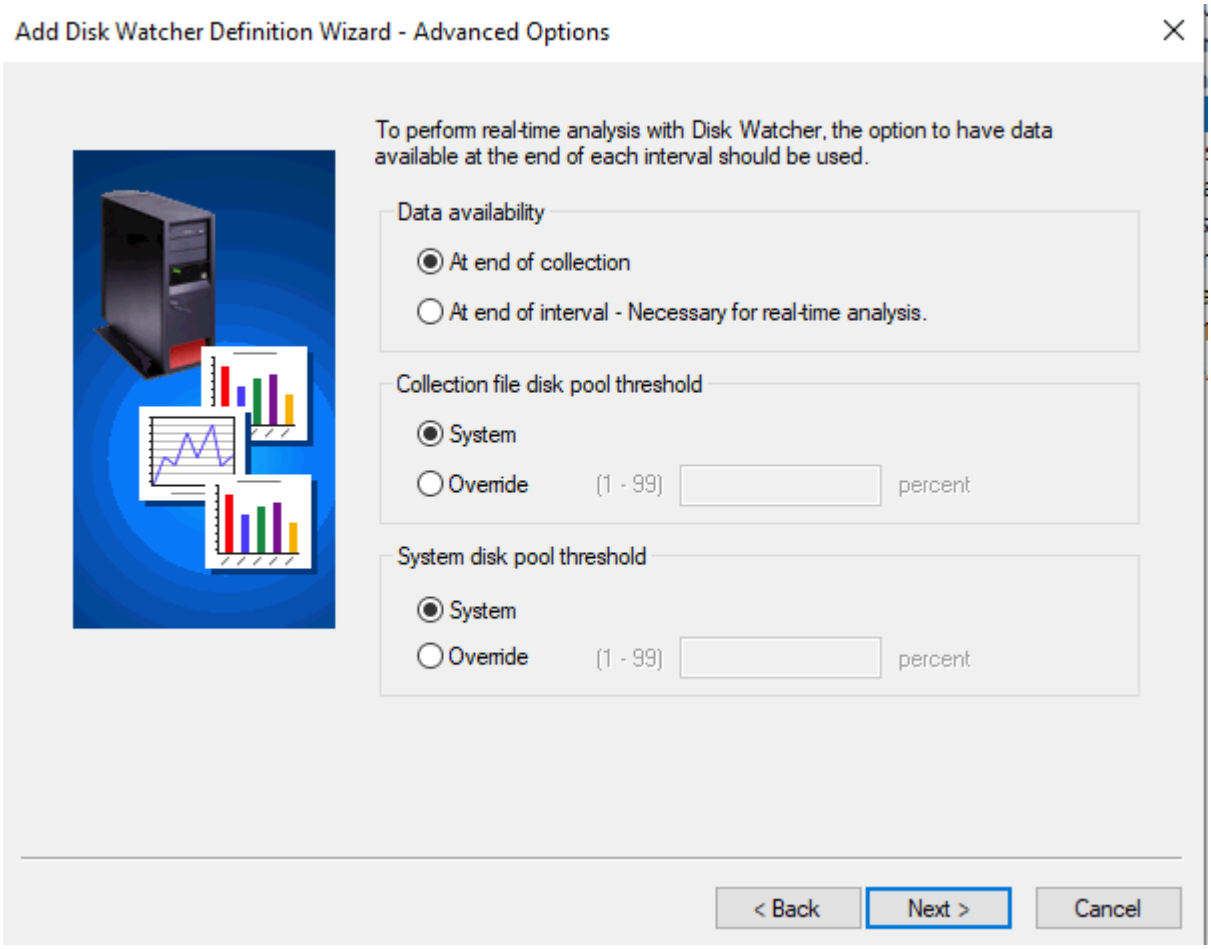

Add Disk Watcher Definition Wizard – Advanced Options

4.1.3 Disk Selection Options

This panel allows the definition to specify from which disk pools (ASPs), disk units or memory pools associated with disk I/Os should data be captured.

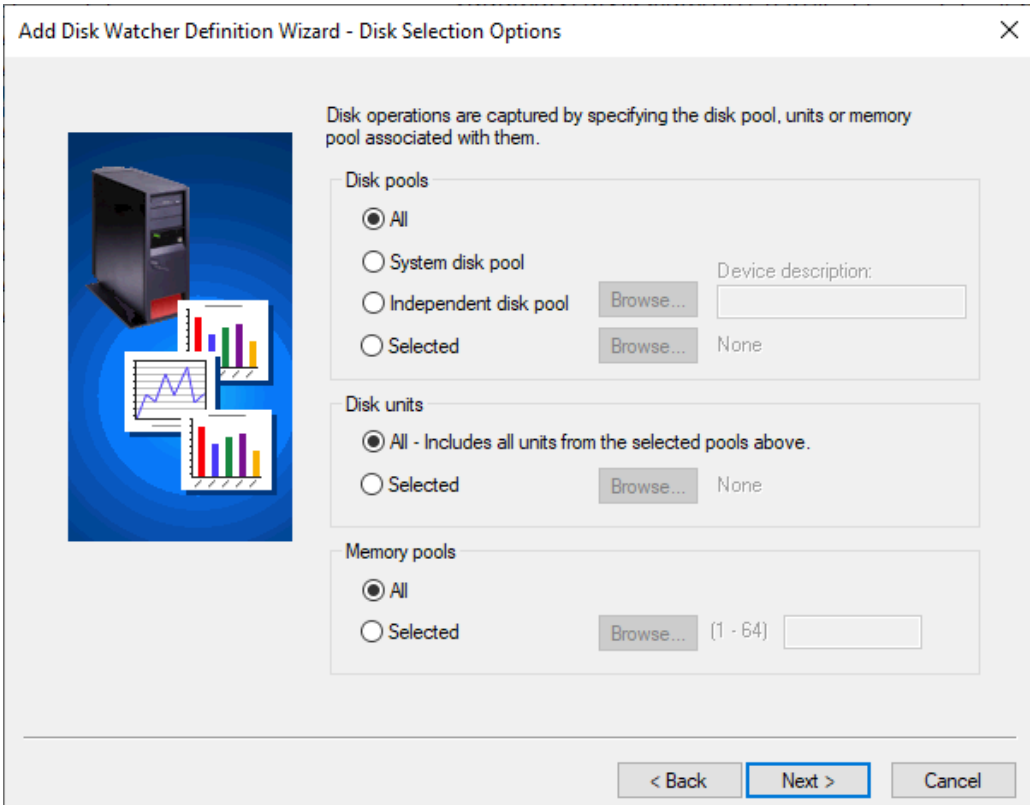

Add Disk Watcher Definition Wizard – Disk Selection Options

4.1.4 Summary

The final screen of this interface will list all selections made and provide the remote command string that will be used to create the definition after pressing the Finish button.

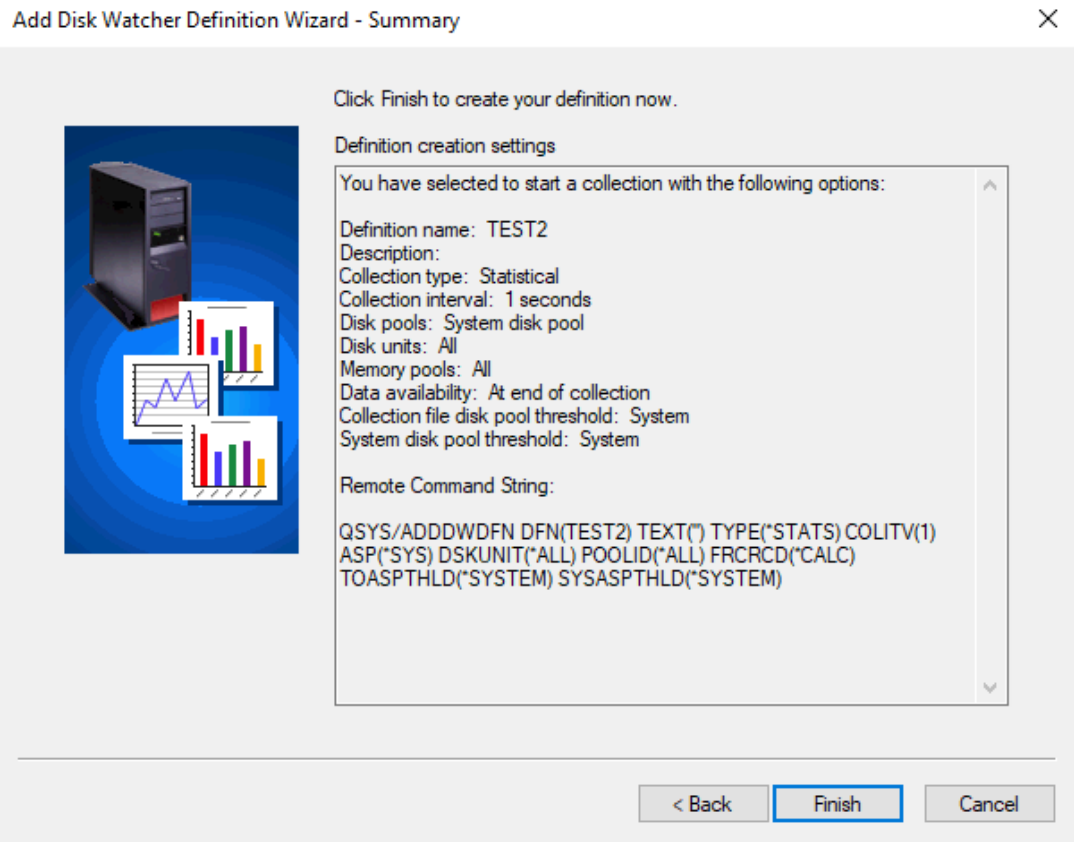

Add Disk Watcher Definition Wizard – Summary

4.2 Properties

Double-clicking on a definition or using the Properties menu from the Disk Watcher Definitions View displays all the parameters that were used when creating the definition. An example of this interface is:

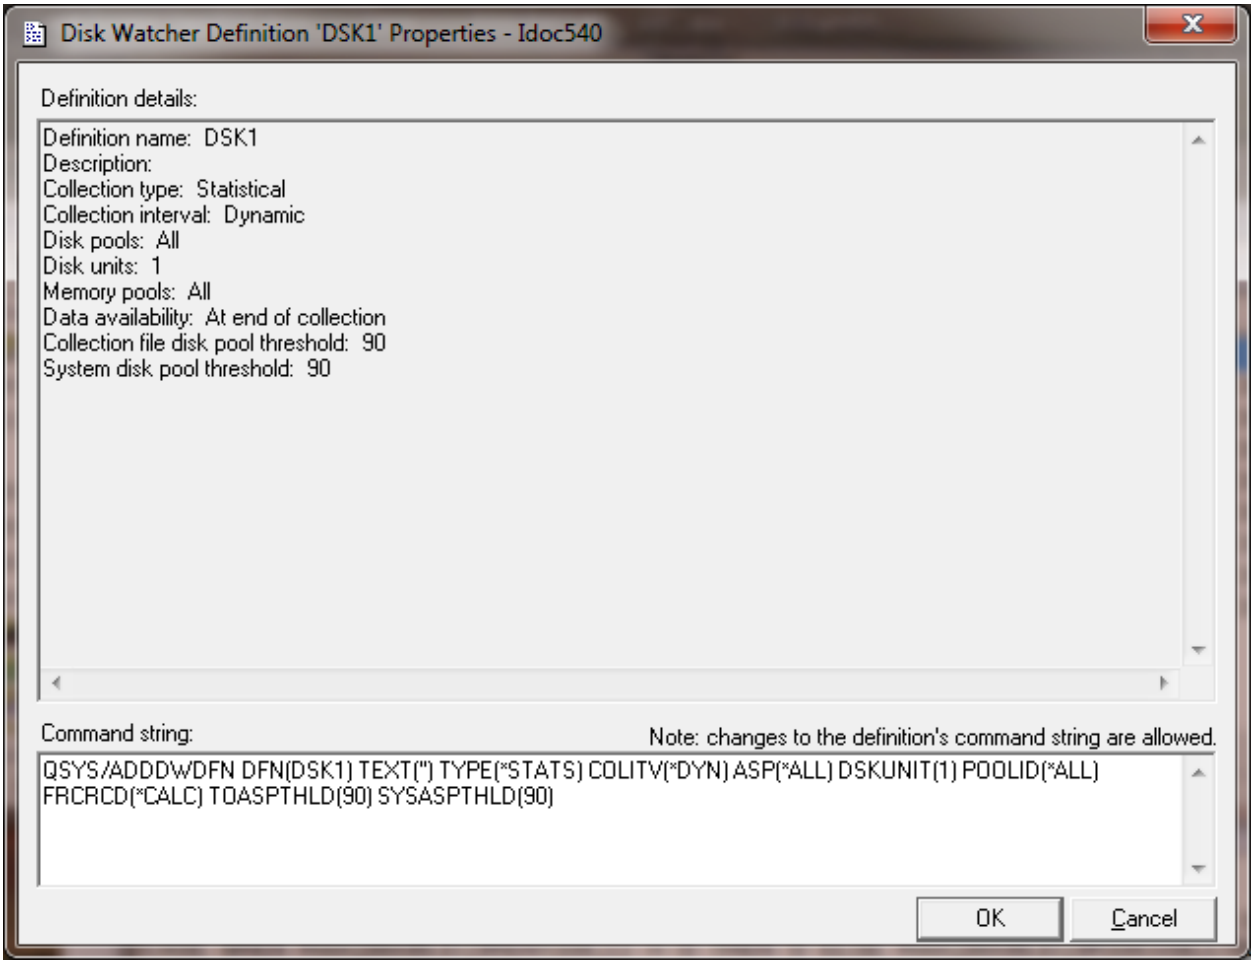

Disk Watcher Definition Properties

Advanced users can change the command string that defines the definition if desired. If changes have been made to the command string, pressing the OK button will remove the existing definition from the system and replace it using the command string specified.

Note: IBM-supplied definitions cannot be changed.

5 Start Disk Watcher Wizard

Disk Watcher provides the capability to collect detailed information about all disk IOs occurring on the system.

This section covers the creation of a collection using the Start Disk Watcher Wizard. The Wizard is accessible via the Start Collection menu on the Disk Watcher or library folder icons. This Wizard guides the user step by step through the process of creating a collection. Each page is covered in detailed within the next sections.

Tip: If Disk Watcher data already exists in the library it must match the currently installed OS release of IBM i or you will be unable to collect more data in that library. You cannot combine data of different releases in the same library.

5.1 Welcome

The Welcome page in the Start Disk Watcher Wizard introduces the user to the wizard and offers information about what it will do.

Tip: Starting a collection requires a definition. Use the Add Disk Watcher Definition Wizard first if you do not wish to use the IBM-supplied definitions.

5.2 Definition Selection

This screen allows the user to specify the type of definition to use (IBM or user-defined) and select the definition name.

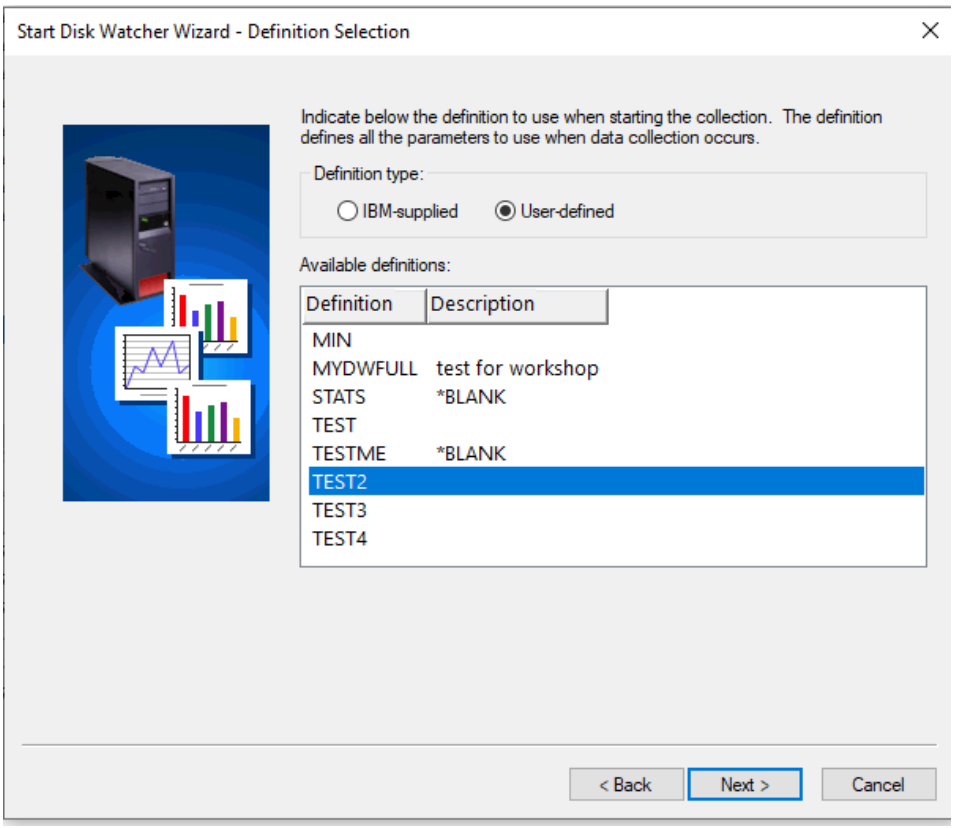

Start Disk Watcher Wizard – Definition Selection

5.3 Basic Options

The Basic Options Page allows the user to specify the collection name, library, collection interval as well as scheduling options. The following is an example of this interface:

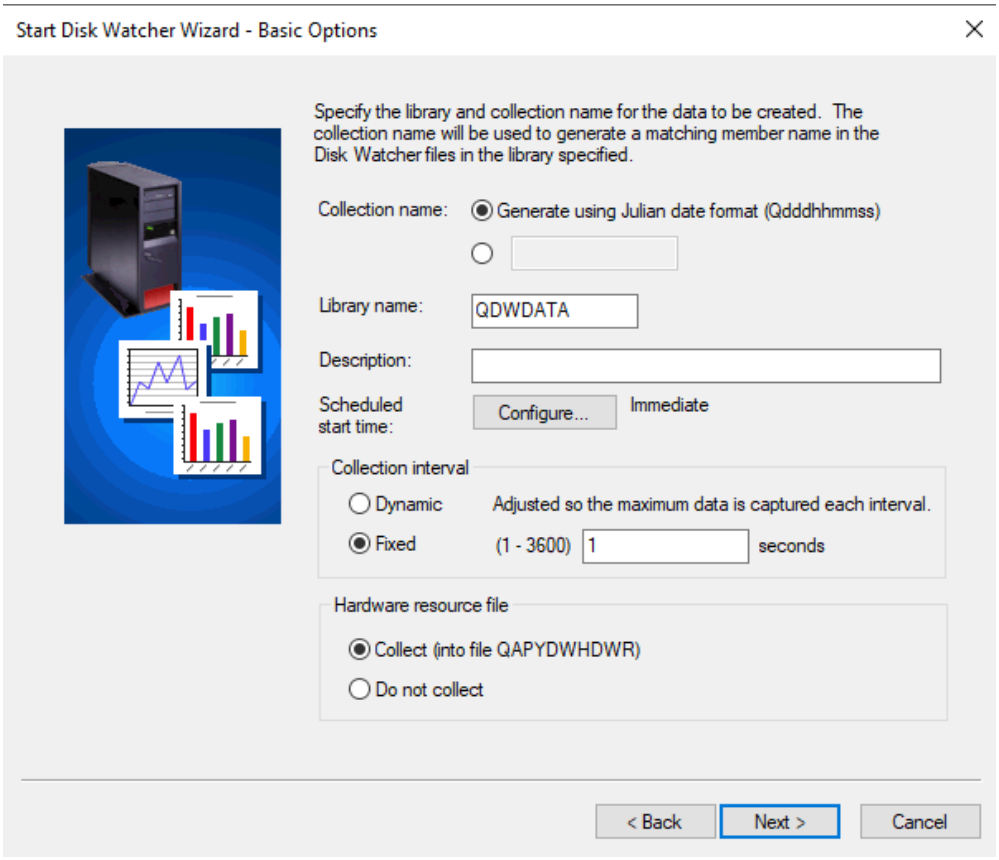

Start Disk Watcher Wizard – Basic Options

5.4 Scheduling Options

This page allows the user to determine a specific date and time for the collection to begin collecting data. By clicking the checkbox, the user can optionally include a date/time to schedule the collection. This option will create a scheduled job on the system.

Use the iDoctor Scheduled Jobs window to check the status of scheduled iDoctor jobs on the system. Access that window by right clicking the Disk Watcher icon in the Disk Watcher component view.

Tip: To configure the default scheduled time (number of days and hours in advance) preference, see the Preferences -> Scheduling interface.

An example of this page of the Wizard is:

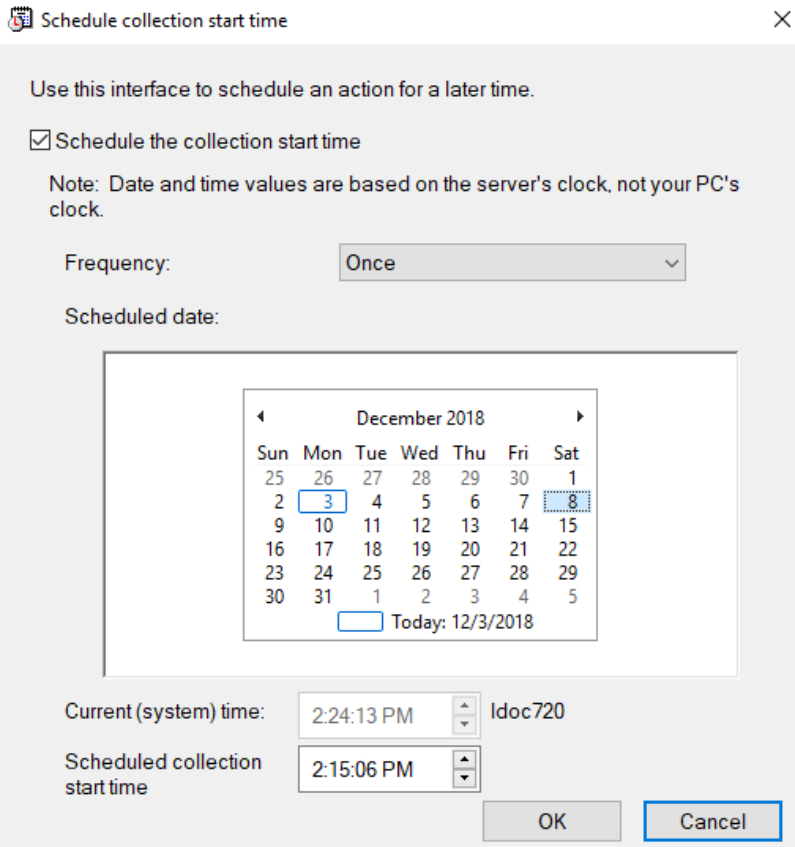

Start Disk Watcher Wizard – Schedule Collection Start Time

5.5 Termination

The Termination Page allows the user to specify what conditions should cause the collection to end. Whichever option is satisfied first, will cause the collection to end immediately.

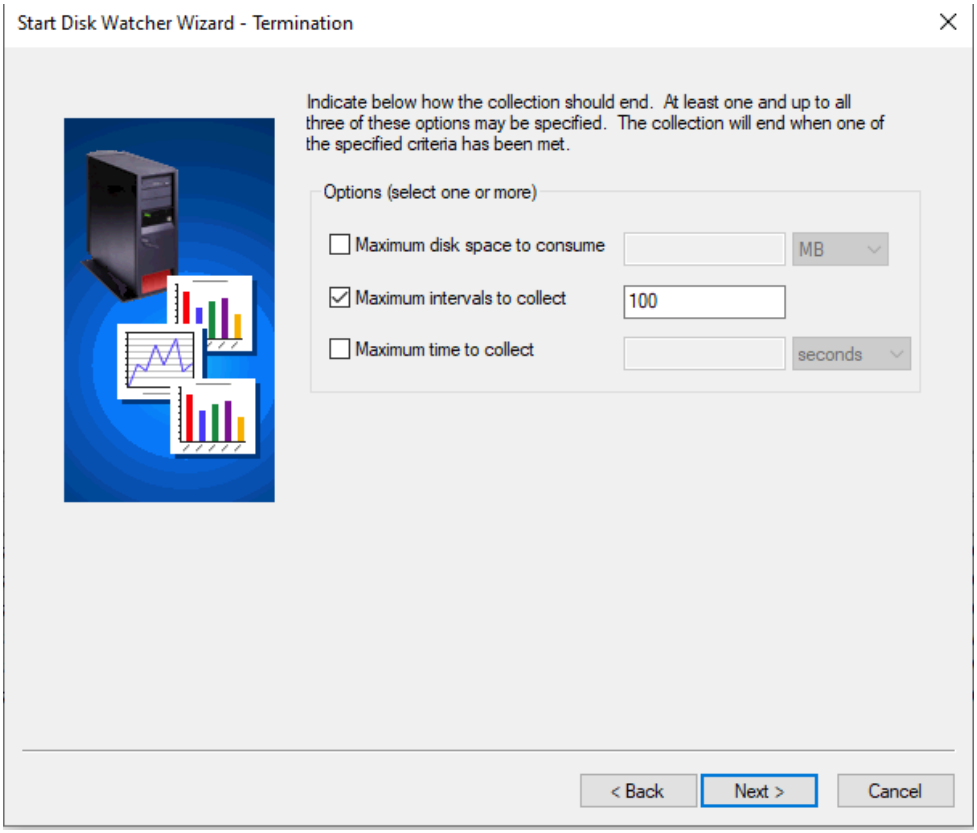

Start Disk Watcher Wizard – Termination

5.6 Summary

This page provides complete details about all selections made in the wizard. If anything listed doesn't look right, use the Back button to go back and make any changes necessary. After clicking 'Finish' a STRDW command will be issued to start the collection. This command is listed at the bottom this page and can be copied to a green screen session and modified if necessary.

After the collection is started will take several seconds before anything appears in the GUI while the collection is initialized. Use F5 to refresh the list of collection in the collection library to work with the new collection.

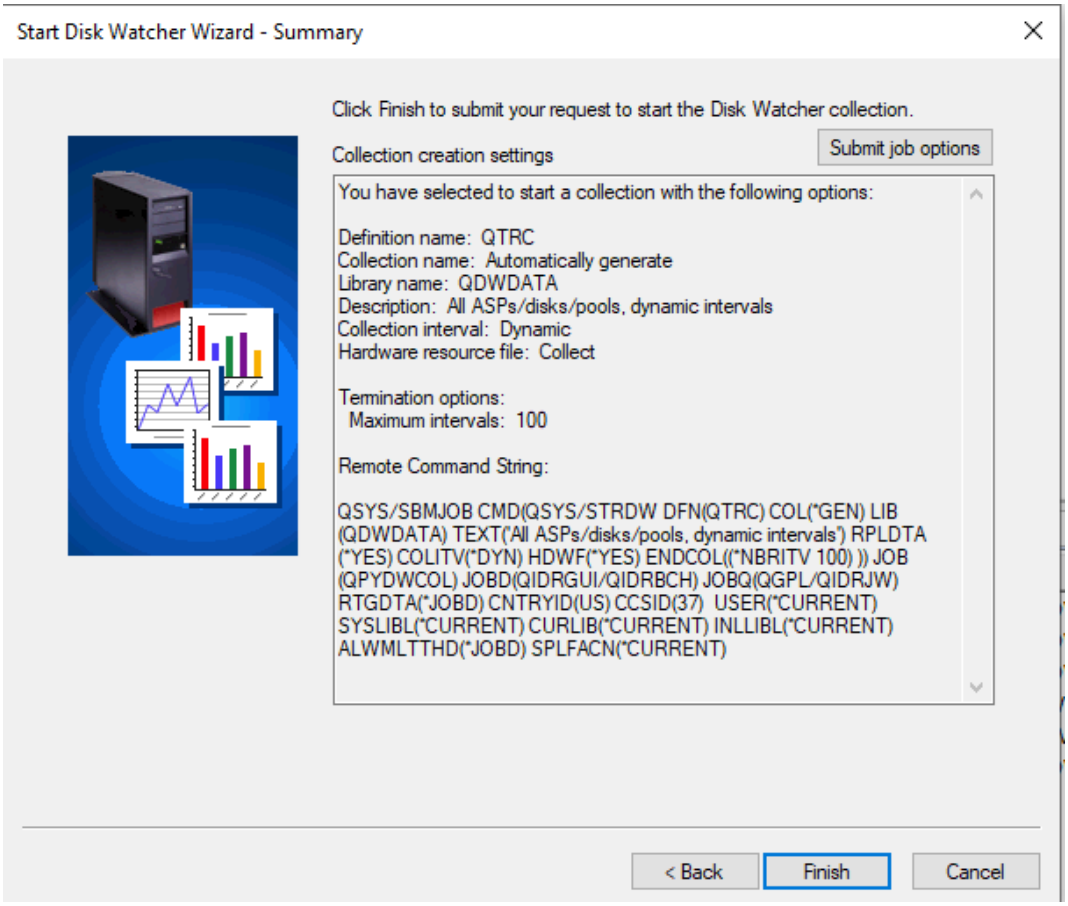

Start Disk Watcher Wizard – Finish

The following section lists the parameters available on this interface:

Option	Description
Submit job options	This button allows you to change parameters on the SMBJOB command. For more information see the help text for the SBMJOB command on the IBM i. 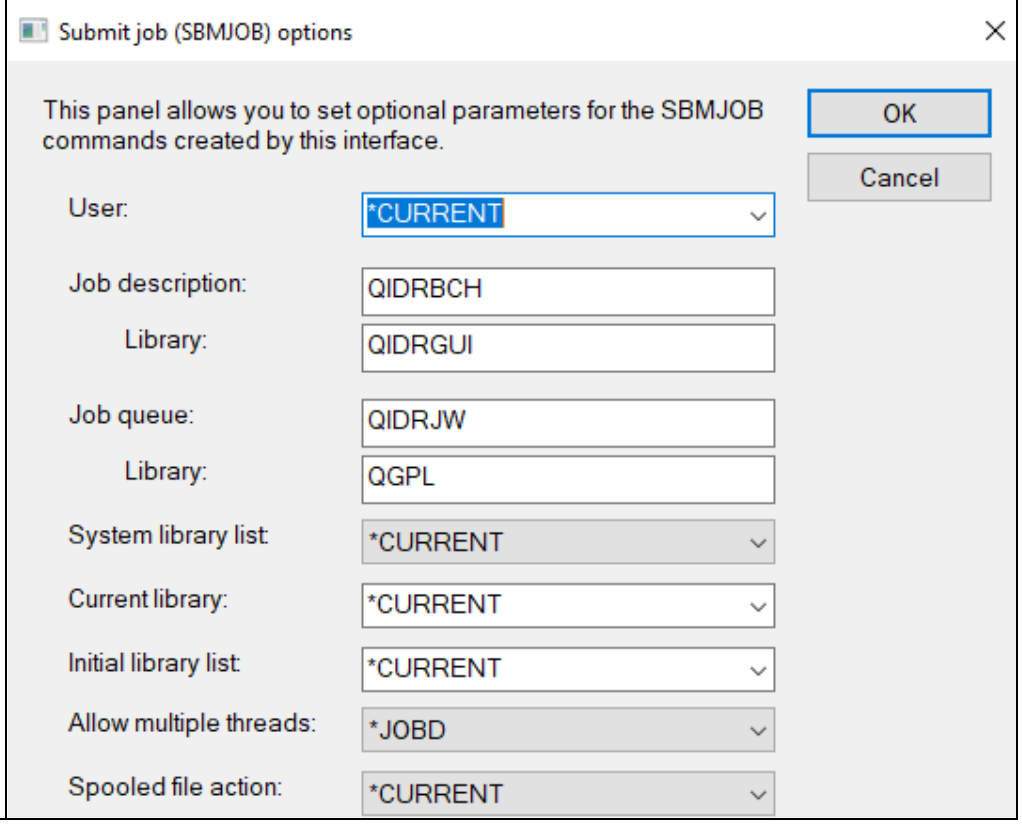

6 Libraries Folder

This folder contains the libraries on the system that contain Disk Watcher data. Specifically, these are the libraries containing file QAPYDWRUNI. The list displays each library's name and description. By expanding a library in the tree, you will see the collections that exist within it.

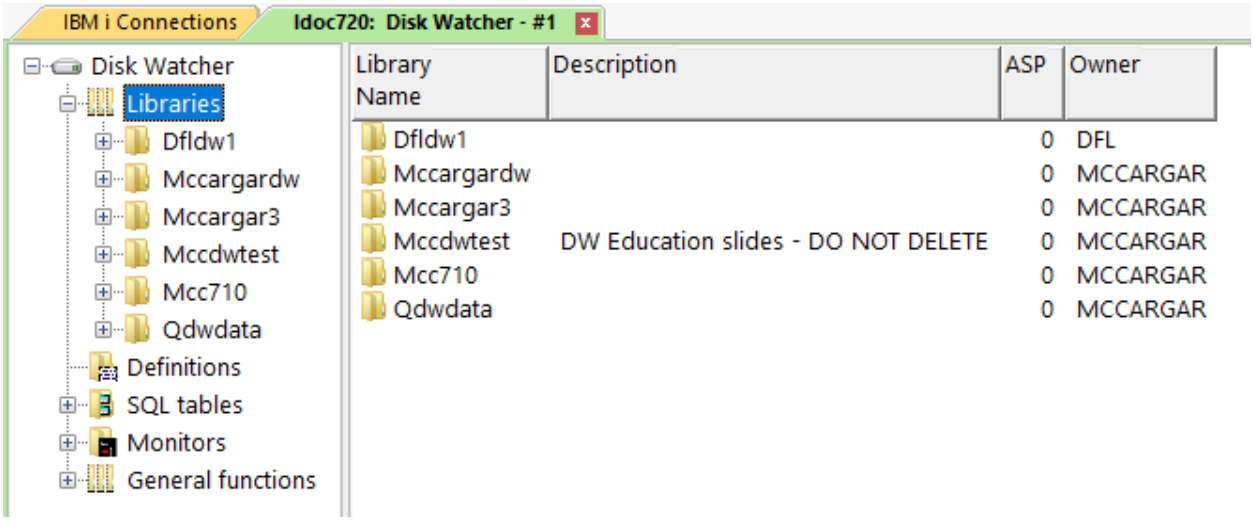

Libraries in the Disk Watcher Component View

For more information on this, visit the Main Window PDF documentation on the Libraries Folder and Library Folders.

6.1 Menu Options

The following menu options are available by right clicking on a library in the component view.

Menu Item	Description
Start Collection...	This menu will open the Start Disk Watcher Wizard where the user can define and run a collection.

Additional menu options that are common to all library folders in iDoctor are discussed Main Window PDF documentation.

7 Monitors

Disk Watcher monitors allow for 24x7 collection of Disk Watcher data on a system. They run continuously storing only the most recent collections desired. Disk Watcher monitors will run until ended manually by the user. Monitors can be held and released if the user wishes to stop collecting data for now and then continue collection again later. Monitors can also be scheduled to start and end at the desired times.

Once a monitor has been started and ended, it must be restarted using the Restart Monitor option. You cannot use the Start New Monitor option to restart an existing monitor.

A Monitors folder is provided in Disk Watcher to allow the user to work with the monitors that exist on the current system. For more information about monitors, see the section on Monitors in the Main Window PDF documentation.

8 SQL Tables

This folder contains all the SQL tables that exist on the system generated by Disk Watcher analyses.

The folder exists in 3 places and each will filter the contents appropriately based on where it is located:

Location	Description
Under Disk Watcher	Entire system
Under a library	All collections in the library
Under a collection	Only this collection

For more information see the SQL Tables section in the Main Window PDF documentation.

9 Collections

Moving down the tree within each Library folder are one or more collections that have been created or are currently being created.

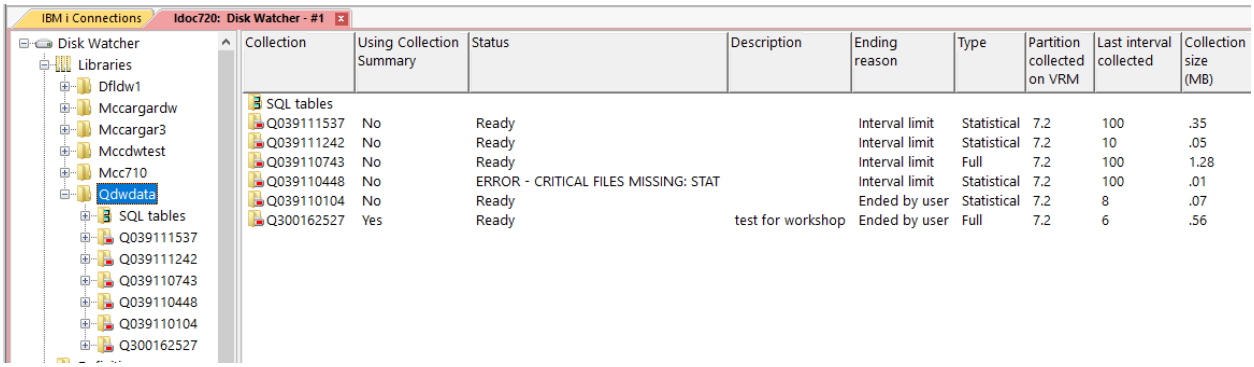

Disk Watcher Collections in a Library

Some of the less obvious columns shown in a list of collections are described below:

Column	Description
Using Collection Summary	This column indicates if the Collection Summary analysis has been ran. This is required to produce trace-based graphs in Disk Watcher. It is only applicable to collections that contain Trace data (either Full or Trace.)
Status	This indicates what files are missing in most components. It is normal for some files to be missing in all components. Only if the status indicates: "ERROR – CRITICAL FILES MISSING", then will the collection be unusable. Place your mouse pointer over this column to get more information about the missing files and which reports they apply to. **Note:** Use the Refresh Status menu option on the collection to update the status if it is incorrect.
Ending reason	This field indicates what caused the collection to end. There are several possible reasons a collection may end as described below: Size limit – The collection exceeded the maximum disk space allowed as described in the definition. Interval limit – The collection stopped when the maximum intervals to collect was met. Time limit – The collection stopped when the maximum time limit to collect was met. ASP limit – The system ASP limit as defined in SLIC service tools has been exceeded causing the collection to end. Ended by user – Disk Watcher detected that the user ended the collection manually.
Type	The type of collection created. The valid types are: **Full** – includes both Trace and Statistical data **Statistical** – does not include information about all I/Os, only higher-level summary data. **Trace** – includes information about each I/O that occurred and optionally the causing program, object or memory pool.

9.1 Menu Options

The table below outlines the different types of operations that may be performed by right clicking on a collection within Disk Watcher.

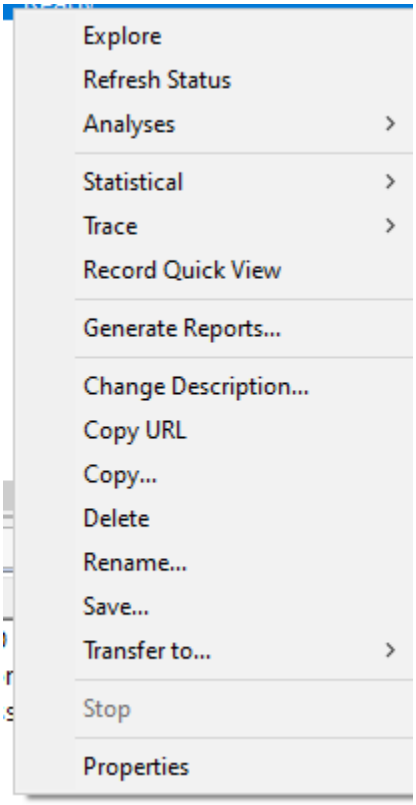

Collection popup-menu

Menu Item	Description
Explore	Show the contents of the collection.
Refresh Status	In some situations, the Status column may indicate files are missing incorrectly. This option is used to refresh the collections cache for the selected collection(s) to be sure that the files are truly missing. This also can be used in cases where report folders are missing.
Analyses -> Run Trace Summary	Use this option to kick off the required analysis for Trace-based collections.
Statistical	Provides several graphs over the statistical data generated by Disk Watcher
Trace	Provides several graphs over the trace data generated by Disk Watcher. **Note:** This option is only available if trace data has been collected and the collection has been summarized.
Record Quick View	Lists the information about the selected collection(s) vertically in a new window.
Generate Reports…	This option can be used to build a report of the desired set of Disk Watcher tables and graphs. The report consists of a screenshot of each graph along with its title and collection information. The reports are built into a HTML page and displayed in the web browser when completed.
Copy URL	Creates a link to the component, library and collection that can be accessed later, or sent to another user.
Copy...	Copies one or more collections to another library. Selecting multiples is only available from the list side of the component view.
Delete...	Deletes a collection. Select multiple collections in order to delete more than one at a time. Selecting multiples is only available from the list side of the component view.
Save	Saves the selected collections to a save file on the system. The save file will be added to the list under the Saved collections folder.
Transfer to...	FTP one or more collections to another system. Selecting multiples is only available from the list side of the component view.
Stop	Stops an active collection.
Properties	Use this menu to display the property pages for the collection. The property pages provide quick access to additional summary information about the collection.

Most of these options are discussed in more detail in the Main Window PDF documentation.

9.2 Stop

An active collection can be stopped by using the Stop menu found by right clicking on a collection within the Disk Watcher component view. This action will issue the ENDDW command which by default will take 30 seconds to end the collection.

9.3 Properties

This section covers the property pages for a collection. Access the property pages by right clicking on a collection and choosing the Properties menu.

9.3.1 General

The General property page provides basic information about the collection such as when it was created.

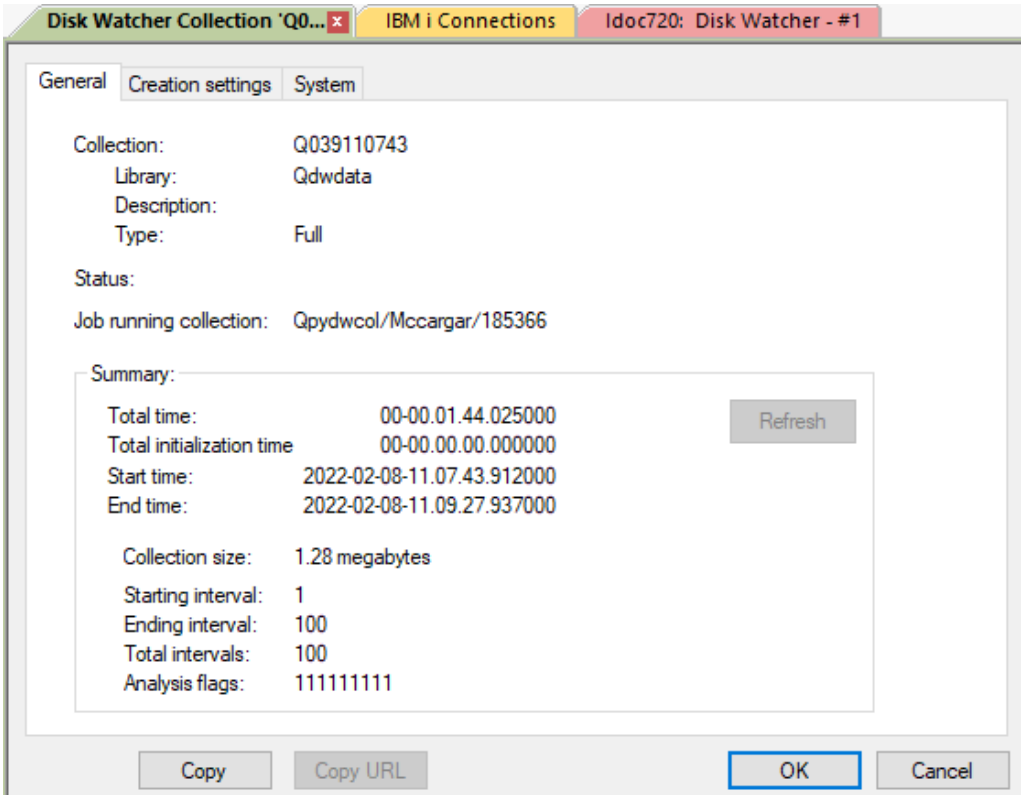

Collection Properties - General

The following information is displayed on the General property page:

Option	Field Description
Collection	Name of the collection. This matches the member name used in the QAPYDW* files on the server in the library specified.
Library	Library the collection resides in.
Description	Description of the collection.
Type	The type of collection indicating the type(s) of data captured: Statistical, Trace or Full (which is both Statistical and Trace)
Status	The status of the collection. This could indicate if the job running the collection failed or that the collection is ready for analysis.
Job running collection	Displays the name of the job that created or is currently creating the collection. If the job log is available a button will be shown to display it.

Option	Field Description
Total time	Displays the total run time of the collection in timestamp format.
Total initialization time	Displays the estimated initialization time for the collection in timestamp format. This is an estimate of the amount of time it took between the collection being started and the 1st interval of data being collected
Start time	The time the collection was started.
End time	The time the collection ended (if it has ended).
Collection size	The total size of the collection. This number does not include any summary files generated.
Starting interval	The 1st interval number detected in the collection.
Ending interval	The last interval number detected in the collection.
Total intervals	The total number of intervals found in the collection.
Analysis flags	This is for IBM/debug use.

9.3.2 Creation Settings

The Creation settings property page provides details about the parameters that were used when creating the collection.

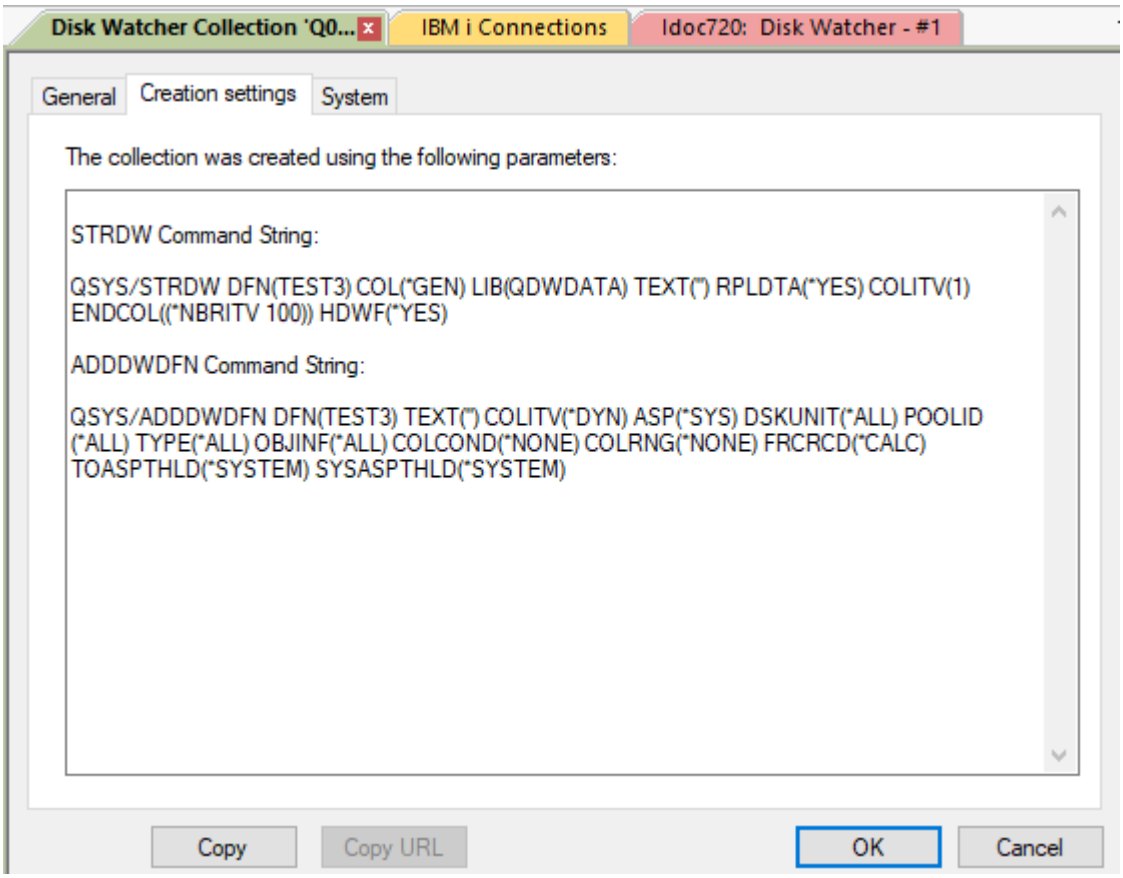

Collection Properties – Creation Settings

The information shown on this window contains 2 command strings used to create this collection.
1 command string for STRDW and the other for ADDDWDFN.

9.3.3 System

The system property page displays details about the system the collection was created on. This information includes the type, model, operating system VRM and the number of processors.

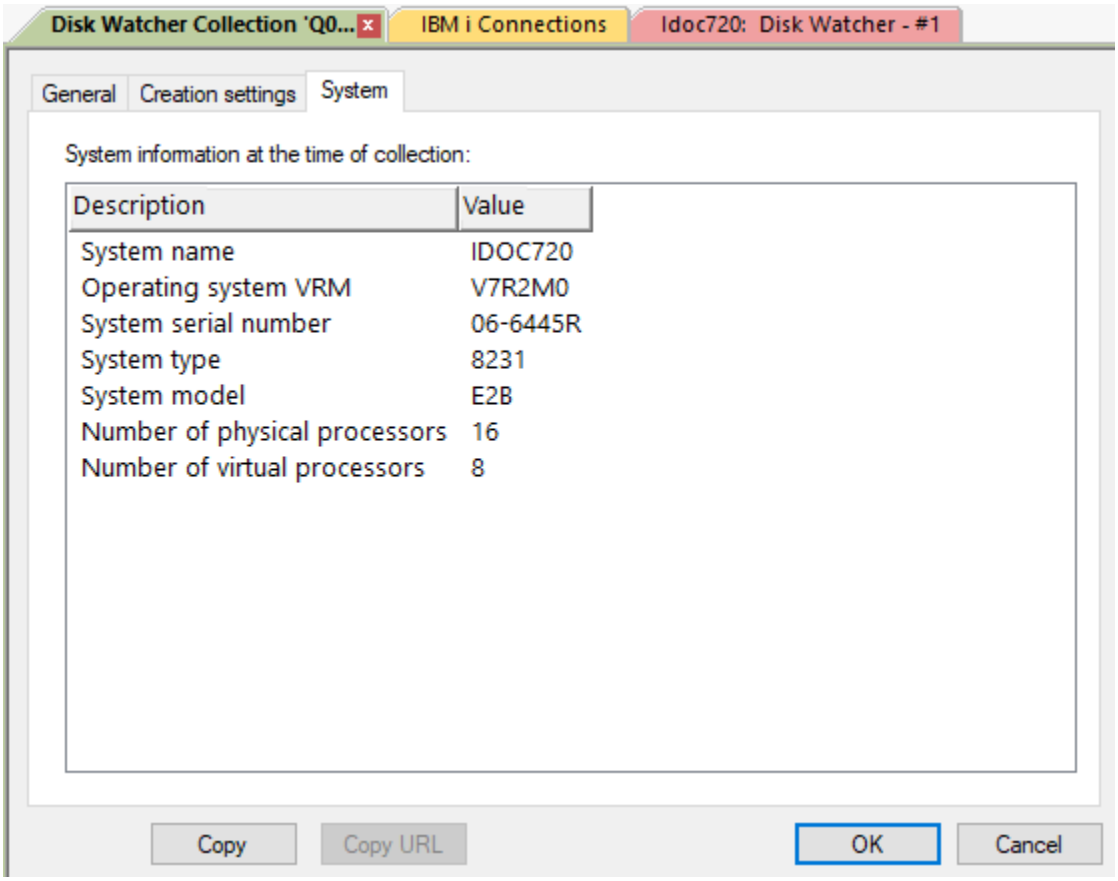

Collection Properties – System

10 Analyses

The available Analyses in Disk Watcher and what they provide is described in this section.

All analyses are written as SQL stored procedures and are initiated from the Analyses menu after selecting one or more collections and right-clicking. Each analysis has a 'fast path' option that allows it to be ran without visiting the Analyze Collection window.

10.1 Analyze Collection Window

The Analyze Collection window presents the user with a list of available analyses that can be ran over the currently selected collection(s). It is opened using the Analyses -> Analyze Collection... menu.

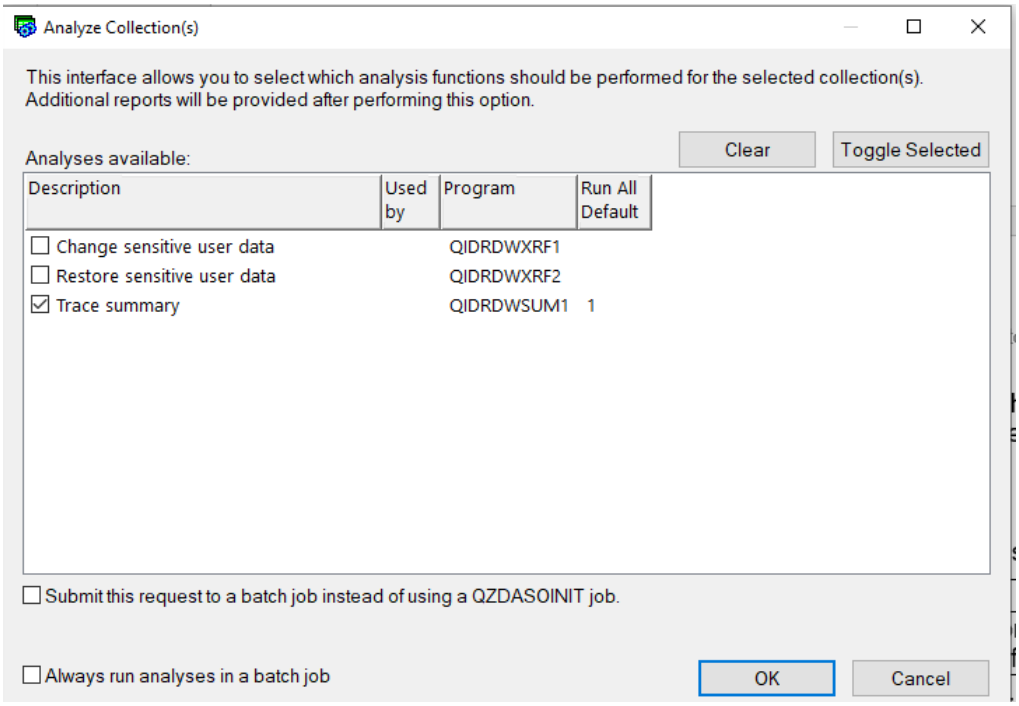

Disk Watcher Analyze Collection(s) Window

Each available analysis is presented to the user on this screen.

The controls on this interface and what they do is described in more detail in the Job Watcher documentation.

10.2 Change/Restore sensitive user data

These options are used to hide or replace potentially sensitive data in Disk Watcher. It will update the following things in the Disk Watcher files:

1) User programs names and library names
2) User module names
3) User program/module procedure names
4) System name and system serial number
5) User job names, current user profiles and wait object names.

Note: This data is changed but not lost, it can be restored again using the Restore sensitive user data analysis. This will update the Disk Watcher files changed back to their original values.

10.2.1 SQL Tables Generated

The list of SQL tables generated by the analysis is shown below:

Table description	SQL table
PGM cross reference file	DWXRFPGMR_<MBR>
RUNI cross reference file	DWXRFRUNI_<MBR>
JOB cross reference file	DWXRFJOB_<MBR>
OBJ cross reference file	DWXRFOBJ_<MBR>

10.3 Trace Summary

The Trace Summary analysis summarizes the disk IOs into 1 second intervals in order to make the graphing done in iDoctor more feasible.

Note: This analysis is required in order for the Trace menu and graphs to appear.

10.3.1 SQL Tables Generated

The list of SQL tables generated by the analysis are shown below:

Table description	SQL table
Trace 1 second intervals	QAIDRDWINTI_<MBR>
Trace statistics	QAIDRDWTSUM_<MBR>
Summarized by interval, ASP	QAIDRDWTSUM2_<MBR>

11 Columns

This section covers the various columns that appear in Disk Watcher and a brief description of what they mean. These metrics appear in all types of reports available which are discussed in the later sections.

These are derived using SQL statements from the QAPYDW* files that exist in the IBM i.

Column	Description	Comments
RDRATE	Reads per second	Disk read operations per second
WRTRATE	Writes per second	Disk write operations per second
OTHRATE	Other I/Os per second	Disk operations other than reads + writes per second. These types of operations are documented in the TCIOFUNCTN column here and would include any operations except 16, 20, 32 and 36. **Note:** These are extremely rare to see in normal performance data and almost always never visible on graphs. For that reason, the rest of the "Other" metrics are not listed here.
RDAVGRSP	Average read response time (ms)	Read time divided by number of reads This **includes** deferred queue wait times.
RDAVGSRVRSP	Average read service time (ms)	Read time divided by number of reads This **excludes** deferred queue wait times.
WRTAVGRSP	Average write response time (ms)	Write time divided by number of writes This **includes** deferred queue wait times.
WRTAVGSRVRSP	Average write service time (ms)	Write time divided by number of writes This **excludes** deferred queue wait times.
AVGKBPERRD	Average read size (kilobytes)	This is the average size of each read operation.
AVGKBPERWRT	Average write size (kilobytes)	This is the average size of each write operation.
RMBPERSEC	Reads per second (megabytes)	In the QAPYDWSTAT file this is the 13 STRDPCnn (read page count) columns added together multiplied by 4.
WMBPERSEC	Writes per second (megabytes)	In the QAPYDWSTAT file this is the 13 STWRTPCnn (write page count) columns added together multiplied by 4.

RDTIME_nn	Reads total response time (seconds)	These columns show the total read time + total deferred queue read time added together for the average response time range indicated in the description.
WRTTIME_nn	Writes total response time (seconds)	These columns show the total write time + total deferred queue write time added together for the average response time range indicated in the description.
RDSRVTIME_nn	Reads service time (seconds)	These columns show the read time **(excluding deferred time)** for the average response time range indicated in the description.
WRTSRVTIME_nn	Writes service time (seconds)	These columns show the write time **(excluding deferred time)** for the average response time range indicated in the description.

TOTRD_nn	Reads (thousands)	This is total reads for the average response time range indicated in the description.
TOTWRT_nn	Writes (thousands)	This is total writes for the average response time range indicated in the description.
RDSRVTIME_nn	Reads service time (seconds)	These columns show the read time **(excluding deferred time)** for the average response time range indicated in the description.
WRTSRVTIME_nn	Writes service time (seconds)	These columns show the write time **(excluding deferred time)** for the average response time range indicated in the description.
RDRATE_nn	Reads per second	This is reads per second for the average response time range indicated in the description.
WRTRATE_nn	Writes per second	This is writes per second for the average response time range indicated in the description.

TOTREADMB	Total reads (megabytes)	This is the total size of all reads.
TOTWRTMB	Total writes (megabytes)	This is the total size of all writes.
RDSRVTIME_nn	Reads service time (seconds)	These columns show the read time **(excluding deferred time)** for the average response time range indicated in the description.
WRTSRVTIME_nn	Writes service time (seconds)	These columns show the write time **(excluding deferred time)** for the average response time range indicated in the description.
DQRATIO	DQ time ratio	This is deferred queue time / service time. This metric appears under the advanced graphs.
DQWRTRATIO	DQ write time ratio	This is deferred queue write times / service write times. This metric appears under the advanced graphs.
DQRDRATIO	DQ read time ratio	This is deferred queue read times / service times. This metric appears under the advanced graphs.

12 Situation – I/Os were missed

In any overview or selection over time graph you may possibly see a red background color. This indicates a time where I/Os were missed and not captured by the STRDW command.

This is likely a time where Disk Watcher itself was overwhelmed and could not capture all data.

An example of this follows:

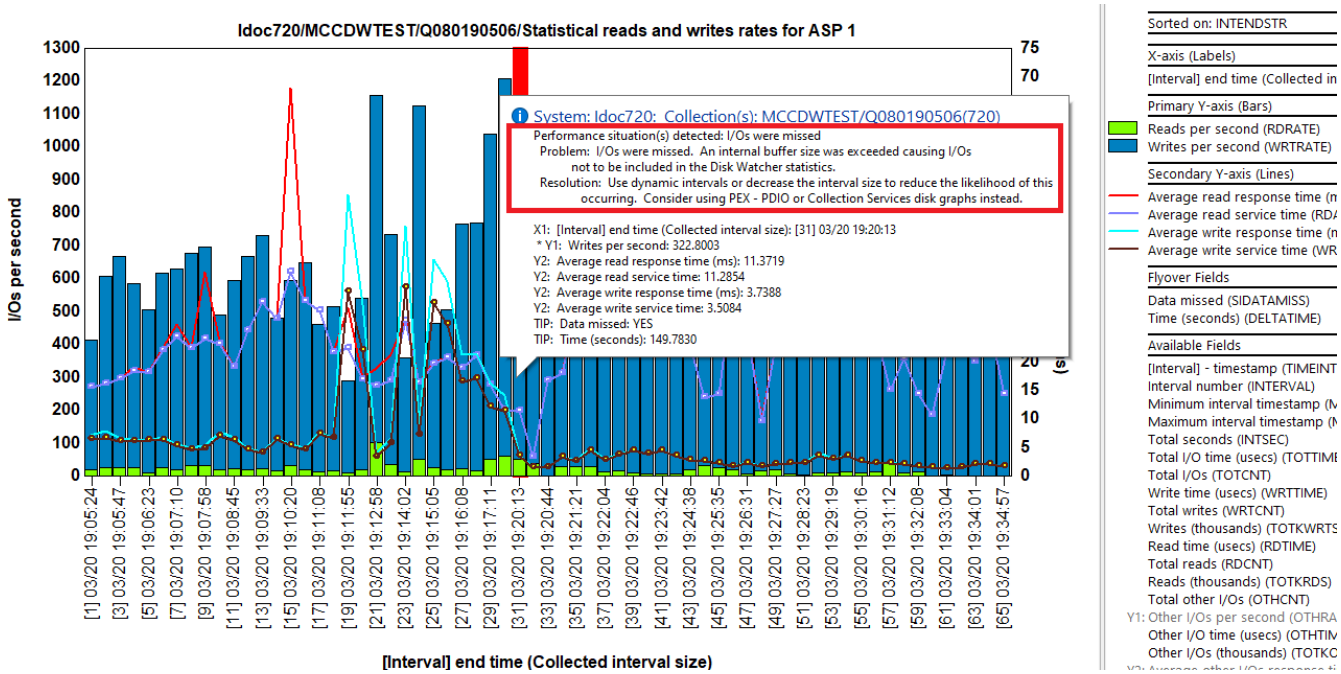

I/Os were Missed Situation Flyover

13 Statistical graphs

There are 12 statistical graphs in Disk Watcher, and these can appear in any of the 3 common formats: Overview, Rankings or Selection over time. Additional graphs are also available under Statistical -> Advanced.

These graphs are built using the statistical data captured by Disk Watcher in file QAPYDWSTAT.

Typically, users start with the desired Overview graph, then select time periods and drill-down into rankings. From rankings, a user will drill down again into selection over time or use the Detail reports menu for additional options.

Tip: The most used graphs are listed first in this list.

If you need help with definitions of specific metrics, see the previous section or contact iDoctor support.

The Statistical folder contains the following:

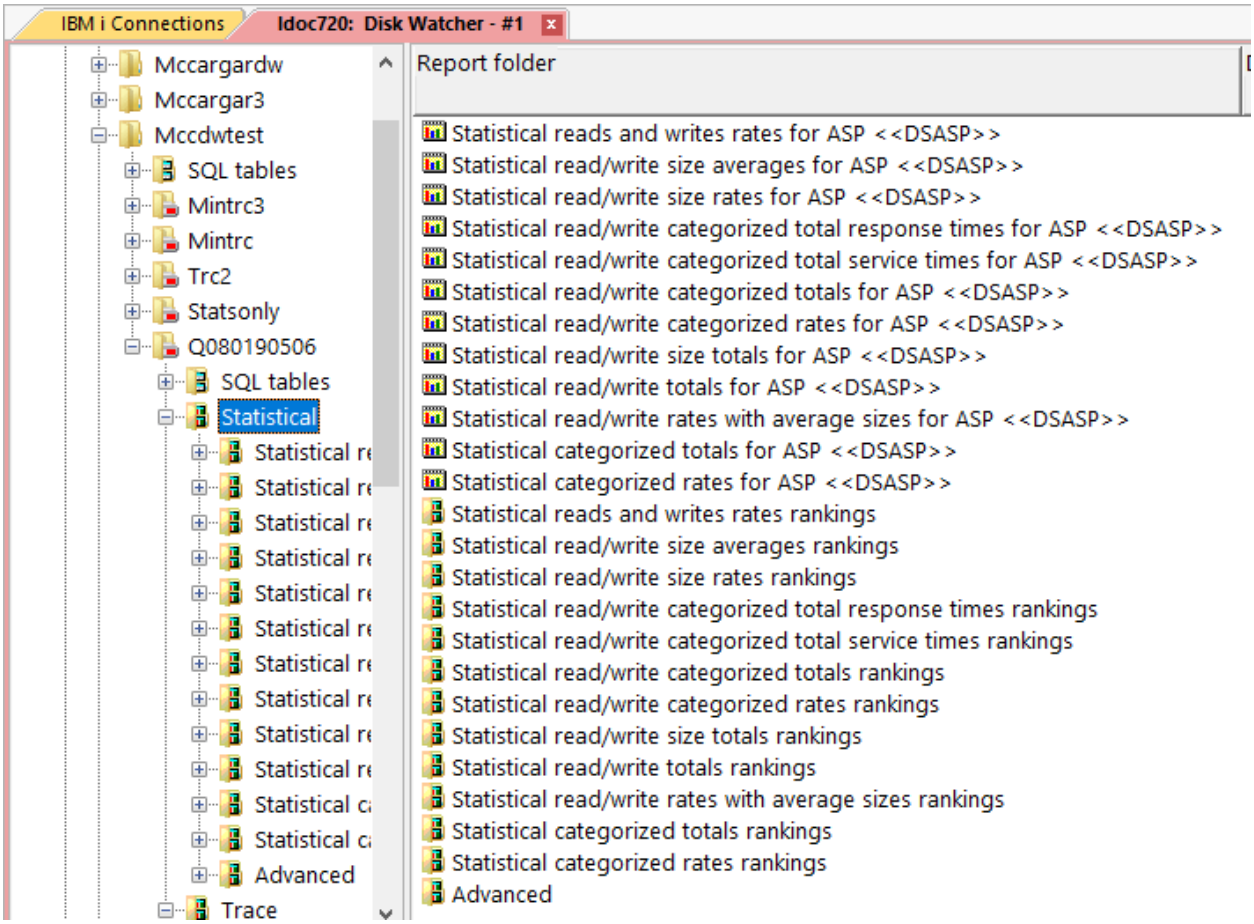

Disk Watcher – Statistical folder

13.1 Statistical reads and writes rates for ASP <<DSASP>>

This graph displays **I/O rates per second** with average response times and service times on the secondary Y-axis.

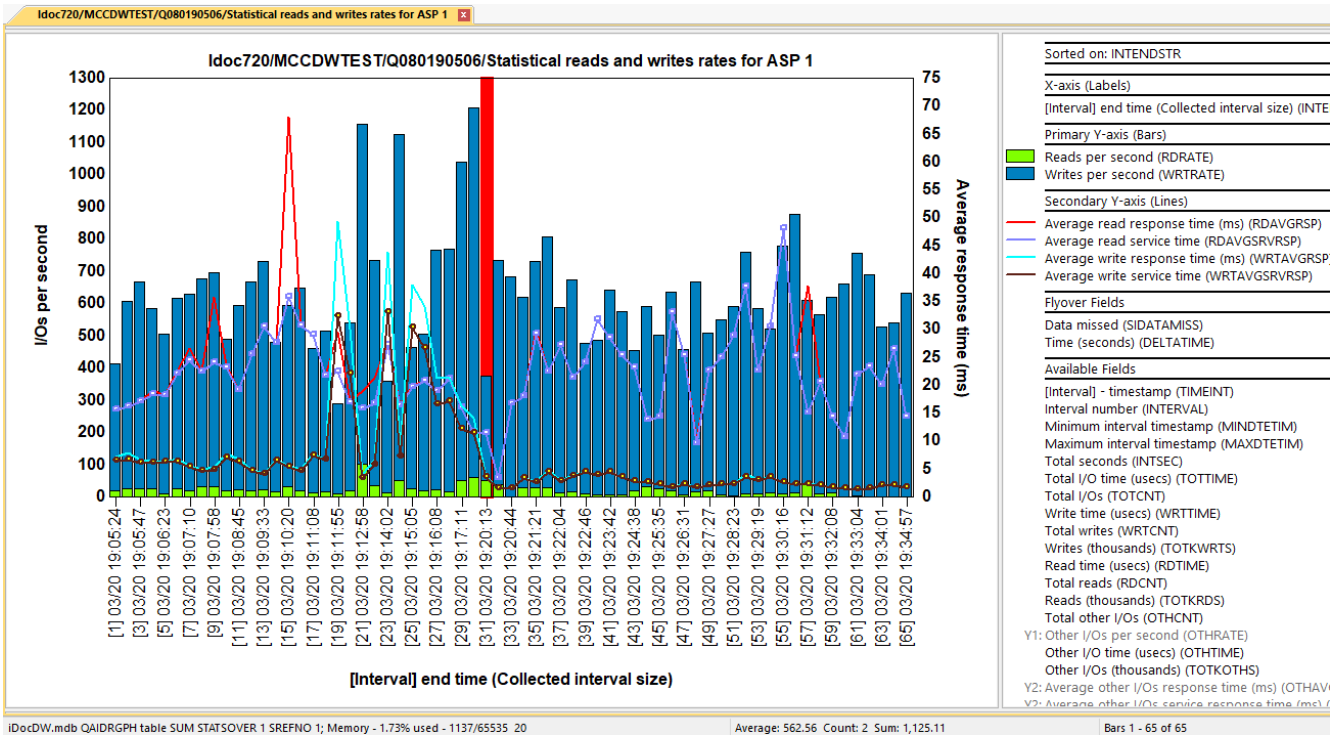

13.2 Statistical read/write size averages for ASP <<DSASP>>

This graph displays **average I/O sizes (in kilobytes)** with average response times and service times on the secondary Y-axis.

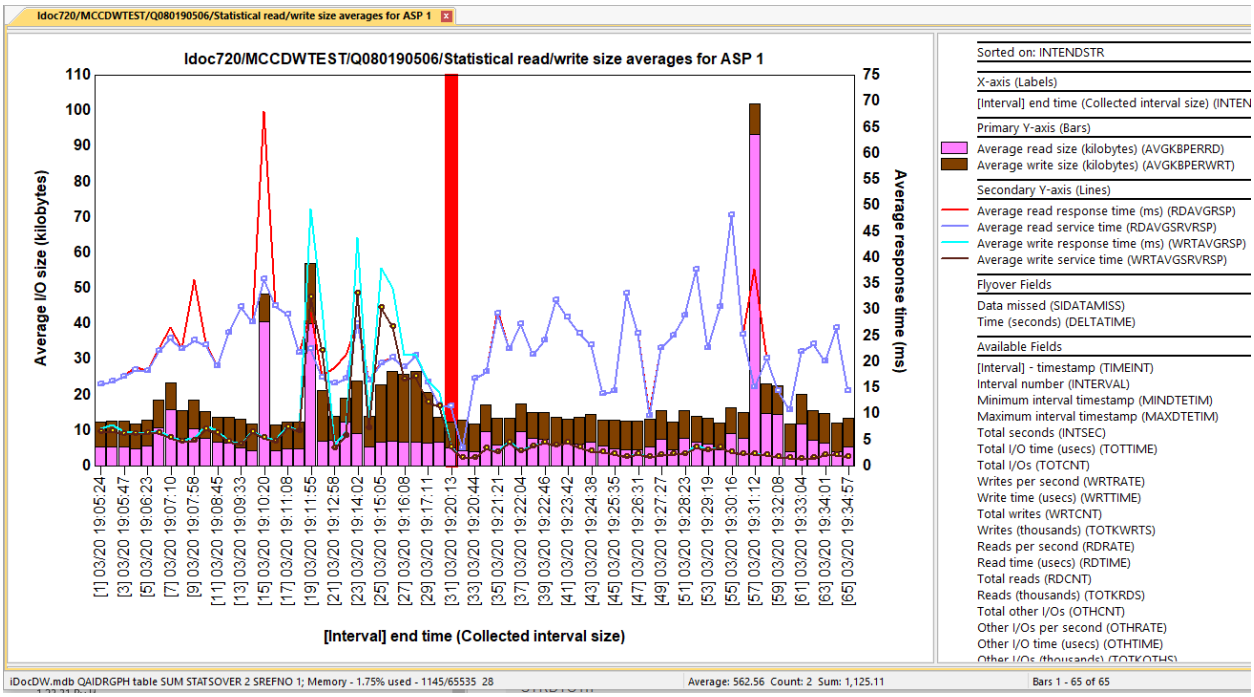

13.3 Statistical read/write size rates for ASP <<DSASP>>

This graph displays the **total I/O size for all operations occurring per second** with average response times and service times on the secondary Y-axis.

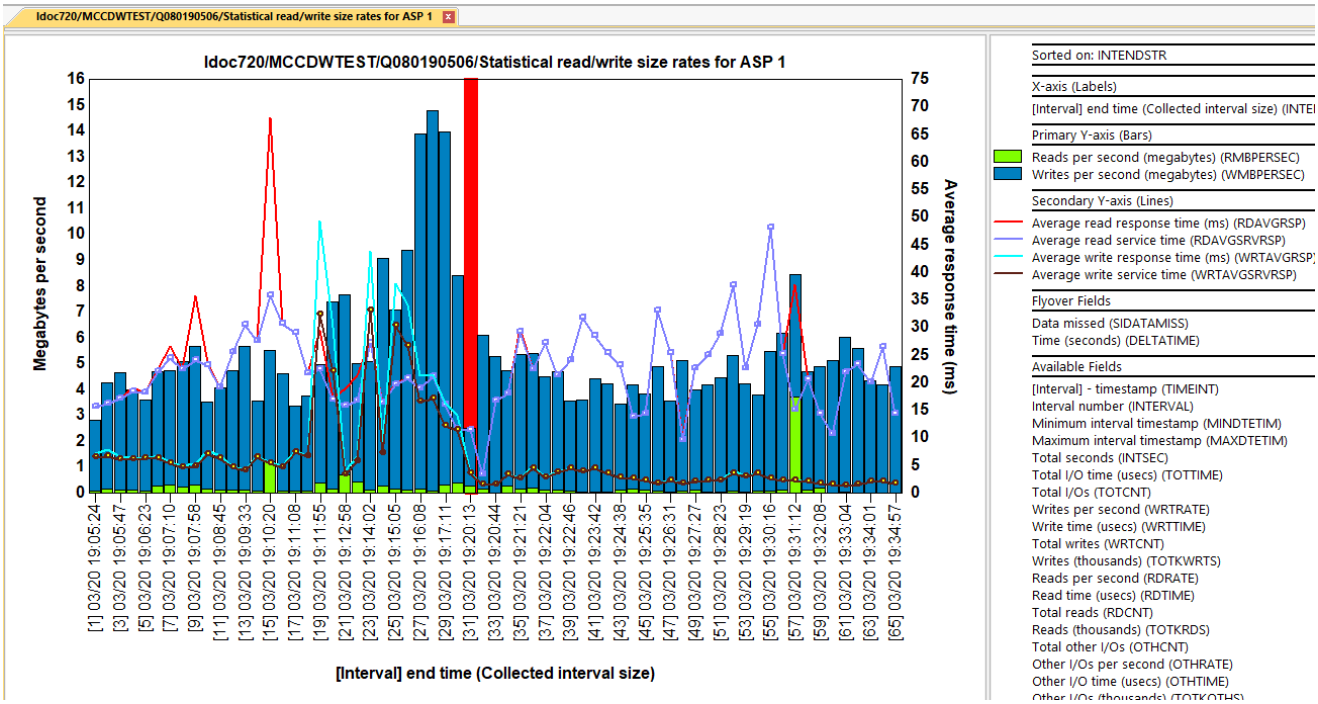

13.4 Statistical read/write categorized total response times for ASP <<DSASP>>

This graph displays the **total I/O times** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

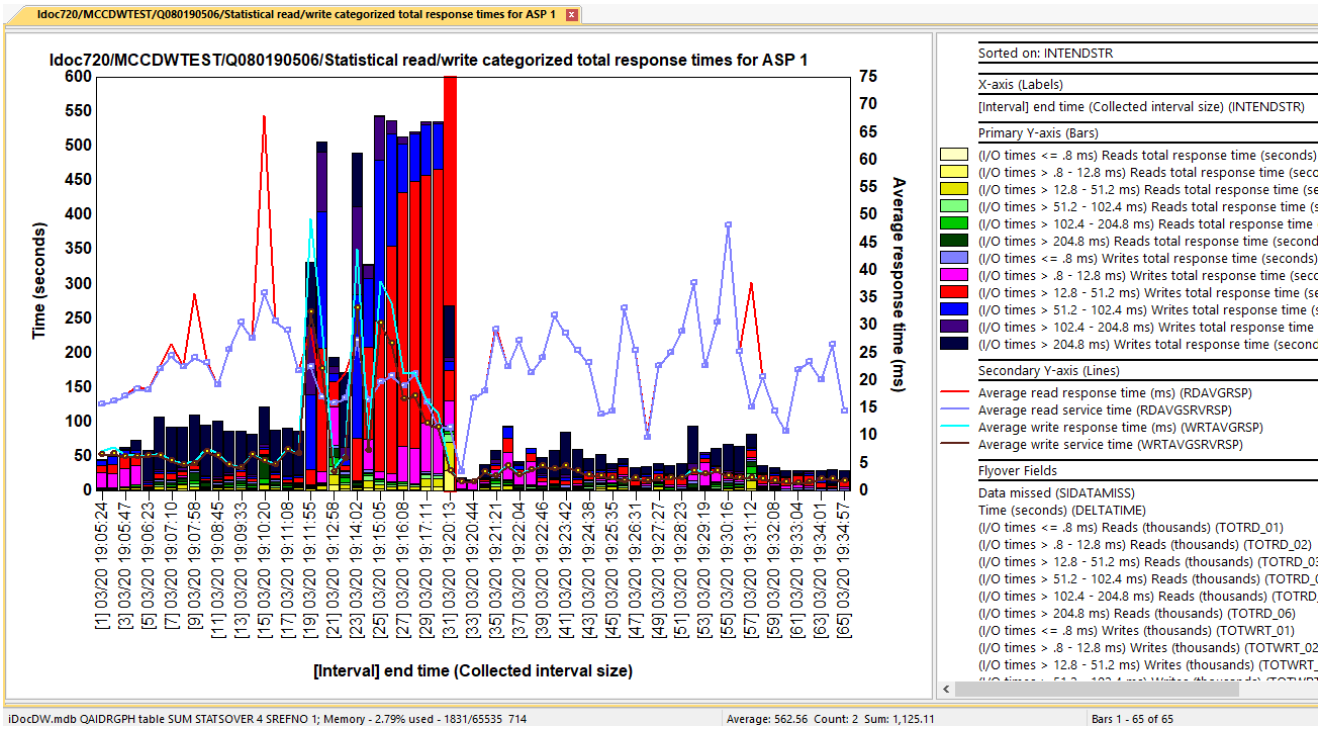

13.5 Statistical read/write categorized total service times for ASP <<DSASP>>

This graph displays the **total service times** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

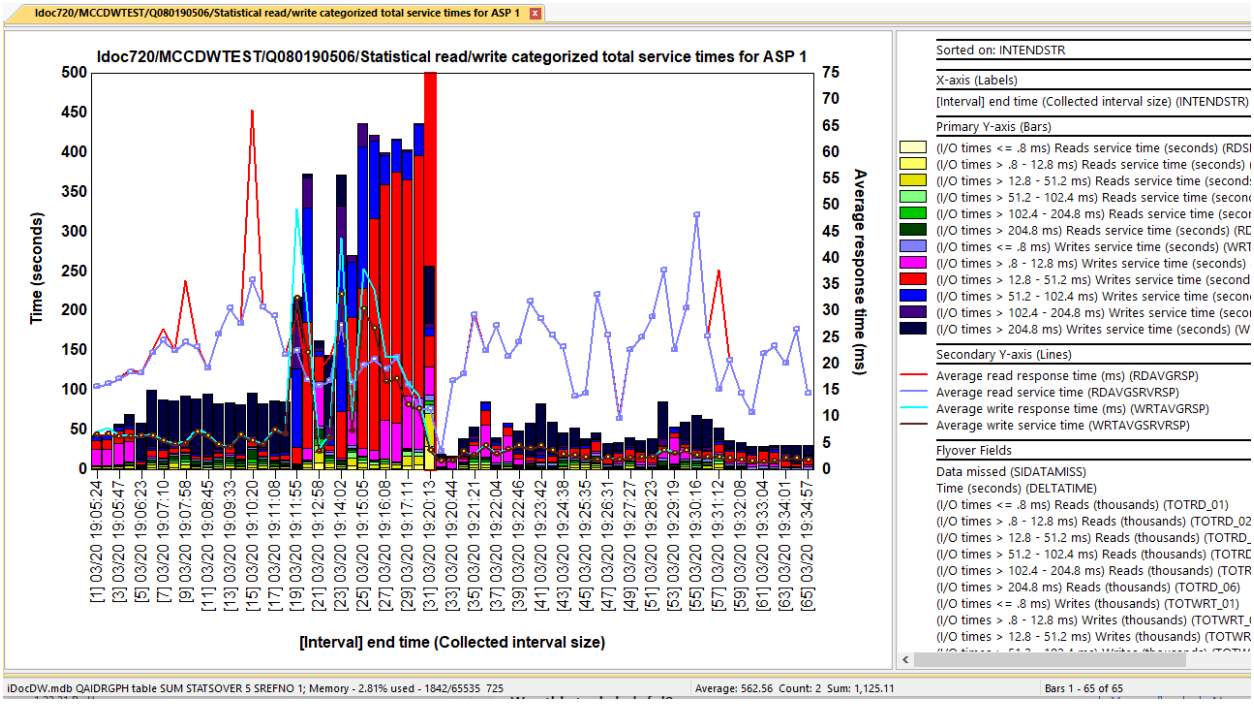

13.6 Statistical read/write categorized totals for ASP <<DSASP>>

This graph displays the **total operations** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

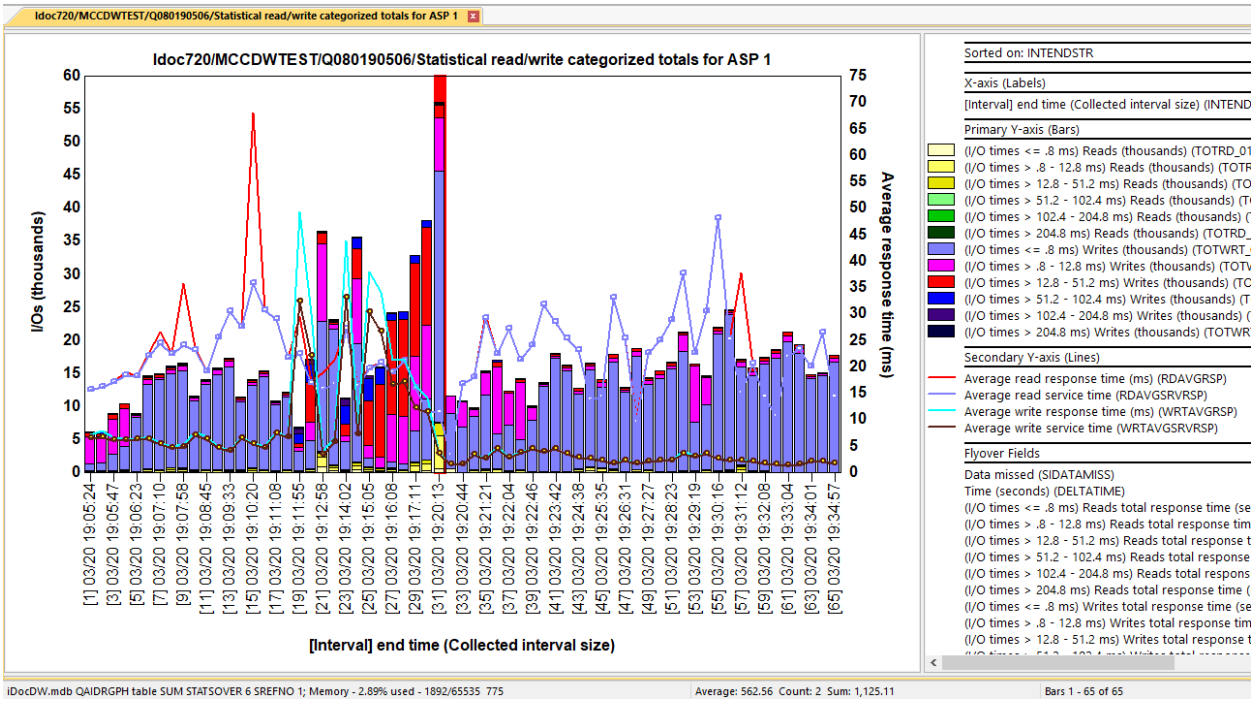

13.7 Statistical read/write categorized rates for ASP <<DSASP>>

This graph displays the **I/Os per second** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

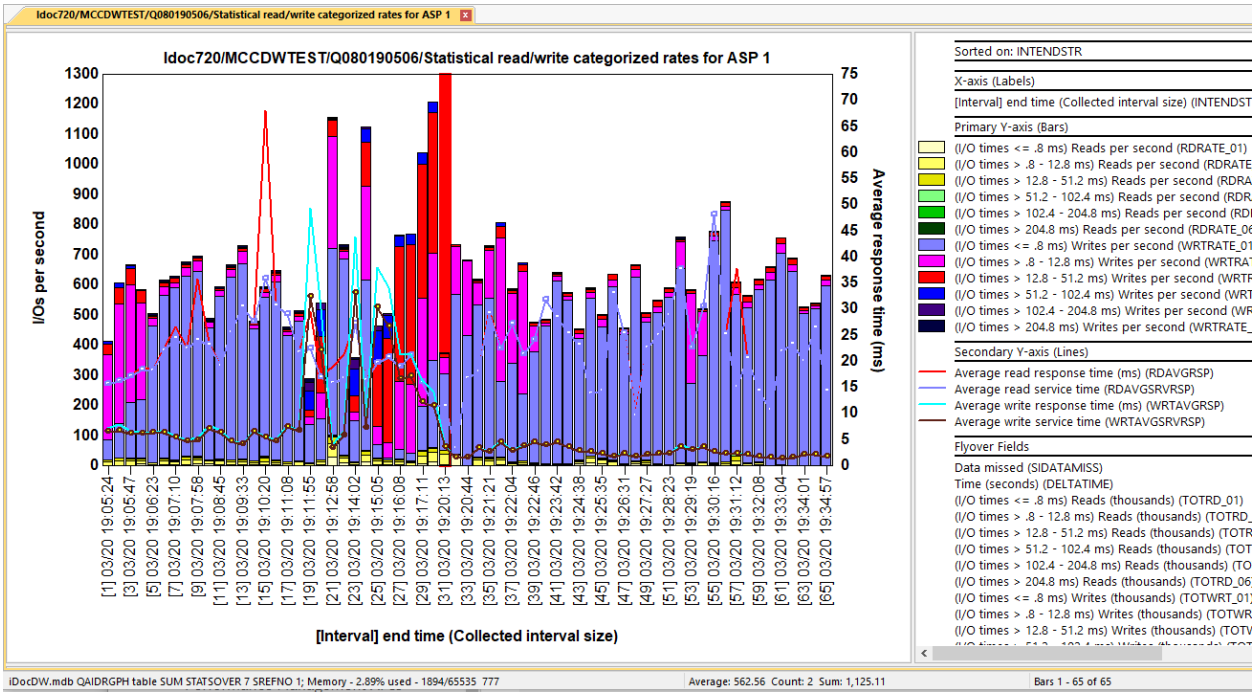

13.8 Statistical read/write size totals for ASP <<DSASP>>

This graph displays the **total I/O size for each of reads, writes and others** with average response times and service times on the secondary Y-axis.

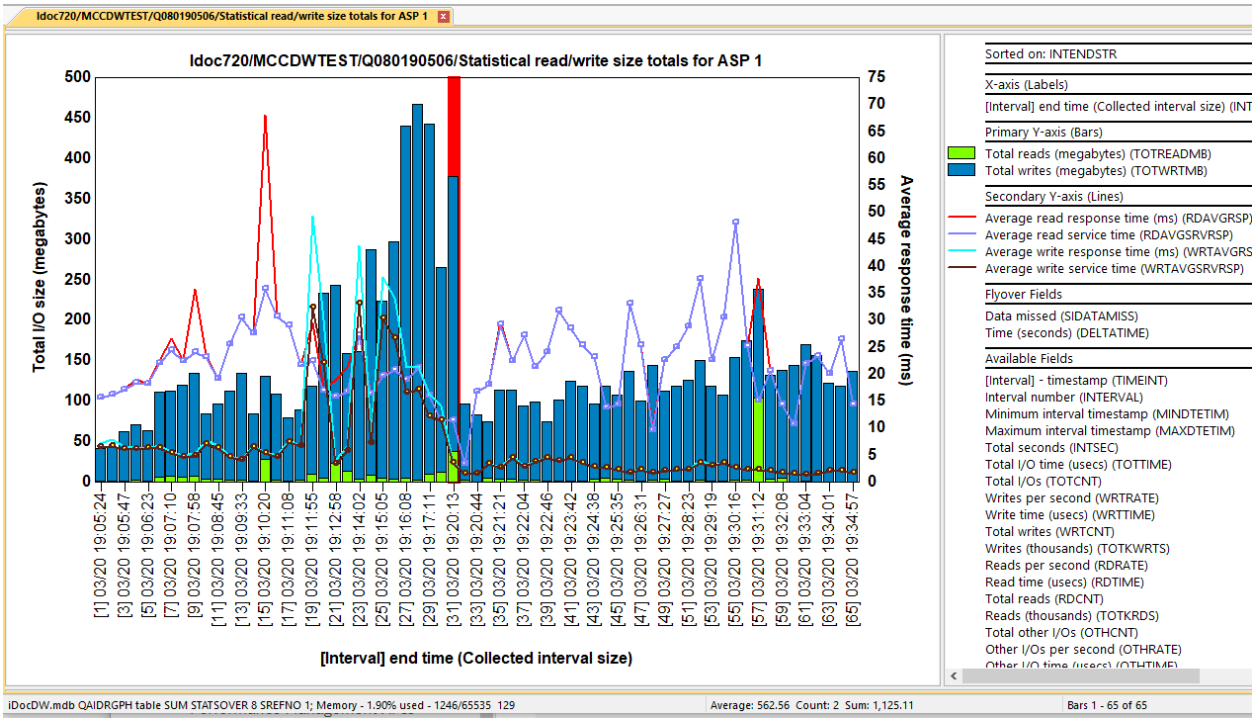

13.9 Statistical read/write totals for ASP <<DSASP>>

This graph displays **I/O totals** with average response times and service times on the secondary Y-axis.

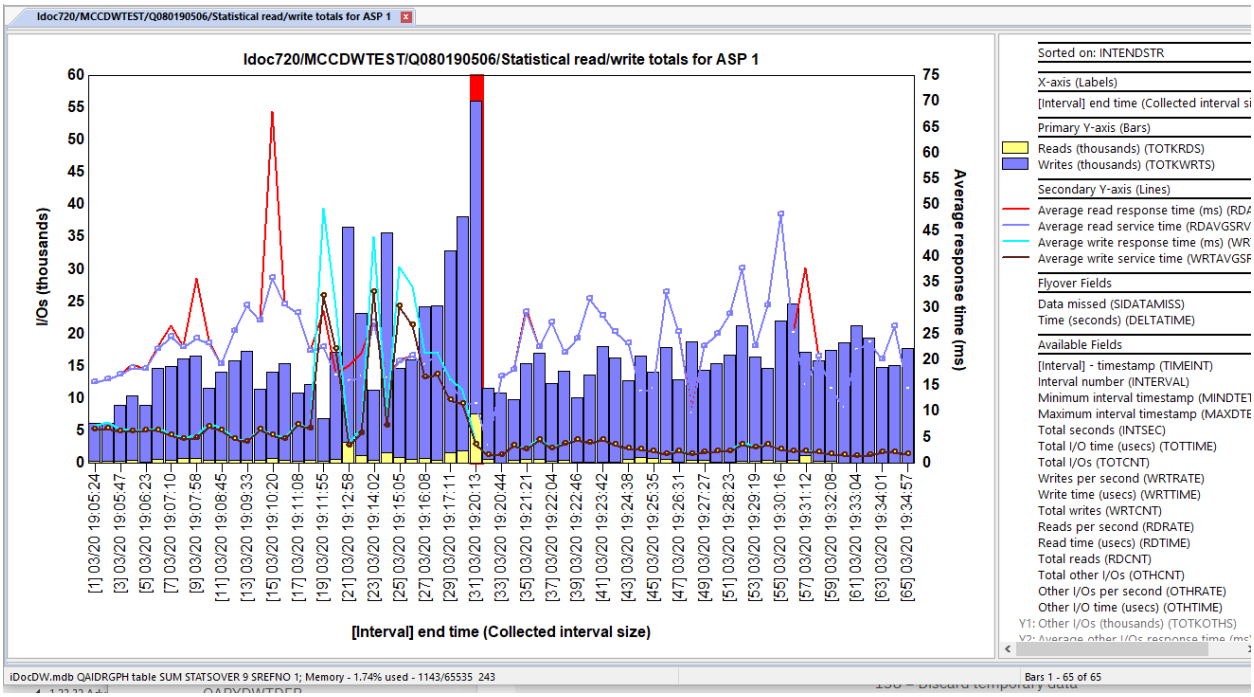

13.10 Statistical read/write rates with average sizes for ASP <<DSASP>>

This graph displays **I/O rates per second** with **average I/O sizes** on the secondary Y-axis.

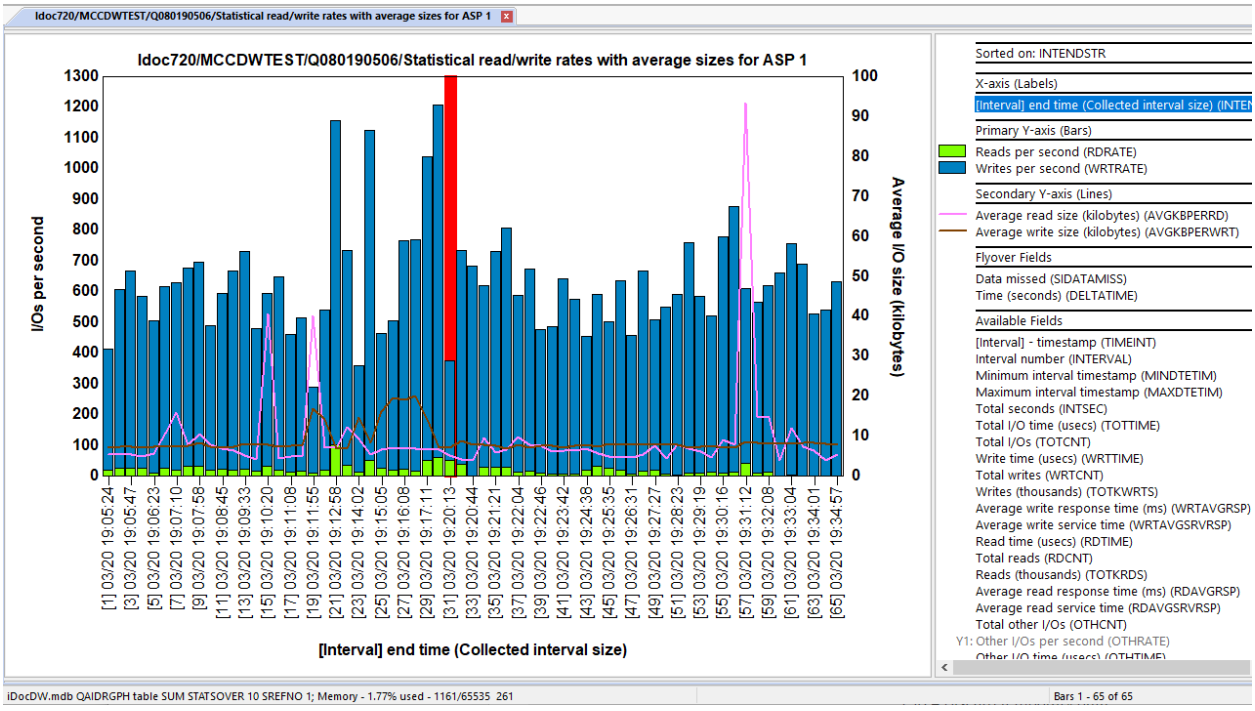

13.11 Statistical categorized totals for ASP <<DSASP>>

This graph combines all 3 I/O types together and shows **total I/Os** with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets

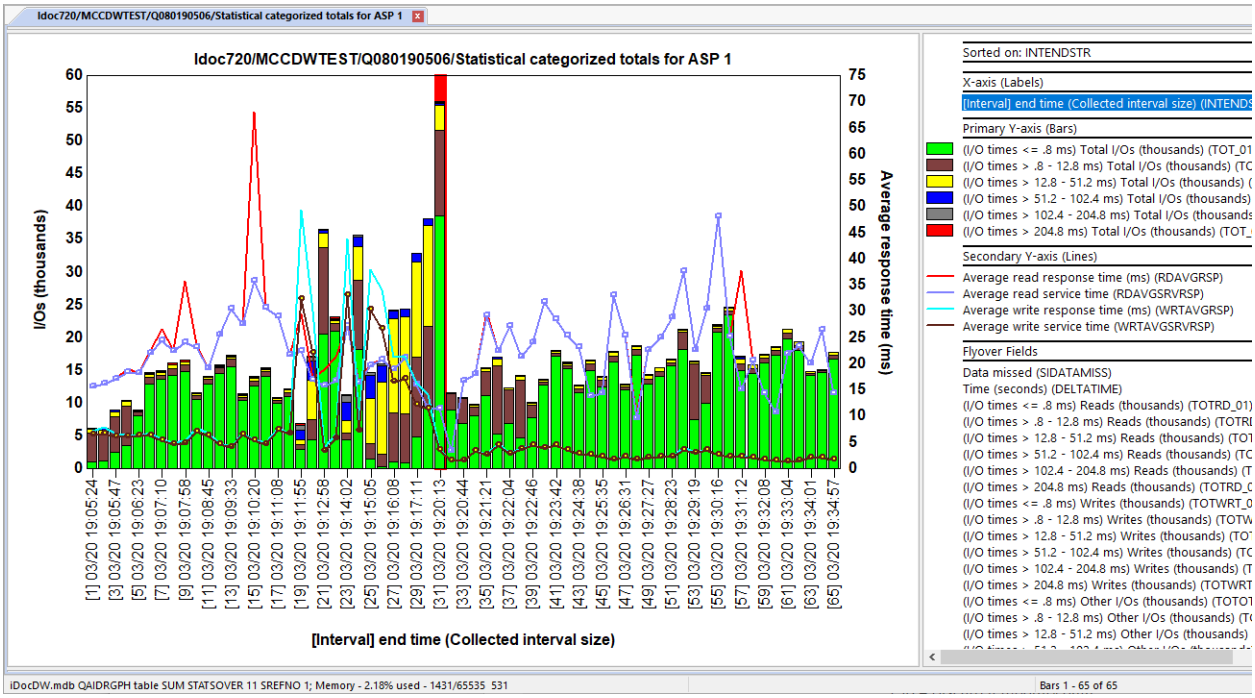

13.12 Statistical categorized rates for ASP <<DSASP>>

This graph combines all 3 I/O types together and shows **I/Os per second** with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets

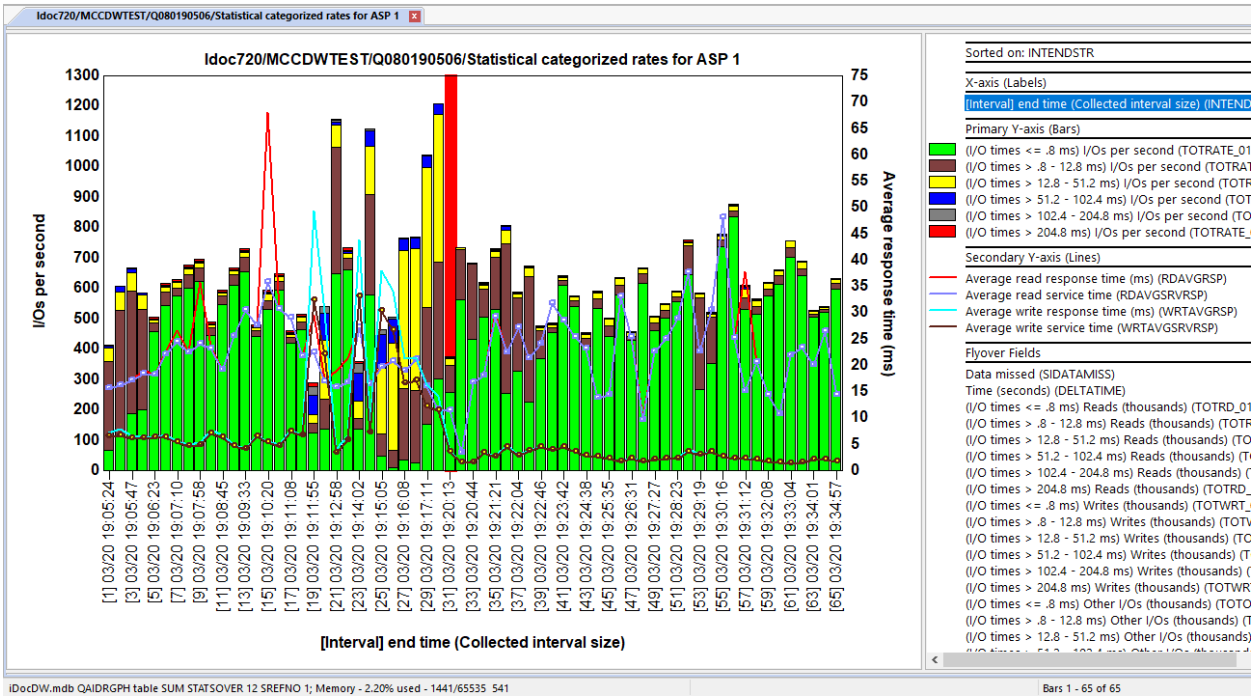

13.13 Rankings graphs

The same set of overview graphs also exist as rankings graphs. For statistical data, the following disk groupings are available:

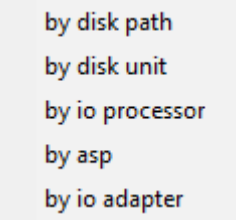

Disk groupings

Access to these can be done either using the appropriate Rankings graph options at the collection-level or via a drill down from an Overview graph.

The typical method to drill down into rankings is to select a time period and right-click as follows:

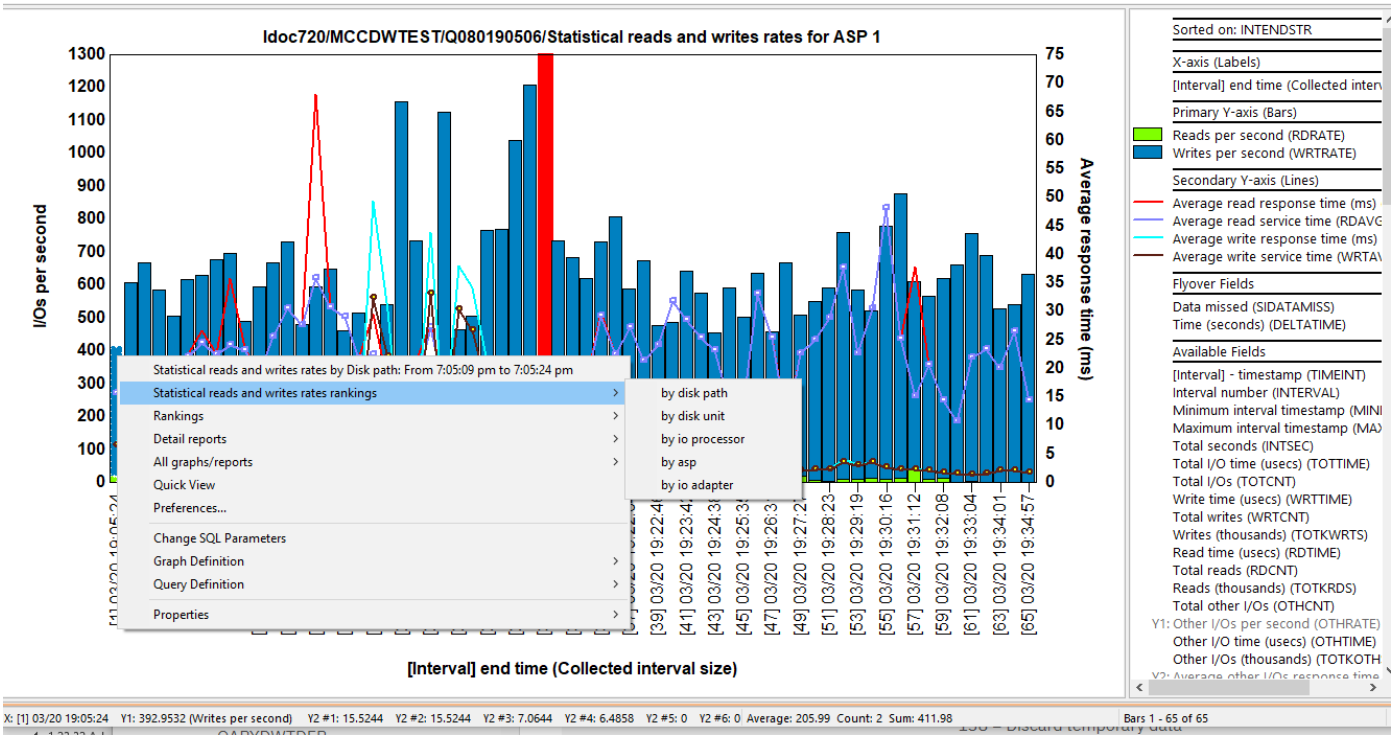

Drilling down into Statistical reads and writes rankings

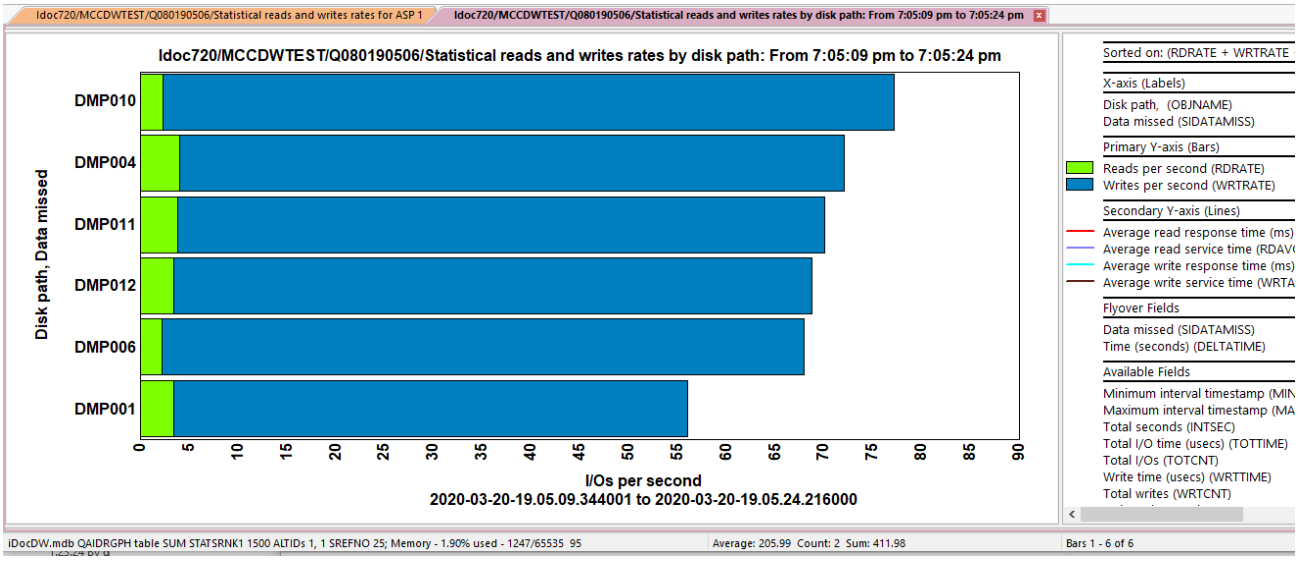

Statistical reads and writes rates by disk path

Rankings graphs by default are viewed with horizontal bars. But the user may press the toggle graph format toolbar button to easily switch these to vertical bars if desired. This will allow for the secondary Y-axis to become visible.

Further drill-downs to selection over time graphs are available using a right-click on the desired grouping:

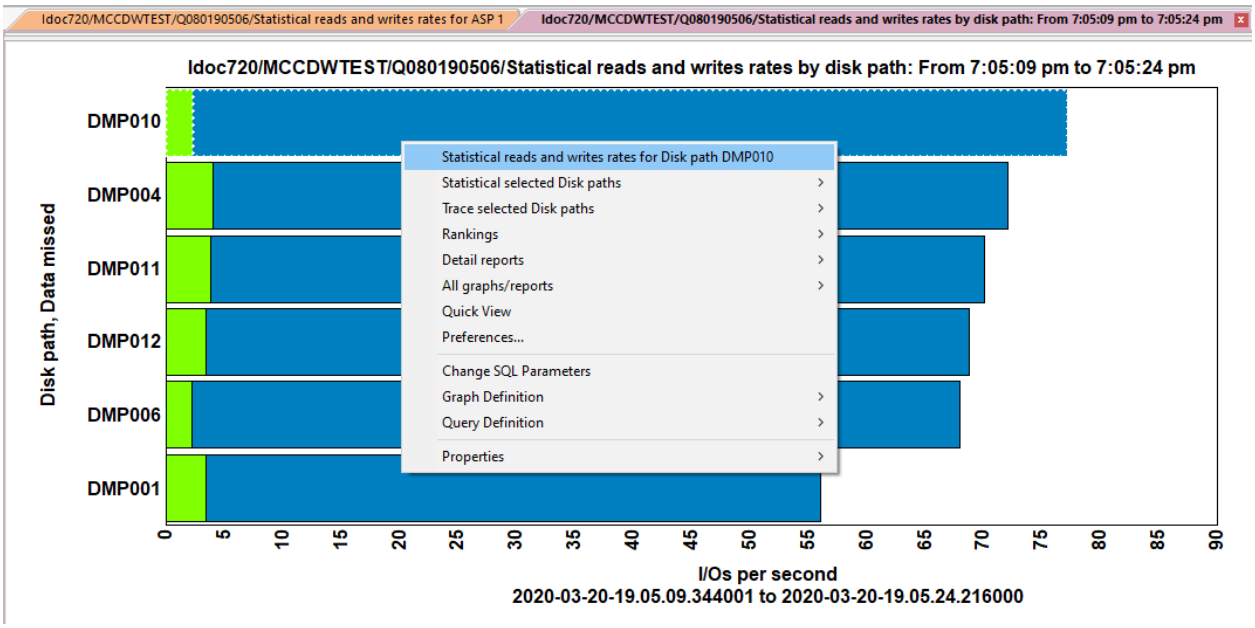

Drilling down to Selection over time graph

13.14 Selection over time graphs

The same set of overview graphs also exist as selection over time graphs. These are only accessible as a drill-down from rankings graphs.

These graphs are almost always identical to the overview graphs, except they are filtered on whatever selection has been made.

An example follows below:

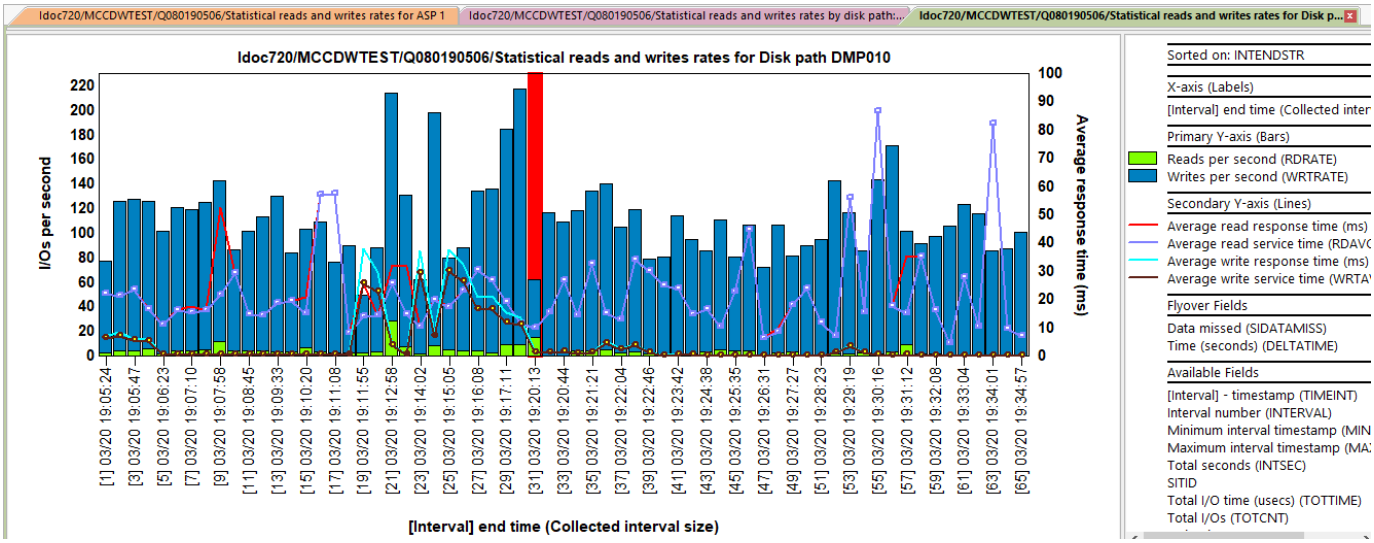

Statistical reads and writes rates for Disk Path DMP010

13.15 Advanced

This folder contains graphs rarely needed and intended for more advanced users or IBM support.

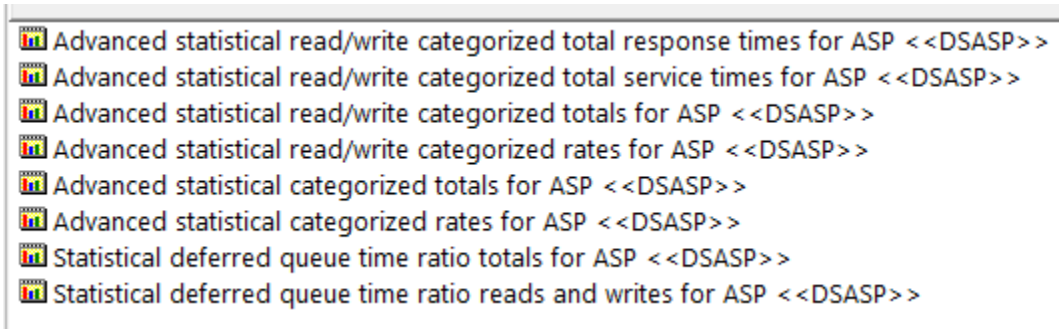

Statistical -> Advanced folder

These are like the graphs previously described except the categorized graphs show 13 buckets per I/O type instead of 6.

An example follows below:

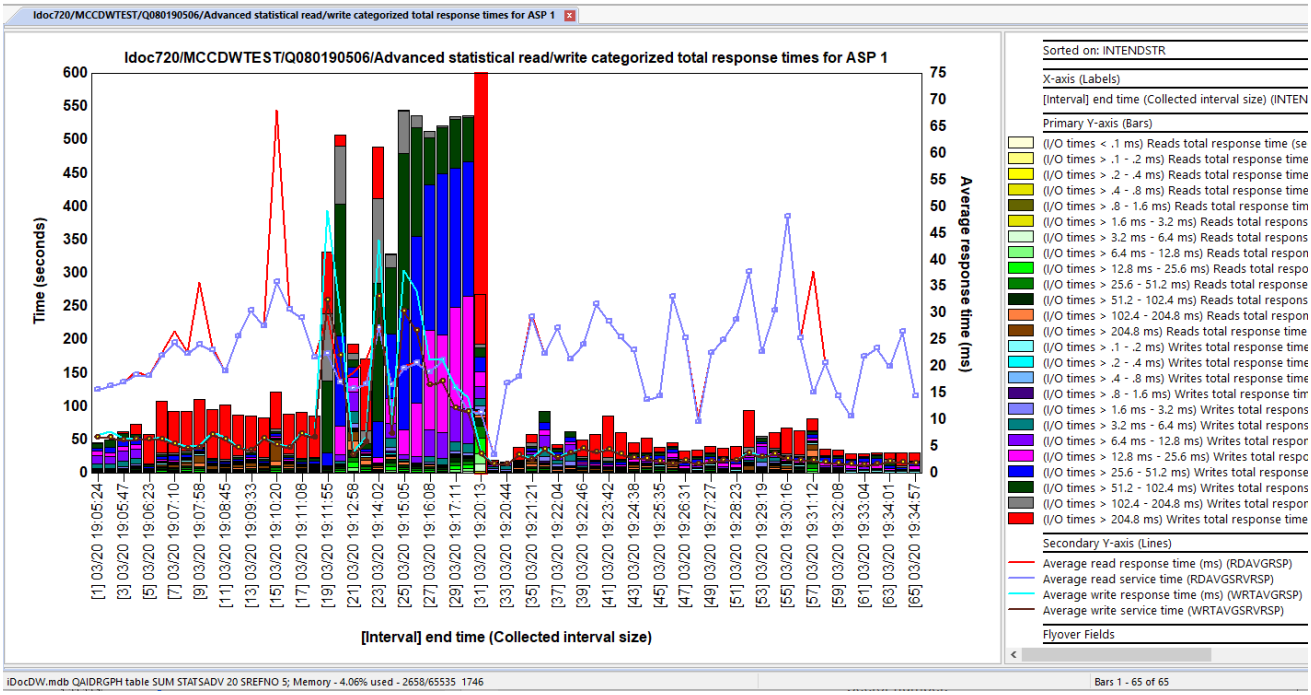

Advanced statistical read/write categorized total response times for ASP 1

13.15.1 Statistical deferred queue time ratio totals for ASP 1

This graph displays the ratio of deferred queue time / service time as well as response and service times on the secondary Y-axis.

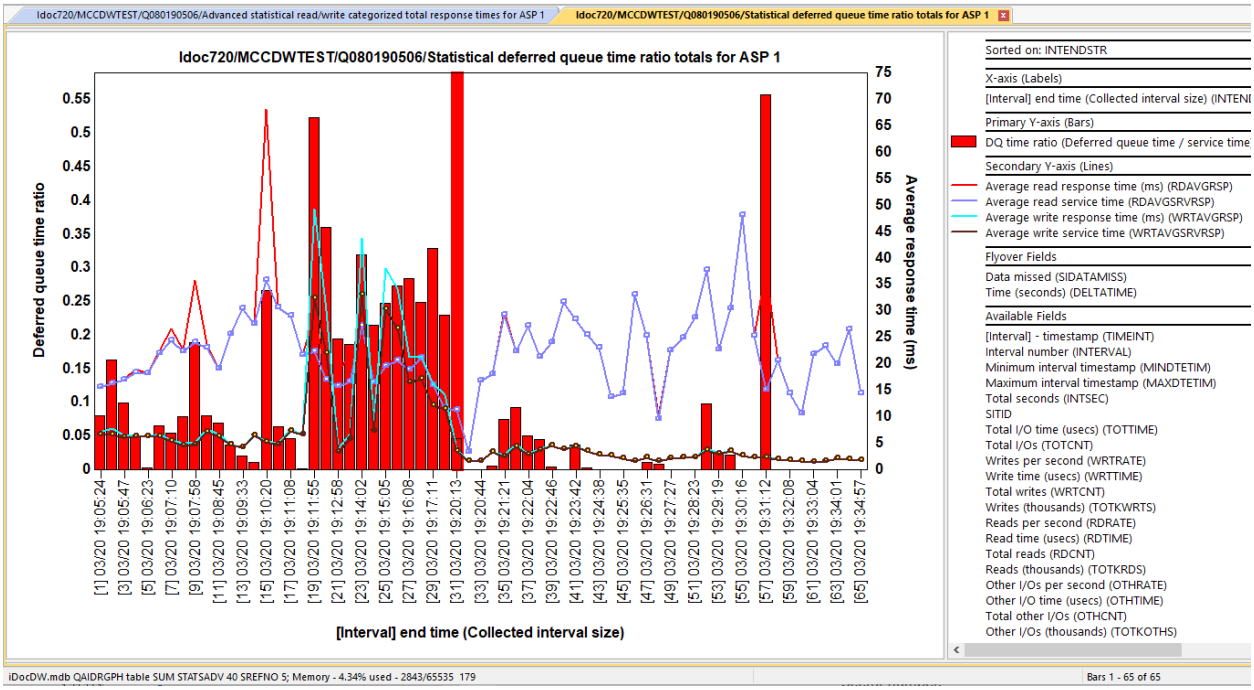

13.15.2 Statistical deferred queue time ratio reads and writes for ASP <<DSASP>>

This is like the previous graph but breaks out the reads, writes and others into separate buckets.

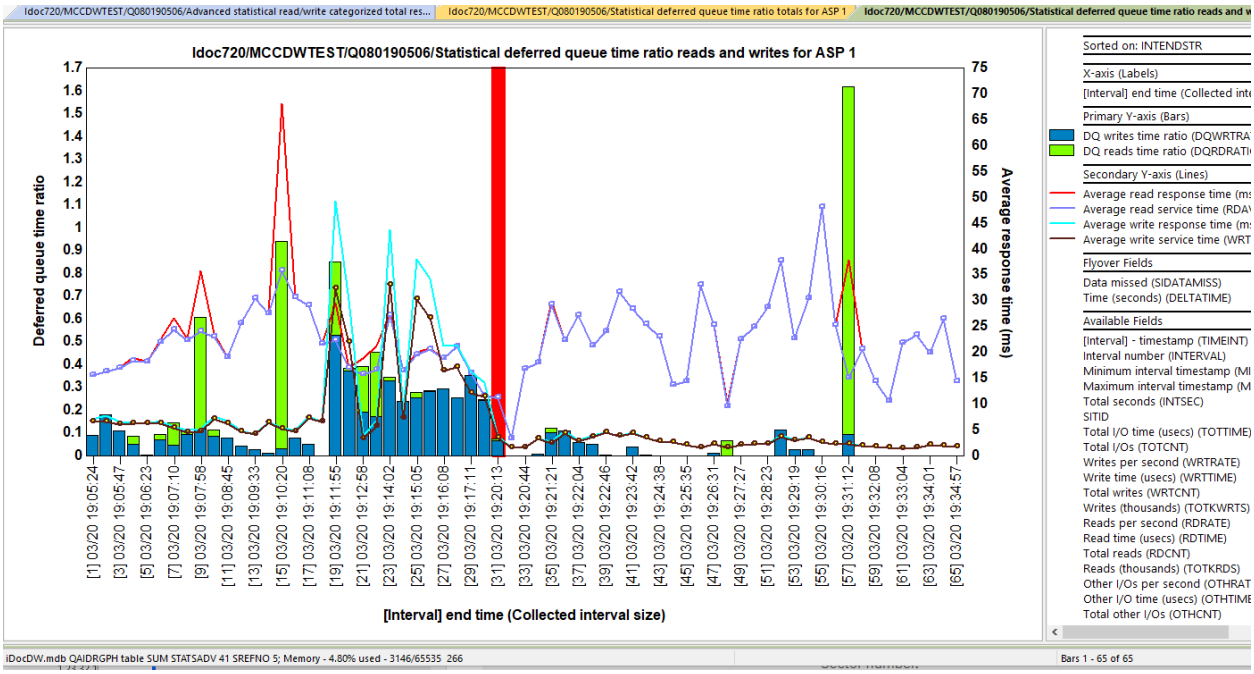

14 Trace graphs

There are 12 trace graphs in Disk Watcher, and these can appear in any of the 3 common formats: Overview, Rankings or Selection over time. Additional graphs are also available under Trace -> Advanced.

These graphs are built using the trace data captured by Disk Watcher in file QAPYDWTRC. Additional SQL tables are built after running the Trace Summary analysis which is required for these graphs to appear.

Typically, users start with the desired Overview graph, then select time periods and drill-down into rankings. From rankings, a user will drill down again into selection over time or use the Detail reports menu for additional options.

Tip: The most used graphs are listed first in this list.

If you need help with definitions of specific metrics, see the previous section or contact iDoctor support.

The Trace folder contains the following:

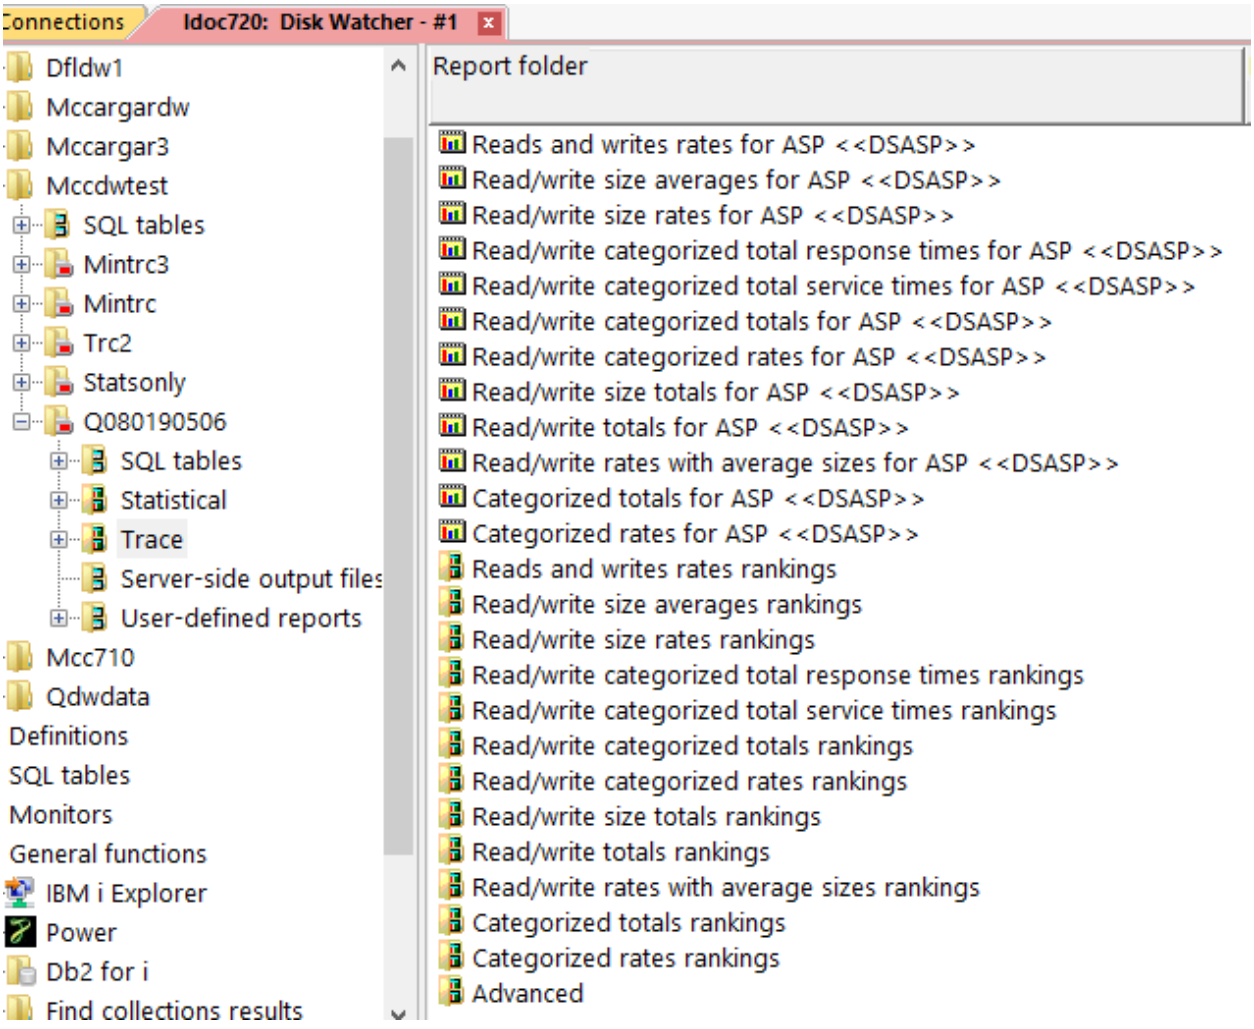

Disk Watcher – Trace folder

14.1 Reads and writes rates for ASP <<DSASP>>

This graph displays **I/O rates per second** with average response times and service times on the secondary Y-axis.

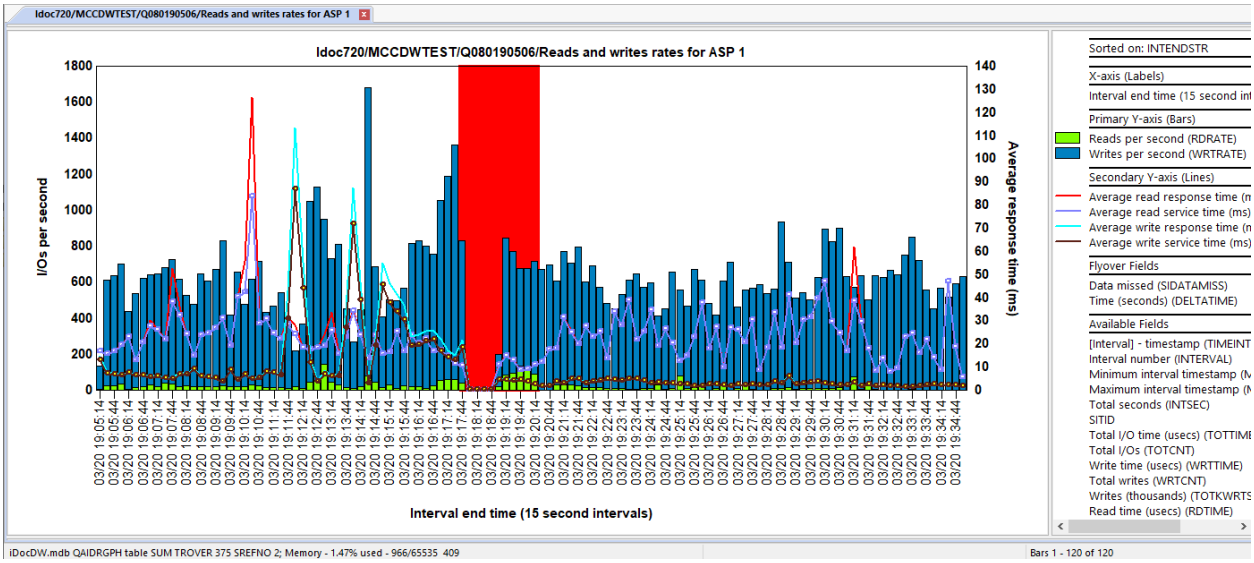

14.2 Read/write size averages for ASP <<DSASP>>

This graph displays **average I/O sizes (in kilobytes)** with average response times and service times on the secondary Y-axis.

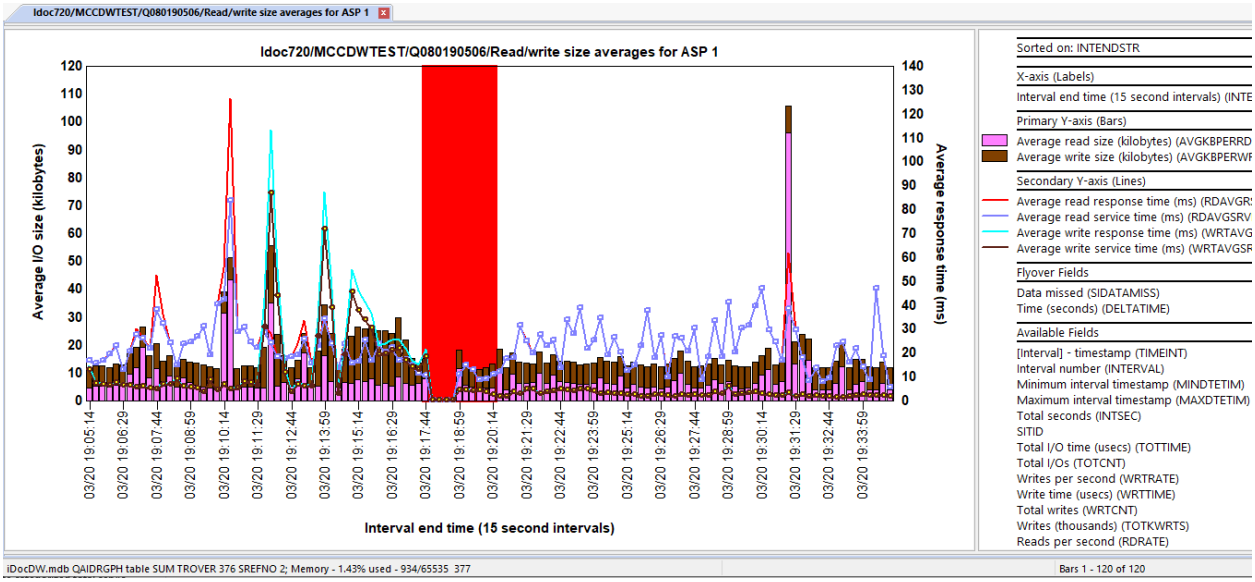

14.3 Read/write size rates for ASP <<DSASP>>

This graph displays the **total I/O size for all operations occurring per second** with average response times and service times on the secondary Y-axis.

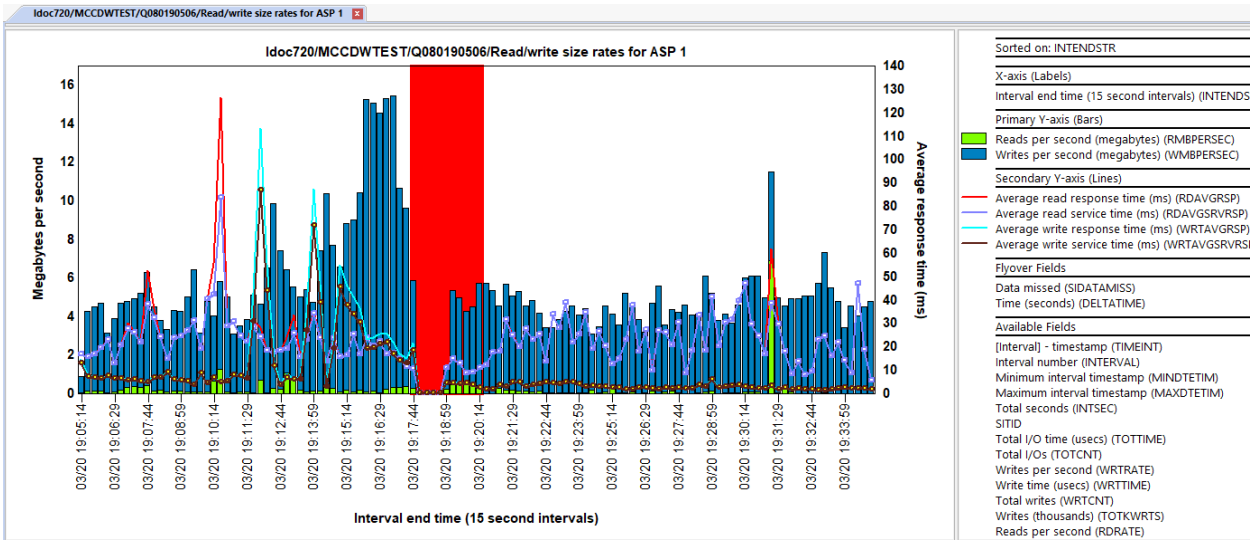

14.4 Read/write categorized total response times for ASP <<DSASP>>

This graph displays the **total I/O times** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

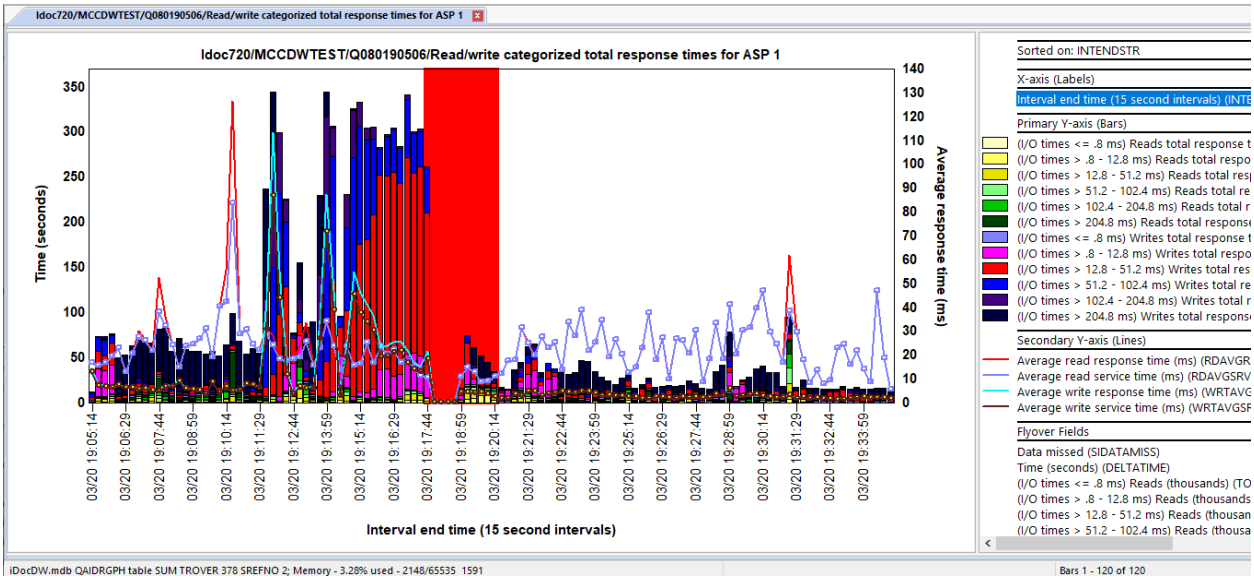

14.5 Read/write categorized total service times for ASP <<DSASP>>

This graph displays the **total service times** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

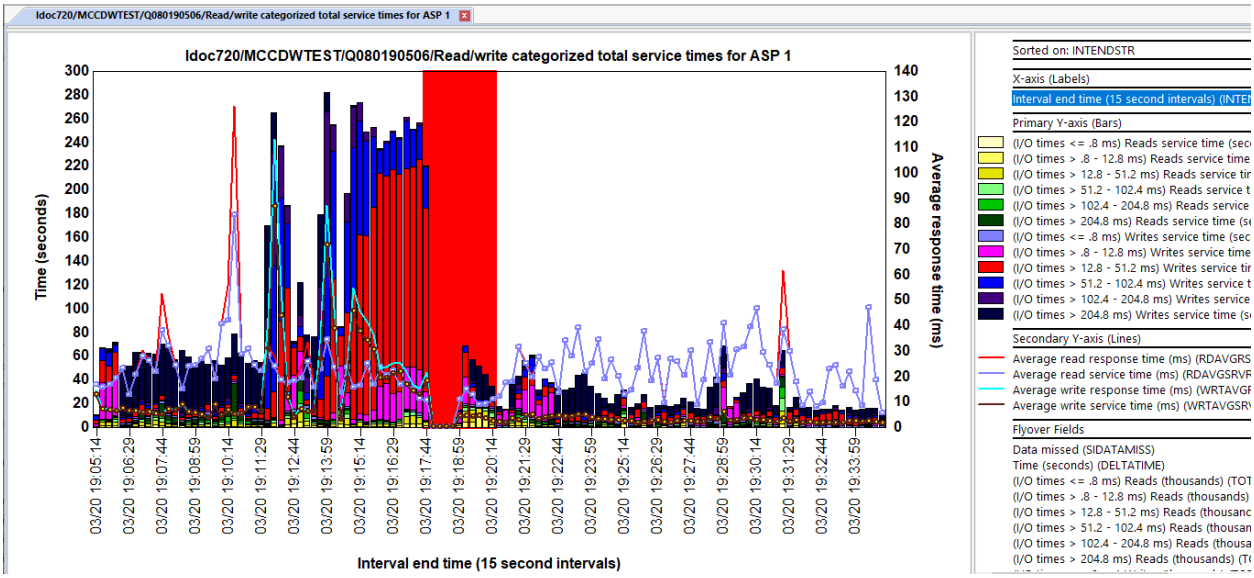

14.6 Read/write categorized totals for ASP <<DSASP>>

This graph displays the **total operations** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

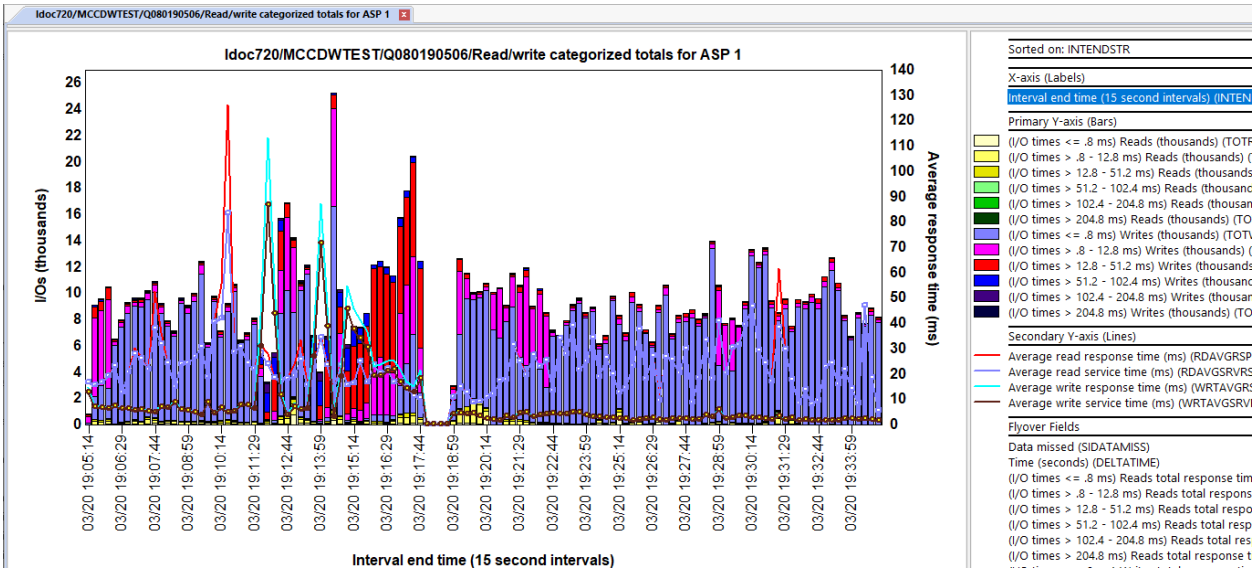

14.7 Read/write categorized rates for ASP <<DSASP>>

This graph displays the **I/Os per second** for each of reads, writes and others with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets for each I/O type with their ranges indicated on the description for each.

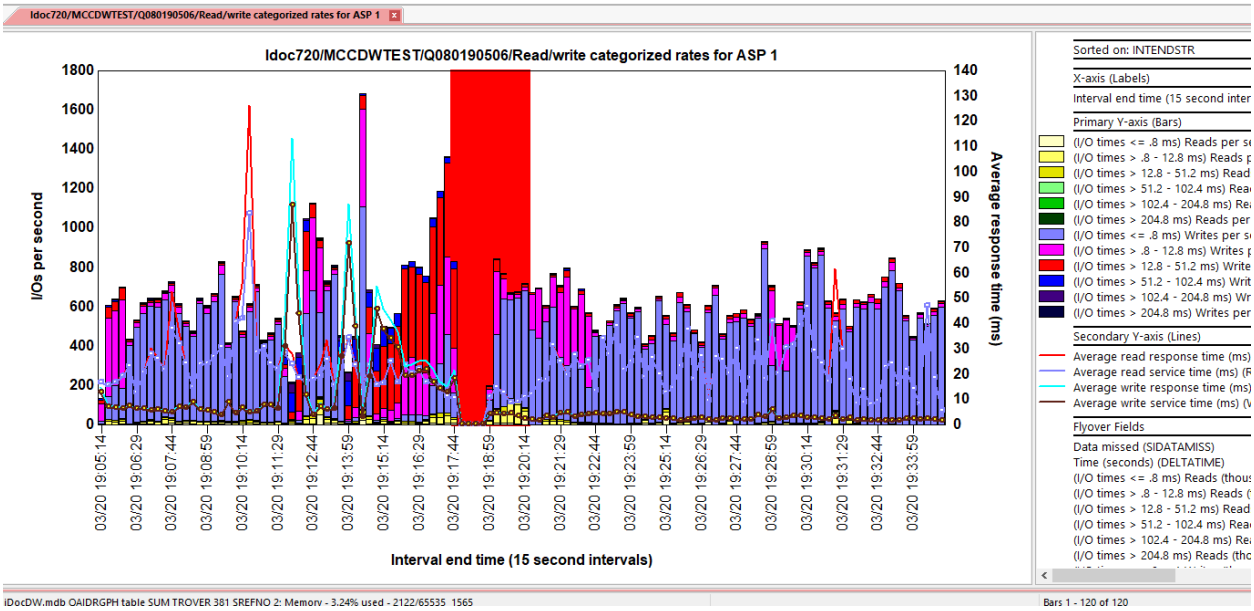

14.8 Read/write size totals for ASP <<DSASP>>

This graph displays the **total I/O size for each of reads, writes and others** with average response times and service times on the secondary Y-axis.

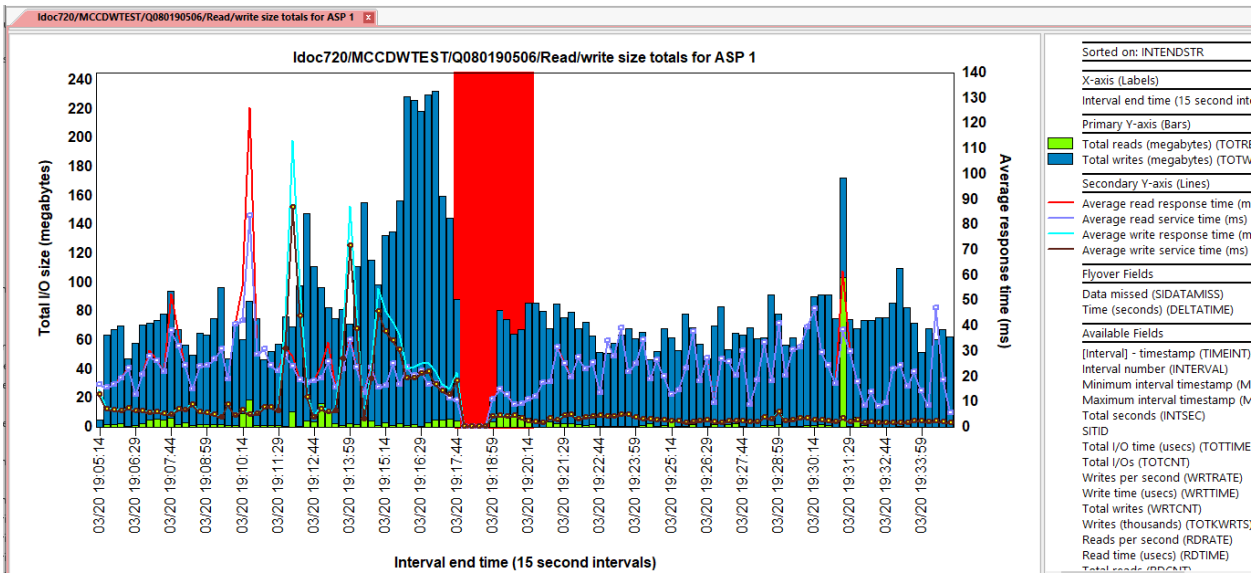

14.9 Read/write totals for ASP <<DSASP>>

This graph displays **I/O totals** with average response times and service times on the secondary Y-axis.

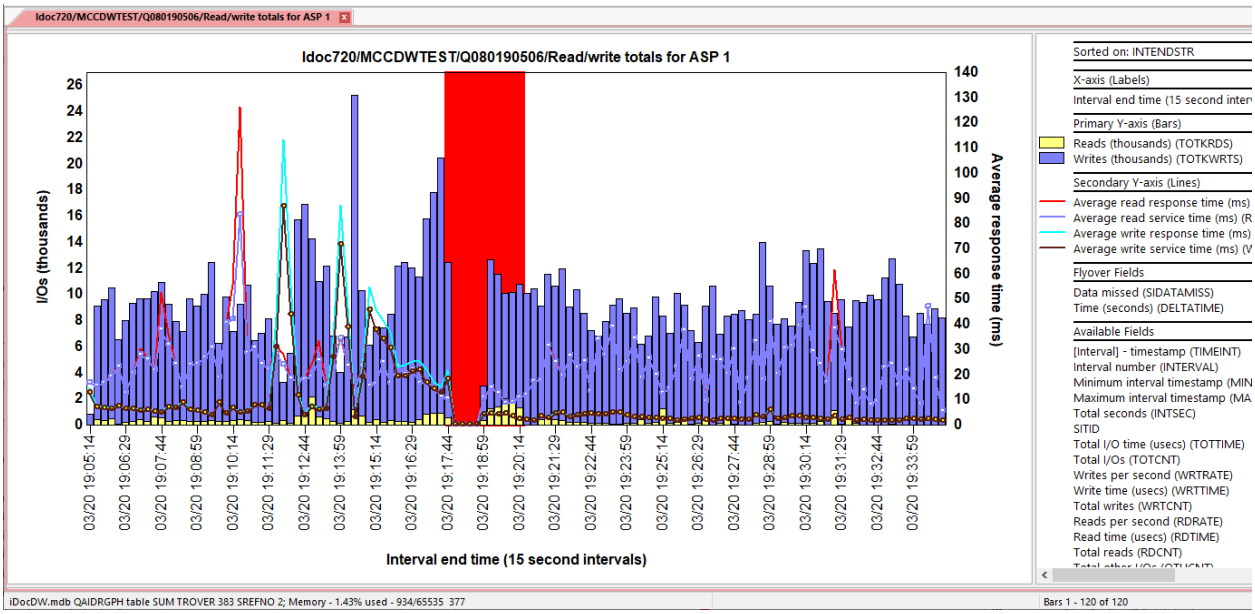

14.10 Read/write rates with average sizes for ASP <<DSASP>>

This graph displays **I/O rates per second** with **average I/O sizes** on the secondary Y-axis.

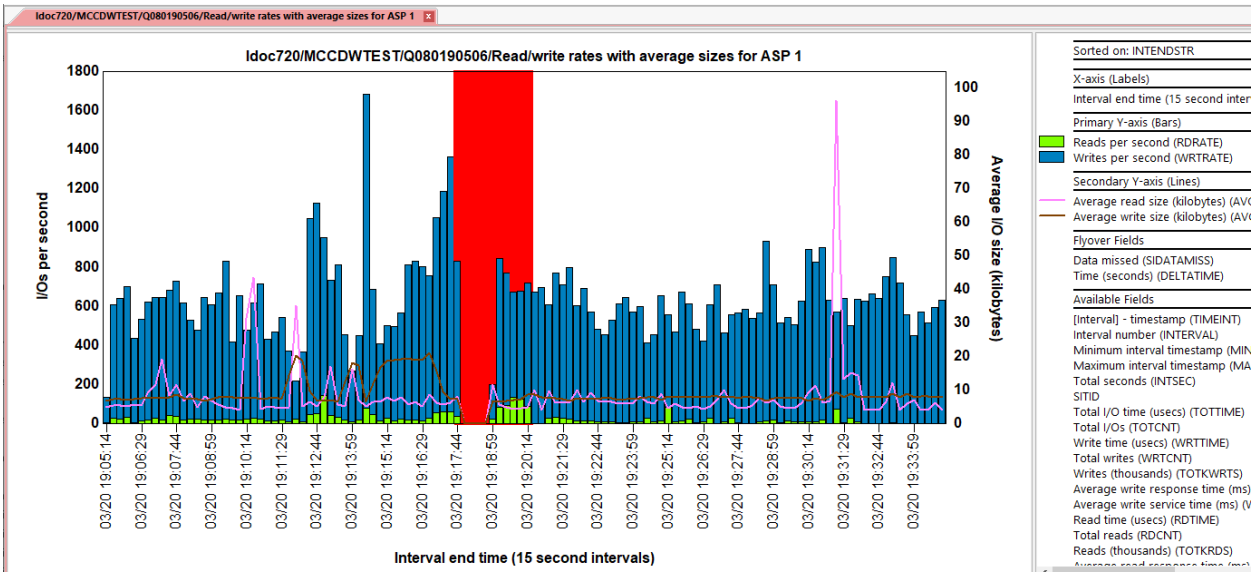

14.11 Categorized totals for ASP <<DSASP>>

This graph combines all 3 I/O types together and shows **total I/Os** with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets

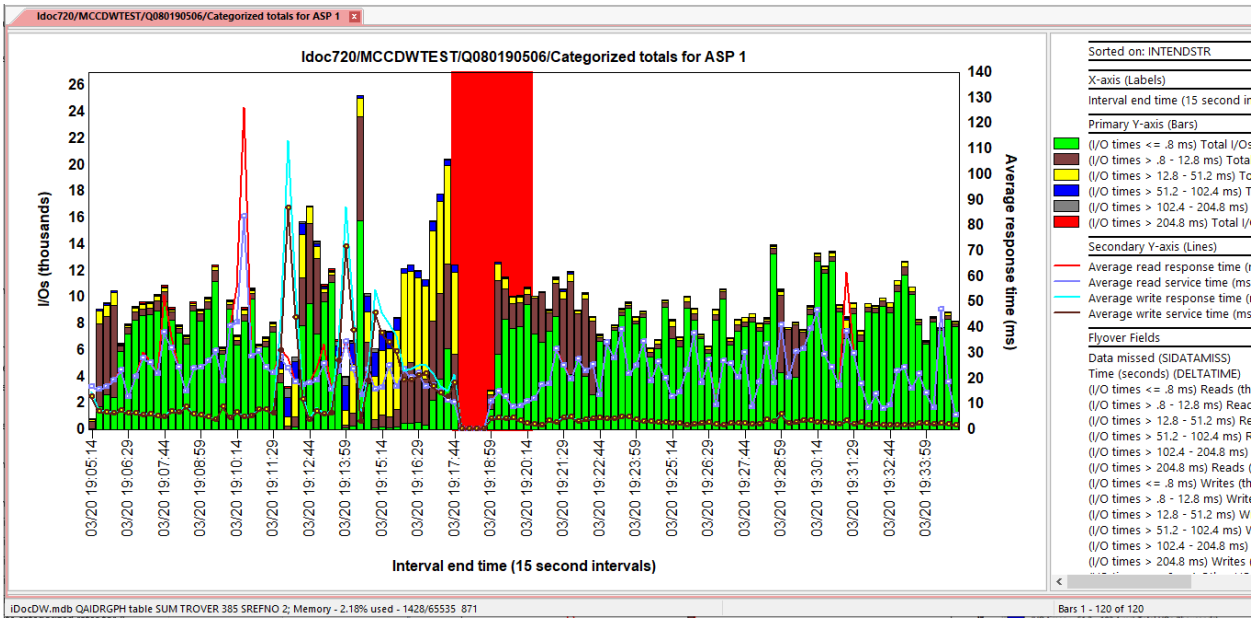

14.12 Categorized rates for ASP <<DSASP>>

This graph combines all 3 I/O types together and shows **I/Os per second** with average response times and service times on the secondary Y-axis. These metrics are grouped into 6 average response time buckets

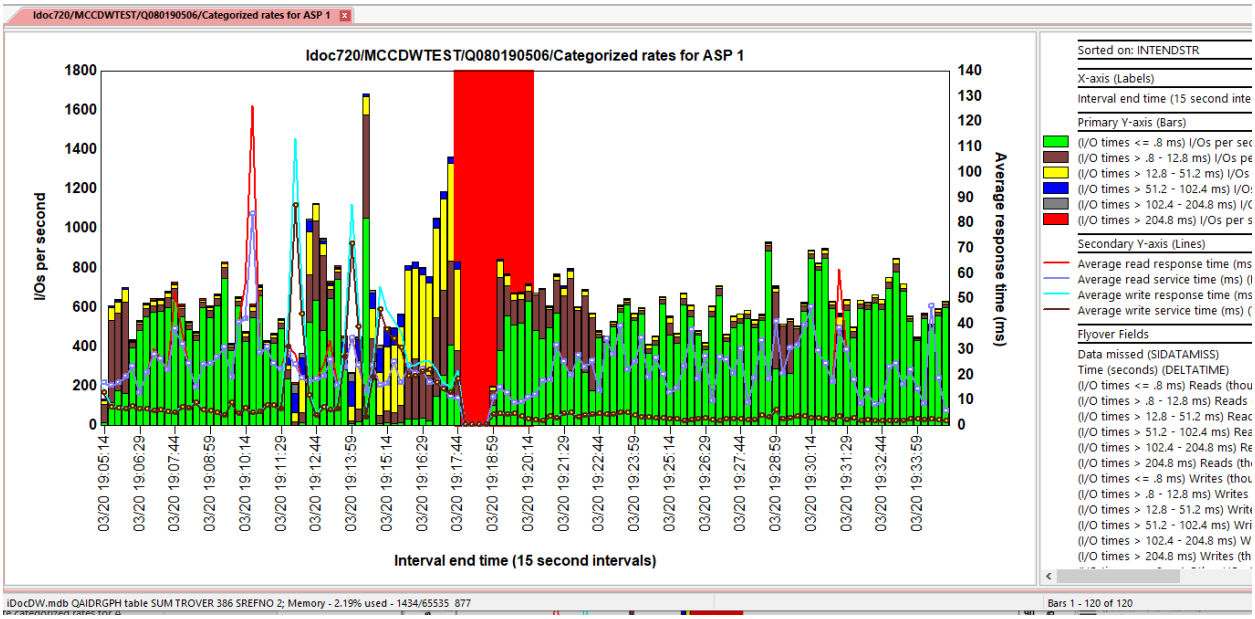

14.13 Rankings graphs

The same set of overview graphs also exist as rankings graphs. For trace data, the following disk groupings are available:

by generic object
by disk path
ASP <<DSASP>> by disk unit
by io processor
by asp
by object
by program
by program - procedure
by thread
by ifs path
by object context
by current user
by io adapter
by disk serial number
by disk type
by lun id-unit
by lun id-unit-path
by lss id
by thread || object
by thread || program
by thread || unit
by thread || unit || object
by generic job
by generic job || object
by generic job || program
by generic job || unit
by generic job || unit || object
by generic job || unit || object || program
by object type
by segment type

Disk groupings

Note: The options here will vary based on what types of additional data were collected.

Access to these can be done either using the appropriate Rankings graph options at the collection-level or via a drill down from an Overview graph.

The typical method to drill down into rankings is to select a time period and right-click as follows:

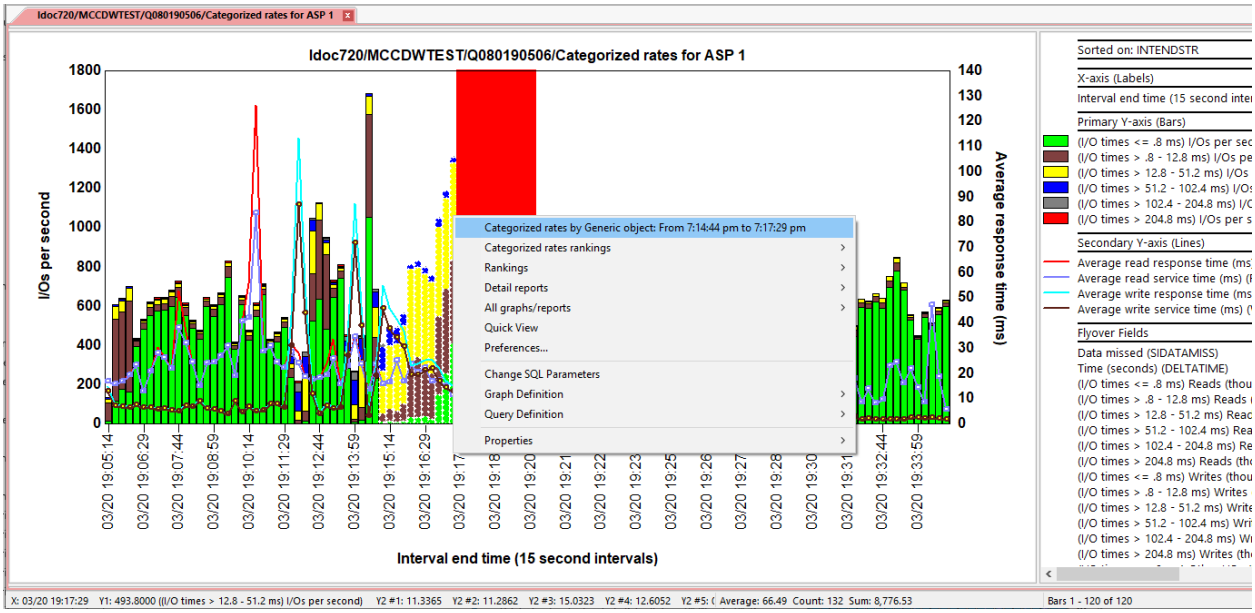

Drilling down into Categorized rates by generic object

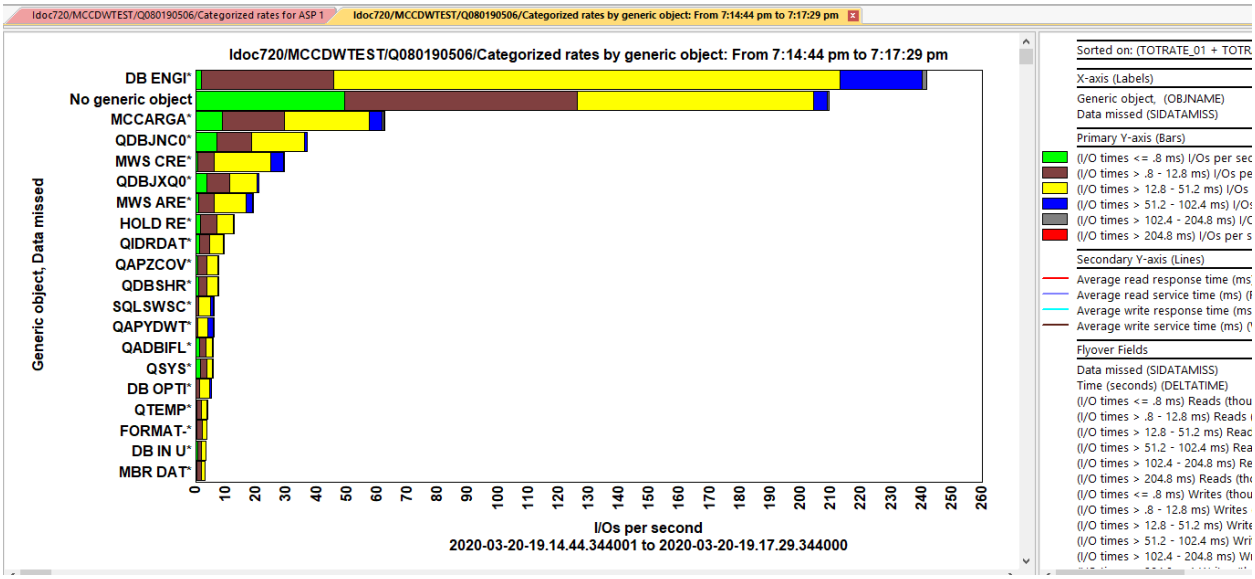

Categorized rates by generic object

Rankings graphs by default are viewed with horizontal bars. But the user may press the toggle graph format toolbar button to easily switch these to vertical bars if desired. This will allow for the secondary Y-axis to become visible.

Further drill-downs to selection over time graphs are available using a right-click on the desired grouping:

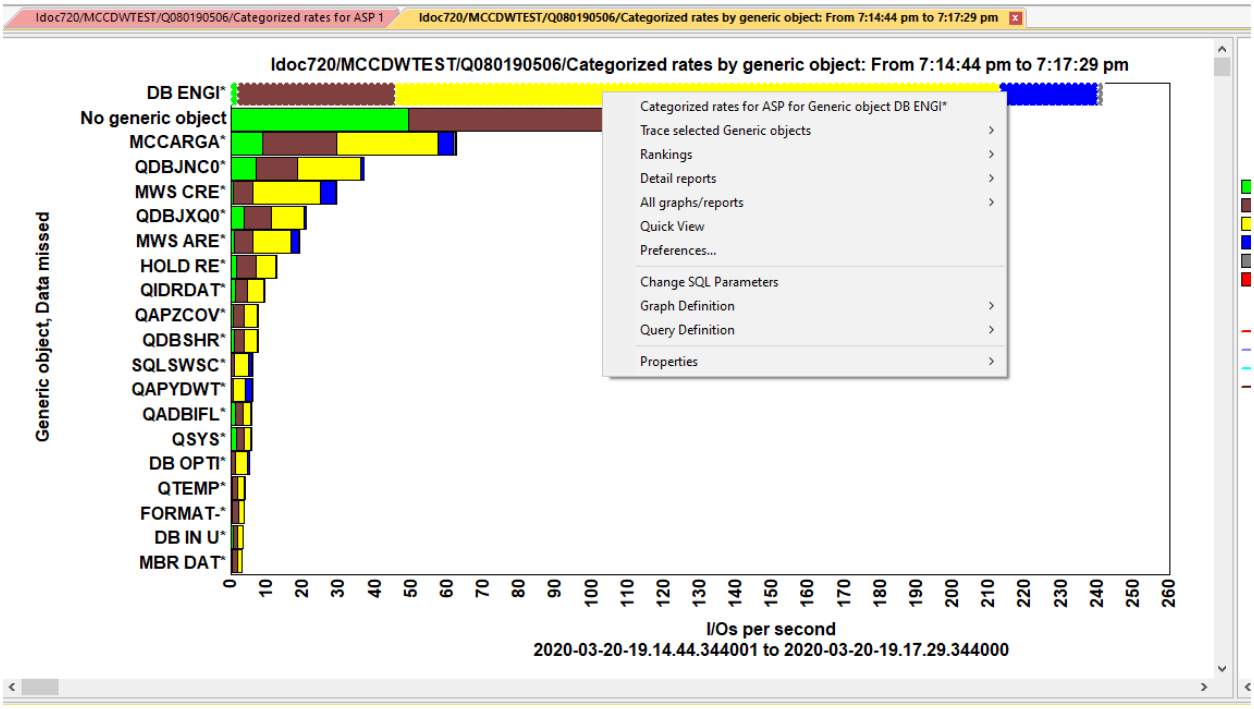

Drilling down to Selection over time graph

14.14 Selection over time graphs

The same set of overview graphs also exist as selection over time graphs. These are only accessible as a drill-down from rankings graphs.

These graphs are almost always identical to the overview graphs, except they are filtered on whatever selection has been made.

An example follows below:

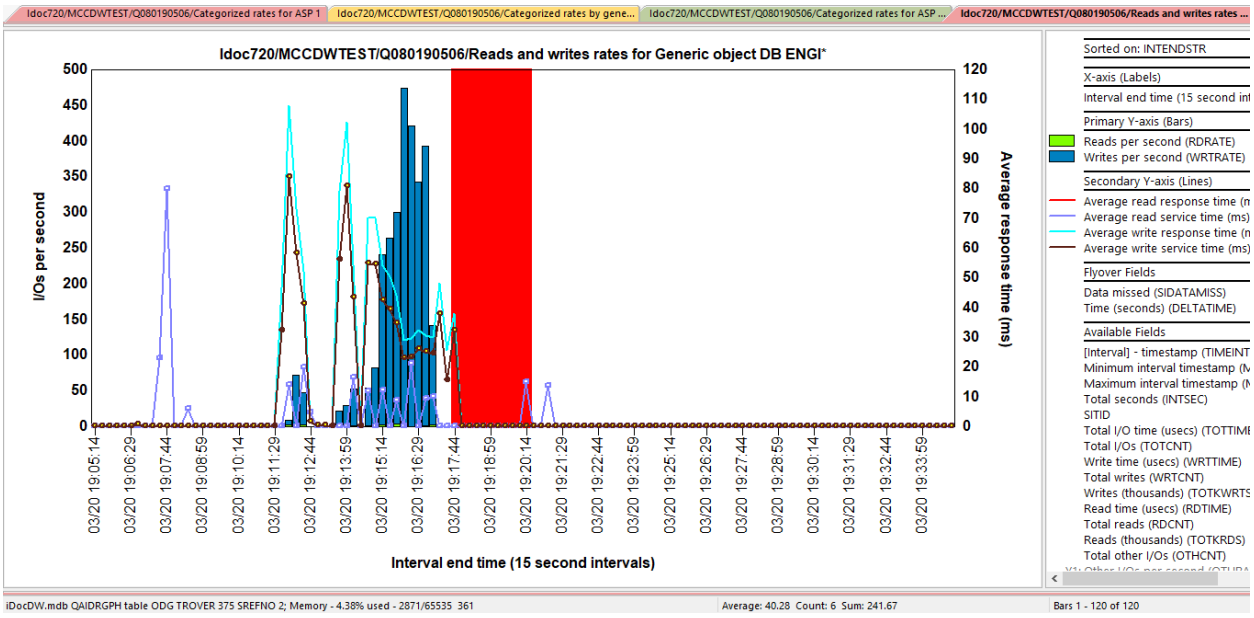

*Reads and writes rates for Generic object DB ENGI**

14.15 Advanced

This folder contains graphs rarely needed and intended for more advanced users or IBM support.

- Advanced read/write categorized total response times for ASP <<DSASP>>
- Advanced read/write categorized total service times for ASP <<DSASP>>
- Advanced read/write categorized totals for ASP <<DSASP>>
- Advanced read/write categorized rates for ASP <<DSASP>>
- Advanced categorized totals for ASP <<DSASP>>
- Advanced categorized rates for ASP <<DSASP>>
- Deferred queue time ratio totals for ASP <<DSASP>>
- Deferred queue time ratio reads and writes for ASP <<DSASP>>

Trace -> Advanced folder

These are like the graphs previously described except the categorized graphs show 13 buckets per I/O instead of 6.

An example follows below:

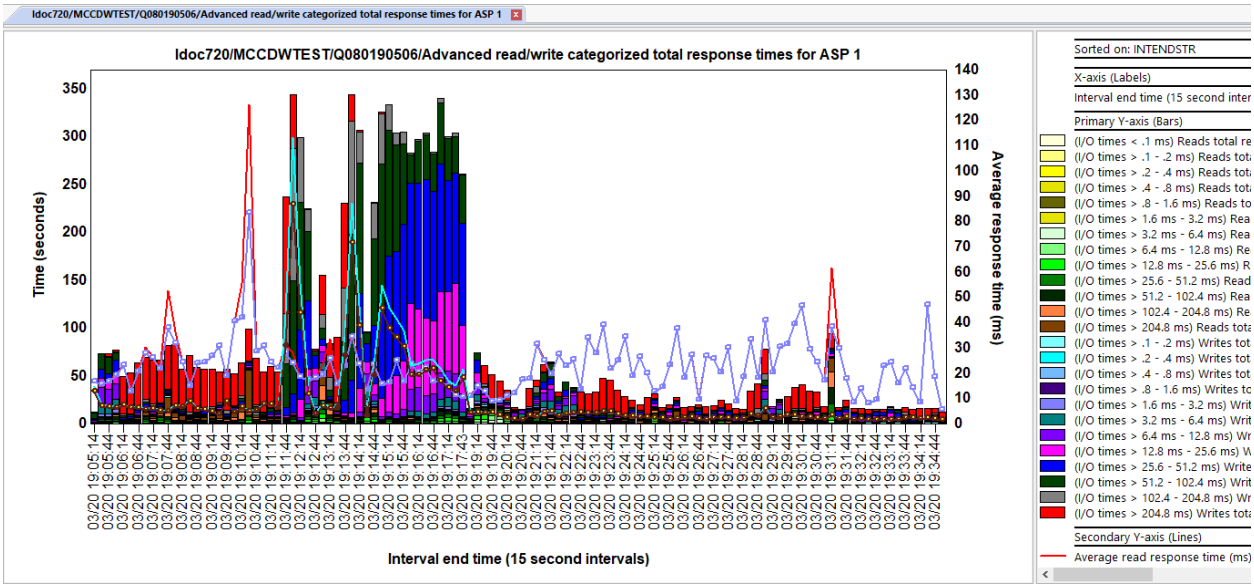

Advanced read/write categorized total response times for ASP 1

15 Detail reports

These reports are drill-down options from either the statistical or trace graphs. They show tables and report more detailed information about the disk I/Os occurring for the time period or selected job/object of interest.

An example of the Detail reports menu follows:

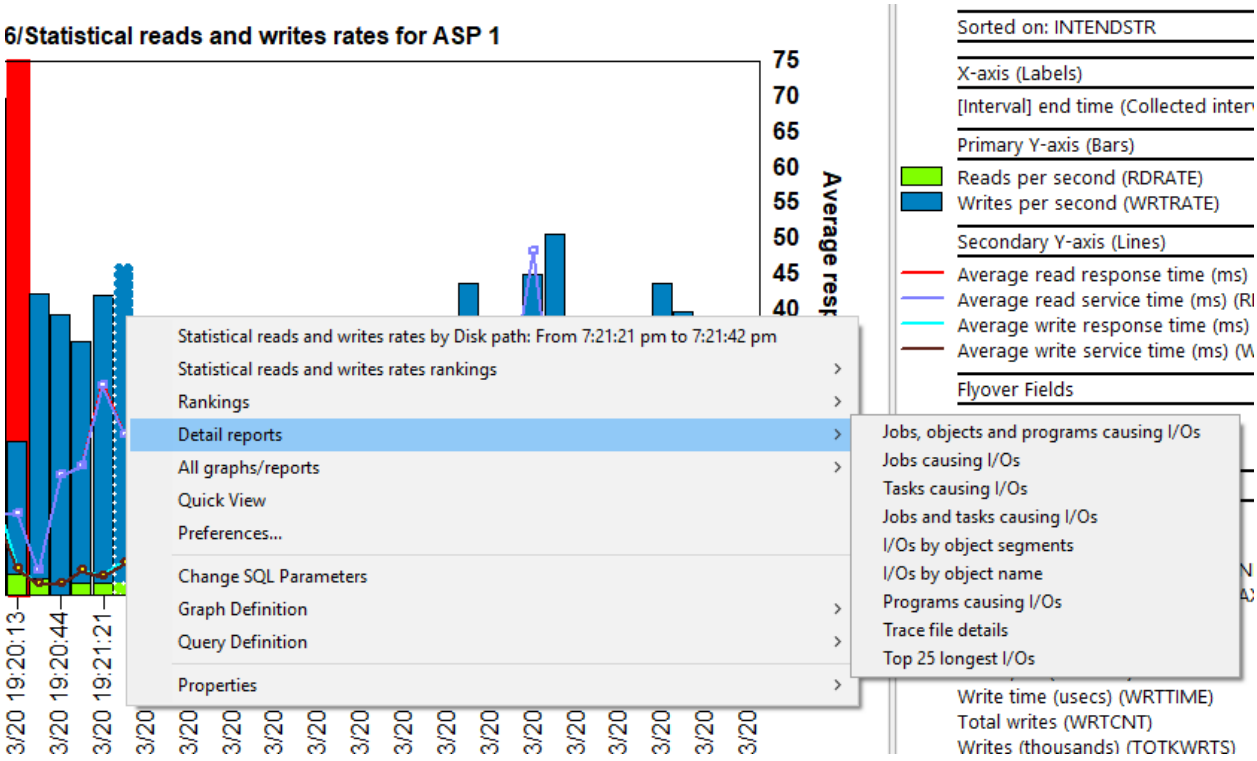

An example follows:

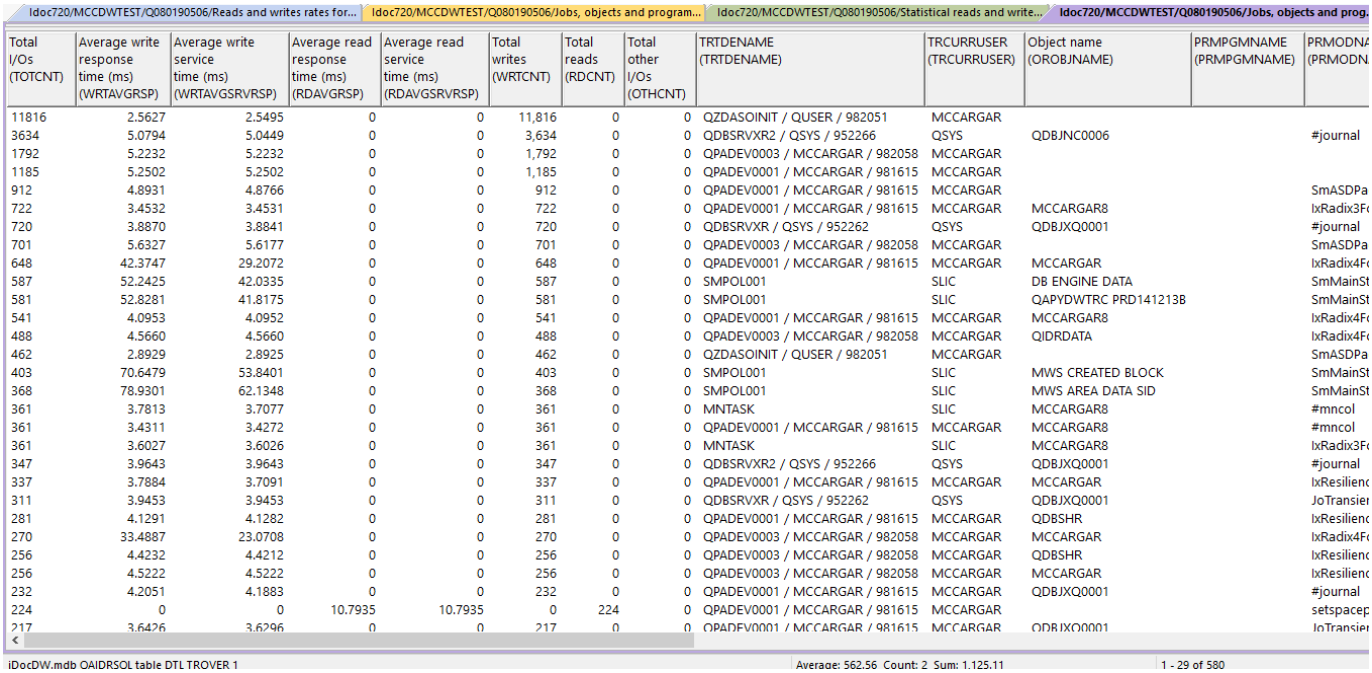

Idoc720/MCCDWTEST/Q080190506/Reads and writes rates for... | Idoc720/MCCDWTEST/Q080190506/Jobs, objects and program... | Idoc720/MCCDWTEST/Q080190506/Statistical reads and write... | Idoc720/MCCDWTEST/Q080190506/Jobs, objects and prog...

Total I/Os (TOTCNT)	Average write response time (ms) (WRTAVGRSP)	Average write service time (ms) (WRTAVGSRVRSP)	Average read response time (ms) (RDAVGRSP)	Average read service time (ms) (RDAVGSRVRSP)	Total writes (WRTCNT)	Total reads (RDCNT)	Total other I/Os (OTHCNT)	TRTDENAME (TRTDENAME)	TRCURRUSER (TRCURRUSER)	Object name (OROBJNAME)	PRMPGMNAME (PRMPGMNAME)	PRMODNA (PRMODN
11816	2.5627	2.5495	0	0	11,816	0	0	QZDASOINIT / QUSER / 982051	MCCARGAR			
3634	5.0794	5.0449	0	0	3,634	0	0	QDBSRVXR2 / QSYS / 952266	QSYS	QDBJNC0006		#journal
1792	5.2232	5.2232	0	0	1,792	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR			
1185	5.2502	5.2502	0	0	1,185	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR			
912	4.8931	4.8766	0	0	912	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR			SmASDPa
722	3.4532	3.4531	0	0	722	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	MCCARGAR8		lxRadix3F
720	3.8870	3.8841	0	0	720	0	0	QDBSRVXR / QSYS / 952262	QSYS	QDBJXQ0001		#journal
701	5.6327	5.6177	0	0	701	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR			SmASDPa
648	42.3747	29.2072	0	0	648	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	MCCARGAR		lxRadix4F
587	52.2425	42.0335	0	0	587	0	0	SMPOL001	SLIC	DB ENGINE DATA		SmMainSt
581	52.8281	41.8175	0	0	581	0	0	SMPOL001	SLIC	QAPYDWTRC PRD141213B		SmMainSt
541	4.0953	4.0952	0	0	541	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	MCCARGAR8		lxRadix4F
488	4.5660	4.5660	0	0	488	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR	QIDRDATA		lxRadix4F
462	2.8929	2.8925	0	0	462	0	0	QZDASOINIT / QUSER / 982051	MCCARGAR			SmASDPa
403	70.6479	53.8401	0	0	403	0	0	SMPOL001	SLIC	MWS CREATED BLOCK		SmMainSt
368	78.9301	62.1348	0	0	368	0	0	SMPOL001	SLIC	MWS AREA DATA SID		SmMainSt
361	3.7813	3.7077	0	0	361	0	0	MNTASK	SLIC	MCCARGAR8		#mncol
361	3.4311	3.4272	0	0	361	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	MCCARGAR8		#mncol
361	3.6027	3.6026	0	0	361	0	0	MNTASK	SLIC	MCCARGAR8		lxRadix3F
347	3.9643	3.9643	0	0	347	0	0	QDBSRVXR2 / QSYS / 952266	QSYS	QDBJXQ0001		#journal
337	3.7884	3.7091	0	0	337	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	MCCARGAR		lxResilienc
311	3.9453	3.9453	0	0	311	0	0	QDBSRVXR / QSYS / 952262	QSYS	QDBJXQ0001		JoTransier
281	4.1291	4.1282	0	0	281	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	QDBSHR		lxResilienc
270	33.4887	23.0708	0	0	270	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR	MCCARGAR		lxRadix4F
256	4.4232	4.4212	0	0	256	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR	QDBSHR		lxResilienc
256	4.5222	4.5222	0	0	256	0	0	QPADEV0003 / MCCARGAR / 982058	MCCARGAR	MCCARGAR		lxResilienc
232	4.2051	4.1883	0	0	232	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	QDBJXQ0001		#journal
224	0	0	10.7935	10.7935	0	224	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR			setspacep
217	3.6426	3.6296	0	0	217	0	0	QPADEV0001 / MCCARGAR / 981615	MCCARGAR	QDBJXQ0001		JoTransier

iDocDW.mdb QAIDRSQL table DTL TROVER 1 | Average: 562.56 Count: 2 Sum: 1,125.11 | 1 - 29 of 580

Jobs, objects and programs causing I/Os

16 Disk Watcher Commands

This section describes green screen commands and functionality related to Disk Watcher in the OS and included with iDoctor.

16.1 IBM i Disk Watcher Commands

At 6.1 and higher the Disk Watcher commands were added to IBM i. The following commands are part of the Disk Watcher support included in library QSYS:

16.1.1 ADDDWDFN

Adds a Disk Watcher definition to the system.
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/adddwdfn.htm

16.1.2 STRDW

Starts a Disk Watcher collection
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/strdw.htm

16.1.3 RMVDWDFN

Deletes a Disk Watcher definition.
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/rmvdwdfn.htm

16.1.4 ENDDW

Ends a Disk Watcher collection.
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/enddw.htm

16.1.5 CPYPFRCOL

Copy a collection to another location.
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/cpypfrcol.htm

16.1.6 CVTPFRCOL

This command will convert a collection created at a previous release to the current release.
At 6.1 this command only supports Collection Services collections. At 7.1 and higher it also supports Disk Watcher, Job Watcher and PEX.

Note: Using this command is NOT recommended if planning to use the data with iDoctor. iDoctor will handle your data regardless of the release it was created on (7.1 and higher)
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/cvtpfrcol.htm

16.1.7 DLTPFRCOL

Deletes one or more collections (either PEX, CS, DW, JW, etc)
Note: This will NOT delete iDoctor created SQL tables, so it not recommended. Use the QIDRWCH/DLTDWCOL instead.

http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/dltpfrcol.htm

16.1.8 SAVPFRCOL

Saves one or more collections to a SAVF (either PEX, CS, DW, JW, etc)
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/savpfrcol.htm

16.1.9 RSTPFRCOL

Restore one or more collections from a SAVF (either PEX, CS, DW, JW, etc)
http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/cl/rstpfrcol.htm

16.2 iDoctor Disk Watcher Commands

The following commands are included with Disk Watcher in library QIDRWCH.

16.2.1 CPYDWCOL

This command is used to copy a Disk Watcher collection. This action will copy every member matching the collection name from the "from" library to the "to" library. If desired the collection can be renamed by setting the TOCOL parameter with a new name.

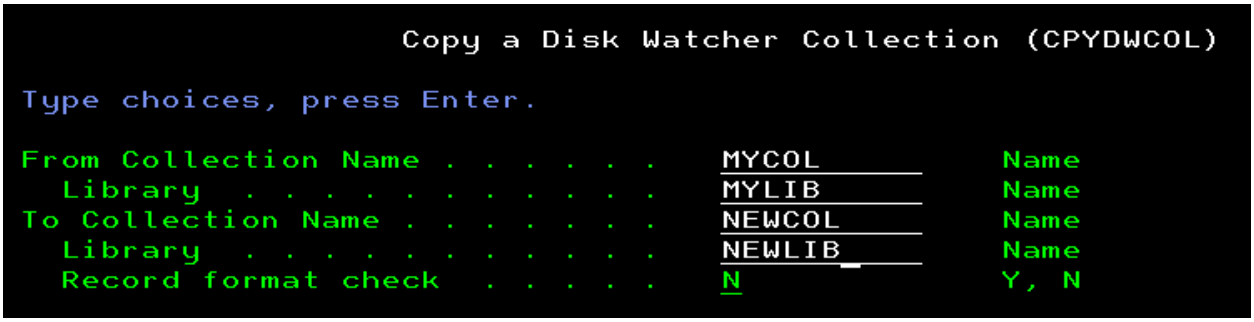

Note: CPYPFRCOL can also perform this action, but this command is still provided to avoid possible problems where CPYPFRCOL will not allow a collection to be copied based on VRM differences of the database files.

Note: This command will NOT copy the iDoctor created SQL Tables produced by iDoctor analyses.

Keyword	Description	Choices
FROMCOL	From collection name	Name
FROMLIB	From library name	Name
TOCOL	To collection name	Name
TOLIB	To library name	Name
FMTOPT	Record format check This option allows the Disk Watcher collection to be copied even if the files in the 'To' library do not match the files in the 'From' library. - If the value is 'N' then CPYF will use FMTOPT(*NOCHK) - If the value is 'Y' then CPYF will use FMTOPT(*MAP *DROP)	N, Y

16.2.2 DLTDWCOL

This command is used to delete a Disk Watcher collection from a user's library on a system. This action will remove the member matching the collection name from every QAPYDW* file found in the library specified.

This command also removes all iDoctor created SQL Tables associated with this collection.

Keyword	Description	Choices
WCHNAME	Collection name	Name
LIB	library name	Name

16.2.3 DLTDWMON

This command is used to delete a Disk Watcher monitor and all collections it contains on the current system. If the monitor is still running, the monitor will first be ended using the ENDDWMON command.

The record in file QUSRSYS/QAIDRDWM2 that identifies the existence of this monitor will also be removed by this command.

Keyword	Description	Choices
MONITOR	Monitor name	Name

16.2.4 ENDDWMON

This command is used to end the job running a Disk Watcher monitor and any jobs currently running collections within the monitor.

The **Ending option** parameter can be used to delay ending the monitor until the currently collection's normal duration completes.

Keyword	Description	Choices
MONITOR	Monitor name	Character value
COLLIB	Ending option This parameter determines if the monitor and any active collection jobs should end immediately or after the currently running collection completes. Use the *DELAYED option to indicate the monitor should remain active and then end once the current collection completes.	*IMMED, *DELAYED
DFNNAME	Run default analyses This option, if set to *YES, will automatically run the Trace summary analysis on any collections that contain trace data after each collection completes. In some cases, this can be resource intensive.	*YES, *NO

16.2.5 HLDDWMON

This command is used to hold a Disk Watcher monitor. A held monitor will not delete any old collections or create any new collections until it is released using the RLSDWMON command.

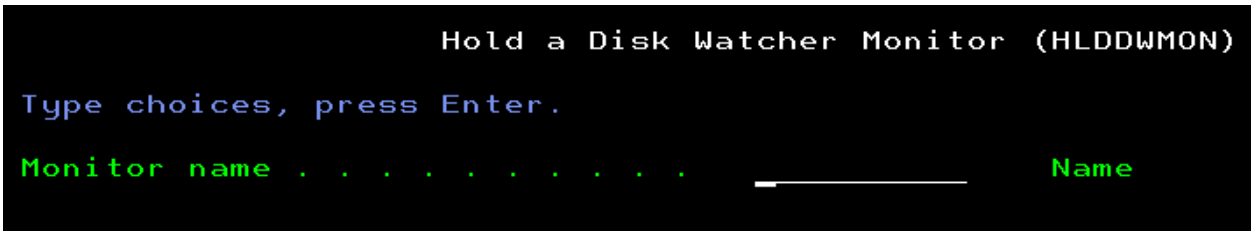

Keyword	Description	Choices
MONITOR	Monitor name	Name

16.2.6 RLSDWMON

This command is used to release a monitor that is currently in a held state. Once released the monitor will continue to create new collections and delete old collections normally.

Keyword	Description	Choices
MONITOR	Monitor name If the monitor specified is not already held, the command will do nothing.	Name

16.2.7 STRDWMON

This command is used to start or restart a Disk Watcher monitor. Use DLTDWMON to delete an existing monitor.

Note: This command should be submitted to batch using the SBMJOB command and not be ran interactively.

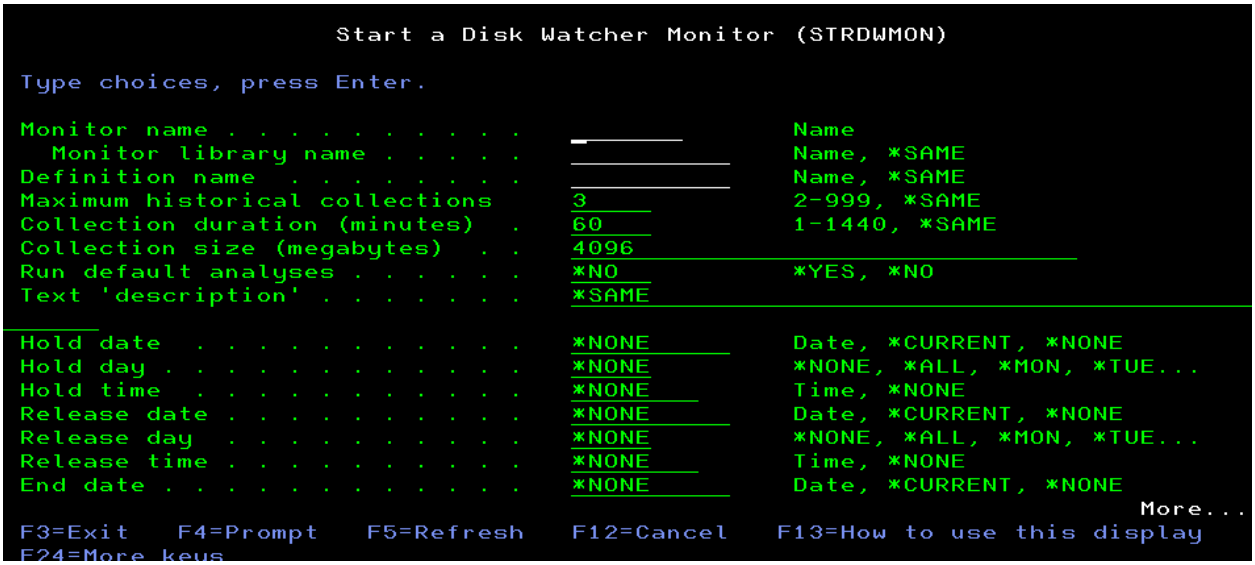

A monitor is a set of collections that continuously collect data over a system overwriting the oldest collection when it creates a new collection.

Monitors are built from a Disk Watcher definition which are stored in file QUSRSYS/QAPYDWDFN. A definition is a ADDDWDFN command string that indicates the parameters the monitor should use in its collections. Definitions can be created using the iDoctor client or by using the ADDDWDFN command.

The maximum historical collections parameter (COLNS) determines how many collections should be saved at one time.

A record in file QUSRSYS/QAIDRDWM2 that identifies the existence and status of the monitor is created and updated by this command.

Keyword	Description	Choices
MONITOR	Monitor name	Name
COLLIB	Monitor library name	Name, *SAME
DFNNAME	Definition name	Name, *SAME
COLNS	Maximum historical collections This is the maximum number of collections that should exist in the library. It does not count/include the currently running collection.	2-999, *SAME
STRGAP	Collection duration (in minutes) This value is how much time should elapse between each collection being submitted.	1-1440, *SAME
MAXSIZE	Collection maximum size (megabytes) This value indicates the maximum amount of data that each collection should contain. **Note:** If this value is reached before the collection duration elapses, then NO data will be captured until the next collection in the monitor starts.	Character value, (10 digits or less)
RUNSUM	Run default analyses This option, if set to *YES, will automatically run the Trace summary analysis on any collections that contain trace data after each collection completes. In some cases, this can be resource intensive.	*YES, *NO
TEXT	Description	
HLDDATE	Hold date Specifies the date on which the monitor is to be held. **Note:** The HLDDATE and HLDDAY parameters are mutually exclusive.	Date, *CURRENT, *NONE
HLDDAY	Hold day Specifies the day of the week on which the monitor is held. **Note:** The HLDDATE and HLDDAY parameters are mutually exclusive.	*ALL *NONE *MON *TUE *WED *THU *FRI *SAT *SUN
HLDTIME	Hold time Specifies the time of day to hold the monitor.	Time, *NONE
RLSDATE	Release date Specifies the date on which the monitor is to be released. **Note:** The RLSDATE and RLSDAY parameters are mutually exclusive.	Date, *CURRENT, *NONE
RLSDAY	Release day Specifies the day of the week on which the monitor is released. **Note:** The RLSDATE and RLSDAY parameters are mutually exclusive.	*ALL *NONE *MON *TUE *WED *THU *FRI *SAT *SUN
RLSTIME	Release time Specifies the time of day to release the monitor.	Time, *NONE
ENDDATE	End date Specifies the date on which the monitor is to be ended.	Date, *CURRENT, *NONE
ENDTIME	End time Specifies the time of day to end the monitor.	Time, *NONE
OVRLAP	This value is no longer used.	1-600, *SAME

17 IBM i Disk Watcher Files

For more information on this topic please visit the IBM Knowledge Center link below:

http://www-01.ibm.com/support/knowledgecenter/ssw_ibm_i_74/rzahx/rzahxdwdatafiles.htm

18 iDoctor Disk Watcher Files

This section discusses the iDoctor created SQL tables and physical files related to the Disk Watcher component.

18.1 QAIDRDWM2

This physical file exists in QUSRSYS and contains just a single member. It is used to keep track of the Disk Watcher monitors that exist on a system and is shipped in the QIDRWCH library.

Note: This file won't exist if a Disk Watcher monitor has never been created.

Column	Description	Attribute
MONITOR	Monitor name	C (7)
LIB	Monitor library	C (10)
DFNNAME	Definition name	C (10)
COLNS	Number of collections This is the maximum number of collections that should exist in the library. It does not count/include the currently running collection.	B (4)
STRGAP	Collection duration (in minutes) This value is how much time should elapse between each collection being submitted.	B (4)
STATUS	Status (of the monitor) Possible values: 'A' – Active 'E' – Ended (not active) 'H' - Held	C (1)
STARTTIME	Start time	Timestamp
COLLJOB	Monitor job The name of job that ran or is currently running the monitor.	C (28)
DESC	Description	C (50)
ACTIVEIDX	Active collection index This value from 001 to 999 indicates the current or last collection running. That collection name will be the monitor name with this value concatenated to it.	C (3)
STROVRLAP	This value is obsolete and no longer used.	B (4)
MAXSIZE	Collection maximum size (megabytes) This value indicates the maximum amount of data that each collection should contain. **Note:** If this value is reached before the collection duration elapses, then NO data will be captured until the next collection in the monitor starts.	C (10)
RES1	Reserved field #1	C (10)
RES2	Reserved field #1	C (10)

18.2 QAIDRDWINTI_<MBR>

This SQL table is created by the Trace summary analysis and acts as a basic time summary of the trace data into 1 second intervals. It is used for in most trace-based graphs in Disk Watcher.

Note: Time periods where no data was collected are also included in this file for ease of graphing the data.

Column	Description	Attribute
INTERVAL	Interval number	D (15,0)
TOTIOS	Total I/Os	B (4)
RISTARTTOD	End of interval time of day This is the time that the collection interval began ending.	Timestamp
MINIO	First I/O usecs since collection start Within the 1 second interval time period, this is exact time that the 1st I/O occurred since the collection started in microseconds.	B (8)
MAXIO	Last I/O usecs since collection start Within the 1 second interval time period, this is exact time that the final I/O occurred since the collection started in microseconds.	B (8)
STDATAMISS	This is used to indicate if any I/Os were NOT captured by Disk Watcher in this interval. It will be set to '1' if data was not collected that should have been.	B (4)

18.3 QAIDRDWTSUM_<MBR>

This SQL table is created by the Trace summary analysis and acts as a detailed summary of the trace data into 1 second intervals. It is used for the trace-based rankings and selection over time graphs.

This file contains 1 record per disk path (with I/Os captured) per 1 second interval per object segment, pgm/procedure, and taskcount.

13 I/O time buckets are used in this table which have field names numbered 1-13. These I/O time duration buckets are hard coded as the following:

Bucket	Description
1	< 100 microseconds
2	100 – 199 microseconds
3	200 – 399 microseconds
4	400 – 799 microseconds
5	800 – 1599 microseconds
6	1.6 – < 3.2 milliseconds
7	3.2 – < 6.4 milliseconds
8	6.4 – < 12.8 milliseconds
9	12.8 – < 25.6 milliseconds
10	25.6 – < 51.2 milliseconds
11	51.2 – < 102.4 milliseconds
12	102.4 – 204.8 milliseconds
13	>204.8 milliseconds

Column	Description	Attribute
INTERVAL	Interval number	B (8)
DEVRNAME	Disk path	C (10)
STASPNUM	ASP number	B (4)
STUNITNUM	Disk unit number	B (4)
TCSEGKEY	Segment key This is the identifier in column SEGKEY in file QAPYDWOBJR (object resolution data) that identifies the object or segment associated with these I/O operations.	H (8)
TCPROCKEY	Program or procedure key This is the identifier in column PROCKEY in file QAPYDWPGMR (program/procedure resolution data) that identifies the program or procedure associated with these I/O operations.	H (8)
TCTASKCNT	Task count This is the identifier in column TSKCNT in file QAPYDWTDER (TDE resolution data) that uniquely identifies the job/thread or system task associated with these I/O operations.	B (8)
TOTCNT	Total I/Os	B (4)
TOTTIME	Total I/O time (usecs) **Note: This time includes the deferred queue time if any.**	B (8)
DQTOTTIME	Total deferred queue time (usecs) The total time operations spent waiting on the deferred I/O queue (in microseconds).	B (8)
TOTPCNT	Total page counts This is the number of pages in all I/O requests.	B (8)
WRTCNT	Total writes	B (4)
WRTTIME	Total write time (usecs)	B (8)
DQWRTTIME	Total deferred queue write time (usecs)	B (8)
WRTPCNT	Total write page counts	B (8)

RDCNT	Total reads	B (4)
RDTIME	Total read time (usecs)	B (8)
DQRDTIME	Total deferred queue read time (usecs)	B (8)
RDPCNT	Total read page counts	B (8)
OTHCNT	Total other I/Os	B (4)
OTHTIME	Total other time (usecs)	B (8)
DQOTHTIME	Total deferred queue other I/O time (usecs)	B (8)
OTHPCNT	Total other I/O page counts	B (8)

STWRTTOT1-13	Total write time (usecs)	B (8)
STWRTDQ1-13	Total write deferred queue time (usecs)	B (8)
STWRTC1-13	Total writes	B (4)
STWRTPC1-13	Total write page counts	B (8)

STRDTOT1-13	Total read time (usecs)	B (8)
STRDDQ1-13	Total read deferred queue time (usecs)	B (8)
STRDC1-13	Total reads	B (4)
STRDPC1-13	Total read page counts	B (8)

STOTHTOT1-13	Total other I/O time (usecs)	B (8)
STOTHDQ1-13	Total other I/O deferred queue time (usecs)	B (8)
STOTHC1-13	Total others I/Os	B (4)
STOTHPC1-13	Total other I/O page counts	B (8)

18.4 QAIDRDWTSUM2_<MBR>

This SQL table is created by the Trace summary analysis and acts as a summary of the trace data into 1 second intervals. It is used in the Trace overview graphs in Disk Watcher.

This file contains 1 record per ASP (with I/Os occurring) and 1 second interval.

See the previous section for more information about the columns defined in this file.