IBM iDoctor for IBM i
iDoctor Monitors Guide

IBM iDoctor for IBM i Development Team
11 October 2016

Licensed Materials - Property of IBM
© Copyright International Business Machines Corporation 2010, 2016. All rights reserved.

Abstract
Guide to running the iDoctor Monitors. Monitors provide 24x7 collection support of Job Watcher, Disk Watcher and PEX statistics for detailed performance investigation data when you need it most.

Changes
19 June 2006 –Initial Creation
11 Dec 2008 – Updated to match GUI changes.
12 Jan 2009 – Added mention of new support for PEX Analyzer Monitors at 6.1. Added max collection size parameter at 6.1.
17 Feb 2009 – Added comments on restarting monitors for scheduling or after an IPL.
20 May 2009 – Updates for latest GUI changes. Added Appendix B (describes modification of system startup program to automatically restart a monitor after an IPL).
11 Feb 2010 – Updated for latest changes (RSTJWMON removed)
11 Oct 2016 – Updated for latest changes

Introduction

The purpose of this document is to demonstrate how to define and start an iDoctor Monitor. The following are some key features that I think makes Monitors a key tool in performance analysis:

1. Allows continuous 24x7 collections.
2. Collections can overlap so no activity is lost.
3. Users can define the number of historical collections retained.
4. Can be started, paused, resumed, and ended using either the iDoctor GUI or by using traditional “green screen” commands. In this document we are going to focus on using the iDoctor GUI, in appendix A there will be samples of the “green screen” commands.
5. Multiple monitors can be active at the same time.

The support for Job Watcher , Disk Watcher and PEX Monitors is available in IBM i version V5R3M0 and higher.

Starting a Monitor

You first want to start the iDoctor GUI and launch one of the components that support monitors by double-clicking the desired system as shown in figure 1. If you need to add a connection to your system first, then right-click the list of connections and choose the Add Connection… menu. This example will show how to start a Job Watcher monitor.

Within Job Watcher, expand the Job Watcher icon and then (left) click Monitors in order to work with the monitors available on the current system as shown in figure 2.

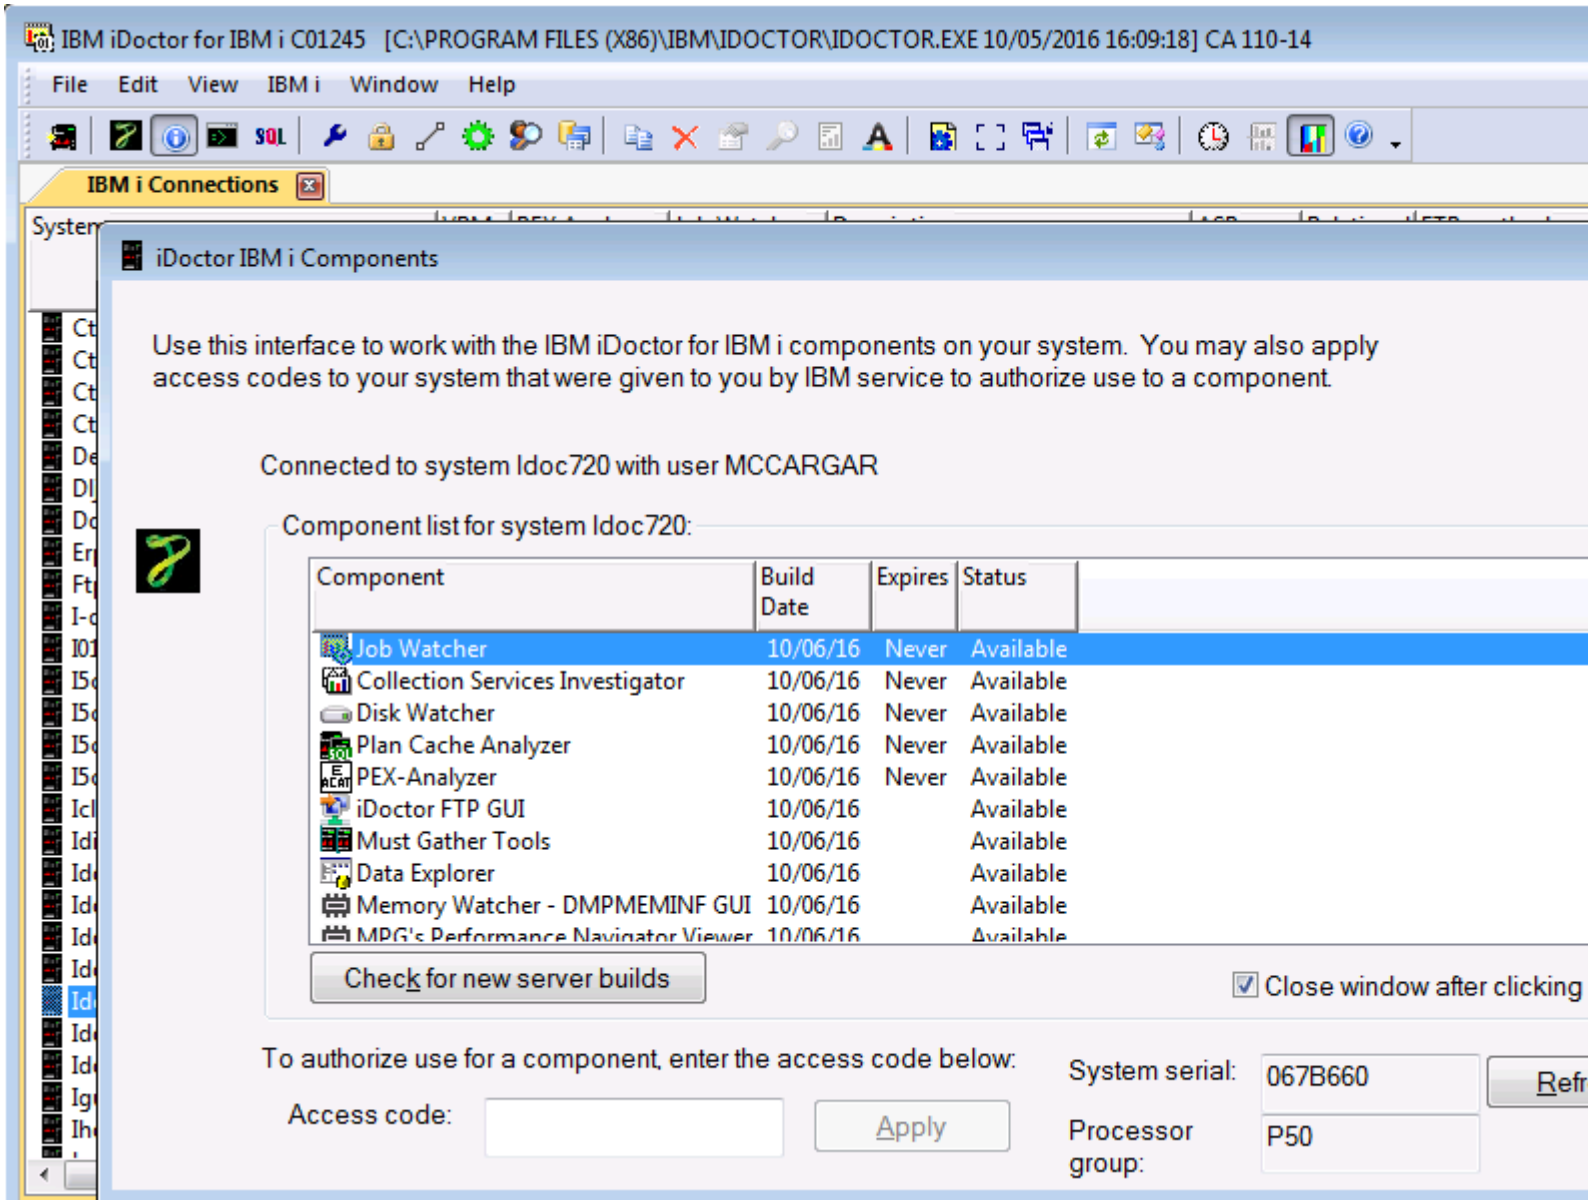

Figure 1: Starting Job Watcher

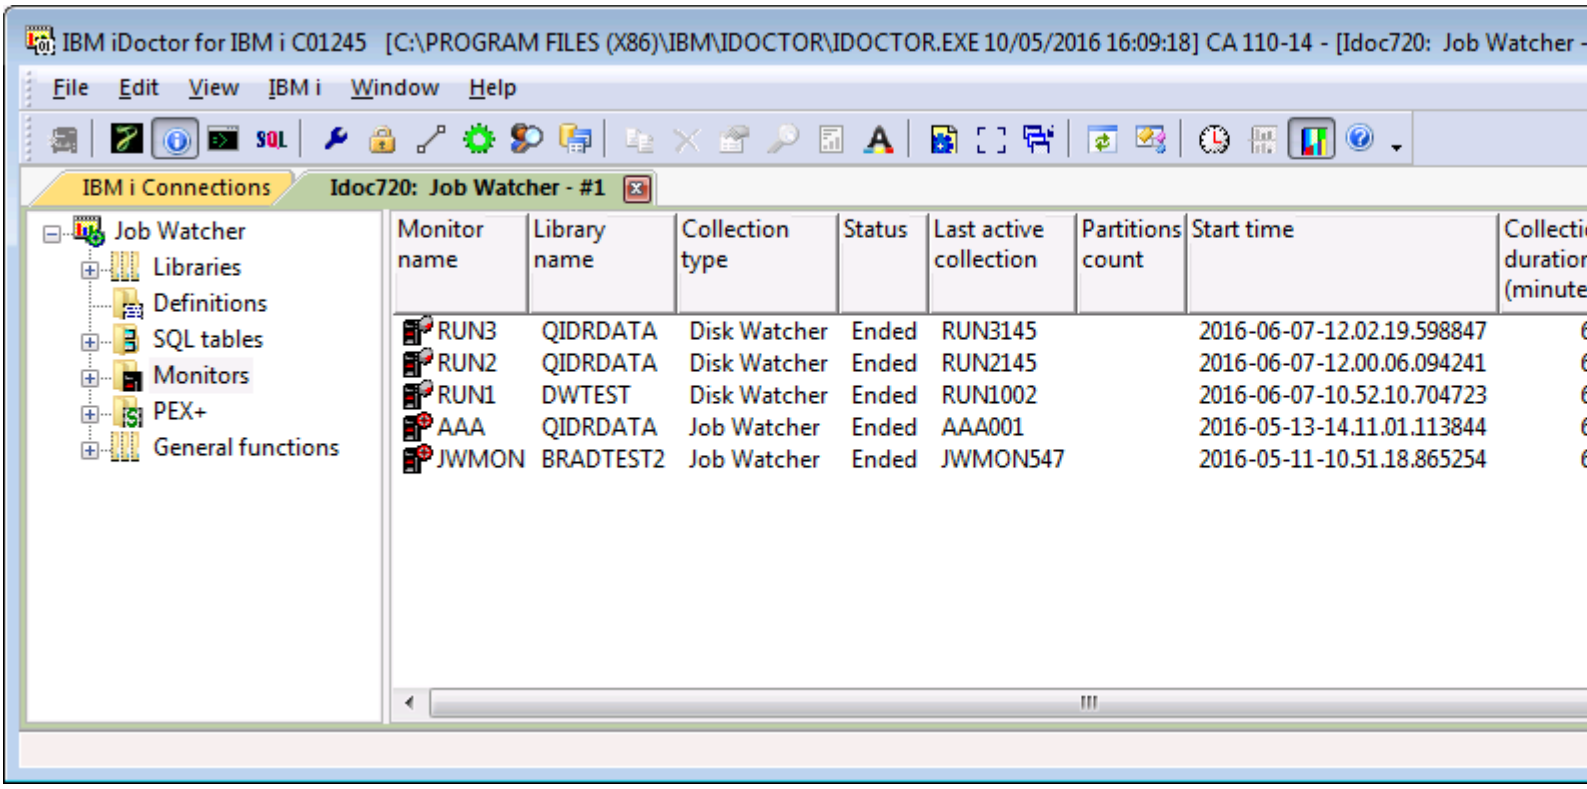

Figure 2: The Monitors Folder in Job Watcher

Any existing monitors will be listed in this window. Otherwise the window will be blank. Create a new monitor by right-clicking the Monitors folder and selecting the option, Start Monitor. Figure 3 shows an example of this menu.

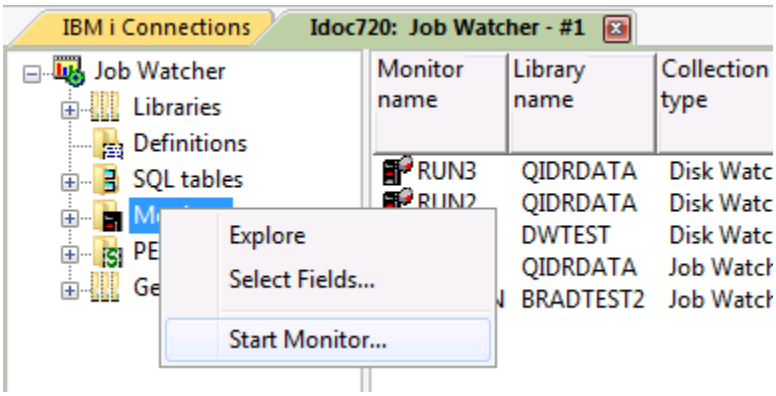

Figure 3: Starting a monitor

This interface allows you to start a Monitor.

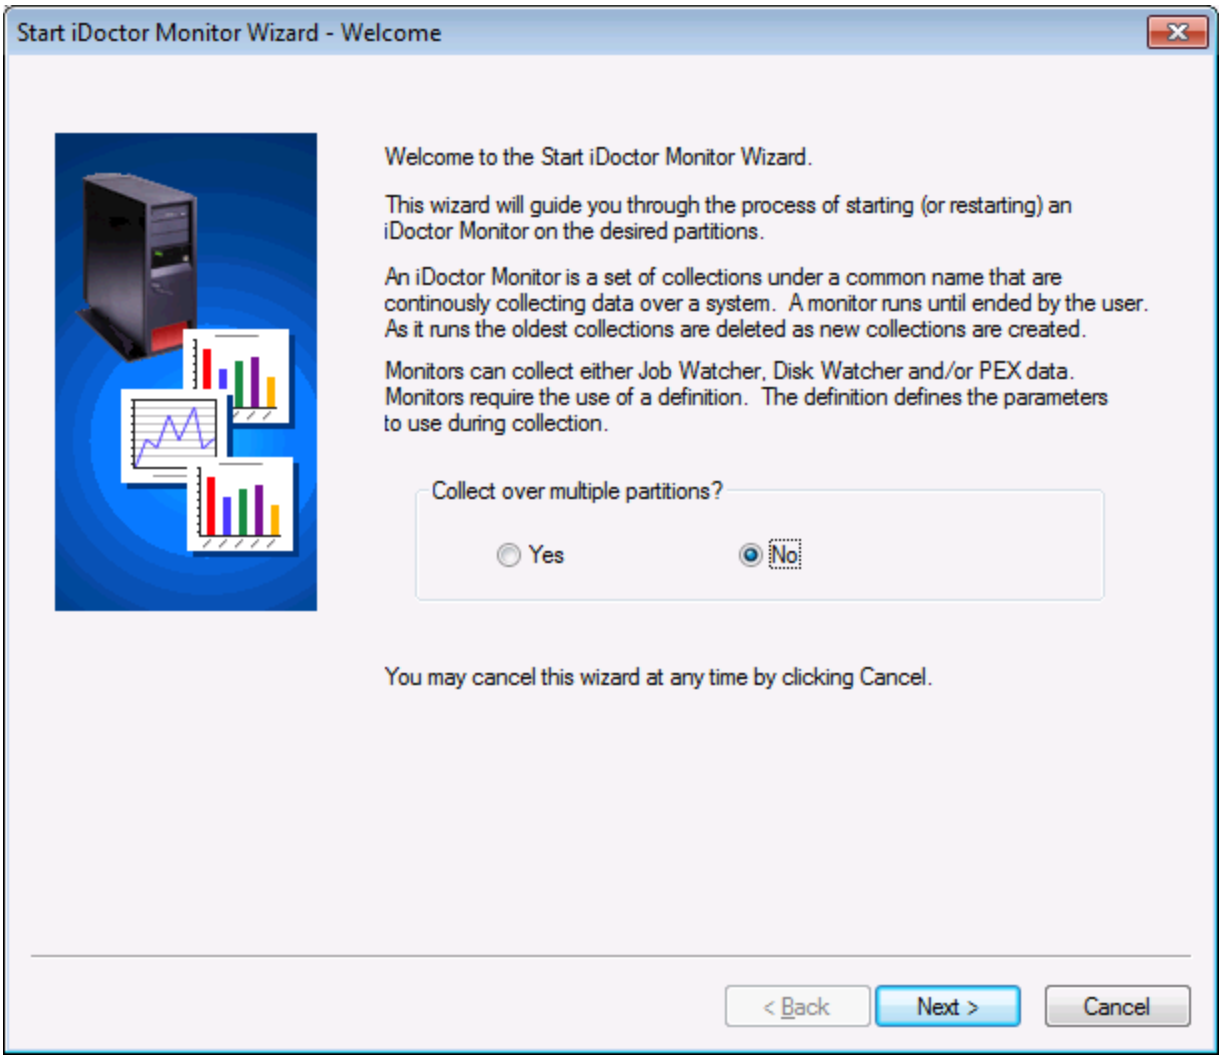

Figure 4: Welcome screen

When starting a monitor we must provide some required parameters as shown in figure 5.

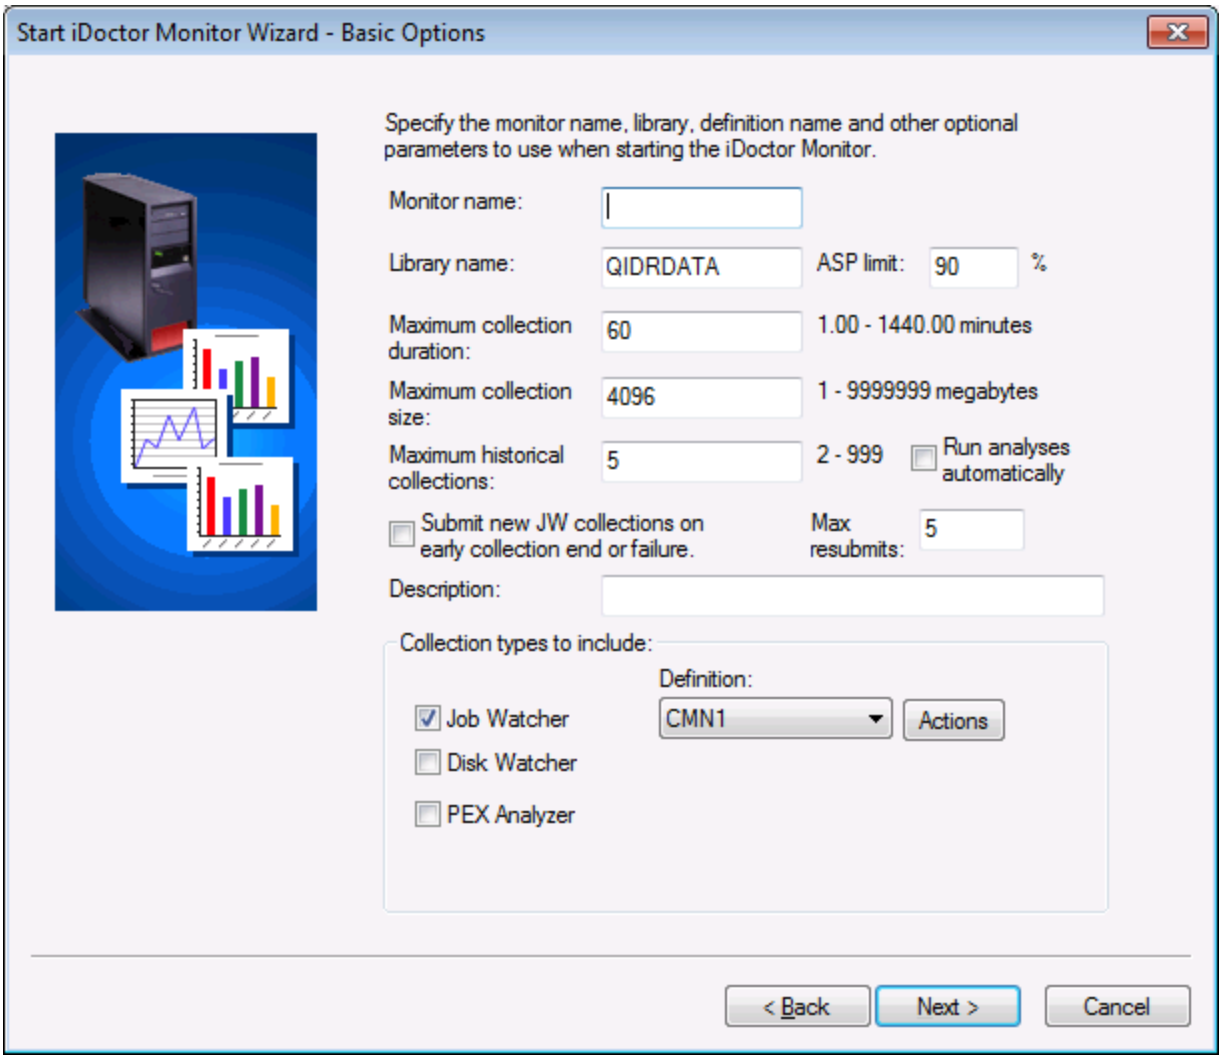

Figure 5: Start iDoctor Monitor Wizard – Basic Options

The following is a description of each of these required parameters:

- **Monitor name**: This is the name we give our monitor and our collection members will also start with this name when creating a new collection.

- **Library**: This is the library where we want to create our JW collection data.

- **ASP limit**: This value provides the capability to control the maximum amount of ASP usage allowed before the monitor will end. The ASP is based on the library's ASP. The default value shown is 90% and any changes made will update data area QMONASPLMT in QIDRWCH and be shown on this screen as the default value during future uses.

- **Maximum collection duration**: How long do we want each collection to run for?
  - **Note:** Definitions do not contain any parameters that determine when the collection will end. This is controlled by the STRJW command which is called by the monitor commands.

- **Maximum collection size**: This parameter indicates the maximum size (disk space) for each collection in the monitor. If the amount of data collected exceeds this value in a collection, the collection will end prematurely and there will be gaps in the monitor data. The next monitor won't start until the Maximum collection duration is reached.

- **Maximum historical collections**: How many collections do you want keep? The JW monitor allows you to keep up to 999 historical collections. You will want to set this value high enough so the data is available when needed but not too high that you are consuming a lot disk space unnecessarily.

- **Run analyses automatically:** This option if checked will automatically run all "default" iDoctor analyses on each collection created by the monitor. The default analyses are the most commonly used analyses and can be identified by using the Analyses -> Analyze collection… menu on a collection and looking at the "Run all default" column in the window. Keep in mind that in some cases this runs SQL statements that could be resource intensive (CPU/disk) so you may not want to do this on production systems.

- **Submit new JW collections on early collection end or failure**: If the Job Watcher monitor detects that new collections are not being created, this option will allow additional attempts to create collections to happen. In some situations like if a system save is occurring (which can cause the STRJW command to fail) using this option will minimize gaps in the data collected.

- **Max resubmits:** Use this option to indicate how many times additional Job Watcher collections will attempt to be created on each iteration of the monitor. **Note:** This option only applies when the "Submit new JW collections on early collection end or failure" is checked.

- **Description:** You can give your monitor a description here.

- **Collection types to include**: This area allows you to indicate if Disk Watcher, PEX and/or Job Watcher should be collected (or all). The definition to use for each and the available actions to work with the definitions are available by clicking the Actions button.

- **Definition**: This is the name of the definition we are going to use. The definition provides all of the parameters used when running collections in the monitor. Clicking the Actions button provides several options as shown in the figure below:

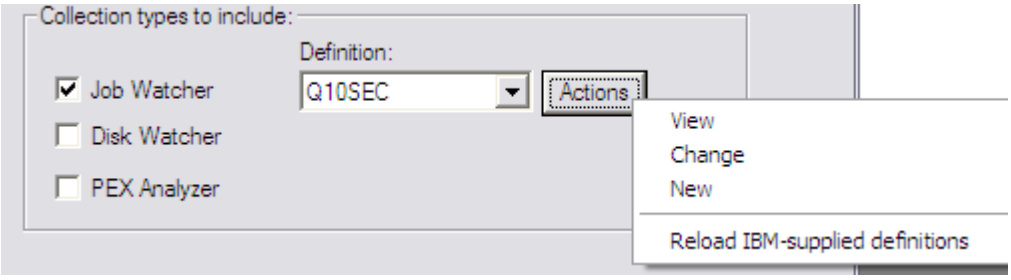

  - **View** – Display the definition shown in the drop down list.
  - **Change** – Allows you to modify the definition within the Add Job Watcher Definition Wizard

- **New** – Allows you to create a new definition using the Add Job Watcher Definition Wizard.
- **Reload** – Rebuilds the list of IBM-supplied definitions. This may be necessary if the Q* definitions aren't showing correctly because there were removed.

When done with this screen click Next to move on to the scheduled options.

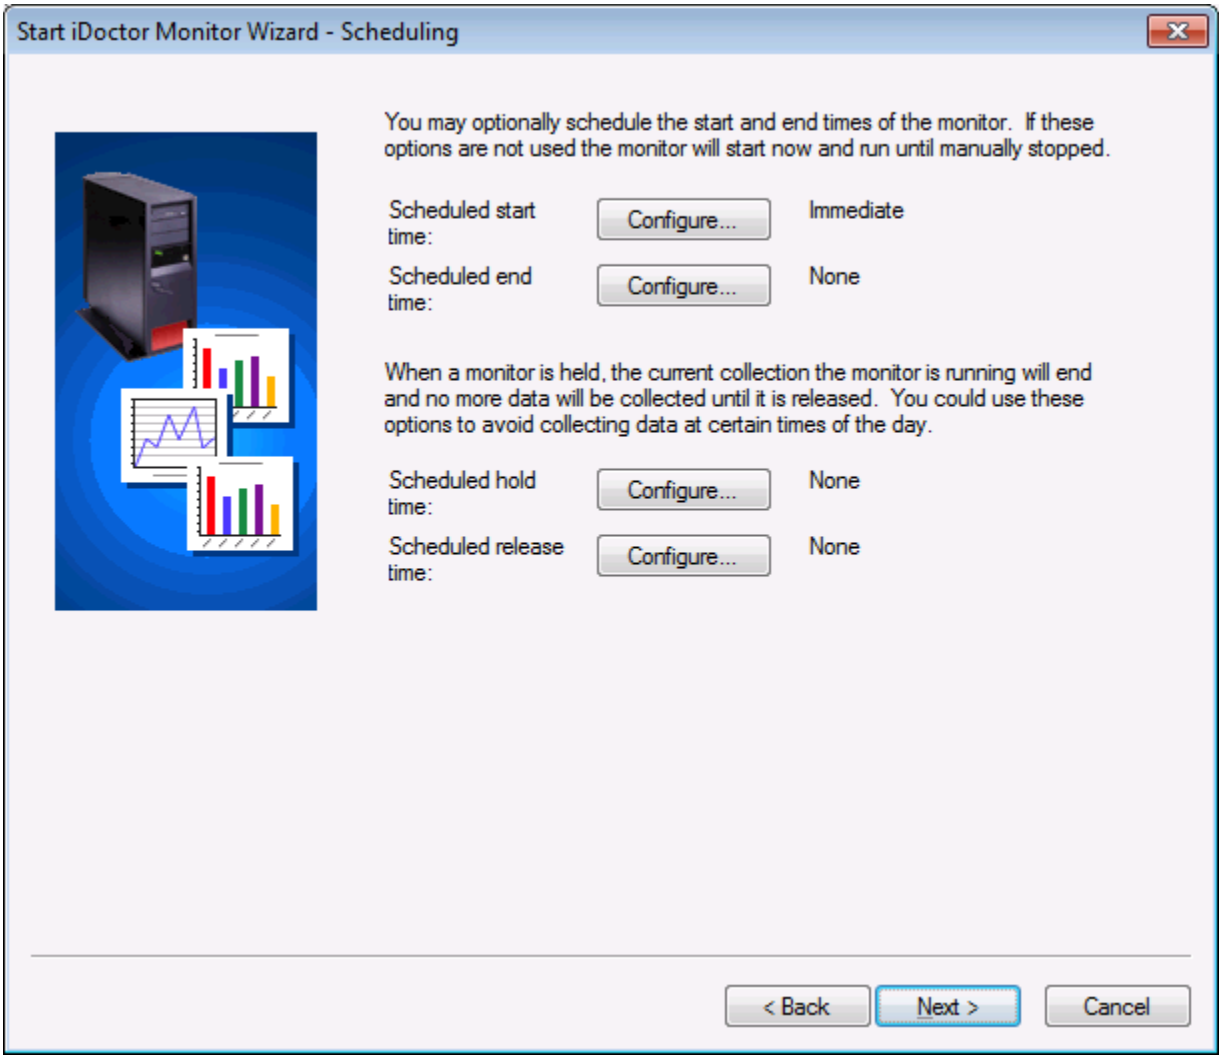

Figure 6: Start iDoctor Monitor Wizard – Scheduling

To automate when the monitor(s) run there is also a built in function to create job scheduled entries to:

- **Scheduled start time:** When to start the monitor.
- **Scheduled end time:** When to end the monitor.
- **Scheduled hold time:** By scheduling a hold time you can temporarily stop the monitor from starting additional collections. This will also end any collections currently running.

- **Scheduled release time:** Schedule a release time to restart the monitor. The collection name will be incremented, to not overwrite any historical collections.

The example shown in figure 7 is the Schedule monitor start time window. This window will be the same whether you are scheduling the start, end, hold or release of a Job Watcher monitor. You can see from figure 7 you have the option to create an entry to run the job schedule entry once, weekly or daily. Examples of each of these options are included in the following pages.

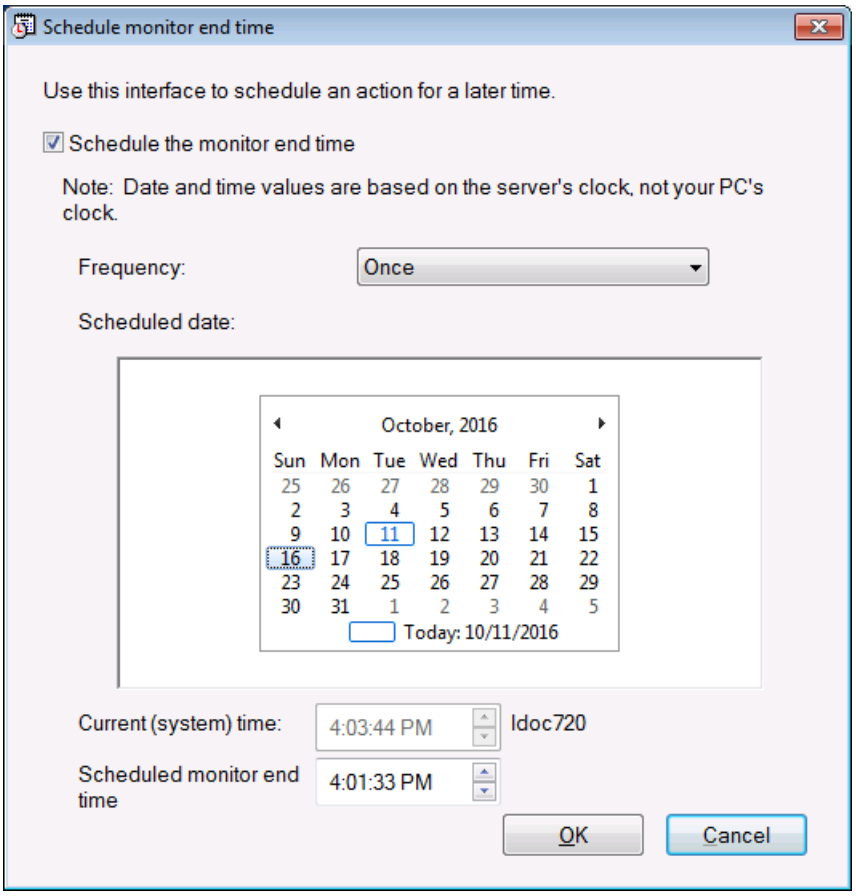

Figure 7: Scheduling interface

Restarting a Job Watcher monitor after an IPL requires the use of the STRJWMON command**.** Also remember to start the Job Watcher subsystem **STRSBS SBSD(QSYS/QIDRJW)**

If you select the option to run the scheduled entry Weekly the window will look like that shown in figure 8. Pick the day of the week and time you want to the scheduled entry to run. As you probably would expect you can only pick one day of the week.

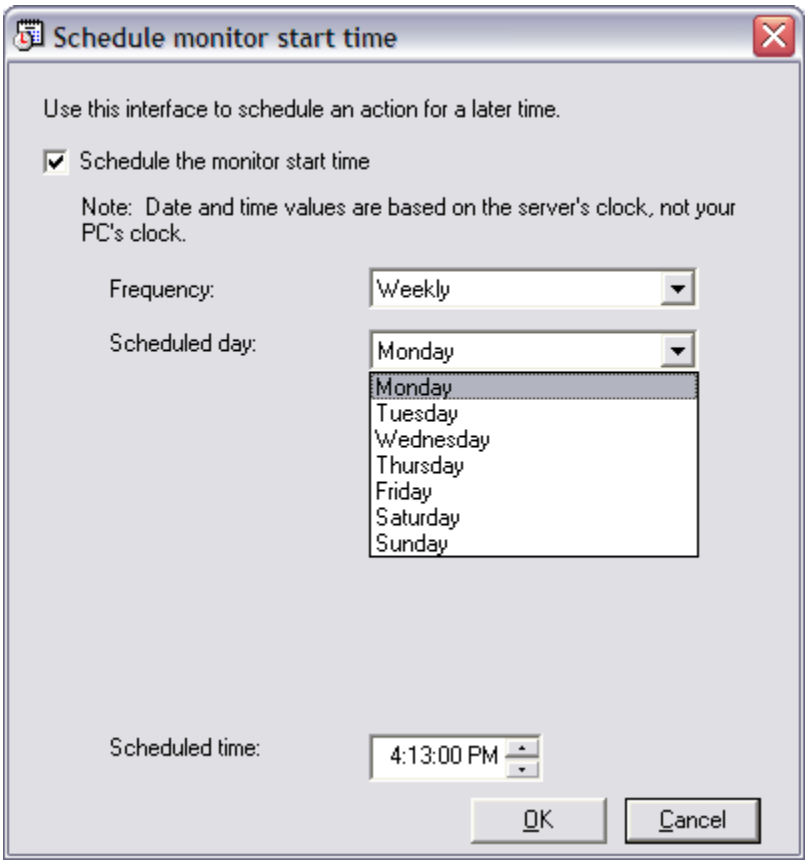

Figure 8: Schedule Monitor – Weekly

When you select the option to run the scheduled entry Daily you will only need to specify the time that you want the scheduled entry to run. See figure 9 for an example.

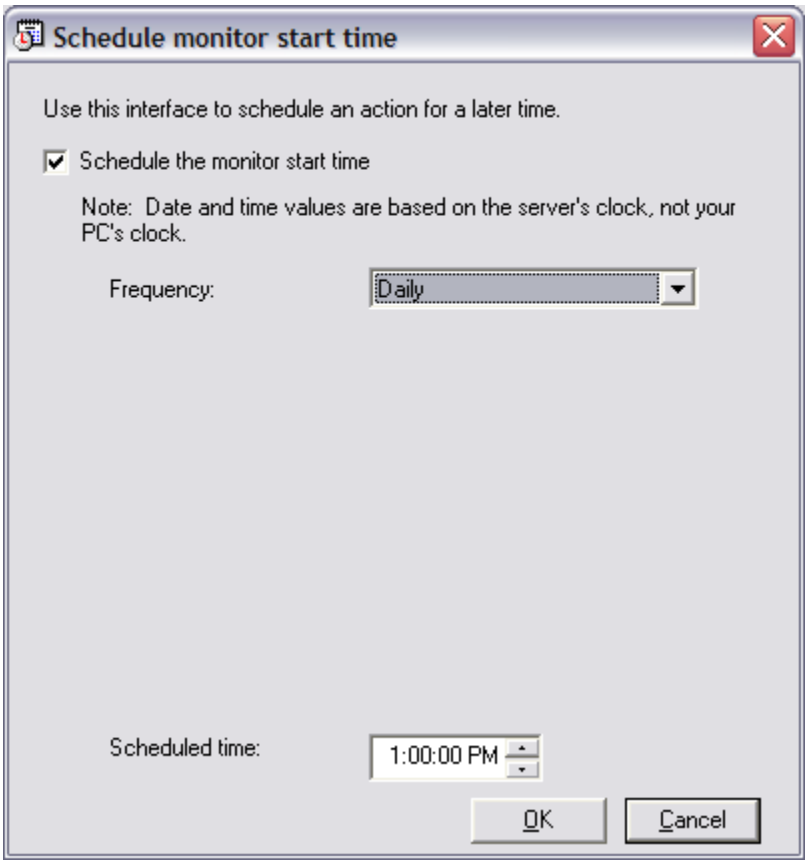

Figure 9: Schedule Monitor – Daily

You can view JW monitor scheduled entries by using the iDoctor GUI, by clicking the Scheduled jobs folder under General functions -> Work management
Or you can use the “green screen” command WRKJOBSCDE.

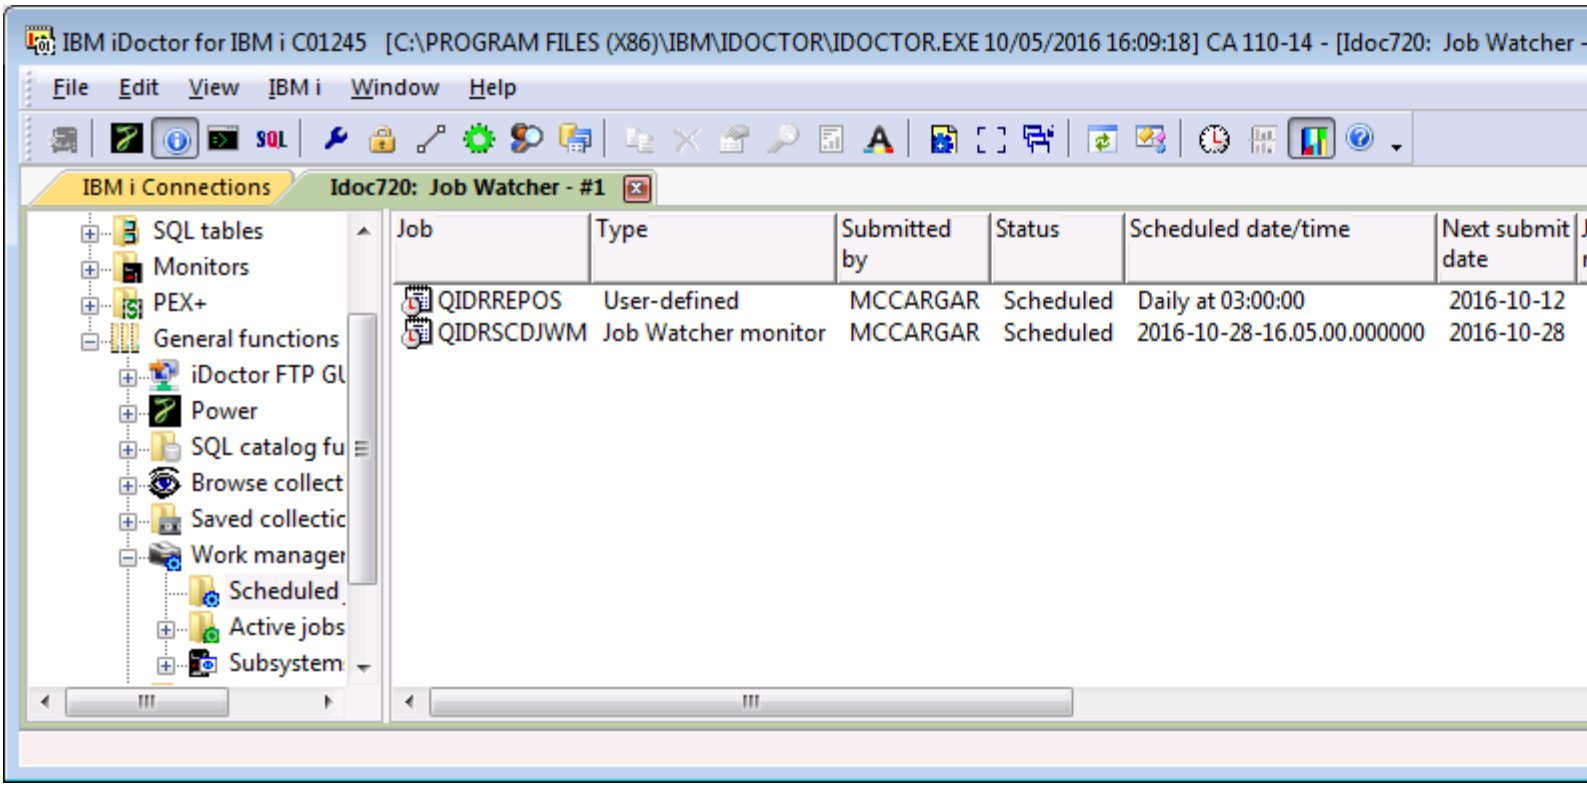

Figure 10: General functions -> Work management -> Scheduled jobs folder

On the last page of the Wizard, the options selected are presented to the user for final confirmation before the monitor is started or scheduled.

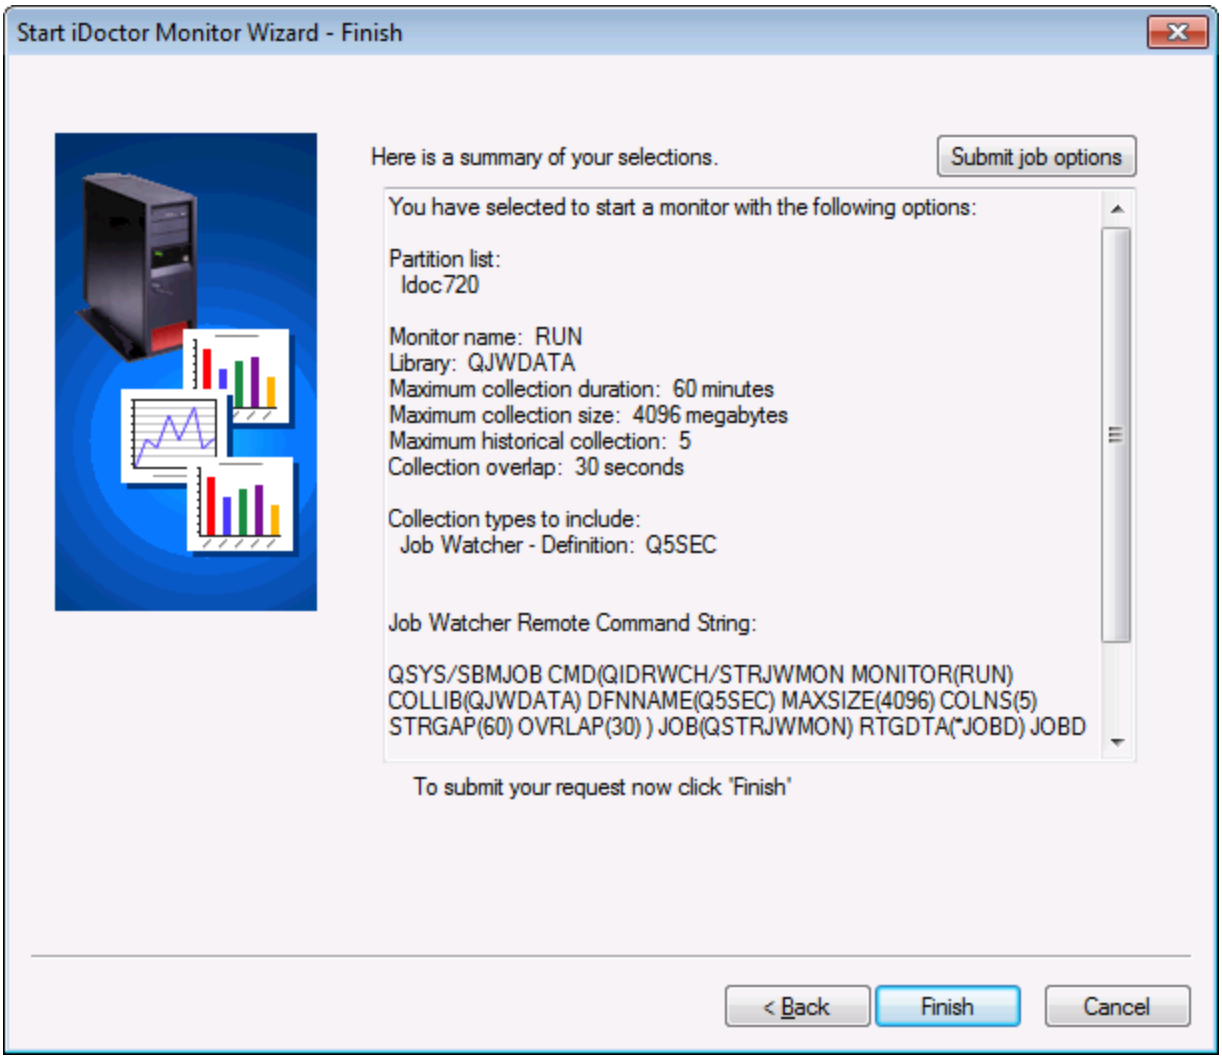

Figure 12: Start iDoctor Monitor Wizard - Finish

In the example in figure 12, I have configured the monitor RUN to collect the data in library QJWDATA. I am using the iDoctor supplied JW definition Q5SEC. You will also notice that the monitor will start immediately. By clicking on the Finish button the monitor will start and submit the first Job Watcher collection.

Next refresh the screen by pressing F5 or use the View -> Refresh selected menu.

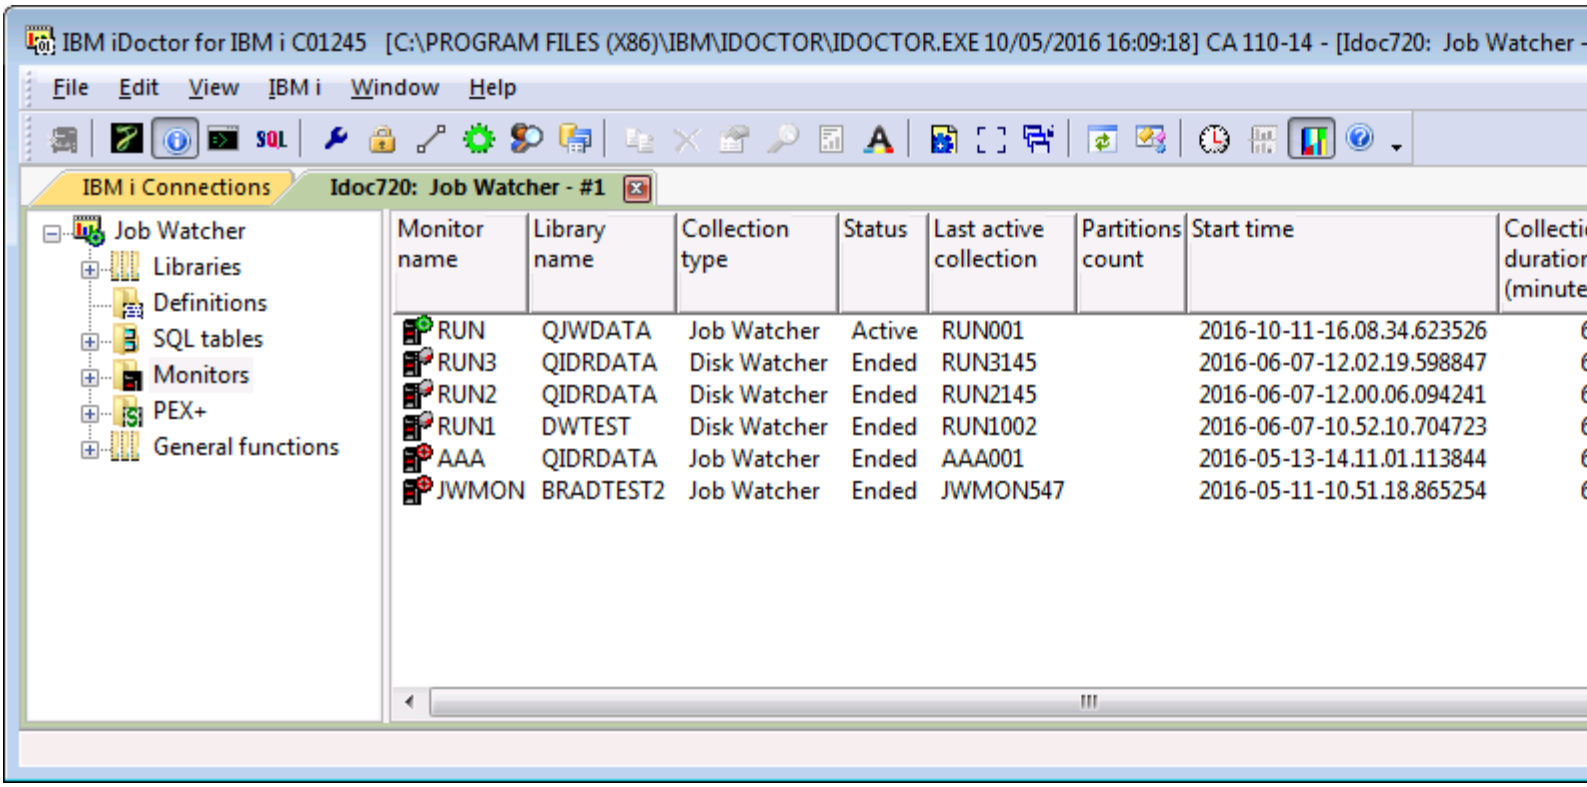

Figure 13: Job Watcher Monitor status
Figure 13: Job Watcher Monitor status

Note the name of the last active collection; as each new collection is started the name will increment.

As a Monitor runs, the collections it contains are ended by the Monitor job. The monitor will also check for the existence of file/member QAPYJWRUNI. This may cause messages to show up in the message queue, but this is **normal** and is not cause for concern.

Working with Job Watcher Monitors

Once we have a JW monitor created there are a number of options available to manage the monitor and the data it creates. Right-click a monitor to work with the options available for it. Double-click a monitor to drill in and work with the collections it contains. Figure 14 shows an example of these options and a description of each. The following is a description of these options:

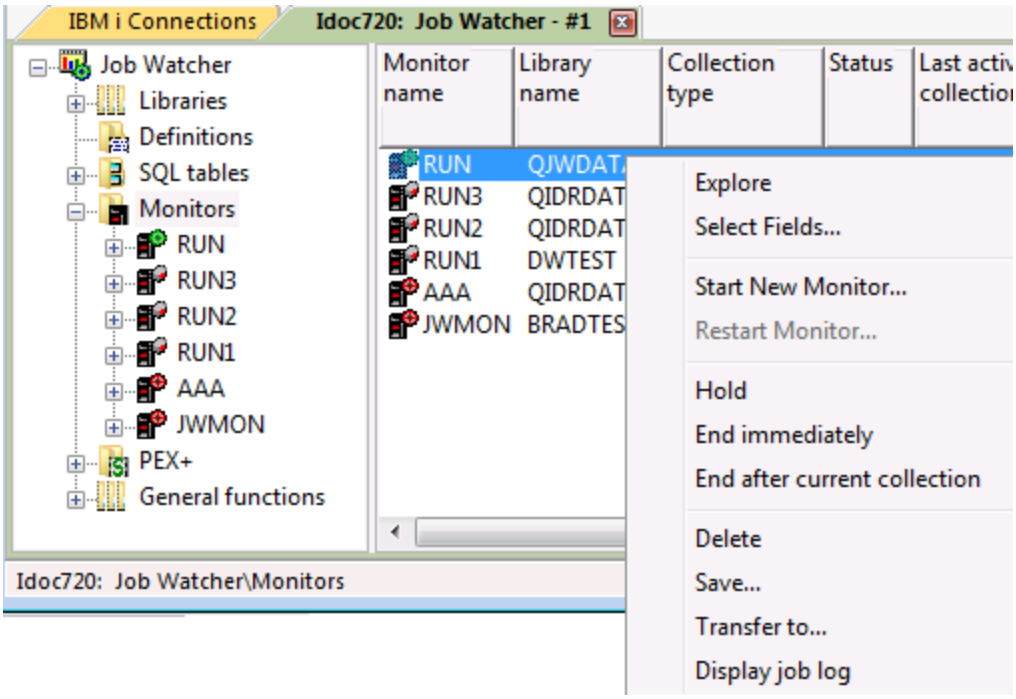

Figure 14: Job Watcher Monitor options

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

Popup Menu	Description
Explore	This option allows you to view the collections that are contained within the selected monitor.
Select Fields…	This option lets you modify the columns shown for the list of collections within the monitor.
Start New Monitor	Opens the Start iDoctor Monitor Wizard in order to create a new monitor.
Restart Monitor	Opens the Start iDoctor Monitor Wizard to restart the selected monitor. This option is only enabled if 1 monitor is selected and it is NOT already running.
Hold or Release	This option allows the selected monitor to be held (or released if it is held). If held the active collection will be ended immediately and no more collections will be started until the monitor is released. When a monitor is released it will begin collecting data again.
End immediately	This option will end the monitor and all active collections defined within immediately.
End after current collection	This option will end the monitor once the current collection running completes.
Delete	This option will remove the monitor and all collections contained within them from the system.
Save	Provides the capability to save all collections within the monitor to a save file. Collections saved using the iDoctor GUI are accessible using the General Functions -> Saved Collections folder.
Transfer to	This option allows a user to save and then transfer a collection to another system, to the PC or IBM (Ecurep/testcase.)
Display job log	Displays the job log for the selected monitor. This option is only enabled if 1 monitor is selected and the job log exists.

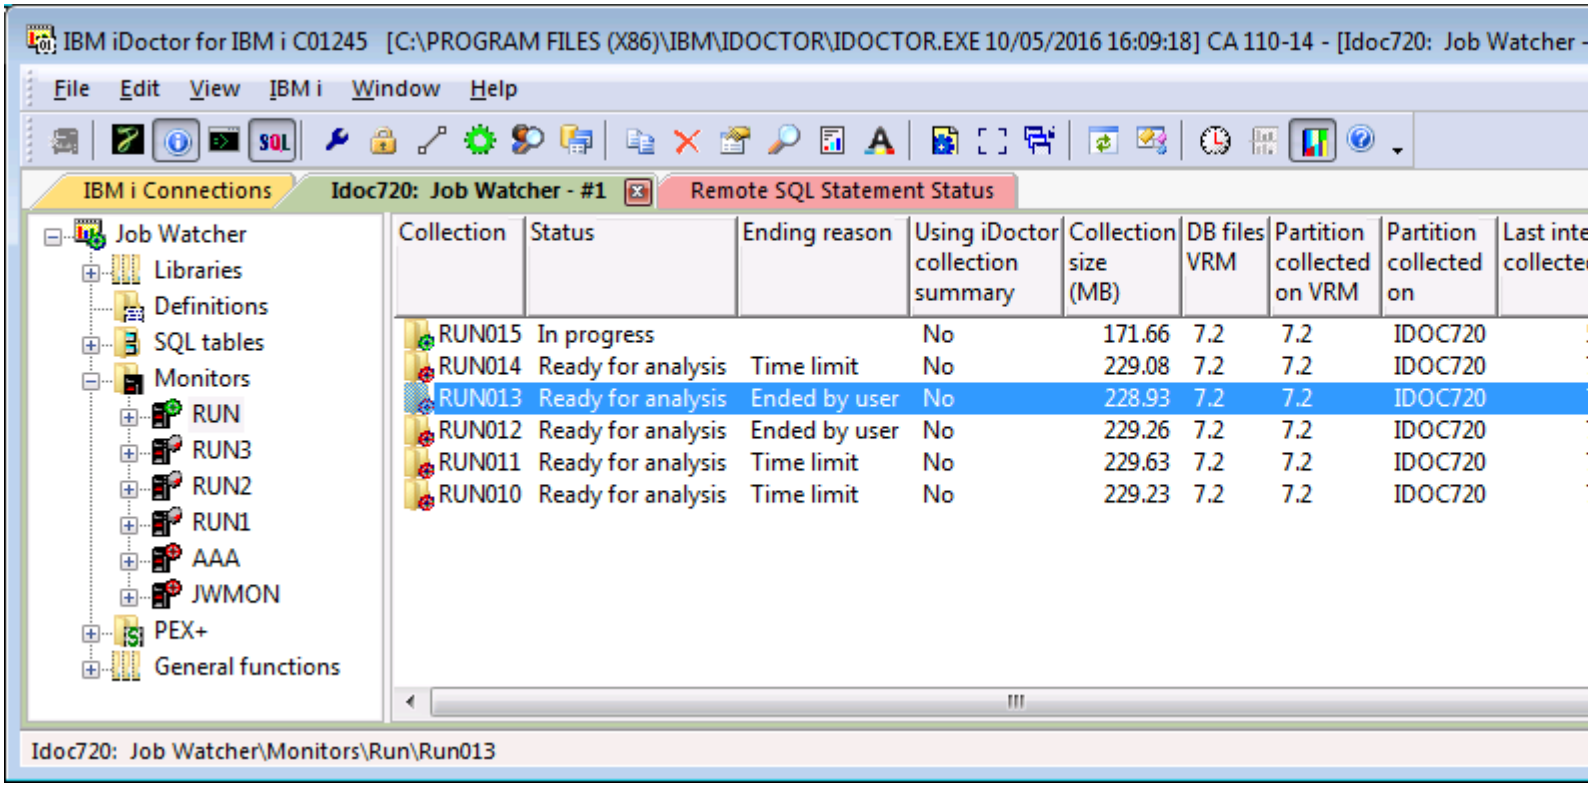

Figure 15: Work with Job Watcher collections

As the Job Watcher monitor continues to run there will always be a number of historical collections available. The number of historical collections will depend on what we specified when we first started the monitor. In the example used in this document we specified 5 historical collections. You can begin your analysis of any of these collections at any time once at least 1 interval of data has been collected. Figure 16 shows a list of the graph folders available within Job Watcher. The description of these folders and what graphs to use is outside the scope of this document. For more information, see the iDoctor documentation on our website.

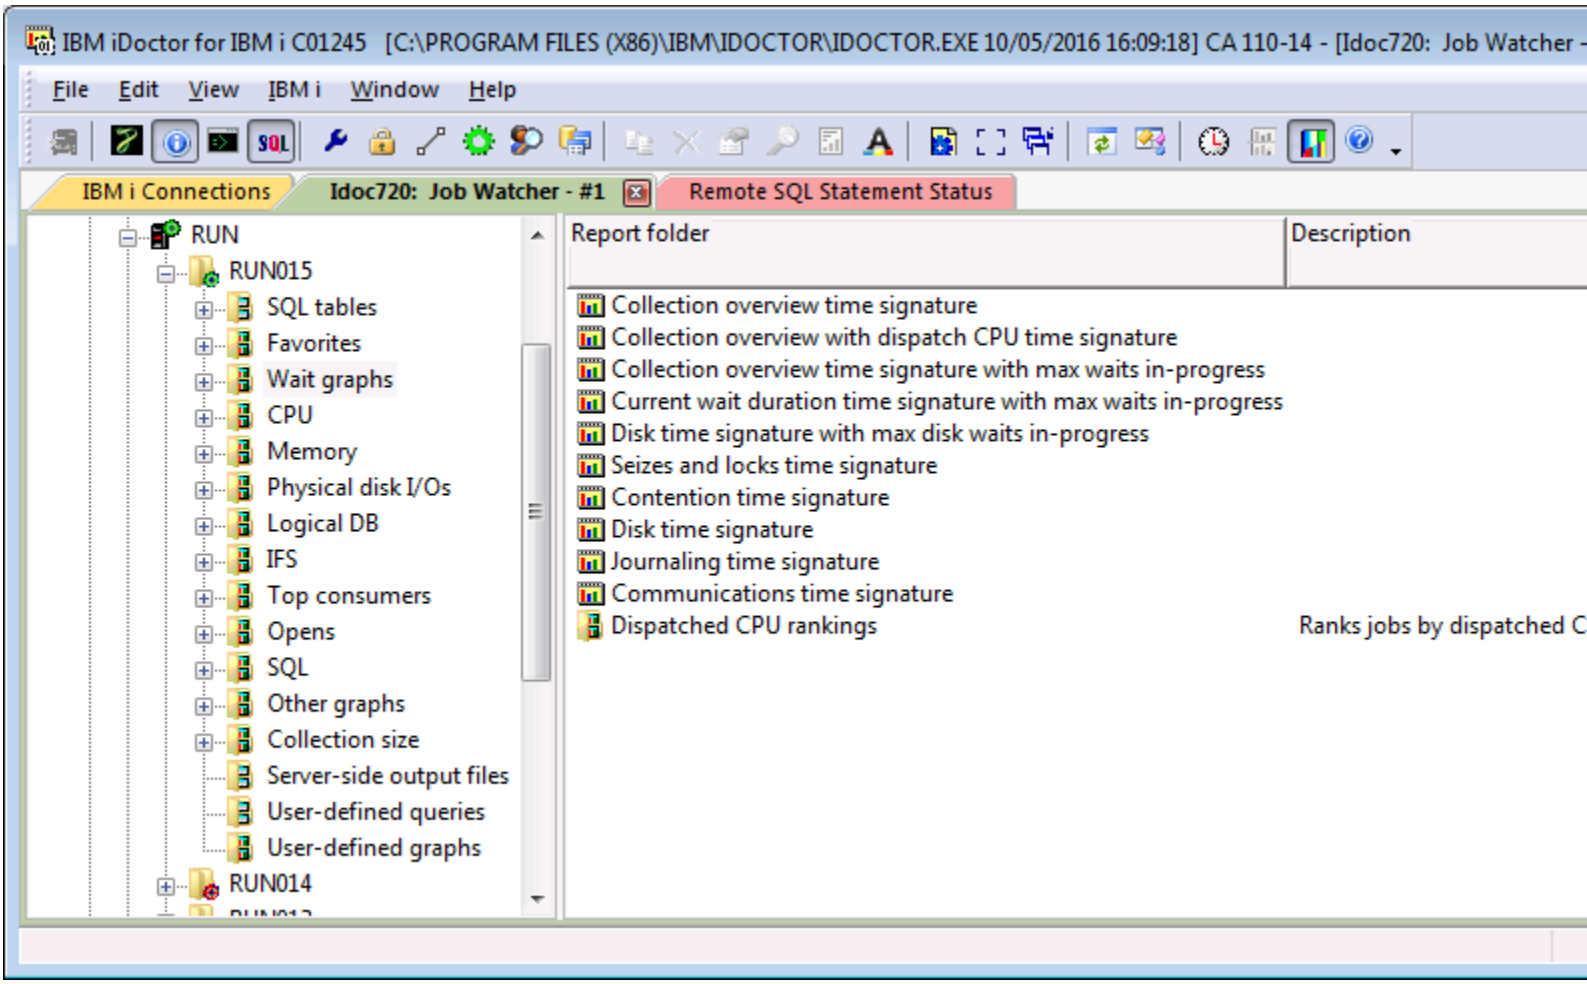

Figure 16: Graphs available within a Job Watcher collection (within a monitor)

If you have any questions or suggestions about any of the iDoctor products you can either direct these to the iDoctor Community at https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUuid=5f260572-0786-4dee-889d-ae7adc18944f or email us at idoctor@us.ibm.com

Appendix A - Monitor Commands

The following commands are found in the library QIDRWCH and can be entered from an IBM i command line or also submitted to batch using the SBMJOB command.

Add Job Watcher definition – ADDJWDFN (found in QSYS)
Start Job Watcher – STRJW (found in QSYS)
Start Job Watcher monitor – STRJWMON

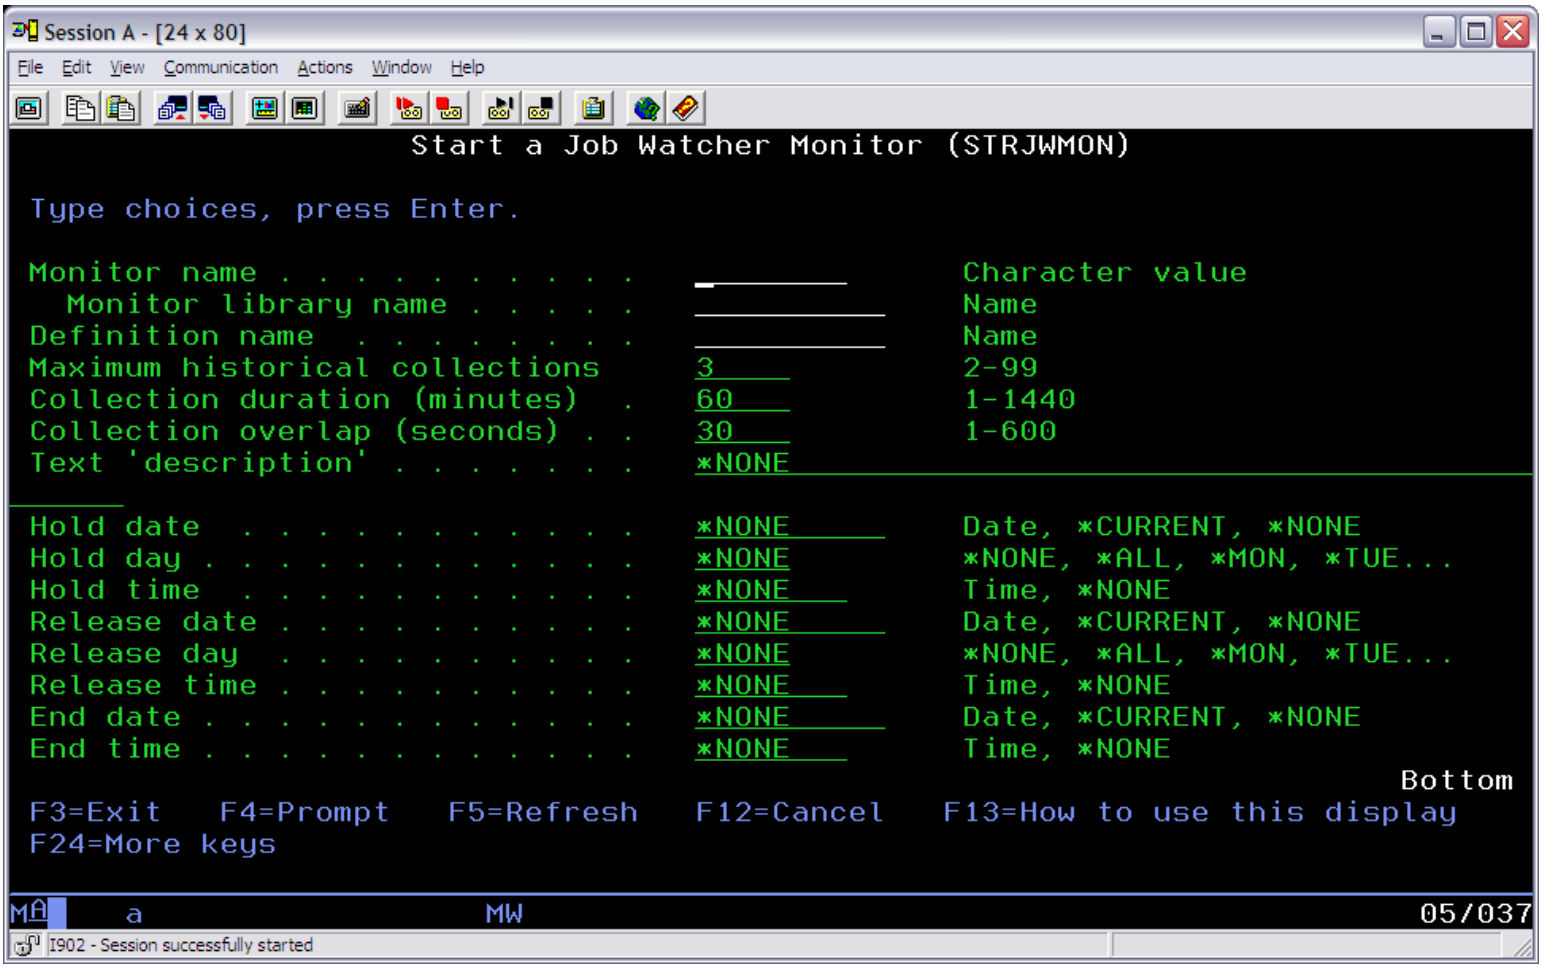

Hold Job Watcher monitor - HLDJWMON

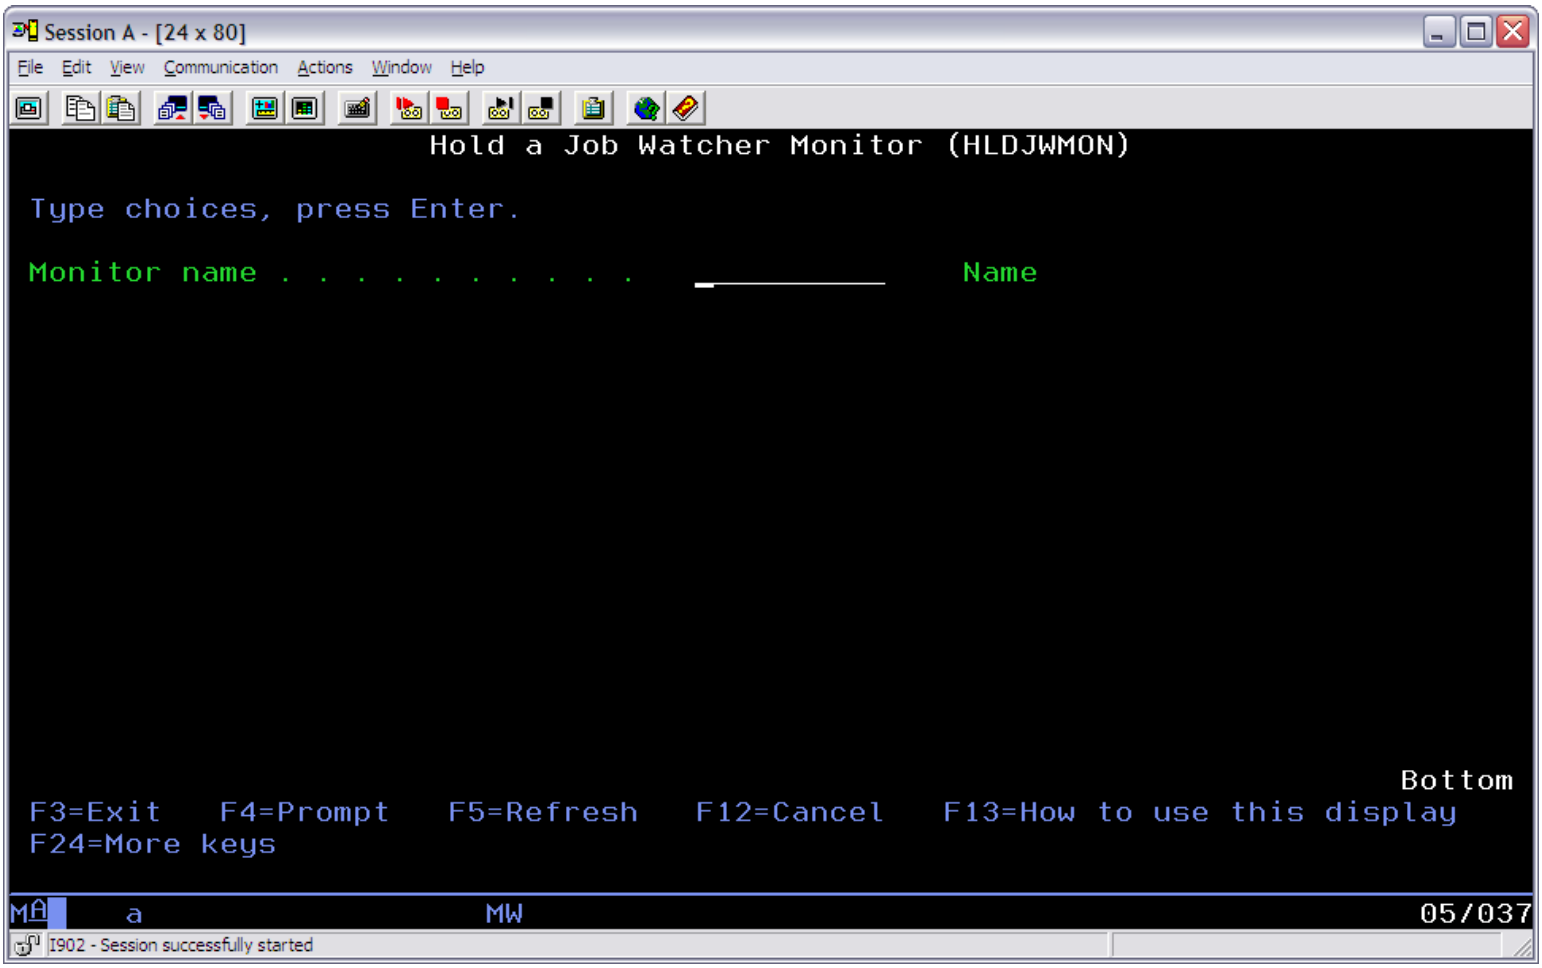

Release Job Watcher monitor – RLSJWMON

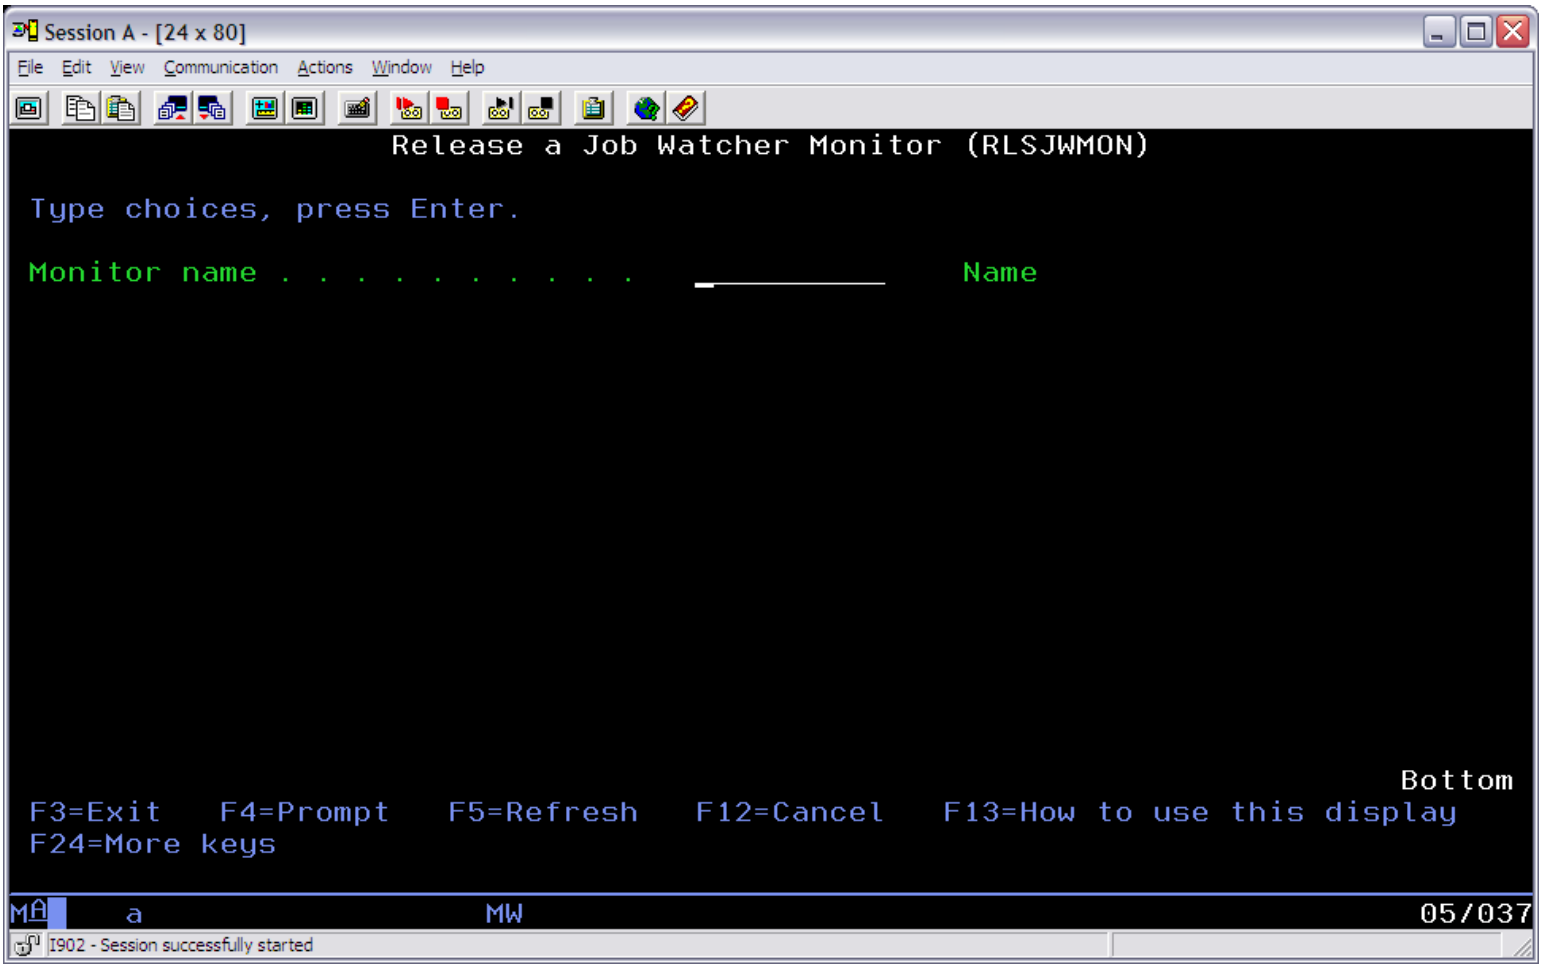

End Job Watcher monitor - ENDJWMON

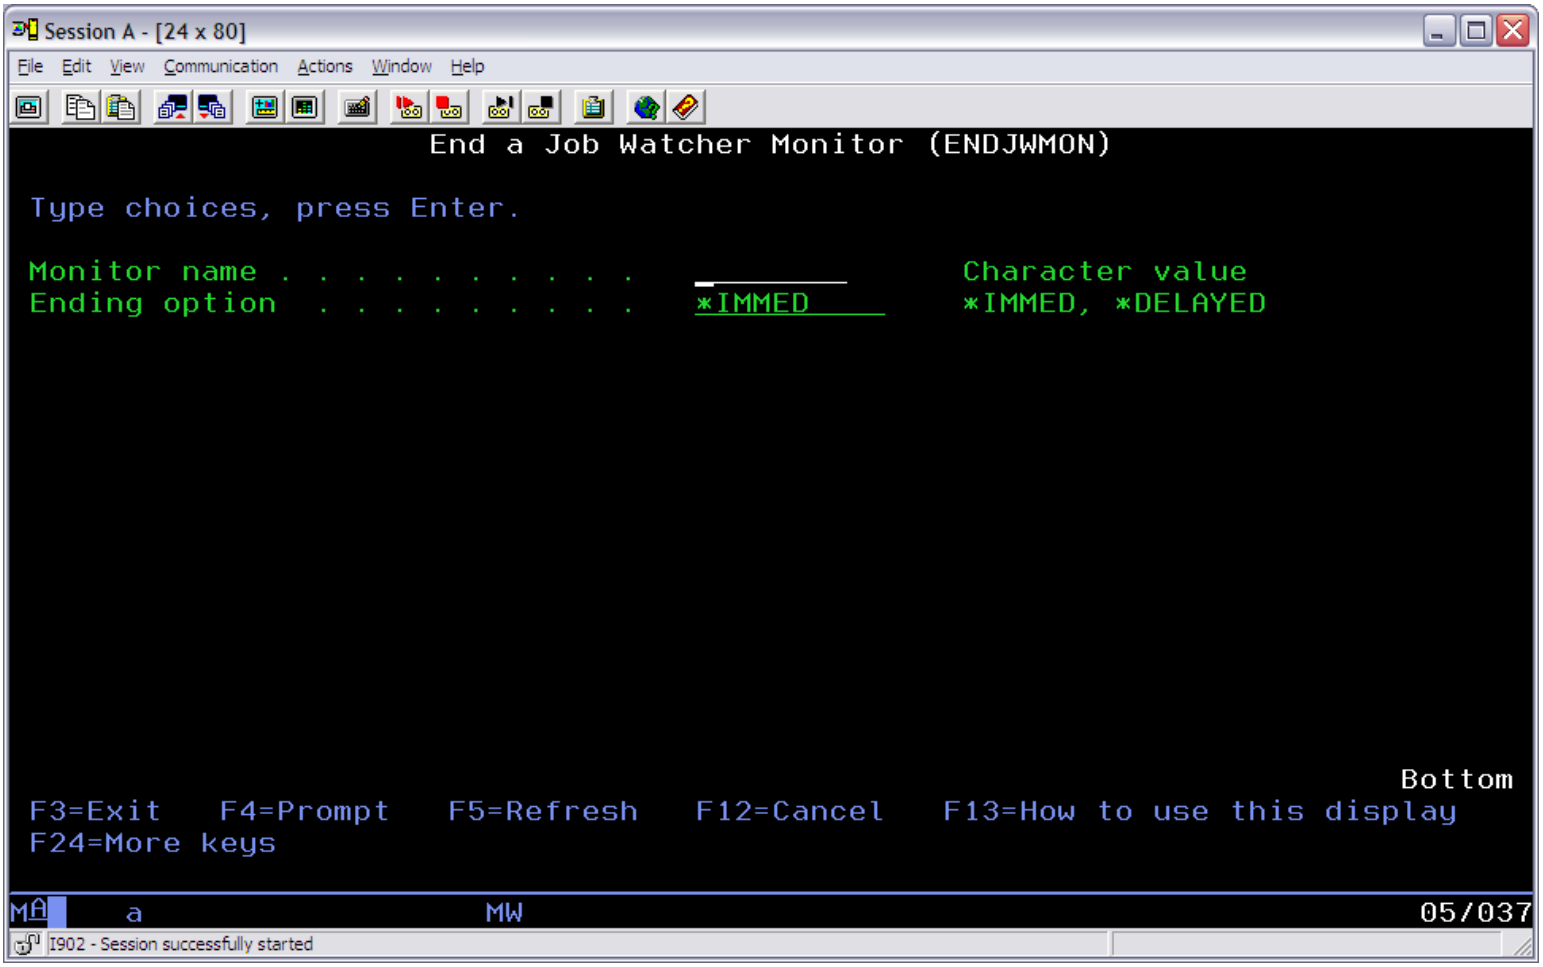

Delete Job Watcher monitor - DLTJWMON

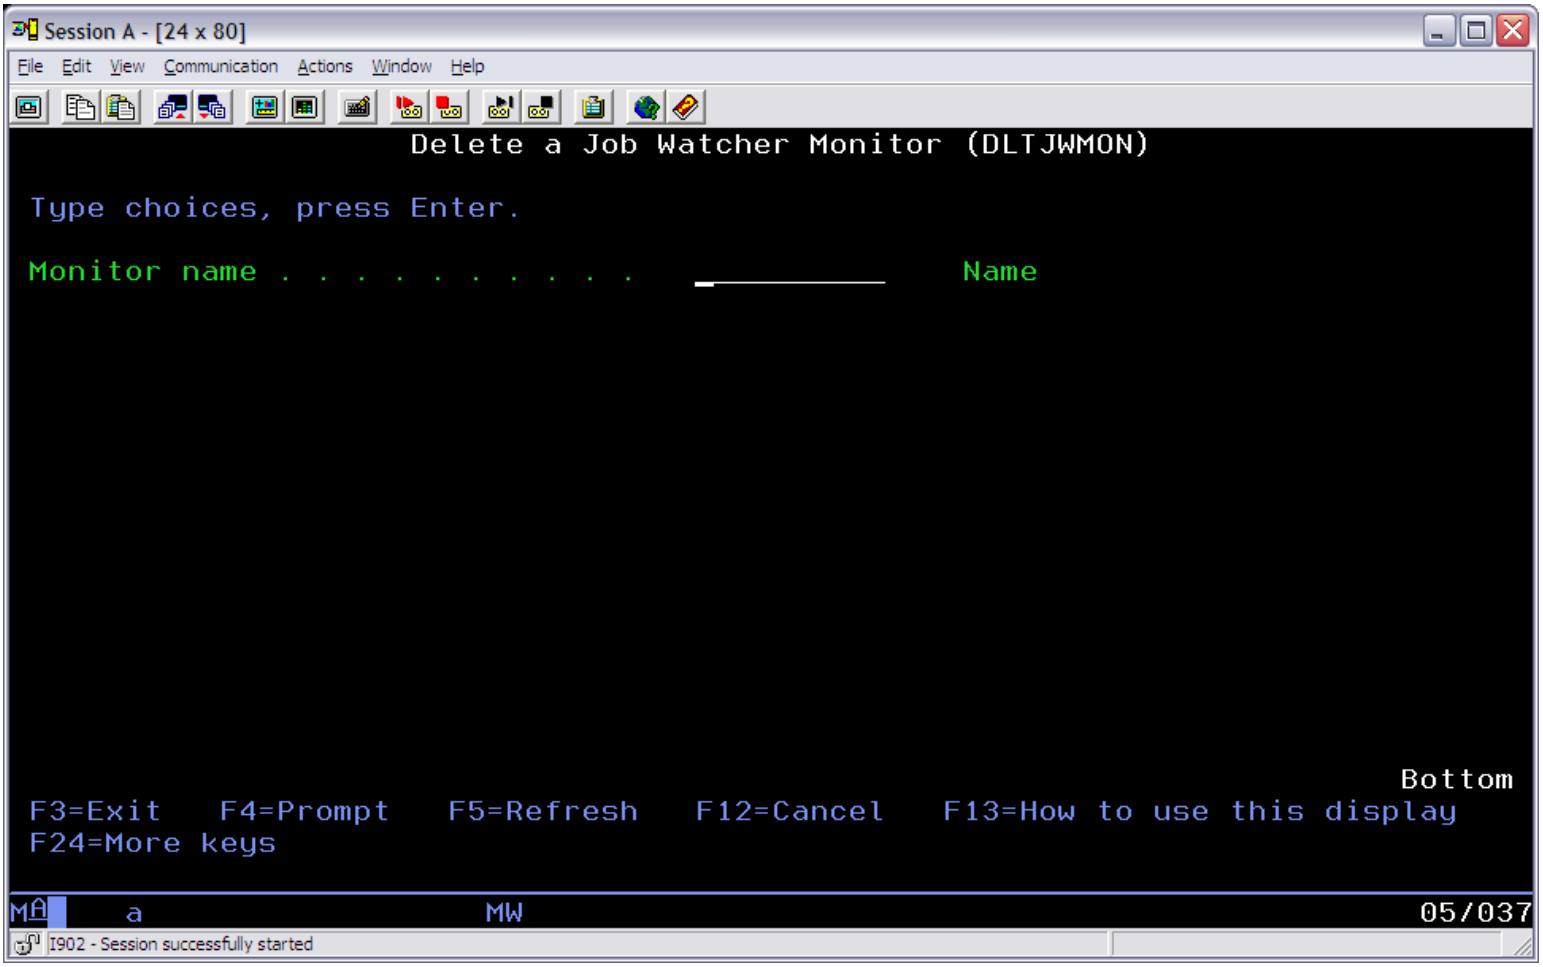

The delete Job Watcher monitor command will delete the monitor as well as any Job Watcher collection data collected by this monitor. If the monitor is active the ENDJWMON command will be first issued automatically to end the monitor.

PEX Monitors use commands:
STRPAMON, RLSPAMON, HLDPAMON, ENDPAMAON, DLTPAMON

Disk Watcher Monitors use commands:
STRDWMON, RLSDWMON, HLDDWMON, ENDDWMON, DLTDWMON

Appendix B – Automatically restarting a monitor after an IPL

In order to automatically restart an iDoctor Monitor after an IPL, it is necessary to add the required commands to the system startup program. To determine the system startup program, use this command:

DSPSYSVAL SYSVAL(QSTRUPPGM)

The STRJWMON, STRDWMON, STRPAMON commands are used to start or restart the desired monitor. These commands offer a parameter value of *SAME for most parameters which allows you to simply reuse the parameters the monitor had when it was initially created.

An example of one of the commands is:

QSYS/SBMJOB CMD(QIDRWCH/STRJWMON MONITOR(MYMON) COLLIB(MYLIB) DFNNAME(DFN1) COLNS(*SAME) STRGAP(*SAME) MAXSIZE(*SAME) OVRLAP(*SAME) TEXT(*SAME)) JOB(QSTRJWMON) RTGDTA(*JOBD) JOBD(QIDRGUI/QIDRBCH) JOBQ(QGPL/QIDRJW) OUTQ(*CURRENT) MSGQ(*NONE)

This SBMJOB command will restart the Job Watcher monitor MYMON in library MYLIB using definition DFN1. The job will be named QSTRJWMON. A command like this one will need to be added to the system startup program.

For more information on changing the IPL startup program: https://www-304.ibm.com/support/docview.wss?uid=nas8N1019476

Trademarks and Disclaimers

© IBM Corporation 1994-2006. All rights reserved.
References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AS/400	e-business on demand	i5/OS
Eserver	IBM	OS/400
eServer	IBM (logo)	System i5
	iSeries	DB2

Rational is a trademark of International Business Machines Corporation and Rational Software Corporation in the United States, other countries, or both.
Intel, Intel Logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.