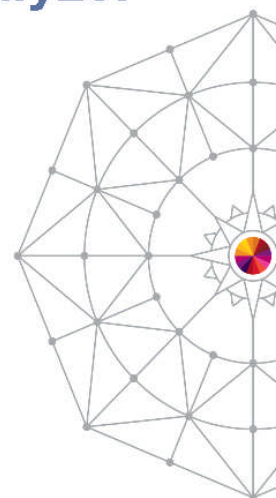# System z Batch Network Analyzer (zBNA) Tool Hands-on Lab

John Burg
Brad Snyder
Valerie Spencer

IBM

August 7, 2014
Session Number 15671

zBNA Lab Guide

The purpose of this zBNA Lab is to provide an exercise in running the zBNA tool; utilizing its functions to successfully complete a simple Batch analysis.

In this exercise you will complete the following tasks:

1) Explore the Main Screen
   Start zBNA and load in two data files
2) Filter Data
   - Use the job filtering capabilities (CPU time, Service classes, exclude jobs, key jobs and job masking) to select a subset of candidate Batch jobs
   - Save as zBNA File
   - Filter Top Program Pct
   - Load Step level records, and drill down into the Step details
3) Display a Graph and Create Reports
   Display the job subset created with the filters
4) Display SMF 42(6) DASD Dataset Analysis
   - Job/Dataset Report
   - Top 10 Dataset Report
5) Perform Alternate Processor Analysis
   Assess the impact of an alternate CPU technology
6) Explore zEDC Compression
   Identify data sets that will benefit from moving to zEDC cards
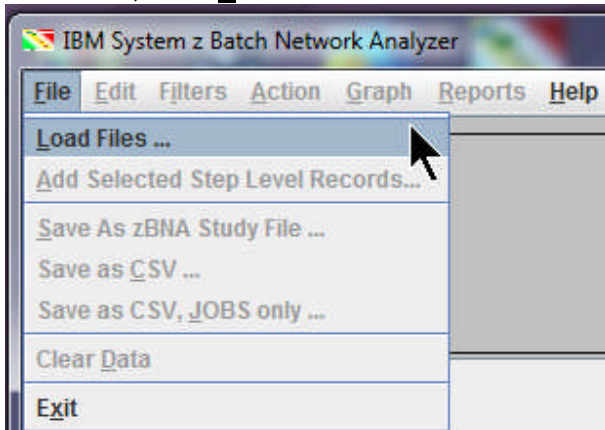7) Save the final zBNA file

zBNA Lab Guide

## Task 1 - **Exploring the Main Screen**

1. To start the System z Batch Network Analyzer (zBNA), first double-click the icon.
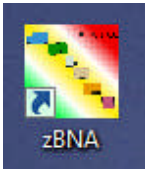
2. Click *File*, then *Load Files …*

3. If this is your first time using the zBNA tool, select the SMF70 (.edf) and z/OS SMF (.dat) files by clicking the appropriate *Browse* buttons. Navigate to **C:\CPSTOOLS\zBNA.** Both files are required to be loaded together. Note that a previously saved study file (.zBNA) is required to use the *Browse For zBNA File* button, in addition to the original SMF70 (.edf) and z/OS SMF (.dat) files.
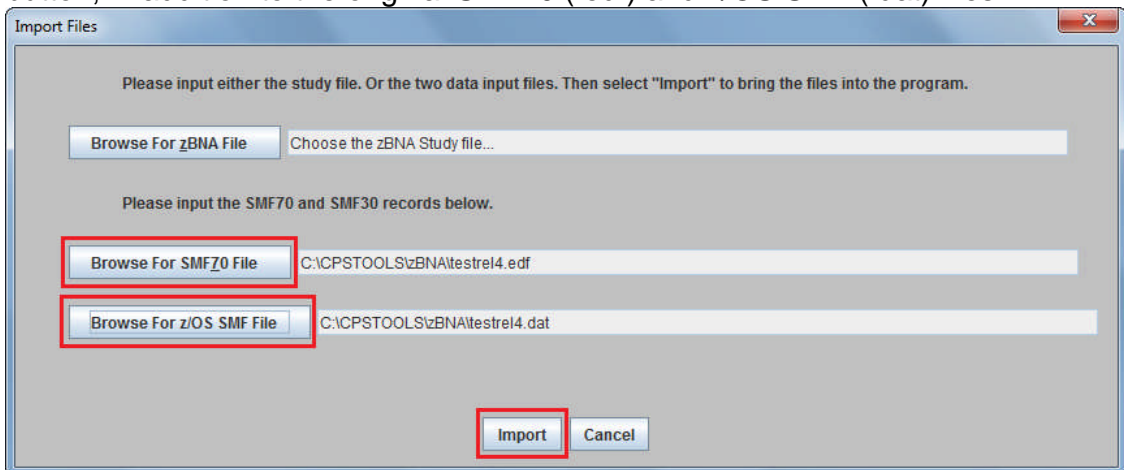
    The SMF70 file name is **testrel4.edf** and **testrel4.dat** for the z/OS SMF one. Click ***Import***.

4. The zBNA tool will load the desired data in tabular format with job information displayed. At the bottom of the panel the messages indicate that **5147 jobs** have been loaded from **JOB end records (SMF 30 subtype 5)**.

5. Individual jobs may be selected with a single click. Right-clicking the first job, **M373Q3S**, displays a menu. Select ***Show Step Details*** to display the Step details. However, at this point, zBNA will only display Job End record information (not Step Detail) because the SMF 30 subtype 4 data has not been loaded. Once filtering is completed later, the Step Detail records will be loaded.

6. The Job Information panel displays the specific job information at the top of the **Steps** table.

Note: The job details will be displayed once you have performed **File**, **Add Selected Step Level Records** (performed after the **Filtering** process is complete) on the zBNA main panel.
In the **Step Number** column "**7 Total**" refers to the total number of steps in this job, **M373Q3S**. Also, notice that there is a scroll bar so that all of the various fields can be seen. Click *OK* to return to the main panel.

7. Jobs may be sorted by any parameter on the screen in both ascending and descending order, simply by clicking on the corresponding column header. Click the *CPU Time* column twice to sort from the largest to smallest values. Also note that the number of jobs in the screen, displayed in the bottom left-hand corner, is still currently **5147 jobs**.

zBNA Lab Guide

Task 2 - **Filtering Data**

1. There can be data from hundreds or thousands of jobs. It is often necessary to filter the jobs based on some criteria to reduce the number to be more manageable for the analysis, and Filters can help reduce the number. To apply a filter, click **Filters, Set Table Filters**…
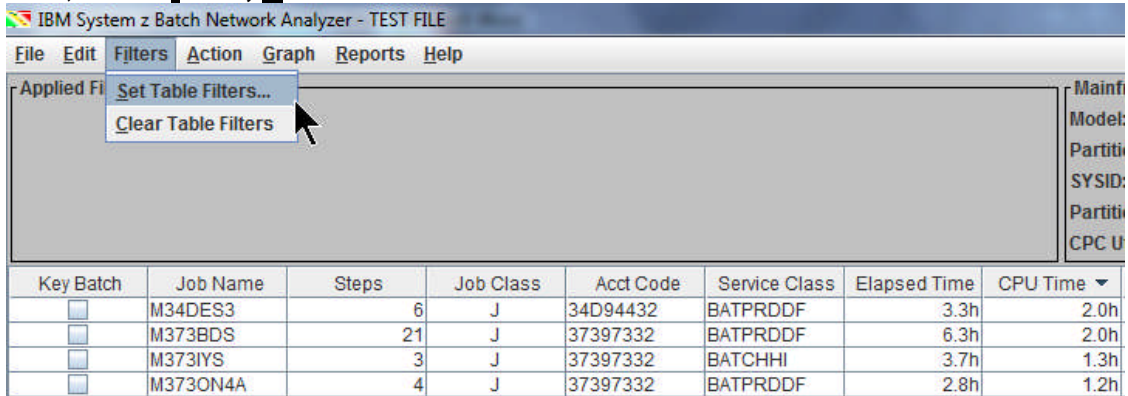
2. Start by setting the **GCP Time** to **10 seconds**, which will filter out jobs that took less than 10 seconds of CPU during the job's elapsed time. Clicking on another option or pressing Tab will allow the changes to take effect. The purpose of setting the filter is to remove a number of jobs that took an extremely small amount of CPU resource, so that the focus on the analysis can be on the jobs with the most amount of impact.
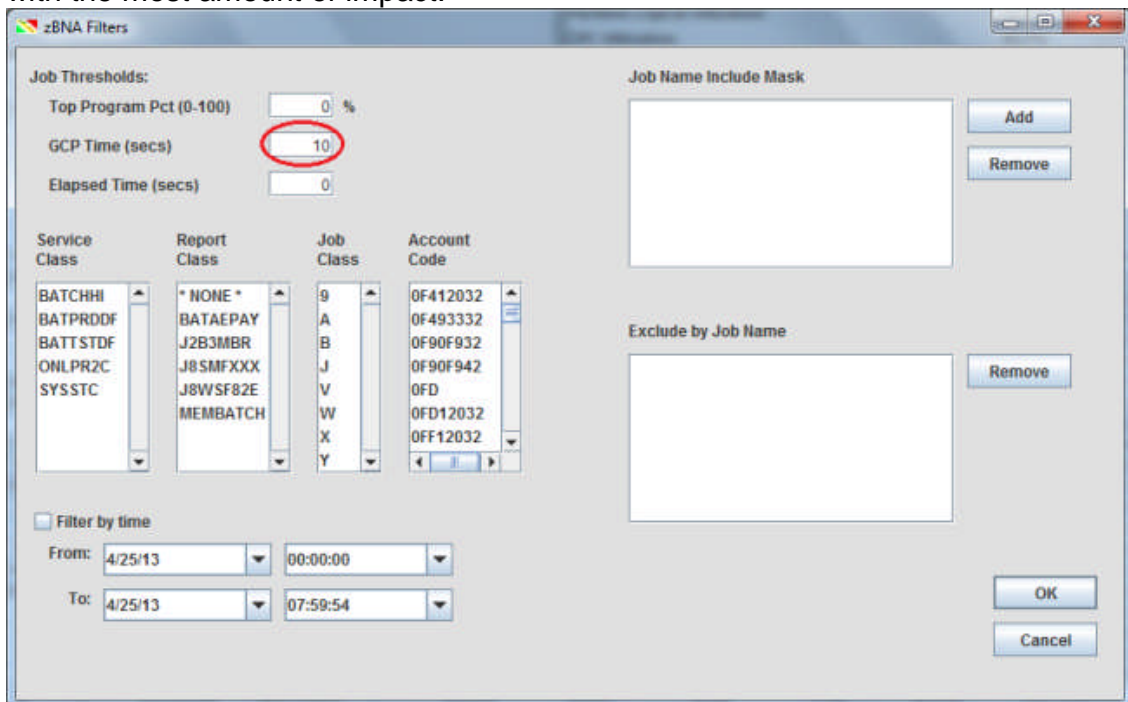
3. Next, filter by the Service Class name. This allows one to filter on the WLM construct that is already aligned to business importance and classification. Multiple Service Classes may be selected by holding the Control key while clicking the desired service class names. Select **BATCHHI**, **BATPRDDF**, and **BATTSTDF**. Note that there are now **938** jobs in the table. To remove a selection, hold Control and click it again. Similarly, one can filter by Report Class, Job Class, or Account Code, if desired.

4. Job names may also be filtered by clicking **Add**. Specific jobs can be named, or only parts of the name may be used, followed by an asterisk, which will match any number of characters. Please add **M4*** and **M3***, as separate entries, to the Job Name Mask. Click **OK** after keying in each Job Name Mask.

5. As shown below, **M4*** will find all jobs starting with M4, and **M3*** will find all jobs starting with M3. Note that there are now **874** jobs on the main panel. Click *OK* to return to the main panel.

6. Even with a filtered list there may be jobs that should not be included. These can be excluded from the analysis. Select the line for job **M373DVF** and right-click; select ***Exclude Data*** to remove it from the table.

7. Note that returning to the zBNA Filters panel shows that job in the Exclude by Job Name list. There is one less job, now **873 Jobs**.

Click *OK*.

8. If there are key jobs that you would like to focus on, select those in the **Key Batch** column. These will **always** be included in the analysis regardless of the job filter definitions. Select the following jobs as key: **M373BJ5**, **M402GX3L**, and **M3E0IKSN**.

Now sort the *EXCPs* column in descending order to view the values from largest to smallest.

Let's find job **M4E07B1H**, which has **3,028,474 EXCPs,** in the table. You can either slowly scroll down the table to job **M4E07B1H** or use the **Edit**, **Find** (Ctrl+F) function (the search argument **must be** upper case). Click the **Key Batch** checkbox. Note that the other three *Key Batch* jobs are still selected, however, they just are not in this view since we performed the sort by EXCPs. **NOTE**: This technique of identifying jobs as *Key Batch* can be used to keep known jobs always in the analysis (e.g. critical path jobs, high importance, etc.) so that other filtering techniques do not inadvertently remove them. There is also the ability to separately report on these "Key" jobs.

At this point, let's stop and save the current filters that have been set along with the four jobs that are identified as key in a zBNA study file.
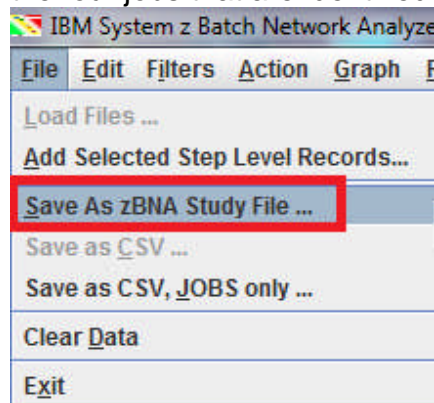
Name the file **testrel4** (".zBNA" will automatically be appended to the file name), and click **Save**.

9. Return to the zBNA Filters panel and set the ***Top Program Pct*** to **10%**, which will only include jobs where a Top Program is 10% or greater. Note that there are now only **36 jobs** in the table, including the four that we selected as Key Batch jobs. Click ***OK.***

10. Let's add the job step data (SMF Type 30 subtype 4 records). Click **File**, **Add Selected Step Level Records**.

11. The main zBNA panel is redisplayed. Now a job can be drilled down to show the step level details. (Note that the message "**Only JOB end records (type 30 subtype 5) have been loaded**" **is no longer displayed** in the information bar).

Let's sort on the **Elapsed Time** column so that the longest running job is the first one displayed in the table.

Job **M373BDS** is the longest running job in this filtered set. You can see that the elapsed time is **6.3 hours** and had **21 Steps**. Right click on that job, and select **Show Step Details**. **Note**: Double clicking in the job row will perform the same task.

12. The details on the steps are displayed. One row per each Step is provided, and all the columns for the Job level are provided for each Step. Remember to use both the vertical and horizontal scroll bars to view all of the information.

The detailed information on each step of the job includes:
- start/end time and date
- step name
- program name
- step number
- sub type
- job class
- account code
- service class
- report class
- elapsed time
- CPU time
- zAAP time
- zIIP time
- IIP CP time
- EXCPs
- CPU intensity
- Top Program
- Top PGM %

These step level fields may be useful once you've identified a job that you want to reduce the elapsed time, because you'll be able to identify the step and program level resources to know where to focus tuning or alternative technology.
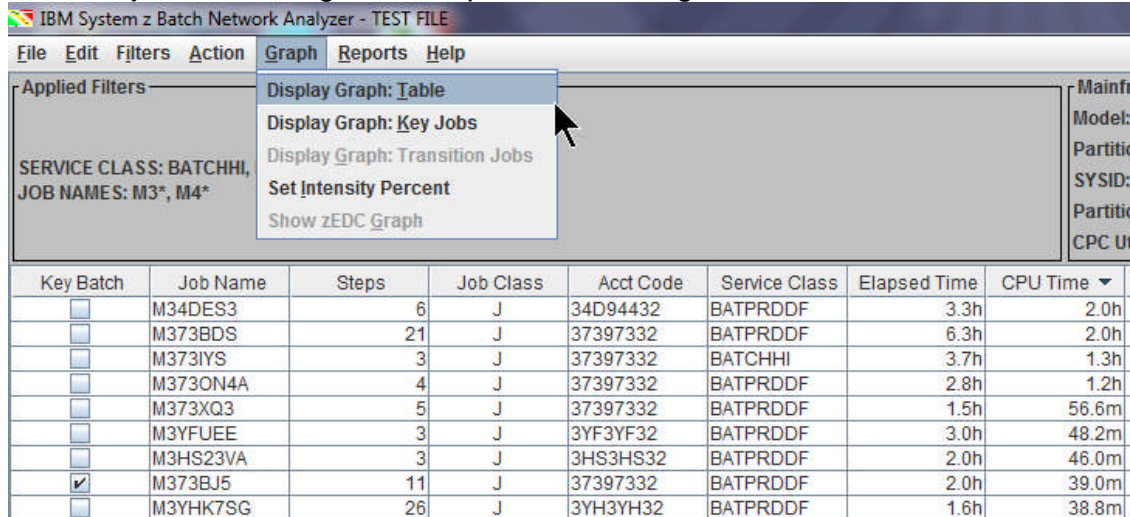
Click OK to return to the zBNA main panel.

zBNA Lab Guide

Task 3 - **Displaying a Graph**

1. The data in the table on the main zBNA panel may also be displayed in a graph format by selecting **_Graph_** then **_Display Graph: Table_**. This will graph the selected jobs remaining from the previous filtering.

2. Each job appears on its own line of the graph.
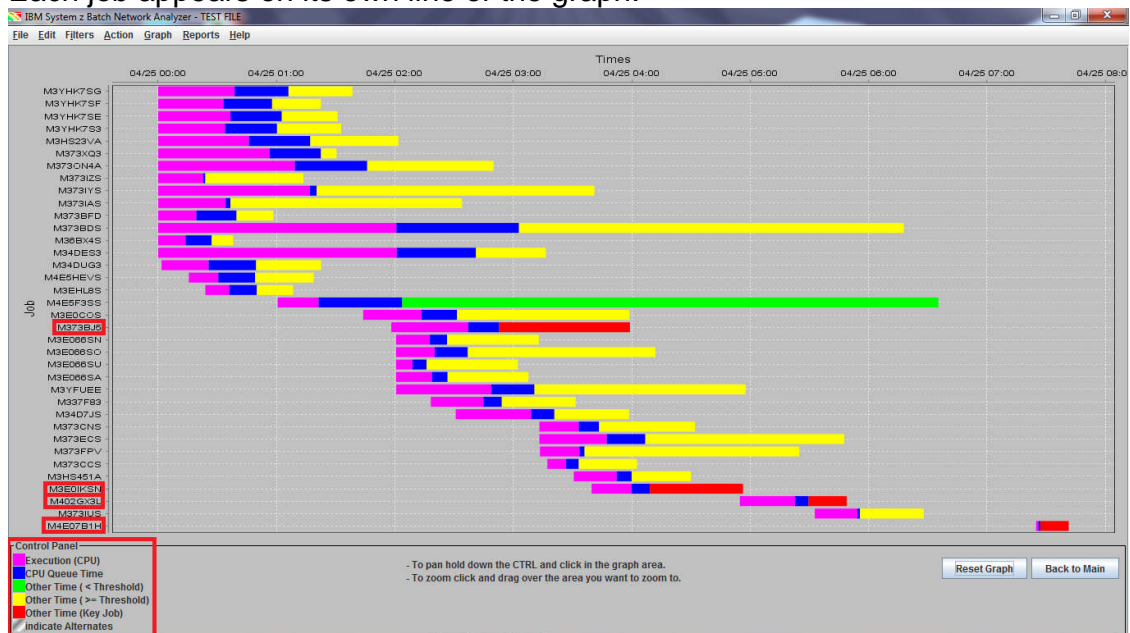
The Elapsed Time for a job is the sum of **CPU Time + CPU Queue Time + Other Time**. Other time is all other time, and is typically comprised of I/O time. The sum of the 3 components is placed on the X axis when the Job's Elapsed Time occurred in the interval, but they represent the % of time spent in each component (e.g. the actual CPU Time does not all occur at the beginning of the job).

The legend for the graph appears in the bottom left corner.

- Pink, **Execution (CPU Time**, shows the measured CPU time for a job.
- Blue, **CPU Queue Time**, represents the estimated CPU wait time for a job, which is calculated from the *RMF Service class waiting for dispatch* field.
- **Other Time**, a green bar signifies that the job's CPU execution time is less than 10% (default value for **Set Intensity Percent**) of the job's duration.
- **Other Time**, a yellow bar signifies that the CPU execution time is more than 10% (default value for **Set Intensity Percent**) of the duration.
- **Other Time**, a red bar signifies *Key batch* jobs.

3. Clicking and dragging an area on the graph will zoom the graph in to that area.
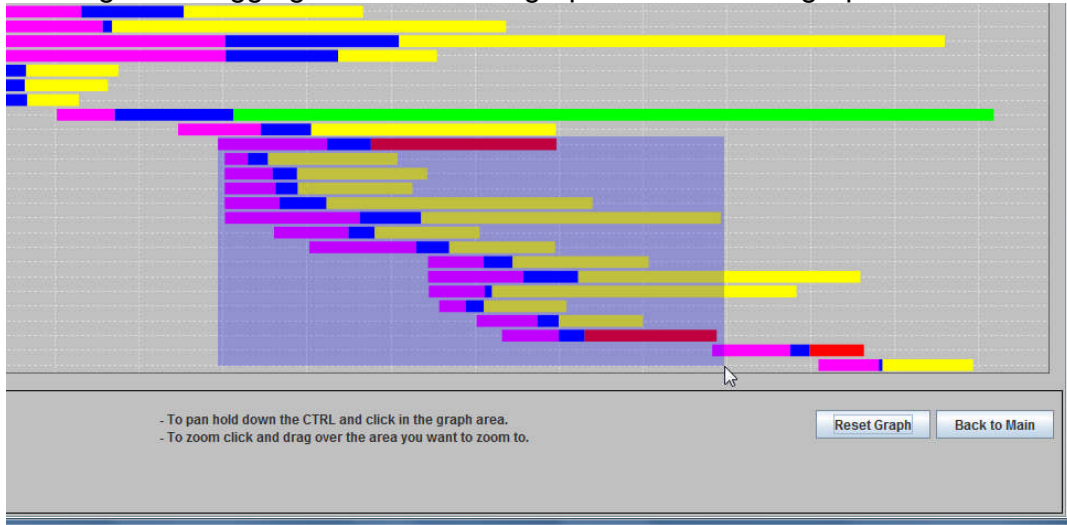
4. Holding Control allows the user to pan across the graph. The cursor will become a cross when this is happening.
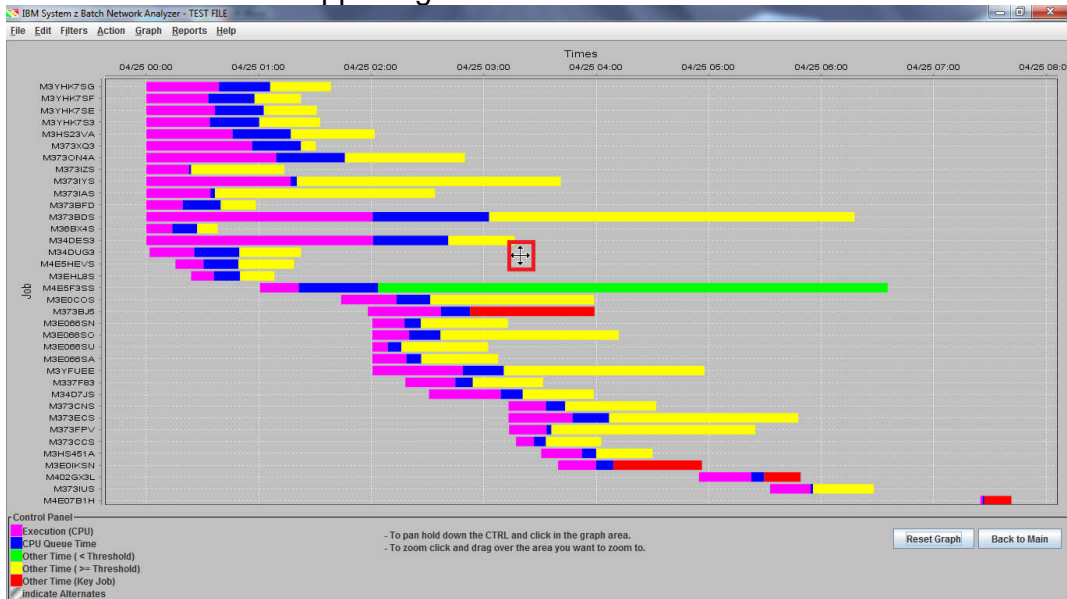
Click **Reset Graph** to show the original graph.

zBNA Lab Guide

5. Hold the mouse over a job to show the Job information.
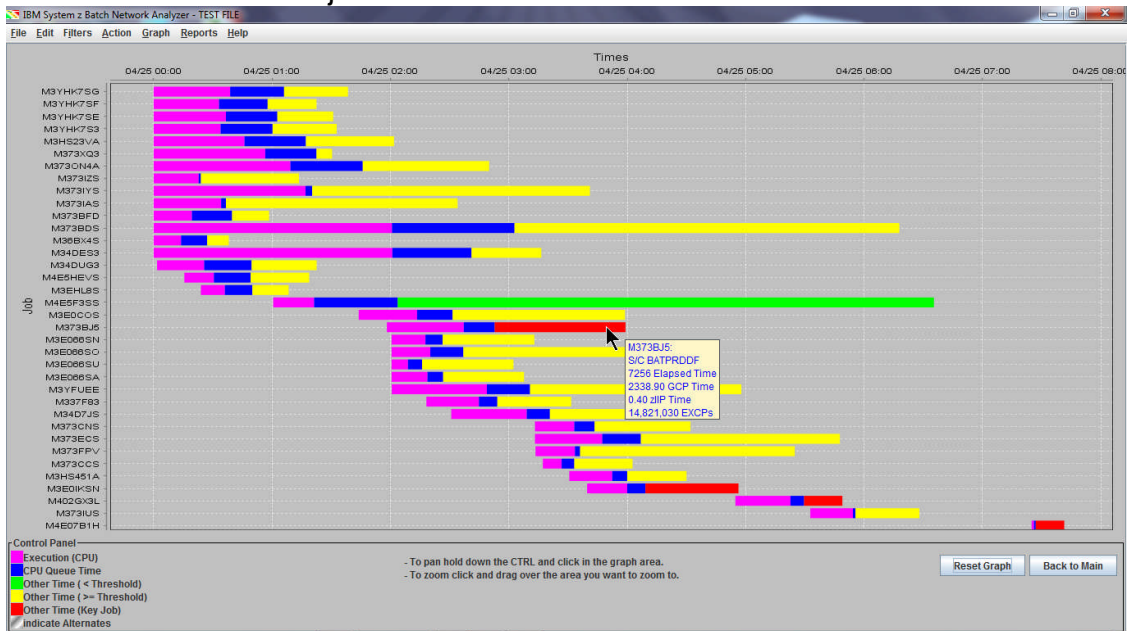
Further detail is available for each job by right-clicking and selecting *Show Step Details*. Right-click **M373BJ5** (the first Key job with Red Other time) and click ***Show Step Details***.
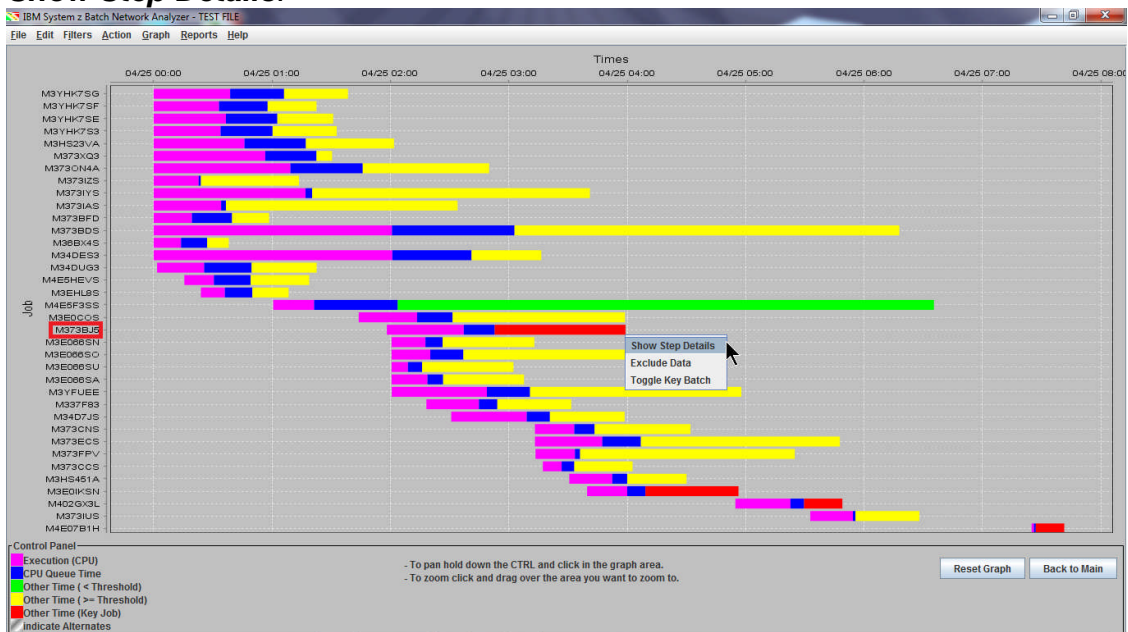
6. The same Job Step panel that is accessible from the main panel displays.
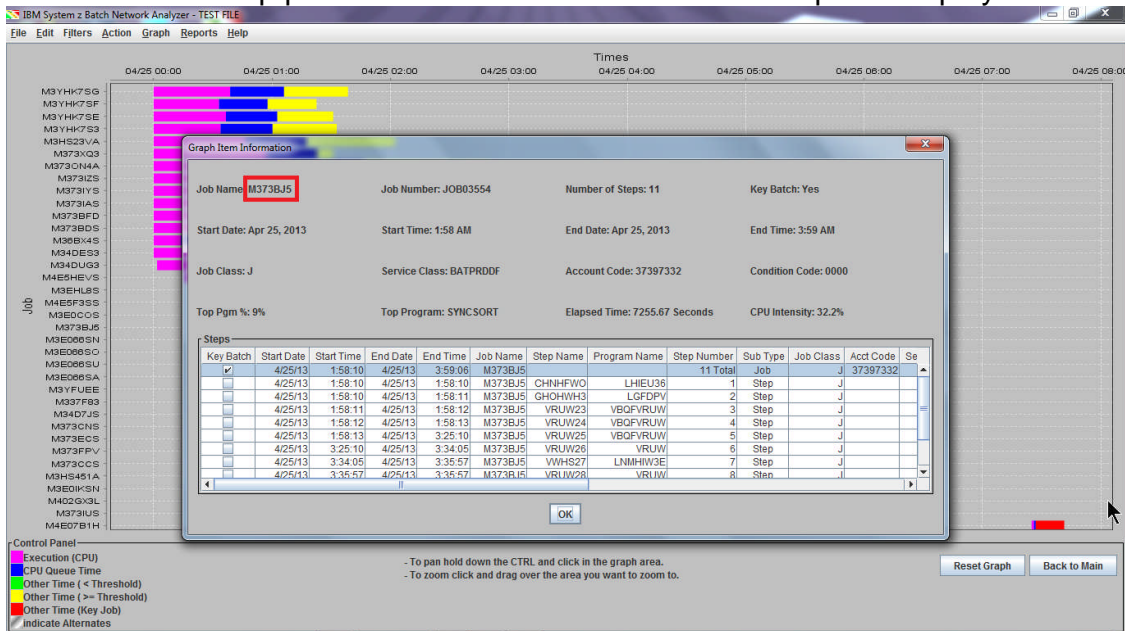
Click **OK** to return to the graph.

7. A graph report can automatically be created by using the **_Reports_** menu while displaying the graph. Click **_Generate Graph Report_** then select what job attribute (Start Time, Job Name, CPU Time, Elapsed Time) you would like the data sorted in the table that is included in the report. Select "**Sort By Start Time**".
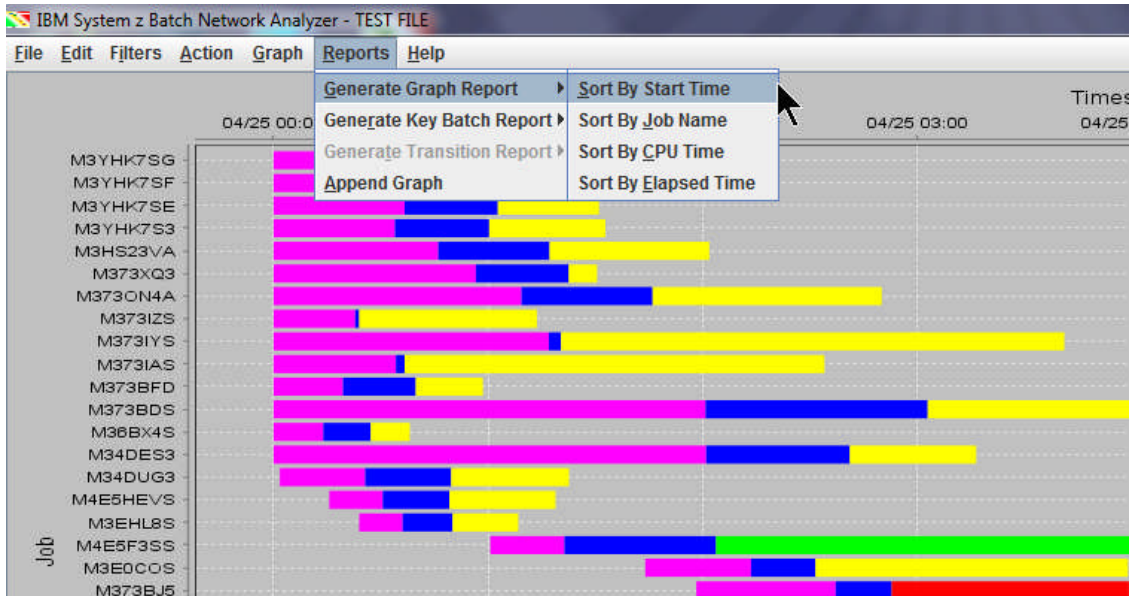
This will prompt you to save the report as an HTML file. Key in a file name, e.g. "**TEST_Report**.htm" and click **_Save_**.

8. Open the file (**TEST_Report.htm**) in an internet browser. After a legal
disclaimer, the report will show the filters that were used and the resulting table.
Key batch jobs are in bold. There is one line for each job, and at the very bottom
there is a "Total" job line that is the sum of the resources used for all the Filtered
jobs.

Filters

Type	Filter
Top Percent	Greater than 10%
CPU Time	Greater than 10.0
Service Class	Must be BATCHHI, BATPRDDF or BATTSTDF
Job Names	Must match M4* or M3*
Exclude	Excluded from analysis: M373DVF(JOB27670)

Data

There are 36 jobs in the following table.

Line	Key	Job Name	Program Name	Start	End	Steps	Job Class	Acct Code	Serv Class	Elapsed Time	CPU Time	Top Program	Top Pgm %
1		M3YHK7SG		4/25/13 12:00 AM	4/25/13 1:38 AM	26	J	3YH3YH32	BATPRDDF	5,908	2,327	DSNECP10	62
2		M3YHK7SF		4/25/13 12:00 AM	4/25/13 1:22 AM	26	J	3YH3YH32	BATPRDDF	4,950	1,988	DSNECP10	63
3		M3YHK7SE		4/25/13 12:00 AM	4/25/13 1:30 AM	26	J	3YH3YH32	BATPRDDF	5,458	2,198	DSNECP10	64
4		M3YHK7S3		4/25/13 12:00 AM	4/25/13 1:32 AM	26	J	3YH3YH32	BATPRDDF	5,565	2,035	DSNECP10	62
5		M3HS23VA	LNMHIW3D	4/25/13 12:00 AM	4/25/13 2:01 AM	3	J	3HS3HS32	BATPRDDF	7,308	2,763	DSNECP10	49
36	X	M4E07B1H		4/25/13 7:24 AM	4/25/13 7:41 AM	132	B	4E595732	BATCHHI	991	72	IEFIIC	0
		Total								239,325	72,959		

Note: A report may also be generated solely for key batch jobs by selecting
Generate Key Batch Report on the **Reports** menu. The following is included in
the report.

Key Batch Jobs

These are jobs that the user specifically selected for display.

There are 4 jobs in the following table.

Line	Key	Job Name	Program Name	Start	End	Steps	Job Class	Acct Code	Serv Class	Elapsed Time	CPU Time	Top Program	Top Pgm %
1	X	M373BJ5		4/25/13 1:58 AM	4/25/13 3:59 AM	11	J	37397332	BATPRDDF	7,255	2,339	SYNCSORT	9
2	X	M3E0IKSN		4/25/13 3:39 AM	4/25/13 4:56 AM	4	J	3E09E032	BATPRDDF	4,601	1,218	DSNECP10	8
3	X	M402GX3L		4/25/13 4:54 AM	4/25/13 5:49 AM	17	J	40242032	BATPRDDF	3,252	1,674	ENGEXE	4
4	X	M4E07B1H		4/25/13 7:24 AM	4/25/13 7:41 AM	132	B	4E595732	BATCHHI	991	72	IEFIIC	0
		Total								16,099	5,301		

9. When the graph report is initially generated, the graph is not present. *To include the graph in the report, click **Reports**, **Append Graph**.* You will be prompted to select the previously saved report file. Then click **Save**, and the graph will be appended to the report.
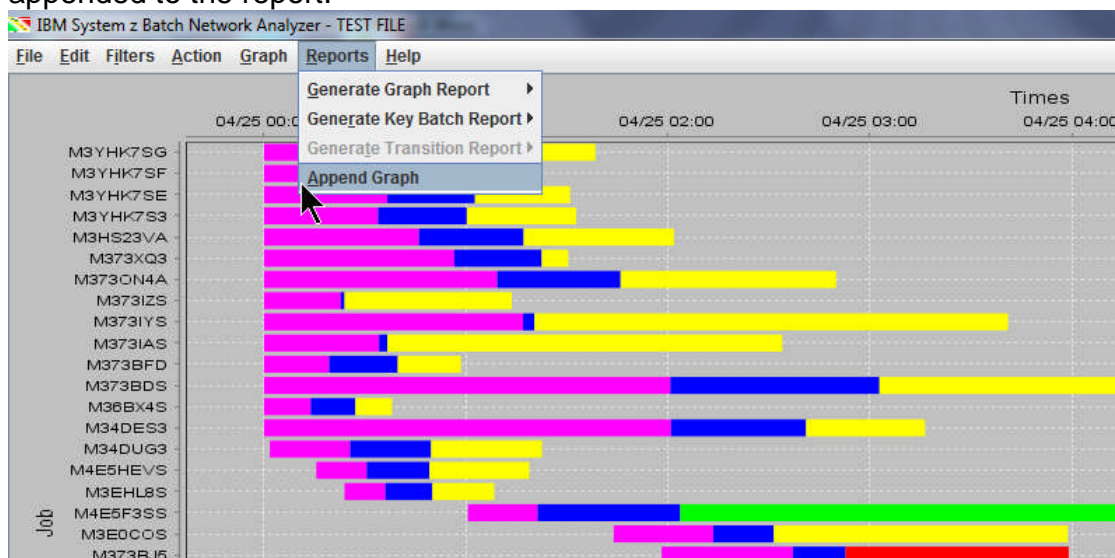
Click **Back to Main** to return to the zBNA main panel.

10. Reload (or Refresh via F5) the report, which will now include the graph positioned below the job table.

33	X	M3E0IKSN	4/25/13 3:39 AM	4/25/13 4:56 AM	4	J	3E09E032	BATPRDDF	4,601	1,218	DSNECP10	8
34	X	M402GX3L	4/25/13 4:54 AM	4/25/13 5:49 AM	17	J	40242032	BATPRDDF	3,252	1,674	ENGEXE	4
35		M373IUS	4/25/13 5:32 AM	4/25/13 6:28 AM	14	J	37397332	BATCHHI	3,315	1,296	DSNECP10	24
36	X	M4E07B1H	4/25/13 7:24 AM	4/25/13 7:41 AM	132	B	4E595732	BATCHHI	991	72	IEFIIC	0
		Total							239,325	72,959		
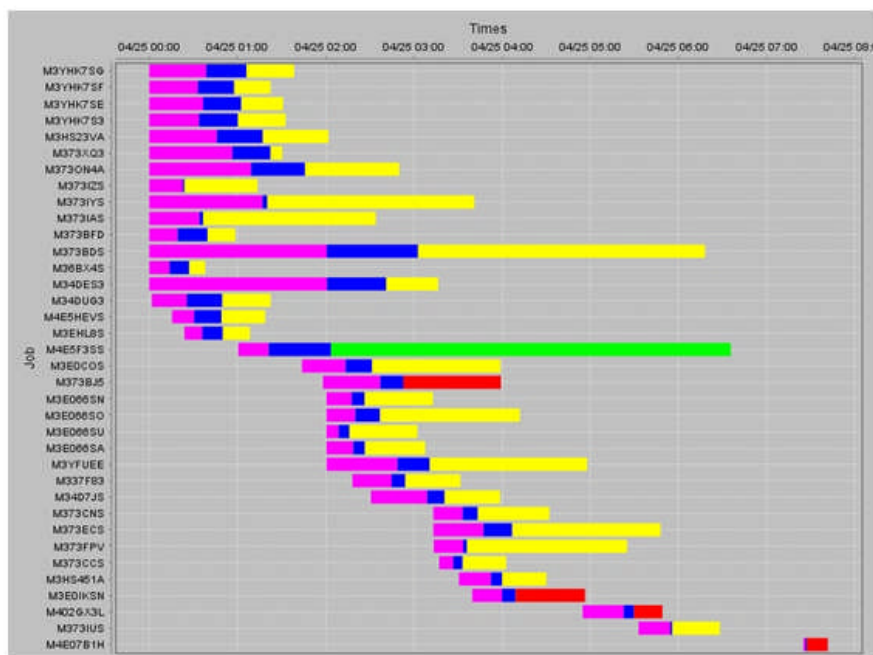
(magenta)	Execution
(blue)	Wait for Execution
(green)	Other Time (< Threshold)
(yellow)	Other Time (>= Threshold)
(red)	Other Time (Key Job)
(hatched)	indicate Alternates

zBNA Lab Guide

Task 4 – __Reviewing DASD Data Set Information__

1. The SMF Type 42 subtype 6 records are required to view the DASD data set I/O information. They are loaded into zBNA via the .dat file when the job step level data is added.

A job must be selected to display the data set information. Let's focus on one of the jobs identified as key batch. Double click on the **Key Batch** header to sort that column. Right click on the job, **M4E07B1H**, and select **Job Dataset Report** (this option is also available on the **Action** menu). zBNA reads the SMF Type 42 subtype 6 data.
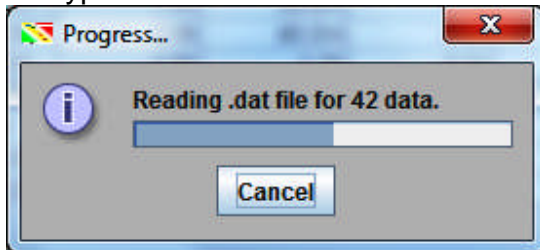
2. The zBNA **Job Dataset Report** panel displays the data sets for job **M4E07B1H**.

Step	Step Number	DSN	Total IOTime	IO Count	Response Time	Queue Time	Pending Time	Connect Time	Disconr Time
S4E5N227	92	I4E5SEY.M4E57B1S.SOQDVSG.LQGHA	188.0s	1879622	0.1	0.0	0.0	0.0	
S4E5H22E	76	I4E5SE.M4E57B1S.PHD.HAWUDFW.J2439Y22	42.1s	619	68.0	0.0	0.1	34.6	
S4E0T8A4	66	Y325.L576.WPV	25.0s	249682	0.1	0.0	0.0	0.0	
S4E03FQG	44	I4E0SEY.M4E07B1S.HAW2KLS.GDWD	22.5s	7746	2.9	0.0	0.0	2.8	
S4E5N27G	91	I4E5SE.VRUWILOH.M4E57B1S.J2421Y22	19.8s	738	26.8	0.0	0.0	20.7	
S4E5H22E	76	I4E5SE.SE5H2233.M4E57B1S	19.5s	698	28.0	0.0	0.0	21.5	
S4E03FQ7	36	VBV35337.W294677.UD222.M4E07B1H.U2910380	15.7s	83	189.0	0.0	1.4	159.3	
S4E5N26F	82	I4E5SE.SE5N226F.M4E57B1S	15.6s	10401	1.5	0.0	0.0	1.4	
S4E5N24E	75	I4E5SE.SE5N2233.M4E57B1S	13.2s	145	90.7	0.0	0.0	84.2	
S4E5N27E	89	I4E5SE.HAWUDFW.M4E57B1S.ILOH	12.8s	3276	3.9	0.0	0.0	2.5	
S4E5N227	92	I4E5SE.VRUWILOH.M4E57B1S.J2421Y22	8.4s	5249	1.6	0.0	0.0	1.5	
S4E03FQJ	47	I4E0SEY.M4E07B1S.HAW2KLS.LQGHA	8.4s	83547	0.1	0.0	0.0	0.0	

Job Details:
Job Name: M4E07B1H   Key Batch: Yes   Elapsed Time: 991.79 Seconds   CPU Intensity: 7.2%
Start Date: Apr 25, 2013   Start Time: 7:24 AM   End Date: Apr 25, 2013   End Time: 7:41 AM

Be sure to use the scroll bars to get a complete view of the job details. Sort the **Total IO Time** column in descending order so that the data set with the most IO time is positioned in the first row.

3. Right click on **I4E5SEY.M4E57B1S.SOQDVSG.LQGHA**, and select **Get the Life of this Dataset**.

Step	Step Number	DSN	Total IOTime	IO Count	Response Time	Queue Time	Pending Time	Connect Time	Disconr Time
S4E5N227	92	I4E5SEY.M4E57B1S.SOQDVSG.LQGHA	188.0s	1879622	0.1	0.0	0.0	0.0	
S4E5H22E	76	I4E5SE.M4E57B1S.PHD.HAWUDFW.J24 [Get the Life of this Dataset]		619	68.0	0.0	0.1	34.6	
S4E0T8A4	66	Y325.L576.WPV	25.0s	249682	0.1	0.0	0.0	0.0	
S4E03FQG	44	I4E0SEY.M4E07B1S.HAW2KLS.GDWD	22.5s	7746	2.9	0.0	0.0	2.8	
S4E5N27G	91	I4E5SE.VRUWILOH.M4E57B1S.J2421Y22	19.8s	738	26.8	0.0	0.0	20.7	
S4E5H22E	76	I4E5SE.SE5H2233.M4E57B1S	19.5s	698	28.0	0.0	0.0	21.5	
S4E03FQ7	36	VBV35337.W294677.UD222.M4E07B1H.U2910380	15.7s	83	189.0	0.0	1.4	159.3	
S4E5N26F	82	I4E5SE.SE5N226F.M4E57B1S	15.6s	10401	1.5	0.0	0.0	1.4	
S4E5N24E	75	I4E5SE.SE5N2233.M4E57B1S	13.2s	145	90.7	0.0	0.0	84.2	
S4E5N27E	89	I4E5SE.HAWUDFW.M4E57B1S.ILOH	12.8s	3276	3.9	0.0	0.0	2.5	
S4E5N227	92	I4E5SE.VRUWILOH.M4E57B1S.J2421Y22	8.4s	5249	1.6	0.0	0.0	1.5	
S4E03FQJ	47	I4E0SEY.M4E07B1S.HAW2KLS.LQGHA	8.4s	83547	0.1	0.0	0.0	0.0	

Job Details:
Job Name: M4E07B1H   Key Batch: Yes   Elapsed Time: 991.79 Seconds   CPU Intensity: 7.2%
Start Date: Apr 25, 2013   Start Time: 7:24 AM   End Date: Apr 25, 2013   End Time: 7:41 AM

4. zBNA reads the .dat file that is loaded for the SMF 42 then the SMF 30 data. It searches through **all Jobs (5147)**, not just the Filtered Jobs. When it finishes the process, the **zBNA: Life of a Dataset** panel is displayed.

The job names using this data set are shown. Use the scroll bar to view all of the data, and the columns can be sorted.

In this case, **Job M4E07B1H** is the only job that accessed the data set; in Steps 88 and 92. Step 92 has the most **Total IO Time**, 188 seconds. The response time is very low. If you scroll to the right, in the column **Type**, you'll see it is a "**KSDS Index**". While not currently provided in zBNA, one could investigate SMF 64s and consider increasing LSR / NSR buffers to hold Index Set and potentially eliminate ~3 Minutes of I/O time, which would be approximately 18% of the Job's elapsed time (16.5 minutes).

Click **OK** until the zBNA main panel is displayed.

zBNA Lab Guide

5. Click the **Action** menu then **Top 10 Dataset Report**.

zBNA displays an information panel showing that it is reading the SMF 42 (6) then SMF 30 data from the loaded .dat file. The **zBNA: Top 10 Data Sets** panel is displayed.
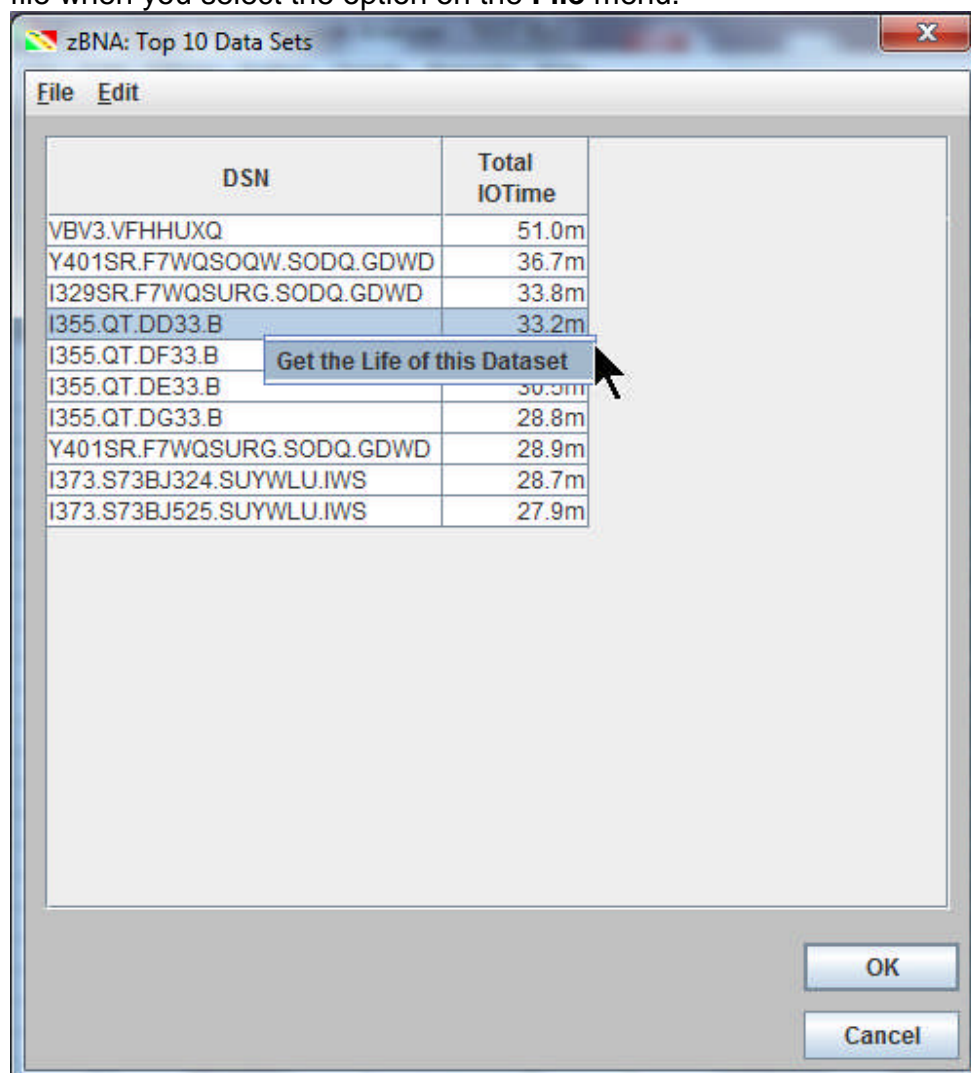
zBNA Lab Guide

The purpose is to show where the most I/O time is, over the entire interval and regardless of who is accessing the dataset. Then looking at the characteristics, technology options can be evaluated to improve the response time, and thus the elapsed times of the jobs/online applications that are accessing it. In this case, it appears that 4 data sets starting with **I335.QT.** are the 4th through 7th Top data sets. Perhaps they are clones that we enabled for parallel processing? We'll investigate one of these files.

6. The Top 10 data sets are displayed, and the information can be written to a CSV file when you select the option on the **File** menu.

Right click on the **I355.QT.DD33.B** data set then **Get the Life of this Dataset**. After zBNA reads the SMF 42 and 30 data in the .dat file, the **zBNA: Life of a Dataset** panel is displayed.

zBNA Lab Guide

7. The job details are shown for the **I355.QT.DD33.B** data set. You can see that multiple different Jobs access this data set throughout the Batch interval.

Now we want to see which Jobs have the most IO Time. Perform a sort on the **Total IO Time** column in descending order.

We can see that many of the Jobs have Response times in the 2 - 6 MS range. Based on this, perhaps an investigation of I/O technology to reduce I/O response times should be a follow-on action.
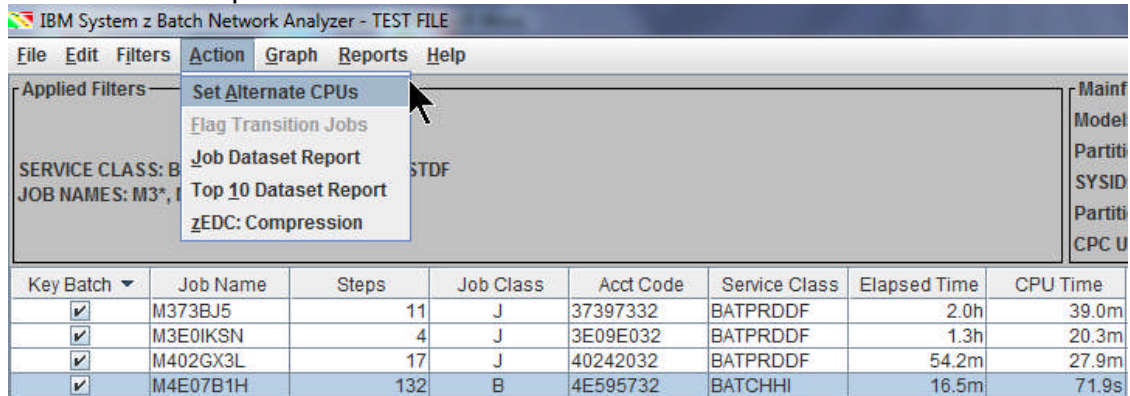
Click **OK** until the zBNA main panel is displayed.

zBNA Lab Guide
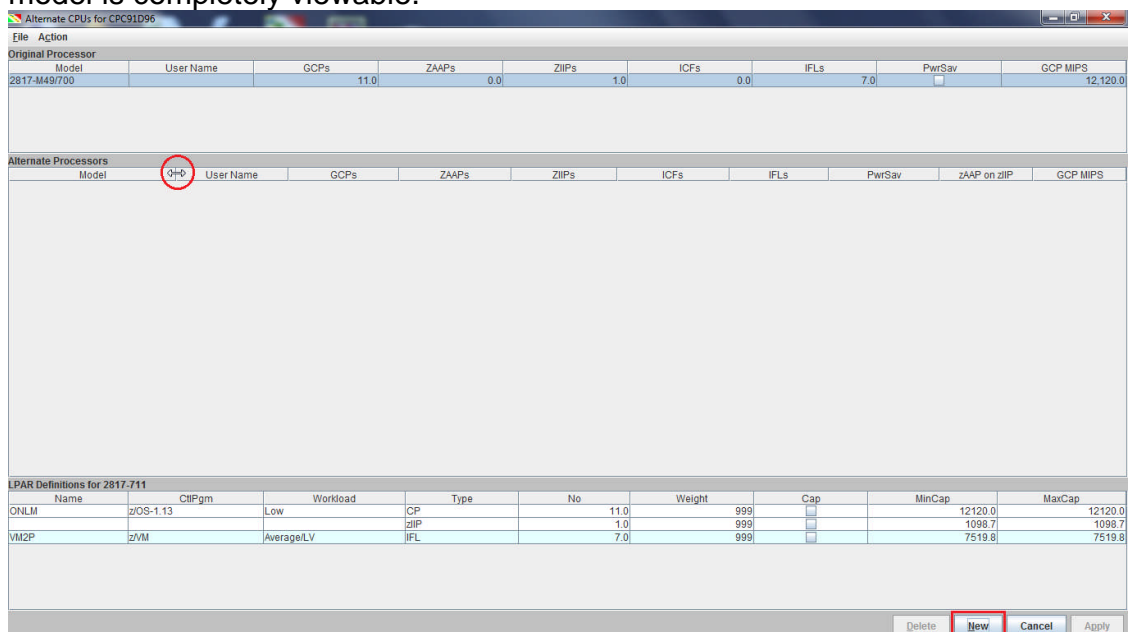
Task 5 – __Performing an Alternate Processor Analysis__

1. Now we will view a "what-if" scenario by selecting an alternate processor to "execute" the same batch jobs. **Click *Action, Set Alternate CPUs*** to load the Alternate CPUs panel.

2. Maximize the **Alternate CPUs** window to show all of the columns. Then expand the **Model** column in the **Alternate Processors** table so that the name of each model is completely viewable.

Click ***New***.

3. A drop-down menu will appear that allows you to select the new processor. In this example we are going to select a processor with less total capacity and also less capacity per engine. Select the **zEC12/600** family, and then the **2827-607**.

   In this example, we are selecting a zEC12 607 subcapacity model versus the current z196 711 full capacity model. (Perhaps they have a zEC12 607 and are considering migrating these jobs to that processor, and they want to understand the impact to elapsed time changes versus their required Batch completion time.

Click **Apply** to view the hypothetical scenario with this new processor.

4. From the main zBNA panel, use **Graph, Display Graph: _Table_** to display the graph to see that each row now contains a gray striped one below it. This second row shows the same jobs, however, the total times are estimated as if the jobs were run on the alternative new processor.
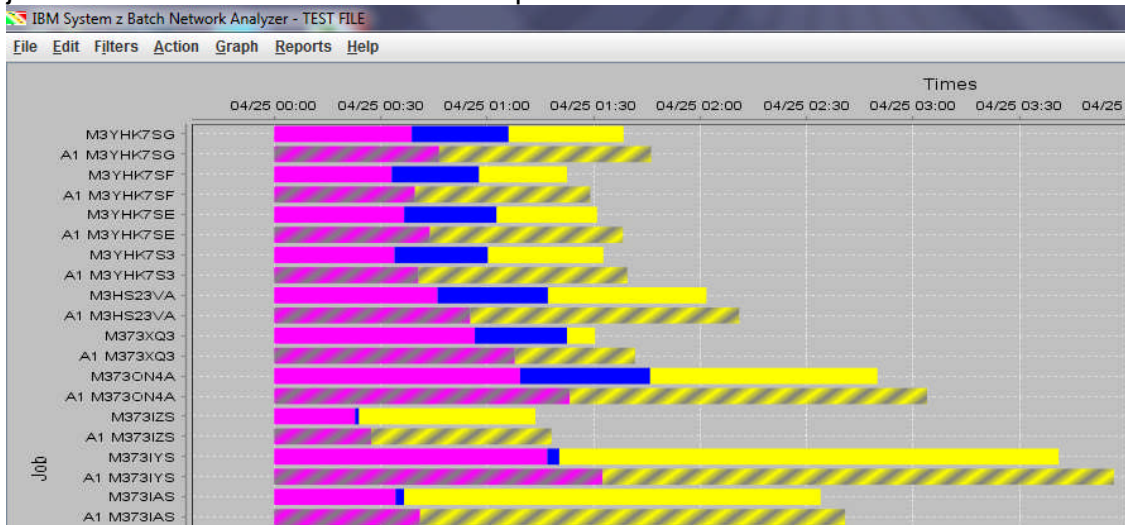
5. We can generate a new report that includes the alternate processor details or we can append to the one previously saved in **Task 3**. To do this, click on the **Reports** menu, then select **Generate Graph Report**, **Sort By Start Time**. This will prompt you to save the report as an HTML file. Since we have already created a report, you can select that file. zBNA will display the following message.
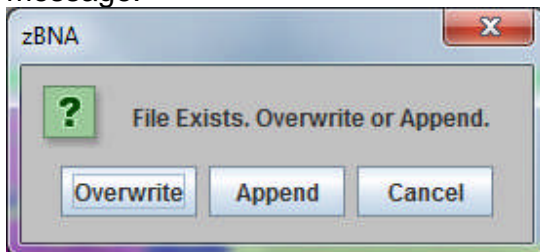
Click **Append** to add to the end of the previously generated file.
Note: Click **Overwrite** to replace the file. Click **Cancel** to return to the graph.

6. The report will now include the alternate processor, as well as the estimated run-time in the table for this new processor.

The processors considered in this analysis are the following:

Note: There is no effort to determine if the alternate processor has the total capacity to run this workload. The analysis is simply comparing the single engine speed of base versus the alternate processor.

Name	Processor	Single GCP Mips	Ratio
Base (B)	2817-711	1,102	
Alternate 1 (A1)	2827-607	914	-17.0%

The analysis follows:

Data

There are 36 jobs in the following table.

Name	Line	Key	Job Name	Program Name	Start	End	Steps	Job Class	Acct Code	Serv Class	Elapsed Time	CPU Time	Top Program	Top Pgm %
B	33	X	M3E0IKSN		4/25/13 3:39 AM	4/25/13 4:56 AM	4	J	3E09E032	BATPRDDF	4,601	1,218	DSNECP10	8
A1	33	X	M3E0IKSN		4/25/13 3:39 AM	4/25/13 5:00 AM	4	J	3E09E032	BATPRDDF	4,851(5.4%)	1,468		
B	34	X	M402GX3L		4/25/13 4:54 AM	4/25/13 5:49 AM	17	J	40242032	BATPRDDF	3,252	1,674	ENGEXE	4
A1	34	X	M402GX3L		4/25/13 4:54 AM	4/25/13 5:54 AM	17	J	40242032	BATPRDDF	3,595(10.5%)	2,018		
B	35		M373IUS		4/25/13 5:32 AM	4/25/13 6:28 AM	14	J	37397332	BATCHHI	3,315	1,296	DSNECP10	24
A1	35		M373IUS		4/25/13 5:32 AM	4/25/13 6:32 AM	14	J	37397332	BATCHHI	3,581(8.0%)	1,561		
B	36	X	M4E07B1H		4/25/13 7:24 AM	4/25/13 7:41 AM	132	B	4E595732	BATCHHI	991	72	IEFIIC	0
A1	36	X	M4E07B1H		4/25/13 7:24 AM	4/25/13 7:41 AM	132	B	4E595732	BATCHHI	1,006(1.5%)	87		
B			Total								239,325	72,959		
A01			Total								254,106(6.2%)	87,740		

In this case we can see that the Alternate Processor had a **Ratio** of **-17%** **Single GCP MIPS**, resulting in slightly increased CPU and Elapsed times compared to the current processor for each job.

7. Let's save the study as a zBNA file, click *File*, *Save As zBNA Study File*. This saves a .zBNA file containing the current filters and settings including the key batch jobs. However, when you load the .zBNA file, the original SMF70 and SMF30 files will still be needed.
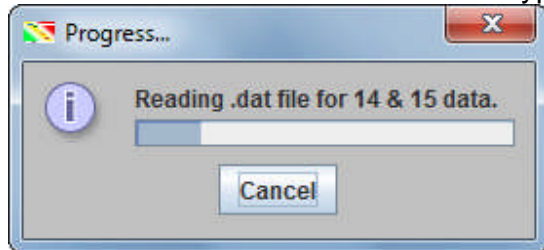
zBNA Lab Guide

Task 6 – <u>Exploring zEDC Compression</u>

1. To use the zBNA zEDC Compression function, SMF Type 14 and 15 (Input/Output Data Set Close) Records must be included in the "**.dat**" file. Click *Action*, *zEDC: Compression* on the main zBNA menu.

zBNA reads the data from the SMF Type 14 and 15 records.

2. The **zEDC Top Data Sets** panel displays after the SMF Type 14 and 15 records have been loaded.

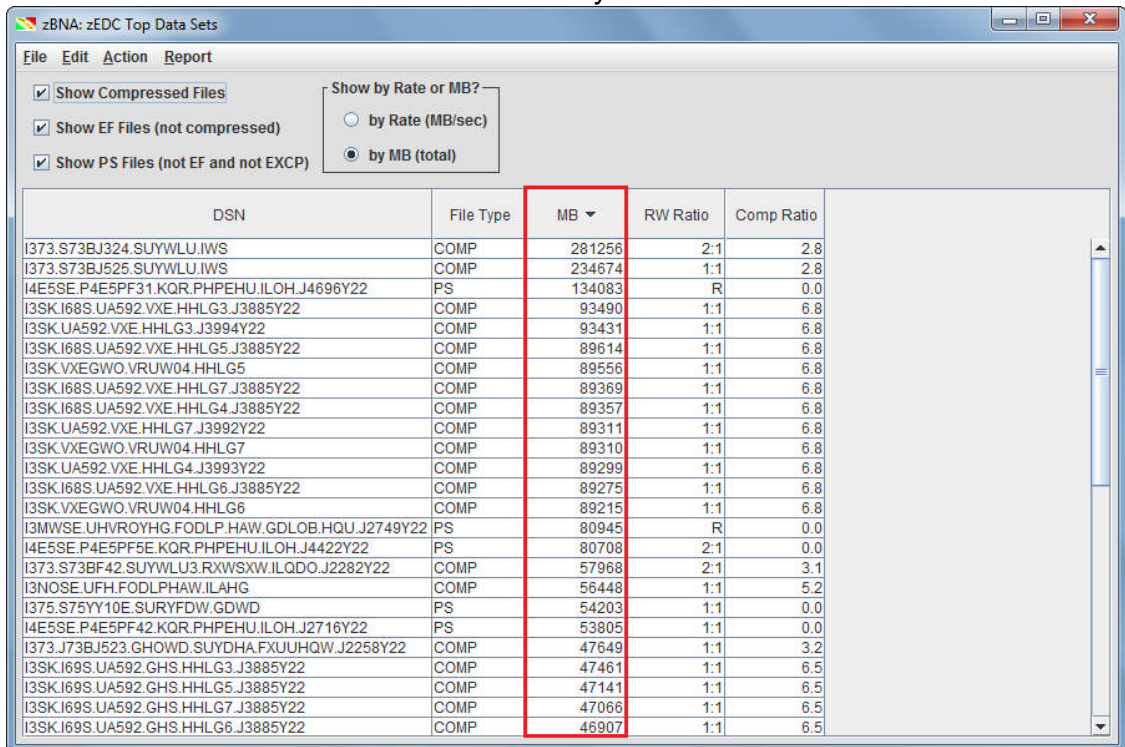| DSN	File Type	MB	RW Ratio	Comp Ratio
I373.S73BJ324.SUYWLU.IWS	COMP	281256	2:1	2.8
I373.S73BJ525.SUYWLU.IWS	COMP	234674	1:1	2.8
I3SK.I68S.UA592.VXE.HHLG3.J3885Y22	COMP	93490	1:1	6.8
I3SK.UA592.VXE.HHLG3.J3994Y22	COMP	93431	1:1	6.8
I3SK.I68S.UA592.VXE.HHLG5.J3885Y22	COMP	89614	1:1	6.8
I3SK.VXEGWO.VRUW04.HHLG5	COMP	89556	1:1	6.8
I3SK.I68S.UA592.VXE.HHLG7.J3885Y22	COMP	89369	1:1	6.8
I3SK.I68S.UA592.VXE.HHLG4.J3885Y22	COMP	89357	1:1	6.8
I3SK.UA592.VXE.HHLG7.J3992Y22	COMP	89311	1:1	6.8
I3SK.VXEGWO.VRUW04.HHLG7	COMP	89310	1:1	6.8
I3SK.UA592.VXE.HHLG4.J3993Y22	COMP	89299	1:1	6.8
I3SK.I68S.UA592.VXE.HHLG6.J3885Y22	COMP	89275	1:1	6.8
I3SK.VXEGWO.VRUW04.HHLG6	COMP	89215	1:1	6.8
I373.S73BF42.SUYWLU3.RXWSXW.ILQDO.J2282Y22	COMP	57968	2:1	3.1
I3NOSE.UFH.FODLPHAW.ILAHG	COMP	56448	1:1	5.2
I373.J73BJ523.GHOWD.SUYDHA.FXUUHQW.J2258Y22	COMP	47649	1:1	3.2
I3SK.I69S.UA592.GHS.HHLG3.J3885Y22	COMP	47461	1:1	6.5
I3SK.I69S.UA592.GHS.HHLG5.J3885Y22	COMP	47141	1:1	6.5
I3SK.I69S.UA592.GHS.HHLG7.J3885Y22	COMP	47066	1:1	6.5
I3SK.I69S.UA592.GHS.HHLG6.J3885Y22	COMP	46907	1:1	6.5
I3SK.I69S.UA592.GHS.HHLG4.J3885Y22	COMP	46844	1:1	6.5
I3SK.I05S.VXE.IXOOILOH.HHLG3.J3891Y22	COMP	46678	R	6.8
I3SK.I05S.VXE.IXOOILOH.HHLG5.J3891Y22	COMP	44740	R	6.8
I3SK.I05S.VXE.IXOOILOH.HHLG7.J3891Y22	COMP	44618	R	6.8
I3SK.I05S.VXE.IXOOILOH.HHLG4.J3891Y22	COMP	44612	R	6.8

These are the data sets that zBNA has calculated are the top zEDC Compression candidates. **Note**: By default, the list is ordered by the top data sets, according to MB, in each of the compression categories.

The purpose of providing the Top Data Sets is to identify which ones will provide the most impact/benefit from zEDC compression, and may provide a starting point for which ones to implement first. Also, you can drill down further on a data set by right clicking on its name and selecting **Get the Life of this Dataset**.
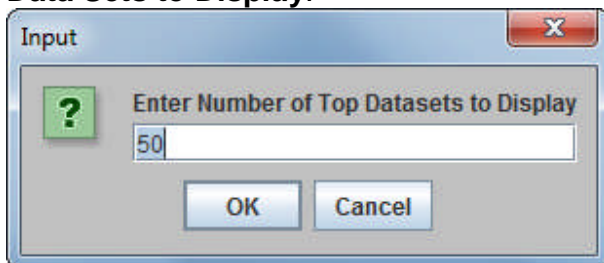
3. Double click on **MB** to sort the list of files by this characteristic.

Now the top data sets are listed by MB no matter which compression category they are in.

4. Let's change the number of top data sets to display. Use **Action**, **Number of Data Sets to Display**.

Change the default value of 50 to **10**. Click **OK** to redisplay the **zEDC top Data Sets** panel.

5. To show the zEDC graph, click **Action**, **Show Graph**.
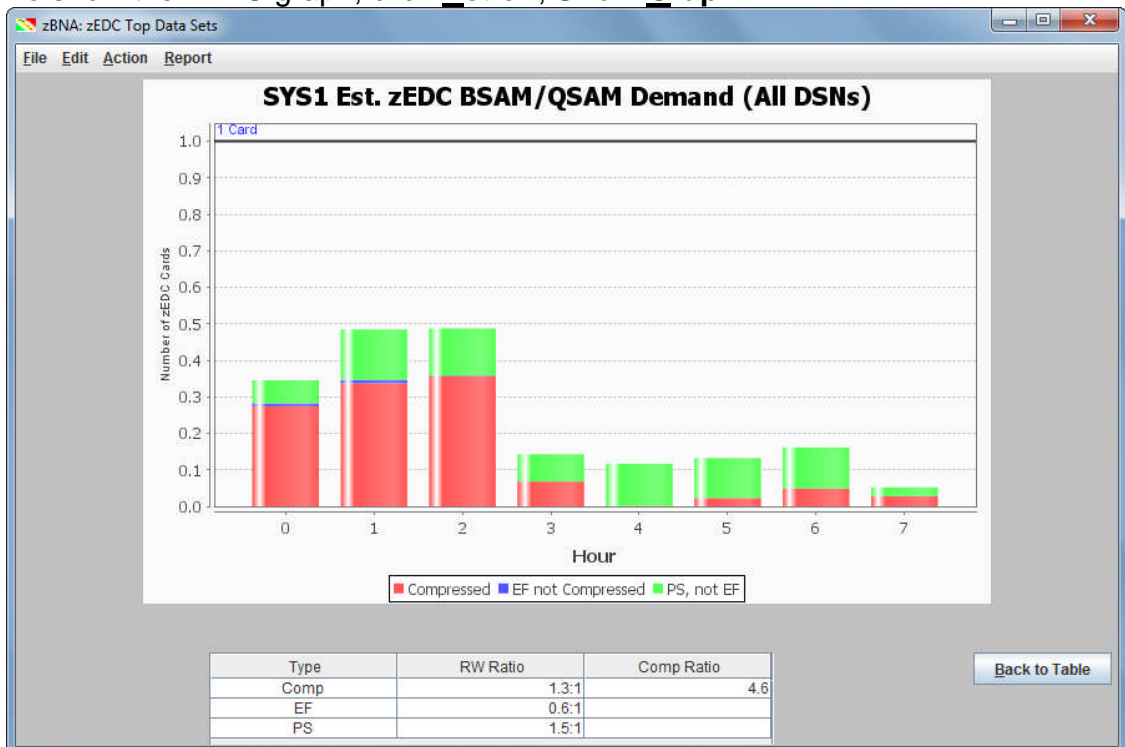
This graph shows the estimated number of zEDC cards by hour needed to support the workload for **all data sets** that met the criteria in the interval. With this graph you can see the peak time and how many cards are required from a capacity perspective. Save this data and graphic image to a zBNA report file by clicking **Report**, **Output Graph**. Input "**zEDCgraph**" for the file name, and click **Save**. Both the ".htm" and ".jpg" files are generated.

You have successfully completed all the tasks in running the zBNA Lab.