

IBM i Performance Data Investigator Hands-On lab



Gottfried Schimunek

**Senior IT Architect
Application Design
IBM STG Software
Development
Lab Services**

**3605 Highway 52 North
Rochester, MN 55901**

**Tel 507-253-2367
Fax 845-491-2347**

IBM ISV Enablement

Gottfried@us.ibm.com

Review Performance Data

Estimated time

One to one hour and a half

What this exercise is about

This lab introduces you to the Performance Data Investigator (PDI) feature in IBM Systems Director Navigator for i. You will use PDI to review performance data collected previously on IBM i systems. You will learn the key features of PDI used in analyzing Collection Services performance data and learn the basic methodology needed to analyze system performance on IBM i.

What you should be able to do

At the end of this exercise, you should be able to:

- Use Systems Director Navigator for i to access the performance tasks in Performance Data Investigator (PDI).
- Use the Health Indicators data perspective of PDI to review the health of the key system resources at a high level.
- Use the Collection Services data perspective of PDI to review resource utilization and use drill down features to analyze memory usage and CPU usage by threads and tasks.
- Use the navigation features and filtering and sorting options of PDI to analyze performance data more effectively.
- Learn how to export performance data to a spreadsheet for further analysis.

Requirements

-

Instructor exercise overview

Exercise instructions

Section 1. Access IBM Systems Director Navigator for IBM i


This section of the lab uses the browser interface through IBM Systems Director Navigator for IBM i which is supported by IBM i 6.1 and later releases, to graphically review Collection Services. For this lab, use the data in

- **Library: OL66**
- **Performance Collection: TEXAS and CHICAGO**.....

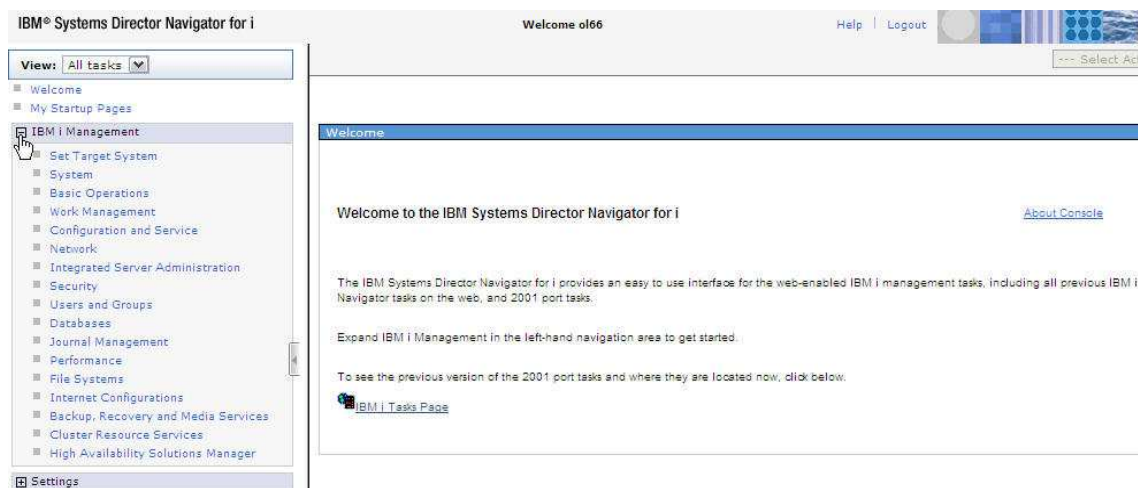
- ___ 1. From a Web Browser (like Internet Explorer or Mozilla Firefox) enter the following URL: <http://<systemname>:2001>.

A minor section heading

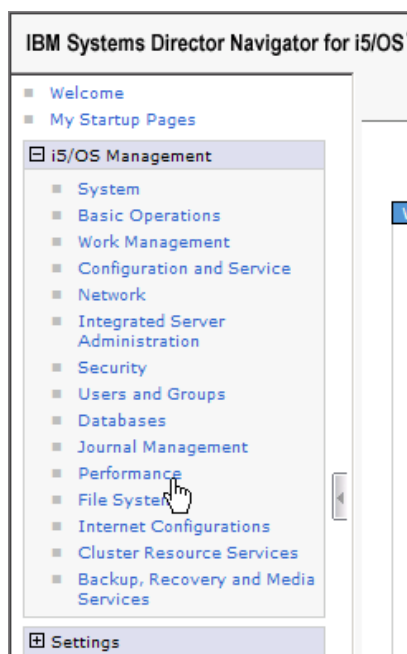
- ___ 2. You will see the following sign-on panel. Enter the user id and password assigned to you, and select the **Log in** button.



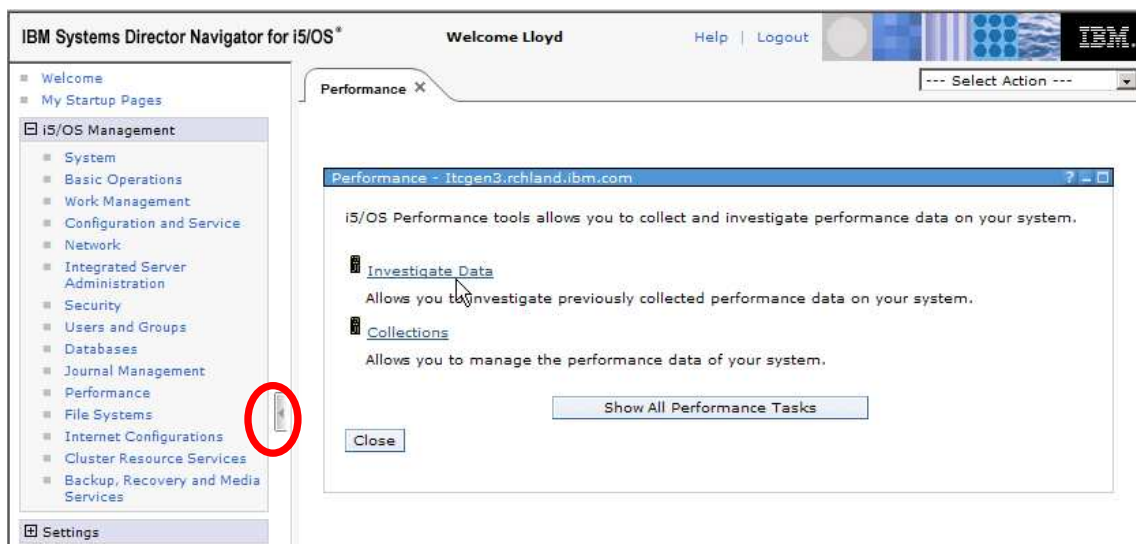
- ___ 3. Expand the **IBM i (i5/OS) Management** container on the left: Note: If using an i 6.1 system, you will see “i5/OS”. If using an i 7.1 system, you will see “IBM i” on the following screen.



4. You will see all the functions available through IBM System Director Navigator for IBM i (i5/OS). Select the **Performance** option from the navigation pane on the left:



5. On the next dialog, select **Investigate Data**:



Note: Click on the “Hide Navigation” tab on the left of the Performance display window. This will expand the area available to display graphics.

___ 6.

The Perspectives dialog is displayed. (In this lab we will use **Health Indicators and Collection Services** data only. Note that there are additional options to review performance data from Job Watcher, Disk Watcher, and Performance Explorer. You also have the option to save your own custom perspective containing the views that you find most valuable.)

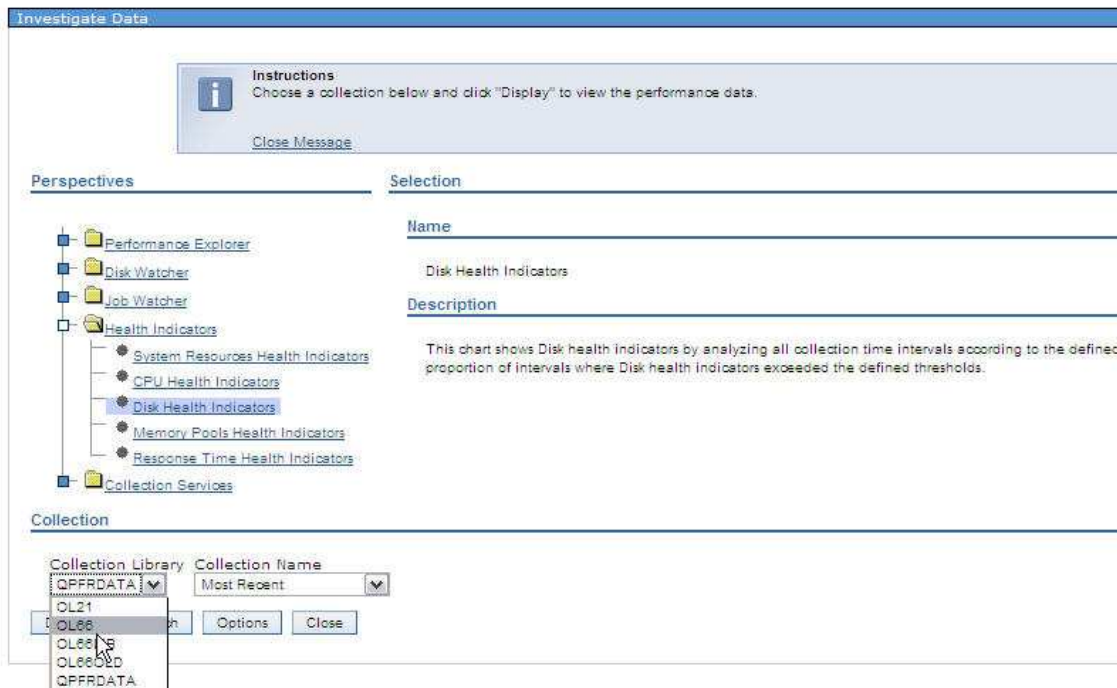
Click the **Health Indicators** link in the Perspectives dialog and then select the **Disk Health Indicators** link.



Section 2. Select the Health Indicators Performance data

This section of the lab uses the browser interface through IBM Systems Director Navigator for IBM i to select the data set for review. (Your IBM i environment may have multiple data collections).

- ___ 7. . Use the **Collection Library** pull down to identify the Collection library assigned to this lab:



- ___ 8. Next, use the **Collection Name** pull down to select the first collection assigned for this lab:

Instructions
Choose a collection below and click "Display" to view the performance data.
[Close Message](#)

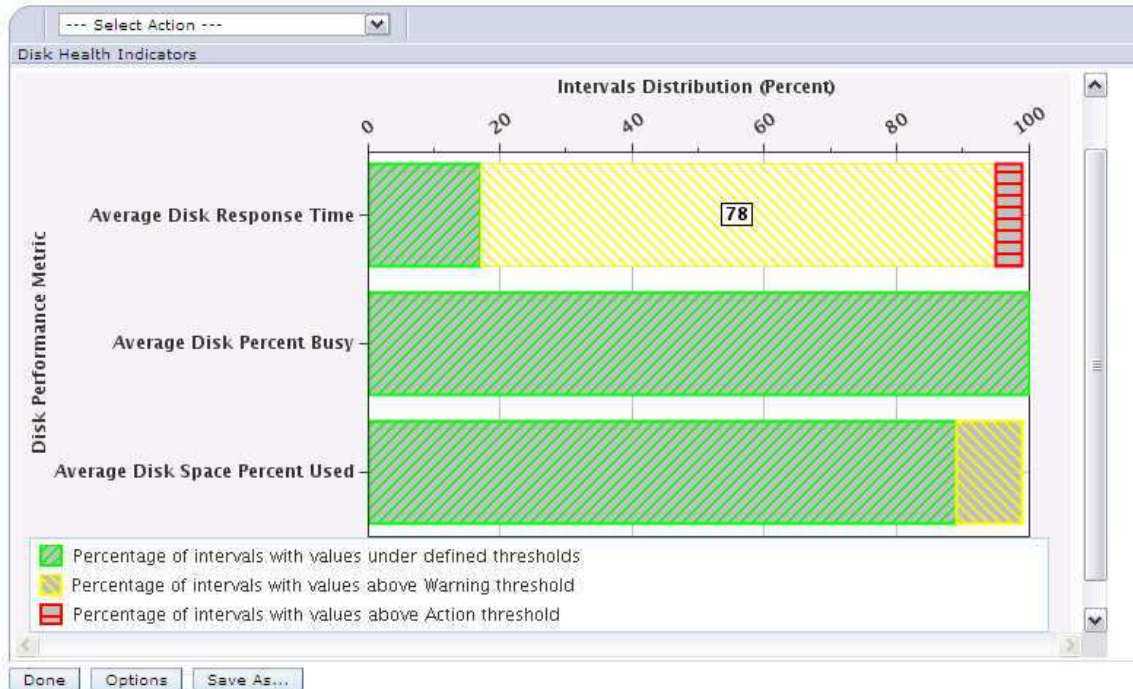
Perspectives	Selection				
<ul style="list-style-type: none"> Performance Explorer Disk Watcher Job Watcher Health Indicators <ul style="list-style-type: none"> System Resources Health Indicators CPU Health Indicators Disk Health Indicators Memory Pools Health Indicators Response Time Health Indicators Collection Services 	<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Name</th> </tr> </thead> <tbody> <tr> <td>Disk Health Indicators</td> </tr> </tbody> </table> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>This chart shows Disk health indicators by analyzing all collection proportion of intervals where Disk health indicators exceeded the</td> </tr> </tbody> </table>	Name	Disk Health Indicators	Description	This chart shows Disk health indicators by analyzing all collection proportion of intervals where Disk health indicators exceeded the
Name					
Disk Health Indicators					
Description					
This chart shows Disk health indicators by analyzing all collection proportion of intervals where Disk health indicators exceeded the					
Collection					
<table style="width: 100%;"> <tr> <td style="width: 30%;"> Collection Library OL68 </td> <td style="width: 70%;"> Collection Name Most Recent </td> </tr> <tr> <td style="text-align: center;"> <input type="button" value="Display"/> <input type="button" value="Search"/> </td> <td style="border: 1px solid gray;"> Most Recent All MDL810 (*CSFILE) CHICAGO (*CSFILE) TEXAS (*CSFILE) </td> </tr> </table>		Collection Library OL68	Collection Name Most Recent	<input type="button" value="Display"/> <input type="button" value="Search"/>	Most Recent All MDL810 (*CSFILE) CHICAGO (*CSFILE) TEXAS (*CSFILE)
Collection Library OL68	Collection Name Most Recent				
<input type="button" value="Display"/> <input type="button" value="Search"/>	Most Recent All MDL810 (*CSFILE) CHICAGO (*CSFILE) TEXAS (*CSFILE)				

- ___ 9. Click **Display** to display the disk health indicators for this collection.

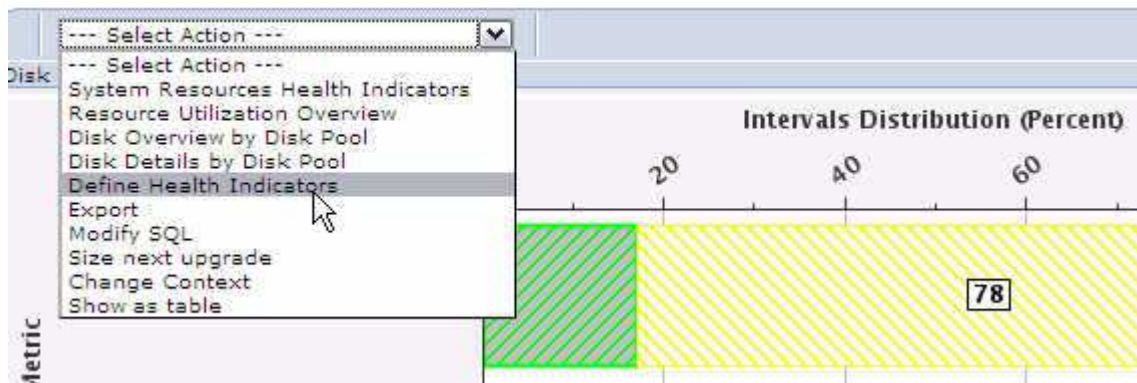
Section 3. Review Performance Health Indicators

This section of the lab uses the Performance Data Investigator to review performance health indicators for disk on the IBM i system.


- ___ 10. On the following pane, notice the Average Disk Response Time shows a high percentage of time when it exceeded the warning and action threshold. Click the Warning (Yellow) portion of the Average Disk Response Time bar and notice the percentage of time the system exceeded the warning threshold appears in the graph.



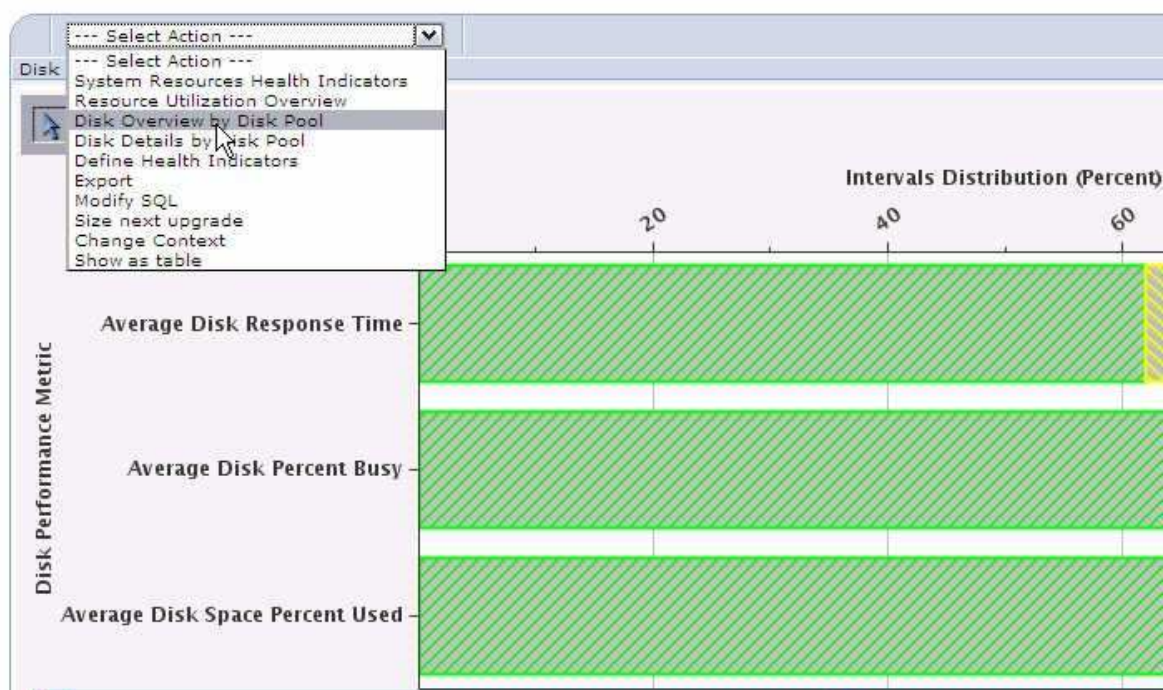
11. At this time, we do not know if the average disk response time is an issue until the Health Indicator Definitions are reviewed to determine how the warning and action thresholds are defined. Use the Action pull down at the top of the chart to select **Define Health Indicators** as shown here:



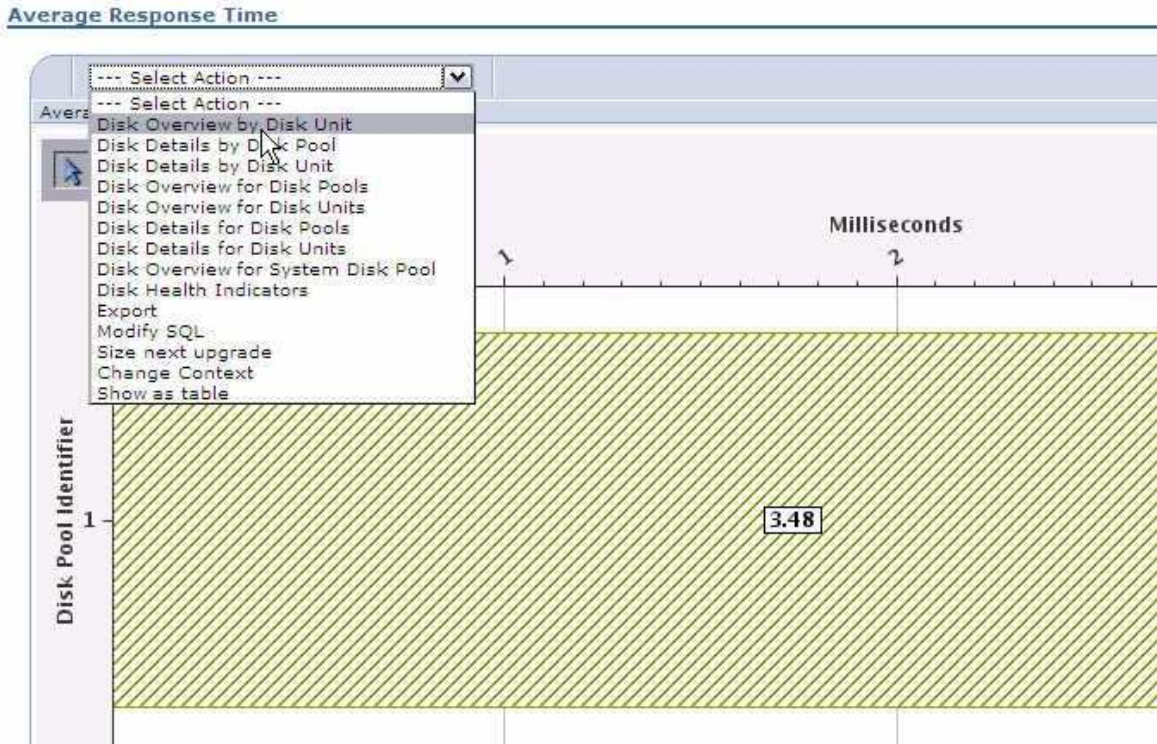
12. Review the Disk Health Indicator thresholds by clicking the **Average Disk Response Time** indicator. Notice the Warning threshold is set to 2 milliseconds and the Action threshold is set to 6 milliseconds. (Note: Milliseconds is not shown here! The development team is aware of this omission and we expect this to be fixed in a future release.)




- ___ 13. Change the Warning threshold to 4 milliseconds and click **Ok** at the bottom of this dialog. Notice that the average disk response time is now acceptable over 60% of the time, but the system still exceeds the Warning threshold over 30% of the time and exceeds the Action threshold about 4% of the time. Again, you can click each section of the bar graph to see how much time is spent at each level.
- ___ 14. Now, let's investigate the disk resources in more detail. Click the Action pull down at the top of the graph and select **Disk Overview by Disk Pool** as shown here:




- ___ 15. This system has only one disk pool (also called Auxiliary Storage Pool or ASP) so the graphs are not very interesting. However, it does show that the average disk response time and the average disk busy rate for the entire data collection when you click the bar graphs. Now drill down to view the disk units by selecting **Disk Overview by Disk Unit** from the Action pull down.



16. Notice the disk units with the longest average response time appear in the top graph. You can also scroll down to see the Disk Busy rate for the disk units. Notice some disk units are averaging over 7 millisecond response time.



- 17. Now, let's view detailed statistics for each disk unit by selecting Disk Details by Disk Unit from the Action pull down.



- 18. You should see the detailed view as shown below. Use the scroll bars to view the statistics for each disk unit.

Select	Disk Unit Name	Average Response Time (Milliseconds)	Percent Disk Busy	Reads Per Second	Writes Per Second	Operations Per Second
<input type="checkbox"/>	0054	7.37	7.62	10.89	10.93	21.82
<input type="checkbox"/>	0014	7.26	9.36	13.31	11.35	24.66
<input type="checkbox"/>	0060	7.06	7.99	12.35	12.04	24.39
<input type="checkbox"/>	0008	6.45	8.04	12.62	12.04	24.66
<input type="checkbox"/>	0052	6.44	7.9	11.76	10.03	21.79
<input type="checkbox"/>	0012	6.43	6.53	10.72	10.4	21.13
<input type="checkbox"/>	0042	5.91	9.47	12.45	27.1	39.55
<input type="checkbox"/>	0058	5.86	7.6	10.8	12.45	23.25
<input type="checkbox"/>	0032	5.22	8.58	10.68	31.72	42.4
<input type="checkbox"/>	0019	4.72	7.62	11.77	13.96	25.72
<input type="checkbox"/>	0059	4.65	7.98	10.98	16.83	27.81
<input type="checkbox"/>	0034	4.51	6.93	10.04	13.93	23.98

Total: 59 Filtered: 59

- ___ 19. Notice it can be difficult to view the data in this format. You may want to use the Export option from the Action pull down to export the data to a file. Note: We might also want to do more detailed performance analysis at this point to investigate why the disk response times do not meet expectations. This is beyond the scope of this lab, but one of the purposes of this lab was to show the capability that exists in the Performance Data Investigator for more detailed analysis if necessary. Note: You will get a chance to use the export function later in this lab and perform more detailed analysis in a future lab exercise.
- ___ 20. Now, let's end our disk health analysis and return to where we began. This can be done quickly by using the History tab at the top of the dialog. Click **History** and then select **Home** to return to the initial Health Indicators view.



- ___ 21. You may want to review the Memory Pools and Response Time health indicators and drill down in more detail as was done for the Disk health indicators. However, in order to complete this lab in the time allotted, this is left as a future exercise for the student.

Section 4. Review Collection Services Data Perspective

This section of the lab uses the browser interface through IBM Systems Director Navigator for IBM i to work with **Collection Services** to review more performance data.


- ___ 22. From the Perspectives pane, select **Collection Services** link (Notice several categories of data are available for investigation in this view. Also notice the Collection Library and Collection name you selected earlier are displayed near the bottom of the pane.)

Next, select **Resource Utilization Overview** as shown here:

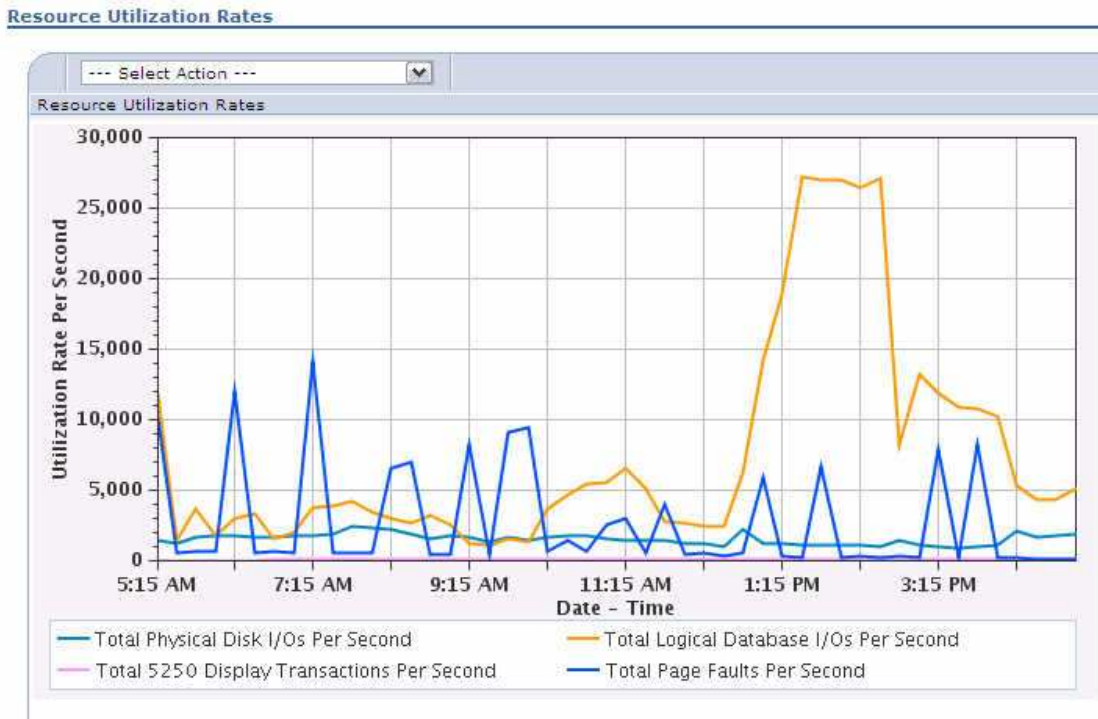


- ___ 23. Click the **Display** button at the bottom of the pane to display the resource utilization overview. A graph of Percent Disk Busy, Disk Space, and CPU Utilization appears at the top of this pane.

(Note: You may not always see the entire pane shown in this lab. There may be two vertical slide-bars on the right-side in addition to the horizontal slide bar at the bottom. To see all of the graphics/legends, you may have to use all the slide-bars).



- ___ 24. Use the outer slider bar to scroll down and view a second graph titled **Resource Utilization Rates**. Review the metrics that appear on each graph.



- ___ 25. Scroll up to the top of the pane and notice the navigation options that appear in upper left of each graph which looks like the following:



These options allow you to manipulate and view the graphs in different ways. Here is the purpose of each option from left to right:

Select - allows you to select specific data points on the graph

Pan – enables you to slide the window over the graph to view a different portion of the graph


Tool Tip – shows the value of a specific data point as you hover over the point

Zoom Region – allows you to draw a region on the graph to zoom into that portion of the graph


Zoom Out – zoom out after zooming into a specific region

Full Zoom Out – shows the full range of data points available on the graph

- ___ 26. Click the Full Zoom Out option (on the far right above) to show the entire Resource Utilization Percentages graph. The full graph should appear as shown here:

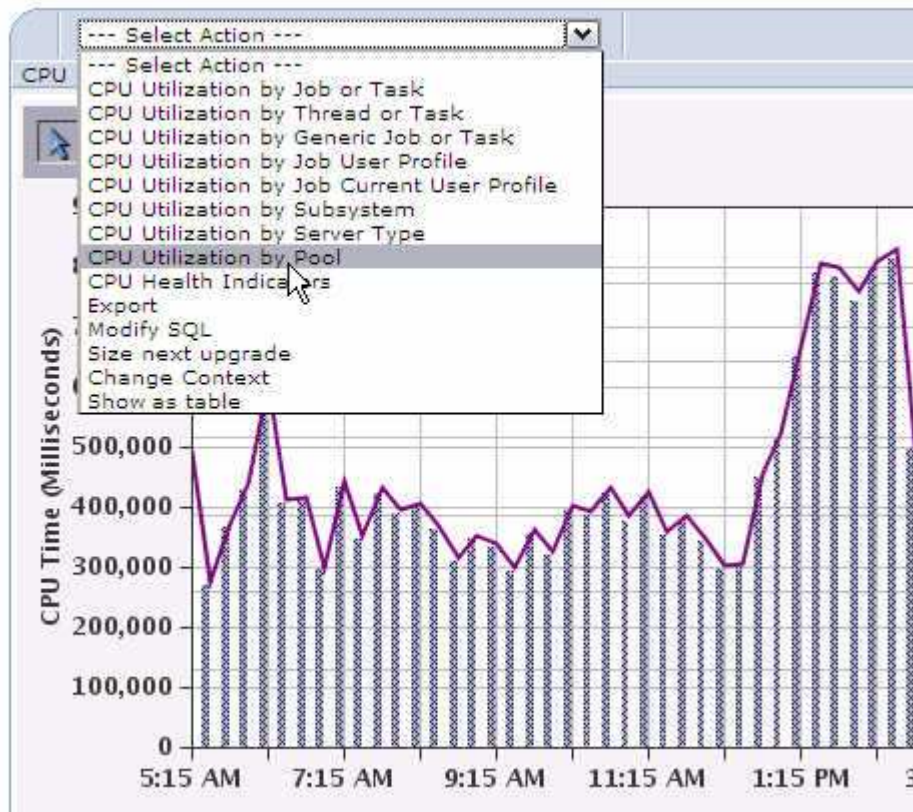


- ___ 27. Drill down to review CPU utilization in more detail by selecting **CPU Utilization Overview** from the Action pull down.




- ___ 28. Review the graph which shows CPU by milliseconds and by utilization rates. From the Action pull down, select **CPU Utilization by Pool** to investigate which pools contain jobs that are using the most CPU.

CPU Utilization Overview

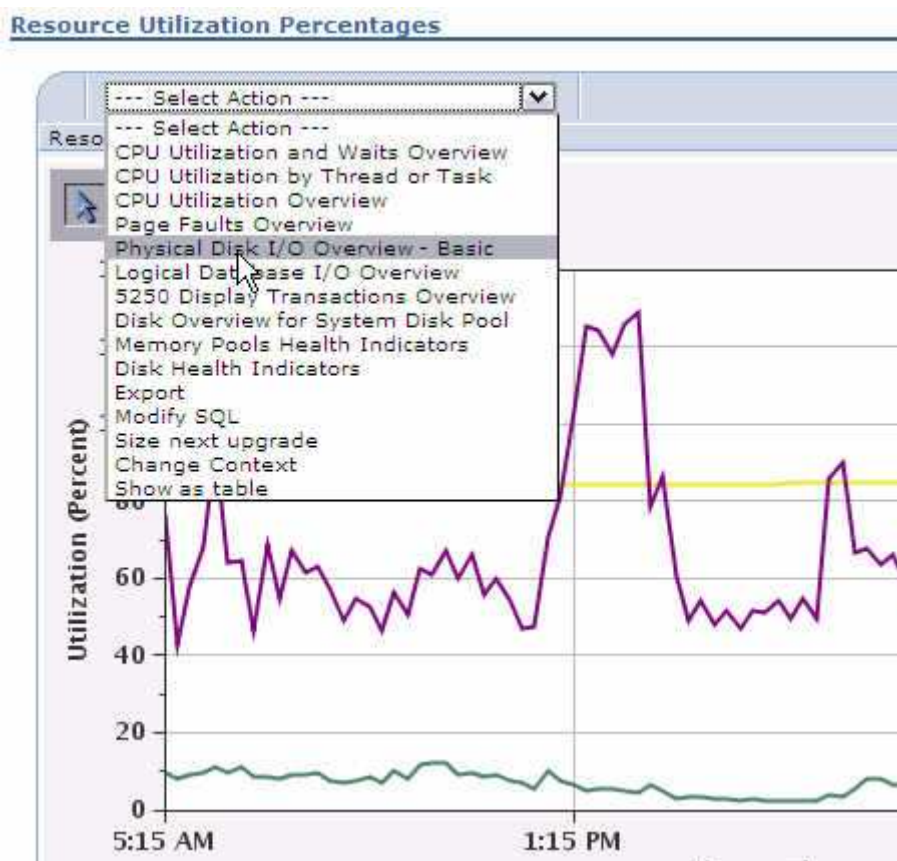


___ 29. Review the graph and notice that most of the CPU is being consumed by jobs running in Pool 2.


CPU Utilization by Pool



- ___ 30. Click **Done** on the CPU Utilization by Pool graph and then click **Done** again to close the CPU Utilization Overview graph. You will return to the Resource Utilization Percentages graph.
- ___ 31. Review disk I/O data by selecting **Physical Disk I/O Overview – Basic** from the Action pull down.




- ___ 32. From the Physical Disk I/O graph click the Tool Tips navigation option (circled below). Then move the cursor over a data point on the graph to see detailed information about the data point.




___ 33. Click Done to return to the Resource Utilization Percentages graph.

___ 34. Review main storage faulting by selecting **Page Faults Overview** from the Action pull down.


Resource Utilization Percentages




___ 35. Review the page faults graph which shows total faults and faults per second.



___ 36. Now, we will drill down and view which pools are doing the most faulting. Select **Page Faults by Pool** from the Action pull down.

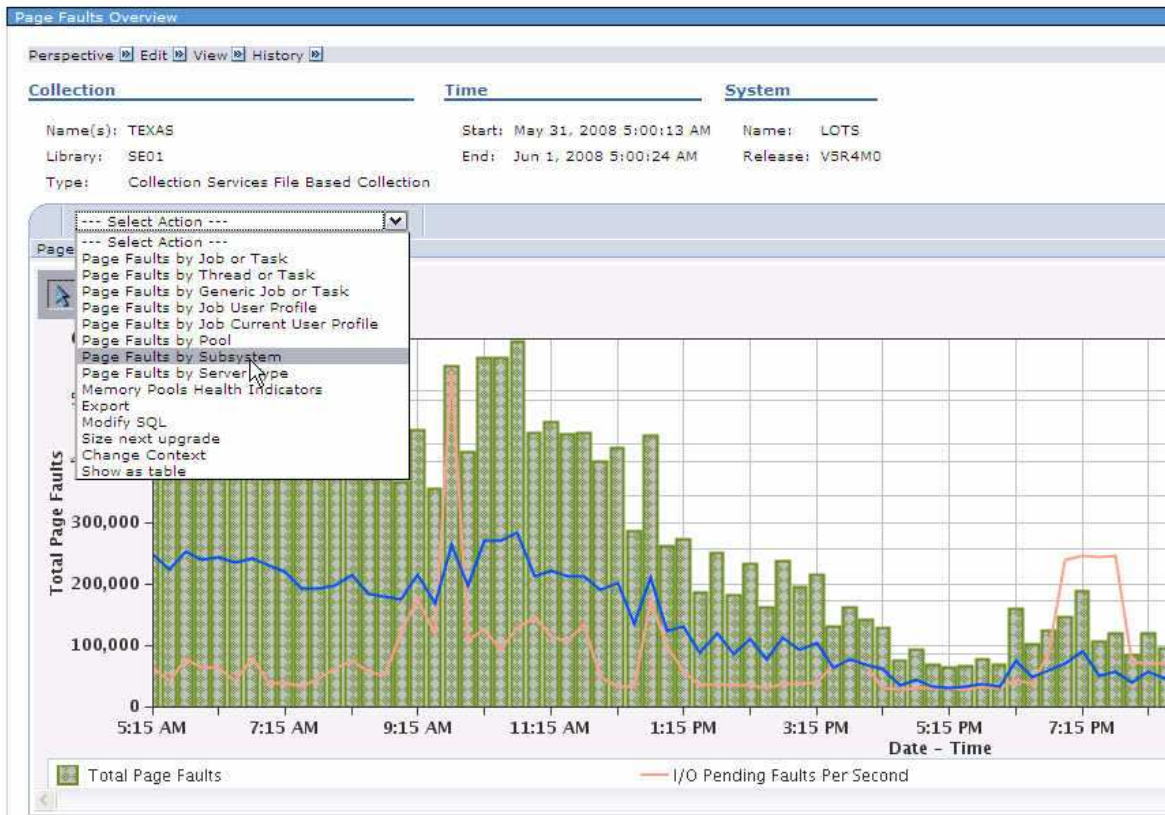


___ 37. A graph of page faults for each pool appears here. This graphs shows the page faults per second per individual pool. Here we can quickly learn that most of the faulting is occurring in Pool 2 on this system.

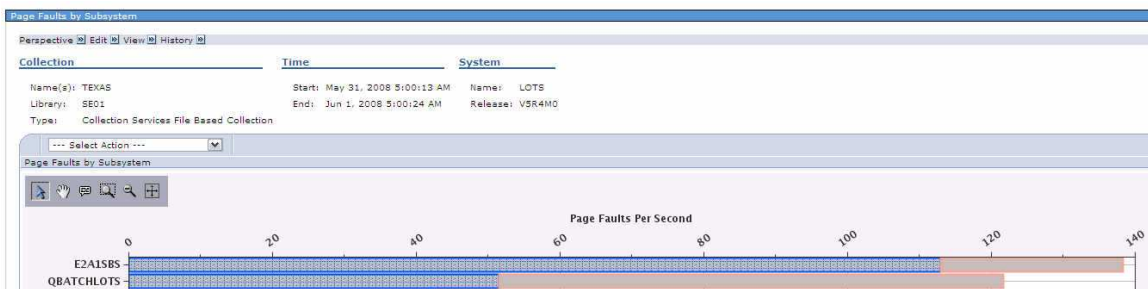


___ 38. Click Done to return to the Page Faults Overview graph.


___ 39. Now, we will look at faulting in another way by viewing which subsystems are doing the most faulting. Select **Page Faults by Subsystem** from the Action pull down.



___ 40. A graph of subsystems with the highest number faults appears here. This shows the average faulting rate for the entire collection.



___ 41. Now, we can examine faulting for a single subsystem. Select **Page Faults for One Subsystem** from the Action pull down.



- ___ 42. Enter the subsystem name E2A1SBS and click **OK** as shown below. This is the highest faulting subsystem.


Variable

Variable	Description	Value	Required
JESSYS	Subsystem Name	E2A1SBS	Yes
MINDTECEN	Century Digit		No
MINDTETIM	Interval Date And Time		No
MAXDTECEN	Century Digit		No
MAXDTETIM	Interval Date And Time		No
Collection Library		SE01	Yes
Collection Name		TEXAS	Yes

Page 1 of 1 Total: 7 Displayed: 7

OK Cancel

- ___ 43. Review the graph which shows faulting data for subsystem E2A1SBS. This graph shows total page faults and faults per second.

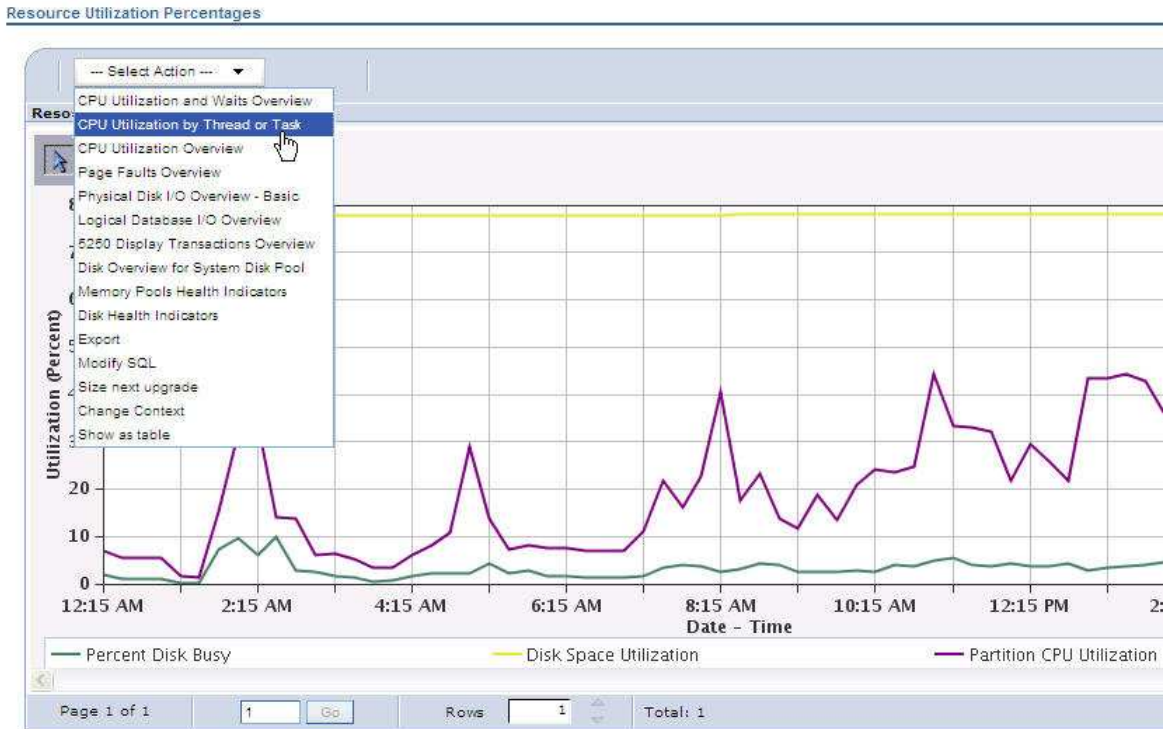


- ___ 44. Click the History tab at the top of the pane and select Home to return to the Investigate Data pane.
- ___ 45. Now we will analyze performance of a second collection. Review the Resource Utilization Overview of the second collection by selecting it as shown below.

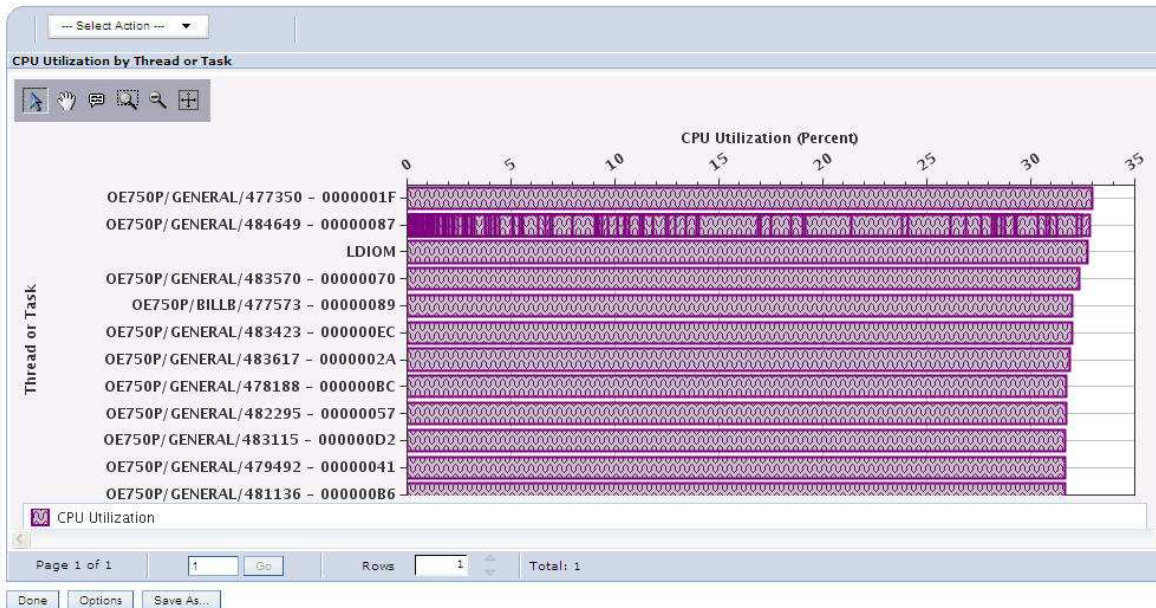


- ___ 46. Click the **Display** button at the bottom of the pane to display the resource utilization overview for the new collection.
- ___ 47. From the Resource Utilization graph, let's drill down and look at CPU utilization by individual threads. Select **CPU Utilization by Thread or Task** from the Action pull down as shown below.

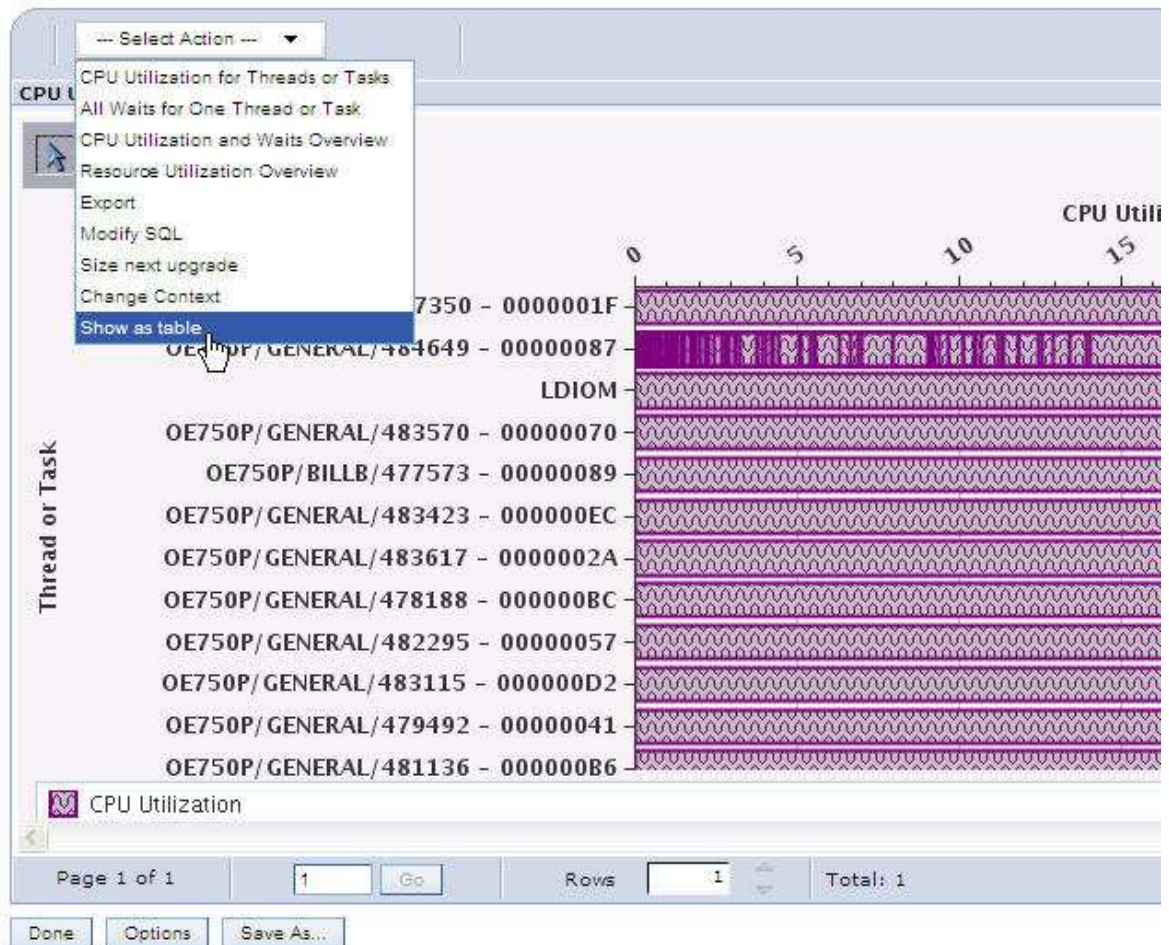
Note: This task could take several seconds or even minutes because this is a query against every thread and task running on the system when the data was collected.



48. After the query completes, you will see a list of threads/tasks by CPU utilization as shown here.



49. More information about these jobs can be viewed by looking at the performance data in table format. View the table data by selecting **Show as table** from the Action pull down.



- ___ 50. The CPU and related performance data appears in a table as shown. If you would like to review the data in more detail, you can use the scroll bar to see additional metrics in the table or scroll down to see additional threads and tasks.

Select	Name ^	Job User ^	Job Number ^	Thread Identifier ^	CPU Utilization (Percent) ▼	CPU Time (Milliseconds) ^	Scaled CPU Utilization (Percent) ^	Scaled CPU Time (Milliseconds)
<input type="checkbox"/>	OE750P	GENERAL	477350	0000001F	32.97	69239	0	
<input type="checkbox"/>	OE750P	GENERAL	484649	00000087	32.85	70960	0	
<input type="checkbox"/>	LDIOM				32.78	983	0	
<input type="checkbox"/>	LDIOM				32.49	975	0	
<input type="checkbox"/>	OE750P	GENERAL	483570	00000070	32.35	66975	0	
<input type="checkbox"/>	LDIOM				32.22	158501	0	
<input type="checkbox"/>	OE750P	BILLB	477573	00000089	32.01	47060	0	
<input type="checkbox"/>	OE750P	GENERAL	483423	000000EC	32.01	67227	0	
<input type="checkbox"/>	OE750P	GENERAL	483617	0000002A	31.9	66982	0	
<input type="checkbox"/>	OE750P	GENERAL	478188	000000BC	31.78	63841	0	
<input type="checkbox"/>	OE750P	GENERAL	482295	00000057	31.74	67602	0	
<input type="checkbox"/>	OE750P	GENERAL	483115	000000D2	31.7	68465	0	

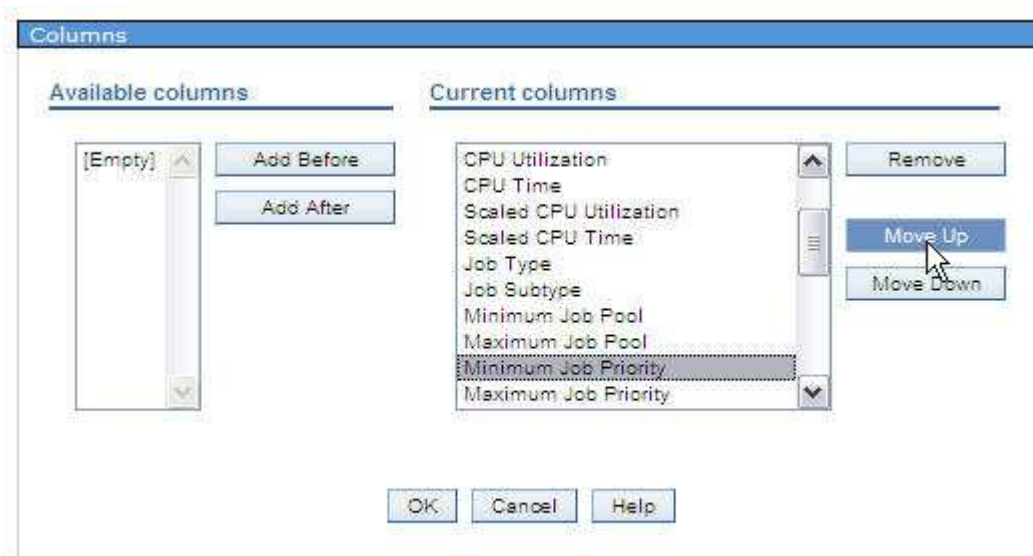
Total: 20646 Filtered: 20646

51. Let's reorder how the data is displayed and bring the priority of each job closer to the CPU metrics. Select **Columns...** from the Action pull down to begin this process.

Select	Name ^	Job User ^	Job Nu	on t)	CPU Time (Milliseconds) ^
<input type="checkbox"/>	OE750P	GENERAL	477350	32.97	69239
<input type="checkbox"/>	OE750P	GENERAL	484649	32.85	70960
<input type="checkbox"/>	LDIOM			32.78	983
<input type="checkbox"/>	LDIOM			32.49	975
<input type="checkbox"/>	OE750P	GENERAL	483570	32.35	66975
<input type="checkbox"/>	LDIOM			32.22	158501
<input type="checkbox"/>	OE750P	BILLB	477573	32.01	47060
<input type="checkbox"/>	OE750P	GENERAL	483423	32.01	67227
<input type="checkbox"/>	OE750P	GENERAL	483617	31.9	66982
<input type="checkbox"/>	OE750P	GENERAL	478188	31.78	63841
<input type="checkbox"/>	OE750P	GENERAL	482295	31.74	67602
<input type="checkbox"/>	OE750P	GENERAL	483115	31.7	68465

Total: 20646 Filtered: 20646

- ___ 52. Scroll down through the list in Current Columns and select Minimum Job Priority. This is the minimum priority number a thread can run at which means it is the highest priority for that thread. Click the **Move Up** button 6 times to move the Minimum Job Priority column right after the CPU Time column. Then click **Ok**.



Note: An alternative way to do this is to click **Remove** to remove the Minimum Job Priority column from the list of current columns. Then click on CPU Time and then the **Add After** button to add this column after CPU Time.

- ___ 53. The minimum job priority now appears in the initial view as shown. Now click on the CPU Time column TWICE as shown below to sort the CPU Time used in descending order. (Clicking only once will sort by ascending order). We now have threads and tasks that used the most CPU at the top of the list. This will be our secondary sort column

Select	Name ^	Job User ^	Job Number ^	Thread Identifier ^	CPU Utilization (Percent) ▾	CPU Time (Milliseconds) ↗	Minimum Job Priority ^
<input type="checkbox"/>	OE760P	GENERAL	477350	0000001F	32.97	69239	050
<input type="checkbox"/>	OE750P	GENERAL	484649	00000087	32.85	70960	050
<input type="checkbox"/>	LDIOM				32.78	983	000
<input type="checkbox"/>	LDIOM				32.49	975	000
<input type="checkbox"/>	OE750P	GENERAL	483570	00000070	32.35	66975	050
<input type="checkbox"/>	LDIOM				32.22	158601	000
<input type="checkbox"/>	OE760P	BILLB	477573	00000089	32.01	47060	050
<input type="checkbox"/>	OE750P	GENERAL	483423	000000EC	32.01	67227	050
<input type="checkbox"/>	OE760P	GENERAL	483617	0000002A	31.9	66982	050
<input type="checkbox"/>	OE750P	GENERAL	478188	000000BC	31.76	63841	050
<input type="checkbox"/>	OE750P	GENERAL	482295	00000057	31.74	67602	050
<input type="checkbox"/>	OE760P	GENERAL	483115	000000D2	31.7	68465	050

Total: 20646 Filtered: 20646

54. Now sort by Minimum Job Priority. Click on the up arrow ONCE for that column as shown below to sort job priority in ascending order. This is our primary sort column.

Select	Name ^	Job User ^	Job Number ^	Thread Identifier ^	CPU Utilization (Percent) 2 ▾	CPU Time (Milliseconds) 1 ↗	Minimum Job Priority ^
<input type="checkbox"/>	CFINT01				1.97	5122456	000
<input type="checkbox"/>	ZIECTL	QPGMR	477094	00000070	1.33	3042199	040
<input type="checkbox"/>	ODSM600PAA	QSYSOPR	477101	00000068	1.31	2827229	020
<input type="checkbox"/>	LDIOM				29.36	1936447	000
<input type="checkbox"/>	TP4	APLUSCTL6	477082	0000004A	0.69	1525034	050
<input type="checkbox"/>	PETRABS1	PETRAB	477422	00000077	3.99	978339	020
<input type="checkbox"/>	IMRSTP	BILLB	481173	000000B7	27.65	968979	050
<input type="checkbox"/>	OSTRANP	APCRMCTL6	477092	0000009E	0.42	957154	050
<input type="checkbox"/>	IMRSTP	BILLB	480479	000000C6	29.93	953679	020
<input type="checkbox"/>	IMRSTP	BILLB	480119	00000086	28.2	951761	050
<input type="checkbox"/>	APOSBKR6	CHUCKP	477093	00000049	0.37	846603	040
<input type="checkbox"/>	TP3	APLUSCTL6	477081	0000000D	0.3	638608	050

Total: 20646 Filtered: 20646

- ___ 55. We now have the data sorted for each priority by CPU time used. You can scroll down to see the data for each job as shown, but this can take an extensive amount of time with this browser-based interface.

Select	Name	Job User	Job Number	Thread Identifier	CPU Utilization (Percent)	CPU Time (Milliseconds)	Minimum Job Priority	Scaled CPU Utilization (Percent)
<input type="checkbox"/>	CFINT01				1.97	5122458.000		0
<input type="checkbox"/>	LDIOM				29.36	1836447.000		0
<input type="checkbox"/>	CFINT03				0.09	227709.000		0
<input type="checkbox"/>	CFINT02				0.07	189413.000		0
<input type="checkbox"/>	LDIOM				32.22	158501.000		0
<input type="checkbox"/>	QPFPRADJ	QSYS	432807	00000001	0.08	145228.000		0
<input type="checkbox"/>	IOCMETHERLIND01				0.04	101323.000		0

- ___ 56. When doing performance analysis or capacity planning, you will often discover that many of the threads/tasks in the collection will use very little amounts of CPU. Sometimes threads are started and never do any meaningful work. To reduce the number of threads and tasks that are analyzed, let's filter out or remove all threads and tasks that have used less than 10 milliseconds of CPU time. Select **Table Actions >> Show Filter Row** as shown below to start the filtering process.

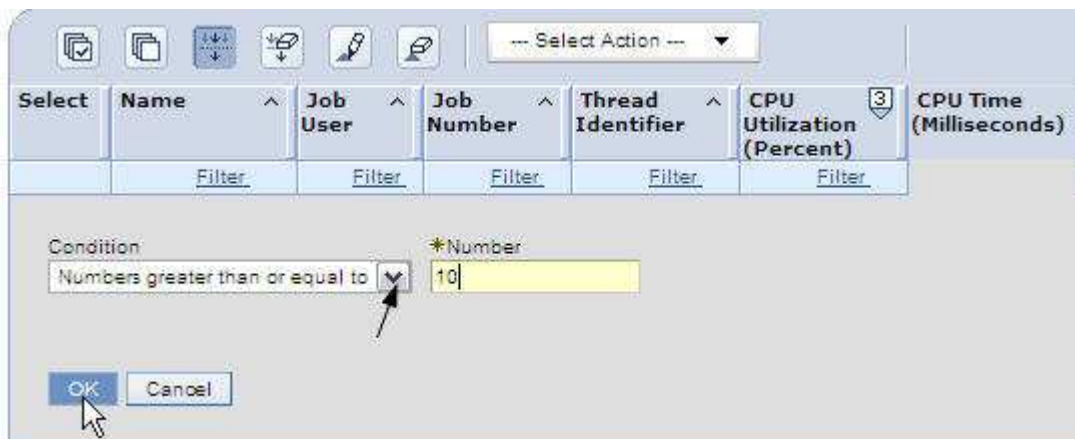
Select	Name	Job User	Job Number	Thread Identifier	CPU Utilization (Percent)	CPU Time (Milliseconds)	Minimum Job Priority	Scaled CPU Utilization (Percent)
<input type="checkbox"/>	CFINT01				1.97	5122458.000		0
<input type="checkbox"/>	LDIOM				29.36	1836447.000		0
<input type="checkbox"/>	CFINT03				0.09	227709.000		0
<input type="checkbox"/>	CFINT02				0.07	189413.000		0
<input type="checkbox"/>	LDIOM				32.22	158501.000		0
<input type="checkbox"/>	QPFPRADJ	QSYS	432807	00000001	0.08	145228.000		0
<input type="checkbox"/>	IOCMETHERLIND01				0.04	101323.000		0
<input type="checkbox"/>	LDIOM				0.02	89226.000		0
<input type="checkbox"/>	LDIOM				0.02	83305.000		0
<input type="checkbox"/>	IOCMETHERLIND00				0.02	78737.000		0
<input type="checkbox"/>	QYPSJSVR	QYPSJSVR	433032	00000020	0.02	39046.000		0
<input type="checkbox"/>	LDIOM				9.14	34003.000		0

Total: 20646 Filtered: 20646

- ___ 57. Now click the Filter link in the CPU Time column as shown.

Select	Name	Job User	Job Number	Thread Identifier	CPU Utilization (Percent)	CPU Time (Milliseconds)	Minimum Job Priority
<input type="checkbox"/>	CFINT01				1.97	122466	000
<input type="checkbox"/>	LDIOM				29.36	1836447	000
<input type="checkbox"/>	CFINT03				0.09	227708	000

58. Select **Numbers greater than or equal to** from the Condition pull down. Enter 10 in the Number field and Click Ok as shown. This will filter the threads and tasks and keep only those that used 10 milliseconds or more of CPU time.



Notice when the list is displayed again, look near the bottom of the pane to see that the total number of entries was 20,646 but the new filtered number is 9,647 which means we eliminated almost 11,000 threads and tasks just by filtering out those that used less than 10 milliseconds of CPU.

59. Now let's export the remaining CPU data to .csv format which can be easily used with Excel. Select **Table Actions >> Select All** to select all remaining data for the export operation.

Select	Name	Job User	CPU Utilization (percent)	CPU Time (Milliseconds)	Minimum Job Priority
<input type="checkbox"/>	CFINT01		1.97	5122458	000
<input type="checkbox"/>	LDIOM		29.38	1836447	000
<input type="checkbox"/>	CFINT03		0.09	227709	000
<input type="checkbox"/>	CFINT02		0.07	189413	000
<input type="checkbox"/>	LDIOM		32.22	158501	000
<input type="checkbox"/>	QFFRADJ	QSYS		145228	000
<input type="checkbox"/>	IOCMETHERLIND01			101323	000
<input type="checkbox"/>	LDIOM			99226	000
<input type="checkbox"/>	LDIOM			93305	000
<input type="checkbox"/>	IOCMETHERLIND00		0.03	78737	000

__ 60. Select Export from the Action pull down.

Select	Name	Job User	CPU Utilization (percent)	CPU Time (Milliseconds)	Minimum Job Priority
<input checked="" type="checkbox"/>	CFINT01		1.97	5122458	000
<input checked="" type="checkbox"/>	LDIOM		29.38	1836447	000
<input checked="" type="checkbox"/>	CFINT03		0.09	227709	000
<input checked="" type="checkbox"/>	CFINT02		0.07	189413	000
<input checked="" type="checkbox"/>	LDIOM		32.22	158501	000
<input checked="" type="checkbox"/>	QFFRADJ	QSYS		145228	000
<input type="checkbox"/>	IOCMETHERLIND01			101323	000
<input type="checkbox"/>	LDIOM			99226	000
<input type="checkbox"/>	LDIOM			93305	000
<input type="checkbox"/>	IOCMETHERLIND00		0.03	78737	000

__ 61. Select .csv format and **All data** and click Ok as shown below.

Title

CPU Utilization by Thread or Task

Format

Comma Delimited (*.csv)

Data Range

All data
 Displayed data
 User-defined range:

Data Series

Name
 Job User
 Job Number
 Thread Identifier
 CPU Utilization

First Record Number: 1 1,2,3...9,647
 Last Record Number: 9647 1,2,3...9,647

OK Cancel

62. When the browser Save dialog appears, save the file to your laptop. Name the file OL66Lab1.csv. Note: if your browser does not allow pop-ups to be displayed, you may need to allow pop-ups from this IBM i system.

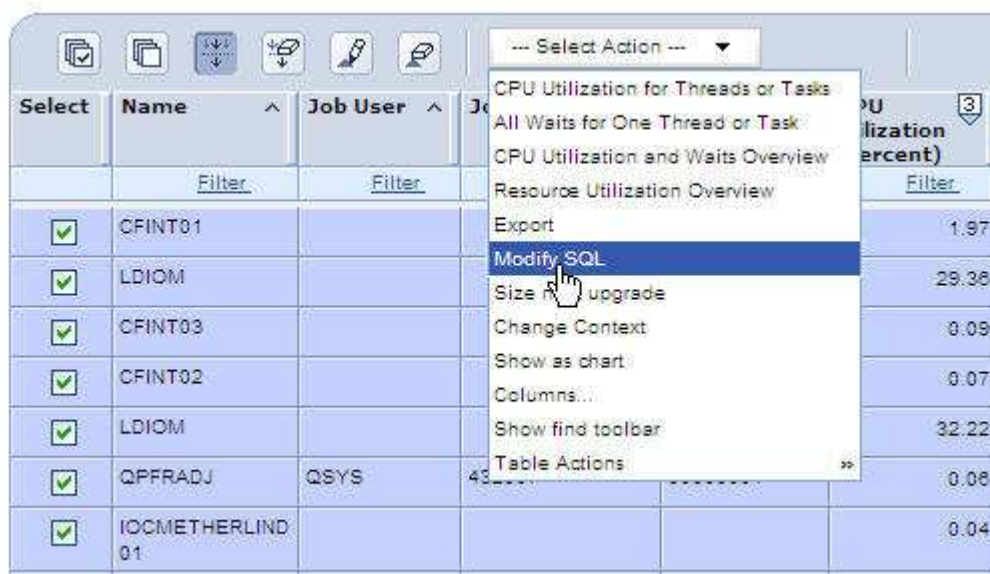
63. Open the file you just saved and review the data that was exported here. You may need to adjust the column widths to see the data as shown.

	A	B	C	D	E	F	G
1	CPU Utilization by Thread or Task						
2	Name	Job User	Job Number	Thread Identifier	CPU Utilization	CPU Time	Minimum Job Priority
3	CFINT01				1.97	5122456	0
4	LDIOM				29.36	1836447	0
5	CFINT03				0.09	227709	0
6	CFINT02				0.07	189413	0
7	LDIOM				32.22	158501	0
8	QPFRAJ	QSYS	432807	1	0.06	145228	0
9	IOCMETHERLIND 01				0.04	101323	0
10	LDIOM				30.35	89226	0
11	LDIOM				11.24	83305	0
12	IOCMETHERLIND 00				0.03	78737	0

64. Scroll down through the rows of the table paying attention to columns F and G, CPU Time and Minimum Job Priority. You can see that approximately 1540 entries have a priority of 0 which indicates it is a system thread/task. If you continue to scroll down, you will find thousands of priority 10 jobs. The priority 20 jobs begin in row 6068. Notice each task is sorted by CPU Time Used within each priority as

you would expect since this is how we sorted the data in Performance Data Investigator.

- ___ 65. From here, those with a minimal level of Excel expertise could graph the CPU time used for system jobs (priority 0), high priority jobs (i.e. priority 1-20), medium priority jobs (i.e. priority 21-40), and low priority jobs (41-99). This will be left as an optional exercise for the student to be done at your discretion.
- ___ 66. Return to the Performance Data Investigator window and select Modify SQL from the Action pull down as shown.



- ___ 67. Inspect the SQL that gets created to produce the view shown (and appreciate all the work that is done for you!!!). Later in the class you will learn where the columns (fields) for this query are documented. Depending on your knowledge of SQL, you might wish to copy and/or modify the query to produce another view of the data. However, the sort and filter options that you just used gives you a simpler way to modify the view and will meet your needs in many situations.

SQL Statement

Reset

```

SELECT JBNAME, JBUSER, JENBR, JETHID, CASE WHEN TOTSEC = 0 THEN 0 ELSE CPUTOT/(DOUBLE(TOTSEC) + 10
- MAXNUMVP) END AS PCTCPU, CPUTOT, MINJBPRTY, CASE WHEN TOTSEC = 0 THEN 0 ELSE
SCPUTOT/(DOUBLE(TOTSEC) + 10 - MAXNUMVP) END AS PCTSCPU, JBTYP, JBSTYP, MINJBPOOL,
MAXJBPOOL, MAXJBPRTY, MINJBSSYS, MAXJBSSYS, MINJBCUSR, MAXJBCUSR, JBSJNM, JBSJNB,
MINJBSVRT, MAXJBSVRT, CASE WHEN JBTYPE = 'V' THEN JBNAME ELSE TRIM(JBNAME) || '/' || TRIM(JBUSER)
|| '/' || TRIM(JENBR) || ' - ' || TRIM(JETHID) END AS FULLNAME, JBSLICNT, JBNAME, JWIDE
FROM (
SELECT JBNAME, JBUSER, JENBR, JETHID, JBTYPE, JBSTYP, MIN(JBPOOL) AS MINJBPOOL, MIN(JBPRTY) AS
MINJBPRTY, MIN(JBCUSR) AS MINJBCUSR, MIN(JBSSYS) AS MINJBSSYS, MAX(JBPOOL) AS MAXJBPOOL,
MAX(JBPRTY) AS MAXJBPRTY, MAX(JBCUSR) AS MAXJBCUSR, MAX(JBSSYS) AS MAXJBSSYS, JBSJNM, JBSJNB,
JBSJNB, MIN(JBSVRT) AS MINJBSVRT, MAX(JBSVRT) AS MAXJBSVRT, MAX(JBSLICNT) AS JBSLICNT, MAX(JBNAME)
AS JBNAME, SUM(JBCPU) AS CPUTOT, SUM(JBSCPU)/1000 AS SCPUTOT, MAX(NUMVP) AS MAXNUMVP, ROUND(CASE
WHEN MIN(DTETIM) = MAX(DTETIM) THEN MAX(INTSEC) ELSE TIMEDIFF(2, CAST(CAST(CASE WHEN
MAX(DTECEN) = '1' THEN '20' ELSE '19' END || SUBSTR(MAX(DTETIM), 1, 2) || '-' ||
SUBSTR(MAX(DTETIM), 3, 2) || '-' || SUBSTR(MAX(DTETIM), 5, 2) || '-' || SUBSTR(MAX(DTETIM), 7, 2)
|| '-' || SUBSTR(MAX(DTETIM), 9, 2) || '-' || SUBSTR(MAX(DTETIM), 11, 2) || '.000000' AS TIMESTAMP)
- MIN(CAST(CASE WHEN DTECEN = 1 THEN '20' ELSE '19' END || SUBSTR(DTETIM, 1, 2) || '-' ||
SUBSTR(DTETIM, 3, 2) || '-' || SUBSTR(DTETIM, 5, 2) || '-' || SUBSTR(DTETIM, 7, 2) || '-' ||
SUBSTR(DTETIM, 9, 2) || '-' || SUBSTR(DTETIM, 11, 2) || '.000000' AS TIMESTAMP) - (INTSEC *
1000000) MICROSECONDS) AS CHAR(22))) END, 6) AS TOTSEC, JWIDE
FROM (

```

 Allow collection choice

OK

Cancel

___ 68. Click **Cancel** on the above dialog and click **Logout** (upper right in your browser) to end this PDI session.

End of exercise