

Section 3: Specifying Guidance Variables

The CPEXPRT.USOURCE(DASGUIDE) partitioned data set member contains variables to establish the overall guidance for the DASD Component. You modify the variables in the DASGUIDE member whenever you wish to change the guidance to CPExpert. This chapter describes these variables, how the variables are used, and what should be specified for the variables.

The variables in the DASGUIDE module can be viewed as "data selection and presentation" variables and "analysis control" variables. These two types of control variables are discussed separately.

The data selection and presentation variables allow you to select particular time intervals to be analyzed, and allow you to specify how the results from the analysis are to be presented.

The analysis control variables allow you to control the analysis the DASD Component will perform. These variables specify whether expanded analysis is to be performed, identify critical workloads, specify critical data sets, exclude devices from analysis, control CPExpert's analysis of VSAM data set problems or potential problems, etc.

Please do not allow CPExpert to perform analysis or produce reports which are meaningless in your environment. If the analysis and reports produced by CPExpert do not meet your needs, alter the guidance to CPExpert. If the guidance is insufficient, please call Computer Management Sciences at (703) 922-7027 (or e-mail Don_Deese@cpexpert.com) so we can make changes to improve CPExpert for you!

Chapter 1: Data Selection and Presentation Variables

The data selection and presentation variables allow you to select particular time intervals to be analyzed, and allow you to specify how the results from the analysis are to be presented. This chapter describes these variables.

Exhibit 3-1 illustrates the data selection and presentation variables contained in USOURCE(DASGUIDE).

```

*****;
*   DATA SELECTION AND PRESENTATION VARIABLES   ;
*****;
%LET DASDATES  =01FEB1991;   * START DATE FOR DATA ANALYSIS   ;
%LET DASTIMES  =08:00:00;   * START TIME FOR DATA ANALYSIS   ;
%LET DASDATEE  =31DEC9999;   * END DATE FOR DATA ANALYSIS     ;
%LET DASTIMEE  =16:00:00;   * END TIME FOR DATA ANALYSIS     ;
%LET DASDAT2E  =0;         * DEFAULT SECOND SELECTION DATE - END   ;
%LET DASDAT2S  =0;         * DEFAULT SECOND SELECTION DATE - START  ;
%LET DASTIM2E  =0;         * DEFAULT SECOND SELECTION TIME - END    ;
%LET DASTIM2S  =0;         * DEFAULT SECOND SELECTION TIME - START   ;
%LET SHIFT     =Y;         * START AND END TIMES REFER TO A SHIFT   ;
%LET SYSTEM    =*ALL;      * SPECIFY SYSTEM TO PROCESS (*ALL=ALL)   ;
%LET SYSTEMn   =system;    * PROCESS SYSTEMn (n = 1-9)         ;
%LET SYSPLEX   =*ALL;      * SPECIFY SYSPLEX TO PROCESS (*ALL=ALL)  ;
%LET CONFIG    =&CPEWORK;  * SAS LIBRARY FOR I/O CONFIGURATION DATA ;
%LET SASODS    =N;         * CONTROLS WHETHER SAS ODS IS USED   ;
%LET PATH      =;         * PATH FOR ODS OUTPUT                 ;
%LET FRAME     =DASFRAME;  * GENERIC ODS FRAME NAME             ;
%LET CONTENTS  =DASDCONT;  * GENERIC ODS CONTENTS NAME          ;
%LET BODY      =DASDBODY;  * GENERIC ODS BODY NAME              ;
%LET STYLE     =;         * ODS HTML STYLE OPTION                 ;
%LET PDFODS    =N;         * CONTROLS WHETHER SAS PDF IS USED   ;
%LET LINKPDF   =;         * LINK TO CPEXPRT DOCUMENTATION        ;
%LET URL       =N;         * CONTROLS .HTM IN SAS ODS FRAME OUTPUT  ;
%LET VERBOSE   =V;         * RESULTS: VERBOSE/CONCISE/SUMMARY   ;
%LET MAXRULES  =1000;     * PRODUCE -REPORT TOO BIG- GUIDANCE ;

```

DATA SELECTION AND PRESENTATION VARIABLES

EXHIBIT 3-1

Chapter 1.1: DASDATES and DASTIMES variables

The DASDATES and DASTIMES variables specify the start date and start time, respectively, for the interval of SMF data the DASD Component is to analyze. These variables (in conjunction with the DASDATEE and DASTIMEE variables) allow you to select specific periods of data to analyze. The following example shows how to specify that data selection should start at 08:00:00 on March 6, 2001:

```
%LET DASDATES = 06MAR2001; * START DATE FOR DATA ANALYSIS;  
%LET DASTIMES = 08:00:00; * START TIME FOR DATA ANALYSIS;
```

Chapter 1.2: DASDATEE and DASTIMEE variables

The DASDATEE and DASTIMEE variables specify the end date and end time, respectively, for the interval of SMF data the DASD Component is to analyze. The following example shows how to specify that data selection should end at 17:00:00 on March 8, 2001:

```
%LET DASDATEE = 08MAR2001 * END DATE FOR DATA ANALYSIS;  
%LET DASTIMEE = 17:00:00; * END TIME FOR DATA ANALYSIS;
```

Chapter 1.3: DASDAT2S and DASTIM2S variables

The DASDAT2S and DASTIM2S variables are optional. These variables specify the start date and start time, respectively, for a second interval of SMF data the DASD Component is to analyze. These variables (in conjunction with the optional DASDAT2E and DASTIM2E variables) allow you to select a second period of data to analyze, in addition to the period specified by the DASDATES/DASTIMES and DASDATEE/DASTIMEE selection variables. The following example shows how to specify that a second period of data selection should start at 20:00:00 on March 6, 2001:

```
%LET DASDAT2S = 06MAR2001 * START DATE FOR DATA ANALYSIS;  
%LET DASTIM2S = 20:00:00; * START TIME FOR DATA ANALYSIS;
```

Chapter 1.4: DASDAT2E and DASTIM2E variables

The **DASDAT2E** and **DASTIM2E** variables are **optional**. These variables specify the end date and end time, respectively, for a second interval of SMF data the DAS Component is to analyze. These variables (in conjunction with the optional **DASDAT2S** and **DASTIM2S** variables) allow you to select a second period of data to analyze, in addition to the period specified by the **DASDATES/DASTIMES** and **DASDATEE/DASTIMEE** selection variables. The following example shows how to specify that a second period of data selection should end at 22:00:00 on March 6, 2001:

```
%LET DASDAT2E = 06MAR2001; * END DATE FOR DATA ANALYSIS;  
%LET DASTIM2E = 22:00:00; * END TIME FOR DATA ANALYSIS;
```

Chapter 1.5: SHIFT variable

The **SHIFT** variable is used with the **DASDATES**, **DASTIMES**, **DASDATEE**, and **DASTIMEE** variables. The **SHIFT** variable allows you indicate how the time-selection variables should be used.

- If the **SHIFT** variable is "N", the time-selection will be based upon the absolute start and end dates/times specified. For example, if you wish CPEXpert to process **all** data during a week, the start date and start time would be specified as the beginning of the week, and the end date and end time would be specified as the end of the week. You would specify "%LET SHIFT = N;" to process each 24-hour day. In the example shown above, data would be processed from 08:00:00 on 4 March until 17:00:00 on 8 March.
- If the **SHIFT** variable is "Y", the time-selection will be based upon the start and end dates, and the start and end times within each selected date. In the example shown above, perhaps you wished to process only the daily shift beginning at 08:00:00 and ending at 17:00:00. You would specify "%LET SHIFT = Y;" to process only the identified shift data, during the selected dates.

Chapter 1.6: SYSTEM variable

The **SYSTEM** variable is used to specify whether all systems in the performance data base should be evaluated, or to select a specific system identification to be evaluated.

Some users have data from multiple systems in their performance data base. For many of these users, or for users who have data for a single system represented in their

performance data base, the default "*"ALL" will be appropriate. No change of the SYSTEM variable would be required for these users.

However, some users who have data from multiple systems may wish to evaluate only a single system with the parameters specified in this member of DASGUIDE. For example, they might be temporarily interested in evaluating the performance of only an "important" system (such as a major production system) and not be interested in evaluating the performance of other systems with data in the performance data base. This evaluation can be accomplished by changing the SYSTEM variable to specify the system identification to be evaluated. For example, to specify that only data from SYS1 should be evaluated, specify:

```
%LET SYSTEM = SYS1 ; * PROCESS ONLY DATA FROM SYS1;
```

In another situation, a CPEXPERT user might wish to evaluate different systems with different DASGUIDE parameters. These different evaluations can be accomplished by different executions of the DASD Component. For each execution of the DASD Component, the USOURCE DD statement would be changed to reference different USOURCE libraries. Each USOURCE library would contain guidance members with appropriate guidance variables. The SYSTEM variable for each DASGUIDE guidance member would specify the system identification to which the guidance applied.

If you are using MICS as your performance data base, the SYSTEM variable refers to the MICS system identifier, rather than the SMF system identifier.

Chapter 1.7: SYSTEMn variable(s)

The SYSTEMn variable(s) are used to select multiple systems to be evaluated.

As described in the SYSTEM guidance variable discussion above, some sites have data from multiple systems in their performance data base. These sites can process data from all systems by specifying %LET SYSTEM=ALL; in USOURCE(CICGUIDE), or can select a specific system to process by specifying %LET SYSTEM=system; in USOURCE(CICGUIDE), where "system" is the system identification of the system to be processed.

Some sites have data from multiple systems in their performance data base and do not want to process all systems, but do wish to process more than one system. For example, some systems might be production systems and some might be test systems. For these

sites, the **SYSTEMn** guidance variable can be used to select more than one specific system to analyze.

The **SYSTEM** guidance variable can be used to select data from only one system to analyze, and the **SYSTEMn** guidance variable(s) can be used to select up to 9 additional systems to analyze. For example, if you wish to analyze data from four systems (named **SYSA**, **SYSB**, **SYSC**, AND **SYSX**) in a single execution of the DASD Component, specify:

```
%LET SYSTEM = SYSA ; * PROCESS DATA FROM SYSA;  
%LET SYSTEM1 = SYSB ; * PROCESS DATA FROM SYSB;  
%LET SYSTEM2 = SYSC ; * PROCESS DATA FROM SYSC;  
%LET SYSTEM3 = SYSX ; * PROCESS DATA FROM SYSX;
```

Chapter 1.8: SYSPLEX variable

The **SYSPLEX** variable is used to specify whether data from each sysplex in the performance data base should be evaluated, or whether CPEXpert should select data for a specific sysplex to be evaluated.

Some users have data from more than one sysplex in their performance data base. For many of these users, or for users who have data for a single sysplex represented in their performance data base, the default **"*ALL"** will be appropriate. No change of the **SYSPLEX** variable would be required for these users.

However, some users who have data from more than one sysplex may wish to evaluate only a single sysplex with the parameters specified in this member of **DASGUIDE**. This evaluation can be accomplished by changing the **SYSPLEX** variable to specify the sysplex to be evaluated. The following example shows how to specify that data only from **PRODPLEX** should be evaluated:

```
%LET SYSPLEX = PRODPLEX ; * PROCESS DATA ONLY FROM PRODPLEX;
```

Chapter 1.9: CONFIG variable

The **CONFIG** variable allows organizations to “save” the model of the I/O configuration that is created by CPEXpert into a SAS library, described by the **CONFIG DDNAME**. The default value for the **CONFIG** guidance variable is **%LET CONFIG=WORK;** which causes

the I/O configuration model to be placed into the SAS WORK temporary library. You can cause the DASD Component to place the I/O configuration into a permanent SAS library by using the CONFIG guidance variable. The I/O configuration model created by CPEXpert contains not only the I/O configuration (physical channels, logical channels, control units, and devices attached to each system), but also contains key performance metrics related to each aspect of the I/O configuration. This option can be useful for reporting I/O activity for capacity planning. If you wish to save the I/O configuration model produced by CPEXpert, specify **%LET CONFIG=ddname;**, where "ddname" is the name of JCL DD statement that describes the SAS library where the I/O configuration model is to be placed.

Chapter 1.10: SAS Output Delivery System

Output from CPEXpert is created using Basic SAS statements. This Basic SAS output is designed for a standard SAS printer (line) format. With SAS Release 8, SAS users can use the SAS Output Delivery System to create output that is formatted in Hypertext Markup Language (HTML). This output can be browsed with Internet Explorer, Netscape, or any other browser that fully supports the HTML 3.2 tag set.

The CPEXpert WLM Component, DB2 Component, CICS Component, and DASD Component support the SAS ODS features.

Please reference the *CPEXpert Installation Guide* for more detailed information about using the SAS ODS feature of CPEXpert.

Chapter 1.11: VERBOSE variable

The VERBOSE variable provides a control on the amount of narrative that the DASD Component lists with each rule result. Some installations prefer to produce only concise findings, and evaluate the results when the findings are significant. Other installations wish to produce expanded findings, and evaluate the results on a daily basis. You can use the VERBOSE variable to control the amount of narrative, depending upon your preferences. The options with the VERBOSE variable are:

- V** = Print verbose comments related to each rule that was produced during the evaluation of DASD performance. The verbose comments describe the rule, provide key information associated with the rule, and may provide a specific reference related to the rule.
- C** = Print "concise" results about the DASD devices. Each "rule record" created by the DASD Component has a "level" associated with it. "Level 1" records are those records describing basic information about the DASD devices with the most potential for performance improvement. "Level 2" records are those

records which analyze the causes of poor performance and suggest ways in which improvements may be made. When the **C** verbose option is specified, only Level 1 information is printed.

S = Print “summary” results about the DASD devices. When the **S** verbose option is specified, only DAS000 (sysplex-wide intensity) and DAS050 (system intensity) information is printed.

Chapter 1.12: MAXRULES variable

A very large amount (several hundred pages) of output can be produced during some executions of the DASD Component. Such a large amount of output is not easy to analyze, and tends to cause the report to be rejected by analysts or management. Fortunately, there are several options in the DASD Component that can be used to reduce the amount of output. New users of the DASD Component might not be aware of these options, and might not be favorably impressed with the DASD Component since they would perceive the DASD Component output to be useless. This is not a desirable result!

The MAXRULES variable is used to control whether CPEXpert detects that “a large amount of output” was produced. CPEXpert produces suggestions on how to reduce the output to a more manageable size when the number of rules (and supporting data) exceed the value of the MAXRULES guidance variable.

The default value of the MAXRULES guidance variable is 1000, which should be approximately 50 pages of output (basic rules have an average of 10 lines per description, while supporting rules have only one line for the data). You can alter the MAXRULES guidance variable and cause CPEXpert to produce the suggestions for ways to reduce the output when either more or less rules (and supporting data) are produced. For example, to indicate that the suggestions on ways to reduce the output should be printed only when more than 2000 rules were produced, specify:

```
%LET MAXRULES=2000 ; * PRODUCE “REPORT TOO BIG” GUIDANCE;
```

Chapter 2: Analysis Control Variables

The analysis control variables allow you to control how the DASD Component analyzes your DASD configuration. These variables can be used to (1) specify whether expanded analysis is to be performed based upon workload categories, (2) specify whether expanded analysis is to be performed based upon data set names, (3) exclude volumes from analysis altogether, or (4) select specific volumes for analysis.

The analysis control variables are optional. You do not have to specify any analysis control variables unless you wish to alter the basic processing performed by the DASD. This chapter describes the analysis control variables associated with the DASD Component. Exhibit 3-2 illustrates the analysis control variables contained in USOURCE(DASGUIDE). Please note that these are basic analysis guidance variables. Additional guidance variables are listed in other chapters in this section.

```

*****;
*   ANALYSIS GUIDANCE VARIABLES                               ;
*****;
%LET ANALYZE      = 3;   * DEFAULT FOR NUMBER OF DEVICES TO ANALYZE      ;
%LET DASDEXCP    =100;  * MINIMUM NUMBER OF EXCP TO SELECT (DASD30DD)    ;
%LET DASDSN      = N;   * ANALYZE CRITICAL DATA SETS?                    ;
%LET EVALDASD    =Y;    * OPTION TO PRODUCE ONLY I/O CONFIGURATION        ;
%LET EXCLUDE     = N;   * DEFAULT FOR EXCLUDING VOLUMES                  ;
%LET GOALMODE    = Y;   * OPERATING IN GOAL MODE                          ;
%LET LIST42DS    =10;   * LIST ONLY TOP NUMBER OF DATA SETS             ;
%LET LISTALL     =Y;    * LIST DATA FOR ALL RMF INTERVALS;              ;
%LET LOVED1      =N;    * PERFORM 'LOVED ONE' ANALYSIS?                   ;
%LET MIN42PCT    =0.1;  * EXCLUDE REPORTING LOW-ACTIVITY DATA SETS      ;
%LET MINIORT     =0     * MINIMUM I/O RATE TO ANALYZE                     ;
%LET MINIOWT     =0;    * MINIMUM I/O INTENSITY (OR WEIGHT) TO ANALYZE   ;
%LET MINRESP     =0;    * MINIMUM I/O RESPONSE TO ANALYZE                ;
%LET REPORT      =11;   * NUMBER OF TOTAL VOLUMES TO REPORT              ;
%LET SELECT      = Y;   * DEFAULT FOR VOLUME SELECTION                   ;
%LET SHARED      = N;   * DEFAULT FOR PROCESSING SHARED SYSTEM DATA     ;
%LET TYPE30DD    = N;   * MXG TYPE30 MODIFICATION (Y = AVAILABLE)         ;
%LET TYPE42DS    = N;   * MXG TYPE42DS (Y = AVAILABLE)                   ;
*****;

```

ANALYSIS GUIDANCE VARIABLES

EXHIBIT 3-2

Chapter 2.1: Number of devices to analysis: ANALYZE variable

The optional ANALYZE variable is used to specify the number of devices which CPEXpert should analyze in detail.

As explained in Section 1, most Data Base Administrators are able to address a limited number of DASD problems each day. The DASD Component was designed to identify the most serious DASD problems. The initial design identified the device with the most potential for improvement, and analyzed this device in detail.

Some users of CPEXpert felt that the "identify the device with the most potential for improvement" design philosophy was unnecessarily restrictive. These users wanted to specify the number of devices which should be analyzed in detail. The ANALYZE variable meets this requirement.

You can direct CPEXpert to analyze up to 999 DASD devices in a single execution, by using the ANALYZE variable to specify the number of devices to analyze in detail. CPEXpert will provide a detailed analysis of all devices whose device response time exceeded the average device response time for their device type, up to the number specified by the ANALYZE variable. For example, if you wish CPEXpert to analyze the top 25 devices with the worst performance, specify **%LET ANALYZE=25;** in USOURCE(DASGUIDE).

You likely will receive an unwieldy amount of output if you specify a large number of devices using the ANALYZE variable. One way to limit the output is to specify **%LET VERBOSE=S** so that only device summary results will be printed.

Chapter 2.2: Exclude devices with low activity: DASDEXCP variable

The modification to MXG or NeuMICS includes a MINIO parameter which specifies the minimum number of I/O operations to any data set within a particular SMF Type 30(Interval) record. The CPEXpert modification will ignore any data sets encountered if the number of I/O operations is below the MINIO value. The default for the MINIO parameter is MINIO=100, indicating that data sets will be included in the DASD30DD data set generated by the CPEXpert modification if at least 100 EXCPs were directed to the data set.

This screening is based on data set, and is not based on the total I/O activity to the device. Consequently, some devices might have data represented in the DASD30DD data set (generated by the CPEXpert modification), but these devices might have very low activity. For example, a device could be represented in the DASD30DD with only 100 EXCPs per SMF Type 30 recording interval if only a single data set were referenced 100 times.

It normally is not interesting to consider devices with a very low activity, and many sites have a large number of such devices in the DASD30DD data set created by CPEXpert. To reduce processing time, the DASDEXCP guidance variable can be used to exclude such low-activity volumes.

The default value for the DASDEXCP guidance variable is 100 (the same as the MINIO variable in the CPEXpert modification) that creates the DASD30DD data set. If you wish CPEXpert to ignore devices with low activity, change to DASDEXCP guidance to a larger value. For example, if you wish CPEXpert to consider only those devices with more than 1000 EXCPs per SMF Type 30 recording interval, specify **%LET DASDEXCP=1000;** in USOURCE(DASGUIDE).

Chapter 2.3: Analyze using response objectives: DASDSN variable

The DASGUIDE **DASDSN** guidance variable specifies whether to perform expanded analysis based upon data set response objectives. Specify **%LET DASDSN = Y**; to tell CPExpert to perform expanded analysis based upon data set response objectives.

Please refer to Section 2 (Chapter 4) for operational notes regarding changing your specification of critical data sets or changing response objectives associated with the data sets.

Please refer to Section 3 (Chapter 5) for a description of the procedures to follow if you wish to perform analysis based on data set response objectives.

Chapter 2.4: Produce only I/O configuration: EVALDASD variable

The EVALDASD guidance variable allows organizations to suppress the evaluation of DASD I/O problems. This option is for those users who wish to use the DASD Component simply to produce SAS data sets representing the I/O configuration (and associated data) for each system in the sysplex, and a SAS data set for the overall sysplex.

The DASD Component processes SMF data (in a performance data base) to create a single record for each RMF measurement interval, that contains information about the channel paths, controllers, and device for each volume. This information is interesting to some users for management reporting purposes, or for data base administrator reporting purposes. These users do not necessarily want to use the DASD Component to analyze DASD problems, since their sole use of the DASD Component is to have an easy-to-use tool that creates the I/O configuration information. This option can be useful for reporting I/O activity for capacity planning.

- If the EVALDASD guidance variable is "Y" (which is the default value), the DASD Component will create the model of the I/O configuration and evaluate DASD performance problems represented by this model.
- If the EVALDASD guidance variable is "N", the DASD Component will create the model of the I/O configuration and place the results in the SAS library described by the optional **CONFIG DDNAME** DD statement. The DASD Component will then issue an **ABORT RETURN 100**; statement within the SAS coding. This statement results in SAS aborting the SAS session, with a return code of 100.

Chapter 2.5: Excluding volumes from analysis: EXCLUDE variable

The EXCLUDE guidance variable is used to tell CPEXpert whether you wish to exclude volumes from analysis. The EXCLUDE variable acts as a "switch" to control whether CPEXpert processes the DASGUIDE member searching for volumes to exclude from analysis. The point of having a "switch" variable is that some installations may wish to regularly exclude volumes from analysis, but periodically analyze all volumes.

If the EXCLUDE variable is **N**, CPEXpert will not exclude any volumes from analysis. If the EXCLUDE variable is **Y**, CPEXpert will process the DASGUIDE member to identify all volumes to be excluded.

Please refer to Chapter 3 for more information on excluding volumes from analysis.

Chapter 2.6: Specifying data sets to list: LIST42DS variable

If SMF Type 42 (Data Set Statistics) information is available¹ in a MXG performance data base, the DASD Component will process MXG TYPE42DS to select data sets that were referenced during RMF measurement intervals in which the poorly performing devices exceeded the average performance. There can be many data sets referenced on the devices (hundreds or even thousands of data sets can be referenced). It is not helpful to have a large number of data sets listed. Consequently, CPEXpert uses the **LIST42DS** variable to limit the number of data sets listed in any RMF measurement interval, for a particular device. Only the number of data sets specified by the LIST42DS variable will be listed individually, and any remaining data sets will be summarized and listed on a single line.

The default value for the LIST42DS variable is 10, indicating that 10 data sets will be listed for each RMF interval in which a poorly-performing device had a performance problem.

You can alter this number of data sets listed by using the LIST42DS guidance variable. For example, if you wish CPEXpert to list only 5 data sets for each RMF interval in which a poorly-performing device had a performance problem, specify **%LET LIST42DS=5;** in USOURCE(DASGUIDE).

Chapter 2.7: Perform "loved one" analysis: LOVED1 variable

The LOVED1 guidance variable is used to control whether the DASD Component performs "loved one" analysis.

¹**%LET TYPE42DS = Y;** must be specified in USOURCE(GENGUIDE) or USOURCE(DASGUIDE) to advise CPEXpert that TYPE42DS is available.

IBM speakers at professional conferences have long advocated the performance analysis principle of "Know your loved ones and always have someone to kick around." As described in Section 1 (Chapter 5.2) and Section 2 (Chapter 5), you can identify workload categories (such as TSO, CICS, BATCH, etc.) to the DASD Component. The DASD Component can then analyze your DASD configuration from the view of the workload categories, focusing on the performance provided to your "loved one" workload.

Analysis of critical (or "loved one") workload is possible only if you accomplish the following:

- You have installed the modification to MXG or MICS to collect application-related information as described in Section 2 of this User Manual. The DASD Component cannot perform expanded analysis based on workload categories unless it has information relating DASD use to specific performance groups or service classes, jobs, and job steps. This information is extracted from the SMF Type 30(DD) records by the modifications described in Chapter 1 and Chapter 2 of Section 2.
- You have defined workload categories to the DASD Component. At least one workload category (the most critical category) must be defined using the process described in Chapter 3 of Section 2.
- You identify the critical (or "loved one") workload to the DASD Component using the **LOVED1** guidance variable.
- You specify **%LET TYPE30DD = Y;** in the USOURCE(GENGUIDE) CPEXpert general guidance module. This tells CPEXpert that you have installed the modification to MXG or MICS to collect application-related information.

If you accomplish the above steps, the DASD Component will perform an expanded analysis of your DASD configuration from the perspective of the "loved one" workload category. That is, only devices and channel paths used by the critical workload will be analyzed. This analysis can be quite useful to determine which DASD devices are providing the worst performance from the perspective of the critical workload.

The **LOVED1** guidance variable specifies whether to perform expanded analysis of workload categories, and specifies the most critical (or "loved one") workload.

For example, if your critical workload is CICS production executing in Service Class CICSAOR1, Service Class CICSAOR2, and Service Class CICSAOR3 you might define the CICS workload as illustrated in Exhibit 2-9 in Section 2. You could identify the CICS workload as your most critical workload as:

```
%LET LOVED1 = CICS ; * CICS IS THE CRITICAL WORKLOAD;
```

In this example, the DASD Component will analyze all DASD use by Service Class CICSOR1, CICSOR2, and CICSOR3. The DASD Component will isolate the devices providing the worst performance from the view of this CICS workload, identify likely causes of poor performance, identify the job and job steps **in other service classes** which **probably** cause performance to be poor, and suggest ways to correct the problems. A similar analysis would be performed for performance groups, if you are operating in Compatibility Mode.

A significant phrase was used in the above paragraph: "probably cause performance to be poor". Please review the discussion associated with appropriate rules in Appendix A to appreciate why the identification might not be precise, but might be only "probable". In brief summary, the identification is "probable" rather than "precise" because of synchronization problems between SMF Type 30 data and SMF Type 70(series) data, unless you have specified that the SMF/RMF recording be synchronized. This issue is discussed in some detail in Section 5 and in appropriate rules in Appendix A.

The DASD Component will not perform "loved one" analysis if the LOVED1 macro variable is null (this is the default specification).

The "loved one" analysis is incompatible with the data set response objective analysis². That is, if you perform "loved one" analysis, you cannot perform an analysis of response objectives for critical data sets in the same execution of CPEXpert. You could perform "loved one" analysis in one execution of the DASD Component and perform data set response objective analysis in another execution of the DASD Component, processing the same data in your performance data base.

Chapter 2.8: Exclude reporting low-activity data sets: MIN42PCT variable

If SMF Type 42 (Data Set Statistics) information is available³ in a MXG performance data base, the DASD Component will process MXG TYPE42DS to select data sets that were referenced during RMF measurement intervals in which the poorly performing devices exceeded the average performance. There can be many data sets referenced on the devices (hundreds or even thousands of data sets can be referenced). It is not helpful to have a large number of data sets listed. Additionally, some data sets might have a low I/O activity and would not be interesting to analyze. Consequently, CPEXpert uses the **MIN42PCT** variable to limit the number of data sets listed in any RMF measurement interval, for a particular device. Only data sets having a percent of activity for the total

²This design decision was made because there would be significant problems with interpreting the results from CPEXpert if both types of analysis were performed simultaneously, and because a large coding effort would be required to perform the concurrent analysis. Please call Computer Management Sciences if you have a strong requirement for concurrent analysis.

³%LET TYPE42DS = Y; was specified in USOURCE(GENGUIDE).

volume greater than the percent specified by the MIN42PCT variable will be listed individually, and any remaining data sets will be summarized and listed on a single line.

The default value for the MIN42PCT variable is 0.1, indicating that data sets will not be listed individually for any RMF interval unless the data set intensity of access (I/O rate * response time) was greater than 0.1% of the total volume intensity of access. You can alter this percent of data sets listed by using the MIN42PCT guidance variable.

For example, if you specified **%LET MIN42PCT = 25;** in USOURCE(DASGUIDE), a maximum of 4 data sets would be listed in any RMF interval since no more than 4 data sets could have 25% or higher access intensity.

Please note that, regardless of the data set access intensity of any particular data set, only the number of data sets specified by the **LIST42DS** guidance variable (described earlier) will be listed. This means that there are two ways to limit the number of data sets listed: (1) the LIST42DS which limits the number of data sets listed in any RMF interval, and (2) the **MIN42PCT** guidance variable which limits the data sets listed to those that exceed the specified percent.

Chapter 2.9: List data for all RMF intervals: LISTALL variable

Once the DASD Component has selected one or more devices to analyze, it analyzes device-related data for all RMF intervals. This analysis determines whether each device has performance problems, and determines the most significant problems during the RMF interval. When the DASD Component produces the results of its analysis, it shows device information for each RMF interval analyzed, regardless of whether a performance problem existed with the device. This listing of all RMF intervals is the default, since it presents a continuous view of the device performance for all RMF intervals in the performance data base.

Some users analyze DASD problems for many systems in a single execution of the DASD Component, and they collect RMF data on short time intervals (for example, 15 minute RMF intervals or shorter). The default listing of data for all RMF intervals can produce very large reports for these users.

The LISTALL variable can be used to suppress the listing of RMF intervals in which no performance problem was detected with a device. You can suppress listing all RMF intervals by specifying **%LET LISTALL=N;** in USOURCE(DASGUIDE). With this specification, the DASD Component will list only those intervals in which performance problems were detected.

Chapter 2.10: Minimum I/O rate to analyze: MINIORT variable

Before the DASD Component selects a device for analysis, the code determines whether the device has an I/O rate greater than the MINIORT value. This feature allows user to exclude low-activity devices from analysis. The default value for the MINIORT variable is 0, indicating that all devices will be candidates for analysis, regardless of their I/O rate. You can alter this I/O rate by using the MINIORT guidance variable.

For example, if you specified **%LET MINIORT = 1;** in USOURCE(DASGUIDE), only devices with a minimum I/O rate greater than 1 I/O per second would be selected for analysis.

Chapter 2.11: Minimum I/O rate to analyze: MINIOWT variable

Before the DASD Component selects a device for analysis, the code determines whether the device has an I/O “intensity” (or weight) greater than the MINIOWT value. This feature allows user to exclude low-activity devices from analysis. The default value for the MINIOWT variable is 0, indicating that all devices will be candidates for analysis, regardless of their I/O intensity. You can alter this I/O rate by using the MINIOWT guidance variable.

For example, if you specified **%LET MINIOWT = 1000;** in USOURCE(DASGUIDE), only devices with a minimum I/O intensity greater than 1000 would be selected for analysis.

Chapter 2.12: Minimum I/O response to analyze: MINRESP variable

Before the DASD Component selects a device for analysis, the code determines whether the device has an I/O response greater than the MINRESP value. This feature allows user to exclude low-activity devices from analysis. The default value for the MINRESP variable is 0, indicating that all devices will be candidates for analysis, regardless of their I/O response. You can alter this I/O rate by using the MINRESP guidance variable. The MINRESP variable is specified in millisecond units.

For example, if you specified **%LET MINRESP = 5;** in USOURCE(DASGUIDE), only devices with a minimum I/O response greater than milliseconds would be selected for analysis.

Chapter 2.13: Number of volumes to report: REPORT variable

The DASD Component analyzes performance information about the device(s) with the most performance improvement potential, based on ANALYZE guidance variable described earlier. To put the performance improvement potential of these devices in perspective relative to the overall sysplex or system, the DASD Component also produces

a brief summary of relevant characteristics of other devices, with the list ranked in descending order by device access intensity (I/O RATE * average I/O response). The summary by sysplex is produced in Rule DAS000 and the summary by system is produced in Rule DAS050.

The number of devices for which summary information is produced is controlled by the **REPORT** guidance variable. The default for the REPORT guidance variable is 10, indicating that information is produced for the top 10 devices on a system basis. The value of the REPORT variable is doubled for the sysplex view.

Chapter 2.14: Selecting volumes to analyze: SELECT variable

The SELECT guidance variable is used to tell CPEXpert whether you wish to select specific volumes for analysis. The SELECT variable acts as a "switch" to control whether CPEXpert processes the DASGUIDE member searching for volumes to select for analysis. The point of having a "switch" variable is that some installations may wish to regularly analyze all volumes, but periodically select specific volumes for analysis.

If the SELECT variable is **N**, CPEXpert will not select specific volumes for analysis, but will analyze all volumes after applying data selection criteria. If the SELECT variable is **Y**, CPEXpert will process the DASGUIDE member to identify the volumes to be analyzed.

Please refer to Chapter 4 for more information on selecting volumes to analyze.

Chapter 2.15: Analyze shared DASD Conflicts - SHARED variable

CPEXpert can perform an analysis of conflicts between DASD shared between systems or MVS images. The shared DASD analysis is accomplished when you perform the following steps:

- You collect and process SMF data for each system to be analyzed. The SMF data must be placed into your standard performance data base. This step is not specific for CPEXpert; it is the normal processing you would do. The step is listed first just for completeness.
- **You should ensure that the system clocks on all systems are set at roughly the same time.** CPEXpert will examine the SMF data contained in your performance data base, based upon the SMF time associated with the measurement data. Shared DASD problems between systems will be detected based on this SMF time. If the SMF times are significantly different, the analysis might not properly identify conflicts or conflicting applications.

You should not worry that the system clocks be set at **exactly** the same time. The analysis performed by CPEXpert is not intended to identify an isolated performance problem. Rather, CPEXpert attempts to identify those problems which **continually** cause shared DASD performance problems. If a shared DASD problem is continual, the problem will be reflected in multiple SMF recording intervals. Consequently, it is not essential that the SMF times be exactly matched between systems.

If the shared devices are in the same sysplex, this step will be performed automatically by the sysplex timer. The step is listed just for completeness.

- You specify **%LET SHARED = Y;** in USOURCE(DASGUIDE) to tell CPEXpert to perform analysis of shared DASD.
- If you wish the DASD Component to identify applications causing shared DASD conflicts, you must perform the following additional steps:
 - Install the modification to MXG or MICS to collect data set-related information as described in Section 2 of this User Manual.
 - Specify **%LET TYPE30DD = Y;** in the USOURCE(GENGUIDE) general guidance module. This tells CPEXpert that you have installed the modification to MXG or NeuMICS to collect application-related information.

Chapter 2.16: SMF Type 30 modification installed - TYPE30DD variable

The **TYPE30DD** guidance variable tells CPEXpert whether you have installed the CPEXpert modification to MXG or MICS code to collect SMF Type 30 (Data Definition) information⁴.

This modification allows CPEXpert to relate DASD performance information contained in SMF Type 74 records to specific service classes or performance groups, and to relate the Type 74 information to specific jobs or job steps.

The TYPE30DD statement should be specified as **%LET TYPE30DD = Y;** if you have installed the CPEXpert modification to collect Type 30 (Data Definition) information for the system(s) being analyzed.

⁴The TYPE30DD guidance variable is normally specified when installing CPEXpert and is discussed in the *CPEXpert Installation Guide*. The TYPE30DD guidance variable is included in this document simply to remind you that the SMF Type 30 information is required if you wish to have CPEXpert show applications that access devices with poor performance, wish to perform "loved one" analysis, or wish to perform critical data set analysis. The TYPE30DD is not included in the USOURCE(DASGUIDE) member released with CPEXpert, since a specification in USOURCE(DASGUIDE) would override any specification in USOURCE(GENGUIDE).

Chapter 2.17: SMF Type 42 (Data Set Statistics)- TYPE42DS variable

The **TYPE42DS** guidance variable tells CPEXpert whether SMF Type 42 (Data Set) records are available for analysis⁵. If you collect the SMF Type 42 (Data Set) records, CPEXpert can report data set information⁶ for data sets residing on volumes with poor performance.

The default value for the TYPE42DS guidance variable is "N", indicating that the TYPE42DS data set is not available for analysis. You can change this guidance if Type 42 (Data Set) records are available, and if you wish CPEXpert to report data set information, by specifying **%LET TYPE42DS=Y;** in USOURCE(DASGUIDE).

⁵The TYPE42DS data set statistics information is available only with a MXG performance data base.

⁶The TYPE42DS guidance variable is normally specified when installing CPEXpert and is discussed in the *CPEXpert Installation Guide*. The TYPE42DS guidance variable is included in this document simply to remind you that the SMF Type 42 information is required if you wish to have CPEXpert show data sets that are accessed on devices with poor performance. The TYPE42DS is not included in the USOURCE(DASGUIDE) member released with CPEXpert, since a specification in USOURCE(DASGUIDE) would override any specification in USOURCE(GENGUIDE).

Chapter 3: Excluding volumes from analysis

For a variety of reasons, you may wish to exclude certain volumes from analysis. The most common reason is that CPEXpert repeatedly identifies problems with particular volumes, but (1) you do not wish to make changes to correct the problems, (2) you are unable to make changes (because of application requirements or political realities), or (3) the DASD Component analysis is "flawed" because of data problems or data averaging. Whatever the reason, you may wish to exclude certain volumes from analysis.

Excluding volumes from analysis is optional; you do not have to exclude any volumes from analysis. If you do **not** exclude volumes from analysis, CPEXpert will analyze every DASD VOLSER encountered in SMF Type 74 data after applying data selection criteria.

If you wish to exclude volumes from analysis, you must (1) use the **EXCLUDE** guidance variable and (2) define the volumes to exclude.

Chapter 3.1: EXCLUDE variable

The EXCLUDE guidance variable is used to tell CPEXpert whether you wish to exclude volumes from analysis. The EXCLUDE variable acts as a "switch" to control whether CPEXpert processes the DASGUIDE member searching for volumes to exclude from analysis. The point of having a "switch" variable is that some installations may wish to regularly exclude volumes from analysis, but periodically analyze all volumes.

If the EXCLUDE variable is **N**, CPEXpert will not exclude any volumes from analysis. If the EXCLUDE variable is **Y**, CPEXpert will process the DASGUIDE member to identify all volumes to be excluded.

Chapter 3.2: Defining volumes to exclude

You may exclude any number of volumes from analysis by coding the name of the volume in the "EXCLUDE VOLUME" portion of the DASGUIDE Module. You simply enter the VOLSER of each volume to be excluded after the "EXCLUDE" statement.

- The volumes to be excluded are entered one VOLSER per line or multiple VOLSERs per line.
- If you include multiple VOLSERs on a single line, you must separate them by blanks or commas.
- You can include comments on the line by placing "*" before the comment.

- You can use "generic" exclude logic by simply listing the first "n" characters of the VOLSERS you wish to exclude. For example, to exclude all VOLSERS beginning with ABC, simply specify ABC as the VOLSER to exclude.

CPEXpert will process the DASGUIDE Module, searching for the "EXCLUDE" statement. Any VOLSERS between the "EXCLUDE" statement and the "*/" statement will be placed into global SAS macro variables and may be excluded from analysis.

Exhibit 3-3 was extracted from the distributed USOURCE(DASGUIDE) module, and illustrates how to exclude volumes from analysis.

```
*****;  
* OPTIONAL GUIDANCE TO EXCLUDE VOLUMES FROM ANALYSIS;  
*****;  
* DO NOT REMOVE OR ALTER THE FOLLOWING SAS MACRO COMMENT LINE!;  
/* EXCLUDE THE FOLLOWING VOLSER FROM ANALYSIS:  
EXCLUDE  
MVS21AA * SAMPLE VOLSER DEFINITION  
CHKPT01 * SAMPLE VOLSER DEFINITION  
CHPPT02 * SAMPLE VOLSER DEFINITION  
SPOOL1 , SPOOL2 , SPOOL3  
* DO NOT REMOVE THE FOLLOWING MACRO COMMENT LINE!  
*/
```

EXCLUDING SPECIFIC VOLUMES FROM ANALYSIS

EXHIBIT 3-3

Chapter 4: Selecting specific volumes for analysis

You may wish to select specific volumes for analysis by the DASD Component, irrespective of whether these volumes are among the "worst" performing devices. For example, you may wish to examine only volumes containing specific application data sets.

Selecting volumes for analysis is optional; you do not have to select specific volumes for analysis. If you do not select specific volumes for analysis, CPEXpert will analyze every DASD VOLSER encountered in SMF Type 74 data after applying data selection criteria. If you **do** select specific volumes for analysis, CPEXpert will report on only those volumes (but CPEXpert will analyze the performance of the volumes within the context of the overall I/O configuration).

If you wish to select volumes from analysis, you must (1) use the **SELECT** guidance variable and (2) define the volumes to select.

Chapter 4.1: SELECT variable

The SELECT guidance variable is used to tell CPEXpert whether you wish to select specific volumes for analysis. The SELECT variable acts as a "switch" to control whether CPEXpert processes the DASGUIDE member searching for volumes to select for analysis. The point of having a "switch" variable is that some installations may wish to regularly analyze all volumes, but periodically select specific volumes for analysis.

If the SELECT variable is **N**, CPEXpert will not select specific volumes for analysis, but will analyze all volumes after applying data selection criteria. If the SELECT variable is **Y**, CPEXpert will process the DASGUIDE member to identify the volumes to be analyzed.

Chapter 4.2: Defining volumes to analyze

You may select any number of volumes for specific analysis by coding the name of the volume in the "SELECT VOLSER" portion of the DASGUIDE Module. You simply enter the VOLSER of each volume to be excluded after the "SELECT" statement.

- The volumes to be selected are entered one VOLSER per line or multiple VOLSERs per line.
- If you include multiple VOLSERs on a single line, you must separate them by blanks or commas.
- You can include comments on the line by placing "*" before the comment.

- You can use "generic" select logic by simply listing the first "n" characters of the VOLSERs you wish to select. For example, to select all VOLSERs beginning with DEF, simply specify DEF as the VOLSER to select.

CPEXpert will process the DASGUIDE Module, searching for the "SELECT" statement. Any VOLSERs between the "SELECT" statement and the "*/" statement will be placed into global SAS macro variables and will be the only volumes CPEXpert analyzes in detail (note that CPEXpert will continue to process data describing the entire I/O configuration so that the analysis is done within the context of your configuration).

Exhibit 3-4 was extracted from the distributed USOURCE(DASGUIDE) module, and illustrates how to select specific volumes for analysis.

You should note that CPEXpert still applies the "worst volume" philosophy when it reports the results from the analysis. That is, only volumes exceeding the average of those selected will be analyzed in detail. You might be confused when you examine the "next worst volumes" report, unless you keep this in mind. (Give us feedback if you don't like this approach. If it is unsatisfactory, we can change the design.)

```
*****;
* OPTIONAL GUIDANCE TO SELECT SPECIFIC VOLUMES FOR ANALYSIS;
*****;
* DO NOT REMOVE OR ALTER THE FOLLOWING SAS MACRO COMMENT LINE!;
/* SELECT THE FOLLOWING VOLSER FOR ANALYSIS:
SELECT
PAGE21      * SAMPLE VOLSER DEFINITION
PAGE22      * SAMPLE VOLSER DEFINITION
PAGE23      * SAMPLE VOLSER DEFINITION
PAGE24      * SAMPLE VOLSER DEFINITION
* DO NOT REMOVE THE FOLLOWING MACRO COMMENT LINE!
*/
```

SELECTING SPECIFIC VOLUMES FOR ANALYSIS

EXHIBIT 3-4

Chapter 5: Analyzing response objectives for critical data sets

As described in Section 1 (Chapter 5.3: Expanded analysis - Specific Data Sets) and Section 2 (Chapter 4: Defining Critical Data Sets), you can identify critical data sets to the DASD Component and specify a response objective for these data sets. The DASD Component can then analyze the performance of your DASD configuration, in one of two ways: (1) analyzing data set response based on information from TYPE42DS statistics or (2) analyzing data based on information in TYPE14/15 and using the CPEXpert modification to MXG or NeuMICS as the SMF Type30 records are processed.

Chapter 5.1: Analysis based on TYPE42DS

This method can provide comprehensive information, without requiring the modification to MXG. This method **is not applicable** to NeuMICS, since NeuMICS does not retain sufficient information related to SMF Type 42 (Data Set Statistics).

MXG creates TYPE42DS from the SMF Type 42 (Subtype 6 - data set I/O statistics) records created by SMS. The TYPE42DS file contains I/O access characteristics information, at the data set level. CPEXpert extracts data set information and data set response statistics from TYPE42DS, and compares these to the response objectives you have made in USOURCE(DASGUIDE) about critical data sets. When any data set response time exceeds the specified objective, the DASD Component selects the data set and the volume it resides on for detailed analysis.

The following steps are necessary to implement the critical data set analysis based on TYPE42DS records:

- Specify **%LET DASDSN = Y;** in USOURCE(DASGUIDE) to tell the DASD Component to analyze DASD performance based upon data set name.
- Specify **%LET TYPE42DS = Y;** in the USOURCE(GENGUIDE) CPEXpert general guidance module. This tells CPEXpert that you have SMF Type 42 (Data Set Statistics) available in your performance data base.
- Define critical data sets to the DASD Component and specify a response objective for each data set. This process is described in Section 2 (Chapter 6).

Chapter 5.2: Analysis based on TYPE14/15 and CPEXpert modification

This method is a bit more involved, but can be used regardless of whether data sets are managed by SMS, and this method can be used regardless of whether you use MXG or NeuMICS to create your performance data base.

With this method, the DASD Component analyzes SMF Type 14/15 records to extract data set names that correspond to the critical data sets that you have identified in USOURCE(DASGUIDE). The CPExpert modification to MXG or NeuMICS extracts DD information as SMF Type 30 records are processed. The DASD Component then correlates the data set information with the DD information, to determine whether critical data sets exceed the specified response objective. When any data set responses time exceeds the specified objective, the DASD Component selects the data set and the volume it resides on for detailed analysis.

The following steps are necessary to implement the critical data set analysis based on SMF Type 14/15, Type 30, and Type 74

- Install the modification to MXG or NeuMICS to collect data set-related information as described in Section 2 of this User Manual. The DASD Component cannot perform expanded analysis based on specific data sets unless it has information relating DASD use to specific data sets, service classes, jobs, and job steps. This information is extracted from the SMF Type 30(DD) records by the modifications described in Section 2 (Chapter 3 and Chapter 4).
- Define critical data sets to the DASD Component and specified a response objective for each data set. This process is described in Section 2 (Chapter 6).
- Specify **%LET DASDSN = Y;** in USOURCE(DASGUIDE) to tell the DASD Component to analyze DASD performance based upon data set name.
- Specify **%LET TYPE30DD = Y;** in the USOURCE(GENGUIDE) CPExpert general guidance module. This tells CPExpert that you have installed the modification to MXG or MICS to collect application-related information.
- Collect SMF Type 14/15 records and execute the DAS1415 module (described in Section 4 of this manual) to process the SMF Type 14/15 records. You must execute the DAS1415 module **before** you execute your normal daily update of your performance data base.

Chapter 6: Analyzing VSAM data sets

VSAM data set activity typically accounts for a large percent of I/O activity (more than 70% at some sites). Tuning of a few files or correcting common problems often can result in significantly improved performance (IBM benchmarks show up to 90% improvement resulting from some simple changes).

Analysis of VSAM data sets is not applicable to performance data bases created with NeuMICS, since the required SMF Type 42(Data Set statistics) and Type 64 (VSAM statistics) are not available.

With CPExpert Release 12.2, the DASD Component was enhanced to provide a rudimentary analysis of common VSAM problems. Additional analysis will be added in future enhancements to the DASD Component.

The DASD Component can optionally analyze VSAM data set performance problems or potential problems, only if the MXG TYPE64 file is available (and CPExpert is provided with the **%LET TYPE64=Y**; guidance variable). Additionally, most analysis depends on having the MXG TYPE42DS file available (and CPExpert is provided with the **%LET TYPE42DS=Y**; guidance variable).

Most sites have many VSAM data sets. Some of the VSAM data sets are open for a long time, but most are open for only a short interval. Some of the VSAM data sets have a significant amount of I/O activity, while other VSAM data sets have very low activity. These characteristics can result in very large reports produced by the DASD Component, since the DASD component can report on problems with each VSAM data set.

Options are provided with the VSAM analysis to summarize findings for VSAM data sets, to ignore VSAM data sets that are open for a short time, or to ignore VSAM data sets that have less than a specified amount of I/O activity. Additionally, each finding is based on guidance variables. These guidance variables can be altered so CPExpert will produce results only when more restrictive conditions are met. All these options can be (and should be) used to reduce the output from the DASD Component's to a moderate size.

This chapter describes the guidance variables that are available for analyzing performance problems of VSAM data sets. Exhibit 3-5 illustrates the VSAM analysis guidance variables contained in USOURCE(DASGUIDE) Component.

```

*****;
*   VSAM ANALYSIS GUIDANCE VARIABLES                               ;
*****;
%LET ANALVSAM  =BAD; * VSAM ANALYSIS OPTION                       ;
%LET CASPLITS  = 10; * EXCESSIVE CONTROL AREA SPLITS              ;
%LET DIRINDEX  =25; * PCT DIRECT ACCESS TO VSAM INDEX COMPONENT   ;
%LET EXTENTS   = 1; * EXCESSIVE SECONDARY EXTENTS                  ;
%LET LSRSEQ    = 50; * LSR SEQUENTIAL ACCESS DOMINATE              ;
%LET MXEXTENT  =225; * MAXIMUM EXTENTS INDICATING POTENTIAL PROBLEM ;
%LET NSRDIR    = 50; * NSR DIRECT ACCESS DOMINATE                  ;
%LET OPENTIME  =900; * MINIMUM VSAM DATA SET OPEN TIME FOR ANALYSIS ;
%LET PCTDIR    =80; * PERCENT DIRECT VSAM ACCESSES                 ;
%LET PCTSEQ    =80; * PERCENT SEQUENTIAL VSAM ACCESSES            ;
%LET VSAMEXCL  = NO; * OPTION: EXCLUDE SELECTED VSAM DATA SETS   ;
%LET VSAMIO    = 100; * VSAM I/O ACTIVITY SIGNIFICANT TO ANALYZE  ;
%LET VSAMSMRY  = NO; * VSAM DATA SET SUMMARIZED?                  ;
*****;

```

VSAM ANALYSIS GUIDANCE VARIABLES

EXHIBIT 3-5

Chapter 6.1: Controlling analysis of VSAM: ANALVSAM variable

The DASD Component optionally analyzes VSAM data sets to identify performance problems or potential problems if SMF Type 64 information is available, and if **%LET TYPE64=Y**; has been specified in **USOURCE(DASGUIDE)**⁷.

For flexibility, there are four options that control whether and how the VSAM data set analysis is performed. These options are specified by the ANALVSAM guidance variable:

- **%LET ANALVSAM = BAD; is specified in USOURCE(DASGUIDE)**. This is the default specification. When **%LET ANALVSAM=BAD**; is specified, the analysis is done **only** for VSAM data sets that reside on the devices that have been selected for analysis because the devices are the “worst performing” DASD devices in the configuration.

⁷ Most of the VSAM analysis also requires TYPE42DS information. Analysis that requires TYPE42DS information will be suppressed unless **%LET TYPE42DS=Y**; has been specified in either **USOURCE(GENGUIDE)** or **USOURCE(DASGUIDE)**.

- **%LET ANALVSAM = ALL; is specified in USOURCE(DASGUIDE).** This option directs CPEXpert to analyze all VSAM data sets in the configuration, regardless of which devices the VSAM data sets reside. CPEXpert will analyze VSAM performance for all volumes, after analyzing those volumes whose performance is worse than average. The basic data selection criteria (for example, date and time, system, etc.) will be applied before performing the VSAM analysis.
- **%LET ANALVSAM = ONLY; is specified in USOURCE(DASGUIDE).** This option directs CPEXpert to analyze **only** VSAM data sets, based on SMF Type 64 and (optionally) SMF Type 42 (Data Set Statistics) records. When the ONLY option is selected, CPEXpert will analyze VSAM performance without performing the basic DASD analysis of worst performing volumes. Exercising this option eliminates the processing associated creating a model of the I/O configuration and analyzing device performance, and CPEXpert will ignore devices with non-VSAM data sets. This option will be particularly useful when new applications are produced and these applications use VSAM data sets.
- **%LET ANALVSAM = NO; is specified in USOURCE(DASGUIDE).** This option directs CPEXpert to eliminate the VSAM analysis. This option might be selected after CPEXpert has analyzed VSAM data set performance, and you have made all of the changes that you intend to implement.

Chapter 6.2: Excessive Control Area splits: CASPLITS variable

CPEXpert examines the SMF Type 64 information contained in the MXG TYPE64 data set to identify VSAM data sets that have excessive Control Area splits. CPEXpert sums the ACCASPLT variable (the number of CA splits since the data set was created) and the CASPLITS variable (the number of CA splits with the current OPEN). CPEXpert compares this sum with the **CASPLITS** guidance variable in USOURCE(DASGUIDE). CPEXpert produces Rule DAS600 when the total number of CA splits exceeds the value specified by the **CASPLITS** guidance variable.

The default value for the **CASPLITS** guidance variable is 10, indicating that CPEXpert should produce Rule DAS600 when a VSAM data set experienced more than 10 CA splits. You should alter this guidance variable if you feel that Rule DAS600 is produced too often, or if you do not wish to take action when only 10 CA splits occur. For example, if you wish CPEXpert to produce Rule DAS600 only when more than 50 CA splits occur for any VSAM data set, specify **%LET CASPLITS=50;** in USOURCE(DASGUIDE).

Chapter 6.3: Percent direct access to VSAM index component: DIRINDEX variable

CPEXpert examines the SMF Type 64 information contained in MXG TYPE64 data set to identify VSAM KSDS or VRRDS data sets that have insufficient buffers assigned to the index component. CPEXpert uses the TYPE42DS information to compute the percent of direct accesses to the VSAM data set, using the following algorithm:

$$\text{Percent direct accesses} = \frac{S42AMDRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces *Rule DAS621* under the following conditions:

- The TYPE42DS S42DSBUF variable showed that NSR was used for KSDS or VRRDS VSAM data sets, **and**
- The percent of direct accesses for the index component was greater than **DIRINDEX** guidance variable in USOURCE(DASGUIDE), **and**
- The number of buffers (the MXG BUFDRNO variable) assigned to the index component was less than the number of index levels (the MXG ACCLEVEL variable) for the VSAM data set.

CPEXpert produces *Rule DAS622* under the following conditions:

- The TYPE42DS S42DSBUF variable showed that NSR was used for KSDS or VRRDS VSAM data sets, **and**
- The percent of direct accesses for the index component was greater than **DIRINDEX** guidance variable in USOURCE(DASGUIDE), **and**
- The STRNO specification in the ACB (the MXG ACBSTRNO) was greater than one (indicating that concurrent accesses had been specified for direct processing), **and**
- The number of buffers (the MXG BUFDRNO variable) assigned to the index component was less than the ACBSTRNO value, plus 1.

The default value for the DIRINDEX guidance variable is 25%, so CPEXpert will produce Rule DAS621 or DAS622 for NSR VSAM data sets when more than 25% of the accesses were direct for the index component and the number of buffers assigned to the index component was less than the number of index levels (for rule DAS621) or less than the

STRNO value in the ACB, plus 1 (for Rule DAS622). You can change the DIRINDEX guidance variable if you feel that Rule DAS621 or Rule DAS622 are produced too often, or if you do not wish to take action when to increase the number of buffers assigned to the VSAM data sets listed.

For example, if you wish CPEXpert to produce Rule DAS621 or Rule DAS622 only when more than 50% of the accesses for an index component were direct, specify **%LET DIRINDEX=50;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS621 and Rule 622 by specifying **%LET DIRINDEX=100;** in USOURCE(DASGUIDE), since the percent cannot exceed 100.

Chapter 6.4: Excessive EXTENTS were allocated: EXTENTS variable

CPEXpert examines the SMF Type 64 information contained in MXG TYPE64 data set to identify VSAM KSDS or VRRDS data sets that have excessive secondary allocations.

CPEXpert compares NREXTNTS variable (the number of secondary extents in the VSAM data set this OPEN) with the **EXTENTS** guidance variable in USOURCE(DASGUIDE). CPEXpert produces Rule DAS604 when the NREXTENT (the total number of extents) is greater than one, and the number of secondary extents allocated for this OPEN exceeds the value specified by the EXTENTS guidance variable.

The default value for the EXTENTS guidance variable is 0, indicating that CPEXpert should produce Rule DAS604, Rule DAS605, or Rule DAS606 (depending on data encountered) when any secondary extent was allocated for the VSAM data sets listed. You can change the EXTENTS guidance variable if you feel that Rule DAS604 is produced too often, or if you do not wish to take action when secondary extents are allocated. For example, if you wish CPEXpert to produce Rule DAS604 only when more than 5 secondary extents were allocated for VSAM data sets, specify **%LET EXTENTS=5;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS635 by specifying **%LET EXTENTS=255;** in USOURCE(DASGUIDE), since VSAM cannot allocate 255 extents.

Chapter 6.5: Specifying LSR sequential access domination: LSRSEQ variable

With Local Shared Resource (LSR), VSAM buffers normally are shared among VSAM data sets accessed by tasks in the same address space. Since LSR is oriented toward shared (and direct) access, there is an expectation that a record might be re-used. Consequently, buffer management algorithms retain buffers as long as possible, using a least-recently used (LRU) algorithm, after a record is processed from the LSR buffers. LSR is not suited for applications that use sequential or skip sequential as their primary access mode, because there is no read-ahead algorithm with LSR, and there is no inherent overlap of

I/O and CPU processing. Consequently, using LSR for sequential access processing could degrade rather than improve performance.

After applying the screening criteria specified for VSAM data sets, and extracting SMF Type 64 information for those VSAM data sets, CPEXpert examines SMF Type 42 (Data Set Statistics) information for the selected VSAM data sets. CPEXpert uses the TYPE42DS information to compute the percent of sequential accesses to the VSAM data set, using the following algorithm:

$$\text{Percent sequential accesses} = \frac{S42AMSRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces Rule DAS635 under the following conditions:

- The TYPE42DS S42DSBUF variable showed that LSR was used for KSDS or VRRDS VSAM data sets, and
- The percent of sequential accesses for the data component was greater than the value specified for the **LSRSEQ** guidance variable in USOURCE(DASGUIDE).

The default value for the LSRSEQ guidance variable is 75%, so CPEXpert will produce Rule DAS635 when LSR was specified as the buffering technique, and more than 75% of the accesses were sequential for the data component.

You can change the percent of sequential access that CPEXpert uses to determine whether to produce Rule DAS635 by altering the LSRSEQ guidance variable. For example, if you wish CPEXpert to produce Rule DAS635 only when more than 90% of the accesses to a LSR buffer pool were sequential, specify **%LET LSRSEQ=90;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS635 by specifying **%LET LSRSEQ=100;** in USOURCE(DASGUIDE), since the percent cannot exceed 100.

Chapter 6.6: Specifying maximum extents: MXEXTENT variable

CPEXpert examines the SMF Type 64 information contained in the MXG TYPE64 data set to identify VSAM data sets are in danger of reaching the maximum allowed number of extents. CPEXpert compares NREXTENT variable (the total number of extents in the VSAM data set) with the **MXEXTENT** guidance variable in USOURCE(DASGUIDE). CPEXpert produces Rule DAS607 when the NREXTENT is greater than the value specified by the MXEXTENT guidance variable **and** at least one extent was allocated during the

current OPEN of the data set (CPEXpert uses the NREXTNTS variable in TYPE64 for this decision).

The default value of the MXEXTENT guidance variable is 225, indicating that CPEXpert should produce Rule DAS607 when at least 225 extents have been allocated for a VSAM data set. Since the maximum allowable is 255, the default value provides a threshold at which CPEXpert provides notification that there is a potential problem.

You can alter this threshold by using the MXEXTENT guidance variable. For example, if you wish to be notified when the number of extents reach 200 (and at least one extent was required with the current OPEN), specify `%LET MXEXTENT=200;` in USOURCE(DASGUIDE).

Chapter 6.7: Specifying NSR direct access domination: NSRDIR variable

Non-shared resource (NSR) is the default VSAM buffering technique. VSAM data sets with NSR buffering can be accessed sequentially or direct (or both). However, NSR is suited for sequential processing because, if the data set access is sequential, the buffers are managed with a read-ahead algorithm. The read-ahead algorithm provides overlap of I/O and CPU processing and is efficient for sequential accesses. Since NSR is oriented toward sequential access, there is no expectation that a record will be re-used (as might exist with direct processing). Consequently, once a record is processed from the NSR buffers, the buffer is likely to be reclaimed for another record read from DASD.

NSR is not suited for direct processing, although NSR often **is** used for direct processing because it is easy to use and is the default buffering technique. Nonetheless, performance can be significantly improved if LSR is used for direct processing of VSAM data sets.

After applying the screening criteria specified for VSAM data sets, and extracting SMF Type 64 information for those VSAM data sets, CPEXpert examines SMF Type 42 (Data Set Statistics) information for the selected VSAM data sets. CPEXpert uses the TYPE42DS information to compute the percent of direct accesses to the VSAM data, using the following algorithm:

$$\text{Percent direct accesses} = \frac{S42AMDRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces Rule DAS625 under the following conditions:

- The TYPE42DS S42DSBUF variable showed that NSR was used for KSDS or VRRDS VSAM data sets, and
- The percent of direct accesses for the data component was greater than the value specified for the NSRDIR guidance variable in USOURCE(DASGUIDE).

The default value for the NSRDIR guidance variable is 75%, so CPExpert will produce Rule DAS625 when NSR was specified as the buffering technique, and more than 75% of the accesses were direct for the data component.

You can change the percent of sequential access that CPExpert uses to determine whether to produce Rule DAS635 by altering the NSRDIR guidance variable. For example, if you wish CPExpert to produce Rule DAS625 only when more than 90% of the accesses to a VSAM data set with NSR were sequential, specify **%LET NSRDIR=90;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS625 by specifying **%LET NSRDIR=100;** in USOURCE(DASGUIDE), since the percent cannot exceed 100.

Chapter 6.8: Minimum VSAM open time: OPENTIME variable

Many installations have a large number of VSAM data sets that are open for a relatively short time, and analyzing these data sets would unnecessarily clutter the reports produced by the DASD Component. The **OPENTIME** guidance variable allows installations to control which data sets are analyzed, based on the amount of elapsed time that the data sets are open.

The SMF variables reflecting when the data set was opened (SMF30TM and SMF30DT) were available with z/OS Version 1 Release 1 (z/OS V1R1). For versions of MVS prior to z/OS V1R1, the default value for the OPENTIME guidance variable is zero, and this value cannot be changed (since CPExpert would have nothing in SMF to compare against). Beginning with z/OS V1R1, the default value for the OPENTIME guidance variable is 300, indicating that CPExpert should ignore⁸ VSAM data sets that are open for less than 300 seconds (or 5 minutes).

If you are running MVS with z/OS V1R2 or subsequent release, and you wish CPExpert to ignore VSAM data sets that are open for different elapsed time, you can change the OPENTIME guidance variable to a different number of seconds. For example, if you wish CPExpert to analyze only those VSAM data sets that are open for more than 30 minutes, specify **%LET OPENTIME=1800;** in USOURCE(DASGUIDE).

⁸CPExpert does NOT ignore any VSAM data set based on open time if **%LET VSAMSMRY=Y;** has been specified in USOURCE(DASGUIDE), until all activity for each VSAM data set has been summarized. If **%LET VSAMSMRY=Y;** has been specified, the guidance is applied to the summarized data.

Chapter 6.9: Specifying percent direct access for control interval size: PCTDIR variable

A *Control Interval* is a continuous area of direct access storage that VSAM uses to store logical records. The size of Control Intervals can vary from one VSAM data set to another, but all the Control Intervals within the data portion of a particular data set must be the same length. The type of processing that is used should guide the choice of control interval size. When direct processing accounts for most of the accesses, a small data Control Interval is preferable. This is because only one record is retrieved at a time with direct processing. If a large Control Interval is specified, unnecessary I/O overhead is incurred reading the excess information. IBM suggests that a 4096 byte Control Interval normally would be appropriate for direct processing.

After applying the screening criteria specified for VSAM data sets, and extracting SMF Type 64 information for those VSAM data sets, CPEXpert examines SMF Type 42 (Data Set Statistics) information for the selected VSAM data sets. CPEXpert uses the TYPE42DS information to compute the percent of accesses to the data component of the VSAM data set that were direct, using the following algorithm:

$$\text{Percent sequential accesses} = \frac{S42AMSRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces Rule DAS611 when the percent of direct accesses for the data component was greater than the **PCTDIR** guidance variable in USOURCE(DASGUIDE), and the Control Interval size was greater than 4096 bytes. Additionally, CPEXpert verifies that the maximum logical records (maximum LRECL) is less than 50% of the Control Interval size. This verification is done to make sure that Rule DAS611 is not produced for VSAM data sets that have spanned records.

The default value for the PCTDIR guidance variable is 80%, so CPEXpert will produce Rule DAS611 when more than 80% of the accesses were direct for the data component, and the Control Interval size was more than 4096 bytes. You can alter this algorithm by changing the PCTDIR guidance variable in USOURCE(DASGUIDE). For example, if you wish CPEXpert to produce Rule DAS611 only when more than 90% of the accesses to the data component were direct, specify **%LET PCTDIR=90;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS611 by specifying **%LET PCTDIR=100;** in USOURCE(DASGUIDE), since the percent cannot exceed 100.

Chapter 6.10: Specifying percent sequential access for control interval size: PCTSEQ variable

A *Control Interval* is a continuous area of direct access storage that VSAM uses to store logical records. The size of Control Intervals can vary from one VSAM data set to another, but all the Control Intervals within the data portion of a particular data set must be the same length. The type of processing that is used should guide the choice of control interval size. When sequential processing accounts for most of the accesses, a large data Control Interval would normally be a good choice. This is because multiple records can be read into buffers and processed sequentially. For example, given a 16KB data buffer space, it is better to read two 8 KB Control Intervals with one I/O operation than four 4 KB Control Intervals with two I/O operations.

After applying the screening criteria specified for VSAM data sets, and extracting SMF Type 64 information for those VSAM data sets, CPEXpert examines SMF Type 42 (Data Set Statistics) information for the selected VSAM data sets. CPEXpert uses the TYPE42DS information to compute the percent of sequential accesses to the data component of the VSAM data set, using the following algorithm:

$$\text{Percent sequential accesses} = \frac{S42AMSRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces Rule DAS610 when the percent of sequential accesses for the data component was greater than the **PCTSEQ** guidance variable in USOURCE(DASGUIDE), and the Control Interval size was less than 8192 bytes.

The default value for the PCTSEQ guidance variable is 80%, so CPEXpert will produce Rule DAS610 when more than 80% of the accesses were sequential for the data component, and the Control Interval size was less than 8192 bytes. You can alter this algorithm by changing the PCTSEQ guidance variable in USOURCE(DASGUIDE). For example, if you wish CPEXpert to produce Rule DAS610 only when more than 90% of the accesses to the data component were sequential, specify **%LET PCTSEQ=90;** in USOURCE(DASGUIDE). You can completely suppress Rule DAS610 by specifying **%LET PCTSEQ=100;** in USOURCE(DASGUIDE), since the percent cannot exceed 100.

Chapter 6.11: Excluding VSAM data sets: VSAMEXCL variable

The **VSAMEXCL** guidance variable tells CPEXpert whether you wish to exclude certain VSAM data sets from analysis. The VSAMEXCL variable acts as a "switch" to

control whether CPEXpert processes the DASGUIDE member searching for VSAM data set names to exclude from analysis. The point of having a "switch" variable is that some installations may wish to regularly exclude VSAM data sets from analysis, but periodically analyze all VSAM data sets.

If the VSAMEXCL variable is **N**, CPEXpert will not exclude any VSAM data sets from analysis⁹. If the VSAMEXCL variable is **Y**, CPEXpert will process the DASGUIDE member to identify all VSAM data sets to be excluded.

Please refer to Chapter 7 for more information on excluding VSAM data sets from analysis.

Chapter 6.12: Specifying significant VSAM I/O activity: VSAMIO variable

Many installations have a large number of VSAM data sets that have little I/O activity, and analyzing these data sets would unnecessarily clutter the reports produced by the DASD Component. The **VSAMIO** guidance variable allows installations to control which data sets are analyzed, based on the amount of I/O activity of the data sets. CPEXpert examines the EXCP variable (number of EXCPs) in MXG TYPE64 records. CPEXpert ignores¹⁰ all records that have less than the amount of I/O activity specified in the VSAMIO guidance variable.

The default value for the VSAMIO guidance variable is 100, indicating that CPEXpert should ignore VSAM data sets that have less than 100 I/O operations. This low value was selected as a default so that you could appreciate potential problems with most VSAM data sets in your installation.

If you wish CPEXpert to ignore VSAM data sets that have a different I/O activity, you can change the VSAMIO guidance variable. For example, if you wish CPEXpert to analyze only those VSAM data sets that have more than 1000 EXCPs during each OPEN, specify **%LET VSAMIO=1000;** in USOURCE(DASGUIDE).

⁹Please note that the ANALVSAM guidance variable also acts to exclude VSAM data sets from analysis. If **%LET ANALVSAM=BAD;** is specified in USOURCE(DASGUIDE), only VSAM data sets residing on poorly performing volumes will be analyzed.

¹⁰CPEXpert does NOT ignore any VSAM data set based on I/O activity if **%LET VSAMSMRY=Y;** has been specified in USOURCE(DASGUIDE), until all activity for each VSAM data set has been summarized. If **%LET VSAMSMRY=Y;** has been specified, the guidance is applied to the summarized data.

Chapter 6.13: Summarizing VSAM activity: VSAMSMRY variable

Many installations have a large number of VSAM data sets that are open for a very short time, or have VSAM data sets that exhibit little I/O activity. Analyzing these data sets would unnecessarily clutter the reports produced by the DASD Component and produce a large amount of output¹¹. The **VSAMSMRY** guidance variable can be used to cause CPEXpert to summarize VSAM findings.

If **%LET VSAMSMRY=N**; is specified in USOURCE(DASGUIDE), CPEXpert will list every occasion in which a finding applies to a specific VSAM data set. This listing will show the time that the VSAM information was written to SMF, based on merging of MXG TYPE64 and TYPE42DS records. The information will show the JOB NAME using the VSAM data set, and other information associated with the rule. Since many VSAM data sets are OPENed and CLOSEd frequently (particularly with CLOSE=T), a very large number of lines of output can be produced. However, this level of output might be useful if VSAM problems are to be correlated with specific times of poor performance.

If **%LET VSAMSMRY=Y**; is specified in USOURCE(DASGUIDE), CPEXpert will summarize information for each VSAM data set and apply the VSAM analysis algorithms to the summarized information. The listing will show the time that the **last** VSAM information was written to SMF, based on merging of MXG TYPE64 and TYPE42DS records. The information will show the **last** JOB NAME using the VSAM data set, and a summary of other information associated with the rule. This level of output is useful for analyzing whether problems occur with VSAM data sets, over the entire measurement interval retained in the performance data base.

¹¹ Hundreds of pages of output might be created by CPEXpert if every VSAM data set were analyzed each time that it was OPENed.

Chapter 7: Excluding VSAM data sets from analysis

For a variety of reasons, you may wish to exclude certain VSAM data sets from analysis. The most common reason is that CPEXpert repeatedly identifies problems with particular VSAM data sets, but (1) you do not wish to make changes to correct the problems, (2) you are unable to make changes (because of application requirements or political realities), or (3) the DASD Component analysis is "flawed" because of data problems or data averaging. Whatever the reason, you may wish to exclude certain VSAM data sets from analysis.

Excluding VSAM data sets from analysis is optional; you do not have to exclude any VSAM data sets from analysis. If you do **not** exclude VSAM data sets from analysis, CPEXpert will analyze every VSAM data set encountered in SMF Type 64 after applying data selection criteria.

If you wish to exclude VSAM data sets from analysis, you must (1) use the **VSAMEXCL** guidance variable and (2) define the VSAM data sets to exclude.

Chapter 7.1: VSAMEXCL variable

The VSAMEXCL guidance variable is used to tell CPEXpert whether you wish to exclude VSAM data sets from analysis. The VSAMEXCL variable acts as a "switch" to control whether CPEXpert processes the DASGUIDE member searching for VSAM data sets to exclude from analysis. The point of having a "switch" variable is that some installations may wish to regularly exclude VSAM data sets from analysis, but periodically analyze all VSAM data sets.

If the VSAMEXCL variable is **N**, CPEXpert will not exclude any VSAM data sets from analysis. If the VSAMEXCL variable is **Y**, CPEXpert will process the DASGUIDE member to identify all VSAM data sets to be excluded.

Chapter 7.2: Defining VSAM data sets to exclude

You may exclude any number of VSAM data sets from analysis by coding the name of the VSAM data set in the "EXCLUDE SPECIFIC VSAM DATA SETS FROM ANALYSIS" portion of the DASGUIDE Module. You simply enter the VSAM data set name of each VSAM data set to be excluded after the `/* EXCLUDE VSAM DATA SETS FROM ANALYSIS` statement.

- The VSAM data sets to be excluded are entered one data set name per line or multiple data set names per line.

- If you include multiple data set names on a single line, you must separate them by blanks or commas.
- You can include comments on the line by placing " *" before the comment. NOTE that a blank character MUST precede the asterisk.
- You can use "generic" exclude logic by simply listing the first "n" characters of the VSAM data set names you wish to exclude, followed by an asterisk. For example, to exclude all VSAM data set names beginning with "D10.RMF.MONITOR3" simply specify "D10.RMF.MONITOR3*" as the data set name to exclude.

CPEXPERT will process the DASGUIDE Module, searching for the "/* EXCLUDE VSAM DATA SETS FROM ANALYSIS" statement. Any VSAM data sets between the "/* EXCLUDE VSAM DATA SETS FROM ANALYSIS" statement and the "*/" statement will be placed into global SAS macro variables and may be excluded from analysis.

Exhibit 3-6 was extracted from the distributed USOURCE(DASGUIDE) module, and illustrates how to exclude VSAM data sets from analysis.

```

*****;
* OPTIONAL GUIDANCE TO EXCLUDE SPECIFIC VSAM DATA SETS FROM ANALYSIS ;
*****;
* DO NOT REMOVE OR ALTER THE FOLLOWING SAS MACRO COMMENT LINE! ;
/* EXCLUDE VSAM DATA SETS FROM ANALYSIS
DSN=RLSADSW.VF05D.ITEMACT.DATA
DSN=RLSADSW.VF07D.ITEMACT
DSN=D10.RMF.MONITOR3*
* DO NOT REMOVE THE FOLLOWING MACRO COMMENT LINE!
*/
*****;

```

EXCLUDING VSAM DATA SETS FROM ANALYSIS

EXHIBIT 3-6

Section 3: Specifying Guidance Variables	3-1
Chapter 1: Data Selection and Presentation Variables	3-2
Chapter 1.1: DASDATES and DASTIMES variables	3-3
Chapter 1.2: DASDATEE and DASTIMEE variables	3-3
Chapter 1.3: DASDAT2S and DASTIM2S variables	3-3
Chapter 1.4: DASDAT2E and DASTIM2E variables	3-4
Chapter 1.5: SHIFT variable	3-4
Chapter 1.6: SYSTEM variable	3-4
Chapter 1.7: SYSTEMn variable(s)	3-5
Chapter 1.8: SYSPLEX variable	3-6
Chapter 1.9: CONFIG variable	3-6
Chapter 1.10: SAS Output Delivery System	3-7
Chapter 1.11: VERBOSE variable	3-7
Chapter 1.12: MAXRULES variable	3-8
Chapter 2: Analysis Control Variables	3-9
Chapter 2.1: Number of devices to analysis: ANALYZE variable	3-10
Chapter 2.2: Exclude devices with low activity: DASDEXCP variable	3-11
Chapter 2.3: Analyze using response objectives: DASDSN variable	3-12
Chapter 2.4: Produce only I/O configuration: EVALDASD variable	3-12
Chapter 2.5: Excluding volumes from analysis: EXCLUDE variable	3-13
Chapter 2.6: Specifying data sets to list: LIST42DS variable	3-13
Chapter 2.7: Perform "loved one" analysis: LOVED1 variable	3-13
Chapter 2.8: Exclude reporting low-activity data sets: MIN42PCT variable	3-15
Chapter 2.9: List data for all RMF intervals: LISTALL variable	3-16
Chapter 2.10: Minimum I/O rate to analyze: MINIORT variable	3-17
Chapter 2.11: Minimum I/O rate to analyze: MINIOWT variable	3-17
Chapter 2.12: Minimum I/O response to analyze: MINRESP variable	3-17
Chapter 2.13: Number of volumes to report: REPORT variable	3-17
Chapter 2.14: Selecting volumes to analyze: SELECT variable	3-18
Chapter 2.15: Analyze shared DASD Conflicts - SHARED variable	3-18
Chapter 2.16: SMF Type 30 modification installed - TYPE30DD variable	3-19
Chapter 2.17: SMF Type 42 (Data Set Statistics)- TYPE42DS variable	3-20
Chapter 3: Excluding volumes from analysis	3-21
Chapter 3.1: EXCLUDE variable	3-21
Chapter 3.2: Defining volumes to exclude	3-21
Chapter 4: Selecting specific volumes for analysis	3-23
Chapter 4.1: SELECT variable	3-23

Chapter 4.2: Defining volumes to analyze	3-23
Chapter 5: Analyzing response objectives for critical data sets	3-25
Chapter 5.1: Analysis based on TYPE42DS	3-25
Chapter 5.2: Analysis based on TYPE14/15 and CPExpert modification	3-25
Chapter 6: Analyzing VSAM data sets	3-27
Chapter 6.1: Controlling analysis of VSAM: ANALVSAM variable	3-28
Chapter 6.2: Excessive Control Area splits: CASPLITS variable	3-29
Chapter 6.3: Percent direct access to VSAM index component: DIRINDEX variable	3-30
Chapter 6.4: Excessive EXTENTS were allocated: EXTENTS variable	3-31
Chapter 6.5: Specifying LSR sequential access domination: LSRSEQ variable	3-31
Chapter 6.6: Specifying maximum extents: MXEXTENT variable	3-32
Chapter 6.7: Specifying NSR direct access domination: NSRDIR variable .	3-33
Chapter 6.8: Minimum VSAM open time: OPENTIME variable	3-34
Chapter 6.9: Specifying percent direct access for control interval size: PCTDIR variable	3-35
Chapter 6.10: Specifying percent sequential access for control interval size: PCTSEQ variable	3-36
Chapter 6.11: Excluding VSAM data sets: VSAMEXCL variable	3-36
Chapter 6.12: Specifying significant VSAM I/O activity: VSAMIO variable	3-37
Chapter 6.13: Summarizing VSAM activity: VSAMSMRY variable	3-38
Chapter 7: Excluding VSAM data sets from analysis	3-39
Chapter 7.1: VSAMEXCL variable	3-39
Chapter 7.2: Defining VSAM data sets to exclude	3-39