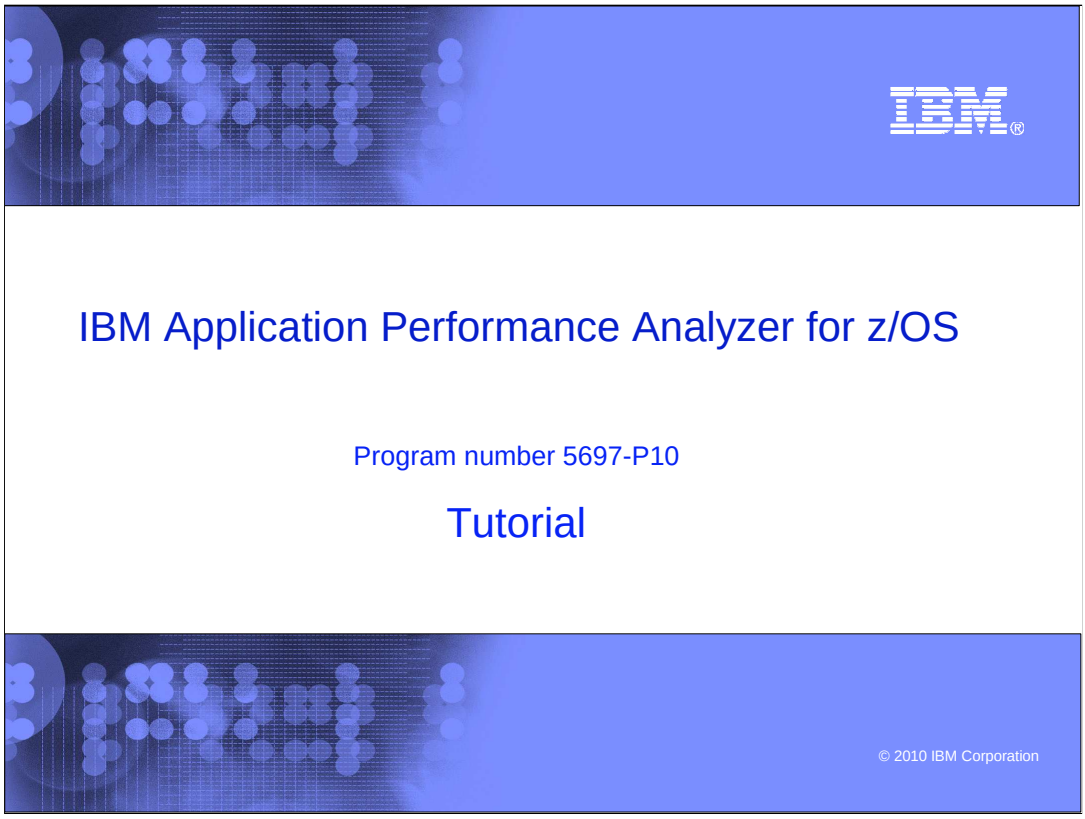

This is the tutorial for IBM's Application Performance Analyzer for z/OS®, one of the IBM zSeries® problem determination tools.

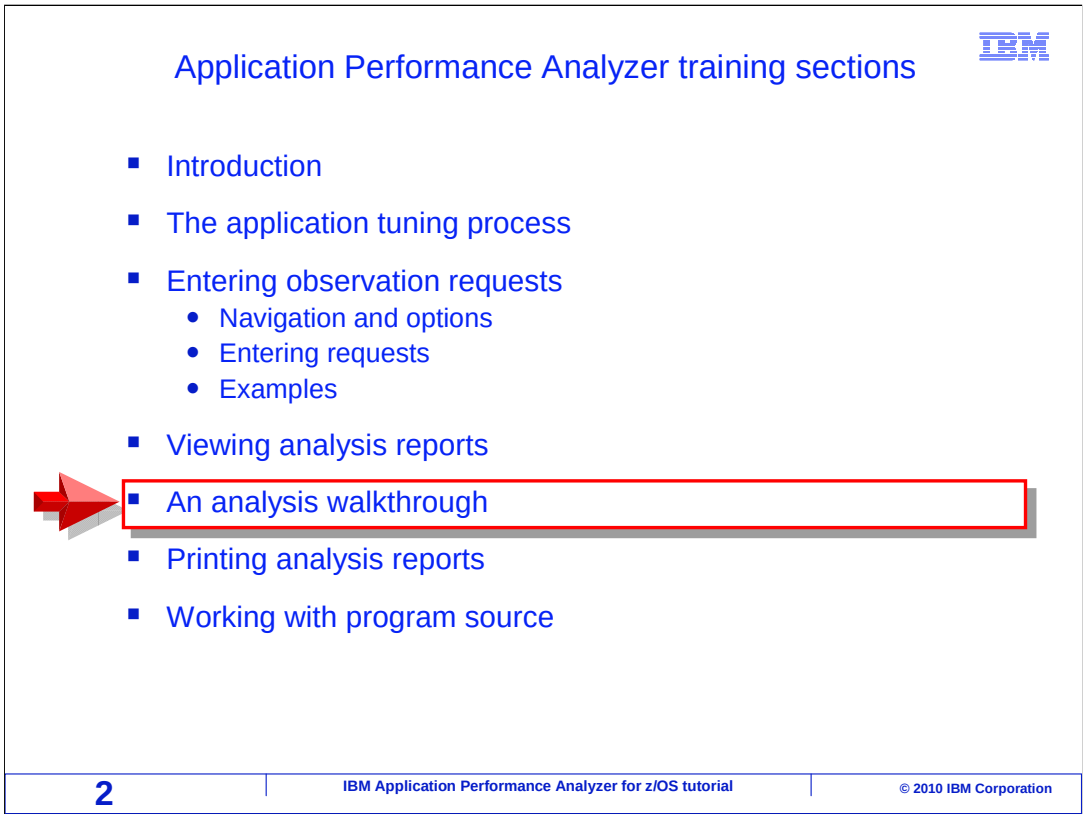

In this section: “An analysis walkthrough”, there is a complete walkthrough of entering an observation request, monitoring an application, and analyzing its performance.

Using APA – an example

- **Scenario: APA is used to analyze the performance of a batch job**
  - The job is running too long, and is using a lot of CPU time
- **Use APA to answer the question:**
  - **Where is the application spending its time?**
- **Based on what you find …**
  - **Can you change the application to make it run faster?**

In this example, APA is used to analyze the performance of a batch job. The job in question is already running, and seems to be using an excessive amount of CPU time. You can use APA to answer the question: “where is the application spending its time?” And based on what you find, you can come up with ideas to answer the question “can you change the application to make it run faster?”

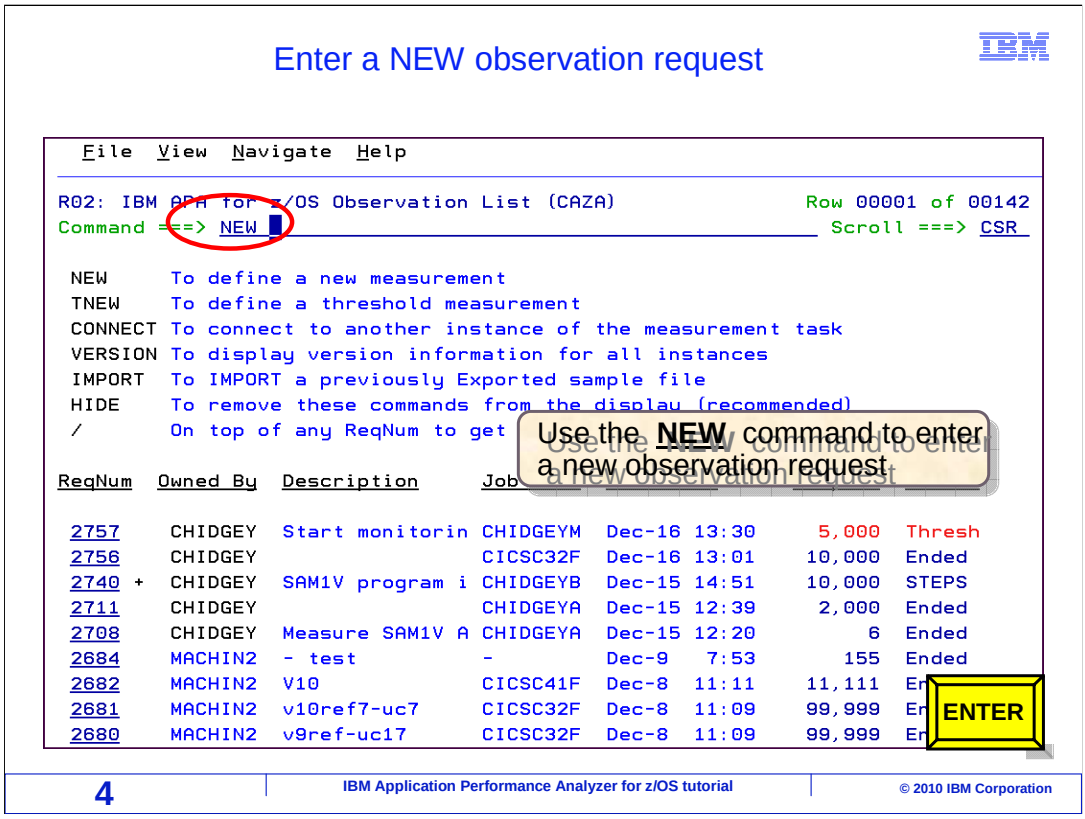

The job to monitor is already running, and you have noticed that it is running too long. You will enter a new observation request.

Just type "NEW" on the command line, and press Enter.

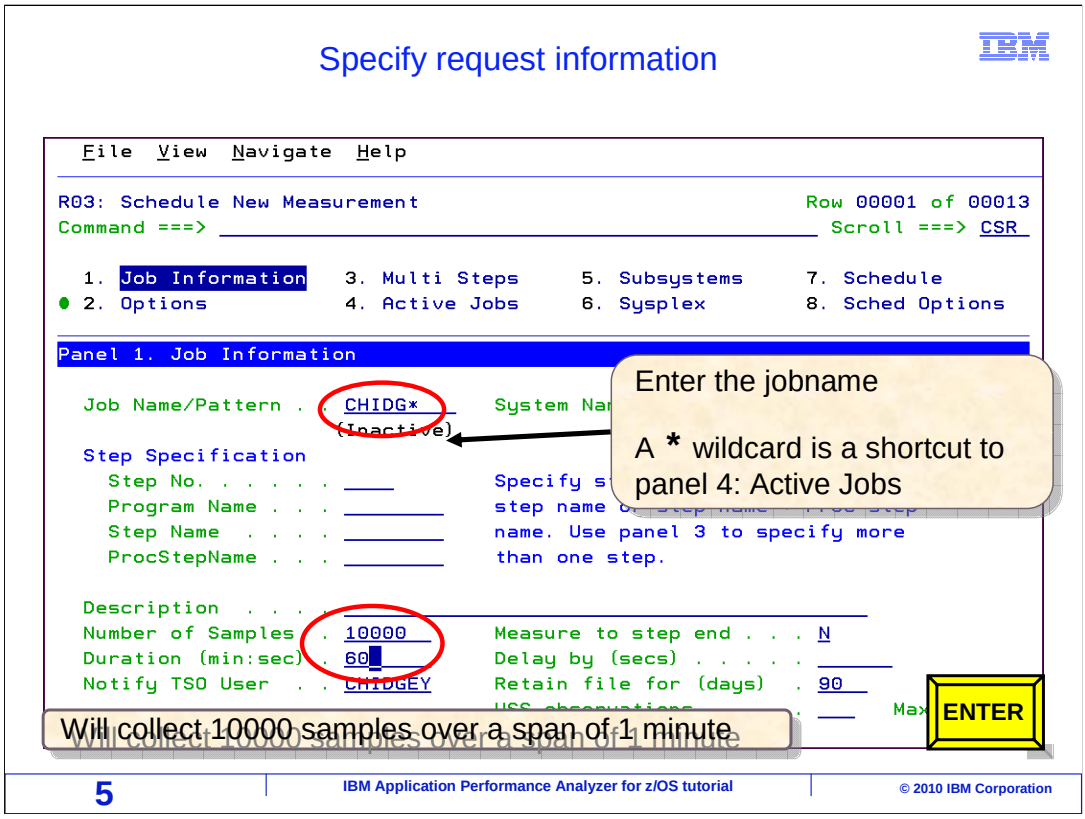

Panel number one, “Job information”, is always displayed first. At a minimum, you will want to enter the job name, the number of samples, and the monitoring duration.

It is requested that APA should take 10,000 samples over a duration of 60 seconds. The job name is required. The first few characters of the job name followed by an * as a wildcard is entered. Remember that when you use an * wildcard in the job name, that is a shortcut to panel number four: Active jobs. Enter.

Here is panel four, active jobs, where it displays a list of jobs running on the system that match the job name. Remember that when you select a job from panel four, APA will start monitoring the job immediately, rather than wait for the next time that it starts.

To select the job you want to monitor, use an “S” line command. Press Enter.

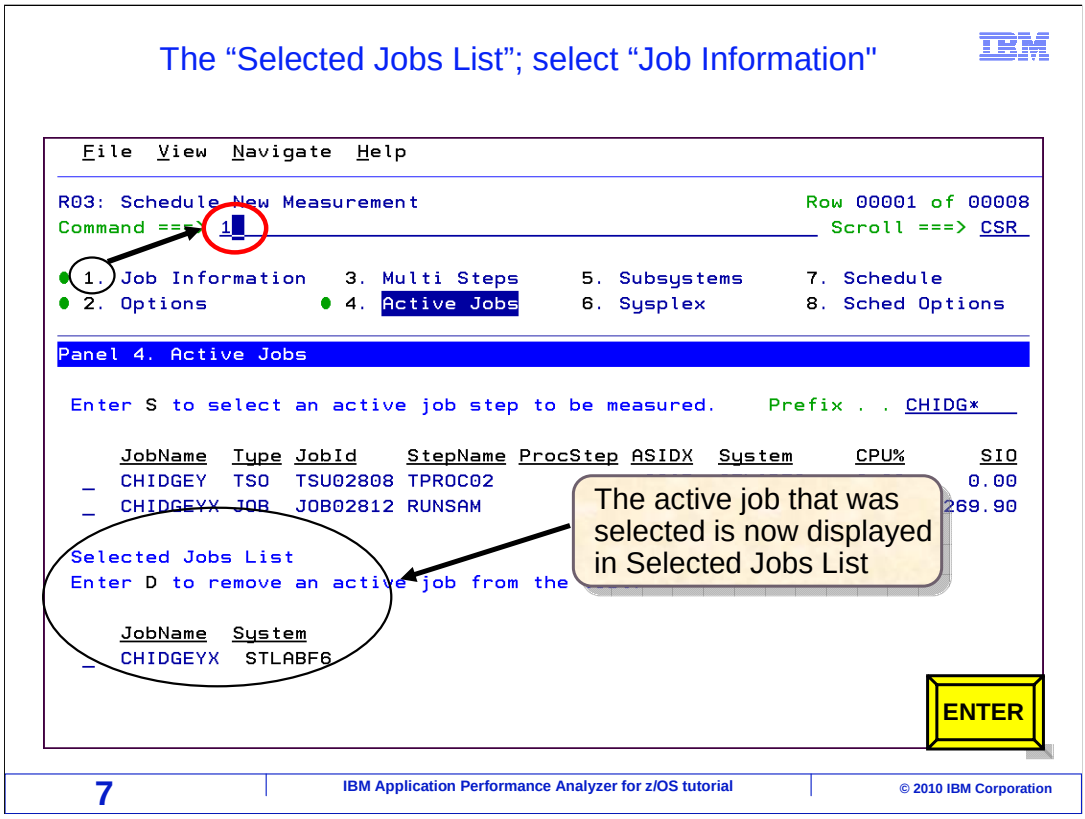

Notice that the active job you selected is now displayed in the lower half of the panel in Selected Jobs List.

Remember, from this panel, you can go directly to any of the other panels. Do not forget to go to panel two if you need to turn on any of the data extractors for CICS®, DB2®, IMS™, MQSeries®, or Java™. But in this case, the application does not need any of those extractors, so you will go directly to panel one, Job Information, to submit the request.

Enter one on the command line and press Enter.

Start the observation session

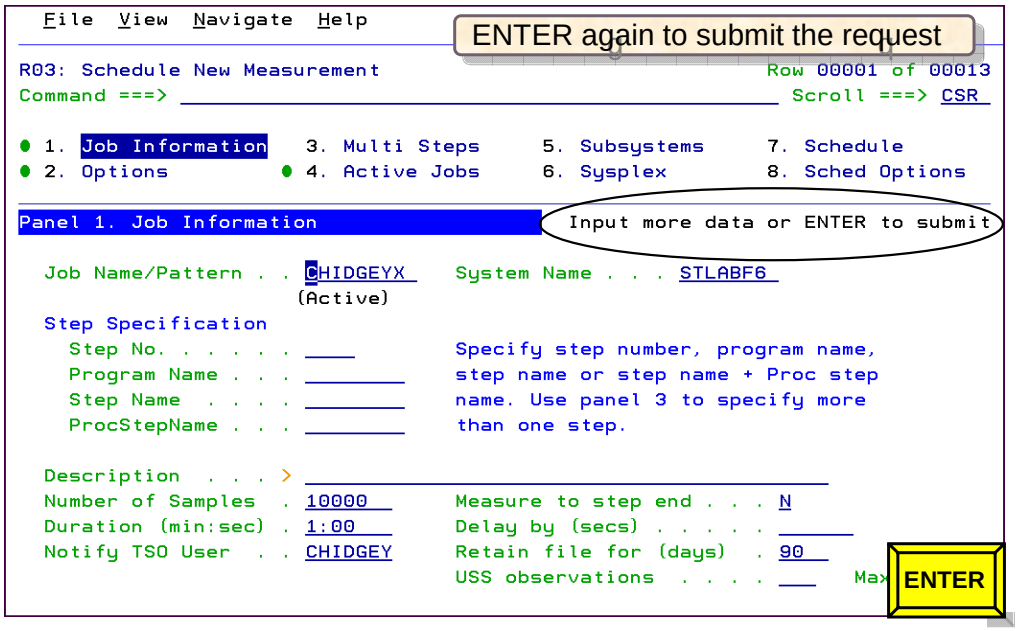

When you get this message: "input more data or enter to submit", then APA is ready to accept the request. Press Enter to submit the request to APA.

A new request was added

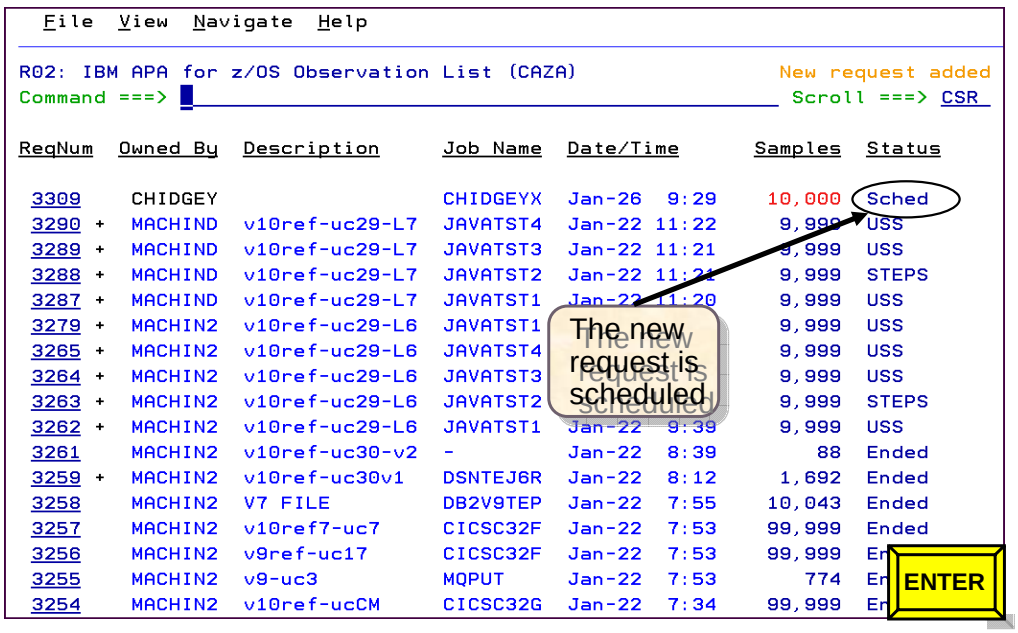

The new request is displayed at the top of the observation list. Notice that its status is “scheduled”, but it will change to "active" immediately. Press Enter to refresh the list.

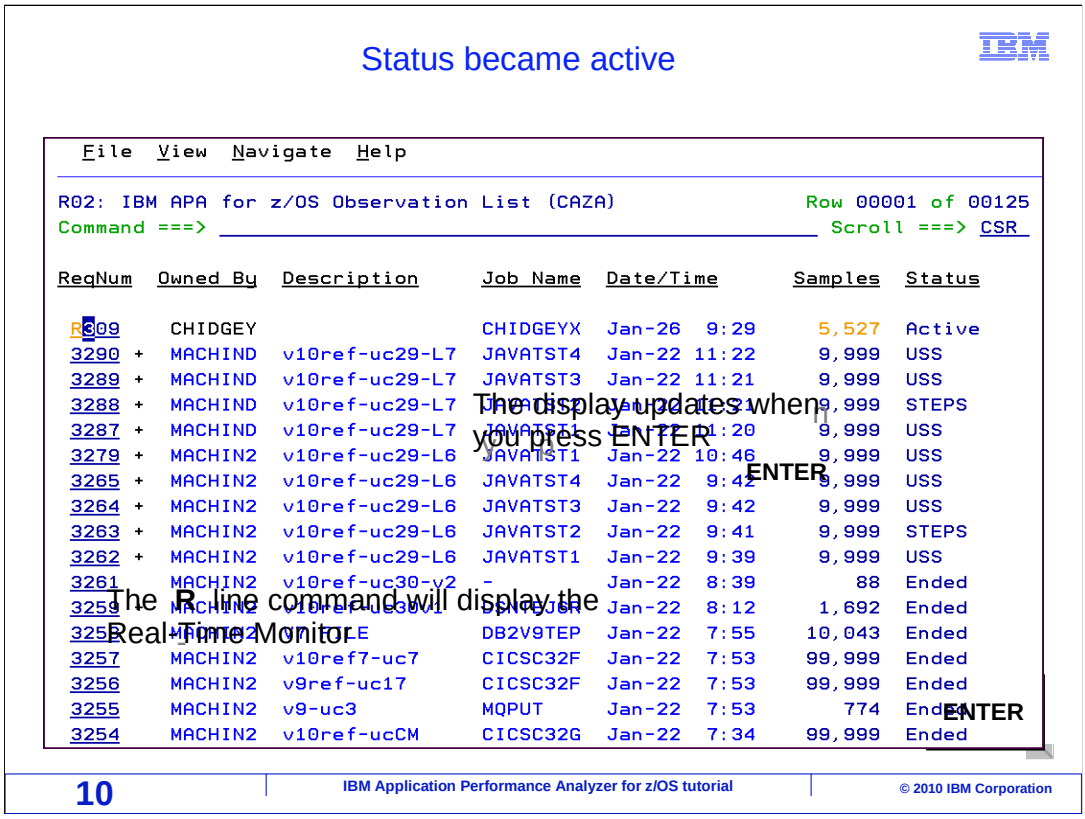

Now the status has changed to "active". That means that APA is now monitoring the application. It will monitor the application for 60 seconds, since that is what was requested. The samples column displays how many samples APA has taken so far. In this example, 10,000 samples were requested, and so far it has taken over 5,500 samples.

You can view real-time reports while APA is monitoring. Use the "R" line command to display the real-time monitor, and press Enter.

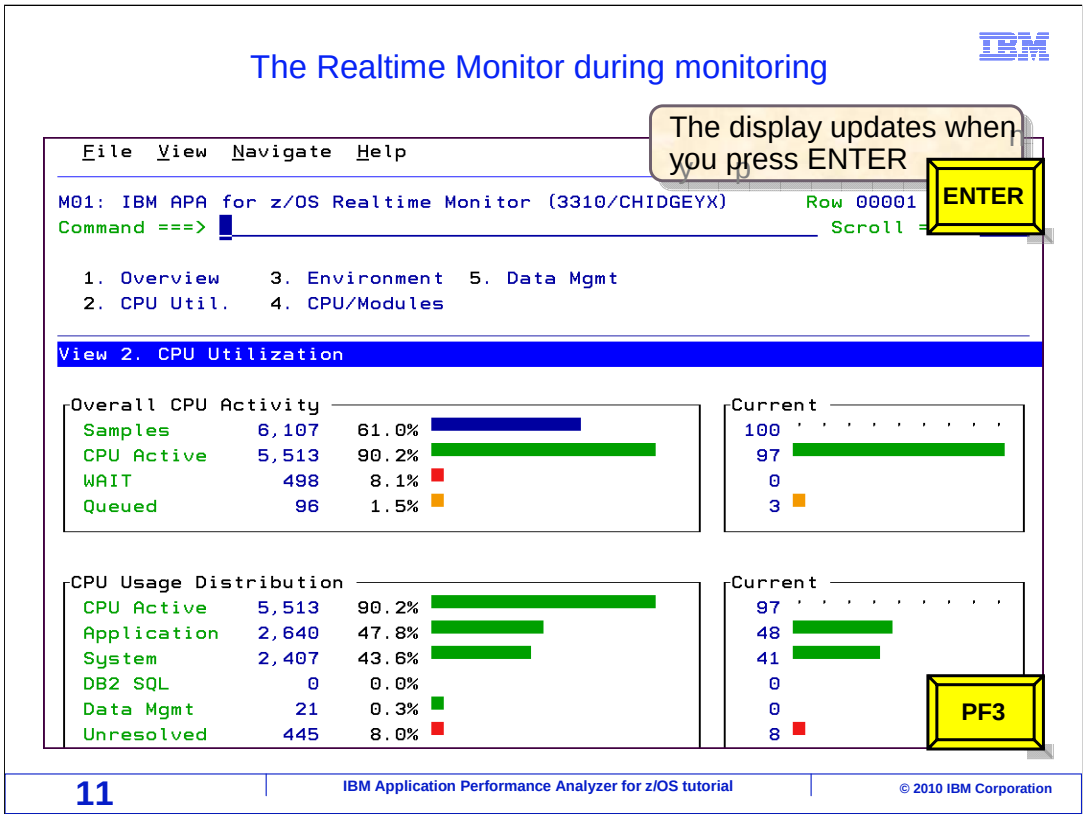

The real-time monitor is displayed for the observation. A lot of information is available during real-time, although the real-time monitor does not provide nearly as much detail as the performance reports that you can view after the session is complete. On this screen, there is an overview of the application's performance characteristics, CPU utilization statistics, and information about the environment. There is also a list of modules that APA discovered during its sampling, and file and data management information. To refresh the display, press Enter.

If you want, you can continue to display the real-time report for as long as APA continues to monitor. Or you can exit at any time by pressing PF3.

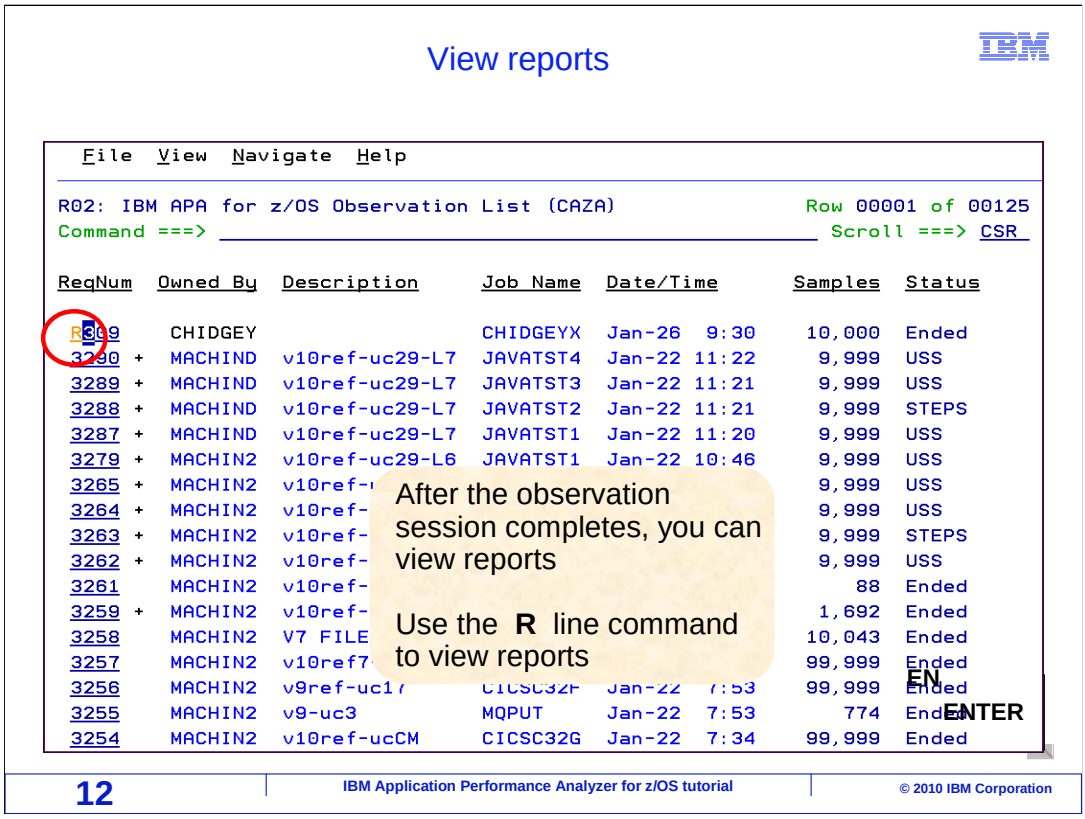

In this example, a duration of 60 seconds was requested. When the monitoring duration elapsed, monitoring stopped and APA sent a notification that sampling was complete. Press Enter to clear the notification message.

APA writes the data that it collected to measurement data sets, and the measurement session is complete. Notice that the status changed to "ended".

After an observation session is complete, you can access the reports with an “R” line command. Understand that the “R” line command means “Real-time report” while a session is active, and “Reports” when a session is complete. Enter.

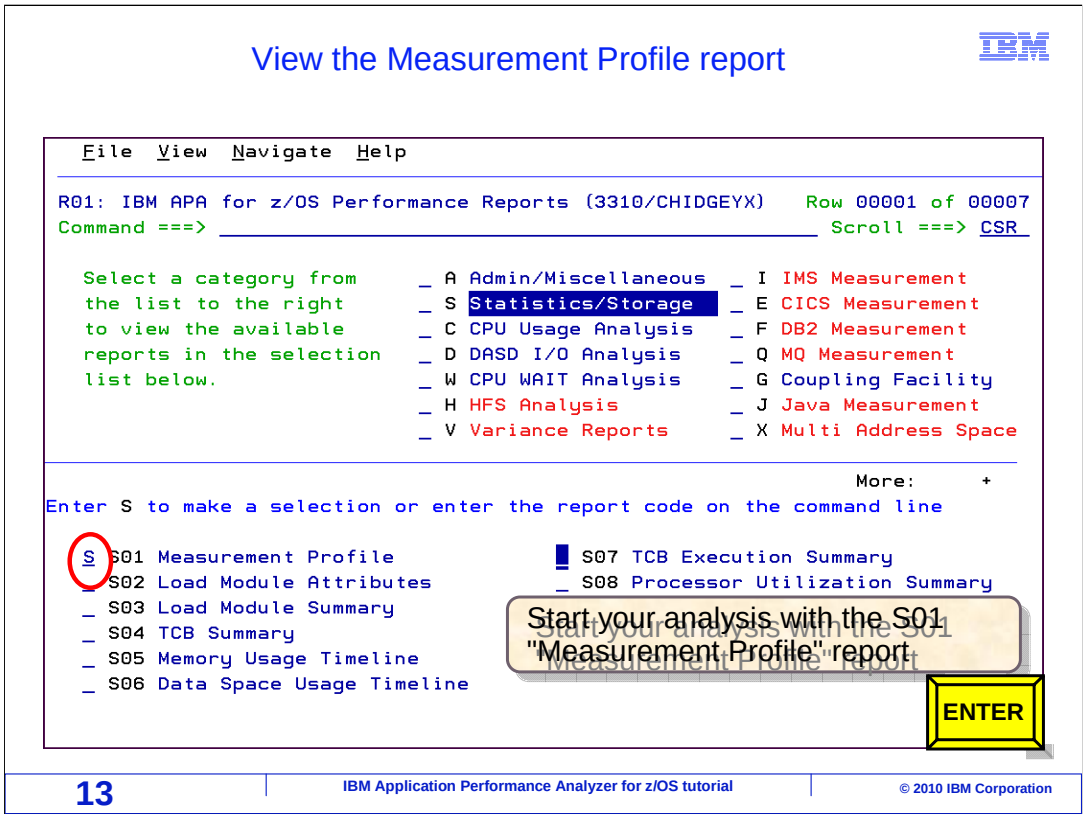

As a general recommendation, always begin your analysis with the S01 "measurement profile" report. Notice that right now, in the upper part of the panel, the “statistics/storage” category is selected. Because of that, the statistics and storage reports are shown on the lower part of the panel.

Here, the S01 report is selected. Enter.

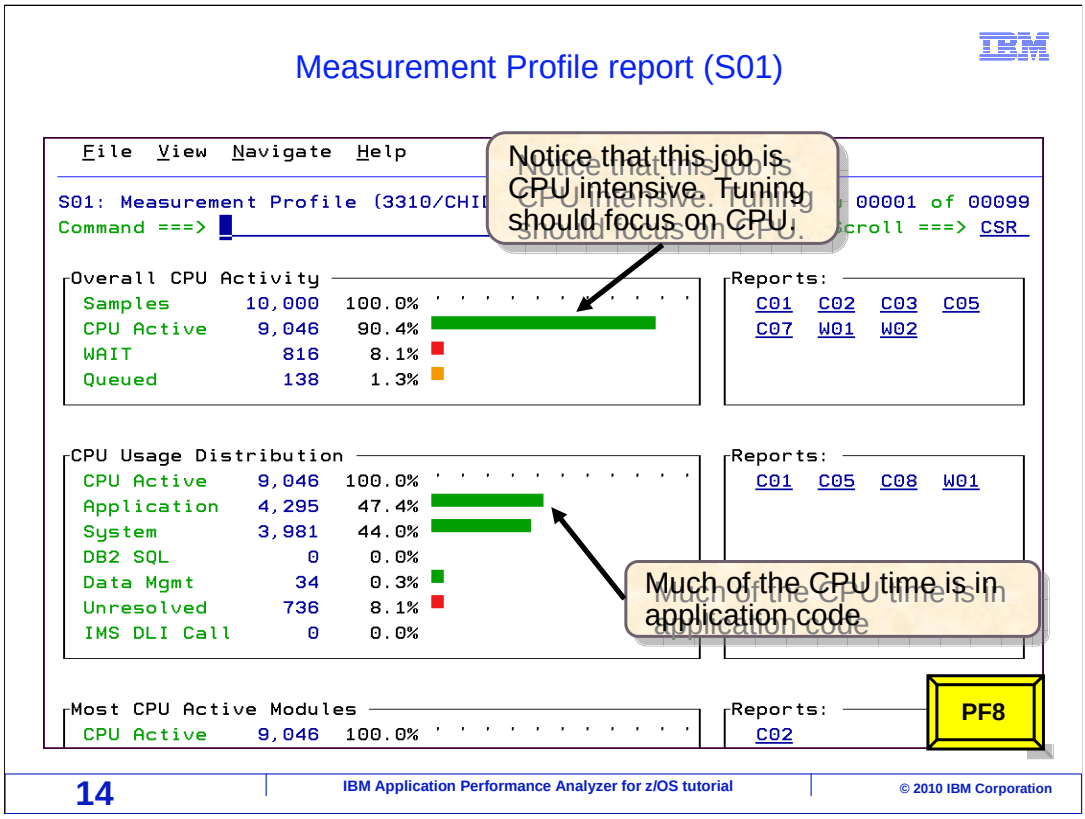

The S01 measurement profile report shows you overall statistics and application activity. Typically, the first thing you will want to look at is the "overall CPU activity" section. It shows you the percentage of time the application spent in each of the three states: CPU active, wait, or queued.

In this example, notice that the job is CPU intensive. Most of the non-queued time is spent in the CPU active state. In this case, about 90% of the time. Also notice the next section, CPU usage distribution. It shows you how that 90% is comprised. In this case, over 47% of the CPU time was spent executing application programs. That is, of the 90% CPU active time, 47% of it was spent in applications. Just so you know, that is a relatively high number.

Most healthy applications spend a much larger percentage of their time executing system code, that is, system level programs that are performing activities at the request of application programs. The fact that the application percentage is so high can be an indication of a loop. Scroll forward in the report with a PF8.

Why use the Measurement Profile report (S01)?

- Start your analysis with the S01 report
- This report tells you whether most of the time was spent in CPU or WAIT (or a lot of both)
  - Based on this information, you can focus your tuning effort on CPU or WAIT

Excerpt of the APA "Measurement Profile" report

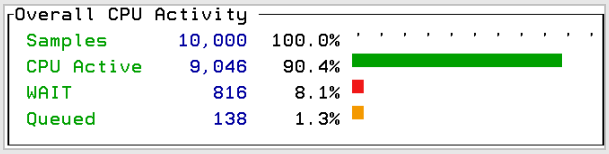

- The S01 report also provides information about the "Quality" of the measurement and APA overhead

Remember, you will typically want to start your analysis with the S01 report. That is because this report tells you whether most of the time was spent in CPU or wait, and based on that, you can focus your efforts on the right thing. The trick to doing performance analysis is to identify where the application is spending most of its time, and then understanding why. So understanding whether an application spends more time in a CPU or wait state tells you how to begin your research.

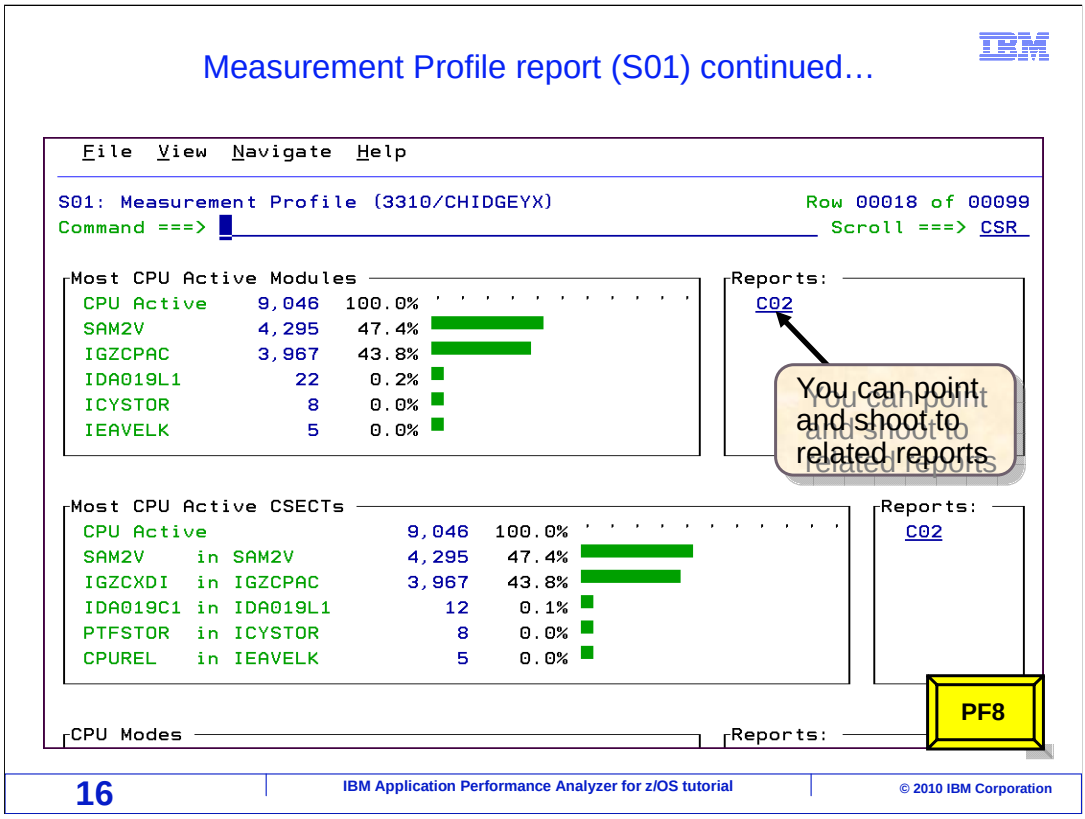

Back to the S01 report. The next section shows a breakdown of CPU activity by CSECTs. Remember that 90% of the CPU activity was spent executing application instructions. This section shows what percentage of that 90% CPU activity was used by individual CSECTs. In this example, 47% of the 90% CPU activity was spent in CSECT SAM2V which is the application. So your tuning efforts might want to focus in on that module. Scroll forward again. PF8.

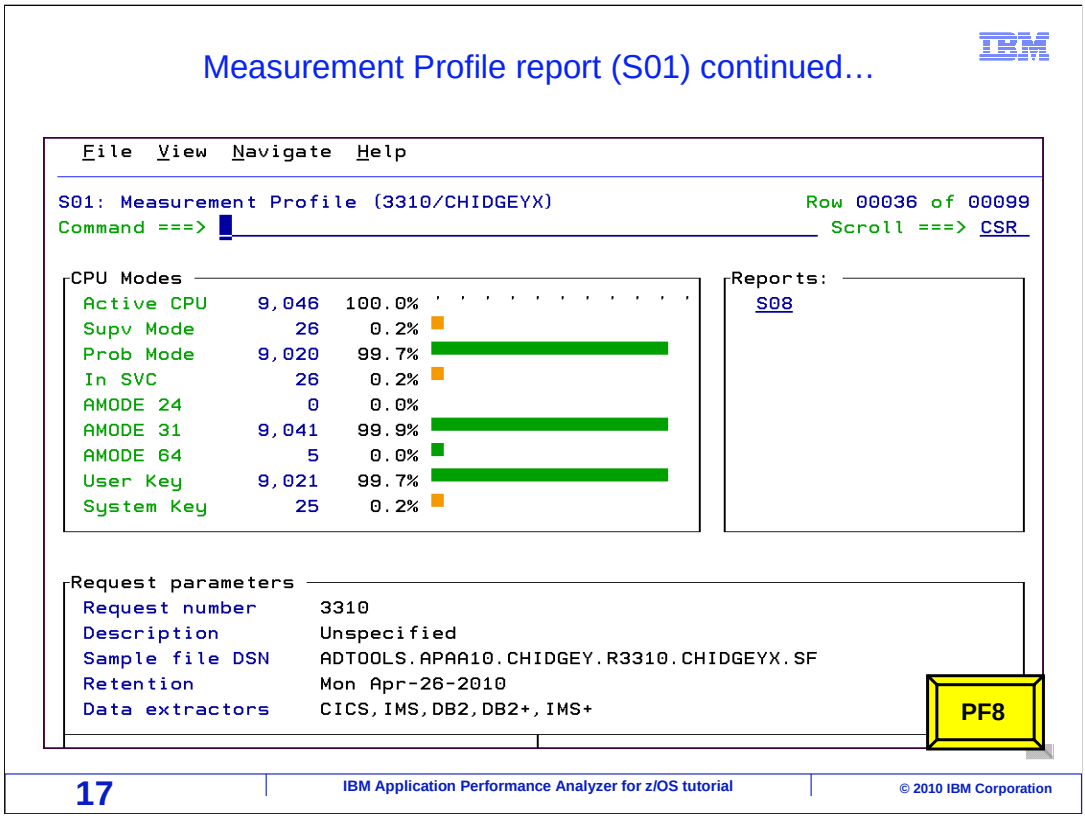

The next section is CPU modes. There are several CPU modes and addressing modes that an application can be running in at any given point in time, and this section shows you the percentages. Most applications spend the bulk of their time in problem mode, which is the case in this example. Supervisor mode is typically time spent in a special mode required by certain system-level functions. Notice that on the right side of the report, there are some point and shoot fields for other reports. These are reports you can access if you want to display more detailed, related information. PF8 again to scroll forward.

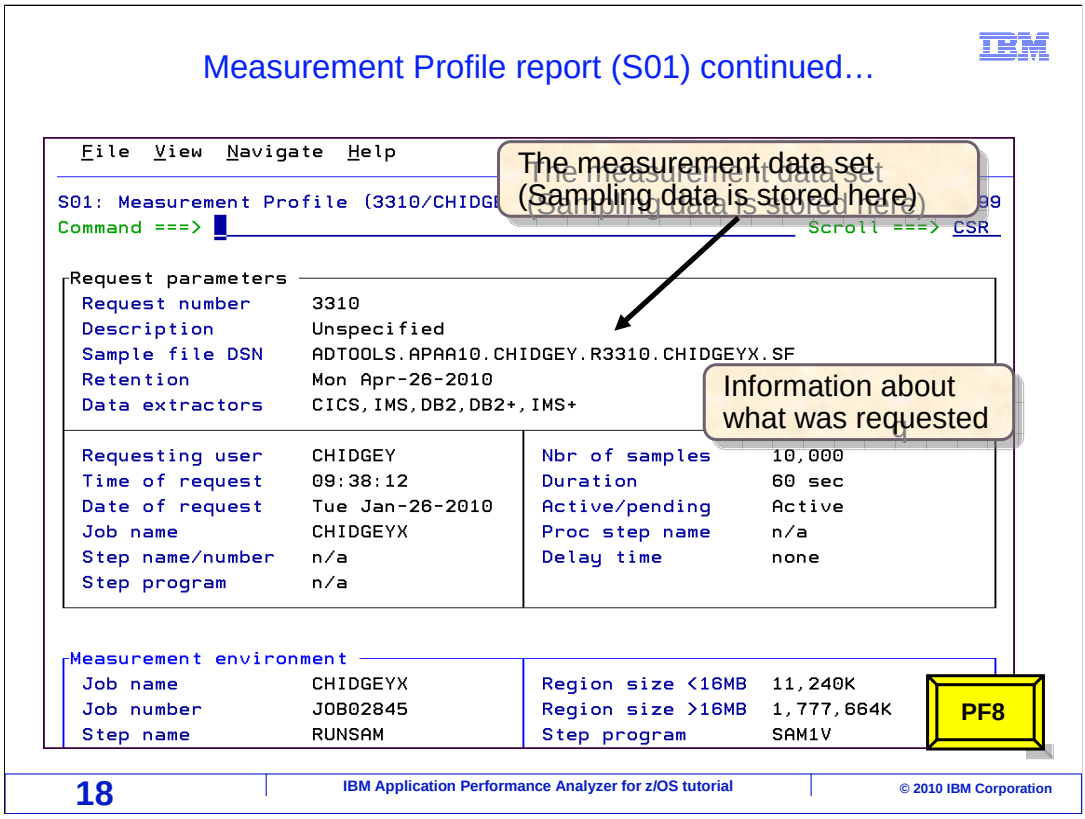

Next is the "request parameters" section. This section of the report displays the parameters specified to APA when the observation request was entered. It can be handy to have this documented, so that you can know, for example, the name of the job, number of samples, and duration requested. PF8 to scroll forward.

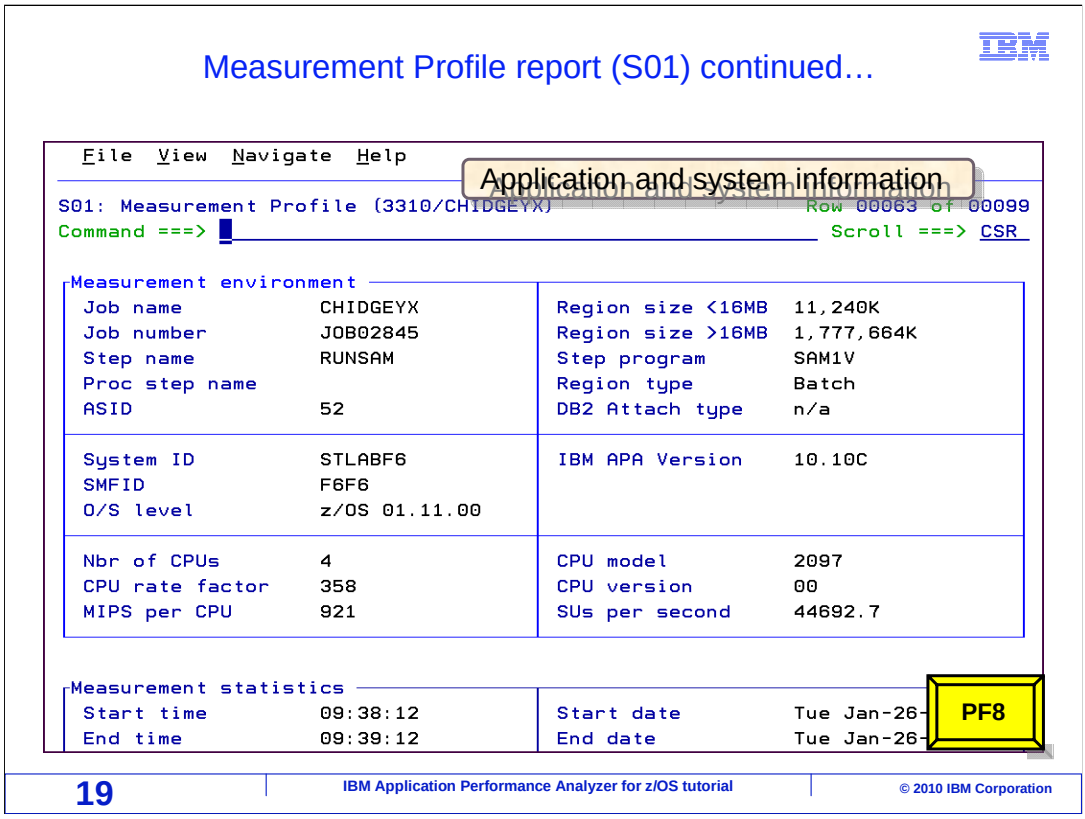

That brings you to the “measurement environment" section. This shows information about the job that was monitored and the system it was running on. This section of the report displays the job name, job number, step name, and region size, in addition to the z/OS system ID, and information about the processor. PF8.

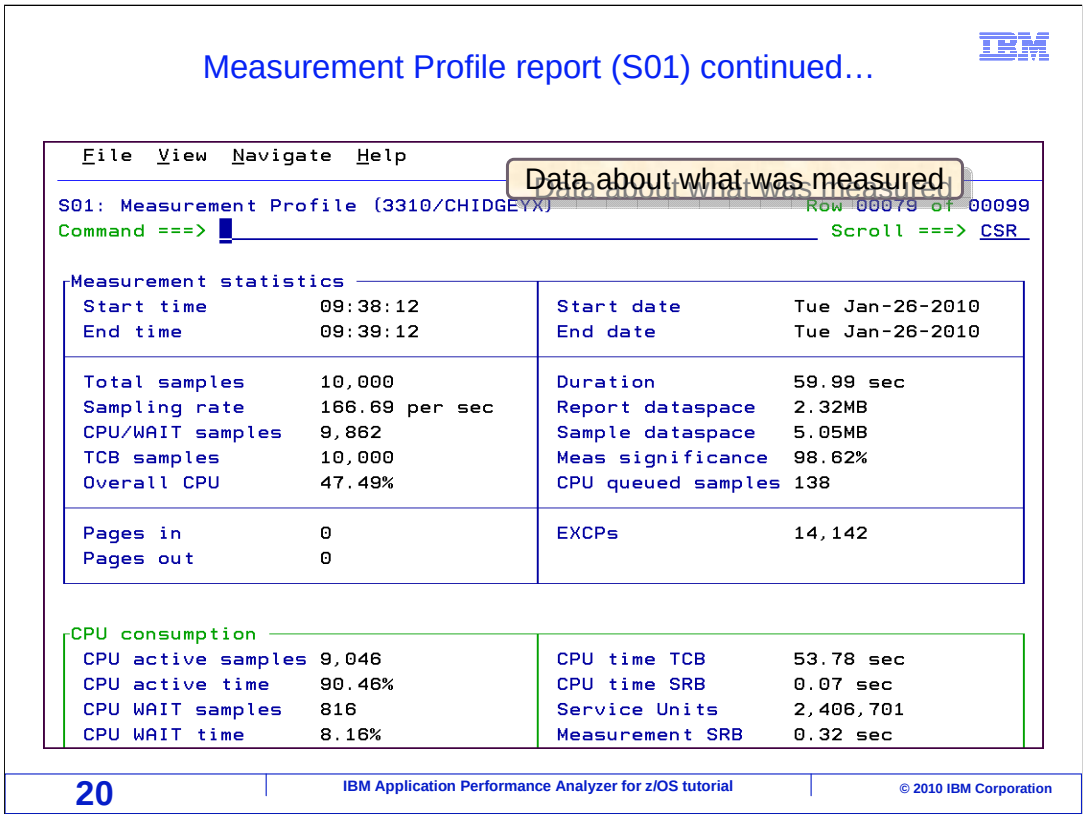

Next is the measurement statistics section. One of the earlier sections was "request parameters", which contained information about what was requested. This section shows information about what was actually collected. It shows the start and end time of APA's monitoring, how many total samples were taken, the sampling rate, and the actual measurement duration. It is possible that the requested duration and number of samples can be different from the actual. For example, say that you requested a monitoring duration of 60 seconds, but then the job only ran for 30 seconds. That will cause the total duration to be only 30 seconds, and only about half of the requested number of samples will have been taken. And that is the type of information reported here in this section.

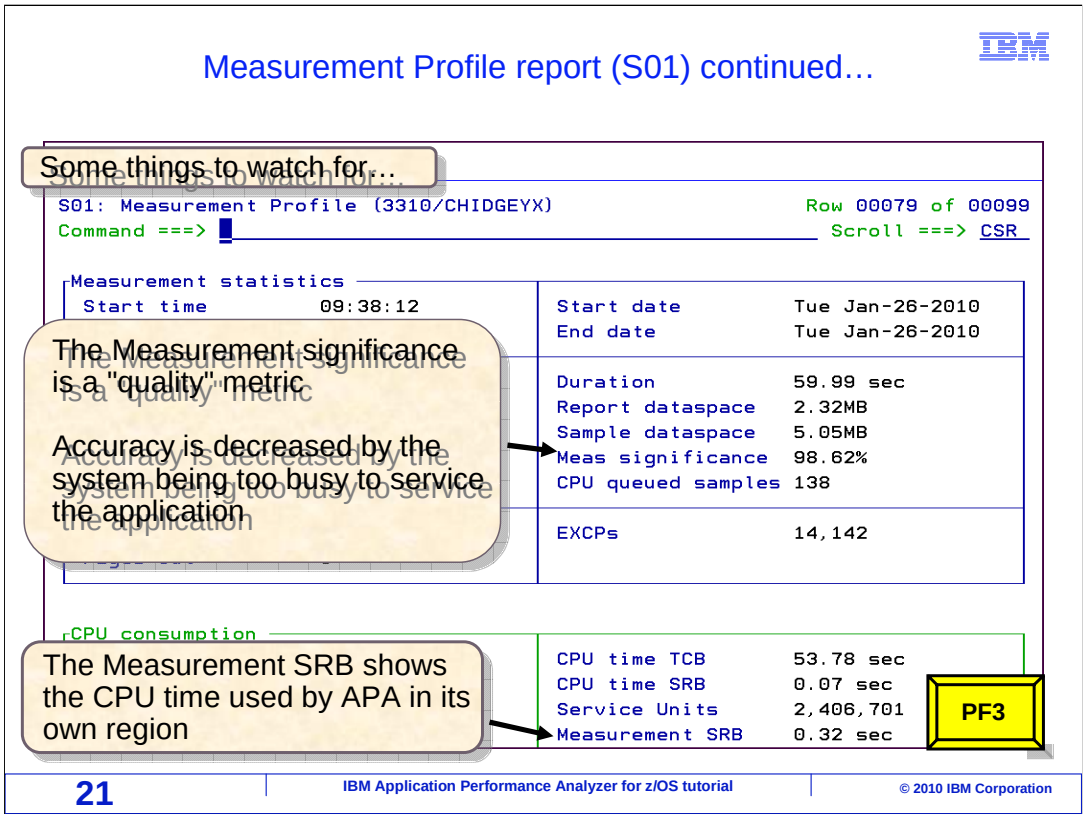

Here are a couple of things to watch for in the measurement profile report. The measurement significance is a quality metric. The accuracy of APA's reports is decreased by the system being too busy to service the application.

In a worst-case scenario, the system can be so busy processing other tasks that the application is in a queued state 100% of the time. Therefore, no CPU or wait time is incurred. Obviously, if that happens, there is not any meaningful information in APA's reports.

The measurement significance shows the percentage of time that the application was not in the queued state. The higher the number, the better, although sometimes when you are running on an extremely busy system, the number might always be somewhat low and there might not be anything you can do about that.

It is also good to know about another statistic called the measurement SRB. It shows APA’s "overhead", and it is CPU time used by APA in its own region to collect the performance data. Once you are done with this report, press PF3 to return to the report menu.

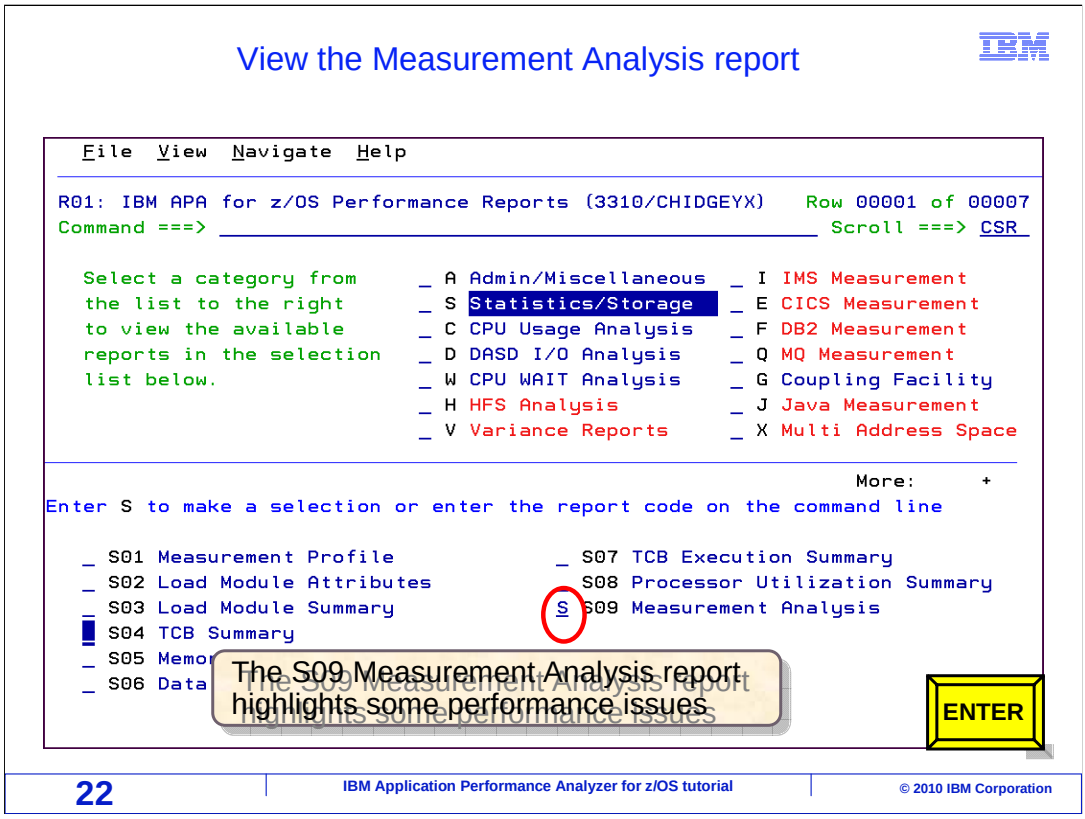

What information is available in the “S09 Measurement Analysis” reports?

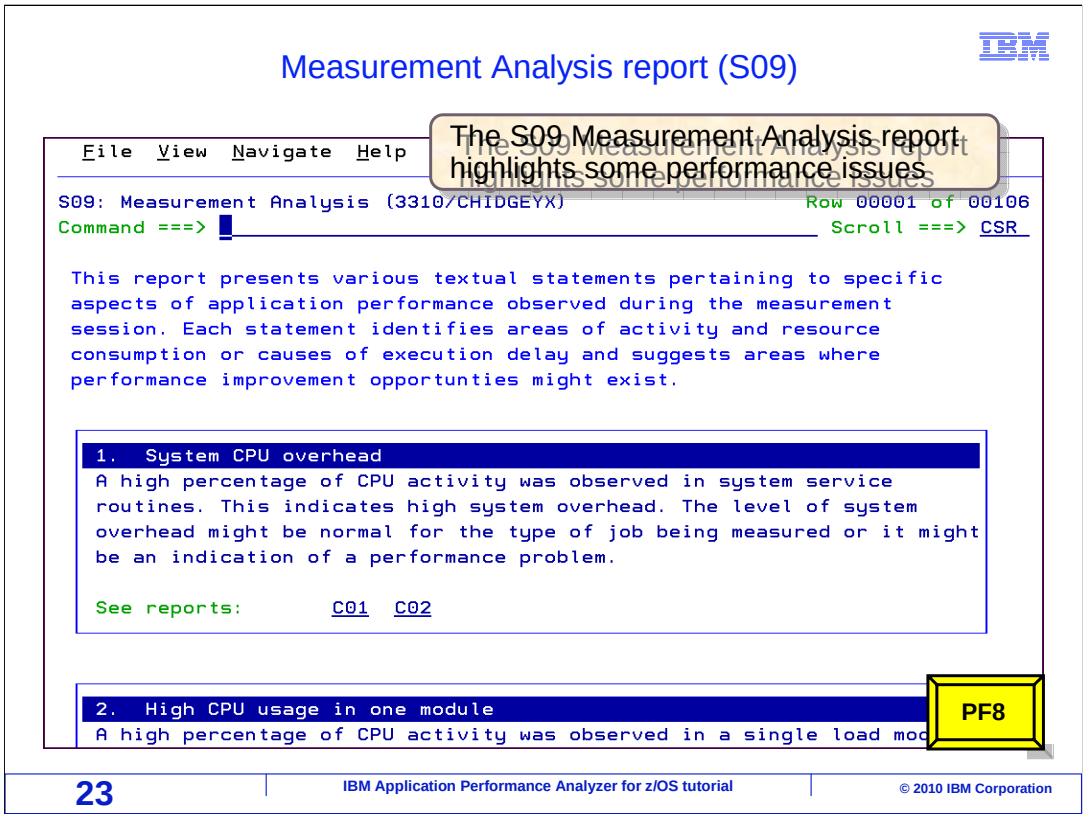

S09 presents statements, each representing an observation made about some aspect of the measured job. The purpose of each of these observations is to provide a synoptic analysis of an area of resource usage and, in some cases, suggest where some performance improvement opportunities might exist.

In this example, the focus is on CPU consumption.

Analyze these observations in the context of how you expect the measured job to perform. Some of the statements in this report might draw your attention to aspects of resource consumption that is perfectly normal for the job. For example, high CPU consumption might be noted in a certain module in a situation where you actually expect high CPU usage in that module. To display more scroll forward with PF8.

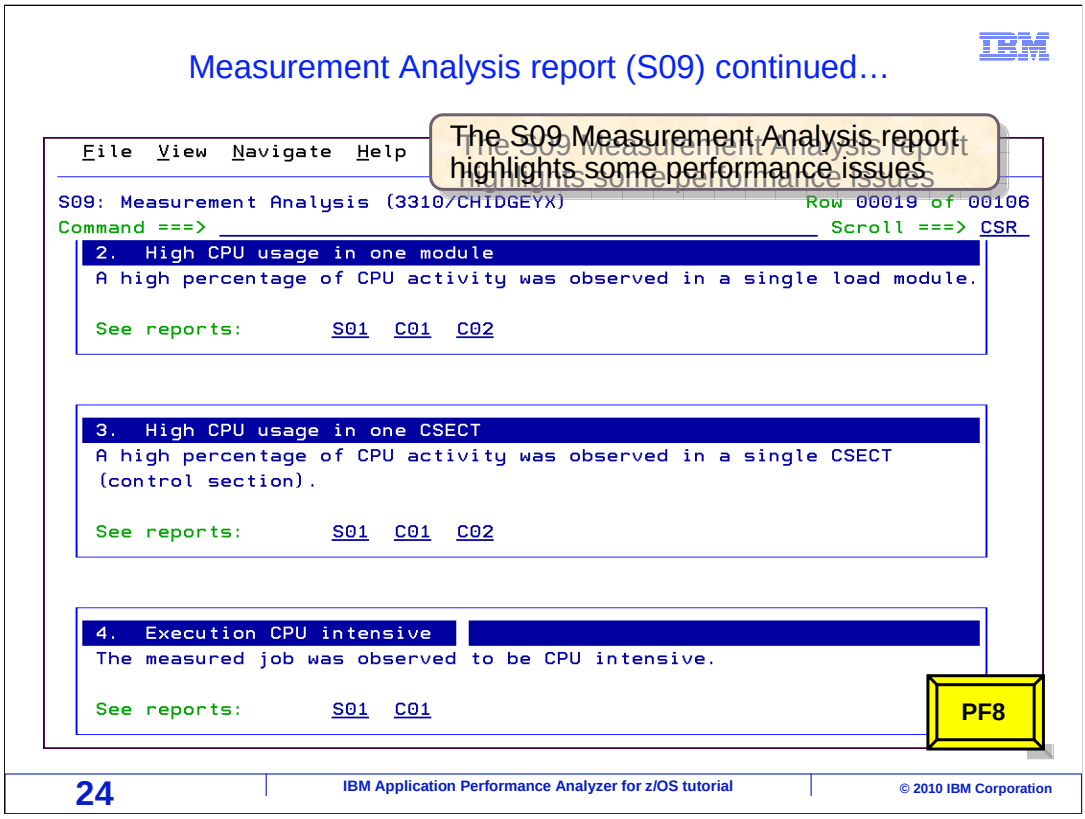

In this example, the Measurement Analysis has indicated that there is High CPU consumption in a single load module. Message three narrowed it down to a single CSECT. This report also suggests specific reports where you will find detailed related information. Press PF8 to go down.

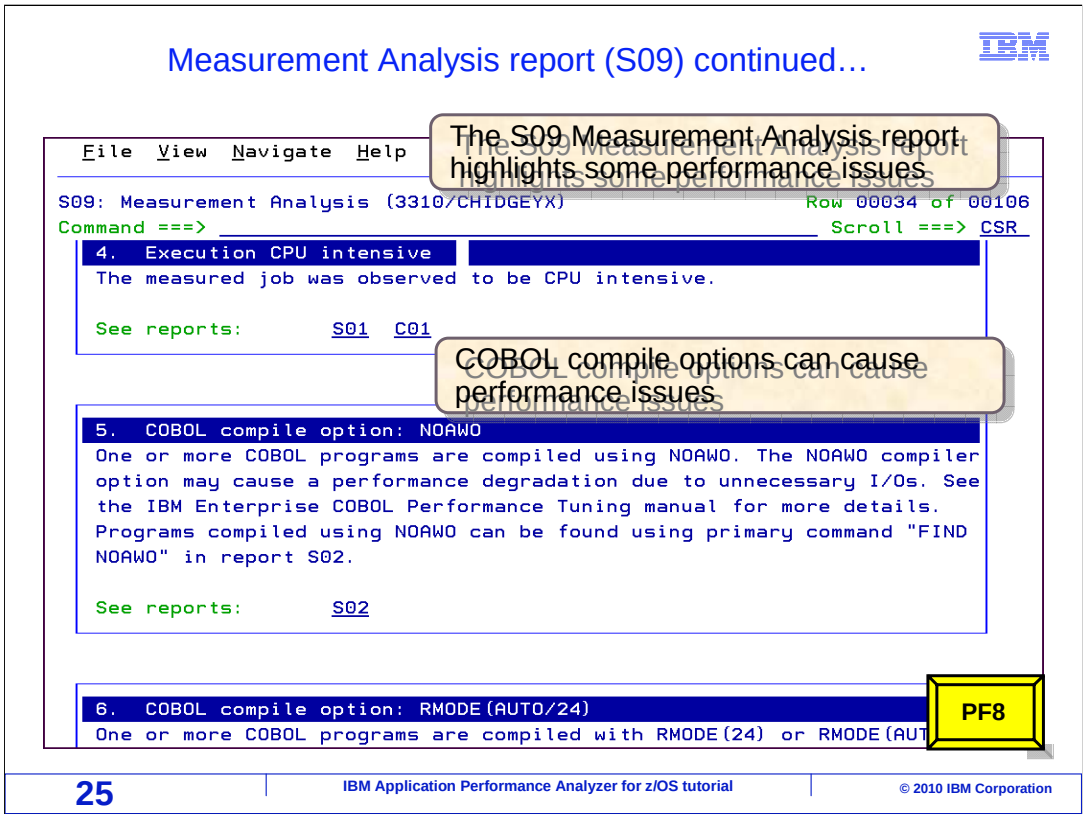

This report will also indicate whether the job is CPU intensive or wait intensive. In this example, the report indicated that the job is CPU intensive as identified in message four. Also, the S09 report will highlight any possible performance issues. Sometimes certain compile options can cause these performance issues to occur. The S09 report will highlight these compile options and provide information regarding their impact. Press PF8 to scroll through the rest of this report.

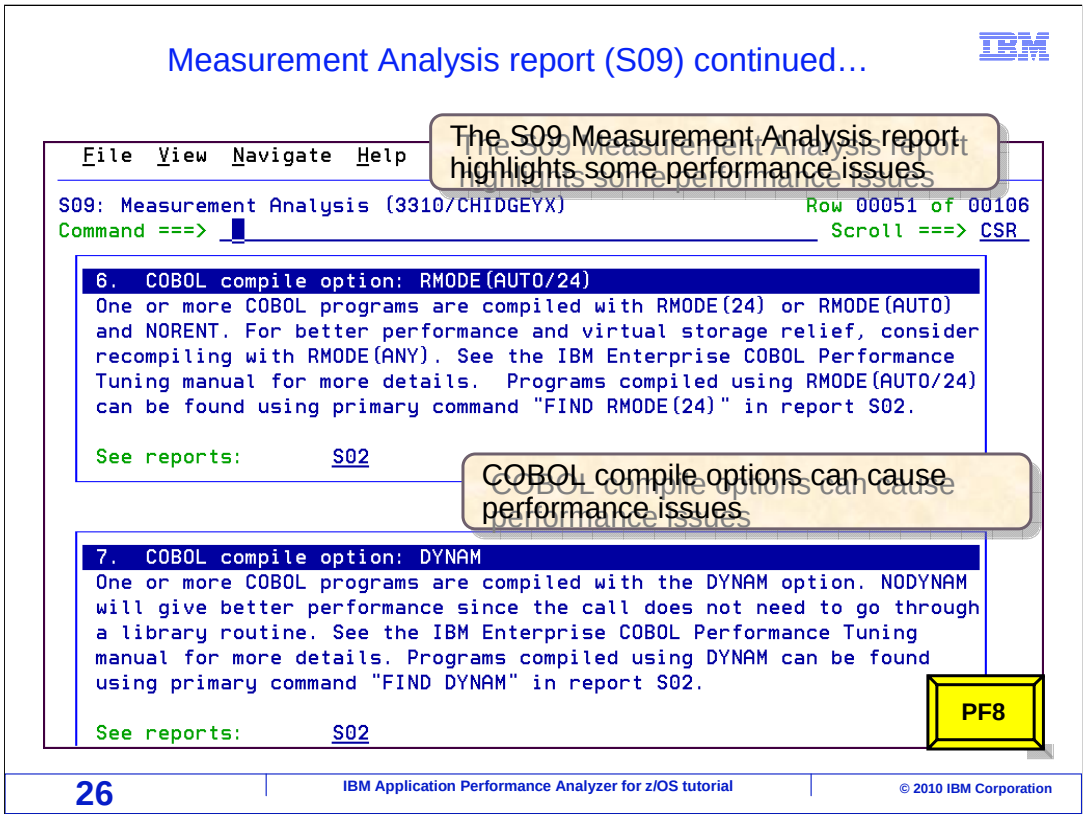

Here are additional compile options that can affect the performance of the application. PF8 to scroll down.

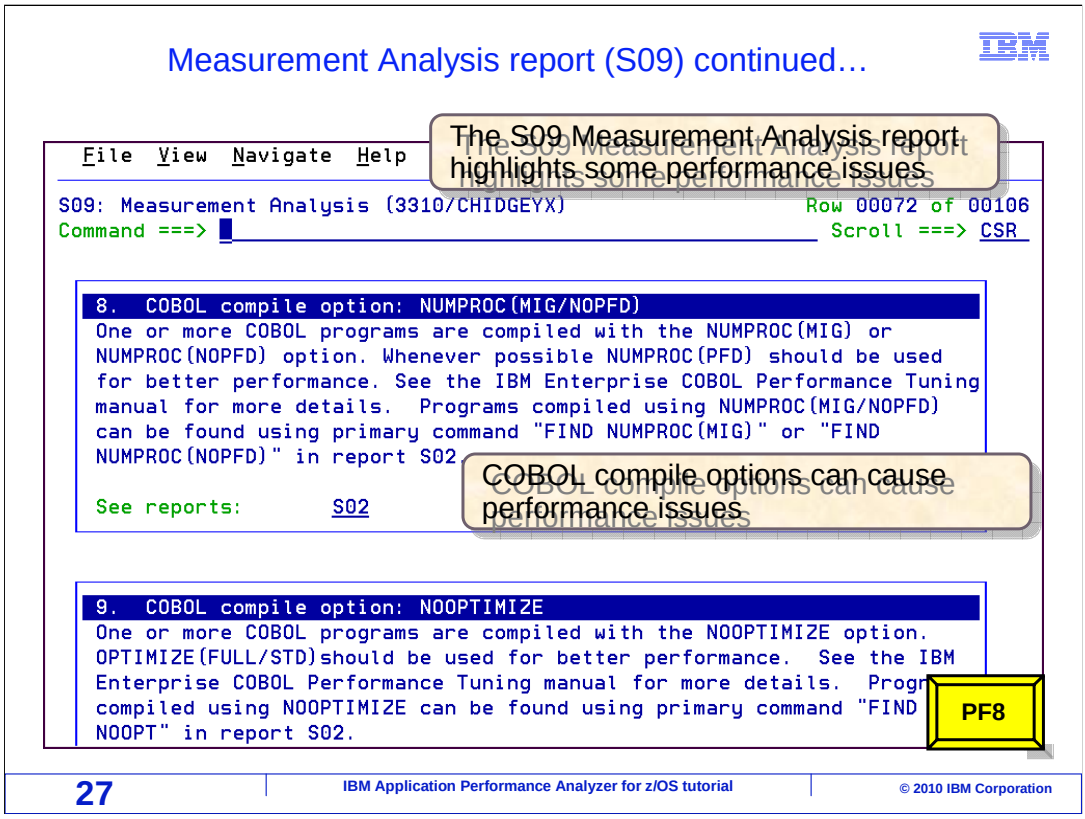

More compile options are highlighted. PF8 to scroll down.

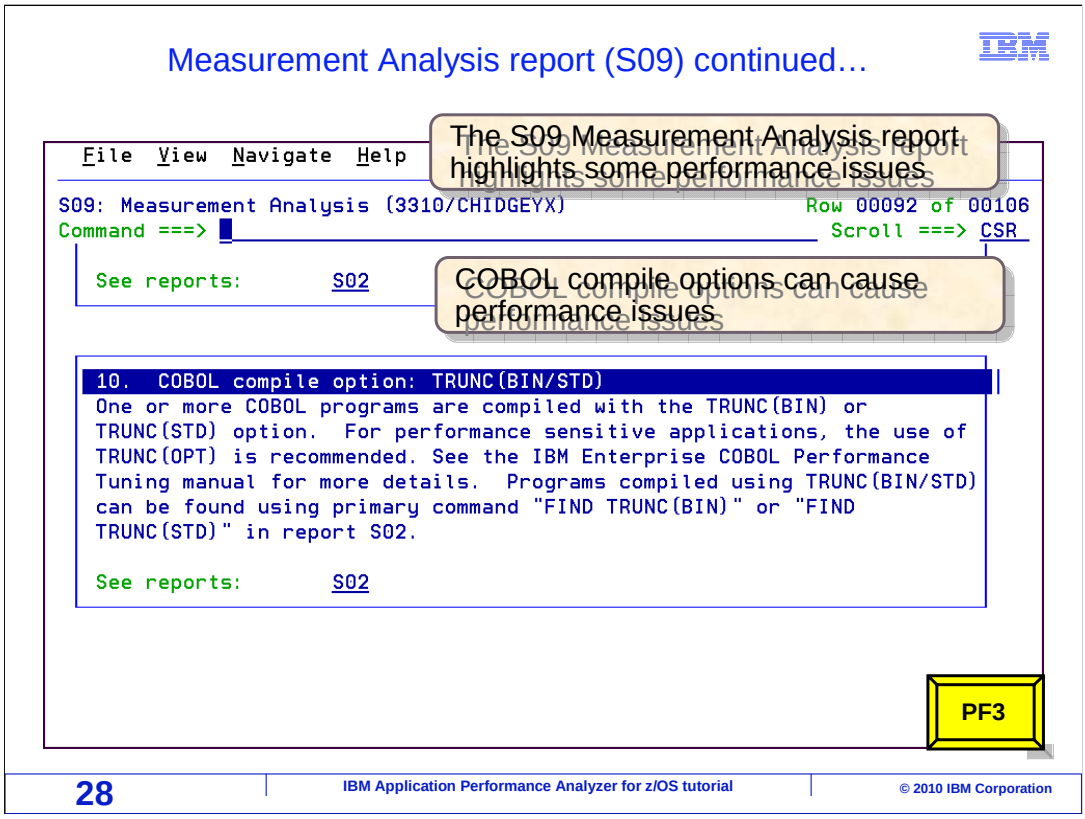

Now you are at the end of the report. Using the recommendations and focusing on the modules indicated in this report could help you significantly improve your applications performance. PF3 to exit the report.

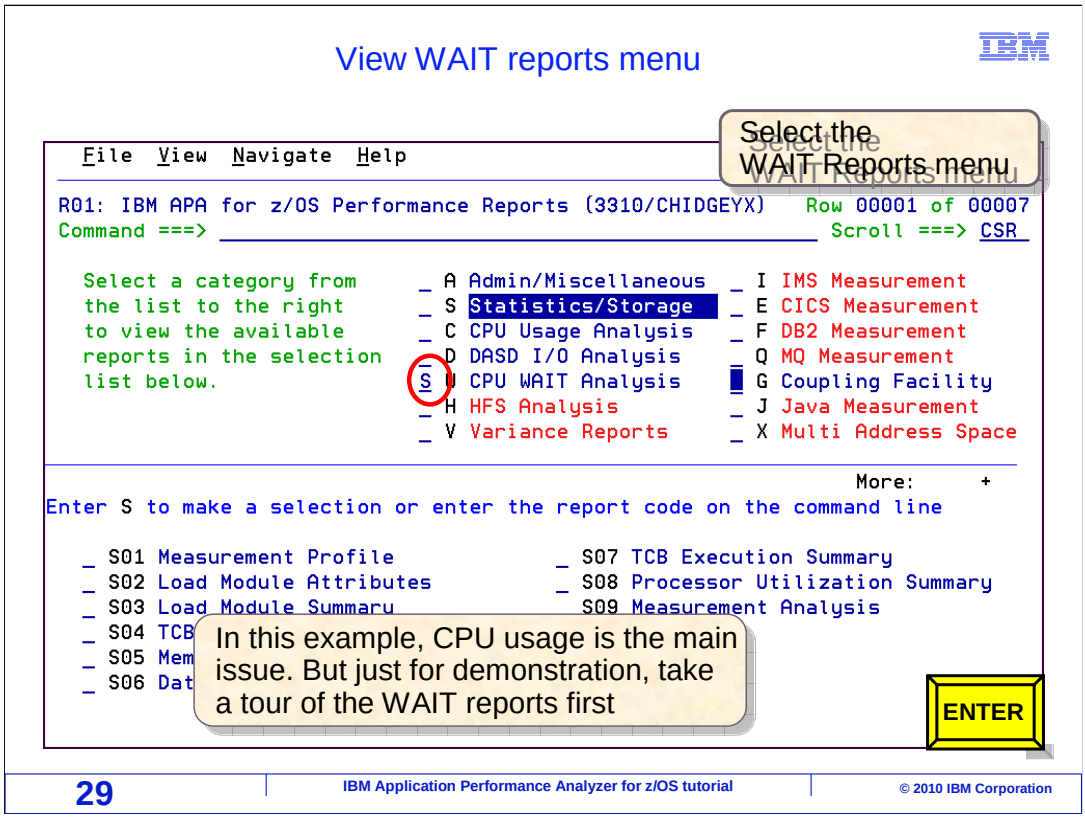

In this example, CPU usage is the main issue, and that is where the focus should be in your performance analysis efforts. But just for demonstration, you are guided through the information that is available in the “wait” reports.

To select the “wait” reports category, place an ”S” next to it and Press Enter.

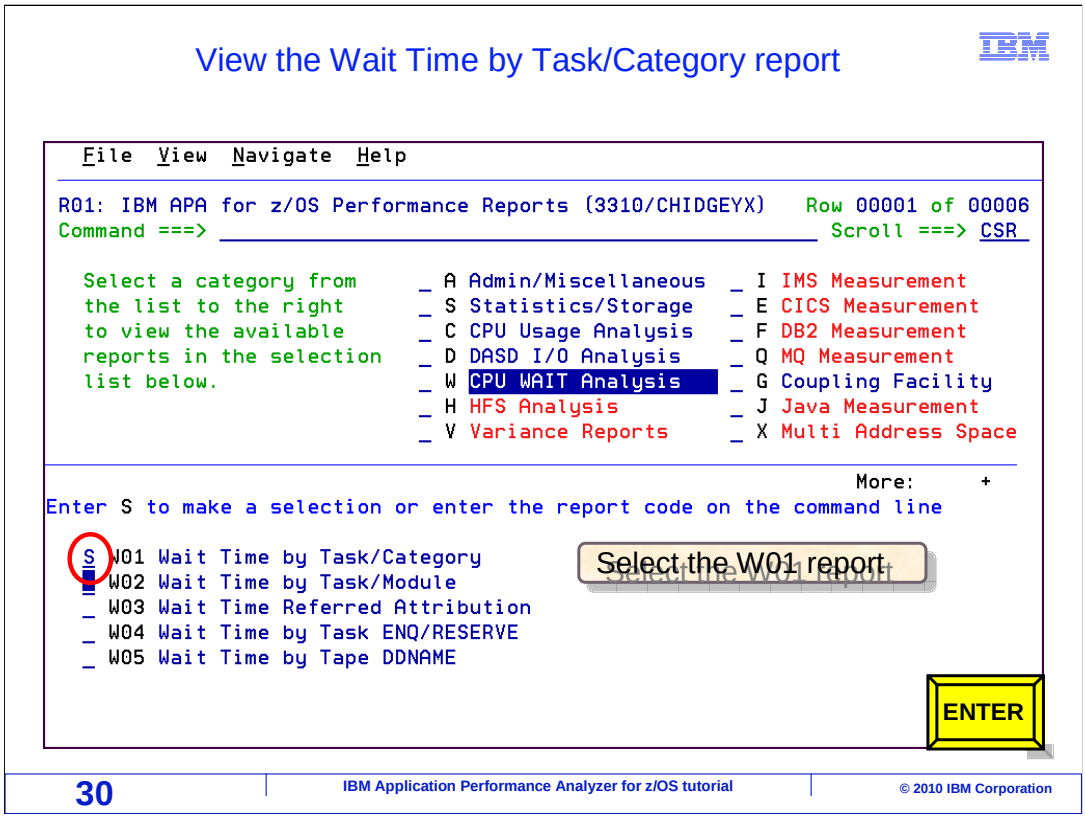

And now the ‘wait’ reports are shown. In this example the ‘W01 wait time by task/category’ report is selected.

You selected it with an “S” and Press Enter.

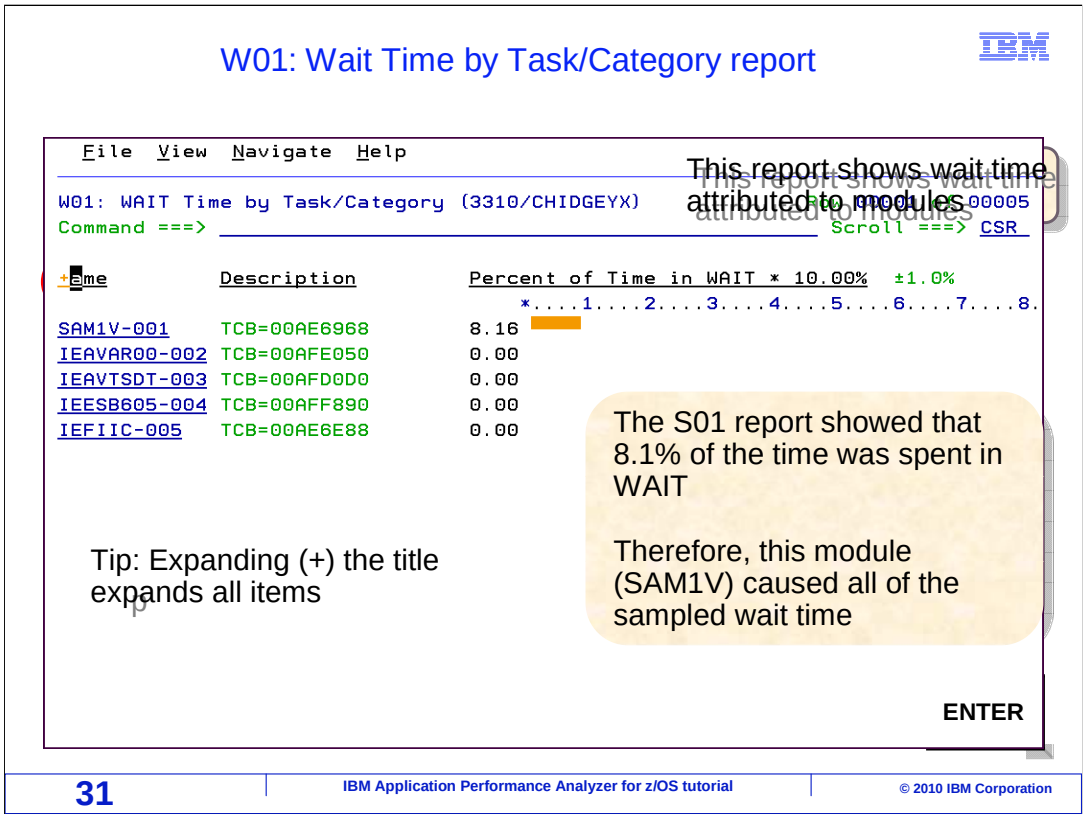

This report shows wait time that is attributed to modules. On the left is a list of program modules, and on the right the percentage of wait time caused by each one is shown. You might recall that the S01 report showed that a total of 8.1% of the time was spent in wait. In this report, that number is broken down by program. Notice that one module, SAM1V, caused all of the sampled wait time.

To expand all of the items, use a “+” line command in the column heading, and press Enter.

W01: Wait Time by Task/Category report expanded

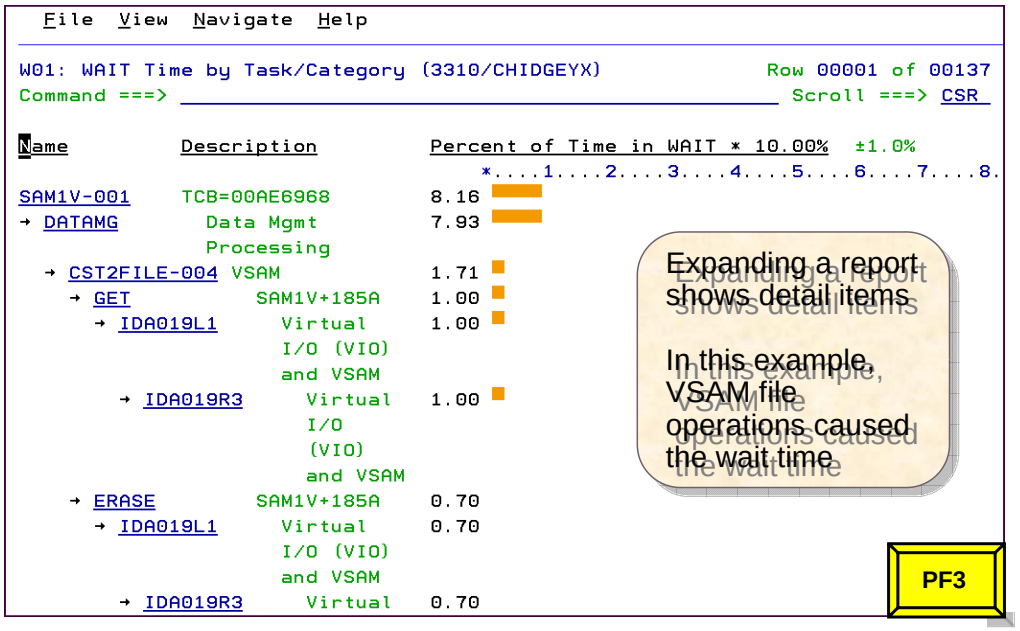

Now all of the items have been fully expanded. And notice that you can identify the individual operations in the program that caused wait time. Right under the program name is a line item for data management processing. Notice that almost all of the wait time for this program was caused by data management processing. About 7.93%.

And it is broken down further. Under data management, it shows a DD name: CST2FILE. It looks like most of the wait time was spent processing this file. Below that is a line item for GET. Notice that about 1% of the time was spent waiting because of the GET. And further down is an “ERASE” which caused .7% also..

That is good information, since it shows you exactly which programs are causing the most wait time, and what operations the program is doing to cause that wait time. In a moment, you will learn how to drill down into the program source to access the actual source statements that caused this wait time. PF3 to return.

View the W03 report

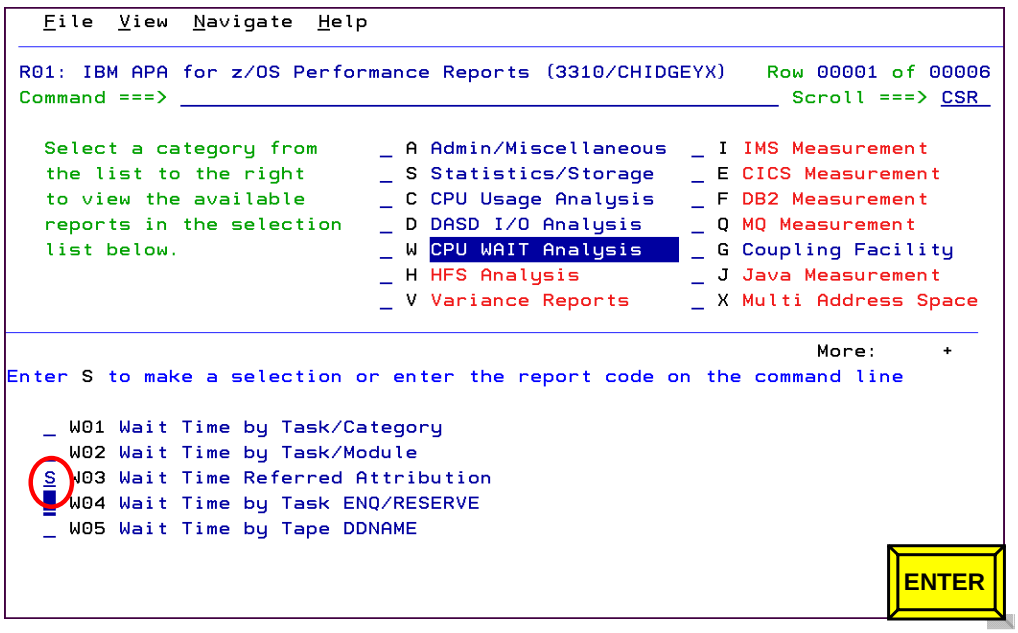

Here the ‘W03 wait time referred attribution’ report is selected.

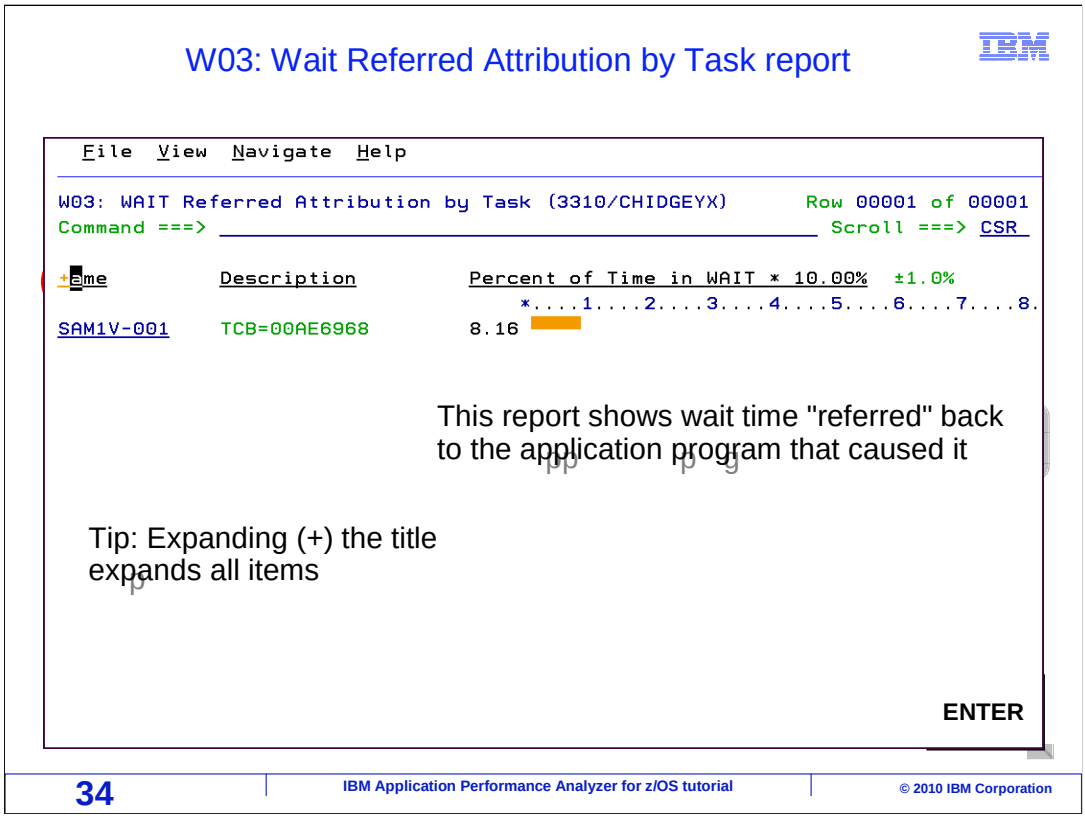

This report shows you how wait time is broken down by tasks and programs. It is the same amount of wait time that was indicated in the W01 report, but sorted a different way.

To expand the report to its deepest level, type a “+” line command in the column heading. Press Enter.

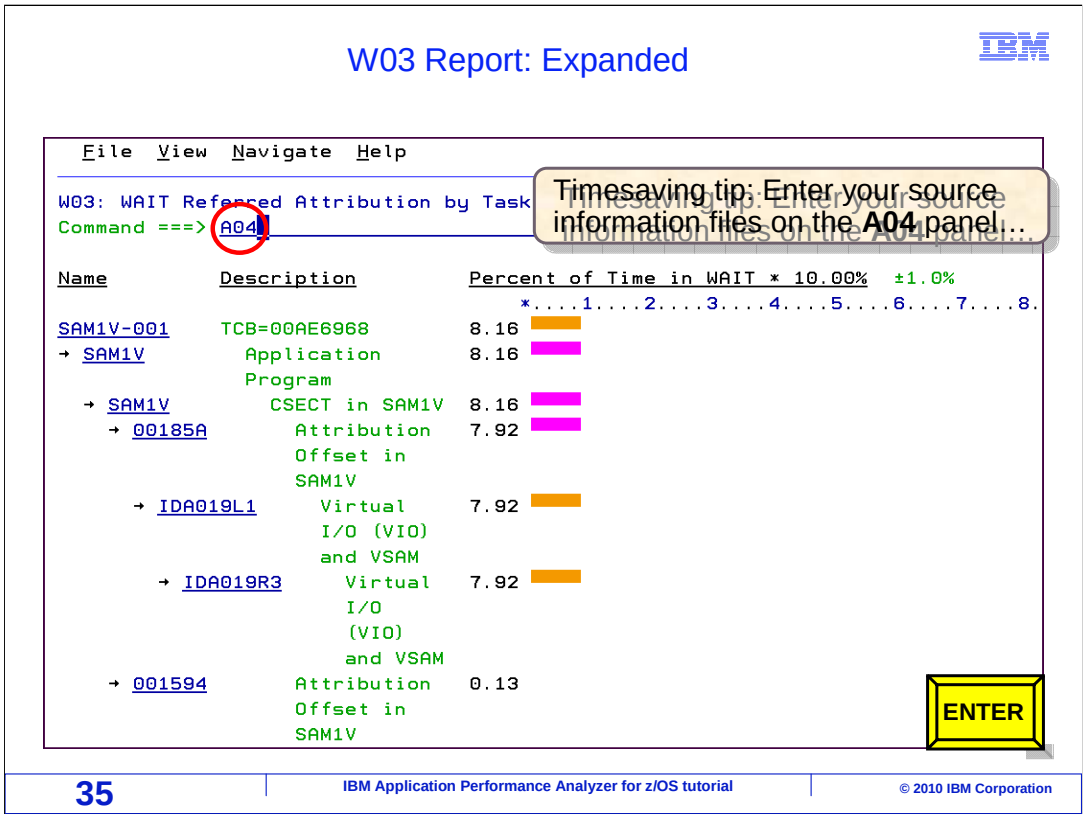

The first line shows the task called SAM1V-1. Under that task, is the program, SAM1V. The task was responsible for all of the wait time, 8.16%, and module SAM1V within the task was responsible for all of it. Within the Sam1V module is a CSECT named SAM1V. When a program CSECT is shown, you can drill down into program source code.

But before you do that, you should tell APA where to find your source information. You do that with the “A04” primary command. Just type “A04” on the command line and press Enter.

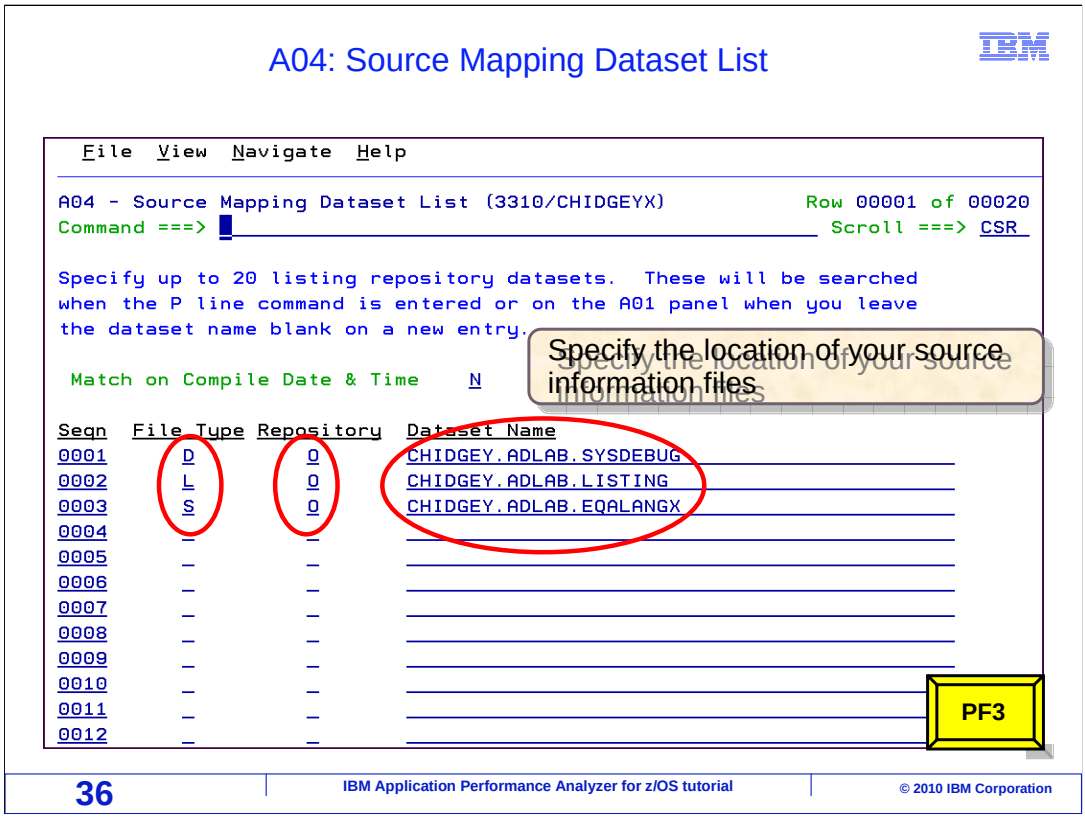

This will bring up the A04 Source Mapping Dataset List panel. Use this panel to specify the location of your source information files. If you do this, APA will automatically search the list for your source. In this example, the libraries that contain source information is specified. You can learn more about source information files in the “Working with Program Source” section of this tutorial. PF3 back.

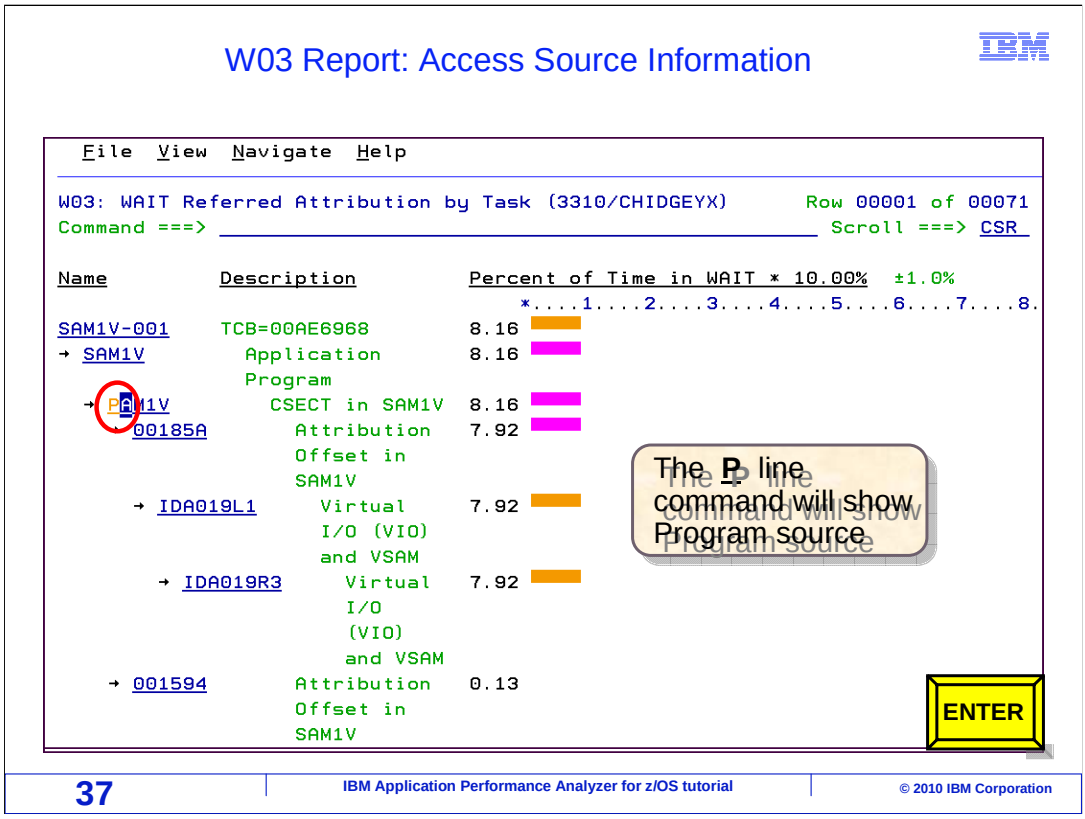

Remember that when a program CSECT is shown, you can drill down into program source code.

You do that with the “P” (for program) line command. Here, “P” is typed on the entry for CSECT SAM1V to access the source, and Enter is pressed.

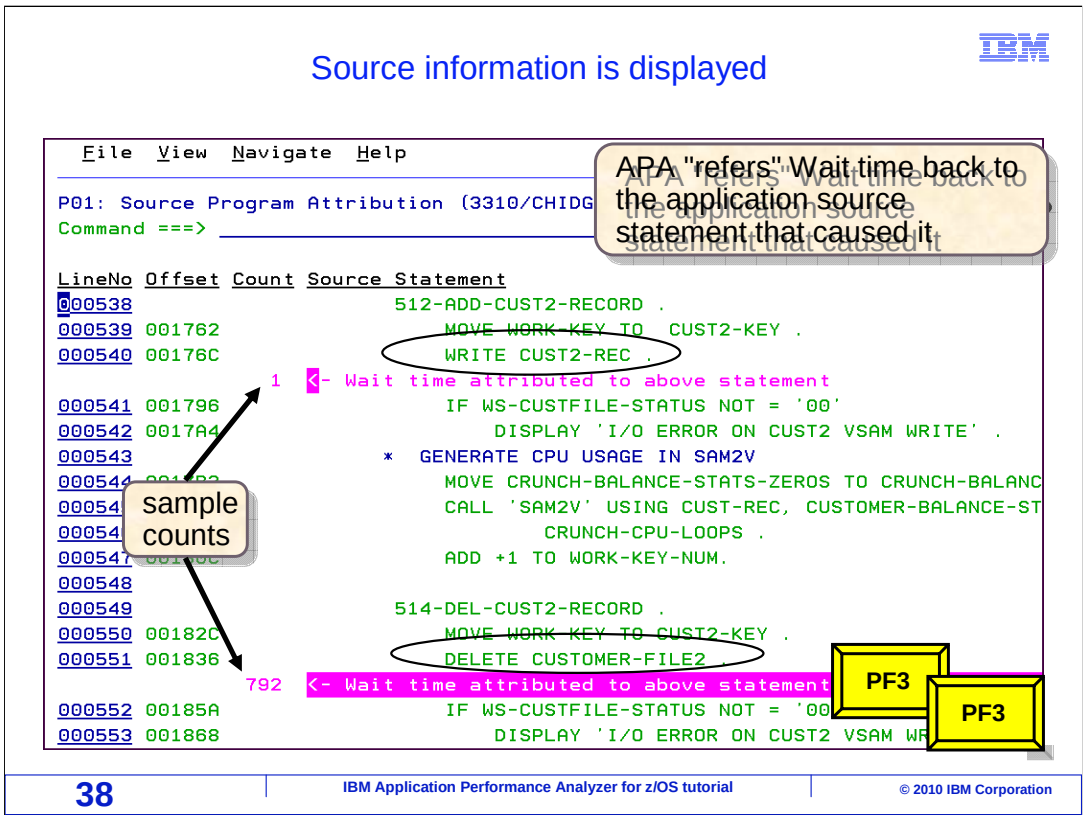

Now, APA has drilled down into the program source information using data from the source mapping file. This is a very important feature of APA. In this report, it refers wait time back to the application's source statements that caused it. This is very important, because when it comes to wait time, there is typically some system-level program that executes the code that actually causes the wait. But the system-level programs are running because they are doing something that was requested by the application.

When APA samples an application in a wait state, it traverses the call chain back to the application program. Because of this capability, you can identify exactly which of the programs statements caused the wait time, and this is extremely valuable information. In this example, there are two program statements that caused wait time. A ‘write’ statement, and a ‘delete’ statement. Notice the sample counts. APA collected one sample in wait state because of the ‘write’ statement. That is very insignificant because of the low sample count. However, it collected 792 samples in wait state because of the ‘delete’ statement.

If you want to tune this application to reduce wait time, you should probably focus on the delete statement, since most of the wait time is attributed to that statement. In this example, PF3 is hit a couple of times to get back to the report menu.

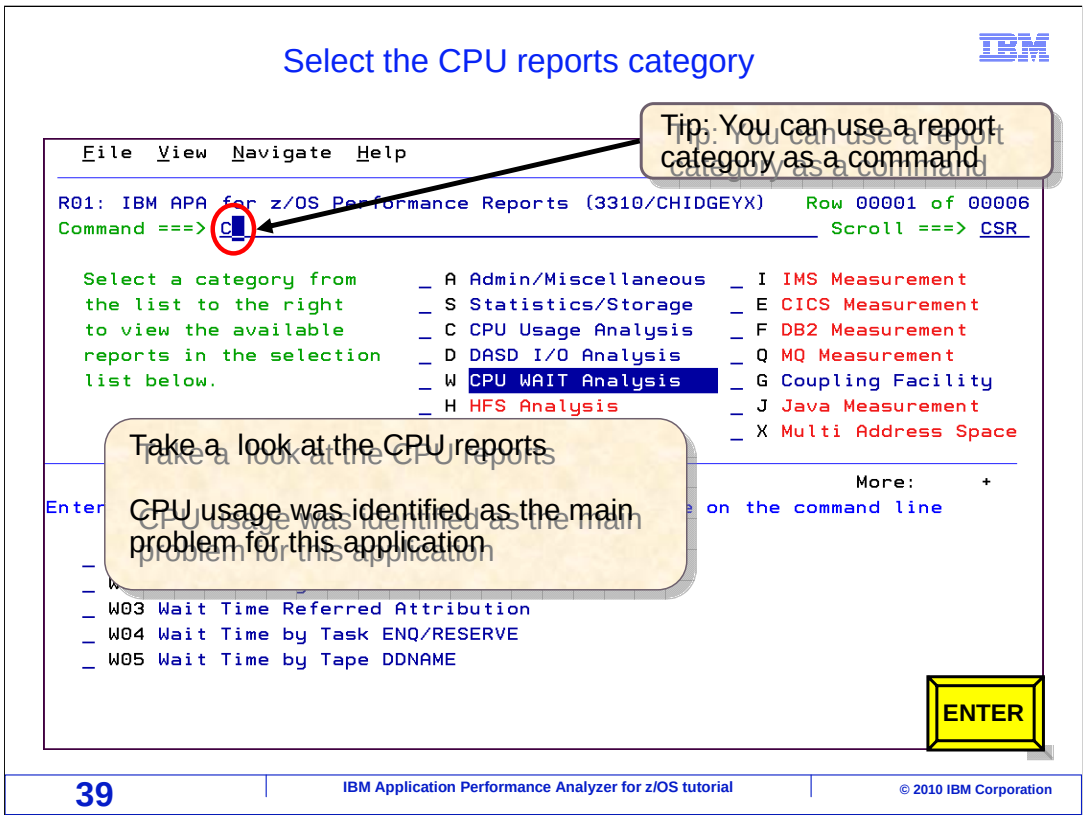

At this point, you have explored a couple of the more important reports that can be used to find out where an application spent its 'wait' time. Next, you will explore some of the CPU reports. Remember that this application spent about 88% of its time in CPU, so CPU is the main performance issue in this example. You can switch over to the CPU reports by typing "S" next to the CPU category. Or, here is another way to select a category.

Just type the first letter of a report on the command line. For instance, type "C" (for CPU) on the command line, and press Enter.

Select the CPU Usage By Category report

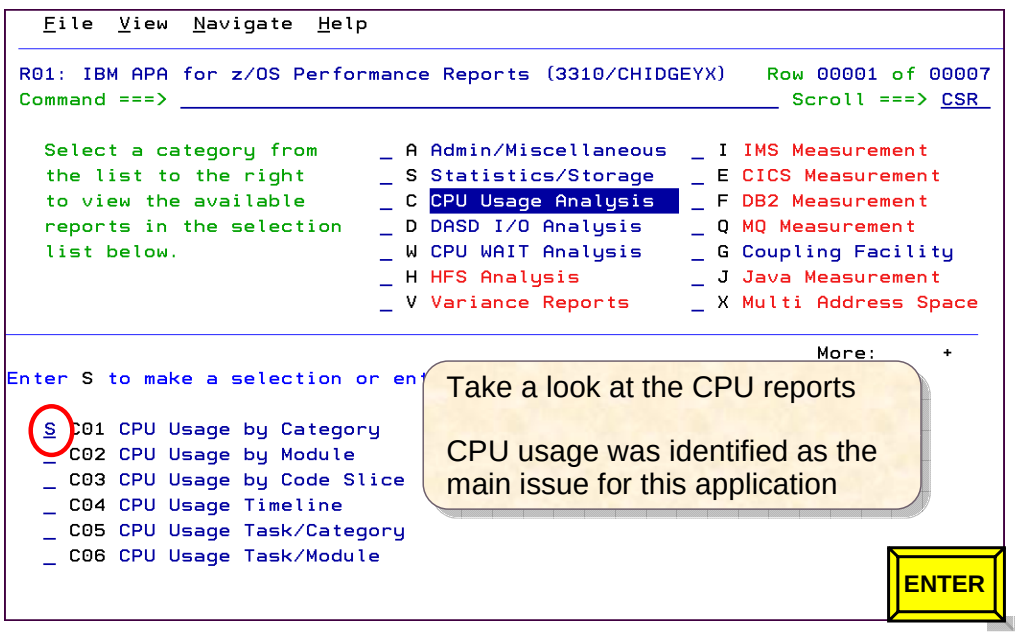

That shows the CPU reports. You can start the CPU research by examining the ‘C01: CPU Usage by Category’ report.

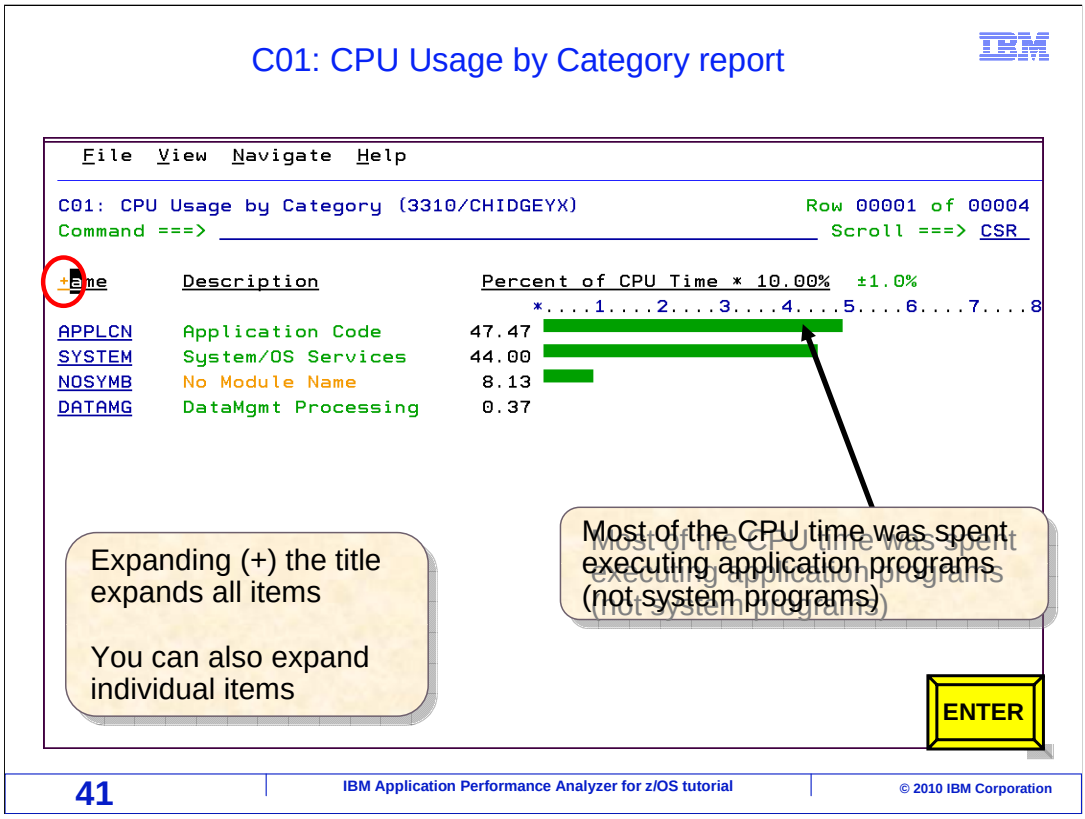

This report breaks CPU time down into categories. In this example, you know that nearly all of the CPU time was spent executing application programs, since most of the time is attributed to the application category. There are four categories listed on this report:: application, system services, no module name and data management. There are other categories, but only these four were sampled in this application. It is interesting that such a large percentage of the CPU time was spent in the application category. That means that most of the CPU time was spent actually executing the instructions within the application programs. It is typical to have more time spent in other categories. For example, a lot of CPU time might be spent performing data management operations. The fact that so much time was spent in application code is an indicator that this application might have been in a loop.

To expand all of the entries, type a “+” line command in the column heading. Press Enter.

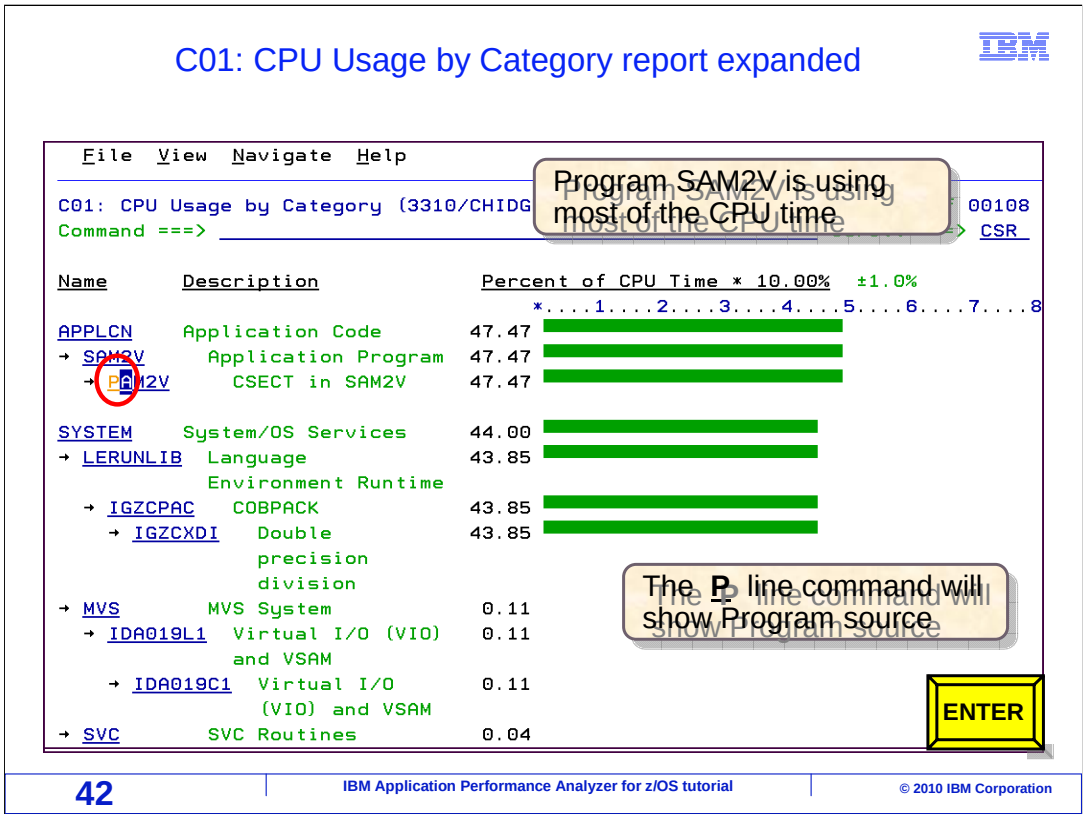

With everything expanded, you can identify which application programs used all that CPU time. In this case, all of the CPU time was used by just one program: SAM2V. You can go straight to the source code to identify which statements caused the most CPU time.

Here, a “P” (for program) line command is typed next to the CSECT entry for the program and enter is pressed.

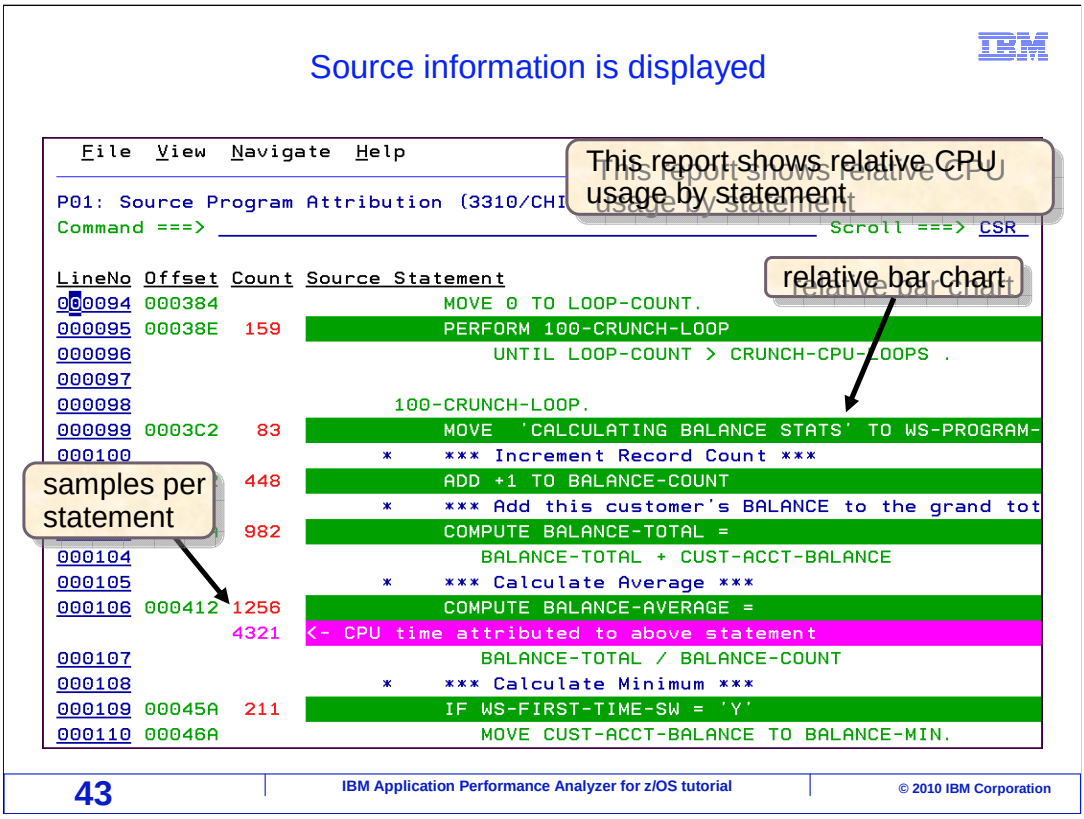

Here are the program statements that used the most CPU time. There is a bar chart overlaid on top of the program, to give you visual clues as to which statements were the heavy hitters. The bar chart is helpful, but it is more important to reference the sample counts, or the samples per statement. These show the relative impact of each statement. Notice that most of the CPU time was spent in the same area, and two of the statements (both COMPUTE) have the bulk of the responsibility.

But there are statements that did NOT show any samples. If a statement does not have a sample count, does that mean that it did not run at all? The answer is… maybe. Remember that APA is taking samples at fixed intervals, it is not monitoring every single instruction in the program.

As a result, if a program statement does not use much resource, there is a good chance that it is not running when APA takes a sample, and it will not have a sample count. However, that is OK. Because if a statement did not use much resource, when it comes to performance analysis, it really does not matter if it ran or not.

The idea is to identify the statements that use the most resource. And statistically speaking, those are the statements with the highest sample counts. The same thing can hold true for programs. Subroutines that run very quickly might not be caught in the act of running by APA. Again, that is OK, because those programs did not use much CPU or wait time. .

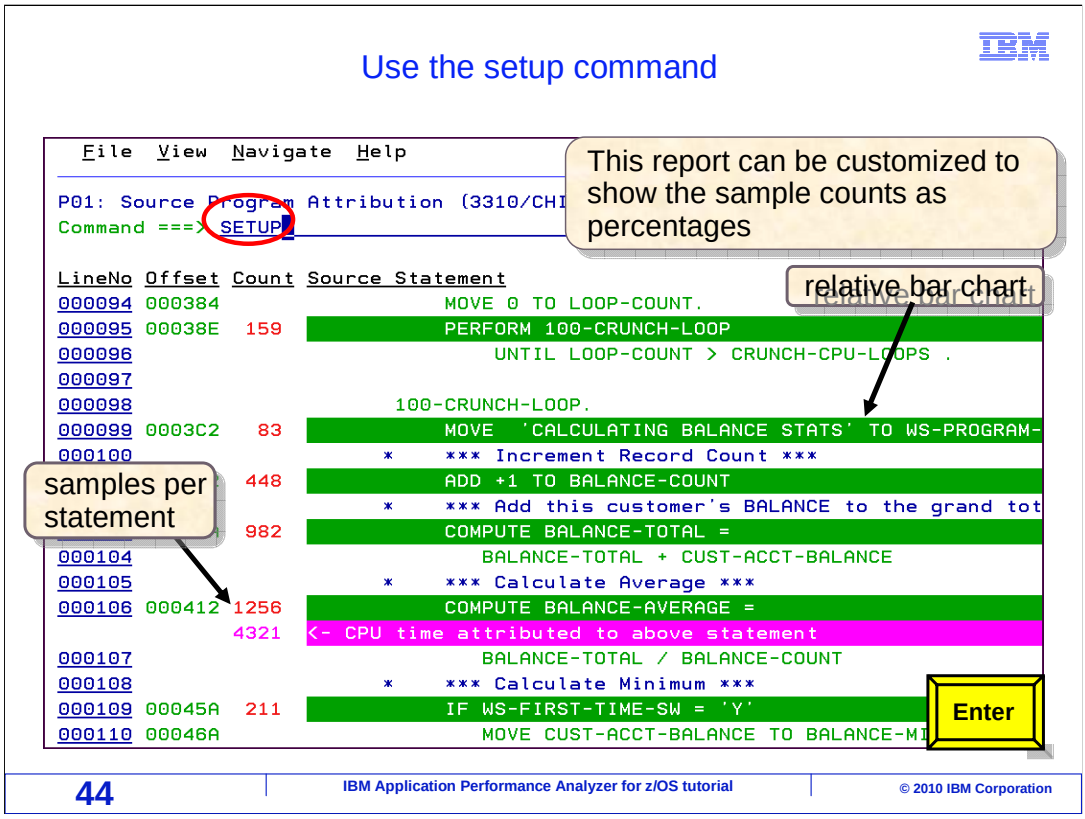

Remember that many of the reports can be customized. In this example the setup primary command is used to customize this report. Enter.

Customizing the Source Program Mapping

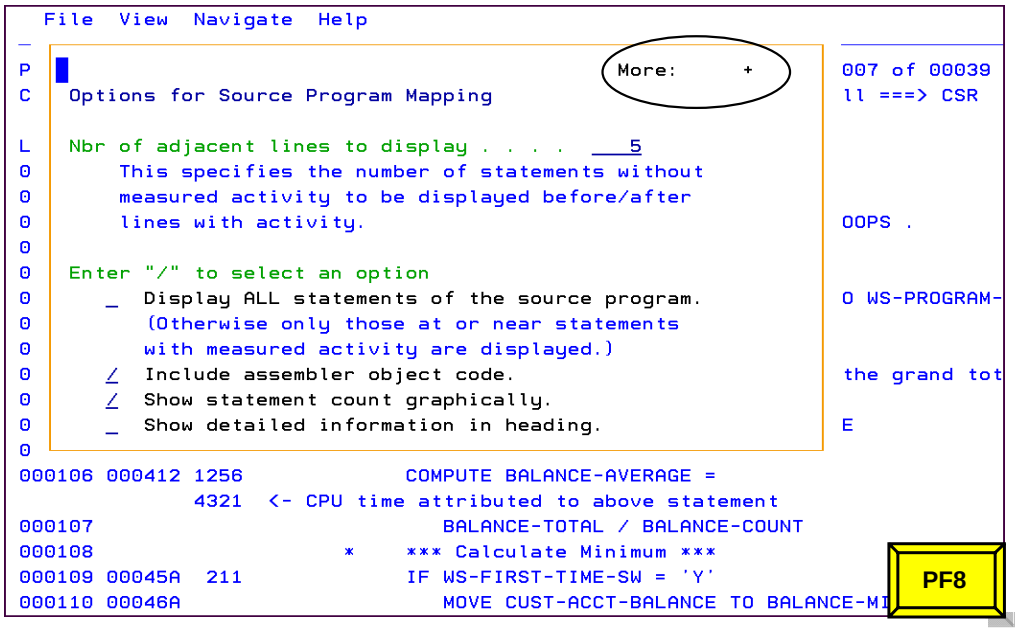

What customizations can be done depends on the report or panel where the setup command was issued. In the Source program attribution report, several options are available for customization.

PF8 to see more options.

Customizing the Source Program Mapping

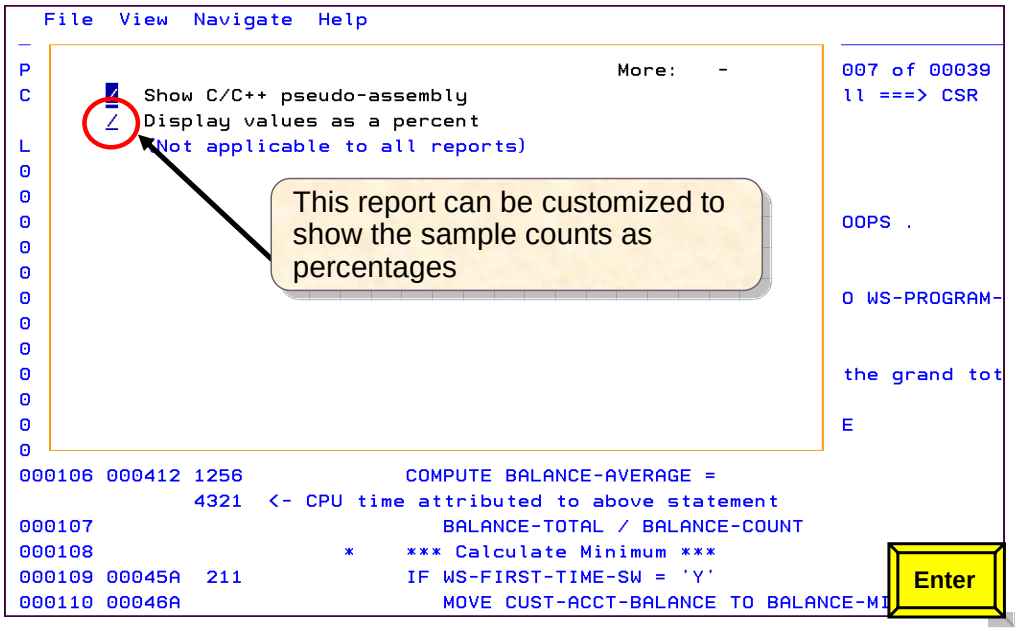

Here, a slash is placed by the option to display the sample values as percentages. Enter.

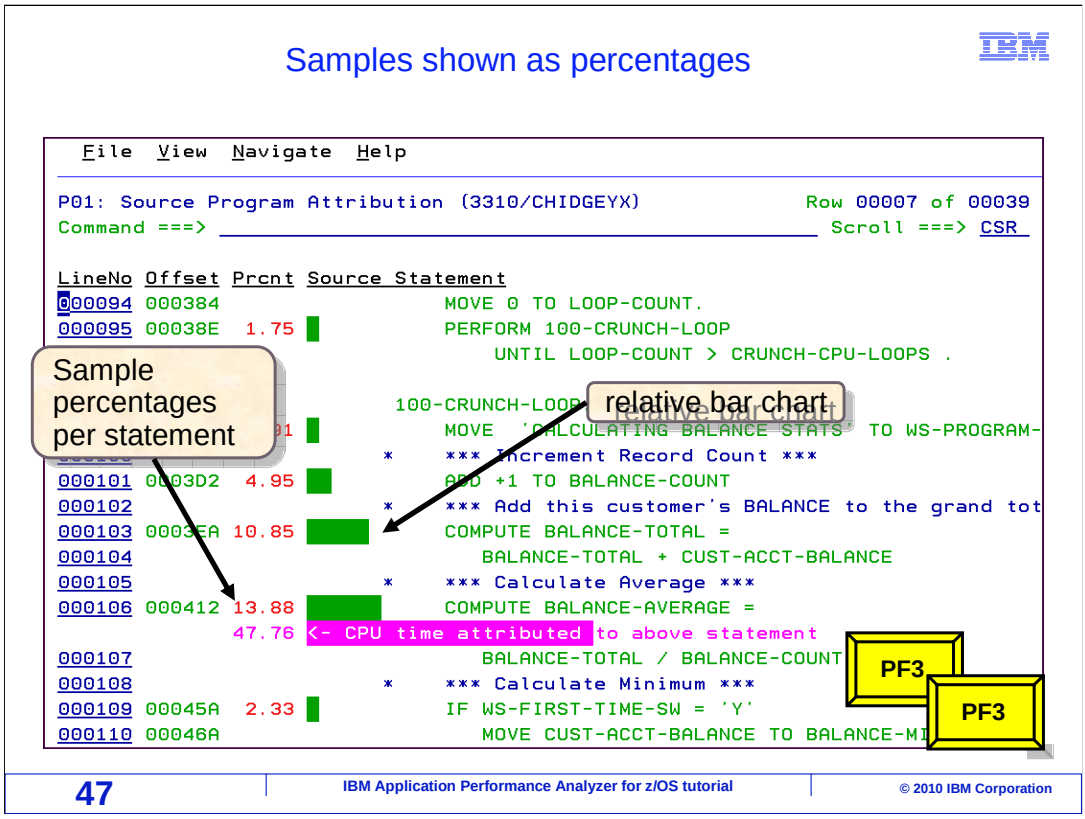

Now, instead of showing the number of samples taken while a statement was executing. This report shows the samples as percentages of the total sample count taken while the application was executing instructions. Now press PF3 a couple of times to exit this report and get back to the report menu..

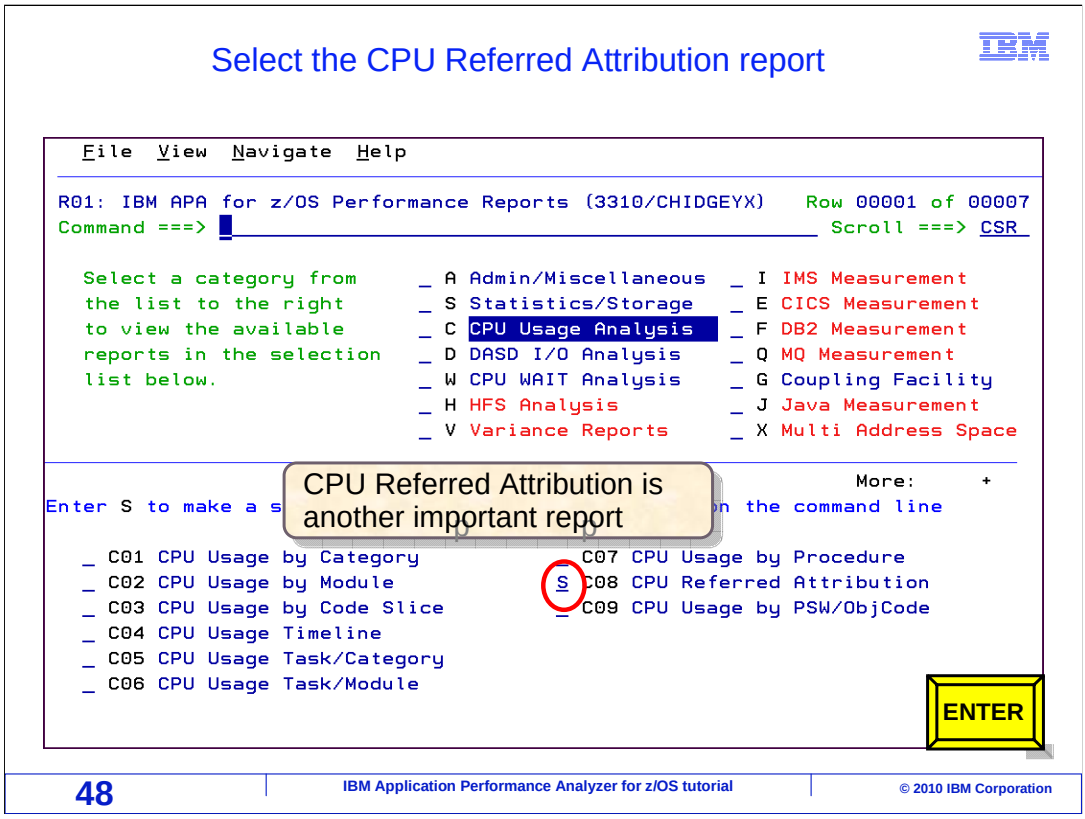

There is another report that is very important when you are researching CPU time in an application. It is the ‘C08 CPU referred attribution’ report. You can select it from the menu with an “S” line command. Enter.

C08: CPU Usage Referred Attribution report

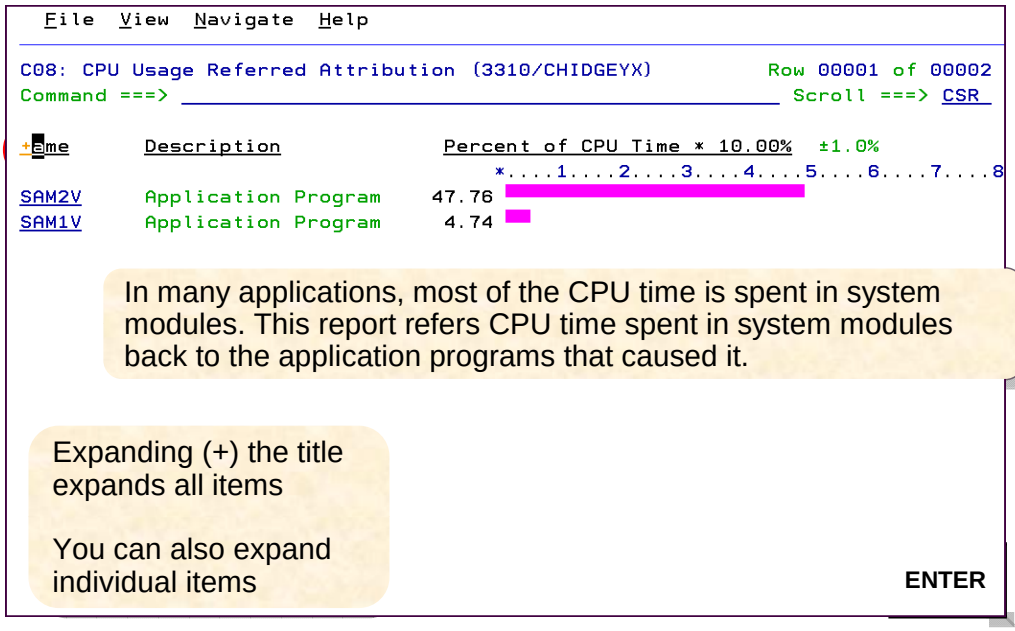

This report shows you how CPU time spent in system modules is referred back to the application program or programs that caused them to run. In many applications, most of the CPU time is spent in system modules, so this can be a very important report. Here, a “+” line command is entered in the column heading to expand all entries. Enter.

C08: CPU Usage Referred Attribution report expanded

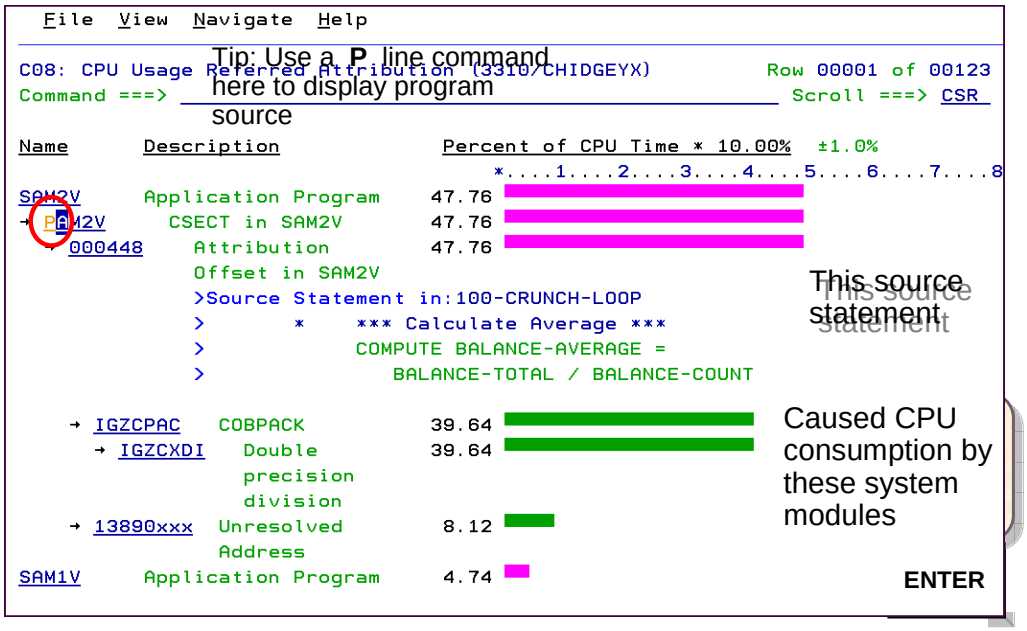

Now all of the entries are expanded. Notice that there are system level programs listed under program SAM2V. These are all programs that are run due to service requests made by SAM2V. This is extremely important information, because it tells you exactly which statement caused the lower-level system routines to run. In this example, the ‘Compute’ caused an excessive amount of CPU time to be used by system-level routines.

If you use a “P” (for program) line command next to the CSECT entry for the application program , you can drill down into the program source. Hit Enter.

Tuning the application

- At this point, there is enough information to start tuning the application
  - The application spent most of its time in CPU Active
  - Most of the CPU time was spent executing instructions in application program SAM2V or system modules invoked by SAM2V
  - These source statements below caused most of the CPU time

```
000100                      *    *** Increment Record Count ***
000101 0003D2  4.95 ▓▓▓          ADD +1 TO BALANCE-COUNT
000102                      *    *** Add this customer's BALANCE to the grand tot
000103 0003EA 10.85 ▓▓▓▓▓        COMPUTE BALANCE-TOTAL =
000104                              BALANCE-TOTAL + CUST-ACCT-BALANCE
000105                      *    *** Calculate Average ***
000106 000412 13.88 ▓▓▓▓▓▓       COMPUTE BALANCE-AVERAGE =
           47.76 <- CPU time attributed to above statement
000107                              BALANCE-TOTAL / BALANCE-COUNT
000108                      *    *** Calculate Minimum ***
000109 00045A  2.33 ▓            IF WS-FIRST-TIME-SW = 'Y'
000110 00046A                       MOVE CUST-ACCT-BALANCE TO BALANCE-MIN.
```


At this point, there is enough information to start tuning the application. You have learned that this application spent most of its time in the CPU active state. Most, in fact, nearly all, of the CPU time was spent executing instructions in an application program called SAM2V. These source statements caused most of the CPU time. Statements 101 through 110 are the primary culprits.

To recap how you can use APA to analyze an application's performance. Remember that the first order of business is to identify whether most of the time is being spent in CPU or wait state. After that, you know whether to focus your efforts on researching CPU or wait. Then find the program or programs that are causing most of the time, and if possible, drill down into the application source code to determine which statements are responsible.

Tuning the application

IBM

- What can be improved in this application?
- Statements 101 through 110 use a lot of CPU because they are inside an iterative loop, and do not need to be.
  - Improvement: move them outside the loop
- Many of the numeric work variables are in character format, such as BALANCE-COUNT and BALANCE-TOTAL
  - Improvement: define them as packed or binary numeric

```
000100                       *     *** Increment Record Count ***
000101 0003D2  4.95                ADD +1 TO BALANCE-COUNT
000102                       *     *** Add this customer's BALANCE to the grand tot
000103 0003EA 10.85                COMPUTE BALANCE-TOTAL =
000104                                BALANCE-TOTAL + CUST-ACCT-BALANCE
000105                       *     *** Calculate Average ***
000106 000412 13.88                COMPUTE BALANCE-AVERAGE =
              47.76 <- CPU time attributed to above statement
000107                                BALANCE-TOTAL / BALANCE-COUNT
000108                       *     *** Calculate Minimum ***
000109 00045A  2.33                IF WS-FIRST-TIME-SW = 'Y'
000110 00046A                         MOVE CUST-ACCT-BALANCE TO BALANCE-MIN.
```


Using this example, what can be improved in this application? Statements 101 through 110 used a lot of CPU time because they are inside an iterative loop, and they do not need to be. An improvement can be to move them outside the loop. Also, a lot of the numeric work variables are defined in character format. That causes the program to convert them into a numeric format every time they are compared or used in a computation. It is more efficient to define them, for example, as packed decimal or binary. Just so you know, making those changes to this example program reduced the CPU time by more than half.

At this point, you have learned how you can use APA to research performance problems in your applications. There are a lot of reports available, and so far you have only gone through a few of the most important ones. However, now you should be able to navigate to reports, and research your applications using which ever reports is helpful based on your applications.

-

That is the end of this section, a walkthrough of a performance analysis.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_APAv10s05AnalysisWalkthrough.ppt

This module is also available in PDF format at: ../APAv10s05AnalysisWalkthrough.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, CICS, DB2, IMS, MQSeries, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

Java, and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.