Advanced Technical Support
At the Washington Systems Center

Using zCP3000 for Performance Analysis

Ray Wicks
RayWicks@us.ibm.com

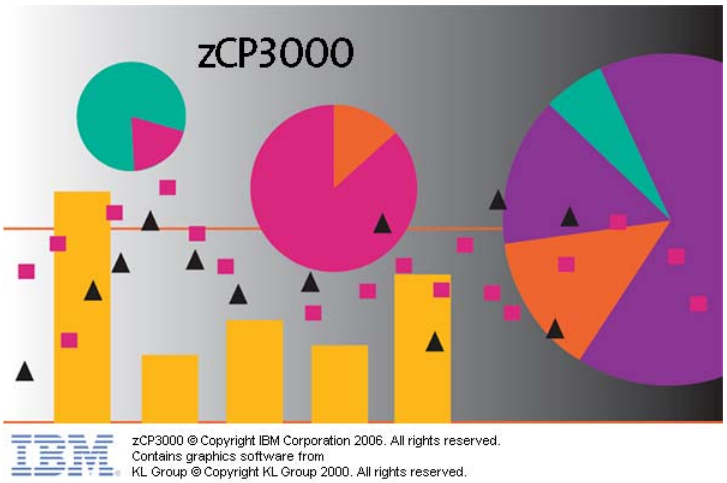

zCP3000 is a tool for performance analysis and capacity planning on zSeries. The purpose of the performance analysis phase of the process is to determine the state of the configuration prior to capacity planning. Some of the questions that need to be answered prior to capacity planning are as follows.

- Is the hardware configuration correctly specified?
- Where, if any, are the current bottle necks? Clearly if there's a DASD bottleneck, upgrading the processor may not provide the expected capacity improvement.
- Is there a significant latent demand? One expects that the workload moved to a new processor to immediately remain the same. Latent demand creates a burst of new work when the capacity of the environment improves. It should not be unexpected.
- Does the modeled sample interval represent the business environment? Should there be more than one model interval selected?

Overall Process

- Load EDF Files
- Review Intervals
- Review Physical & Logical Configuration
- Analysis
    - Enterprise - Processor
        - CEC (CPC)
        - Partition (SYSid)
        - Workload (application, service class, or collection of service classes)
    - Enterprise – BCU (Basic Configurable Unit)
        - Controller
        - Paths
        - DASD
        - Data Set
    - Enterprise Sysplex

Loading Files

Data is loaded from data files designed for zCP3000. The files come from z/OS and z/VM. Both CP2KEXTR and CP2KVMXT (the output from these programs were also for CP2000) produce .EDF files: One file per system image (z/OS) or collection of Virtual Machines (z/VM).

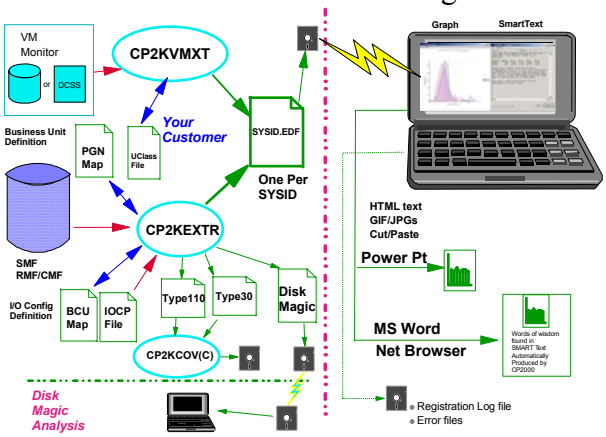

CP2KEXTR uses an extensive set of SMF record and produces other files for use in other types of analyses. This is explained in the documentation for CP2KEXTR.

Since multiple files (data from multiple partitions) can be processed in zCP3000, there are some restrictions and choices when these files are loaded.

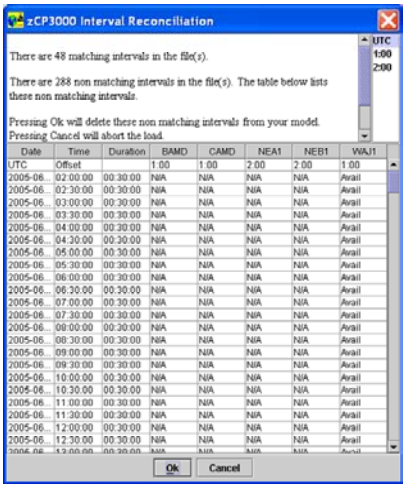

First off, zCP3000 checks that the Dates, times, and durations from all the .EDF files match. If zCP3000 finds intervals in any .EDF file that are not in all the .EDF files, the window at the left will appear identifying the intervals that zCP3000 intends to delete. If you press OK, these intervals will be deleted from the model.

One oddity is found in the matching of dates and times. The Dates and times in the .EDF file are the local times for the partition. We know of course that one partition could be supporting Paris (UTC +1) and New York (UTC -5 where UTC is Coordinated Universal Time or Zulu time with Greenwich, U.K. as UTC 0.) You will notice that zCP3000 has identified two different UTC values in the data. The user has to set the preferred *local time* for the analysis.

If the window does not appear, everything matches.

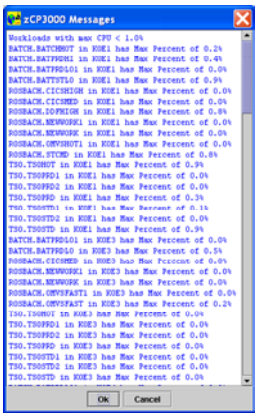

Quite often the number of workloads in the EDF files far exceeds the number of active workloads. The installation may have one Work Load Manager (WLM) specification for the entire installation. This means that there may be a number of workloads that are 0 resource consumers for the sample set. There may also be a number of workloads which could be insignificant consumers.

The insignificant workloads are defined as those whose maximum CPU consumption for any interval was less than 1%. Press OK and these will be deleted from the model.

Why delete workloads? Foremost, this keeps the workload graphs simpler. Nothing messes up a graph like a lot of workloads in the

legend and nothing in the graph.
When might you retain a small workload? It might me a *significant* workload to an important person or the small CPU workload may have a non trivial storage usage.
If you decide to keep all the workloads, you can go through the arduous task of deleting them afterwards. Not usually done.

Interval Selection

When the EDF files are finally loaded, zCP3000 chooses the 90th percentile as the default interval (if there are enough intervals). This is the 90th percentile of the CP data from all the CECs. You can reset the interval to any one you want. You can also influence the choice by checking zAAP, ICF, or IFL data to include any combination of processor time.

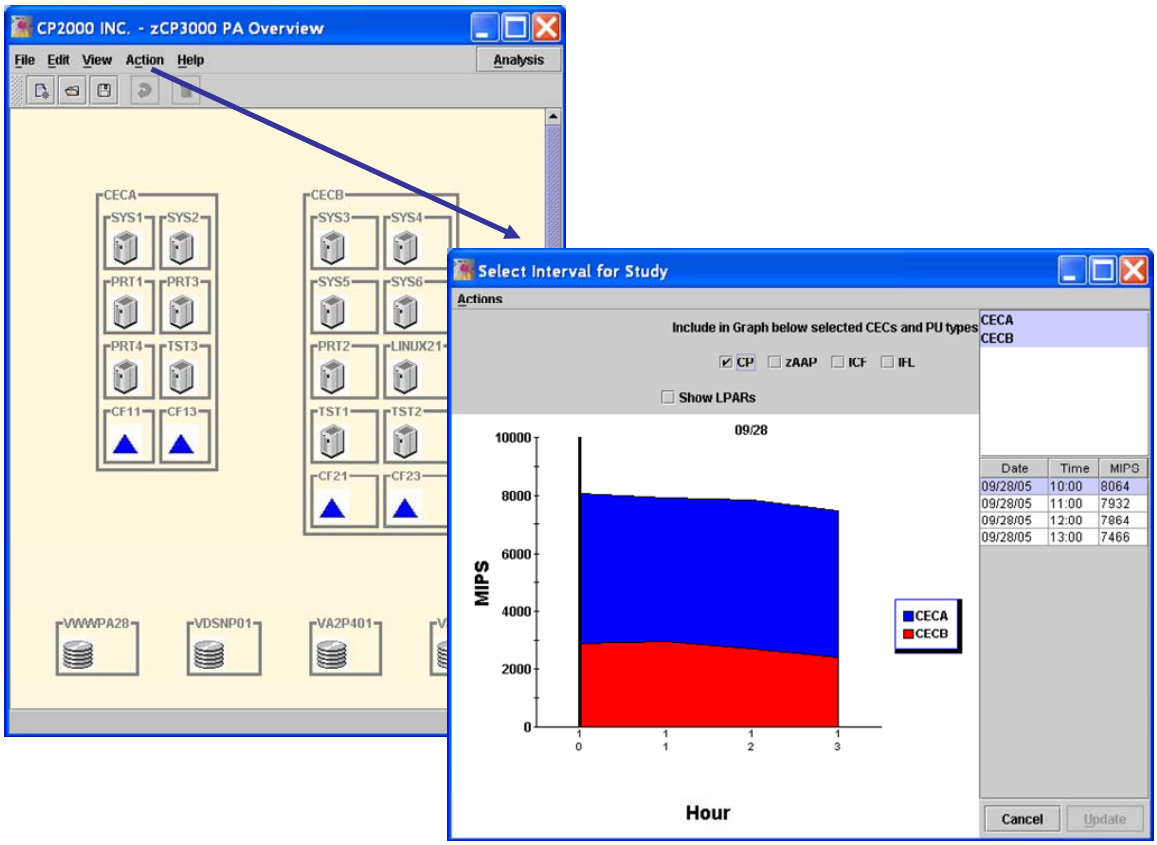

The interval chosen will be used globally to describe the resource demand behavior for all the CECs, system images (SYSIDs), and workloads. Hence it is important to review the data carefully to insure the right workloads are active and is in the proper proportion for the selected interval. If you must model separate intervals for separate CECs, you will have to build a separate model for each CEC.
A specific interval is the method of choice in capacity planning. You want to see a large enough number of intervals to view the workload behavior. But you also want a single interval to model. The use of an average across many intervals smoothes out the resource interdependence too much.

Configuration Specification

The input to zCP3000 is ordinarily EDF files. The EDF files come from our CP2KEXTR program for z/OS and CP2KVMXT for z/VM. Each file contains detailed information on the CEC and one partition's workload. When you load the EDF files into zCP3000 (using the File load), the PA Overview window shows each CEC partition structure for the

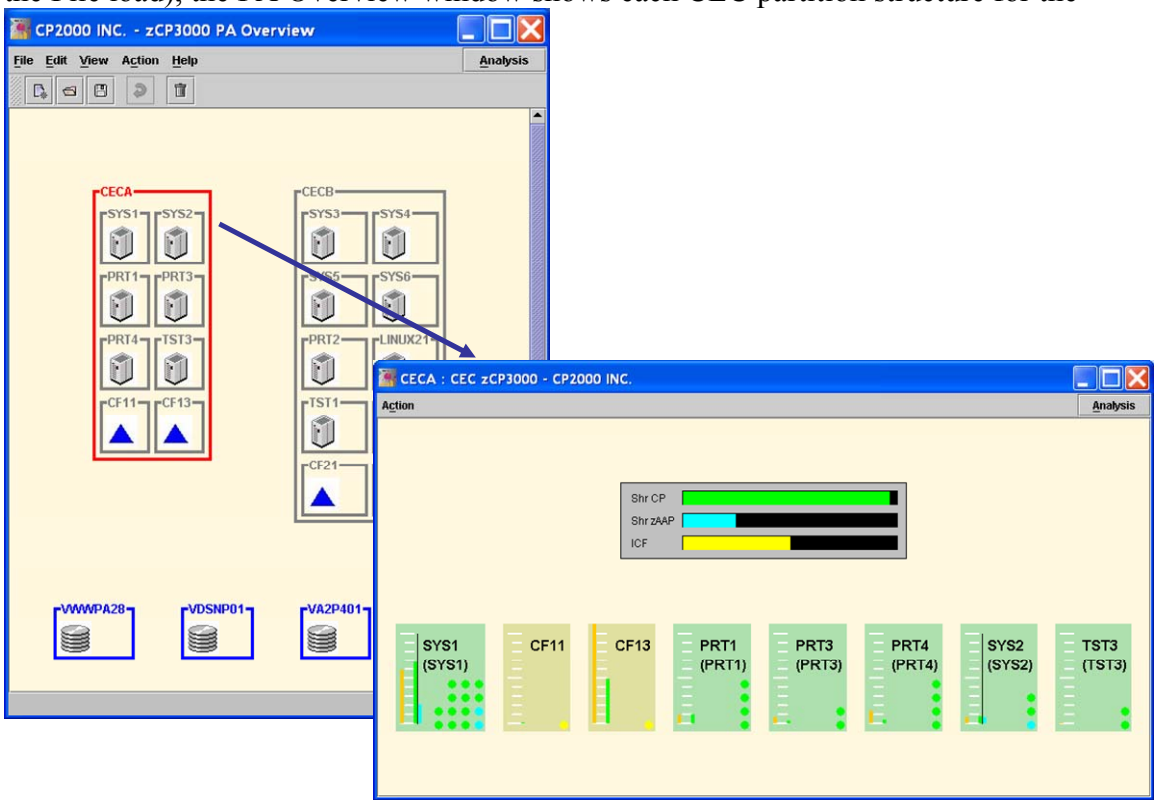

partition information in the EDF files along with the DASD configuration. The partition definition data in the EDF file is much more extensive than simply the one partition's information. If you double click on a CEC, the current configuration information is displayed as seen on the right above.

In many cases, there are more partitions in the CEC than EDFs in the input. For example, there's no EDF from an internal coupling facility. However, there is information in the EDF about the logical configuration and utilization. There may even be other z/OS partitions for which no EDF was provided. These will automatically be generated by zCP3000. On the CEC window you can visually review the configuration for the selected interval. Return to the Overview window to examine the logical configuration.

Right click on any CEC and select the Definition item from the pop up.

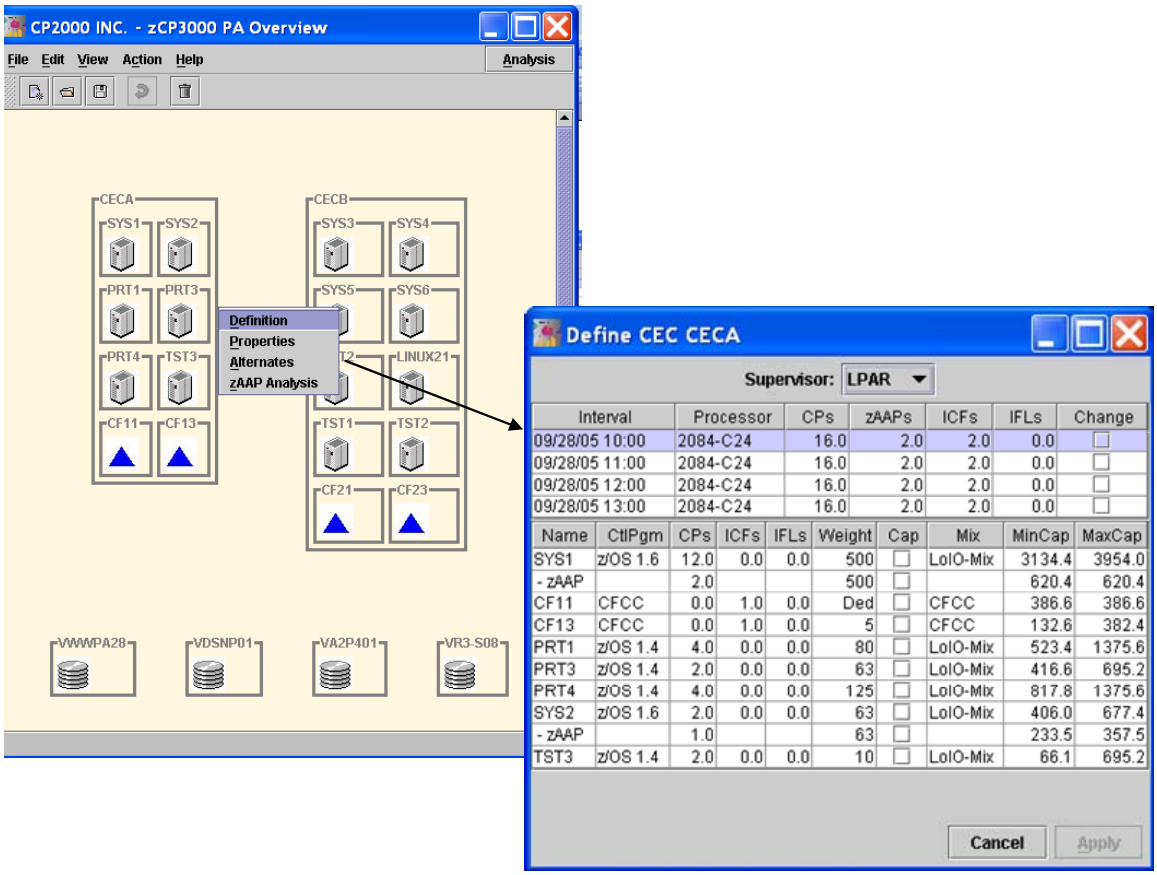

What appears is a complete description of the logical configuration. (The logical configuration can be interval dependent. IRD and WLM can change the logical configuration in an attempt to meet the objectives specified.) On the top is the physical processing unit configuration by PU type. zCP3000 makes an attempt to recognize the different PU types configured in each partition. However, often the PU type cannot be distinguished. For example, an IFL PU often appears as an ICF. You can edit the configuration on this window. You can change the physical configuration (on top) and logical configuration, control program and workload mix on the bottom. This information comes from the data in the EDF file.

Enterprise Processor Analysis

On the PA Overview window, press analysis. A list of analysis will be displayed. Each analysis at any level has both a graphic and text. The text does not merely describe the graph; it usually provides intelligent commentary about the specific data shown in the graph using the graph data to illustrate the meaning. We call this SmartText.

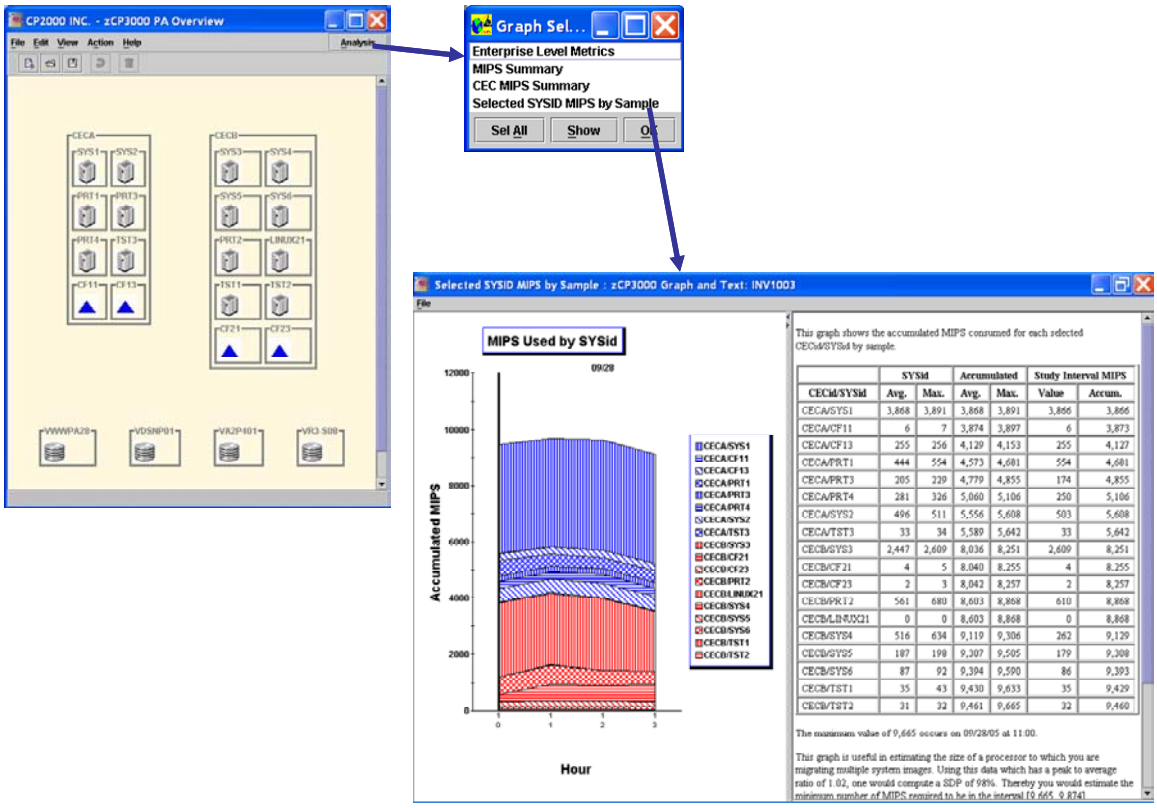

CECid/SYSid	SYSid Avg.	SYSid Max.	Accumulated Avg.	Accumulated Max.	Study Interval MIPS Value	Study Interval MIPS Accum.
CECA/SYS1	3,868	3,891	3,868	3,891	3,866	3,866
CECA/CF11	6	7	3,874	3,897	6	3,873
CECA/CF13	255	256	4,129	4,153	255	4,127
CECA/PRT1	444	554	4,573	4,681	554	4,681
CECA/PRT3	205	229	4,779	4,855	174	4,855
CECA/PRT4	281	326	5,060	5,106	250	5,106
CECA/SYS2	496	511	5,556	5,608	503	5,608
CECA/TST3	33	34	5,589	5,642	33	5,642
CECB/SYS3	2,447	2,609	8,036	8,251	2,609	8,251
CECB/CF21	4	5	8,040	8,255	4	8,255
CECB/CF23	2	3	8,042	8,257	2	8,257
CECB/PRT2	561	680	8,603	8,868	610	8,868
CECB/LINUX21	0	0	8,603	8,868	0	8,868
CECB/SYS4	516	634	9,119	9,306	262	9,129
CECB/SYS5	187	198	9,307	9,505	179	9,308
CECB/SYS6	87	92	9,394	9,590	86	9,393
CECB/TST1	35	43	9,430	9,633	35	9,429
CECB/TST2	31	32	9,461	9,665	32	9,460

In this sample, the analysis displays the MIPS consumed by partition over the samples found in the EDF file. The text can be very interesting. Often we forget why the graph is important… the SmartText reminds us. The graph and SmartText can be saved in a HTML document. The document will be build with a cover and preface. As you see analyses that appear interesting, these can be dynamically included in the document as you review each analysis. This document can later be formatted with MS Word or your favorite word processor as a particularly impressive customer deliverable.

To get to the CEC window, double click on a CEC and press analysis. Here also, as with most windows in zCP3000, there's an analysis button.

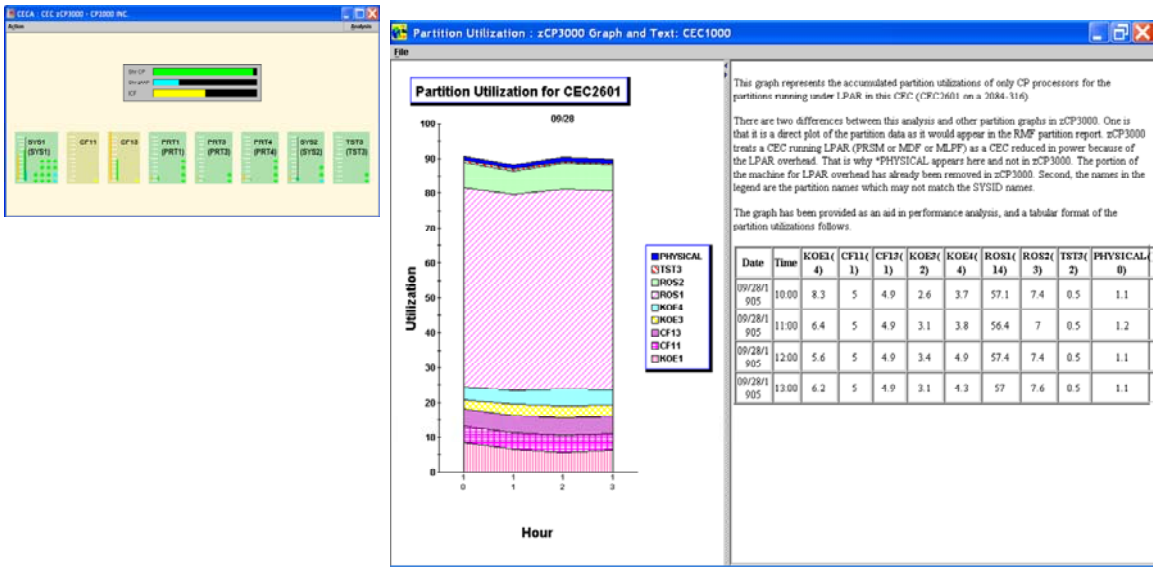

This is one such analysis. It shows the utilization for each partition defined in the CEC. As with the processor report in RMF, the physical utilization (part of the PR/SM overhead) is shown. You should step through all the analyses at each level to familiarize yourself with the contents of each.

If you double click on one of the partitions, the System information table appears.

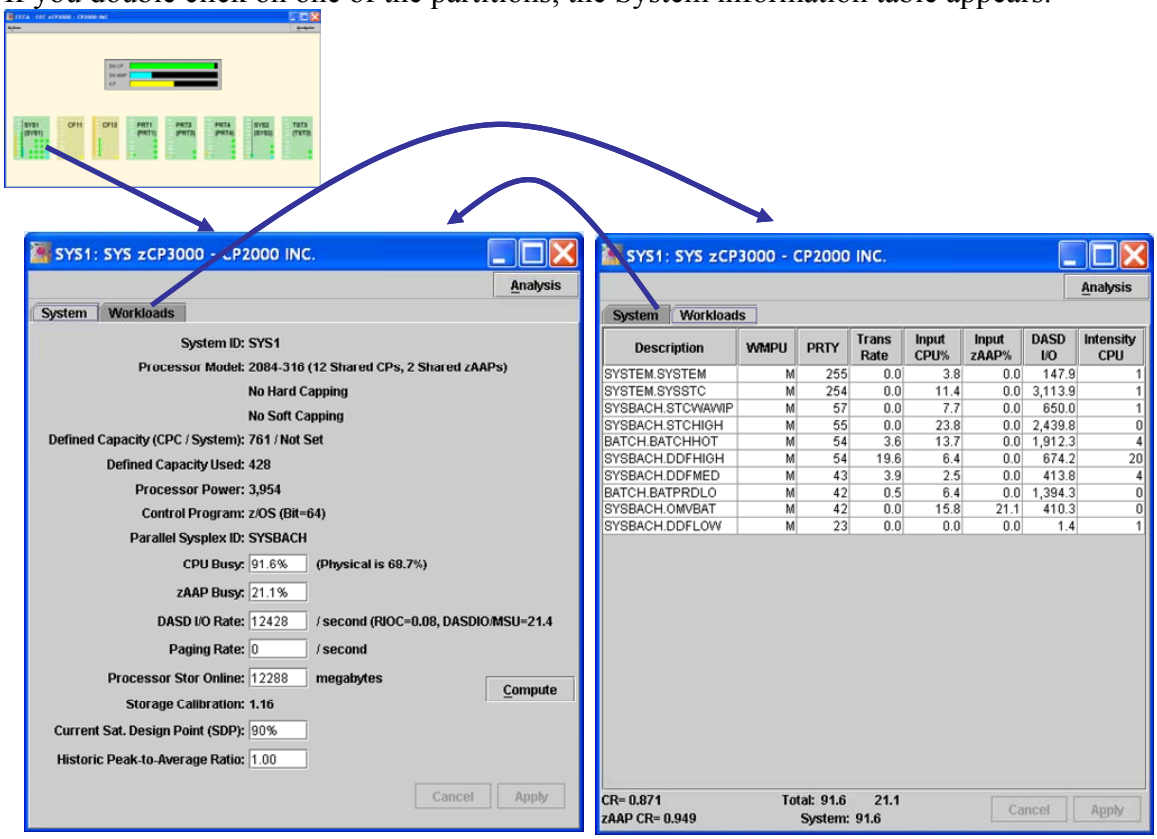

From here you can review the partition (system image) information or by selecting the Workload tab, the workload details. Each of these windows (shown together above), have an analysis button.

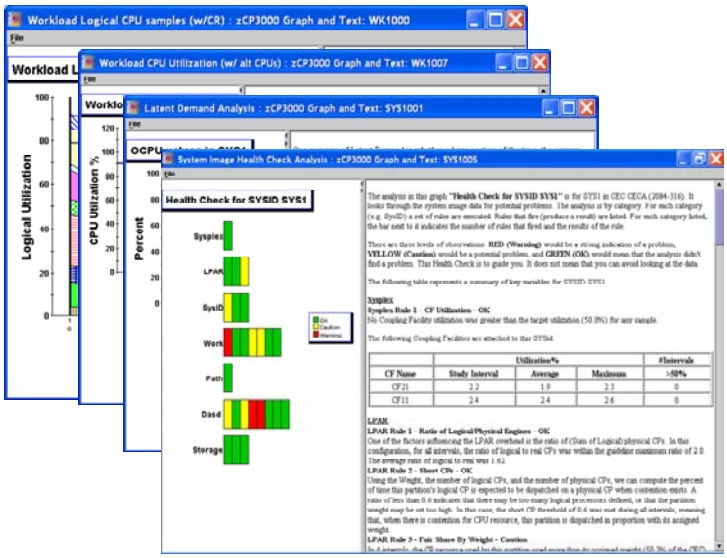

The analyses found off these two windows are particularly important in the performance analysis process. These should be reviewed carefully. As a high level indicator, the performance health of the system image, look at the Heath Check Analysis on the System window. This analysis puts the system level data (system performance variables, workload variables, and I/O information through a set of rules which assess the acceptability of the variable.

On the workload table, you can right click on a workload to delete it. You can also change the capture ratio methodology to I/O or System. Setting it to one will show the uncaptured time in the analyses.

If you use the workload window and double click on any workload, you can examine the workload details. Note that some important variables may not be in the EDF (because it may not be in RMF). For example, the number of transactions may not be available for a CICS service class. Using the transaction rate and CPU%, zCP3000 can figure out the MIPS/transaction. Without the transaction rate… no can do. However, you could provide an estimated MIPS/transaction and zCP3000 will compute the other since the three are related.

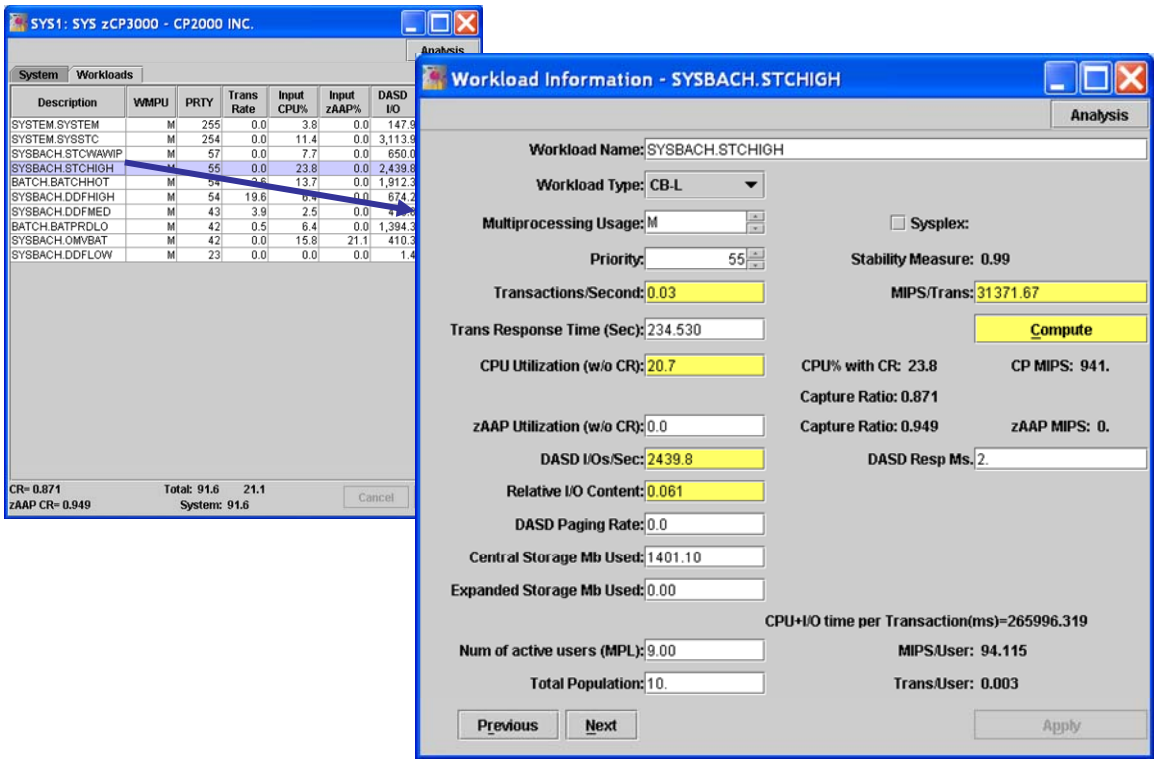

Enterprise I/O Analysis

In you now return to the Overview window and double click on any BCU, the BCU window will appear.

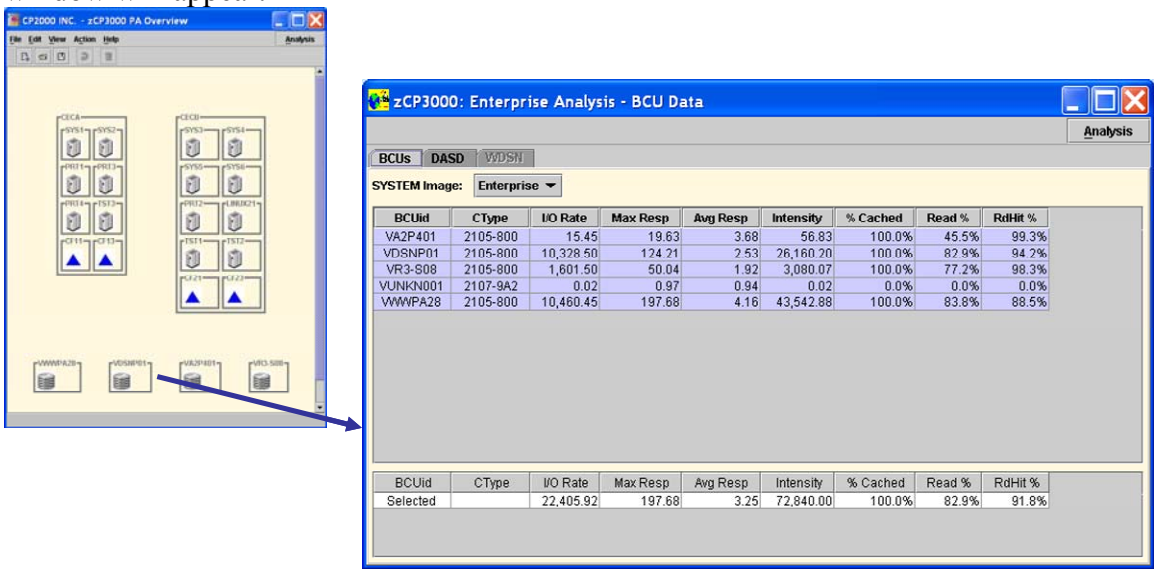

On this window you can review the BCU information for the Enterprise (averaged across all the system image usage data) or you can view the data from a specific partition. (Use the drop down list near the top which defaults to Enterprise.) Of course there's an analysis button here too.

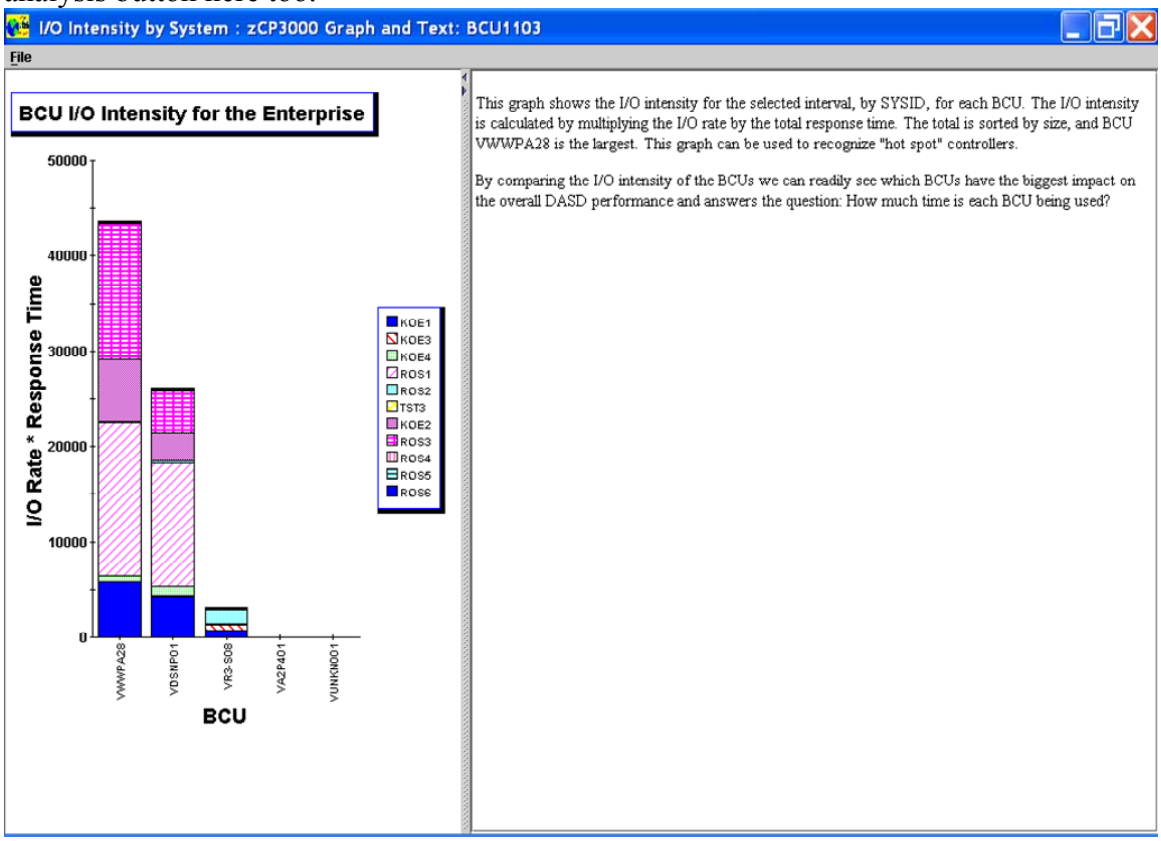

This analysis shows the I/O intensity for each BCU by partition. An easy way to answer the question: "Who's using which BCU?" This is very nice.

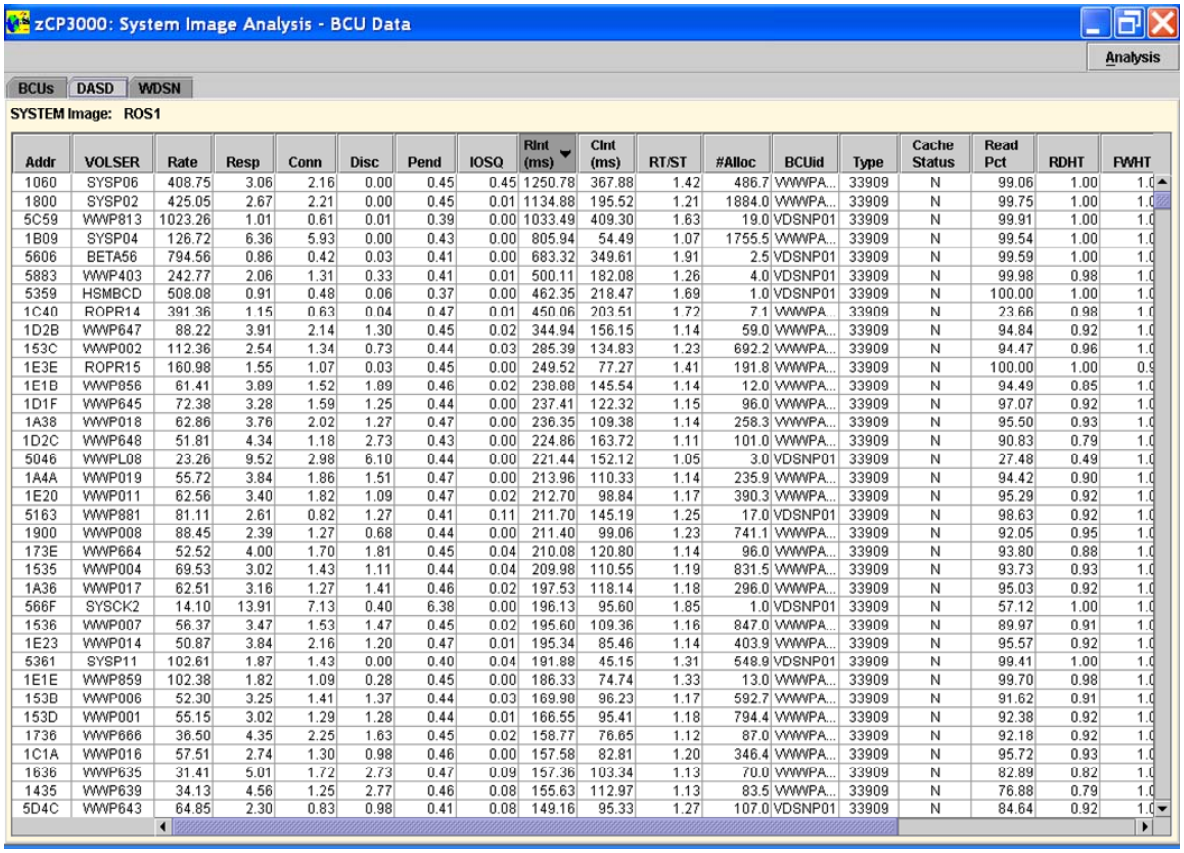

Addr	VOLSER	Rate	Resp	Conn	Disc	Pend	IOSQ	Rint (ms)	Clnt (ms)	RT/ST	#Alloc	BCUid	Type	Cache Status	Read Pct	RDHT	FWHT
1060	SYSP06	408.75	3.06	2.16	0.00	0.45	0.45	1250.78	367.88	1.42	486.7	VWWPA...	33909	N	99.06	1.00	1.0
1800	SYSP02	425.05	2.67	2.21	0.00	0.45	0.01	1134.88	195.52	1.21	1884.0	VWWPA...	33909	N	99.75	1.00	1.0
5C59	WWP813	1023.26	1.01	0.61	0.01	0.39	0.00	1033.49	409.30	1.63	19.0	VDSNP01	33909	N	99.91	1.00	1.0
1B09	SYSP04	126.72	6.36	5.93	0.00	0.43	0.00	805.94	54.49	1.07	1755.5	VWWPA...	33909	N	99.54	1.00	1.0
5606	BETA56	794.56	0.86	0.42	0.03	0.41	0.00	683.32	349.61	1.91	2.5	VDSNP01	33909	N	99.59	1.00	1.0
5883	WWP403	242.77	2.06	1.31	0.33	0.41	0.01	500.11	182.08	1.26	4.0	VDSNP01	33909	N	99.98	0.98	1.0
5359	HSMBCD	508.08	0.91	0.48	0.06	0.37	0.00	462.35	218.47	1.69	1.0	VDSNP01	33909	N	100.00	1.00	1.0
1C40	ROPR14	391.36	1.15	0.63	0.04	0.47	0.01	450.06	203.51	1.72	7.1	VWWPA...	33909	N	23.66	0.98	1.0
1D2B	WWP647	88.22	3.91	2.14	1.30	0.45	0.02	344.94	156.15	1.14	59.0	VWWPA...	33909	N	94.84	0.92	1.0
153C	WWP002	112.36	2.54	1.34	0.73	0.44	0.03	285.39	134.83	1.23	692.2	VWWPA...	33909	N	94.47	0.96	1.0
1E3E	ROPR15	160.98	1.55	1.07	0.03	0.45	0.00	249.52	77.27	1.41	191.8	VWWPA...	33909	N	100.00	1.00	0.9
1E1B	WWP856	61.41	3.89	1.52	1.89	0.46	0.02	238.88	145.54	1.14	12.0	VWWPA...	33909	N	94.49	0.85	1.0
1D1F	WWP645	72.38	3.28	1.59	1.25	0.44	0.00	237.41	122.32	1.15	96.0	VWWPA...	33909	N	97.07	0.92	1.0
1A38	WWP018	62.86	3.76	2.02	1.27	0.47	0.00	236.35	109.38	1.14	258.3	VWWPA...	33909	N	95.50	0.93	1.0
1D2C	WWP648	51.81	4.34	1.18	2.73	0.43	0.00	224.86	163.72	1.11	101.0	VWWPA...	33909	N	90.83	0.79	1.0
5046	WWPL08	23.26	9.52	2.98	6.10	0.44	0.00	221.44	152.12	1.05	3.0	VDSNP01	33909	N	27.48	0.49	1.0
1A4A	WWP019	55.72	3.84	1.86	1.51	0.47	0.00	213.96	110.33	1.14	235.9	VWWPA...	33909	N	94.42	0.90	1.0
1E20	WWP011	62.56	3.40	1.82	1.09	0.47	0.02	212.70	98.84	1.17	390.3	VWWPA...	33909	N	95.29	0.92	1.0
5163	WWP881	81.11	2.61	0.82	1.27	0.41	0.11	211.70	145.19	1.25	17.0	VDSNP01	33909	N	98.63	0.92	1.0
1900	WWP008	88.45	2.39	1.27	0.68	0.44	0.00	211.40	99.06	1.23	741.1	VWWPA...	33909	N	92.05	0.95	1.0
173E	WWP664	52.52	4.00	1.70	1.81	0.45	0.04	210.08	120.80	1.14	96.0	VWWPA...	33909	N	93.80	0.88	1.0
1535	WWP004	69.53	3.02	1.43	1.11	0.44	0.04	209.98	110.55	1.19	831.5	VWWPA...	33909	N	93.73	0.93	1.0
1A36	WWP017	62.51	3.16	1.27	1.41	0.46	0.02	197.53	118.14	1.18	296.0	VWWPA...	33909	N	95.03	0.92	1.0
566F	SYSCK2	14.10	13.91	7.13	0.40	6.38	0.00	196.13	95.60	1.85	1.0	VDSNP01	33909	N	57.12	1.00	1.0
1536	WWP007	56.37	3.47	1.53	1.47	0.45	0.02	195.60	109.36	1.16	847.0	VWWPA...	33909	N	89.97	0.91	1.0
1E23	WWP014	50.87	3.84	2.16	1.20	0.47	0.01	195.34	85.46	1.14	403.9	VWWPA...	33909	N	95.57	0.92	1.0
5361	SYSP11	102.61	1.87	1.43	0.00	0.40	0.04	191.88	45.15	1.31	548.9	VDSNP01	33909	N	99.41	1.00	1.0
1E1E	WWP859	102.38	1.82	1.09	0.28	0.45	0.00	186.33	74.74	1.33	13.0	VWWPA...	33909	N	99.70	0.98	1.0
153B	WWP006	52.30	3.25	1.41	1.37	0.44	0.03	169.90	96.23	1.17	592.7	VWWPA...	33909	N	91.62	0.91	1.0
153D	WWP001	55.15	3.02	1.29	1.28	0.44	0.01	166.55	95.41	1.18	794.4	VWWPA...	33909	N	92.38	0.92	1.0
1736	WWP666	36.50	4.35	2.25	1.63	0.45	0.02	158.77	76.65	1.12	87.0	VWWPA...	33909	N	92.18	0.92	1.0
1C1A	WWP016	57.51	2.74	1.30	0.98	0.46	0.00	157.58	82.81	1.20	346.4	VWWPA...	33909	N	95.72	0.93	1.0
1636	WWP635	31.41	5.01	1.72	2.73	0.47	0.09	157.36	103.34	1.13	70.0	VWWPA...	33909	N	82.89	0.82	1.0
1435	WWP639	34.13	4.56	1.25	2.77	0.46	0.08	155.63	112.97	1.13	83.5	VWWPA...	33909	N	76.88	0.79	1.0
5D4C	WWP643	64.85	2.30	0.83	0.98	0.41	0.08	149.16	95.33	1.27	107.0	VDSNP01	33909	N	84.64	0.92	1.0

If you select a specific system image from the drop down list and press DASD, you'll get a complete overview of the DASD data for that system.

You can sort any column by clicking on the header. Above you see the DASD sorted by Response Time Intensity. The analyses from here will display all the DASD. You can also select some number of the actuators (top ten?) and the analyses will be for only the selected actuators.

If you have SMF 42.6 type data in the EDF file, the Data Set button (WDSN) will be enabled. This will display the available data set information for the selected actuators.

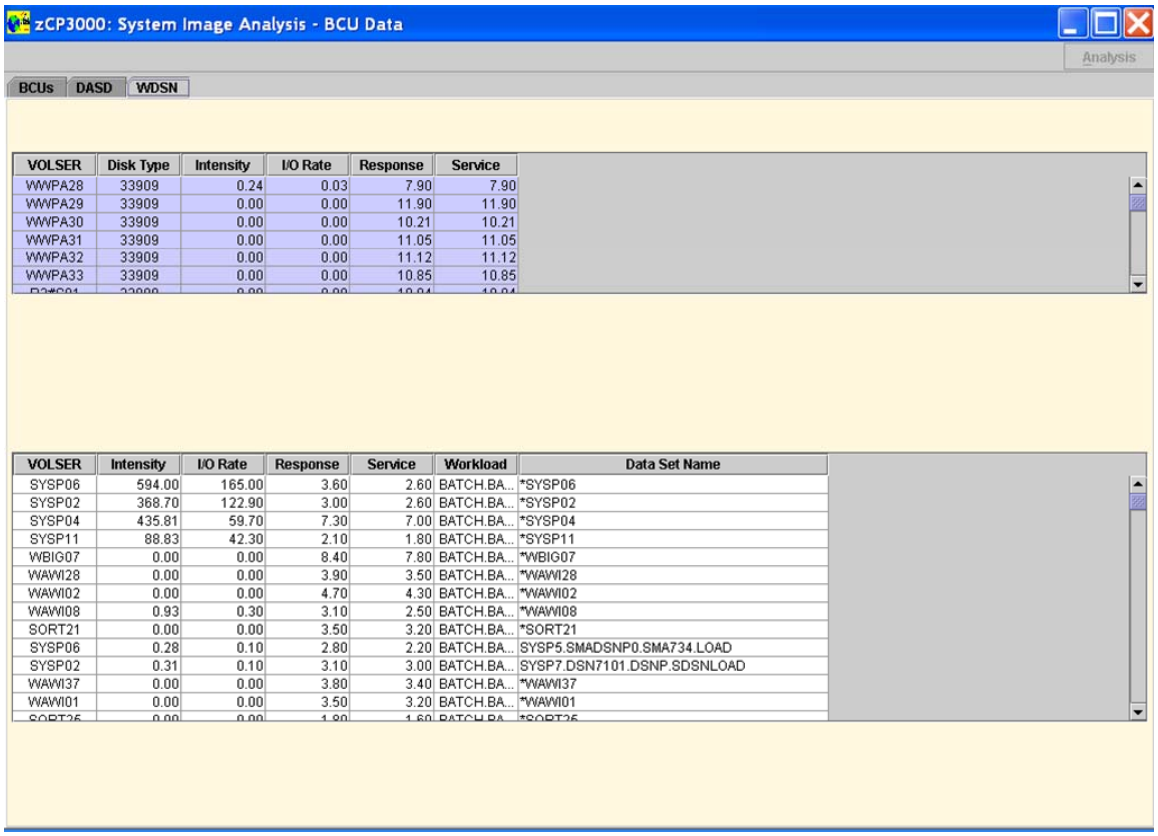

So, not only can you identify the problem actuators, you can identify the datasets on the actuator and the service class that is using the dataset. You can identify the problem and the application causing or suffering the pain… if you have the SMF 42.6 data.

Enterprise Sysplex Analysis

The Sysplex view of the enterprise is the most encompassing. You view the logical structure of the application. The communication between systems participating in the Sysplex is via locks, lists, and structures in the coupling facility (CF). Double click on any CF and the information appears.

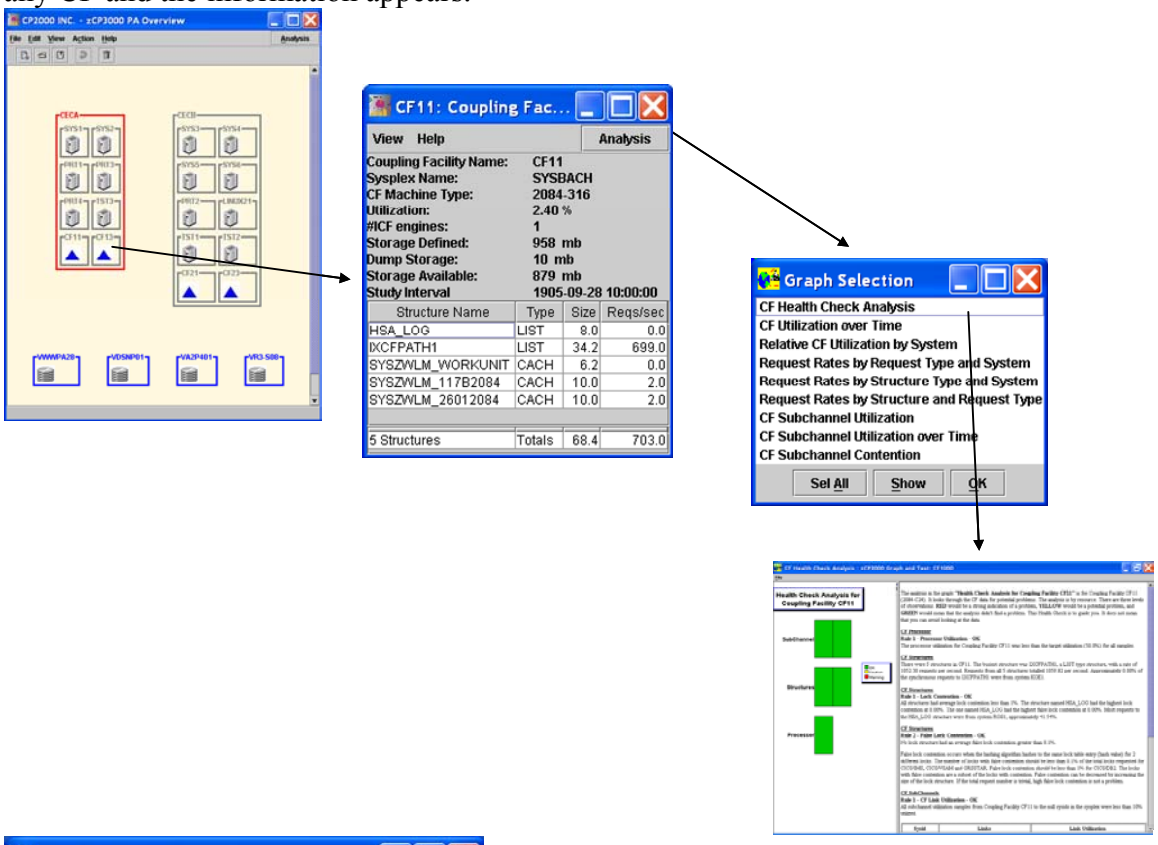

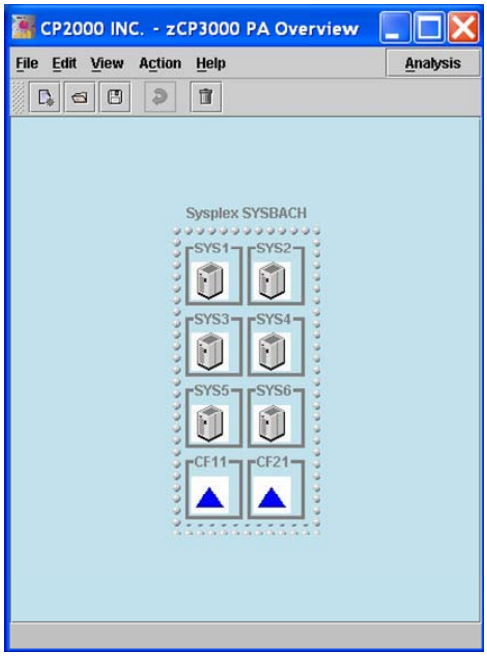

The analysis illustrated here is the CF Health Check. It is particularly useful in getting an overall view of the CF behavior.
You can also obtain a logical view of the Sysplex by means of the View item on the menu bar. This shows the Sysplex as a logical structure. If the enterprise had multiple Sysplexes, you would see them separated.

Alternate Processors

After reviewing the performance data, you can take a first look at the impact of a migration from the existing CEC model to a new model. This is done on the Overview window. Right click on a CEC, click on Alternates and select Alternatives.

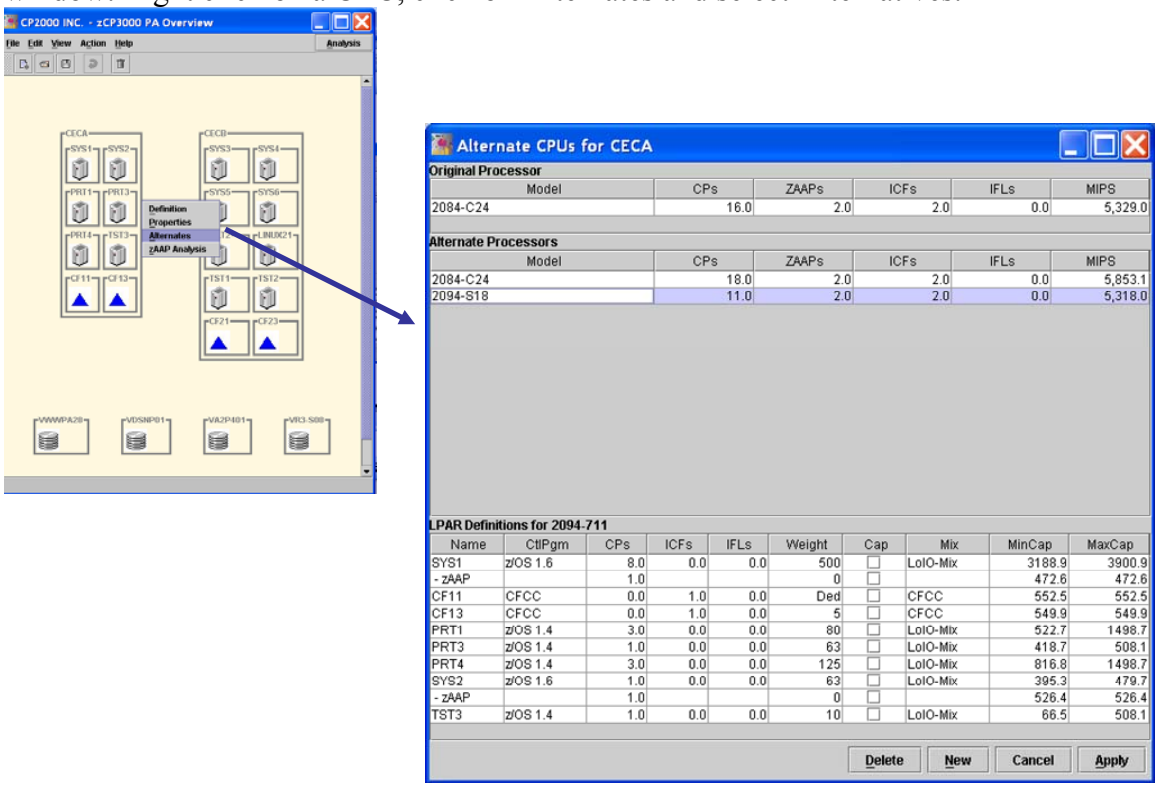

Processors are added to the alternatives list by pressing New. Since the alternative model may have a different number of PUs, zCP3000 rescales the logical configuration to match the base processor model. The logical configuration for each model can be inspected by selecting that model. You can also change the logical configuration if you are not satisfied with zCP3000's choice.

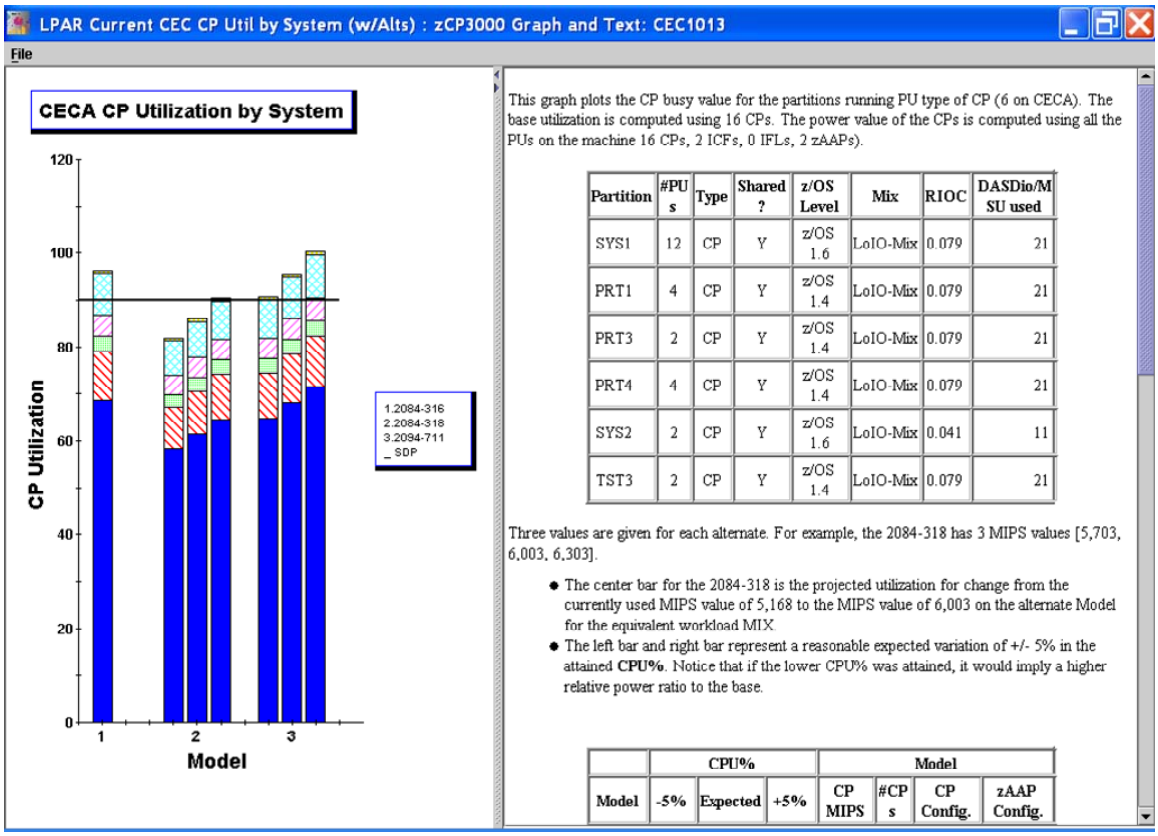

Once you have selected the models you want, press Apply to exit. This saves what you have done. Then go to the CEC window and look for analyses which compare the base model to the alternatives.

This analysis shows the base processor utilization (for the selected interval) and the projected utilization of the **exact** same work on the alternatives. You'll see three bars for each alternative. The center bar is the utilization scaled to the MIPS rating of the alternative. The bar to each side is the view if the MIPS rating were + or – 5% of the expected value. This is a warning that the MIPS rating is really not a single number but is an expected value with a range.

And more

On the Overview window there are a number of additional functions you can perform. Using the Actions menu item:

- You can create a small document containing the key graphs found in zCP3000.
- You can set the Saturation Design Point (SDP) for all system images. The default is 90%.
- zCP3000 uses the zPCR default processor 2084-301 to compute ITRRs. It then scales the ITRR to 450 MIPS. You can reset the scaling processes to match the MIPS requirement of the customer.
- And when you are ready, you can begin the Capacity Planning process. In this step you will specify the growth for the workloads. Before you enter this phase, you should save the model (in a file with qualifier .3pa). This is the subject of yet another zCP3000 monograph.