


CIMS Lab, Inc.



CIMS Capacity Planner

User Guide

Version 6.0

CIMS Lab Publication Number: CP-UG-600-01

Published 02/06/06

Information in this guide is subject to change without notice and does not constitute a commitment on the part of CIMS Lab, Inc. It is supplied on an "as is" basis without any warranty of any kind, either explicit or implied. Information may be changed or updated in this guide at any time.

Copyright Information

CIMS is ©copyright 1974–2006 by CIMS Lab, Inc. and its subsidiaries. This guide is ©copyright 1974–2006 by CIMS Lab, Inc. and its subsidiaries and may not be reproduced in whole or in part, by any means, without the written permission of CIMS Lab, Inc. and its subsidiaries.

Trademarks

The following are trademarks of International Business Machines Corporation in the United States, other countries, or both:

Candle	IMS
CICS	MVS
CL/SUPERSESSION	OMEGAMON
DB2	OS/390
IBM	VSAM

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Mailing Address

CIMS Lab, Inc.
3013 Douglas Blvd., Suite 120
Roseville, CA 95661-3842

Table of Contents



Preface

Philosophy	xv
Contacting the CIMS Lab	xvi
About This Guide	xvi
Conventions	xviii
Related Publications	xviii

1 • CIMS Capacity Planner Workload Analysis Subsystem

Introduction	1-3
Data Reduction	1-3
Processing Live SMF Files	1-3
Processing SMF History Files	1-3
SMF/RMF Records Required by CIMS Capacity Planner	1-4
Reports Produced During Data Reduction	1-5
Operating JCL for Data Reduction	1-5
Data Reduction Input Parameters	1-7
Ad-Hoc Report Request Parameters	1-12
Sample Jobstreams for Data Reduction	1-14
Special Considerations for MVS Release 5.1 (and later) Goal Mode	1-17
Workload Reports	1-18
Ad-Hoc Reports	1-18
Summary Analysis Report	1-20
JOBNAME: CPU Activity Report	1-21
Performance Reports	1-23
Workload Statistics Reports	1-24
Workload Exception Reports	1-27
Workload DASD Reports	1-28
Workload Graphs	1-29
Central/Expanded/Auxiliary Storage Activity Graphs (ESA or later only)	1-33
Operating JCL for Workload Batch Reports	1-34
ONLINE	1-34
INDEX	1-34
CPPRPARAM	1-34
CPPRERT	1-34
SYSPRINT	1-35
SYSNAP	1-35
EXCLUDE/INCLUDE	1-35
SYSIN	1-36

Input Parameter Statements for Batch Reports	1-38
Common Parameters	1-38
Report Statements	1-42
Exception Reports	1-46
Graph Statements	1-47
General Report Statements	1-51
Additional Report Options	1-52
Trends Analysis	1-56
Approach	1-56
Operating JCL for Trends Reports	1-57
Input Parameter Statements for Trends Reports	1-58
Sample JOBSTREAMS for Trends Reports	1-60
ISPF/PDF Workload Subsystem Interface	1-61
On-line Graphs	1-61
On-line Reports	1-63
Summary Analysis Report	1-66
Trends Analysis Report	1-66
Exception Analysis Reports	1-66
Ad Hoc Reports	1-67
Batch Reports	1-69
Batch GDDM Graphs	1-69
Presentation Graphics	1-74
Graph Production Procedures	1-74
Naming Conventions	1-75
Graph Periods	1-75
Graph Descriptions	1-75
Graphic Report Tables	1-76

2 • DASM SUBSYSTEM

DASM Subsystem Overview	2-3
DATA REDUCTION	2-3
SMF Data	2-3
VTOC Data	2-3
DCOLLECT Data	2-4
Reports Produced During Data Reduction	2-6
DASM REPORT DATA	2-11
PARMLIB Data	2-11
DASM REPORTS	2-12
DASD Space Allocation—By Device Type (SSA1DASM or SSA1DASH)	2-12
DASD Space Allocation—By Device Volume (SSA1DASM or SSA1DASH)	2-12
DASD Volume Map by CCHH (SSA1DASM)	2-13
Volume Map by DSNAME (SSA1DASM)	2-13
Volume Activity Report (SSA1RPT)	2-13
Data Set Activity Report (SSA1RPT)	2-13
Data Set Detail Report (SSA1RPT)	2-13
DASM Volume Summary Report (SSA1DASV)	2-14

DASM Selected VSAM Summary Report (SSA1VSAM)	2-15
DASM VTOC Selected DSN on Volume Report (SSA1DASM)	2-16
DASM VTOC by CCHH on Volume Report (SSA1DASM)	2-18
DASM VTOC FREESPACE on Volume Report (SSA1DASM)	2-19
DASM VTOC Detailed Volume Space Summary Report (SSA1DASM)	2-20
DASM VTOC Quick Volume Space Summary Report (SSA1DASM)	2-21
DASM X37 Candidate List (SSA1DASM)	2-22
DASM Re-blocking Candidate List (SSA1DASM)	2-24
DASM Archival Candidate List (SSA1DASM)	2-26
DASM Release Candidate List (SSA1DASM)	2-27
OPERATING JCL	2-29
DD Statements	2-29
Parameters	2-31
Additional Report Options	2-35
SAMPLE JOBSTREAMS	2-39
Scanning the DASD Farm	2-39
Saving the DASM Tables	2-40
Reporting from the Online Tables	2-41
ISPF/PDF DASM Subsystem Interface	2-42
Reports	2-42
Presentation Graphics	2-43
Graph Production Procedures	2-43
Naming Conventions	2-44
Presentation Graphics Distribution Disk	2-46
Graph Periods	2-47
Graph Descriptions	2-47
Graphic Report Tables	2-48
3 • CICS SUBSYSTEM	
Subsystem Overview	3-3
CICS Workload and Performance Reports	3-3
CICS Summary Analysis Report	3-3
CICS Response Performance Report	3-4
CICS Transaction Statistics Report	3-4
CICS Terminal Statistics Report	3-6
CICS Exception Analysis	3-6
CICS Ad Hoc Report	3-7
CICS Transaction Activity List	3-8
CICS Terminal Activity List	3-9
CICS System Overview	3-10
CICS Graphs	3-11
CICS Transaction Response Graph	3-11
CICS Terminal Response Graph	3-11
CICS Transaction Activity Graph	3-12
CICS Terminal Activity Graph	3-12
CICS CPU Activity Graph	3-12

CICS I/O Activity Graph	3-12
CICS Transaction Profile	3-12
CICS Trends Analysis Reports	3-13
Operating JCL	3-14
CICS Subsystem Installation JCL	3-14
Data Reduction JCL	3-14
Report Production JCL	3-27
Additional Report Options	3-43
ISPF/PDF CICS Subsystem Interface	3-47
On-line Graphs	3-47
On-line Reports	3-48
CICS Summary Analysis Report	3-49
CICS Trends Analysis Report	3-50
CICS Exception Analysis Report	3-50
CICS Ad Hoc Report	3-50
Batch Reports	3-51
Batch GDDM Graphs	3-51
Presentation Graphics Interface	3-56
Graph Production Procedures	3-56
Naming Conventions	3-57
Presentation Graphics Distribution Disk	3-59
Graph Periods	3-60
Graph Descriptions	3-60
Graphic Report Tables	3-61

4 • IDMS SUBSYSTEM

Subsystem Overview	4-3
IDMS Workload and Performance Reports	4-3
IDMS Summary Analysis Report	4-3
IDMS Response Performance Report	4-4
IDMS Transaction Statistics Report	4-4
IDMS Terminal Statistics Report	4-6
IDMS Exception Analysis Report	4-6
IDMS Ad Hoc Report	4-7
IDMS Graphs	4-8
IDMS Transaction Response Graph	4-8
IDMS Transaction Activity Graph	4-8
IDMS Terminal Activity Graph	4-8
IDMS CPU Activity Graph	4-9
IDMS D/B Activity Graph	4-9
IDMS I/O Activity Graph	4-9
IDMS Transaction Profile	4-9
IDMS Trends Analysis Report	4-10
Operating JCL	4-11
IDMS Subsystem Installation JCL	4-11
Data Reduction JCL	4-11

Report Production JCL	4-21
IDMS TRENDS ANALYSIS	4-34
Additional Report Options	4-36
ISPF/PDF IDMS Subsystem Interface	4-39
On-line Graphs	4-40
On-line Reports	4-40
IDMS Summary Analysis Report	4-41
IDMS Trends Analysis Report	4-42
IDMS Exception Analysis Report	4-42
IDMS Ad Hoc Report	4-42
Batch Reports	4-43
Batch GDDM Graphs	4-43
Presentation Graphics Introduction	4-48
Graph Production Procedures	4-48
Naming Conventions	4-49
Presentation Graphics Distribution Disk	4-51
Graph Periods	4-51
Graph Descriptions	4-52
GRAPHIC REPORT TABLES	4-52
5 • Network Subsystem	
Subsystem Overview	5-3
Network Workload and Performance Reports	5-3
Network Summary Analysis Report	5-3
Network Response Performance Report	5-4
Network Terminal Statistics Report	5-5
Network Exception Analysis	5-6
Network Physical Exception Analysis	5-6
Network Ad Hoc Report	5-6
Network Physical Line Statistics Report	5-7
Network Physical Line Activity Report	5-8
Cluster Controller Statistics Report	5-9
Cluster Controller Activity Report	5-10
Network Graphs	5-11
Network Transaction Response Graph	5-11
Network Terminal Activity Graph	5-12
Network Transaction Activity Graph	5-12
Network Inbound Message Activity Graph	5-12
Network Inbound Message Traffic Graph	5-13
Network Outbound Message Traffic Graph	5-13
Network Total Traffic Graph	5-13
Network Outbound Message Activity Graph	5-14
Network Physical Line Capacity Graph	5-14
Network Physical Line Inbound Traffic Graph	5-15
Network Physical Line Outbound Traffic Graph	5-15
Network Physical Line Retransmitted Traffic Graph	5-16

Network Physical Line Total Traffic Graph	5-17
Network Processor Activity Graph	5-17
Network Processor Slowdown Count Graph	5-18
NCP Free Buffer High Watermark Graph	5-18
NCP Free Buffer Low Watermark Graph	5-19
NCP Channel Hold Queue Length Graph	5-19
Cluster Controller Inbound PIU Graph	5-19
Network Cluster Controller Outbound PIU Graph	5-20
Cluster Controller Inbound Traffic Graph	5-20
Cluster Controller Outbound Traffic Graph	5-21
Cluster Controller Retransmitted Traffic Graph	5-22
Cluster Controller Total Traffic Graph	5-22
Network Terminal Profile	5-23
Network Trends Analysis Reports	5-24
Operating JCL	5-24
Network Subsystem Installation JCL	5-25
Data Reduction JCL	5-25
Report Production JCL	5-40
Additional Report Options	5-60
NETWORK ISPF/PDF INTERFACE	5-62
On-line Graphs	5-63
On-line Reports	5-63
Batch GDDM Graphs	5-66
Presentation Graphics	5-71
Graph Production Procedures	5-71
Naming Conventions	5-72
Presentation Graphics Distribution Disk	5-74
Graph Periods	5-75
Graph Descriptions	5-75
Graphic Report Tables	5-75

6 • IMS SUBSYSTEM

Subsystem Overview	6-2
IMS Workload and Performance Reports	6-2
IMS Summary Analysis Report	6-2
IMS Response Performance Report	6-3
IMS Transaction Statistics Report	6-3
IMS Terminal Statistics Report	6-4
IMS Exception Analysis	6-4
IMS Graphs	6-5
IMS Transaction Response Graph	6-5
IMS Transaction Activity Graph	6-5
IMS Terminal Activity Graph	6-5
IMS CPU Activity Graph	6-5
IMS I/O Activity Graph	6-6
IMS Trends Analysis Reports	6-6

Operating JCL	6-7
IMS Subsystem Installation JCL	6-7
Data Reduction	6-7
Report Production	6-16
Additional Report Options	6-28
ISPF/PDF IMS Subsystem Interface	6-30
On-line Graphs	6-31
On-line Reports	6-31
Batch GDDM Graphs	6-33
Presentation Graphics	6-37
Graph Production Procedures	6-37
Naming Conventions	6-38
Presentation Graphics Distribution Disk	6-40
Graph Periods	6-41
Graph Descriptions	6-41
Graphic Report Tables	6-41
7 • DB2 SUBSYSTEM	
Subsystem Overview	7-3
DB2 Workload and Performance Reports	7-3
DB2 Subsystem-Wide Reports	7-3
DB2 Connect Name Reports—by Plan	7-4
AUTH-ID Reports	7-9
DB2 Ad Hoc Reports	7-13
DB2 Trends Reports	7-16
Data Reduction	7-17
Processing SMF Records	7-18
Ad Hoc Request	7-19
Data Reduction JCL	7-20
CIMS Capacity Planner DB2 Data Reduction Parameters	7-21
DB2 Reporting	7-24
DB2 Reporting DD Statements	7-24
DB2 Reporting Parameters	7-26
Report Job Streams	7-36
DB2 ISPF/PDF INTERFACE	7-38
On-line Reports	7-38
Additional Report Options	7-39
GENERIC ELEMENT MASK	7-40
INCLUSIVE ELEMENT MASK	7-40
EXCLUSIVE ELEMENT MASK	7-41
Presentation Graphics	7-41
Graph Production Procedures	7-42
Naming Conventions	7-42
Presentation Graphics Distribution Disk	7-44
Graph Periods	7-45
Graph Descriptions	7-45

DB2 Presentation Graphics JCL	7-45
Graphic Report Tables	7-46
8 • GENERIC SUBSYSTEM	
Introduction	8-4
Components of the Generic Subsystem	8-5
Functional Description	8-8
Data Reduction—SSA1GENW	8-8
Report Generation	8-15
Generic Subsystem Operation	8-18
Data Reduction	8-18
User Table Specification	8-34
Ad Hoc Reports	8-38
Batch Reports	8-40
Generating Reports From CIMS Capacity Planner Maintained Tables	8-45
Generic Subsystem JCL (Data Reduction)	8-45
SYSPRINT	8-45
SYSUT1	8-45
CPPRERT	8-45
CPPRPARM	8-46
INDEX	8-46
ONLINE	8-46
SYSNAP	8-46
SYSUT3	8-46
SYSIN	8-46
INPUT PARAMETER STATEMENTS FOR DATA REDUCTION	8-47
SELECTED SYSTEM	8-47
SMFILE	8-47
BEGIN DATE	8-48
END DATE	8-48
BEGIN TIME	8-48
END TIME	8-49
GENERIC ELEMENT KEY	8-49
GEOGRAPHIC LOCATION	8-49
LOCAL HOLIDAYS	8-50
DEBUGON	8-50
SAMPLE JOB STREAMS FOR DATA REDUCTION	8-50
OPERATING JCL FOR BATCH REPORTS	8-54
SYSPRINT	8-54
CPPRERT	8-55
CPPRPARM	8-55
INDEX	8-55
ONLINE	8-55
EXCLUDE/INCLUDE	8-55
SYSNAP	8-56
SYSIN	8-56

INPUT PARAMETER STATEMENTS FOR BATCH REPORTS	8-56
SELECTED SYSTEM	8-56
GENERIC STATISTICS REPORT	8-57
GENERIC DAILY OVERVIEW	8-57
GENERIC ELEMENT ACTIVITY LIST	8-57
BEGIN DATE	8-58
END DATE	8-58
BEGIN TIME	8-58
END TIME	8-59
PRIME SHIFT FIRST HOUR	8-59
LATE SHIFT FIRST HOUR	8-59
SELECTED DAY	8-59
GEOGRAPHIC LOCATION	8-60
DEBUGON	8-60
LOCAL HOLIDAYS	8-60
TABLEDEF STATEMENTS FOR BATCH REPORTS	8-61
USER TABLE NUMBER	8-61
OUTPUT ELEMENT CELL	8-61
SAMPLE JOB STREAMS FOR BATCH REPORTS	8-62
PRODUCING BASIC PC GRAPHICS DATA POINT MEMBERS	8-63
OPERATING JCL FOR BASIC GRAPH DATA POINTS	8-63
SYSPRINT	8-63
CPPRERT	8-64
CPPRPARM	8-64
INDEX	8-64
ONLINE	8-64
HGDLIB	8-64
EXCLUDE/INCLUDE	8-64
SYSNAP	8-65
SYSIN	8-65
INPUT PARAMETER STATEMENTS FOR BASIC GRAPH DATA POINTS	8-66
FILTER	8-66
SELECTED SYSTEM	8-66
BEGIN DATE	8-67
END DATE	8-67
PRIME SHIFT FIRST HOUR	8-67
LATE SHIFT FIRST HOUR	8-68
SELECTED DAY	8-68
GEOGRAPHIC LOCATION	8-68
DEBUGON	8-68
LOCAL HOLIDAYS	8-69
GRAPH PERIOD	8-69
TABLEDEF Statements for Basic Graph Data Points	8-69
USER TABLE NUMBER	8-69
OUTPUT ELEMENT CELL	8-70
SAMPLE JOBSTREAMS FOR BASIC CELL GRAPH DATA POINTS	8-71

9 • Operations Analysis Subsystem

Subsystem Overview	9-2
Shift Turnover Report	9-2
Batch Window Analysis	9-4
What Is a Batch Window	9-4
Events Inside a Batch Window	9-5
Data Collection	9-6
Reports	9-6
Key Capacity Factors Analysis	9-9
Which Factors are Analyzed	9-10
Data Collection	9-10
Reports	9-11
Application Profile Analysis	9-13
Which Values are Maintained	9-13
Data Collection	9-14
Application Profile Report	9-14

10 • Model 204 Subsystem

Subsystem Overview	10-3
Model 204 Workload Analysis and Performance Reports	10-3
Model 204 Summary Analysis Report	10-3
Model 204 Response Performance Report	10-4
Model 204 Transaction Statistics Report	10-4
Model 204 Terminal Statistics Report	10-7
Model 204 Exception Analysis	10-7
Model 204 Transaction Activity List	10-8
Model 204 Ad Hoc Report	10-8
Model 204 Workload and Performance Graphs	10-9
Model 204 Transaction Response Graph	10-9
Model 204 Terminal Response Graph	10-9
Model 204 Transaction Activity Graph	10-10
Model 204 Terminal Activity Graph	10-10
Model 204 CPU Activity Graph	10-10
Model 204 I/O Activity Graph	10-10
Model 204 User ID Profile	10-11
Model 204 Trends Analysis Report	10-11
Operating JCL and Parameters	10-12
Model 204 Subsystem Installation JCL	10-12
Data Reduction JCL	10-12
Additional Report Options	10-32
ISPF/PDF Model 204 Subsystem Interface	10-34
On-line Graphs	10-35
On-line Reports	10-35
Model 204 Summary Analysis Report	10-36
Model 204 Trends Analysis Report	10-37
Model 204 Exception Analysis Report	10-37

Model 204 Ad Hoc Report	10-37
Batch Reports	10-38
Batch GDDM Graphs	10-38
Presentation Graphics Interface	10-43
Graph Production Procedures	10-44
Naming Conventions	10-44
Presentation Graphics Distribution Disk	10-46
Graph Periods	10-48
Graph Descriptions	10-48
Graphic Report Tables	10-48
11 • Tape Subsystem	
Tape Subsystem Overview	11-2
Tape Volume Activity Reports	11-2
Tape Volume Activity Graphs	11-3
Tape Drive Activity Reports	11-3
Tape Drive Activity Graphs	11-4
Input Data Supported	11-5
SMF Data (SSA1WKLD)	11-5
Tape Support for VTS	11-6
Tape Management System Control Information (SSA1TAPW)	11-6
The Reporting Component	11-7
Creating Tape Volume Ad Hoc Reports	11-7
Creating Tape Volume Activity Reports	11-8
Creating Tape Drive Activity Reports	11-8
The PC Graphing Component	11-9
Graph Periods	11-9
Graph Descriptions	11-9
Tape Volume Activity Graphs	11-9
Using Batch JCL from the Distribution CNTL Library	11-9
Using the Batch Panel in the ISPF/PDF Interface	11-9
Tape Drive Activity Graphs	11-10
Using Batch JCL from the Distribution CNTL Library	11-10
Using the Batch Panel in the ISPF/PDF Interface	11-10
Tape Ownership Graphs	11-10
Using Batch JCL from the Distribution CNTL Library	11-10
Using the Batch Panel in the ISPF/PDF Interface	11-10

12 • UNIX Reporting Subsystem

Overview	12-2
Input Data Supported	12-2
Using Standard UNIX Reports as Input	12-3
Identifying the UNIX Server	12-4
Consolidating Multiple SAR Files	12-5
Processing the Consolidated Output File	12-5
Using PARMLIB Member UNIXSRVR	12-6
Format of the UNIXSRVR Member Record	12-6
Creating or Editing a UNIXSRVR Member Record	12-7
The Reporting Component	12-8
Generating Reports and Graphs	12-8
UNIX Summary Report	12-9
UNIX Statistics Reports	12-10
UNIX Technical Graph Reports	12-12
UNIX Server Availability Reports	12-13
Tracking UNIX File System Capacity	12-14

Index



Preface

As companies continue to integrate computer technology into their business operations, it becomes increasingly important to properly administer the IT function, particularly with respect to performance and cost. And the best way to control costs is to plan for them.

CIMS is a comprehensive, flexible software solution that consolidates a wide variety of data for multiple operating systems into a single file that may be accessed from either the mainframe or a workstation. Simply put, CIMS is an essential component of an effective management system.

The CIMS Capacity Planner comprises several computer utilization and performance reporting subsystems. This product generates reports and graphs that deliver information necessary to evaluate the operation of a data center whose primary operating system is IBM® OS/390®.

Philosophy

CIMS is focused on meeting the financial, resource and capacity planner reporting requirements of Information Services Departments. CIMS has evolved with corporate IT management requirements. Focused commitment to client service and support sets CIMS apart from competing products. Our goal is to provide the best chargeback, resource reporting and capacity planning software in the world at the lowest possible cost to our customers.

The CIMS Lab strongly believes in and executes the concept of continuous product improvement. Customers have access to CIMS product development personnel to ensure that customer feedback and other critical issues are incorporated into the next release of the product.

Contacting the CIMS Lab

You can contact us with any questions or problems you have. Please use one of the methods below to contact us.

For product assistance or information, contact:

USA & Canada, toll free	(800) 283-4267
International	(916) 783-8525
FAX	(916) 783-2090
World Wide Web	www.cimslab.com

Our Mailing Address is:

CIMS Lab, Inc.
 3013 Douglas Blvd., Suite 120
 Roseville, CA 95661-3842

About This Guide

The *CIMS Chargeback CICS User Guide* consists of thirteen chapters, each describing a description of a specific subsystem upon which the CIMS Capacity Planner reports.

This guide assumes that the appropriate *CIMS Capacity Planner* components have been installed at your site. The instructions for installing the product are in the *CIMS Capacity Planner Installation and Getting Started Guide*.

Ch. No.	Chapter Name	Content Description
1	Workload Analysis Subsystem	Introduces you to the functions and features of the Workload Analysis Subsystem.
2	DASM Subsystem	Introduces you to the functions and features of the DASM Subsystem.
3	CICS Subsystem	Introduces you to the functions and features of the CICS® Subsystem.
4	IDMS Subsystem	Introduces you to the functions and features of the IDMS Subsystem.
5	Network Subsystem	Introduces you to the functions and features of the Network Subsystem.
6	IMS Subsystem	Introduces you to the functions and features of the IMS Subsystem.
7	DB2 Subsystem	Introduces you to the functions and features of the DB2® Subsystem.
8	Generic Subsystem	Introduces you to the functions and features of the Generic Subsystem.

Ch. No.	Chapter Name	Content Description
9	Operations Analysis Subsystem	Introduces you to the functions and features of the Accounting Subsystem.
10	Model 204 Subsystem	Introduces you to the functions and features of the Model 204 Subsystem.
11	Tape Subsystem	Introduces you to the functions and features of the Tape Subsystem.
12	UNIX Reporting Subsystem	Introduces you to the functions and features of the UNIX® Reporting Subsystem.

Conventions

Some or all of the following conventions appear in this guide:

Symbol or Type Style	Represents	Example
Bold	a new term	...called a source object .
<i>alternate color</i>	(online only) hotlinked cross-references to other sections in this guide; if you are viewing this guide online in PDF format, you can click the cross-reference to jump directly to its location	...see <i>Chapter 3, Data Migration</i> .
<i>Italic</i>	words that are emphasized	...the entry <i>after</i> the current entry...
	the titles of other documents	<i>CIMS Capacity Planner Reference Guide</i>
	syntax variables	COPY <i>filename</i>
Monospace	directories, file names, command names, computer code	&HIGHLVL.SRCLIB
	computer screen text, system responses, command line commands	Copy file? Y/N
Monospace bold	what a user types	...enter RUN APP.EXE in the Application field
< >	the name of a key on the keyboard	Press <Enter>.
▶	choosing a command from a cascading menu	File ▶ Import ▶ Object

Related Publications

As you use this guide, you might find it helpful to have these additional books available for reference:

- *CIMS Capacity Planner Installation and Getting Started Guide*
- *CIMS Capacity Planner Reference Guide*



CIMS Capacity Planner Workload Analysis Subsystem

Introduction	1-3
Data Reduction	1-3
Processing Live SMF Files	1-3
Processing SMF History Files	1-3
SMF/RMF Records Required by CIMS Capacity Planner	1-4
Reports Produced During Data Reduction	1-5
Operating JCL for Data Reduction	1-5
Data Reduction Input Parameters	1-7
Ad-Hoc Report Request Parameters	1-12
Sample Jobstreams for Data Reduction	1-14
Special Considerations for MVS Release 5.1 (and later) Goal Mode	1-17
Workload Reports	1-18
Ad-Hoc Reports	1-18
Summary Analysis Report	1-20
JOBNAME: CPU Activity Report	1-21
Performance Reports	1-23
Workload Statistics Reports	1-24
Workload Exception Reports	1-27
Workload DASD Reports	1-28
Workload Graphs	1-29
Central/Expanded/Auxiliary Storage Activity Graphs (ESA or later only)	1-33
Operating JCL for Workload Batch Reports	1-34
ONLINE	1-34
INDEX	1-34
CPPRPARM	1-34
CPPRERT	1-34
SYSPRINT	1-35

SYSNAP	1-35
EXCLUDE/INCLUDE	1-35
SYSIN	1-36
Input Parameter Statements for Batch Reports	1-38
Common Parameters	1-38
Report Statements	1-42
Exception Reports	1-46
Graph Statements	1-47
General Report Statements	1-51
Additional Report Options	1-52
Trends Analysis	1-56
Approach	1-56
Operating JCL for Trends Reports	1-57
Input Parameter Statements for Trends Reports	1-58
Sample JOBSTREAMS for Trends Reports	1-60
ISPF/PDF Workload Subsystem Interface	1-61
On-line Graphs	1-61
On-line Reports	1-63
Summary Analysis Report	1-66
Trends Analysis Report	1-66
Exception Analysis Reports	1-66
Ad Hoc Reports	1-67
Batch Reports	1-69
Batch GDDM Graphs	1-69
Presentation Graphics	1-74
Graph Production Procedures	1-74
Naming Conventions	1-75
Graph Periods	1-75
Graph Descriptions	1-75
Graphic Report Tables	1-76

Introduction

The CIMS Capacity Planner Workload Analysis Subsystem consists of the functions required to maintain the Workload portion of the Performance Data Base and to generate a wide variety of reports related to the workload and the performance of the your Data Center. The reports are useful to all levels of your management and technical staff in support of a number of critical activities ranging from Performance Evaluation to Performance Tuning and Capacity Planning.

Data Reduction

The Workload Analysis Data Reduction phase accepts SMF/RMF data from either the live cluster (SYS1.MANx) or archived SMF files. The SMF/RMF data is reduced into a set of Consolidated Tables (Performance Data Base) that are stored on DASD. These Consolidated Tables are later used by the various report generation programs to produce the reports described herein.

The Workload Subsystem Data Reduction process has been considerably shortened by changing the default to `DSNAME RECORDS=IGNORE`. This change has been shown to reduce the data reduction CPU time by up to 90%. With the advent of Type 42 records (SMS) the type 14, 15, 17, 18 and 64 records are no longer as useful as they were previously. It was also discovered that Table 002 consumed as much as 35% of the space used within the WKLD database.

Should you wish to continue the build Table 002, simply add a statement to the `SYSIN` which states:

```
DSNAME RECORDS=INCLUDE
```

Processing Live SMF Files

The live SMF files from `SYS1.MANx` are VSAM clusters that are used by the MVS operating system to store system measurement information. Periodically, these VSAM clusters become full and must be unloaded to a sequential (QSAM) data set, usually on magnetic tape. During the unloading process, the CIMS Capacity Planner program `SSA1WKLD` is used to reduce the workload data into the CIMS Capacity Planner Consolidated Tables mentioned above. This process does not replace or eliminate the need to unload the SMF/RMF data to magnetic tape. The CIMS Capacity Planner tables do not replicate raw SMF/RMF data.

Processing SMF History Files

The data reduction program can operate against historical SMF/RMF data (already dumped from `MANx`) collected on magnetic tape or other media. This procedure does not yield up-to-the-minute data for the Consolidated Tables, but it does provide a way to accumulate historical data in the CIMS Capacity Planner tables that was collected prior to the time CIMS Capacity Planner was installed. By reducing the historical SMF/RMF data, the CIMS Capacity Planner Consolidated Tables will reflect the historical trends prior to the period when CIMS Capacity Planner was installed and immediately enhance the value of CIMS Capacity Planner to your installation.

SMF/RMF Records Required by CIMS Capacity Planner

The Data Reduction Program presumes that RMF is installed and that the following SMF and RMF records are being collected and written to the SMF Log:

SMF Records

0	(IPL)
6	(JESx Printer Record)
14	(Non-VSAM Input Close Record)
15	(Non-VSAM Output Close Record)
17	(Non-VSAM Data Set Delete Record)
18	(Non-VSAM Rename Record)
21	(Tape Error Statistics by Volume)
30	(Job/Step Statistics/Interval Accounting Record)
32	(TSO/E Command Statistics Record)
64	(VSAM Close Statistics Record)

RMF Records

70	(CPU Utilization Record)
71	(Paging Activity Record)
72	(Workload Record)
73	(Channel Activity Record)
74	(Device Activity Record)
75	(Page/Swap Data Set Activity Record)
77	(Enqueue Conflict Record)

In order to ensure that the utilization data is reported within the correct timeframe, it is advisable to produce the SMF Interval Accounting records at 30 minute intervals.

In order to maintain reasonably current levels of data in the CPPR Consolidated Tables, we recommend that the data reduction programs be run on a daily basis while the SMF / RMF files are being unloaded to magnetic tape. It is a relatively simple task to include the data reduction program into the regularly run jobstream that is used for this purpose.

Reports Produced During Data Reduction

All the reports specified under the Workload Reports section are produced by the Data Reduction program. In addition, several Ad-Hoc reports can be generated only by the Data Reduction program while the SMF/RMF data is being processed. The Ad-Hoc reports include:

- DSNNAME
- JOBNAME
- PGMNAME
- TSO USERID
- Ad-Hoc reports are illustrated in the report descriptions. See [Workload Reports on page 1-18](#).

Operating JCL for Data Reduction

The CIMS Capacity Planner JCL has been designed to conform to the JCL standards for IBM Utility Programs wherever possible. At this time, two forms of SMF/RMF input are acceptable to the Data Reduction program, that can come from any one of the following sources:

- Historical (QSAM) SMF/RMF Data
- Live (VSAM) SMF/RMF Data
- Active (the active SMF cluster)

Each DD Statement required to execute the Workload Data Reduction Program is described below:

ONLINE

The **ONLINE** DD Statement specifies the **DSNAME** and **DISPosition** of the CIMS Capacity Planner Performance Data Base containing the Consolidated Tables.

INDEX

The **INDEX** DD Statement specifies the **DSNAME** and the **DISPosition** of the data set containing the index to the CIMS Capacity Planner Performance Data Base.

CPPRERT

The **CPPRERT** DD Statement is used to describe the CIMS Capacity Planner Element Registration file containing selected data related to CPU type, CPU speed, SMFID's, etc.

CPPRPARM

The **CPPRPARM** DD Statement is used to specify the **CPPR** Parameter Library. The program will determine which members are referenced within the parameter library.

SYSUT3

The `SYSUT3` DD Statement describes a BSAM data Set used to hold a working copy of the `ONLINE` file during the data reduction processing. It must be larger than the `ONLINE` data set. Beginning with the CIMS Capacity Planner Version 5.1, there was a change in the way that the online file is processed during Data Reduction. Rather than move the entire Online file to `SYSUT3` to do updates, the CIMS Capacity Planner now copies only the tables to be updated to `SYSUT3` and then moves them back to the online file at the successful conclusion of the Data Reduction process. If you run out of space when moving the updated tables from `SYSUT3` back to the Online file (B37), the Index file reflects the pre-update state of the Online file. If you run out of space when rewriting the Index file, the `SSA1INDX` utility can be used to rebuild the Index from the Online file.

SYSUT1

The `SYSUT1` DD Statement describes the Historical SMF/RMF input file. The only parameters required are the data set name and the disposition.

SYSPRINT

The `SYSPRINT` DD Statement describes the report output file. Unless otherwise specified, the DCB characteristics are:

```
RECFM=FBA, LRRCL=133, BLKSIZE=133
```

SYSNAP

The `SYSNAP` DD Statement specifies a print file that is used for printing diagnostic hexadecimal snap dumps when operating in the debugging mode (See `DEBUGON` in the parameter descriptions).

SYSIN

The `SYSIN` DD Statement specifies the CIMS Capacity Planner parameter file that is used to specify the parameters that control many significant features of the data reduction and reporting functions of CIMS Capacity Planner.

Data Reduction Input Parameters

The CIMS Capacity Planner provides many parameters to give you a significant degree of control and to simplify its use.

The parameters that apply to the Workload Data Reduction phase are described below. Parameters that control the reporting functions are described under *Input Parameter Statements for Batch Reports on page 1-38*.

The following descriptions of the data reduction parameters specify the functions controlled by each parameter, the options or data that must be provided, and the default values of each parameter.

The parameters need not be specified whenever the default settings are acceptable.

SELECTED SYSTEM

This parameter is used to identify the SID of the System for which SMF/RMF data is to be analyzed. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Up to five SIDs can be specified in the current version of the Workload Analysis Program. For example, if a tape contains SMF/RMF data for a variety of systems, this parameter can be used to specify the systems for which data are to be reduced in a single run.

If multiple SIDs are indicated, all parameters pertain to each of the selected systems. That is, all other parameters (beginning and ending times and dates, language, output filtering) apply equally to all systems being analyzed.

Note • For storage requirements, figure that each SID specified will require approximately 800K of Virtual Storage. For example, if five systems are to be analyzed in a single run, the REGION size for the run should be 4096K.

Options

This parameter has several options. If the program is to produce a single report for a single system (SID), specify:

- The SID itself (e.g, `SELECTED SYSTEM=MVS1`)
- An '*', indicating the SID for the system on which the program is being executed (e.g, `SELECTED SYSTEM=*`)
- The characters '####', indicating the first SID that is encountered (e.g, `SELECTED SYSTEM=####`)

If multiple (up to five) Systems are to be analyzed in a single run, specify the parameter as follows:

- Each specific SID, separated by commas (e.g, `SELECTED SYSTEM=MVS1,XAS1,MVS2`)
- The parameter "INCLUDE(n)", where the "n" is the number of unique SIDs present on the tape (up to five). The Workload Analysis Program will dynamically determine which SIDs to select (e.g, `SELECTED SYSTEM=INCLUDE(3)`)

Default Option

`SELECTED SYSTEM=*`

BEGIN DATE

This parameter is used to specify the earliest date for which an SMF/RMF record is to be selected for analysis.

Specify Date

Specify a beginning date in one of the following forms:

- YYYYDDD (Standard Julian format)
- MM/DD/YYYY (USA Standard Gregorian format)
- DD.MM.YYYY (European Standard Gregorian format)
- *, *-n (relative day format where *=today)

Note • All dates must be specified in one of the above formats.

Default Date

BEGIN DATE=00001

END DATE

This parameter is used to specify the last date for which an SMF/RMF record is to be selected for analysis.

Specify Date

Specify an ending date in one of the formats described above.

Default Date

END DATE=2099365

REPORT LANGUAGE

This parameter describes the language that is to be used for the narrative sections of the report.

Options

ENGLISH, DEUTSCH or any combination of the languages, separated by commas.

If multiple languages and multiple SIDs are specified, the order of the reports are as follows:

- The first selected system (SID) in the primary language
- The first selected system (SID) in the secondary and succeeding languages (if more than two are specified)
- The second selected system (SID) in the primary language
- The second selected system (SID) in the secondary and succeeding languages (if more than two are specified)
- The third selected system (SID) in the primary language (and so forth)

Default Option

REPORT LANGUAGE=ENGLISH

BEGIN TIME

This parameter is used to specify a time of the day before which no SMF/RMF data is to be selected for analysis. For example, if specific shift-related activity were being analyzed, the beginning of the shift would be specified by `BEGIN TIME` and the end of the shift would be specified by `END TIME`.

Specify Time

The beginning time, in the form `HH.MM.SS`

Default Time

`BEGIN TIME=00.00.00`

END TIME

This parameter is used to specify a time of the day after which no SMF/RMF data is to be selected for analysis. See the discussion of `BEGIN TIME` above.

Specify Time

The ending time, in the form `HH.MM.SS`

Default Time

`END TIME=24.00.00`

SMFILE

This parameter describes the format of the SMF/RMF input data. If `SYSUT1` is used as input, this parameter should specify "HISTORY". If the Active cluster is to be used as input, specify "ACTIVE".

Options

Either `HISTORY` or `ACTIVE` (e.g, `SMFILE=ACTIVE`)

Default Option

`SMFILE=HISTORY`

DUMP SMF STATISTICS

This parameter determines whether or not the `WKLD` program will produce a matrix showing the counts for each SMF record type encountered during the run.

Options

`YES` or `NO`. (e.g, `DUMP SMF STATISTICS=YES`)

Default Option

`DUMP SMF STATISTICS=NO`

SMF RECORD SET

This Workload Subsystem option, when specified with the Workload Data Reduction program (SSA1WKLD) is used to spool a set of SMF records to an auxiliary file for later processing. Previous versions of CIMS Capacity Planner allowed you to specify SYSUT2=YES along with a FILTER statement listing up to 16 individual SMF record types. This combination caused CIMS Capacity Planner to send the specified set of SMF records to the SYSUT2 output file for later processing.

The new SMF RECORD SET statement is an improvement over the SYSUT2 approach. For each occurrence of a SMF RECORD SET statement, a file is designated (SYSUT0xx, where the "xx" is a sequential number from 00 to 09) that is used to contain the set of records indicated by the list that follows the statement.

Note • A maximum of 10 such statements are permitted.

For example, suppose that the SYSIN for SSA1WKLD contained the following five statements in this order:

```
SMF RECORD SET=110
SMF RECORD SET=100,101
SMF RECORD SET=28,39
SMF RECORD SET=30
SMF RECORD SET=240
```

This combination would result in five output files being created, whose DDNAMEs are as follows:

- SYSUT000 would hold CICS Monitoring Facility Records (110)
- SYSUT001 would hold DB2 Records (100,101)
- SYSUT002 would contain Network Records (28,39)
- SYSUT003 would contain Accounting records (Type 30)
- SYSUT004 would contain User records (240)

SCANONLY

This parameter is used to cause the Workload data reduction program to read the SMF/RMF input data without performing any file maintenance against the CIMS Capacity Planner Performance Data Base. SCANONLY is generally used in combination with the SMFRECORDSET= or FILTER= parameters to extract SMF/RMF data to be passed to other CIMS Capacity Planner data reduction programs such as CICS, IDMS, IMS, etc.

DEBUGON

This parameter causes trace information to be written out to diagnose problems with CIMS Capacity Planner. This parameter is generally used only when working with CIMS Lab personnel in diagnosing problems.

CPPRPASS

This parameter is used to supply the encoded CPU serial number of the CPUs on which CIMS Capacity Planner operates. The password is supplied by the CIMS Lab at the time the CIMS Capacity Planner software is distributed. The CPPRPASS parameter can be specified in the GLOBAL member of the CIMS Capacity Planner Parmlib (recommended method) or in all report steps as a //SYSIN Parameter.

Ad-Hoc Report Request Parameters

Four Ad-Hoc reports can be produced during the Data Reduction phase. The following paragraphs describe the parameters necessary to request these Ad-Hoc reports.

JOBNAME

This parameter is used to select all references to a given job during the measured period. A report is produced showing the date and time the job entered the system, the date and time the job completed execution, and the amount of CPU time used and the total number of DASD EXCPs issued during the job.

Generic JOBNAME specification is supported. For example, if you wanted to see all occurrences of jobs whose names begin with S01M, you would specify:

```
JOBNAME=S01M*
```

Specify Jobname

JOBNAME, up to 8 characters (e.g, JOBNAME=PAYROLL3)

Default Jobname

(no default)

PGMNAME

This parameter is used to select all references to a given Program during the measured period. A report is produced showing the date and time the Program entered the system, the date and time the Program completed execution, and the amount of CPU time used and the total number of DASD EXCPs issued during the step.

Generic PGMNAME specification is supported. For example, if you wanted to see all occurrences of Programs beginning with S01M and followed by any other qualifier, you would specify:

```
PGMNAME=S01M*
```

It is also possible to select all programs that executed in a specific Performance Group (PGN). In this case, specify the decimal PGN instead of a Program Name. The PGN must be three numeric positions (use leading zeroes as necessary). For example, to produce an Ad Hoc report showing all programs that executed in Performance Group number 13, you could specify:

```
PGMNAME=013
```

Specify Program Name

PGMNAME, up to 8 characters (e.g, PGMNAME=PAYR010)

Default Program Name

(no default)

USERID

This parameter is used to select all references to a given TSO USERID during the measured period. A report is produced showing the date and time the TSO User logged on, the date and time the TSO User logged off, and the amount of CPU time used and the total number of transactions (commands) executed during the TSO session.

Generic USERID specification is supported. For example, if you wanted to see all occurrences of USERIDs beginning with S01M and followed by any other qualifier, you would specify:

```
USERID=S01M*
```

Specify TSO USERID

TSO USERID, up to 8 characters (e.g, USERID=VSSA1)

Default TSO USERID

(no default)

DSNAME

The DSNAME parameter is used to limit the selection of records to those records containing a specified data set name or a generic data set name. DSNAME= SYS1 . PROCLIB would, for example, limit the selection to only those records containing the data set SYS1 . PROCLIB while DSNAME=SYS1* would result in the selection of all records pertaining to data sets beginning with the characters "SYS1" providing any other selection criteria are satisfied.

Specify Data Set Name

Data set name, up to 44 characters (e.g, DSNAME=SYS3 . PAYROLL)

or

Generic data set name, dsroot* (e.g, DSNAME=SYS3 . PAYR*)

Default Data Set Name

(no default)

Sample Jobstreams for Data Reduction

This section contains several sample sets of JCL and CIMS Capacity Planner parameters for a variety of Workload Data Reduction processing options.

Gathering Data from an Active SMF Cluster

The following is an example of a jobstream that will process SMF/RMF data from the active SMF cluster into the Consolidated Table File (ONLINE) and, at the same time, will produce a Summary report. Job, Program and TSO Statistics reports will also be produced.

The reports are produced in English. The program will only process records that have been produced by the IP01 System. Input is furnished through the SYSMANO DD Statement. All records from the beginning of October 2005 through the end of October 2005 are processed.

```
//SMFJOB JOB (aaaaaaaa)
//STA EXEC PGM=SSAIWKLD,REGION=4096K
//SYSMANO DD DISP=SHR,DSN=SYS1.MANO
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//INDEX DD DSN=CPPR.INDEX,DISP=SHR
//ONLINE DD DSN=CPPR.ONLINE,DISP=SHR
//SYSUT3 DD DISP=(,DELETE),UNIT=SYSDA,SPACE=(CYL,(30,5))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
SMFILE=ACTIVE
BEGIN DATE=10/01/2005
END DATE=10/31/2005
REPORT LANGUAGE=ENGLISH
SUMMARY REPORT=YES
JOB STATISTICS REPORT=YES
PROGRAM STATISTICS REPORT=YES
TSO USER STATISTICS=YES
TSO COMMAND STATISTICS REPORT=YES
```


Gathering Data from History

The following is an example of a jobstream that will process historical (magnetic tape resident) SMF/RMF data into the Consolidated Table File (ONLINE) and, at the same time, will produce a Summary Report for all Tuesdays in October. Job, Program and TSO Statistics reports will also be produced.

The reports are produced twice: the first time in English and the second time in Deutsche. At the same time, the program will write RMF type 70 (CPU Activity) and RMF type 71 (Paging Activity) records to the SYSUT2 dataset. The program will process only records that have been produced by the IP01 System. Input is furnished through the SYSUT1 DD Statement.

```
//SMFJOB JOB (aaaaaaaa)
//STEP1 EXEC PGM=SSA1WKLD
//SYSUT1 DD DISP=SHR,DSN=HISTORY.SAVE.SMF(-2)
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//INDEX DD DSN=CPPR.INDEX,DISP=SHR
//ONLINE DD DSN=CPPR.ONLINE,DISP=SHR
//SYSUT2 DD DSN=TEMP.SAVE.SMF,DISP=SHR
//SYSUT3 DD DISP=(,DELETE),UNIT=SYSDA,SPACE=(CYL,(30,5))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
SELECTED DAY=TUESDAY
BEGIN DATE=2005032
END DATE=2005061
SYSUT2=YES
FILTER=70,71
REPORT LANGUAGE=ENGLISH,GERMAN
SUMMARY REPORT=YES
JOB STATISTICS REPORT=YES
PROGRAM STATISTICS REPORT=YES
TSO USER STATISTICS=YES
TSO COMMAND STATISTICS REPORT=YES
```

Producing a Selected DSNAME Report

The following is an example of a jobstream that will scan through the historical (magnetic tape resident) SMF/RMF data and will produce a Selected DSNAME Report for all Datasets whose name begins with the four characters "SYS1". The date range for the report will include the 22nd and the 23rd of May, 2005. No ONLINE database tables will be updated.

```
//SMFJOB JOB (aaaaaaaa)
//STEP1 EXEC PGM=SSA1WKLD
//SYSUT1 DD DISP=SHR,DSN=HISTORY.SAVE.SMF(-2)
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//INDEX DD DUMMY
//ONLINE DD DUMMY
//SYSUT3 DD DISP(,DELETE),UNIT=SYSDA,SPACE=(TRK,(1,1))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
BEGIN DATE=05/22/2005
END DATE=05/23/2005
SCANONLY
SUMMARY REPORT=NO
DSNAME=SYS1*
```

Producing a Selected TSO USERID Report

The following is an example of a jobstream that will scan through the historical (magnetic tape resident) SMF/RMF data and will produce a Selected TSO USERID Report for all TSO Users whose USERID begins with the four characters "S01M". The date range for the report will include the 22nd and the 23rd of May, 2005.

```
//SMFJOB JOB (aaaaaaaa)
//STEP1 EXEC PGM=SSA1WKLD
//SYSUT1 DD DISP=SHR,DSN=HISTORY.SAVE.SMF(-2)
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//INDEX DD DUMMY
//ONLINE DD DUMMY
//SYSUT3 DD DISP(,DELETE),UNIT=SYSDA,SPACE=(TRK,(1,1))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
BEGIN DATE=05/22/2005
END DATE=05/23/2005
SUMMARY REPORT=NO
USERID=S01M*
```

Special Considerations for MVS Release 5.1 (and later) Goal Mode

MVS Release 5.1 Compatibility Mode requires no special considerations because, as the name implies, there is very little difference from the previous release. When running in Goal Mode, however, several things change. For one thing, there are no longer any Performance Group Numbers. Instead the workload is managed by Service Classes.

In order to maintain compatibility with earlier releases, CIMS Capacity Planner has incorporated the `&sid.SVCL PARMLIB` member that associates Service Classes with pseudo-PGNs. A pseudo-PGN is simply a Performance Group Number that can be referenced in the local member of `PARMLIB` in order to classify the Workload by group. A new table, Table 070, is built that contains Service Units by Service Class Name, so you can evaluate the Workload by Service Class as well.

The `&sid.SVCL PARMLIB` member

The new `PARMLIB` member follows CIMS Capacity Planner conventions for `PARMLIB` members. Any line beginning with an `"*"` is considered to be a comment. Up to 64 categories of information (pseudo-PGNs) can be specified. Each pseudo-PGN group begins with a numbered line where the first two columns are numeric, the third column is a period, the fourth column is a blank and the fifth through eighth columns contain a three digit numeric value that represents the pseudo-PGN.

The second and succeeding lines in the group must begin with a blank in column 1 and contain one or more Service Class Names, separated by commas. The last name in a group must not be followed by a comma. Wildcards are allowed. For example, if you want to include all Service Class Names beginning with `BAT`, specify `BAT*` as the Service Class Name within a group definition.

For a sample of the `&sid.SVCL PARMLIB` member, please refer to the `CPPRSVCL` member in the distributed `CPPR.PARMLIB`.

Workload Reports

The Workload subsystem produces a variety of reports. Several of the reports (Ad-Hoc) are only available from the Data Reduction program—SSA1WKLD, while the remainder of the reports are produced by the report generation program SSA1RPT. The SSA1WKLD program produces its reports from the SMF/RMF input data and reports upon portions of the input data that are not retained in the CIMS Capacity Planner Performance Data Base. The SSA1RPT program produces its reports from the data contained in the Performance Data Base. The reports produced by the Report Generation program can be run at any time. Report generation generally takes only a few seconds.

The first three reports discussed are produced by the Data Reduction program. The remainder are produced by the Report Generation program - SSA1RPT.

Ad-Hoc Reports

The CIMS Capacity Planner generates several ad-hoc reports which enable reporting at the detailed interval level, against jobs, programs, or TSO usersids.

JOBNAME Report (Ad-Hoc)

The JOBNAME REPORT is requested by including a JOBNAME= parameter as a SYSIN parameter in the SSA1WKLD data reduction Jobstream. The JOBNAME parameter consists of up to eight characters. A generic JOBNAME is supported by including an asterisk after the root of the JOBNAME. For example, you can select all jobs beginning with SOM1 by specifying JOBNAME=SOM1*.

The selection of jobs to be included in this report is limited by the BEGIN DATE, END DATE, BEGIN TIME, and END TIME parameters. The JOBNAME report includes the following data:

- Date and time the job entered the system
- Date and time the job finished execution
- CPU time used by the job
- Total number of DASD accesses issued

Two JOBNAME Detail Ad Hoc reports are also available. The DDNAME Detail report shows, in addition to the information contained in the JOBNAME Ad Hoc Report, specific information pertaining to the DDNAME of the specified job. The S/U Detail report shows, in addition to the information contained in the JOBNAME Ad Hoc Report, specific information pertaining to the Service Units of the specified job.

PGMName Report (Ad-Hoc)

The PGMNAME REPORT is requested by including a PGMNAME= parameter as a SYSIN parameter in the SSA1WKLD data reduction Jobstream. The PGMNAME parameter consists of up to eight characters. A generic PGMNAME is supported by including an asterisk after the root of the PGMNAME. For example, you can select all programs beginning with SOM1 by specifying PGMNAME=SOM1*.

The selection of Programs to be included in this report is limited by the BEGIN DATE, END DATE, BEGIN TIME, and END TIME parameters.

The PGMNAME report includes the following data:

- Date and time the job entered the system
- Date and time the job finished execution
- CPU time used by the job
- Total number of DASD accesses issued

TSO USERID Report (Ad-Hoc)

The TSO USERID REPORT is requested by specifying the USERID= parameter as SYSIN parameter in the Data Reduction program Jobstream. The USERID consists of up to eight characters. A generic USERID is supported by including an asterisk after the root of the USERID. For example, all records with references to a USERID of TS01 would be requested by specifying USERID=TS01*.

The selection of the records from the input data is also limited by the BBEGIN DATE, END DATE, BEGIN TIME, and END TIME parameters.

The TSO USERID report contains the following data pertaining to each record selected:

- Time the TSO User logged on
- Time the TSO User logged off
- Amount of CPU time used
- Number of TSO commands executed

Summary Analysis Report

The SUMMARY ANALYSIS REPORT provides a narrative summary of the data encountered during the measured period. This report is especially useful for quickly evaluating the magnitude of the workload and the level of service provided by the data center during the measured period. It is also useful at the technical level to provide some basic indicators of developing performance problems and to identify areas to examine to solve current performance problems.

The following topics are addressed in the report:

- The dates included in the measured period
- Peak and average TSO activity including session statistics, average response time, average number of prime-time sessions, and the total number of sessions
- Peak paging activity
- Total jobs executed, including a breakdown by Class and Shift
- Total number of TSO transactions processed, including a breakdown by transaction size and shift.
- CPU Utilization by shift for each component of the workload
- Average paging activity by shift for each major component of the workload
- Average DASD device busy percentage by shift for the 10 busiest devices
- Average DASD device delay time in milliseconds by shift for the 10 devices experiencing the highest delays
- Average I/O service times in milliseconds by shift for the 10 devices experiencing the longest service times
- Average DASD I/O activity by shift for the 10 busiest strings
- Average channel busy percentage by shift for the 10 busiest channels

JOBNAME: CPU Activity Report

This report details all jobs that executed during the measured period, reporting the time that each job was resident and selected data related to the resource consumption. All values are average daily values. The reported data include:

System-Wide Statistics

TIMEFRAM	The average number of seconds per day included in the measured period. If the BEGIN TIME:END TIME parameters are omitted, the entire day would be included and the TIMEFRAM would be 86,400 seconds (24x60x60).
UP-TIME	The time in seconds that the CPU was available (daily average).
CPU(nn)	The time in seconds that the CPU was busy during the measured period (daily average).
NON-SYS	The number of seconds of CPU time expended in processing programs in problem state during the measured period (daily average).
OTHER	The non-system CPU time not accounted for by individual jobs.
SYSTEM	The number of seconds of CPU time consumed by the Operating System during the measured period.
WAIT	The number of seconds that the CPU was in the WAIT State during the measured period.

Bear in mind that the total number of CPU seconds available during the day is a function of the number of physical CPU engines installed. Thus, for a 2 engine SYSPLEX, if the machine was up 100% of the time, the total UP-TIME would be 172,800 seconds.

Job Related Statistics

JOB NAME

TOTAL SECNDS The number of CPU seconds consumed by the application (TCB + SRB)

% UP TIME The amount of CPU time consumed by the job expressed as a percent of the measured period:

$$\% \text{ UP TIME} = \frac{\text{TOTAL SECNDS}}{\text{UP-TIME}} \times 100$$

% CPU TIME The amount of CPU time consumed by the program expressed as a percentage of the time the CPU was busy:

$$\% \text{ CPU TIME} = \frac{\text{TOTAL SECNDS}}{\text{UP-TIME}} \times 100$$

Job Related Statistics

% NON-SYSTEM The amount of CPU time consumed by the job expressed as a percentage of the NON-SYSTEM time:

$$\% \text{ NON-SYSTEM} = \frac{\text{TOTAL SECNDS}}{\text{NON-SYSTEM}} \times 100$$

TIME ACTIVE The total time in seconds that the job was resident in memory.

% AVAILABLE The amount of time the job was available for processing expressed as a percentage of the TIMEFRAM.

$$\% \text{ AVAILABLE} = \frac{\text{TOTAL SECNDS}}{\text{TIMEFRAM}} \times 100$$

This report has an option (AVERAGES=NONE) that produces a list of all jobs that were executed during the named period (BEGIN DATE:END DATE) during the stated time (BEGIN TIME:END TIME), whether or not they consumed a significant percentage of CPU time. Furthermore, the report shows the total amount of CPU time used, in minutes, as opposed to a daily average value (the default for the Linear List format report).

Performance Reports

CIMS Capacity Planner produces two reports that present the batch job turnaround performance for various classifications of jobs and the TSO response performance for several classes of TSO transactions.

BATCH PERFORMANCE REPORT

The BATCH PERFORMANCE REPORT (Batch Turnaround Performance) provides a summary of the turnaround time experienced by batch jobs within several sets of job classifications expressed as a percentage of the number of jobs within each Category. A count of the number of jobs within each Category is also provided. The job classifications are Short, Medium, and Long. Within each of these classifications the turnaround times are broken down by the Region Sizes and whether or not they use any Magnetic Tapes. The Response Times are the difference in the time when the jobs are received into the input queue until they terminate and are available for output. The time that a job is awaiting output is not considered to be part of the turnaround time.

TSO PERFORMANCE REPORT

The TSO PERFORMANCE REPORT (TSO Response Performance) provides a summary of the TSO response time for Short, Medium, Long, and Long+ (extra long) transactions expressed as a percentage of the total transactions in each transaction Category. The above transaction categories are reported by Shift for each of the TSO Performance Groups.

The transaction categories are defined by performance period specifications defined in `SYS1.PARMLIB(1EA1PSxx)`:

- **SHORT:** transactions that completed in period 1
- **MEDIUM:** transactions that completed in period 2
- **LONG:** transactions that completed in period 3
- **LONG+:** transactions that completed after period 3

Workload Statistics Reports

CIMS Capacity Planner produces a set of Workload Statistics Reports that are designed to indicate the major consumers of CPU resources. The reports are available for batch jobs, programs, TSO USERID, TSO Command, printer and performance group. Each of the reports are described below.

JOB STATISTICS REPORT

The JOB STATISTICS REPORT provides the following information related to batch jobs:

- The top forty jobs in terms of elapsed (residency) time
For each job listed, the report details the JOBNAME, the total amount of time that the job was resident, and the percentage of the measured period that the job was resident.
- The top forty jobs in terms of their frequency of execution.
For each job listed, the report details the JOBNAME, the number of times that the job, or another job with the same name, was executed, and the percentage of the total mix represented by the job.
- The top forty jobs in terms of the amount of CPU time they consumed during the measured period.
For each job listed, the report details the amount of CPU time consumed and the percentage of the total CPU time consumed during the measured period.

Specific jobs can be included or excluded from the report through the use of the //INCLUDE or the //EXCLUDE DD Statements and their parameters that are described in the parameter section that follows the report descriptions.

PROGRAM STATISTICS REPORT

The PROGRAM STATISTICS REPORT provides the following information related to the programs exhibiting the highest Residence time, Frequency of execution, and CPU time.

- The top forty programs in terms of elapsed (residency) time.
For each program listed, the report details the program name and the total amount of time, in hours, minutes, and seconds, that the program was resident.
- The top forty programs in terms of their frequency of execution.
For each program listed, the report details the program name, the number of times that the program was executed, and the percentage of the total mix represented by the program.
- The top forty programs in terms of the amount of CPU time that they consumed.
For each program listed, the report details the amount of CPU time consumed and its percentage of the total CPU time consumed, in hours, minutes, and seconds, during the measured period.

Specific programs can be included or excluded from the report through the use of the //INCLUDE or the //EXCLUDE DD Statements and their parameters that are described in the parameter section following the report descriptions.

TSO USER STATISTICS REPORT

The TSO USER STATISTICS REPORT provides the following information related to the TSO Users exhibiting the highest Connect time, number of TSO Commands executed, and CPU time.

- The top forty TSO Users in terms of Connect time

For each TSO User listed, the report details the TSO USERID, the total amount of active time, in hours, minutes, and seconds, and the percentage of the measured period that the user was active.

- The top forty TSO Users in terms of the number of TSO transactions executed.

For each User listed, the report details the USERID, the number of TSO commands executed, and the percentage of the total TSO transactions submitted.

- The top forty TSO Users in terms of the amount of CPU time that they consumed.

For each TSO User listed, the report details the amount of CPU time, in hours, minutes, and seconds, that they consumed and their percentage of the total TSO CPU time consumed during the measured period.

Specific TSO Users can be included or excluded from the report through the use of the `//INCLUDE` or the `//EXCLUDE DD` Statements that are described in the parameter description section that follows the report descriptions.

TSO COMMAND STATISTICS REPORT

The TSO COMMAND STATISTICS REPORT provides the following information related to the TSO commands exhibiting the highest usage.

- The top forty TSO commands in terms of frequency of use

For each TSO command listed, the report details the TSO command, the number of times that the command was issued, and its percentage of all TSO commands issued.

- The top forty TSO commands in terms of the amount of CPU time that they consumed

For each TSO command listed, the report details the amount of CPU time, in hours, minutes, and seconds, that it consumed and its percentage of the total TSO CPU time consumed during the measured period.

Specific TSO Commands can be included or excluded from the report through the use of the `//INCLUDE DD` and/or the `//EXCLUDE DD` Statements that are described in the parameter description section that follows the report descriptions.

PRINTER STATISTICS REPORT

The PRINTER STATISTICS REPORT provides the following information related to the printers that exhibit the highest activity in terms of the number of lines, pages, and data sets printed.

- The top forty printers in terms of the number of Lines printed

For each printer listed, the report details the number of lines printed and its percentage of the total print workload.

- The top forty printers in terms of the number of pages printed

For each printer listed, the report details the printer ID, the number of pages printed and the its percentage of all pages printed.

- The top forty printers in terms of the number of print data sets that they processed.

For each printer listed, the report details the printer id, the number of print data sets processed, and its percentage of all print data sets processed.

Specific printers can be included or excluded from the report through the use of the `//INCLUDE` or the `//EXCLUDE DD` Statements that are described in the parameter description section that follows the report descriptions.

PRINTER ACTIVITY REPORT

The Printer Linear List shows actual printer activity (lines printed, pages printed, datasets processed) for all printers in the data center. This report supplements the Printer Statistics Report which shows the top 40 printers in each of the above three categories.

The Printer Activity Report is invoked through the `SSA1RPT` program. Sample JCL is available in member `DWKLREPT` in the CIMS Capacity Planner job library.

PGN SERVICE UNIT STATISTICS REPORT

The PGN Service Unit Statistics Report provides the following information related to Performance Groups as defined in the IPS:

- The top forty PGNs in terms of CPU Service Units
- The top forty PGNs in terms of MSO Service Units
- The top forty PGNs in terms of IOC Service Units
- The top forty PGNs in terms of Paging Activity

Specific Performance Groups Numbers can be included or excluded from the report through the use of the `//INCLUDE` or `//EXCLUDE DD` Statements. In either case, the performance group to be included/excluded must be preceded by `PGN` followed by 5 numeric digits (`PGN00001`, `PGN00100`) or by `PGN` plus a wild card (`PGN001*`) to include/exclude a range.

Workload Exception Reports

CIMS Capacity Planner produces several reports which identify exceptional conditions affecting processor, DASD, and channels. Each of the reports is described below.

PROCESSOR EXCEPTION ANALYSIS REPORT

The PROCESSOR EXCEPTION ANALYSIS Report provides exception data related to each fifteen-minute period during which the CPU utilization and/or the paging rates exceed the predefined levels specified in the CPPR.PARMLIB member &sidXCPT.

The following information is included in the report for each fifteen-minute period during which the CPU utilization or the paging rates exceed the specified limits:

- Date and time of the fifteen-minute period
- The type of limit that was exceeded (either BUSY or PAGES)
- The CPU number if the CPU busy limit was exceeded
- The amount of CPU time consumed (% BUSY)
- The number of pages per second if the paging rate was exceeded

DASD EXCEPTION ANALYSIS REPORT

The DASD EXCEPTION ANALYSIS REPORT provides exception data related to each one-hour period during which the percentage busy by device, queue length exceptions by device, or I/O service times by device exceed a set of pre-established limits. The limits are specified in the CPPR.PARMLIB member &sidXCPT.

The following information is included in the report for each one-hour period during which one or more of the above limits are exceeded.

- Date and time during which the limits were exceeded
- The type of limit(s) exceeded
- Device busy

Whenever a device is exceeds the specified limit, the type of exception is listed (BUSY) along with the percent of the period that the device was busy, and the volume serial number.

- Queue Length

Whenever the queue length is exceeded for a device, the type of exception is listed (QLEN) along with the average queue length, and the volume serial number.

CHANNEL EXCEPTION ANALYSIS REPORT

The CHANNEL EXCEPTION ANALYSIS REPORT provides exception data related to the data channel traffic whenever the percentage of time that a channel is busy exceeds a pre-specified limit within any given one-hour period. The Pre-specified limit is contained in the CPPER.PARMLIB member &sidXCPT.

Whenever an exception is detected, the following information is listed on the report:

- Date and time period during which the norm is exceeded
- BUSY % indicates that the channel busy percent limit was exceeded
- The three-digit channel (CHPID) identifier of each channel exceeding the reporting threshold
- The percentage of the time that the channel was busy

Workload DASD Reports

CIMS Capacity Planner produces reports which highlight I/O activity in the DASD farm. The reports are described below.

DASD ACTIVITY REPORT

This report shows, for all volumes in the DASD farm, the daily average values for the following categories:

- SIOs per Second
- Queue Delay in Milliseconds
- Percent Busy
- Average I/O Service Time exclusive of Queue Delay in Milliseconds

In addition, the values are summarized by LCU (Logical Control Unit).

The time frame for the report is bounded through the use of BEGIN TIME and END TIME parameter statements.

DASD DETAIL REPORT

The DASD DETAIL REPORT shows, by day, the total number of SIO incidents for each specific volume in the DASD farm. The SIO count is further broken down by shift. The values are presented by logical string.

Workload Graphs

CPPR produces a number of reports that graphically depict selected elements of the workload. These graphs are produced on the system printer as part of the workload report generation program. Each graph is individually selectable using the various `SYSIN` parameters included in the `SSA1RPT` jobstream.

CPU ACTIVITY GRAPH

The CPU ACTIVITY GRAPH is produced by day of the week such as Monday, Tuesday, etc. In the event that the reporting period, as defined by the `BEGIN DATE` and the `END DATE`, contains more than one Monday, for example, the graph produced for Monday would include more than one days' worth of data. If a list of specific days is not specified, the CPU ACTIVITY GRAPHS for all days (Sunday through Saturday) are produced.

When more than one day's worth of data is combined into a single graph, the CPU BUSY percentages for each 15 minute period throughout the day are averaged for reporting purposes. The high water mark average value is also retained so that the graph will present both the average and maximum CPU percent utilization for each 15 minute period throughout each day. The average and maximum percent utilization values are presented on each line of the graph.

The graph is presented in a horizontal bar chart format with the horizontal axis representing the CPU utilization percent and the vertical axis representing the time of day in 15 minute intervals. The data are presented in four groups of six hours each.

The peak periodic CPU Utilization value is represented by a line of dashes (-) while the average values are represented by a set of asterisks (*). When no dashes are visible, it means that the average and peak values are identical or nearly so.

For the various parameters that apply to the CPU activity graph, see the parameters under the report generation parameter section following the report descriptions.

TSO ACTIVITY GRAPH

The TSO ACTIVITY GRAPH is produced by day of the week such as Monday, Tuesday, etc. In the event that the reporting period, as defined by the `BEGIN DATE` and the `END DATE` contains more than one Monday, for example, the graph produced for Monday would include more than one days' worth of data. If a list of specific days is not specified, the TSO ACTIVITY GRAPHS for all days (Sunday through Saturday) are produced.

When more than one day's worth of data is combined into a single graph, the TSO Activity values for each 15 minute period throughout the day are averaged for reporting purposes. The high water mark value is also retained so that the graph will present both the average and maximum value for each 15 minute period throughout each day. The average and maximum values are presented on each line of the graph.

The graph is presented in a horizontal bar chart format with the horizontal axis representing TSO Users Active and the vertical axis representing the time of day in 15 minute intervals. The data are presented in four groups of six hours each.

The maximum TSO Activity value is represented by a line of dashes (-) while the average values are represented by a set of asterisks (*). When no dashes are visible, it means that the average and peak values are identical or virtually so.

For the various parameters that apply to the TSO ACTIVITY GRAPH, see the parameters under the report generation parameter section following the report descriptions.

PAGING ACTIVITY GRAPH

The PAGING ACTIVITY GRAPH is produced by day of the week such as Monday, Tuesday, etc. In the event that the reporting period, as defined by the BEGIN DATE and the END DATE contains more than one Monday, for example, the graph produced for Monday would include more than one days' worth of data. If a list of specific days is not specified, the PAGING ACTIVITY GRAPHS for all days (Sunday through Saturday) are produced.

When more than one day's worth the data is combined into a single graph, the paging rates for each 15 minute period throughout the day are averaged for reporting purposes. The peak value for each 15 minute period is retained so that the graph will present both the average and peak paging rates for each 15 minute period throughout each day. The average and peak paging rates are printed on each line of the graph.

The graph is presented in a horizontal bar chart format with the horizontal axis representing the paging rates (pages per second) and the vertical axis representing the time of day in 15 minute intervals. The data are presented in four groups of six hours each.

The maximum paging rate is represented by a line of dashes (-) while the average values are represented by a set of asterisks (*). When no dashes are visible, it means that the average and peak values are identical or virtually so.

For the various parameters that apply to the PAGING ACTIVITY GRAPH, see the parameters under the report generation parameter section following the report descriptions.

CHANNEL ACTIVITY GRAPH

The CHANNEL ACTIVITY GRAPH presents the average and peak channel utilization (% busy) for each day of the week by channel for the reporting period included within the BEGIN DATE and END DATE parameters specified with the report request. The data are reported by shift for each data channel. A performance threshold is drawn on each chart to highlight when any channel activity exceeds the system performance threshold.

The data are presented in a vertical bar chart format with the horizontal axis depicting the days of the week by data channel and the vertical axis denoting the channel utilization levels.

The average channel busy percentages are determined by averaging the channel busy percentages for each of the 15 minute periods throughout each shift for the days specified in the report request. The peak values are determined by retaining the maximum of all channel busy percentages during each shift for each day.

Each request for the CHANNEL ACTIVITY GRAPH will result in three sets of graphs being produced (one set for each shift). When the reporting period exceeds one week, multiple days are averaged (multiple Mondays are taken together and reported as Monday's average and peak values, and so on throughout the remaining days of the week). The weekly average and peak values are reported for each data channel.

Average values are represented by asterisks "*", while peak values are represented by vertical bars "|." When the average values are equal to the peak values, the vertical bars are overlaid with asterisks. The absence of any vertical bars indicates that the peak and the average values are either equal or nearly so.

CPU RATIO GRAPH

The CPU Ratio Graph shows the percent CPU busy by hour broken down into the workload categories as specified in the Local member of the CPPR.PARMLIB.

PAGING RATIO GRAPH

The Paging Ratio Graph shows paging activity (Central to Auxiliary and Auxiliary to Central) in terms of pages per second, by hour of the day, broken down into the workload categories as specified in the Local member of CPPR.PARMLIB.

DASD I/O SERVICE TIME GRAPH

The DASD I/O SERVICE TIME GRAPH presents the average and peak I/O service times in milliseconds for each day of the week by DASD VOLSER/Channel address for the reporting period included within the BEGIN DATE and END DATE parameters specified with the report request. The data are reported by shift for each device. A performance threshold is drawn on each chart to highlight when any service times exceed the system performance threshold.

The data are presented in a vertical bar chart format with the horizontal axis depicting the days of the week by device and with the vertical axis denoting the service times.

The average service times are determined by averaging the service times for each of the 15 minute periods throughout each shift for the days specified in the report request. The peak values are determined by retaining the maximum of all service times encountered during each shift for each day.

Each request for the DASD I/O SERVICE TIME GRAPH will result in three sets of graphs being produced (one set for each shift). When the reporting period exceeds one week, multiple days are averaged (multiple Mondays are taken together and reported as Monday's average and peak values, and so on throughout the remaining days of the week). The weekly average and peak values are reported for each data channel.

Average values are represented by asterisks "*", while peak values are represented by vertical bars "|". When the average values are equal to the peak values, the vertical bars are overlaid with asterisks. The absence of any vertical bars indicates that the peak and the average values are either equal or nearly so.

DASD DEVICE ACTIVITY GRAPH

The DASD DEVICE ACTIVITY GRAPH presents the average and peak loads in I/O operations per second for each day of the week for each DASD string by VOLSER/Channel address within the reporting period defined by the BEGIN DATE and END DATE parameters specified with the report request. The data are reported by shift for each DASD string.

The data are presented in a vertical bar chart format with the horizontal axis depicting the days of the week by device and with the vertical axis denoting the number of I/O operations per second for each device.

The average I/Os per second are determined by averaging the number of I/Os per second for each of the 15 minute periods throughout each shift for the days specified in the report request. The peak values are determined by retaining the maximum average of all I/Os per second encountered during each shift for each day.

Each request for the DASD DEVICE ACTIVITY GRAPH will result in three sets of graphs being produced (one set for each shift). When the reporting period exceeds one week, multiple days are averaged (multiple Mondays are taken together and reported as Monday's average and peak values, and so on throughout the remaining days of the week). The weekly average and peak values are reported for each DASD string.

Average values are represented by asterisks "*", while peak values are represented by vertical bars "|". When the average values are equal to the peak values, the vertical bars are overlaid with asterisks. The absence of any vertical bars indicates that the peak and the average values are either equal or nearly so.

DASD QUEUE DELAY GRAPH

The DASD QUEUE DELAY GRAPH presents the average and peak DASD queue delay times in milliseconds for each day of the week for each DASD string by VOLSER/Channel address for the reporting period included within the BEGIN DATE and END DATE parameters specified with the report request. The data are reported by shift for each DASD string.

The data are presented in a vertical bar chart format with the horizontal axis depicting the days of the week by device and with the vertical axis denoting the queue delays.

The average queue delay times are determined by averaging the queue delay times for each of the 15 minute periods throughout each shift for the days specified in the report request. The peak values are determined by retaining the maximum of all queue delay times encountered during each shift for each day.

Each request for the DASD QUEUE DELAY GRAPH will result in three sets of graphs being produced (one set for each shift). When the reporting period exceeds one week, multiple days are averaged (multiple Mondays are taken together and reported as Monday's average and peak values, and so on throughout the remaining days of the week). The weekly average and peak values are reported for each DASD string.

Average values are represented by asterisks "*", while peak values are represented by vertical bars "|". When the average values are equal to the peak values, the vertical bars are overlaid with asterisks. The absence of any vertical bars indicates that the peak and the average values are either equal or nearly so.

DASD DEVICE BUSY GRAPH

The DASD DEVICE BUSY GRAPH presents the average and peak loads in percent busy for each day of the week for each DASD string by VOLSER/Channel address for the reporting period included within the BEGIN DATE and END DATE parameters specified with the report request. The data are reported by shift for each DASD string.

The data are presented in a vertical bar chart format with the horizontal axis depicting the days of the week by device and with the vertical axis denoting the percentage of time each device was busy.

The average device busy percentages are determined by averaging the device busy percentages for each of the 15 minute periods throughout each shift for the days specified in the report request. The peak values are determined by retaining the maximum of all device busy percentages encountered during each shift for each day.

Each request for the DASD DEVICE BUSY GRAPH will result in three sets of graphs being produced (one set for each shift). When the reporting period exceeds one week, multiple days are averaged (multiple Mondays are taken together and reported as Monday's average and peak values, and so on throughout the remaining days of the week). The weekly average and peak values are reported for each DASD string.

Average values are represented by asterisks "*", while peak values are represented by vertical bars "|". When the average values are equal to the peak values, the vertical bars are overlaid with asterisks. The absence of any vertical bars indicates that the peak and the average values are either equal or nearly so.

Central/Expanded/Auxiliary Storage Activity Graphs (ESA or later only)

CIMS Capacity Planner produces five graphs that depict the paging activity between Central, Expanded, and Auxiliary storage. They are:

CENTRAL-TO-EXPANDED STORAGE ACTIVITY GRAPH

This graph shows Page Move activity from Central Storage to Expanded Storage in terms of Pages per second. The average value for each fifteen minute period is shown, along with the peak value for that fifteen minute period.

EXPANDED-TO-CENTRAL STORAGE ACTIVITY GRAPH

This graph shows Page Move activity from Expanded Storage to Central Storage in terms of Pages per second. The average value for each fifteen minute period is shown, along with the peak value for that fifteen minute period.

EXPANDED-TO-AUXILIARY STORAGE ACTIVITY GRAPH

This graph shows Page Out activity from Expanded Storage to Auxiliary Storage through Central Storage in terms of Pages per second. The average value for each fifteen minute period is shown, along with the peak value for that fifteen minute period.

CENTRAL-TO-AUXILIARY STORAGE ACTIVITY GRAPH

This graph shows Page Out activity from Central Storage to Auxiliary Storage in terms of Pages per second. The average value for each fifteen minute period is shown, along with the peak value for that fifteen minute period.

AUXILIARY-TO-CENTRAL STORAGE ACTIVITY GRAPH

This graph shows Page In activity from Auxiliary Storage to Central Storage in terms of Pages per second. The average value for each fifteen minute period is shown, along with the peak value for that fifteen minute period.

Operating JCL for Workload Batch Reports

The operating JCL has been designed to conform to system utility JCL, following IBM conventions wherever possible. Input comes from the Consolidated Table File (ONLINE). The DD Statements required for program execution are described below.

ONLINE

This DD Statement specifies the data set name and the disposition of the Consolidated Table File (Performance Data Base)

INDEX

This DD Statement specifies the data set name and disposition of the CPPR Table Index dataset.

CPPRPARM

This DD Statement specifies the data set name and the disposition of the CPPR Parameter Library.

CPPRERT

This DD Statement specifies the dataset name and the disposition of the CPPR Element Registration Table dataset.

SYSPRINT

This describes the report output file. Unless otherwise specified, the DCB characteristics are as follows:

```
RECFM=FBA, BLKSIZE=133, LRECL=133
```

The SYSPRINT data set is generally specified as SYSOUT=* and routed to the system printer after the condition codes are examined and the output has been reviewed at the terminal.

SYSNAP

This DD Statement specifies an optional print file for diagnostic hexadecimal snap dumps (see `DEBUGON` below)

EXCLUDE/INCLUDE

This is the Element Exclusion or Inclusion File. It is used by the Report Program to screen out or specifically select elements from the Statistics Reports. The format of the input statements is as follows:

- An element can be up to eight characters in length
- Multiple elements can be included in a single record
- Elements included in a single record must be separated by commas
- Elements can be terminated by a "wildcard" character (*), indicating that all elements that match the characters up to the "wildcard" are to be excluded from the reports.

The `EXCLUDE` or `INCLUDE` data set is usually specified using in-stream input data. Either an `EXCLUDE` or `INCLUDE` data set (but not both) may be specified.

For example, the following input statement to the Exclusion file:

```
//EXCLUDE DD *  
JES2,TSOU*,CICS161T,ALLOCAS  
//
```

would indicate that the Statistics Reports would not include JES2; any Job name, Program Name, TSO USERID or TSO Command beginning with the four characters "TSOU". Any job name, program name, user id, or TSO command of CICS161T or ALLOCAS would also be excluded.

The exclusion/inclusion file must have an `LRECL` of 80. Only the first 72 columns of the statement are interpreted, leaving columns 73-80 for sequence numbers.

SYSIN

This is the parameter input file. For detailed information regarding the input parameters, see *Input Parameter Statements for Batch Reports on page 1-38*.

Sample JCL

The following is an example of a Jobstream that will produce several of the reports supported by the CIMS Capacity Planner Workload Subsystem. This sample JCL is found in the CIMS Capacity Planner job library under the member name DWKLREPT.

```
//SSA1REPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1RPT, REGION=5000K, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD, DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.Vnnn.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSMSG DD SYSOUT=*
//SYSIN DD *
BEGIN DATE=01/01/2005
END DATE=01/31/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
SELECTED SYSTEM=*
CPU ACTIVITY GRAPH=YES
JOBNAME:CPU ACTIVITY REPORT=YES
PGN SERVICE UNIT STATISTICS REPORT=YES
PAGING ACTIVITY GRAPH=YES
TSO ACTIVITY GRAPH=YES
DASD ACTIVITY GRAPH=YES
DASD DEVICE BUSY GRAPH=YES
DASD I/O SERVICE TIME GRAPH=YES
DASD QUEUE DELAY GRAPH=YES
DASD ACTIVITY REPORT=YES
PRINTER ACTIVITY REPORT=YES
CHANNEL ACTIVITY GRAPH=YES
CPU RATIO GRAPH=YES
PAGING RATIO GRAPH=YES
TSO PERFORMANCE REPORT=YES
BATCH PERFORMANCE REPORT=YES
JOBCLASS=T
DASD DETAIL REPORT=YES
JOB STATISTICS REPORT=YES
PROGRAM STATISTICS REPORT=YES
PRINTER STATISTICS REPORT=YES
TSO USER STATISTICS REPORT=YES
TSO COMMAND STATISTICS REPORT=YES
PROCESSOR EXCEPTION ANALYSIS=YES
DASD EXCEPTION ANALYSIS=YES
CHANNEL EXCEPTION ANALYSIS=YES
*
***** THE FOLLOWING PARAMETERS PRODUCE DASD REPORTS
*
```

```
***PRODUCE THE DSNAME-BY-VOLUME REPORT (10 BUSIEST PLUS MVSDLB)
*
VOLUME ACTIVITY REPORT=YES
VOLUME SELECTION CRITERIA=TOP10
SELECTED VOLUME=MVSDLB
*
*** PRODUCE THE DSNAME-BY-DATA CENTER REPORT
*
DSNAME ACTIVITY REPORT=YES
*
*** PRODUCE THE DSNAME-BY-SELECTION CRITERIA REPORT
*
DSNAME DETAIL REPORT=YES
DSNAME=SYS2.CPPR*
VOLSER=SYS83*
*
*** THE FOLLOWING REPORTS ARE FOR ESA SYSTEMS
*
CENTRAL TO EXPANDED STORAGE ACTIVITY GRAPH=YES
EXTENDED TO CENTRAL STORAGE ACTIVITY GRAPH=YES
EXPANDED TO AUXILIARY STORAGE ACTIVITY GRAPH=YES
CENTRAL TO AUXILIARY STORAGE ACTIVITY GRAPH=YES
AUXILIARY TO CENTRAL STORAGE ACTIVITY GRAPH=YES
/* * * IF YOU WISH TO LIMIT THE DASD DEVICES LISTED IN THE
/* * * SUMMARY REPORT OR THE DASD ACTIVITY REPORT, SPECIFY
/* * * THE VOLSER OF THOSE VOLUMES YOU WISH TO INCLUDE/EXCLUDE
/* * * IN AN INCLUDE/EXCLUDE STREAM. FOR EXAMPLE, TO EXCLUDE /* * * ALL VOLUMES
BEGINNING WITH MVS:
//EXCLUDE DD *
MVS*
```

Input Parameter Statements for Batch Reports

Common Parameters

The following input statements set common parameters which apply to all reports or graphs in this run.

SELECTED SYSTEM

This parameter is used to identify the SID of the System for which SMF/RMF data is to be reported. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Up to five SIDs can be specified in the current version of the Workload Analysis Program. If multiple SIDs are indicated, all parameters pertain to each of the selected systems. That is, all other parameters (beginning and ending times and dates, language, output filtering) apply equally to all systems being analyzed.

Note • For storage requirements, figure that each SID specified will require approximately 800K of Virtual Storage. For example, if five systems are to be analyzed in a single run, the REGION size for the run should be 4096K.

Options

This parameter has many options. If the program is to produce a single report for a single system (SID), specify:

- The SID itself (e.g, `SELECTED SYSTEM=MVSA`)
- An '*', indicating the SID for the system on which the program is being executed (e.g, `SELECTED SYSTEM=*`)

If multiple (up to five) Systems are to be analyzed in a single run, specify the parameter as follows:

- Each specific SID, separated by commas (e.g, `SELECTED SYSTEM=MVSA,MVSB,MVSC`)

Default Option

`SELECTED SYSTEM=*`

BEGIN DATE

This parameter is used to specify the earliest date for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify a beginning date in one of the following forms:

- YYYYDDD (Standard Julian format)
- MM/DD/YYYY (USA Standard Gregorian format)
- DD.MM.YYYY (European Standard Gregorian format)
- * or *-n (relative date where * is today)

Note • All dates must be specified in one of the above formats.

Default Date

BEGIN DATE=00001

END DATE

This parameter is used to specify the last date for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify an ending date.

Default Date

END DATE=2099365

PRIME SHIFT FIRST HOUR

This parameter is used to specify the first hour of Prime Shift for reporting purposes.

Specify Hour

Specify an hour in the 24 hour format (e.g, PRIME SHIFT FIRST HOUR=7)

Default Hour

PRIME SHIFT FIRST HOUR=8

LATE SHIFT FIRST HOUR

This parameter is used to specify the first hour of Late Shift for reporting purposes.

Specify Hour

Specify an hour in the 24 hour format (e.g, LATE SHIFT FIRST HOUR=20).

Default Hour

LATE SHIFT FIRST HOUR=16

SELECTED DAY

This parameter is used to specify a specific day of the week for which reports are to be generated. The day is specified by name (e.g, MONDAY). Multiple days can be specified, separated by commas (e.g, MONDAY , WEDNESDAY).

Specify Day(s)

Specify one or more names of days, or ALL DAYS (e.g, SELECTED DAY=MONDAY , FRIDAY).

Default Day(s)

SELECTED DAY=ALL DAYS

LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter does not cause any data to be removed from the Performance Data Base.

Specify Hour

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

REPORT LANGUAGE

This parameter describes the language that is to be used for the narrative sections of the report.

Specify Report Language

ENGLISH, DEUTSCH or any combination of the languages, separated by commas.

If multiple languages and multiple SIDs are specified, the order of the reports are as follows:

- The first selected system (SID) in the primary language
- The first selected system (SID) in the secondary and succeeding languages (if more than two are specified)
- The second selected system (SID) in the primary language
- The second selected system (SID) in the secondary and succeeding languages (if more than two are specified)
- The third selected system (SID) in the primary language (and so forth)

Default Language

REPORT LANGUAGE=ENGLISH

GEOGRAPHIC LOCATION

This parameter is used to determine the format of printed numerics (commas and decimal points) as well as the format of the date that is printed within the reports.

Specify Location

GEOGRAPHIC LOCATION=USA (date=mm/dd/yy)

GEOGRAPHIC LOCATION=EUROPE (date=dd.mm.yy)

GEOGRAPHIC LOCATION=GREAT BRITAIN (date=dd/mm/yy)

Default Location

GEOGRAPHIC LOCATION=USA

DETAIL

This parameter turns on or turns off the detail information that describes specific volumes, percentages, counts, etc. during the generation of reports.

Options

Either YES or NO.

Default Option

DETAIL=YES

DEBUGON

This parameter is used to request SNAP dumps at various points in the program. It should only be used if problem determination is required.

Options

None.

Default Option

None.

LOCAL HOLIDAYS

This parameter tells the program whether or not to exclude holidays when preparing a set of reports. It is used in conjunction with the PARMLIB member HOLIDAYS that contain a list of all local holidays.

Options

LOCAL HOLIDAYS=EXCLUDE (don't include holidays) or

LOCAL HOLIDAYS=INCLUDE

Default Option

LOCAL HOLIDAYS=INCLUDE

Report Statements

The following input statements generate the described reports or allow further modification.

SUMMARY ANALYSIS REPORT

This parameter tells the program whether or not to produce the Summary report.

Options

SUMMARY REPORT=YES (produce the report) or

SUMMARY REPORT=NO (don't produce it)

Default Option

SUMMARY REPORT=YES

JOB STATISTICS REPORT

This parameter tells the program whether or not to produce the Job Statistics report.

Options

JOB STATISTICS REPORT=YES (produce the report) or

JOB STATISTICS REPORT=NO (don't produce it)

Default Option

JOB STATISTICS REPORT=NO

PROGRAM STATISTICS REPORT

This parameter tells the program whether or not to produce the Program Statistics report.

Options

PROGRAM STATISTICS REPORT=YES (produce the report) or

PROGRAM STATISTICS REPORT=NO (don't produce it)

Default Option

PROGRAM STATISTICS REPORT=NO

TSO USER STATISTICS REPORT

This parameter tells the program whether or not to produce the TSO User Statistics report.

Options

TSO USER STATISTICS REPORT=YES (produce the report) or

TSO USER STATISTICS REPORT=NO (don't produce it)

Default Option

TSO USER STATISTICS REPORT=NO

TSO COMMAND STATISTICS REPORT

This parameter tells the program whether or not to produce the TSO Command Statistics report.

Options

TSO COMMAND STATISTICS REPORT=YES (produce the report) or

TSO COMMAND STATISTICS REPORT=NO (don't produce it)

Default Option

TSO COMMAND STATISTICS REPORT=NO

PRINTER STATISTICS REPORT

This parameter tells the program whether or not to produce the Printer Statistics report.

Options

PRINTER STATISTICS REPORT=YES (produce the report) or

PRINTER STATISTICS REPORT=NO (don't produce it)

Default Option:

PRINTER STATISTICS REPORT=NO

BATCH PERFORMANCE REPORT

This parameter tells the program whether or not to produce the Batch Performance Objectives report.

Options

BATCH PERFORMANCE REPORT=YES (produce the report) or

BATCH PERFORMANCE REPORT=NO (don't produce it)

Default Option

BATCH PERFORMANCE REPORT=NO

BATCH TOTALS

This parameter tells the program whether to convert the totals in the Batch Performance Report to percentages (PERCENT) or to leave them as actual numbers (NUMERIC).

Options

BATCH TOTALS=PERCENT (convert the totals to percentages)

BATCH TOTALS=NUMERIC (don't convert the totals)

Default Option

BATCH TOTALS=PERCENT

JOBCLASS

This parameter allows you to select a specific Job Class or Job Classes that are used to produce the Batch Performance Objectives report. If no Job Class selection parameter is specified, all non-numeric (A-Z) Job Classes are selected. One or more Job Classes can be specified. Job Classes must NOT be separated by commas.

Options

JOBCLASS=ABCEFG...

(select specific Job Classes)

Default Option

(all non-numeric Job Classes are selected)

TSO PERFORMANCE REPORT

This parameter tells the program whether or not to produce the TSO Performance Objectives report.

Options

TSO PERFORMANCE REPORT=YES (produce the report) or

TSO PERFORMANCE REPORT=NO (don't produce it)

Default Option

TSO PERFORMANCE REPORT=NO

DASD DETAIL REPORT

This parameter is used to request a detailed DASD Device Activity report. The report shows actual SIO/SSCH counts by day for each DASD Device. The report is printed only if this parameter is specified.

Options

DASD DETAIL REPORT=YES or DASD DETAIL REPORT=NO

Default Option

DASD DETAIL REPORT=NO

CANCEL TRANSIENT OPTION

This parameter indicates whether or not Transient Datasets (i.e, datasets that have been allocated and deleted in a 24 hour period) are to be listed individually by DSNAME or are to be collected together under the common heading "TRANSIENT" on the Volume Activity Report.

Options

CANCEL TRANSIENT OPTION=YES, or

CANCEL TRANSIENT OPTION=NO

Default Option

CANCEL TRANSIENT OPTION=NO

JOBNAME: CPU ACTIVITY REPORT

This parameter specifies whether or not the CPU Activity report is to be produced.

Options

CPU ACTIVITY REPORT=YES (produce the report) or

CPU ACTIVITY REPORT=NO (don't produce it)

Default option

CPU ACTIVITY REPORT=YES

DASD ACTIVITY REPORT

This parameter is used to request a detailed DASD Activity Report. The report shows average SIO/SSCH counts per second for each DASD Device along with % busy, average I/O service times and average queue delay. The report is printed only if this parameter is specified.

Options

DASD ACTIVITY REPORT=YES or DASD ACTIVITY REPORT=NO

Default option

DASD ACTIVITY REPORT=NO

PGN SERVICE UNIT STATISTICS REPORT

This parameter is used to request a PGN Service Unit Statistics report that shows, for the top 40 PGNs in each category, CPU Service Units, MSO Service Unit, IOC Service Units and Page In Activity.

Options

PGN SERVICE UNIT STATISTICS REPORT=YES (produce the report)

PGN SERVICE UNIT STATISTICS REPORT=NO (don't produce the report)

Default option

PGN SERVICE UNIT STATISTICS REPORT=NO

DASD DETAIL REPORT

This parameter is used to request a DASD Detail report that shows, for each day in the period, the total number of SIO/SSCHs for each DASD Volume.

Options

DASD DETAIL REPORT=YES (produce the report)

DASD DETAIL REPORT=NO (don't produce the report)

Default option

DASD DETAIL REPORT=NO

Exception Reports

The following statements generate the described exception reports.

PROCESSOR EXCEPTION ANALYSIS

This parameter tells the program whether or not to produce the Processor Exception analysis report.

Options

PROCESSOR EXCEPTION ANALYSIS=YES (produce the report) or

PROCESSOR EXCEPTION ANALYSIS=NO (don't produce it)

Default option

PROCESSOR EXCEPTION ANALYSIS=NO

DASD EXCEPTION ANALYSIS

This parameter tells the program whether or not to produce the DASD Exception analysis report.

Options

DASD EXCEPTION ANALYSIS=YES (produce the report) or

DASD EXCEPTION ANALYSIS=NO (don't produce it)

Default option

DASD EXCEPTION ANALYSIS=NO

CHANNEL EXCEPTION ANALYSIS

This parameter tells the program whether or not to produce the Channel Exception analysis report.

Options

CHANNEL EXCEPTION ANALYSIS=YES (produce the report) or

CHANNEL EXCEPTION ANALYSIS=NO (don't produce it)

Default option

CHANNEL EXCEPTION ANALYSIS=NO

Graph Statements

The following input statements generate the described graphs.

CPU ACTIVITY GRAPH

This parameter tells the program whether or not to produce the CPU Activity Graphic report.

Options

CPU ACTIVITY GRAPH=YES (produce the report) or

CPU ACTIVITY GRAPH=NO (don't produce it)

Default option

CPU ACTIVITY GRAPH=NO

PAGING ACTIVITY GRAPH

This parameter tells the program whether or not to produce the Paging Activity Graphic report.

Options

PAGING ACTIVITY GRAPH=YES (produce the report) or

PAGING ACTIVITY GRAPH=NO (don't produce it)

Default option

PAGING ACTIVITY GRAPH=NO

TSO TRANSACTION ACTIVITY GRAPH

This parameter tells the program whether or not to produce the TSO Transaction Activity Graphic report.

Options

TSO TRANSACTION ACTIVITY GRAPH=YES (produce the report) or

TSO TRANSACTION ACTIVITY GRAPH=NO (don't produce it)

Default option

TSO TRANSACTION ACTIVITY GRAPH=NO

TSO TRANSACTION RESPONSE GRAPH

This parameter tells the program whether or not to produce the TSO Transaction Response Graphic report.

Options

TSO TRANSACTION RESPONSE GRAPH=YES (produce the report) or

TSO TRANSACTION RESPONSE GRAPH=NO (don't produce it)

Default option

TSO TRANSACTION RESPONSE GRAPH=NO

TSO ACTIVITY GRAPH

This parameter tells the program whether or not to produce the TSO Activity Graphic report.

Options

TSO ACTIVITY GRAPH=YES (produce the report) or

TSO ACTIVITY GRAPH=NO (don't produce it)

Default Option:

TSO ACTIVITY GRAPH=NO

DASD ACTIVITY GRAPH

This parameter tells the program whether or not to produce the DASD Activity Graphic report.

Options

DASD ACTIVITY GRAPH=YES (produce the report) or

DASD ACTIVITY GRAPH=NO (don't produce it)

Default option

DASD ACTIVITY GRAPH=NO

DASD QUEUE DELAY GRAPH

This parameter tells the program whether or not to produce the DASD Queue Delay Graphic report.

Options

DASD QUEUE DELAY GRAPH=YES (produce the report) or

DASD QUEUE DELAY GRAPH=NO (don't produce it)

Default option

DASD QUEUE DELAY GRAPH=NO

DASD I/O SERVICE TIME GRAPH

This parameter tells the program whether or not to produce the DASD I/O Service Time Graphic report.

Options

DASD I/O SERVICE TIME GRAPH=YES (produce the report) or

DASD I/O SERVICE TIME GRAPH=NO (don't produce it)

Default option

DASD I/O SERVICE TIME GRAPH=NO

DASD DEVICE BUSY GRAPH

This parameter tells the program whether or not to produce the DASD Device Busy Graphic report.

Options

DASD DEVICE BUSY GRAPH=YES (produce the report) or

DASD DEVICE BUSY GRAPH=NO (don't produce it)

Default option

DASD DEVICE BUSY GRAPH=NO

CHANNEL ACTIVITY GRAPH

This parameter tells the program whether or not to produce the Channel Activity Graphic report.

Options

CHANNEL ACTIVITY GRAPH=YES (produce the report) or

CHANNEL ACTIVITY GRAPH=NO (don't produce it)

Default option

CHANNEL ACTIVITY GRAPH=NO

CPU RATIO GRAPH

This parameter tells the program whether or not to produce the CPU Ratio Graphic report.

Options

CPU RATIO GRAPH=YES (produce the report) or

CPU RATIO GRAPH=NO (don't produce it)

Default option

CPU RATIO GRAPH=NO

PAGING RATIO GRAPH

This parameter tells the program whether or not to produce the Paging Ratio Graphic report.

Options

PAGING RATIO GRAPH=YES (produce the report) or

PAGING RATIO GRAPH=NO (don't produce it)

Default option

PAGING RATIO GRAPH=NO

CENTRAL-TO-EXPANDED STORAGE ACTIVITY GRAPH

This parameter is used to request the graph that shows Page Moves from Central to Expanded Storage.

Options

CENTRAL TO EXPANDED STORAGE ACTIVITY GRAPH=YES (produce the graph)

CENTRAL TO EXPANDED STORAGE ACTIVITY GRAPH=NO (don't produce the graph)

Default option

CENTRAL TO EXPANDED STORAGE ACTIVITY GRAPH=NO

EXPANDED-TO-CENTRAL STORAGE ACTIVITY GRAPH

This parameter is used to request the graph that shows Page Moves from Expanded to Central Storage.

Options

EXPANDED TO CENTRAL STORAGE ACTIVITY GRAPH=YES (produce the graph)

EXPANDED TO CENTRAL STORAGE ACTIVITY GRAPH=NO (don't produce the graph)

Default option

EXPANDED TO CENTRAL STORAGE ACTIVITY GRAPH=NO

EXPANDED-TO-AUXILIARY STORAGE ACTIVITY GRAPH

This parameter is used to request the graph that shows Page Outs from Expanded to Auxiliary Storage through Central Storage.

Options

EXPANDED TO AUXILIARY STORAGE ACTIVITY GRAPH=YES (produce the graph)

EXPANDED TO AUXILIARY STORAGE ACTIVITY GRAPH=NO (don't produce the graph)

Default option

EXPANDED TO AUXILIARY STORAGE ACTIVITY GRAPH=NO

CENTRAL-TO-AUXILIARY STORAGE ACTIVITY GRAPH

This parameter is used to request the graph that shows Page Outs from Central to Auxiliary Storage.

Options

CENTRAL TO AUXILIARY STORAGE ACTIVITY GRAPH=YES (produce the graph)

CENTRAL TO AUXILIARY STORAGE ACTIVITY GRAPH=NO (don't produce the graph)

Default option

CENTRAL TO AUXILIARY STORAGE ACTIVITY GRAPH=NO

AUXILIARY-TO-CENTRAL STORAGE ACTIVITY GRAPH

This parameter is used to request the graph that shows Page Ins from Auxiliary to Central Storage.

Options

AUXILIARY TO CENTRAL STORAGE ACTIVITY GRAPH=YES (produce the graph)

AUXILIARY TO CENTRAL STORAGE ACTIVITY GRAPH=NO (don't produce the graph)

Default option

AUXILIARY TO CENTRAL STORAGE ACTIVITY GRAPH=NO

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems.

TOP LABEL

Example

TOP LABEL=YES

This statement causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time:End Time as specified
- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

Example

FOOT1=UP TO 64 CHAR FOOTNOTE

This statement is used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks), that is included on all reports where there is room.

FOOT2

Example

FOOT2=UP TO 64 CHAR 2ND FOOTNOTE

This statement is used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks), that is included on all reports where there is room.

Additional Report Options

It is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it is useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report is prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****" or replaced by substitution - see below). Finally, the statement:

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE or EXCLUDE facilities, however it does not preclude their use. The combination of the various EXCLUDE/INCLUDE features in conjunction with the Generic Element Mask is discussed below.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a GENERIC ELEMENT MASK=11110011. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL**. If GENERIC ELEMENT MASK=11111111 were to be specified, all records would be summarized into a single entry or be subject to substitution as discussed below.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the report processing programs to intercept the key of each element, to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The `INCLUSIVE ELEMENT MASK` is used with an `INCLUDE` data set specified in the JCL by the `//INCLUDE DD` statement. The format of the `INCLUDE` entries are:

```
//INCLUDE DD *  
xxxxxxxx,xxxxxxxx,xxxxxxxx,xxxxxxxx, . . . .
```

where `xxxxxxxx` is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the `INCLUDE` data set would contain "????23??. The "?" characters in the `INCLUDE` selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The `EXCLUSIVE ELEMENT MASK` keyword phrase causes the report processing programs to intercept the key of each element, to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the `EXCLUDE` selection module. In this case, the key is restored to its original value once the `EXCLUDE` selection processing is completed.

The `EXCLUSIVE ELEMENT MASK` parameter differs from the `GENERIC ELEMENT MASK` in that the `EXCLUSIVE ELEMENT MASK` is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the `GENERIC ELEMENT MASK` is used.

The mask specification procedure is the same as for the `GENERIC ELEMENT MASK` described above.

The `EXCLUSIVE ELEMENT MASK` is used with an `EXCLUDE` data set specified in the JCL by the `//EXCLUDE DD` statement. The format of the `EXCLUDE` entries are:

```
//EXCLUDE DD *  
xxxxxxxx,xxxxxxxx,xxxxxxxx,xxxxxxxx, . . . .
```

where `xxxxxxxx` is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the `EXCLUDE` data set would contain "????23??. The "?" characters in the `EXCLUDE` selection statements correspond to the "*" positions in the modified keys.

Use of GENERIC ELEMENT MASK with EXCLUDE/INCLUDE and/or INCLUSIVE or EXCLUSIVE ELEMENT MASK

The following discussion proscribes the rules for using the various element masks and the order in which selection and substitution is performed.

- EXCLUSIVE ELEMENT MASK and the INCLUSIVE ELEMENT MASK are mutually exclusive as are the EXCLUDE and the INCLUDE DD statements.
- EXCLUSIVE ELEMENT MASK and INCLUSIVE ELEMENT MASK must be used in conjunction with the EXCLUDE or INCLUDE DD statement respectively to define the key elements that are to be either excluded or included. Without the EXCLUDE or INCLUDE lists defined through the use of the DD statements, the EXCLUSIVE and the INCLUSIVE ELEMENT MASKs are meaningless.
- The EXCLUDE and INCLUDE lists defined through the use of the EXCLUDE and INCLUDE DD statements can be used in conjunction with the GENERIC ELEMENT MASK to limit the data selected for inclusion into the reports irrespective of whether an EXCLUSIVE or INCLUSIVE ELEMENT MASK is specified.
- EXCLUSIVE and INCLUSIVE ELEMENT MASKING is performed prior to the EXCLUDE/INCLUDE logic.
- GENERIC ELEMENT MASKING takes place after the EXCLUDE/INCLUDE logic is performed.
- If the GENERIC ELEMENT MASK is specified as all "1"s (11111111), then:
 - If a CPPR.PARMLIB member &sidJGRP is present, the substitutions are made based upon the data contained in &sidJGRP.
 - If the &sidJGRP member is not present in CPPR.PARMLIB, then the entire key is set to "*"s and all records would be summarized into a single entry.

Please refer to Appendix A for a description of the CPPR.PARMLIB member &sidJGRP. A sample of the &sidJGRP member is found in the CPPR.PARMLIB as member CPPRJGRP.

Trends Analysis

Approach

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. The elements that have been chosen for the comparison are the following:

- CPU Utilization Statistics, including a breakdown by the installation-defined categories discussed in Section 3 above (TSO, BATCH, O/S, ONLINE, DATABASE, NETWORK)
- Paging Activity Statistics, including a breakdown by the installation-defined categories (TSO, BATCH, ONLINE, DATABASE, NETWORK)
- TSO Related statistics, including average prime shift response, number of active users and total trivial transactions
- The average number of Batch jobs through the system each day
- DASD Statistics, including (for all DASD Devices greater than 1% busy) Device percent busy, In-storage queue delay factor, DASD I/O Service time and, for all DASD in the Data center, the average number of I/O operations per second during prime shift.
- For all channels attached to a system, the average percent busy during prime shift

This information is gathered and presented in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right hand side of the report
- As a series of data points suitable for graphing or further analysis with workstation (PC-based) programs. The manner in which the data is down loaded is left to you. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Operating JCL for Trends Reports

The operating JCL has been designed to conform to IBM system utility JCL conventions wherever possible. Input comes from the Consolidated Table File (ONLINE).

ONLINE

This DD Statement specifies the data set name and the disposition of the permanent file used to save (ONline Performance Data Base) activity records. It is built by the Data Reduction Program (SSA1WKLD).

The EXEC statement references:

- PGM=SSA1TRND for the Trends Analysis Summary program
- PGM=SSA1TGRF for the Trends Analysis Graph program

The DD Statements required for program execution are described below.

INDEX

This DD Statement specifies the data set name and disposition of the CPPR ONLINE Table Index dataset.

CPPRERT

This DD Statement specifies the data set name and the disposition of the CPPR Element Registration Table dataset.

CPPRPARM

This DD Statement specifies the data set name and disposition of the CPPR Parameter Library.

HGDLIB

This DD Statement specifies the data set name and the disposition of the Partitioned Data set used to save the members produced by the Trends Analysis Graphics program (SSA1TGRF). These members are later down loaded to a Personal Computer and used as input to the Harvard Graphics (or other graphing) program to build the actual graphs.

SYSPRINT

This DD statement specifies the report output file. Unless otherwise specified, the DCB characteristics are as follows:

RECFM=FBA

BLKSIZE=133

LRECL=133

SYSNAP

This DD Statement specifies an optional print file for diagnostic hexadecimal snap dumps.

SYSIN

This is the parameter input file. For detailed information regarding the input parameters, see the section following.

Input Parameter Statements for Trends Reports

Selected System

This parameter is used to identify the SID of the System for which Trends are to be analyzed. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Options

This parameter has the following options:

- The SID itself (e.g, `SELECTED SYSTEM=MVSA`)
- An '*', indicating the SID for the system on which the program is being executed (e.g, `SELECTED SYSTEM=*`)

Default Option

`SELECTED SYSTEM=*`

1st Period Begin Date

This parameter is used to specify the beginning date of the first period (for the Trends Analysis Summary Program) or the beginning date for the entire period (for the Trends Analysis Graph Program) for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify a beginning date in one of the following forms:

- YYYYDDD (Standard Julian format)
- MM/DD/YYYY (USA Standard Gregorian format)
- DD.MM.YYYY (European Standard Gregorian format)
- *, *-n (relative day format where *=today)

Note • All dates must be specified in one of the above formats.

Default Date

`1ST PERIOD BEGIN DATE=00001`

1st Period End Date

This parameter is used to specify the ending date of the first period (for the Trends Analysis Summary Program) or the ending date for the entire period (for the Trends Analysis Graph Program) for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify an ending date in one of the formats specified above.

Default Date

1ST PERIOD END DATE=2099365

2nd Period Begin Date

This parameter is used to specify the beginning date of the second period (for the Trends Analysis Program) for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify a beginning date in one of the formats illustrated above

Default Date

2ND PERIOD BEGIN DATE=00001

2nd Period End Date

This parameter is used to specify the ending date of the second period (for the Trends Analysis Summary Program) for which Consolidated Table information is to be selected for analysis.

Specify Date

Specify an ending date.

Default Date

2ND PERIOD END DATE=2099365

Prime Shift First Hour

This parameter is used to specify the first hour of Prime Shift for analysis purposes.

Specify Hour

Specify an hour using the 24 hour system (e.g, PRIME SHIFT FIRST HOUR=7)

Default Hour

PRIME SHIFT FIRST HOUR=8

Late Shift First Hour

This parameter is used to specify the first hour of Late Shift for analysis purposes.

Specify Hour

Specify an hour using the 24 hour system (e.g, LATE SHIFT FIRST HOUR=17)

Default Hour

LATE SHIFT FIRST HOUR=16

Sample JOBSTREAMS for Trends Reports

Trends Analysis Report

The following is an example of a Jobstream that would produce the Trends Analysis Report, comparing the period beginning on the 13th of June and ending on the 17th of June to the period beginning on the 11th of January and ending on the 15th of January. Prime shift is defined as beginning at 7 in the morning and continuing up to 7 at night. The system for which the comparison is to be made is IP01.

```
//SMFJOB JOB (aaaaaaaa)
//STA EXEC PGM=SSAITRND,REGION=4096K
//STEPLIB DD DSN=CPPR.LOADLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//INDEX DD DSN=CPPR.INDEX,DISP=SHR
//ONLINE DD DSN=CPPR.ONLINE,DISP=SHR
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
1ST PERIOD BEGIN DATE=01/11/2005
1ST PERIOD END DATE=01/15/2005
2ND PERIOD BEGIN DATE=06/13/2005
2ND PERIOD END DATE=06/17/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
```

ISPF/PDF Workload Subsystem Interface

The CIMS Capacity Planner Workload subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM Graphs, generate the full range of batch reports to be printed on the system printer, and generate input to the PC Presentation Graphics subsystem (Harvard Graphics or Excel).

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, CICS, etc.).

On-line Graphs

The following Workload graphs are available on-line through the ISPF Interface:

CPU and Paging

The CPU and Paging graph depicts the proportional use of a CPU by the various elements within a selected system (System Control Program, TSO, and Batch processing), along with a line indicating the paging rate in number of pages per second. The graph shows the average utilization per hour within the range of dates specified in the graph request.

DASD and CPU Graph

The DASD and CPU graph depicts, for one or more selected MVS systems, the percent busy for a given device or set of devices, combined with the average percent CPU utilization per hour within the range of dates specified in the graph request.

CPU Percent Busy

The CPU Percent Busy graph depicts CPU activity by the various system components (System Control Program, Batch, TSO, On-line, Data Base, and Network) as the percent of the time that the CPU is processing each of the components of the workload. The graph will show the average utilization percentages by the hour, day, week, or month (depending upon the period requested) within the range of dates specified in the graph request.

Paging Activity

The Paging Activity graph is a stacked bar graph that depicts the average paging rates in pages per second by the system components (System Control Program, Batch and Started Tasks, TSO, On-line, Data Base, and Network) for the selected system within the range of dates specified in the graph request.

TSO Response

The TSO Response graph depicts the average response time for TSO trivial transactions for a selected system during the Prime Shift within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

Batch Throughput

The Batch Throughput graph depicts the total number of jobs processed by a selected system, divided into the six groups specified in the "&sidJOBS" member defined in the CPPR Parmlib (see the *CIMS Capacity Planner Installation and Getting Started Guide*) within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

DASD Percent Busy

The DASD Percent Busy graph depicts the percent of the time that each DASD device was busy for a selected system within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

Channel Percent Busy

The Channel Percent Busy graph depicts the percent of the time each channel was in use for a selected system within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

Channel and CPU

The Channel and CPU graph depicts the percent of the time that a selected channel was busy along with the percent of the time the CPU was busy for a selected system within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

ESA CS-to-ES Activity

The ESA CS-to-ES Activity graph depicts the average number of pages moved from central storage to expanded storage per second in a specified ESA system within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

ESA ES-to-AUX Activity

The ESA CS-to-AUX Activity graph depicts the average number of pages moved from expanded storage to auxiliary storage per second in a specified ESA system within the range of dates specified in the graph request. The graph will show the Prime Shift averages by hour, day, week, or month, depending upon the period specified in the graph request.

ESA Storage Activity

The ESA Storage Activity graph is a bar chart that shows Page Movement within an ESA environment in terms of Private Reads, Private Writes, Hyperspace Reads, Hyperspace Writes, VIO Reads and VIO Writes. All values are given in Pages per second.

PR/SM LPAR Activity

This bar graph shows the percent busy of the PR/SM SYSPLEX broken down by LPAR. PR/SM overhead is also shown as a separate category.

Jobname: CPU Activity

This pie chart shows the percent of the CPU that is being used by different address spaces. A variable allows you to specify the percent busy threshold (e.g, 3%) above which individual address spaces are shown. For example, if you specify 3%, all address spaces that consumed, on average, 3 or more percent of the CPU is shown. The time frame is bounded by specifying a Begin Time/End Time pair. This graph relies on the presence of SMF Type 30 Interval Accounting Records in the original Input to the SSA1WKLD Data Reduction program.

On-line Reports

The following reports are produced on-line through the CPPR ISPF Interface.

Statistical Reports

Job Statistics Report

The Job Statistics Report provides a list of the 40 jobs that are run most frequently, exhibit the highest residency time, and consume the most CPU time during the measured period.

- Elapsed time:

The report identifies the 40 jobs that exhibit the highest amount of residence time in the CPU. It lists the job name, the total residency time, in hours, minutes, and seconds, and the relative percent of the time that the Job was resident on the mainframe.

- Execution frequency:

The report lists the 40 jobs that were run most frequently during the measured period. It lists the job name, the number of times that the job was executed, and the relative percent of the total number of jobs executed.

- CPU time:

The report lists the 40 jobs that consumed the most CPU time during the measured period. The report lists the job names, the total CPU time consumed, in hours, minutes, and seconds, and the respective percent of the total CPU time consumed by each job.

Program Statistics Report

The Program Statistics Report provides a list of the 40 programs that are run most frequently, exhibit the highest residency time, and consume the most CPU time during the measured period.

■ **Elapsed time:**

The report identifies the 40 programs that exhibit the highest amount of residence time in the CPU. It lists the program name and the total residency time.

■ **Execution frequency:**

The report lists the 40 programs that were run most frequently during the measured period. It lists the program name, the number of times that the program was executed, and the relative percent of the total number of programs executed.

■ **CPU time:**

The report lists the 40 programs that consumed the most CPU time during the measured period. The report lists the program name, the total CPU time consumed, and the respective percent of the total CPU time consumed, in hours, minutes, and seconds, by each program.

TSO Command Statistics Report

The TSO Command Statistics Report provides a list of the 40 TSO commands that are most frequently executed and consume the most CPU time during the measured period.

■ **Execution frequency:**

The report lists the 40 TSO commands that were run most frequently during the measured period. It lists the command name, the number of times that the command was executed, and the relative percent of the total number of commands executed.

■ **CPU time:**

The report lists the 40 TSO commands that consumed the most CPU time during the measured period. The report lists the command name, the total CPU time consumed, and the respective percent of the total CPU time consumed, in hours, minutes, and seconds, by each command.

TSO User Statistics Report

The TSO User Statistics Report provides a list of the 40 TSO Users who executed most TSO commands, exhibit the highest connect time, and consume the most TSO CPU time during the measured period.

- **Connect time:**

The report identifies the 40 TSO users that consumed the highest amount of TSO connect time. It lists the User ID, the amount of time you were connected, in hours, minutes, and seconds, and the relative percentage of the TSO connect time consumed by all the TSO users.

- **Commands executed:**

The report lists the 40 TSO users that executed the highest number of TSO commands during the measured period. It lists the User ID, the number of commands that the program was executed, and the relative percent of the total number of TSO commands executed.

- **CPU time:**

The report lists the 40 TSO users that consumed the most TSO CPU time during the measured period. The report lists the User ID, the total CPU time consumed, in hours, minutes, and seconds, and the respective percentage of the total amount of CPU time consumed by each user.

Printer Statistics Report

The Printer Statistics Report provides a list of the 40 printers that print the most lines, print the most pages, and print the largest number of data sets during the measured period.

- **Lines printed:**

The report identifies the 40 printers that printed the highest number of lines. It lists the printer name, the total number of lines printed, and the relative percentage of the total lines printed by all printers.

- **Pages printed:**

The report lists the 40 printers that printed the most pages during the measured period. It lists the printer name, the number of pages printed, and the relative percent of the total number of pages printed.

- **Data sets printed:**

The report lists the 40 printers that processed the most printer data sets during the measured period. The report lists the printer name, the number of data sets printed, and the respective percent of the total number of data sets printed.

Summary Analysis Report

The on-line Summary Analysis report provides the following information related to the workload:

- TSO LOGON activity and TSO average response time
- Peak paging Activity
- Total jobs executed including a breakdown by class and shift
- Total TSO transactions executed including a breakdown by transaction size and shift
- Percent CPU busy by shift
- Percent CPU busy by shift for weekdays only
- Percent paging activity by shift
- Average DASD device busy percentage by shift (top 10 DASD devices)
- Average DASD delay time by shift in milliseconds (top 10 DASD devices)
- Average DASD I/O activity per string by shift (top 10 DASD strings)
- Average channel busy percentage by shift (top 10 channels)

Trends Analysis Report

The on-line Trend Analysis report compares the prime shift workload for two time periods specified by you, showing changes in the workload as percentage values for CPU utilization, paging activity, TSO activity, number of batch jobs executed, DASD utilization, and channel utilization.

Exception Analysis Reports

Processor Exception Report

The Processor Exception analysis report shows, for the period selected, any 15 minute periods during which the CPU busy percentage or the paging rate exceeded the limits specified in the CPPR.PARMLIB member &sidXCPT.

DASD Exception Report

This report shows, for the period selected, any one hour intervals in which the DASD performance limits specified in the CPPR.PARMLIB member &sidXCPT were exceeded. The limits that are screened are:

- DASD queue delay (in milliseconds)
- DASD percentage busy
- DASD service time (in milliseconds)

Channel Exception Report

This report shows, for the specified period, any one-hour intervals during which the channel busy percentage limits specified in the `CPPR.PARMLIB` member `&sidXCPT` were exceeded.

Ad Hoc Reports

The Ad Hoc reports provide the ability to extract various detail data directly from SMF files. This differs from the other `CPPR` reporting functions that take their data from the Performance Data Base. The Ad Hoc report requests are constructed by the `ISPF` interface, but are submitted as batch jobs, unlike the other reports submitted through the on-line facility that are produced within your own `TSO` address space and displayed at the terminal.

Job Name Ad Hoc Report

This report contains detail data extracted from the SMF input that relates to a set of specified Jobs or Started Tasks that are selected by specifying the job names to be selected and the dates that bound the period to be considered.

The information reported is:

- System ID (SMF SID)
- Job Name
- Date-in
- Date-out
- Time-in
- Last completion code
- CPU time (seconds)
- DASD EXCP's
- Job class
- Accounting data

Program Name Ad Hoc Report

This report contains detail data extracted from the SMF input that relates to a set of specified Programs that are selected by specifying the Programs to be selected and the dates that bound the period to be considered. The selected programs can be from the execution of batch jobs, started tasks, or TSO commands.

The information reported is:

- System ID (SMF SID)
- Job name
- Program name
- Performance group
- Date-in
- Date-out
- Time-in
- Time-out
- Completion code
- CPU time (seconds)
- DASD EXCP's
- Paging activity
- Job class

TSO User Ad Hoc Report

This report contains detail data extracted from the SMF input that relates to a set of TSO users that are selected by specifying the TSO User ID(s) to be selected and the dates that bound the period to be considered. The information reported is:

- System ID (SMF SID)
- TSO User ID
- Date and time logged-on
- Date and time logged-off
- CPU time (seconds)
- # TSO command executed

Data Set Name Ad Hoc Report

This report contains detail data extracted from the SMF input that relates to a group of data sets that are selected by specifying the data set name(s) to be selected and the dates that bound the period to be considered. The information reported is:

- System ID (SMF SID)
- Data set name
- Date and time
- Job name or TSO USERID
- Data Set Disposition
- Data set function that took place
- Volume serial number
- Record format
- Logical record length
- Data set organization
- EXCP count

Batch Reports

A Batch Job submission panel is provided so that the full range of batch reports and graphs described earlier can be submitted through the ISPF Interface. This considerably simplifies the operation of CPPR.

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch jobs instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch job
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- Saving the Chart as a GDF File (Batch)

These options are explained in detail below.

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual workstation-based charts. As originally designed, the data members are created by Batch jobs, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through the HGDLIB. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch job is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of several key phrases.

```
GDDM BATCH PROCESS=YES
```

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

```
GDDM PRINTER NAME=pppppppp
```

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQ DD Statement in the JCL.

Several members have been added to the Distribution job file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch job. The member named "GDDMBATW" shows how to produce any of the Workload GDDM Graphs in Batch.

GDDMBAT CNTL Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (CPPR201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST0001) WHICH MAY THEN BE
/* SENT TO A GDDM PRINTER (PRRINT01).
/* * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR, DSN=SYS1.GDDMSYM
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR, DSN=SYS1.GDDMGDF
/* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR, DSN=SYS1.GDDM.REQUEST.QUEUE
//ADMCDATA DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00) '
//SYSIN   DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*          | | | % CHARACTER LENGTH
*          | | % CHARACTER WIDTH
*          | HORIZONTAL OFFSET
*          | VERTICAL OFFSET
//INCLUDE DD *
CPPR201
```

GDDMBATW CNTL Member

```
//SSAGDDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR ALL OF
/* THE SUPPORTED WORKLOAD CHARTS AND SENDING THEM TO A PRINTER.
/* * * * * *
//STO      EXEC PGM=IKJEFT01,REGION=5000K,DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//        DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
/* THE FOLLOWING DDNAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING DDNAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN  DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD01)' /* CPU AND PAGING /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD02)' /* DASD AND CPU /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD03)' /* CPU BUSY /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD04)' /* PAGING ACTIVITY /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD05)' /* TSO RESPONSE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD06)' /* BATCH THROUGHPUT /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD07)' /* DASD BUSY /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD08)' /* CHANNEL BUSY /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD09)' /* CHANNEL AND CPU /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD0a)' /* ESA CS=>ES /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD0b)' /* ESA ES=>AUX /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD0c)' /* ESA STORAGE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD0d)' /* PR/SM LPAR /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD0e)' /* JOBNAME:CPU /*
//SYSIN    DD *
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*          | | | % CHARACTER LENGTH
*          | | % CHARACTER WIDTH
*          | HORIZONTAL OFFSET
*          VERTICAL OFFSET
```

Specifying GDDM Batch Print Controls

It is possible to specify certain Print Controls that are to take effect when printing a chart in Batch. The key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter allows you to specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

You might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. In order to invoke these local templates in a Batch job, a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter is added to the SYSIN, where the "nnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=dddddddd
```

where the "dddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeeee
```

where the "eeeeeeee" is the name under which the graph is saved in the GDF library.

Presentation Graphics

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics (release 2.3 or 3.0), plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, which describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on the CIMS product CD. The templates must be installed on the local workstation using the SETUP.EXE program located in the CIMSPPR/HARGRAPH/DISK1 directory.

Graph Production Procedures

The process of generating the CIMS Capacity Planner WORKLOAD graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in a subsequent section.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.
- The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters in the form:

"&sidpiiq" where:

"&sid" is the SMF SID for the MVS system to which the data pertains

"p" is the time period specified for the graph. The possible values are:

"1" - Hourly

"2" - Daily

"3" - Weekly

"4" - Monthly

"iiq" is the graph identifier that consists of two alpha-numeric characters. The "piiq" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under paragraph 8.2.3. The qualifier "q" might or might not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "CSV" is appended to conform to the following format:

&sidpiiq.CSV, where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Graph Periods

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide* for a description of the hourly, daily, weekly, and monthly graph periods

Graph Descriptions

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide* for descriptions of each of the graphs produced by the CIMS Capacity Planner Workload Subsystem.

Graphic Report Tables

The following tables summarize the relationship of members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-based graph program.

Workload Subsystem Hourly Graphs

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGHY	&sid124Y	(none)	TSO Volume by PGN
	&sid125Y	(none)	TSO Average Response by PGN
SSA1TGRA	MISP11Q (this is an XFRLIB member)	CPUTABLE	CPU Utilization in MIPS by SMF SID
SSA1TGRD	&sid133	&sidJGRP	CPU by Organization
SSA1TGRE	&sid1E0P	&sid1CPUV	CPU Busy - Prime Shift
SSA1TGRF	&sid101	(none)	CPU Utilization by Category
	&sid102	(none)	CPU Utilization by Category 1
	&sid103	(none)	CPU Utilization by Category 2
	&sid104	(none)	CPU Utilization by Category 3
	&sid105	(none)	CPU Utilization by Category 4
	&sid106	(none)	CPU Utilization by Category 5
	&sid107	(none)	CPU Utilization by Category 6
	&sid110	(none)	Paging Activity for Category
	&sid111	(none)	Paging Activity for Category 1
	&sid112	(none)	Paging Activity for Category 2
	&sid113	(none)	Paging Activity for Category 3
	&sid114	(none)	Paging Activity for Category 4
	&sid115	(none)	Paging Activity for Category 5

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRF (continued)	&sid120	(none)	TSO Activity
	&sid121	(none)	TSO Response
	&sid122	(none)	TSO Transaction Volume
	&sid130	(none)	Batch Volume
	&sid1301	&sidJOBS	Selected Batch Volume
	&sid1302	&sidJOBS	Selected Batch Volume
	&sid1303	&sidJOBS	Selected Batch Volume
	&sid1304	&sidJOBS	Selected Batch Volume
	&sid1305	&sidJOBS	Selected Batch Volume
	&sid1306	&sidJOBS	Selected Batch Volume
	&sid131	(none)	Batch Turnaround
	&sid1311	&sidJOBS	Selected Batch Turnaround
	&sid1312	&sidJOBS	Selected Batch Turnaround
	&sid1313	&sidJOBS	Selected Batch Turnaround
	&sid1315	&sidJOBS	Selected Batch Turnaround
	&sid1316	&sidJOBS	Selected Batch Turnaround
	&sid140	(none)	DASD Device Busy
	&sid141	(none)	DASD Queue Delay
	&sid142	(none)	DASD I/O Service Time
	&sid143	(none)	DASD I/O Activity
&sid150	(none)	Channel Busy %	
SSA1TGRH	&sid123	(none)	TSO Volume Response by PGN

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRJ	&sid11M	&sidSVCN and CPutable	MIPS by Service Class Category (SVCN) with Capture Ratio as a Separate Column
	&sid11N		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU S/U
	&sid110		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of IOC S/U
	&sid11P		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU and IOC S/U
SSA1TGRK	&sid136	(none)	Task Schedule
	&sid137	(none)	Task Schedule %
SSA1TGRR	&sid10C	&sidPGNS	Private Storage Usage by PGN
	&sid11A	&sidPGNS	CPU SUs %
	&sid11B	&sidPGNS	MSO SUs
	&sid11C	&sidPGNS	IOC SUs
	&sid11D	&sidPGNS	CPU SUs Count
	&sid11E	&sidPGNS	MSO SUs Count
	&sid11F	&sidPGNS	IOC SUs
	&sid11G	&sidPGNX	% of CPU SUs by PGN
	&sid11H	&sidPGNX	% of MSO SUs by PGN
	&sid11I	&sidPGNX	% of IOC SUs by PGN
SSA1TGRS	&sid108	(none)	JOBNAME:CPU Activity
SSA1TGRU	&sid1X5	&sidCPUV	PR/SM Busy by Range
	&sid1X6	(none)	PR/SM Busy by LPAR
	&sid1X7	(none)	PR/SM SYSPLEX CPU Hi, Av, Lo % Busy
	&sid1X8	(none)	PR/SM SYSPLEX Peak, Prime, Off Shift average % Busy

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRV	&sid135	(none)	Task Availability
SSA1TGRX	&sid1X0	(none)	CS:ES
	&sid1X1	(none)	ES:AUX
	&sid1X2	(none)	ES Storage Activity
	&sid1X3	(none)	Central Storage Map
	&sid1X4	(none)	Expanded Storage Map
SSA1WKLK	&sid1G6	(none)	Performance Index
	&sid1G7	(none)	Transactions Processed
	&sid1G8	(none)	S/U(K) Used
	&sid1G9	(none)	Average Response

&sid is the SMF System ID (e.g, CPPR)

Workload Subsystem Daily Graphs

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGHY	&sid224Y	(none)	TSO Volume by PGN
	&sid225Y	(none)	TSO Average Response by PGN
SSA1TGRA	MISP11Q (this is an XFRLIB member)	CPUTABLE	CPU Utilization in MIPS by SMF SID
SSA1TGRD	&sid233	&sidJGRP	CPU by Organization
SSA1TGRE	&sid2E0P	&sid1CPUV	CPU Busy - Prime Shift
SSA1TGRF	&sid201	(none)	CPU Utilization by Category
	&sid202	(none)	CPU Utilization by Category 1
	&sid203	(none)	CPU Utilization by Category 2
	&sid204	(none)	CPU Utilization by Category 3
	&sid205	(none)	CPU Utilization by Category 4
	&sid206	(none)	CPU Utilization by Category 5
	&sid207	(none)	CPU Utilization by Category 6
	&sid210	(none)	Paging Activity by Category
	&sid211	(none)	Paging Activity by Category 1
	&sid212	(none)	Paging Activity by Category 2
	&sid213	(none)	Paging Activity by Category 3
	&sid214	(none)	Paging Activity by Category 4
	&sid215	(none)	Paging Activity by Category 5

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRF (continued)	&sid220	(none)	TSO Activity
	&sid221	(none)	TSO Response
	&sid222	(none)	TSO Transaction Volume
	&sid230	(none)	Batch Volume
	&sid2301	&sidJOBS	Selected Batch Volume
	&sid2302	&sidJOBS	Selected Batch Volume
	&sid2303	&sidJOBS	Selected Batch Volume
	&sid2304	&sidJOBS	Selected Batch Volume
	&sid2305	&sidJOBS	Selected Batch Volume
	&sid2306	&sidJOBS	Selected Batch Volume
	&sid231	(none)	Batch Turnaround
	&sid2311	&sidJOBS	Selected Batch Turnaround
	&sid2312	&sidJOBS	Selected Batch Turnaround
	&sid2313	&sidJOBS	Selected Batch Turnaround
	&sid2314	&sidJOBS	Selected Batch Turnaround
	&sid2315	&sidJOBS	Selected Batch Turnaround
	&sid2316	&sidJOBS	Selected Batch Turnaround
	&sid240	(none)	DASD Device Busy
	&sid241	(none)	DASD Queue Delay
	&sid242	(none)	DASD I/O Service Time
	&sid243	(none)	DASD I/O Activity
	&sid250	(none)	Channel Busy %
	SSA1TGRH	&sid223	(none)

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRJ	&sid21M	&sidSVCN and CPUTABLE	MIPS by Service Class Category (SVCN) with Capture Ratio as a Separate Column
	&sid21N		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU S/U
	&sid210		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of IOC S/U
	&sid21P		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU and IOC S/U
SSA1TGRK	&sid236	(none)	Task Schedule
	&sid237	(none)	Task Schedule %
SSA1TGRR	&sid20C	&sidPGNS	Private Storage Usage by PGN
	&sid21A	&sidPGNS	CPU SUs %
	&sid21B	&sidPGNS	MSO SUs
	&sid21C	&sidPGNS	IOC SUs
	&sid21D	&sidPGNS	CPU SUs Count
	&sid21E	&sidPGNS	MSO SUs Count
	&sid21F	&sidPGNS	IOC SUs
	&sid21G	&sidPGNX	% of CPU SUs by PGN
	&sid21H	&sidPGNX	% of MSO SUs by PGN
	&sid21I	&sidPGNX	% of IOC SUs by PGN
SSA1TGRS	&sid208	(none)	JOBNAME:CPU Activity
SSA1TGRU	&sid2X5	&sidCPUV	PR/SM Busy by Range
	&sid2X6	(none)	PR/SM Busy by LPAR
	&sid2X7	(none)	PR/SM SYSPLEX CPU Hi, Avg, Lo % Busy
	&sid2X8	(none)	PR/SM SYSPLEX Peak, Prime, Off Shift Average % Busy

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRV	&sid235	(none)	Task Availability
SSQ1TGRX	&sid2X0	(none)	CS:ES
	&sid2X1	(none)	ES:AUX
	&sid2X2	(none)	ES Storage Activity
	&sid2X3	(none)	Central Storage Map
	&sid2X4	(none)	Expanded Storage Map
SSA1WKLG	&sid2G6	(none)	Performance Index
	&sid2G7	(none)	Transactions Processed
	&sid2G8	(none)	S/U(K) Used
	&sid2G9	(none)	Average Response

&sid is the SMF System ID (e.g, CPPR)

Workload Subsystem Weekly Graphs

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGHY	&sid324Y	(none)	TSO Volume by PGN
	&sid325Y	(none)	TSO Average Response by PGN
SSA1TGRA	MISP11Q (this is an XFRLIB member)	CPUTABLE	CPU Utilization in MIPS by SMF SID
SSA1TGRD	&sid333	&sidJGRP	CPU by Organization
SSA1TGRE	&sid3E0E	&sidCPUV	CPU Busy - Early Shift
	&sid3E0P	&sidCPUV	CPU Busy - Prime Shift
	&sid3E0L	&sidCPUV	CPU Busy - Late Shift
SSA1TGRF	&sid301	(none)	CPU Utilization by Category
	&sid302	(none)	CPU Utilization by Category 1
	&sid303	(none)	CPU Utilization by Category 2
	&sid304	(none)	CPU Utilization by Category 3
	&sid305	(none)	CPU Utilization by Category 4
	&sid306	(none)	CPU Utilization by Category 5
	&sid307	(none)	CPU Utilization by Category 6
	&sid310	(none)	Paging Activity by Category
	&sid311	(none)	Paging Activity by Category 1
	&sid312	(none)	Paging Activity by Category 2
	&sid313	(none)	Paging Activity by Category 3
	&sid314	(none)	Paging Activity by Category 4
	&sid315	(none)	Paging Activity by Category 5

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRF (continued)	&sid320	(none)	TSO Activity
	&sid321	(none)	TSO Response
	&sid322	(none)	TSO Transaction Volume
	&sid330	(none)	Batch Volume
	&sid3301	&sidJOBS	Selected Batch Volume
	&sid3302	&sidJOBS	Selected Batch Volume
	&sid3303	&sidJOBS	Selected Batch Volume
	&sid3304	&sidJOBS	Selected Batch Volume
	&sid3305	&sidJOBS	Selected Batch Volume
	&sid3306	&sidJOBS	Selected Batch Volume
	&sid331	(none)	Batch Turnaround
	&sid3311	&sidJOBS	Selected Batch Turnaround
	&sid3312	&sidJOBS	Selected Batch Turnaround
	&sid3313	&sidJOBS	Selected Batch Turnaround
	&sid3314	&sidJOBS	Selected Batch Turnaround
	&sid33115	&sidJOBS	Selected Batch Turnaround
	&sid33116	&sidJOBS	Selected Batch Turnaround
	&sid340	(none)	DASD Device Busy
	&sid341	(none)	DASD Queue Delay
	&sid342	(none)	DASD I/O Service Time
	&sid343	(none)	DASD I/O Activity
	&sid350	(none)	Channel Busy %
	SSA1TGRH	&sid323	(none)

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRJ	&sid31M	&sidSVCN and CPutable	MIPS by Service Class Category (SVCN) with Capture Ratio as a Separate Column
	&sid31N		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU S/U
	&sid31O		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of IOC S/U
	&sid31P		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU and IOC S/U
SSA1TGRK	&sid336	(none)	Task Schedule
	&sid337	(none)	Task Schedule %
SSA1TGRR	&sid31A	&sidPGNS	CPU SUs %
	&sid31B	&sidPGNS	MSO SUs
	&sid31C	&sidPGNS	IOC SUs
	&sid31D	&sidPGNS	CPU SUs Count
	&sid31E	&sidPGNS	MSO SUs Count
	&sid31F	&sidPGNS	IOC SUs Count
	&sid31G	&sidPGNX	% of CPU SUs by PGN
	&sid31H	&sidPGNX	% of MSO SUs by PGN
SSA1TGRS	&sid308	(none)	JOBNAME:CPU Activity
	&sid31I	&sidPGNX	% of IOC SUs by PGN
SSA1TGRU	&sid3X5	&sidCPUV	PR/SM Busy by Range
	&sid3X6	(none)	PR/SM Busy by LPAR
	&sid3X7	(none)	PR/SM SYSPLEX CPU Hi, Avg, Lo % Busy
	&sid3X8	(none)	PR/SM SYSPLEX Peak, Prime, Off Shift average % Busy
SSA1TGRV	&sid335	(none)	Task Availability

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRX	&sid3X0	(none)	CS:ES
	&sid3X1	(none)	ES:AUX
	&sid3X2	(none)	ES Storage Activity
	&sid3X3	(none)	Central Storage Map
	&sid3X4	(none)	Expanded Storage Map
SSA1WKLG	&sid3G6	(none)	Performance Index
	&sid3G7	(none)	Transactions Processed
	&sid3G8	(none)	S/U(K) Used
	&sid3G9	(none)	Average Response

&sid is the SMF System ID (e.g, CPPR)

Workload Subsystem Monthly Graphs

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSAITGHY	&sid424Y	(none)	TSO Volume by PGN
	&sid425Y	(none)	TSO Average Response by PGN
SSAITGRA	MISP11Q (this is an XFRLIB member)	CPUTABLE	CPU Utilization in MIPS by SMF SID
SSAITGRD	&sid433	&sidJGRP	CPU by Organization
SSAITGRE	&sid4E0E	&sidCPUV	CPU Busy - Early Shift
	&sid4E0P	&sidCPUV	CPU Busy - Prime Shift
	&sid4E0L	&sidCPUV	CPU Busy - Late Shift
SSAITGRF	&sid401	(none)	CPU Utilization by Category
	&sid402	(none)	CPU Utilization by Category 1
	&sid403	(none)	CPU Utilization by Category 2
	&sid404	(none)	CPU Utilization by Category 3
	&sid405	(none)	CPU Utilization by Category 4
	&sid406	(none)	CPU Utilization by Category 5
	&sid407	(none)	CPU Utilization by Category 6
	&sid410	(none)	Paging Activity by Category
	&sid411	(none)	Paging Activity by Category 1
	&sid412	(none)	Paging Activity by Category 2
	&sid413	(none)	Paging Activity by Category 3
	&sid414	(none)	Paging Activity by Category 4
	&sid415	(none)	Paging Activity by Category 5

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRF (continued)	&sid420	(none)	TSO Activity
	&sid421	(none)	TSO Response
	&sid422	(none)	TSO Transaction Volume
	&sid430	(none)	Batch Volume
	&sid4301	&sidJOBS	Selected Batch Volume
	&sid4302	&sidJOBS	Selected Batch Volume
	&sid4303	&sidJOBS	Selected Batch Volume
	&sid4304	&sidJOBS	Selected Batch Volume
	&sid4305	&sidJOBS	Selected Batch Volume
	&sid4306	&sidJOBS	Selected Batch Volume
	&sid431	(none)	Batch Turnaround
	&sid4311	&sidJOBS	Selected Batch Turnaround
	&sid4312	(none)	Batch Turnaround
	&sid4313	(none)	Batch Turnaround
	&sid4314	(none)	Batch Turnaround
	&sid4315	(none)	Batch Turnaround
	&sid4316	(none)	Batch Turnaround
	&sid440	(none)	DASD Device Busy
	&sid441	(none)	DASD Queue Delay
	&sid442	(none)	DASD I/O Service Time
	&sid443	(none)	DASD I/O Activity
	&sid450	(none)	Channel Busy %
	SSA1TGRH	&sid423	(none)

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRJ	&sid41M	&sidSVCN and CPUTABLE	MIPS by Service Class Category (SVCN) with Capture Ratio as a Separate Column
	&sid41N		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU S/U
	&sid41O		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of IOC S/U
	&sid41P		MIPS by Service Class Category (SVCN) Capture Ratio Spread by % of CPU and IOC S/U
SSA1TGRK	&sid435	(none)	Task Availability
	&sid436	(none)	Task Schedule
	&sid437	(none)	Task Schedule %
SSA1TGRR	&sid41A	&sidPGNS	CPU SUs %
	&sid41B	&sidPGNS	MSO SUs
	&sid41B	&sidPGNS	IOC SUs
	&sid41D	&sidPGNS	CPU SUs Count
	&sid41E	&sidPGNS	MSO SUs Count
	&sid41F	&sidPGNS	IOC SUs Count
	&sid41G	&sidPGNX	% OF CPU SUs by PGN
	&sid41H	&sidPGNX	% OF MSO SUs by PGN
	&sid41I	&sidPGNX	% OF IOC SUs by PGN
SSA1TGRS	&sid408	(none)	JOBNAME:CPU Activity
SSA1TGRU	&sid4X5	&sidCPUV	PR/SM Busy by Range
	&sid4X6	(none)	PR/SM Busy by LPAR
	&sid4X7	(none)	PR/SM SYSPLEX CPU Hi, Avg, Lo % Busy
	&sid4X8	(none)	PR/SM SYSPLEX Peak, Prime, Off Shift Average % Busy

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1TGRX	&sid4X0	(none)	CS:ES
	&sid4X1	(none)	ES:AUX
	&sid4X2	(none)	ES Storage Activity
	&sid4X3	(none)	Central Storage Map
	&sid4X4	(none)	Expanded Storage Map
SSA1WKLG	&sid4G6	(none)	Performance Index
	&sid4G7	(none)	Transactions Processed
	&sid4G8	(none)	S/U(K) Used
	&sid4G9	(none)	Average Response

&sid is the SMF System ID (e.g, CPPR)

DASM SUBSYSTEM

DASM Subsystem Overview	2-3
DATA REDUCTION	2-3
SMF Data	2-3
VTOC Data	2-3
DCOLLECT Data	2-4
Reports Produced During Data Reduction	2-6
DASM REPORT DATA	2-11
PARMLIB Data	2-11
DASM REPORTS	2-12
DASD Space Allocation—By Device Type (SSA1DASM or SSA1DASH)	2-12
DASD Space Allocation—By Device Volume (SSA1DASM or SSA1DASH)	2-12
DASD Volume Map by CCHH (SSA1DASM)	2-13
Volume Map by DSNAME (SSA1DASM)	2-13
Volume Activity Report (SSA1RPT)	2-13
Data Set Activity Report (SSA1RPT)	2-13
Data Set Detail Report (SSA1RPT)	2-13
DASM Volume Summary Report (SSA1DASV)	2-14
DASM Selected VSAM Summary Report (SSA1VSAM)	2-15
DASM VTOC Selected DSN on Volume Report (SSA1DASM)	2-16
DASM VTOC by CCHH on Volume Report (SSA1DASM)	2-18
DASM VTOC FREESPACE on Volume Report (SSA1DASM)	2-19
DASM VTOC Detailed Volume Space Summary Report (SSA1DASM)	2-20
DASM VTOC Quick Volume Space Summary Report (SSA1DASM)	2-21
DASM X37 Candidate List (SSA1DASM)	2-22
DASM Re-blocking Candidate List (SSA1DASM)	2-24
DASM Archival Candidate List (SSA1DASM)	2-26
DASM Release Candidate List (SSA1DASM)	2-27

OPERATING JCL	2-29
DD Statements	2-29
Parameters	2-31
Additional Report Options	2-35
SAMPLE JOBSTREAMS	2-39
Scanning the DASD Farm	2-39
Saving the DASM Tables	2-40
Reporting from the Online Tables	2-41
ISPF/PDF DASM Subsystem Interface	2-42
Reports	2-42
Presentation Graphics	2-43
Graph Production Procedures	2-43
Naming Conventions	2-44
Presentation Graphics Distribution Disk	2-46
Graph Periods	2-47
Graph Descriptions	2-47
Graphic Report Tables	2-48

DASM Subsystem Overview

The DASM Subsystem provides a series of reports that assist you in evaluating various aspects of how the DASD space within the Data Center is being allocated and used. To provide the required information, the DASM Subsystem provides:

- The names of the most frequently accessed DASD data sets along with their associated DASD devices
- The physical and logical manner in which the DASD volumes are laid out
- The most active DASD volumes in terms of data set access
- The identity of the organizations owning space on a specified set of DASD volumes
- The information required to determine whether a given volume or set of volumes is being used inefficiently to store information by a user group
- A variety of DASD Pool related information that considers the DASD volumes as members of a set of generic Pools and then reports on the availability and allocation characteristics of each individual DASD Pool.

DATA REDUCTION

The DASM Subsystem uses data derived from both SMF and the VTOCs of the various devices making up the DASD farm. Optionally, you might want to use DCOLLECT data in place of the VTOC data.

SMF Data

The DASM Subsystem uses data contained in the type 14, 15, 64, and 74 SMF records. The reduction of the SMF data is handled by the CIMS Capacity Planner Workload data reduction program at the time the Workload data is being reduced. The required data is written to the Online Performance Data Base for reference by the DASM Reporting Subsystem.

VTOC Data

The data set allocation data is derived by dynamically scanning the VTOCs of the volumes within the DASD farm. The allocation data is not automatically written to the Online Performance Data Base because you can choose to run the allocation reports more frequently than is necessary to maintain trending and historical DASD allocation data. Provision is made, however, for writing the DASD allocation data to the Online Performance Data Base on demand using the COMMIT parameter, which is described later. We recommend that the allocation data be committed to the Online Performance Data Base on a weekly basis, preferably on the same day of each week.

The following is an example of the Job stream to be used to scan the VTOCs and COMMIT the allocation tables to the Performance Data Base.

```
//jobname JOB (valid job card data)
//ST1      EXEC PGM=SSA1DASM,REGION=4096K,TIME=60
//STEPLIB DD DISP=SHR,DSN=CPPR.LOAD
//CPPRPARM DD DISP=SHR,DSN=CPPR.PARMLIB
//ONLINE   DD DISP=SHR,DSN=CPPR.ONLINE
//INDEX    DD DISP=SHR,DSN=CPPR.INDEX
//CPPRERT  DD DISP=SHR,DSN=CPPR.CPPRERT
//SYSUT1   DD DISP=SHR,DSN=CPPR.SMF.INPUT
//SYSUT3   DD SPACE=(CYL,(10,8)),UNIT=SYSDA,
//          DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=*
COMMIT
```

DCOLLECT Data

Optionally, you can derive the data set allocation data by processing DCOLLECT output data. DCOLLECT is an IDCAMS function that produces a sequential output stream consisting of several types of records that describe the manner in which DASD storage is being used. DFHSM statistics (migration to level 1 or 2) are included as well.

If DCOLLECT is used at a particular installation and you prefer DCOLLECT data over VTOC data, you can reduce the DCOLLECT data into the DASM Performance Database with the program named SSA1DCLW. For a sample of the JCL, refer to the DASMCOLW distributed CNTL library.

The following is an example of a Job that reduces DCOLLECT data into the DASM Performance Database.

```
//SSADASM JOB (...),'SSA',CLASS=A,MSGCLASS=X
//*****
//* THIS JOB PROCESSES DCOLLECT INPUT TO PRODUCE DASM TABLES AND
//* OPTIONALLY TO PRODUCE AN AD HOC REPORT
//*****
//SCAN EXEC PGM=SSA1DCLW,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//ONLINE  DD DSN=&PREFIX.CPPR.VNNN.ONLINE.DASM,DISP=SHR
//INDEX   DD DSN=&PREFIX.CPPR.VNNN.INDEX.DASM,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
```

```
//SYSUT1 DD DISP=SHR,DSN=FLAT.FILE.FROM.DCOLLECT
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//*****
//* IF YOU WANT TO INCLUDE A SPECIFIC SET OF VOLUMES IN THE
//* SCAN, USE:
//* //INCVOLS DD *
//* IF YOU WANT TO EXCLUDE A SPECIFIC SET OF VOLUMES FROM THE
//* SCAN, USE:
//* //EXCVOLS DD *
//* - OR -
//* //VOLLIST DD *
//*****
//* IN THIS CASE, DONT EXCLUDE ANY VOLUMES FROM THE PROCESS
//*****
//VOLLIST DD DUMMY
//*****
//* IGNORE
//* IGNORE UNDESIREED HIGH LEVEL QUALIFIERS
//*****
//* $IGN
//DSNLIST DD *
SYS*
/*
//SYSIN DD *
DSNAME SELECTION=EXCLUDE
DASD MAPPING REPORT=YES
OWNER ALLOCATION THRESHOLD=45 /* ONLY PRINT OWNERS WITH ALLOCATIONS
* GREATER THAN 45 MEGABYTES
DEVICE TYPE REPORT=YES
VOLUME ALLOCATION REPORT=YES
* IF YOU DONT WANT TO STORE THE TABLES, COMMENT OUT THE FOLLOWING:
COMMIT
* IF YOU WANT AN AD HOC REPORT FOR A DATASET, UNCOMMENT THE FOLLOWING:
*DSNAME=&PREFIX.CPPR*
```

```
* IF YOU WANT TO SEE UNUSED SPACE BY VSAM ALSO, UNCOMMENT THE
* FOLLOWING STATEMENT. BE AWARE THAT THE PROCESS WILL BE MUCH
* LONGER (@10X) AS THE CATALOG MUST BE ACCESSED FOR EACH VSAM ENTRY
*DASM VSAM STATISTICS=YES
```

Reports Produced During Data Reduction

Both the SMF Data Reduction program (SSA1WKLD) and the DCOLLECT Data reduction program (SSA1DCLW) produce ad hoc reports. These reports are described in the following sections.

DSNAME Ad Hoc Report

The SSA1WKLD program produces the DSNAME Ad Hoc report from type 14, 15, and 64 records.

The DSNAME Ad Hoc report, which can process any specified active SMF cluster as well as the archived SMF data, is particularly useful to quickly determine which User/Job might have allocated, deleted, or updated a particular data set. When the system has experienced poor response time, you can also examine the activity of the various data sets to get an indication of what is causing the problem.

DSNAME Ad Hoc Report SYSIN Parameters

Use the following unique key phrases as SYSIN parameters with the DSNAME Ad Hoc report to define the data that appears in the report.

DSNAME

The DSNAME parameter limits the selection of records to those containing a specified data set name or a generic data set name. For example, DSNAME= SYS1.PROCLIB would limit the selection to only those records containing the data set SYS1.PROCLIB while DSNAME=SYS1* selects all records pertaining to data sets beginning with the characters "SYS1" providing any other selection criteria are satisfied.

The data set name can be either of the following:

- Data set name, up to 44 characters (e.g, DSNAME=SYS3.PAYROLL)

or

- Generic data set name, dsnroot* (e.g, DSNAME=SYS3.PAYR*)

There is no default data set name.

- VOLSER** The VOLSER parameter is used to limit the selection of data sets reported to those residing on a specified volume.
- The volume serial number can be up to 6 characters (e.g, VOLSER=MVS006).
- There is no default VOLSER.
- BEGIN DATE and END DATE** The BEGIN DATE and END DATE parameters limit the selection of records to those written by SMF during the period specified by the begin and end date.
- The begin and end dates can be any of the following values:
- MM/DD/YY
 - MM/DD/YYYY
 - DD.MM.YY
 - DD.MM.YYYY
 - YYDDD
 - YYYYDDD
- The default value is 00001.

Example

Specifying the following in the SYSIN file would produce a report for all data set records written between June 2, 2005 to June 15, 2005 inclusive, containing data set names beginning with the characters "SYS1" providing they reside on the volume MVS203.

```
DSNAME=SYS1*  
VOLSER=MVS203  
BEGIN DATE=06/02/2005  
END DATE=06/15/2005
```

DSNAME Ad Hoc Report Data

The following data is displayed for each record selected for inclusion into the report as appropriate for the record type:

- The DSNAME
- The SMF SID
- The date the data set was closed
- The time the data set was closed
- The name of the job that used the data set
- Data set disposition
- Whether the data set was used as INPUT or OUTPUT and whether the data set was renamed or deleted
- The volume serial on which the data set resided
- The DCB characteristics of the data set
- The number of physical accesses

Operating JCL for DSNAME Ad Hoc Report

The following job stream is recommended to produce the DSNAME Ad Hoc report outside the routine Workload data reduction processing procedure.

```
//JOBNAME JOB . . .
//ST1 EXEC PGM=SSA1WKLD,REGION=4096K,TIME=60
//STEPLIB DD DSN=CPPR.LOAD,DISP=SHR
//CPPRPARM DD DISP=SHR,DSN=CPPR.PARMLIB
//ONLINE DD DUMMY
//INDEX DD DUMMY
//CPPRERT DD DISP=SHR,CPPR.CPPRERT
//SYSUT1 DD DISP=SHR,DSN=INPUT.CPPR.SMF
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,
// DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
```

DCOLLECT Ad Hoc Report

The SSA1DCLW program produces the DCOLLECT Ad Hoc report that shows, for a given storage group, the volume serial numbers for the volumes included in the group and the data set names that reside on the volumes. For each data set name, the following is shown: the amount of space allocated in kilobytes, the amount of space used in kilobytes, and the date the data set was last referenced.

DCOLLECT Ad Hoc Report SYSIN Parameters

Use the following unique key phrases as SYSIN parameters with the DCOLLECT Ad Hoc report to define the data that appears in the report.

DAYS SINCE LAST REFERENCE

The DAYS SINCE LAST REFERENCE parameter limits the selection of data sets to those referenced within a specified number of days (maximum of 999).

STORAGE GROUP NAME

The STORAGE GROUP NAME parameter limits the selection of data sets to those containing a specified storage group name or a generic storage group name.

DSNAME

The DSNAME parameter limits the selection of data sets to those containing a specified data set name or a generic data set name. For example, DSNAME= SYS1.PROCLIB would limit the selection to only those records containing the data set SYS1.PROCLIB while DSNAME=SYS1 * selects all records pertaining to data sets beginning with the characters "SYS1" providing any other selection criteria are satisfied.

The data set name can be either of the following:

- Data set name, up to 44 characters (e.g, DSNAME=SYS3.PAYROLL)

or

- Generic data set name, dsnroot* (e.g, DSNAME=SYS3.PAYR*)

There is no default data set name.

DETAIL

The DETAIL parameter determines whether the ad hoc report shows both data set name and volser information, or just data set or volser data.

The values for this parameter are:

- YES (data set name and volser information is provided in the report)
- VOL (only volume totals are provided in the report)
- DSN (only data set information is provided in the report)

In addition, the //INCVOLS or //EXCVOLS DD statement can be used to limit the volumes included in the report.

Operating JCL for DCOLLECT Ad Hoc Report

The following job stream is recommended to produce the DCOLLECT Ad Hoc report outside the routine DCOLLECT data reduction processing procedure.

```
//CIMS03DC JOB (CIM0),'CIMSLAB',MSGCLASS=X,
//          MSGLEVEL=(1,1),NOTIFY=CIMS03,
//          CLASS=A
//*****
//*          THIS STEP PROCESSES DCOLLECT DATA
//*****
//ST2 EXEC PGM=SSA1DCLW,REGION=OM
//STEPLIB DD DSN=CIMS03.CPPRDEV.V520.LOADLIB,DISP=SHR
//CIMPASS DD DISP=SHR,DSN=CIMS03.CPPRDEV.V520.CNTL(CIMSNUM)
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUT1 DD DISP=SHR,DSN=CIMS.BACKUP.SAMPDATA.DCOLLECT.DATA.PDMLL
//SYSUT3 DD DISP=(,PASS),UNIT=SYSDA,SPACE=(CYL,(1,1))
//ONLINE DD DUMMY
//CPPRPARM DD DISP=SHR,DSN=CIMS03.CPPRDEV.PARMLIB
//CPPRERT DD DISP=SHR,DSN=SSA.CPPRXDM.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSADHOC DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//DSNLIST DD DUMMY
//INCVOLS DD *
DEV142
//SYSIN DD *
DAYS SINCE LAST REFERENCE=120
SELECTED SYSTEM=PT11
DSNAME SELECTION=EXCLUDE
DASD MAPPING REPORT=YES
*STORAGE GROUP NAME=*
STORAGE GROUP NAME=SGDEV
*DSNAME=SYS*
*DETAIL=YES
*DETAIL=VOL
DETAIL=DSN
```


DASM REPORT DATA

As previously mentioned, the DASM Subsystem uses data collected from both SMF/RMF and the VTOCs of the volumes in the DASD farm. It also uses certain control data residing in two CIMS Capacity Planner PARMLIB members to determine data set ownership and to determine the configuration of your DASD Pools.

PARMLIB Data

The DASM Subsystem uses control information from two CIMS Capacity Planner PARMLIB members.

DASDPOOL Parmlib Member

The DASDPOOL member of the CIMS Capacity Planner Parameter library is used to specify the composition of up to six DASD Pools.

The DASD Pool member is laid out as follows:

- A total of six generic pools are available
- Within each pool, you can specify the following information:

The number assigned to the pool (1-6). This pool number must be located in the first column and must be followed by a period in column 2. No additional data, other than comments, can be included in the first statement of a DASD Pool set.

The statement containing the pool number is immediately followed with a list of the volumes assigned to that pool. The first column of the volume list statements must contain a blank. Multiple volume serial numbers are separated by commas. You can specify generic volume serial numbers using the "*" as a wild card character. For example, MVS* causes all volumes with serial numbers beginning with MVS to be considered as members of the pool being specified.

For additional information, refer to the Installation Instructions contained in the *CIMS Capacity Planner Installation Guide*.

&sid.DSNX Parmlib Member

The &sid.DSNX member of the CIMS Capacity Planner parameter library contains the information required to determine the owners of the data sets stored on the DASD volumes in those cases where the owner is not specified by the first level of qualification in the data set name.

For a detailed description of this feature and detailed instructions in setting up the &sid.DSNX member, refer to the Installation Instructions contained in the *CIMS Capacity Planner Installation Guide*.

STGROUPS Parmlib Member

The STGROUPS member of the CIMS Capacity Planner parameter library maintains a correspondence between VOLSER and storage groups. This member is created by the SSA1DASX program (reference CNTL member DASMDASX) and may be used by the SSA1DSTG program to create the &sid.24A–&sid.24G HGDLIB members (reference CNTL member DASMDSTG).

DASM REPORTS

You can produce a variety of reports using the DASM Subsystem. All the reports are selectable and are produced dynamically whenever the DASM Subsystem is executed (SSA1DASM). The same reports can be produced up to the last time that the DASM tables are COMMITed to the Online Performance Data Base by executing SSA1DASH or SSA1RPT. The DASM reports are discussed in detail below.

DASD Space Allocation–By Device Type (SSA1DASM or SSA1DASH)

This report indicates how much DASD space is allocated to each specific "OWNER" by device type on all the DASD devices attached to the processor on which the program is executed.

If SSA1DASM is being executed, the report data is obtained by scanning the VTOCs on the DASD devices attached to the system on which the program is executed. When running SSA1DASM, you can commit the allocation tables that are used to prepare the report to the On-line Performance Data Base using the SYSIN parameter COMMIT.

If SSA1DASH is run, the data used to construct the report is taken from the On-line Performance Data Base and is current up to the point that the data was most recently COMMITted.

The report is presented in columnar format with a column for User ID and separate columns for each type of device detected during the process of scanning the DASD devices. For each user, the number of tracks allocated to each type of device is reported in the appropriate column.

DASD Space Allocation–By Device Volume (SSA1DASM or SSA1DASH)

This report indicates how much DASD space is allocated to each specific OWNER by volume, or set of volumes, on all the DASD devices attached to the processor on which the program is executed.

If SSA1DASM is being executed, the report data is obtained by scanning the VTOCs on the DASD devices attached to the system on which the program is executed. When running SSA1DASM, you can use the SYSIN parameter COMMIT to commit the allocation tables that are used to prepare the report to the Online Performance Data Base the.

If SSA1DASH is run, the data used to construct the report is taken from the Online Performance Data Base and is current up to the point that the data was most recently COMMITted.

The report is presented in columnar format with a column for User Id and separate columns for each type of device detected during the process of scanning the DASD devices. For each user, the number of tracks allocated to each type of device is reported in the appropriate column.

DASD Volume Map by CCHH (SSA1DASM)

This report details how the space on a DASD volume is allocated, sorted by the address of the space on the volume (CCHH).

The data used to produce the DASD Volume Map is taken from the VTOC of the volume being mapped.

Volume Map by DSNAME (SSA1DASM)

This report details how space on a DASD volume is allocated by DSNAME.

The data used to produce the DASD Volume Map is taken from the VTOC of the volume being mapped

Volume Activity Report (SSA1RPT)

The Volume Activity Report shows the data set activity by volume for the 50 most heavily used data sets on each reported volume. The data set listing for each volume is ordered by the number of I/O operations during the measured period.

The data used to produce the Volume Activity Report is taken from the On-line Performance Data Base.

Data Set Activity Report (SSA1RPT)

The Data Set Activity Report shows the 25 most active data sets within the data center for a specified reporting period in terms of the average daily activity per data set.

The data used to produce the Data Set Activity Report is taken from the On-line Performance Data Base.

Data Set Detail Report (SSA1RPT)

This report provides detailed I/O activity information regarding datasets whose names conform to the specifications provided in the "DSNAME" and "VOLSER" parameters. The I/O count for each of the selected datasets is shown for each date in the date range (Begin Date/End Date) along with the actual dataset name and the volume serial number upon which the dataset resides.

DASM Volume Summary Report (SSA1DASV)

This report provides summarized information with respect to the individual volumes in the DASD farm. The data comes from the DASM Performance Database. A COMMIT must have been done in order for the report to produce any information. The following lists the columns of information available on the Volume Summary Report:

DEVICE	The type of device (3380, 3390)
TRACKS TOTAL	The total number of tracks on the volume
TRACKS FREE	The total number of free tracks on the volume
FREE(%)	The percentage of space on the volume that is available for allocation
TRACKS IDLE	The total number of tracks that, although allocated, are not in use
IDLE(%)	The percentage of idle tracks as a function of allocated space

An example of the JCL used to produce the Volume Summary Report is shown below:

```
//SSADASM JOB (...),'SSA',CLASS=A,MSGCLASS=X
//*****
//*   THIS JOB IS USED TO BUILD THE DASD VOLUME SUMMARY REPORT
//*****
//VOLUMSUM EXEC PGM=SSA1DASV
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.VNNN.ONLINE.DASM,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.VNNN.INDEX.DASM,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSIN DD *
SELECTED SYSTEM=*
BEGIN DATE=*
END DATE=*
```

DASM Selected VSAM Summary Report (SSA1VSAM)

This report provides summarized information with respect to a specific VSAM Cluster or set of Clusters. The data comes from the System Catalog. The following lists the information available on the VSAM Summary Report:

LOGICAL REC-LEN	The logical record length for the cluster
HIGH-ALLOC-RBA	The highest allocated Relative Byte Address for the cluster
HIGH-USED-RBA	The highest used Relative Byte Address for the cluster
CISIZE	The Control Interval size
CI/CA	The number of Control Intervals per Control Area
CYLINDERS ALLOC	The number of cylinders allocated to the cluster
CYLINDERS USED	The number of cylinders actually used
% SPACE USED	The percentage of allocated space that was used

An example of the JCL used to produce the Selected VSAM Summary Report is shown below:

```
//SSADASM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
//ST1 EXEC PGM=SSA1VSAM, REGION=4096K, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB, DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ONLINE DD DUMMY
//INDEX DD DUMMY
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB, DISP=SHR
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.VNNN.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD DUMMY
//*****
//* IF YOU WANT TO INCLUDE A SPECIFIC SET OF VOLUMES IN THE
//* SCAN, USE:
//* //INCVOLS DD *
//* IF YOU WANT TO EXCLUDE A SPECIFIC SET OF VOLUMES FROM THE
//* SCAN, USE:
//* //EXCVOLS DD *
//* - OR -
//* //VOLLIST DD *
//*****
```

```
//VSAMIN DD *  
GVSAM.APPAYMTS.DATA  
GVSAM.APPAYMTS.INDEX  
GVSAM.BMCAGE.ALT.DATA  
GVSAM.BMCAGE.ALT.INDEX  
GVSAM.BMCAGE.DATA  
GVSAM.BMCAGE.INDEX
```

DASM VTOC Selected DSN on Volume Report (SSA1DASM)

This report provides data set characteristics information with respect to selected data sets on specific Volumes. The data comes from the VTOC, either using CVAF or the DASM Scan. The following lists the columns of information available on the VTOC Selected DSN on Volume Report:

DSNAME	The name of the data set
C C H H	The address of the beginning of the extent
C C H H	The address of the end of the extent
CREATED	The date the data set was created
LAST REF	The date the data set was last referenced
TRKALLOC	The number of tracks allocated to the extent
TRK USED	The number of tracks used by the entire data set
TRK IDLE	The number of tracks idle in the allocated space
DSORG	The data set organization
LRECL	The logical record length of the data set
BLKSZ	The block size of the data set
RECFM	The record format of the data set

An example of the JCL used to produce the VTOC Selected DSN on Volume Report is shown below:

```
//SSAVTOC JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * *
/* THIS VARIATION ON THE VTOC PRINT UTILITY PRINTS ONLY THE
/* DSNAME ORDERED LISTING OF THE VOLUMES BEGINNING WITH MVS
/* AND SELECTS DSNAMEs BEGINNING WITH &PREFIX.CPPR
/* * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*) *
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGs DD SYSOUT=*
//SYSIN DD *
SELECTED VOLUME=MVS*
SELDSNMSK /* ALLOW THE DSNAME TO BE WILDCARDED */
SELVOLMSK /* ALLOW THE VOLSER TO BE WILDCARDED */
DSNAME=&PREFIX.CPPR*
CCHVOLMAP /* DONT PRINT THE ADDRESS VOLUME MAP */
* IN ORDER TO SEE VSAM DETAILS, UNCOMMENT THE FOLLOWING STATEMENT
*DASM VSAM STATISTICS=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAf
```

This report is also available online through the DASM section of the CIMS Capacity Planner ISPF/PDF interface.

DASM VTOC by CCHH on Volume Report (SSA1DASM)

This report provides a listing of the actual space allocation by address (CCHH) on specific Volumes. The data comes from the VTOC, either using CVAF or the DASM Scan. The following lists the columns of information available on the VTOC by CCHH on Volume Report:

DSNAME	The name of the data set occupying the extent
C C H H	The address of the beginning of the extent
C C H H	The address of the end of the extent
TRKALLOC	The number of tracks in the extent
DSORG	The organization of the data set occupying the extent
LRECL	The logical record length of the data set
BLKSZ	The block size of the data set
RECFM	The record format of the data set

An example of the JCL used to produce the VTOC by CCHH on Volume Report is shown below:

```
//SSAVTOC JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
/* THIS VARIATION ON THE VTOC PRINT UTILITY PRINTS ONLY THE
/* ADDRESS ORDERED LISTING OF THE VTOC FOR MVSRES
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM, REGION=2048K, TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB, DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED VOLUME=MVSRES
DSNVOLMAP /* DONT PRINT THE DSNAME VOLUME MAP */
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

This report is also available online through the DASM section of the CIMS Capacity Planner ISPF/PDF interface.

DASM VTOC FREESPACE on Volume Report (SSA1DASM)

This report provides a list of all Free Space extents on specific Volumes. The data comes from the VTOC, either using CVAF or the DASM Scan. The following lists the columns of information available on the VTOC FREESPACE on Volume Report:

C C H H	The address of the beginning of the free space extent
C C H H	The address of the end of the free space extent
TRKSFREE	The total number of free tracks in the extent

An example of the JCL used to produce the VTOC FREESPACE on Volume Report is shown below:

```
//SSAVTOC JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * *
/* THIS VARIATION ON THE VTOC PRINT UTILITY PRINTS ONLY THE
/* FREE SPACE EXTENTS OF THE VTOC FOR MVSRES
/* * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED VOLUME=MVSRES
DSNVOLMAP /* DONT PRINT THE DSNAME VOLUME MAP */
FREESPACE ONLY /* ONLY SHOW FREE SPACE */
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

This report is also available online through the DASM section of the CIMS Capacity Planner ISPF/PDF interface.

DASM VTOC Detailed Volume Space Summary Report (SSA1DASM)

This report provides detailed volume summary information on specific Volumes. The data comes from the VTOC, either using CVAF or the DASM Scan. The following lists the columns of information available on the VTOC Detailed Volume Space Summary Report:

TRKS/VOL	The total number of tracks on the volume
TRK FREE	The total number of tracks on the volume that are unallocated
TRKALLOC	The total number of allocated tracks on the volume
TRK USED	The total number of allocated tracks that are used
TRK IDLE	The total number of tracks that, although allocated, are not in use

An example of the JCL used to produce the VTOC Detailed Volume Summary Report is shown below:

```
//SSAVTOC JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
/* THIS VARIATION ON THE VTOC PRINT UTILITY PRINTS ONLY THE
/* VOLUME SPACE SUMMARY FOR ALL VOLUMES BEGINNING WITH MVS
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM, REGION=2048K, TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB, DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED VOLUME=MVS*
SELVOLMSK /* ALLOW THE VOLSER TO BE WILDCARDED */
VOLUME SPACE SUMMARY
CCHVOLMAP /* DONT PRINT THE ADDRESS VOLUME MAP */
* IN ORDER TO SEE VSAM DETAILS, UNCOMMENT THE FOLLOWING STATEMENT
*DASM VSAM STATISTICS=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

This report is also available online through the DASM section of the CIMS Capacity Planner ISPF/PDF interface.

DASM VTOC Quick Volume Space Summary Report (SSA1DASM)

This report provides a quick volume space summary report on specific Volumes. The data comes from the VTOC, using the LSPACE command. The following lists the columns of information available on the VTOC Quick Volume Summary Report:

VOLUME SERIAL	The VOLSER of the volume
TOTAL TRACKS	The total number of tracks on the volume
TRACKS ALLOC	The total number of allocated tracks on the volume
TRACKS FREE	The total number of unallocated tracks on the volume
CONTIG FREE TRK	The largest number of contiguous free tracks
# FREE EXTENTS	The number of free space extents
# FREE DSCBS	The number of free data set control blocks in the VTOC
ADR: FREE ALT TRKS	The address of the volume (ccua) and the free alternate tracks remaining
FRAGMENT INDEX	An indicator of how fragmented the free space on the volume is: a value of 0 means no fragmentation (all free space in one contiguous extent)

An example of the JCL used to produce the VTOC Quick Volume Summary Report is shown below:

```
//SSAVTOC JOB (...),SSA:CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * *
/* THIS VARIATION ON THE VTOC PRINT UTILITY PRINTS ONLY THE
/* QUICK VOLUME SUMMARY FOR ALL VOLUMES BEGINNING WITH MVS
/* OR SPO AND, OPTIONALLY, SORTED BY COLUMN
/* * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGSD DD SYSOUT=*
//SYSIN DD *
SELECTED VOLUME=MVS*,SPO*
LSPACE /* INVOKE THE QUICK VOLUME SUMMARY */
```

```
* IF YOU WANT THE DATA SORTED BY COLUMN, USE THE FOLLOWING MODELS:
*DESCENDING SORT COLUMN=8 /* COLUMNS 1-8 ALLOWED */
* OR
*ASCENDING SORT COLUMN=8 /* COLUMNS 1-8 ALLOWED */
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

This report is also available online through the DASM section of the CIMS Capacity Planner ISPF/PDF interface.

DASM X37 Candidate List (SSA1DASM)

This report provides a listing of all data sets that are candidates for x37 Abends on specific Volumes. A data set is presumed to be a candidate for an x37 Abend if it meets the following criteria:

It has already gone into a specified number of secondary extents (See FILTER)

The data set allocation is larger than a specified minimum (See FILTER)

The number of secondary extents allocated exceeds a specified threshold (See FILTER)

The percentage of space used exceeds a specified threshold (See FILTER)

When two or more of the above conditions has been met, the data set is included in the x37 candidate list.

The FILTER key phrase in the SYSIN stream is used to specify thresholds used in determining whether the data set should be included on the list. The format of the FILTER key phrase is as follows:

FILTER=aaa,bbb,ccc,ddd,eee,fff,ggg where

aaa = the minimum data set allocation, in tracks, to consider for inclusion on any list

bbb = the minimum number of secondary extents, regardless of % used, for x37 list

ccc = the minimum number of secondary extents, combined with % used, for x37 list

ddd = % used in combination with (ccc) for x37 list

eee = minimum time un-referenced, in days, for archive list

fff = smallest block size, in K, to avoid being placed on the re-blocking list

ggg = lowest threshold, in % used, to avoid release list

The data comes from the VTOC. The following lists the columns of information available on the DASM X37 Candidate List:

DSNAME	The name of the data set
TRKALLOC	The total number of tracks allocated to that data set
TRK USED	The total number of tracks used by that data set
# EXTENT	The total number of extents allocated to that data set
LREF	The date the data set was last referenced
DSORG	The data set organization
LRECL	The logical record length
BLKSIZE	The block size
RECFM	The record format
VOLSER	The volume serial number of the volume on which the data set resides

An example of the JCL used to produce the DASM X37 Candidate List is shown below::

```
//SSAVTOC JOB (...),SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
FILTER=015,013,006,098,120,004,030
DASM X37 CANDIDATE LIST=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAf
```

DASM Re-blocking Candidate List (SSA1DASM)

This report provides a listing of all data sets that are candidates for re-blocking on specific Volumes. A data set is presumed to be a candidate for re-blocking if it meets the following criteria:

It is larger than a specified allocation size (See FILTER)

The block size is less than a specified threshold (See FILTER)

When the above conditions have been met, the data set is included in the re-blocking candidate list.

The FILTER key phrase in the SYSIN stream is used to specify thresholds used in determining whether the data set should be included in the list. The format of the FILTER key phrase is as follows:

FILTER=aaa,bbb,ccc,ddd,eee,fff,ggg where

aaa = the minimum data set allocation, in tracks, to consider for inclusion on any list

bbb = the minimum number of secondary extents, regardless of % used, for x37 list

ccc = the minimum number of secondary extents, combined with % used, for x37 list

ddd = % used in combination with (ccc) for x37 list

eee = minimum time un-referenced, in days, for archive list

fff = smallest block size, in K, to avoid being placed on the re-blocking list

ggg = lowest threshold, in % used, to avoid release list

The data comes from the VTOC. The following lists the columns of information available on the DASM Re-blocking Candidate List:

DSNAME	The name of the data set
TRKALLOC	The total number of tracks allocated to that data set
TRK USED	The total number of tracks used by that data set
# EXTENT	The total number of extents allocated to that data set
LREF	The date the data set was last referenced

DSORG	The data set organization
LRECL	The logical record length
BLKSIZE	The block size
RECFM	The record format
VOLSER	The volume serial number of the volume on which the data set resides

An example of the JCL used to produce the DASM Re-blocking Candidate List is shown below::

```
//SSAVTOC JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGS DD SYSOUT=*
//SYSIN DD *
FILTER=015,013,006,098,120,004,030
DASM REBLOCKING CANDIDATE LIST=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

DASM Archival Candidate List (SSA1DASM)

This report provides a listing of all data sets that are candidates for Archival from specific Volumes. A data set is presumed to be a candidate for Archival if it meets the following criteria:

It is larger than a specified allocation size (See FILTER)

The number of days since it was last referenced exceeds a specified threshold (See FILTER)

When the above conditions have been met, the data set is included in the Archival candidate list.

The FILTER key phrase in the SYSIN stream is used to specify thresholds used in determining whether the data set should be included on the list. The format of the FILTER key phrase is as follows:

FILTER=aaa,bbb,ccc,ddd,eee,fff,ggg where

aaa = the minimum data set allocation, in tracks, to consider for inclusion on any list

bbb = the minimum number of secondary extents, regardless of % used, for x37 list

ccc = the minimum number of secondary extents, combined with % used, for x37 list

ddd = % used in combination with (ccc) for x37 list

eee = minimum time un-referenced, in days, for archive list

fff = smallest block size, in K, to avoid being placed on the re-blocking list

ggg = lowest threshold, in % used, to avoid release list

The data comes from the VTOC. The following lists the columns of information available on the DASM Archival Candidate List:

DSNAME	The name of the data set
TRKALLOC	The total number of tracks allocated to that data set
TRK USED	The total number of tracks used by that data set
# EXTENT	The total number of extents allocated to that data set
LREF	The date the data set was last referenced

DSORG	The data set organization
LRECL	The logical record length
BLKSIZE	The block size
RECFM	The record format
VOLSER	The volume serial number of the volume on which the data set resides

An example of the JCL used to produce the DASM Archival Candidate List is shown below::

```
//SSAVTOC JOB (...),SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM,REGION=2048K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB,DISP=SHR
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGSGS DD SYSOUT=*
//SYSIN DD *
FILTER=015,013,006,098,120,004,030
DASM ARCHIVAL CANDIDATE LIST=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

DASM Release Candidate List (SSA1DASM)

This report provides a listing of all data sets that are candidates for Release processing on specific Volumes. A data set is presumed to be a candidate for Release processing if it meets the following criteria:

It is larger than a specified allocation size (See FILTER)

The percentage of space used is less than a specified threshold (See FILTER)

When the above conditions have been met, the data set is included in the Release candidate list.

The FILTER key phrase in the SYSIN stream is used to specify thresholds used in determining whether the data set should be included on the list. The format of the FILTER key phrase is as follows:

FILTER=aaa,bbb,ccc,ddd,eee,fff,ggg where

aaa = the minimum data set allocation, in tracks, to consider for inclusion on any list

bbb = the minimum number of secondary extents, regardless of % used, for x37 list

ccc = the minimum number of secondary extents, combined with % used, for x37 list

ddd = % used in combination with (ccc) for x37 list

eee = minimum time un-referenced, in days, for archive list

fff = smallest block size, in K, to avoid being placed on the re-blocking list

ggg = lowest threshold, in % used, to avoid release list

The data comes from the VTOC. The following lists the columns of information available on the DASM Release Candidate List:

DSNAME	The name of the data set
TRKALLOC	The total number of tracks allocated to that data set
TRK USED	The total number of tracks used by that data set
# EXTENT	The total number of extents allocated to that data set
LREF	The date the data set was last referenced
DSORG	The data set organization
LRECL	The logical record length
BLKSIZE	The block size
RECFM	The record format
VOLSER	The volume serial number of the volume on which the data set resides

An example of the JCL used to produce the DASM Release Candidate List is shown below::

```
//SSAVTOC JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
/* * * * * * * * * * * * * * * * * * * * * * *
//SCAN EXEC PGM=SSA1DASM, REGION=2048K, TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.VNNN.LOADLIB, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.VNNN.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.VNNN.PARMLIB, DISP=SHR
```

```
//SYSUDUMP DD SYSOUT=(*)
//SYSNAP   DD SYSOUT=(*)
//SYSPRINT DD SYSOUT=*
//SYSMSGSGS DD SYSOUT=*
//SYSIN    DD *
FILTER=015,013,006,098,120,004,030
DASM RELEASE CANDIDATE LIST=YES
* IF YOUR DASD FARM IS SMS-MANAGED, UNCOMMENT THE FOLLOWING STATEMENT
*SMFILE=CVAF
```

OPERATING JCL

To the extent possible, the operating JCL has been designed to conform to system utility JCL, following IBM conventions. The reports to be produced are selected based upon the program executed and the report selection parameters provided as part of the SYSIN data set.

DD Statements

The DD Statements required by the DASM Subsystem are:

CPPRERT

The CPPRERT DD statement describes the CIMS Capacity Planner Element Registration data set.

CPPRPARM

The CPPRPARM DD statement specifies the CIMS Capacity Planner Parameter Library data set (PDS) to be used by the program specified by the EXEC statement. The referenced program determines which members of the data set are required during the course of executing the program.

DSNLIST

The DSNLIST DD statement specifies a data set that contains a list of high level qualifiers for data set names that are to be excluded (or included, depending on whether DSNAME SELECTION=EXCLUDE or INCLUDE is specified as a SYSIN parameter) during the process of scanning the VTOCs. You can specify the DSNLIST data set as an in-stream data set (`//DSNLIST DD *`) or as a sequential data set. Data set names can be specified as generics or as a full level of qualification. Generic qualification is specified using an asterisk to specify where the character matching is to terminate. For example, the specification `SYS8*` with a DSNAME SELECTION=EXCLUDE would cause all data sets beginning with the name `SYS8` to be excluded from the selection process. Data sets names such as `SYS8SRCE`, `SYS8.PARMLIB`, etc. would be excluded.

HGDLIB

The HGDLIB DD statement specifies a partitioned data set used to store data that is to be passed to the presentation graphics application that runs on a PC under the Harvard Graphics System.

INDEX

This DD statement specifies the index data set that is used to improve performance in accessing the CIMS Capacity Planner Online Performance Data Base.

ONLINE

The ONLINE DD statement specifies the data set containing the CIMS Capacity Planner Online Performance Data Base.

SYSPRINT

The SYSPRINT DD statement describes the data set used to contain the output from the report programs. Unless otherwise specified, its DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

SYSIN

The SYSIN statement is used to specify the options that are to apply during the execution of the data gathering and reporting phases of the DASM programs. You can specify the SYSIN parameters in the JCL stream (SYSIN DD *) or in a sequential data set.

VOLLIST

The VOLLIST DD statement specifies the VOLUMES that are to be excluded from the data gathering and, hence, the reporting phases of the DASM programs. You can specify the volume serial numbers that denote the volumes to be excluded as a full (up to 6 characters) volume serial numbers or generically (using the "*" character). For example, all MVS volumes whose volume serial numbers begin with the characters "MVS" could be excluded by specifying MVS*.

EXCVOLS

You can use the EXCVOLS DD statement instead of the VOLLIST to specify the Volumes that are to be excluded from the data gathering or the reporting phases of the DASM programs. The volume serial numbers that specify the volumes to be excluded can be specified as a full (up to 6 characters) volume serial numbers or generically (using the "*" character). For example, all MVS volumes whose volume serial numbers begin with the characters "MVS" could be excluded by specifying MVS*.

INCVOLS

You can use the INCVOLS DD statement instead of the VOLLIST to specify the Volumes that are to be included in the data gathering or the reporting phases of the DASM programs. The volume serial numbers that specify the volumes to be included can be specified as a full (up the 6 characters) volume serial numbers or generically (using the "*" character). For example, only MVS volumes whose volume serial numbers begin with the characters "MVS" could be included by specifying MVS*.

Parameters

The DASM Subsystem provides a number of run-time options that control which reports are to be produced, specify the source of the data to be used in generating the reports, and limit the data that is to be presented in the reports. For a complete listing of all the parameters available through the SYSIN stream, refer to Chapter 7, Options for Advanced Applications, of the *CIMS Capacity Planner Reference Guide*. This section provides a brief overview of the unique options supported by the DASM subsystem.

DSNAME SELECTION (SSA1DASM)

This parameter is used to specify whether the data set names specified in the DSNLIST data set are to be INCLUDED in or EXCLUDED from the data selected.

Options:

INCLUDE or EXCLUDE

Default Date:

EXCLUDE

SELECTED VOLUME (SSA1DASM)

The SELECTED VOLUME parameter is used to specify the volume to mapped when producing the DASD Volume Map by Address or the DASD Volume Map by DSNAME reports.

Specify Volume:

A single volume serial number in the form vvvvvv

Default Volume:

None

DASD MAPPING REPORT (SSA1DASM)

This parameter is used to specify whether the DASD Space Mapping Reports (DASD Volume Map by CCHHR and DASD Volume Report by DSNAME) are to be produced.

Options:

DASD MAPPING REPORT=YES (produce the report)

DASD MAPPING REPORT=NO (don't produce the report)

Default Option:

DASD MAPPING REPORT=NO

DEVICE TYPE REPORT (SSA1DASM, SSA1DASH)

This parameter is used to specify whether to produce the DASD Space Allocation by Device Type Report.

Options:

DEVICE TYPE REPORT=YES (produce the report)

DEVICE TYPE REPORT=NO (don't produce the report)

Default Option:

DEVICE TYPE REPORT=NO

VOLUME ALLOCATION REPORT (SSA1DASM, SSA1DASH)

This parameter specifies whether the DASD Space Allocation by VOLSER Report is to be produced.

Options:

VOLUME ALLOCATION REPORT=YES (produce the report)

VOLUME ALLOCATION REPORT=NO (don't produce the report)

Default Option:

VOLUME ALLOCATION REPORT=NO

OWNER ALLOCATION THRESHOLD (SSA1DASM, SSA1DASH)

This parameter is used to specify the threshold in Megabytes, to be used by the SSA1DASM and SSA1DASH programs in determining whether to include a given User ID (owner) in the DASD Space Allocation by Device Type Report. For example, it might be useful to determine all users that have allocated more than 45 Megabytes of DASD space.

Options:

OWNER ALLOCATION THRESHOLD=nnn

Where nnn can range from 0 to 255

Default Option:

OWNER THRESHOLD=0

COMMIT (SSA1DASM)

This parameter specifies whether to write the DASM allocation tables to the Online Performance Data Base.

Options:

The COMMIT parameter has no options. The parameter stands by itself, specifying that the DASM VTOC data set allocation data is to be saved in the On-line Performance Data Base.

Default Option:

Not Applicable

VOLUME ACTIVITY REPORT (SSA1RPT)

This parameter is used to specify that the SSA1RPT program is to produce the Volume Activity Report showing the 50 most active data sets on a given DASD volume.

Options:

VOLUME ACTIVITY REPORT=YES (produce the report)

VOLUME ACTIVITY REPORT=NO (don't produce the report)

Default Option:

VOLUME ACTIVITY REPORT=NO

VOLUME SELECTION CRITERIA (SSA1RPT)

This parameter is used by the SSA1RPT program to determine which volumes are to be included in the Volume Activity Report. The options are to specify the TOPnn or NONE. If you specify TOP50, for example, the program will determine the 50 busiest DASD Volumes (in terms of start I/O's) and report on the busiest datasets for these volumes.

Options:

VOLUME SELECTION CRITERIA=TOPnn, or

VOLUME SELECTION CRITERIA= NONE

Default Option:

VOLUME SELECTION CRITERIA=TOP5

Note • If None is specified, the SELECTED VOLUME Parameter must be specified.

SELECTED VOLUME (SSA1RPT)

This parameter explicitly specifies those volumes that are to be selected for inclusion into the Volume Activity Report.

Specify Volume(s):

SELECTED VOLUME=vvvvvv,vvvvvv,... (up to 6 volume serial numbers separated by commas)

Default Volume:

There is no default.

DSNAME ACTIVITY REPORT (SSA1RPT)

This parameter specifies whether the DSNAME Activity Report showing 25 busiest data sets during the specified measurement period should be limited to a specified date range (BEGIN DATE / END DATE).

Options:

DSNAME ACTIVITY REPORT=YES (produce the report)

DSNAME ACTIVITY REPORT=NO (don't produce the report)

Default Option:

DSNAME ACTIVITY REPORT=NO

DSNAME DETAIL REPORT (SSA1RPT)

This parameter, in conjunction with the "DSNAME" and "VOLSER" parameters, produces a report that shows, on a daily basis, the specific I/O activity for each dataset that matches the selection (i.e, DSNAME and VOLSER) criteria.

Options:

DSNAME DETAIL REPORT=YES

DSNAME DETAIL REPORT=NO

Default Option:

DSNAME DETAIL REPORT=NO

Additional Report Options

It is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it might be useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report can be prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****" or replaced by substitution - see below).

Finally, the statement :

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE OR EXCLUDE facilities, however it does not preclude their use. The combination of the various EXCLUDE/INCLUDE features in conjunction with the Generic Element Mask is discussed below.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a GENERIC ELEMENT MASK=11110011. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL*. If GENERIC ELEMENT MASK=11111111 were to be specified, all records would be summarized into a single entry or be subject to substitution as discussed below.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the CICS report processing programs to intercept the key of each element (in this case, the CICS Transaction name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "????23?". The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the CICS report processing programs to intercept the key of each element (in this case, the CICS Transaction name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "????23?". The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

Use of GENERIC ELEMENT MASK with EXCLUDE/INCLUDE and/or INCLUSIVE or EXCLUSIVE ELEMENT MASK

The following discussion proscribes the rules for using the various element masks and the order in which selection and substitution is performed.

- EXCLUSIVE ELEMENT MASK and the INCLUSIVE ELEMENT MASK are mutually exclusive as are the EXCLUDE and the INCLUDE DD statements.
- EXCLUSIVE ELEMENT MASK and INCLUSIVE ELEMENT MASK must be used in conjunction with the EXCLUDE or INCLUDE DD statement respectively to define the key elements that are to be either excluded or included. Without the EXCLUDE or INCLUDE lists defined through the use of the DD statements, the EXCLUSIVE and the INCLUSIVE ELEMENT MASKs are meaningless.
- The EXCLUDE and INCLUDE lists defined through the use of the EXCLUDE and INCLUDE DD statements can be used conjunction with the GENERIC ELEMENT MASK to limit the data selected for inclusion into the reports irrespective of whether an EXCLUSIVE or INCLUSIVE ELEMENT MASK is specified.
- EXCLUSIVE and INCLUSIVE ELEMENT MASKing is performed prior to the EXCLUDE/INCLUDE logic.
- GENERIC ELEMENT MASKing is takes place after the EXCLUDE/INCLUDE logic is performed.
- If the GENERIC ELEMENT MASK is specified as all "1"s (11111111), then:
 - If a CPPR.PARMLIB member &sidDASF is present, the substitutions are made based upon the data contained in &sidDASF.
 - If the &sidDASF member is not present in CPPR.PARMLIB, then the entire key will be set to "*"s and all records would be summarized into a single entry.

See Chapter 1, CIMS Capacity Planner PARMLIB, of the *CIMS Capacity Planner Reference Guide* for a description of the CPPR.PARMLIB member &sidDASF. You can find a sample of the &sidDASF member in the CPPR.PARMLIB as member CPPRDASF.

SAMPLE JOBSTREAMS

Scanning the DASD Farm

The following is an example of a Job stream that dynamically scans the entire DASD farm and produces all the available reports. The volumes whose VOLSER begin with the letters "DR" are excluded. The DSNAMES whose high level qualifiers begin with the characters "SYS8" are also excluded. In addition, the volume labeled SYSRES are mapped.

```
//DASMJOB JOB (aaaaaaaa)
//STA EXEC PGM=SSA1DASM
//SYSPRINT DD SYSOUT=*
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//SYSIN DD *
DASD MAPPING REPORT=YES
DSNAME SELECTION=EXCLUDE
DEVICE TYPE REPORT=YES
VOLUME ALLOCATION REPORT=YES
SELECTED VOLUME=SYSRES
//VOLLIST DD *
DR*
//DSNLIST DD *
SYS8*
```

Saving the DASM Tables

The following is an example of a job stream that dynamically scans the entire DASD farm and saves the allocation in the Online Performance Data Base. Also, a report would be produced showing the space allocated by device type for all owners (with the exception of SYS8*) whose space allocation exceeds 45 Megabytes.

```
//DASMJOB JOB .....
//STA      EXEC PGM=SSAIDASM
//STEPLIB DD DSN=CPPR.LOADLIB,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//INDEX    DD DSN=CPPR.INDEX,DISP=SHR
//ONLINE   DD DSN=CPPR.ONLINE,DISP=SHR
//SYSUT3   DD DISP=(,DELETE),SPACE=(CYL,(30,10)),UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
COMMIT
DASD MAPPING REPORT=YES
DSNAME SELECTION=EXCLUDE
DEVICE TYPE REPORT=YES
OWNER ALLOCATION THRESHOLD=45
//VOLLIST  DD DUMMY
//DSNLIST  DD *
SYS8*
```

Reporting from the Online Tables

The following is an example of a Job stream that scans the Online Performance Database and produces a report showing space allocated by device type on the 1st of June, 2005, for all owners (with the exception of SYS8*) whose space allocation exceeds 45 Megabytes.

```
//DASMJOB JOB (aaaaaaaa)
//STA EXEC PGM=SSA1DASH
//INDEX DD DSN=CPPR.INDEX,DISP=SHR
//ONLINE DD DSN=CPPR.ONLINE,DISP=SHR
//CPPRERT DD DSN=CPPR.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=CPPR.PARMLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DASD MAPPING REPORT=YES
DSNAME SELECTION=EXCLUDE
DEVICE TYPE REPORT=YES
OWNER ALLOCATION THRESHOLD=45
BEGIN DATE=06/01/2005
END DATE=06/01/2005
//DSNLIST DD *
SYS8*
```

ISPF/PDF DASM Subsystem Interface

The CIMS Capacity Planner DASM Subsystem is supported by an ISPF/PDF interface that greatly simplifies running the full range of CIMS Capacity Planner batch reports and generates input to the PC Presentation Graphics Subsystem (Harvard Graphics).

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASM, CICS, etc.).

Reports

You can produce the following reports online through the CIMS Capacity Planner ISPF Interface.

DASM Activity Reports

The DASM report selection panel provides the option of submitting batch jobs to produce the following DASM activity reports:

Volume Activity Report

The Volume Activity is described under paragraph 4.5 of this section.

Data Set Activity Report

The Data Set Activity Report is described under paragraph 4.6 of this section.

Data Set Detail Report

The Data Set Detail Report is described under paragraph 4.7 of this section.

DASM Space Allocation Reports

The space allocation report lets you select up to two space allocation reports:

Space Allocation by Device Type Report

The Allocation by Device Type is described under paragraph 4.1 of this section.

Volume Allocation Report

The Volume Allocation Report is described under paragraph 4.2 of this section.

DASD Volume Mapping Reports

The Volume Mapping Panel lets you produce the following Volume Mapping reports.

Volume Map by CCHH

The DASD Volume Map by CCHH is described under paragraph 4.3 of this section.

Volume Map by DSNAME

The DASD Volume Map by DSNAME is described under paragraph 4.4 of this section.

Ad Hoc Report

The DSNNAME Ad Hoc report lets you extract various detail data directly from the SMF files. This differs from the other CIMS Capacity Planner reporting functions that take their data from the Performance Data Base. The DSNNAME Ad Hoc report requests are constructed by the ISPF Interface, but are submitted as batch jobs. The DSNNAME Ad Hoc report is described in *DSNNAME Ad Hoc Report on page 2-6*.

Presentation Graphics

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the DASD utilization by device, pool, and users as well as the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required utilization and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software plots each set of data points in accordance with the characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Selected aspects of the DASD utilization and trends are supported by the presentation graphics subsystem.

For more information regarding the actual construction of the Presentation Graphics as well as the specific data point members produced by the programs in the DASM subsystem, refer to Chapter 3, CIMS Capacity Planner Graphics, of the *CIMS Capacity Planner Reference Guide*.

Graph Production Procedures

To generate the CIMS Capacity Planner Workload graphs through the Harvard Graphics System

- 1 Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in *Graph Descriptions on page 2-45* below.
- 2 Compile a composite list of the data point generation programs that must be run to produce the required data points.
- 3 Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- 4 Run the required data point generation programs
- 5 Down-load the data points from the host to the PC
- 6 Invoke the Harvard Graphics system on the PC

- 7 Select the proper graph template
- 8 Import the data points into the template

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters:

"&sidpiiq" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly
- "iiq" is the graph identifier and consists of two alpha-numeric characters. The "piiq" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under **Presentation Graphics Distribution Disk on page 2-44**. The qualifier "q" might or might not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpiiq.DAT where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to down-load the data points from HGDLIB to the PC using IND\$FILE. You can modify the BAT files as required to fit the naming conventions of your installation.

The names of the BAT files conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
- "WKLD"—Workload graph data
- "DASM"—DASM graph data
- "CICS"—CICS graph data
- "CICE"—CICS Extension graph data
- "IDMS"—IDMS graph data
- "IDME"—IDMS Extension graph data
- "IMSG"—IMS graph data
- "NETG"—Network graph data
- "NETE"—Network Extension graph data
- "NETL"—Network Physical Lines graph data
- "NETE"—Network Extension graph Data
- "DB2S"—DB2 Subsystem graph data

- "DB2U"—DB2 Connect ID graph data
- "M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M 204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTH- LY	MONTH- LY	MONTH- LY	MONTH- LY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB data point members from the Host processor (assumes the use of IND\$FILE for the file transfer). The naming conventions are as follows:

.BAT File Naming Conventions

The .BAT filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1= Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a subsystem identifier:

- WKLD=Workload
- DASM=DASM
- CICS=CICS
- CICE=CICS Extension

- IDMS=IDMS
- IDME=IDMS Extension
- IMS=IMS
- NETG=Network
- NETE=Network Extension
- NETL=Network Physical Lines
- NETN=Network Physical NCP

For example, the .BAT file used to download the HGDLIB members required to produce the Workload Daily graphs would be named HGD2WKLD.BAT.

.TPL File Naming Conventions

The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1=Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by 2 character graph identifier. In some cases, a 1 character qualifier is added extending the graph identifier to 3 characters. See the Graphic Report Tables for the specific template names.

Graph Periods

See Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for a description of the weekly and monthly graphs periods.

Graph Descriptions

See Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for a description of each of the graphs produced by the CIMS Capacity Planner DASM Subsystem.

Graphic Report Tables

The following tables summarize the relationship of members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-Based graph program.

DASM Subsystem Regular Weekly Graphs

CPPR Program Name	HGDLIB Member Name	CPPRPARM Library Member	Report Description
SSA1DASP	&sid360	DASDPOOL	DASD Space Available
	&sid361	DASDPOOL	DASD Space Allocated
SSA1DASG	&sid362	DASDPOOL	DASD Pool # 1 Alloc. History
	&sid363	DASDPOOL	DASD Pool # 2 Alloc. History
	&sid364	DASDPOOL	DASD Pool # 3 Alloc. History
	&sid365	DASDPOOL	DASD Pool # 4 Alloc. History
	&sid366	DASDPOOL	DASD Pool # 5 Alloc. History
	&sid367	DASDPOOL	DASD Pool # 6 Alloc. History
SSA1DASP	&sid368	DASDPOOL	DASD Space Allocated:Available
SSA1DASJ	&sid369	&sidDASF	DASD Space Wasted
	&sid36A	&sidDASF	DASD Space % Wasted
	&sid36B	&sidDASF	DASD Space Un-referenced
SSA1DAST	&sid36C	&sidDASF	Space Allocated over time by Org
	&sid36D	&sidDASF	Space Wasted over time by Org
	&sid36E	&sidDASF	Space unref 30 days by group
	&sid36F	&sidDASF	Space unref 60 days by group
	&sid36G	&sidDASF	Space unref 90 days by group

&sid is the SMF System ID (e.g, CPPR)

DASM Subsystem Regular Monthly Graphs

CPPR Program Name	HGDLIB Member Name	CPPRPARM Library Member	Report Description
SSA1DASP	&sid460	DASDPOOL	DASD Space Available
	&sid461	DASDPOOL	DASD Space Allocated
SSA1DASG	&sid462	DASDPOOL	DASD Pool # 1 Alloc. History
	&sid463	DASDPOOL	DASD Pool # 2 Alloc. History
	&sid464	DASDPOOL	DASD Pool # 3 Alloc. History
	&sid465	DASDPOOL	DASD Pool # 4 Alloc. History
	&sid466	DASDPOOL	DASD Pool # 5 Alloc. History
	&sid467	DASDPOOL	DASD Pool # 6 Alloc. History
SSA1DASP	&sid468	DASDPOOL	DASD Space Allocated:Available
SSA1DASJ	&sid469	&sidDASF	DASD Space Wasted
	&sid46A	&sidDASF	DASD Space % Wasted
	&sid46B	&sidDASF	DASD Space Un-referenced
SSA1DAST	&sid46C	&sidDASF	Space Allocated over time by Org
	&sid46D	&sidDASF	Space Wasted over time by Org
	&sid46E	&sidDASF	Space unref 30 days by group
	&sid46F	&sidDASF	Space unref 60 days by group
	&sid46G	&sidDASF	Space unref 90 days by group

DASM Subsystem Regular Daily Graphs

CPPR Program Name	HGDLIB Member Name	CPPRPARM Library Member	Report Description
SSA1DASU	&sid240A	(none)	Average % Busy by Specific Volume
	&sid241A		Average Queue Delay by Specific Volume
	&sid242A		Average I/O Service Time by Specific Volume
	&sid243A		Average I/O Counts by Specific Volume
SSA1DASX	&sid26X	(none)	Space Allocated by Storage Group
SSA1DSTG	&sid26A	(none)	Aggregate I/O Count by Storage Group
	&sid26B	(none)	Average Queue Delay by Storage Group
	&sid26C	(none)	Average I/O Connect Time by Storage Group
	&sid26D	(none)	Average I/O Pending Time by Storage Group
	&sid26E	(none)	Average I/O Active Time by Storage Group
	&sid26F	(none)	Average I/O Disconnect Time by Storage Group
	&sid26G	(none)	Aggregate I/O Allocations by Storage Group

CICS SUBSYSTEM

Subsystem Overview	3-3
CICS Workload and Performance Reports	3-3
CICS Summary Analysis Report	3-3
CICS Response Performance Report	3-4
CICS Transaction Statistics Report	3-4
CICS Terminal Statistics Report	3-6
CICS Exception Analysis	3-6
CICS Ad Hoc Report	3-7
CICS Transaction Activity List	3-8
CICS Terminal Activity List	3-9
CICS System Overview	3-10
CICS Graphs	3-11
CICS Transaction Response Graph	3-11
CICS Terminal Response Graph	3-11
CICS Transaction Activity Graph	3-12
CICS Terminal Activity Graph	3-12
CICS CPU Activity Graph	3-12
CICS I/O Activity Graph	3-12
CICS Transaction Profile	3-12
CICS Trends Analysis Reports	3-13
Operating JCL	3-14
CICS Subsystem Installation JCL	3-14
Data Reduction JCL	3-14
Report Production JCL	3-27
Additional Report Options	3-43
ISPF/PDF CICS Subsystem Interface	3-47
On-line Graphs	3-47
On-line Reports	3-48

CICS Summary Analysis Report	3-49
CICS Trends Analysis Report	3-50
CICS Exception Analysis Report	3-50
CICS Ad Hoc Report	3-50
Batch Reports	3-51
Batch GDDM Graphs	3-51
Presentation Graphics Interface	3-56
Graph Production Procedures	3-56
Naming Conventions	3-57
Presentation Graphics Distribution Disk	3-59
Graph Periods	3-60
Graph Descriptions	3-60
Graphic Report Tables	3-61

Subsystem Overview

The CIMS Capacity Planner CICS Subsystem provides a variety of reports and graphs related to the CICS workload and the performance of your CICS system and applications. The CICS reports provide both detail and summary level information designed to meet the needs of both the technical staff and technical management. A variety of graphs are provided through both GDDM and the PC-based Harvard Graphics system. GDDM lets you instantly display a number of CICS graphs on the Host using the CIMS Capacity Planner ISPF interface. The Harvard Graphics system is used primarily to produce professional quality colored graphs that are suitable for presentations to users and management.

CICS Workload and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner CICS reports and describes the contents of each report.

CICS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the CICS workload and how well the CICS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

You can produce the Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the CICS system.
- The average number of active terminals (high watermark)

A terminal is considered active during each 15-minute period throughout the day only if at least one transaction is submitted. If a terminal is logged on throughout three shifts, but transactions are submitted only during two 15 minute periods, the terminal is considered active for only 30 minutes.
- The average number of CICS transactions processed per minute during the prime shift
- The average number of CICS-related EXCPs per second during the prime shift
- The average terminal response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift

- The average percentage of the CPU time consumed by CICS and the CICS applications during each shift within the measured period.
- The average number of transactions processed by shift for the measured period.

The 25 most frequently executed transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions is reported by shift.

CICS Response Performance Report

This report shows, by shift, for a single day or a range of days, how the CICS transactions flowed through the system for a specified CICS region. Each of the 25 most frequently executed CICS transactions are listed individually, by shift, with an indication of the percentage of the response times that fell within certain predefined time-frames. Transactions that are not among the 25 most frequently executed are grouped together and reported under "OTHER". A Totals line is accumulated and written to indicate how all transactions performed as a group.

This is a key report in measuring how well your performance goals are being met within the CPU by the CICS system and the application programs.

The six response time categories under which the transactions are classified are defined by the CPPR>PARMLIB member &sidCICR.

CICS Transaction Statistics Report

This report shows, for a specified CICS region, which transactions were among the top 40 in the following categories:

- Most frequently executed

The transaction codes and the transaction counts of the 40 most frequently executed transactions along with their respective percentages of the total CICS transaction workload. The transactions are listed in descending order by frequency of execution.

- Most cumulative elapsed time

The transaction codes and the total amount of elapsed (residence) time for each of 40 transactions experiencing the highest residency time along with their respective percentages of the total CICS transaction residence time. The transactions are listed in descending order by residency time.

- Most cumulative CPU time

The transaction codes and the total amount of CPU time consumed for each of the 40 transactions that consumed the most CPU time along with their respective percentages of the total CICS CPU time. The transactions are listed in descending order by CPU time.

■ Most File Accesses

The transaction codes and the number of File Access Calls for each of the 40 transactions issuing the most File Access calls along with their respective percentages of all CICS file access activity. The transactions are listed in descending order by file access activity.

■ Most Journal Accesses

The transaction codes and the number of Journal Program calls for each of the 40 transactions issuing the most calls to the Journal Control Program along with their respective percentages of all the journal access activity. The transactions are listed in descending order by the number of journal issued.

■ Most Frequent Requests for Temporary Storage

The transaction codes and the number of Temporary Storage requests for each of the 40 transactions most frequently requesting Temporary Storage along with their respective percentages of the total Temporary Storage calls for all the CICS transactions. The listing is produced in descending order by the number of Temporary Storage requests.

■ Most Requests for Transient Data

The transaction codes and the number of requests for Transient Data for each of the 40 transactions issuing the most Transient Data requests along with their respective percentages of the total Transient Data requests for all CICS transactions. The listing is produced in descending order by the number of Transient Data requests.

■ Most Paging Activity

The transaction codes and the amount of paging activity for the 40 transactions experiencing the highest number of page faults along with their respective percentage of all paging activity related to CICS transaction processing. The listing is produced in descending order by the amount of paging required.

■ Most Storage Used

The transaction codes and the high watermark (in bytes) for the 40 CICS transactions requiring the most storage. The listing is produced in descending order by the amount of storage used for each transaction.

You can exclude transactions from the Transaction Statistics Report using the EXCLUDE facility (see JCL). Similarly, you can include transactions that would not normally be included using the INCLUDE facility (JCL).

Any counts that exceed 6 digits are expressed in thousands by suffixing the number with the character "K".

CICS Terminal Statistics Report

This report shows which terminals, for a specified CICS Region, were among the top 40 in the following categories:

- **Most Active**

The Terminal ID and the amount of time active (hours:minutes) are reported along with the percentage of the combined active time for all active terminals for the 40 most active terminals. A terminal is considered active within any given 15 minute period only if transactions are received from the terminal. A terminal that is signed-on, but does not submit any transactions is not considered active.

- **Most Transactions Executed**

The Terminal ID and the number of transactions submitted are reported for the 40 terminals submitting the most transactions along with their respective percentages of the total number of transactions submitted during the measured period. The listing is produced in descending order by the number of transactions submitted.

Any counts that exceed 6 digits are expressed in thousands by suffixing the number with the character "K".

You can exclude selected terminals from the report using the EXCLUDE facility (see JCL). Similarly, you can include terminals that do not qualify as among the top forty in the report using the INCLUDE facility (see JCL).

CICS Exception Analysis

The CICS Exception Analysis Report shows, for each fifteen minute period for a specified CICS region, all transactions and terminals whose response times exceeded the pre-defined thresholds specified in the CIMS Capacity Planner Parmlib member &sidXCPT. The Parmlib member &sidXCPT contains a number of threshold values for the installation. This member identifies the performance thresholds for the various components of the overall environment.

You can find additional information related to the &sidXCPT Parmlib member in Chapter 1, CIMS Capacity Planner PARMLIB Overview, in the *CIMS Capacity Planner Reference Guide*.

CICS Ad Hoc Report

The CIMS Capacity Planner CICS Subsystem supports the production of an Ad Hoc report from the original data before portions of the detail data are summarized or discarded. The Ad Hoc report is produced by the Data Reduction program (SSA1CICW).

To produce the CICS Ad Hoc report, specify the report request parameters in the CICS Data Reduction Job stream as part of the SYSIN data set. The parameters that apply to the Ad Hoc report are:

CICS TRANSACTION NAME=

CICS TERMINAL NAME=

The CICS Ad Hoc report contains the following information for each record selected for inclusion in the report:

- The CICS Region Name
- The time that the transaction began execution
- The time that the transaction completed execution
- The transaction ID
- The terminal ID
- The CICS Transaction processing program name
- The elapsed time for the transaction
- The CPU time used
- The Wait time
- The Suspend time
- The number of File Access Calls
- The number of Temporary Storage Calls
- The number of Transient Data Requests
- The number of Paging Operations
- The High Watermark in storage required

In addition to selecting records for inclusion into the Ad Hoc report by specific transaction and terminal names, you can select groups of both transactions and terminals using the wildcard character "*".

For example, you can select all transactions beginning with the characters "TR" by specifying CICS TRANSACTION NAME=TR*. Similarly, you can select all terminals beginning with "L422" by specifying CICS TERMINAL NAME=L422*.

This convention can be extended to include all transactions by specifying CICS TRANSACTION NAME=*. CICS TERMINAL NAME=* would result in the selection of all terminals. If both the transaction and terminal names are specified as "*", all records would be selected.

The selection of records for inclusion into the Ad Hoc report can be further restricted by specifying date and/or time ranges by using the BEGIN DATE=, END DATE=, BEGIN TIME=, and END TIME= parameters.

CICS Transaction Activity List

The CICS Transaction Activity List summarizes the system resources, both physical and logical, consumed in processing each type of transaction included in the report. The period for which the transaction data is summarized is defined by the BEGIN DATE/BEGIN TIME and the END DATE/END TIME parameters of their default values.

The following summary data is reported for each transaction type:

- **Transaction Code**

The transaction code is a field consisting of up to 8 characters that represents a standard CICS transaction code or a synonym substituted by the Generic Element Mask feature (see Additional Report Options, paragraph 5.4 of this section for an explanation of the Generic Element Mask).

- **Transaction Count**

This is the average frequency (daily) that the transaction was executed during the measured period.

- **Average Response Time**

The average daily response time (in seconds and hundredths of seconds) is reported for each type of transaction included in the report.

- **Average CPU Time**

The average daily CPU time consumed (in seconds and hundredths of seconds) is reported for each type of transaction included in the report.

- **Transaction Paging Activity**

The report includes the average number of paging operations (daily) required to service each type of transaction.

- **I/O Activity**

The daily average number of calls to the file manager required to support the processing of each type of transaction is reported.

- **Temporary Storage Activity**

The daily average number of requests for temporary storage required to support the processing of each type of transaction is reported.

- Transient Data Requests

The daily average number of requests for transient data required to support the processing of each type of transaction is reported.

- High Watermark

The high watermark in storage required to process each type of transaction is reported for each type of transaction included in the report.

- Dispatch

The daily average amount of time the transaction was dispatched, in hundredths of a second (response-wait=dispatch)

You can apply the EXCLUDE/INCLUDE DD screens Transaction Activity List Report generation to exclude selected transaction types or specify the specific transaction types that are to be included in the report.

You can use the generic element mask to summarize a group of transaction types into a single transaction or group name (such as department) for summarization purposes.

You can use the EXCLUDE/INCLUDE filter in conjunction with the generic element mask. In this case, the EXCLUDE/INCLUDE filter is applied prior to the generic element mask. See [GENERIC ELEMENT MASK on page 3-44](#) of this section for instructions for using the generic element mask.

CICS Terminal Activity List

The CICS Terminal Activity List provides a summary of the activity of the terminals associated with a specified CICS region during a specified period designated by a set of BEGIN DATE and END DATE parameters. You can further refine the measured period on a daily basis using the BEGIN TIME and END TIME parameters. If you use the BEGIN TIME and END TIME parameters, the measured period is restricted to the specified time period for each day within the specified date range. If you do not specify the BEGIN TIME and END TIME parameters, the time values default to 00:00.00 to 24:00.00.

The values reported might consist of daily averages or the totals for the specified time period. To report the totals rather than the daily average values, include the AVERAGES=NONE parameter in the //SYSIN data.

The following information is reported for all active terminals associated with the specified CICS region during the specified time period:

- Terminal

The Terminal ID of each active terminal

- Transactions

The number of transactions submitted by the terminal. This value can be either a daily average or a total value, depending on whether the AVERAGES=NONE parameter is specified.

■ Percent of Total

The percentage of the total number of transactions represented by the transaction count.

■ Time Active

This value contains either the total time each terminal is active on a daily average basis or a total across multiple days, depending upon whether the AVERAGES=NONE parameter is specified.

■ Percent of Up-time

This value contains the percentage of the CICS region up-time represented by the time the terminal is active. You can calculate this value based upon the daily average up-time or total up-time, depending upon the presence or absence of the AVERAGES=NONE parameter.

■ Average Response

The average response contains the transaction response time computed on a daily average basis or an overall total value, depending on whether the AVERAGES=NONE parameter is specified.

■ Totals

The totals line contains the total values for the Transactions and the Time Active.

CICS System Overview

The CICS System Overview shows the hourly average values for all included CICS transactions (INCLUDE/EXCLUDE applies) for each day in the measured period (BEGIN DATE:END DATE) for the following items of information:

■ Transaction Count**■ Average Response per Transaction****■ Cumulative CPU Time in hundredths of a second****■ Total Dispatch Time in hundredths of a second (Dispatch Time is defined as elapsed time-wait time)****■ Pages per Dispatch Second**

The hours included in the report are determined by BEGIN TIME:END TIME. The default times are 00.00-24.00. In addition to the hourly values for each day, a totals line is provided for each of the values

CICS Graphs

The CICS Subsystem produces a number of graphs on the mainframe that depict the workload and performance of a specified CICS Region.

CICS Transaction Response Graph

The Transaction Response Graph shows, by day of the week, the peak and average transaction response times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the response time (both peak and average). Both the average and peak average values are printed in the middle of each line. The lines, in bar graph format, form the graphic representation of the values. The peak average values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak average values and the average values are either equal or very nearly so.

You can produce the Transaction Response Graph for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced-- one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

CICS Terminal Response Graph

The Terminal Response Graph shows, by day of the week, the peak average and average terminal response times for each 15-minute period through the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the response time (both peak and average). Both the average and peak average values are printed in the middle of each line. The lines, in bar graph format, form the graphic representation of the values. The peak values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak and average values are either equal or very nearly so.

You can produce the Terminal Response Graph for a single day's activity or for several days. When more than one day's activity is graphed, separate graphs are produced-- one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

The Terminal Response Graph differs from the Transaction Response Graph in two ways:

- Not all transaction responses are directed at terminals
- Through the judicious use of the EXCLUDE facility, it is possible to exclude all printers, monitors, and responses to long running transactions from the graph.

CICS Transaction Activity Graph

The Transaction Activity Graph is organized in the same format as the Transaction and Terminal Response Graphs described above. It shows, by day of the week, the average and peak number of transactions submitted from all the CICS terminals during each 15-minute period throughout the day. The horizontal value represents the number of transactions submitted per minute.

CICS Terminal Activity Graph

The Terminal Activity Graph shows, by day of the week, the number of terminals active during each 15-minute period throughout the day. Both the peak and the average values are reported. The values along the horizontal axis represent the number of active terminals.

CICS CPU Activity Graph

The CPU Activity Graph shows, by day of the week, the percentage of the CPU processing capacity that was expended in processing CICS transactions for a specified CICS region. The average and peak percentage values are reported in 15-minute intervals throughout the day. As is the case with the other graphs, you can specify the measured period to include anywhere from a day to several months.

CICS I/O Activity Graph

This graph shows, for a specified CICS region, the amount of EXCP activity related to processing CICS transactions. Both the peak and the average I/O rates are reported. The format of the I/O Activity Graph is identical to the graphs described above, except that the values along the horizontal axis represent the number of I/Os per second related to processing the CICS workload.

CICS Transaction Profile

The Transaction Profile report shows, for a specified CICS region type, how a specified CICS transaction behaved for the following categories:

- How many times per day, on average, it was invoked
- The average response time per invocation of the transaction
- The average CPU time per invocation used to process the transaction
- The average time per invocation that the transaction was on the SUSPEND queue
- The average number of page faults experienced by the transaction per invocation
- The average number of calls to the File Control Program per invocation
- The average number of FCT calls that resulted in physical I/O activity.
- The average number of requests for Temporary Storage per invocation
- The average number of requests for Transient Data per invocation
- The average number of calls to the Journal Control Program per invocation

CICS Trends Analysis Reports

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity- and performance-related elements and providing either a summary or graphic comparison of the values of those elements over time. The CICS elements that are selected for comparison are:

- CPU Utilization Statistics, including a breakdown in terms of CICS overhead and Task related CPU usage
- The number of transactions executed per minute computed as an average during Prime shift
- The average number of EXCPs per second during Prime shift
- The average response time during Prime shift
- The ratio of Early shift to Prime shift transactions
- The ratio of Late shift to Prime shift transactions
- The ratio of Early shift to Prime shift terminals
- The ratio of Late shift to Prime shift terminals

This information is gathered and presented to you in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right-hand side of the report
- In data suitable for graphing with the Harvard Graphics Program once it has been down loaded to a Personal Computer. The manner in which the data is down loaded is left to you. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Operating JCL

The CICS Subsystem functions in much the same manner as the Workload Analysis Subsystem, but there are a few differences that stem from the nature of CICS systems:

- Each CICS region is treated as a separate entity.

When you want to examine the performance of CICS, you normally look at a specific region executing on a specific MVS system. Thus, each of the CICS regions in the data center must be registered for each of the MVS systems under which it can execute.

- Separate ONLINE table files.

You might want to allocate separate ONLINE table files for the CICS tables related to a given MVS system or you might even want to allocate separate ONLINE table files for each separate CICS region. When processing requests that relate to several CICS regions or a single CICS region in concert with the Workload tables, simply concatenate the required table files and build a temporary data base and Index using the CIMS Capacity Planner Utility program—SSA1LOAD.

CICS Subsystem Installation JCL

The JCL required to install the CIMS Capacity Planner CICS Subsystem is described in the *CIMS Capacity Planner User Guide*.

Data Reduction JCL

CICS input data can come from a variety of sources. CIMS Capacity Planner supports the following input formats:

- CICS 110 SMF records
- A history tape created by the Landmark Monitor

Extract programs are supplied for both input record formats. These programs parse the input data and produce a 200 byte record suitable for further processing by either the CIMS Capacity Planner CICS data reduction program (SSA1CICW) or the CIMS Chargeback program (CIMSCMF2).

Regardless of the format of the input, all data must be presented to the CICS data reduction program in chronological order. CIMS Capacity Planner maintains a table in the Performance Database showing the time and date of the last transaction processed. If a transaction is read into the data reduction program that is time stamped earlier than the latest transaction in the table, the transaction file is considered a duplicate of a previously processed file. A message is issued to indicate that duplicate data was found in the input stream and the transaction is discarded. If the user wishes to override this duplicate check logic, it is necessary to place the following control statement in the SYSIN:

```
FORCE CICS INPUT=YES
```

About CICS 110 SMF Records

The CMF2 Extract program (SSA1CMFX) maintains a PDS (CMF2DCTN) that contains mapping information used to parse the 110 records. This mapping information originates with the CICS 110 dictionary records, which are cut when the region is initially brought up or when the mapping information changes. A separate mapping member is maintained for each individual region. The CMF2DCTN PDS may also be used and maintained by the extract facility embedded within the CIMS Capacity Planner Workload Data Reduction program (SSA1WKLD).

CMF2 records cannot be extracted from CICS 110 records until a set of dictionary records is encountered for any given region. The CICS 110 records are discarded until the initial set of dictionary records is found. This problem can be resolved in any of the following ways:

- Because the dictionary records are written each time the region is brought up, locate a file that was written immediately after the region was brought up and begin CMF2 Extract processing with that file. Subsequent SMF data sets containing CICS 110 records for that region will be processed successfully.
- If the records associated with the file containing the dictionary records have already been processed in *CIMS Capacity Planner*, it is possible to populate the CMF2DCTN PDS member for a particular region without actually reducing the CICS data in the Performance Database. Simply pass the dictionary records through the SSA1CMFX program and discard the extracted CMF2 records.
- If all of the CICS regions have exactly the same mapping, it is possible to propagate the mapping information once a single CMF2DCTN PDS member has been built. This is done by replicating the original member, storing it under a new name that matches the VTAM APPLID of the region you want to map, and using ISPF/PDF Edit to change the original VTAM APPLID in the body of the member to the correct name.

About Landmark Monitor Records

The Landmark Monitor records can be processed in their native state, either compressed or uncompressed. To uncompress the Landmark Monitor records, you need to dynamically load a processing routine from the Landmark Monitor LOADLIB. The sample JCL on [page 3-19](#) contains a reference to this LOADLIB.

Processing CICS 110 SMF records

The following is sample JCL that can be used to process a set of CICS 110 SMF records. This JCL is in member DCICSMF in the &PREFIX.CPPR.Vnnn.CNTL library.

Note that the JCL contains three steps:

- The CMF2 Extract program (SSA1CMFX)
- The SORT step to sort the records by APPLID, Date, Time and Transaction
- The CPPR Data reduction program (SSA1CICW)

Note • If you are using CICS Transaction Server 1.1 or later, you no longer need to supply the GMT OFFSET in this step. The adjustment for GMT will be received dynamically from the CICS 110 SMF records.

CICS 110 SMF Record Processing Job Control

```
//CPPRCICS JOB (...),'CIMS',CLASS=A,MSGCLASS=X
//*****
//* THIS STEP EXTRACTS THE C110 CICS DATA INTO A CMF2          *
//* RECORD, SUITABLE FOR PROCESSING BY THE CHARGEBACK SYSTEM    *
//*****
//ST1001 EXEC PGM=SSA1CMFX,REGION=OM
//STEPLIB DD DSN=CIMS.CPPR.V530.LOADLIB,DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//*****
//* THE FOLLOWING LIBRARY CONTAINS THE DICTIONARY RECORDS      *
//*****
//CMF2DCTN DD DISP=SHR,DSN=CIMS.CPPR.CMF2DCTN
//*****
//* THE FOLLOWING FILE CONTAINS THE C110 INPUT RECORDS        *
//*****
//SYSUT1 DD DISP=SHR,DSN=CIMS.SAMPDATA.CICS110.DATA
//*****
//* THE FOLLOWING FILE CONTAINS THE CMF2 OUTPUT RECORDS        *
//*****
//SYSUT2 DD DISP=(,CATLG),UNIT=SYSDA,SPACE=(CYL,(10,10),RLSE),
//          DCB=(LRECL=200,BLKSIZE=27800,RECFM=FB),
//          DSN=CIMS.CPPR.C110XTR1.SYSUT2
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
```



```

//SYMSGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=PROD
//*
/*****
/* THE NEXT STEP SORTS THE OUTPUT BY TRANSACTION          *
/* NAME WITHIN TIME WITHIN DATE WITHIN APPLID             *
/*****
//ST2001 EXEC PGM=SORT,REGION=OM
//*
//SYSOUT DD SYSOUT=*
//*
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK04 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//*
//SORTIN DD DISP=SHR,DSN=*.ST1001.SYSUT2
//*
//SORTOUT DD DSN=CIMS.CPPR.CMF2.SORTED,
//          DISP=(,CATLG),UNIT=SYSDA,SPACE=(CYL,(50,10),RLSE),
//          DCB=(RECFM=FB,LRECL=200,BLKSIZE=27800)
//*
//SYSIN DD *,DCB=BLKSIZE=80
SORT FIELDS=(057,08,CH,A,009,04,PD,A,045,04,CH,A,005,04,BI,A)
/*
/*****
/* THE NEXT STEP REDUCES THE CMF2 RECORDS INTO THE CPPR PDB *
/*****
//ST3001 EXEC PGM=SSA1CICW,REGION=OM
//STEPLIB DD DSN=CIMS.CPPR.V530.LOADLIB,DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ONLINE DD DISP=SHR,
//          DSN=CIMS.CPPR.PROD.ONLINE.CICS
//INDEX DD DISP=SHR,
//          DSN=CIMS.CPPR.PROD.INDEX.CICS
//CPPRPARM DD DISP=SHR,DSN=CIMS.CPPR.PROD.PARMLIB

```

```
//CPPRERT DD DISP=SHR,DSN=CIMS.CPPR.PROD.CPPRERT
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSUT1 DD DISP=SHR,DSN=*.ST2001.SORTOUT
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=PROD

SMFILE=CMF2
DUMP SMF STATISTICS=YES

NO SMF SID=YES
* The following presumes Local time is 6 hours west of GMT

GMT OFFSET=W,6
```

The ST1001 step in the preceding JCL saves the Dictionary records in the PDS referenced by the CMF2DCTN DD statement. This PDS must be allocated prior to the execution of the DCICSMF JCL. You can use the following JCL to allocate the PDS:

```
//CIMSCPPR JOB (...),'CIMS',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/*****
/* This step allocates and initializes the online file *
/* that will hold the scratch pad records (Dictionary) *
/*****
//ST1 EXEC PGM=IEFBR14
//CMF2DCTN DD DISP=(,CATLG),UNIT=SYSDA,SPACE=(TRK,(1,1,25)),
// DCB=(LRECL=3094,BLKSIZE=3094,RECFM=F),
// DSN=&PREFIX.CPPR.PROD.CMF2DCTN
```

As the CMF2 Extract program (SSA1CMFX) encounters Dictionary records, members are built and stored in the CMF2DCTN library for each individual CICS Region named by APPLID. The members are updated as new dictionary records are read.

Processing Landmark Monitor Records

The following is sample JCL that can be used to process a set of Landmark Monitor records. This JCL is in member DCICTMON in the &PREFIX.CPPR.Vnnn.CNTL library.

Note that the JCL contains three steps:

- The CMF2 Extract program (SSA1TMNX)
- The SORT step to sort the records by APPLID, Date, Time and Transaction
- The CPPR Data reduction program (SSA1CICW)

Landmark Monitor Record Processing Job Control

```
//CPPRCICS JOB (...), 'CIMS', CLASS=A, MSGCLASS=X
//*****
//* THE FIRST STEP EXTRACTS THE TMON CICS DATA INTO A CMF2      *
//* RECORD, SUITABLE FOR PROCESSING BY THE CHARGEBACK SYSTEM    *
//*****
//ST1001 EXEC PGM=SSA1TMNX, REGION=OM
//STEPLIB DD DSN=CIMS.CPPR.V530.LOADLIB, DISP=SHR
//          DD DISP=SHR, DSN=LANDMARK.TMON20.TCELOAD
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRPARM DD DISP=SHR, DSN=CIMS.CPPR.PROD.PARMLIB
//CPPRERT DD DISP=SHR, DSN=CIMS.CPPR.PROD.CPPRERT
//*****
//* THE FOLLOWING FILE CONTAINS THE TMON RECORDS                *
//*****
//SYSUT1  DD DISP=SHR, DSN=LANDMARK.HISTORY.RECORDS(-1)
//*****
//* THE FOLLOWING FILE CONTAINS THE CMF2 RECORDS                *
//*****
//SYSUT2  DD DISP=SHR,
//          DSN=CIMS.CPPR.TMONXTR1.SYSUT2,
//          DCB=(LRECL=200, BLKSIZE=27800, RECFM=FB)
//SYSPRINT DD SYSOUT=*
//SYSMSGSGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=PROD
//*
//*****
//* THE NEXT STEP SORTS THE OUTPUT BY TRANSACTION              *
```

```

//* NAME WITHIN TIME WITHIN DATE WITHIN APPLID *
//*****
//ST2001 EXEC PGM=SORT,REGION=OM
//*
//SYSOUT DD SYSOUT=*
//*
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//SORTWK04 DD UNIT=SYSDA,SPACE=(CYL,(50),,CONTIG)
//*
//SORTIN DD DISP=SHR,DSN=*.ST1001.SYSUT2
//*
//SORTOUT DD DSN=CIMS.CPPR.CMF2.SORTED,
// DISP=(,CATLG),UNIT=SYSDA,SPACE=(TRK,(100,10),RLSE),
// DCB=(RECFM=FB,LRECL=200,BLKSIZE=27800)
//*
//SYSIN DD *,DCB=BLKSIZE=80
SORT FIELDS=(057,08,CH,A,009,04,PD,A,045,04,CH,A,005,04,BI,A)
/*
//*****
//* THE NEXT STEP REDUCES THE CMF2 RECORDS INTO THE CPPR PDB *
//*****
//ST3001 EXEC PGM=SSA1CICW,REGION=OM
//STEPLIB DD DSN=CIMS.CPPR.V530.LOADLIB,DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ONLINE DD DISP=SHR,
// DSN=CIMS.CPPR.PROD.ONLINE.CICS
//INDEX DD DISP=SHR,
// DSN=CIMS.CPPR.PROD.INDEX.CICS
//CPPRPARM DD DISP=SHR,DSN=CIMS.CPPR.PROD.PARMLIB
//CPPRERT DD DISP=SHR,DSN=CIMS.CPPR.PROD.CPPRERT
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSUT1 DD DISP=SHR,DSN=*.ST2001.SORTOUT
//SYSRINT DD SYSOUT=*
//SYSMSGG DD SYSOUT=*
//SYSIN DD *

```

```
SELECTED SYSTEM=PROD
SMFILE=CMF2
DUMP SMF STATISTICS=YES
NO SMF SID=YES
```

CIMS Capacity Planner CICS Ad Hoc Report

CIMS Capacity Planner provides you the option of producing one or more reports from the original input data. These reports can contain much of the original detail information that can be discarded or lost from summarization during the data reduction process.

The CICS Ad Hoc report is generated by the CICS data reduction program (SSA1CICW), when all the detail data is available. The Ad Hoc reports are generated on an as-needed basis and it is likely, therefore, that the need for the Ad Hoc reports is determined after the original data are already processed and stored into the Performance Data Base. In such a case, it is important that the tables created by the data reduction program not be stored into the Performance Data Base twice. This is accomplished by specifying DD DUMMY for both the ONLINE and the INDEX data sets.

In general, you can limit the records selected for inclusion in the Ad Hoc reports by including the following parameters in the SYSIN data set in the Job stream:

- a specific transaction ID
CICS TRANSACTION NAME=CS*
- a specific terminal ID
CICS TERMINAL NAME=LO*
- a date range
BEGIN DATE=02/04/2005
END DATE=02/07/2005
- a time period
BEGIN TIME=08.00.00
END TIME=10.00.00
- a specific CICS region
CICSNAME=CICSPROD

Using the data provided above, the Ad Hoc report produced would be limited to data from the CICS Region CICSPROD and would contain only data generated between 08:00 AM on February 4, 2005 and 10:00 AM on February 7, 2005. The inclusion of the data would be further limited to transactions beginning with the characters "CS" originating from terminals whose terminal IDs begin with the characters "LO".

To request a CICS Ad Hoc report, merely specify the terminals and transactions that you want to include in the report. To request that an Ad Hoc report not be produced, remove any selection criteria from the SSA1CICW Job Step.

Data Reduction DD Statements

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library. For OMEGAMON®/CICS processing, the Omegamon library containing the uncompress program must be concatenated to the CIMS Capacity Planner load library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all CICS Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ SYSUT1

The SYSUT1 DD statement specifies the input to the CIMS Capacity Planner Data Reduction Program - SSA1CICW.

■ SYSUT3

The SYSUT3 DD statement is used to specify a temporary data set used to hold a working copy of the Performance Data Base during the data reduction process. It is no longer needed after the performance data is copied back to the ONLINE file at the conclusion of the data reduction phase. The parameters that must be specified are: Disposition, Unit, and Space.

■ SSASPILL

The SSASPILL DD statement is used to specify a temporary data set to be used as a work area to hold data when selected non-IBM history files are being directly processed by CIMS Capacity Planner. The parameters that must be specified are: Disposition, Unit, and Space.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that contains the CICS Ad Hoc report. Unless otherwise specified, the DCB characteristics are:

```
RECFM=FBA
```

```
BLKSIZE=133
```

```
LRECL=133
```

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the CICS data reduction phase are described in the next paragraph.

CIMS Capacity Planner CICS Data Reduction Parameters

The CIMS Capacity Planner Parameters permit you to specify certain run-time options and to provide selected data required by the data reduction program in processing the CICS history files and producing the Ad Hoc report.

■ SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMFSID of the MVS system for which data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID can be specified by this parameter in running the CICS data reduction program.

Options:

The SELECTED SYSTEM= parameter can be specified in several ways:

- The SID itself (e.g. SELECTED SYSTEM=SYS1)
- An "*", indicating the SID for the system on which the program is being executed (e.g. selected system=*)

Default Option:

```
SELECTED SYSTEM=*
```

■ **BEGIN DATE**

The BEGIN DATE= parameter is used to specify the earliest date (inclusive) for which data is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

Default Date:

BEGIN DATE=00001

■ **END DATE**

The END DATE= parameter specifies the last date (inclusive) for which a CICS record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above.

Default Date:

END DATE=99365

■ **DUMP SMF STATISTICS**

The Dump SMF Statistics = parameter specifies whether the input record counts are to be printed for each of the CICS regions processed.

Options:

DUMP SMF STATISTICS = YES (Report the record counts)

DUMP SMF STATISTICS = NO (Don't report the record counts)

Default Option:

DUMP SMF STATISTICS = NO

■ REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Specify Report L:language:

ENGLISH, DEUTSCH or any combination of languages separated by commas.
(e.g, REPORT LANGUAGE=ENGLISH)

Default Language:

REPORT LANGUAGE=ENGLISH

■ FORCE CICS INPUT

The FORCE CICS INPUT parameter governs whether the dates of the CICS history file are checked to avoid detecting duplicate updates. It is useful to process logs that have been skipped.

Options:

FORCE CICS INPUT=YES

FORCE CICS INPUT=NO

Default Option:

FORCE CICS INPUT=NO

■ SMFILE

The SMFILE parameter specifies the type of input that is to be processed by the CICS data reduction program.

Options:

SMFILE=HISTORY (used for SMF records from a QSAM file)

SMFILE=LIVE (used for records from a VSAM SMF cluster)

SMFILE=ACTIVE (used for the currently active SMF cluster)

SMFILE=110J (used for CMF Journal records)

SMFILE=TMON (used for Landmark TMON records - regular and summary)

SMFILE=OMON (used for Omegamon/CICS history records)

SMFILE=JARS (used for JARS/CICS history records)

Default Option:

SMFILE=HISTORY

■ **DEBUGON**

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

Options:

None. The DEBUGON parameter has no operands.

Default Parameter:

None (equivalent to DEBUGOFF)

■ **CICS TRANSACTION NAME**

This parameter specifies the name of a transaction or group of transactions to be included in the CICS Ad Hoc report. The wild card character "*" is supported to specify a group of transactions beginning with a common set of letters.

Specify Transaction Name:

CICS TRANSACTION NAME=PROLL001 (select records related to transaction type PROLL001)

CICS TRANSACTION NAME=PR* (select records related with all transactions beginning with the characters "PR")

Default Name:

None

■ **CICS TERMINAL NAME**

The CICS TERMINAL NAME parameter limits the transactions selected for inclusion into the CICS Ad Hoc report to a specific terminal or group of terminals.

Specify Terminal Name:

CICS TERMINAL NAME=L101T001 (used to specify a single terminal)

CICS TERMINAL NAME=L101* (used to specify all terminals beginning with the characters "L101")

Default Name:

None

Report Production JCL

With the exception of the CICS Ad Hoc Report, all the CICS batch reports are produced by one of the following report generation programs:

- SSA1CICA—CICS Multi-Region Graph
- SSA1CICE—CICS Consolidated Graph
- SSA1CICV—CICS Distributed Response by Application
- SSA1CICY—CICS Special Graph
- SSA1CICT—CICS Trends Analysis/Summary
- SSA1CICG—CICS Trends Analysis Graph

DD Statements

The DD Statements required to execute the above programs are described below:

- STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

- CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

- CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all CICS Job streams.

- INDEX

The INDEX DD statement specifies the data set name and the disposition of the INDEX to the ONLINE (the Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

- ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

- HGDLIB

The HGDLIB DD statement specifies the Partitioned Data Set used to save the data points required to produce the various graphs using the Harvard Graphics system. The data points contained in the members of the HGDLIB PDS are downloaded to a PC and matched up with a set of Harvard Graphics templates that specify how the graphic data is to be presented. The parameters required are the Disposition and the Data Set Name.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that contains the CICS printed reports. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*

■ EXCLUDE/INCLUDE

This is the (optional) Element Exclusion file. The CICS Report Program uses it to force certain elements to be excluded from the reports or to force the inclusion of certain elements in the reports that would not normally be included. When you use the INCLUDE, only the specified elements are included in the reports.

The EXCLUDE/INCLUDE file must consist of 80 character records, 72 of which are interpreted. Columns 73 through 80 can be used for sequence numbers. Multiple elements can be entered in a single record, but they must be separated by commas. The format of the input statement follows:

- An element can be up to eight characters in length.
- Multiple elements can be included in a single record provided that they are separated by commas or blanks.
- Elements can be terminated by the wild card character "*", indicating that all elements that match up to the "wild card" are to be excluded/included.

For example, the following statement in the EXCLUDE file:

```
TSOP,CS*
```

would cause the report program SSACICR to exclude any data related to the transaction named TSOP. It would also cause the exclusion of all transactions beginning with the characters "CS"

The use of the INCLUDE causes the report program to limit the contents of the report to only those elements specified in the INCLUDE data set.

■ INCNAMES

The INCNAMES DD statement specifies an input data set containing a list of CICS regions that are to be INCLUDED in the CICS Trends Analysis. The use of the INCNAMES DD statement and the CICSNAME parameter are mutually exclusive.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the CICS data reduction phase are described in *Additional Report Options on page 3-43* below.

Report Job Streams

CIMS Capacity Planner Batch Reports

The standard CIMS Capacity Planner CICS batch reports are all produced by the CIMS Capacity Planner CICS Report Program—SSA1CICR.

Following is an example of the JCL required to run the report program:

```
//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1CICR, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.CICS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.CICS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
CICSNAME=cicsname
CICS EXCEPTION ANALYSIS=YES
CICS SYSTEM OVERVIEW=YES
CICS PERFORMANCE REPORT=YES
CICS TRANSACTION STATISTICS REPORT=YES
CICS TERMINAL STATISTICS REPORT=YES
CICS SUMMARY REPORT=YES
CICS TRANSACTION RESPONSE GRAPH=YES
CICS TERMINAL RESPONSE GRAPH=YES
```

```
CICS TRANSACTION ACTIVITY GRAPH=YES
CICS TERMINAL ACTIVITY GRAPH=YES
CICS TRANSACTION PROFILE=YES
CICS TRANSACTION NAME=tttttttt
CPU ACTIVITY GRAPH=YES
CICS I/O ACTIVITY GRAPH=YES
//
```

The SYSIN parameters used to request reports and to specify reporting options are described below in *Additional Report Options on page 3-43*.

Member DCICREPT in the PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that can be customized for your installation.

CICS Trends Analysis Summary Program

Following is an example of the JCL required to run the CICS Trends Report Program—SSA1CICT.

```
//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1CICT,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.CICS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.CICS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
CICSNAME=CICSPROD
1ST PERIOD BEGIN DATE=02/16/92
1ST PERIOD END DATE=04/15/92
2ND PERIOD BEGIN DATE=05/16/92
2ND PERIOD END DATE=07/15/92
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
LUNCH BREAK BEGIN HOUR=11
LUNCH BREAK END HOUR=12
** IF YOU WANT TO COMBINE MULTIPLE CICS REGIONS INTO A
```

```

** SINGLE REPORT, REMOVE THE CICSNAME PARAMETER STATEMENT
** FROM THE SYSIN DATA SET AND USE THE INCNAMES DD
** STATEMENT TO SPECIFY A DATA SET CONTAINING THE NAMES OF
** THE REGIONS TO BE INCLUDED.
*//INCNAMES DD *
*CICSPROD
*CICSTEST
CICS FILTERED VALUES=YES /* SEE EXCLUDE BELOW */
** THE FOLLOWING DD STATEMENT IS USED TO SPECIFY A FILE
** CONTAINING A LIST OF TRANSACTIONS THAT ARE TO BE
** EXCLUDED FROM THE TRENDS REPORT
//EXCLUDE DD *
CS*,NMON,TOSP
//

```

The SYSIN parameters used to request the trends reports and to specify the reporting options are described below under [Additional Report Options on page 3-43](#).

Member DCICTRPT in the &PREFIX.CPPR.Vnnn.CNTL library contains a set of model JCL that you can customize for use in your installation.

CICS Transaction Summary Matrix

Following is an example of the JCL required to run the CICS Transaction Summary Matrix Program—SSA1CICH.

```

//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//STA EXEC PGM=SSA1CICH, REGION=4096K, TIME=60
//STEPLIB DD DSN=SYS2.CPPRDEV.Vnnn.LOAD, DISP=SHR
/**STEPLIB DD DSN=SYS2.CPPRDIST.Vnnn.LOADLIB, DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ONLINE DD DSN=SYS2.CPPRXDM.ONLINE, DISP=SHR
//INDEX DD DSN=SYS2.CPPRXDM.INDEX, DISP=SHR
//RPTTABLE DD DSN=SYS2.CPPRXDM.RPTTABLE, DISP=SHR
//RPTINDEX DD DSN=SYS2.CPPRXDM.RPTINDEX, DISP=SHR
//SYSUT3 DD DISP=(,DELETE), UNIT=SYSDA, SPACE=(CYL,(1,1))
//CPPRERT DD DSN=SYS2.CPPRXDM.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=SYS2.CPPRXDM.PARMLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSIN DD *

```

```
BEGIN DATE=05/01/2005
END DATE=05/10/2005
BEGIN TIME=09.00
END TIME=14.00
SELECTED SYSTEM=MVSA
CICSNAME=CICS2
GRAPH PERIOD=WEEKLY
*GRAPH PERIOD=HOURLY
*FILTER=8
FILTER=1
AVERAGES=NONE
```

CICS Presentation Graphics Program

The CIMS Capacity Planner CICS Subsystem provides an extensive number of presentation-quality graphs that depict the CICS workload and performance. You can produce the graphs for individual CICS regions operating under a specific **SELECTED SYSTEM**. The data to be included in the graphs are produced by the CIMS Capacity Planner program (SSA1CICG) and written to selected members of the HGDLIB data set to be down-loaded to a PC and subsequently plotted by the Harvard Graphics presentation graphics system.

You can also produce CICS graphs for all CICS regions that are registered in the Element Registration Table for a specific **SELECTED SYSTEM**, consolidated together, by executing the CIMS Capacity Planner CICS Extension Graphs Program (SSA1CICE)

CICS Graph Program—(SSA1CICG) JCL

Following is an example of the JCL required to operate the SSA1CICG program to generate the CICS graph data points.

```
//SSAGRPH JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1CICG,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.CICS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.CICS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR*
```



```

CICSNAME=CICSPROD
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//

```

Several sets of sample Job streams for producing the CIMS Capacity Planner CICS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

- DCICHGRF—Produce Hourly Graphs
- DCICDGRF—Produce Daily Graphs
- DCICWGRF—Produce Weekly Graphs
- DCICMGRF—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program—SSA1CICG. To tailor the Job streams, merely modify the selection criteria as required.

CICS Consolidated Graph Program—(SSA1CICE) JCL

Following is an example of the JCL required to operate the SSA1CICE program to generate the consolidated CICS graph data points.

```

//SSAGRPH JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1CICE, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.CICS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.CICS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGs DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR*
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//

```

Several sets of sample Job streams for producing the CIMS Capacity Planner consolidated CICS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

DCICHGRE—Produce Hourly Graphs

DCICDGRE—Produce Daily Graphs

DCICWGRE—Produce Weekly Graphs

DCICMGRE—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program - SSA1CICE. To tailor the Job streams, merely modify the selection criteria as required.

CICS Reporting Parameters

■ SELECTED SYSTEM

The `SELECTED SYSTEM=` parameter specifies the SID of the system for which reports are to be produced or trends to be analyzed. This is the SMF identifier that is specified in the `SMFPRMxx` member of `SYS1.PARMLIB`.

Specify SID:

The `SELECTED SYSTEM` parameter has the following options:

The SID itself (e.g., `SELECTED SYSTEM=MVSA`)

An `"*"`, indicating the SID for the system on which the report program is being executed (e.g., `SELECTED SYSTEM=*`)

Default SID:

`SELECTED SYSTEM=*`

■ BEGIN DATE

The `BEGIN DATE=` parameter specifies the beginning of the period for which data in the Performance Data Base is to be collected for analysis and reporting.

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

Default Date:

`BEGIN DATE=00001`

■ END DATE

The END DATE= parameter specifies the last date (inclusive) for which a CICS record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in BEGIN DATE.

Default Date:

END DATE=99365

■ LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter *does not cause* any data to be removed from the Performance Data Base.

Specify Hour:

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

■ LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour:

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

■ REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

■ DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

■ PRIME SHIFT FIRST HOUR

The PRIME SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Options:

Specify the beginning of the Prime Shift

Default Parameter:

PRIME SHIFT FIRST HOUR=08

■ LATE SHIFT FIRST HOUR

The LATE SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Late Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Late Shift

Default Hour:

LATE SHIFT FIRST HOUR=16

■ CICSNAME

The CICSNAME= parameter designates the CICS region for which data is to be collected and reported.

Specify CICS Name:

CICSNAME=any CICS Region Name (VTAM APPLID) that has been registered using the SSA1REGC program (see CICS Installation Procedures in the *CIMS Capacity Planner Installation Guide*).

Default Name:

None

■ CICS SUMMARY REPORT

This parameter specifies whether the CICS Summary Report is to be produced.

Options:

CICS SUMMARY REPORT=YES (produce the report)

CICS SUMMARY REPORT=NO (don't produce the report)

Default Option:

CICS SUMMARY REPORT=NO

■ CICS TRANSACTION STATISTICS REPORT

This parameter specifies whether the CICS Transaction Statistics Report is to be produced.

Options:

CICS TRANSACTION STATISTICS REPORT=YES (produce the report)

CICS TRANSACTION STATISTICS REPORT=NO (don't produce the report)

Default Option:

CICS TRANSACTION STATISTICS REPORT=NO

■ CICS TERMINAL STATISTICS REPORT

This parameter specifies whether the Terminal Statistics Report is to be produced.

Options:

CICS TERMINAL STATISTICS REPORT=YES (produce the report)

CICS TERMINAL STATISTICS REPORT=NO (don't produce the report)

Default Option:

CICS TERMINAL STATISTICS REPORT=NO

■ CICS TRANSACTION ACTIVITY LIST

This parameter specifies whether the Transaction Activity List is to be produced.

Options:

CICS TRANSACTION ACTIVITY LIST=YES (produce the report)

CICS TRANSACTION ACTIVITY LIST=NO (don't produce the report)

Default Option:

CICS TRANSACTION ACTIVITY LIST=NO

■ **CICS PERFORMANCE REPORT**

This parameter specifies whether the CICS Performance Report is to be produced.

Options:

CICS PERFORMANCE REPORT=YES (produce the report)

CICS PERFORMANCE REPORT=NO (don't produce the report)

Default Option:

CICS PERFORMANCE REPORT=NO

■ **CICS TERMINAL ACTIVITY LIST**

This parameter specifies whether the CICS Transaction Activity List report is to be produced.

Options:

CICS TRANSACTION ACTIVITY LIST=YES (produce the report)

CICS TRANSACTION ACTIVITY LIST=NO (don't produce the report)

Default Option:

CICS TRANSACTION ACTIVITY LIST=NO

■ **CICS SYSTEM OVERVIEW**

This parameter specifies whether the CICS System Overview report is to be produced.

Options:

CICS SYSTEM OVERVIEW=YES (produce the report)

CICS SYSTEM OVERVIEW=NO (don't produce the report)

Default Option:

CICS SYSTEM OVERVIEW=NO

■ **CICS CPU ACTIVITY GRAPH**

This parameter specifies whether the CICS CPU Activity Graph is to be produced.

Options:

CICS CPU ACTIVITY GRAPH=YES (produce the Graph)

CICS CPU ACTIVITY GRAPH=NO (don't produce the report)

Default Option:

CICS CPU ACTIVITY GRAPH=NO

■ CICS I/O ACTIVITY GRAPH

This parameter specifies whether the CPU I/O Activity Graph is to be produced.

Options:

CICS I/O ACTIVITY GRAPH=YES (produce the graph)

CICS I/O ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

CICS I/O ACTIVITY GRAPH=NO

■ CICS TERMINAL ACTIVITY GRAPH

This parameter specifies whether the CICS Terminal Activity Graph is to be produced.

Options:

CICS TERMINAL ACTIVITY GRAPH=YES (produce the graph)

CICS TERMINAL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

CICS TERMINAL ACTIVITY GRAPH=NO

■ CICS TRANSACTION ACTIVITY GRAPH

This parameter specifies whether the CICS Transaction Activity Graph is to be produced.

Options:

CICS TRANSACTION ACTIVITY GRAPH=YES (produce the graph)

CICS TRANSACTION ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

CICS TRANSACTION ACTIVITY GRAPH=NO

■ CICS TRANSACTION RESPONSE GRAPH

This parameter specifies whether the CICS Transaction Response Graph is to be produced.

Options:

CICS TRANSACTION RESPONSE GRAPH=YES (produce the graph)

CICS TRANSACTION RESPONSE GRAPH=NO (don't produce the graph)

Default Option:

CICS TRANSACTION RESPONSE GRAPH=NO

■ **CICS TERMINAL RESPONSE GRAPH**

This parameter specifies whether the CICS Terminal Response Graph is to be produced.

Options:

CICS TERMINAL RESPONSE GRAPH=YES (produce the graph)

CICS TERMINAL RESPONSE GRAPH=NO (don't produce the graph)

Default Option:

CICS TERMINAL RESPONSE GRAPH=NO

■ **CICS TERMINAL ACTIVITY GRAPH**

This parameter specifies whether the CICS Terminal Activity Graph is to be produced.

Options:

CICS TERMINAL ACTIVITY GRAPH=YES (produce the graph)

CICS TERMINAL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

CICS TERMINAL ACTIVITY GRAPH=NO

■ **CICS EXCEPTION ANALYSIS**

This parameter specifies whether the CICS Exception Analysis Report is to be produced.

Options:

CICS EXCEPTION ANALYSIS=YES (produce the report)

CICS EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

CICS EXCEPTION ANALYSIS=NO

■ **CICS TRANSACTION PROFILE**

This parameter specifies whether the CICS Transaction Profile Report is to be produced.

Options:

CICS TRANSACTION PROFILE=YES (produce the report)

CICS TRANSACTION PROFILE=NO (don't produce the report)

Default Option:

CICS TRANSACTION PROFILE=NO

■ CICS TRANSACTION NAME

This parameter specifies whether the CICS Transaction Name Report is to be produced.

Options:

CICS TRANSACTION NAME=YES (produce the report)

CICS TRANSACTION NAME=NO (don't produce the report)

Default Option:

None

■ 1ST PERIOD BEGIN DATE

This parameter specifies

The beginning of the first of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1CICG)

Or

The beginning date for the entire period to be used for developing the Trends Analysis Graph (SSA1CICT) for which consolidated table information is to be selected for analysis.

Parameters:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

Default Parameter:

1ST PERIOD BEGIN DATE=00001

■ 1ST PERIOD END DATE

This parameter specifies the ending date of the first period for the Trends Analysis Summary Program (SSA1CICT) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1CICG).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

1ST PERIOD END DATE=99365

■ **2ND PERIOD BEGIN DATE**

This parameter specifies the beginning of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1CICT)

Specify Date:

Specify a beginning date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD BEGIN DATE=00001

■ **2ND PERIOD END DATE**

This parameter specifies the ending date of the second period for the Trends Analysis Summary Program (SSA1CICT) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1CICG).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD END DATE=99365

■ **CICS FILTERED VALUES**

This parameter indicates whether the INCLUDE/EXCLUDE file is to be used when producing a set of graph data point members in the HGDLIB

Options:

CICS FILTERED VALUES=YES

CICS FILTERED VALUED=NO

Default Option:

CICS FILTERED VALUES=NO

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems. A new statement

TOP LABEL=YES

causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Time Range as specified

- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

Another statement

```
FOOT1=UP TO 64 CHAR FOOTNOTE
```

can be used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) to be included on all reports where there is room.

A third statement

```
FOOT2=UP TO 64 CHAR 2ND FOOTNOTE
```

can be used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) to be included on all reports where there is room.

Additional Report Options

In addition to the INCLUDE/EXCLUDE options described in *DD Statements on page 3-27*, you can limit the various data elements that are contained in reports or graph data point members using Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it can be useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. You can prepare this type of a report by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx

■ EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****" or replaced by substitution - see below).

Finally, the statement :

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE OR EXCLUDE facilities; however, it does not preclude their use. The combination of the various EXCLUDE/INCLUDE features in conjunction with the Generic Element Mask is discussed below.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a GENERIC ELEMENT MASK=11110011. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL**. If GENERIC ELEMENT MASK=11111111 were to be specified, all records would be summarized into a single entry or be subject to substitution as discussed below.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the CICS report processing programs to intercept the key of each element (in this case, the CICS Transaction name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *  
  
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "???23??. The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the CICS report processing programs to intercept the key of each element (in this case, the CICS Transaction name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *  
  
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "????23??. The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

Use of GENERIC ELEMENT MASK with EXCLUDE/INCLUDE and/or INCLUSIVE or EXCLUSIVE ELEMENT MASK

The following discussion proscribes the rules for using the various element masks and the order in which selection and substitution is performed.

- EXCLUSIVE ELEMENT MASK and the INCLUSIVE ELEMENT MASK are mutually exclusive as are the EXCLUDE and the INCLUDE DD statements.
- EXCLUSIVE ELEMENT MASK and INCLUSIVE ELEMENT MASK must be used in conjunction with the EXCLUDE or INCLUDE DD statement respectively to define the key elements that are to be either excluded or included. Without the EXCLUDE or INCLUDE lists defined through the use of the DD statements, the EXCLUSIVE and the INCLUSIVE ELEMENT MASKS are meaningless.
- The EXCLUDE and INCLUDE lists defined through the use of the EXCLUDE and INCLUDE DD statements can be used in conjunction with the GENERIC ELEMENT MASK to limit the data selected for inclusion into the reports irrespective of whether an EXCLUSIVE or INCLUSIVE ELEMENT MASK is specified.
- EXCLUSIVE and INCLUSIVE ELEMENT MASKing is performed prior to the EXCLUDE/INCLUDE logic.
- GENERIC ELEMENT MASKing takes place after the EXCLUDE/INCLUDE logic is performed.
- If the GENERIC ELEMENT MASK is specified as all "1"s (11111111), then:
 - If a CPPR.PARMLIB member &sidCICF is present, the substitutions are made based upon the data contained in &sidCICF.
 - If the &sidCICF member is not present in CPPR.PARMLIB, then the entire key is set to "*"s and all records are summarized into a single entry.

Please refer to Chapter 3, CIMS Capacity Planner Graphics, in the *CIMS Capacity Planner Reference Guide* for a description of the CPPR.PARMLIB member &sidCICF. You can find a sample of the &sidCICF member in the CPPR.PARMLIB as member CPPRCICF.

ISPF/PDF CICS Subsystem Interface

The CIMS Capacity Planner CICS Subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM Graphs, generate the full range of batch reports to be printed on the system printer, and generate input to the PC Presentation Graphics system (Harvard Graphics). The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, CICS, etc.).

On-line Graphs

The following CICS GDDM graphs are available on-line through the ISPF Interface:

CICS Transaction Response Graph

The Transaction Response Graph depicts the average response time for all transactions run in a selected CICS region, divided into CPU, WAIT, and SUSPEND time. The graph shows the prime shift averages by hour, day, week, or month within the range of dates specified in the report request.

CICS Transaction Volume Graph

The CICS Transaction Volume Graph depicts the average volume for all transactions executed within a selected CICS region. Within the range of dates and the period specified in the report request, the graph shows the prime shift averages by hour, day, week, or month.

CICS Terminal Activity Graph

The CICS Terminal Activity Graph depicts the average number of terminals active for CICS within the range of dates specified in the report request. The graph shows the average number of terminals active by hour, day, week, or month, depending upon which is requested.

On-line Reports

You can produce the following reports on-line through the CIMS Capacity Planner CICS ISPF Interface.

Statistical Reports

CICS Transaction Statistics Report

The Transactions Statistics Report provides for a specified CICS Region, a list of the 40 CICS transactions ranked among the top 40 in the following categories:

- Most frequently executed
- Required the most cumulative residence time
- Consumed the most CPU time
- Issued the most calls to the File Control Program
- Issued the most calls to the Journal Control Program
- Issued the most requests for Temporary Storage
- Issued the most calls for Transient Storage
- Incurred the most Paging Activity
- Experienced the highest Storage Usage (high water mark)

CICS Terminal Statistics Report

The CICS Terminal Statistics Report shows, for a selected CICS region, which terminals were in the top 40 in the following categories:

- Submitted the most transactions
- Accumulated the most active time (a terminal is active during each 15 period throughout the day if it is working. Terminals that are logged on, but not submitting transactions are not considered to be active).

CICS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the CICS workload and how well the CICS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time
- The average number of active terminals (high watermark)
- The average number of CICS transactions processed per minute during the prime shift
- The average number of CICS-related EXCPs per second during the prime shift
- The average terminal response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift
- The average percentage of the CPU time consumed by CICS and the CICS applications during each shift within the measured period
- The average number of transactions processed by shift for the measured period

For a fuller description of the contents of the Summary Analysis Report, refer to the description contained under *CICS Summary Analysis Report on page 3-3* above.

CICS Trends Analysis Report

The on-line Trends Analysis Report shows the trends in selected characteristics of the CICS usage and performance that are indicators of the future CICS workload. Trends are reported for:

- CPU utilization, including a breakdown in terms of CICS overhead and task related usage
- The average number of transactions executed per minute during prime shift
- The average number of EXCP's per second during prime shift
- The average response time during prime shift
- The ratio of early shift to prime shift transactions
- The ratio of late shift to prime shift transactions
- The average ratio of active terminals on early shift compared to prime shift
- The average ratio of active terminals on late shift compared to prime shift

CICS Exception Analysis Report

The CICS online Exception Analysis Report shows, for the period specified in the report request, any 15 minute intervals during which the following predefined limits were exceeded. The limits are specified in the `CPPR.PARMLIB` member `sidXCPT`.

- Average transaction processing time
- Average terminal response time

CICS Ad Hoc Report

The CICS Ad Hoc Report contains the following information for each record selected for inclusion in the report.

- The CICS Region Name
- The time that the transaction began execution
- The time that the transaction completed execution
- The transaction ID
- The terminal ID
- The CICS Transaction processing program name
- The elapsed time for the transaction
- The CPU time used
- The Wait time

- The Suspend time
- The number of File Access Calls
- The number of Temporary Storage Calls
- The number of Transient Data Requests
- The number of Paging Operations
- The High Watermark in storage required

Batch Reports

A Batch Job submission panel is provided so that the full range of CICS batch reports and graphs described earlier under *CICS Workload and Performance Reports on page 3-3*, can be produced through the ISPF Interface. This simplifies the processes of operating CIMS Capacity Planner considerably.

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch JOBS instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch JOB
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- Saving the Chart as a GDF File (Batch)

These options are explained in detail below.

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual Harvard Graphics charts. As originally designed, the data members are created by Batch JOBS, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through Harvard Graphics. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch JOB is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of key phrases:

```
GDDM BATCH PROCESS=YES
```

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

```
GDDM PRINTER NAME=pppppppp
```

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQ DD Statement in the JCL.

Several members have been added to the Distribution CNTL file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch JOB. The member named "GDDMBATC" shows how to produce any of the CICS GDDM Graphs in Batch.

GDDMBAT Parmlib Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (CPPR201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST0001) WHICH MAY THEN BE
/* SENT TO A GDDM PRINTER (PRRINT01).
/* * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR, DSN=SYS1.GDDMSYM
```

```

/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF DD DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00)'
//SYSIN DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PPRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*
* | | | % CHARACTER LENGTH
* | | % CHARACTER WIDTH
* | HORIZONTAL OFFSET
* VERTICAL OFFSET
//INCLUDE DD *
CPPR201
GDDMBATC Parmlib Member
//SSAGDDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
//* * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR ALL OF
/* THE SUPPORTED CICS CHARTS AND SENDING THEM TO A PRINTER.
//* * * * *
//ST1 EXEC PGM=IKJEFT01,REGION=5000K,DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
// DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB

```

```
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//* THE FOLLOWING DDNAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF DD DISP=SHR,DSN=SYS1.GDDMGDF
//* THE FOLLOWING DDNAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD10)' /* CICS RESPONSE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD11)' /* TRANSACTION VOLUME /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD12)' /* TERMINAL VOLUME /*
//SYSIN DD *
CICSNAME=CICSPROD
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
* | | | % CHARACTER LENGTH
* | | % CHARACTER WIDTH
* | HORIZONTAL OFFSET
* VERTICAL OFFSET
```

Specifying GDDM Batch Print Controls

You can specify certain Print Controls to take effect when printing a chart in Batch. Use the key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter allows you to specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

You might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. To invoke these local templates in a Batch JOB, use a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter can be added to the SYSIN, where the "nnnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=ddddddd
```

where the "ddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeee
```

where the "eeeeeee" is the name under which the graph is saved in the GDF library.

Presentation Graphics Interface

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the CICS workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required workload, performance, capacity, and trends data from the Online Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics—release 3.0, plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, which describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and can be used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, can be used to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner CICS graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in *Graph Descriptions on page 3-60* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters:

"&sidpii" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly
- "iiq" is the graph identifier and consists of two alpha-numeric characters. The "pii" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under *Presentation Graphics Distribution Disk on page 3-59*. The qualifier "q" may or may not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are downloaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpii.DAT where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, you can determine the graph template easily. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to down-load the data points from HGDLIB to the PC using IND\$FILE. You can modify the BAT files as required to fit the naming conventions of your installation.

The names of the BAT files conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
 - "WKLD" - Workload graph data
 - "DASM" - DASM graph data
 - "CICS" - CICS graph data
 - "CICE" - CICS Extension graph data
 - "IDMS" - IDMS graph data
 - "IDME" - IDMS Extension graph data
 - "IMSG" - IMS graph data
 - "NETG" - Network graph data
 - "NETE" - Network Extension graph data
 - "NETL" _ Network Physical Lines graph data
 - "NETN" - Network Physical NCP graph data
 - "DB2S" - DB2 Subsystem graph data
 - "DB2U" - DB2 Connect ID graph data
 - "M204" - Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY

Workload	DASM	CICS	IDMS	IMS	Network	DB2	M 204
Hourly		Hourly	Hourly	Hourly	Hourly	Hourly	Hourly
Daily		Daily	Daily	Daily	Daily	Daily	Daily
Weekly	Weekly	Weekly	Weekly	Weekly	Weekly	Weekly	Weekly
Monthly	Monthly	Monthly	Monthly	Monthly	Monthly	Monthly	Monthly

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB data point members from the Host processor (assumes the use of IND\$FILE for the file transfer). The naming conventions are as follows:

.BAT File Naming Conventions

The .BAT filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- - 1= Hourly
- - 2=Daily
- - 3=Weekly
- - 4=Monthly

The graph period identifier is followed by a subsystem identifier:

- - WKLD=Workload
- - DASM=DASM
- - CICS=CICS
- - CICE=CICS Extension
- - IDMS=IDMS
- - IDME=IDMS Extension

- - IMS=IMS
- - NETG=Network
- - NETE=Network Extension
- - NETL=Network Physical Lines
- - NETN=Network Physical NCP

For example, the .BAT file used to download the HGDLIB members required to produce the Workload Daily graphs would be named HGD2WKLD.BAT.

.TPL File Naming Conventions

The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- - 1=Hourly
- - 2=Daily
- - 3=Weekly
- - 4=Monthly

The graph period identifier is followed by a 2-character graph identifier. In some cases, a 1-character qualifier is added, extending the graph identifier to a 3 characters. Refer to the Graphic Report Tables for the specific template names.

Graph Periods

See Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide* for a description of the hourly, daily, weekly, and monthly graph periods.

Graph Descriptions

See Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for descriptions of each of the graphs produced by the CIMS Capacity Planner CICS Subsystem.

Graphic Report Tables

The following tables summarize the relationship of members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-based graph program.

CICS Subsystem Regular Hourly Graphs

These graphs represent specific CICS regions executing within the &sid system. Therefore, CICS Region suffixes are required to uniquely identify a particular CICS region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICA	&sid170Z	(none)	CICS Transaction Response by Region
	&sid171Z	(none)	CICS Transactions Per Minute by Region
	&sid172Z	(none)	CICS CPU % Busy by Region
	&sid173Z	(none)	CICS EXCPs Per Second by Region
	&sid174Z	(none)	CICS Terminals Active by Region
	&sid176Z	(none)	CICS Daily Transaction Volume by Region
	&sid179Z	(none)	CICS Terminal Response by Region
SSA1CICG	&sid170*	(none)	CICS Transaction Response
	&sid171*	(none)	CICS Transaction Volume
	&sid172*	(none)	CICS CPU Utilization
	&sid173*	(none)	CICS I/O Activity
	&sid174*	(none)	CICS Terminal Activity
	&sid175*	(none)	CICS Transactions/Terminal
	&sid176*	(none)	CICS Transactions/Day
	&sid177*	(none)	CICS Transactions/Shift
	&sid178*	(none)	CICS Terminals/Shift
	&sid179*	(none)	CICS Terminal Response

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICY	&sid17A*	(none)	Pages:Dispatch Time
	&sid17B*	(none)	Total Dispatch Time
	&sid17C*	(none)	Total CPU
	&sid17D*	(none)	Application % CPU

*plus a one-character suffix representing the CICS region

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Hourly Extension Graphs

These graphs represent ALL CICS regions executing within the &sid system. Therefore, no CICS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICE	&sid1E1P	(none)	CPU Busy+CICS Prime Shift
	&sid1E3	&sidCICT	CICS Volume by Transaction
	&sid1E4	&sidCICR	CICS Volume by Response
	&sid1E1s	&sidCPUV	CICS Volume:CPU Busy

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Daily Graphs

These graphs represent specific CICS regions executing within the &sid system. Therefore, CICS Region suffixes are required to uniquely identify a particular CICS region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICA	&sid270Z	(none)	CICS Transaction Response by Region
	&sid271Z	(none)	CICS Transactions Per Minute by Region
	&sid272Z	(none)	CICS CPU % Busy by Region
	&sid273Z	(none)	CICS EXCPs Per Second by Region
	&sid274Z	(none)	CICS Terminals Active by Region
	&sid276Z	(none)	CICS Daily Transaction Volume by Region
	&sid279Z	(none)	CICS Terminal Response by Region
SSA1CICG	&sid270*	(none)	CICS Transaction Response
	&sid271*	(none)	CICS Transaction Volume
	&sid272*	(none)	CICS CPU Utilization
	&sid273*	(none)	CICS I/O Activity
	&sid274*	(none)	CICS Terminal Activity
	&sid275*	(none)	CICS Transactions/Terminal
	&sid276*	(none)	CICS Transactions/Day
	&sid277*	(none)	CICS Transactions/Shift
	&sid278*	(none)	CICS Terminals/Shift
	&sid279*	(none)	CICS Terminal Response
SSA1CICY	&sid27A*	(none)	Pages:Dispatch Time
	&sid27B*	(none)	Total Dispatch Time
	&sid27C*	(none)	Total CPU
	&sid27D*	(none)	Application % CPU

* plus a one-character suffix representing the CICS region

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Daily Extension Graphs

These graphs represent *ALL* CICS regions executing within the &sid system. Therefore, no CICS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICE	&sid2E1E	(none)	CPU Busy+CICS Early Shift
	&sid2E1P	(none)	CPU Busy+CICS Prime Shift
	&sid2E1L	(none)	CPU Busy+CICS Late Shift
	&sid2E2	&sidCICR &sidCICO	CICS Response by Org.
	&sid2E3	&sidCICT	CICS Volume by Transaction
	&sid2E4	&sidCICR	CICS Volume by Response
	&sid2E1s	&sidCPUV	CICS Volume:CPU Busy

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Weekly Graphs

These graphs represent specific CICS regions executing within the &sid system. Therefore, CICS Region suffixes are required to uniquely identify a particular CICS region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICA	&sid370Z	(none)	CICS Transaction Response by Region
	&sid371Z	(none)	CICS Transactions Per Minute by Region
	&sid372Z	(none)	CICS CPU % Busy by Region
	&sid373Z	(none)	CICS EXCPs Per Second by Region
	&sid374Z	(none)	CICS Terminals Active by Region
	&sid376Z	(none)	CICS Daily Transaction Volume by Region
	&sid379Z	(none)	CICS Terminal Response by Region
SSA1CICG	&sid370*	(none)	CICS Transaction Response
	&sid371*	(none)	CICS Transaction Volume
	&sid372*	(none)	CICS CPU Utilization
	&sid373*	(none)	CICS I/O Activity
	&sid374*	(none)	CICS Terminal Activity
	&sid375*	(none)	CICS Transactions/Terminal
	&sid376*	(none)	CICS Transactions/Day
	&sid377*	(none)	CICS Transactions/Shift
	&sid378*	(none)	CICS Terminals/Shift
	&sid379*	(none)	CICS Terminal Response
SSA1CICY	&sid37A*	(none)	Pages:Dispatch Time
	&sid37B*	(none)	Total Dispatch Time
	&sid37C*	(none)	Total CPU
	&sid37D*	(none)	Application % CPU

* plus a one-character suffix representing the CICS region

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Weekly Extension Graphs

These graphs represent ALL CICS regions executing within the &sid system. Therefore, no CICS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICE	&sid3E1E	(none)	CPU Busy+CICS Early Shift
	&sid3E1P	(none)	CPU Busy+CICS Prime Shift
	&sid3E1L	(none)	CPU Busy+CICS Late Shift
	&sid3E2	&sidCICR &sidCICO	CICS Response by Org.
	&sid3E3	&sidCICT	CICS Volume by Transaction
	&sid3E4	&sidCICR	CICS Volume by Response
	&sid3E1s	&sidCPUV	CICS Volume:CPU Busy

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Monthly Graphs

These graphs represent specific CICS regions executing within the &sid system. Therefore, CICS Region suffixes are required to uniquely identify a particular CICS region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICA	&sid470Z	(none)	CICS Transaction Response by Region
	&sid471Z	(none)	CICS Transactions Per Minute by Region
	&sid472Z	(none)	CICS CPU % Busy by Region
	&sid473Z	(none)	CICS EXCPs Per Second by Region
	&sid474Z	(none)	CICS Terminals Active by Region
	&sid476Z	(none)	CICS Daily Transaction Volume by Region
	&sid479Z	(none)	CICS Terminal Response by Region
SSA1CICG	&sid470*	(none)	CICS Transaction Response
	&sid471*	(none)	CICS Transaction Volume
	&sid472*	(none)	CICS CPU Utilization
	&sid473*	(none)	CICS I/O Activity
	&sid474*	(none)	CICS Terminal Activity
	&sid475*	(none)	CICS Transactions/Terminal
	&sid476*	(none)	CICS Transactions/Day
	&sid477*	(none)	CICS Transactions/Shift
	&sid478*	(none)	CICS Terminals/Shift
	&sid479*	(none)	CICS Terminal Response
SSA1CICY	&sid47A*	(none)	Pages:Dispatch Time
	&sid47B*	(none)	Total Dispatch Time
	&sid47C*	(none)	Total CPU
	&sid47D*	(none)	Application % CPU

* plus a one-character suffix representing the CICS region

&sid is the SMF System ID (e.g, CPPR)

CICS Subsystem Regular Monthly Extension Graphs

These graphs represent *ALL* CICS regions executing within the &sid system. Therefore, no CICS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICE	&sid4E1E	(none)	CPU Busy+CICS Early Shift
	&sid4E1P	(none)	CPU Busy+CICS Prime Shift
	&sid4E1L	(none)	CPU Busy+CICS Late Shift
	&sid4E2	&sidCICR &sidCICO	CICS Response by Org.
	&sid4E3	&sidCICT	CICS Volume by Transaction
	&sid4E4	&sidCICR	CICS Volume by Response
	&sid4E1s	&sidCPUV	CICS Volume:CPU Busy

&sid is the SMF System ID (e.g, CPPR)

CICS Special Graphs

Because the following graphs do not show an hourly, daily, etc., dimension, the fourth character in the member name (5) does not indicate a dimension,

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1CICF	&sid5C0*	&sidCICF	Response by Organization
	&sid5C1*	&sidCICF	I/O by Organization
	&sid5C2*	&sidCICF	Dispatch by Organization
	&sid5C3*	&sidCICF	CPU Time by Organization
	&sid5C4*	&sidCICF	CICS Volume by Shift by Organization

* plus a one-character suffix representing the CICS region

&sid is the SMF System ID (e.g, CPPR)

IDMS SUBSYSTEM

Subsystem Overview	4-3
IDMS Workload and Performance Reports	4-3
IDMS Summary Analysis Report	4-3
IDMS Response Performance Report	4-4
IDMS Transaction Statistics Report	4-4
IDMS Terminal Statistics Report	4-6
IDMS Exception Analysis Report	4-6
IDMS Ad Hoc Report	4-7
IDMS Graphs	4-8
IDMS Transaction Response Graph	4-8
IDMS Transaction Activity Graph	4-8
IDMS Terminal Activity Graph	4-8
IDMS CPU Activity Graph	4-9
IDMS D/B Activity Graph	4-9
IDMS I/O Activity Graph	4-9
IDMS Transaction Profile	4-9
IDMS Trends Analysis Report	4-10
Operating JCL	4-11
IDMS Subsystem Installation JCL	4-11
Data Reduction JCL	4-11
Report Production JCL	4-21
IDMS TRENDS ANALYSIS	4-34
Additional Report Options	4-36
ISPF/PDF IDMS Subsystem Interface	4-39
On-line Graphs	4-40
On-line Reports	4-40
IDMS Summary Analysis Report	4-41
IDMS Trends Analysis Report	4-42

IDMS Exception Analysis Report	4-42
IDMS Ad Hoc Report	4-42
Batch Reports	4-43
Batch GDDM Graphs	4-43
Presentation Graphics Introduction	4-48
Graph Production Procedures	4-48
Naming Conventions	4-49
Presentation Graphics Distribution Disk	4-51
Graph Periods	4-51
Graph Descriptions	4-52
GRAPHIC REPORT TABLES	4-52

Subsystem Overview

The CIMS Capacity Planner IDMS Subsystem provides a variety of reports and graphs related to the IDMS workload and the performance of your IDMS system and applications. The IDMS reports provide both detail and summary level information designed to meet the needs of both the technical staff and technical management. A variety of graphs are provided through both GDDM and the PC based Harvard Graphics system. GDDM provides the capability to instantly display a number of IDMS graphs on the Host through the use of the CIMS Capacity Planner ISPF interface. The Harvard Graphics system is used primarily to produce professional quality colored graphs that you can use for presentations to users and management.

IDMS Workload and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner IDMS reports and describes the contents of each report.

IDMS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the IDMS workload and how well the IDMS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the IDMS system.

- The average number of active terminals (high watermark)

A terminal is considered active during each 15-minute period throughout the day only if at least one transaction is submitted. If a terminal is logged on throughout three shifts, but transactions are submitted only during two 15 minute periods, the terminal is considered to be active for only 30 minutes.

- The average number of IDMS transactions processed per minute during the prime shift
- The average elapsed time for all IDMS transactions submitted during the prime shift
- The average number of IDMS data base calls made per second during the prime shift
- The average number of IDMS-related EXCPs per second during the prime shift
- The average IDMS terminal response time (in seconds) during the prime shift
- The average ratio of IDMS transactions submitted during the early shift compared to the prime shift

- The average ratio of IDMS transactions submitted during the late shift compared to the prime shift
- The average ratio of IDMS terminals active during the early shift compared to the prime shift
- The average ratio of IDMS terminals active on the late shift compared to the prime shift
- The average percentage of the CPU time consumed by IDMS and the IDMS applications during each shift within the measured period
- The average number of transactions processed by shift for the measured period

The 25 most frequently executed IDMS transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions are reported by shift.

IDMS Response Performance Report

This report shows, by shift, for a single day or a range of days, how the IDMS transactions flowed through the system for a specified IDMS region. Each of the 25 most frequently executed IDMS transactions are listed individually, by shift, with an indication of what percentage of the response times fell within certain predefined time-frames. Transactions that are not among the 25 most frequently executed are grouped together and reported under "OTHER". A totals line is accumulated and written to indicate how all transactions performed as a group.

This is a key report in measuring how well your performance goals are being met within the CPU by the IDMS system and the application programs.

IDMS Transaction Statistics Report

This report shows, for a specified IDMS region, which transactions were among the top 40 in the following categories:

- Most frequently executed

The transaction codes and the transaction counts of the 40 most frequently executed transactions are listed along with their respective percentages of the total IDMS transaction workload. The transactions are listed in descending frequency of execution order.

- Most cumulative elapsed time

The transaction codes and the total amount of elapsed (residence) time for each of 40 transactions experiencing the highest residency time are listed along with their respective percentages of the total IDMS transaction residence time. The transactions are listed in descending residency time order.

- Most cumulative CPU time

The transaction codes and the total amount of CPU time consumed for each of the 40 transactions that consumed the most CPU time along with their respective percentages of the total IDMS CPU time. The transactions are listed in descending CPU time order.

- Most Record Requests

The transaction codes and the number of record request calls are listed for each of the 40 transactions issuing the most record request calls along with their respective percentages of all the record request activity. The transactions are listed in descending order by the number of record request calls.

- Most Data Base Calls

The transaction codes and the number of Data Base Calls are listed for each of the 40 transactions issuing the most Data Base calls along with their respective percentages of all IDMS Data Base activity. The transactions are listed in descending order by Data Base activity.

- Highest Average Record Current of Run Unit

The value "Record Current of Run Unit" indicates the depth within the database that a transaction must go in order to resolve a request. This statistic shows the top 40 transactions in terms of average record current of run unit.

- Most EXCPs

The transaction codes and the number of EXCPs are listed for each of the 40 transactions requesting the most physical I/O activity along with their respective percentages of the total EXCPs issued for all IDMS transactions. The listing is produced in descending order by the number of EXCPs issued.

- Most Terminal Activity

The transaction codes and the amount of terminal activity is listed for the 40 transactions experiencing the highest number of terminal messages along with their respective percentage of all terminal activity related to IDMS transaction processing. The listing is produced in descending order by the amount of terminal activity.

- Most Scratch Requests

This statistic shows the average number of requests for scratch space for the top 40 tasks.

- Most Storage Used

The transaction code and the high watermark (in bytes) is listed for the 40 IDMS transactions requiring the most storage. The listing is produced in descending order by the amount of storage required by each transaction.

Selected transactions can be excluded from the Transaction Statistics Report through the use of the EXCLUDE facility (see JCL). Similarly, you can include transactions that would not normally be included using the INCLUDE facility (JCL).

Whenever any of the reported statistics for a transaction exceed six digits (1 million or more), the numbers are expressed in thousands and suffixed with the character "K".

IDMS Terminal Statistics Report

This report shows which terminals, for a specified IDMS Region, were among the top 40 in the following categories:

- **Most Active**

The Terminal ID and the amount of time active (hours:minutes) are reported along with the percentage of the combined active time for all active terminals for the 40 most active terminals. A terminal is considered active within any given 15 minute period only if transactions are received from the terminal. A terminal that is signed-on, but does not submit any transactions is not considered to be active.

- **Most Transactions Executed**

The Terminal ID and the number of transactions submitted are reported for the 40 terminals submitting the most transactions along with their respective percentages of the total number of transactions submitted during the measured period. The listing is produced in descending order by the number of transactions submitted.

Any counts that exceed 6 digits is expressed in thousands by suffixing the number with the character "K".

You can exclude selected terminals from the report using the EXCLUDE facility (see JCL). Similarly, you can include terminals that might not qualify among the top forty in the report using the INCLUDE facility (see JCL).

IDMS Exception Analysis Report

The IDMS Exception Analysis Report shows, for each fifteen minute period, for a specified IDMS region, all transactions and terminals whose response times exceeded the pre-defined thresholds specified in the CIMS Capacity Planner Parmlib member &sidXCPT. The Parmlib member &sidXCPT contains a number of threshold values for the installation. This member is used to identify the exception reporting thresholds for the various components of the overall environment.

IDMS Ad Hoc Report

The CIMS Capacity Planner IDMS Subsystem supports the production of an Ad Hoc Report from the original data before portions of the detail data are summarized or discarded. The Ad Hoc Report is produced by the Data Reduction program—SSA1IDMW.

To specify that the report is to be produced, you merely specify the report parameters in the Data Reduction Job stream as part of the SYSIN data set. The parameters that apply to the Ad Hoc report are:

IDMS TRANSACTION NAME=

IDMS TERMINAL NAME=

The IDMS Ad Hoc report contains the following information for each record selected for inclusion in the report.

- The IDMS Region Name
- The time that the transaction began execution
- The User ID
- The transaction ID
- The terminal ID
- The IDMS Transaction processing program name
- The elapsed time for the transaction
- The CPU time used
- The Wait time
- The Number of Records Within the Current Run Unit
- The number of Data Base Calls
- The number of Get Storage Calls
- The number of Scratch Requests
- The number of EXCPs
- The High Watermark in storage required

In addition to selecting records for inclusion into the Ad Hoc report by specific transaction and terminal names, you can select groups of both transactions and terminals using the wildcard character "*".

For example, you can select all transactions beginning with the characters "TR" by specifying IDMS TRANSACTION NAME=TR*. Similarly, you can select all terminals beginning with "L422" by specifying IDMS TERMINAL NAME=L422*. This convention can be extended to include all transactions by specifying IDMS TRANSACTION

NAME=*. IDMS TERMINAL NAME=* would result in the selection of all terminals. If both the transaction and terminal names of "*" were specified, all records would be selected.

The selection of records for inclusion into the Ad Hoc Report can be further restricted by specifying date and/or time ranges by using the BEGIN DATE=, END DATE=, BEGIN TIME=, and END TIME= parameters.

IDMS Graphs

The IDMS Subsystem produces a number of graphs on the mainframe that depict the magnitude of the workload and the level of performance of a specified IDMS Region.

IDMS Transaction Response Graph

The Transaction Response Graph shows, by day of the week, the peak and average transaction response times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the response time (both peak and average). In the middle of each line, both the average and peak values are printed. The lines, in bar graph format, form the graphic representation of the values. The peak values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak values and the average values are either equal or very nearly so.

You can produce the Transaction Response Graph for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

IDMS Transaction Activity Graph

The Transaction Activity Graph is organized in the same format as the Transaction and Terminal Response Graphs described above. It shows, by day of the week, the average and peak average number of transactions submitted from all the IDMS terminals during each 15-minute period throughout the day. The values along the horizontal axis represent the number of transactions submitted per minute.

IDMS Terminal Activity Graph

The Terminal Activity Graph shows, by day of the week, the number of terminals during each 15-minute period throughout the day. Both the peak average and the average values are reported. The values along the horizontal axis represent the number of active terminals.

IDMS CPU Activity Graph

The CPU Activity Graph shows, by day of the week, the percentage of the CPU processing capacity that was expended in processing IDMS transactions for a specified IDMS region. The average and peak average percentage values are reported in 15 minute intervals throughout the day. As is the case with the other graphs, you can specify the measured period to include anywhere from a day to several months.

IDMS D/B Activity Graph

This graph shows, for a specified IDMS region, the number of data base service calls handled per second in support of the IDMS applications. The activity shown is the number of data base service calls per second for each 15 minute period throughout each day of the week. Both the daily averages and the peak daily averages are shown using the (-) and (*) notation common to all the graphs.

IDMS I/O Activity Graph

This graph shows, for a specified IDMS region, the amount of EXCP activity related to processing transactions. Both the peak average and the average I/O rates are reported. The format of the I/O Activity Graph is identical to the graphs described above, except that the values along the horizontal axis represent the number of I/Os per second related to processing the IDMS workload.

IDMS Transaction Profile

The Transaction Profile Report shows, for a specified IDMS region type, how a specified IDMS transaction behaved for the following categories:

- How many times per day, on average, it was invoked
- The average response time per invocation of the transaction
- The average CPU time per invocation used to process the transaction
- The average wait time
- The average number of record requests
- The average number of data base calls.
- The average number of records currently within a run unit
- The average number of EXCPs per invocation
- The average number of storage requests per invocation
- The average number of scratch requests per invocation
- The average number of terminal I/Os per invocation

IDMS Trends Analysis Report

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. The IDMS elements that are selected for comparison are:

- CPU Utilization Statistics, including a breakdown in terms of Task Mode and System Mode CPU usage
- The number of transactions executed per minute computed as an average during Prime shift
- The average number of EXCPs per second during Prime shift
- The average response time during Prime Shift
- The ratio of Early shift to Prime shift transactions
- The ratio of Late shift to Prime shift transactions
- The ratio of Early shift to Prime shift terminals
- The ratio of Late shift to Prime shift terminals

This information is gathered and presented to you in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right hand side of the report
- In data suitable for graphing with the Harvard Graphics Program once it has been down loaded to a Personal Computer. The manner in which the data is down loaded is left to you. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Note • The Include/Exclude feature is supported in the IDMS Trends Analysis Report. However, the IDMS Filtered Values=Yes parameter must be specified to cause the Include/Exclude List to be recognized.

Operating JCL

The IDMS Subsystem functions in much the same manner as the Workload Analysis Subsystem, but there are a few differences that stem from the nature of IDMS systems:

- Each IDMS region is treated as a separate entity

When you want to examine the performance of IDMS, you normally look at a specific region executing on a specific MVS system. Thus each of the IDMS regions in the data center must be registered for each of the MVS systems under which it can execute.

- Separate ONLINE table files

You might want to allocate separate ONLINE table files for the IDMS tables related to a given MVS system or you might even want to allocate separate ONLINE table files for each separate IDMS region. When processing requests that relate to several IDMS regions or a single IDMS region in concert with the Workload tables, simply concatenate the required table files and build an Index using the CIMS Capacity Planner Utility program—SSA1INDX.

IDMS Subsystem Installation JCL

The JCL required to install the CIMS Capacity Planner IDMS Subsystem is described in the *CIMS Capacity Planner Installation Guide*.

Data Reduction JCL

The IDMS Subsystem Data Reduction program (SSA1IDMW) processes IDMS Log records, specifically:

- Prior to IDMS Rel 10.2

Statistics (type 6) records using subtype X'1C' records (task-wide statistics)

- IDMS Rel 10.2

Statistics (type 6) records using subtype 230 records

- IDMS Rel 10.2 SMF Records

IDMS Log records can be written to the SMF data set beginning with Rel 10.2 with a user SMF record number.

- IDMS Subtype 02 Records

If the IDMS Integrated Performance Monitor is not installed with Rel 10.2, Task-Wide Statistics records (subtype 02) can be processed instead of subtype 230 records.

- PMDC Short Task Records

- PMDC Long Task Records

Regardless of the format of the input, all data must be presented to the IDMS data reduction program in chronological order. CIMS Capacity Planner maintains a table showing the date and time of the last transaction processed. Should a transaction be read into the data reduction program that is time stamped earlier than the latest transaction in the file, the transaction file is considered a duplicate of a previously processed file and discarded unless the checking feature is overridden by the FORCE IDMS INPUT=YES parameter.

Processing IDMS Log Records

IDMS statistics records are written either to an IDMS Log data set or to the SMF cluster, which is subsequently unloaded to a sequential data set by an installation specified program. The resultant sequential data set, normally in variable length, blocked format is then processed into the CIMS Capacity Planner IDMS Data Reduction Program (SSA1IDMW). An example of the Job stream used to process the IDMS Log is shown below. The example presumes that the IDMS data was stored in the IDMS Log.

```
//SSADIDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner DATA REDUCTION STEP
/*
//ST1      EXEC PGM=SSA1IDMW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1   DD DISP=SHR,DSN=IDMS.LOGTAPE
//SYSUT3   DD DISP=(,PASS),SPACE=(CYL,(10,8)),
//          UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=&sid
DUMP SMF STATISTICS=YES
*      IF YOU WISH TO PROCESS DATA RELATED TO A SPECIFIC
*      IDMS ADDRESS SPACE, SPECIFY:
IDMSNAME=idmsname
*      WHERE idmsname IS THE NAME (APPLID) OF THE IDMS REGION
* * * * *
//
```

Member DIDML102 in the &PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that you might want to customize for your installation.

Processing IDMS SMF Records

If IDMS Log records are written to the SMF cluster, the following JCL can be used to run the IDMS Data Reduction program.

```
//SSAIDMS JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IDMW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1 DD DISP=(OLD,PASS),DSN=&&SRTOUT
//SYSUT3 DD DISP=(,PASS),SPACE=(CYL,(10,8)),
// UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
IDMSNAME=idmsname
SMFILE=SMFH
***** + THIS IS THE USER SMF RECORD NUMBER FOR IDMS
* |
* | + THIS IS THE INTERNAL CENTRAL VERSION #
* | |
* V V
SMF USER RECORD NUMBER=240,06
* IF YOU WANT TO USE OTHER THAN THE TRANSACTION ID AS THE KEY:
*PRIMARY ELEMENT KEY=PROGRAM /* TO USE PROGRAM NAME */
*PRIMARY ELEMENT KEY=USERID /* TO USE USER ID */
//
```

Member **DIDMPSMF** of **&PREFIX.CPPR.Vnnn.CNTL** contains a set of model JCL that you can customize for use in your installation.

Processing IDMS PMDC Records

PMDC from International Software Products produces several records either in the SMF cluster or in the PMDC log. The records that CIMS Capacity Planner processes are the Task records. Long Task records contain a great deal of detailed information and are the preferred record type. CIMS Capacity Planner can also process Short Task records. Because both Long Task records and Short Task records can exist in a single file, it is necessary to specify which of the record types are to be processed.

If an installation produces Long Task records, these records are processed by CIMS Capacity Planner as the default. If only Short Task records are produced, please specify in the SYSIN:

```
PMDC SHORT TASK RECORDS=YES
```

In order to tell CIMS Capacity Planner that PMDC records are to be processed, please specify in the SYSIN:

```
SMFILE=PMDH          /* PMDC RECORDS FROM HISTORY */
```

It is also necessary to tell CIMS Capacity Planner the Internal CV number for which records are to be processed. The internal CV number is located at x'1E' offset from the beginning of the PMDC Task record. Please specify in the SYSIN:

```
FILTER=00,nn         /* nn IS THE INTERNAL CV NUMBER */
```

Following is a sample set of JCL that can be used to reduce data from IDMS Rel 10 or records produced by Rel 10.2 without the Integrated Performance Monitor.

If IDMS Log records are written to the SMF cluster, the following JCL can be used to run the IDMS Data Reduction program.

```
//SSAIDMS JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSAIDMW, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1 DD DISP=(OLD,PASS), DSN=PMDCLOG
//SYSUT3 DD DISP=(,PASS), SPACE=(CYL,(10,8)),
// UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSGG DD SYSOUT=*
//SYSIN DD *
```

```
SELECTED SYSTEM=&sid
IDMSNAME=idmsname
FILTER=240,06
DUMP SMF STATISTICS=YES
*PMDC SHORT TASK RECORDS=YES
//
```

Member DIDMPMDC in the &PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that you might want to customize for your installation by specifying the appropriate parameters and removing the SSA1SPILL DD statement.

CIMS Capacity Planner IDMS Ad Hoc Report

CIMS Capacity Planner gives you the option of producing one or more reports from the original input data. These reports can contain much of the original detail information that can be discarded or lost during summarization during the data reduction process.

The IDMS Ad Hoc Report is generated by the IDMS data reduction program (SSA1IDMW), when all the detail data is available. The Ad Hoc Reports are generated on an as-needed basis and it is likely, therefore, that the need for the Ad Hoc Reports is determined after the original data are already processed and stored into the Performance Data Base. In such a case, it is important that the tables created by the data reduction program not be stored into the Performance Data Base twice. This is accomplished by specifying DD DUMMY for both the ONLINE and the INDEX data sets.

In general, you can limit the records selected for inclusion in the Ad Hoc Reports by including the following parameters in the SYSIN data set in the Job stream:

- a specific transaction ID

```
IDMS TRANSACTION NAME=CS*
```

- a specific terminal ID

```
IDMS TERMINAL NAME=LO*
```

- a date range

```
BEGIN DATE=02/04/2005
```

```
END DATE=02/07/2005
```

- a time period

```
BEGIN TIME=08.00.00
```

```
END TIME=10.00.00
```

- a specific IDMS region

```
IDMSNAME=IDMSPROD
```

Using the data provided above, the Ad Hoc Report produced would be limited to data from the IDMS Region IDMSPROD and would contain only data generated during the period from February 4, 2005 through February 7, 2005, in the time window beginning at 8:00 A.M. and ending at 10:00 A.M. The inclusion of the data would be further limited to transactions beginning with the characters "CS" originating from terminals whose terminal IDs begin with the characters "LO".

To request an IDMS Ad Hoc Report, merely specify the terminals and transactions that you want to include in the report. To request that an Ad Hoc Report not be produced, merely remove any selection criteria from the SSA1IDMW Job Step.

Data Reduction DD Statements

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all IDMS Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the INDEX to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ SYSUT1

The SYSUT1 DD statement specifies the input to the CIMS Capacity Planner Data Reduction Program - SSA1IDMW.

■ SYSUT3

The SYSUT3 DD statement is used to specify a temporary data set used to hold a working copy of the Performance Data Base during the data reduction process. It is no longer needed after the performance data is copied back to the ONLINE file at the conclusion of the data reduction phase. The parameters that must be specified are: Disposition, Unit, and Space.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the IDMS Ad Hoc Report. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the IDMS data reduction phase are described in the next paragraph.

CIMS Capacity Planner IDMS Data Reduction Parameters

The CIMS Capacity Planner Parameters let you specify certain run-time options and to provide selected data required by the data reduction program in processing the IDMS history files and producing the Ad Hoc Report.

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMFSID of the MVS system for which data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID can be specified by this parameter in running the IDMS data reduction program.

Parameters:

The SELECTED SYSTEM= parameter can be specified in several ways:

The SID itself (e.g. SELECTED SYSTEM=SYS1)

An "*", indicating the SID for the system on which the program is being executed (e.g. selected system=*)

Default Parameter:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the earliest date (inclusive) for which data is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which an IDMS record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above.

Default Date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

FORCE IDMS INPUT

The FORCE IDMS INPUT parameter governs whether the dates of the IDMS history file are checked to avoid detecting duplicate updates. It is useful to process logs that have been skipped for one reason or another.

Options:

FORCE IDMS INPUT=YES

FORCE IDMS INPUT=NO

Default Option:

FORCE IDMS INPUT=NO

SMFILE

The SMFILE parameter is used to specify the type of input that is to be processed by the IDMS data reduction program. Input can come from one of the following sources:

- From the IDMS Log file (either pre 10.2 or Release 10.2)
- From the live System SMF Clusters
- From a sequential dump of the SMF file (Release 10.2 only)
- From the IDMS Log file (subtype 02) when the Integrated Performance Monitor is not installed (Release 10.2 only).
- From the PMDC Monitor from International Software Products

Specify Input File Type:

- SMFILE=SMFL (live SMF file subtype 230)
- SMFILE=SMFH (sequential dump of SMF file subtype 230)
- SMFILE=ILOG (IDMS log file subtype 230 or subtype X'1C')
- SMFILE=L102 (IDMS log file subtype 02 Task-Wide Stats only)
- SMFILE=PMDH (PMDH History Tape)

Default File Type:

- SMFILE=ILOG

SMF USER RECORD NUMBER

This parameter is used to specify the User-Assigned SMF Record Number assigned to the IDMS Log records that are to come from the SMF file (either live or history). Because IDMS uses the internal CV number to keep track of IDMS VTAM APPLIDs (as opposed to the IDMSNAME), it is necessary to provide this association to the CIMS Capacity Planner Data Reduction Program also.

Specify User Record Number:

SMF USER RECORD NUMBER=nn,cv where,

"nn" is the user assigned SMF record number (in decimal)

and

"cv" is the internal CV number (in decimal)

Default User Record Number:

None

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

The DEBUGON parameter has no sub-parameters. If it is not included in the SYSIN data set, then the DEBUGON option is set to off.

IDMS TRANSACTION NAME

This parameter is used to specify the name of a transaction or group of transactions to be included in the IDMS Ad Hoc Report. The wild card character "*" is supported to specify a group of transactions beginning with a common set of letters.

Specify Transaction Name:

IDMS TRANSACTION NAME=PROLL001 (select records related to transaction type PROLL001)

IDMS TRANSACTION NAME=PR* (select records related with all transactions beginning with the characters "PR")

Default Transaction Name:

None

IDMS TERMINAL NAME

The IDMS TERMINAL NAME parameter is used to limit the transactions selected for inclusion into the IDMS Ad Hoc Report to a specific terminal or group of terminals.

Specify Terminal Name:

IDMS TERMINAL NAME=L101T001 (used to specify a single terminal)

IDMS TERMINAL NAME=L101* (used to specify all terminals beginning with the characters "L101")

Default Terminal Name:

None

Report Production JCL

With the exception of the IDMS Ad Hoc Report, all the IDMS batch reports are produced by one of the following report generation programs:

SSA1IDMR—IDMS Report Generation

SSA1IDMT—DMS Trends Analysis/Summary

SSA1IDMG—IDMS Trends Analysis Graph

DD Statements

The DD Statements required to execute the above programs are described below:

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all IDMS Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the INDEX to the ONLINE (the Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ HGDLIB

The HGDLIB DD statement specifies the Partitioned Data Set used to save the data points required to produce the various graphs using the Harvard Graphics system. The data points contained in the members of the HGDLIB PDS are downloaded to a PC and matched up with a set of CIMS Capacity Planner graph templates that specify how the graphic data is to be presented. The parameters required are the Disposition and the Data Set Name.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the IDMS printed reports. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*

■ EXCLUDE/INCLUDE

This is the (optional) Element Exclusion file. It is used by the IDMS Report Program to force certain elements to be excluded from the reports or to force the inclusion of certain elements in the reports that would not normally be included. When you use the INCLUDE, only the specified elements are included in the reports.

The EXCLUDE/INCLUDE file must consist of 80 character records, 72 of which are interpreted. Columns 73 through 80 can be used for sequence numbers. Multiple elements can be entered in a single record, but they must be separated by commas. The format of the input statement follows:

- An element can be up to eight characters in length.
- Multiple elements can be included in a single record provided that they are separated by commas or blanks.

- Elements can be terminated by the wild card character "*", indicating that all elements that match up to the "wild card" are to be EXCLUDED/INCLUDED.

For example, the following statement in the EXCLUDE file:

```
TSOP,CS*
```

would cause the report program SSAIDMR to exclude any data related to the transaction named TSOP. It would also cause the exclusion of all transactions beginning with the characters "CS"

The use of the INCLUDE causes the report program to limit the contents of the report to only those elements specified in the INCLUDE data set.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the IDMS data reduction phase are described in *Additional Report Options on page 4-36* below.

Report Job Streams

CIMS Capacity Planner Batch Reports

The standard CIMS Capacity Planner IDMS batch reports are all produced by the CIMS Capacity Planner IDMS Report Program—SSA1IDMR.

Following is an example of the JCL required to run the report program:

```
//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IDMR, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
IDMSNAME=idmsname
IDMS SUMMARY REPORT=YES
IDMS EXCEPTION ANALYSIS=YES
IDMS PERFORMANCE REPORT=YES
```

```

IDMS TRANSACTION STATISTICS REPORT=YES
IDMS TERMINAL STATISTICS REPORT=YES
IDMS TRANSACTION RESPONSE GRAPH=YES
IDMS TERMINAL RESPONSE GRAPH=YES
IDMS TRANSACTION ACTIVITY GRAPH=YES
IDMS TERMINAL ACTIVITY GRAPH=YES
IDMS TRANSACTION PROFILE=YES
IDMS TRANSACTION NAME=tttttttt
CPU ACTIVITY GRAPH=YES
IDMS I/O ACTIVITY GRAPH=YES
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
//

```

The SYSIN parameters used to request reports and to specify reporting options are described below under *Additional Report Options on page 4-36*.

Member DIDMREPT in the PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that can be customized for your installation.

IDMS Trends Analysis Summary Program

Following is an example of the JCL required to run the IDMS Trends Report Program - SSA1IDMT.

```

//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IDMT,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
IDMSNAME=IDMSPROD
1ST PERIOD BEGIN DATE=02/16/2005
1ST PERIOD END DATE=04/15/2005
2ND PERIOD BEGIN DATE=05/16/2005

```

```

2ND PERIOD END DATE=07/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
LUNCH BREAK BEGIN HOUR=11
LUNCH BREAK END HOUR=12
IDMS FILTERED VALUES=YES      /* SEE NOTE FOR //EXCLUDE*/
** THE FOLLOWING DD STATEMENT IS USED TO SPECIFY A FILE
** CONTAINING A LIST OF TRANSACTIONS THAT ARE TO BE
** EXCLUDED FROM THE TRENDS REPORT. THE IDMS FILTERED VALUES=YES
** PARAMETER MUST BE SPECIFIED FOR THE INCLUDE/EXCLUDE LIST TO BE
** RECOGNIZED.
//EXCLUDE DD *
CS*,NMON,TOSP
//

```

The SYSIN parameters used to request the trends reports and to specify the reporting options are described below under [Additional Report Options on page 4-36](#).

Member DIDMTRPT in the &PREFIX.CPPR.Vnnn.CNTL library contains a set of model JCL that you can customize for use in your installation.

IDMS Presentation Graphics Program

The CIMS Capacity Planner IDMS Subsystem provides an extensive number of presentation-quality graphs that depict the IDMS workload and performance. You can produce the graphs for individual IDMS regions operating under a specific SELECTED SYSTEM. The data to be included in the graphs are produced by the CIMS Capacity Planner program (SSA1IDMG) and written to selected members of the HGDLIB data set to be down-loaded to a PC and subsequently plotted by the Harvard Graphics presentation graphics system.

You can also produce IDMS graphs for All IDMS regions that are registered in the Element Registration Table for a specific SELECTED SYSTEM, consolidated together, by executing the CIMS Capacity Planner IDMS Extension Graphs Program (SSA1IDME)

IDMS Graph Program—(SSA1IDMG) JCL

Following is an example of the JCL required to operate the SSA1IDMG program to generate the IDMS graph data points.

```

//SSAGRPH JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IDMG,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS

```

```
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR
IDMSNAME=IDMSPROD
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of sample Job streams for producing the CIMS Capacity Planner IDMS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

DIDMHGRF—Produce Hourly Graphs

DIDMDGRF—Produce Daily Graphs

DIDMWGRF—Produce Weekly Graphs

DIDMMGRF—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program—SSA1IDMG. To tailor the Job streams, merely modify the selection criteria as required.

IDMS Consolidated Graph Program—(SSA1IDME) JCL

Following is an example of the JCL required to operate the SSA1IDME program to generate the consolidated IDMS graph data points.

```
//SSAGRPH JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IDME, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.IDMS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.IDMS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
```



```
//SYSUDUMP DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=CPPR
IDMSNAME=IDMSPROD
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of sample Job streams for producing the CIMS Capacity Planner consolidated IDMS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

- DIDMHGRE—Produce Hourly Graphs
- DIDMDGRE—Produce Daily Graphs
- DIDMWGRE—Produce Weekly Graphs
- DIDMMGRE—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program - SSA1IDME. To tailor the Job streams, merely modify the selection criteria as required.

IDMS Reporting Parameters

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SID of the system for which reports are to be produced or trends to be analyzed. This is the SMF identifier that is specified in the SMFPRMxx member of SYS1.PARMLIB.

Specify &SID:

The SELECTED SYSTEM parameter has the following options:

The SID itself (e.g., SELECTED SYSTEM=MVSA)

An "*", indicating the SID for the system on which the report program is being executed (e.g., SELECTED SYSTEM=*)

Default SID:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the beginning of the period for which data in the Performance Data Base is to be collected for analysis and reporting.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which a IDMS record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in paragraph 5.3.3.2

Default Date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

The DEBUGON parameter has no sub-parameters. The absence of the DEBUGON parameter turns the Debugging features off.

PRIME SHIFT FIRST HOUR

The PRIME SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Prime Shift (e.g, PRIME SHIFT FIRST HOUR=7)

Default Hour:

PRIME SHIFT FIRST HOUR=6

LATE SHIFT FIRST HOUR

The LATE SHIFT FIRST HOUR= parameter specifies the hour of the day on which the LATE Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the LATE Shift (e.g, LATE SHIFT FIRST HOUR=19)

Default Hour:

LATE SHIFT FIRST HOUR=18

LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter *does not cause* any data to be removed from the Performance Data Base.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

IDMSNAME

The IDMSNAME= parameter designates the IDMS region for which data is to be reported.

Specify Region Name:

IDMSNAME=any IDMS Region Name (VTAM APPLID) that has been registered using the SSA1REGD program (see Chapter 1, CIMS Capacity Planner PARMLIB Overview, of the *CIMS Capacity Planner Reference Guide* of this manual).

Default Region Name:

None

IDMS SUMMARY REPORT

This parameter is used to specify whether the IDMS Summary Analysis Report is to be produced.

Options:

IDMS SUMMARY REPORT=YES (produce the report)

IDMS SUMMARY REPORT=NO (don't produce the report)

Default Option:

IDMS SUMMARY REPORT=YES

IDMS TRANSACTION STATISTICS REPORT

This parameter specifies whether the IDMS Transaction Statistics Report is to be produced.

Options:

IDMS TRANSACTION STATISTICS REPORT=YES (produce the report)

IDMS TRANSACTION STATISTICS REPORT=NO (don't produce the report)

Default Option:

IDMS TRANSACTION STATISTICS REPORT=NO

IDMS TERMINAL STATISTICS REPORT

This parameter specifies whether the IDMS Terminal Statistics Report is to be produced.

Options:

IDMS TERMINAL STATISTICS REPORT=YES (produce the report)

IDMS TERMINAL STATISTICS REPORT=NO (don't produce the report)

Default Option:

IDMS TERMINAL STATISTICS REPORT=NO

IDMS PERFORMANCE REPORT

This parameter specifies whether the IDMS Response Performance Report is to be produced.

Options:

IDMS PERFORMANCE REPORT=YES (produce the report)

IDMS PERFORMANCE REPORT=NO (don't produce the report)

Default Option:

IDMS PERFORMANCE REPORT=NO

IDMS CPU ACTIVITY GRAPH

This parameter specifies whether the IDMS CPU Activity Graph is to be produced.

Options:

IDMS CPU ACTIVITY GRAPH=YES (produce the graph)

IDMS CPU ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

IDMS CPU ACTIVITY GRAPH=NO

IDMS EXCEPTION ANALYSIS

This parameter specifies whether the IDMS Exception Analysis Report is to be produced.

Options:

IDMS EXCEPTION ANALYSIS=YES (produce the report)

IDMS EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

IDMS EXCEPTION ANALYSIS=NO

IDMS I/O ACTIVITY GRAPH

This parameter specifies whether the IDMS CPU I/O Activity Graph is to be produced.

Options:

IDMS I/O ACTIVITY GRAPH=YES (produce the graph)

IDMS I/O ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

IDMS I/O ACTIVITY GRAPH=NO

IDMS TERMINAL ACTIVITY GRAPH

This parameter specifies whether the IDMS Terminal Activity Graph is to be produced.

Options:

IDMS TERMINAL ACTIVITY GRAPH=YES (produce the graph)

IDMS TERMINAL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

IDMS TERMINAL ACTIVITY GRAPH=NO

IDMS TRANSACTION RESPONSE GRAPH

This parameter specifies whether the IDMS Transaction Response Graph is to be produced.

Option:

IDMS TRANSACTION RESPONSE GRAPH=YES (produce the graph)

IDMS TRANSACTION RESPONSE GRAPH=NO (don't produce the graph)

Default Option:

IDMS TRANSACTION RESPONSE GRAPH=NO

IDMS TERMINAL RESPONSE GRAPH

This parameter specifies whether the IDMS Terminal Response Graph is to be produced.

Options:

IDMS TERMINAL RESPONSE GRAPH=YES (produce the graph)

IDMS TERMINAL RESPONSE GRAPH=NO (don't produce the graph)

Default Option:

IDMS TERMINAL RESPONSE GRAPH=NO

IDMS TRANSACTION ACTIVITY GRAPH

This parameter specifies whether the IDMS Transaction Activity Graph is to be produced.

Options:

IDMS TRANSACTION ACTIVITY GRAPH=YES (produce the graph)

IDMS TRANSACTION ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

IDMS TRANSACTION ACTIVITY GRAPH=NO

IDMS D/B ACTIVITY GRAPH

This parameter specifies whether the IDMS D/B Activity Graph is to be produced.

Options:

IDMS D/B ACTIVITY GRAPH=YES (produce the graph)

IDMS D/B ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

IDMS D/B ACTIVITY GRAPH=NO

IDMS TRANSACTION PROFILE

This parameter specifies whether the IDMS Transaction Profile Report is to be produced.

Options:

IDMS TRANSACTION PROFILE=YES (produce the report)

IDMS TRANSACTION PROFILE=NO (don't produce the report)

Default Option:

IDMS TRANSACTION PROFILE=NO

IDMS TRANSACTION NAME

This parameter specifies which transaction is to be profiled in the previous report.

Option:

IDMS TRANSACTION NAME=P0132

Default Option:

None

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems. A new statement

TOP LABEL=YES

causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time:End as specified
- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

Another statement

```
FOOT1=UP TO 64 CHAR FOOTNOTE
```

can be used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that are included on all reports where there is room.

FOOT2

A third statement

```
FOOT2=UP TO 64 CHAR 2ND FOOTNOTE
```

can be used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that are included on all reports where there is room.

IDMS TRENDS ANALYSIS

The Trends Analysis portion of the IDMS Subsystem produces either a concise report (SSA1IDMT) or a set of members in the HGDLIB that can be sent to the PC to form a set of presentation graphs (SSA1IDMG or SSA1IDME).

1ST PERIOD BEGIN DATE

This parameter is used to specify the beginning of the first of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1IDMT) or the beginning date for the entire period to be used for developing the Trends Analysis Graph (SSA1IDMG or SSA1IDME) for which consolidated table information is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

```
1ST PERIOD BEGIN DATE=00001
```


1ST PERIOD END DATE

This parameter is used to specify the ending date of the first period for the Trends Analysis Summary Program (SSA1IDMT) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1IDMG).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

1ST PERIOD END DATE=2099365

2ND PERIOD BEGIN DATE

This parameter is used to specify the beginning of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1IDMT)

Specify Date:

Specify a beginning date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD BEGIN DATE=00001

2ND PERIOD END DATE

This parameter is used to specify the ending date of the second period for the Trends Analysis Summary Program (SSA1IDMT).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD END DATE=2099365

IDMS FILTERED VALUES

This parameter is used by the graph program (SSA1IDMG) to produce HGDLIB members taking the INCLUDE/EXCLUDE dataset into consideration. If this parameter is not specified, the entire population of transactions is graphed, regardless of the presence of an INCLUDE/EXCLUDE file.

Specify:

IDMS FILTERED VALUES=YES

IDMS FILTERED VALUES=NO

Default:

IDMS FILTERED VALUES=NO

Additional Report Options

In addition to the INCLUDE/EXCLUDE options described in paragraph 5.3.1, it is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks let you mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it can be useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. You can prepare this type of a report by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 (the default) specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****" or replaced by substitution-see below). Finally, the statement:

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE or EXCLUDE facilities.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, merely specify a GENERIC ELEMENT MASK=00001100. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL*. If GENERIC ELEMENT MASK=11111111 were specified, all records would be summarized into a single entry.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the IDMS report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *
      xxxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "???23??. The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the IDMS report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *
          xxxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "???23??. The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

Use of GENERIC ELEMENT MASK with EXCLUDE/INCLUDE and/or INCLUSIVE or EXCLUSIVE ELEMENT MASK

The following discussion addresses the rules for using the various element masks and the order in which selection and substitution is performed.

- EXCLUSIVE ELEMENT MASK and the INCLUSIVE ELEMENT MASK are mutually exclusive as are the EXCLUDE and the INCLUDE DD statements.
- EXCLUSIVE ELEMENT MASK and INCLUSIVE ELEMENT MASK must be used in conjunction with the EXCLUDE or INCLUDE DD statement respectively to define the key elements that are to be either excluded or included. Without the EXCLUDE or INCLUDE lists defined through the use of the DD statements, the EXCLUSIVE and the INCLUSIVE ELEMENT MASKs are meaningless.
- The EXCLUDE and INCLUDE lists defined through the use of the EXCLUDE and INCLUDE DD statements can be used in conjunction with the GENERIC ELEMENT MASK to limit the data selected for inclusion into the reports irrespective of whether an EXCLUSIVE or INCLUSIVE ELEMENT MASK is specified.
- EXCLUSIVE and INCLUSIVE ELEMENT MASKing is performed prior to the EXCLUDE/INCLUDE logic.
- GENERIC ELEMENT MASKing takes place after the EXCLUDE/INCLUDE logic is performed.

- If the GENERIC ELEMENT MASK is specified as all "1"s (11111111), then:

If a CPPR.PARMLIB member &sidIDMF is present, the substitutions are made based upon the data contained in &sidIDMF.

If the &sidIDMF member is not present in CPPR.PARMLIB, then the entire key is set to "*"s and all records are summarized into a single entry.

Please refer to Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide* for a description of the CPPR.PARMLIB member &sidIDMF. You can find a sample of the &sidIDMF member in the CPPR.PARMLIB as member CPPRIDMF.

ISPF/PDF IDMS Subsystem Interface

The CIMS Capacity Planner IDMS Subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM Graphs, generate the full range of batch reports to be printed on the system printer, and generate input to the PC Presentation Graphics system (Harvard Graphics).

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, IDMS, etc.).

On-line Graphs

The following IDMS GDDM graphs are available on-line through the ISPF Interface:

IDMS Transaction Response Graph

The Transaction Response Graph depicts the average response time for all transactions run in a selected IDMS region, divided into CPU, WAIT, and SUSPEND time. The graph shows the prime shift averages by hour, day, week, or month within the range of dates specified in the report request.

IDMS Transaction Volume Graph

The IDMS Transaction Volume Graph depicts the average volume for all transactions executed within a selected IDMS region. Within the range of dates and the period specified in the report request, the graph shows the prime shift averages by hour, day, week, or month.

IDMS Terminal Activity Graph

The IDMS Terminal Activity Graph depicts the average number of terminals active for IDMS within the range of dates specified in the report request. The graph shows the average number of terminals active by hour, day, week, or month, depending upon which is requested.

On-line Reports

The following reports can be produced on-line through the CIMS Capacity Planner IDMS ISPF Interface.

Statistical Reports

IDMS Transaction Statistics Report

The Transactions Statistics Report provides for a specified IDMS Region, a list of the 40 IDMS transactions ranked among the top 40 in the following categories:

- Most frequently executed
- Required the most cumulative residence time
- Consumed the most CPU time
- Issued the most record requests
- Issued the most data base calls
- Had the highest average record current of run unit
- Issued the most EXCPs
- Issued the most terminal I/O
- Issued the most scratch requests

- Issued the most storage requests
- Required the most storage in bytes (high watermark)

IDMS Terminal Statistics Report

The IDMS Terminal Statistics Report shows, for a selected IDMS region, which terminals were in the top 40 in the following categories:

- Submitted the most transactions
- Accumulated the most active time (a terminal is active during each 15 period throughout the day if it is working. Terminals that are logged on, but not submitting transactions are not considered to be active).

IDMS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the IDMS workload and how well the IDMS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time
- The average number of active terminals (high watermark)
- The average number of IDMS transactions processed per minute during the prime shift
- The average number of IDMS record requests per second during the prime shift
- The average number of data base calls per second during the prime shift
- The average number of IDMS related EXCPs per second during prime shift
- The average terminal IDMS response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift
- The average percentage of the CPU time consumed by IDMS and the IDMS applications during each shift within the measured period.
- The average number of transactions processed by shift for the measured period.

For a full description of the contents of the Summary Analysis Report, please refer to the description contained under *IDMS Summary Analysis Report on page 4-3* above.

IDMS Trends Analysis Report

The on-line Trends Analysis Report shows the trends in selected characteristics of the IDMS usage and performance that are indicators of the future IDMS workload. Trends are reported for:

- CPU utilization, including a breakdown in terms of IDMS overhead and task related usage
- The average number of transactions executed per minute during prime shift
- The average number of EXCP's per second during prime shift
- The average response time during prime shift
- The ratio of early shift to prime shift transactions
- The ratio of late shift to prime shift transactions
- The average ratio of active terminals on early shift compared to prime shift
- The average ratio of active terminals on late shift compared to prime shift

IDMS Exception Analysis Report

The IDMS on-line Exception Analysis Report shows, for the period specified in the report request, any 15 minute intervals during which the following predefined limits were exceeded. The limits are specified in the CPPR.PARMLIB member &sidXCPT.

- Average transaction processing time
- Average terminal response time

IDMS Ad Hoc Report

The IDMS Ad Hoc Report contains the following information for each record selected for inclusion in the report.

- The IDMS Region Name
- The time that the transaction began execution
- The User ID
- The transaction ID
- The terminal ID
- The IDMS Transaction processing program name
- The elapsed time for the transaction
- The CPU time used
- The Wait time

- The Number of Records Within the Current Run Unit
- The number of Data Base Calls
- The number of Get Storage Calls
- The number of Scratch Requests
- The number of EXCPs
- The High Watermark in storage required

For additional information related to the Ad Hoc report, please refer to the IDMS Ad Hoc Report above.

Batch Reports

A Batch Job submission panel is provided so that you can produce the full range of IDMS batch reports and graphs described earlier under paragraph 2.0, IDMS Workload and Performance Reports, through the ISPF Interface. This considerably simplifies the processes of operating CIMS Capacity Planner.

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch JOBS instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch JOB
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- Saving the the Chart as a GDF File (Batch)

These options are explained in detail below.

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual Harvard Graphics charts. As originally designed, the data members are created by Batch JOBS, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through Harvard Graphics. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch JOB is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of several key phrases.

```
GDDM BATCH PROCESS=YES
```

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

```
GDDM PRINTER NAME=pppppppp
```

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQ DD Statement in the JCL.

Several members have been added to the Distribution CNTL file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch JOB. The member named "GDDMBATD" shows how to produce any of the IDMS GDDM graphs in Batch.

GDDMBAT Parmlib Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (CPPR201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST0001) WHICH MAY THEN BE
/* SENT TO A GDDM PRINTER (PRRINT01).
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn. LOADLIB, DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
```

```

//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1 DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF DD DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00)'
//SYSIN DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PPRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*          | | | % CHARACTER LENGTH
*          | | % CHARACTER WIDTH
*          | HORIZONTAL OFFSET
*          VERTICAL OFFSET
//INCLUDE DD *
CPPR201
GDDMBATD Parmlib Member
//SSAGDDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR ALL OF
/* THE SUPPORTED IDMS CHARTS AND SENDING THEM TO A PRINTER.
/* * * * * *
//STO EXEC PGM=IKJEFT01,REGION=5000K,DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
// DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*

```

```

//SYSTSPRT DD  SYSOUT=*
//SYSNAP   DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//CPPRERT  DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE   DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX    DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD  DISP=SHR,DSN=SYS1.GDDMSYM
//ADMCDATA DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD  DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF    DD  DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING DD NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD  DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN  DD  *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD20)' /* IDMS RESPONSE      /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD21)' /* TRANSACTION VOLUME /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD22)' /* TERMINAL VOLUME   /*
//SYSIN    DD  *
IDMSNAME=IDMSPROD
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST0001
GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*           | | | % CHARACTER LENGTH
*           | | % CHARACTER WIDTH
*           | HORIZONTAL OFFSET
*           VERTICAL OFFSET

```

Specifying GDDM Batch Print Controls

It is possible to specify certain Print Controls that are to take effect when printing a chart in Batch. The key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter lets you specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

You might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. In order to invoke these local templates in a Batch JOB, use a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter can be added to the SYSIN, where the "nnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=dddddddd
```

where the "dddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeeee
```

where the "eeeeeeee" is the name under which the graph is saved in the GDF library

Presentation Graphics Introduction

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the IDMS workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics - release 3.0, plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the IDMS workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, which describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and can be used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, you can use to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner IDMS graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of IDMS graphs specified in the list set forth in *Graph Descriptions on page 4-52* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters in the form:

"&sidpii" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly
- "iiq" is the graph identifier and consists of two alpha-numeric characters. The "pii" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under *Presentation Graphics Distribution Disk on page 4-51*. The qualifier "q" might or might not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpii.DAT, where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to down-load the data points from HGDLIB to the PC using IND\$FILE. You can modify the BAT files as required to fit the naming conventions of your installation.

The names of the BAT files all conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies the period (hourly, daily, weekly, or monthly) as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:

"WKLD"—Workload graph data

"DASM"—DASM graph data

"CICS"—CICS graph data

"CICE"—CICS Extension graph data

"IDMS"—IDMS graph data

"IDME"—IDMS Extension graph data

"IMSG"—IMS graph data

"NETG"—Network graph data

"NETE"—Network Extension graph data

"NETL"—Network Physical Lines graph data

"NETN"—Network Physical NCP graph data

"DB2S"—DB2 Subsystem graph data

"DB2U"—DB2 Connect ID graph data

"M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

The PC Floppy Distribution Disk is divided into a number of directories, each of which contains its own sub-directories. The tree structure resembles the following diagram:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M 204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTH- LY	MONTH- LY	MONTH- LY	MONTHLY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB members from the Host processor. The naming conventions are as follows:

- The .BAT Filenames all begin with the three character "HGD" followed by a 1 character graph period indicator (1=Hourly, 2=Daily, 3=Weekly, 4=Monthly) followed by a subsystem identifier (IDMS=IDMS, IDME=IDMS Extension). So, for example, the .BAT file to be used to download the HGDLIB members that would be used to produce the IDMS Daily graphs would be named HGD2IDMS.BAT.
- The Harvard Graphics Template (.TPL files) Filenames all begin with the three characters "HGD" followed by a 1 character graph period indicator (1=Hourly, 2=Daily, 3=Weekly, 4=Monthly) followed by a 2 character graph identifier, followed by a 1 character qualifier in some cases. Please refer to [Graph Descriptions on page 4-52](#) for Template names.

Graph Periods

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*

, for a description of the hourly, daily, weekly and monthly graph periods.

Graph Descriptions

Please reference the *CIMS Capacity Planner Reference Guide* for a description of each of the graphs produced by the CIMS Capacity Planner IDMS Subsystem.

GRAPHIC REPORT TABLES

IDMS Subsystem Hourly Graphs

These graphs represent specific IDMS Regions executing within the &sid system. Therefore, IDMS Region suffixes are required to uniquely identify a particular IDMS Region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDMG	&sid180*	(none)	IDMS Terminal Response
	&sid181*	(none)	IDMS Transaction Volume
	&sid182*	(none)	IDMS CPU Utilization
	&sid183*	(none)	IDMS I/O Activity
	&sid184*	(none)	IDMS Terminal Activity
	&sid185*	(none)	IDMS Transactions/Terminal
	&sid186*	(none)	IDMS Transactions/Day
	&sid189*	(none)	IDMS Transaction Elapsed

* plus a one character suffix representing the IDMS Region

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Extension Hourly Graphs

These graphs represent ALL IDMS Regions executing within the &sid system. Therefore, no IDMS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDME	&sid1E5P	&sidCPUV	CPU Busy+IDMS Prime Shift
	&sid1E7	&sidCPUV	IDMS Volume by Transaction
	&sid1E8	&sidCPUV	IDMS CPU Volume by Response

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Daily Graphs

These graphs represent specific IDMS Regions executing within the &sid system. Therefore, IDMS Region suffixes are required to uniquely identify a particular IDMS Region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDMG	&sid280*	(none)	IDMS Terminal Response
	&sid281*	(none)	IDMS Transaction Volume
	&sid282*	(none)	IDMS CPU Utilization
	&sid283*	(none)	IDMS I/O Activity
	&sid284*	(none)	IDMS Terminal Activity
	&sid285*	(none)	IDMS Transactions/Terminal
	&sid286*	(none)	IDMS Transactions/Day
	&sid287*	(none)	IDMS Transactions/Shift
	&sid288*	(none)	IDMS Terminals/Shift
	&sid289*	(none)	IDMS Transaction Elapsed

* plus a one character suffix representing the IDMS Region

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Extension Daily Graphs

These graphs represent ALL IDMS Regions executing within the &sid system. Therefore, no IDMS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDME	&sid2E5E	&sidCPUV&sidCPUV	CPU Busy+IDMS Early Shift
	&sid2E5P		CPU Busy+IDMS Prime Shift
	&sid2E5L	&sidCPUV&sidIDMR	CPU Busy+IDMS Late Shift
	&sid2E7	&sidIDMO	IDMS Response by Org.
	&sid2E7	&sidIDMT	IDMS Volume by Transaction
	&sid2E8	&sidIDMR	IDMS Volume by Response

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Weekly Graphs

These graphs represent specific IDMS Regions executing within the &sid system. Therefore, IDMS Region suffixes are required to uniquely identify a particular IDMS Region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDMG	&sid380*	(none)	IDMS Terminal Response
	&sid381*	(none)	IDMS Transaction Volume
	&sid382*	(none)	IDMS CPU Utilization
	&sid383*	(none)	IDMS I/O Activity
	&sid384*	(none)	IDMS Terminal Activity
	&sid385*	(none)	IDMS Transactions/Terminal
	&sid386*	(none)	IDMS Transactions/Day
	&sid387*	(none)	IDMS Transactions/Shift
	&sid388*	(none)	IDMS Terminals/Shift
	&sid389*	(none)	IDMS Transaction Elapsed

* plus a one character suffix representing the IDMS Region

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Extension Weekly Graphs

These graphs represent ALL IDMS Regions executing within the &sid system. Therefore, no IDMS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDME	&sid3E5E	&sidCPUV	CPU Busy+IDMS Early Shift
	&sid3E5P	&sidCPUV	CPU Busy+IDMS Prime Shift
	&sid3E5L	&sidCPUV	CPU Busy+IDMS Late Shift
	&sid3E7	&sidIDMR &sidIDMO	IDMS Response by Org.
	&sid3E7	&sidIDMT	IDMS Volume by Transaction
	&sid3E8	&sidIDMR	IDMS CPU Volume by Response

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Monthly Graphs

These graphs represent specific IDMS Regions executing within the &sid system. Therefore, IDMS Region suffixes are required to uniquely identify a particular IDMS Region.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDMG	&sid480*	(none)	IDMS Terminal Response
	&sid481*	(none)	IDMS Transaction Volume
	&sid482*	(none)	IDMS CPU Utilization
	&sid483*	(none)	IDMS I/O Activity
	&sid484*	(none)	IDMS Terminal Activity
	&sid485*	(none)	IDMS Transactions/Terminal
	&sid486*	(none)	IDMS Transactions/Day
	&sid487*	(none)	IDMS Transactions/Shift
	&sid488*	(none)	IDMS Terminals/Shift
	&sid489*	(none)	IDMS Transaction Elapsed

* plus a one character suffix representing the IDMS region

&sid is the SMF System ID (e.g, CPPR)

IDMS Subsystem Extension Monthly Graphs

These graphs represent ALL IDMS Regions executing within the &sid system. Therefore, no IDMS Region suffixes are required.

CPPR Program Name	HGDLIB Member	CPPR PARMLIB Member	Report Description
SSA1IDME	&sid4E5E	&sidCPUV	CPU Busy+IDMS Early Shift
	&sid4E5P	&sidCPUV	CPU Busy+IDMS Prime Shift
	&sid4E5L	&sidCPUV	CPU Busy+IDMS Late Shift
	&sid4E7	&sidIDMR	IDMS Response by Org.
		&sidIDMO	
	&sid4E7	&sidIDMT	IDMS Volume by Transaction
	&sid4E8	&sidIDMR	IDMS CPU Volume by Response

&sid is the SMF System ID (e.g, CPPR)

Network Subsystem

Subsystem Overview	5-3
Network Workload and Performance Reports	5-3
Network Summary Analysis Report	5-3
Network Response Performance Report	5-4
Network Terminal Statistics Report	5-5
Network Exception Analysis	5-6
Network Physical Exception Analysis	5-6
Network Ad Hoc Report	5-6
Network Physical Line Statistics Report	5-7
Network Physical Line Activity Report	5-8
Cluster Controller Statistics Report	5-9
Cluster Controller Activity Report	5-10
Network Graphs	5-11
Network Transaction Response Graph	5-11
Network Terminal Activity Graph	5-12
Network Transaction Activity Graph	5-12
Network Inbound Message Activity Graph	5-12
Network Inbound Message Traffic Graph	5-13
Network Outbound Message Traffic Graph	5-13
Network Total Traffic Graph	5-13
Network Outbound Message Activity Graph	5-14
Network Physical Line Capacity Graph	5-14
Network Physical Line Inbound Traffic Graph	5-15
Network Physical Line Outbound Traffic Graph	5-15
Network Physical Line Retransmitted Traffic Graph	5-16
Network Physical Line Total Traffic Graph	5-17
Network Processor Activity Graph	5-17
Network Processor Slowdown Count Graph	5-18
NCP Free Buffer High Watermark Graph	5-18

NCP Free Buffer Low Watermark Graph	5-19
NCP Channel Hold Queue Length Graph	5-19
Cluster Controller Inbound PIU Graph	5-19
Network Cluster Controller Outbound PIU Graph	5-20
Cluster Controller Inbound Traffic Graph	5-20
Cluster Controller Outbound Traffic Graph	5-21
Cluster Controller Retransmitted Traffic Graph	5-22
Cluster Controller Total Traffic Graph	5-22
Network Terminal Profile	5-23
Network Trends Analysis Reports	5-24
Operating JCL	5-24
Network Subsystem Installation JCL	5-25
Data Reduction JCL	5-25
Report Production JCL	5-40
Additional Report Options	5-60
NETWORK ISPF/PDF INTERFACE	5-62
On-line Graphs	5-63
On-line Reports	5-63
Batch GDDM Graphs	5-66
Presentation Graphics	5-71
Graph Production Procedures	5-71
Naming Conventions	5-72
Presentation Graphics Distribution Disk	5-74
Graph Periods	5-75
Graph Descriptions	5-75
Graphic Report Tables	5-75

Subsystem Overview

The CIMS Capacity Planner Network Subsystem provides a variety of reports and graphs related to the workload and the performance of your network. The Network reports provide both detail and summary level information designed to meet the needs of both the technical staff and technical management. A wide variety of graphs are provided through batch report generation, GDDM, and the PC based Harvard Graphics system. GDDM provides the capability to instantly display a number of Network graphs on the Host through the use of the CIMS Capacity Planner ISPF interface. The Harvard Graphics system is used primarily to produce professional quality colored graphs that are suitable for presentations to users and management.

Network Workload and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner Network reports and describes the contents of each report.

Network Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the Network workload and how well the Network is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the Network.

- The number of active terminals

The peak number of terminals concurrently active during the prime shift over the measured period.

- The average number of transactions processed

The average number of terminal transactions processed during prime shift. If the measured period consists of only one day, the average value is equal to the day's activity.

- Response time

The average terminal response time for transactions submitted during the prime shift over the measured period is reported in seconds (sec.hh).

- Concurrent terminal activity

The average number of terminals active concurrently throughout the prime shift. The value is calculated by summing the average number of terminals active during each 15 minute period throughout the prime shift and dividing the total by the number of 15 minute periods in the prime shift.

- Number of unique terminals accessing the system
- The average number of transactions processed per second during the prime shift
- The average number of outbound messages transmitted per second during the prime shift
- The average length of the inbound messages received during the prime shift
- The average length of the outbound messages transmitted during the prime shift
- The average ratio of transactions during the early shift compared to the prime shift
- The average ratio of transactions during the late shift compared to the prime shift
- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift
- The number of network transactions processed by shift for the measured period.

The 25 most active terminals (inbound) are listed individually by LUNAME. All remaining terminals are reported collectively under the title of "OTHER". The total number of transactions is reported by shift.

- The number of outbound network messages by shift for the measured period.

The 25 most active terminals (outbound) are listed individually. All remaining terminals are reported collectively under the title of "OTHER". The total number of outbound messages is reported by shift.

Network Response Performance Report

This report shows, by shift, for a single day or a range of days, how the NETWORK transactions flowed through the system for a specified VTAM APPLID. Each of the 25 most active terminals is listed individually, by shift, with an indication of what percentage of the response times fell within certain predefined ranges. Terminals that are not among the 25 most active are grouped together and reported under "OTHER". A Totals line is accumulated and written to indicate how all transactions performed as a group. The APPLID is specified through the VTAMNAME parameter.

This is a key report in measuring how well your performance goals are being met within the CPU by the transaction processing application programs and the telecommunications facilities.

The time period column headings at the top of the report are specified through the use of the &sidNETR member of the &PREFIX.CPPR.Vnnn.PARMLIB. Please refer to the *CIMS Capacity Planner Installation Guide* for additional information on setting up the Network Parmlib parameters.

Network Terminal Statistics Report

This report shows, for a specified Network Region (APPLID), which terminals were among the top 40 in the following categories:

- Most active time

The terminal ID, the total active time, and the percent of the measured time period that the terminal was active is listed for the 40 most active terminals. The listing is ordered in descending sequence by the time active. A terminal is considered to be active within each 15 minute period throughout the day only if it transmits a transaction. Terminals that are signed-on, but not actively submitting transactions are not considered to be active.

- Most transaction activity

The terminal ID, the transaction count, and the percentage of all transaction activity is listed for the 40 terminals exhibiting the highest transaction traffic. The listing is ordered in descending sequence by the amount of activity.

- Most outbound message activity

The terminal ID, the transaction count, and the percentage of all transaction activity is listed for the 40 terminals receiving the most outbound message traffic. The listing is ordered in descending sequence by the number of outbound messages.

- Average host response time

The terminal ID and the average host CPU response time for each of 40 terminals experiencing the highest host response times are listed. The terminals are listed in descending order by host response time.

- Average network response time

The terminal ID and the average network response time for each of 40 terminals experiencing the highest network response times are listed. The terminals are listed in descending order by network response time.

- Average total response time

The terminal ID and the average total response time (network + host) for each of 40 terminals experiencing the highest total response times are listed. The terminals are listed in descending order by total response time.

- Average inbound message length

The terminal ID and the average inbound message length are listed for the 40 terminals with the largest average inbound message length. The terminals are listed in descending order by average inbound message length.

- Average outbound message length

The terminal ID and the average outbound message length are listed for the 40 terminals with the largest average outbound message length. The terminals are listed in descending order by average outbound message length.

- Total message traffic

The terminal ID and the total number of Kbytes transmitted and received are listed for the 40 terminals with the most traffic. The terminals are listed in descending order by the number of Kbytes transmitted and received.

The APPLID is specified through the VTAMNAME parameter.

Network Exception Analysis

The Network Exception Analysis Report shows, for each fifteen minute period, for a specified VTAM APPLID, all terminals whose response times exceeded the pre-defined thresholds specified in the CIMS Capacity Planner Parmlib member &sidXCPT. The Parmlib member &sidXCPT contains a number of threshold values for the installation. This member is used to identify the performance thresholds for the various components of the overall environment.

The APPLID is specified through the VTAMNAME parameter.

Network Physical Exception Analysis

The Network Physical Exception Analysis Report shows, for each fifteen minute period, each line that exceeds a specified percentage of its capacity and each NCP that exceeds a pre-defined percentage busy limit. The pre-defined thresholds are specified in the CIMS Capacity Planner Parmlib member &sidXCPT.

Selected lines can be excluded from the report through the use of the EXCLUDE facility. If only certain specified lines are to be considered, they are specified by the INCLUDE facility, in which case only the lines specified through the INCLUDE data set are included in the report.

Network Ad Hoc Report

The CIMS Capacity Planner Network Subsystem supports the production of an Ad Hoc report from the original data before portions of the detail data are summarized or discarded. The Ad Hoc report is produced by the Data Reduction program - SSA1NETW.

To specify that the report is to be produced, you merely specify the report parameters in the Data Reduction Job stream as part of the SYSIN data set. The parameters that apply to the Ad Hoc report are:

VTAM TERMINAL NAME=

The Network Ad Hoc report contains the following information for each record selected for inclusion in the report.

- The VTAM APPLID
- The begin time for the period for which the record was produced.
- The begin date for the period for which the record was produced.

- The end time for the period for which the record was produced.
- The VTAM terminal ID
- The response time (RTM records)
- The inbound message count
- The outbound message count
- The inbound byte count
- The outbound byte count

In addition to selecting records for inclusion into the Ad Hoc report by specific terminal names, groups of terminals are selected through the use of the wildcard character "*".

For example, all terminals beginning with "L422" are selected by specifying VTAM TERMINAL NAME=L422*. This convention can be extended to include all terminals by specifying VTAM TERMINAL NAME=*.

The selection of records for inclusion into the Ad Hoc report can be further restricted by specifying date and/or time ranges by using the BEGIN DATE=, END DATE=, BEGIN TIME=, and END TIME= parameters.

The Ad Hoc report can be limited to a single APPLID if the SYSIN data set for the data reduction program contains a VTAMNAME parameter.

Network Physical Line Statistics Report

This report is available if you are running a monitoring facility like, Netview NPM, that gathers information regarding the physical configuration of the network and passes it on to CIMS Capacity Planner through the SSA1NETW program. Data produced by the NETSPY and NET/MASTER physical monitoring facilities are also supported by the CIMS Capacity Planner Network Subsystem.

The Physical Line Statistics Report shows the busiest 40 lines or a specified line or group of lines (selected through the use of the INCLUDE or EXCLUDE DD statements) in terms of:

- Percent utilization

The average percent of line utilization throughout the prime shift. The lines are listed in descending order by the utilization percentage.

- Baud rate

The stated baud rate for the top 40 lines are listed.

- Total traffic

The total number of Kbytes of data transmitted and the corresponding percentage of the total data communications load are listed for the 40 busiest lines. The list is presented in descending order by the number of Kbytes of data transmitted.

■ Inbound traffic (KB)

The total number of Kbytes of inbound data received is listed for the top 40 lines along with the corresponding percentage of the total inbound traffic. The list is produced in descending order by the number of Kbytes received.

■ Outbound traffic (KB)

The total number of Kbytes of outbound data transmitted is listed for the 40 busiest lines along with their corresponding percentages of the total outbound load.

■ Retransmissions

The number of bytes retransmitted for the 40 lines experiencing the most communications errors, along with their corresponding percentage of the retransmissions.

Network Physical Line Activity Report

The Physical Line Activity Report summarizes several key elements of information related to the communications workload. It summarizes the following information related to each line in the communications configuration:

■ Line ID

■ Average percent busy

The average percent busy is calculated for each line by summing the total percent busy figures for each 15 minute period throughout the day and dividing by 96 to get the average.

■ High watermark percent busy

The high watermark is highest percent busy figure for a 15 minute period during the measured period.

■ Percent of traffic

The percent of the total communications traffic is reported for each communication line in the configuration.

■ Line speed

The baud rate is reported for each line in the communications configuration.

■ Inbound traffic volume

The number bytes of input traffic during the measured period is reported in Kbytes.

■ Outbound traffic volume

The number bytes of outbound traffic during the measured period is reported in Kbytes.

- Retransmissions

The amount of data retransmitted during the measured period in Kbytes.

- Active time (Hours)

The amount of time that each line is active is reported in hours and hundredths (HH.hh).

Cluster Controller Statistics Report

The Cluster Controller Statistics Report provides a summary of the communications workload at the cluster terminal controller level for the 40 cluster controllers experiencing the most traffic. The top 40 cluster controllers or groups of controllers (as defined by INCLUDE/EXCLUDE DD statements) are identified as they relate to the following areas:

- Total traffic (Kbytes)

The 40 cluster controllers with the highest total traffic loads are listed along with the number of Kbytes of data transmitted and received and their relative percentages of the total cluster controller traffic. They are listed in descending sequence by the number of bytes transmitted and received.

- Inbound Traffic (Kbytes)

The 40 busiest cluster controllers are listed along with their total traffic in Kbytes and their relative percentage of all inbound cluster controller traffic. The listing is produced in descending sequence by the amount of inbound traffic.

- Outbound Traffic (Kbytes)

The 40 busiest cluster controllers are listed along with their total traffic in Kbytes and their relative percentage of all outbound cluster controller traffic. The listing is produced in descending sequence by the amount of outbound traffic.

- Retransmissions

The 40 cluster control units requiring the most retransmissions are listed along with their relative percentages of the total retransmissions to all the cluster control units. The listing is ordered in descending sequence by the number of retransmissions required.

Cluster Controller Activity Report

This report provides a summary of the workload supported by each cluster controller in the communications configuration. Selected cluster controllers are eliminated from the report through the use of the EXCLUDE facility. The report is limited to specific cluster controllers by specifying those controllers to be included in the report using the INCLUDE facility. The report provides the following information for each cluster controller:

■ **Control Unit ID**

■ **Inbound physical information units (PIUs)**

The number of inbound PIUs received by the host from each cluster control unit.

■ **Outbound PIUs**

The number of outbound PIUs sent from the host to each cluster control unit.

■ **Percentage of PIUs**

The percentage of the PIUs that each cluster controller contributed to the sum of the PIUs serviced by all the cluster controllers.

■ **Active time**

The amount of time that each cluster controller was active during the measured period is reported in hours and hundredths (HH.hh).

■ **Inbound traffic**

The amount of data transmitted from each cluster controller to the host in Kbytes.

■ **Outbound traffic**

The amount of data transmitted from the host to each cluster controller in Kbytes.

■ **Retransmissions**

The amount of data, in Kbytes, that were retransmitted from the host to each cluster controller.

Network Graphs

The Network Subsystem produces a number of graphs on the mainframe that depict the workload and performance of the network as it relates to specified VTAM APPLIDs or the network as a whole. The graphs are produced by day of the week (Monday, Tuesday, etc.). Whenever the measured period (the period included within the BEGIN DATE and the BEGIN TIME through the END DATE and END TIME) exceeds seven days, at least one graph will contain data pertaining to more than one day. For instance, if the measured period begins on a Monday and runs for 10 days, the resultant graphs for Monday, Tuesday, and Wednesday would contain data pertaining to two days activity, while the remaining four graphs would contain data for a single day's activity.

The majority of the Network graphs, designed to depict the communications workload, the percentage of key resources consumed, or the performance of the network share a common format. Each graph is broken into four sections, each section representing a six hour time period broken into 24 periods of 15 minutes each. The time of day is represented by the vertical axis while the quantities to be graphed are represented along the horizontal axis. The magnitude of the values represented along the horizontal axis are represented by a series of dashes (-) and asterisks (*) in bar graph format. The dashes are used to represent the peak daily average values for the quantities being graphed while the asterisks represent the average values. The peak average values differ from the average values only when multiple days data are combined into a single graph (measured period greater than seven days). Whenever no dashes appear, they have been overlaid by asterisks indicating that the peak and average values are either identical or nearly so. The average and peak values of the quantities being graphed are printed in the middle of each line of the bar graph. The scales used to represent the values along the horizontal axis, except for percentages, are determined dynamically based upon the maximum values of the data being graphed.

Network Transaction Response Graph

The Transaction Response Graph shows, for a specified APPLID, by day of the week, the average and peak daily average network transaction response times for each 15-minute period throughout the day. The terminals for which the response times are graphed are controlled through the use of the INCLUDE/EXCLUDE facility. Response time data for selected terminals are eliminated through the use of the EXCLUDE facility while the terminals to be included are specified through the INCLUDE facility. If no INCLUDE or EXCLUDE list is provided, all terminal data related to the specified APPLID are included in calculating the values shown on the graph.

The APPLID is specified using the VTAMNAME SYSIN parameter.

The Network Transaction Response Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Terminal Activity Graph

The Network Terminal Activity Graph shows, by day of the week, for a specified VTAM APPLID, the average and peak daily average number of active terminals for each 15-minute period throughout the day. Terminals must submit at least one transaction within each 15 minute period to be considered active. The VTAM APPLID is specified through the use of the VTAMNAME parameter.

The terminal data included in calculating the graph values is limited through the use of the EXCLUDE facility or fully defined through the INCLUDE facility.

The Network Terminal Activity Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Transaction Activity Graph

The Network Transaction Activity Graph shows, by day of the week, for a specified VTAM APPLID, the average and peak daily average number of transactions per minute submitted from all the terminals associated with the specified VTAM APPLID (or group of APPLIDs specified using the INCLUDE/EXCLUDE DD statements) during each 15-minute period throughout the day. The VTAM APPLID is specified through the use of the VTAMNAME parameter.

The Network Transaction Activity Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Inbound Message Activity Graph

The Inbound Message Activity Graph shows, for a specified VTAM APPLID, by day of the week, the average number of inbound messages/minute transmitted during each 15-minute period throughout the day. Both the average and peak daily average values are reported.

The VTAM APPLID is specified using the VTAMNAME parameter.

The Network Inbound Activity Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters). The data to be included in calculating the values shown in the graphs is limited by specifying a list of terminals to be excluded from the calculations through the use of the EXCLUDE facility or specifically defining the terminals to be included through the use of the INCLUDE facility.

Network Inbound Message Traffic Graph

The Network Inbound Traffic Graph shows, by day of the week, for a specified VTAM APPLID, the average and peak daily average number of bytes per second transmitted from the host to the terminals associated with the specified VTAM APPLID during each 15-minute period throughout the day. The VTAM APPLID is specified through the use of the VTAMNAME parameter.

The data used in calculating the graph values is limited by excluding selected terminals or groups of terminals through the use of the EXCLUDE facility or fully defined through the use of the INCLUDE facility.

The Network Inbound Message Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Outbound Message Traffic Graph

The Network Outbound Message Traffic Graph shows, for a specified VTAM APPLID, by day of the week, the average number of outbound bytes/second transmitted during each 15 minute period throughout the day. Both the average and the peak daily average values are reported. The VTAM APPLID is specified using the VTAMNAME parameter. The data included in the calculations is limited by excluding selected terminals or groups of terminals through the use of the EXCLUDE facility or fully specified through the use of the INCLUDE facility.

The Network Outbound Message Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Total Traffic Graph

The Network Total Traffic Graph shows, by day of the week, for a specified VTAM APPLID, the average and peak daily average number of bytes per second transferred to and from all the terminals associated with the specified VTAM APPLID during each 15-minute period throughout the day. The VTAM APPLID is specified using the VTAMNAME parameter. The data included in the calculations is limited by excluding selected terminals or groups of terminals through the use of the EXCLUDE facility or fully specifying the terminal data to be included through the use of the INCLUDE facility.

The Network Total Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Outbound Message Activity Graph

The Outbound Message Activity Graph shows, for a specified VTAM APPLID, by day of the week, the average number of outbound messages per minute transmitted during each 15-minute period throughout the day. Both the average and the peak daily average values are reported. The APPLID is specified using the VTAMNAME parameter.

The data included in calculating the graph values are limited by excluding certain terminals or groups of terminals through the use of the EXCLUDE facility or by fully specifying which terminal data is included through the use of the INCLUDE facility.

The Network Outbound Message Activity Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Network Physical Line Capacity Graph

The Network Physical Line Capacity Graph shows, for a specified line or group of lines specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average percentages of the line's (lines') capacity that was utilized. The utilization percentage is shown for each 15 minute period throughout the day.

The data used in calculating the graph values are limited by excluding selected lines through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected lines through the use of the INCLUDE facility. The INCLUDE facility functions differently for the physical line graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the physical line graphs, the INCLUDE causes an individual set of graphs to be produced for each communication line specified in the INCLUDE list.

The Network Physical Line Capacity Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Physical Line Inbound Traffic Graph

The Network Physical Line Inbound Traffic Graph shows, for a specified line or group of lines specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak number of bytes per minute transmitted across the line(s) in transmitting input data to the host. The traffic is shown for each 15 minute period throughout the day.

The data used in calculating the graph values are limited by excluding selected lines through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected lines through the use of the INCLUDE facility. The INCLUDE facility functions differently for the physical line graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the physical line graphs, the INCLUDE causes an individual set of graphs to be produced for each communication line specified in the INCLUDE list.

The Network Physical Line Inbound Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a Network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Physical Line Outbound Traffic Graph

The Network Physical Line Outbound Traffic Graph shows, for a specified line or group of lines specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak number of bytes per minute transmitted across the line(s) in transmitting response data from the host to the terminals. The traffic is shown for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected lines through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected lines through the use of the INCLUDE facility. The INCLUDE facility functions differently for the physical line graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc., specified in the INCLUDE list. In the case of the physical line graphs, the INCLUDE causes an individual set of graphs to be produced for each communication line specified in the INCLUDE list.

The Network Physical Line Outbound Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a Network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Physical Line Retransmitted Traffic Graph

The Network Physical Line Retransmitted Traffic Graph shows, for a specified line or group of lines specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak number of bytes per minute transmitted across the line(s) in retransmitting response data from the host to the terminals. The traffic is shown for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected lines through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected lines through the use of the INCLUDE facility. The INCLUDE facility functions differently for the physical line graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the physical line graphs, the INCLUDE causes an individual set of graphs to be produced for each communication line specified in the INCLUDE list.

The Network Physical Line Retransmitted Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Physical Line Total Traffic Graph

The Network Physical Line Total Traffic Graph shows, for a specified line or group of lines specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of bytes per second transmitted across the line(s) in transmitting inbound and outbound messages to/from the host and retransmitting response data from the host to the terminals. The traffic is shown for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected lines through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected lines through the use of the INCLUDE facility. The INCLUDE facility functions differently for the physical line graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the physical line graphs, the INCLUDE causes an individual set of graphs to be produced for each communication line specified in the INCLUDE list.

The Network Physical Line Total Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Processor Activity Graph

The NCP Activity Graph shows, for a specified NCP or set of NCPs selected through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak percentage capacity used in controlling its associated communications traffic. The utilization percentage is shown for each 15 minute period throughout the day.

Data pertaining to selected NCPs is eliminated from the utilization calculations through the use of the EXCLUDE facility. The data is limited to specific NCPs by specifying the NCPs to be considered in the calculations through the use of the INCLUDE facility. If no NCPs are specified for inclusion or exclusion, then the data pertaining to all NCPs is used in calculating the graph values.

The NCP Activity Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is available only when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Processor Slowdown Count Graph

The NCP Slowdown Count Graph shows, for a specified NCP (selected through the INCLUDE/EXCLUDE facility), by day of the week, the average and peak daily average number of times the NCP entered the slowdown mode due to excessive traffic or insufficient resources. The slowdown counts are shown for each 15 minute period throughout the day.

Data pertaining to selected NCPs is eliminated from the slowdown calculations through the use of the EXCLUDE facility. The data is limited to selected NCPs by specifying the NCPs to be considered in the calculations through the use of the INCLUDE facility. If no NCPs are specified for inclusion or exclusion, then the data pertaining to all NCPs is used in calculating the graph values.

The NCP Slowdown Count Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

NCP Free Buffer High Watermark Graph

The NCP Free Buffer High Watermark Graph shows, for a specified NCP selected through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average values of the maximum number of free buffers available in the NCP during each 15 minute period throughout the day. This is a key measure of whether or not the NCP has sufficient storage to support the communications workload.

The NCP Buffer High Watermark Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

NCP Free Buffer Low Watermark Graph

The NCP Free Buffer Low Watermark Graph shows, for a specified NCP selected through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average values of the minimum number of free buffers available in the NCP for each 15 minute period throughout the day. This is a key measure of whether or not the NCP has sufficient storage to support the communications workload.

The NCP Buffer Low Watermark Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

NCP Channel Hold Queue Length Graph

The NCP Channel Hold Queue Length Graph shows, for a specified NCP selected through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of entries in the channel hold queue for each 15 minute period throughout the day.

The NCP Hold Queue Length Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Cluster Controller Inbound PIU Graph

The Cluster Controller Inbound PIU Graph shows, for a specified cluster controller, by day of the week, the average and peak daily average counts of inbound physical information units that were received from the Network into the host for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc., specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Cluster Controller Inbound PIU Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Cluster Controller Outbound PIU Graph

The Cluster Controller Outbound PIU Graph shows, for a specified cluster controller, by day of the week, the average and peak counts of outbound PIUs that were transmitted from the host to cluster controllers on the network.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most other graphs. The Include facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Network Cluster Controller Outbound PIU Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Cluster Controller Inbound Traffic Graph

The Cluster Controller Inbound Traffic Graph shows, for a specified cluster controller or group of cluster controllers specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of bytes per second transmitted across the line(s) in transmitting input data from the cluster controllers to the host. The values are graphed for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most

other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Cluster Controller Inbound Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Cluster Controller Outbound Traffic Graph

The Cluster Controller Outbound Traffic Graph shows, for a specified cluster controller or group of cluster controllers specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of bytes per second transmitted across the line(s) in transmitting response data from the host to the terminals. The values are graphed for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Cluster Controller Outbound Traffic Graph can be produced for a single day's activity or multiple days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Cluster Controller Retransmitted Traffic Graph

The Cluster Controller Retransmitted Traffic Graph shows, for a specified cluster controller or group of cluster controllers specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of bytes per second transmitted across the line(s) in retransmitting response data from the host to the specified controller(s). The values are graphed for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most other graphs. The Include facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Cluster Controller Retransmitted Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Cluster Controller Total Traffic Graph

The Cluster Controller Total Traffic Graph shows, for a specified cluster controller or group of cluster controllers specified through the INCLUDE/EXCLUDE facility, by day of the week, the average and peak daily average number of bytes per second transmitted across the line(s) in transmitting inbound and outbound messages to/from the host and retransmitting response data from the host to the cluster controllers. The values are graphed for each 15 minute period throughout the day.

The data used in calculating the graph values is limited by excluding selected cluster controllers through the use of the EXCLUDE facility. Separate sets of graphs are produced for selected cluster controllers through the use of the INCLUDE facility. The INCLUDE facility functions differently for the cluster controller graphs than for most other graphs. The INCLUDE facility normally limits the data used for calculating the values displayed on a graph to a specified set of lines, terminals, etc., specified in the INCLUDE list. In the case of the cluster controller graphs, the INCLUDE causes an individual set of graphs to be produced for each cluster controller specified in the INCLUDE list.

The Cluster Controller Total Traffic Graph can be produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

This graph is only available when a network monitoring facility such as Netview NPM, the NETSPY monitor, or the NET/MASTER monitor gathers information related to the physical configuration and utilization of the network and passes it to CIMS Capacity Planner through the SSA1NETW data reduction program.

Network Terminal Profile

The Network Terminals Profile reports key categories of information related to the workload processed by a specified terminal. The categories are:

- The average traffic per day in Kbytes throughout the specified period
- The average number of transactions submitted per day throughout the specified period
- The average number of messages received by the terminal per day during the specified period
- The average host response time for all transactions submitted
- The average network response time for all transactions submitted from the terminal during the specified time period
- The average length of the transactions submitted from the terminal
- The average length of the responses received at the terminal

The Network Terminal Profile is produced for a specific terminal that must be specified by Terminal ID specified using the INCLUDE facility running under a specified APPLID specified through the VTAMNAME parameter.

The Network Terminal Profile is produced in a bar graph format with a scale that runs from 0 to 100 along the vertical axis and the seven categories of data listed above along the horizontal axis. The graphical representation of the of each category of information in the profile originates at the horizontal axis, just above each of the seven category descriptions and extends upward in bar graph format.

The category descriptions consist of an abbreviated title, a scaling factor applied to the reading along the vertical scale, and the value for each category.

Network Trends Analysis Reports

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. The Network elements that are selected for comparison are:

- The number of transactions executed per minute computed as an average during Prime shift
- The average response time during Prime Shift
- The average input transaction size during Prime Shift
- The average output message size during Prime Shift
- The ratio of Early shift to Prime shift transactions
- The ratio of Late shift to Prime shift transactions
- The ratio of Early shift to Prime shift terminals
- The ratio of Late shift to Prime shift terminals

This information is gathered and presented to you in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right hand side of the report
- In data suitable for graphing with the Harvard Graphics Program once it has been down loaded to a Personal Computer. The manner in which the data is down loaded is left to you. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Operating JCL

The Network Subsystem functions in much the same manner as the Workload Analysis Subsystem, but there are a few differences that stem from the nature of Network systems:

- Each Network region is treated as a separate entity.

When you want to examine network performance, you normally look at a specific VTAM APPLID. Thus each of the network regions in the data center must be registered for each of the MVS systems under which it executes.

- Separate ONLINE table files.

You might want to allocate separate ONLINE table files for the Network tables or you might even want to allocate separate ONLINE table files for each separate VTAM APPLID. When processing requests that relate to several network regions or a single network region in concert with the Workload tables, simply concatenate the required table files and build an Index using the CIMS Capacity Planner Utility program—SSA1LOAD.

Network Subsystem Installation JCL

The JCL required to install the CIMS Capacity Planner Network Subsystem is described in the *CIMS Capacity Planner Installation Guide*.

Data Reduction JCL

Network input data comes from from a variety of sources. The Network Subsystem is capable of processing input from the following VTAM monitors:

- NETVIEW

NETVIEW, from IBM, produces Session Monitor and Response Time Monitor (RTM) data in the form of type 39 SMF records that reside in either the SMF MANx clusters or in a historical repository of SMF records. The type 39 records are processed by the SSA1NETW data reduction program.

- NETVIEW PERFORMANCE MONITOR (NPM)

NPM, from IBM, produces Session Monitor and RTM data in the form of type 28 SMF records that reside in the SMF MANx clusters, in a historical repository of SMF records, or in a VSAM cluster in the form of sequential records. The type 28 records are processed by the SSA1NETW data reduction program.

- NETSPY

NETSPY, from Duquesne Systems produces a file of monitor records that are passed as input to the Network Subsystem. Both the Session Monitor and the RTM data produced by NETSPY are processed by the SSA1NETW data reduction program.

- NET/MASTER

NET/MASTER, from Systems Center, produces Session Monitor records in the form of type 39 SMF records and RTM records in the form of type 28 SMF records, both of which reside either in the SMF MANx clusters or in a historical repository of SMF records. NET/MASTER can also produce a user-specified SMF record number.

The data are reduced into a series of Consolidated Tables that are stored on DASD. The data contained in the Consolidated Tables are then used in generating the various reports and graphs.

Regardless of the source of the input data, all data must be processed in chronological order. CIMS Capacity Planner maintains a table showing the date and time of the last record processed. Should a transaction with an earlier date and time be presented to the data reduction program, the record is considered a duplicate and discarded.

DD Statements

The following DD statements are used in Network Data Reduction jobstreams.

STEPLIB

The STEPLIB DD statement describes the data set name and the disposition of the CIMS Capacity Planner load library.

CPPRPARM

The CPPRPARM DD statement contains the data set name and the disposition of the CIMS Capacity Planner parameter library.

CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table.

INDEX

The INDEX DD statement specifies the data set name and the disposition of the data set containing the index to the ONLINE data set. The INDEX data set considerably reduces the amount of time required to run the reports.

ONLINE

The ONLINE DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Consolidated Performance Data Base. When running the Network Ad Hoc report, it is advisable to use a DD DUMMY definition of the ONLINE file to avoid attempting to enter duplicate data into the CIMS Capacity Planner Performance Data Base.

SYSUT1

The SYSUT1 DD statement describes the data set name and the disposition of the data set containing the data to be used as input to the CIMS Capacity Planner Network data reduction program. The program to be executed is specified below as a function of the type and source of the input.

SYSUT3

The SYSUT3 DD statement specifies a temporary data set used to contain a temporary version of the ONLINE data set during the time the Performance Data Base is being updated. The parameters that must be specified are: Unit, Disposition, and Space.

SSASPILL

The SSASPILL DD statment specifies a temporary data set used to hold data related to VTAM APPLIDs that are not being processed during the current pass. It is processed in successive passes, depending upon the selection criteria specified through the parameters input through the SYSIN Data Set. The parameters that must be specified are: Unit, Space, and Disposition.

SYSNAP

The SYSNAP data set is used only in conjunction with the CIMS Capacity Planner parameter DEBUGON. It specifies a file to contain SNAP dumps. It is usually specified as SYSOUT=*

SYSMSG

The SYSMSG DD statement specifies a data set to which diagnostic messages are written.

SYSUDUMP

The SYSUDUMP DD statement describes a standard MVS abend dump print data set. It is usually specified as SYSOUT=*

SYSPRINT

The SYSPRINT DD statement specifies the data set to which the Network Ad Hoc report is directed, It is usually specified as SYSOUT=*, but could be defined as a data set. Its DCB attributes are:

DSORG=PS

RECFM=FBA

LRECL=133

SYSIN

The SYSIN DD statement specifies the data set containing the parameters that control the execution of the data reduction program. The parameters specify the SMF SID of the system for which data is to be reduced, the APPLID/s to be processed, the source of the input, and the various parameters used to limit the data contained in the Ad Hoc Report. A complete list of the SYSIN parameters is specified below under *Data Reduction Input Parameters on page 5-27*. The SYSIN parameters are generally specified in the job stream (SYSIN=*).

Data Reduction Input Parameters

Several parameters are provided to control the operation of the Network Data Reduction programs and the selection criteria for including data into the Network Ad Hoc report.

SELECTED SYSTEM

The SELECTED SYSTEM parameter is used to specify the SMF SID of the system for which network data is to be reduced. This is the SMF System Identifier specified in the SMFPRMxx member of SYS1.PARMLIB.

Specify System ID:

The SELECTED SYSTEM parameter is restricted to a single entry in the Network data reduction programs due to the memory requirements of the software and the tables. The parameter is specified using the following options:

- The SID itself (e.g, SELECTED SYSTEM=MVS1)
- An asterisk "*", indicating that the SMF SID of the MVS system on which the data reduction program is operating is to be used (e.g, SELECTED SYSTEM=*).

Default System ID:

SELECTED SYSTEM=*

VTAMNAME

This parameter is used to specify the Network name of the system for which network data is to be collected. The VTAMNAME is either the VTAM APPLID as specified in the VTAMLST (NETSPY) or the VTAM Resource Name of the VTAM APPLID (NETVIEW and NET/MASTER). If no VTAMNAME is specified, data is collected for all registered VTAM APPLIDs.

Specify VTAMNAME:

VTAMNAME=xxxxxx

Default VTAMNAME:

All registered APPLIDs encountered in the input

SMFILE

The SMFILE parameter specifies the source and form of the input (SYSUT1) to the Network data reduction program. This source must be one of the following:

- NETVIEW type 39 records from the Active SMF Cluster (NVUA)
- NETVIEW type 39 records from the SMF Clusters (NVUL)
- NETVIEW type 39 records from a repository of SMF history (NVUH)
- NETSPY log records from the NETSPY log (NSPY)
- NETSPY log records from the Active SMF Cluster (NSPA)
- NETSPY log records from the SMF Cluster (NSPL)
- NETSPY log records from a repository of SMF history (NSPH)
- NET/MASTER type 39 (or user-specified number) records from the Active SMF Cluster (NMAA)
- NET/MASTER type 39 (or user-specified number) records from the SMF Clusters (NMAL)
- NET/MASTER type 39 (or user-specified number) from a repository of SMF history (NMAH)
- NETVIEW NPM type 28 records from the SMF Active Cluster (NPMA)
- NETVIEW NPM type 28 records from the SMF Clusters (NPML)
- NETVIEW NPM records from a repository of SMF history from the FNMLOGx file (NPMH)
- NETVIEW NPM records extracted from the FNMSESx VSAM cluster (NPMV)

Specify Input Source:

The options available in specifying the SMFILE parameter are:

- SMFILE=NSPY (NETSPY log records)
- SMFILE=NSPA (NETSPY SMF records from the active cluster)
- SMFILE=NSPL (NETSPY SMF records)
- SMFILE=NSPH (NETSPY SMF history records)
- SMFILE=NVUA (NETVIEW SMF records from the active cluster)
- SMFILE=NVUL (NETVIEW SMF records)
- SMFILE=NVUH (NETVIEW SMF history records)
- SMFILE=NMAA (NET/MASTER SMF records from the active cluster)
- SMFILE=NMAL (NET/MASTER live SMF records)
- SMFILE=NMAH (NET/MASTER SMF history records)
- SMFILE=NPMA (NETVIEW NPM SMF records from the active cluster)
- SMFILE=NPML (NETVIEW NPM SMF records)
- SMFILE=NPMH (NETVIEW NPM SMF history records)
- SMFILE=NPMV (NETVIEW NPM SESx records)

Default Input Source:

SMFILE=NSPY

BEGIN DATE

This parameter is used to specify the earliest date that a record should be selected for analysis.

Date Specification:

A date is be specified in any of the following three formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (European Standard Gregorian format)

Default Date Specification:

BEGIN DATE=00001

END DATE

This parameter is used to specify the latest date that a record is selected for analysis.

Date Specification:

See BEGIN DATE above

Default Date Specification:

END DATE=99365

BEGIN TIME

The BEGIN TIME parameter is used to specify the earliest time within the BEGIN DAY that a record is to be selected for analysis.

Time Specification:

The time is specified in hours, minutes, and seconds in the form:

HH.MM.SS

Default Time Specification:

BEGIN TIME=00.00.00

END TIME

The END TIME parameter is used to specify the latest time within the END DATE that a record is to be selected for analysis.

Time Specification:

The time is specified in hours, minutes, and seconds in the form:

HH.MM.SS

Default Time Specification:

BEGIN TIME=24.00.00

REPORT LANGUAGE

This parameter specifies the language that is to be used in preparing the narrative sections of the report.

Specify Report Language:

REPORT LANGUAGE=ENGLISH or

REPORT LANGUAGE=DEUTSCH

Default Language:

REPORT LANGUAGE=ENGLISH

VTAM TERMINAL NAME

This parameter is used to specify a VTAM terminal or group of terminals to be included in the Network Ad Hoc Report. The VTAM TERMINAL NAME parameter in conjunction with the VTAMNAME, BEGIN/END DATE, and the BEGIN/END TIME parameters control the records to be selected for inclusion into the Network Ad Hoc Report.

FILTER

The FILTER parameter is used to specify the user-assigned record number used to designate the NET/MASTER SMF records.

DEBUGON

The DEBUGON parameter has no operands. It is used to activate selected SNAP dumps at specific points in the programs. It should only be used in conjunction with the CIMS Lab technical personnel if problem determination is required.

Processing Network Records

All data reduction processing is performed by one Network Data Reduction program (SSA1NETW). The SSA1NETW program is used to process the data related to the session records (type 39 records) from NETVIEW and NET/MASTER, all NETSPY data, and the type 28 network monitor records produced by NPM and NET/MASTER.

Sample Jobstreams

This section deals with the JCL required to run the Network Data Reduction program. Sample Jobstreams are provided for running the SSA1NETW program. Following is a sample jobstream for processing Network type 39 records residing in a historical SMF data set on tape.

```
//SSADNET JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
/*
//ST1      EXEC PGM=SSA1NETW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1  DD DISP=(OLD,KEEP),DSN=SMF.HIST
//SYSUT3  DD DISP=(,PASS),SPACE=(CYL,(10,8)),UNIT=SYSDA
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8)),UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
```

```

//SYMSGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
DUMP SMF STATISTICS=YES
***** YOU MUST SPECIFY THE VTAM APPLID OR RESOURCE NAME *****
***** IF YOU WANT TO PROCESS A SPECIFIC APPLID *****
***** OTHERWISE, ALL REGISTERED APPLIDS THAT ARE PRESENT *****
***** IN THE INPUT DATA WILL BE PROCESSED. *****
*VTAMNAME=applid
***** NETSPY INPUT *****
*SMFILE=NSPY /* NETSPY FROM THE LOG (THIS IS THE DEFAULT) */
*SMFILE=NSPA /* NETSPY FROM ACTIVE SMF CLUSTER - USE FILTER */
*SMFILE=NSPL /* NETSPY FROM SMF CLUSTER - ALSO USE FILTER */
*SMFILE=NSPH /* NETSPY FROM HISTORY - ALSO USE FILTER */
***** NET/MASTER INPUT *****
*SMFILE=NMAA /* NET/MASTER FROM ACTIVE SMF CLUSTER - USE FILTER */
*SMFILE=NMAL /* NET/MASTER FROM SMF CLUSTER - ALSO USE FILTER */
*SMFILE=NMAH /* NET/MASTER FROM SMF HISTORY - ALSO USE FILTER */
*FILTER=nnn
***** NETVIEW INPUT *****
*SMFILE=NVUA /* NETVIEW FROM ACTIVE SMF CLUSTER */
*SMFILE=NVUL /* NETVIEW FROM SMF CLUSTER */
*SMFILE=NVUH /* NETVIEW FROM SMF HISTORY */
* * * * *
//

```

Model JCL to run SSA1NETW is available in the DNETPROD member of &PREFIX.CPPR.Vnnn.CNTL that you might want to modify for use in your installation.

Processing NETVIEW SMF Records for VTAM APPLIDS

If the NETVIEW program product from IBM is used to monitor network performance, SMF type 39 records are produced that are processed by the CIMS Capacity Planner data reduction program SSA1NETW.

The data reduction program requires that both RTM and Accounting and Availability record segments be produced by NETVIEW. This is accomplished by specifying the following information in the INITMOD statement for NETVIEW:

- AAUINLDM RTM=YES turns on data collection
- AAUINLDM LOG=YES turns on the LOG facility
- AAUINLDM SESSTATS =YES turns on accounting and availability data collection

Sample Jobstream

The following sample jobstream presumes that NETVIEW type 39 records for two VTAM APPLIDS (NETWORKPROD and NETWORKEMAL) are contained in the SMF input file and that a SSASPILL data set is being used to spool the records prior to processing. Note that only a single pass of the SMF file containing the type 39 records is necessary.

```
//SSADNET JOB (...), 'SSA', CLASS=A, MSGCLASS=X
//*JOBPARM S=*
//*
//*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
//*
//ST1      EXEC PGM=SSAINETW, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE  DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1  DD DISP=(OLD, KEEP), DSN=SMF.HIST
//SYSUT3  DD DISP=(, PASS), SPACE=(CYL, (10, 8)),
//          UNIT=SYSDA
//SSASPILL DD DISP=(, PASS), SPACE=(CYL, (30, 8)),
//          UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG  DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=SYS1
SMFILE=NVUH
//
```

If data reduction processing were to be limited to a single VTAM APPLID, like NETWORKPROD, for example, then an additional SYSIN parameter would be added to specify VTAMNAME=NETWORKPROD.

Processing NETSPY Log Records

If Network monitoring records are written to a NETSPY Log File, processing them is straight forward. As in the case of the NETVIEW records above, it is necessary to process a single VTAM APPLID with each pass of the Log File unless a SSASPILL data set is specified. In the example shown below, two VTAM APPLIDs (NETWORKPROD and NETWORKEMAL) have been monitored using NETSPY. Note that with the inclusion of the SSASPILL data set, only one pass of the log is necessary.

Sample Jobstream

```
//SSADNET JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//*
//*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
//*
//ST1      EXEC PGM=SSAINETW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1   DD DISP=(OLD,KEEP),DSN=NETSPY.INPUT.LOG
//SYSUT3   DD DISP=(,PASS),SPACE=(CYL,(10,8)),UNIT=SYSDA
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8)),UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=SYS1
SMFILE=NSPY
//
```

Processing NET/MASTER SMF Records for VTAM APPLIDs

Processing NET/MASTER SMF records is similar to processing NETVIEW records except that NET/MASTER records might have a user-specified record type. In this case, the data reduction program must be informed of the user-specified record type. Record type 232 is used in the example shown below.

Sample Jobstream

```
//SSADNET JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
/*
//ST1      EXEC PGM=SSA1NETW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1  DD DISP=(OLD,KEEP),DSN=SMF.HIST
//SYSUT3  DD DISP=(,PASS),SPACE=(CYL,(10,8)),UNIT=SYSDA
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8)),UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG  DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=SYS1
SMFILE=NMAH
FILTER=232
//
```

Processing NETVIEW NPM Records for VTAM APPLIDs

If the NETVIEW product from IBM is being used in conjunction with NPM to monitor the network performance, NPM Session and RTM records (SMF type 28) are processed by the Network Data Reduction program (SSA1NETW). If the NPM Session and/or RTM records are contained in a FNMSESx data set, they must be extracted and passed to SSA1NETW via SYSUT1.

In order for the NPM records to contain accurate response time information for NETWORK VTAM APPLIDs, it is necessary for you to specify "Definite Response" to NETWORK. This is done in one of the following ways:

- A NETWORK application can post Definite Response itself
- NETWORK can force DEFRESP at the PCT level (MSGINTEG)
- Profiles can specify MSGINTEG to force DEFRESP for all NETWORK.

Definite Response only applies to SNA LUs. Without Definite Response, NPM only tracks transaction transit time into the host (not back to the terminal) and the total transit time will not be accurate.

NPM RTM and Session records might alternatively be written to a FNMSESx data set and later extracted for processing by SSA1NETW. Specify one of the following parameters:

- NPM RTM={SMF|NPMLOG|SESSION} to direct RTM data to a particular file
- NPM SESSION={SMF|NMPLOG|SESSION} to direct session data to a particular file

If RTM and/or Session records are sent to the NPMLOG or to the SESSION file, they are produced in a somewhat different format than the records written to the SMF data set. Therefore, it is necessary to inform SSA1NETW of the record format, which is accomplished through the use of the SMFILE parameter:

- SMFILE=NPMA indicates that NPM records are to be retrieved from the currently active SMF cluster.
- SMFILE=NPML indicates that NPM records are to be retrieved from a live SMF cluster.
- SMFILE=NPML indicate that the NPM records are to be retrieved from a sequential dump of a SMF cluster (or SMF history) or the FNMLOGx file.
- SMFILE=NPMV indicates that the NPM records were written to either the FNMLOGx file or to the NPM Session file and that they have been extracted. They retain the VSAM header tacked on by NPM.

Sample Jobstream

The following sample jobstream presumes that SMF type 28 records are contained in a SMF history data set.

```
//SSADNET JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
/*
//ST1      EXEC PGM=SSA1NETW, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX    DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT  DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1   DD DISP=(OLD,KEEP), DSN=SMF.HIST
//SYSUT3   DD DISP=(,PASS), SPACE=(CYL,(10,8)), UNIT=SYSDA
//SSASPILL DD DISP=(,PASS), SPACE=(CYL,(30,8)), UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=SYS1
SMFILE=NPMH
//
```

Data Reduction for NPM and NET/MASTER Physical Configuration Data

NPM and NET/MASTER, in tracking the communications workload, performance, and utilization of the physical network, such as the network controllers, the lines, and the remote cluster controllers, produces either a type 28 records or, in the case of NET/MASTER, a user-specified SMF record type. These records are processed by the CIMS Capacity Planner data reduction program SSA1NETW to maintain two additional sets of CIMS Capacity Planner tables that are used by the Network Subsystem to produce a variety of reports and graphs that are very useful in measuring the utilization and responsiveness of the physical resources that make up the total network.

NPM Physical Data

NPM Physical Configuration records are written to the SMF data set or to the NPMLOG or SESSION data set. As these records have slightly different formats, it is necessary to inform the data reduction program (SSA1NETW) as to the format of the records that are to be processed. This is by specifying the source of the input data through the use of the SMFILE SYSIN parameter:

- SMFILE=NPMA indicates that the NPM type 28 records are to be retrieved from the currently active SMF cluster.
- SMFILE=MPML indicates that the NPM type 28 records are to be retrieved from one or more specified SMF clusters.
- SMFILE=MPMH indicates that the NPM type 28 records are to be retrieved from a SMF history data set.
- SMFILE=MPMV indicates that the NPM type 28 records were written to the NPMLOG or to the NPM SESSION data set, and that they have been extracted but retain the VTAM header tacked on by NPM.

Sample NPM Jobstream

The following sample jobstream presumes that NPM type 28 records are contained in the SMF input file. Note that no VTAMNAME need be specified, since there is only one physical configuration for any given SMF SID.

```
//SSADNET JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//*
//*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
//*
//ST1      EXEC PGM=SSA1NETW, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE  DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1  DD DISP=(OLD,KEEP), DSN=SMF.HIST
//SYSUT3  DD DISP=(,PASS), SPACE=(CYL,(10,8)), UNIT=SYSDA
//SSASPILL DD DISP=(,PASS), SPACE=(CYL,(30,8)), UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG  DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=SYS1
SMFILE=MPMH
//
```

NET/MASTER Physical Data

NET/MASTER Physical Configuration records are written to the SMF data set or to the NPMLOG or SESSION data set. As these records have slightly different formats, it is necessary to inform the data reduction program (SSA1NETW) as to the format of the records that are to be processed. This is by specifying the source of the input data through the use of the SMFILE SYSIN parameter:

- SMFILE=NMAA indicates that the NET/MASTER type 232 records are to be retrieved from the currently active SMF cluster.
- SMFILE=NMAL indicates that the NET/MASTER type 232 records are to be retrieved from one or more specified SMF clusters.
- SMFILE=NMAH indicates that the NET/MASTER type 232 records are to be retrieved from a SMF history data set.

Sample NET/MASTER Jobstream

The following sample jobstream presumes that NET/MASTER type 232 records are contained in a SMF history input file. Note that no VTAMNAME need be specified, since there is only one physical configuration for any given SMF SID.

```
//SSADNET JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner NETWORK DATA REDUCTION STEP
/*
//ST1      EXEC PGM=SSA1NETW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1  DD DISP=(OLD,KEEP),DSN=SMF.HIST
//SYSUT3  DD DISP=(,PASS),SPACE=(CYL,(10,8)),UNIT=SYSDA
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8)),UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG  DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=SYS1
SMFILE=NMAH
FILTER=232
//
```

Report Production JCL

With the exception of the Ad Hoc Report, all the Network batch reports and graphics are produced by one of the following programs:

- SSA1NETR—Network Report Generation
- SSA1NETT—Network Trends Analysis/Summary
- SSA1NETG—Network Presentation Graphics
- SSA1NETE—Network Subsystem Extension Graphics
- SSA1NETL—Network Line Group Graphics
- SSA1NETN—Network NCP Graphics

DD Statements

The DD Statements required to execute the above programs are described below:

- **STEPLIB**

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

- **CPPRPARM**

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

- **CPPRERT**

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all NETWORK Job streams.

- **INDEX**

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (the Performance Data Base) data set.

- **ONLINE**

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

- **HGDLIB**

The HGDLIB DD statement specifies the Partitioned Data Set used to save the data points required to produce the various graphs using the Harvard Graphics subsystem. The data points contained in the members of the HGDLIB PDS are downloaded to a PC and matched up with a set of CIMS Capacity Planner graph templates that specify how the graphic data is to be presented. The parameters required are the Disposition and the Data Set Name.

■ **SYSPRINT**

The SYSPRINT DD statement specifies the data set that will contain the Network printed reports. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ **SYSNAP**

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ **SYSUDUMP**

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ **EXCLUDE/INCLUDE**

This is the (optional) Element Exclusion file. It is used by the Network Report Program to force certain elements to be excluded from the reports or to force the inclusion of certain elements in the reports that would not normally be included. When using the INCLUDE, only the specified elements are included in the reports.

The EXCLUDE/INCLUDE file must consist of 80 character records, 72 of which are interpreted. Columns 73 through 80 are used for sequence numbers. Multiple elements can be entered in a single record, but they must be separated by commas. The format of the input statement follows:

- An element can be up to eight characters in length.
- Multiple elements can be included in a single record provided that they are separated by commas or blanks.
- Elements can be terminated by the wild card character "*", indicating that all elements that match up to the "wild card" are to be excluded/included.

For example, the following statement in the EXCLUDE file:

```
TSOP,CS*
```

would cause the report program SSA1NETR to exclude any data related to the terminal named TSOP. It would also cause the exclusion of all terminals beginning with the characters "CS"

The use of the INCLUDE causes the report program to limit the contents of the report to only those elements specified in the INCLUDE data set.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Option data to CIMS Capacity Planner.

CIMS Capacity Planner Network Batch Reports

The standard CIMS Capacity Planner Network batch reports are all produced by the CIMS Capacity Planner NETWORK Report Program - SSA1NETR.

Following is an example of the JCL required to run the report program and produce all the reports and graphs supported by the Network batch reporting program:

```
//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1NETR, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
VTAMNAME=applid
***** GENERAL REPORTS *****
NETWORK SUMMARY REPORT=YES
NETWORK EXCEPTION ANALYSIS REPORT=YES
NETWORK PERFORMANCE REPORT=YES
NETWORK TERMINAL STATISTICS REPORT=YES
NETWORK INBOUND ACTIVITY GRAPH=YES
NETWORK OUTBOUND ACTIVITY GRAPH=YES
NETWORK TERMINAL ACTIVITY GRAPH=YES
NETWORK TRANSACTION RESPONSE GRAPH=YES
NETWORK INBOUND TRAFFIC GRAPH=YES
NETWORK OUTBOUND TRAFFIC GRAPH=YES
NETWORK TOTAL TRAFFIC GRAPH=YES
NETWORK TERMINAL PROFILE=YES
```

```
NETWORK TERMINAL NAME=T03SP018
***** CLUSTER CONTROLLER REPORTS *****
CLUSTER CONTROLLER STATISTICS REPORT=YES
CLUSTER CONTROLLER ACTIVITY REPORT=YES
CLUSTER CONTROLLER INBOUND PIU GRAPH=YES
CLUSTER CONTROLLER OUTBOUND PIU GRAPH=YES
CLUSTER CONTROLLER INBOUND TRAFFIC GRAPH=YES
CLUSTER CONTROLLER OUTBOUND TRAFFIC GRAPH=YES
CLUSTER CONTROLLER RETRANSMITTED TRAFFIC GRAPH=YES
CLUSTER CONTROLLER TOTAL TRAFFIC GRAPH=YES
***** PHYSICAL CONFIGURATION REPORTS *****
NETWORK PHYSICAL EXCEPTION ANALYSIS=YES
NETWORK PHYSICAL LINE STATISTICS REPORT=YES
NETWORK PHYSICAL LINE ACTIVITY REPORT=YES
NETWORK PHYSICAL LINE CAPACITY GRAPH=YES
NETWORK PHYSICAL LINE INBOUND TRAFFIC GRAPH=YES
NETWORK PHYSICAL LINE OUTBOUND TRAFFIC GRAPH=YES
NETWORK PHYSICAL LINE RETRANSMITTED TRAFFIC GRAPH=YES
NETWORK PHYSICAL LINE TOTAL TRAFFIC GRAPH=YES
NCP SLOWDOWN TIME COUNT GRAPH=YES
NCP ACTIVITY GRAPH=YES
NCP FREE BUFFER HIGH WATERMARK GRAPH=YES
NCP FREE BUFFER LOW WATERMARK GRAPH=YES
NCP CHANNEL HOLD QUEUE LENGTH GRAPH=YES
//
```

The SYSIN Options used to request reports and to specify reporting options are described below under *Additional Report Options on page 5-60*.

Several members in the &PREFIX.CPPR.Vnnn.CNTL library contain model sets of JCL that can be customized for report generation in your installation. They are:

```
DNETREPT
DNETNPMC
DNETNPML
DNETNPMN
```

Network Trends Analysis Summary Program

Following is an example of the JCL required to run the Network Trends Report Program —SSA1NETT.

```
//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1NETT,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
VTAMNAME=NETWORKPROD
1ST PERIOD BEGIN DATE=02/16/2005
1ST PERIOD END DATE=04/15/2005
2ND PERIOD BEGIN DATE=05/16/2005
2ND PERIOD END DATE=07/15/2005
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
LUNCH BREAK BEGIN HOUR=11
LUNCH BREAK END HOUR=12
VTAM FILTERED VALUES=YES
//*****
//* THE FOLLOWING DDNAME IS USED TO SPECIFY A FILE WHICH MAY
//* OPTIONALLY BE USED TO EXCLUDE CERTAIN UNWANTED TERMINALS
//*****
//EXCLUDE DD *
CS*,T01*
//
```

The SYSIN Options used to request the trends reports and to specify the reporting options are described below under *Network Batch Report Generation Options on page 5-46*.

Member DNETTRPT in the &PREFIX.CPPR.Vnnn.CNTL library contains a set of model JCL that you can customize for use in your installation.

Network Presentation Graphics Program

The CIMS Capacity Planner Network Subsystem provides an extensive number of presentation-quality graphs that depict the Network workload and performance. The graphs are produced for individual Network regions (APPLIDs) operating under a specific SELECTED SYSTEM. The data to be included in the graphs are produced by the CIMS Capacity Planner program (SSA1NETG) and written to selected members of the HGDLIB data set to be down-loaded to a PC and subsequently plotted by the Harvard Graphics presentation graphics system.

Network graphs are also produced for all Network regions that are registered in the Element Registration Table for a specific SELECTED SYSTEM, consolidated together, by executing the CIMS Capacity Planner Network Extension Graphs Program (SSA1NETE).

Network Graph Programs—(SSA1NETW) JCL

Following is an example of the JCL required to operate the SSA1NETW programs to generate the Network graph data points.

```
//SSAGRPH JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1NETW,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.VTAM
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.VTAM
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYMSGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=IP01
VTAMNAME=NETWORKPROD
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of model Job streams for producing the CIMS Capacity Planner Network graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

DNETHGRE—Produce Hourly Graphs

DNETHGRF—Produce Hourly Graphs

DNETHGRL—Produce Hourly Graphs

DNETHGRN—Produce Hourly Graphs

DNETDGRE—Produce Daily Graphs

DNETDGRF—Produce Daily Graphs

DNETWGRE—Produce Weekly Graphs

DNETWGRF—Produce Weekly Graphs

DNETMGRE—Produce Monthly Graphs

DNETMGRF—Produce Monthly Graphs

All ten of the above Job streams contain the JCL required to execute the programs - SSA1NETE, SSA1NETG, SSA1NETL, and SSA1NETN. To tailor the Job streams, merely modify the DD statements and the selection criteria as required.

Network Batch Report Generation Options

SELECTED SYSTEM

The SELECTED SYSTEM= Option is used to specify the SID of the system for which reports are to be produced or trends to be analyzed. This is the SMF identifier that is specified in the SMFPRMxx member of SYS1.PARMLIB.

Options:

The SELECTED SYSTEM Option has the following options:

The SID itself (e.g., SELECTED SYSTEM=MVSA)

An "*", indicating the SID for the system on which the report program is being executed (e.g., SELECTED SYSTEM=*)

Default Option:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= Option is used to specify the beginning of the period for which data in the Performance Data Base is to be collected for analysis and reporting.

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (Relative Day format)

Default Option:

BEGIN DATE=00001

END DATE

The END DATE= Option is used to specify the last date (inclusive) for which Network data is to be included for analysis.

Options:

The format of the Options for the END DATE are identical to the BEGIN DATE Options specified above in *BEGIN DATE on page 5-47*

Default Option:

END DATE=2099365

REPORT LANGUAGE

This Option specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

DEBUGON

This Option is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should only be used when working with the CIMS Lab technical support to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

PRIME SHIFT FIRST HOUR

The PRIME SHIFT FIRST HOUR= Option specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Options:

Specify the beginning of the Prime Shift

Default Option:

PRIME SHIFT FIRST HOUR=8

LATE SHIFT FIRST HOUR

The LATE SHIFT FIRST HOUR= Option specifies the hour of the day on which the LATE Shift begins. It is expressed as an hour using the 24 hour convention.

Options:

Specify the beginning of the Late Shift

Default Option:

LATE SHIFT FIRST HOUR=16

VTAMNAME

The VTAMNAME= Option designates the VTAM region (APPLID) for which data is to collected and reported.

Options:

VTAMNAME=any VTAM Region Name (VTAM APPLID) that has been registered using the SSA1REGN program (see Network Installation Procedures found in the *CIMS Capacity Planner Installation Guide* of this manual).

Default Option:

None

1ST PERIOD BEGIN DATE

This Option is used to specify the beginning data of the first period (for the Trends Analysis Summary Program) or the beginning date for the entire period (for the Trends Analysis Graph Program) for which Consolidated Table Information is to be selected for analysis.

Date Specification:

Specify a date in one of the following three formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (Relative Day format)

Default Date Specification:

1ST PERIOD BEGIN DATE=00001

1ST PERIOD END DATE

This Option is used to specify the ending data of the first period (for the Trends Analysis Summary Program) or the ending date for the entire period (for the Trends Analysis Graph Program) for which Consolidated Table Information is to be selected for analysis.

Date Specification:

Specify a date in one of the following three formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (Relative Day format)

Default Date Specification:

1ST PERIOD END DATE=2099365

2ND PERIOD BEGIN DATE

This Option is used to specify the beginning date of the second period (for the Trends Analysis Summary Program) for which Consolidated Table Information is to be selected for analysis.

Date Specification:

Specify a date in one of the following three formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (Relative Day format)

Default Date Specification:

2ND PERIOD BEGIN DATE=00001

2ND PERIOD END DATE

This Option is used to specify the ending date of the first period (for the Trends Analysis Summary Program) for which Consolidated Table Information is to be selected for analysis.

Date Specification:

Specify a date in one of the following three formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

Default Date Specification:

2ND PERIOD END DATE=2099365

LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter *does not cause* any data to be removed from the Performance Data Base.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

NETWORK SUMMARY REPORT

This Option is used to specify whether or not the Network Summary Report is to be produced.

Options:

NETWORK SUMMARY REPORT=YES (produce the report)

NETWORK SUMMARY REPORT=NO (don't produce the report)

Default Option:

NETWORK SUMMARY REPORT=NO

NETWORK EXCEPTION ANALYSIS

This Option specifies whether or not the Network Exception Analysis Report is to be produced.

Options:

NETWORK EXCEPTION ANALYSIS=YES (produce the report)

NETWORK EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

NETWORK EXCEPTION ANALYSIS=NO

NETWORK TERMINAL STATISTICS REPORT

This Option specifies whether or not the Terminal Statistics Report is to be produced.

Options:

NETWORK TERMINAL STATISTICS REPORT=YES (produce the report)

NETWORK TERMINAL STATISTICS REPORT=NO (don't produce the report)

Default Option:

NETWORK TERMINAL STATISTICS REPORT=NO

NETWORK PERFORMANCE REPORT

This Option specifies whether or not the NETWORK Performance Report is to be produced.

Options:

NETWORK PERFORMANCE REPORT=YES (produce the report)

NETWORKS PERFORMANCE REPORT=NO (don't produce the report)

Default Option:

NETWORK PERFORMANCE REPORT=NO

NETWORK INBOUND ACTIVITY GRAPH

This Option specifies whether or not the Network Inbound Activity Graph is to be produced.

Options:

NETWORK INBOUND ACTIVITY GRAPH=YES (produce the Graph)

NETWORK INBOUND ACTIVITY GRAPH=NO (don't produce the report)

Default Option:

NETWORK INBOUND ACTIVITY GRAPH=NO

NETWORK OUTBOUND ACTIVITY GRAPH

This Option specifies whether not the Network Outbound Activity Graph is to be produced.

Options:

NETWORK OUTBOUND ACTIVITY GRAPH=YES (produce the graph)

NETWORK OUTBOUNDACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

NETWORK OUTBOUND ACTIVITY GRAPH=NO

NETWORK TERMINAL ACTIVITY GRAPH

This Option specifies whether or not the Network Terminal Activity Graph is to be produced.

Options:

NETWORK TERMINAL ACTIVITY GRAPH=YES (produce the graph)

NETWORK TERMINAL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

NETWORK TERMINAL ACTIVITY GRAPH=NO

NETWORK TRANSACTION RESPONSE GRAPH

This Option specifies whether or not the Network Transaction Response Graph is to be produced.

Options:

NETWORK TRANSACTION RESPONSE GRAPH=YES (produce the graph)

NETWORK TRANSACTION RESPONSE GRAPH=NO (don't produce the graph)

Default Option:

NETWORK TRANSACTION RESPONSE GRAPH=NO

NETWORK INBOUND TRAFFIC GRAPH

This Option specifies whether or not the Network Inbound Traffic Graph is to be produced.

Options:

NETWORK INBOUND TRAFFIC GRAPH=YES (produce the graph)

NETWORK INBOUND TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK INBOUND TRAFFIC GRAPH=NO

NETWORK OUTBOUND TRAFFIC GRAPH

This Option specifies whether or not the Network Outbound Traffic Graph is to be produced.

Options:

NETWORK OUTBOUND TRAFFIC GRAPH=YES (produce the graph)

NETWORK OUTBOUND GRAPH=NO (don't produce the graph)

Default Option:

NETWORK OUTBOUND TRAFFIC GRAPH=NO

NETWORK TOTAL TRAFFIC GRAPH

This Option specifies whether or not the Network Total Traffic Graph is to be produced.

Options:

NETWORK TOTAL TRAFFIC GRAPH=YES (produce the graph)

NETWORK TOTAL TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK TOTAL TRAFFIC GRAPH=NO

NETWORK TERMINAL PROFILE

This Option specifies whether or not the Network Terminal Profile Report is to be produced.

Options:

NETWORK TERMINAL PROFILE=YES (produce the graph)

NETWORK TERMINAL PROFILE=NO (don't produce the graph)

Default Option:

NETWORK TERMINAL PROFILE=NO

CLUSTER CONTROLLER INBOUND PIU GRAPH

This Option specifies whether or not the Cluster Controller Inbound PIU graph is to be produced.

Options:

CLUSTER CONTROLLER INBOUND PIU GRAPH=YES (produce the graph)

CLUSTER CONTROLLER INBOUND PIU GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER INBOUND PIU GRAPH=N0

CLUSTER CONTROLLER OUTBOUND PIU GRAPH

This Option specifies whether or not the Cluster Controller Outbound PIU graph is to be produced.

Options:

CLUSTER CONTROLLER OUTBOUND PIU GRAPH=YES (produce the graph)

CLUSTER CONTROLLER OUTBOUND PIU GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER OUTBOUND PIU GRAPH=N0

CLUSTER CONTROLLER INBOUND TRAFFIC GRAPH

This Option specifies whether or not the Cluster Controller Inbound Traffic graph is to be produced.

Options:

CLUSTER CONTROLLER INBOUND TRAFFIC GRAPH=YES (produce the graph)

CLUSTER CONTROLLER INBOUND TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER INBOUND TRAFFIC GRAPH=N0

CLUSTER CONTROLLER OUTBOUND TRAFFIC GRAPH

This Option specifies whether or not the Cluster Controller Outbound Traffic graph is to be produced.

Options:

CLUSTER CONTROLLER OUTBOUND TRAFFIC GRAPH=YES (produce the graph)

CLUSTER CONTROLLER OUTBOUND TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER OUTBOUND TRAFFIC GRAPH=N0

CLUSTER CONTROLLER RETRANSMITTED TRAFFIC GRAPH

This Option specifies whether or not the Cluster Controller Retransmitted Traffic graph is to be produced.

Options:

CLUSTER CONTROLLER RETRANSMITTED TRAFFIC GRAPH=YES (produce the graph)

CLUSTER CONTROLLER RETRANSMITTED TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER RETRANSMITTED TRAFFIC GRAPH=N0

CLUSTER CONTROLLER TOTAL TRAFFIC GRAPH

This Option specifies whether or not the Cluster Controller Total Traffic graph is to be produced.

Options:

CLUSTER CONTROLLER TOTAL TRAFFIC GRAPH=YES (produce the graph)

CLUSTER CONTROLLER TOTAL TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

CLUSTER CONTROLLER TOTAL TRAFFIC GRAPH=N0

CLUSTER CONTROLLER STATISTICS REPORT

This Option specifies whether or not the Cluster Controller Statistics Report is to be produced.

Options:

CLUSTER CONTROLLER STATISTICS REPORT=YES (produce the report)

CLUSTER CONTROLLER STATISTICS REPORT=NO (don't produce the report)

Default Option:

CLUSTER CONTROLLER STATISTICS REPORT=NO

CLUSTER CONTROLLER ACTIVITY REPORT

This Option specifies whether or not the Cluster Controller Activity Report is to be produced.

Options:

CLUSTER CONTROLLER ACTIVITY REPORT=YES (produce the report)

CLUSTER CONTROLLER ACTIVITY REPORT=NO (don't produce the report)

Default Option:

CLUSTER CONTROLLER ACTIVITY REPORT=NO

NETWORK PHYSICAL EXCEPTION ANALYSIS

This Option specifies whether or not the Network Physical Exception Analysis Report is to be produced.

Options:

NETWORK PHYSICAL EXCEPTION ANALYSIS=YES (produce the report)

NETWORK PHYSICAL EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

NETWORK PHYSICAL EXCEPTION ANALYSIS=NO

NETWORK PHYSICAL ACTIVITY REPORT

This Option specifies whether or not the Network Physical Activity Report is to be produced.

Options:

NETWORK PHYSICAL ACTIVITY REPORT=YES (produce the report)

NETWORK PHYSICAL ACTIVITY REPORT=NO (don't produce the report)

Default Option:

NETWORK PHYSICAL ACTIVITY REPORT=NO

NETWORK PHYSICAL STATISTICS REPORT

This Option specifies whether or not the Network Physical Statistics Report is to be produced.

Options:

NETWORK PHYSICAL STATISTICS REPORT=YES (produce the report)

NETWORK PHYSICAL STATISTICS REPORT=NO (don't produce the report)

Default Option:

NETWORK PHYSICAL STATISTICS REPORT=NO

NETWORK PHYSICAL LINE CAPACITY GRAPH

This Option specifies whether or not the Network Physical Line Capacity Graph is to be produced.

Options:

NETWORK PHYSICAL LINE CAPACITY GRAPH=YES (produce the graph)

NETWORK PHYSICAL LINE CAPACITY GRAPH=NO (don't produce the graph)

Default Option:

NETWORK PHYSICAL LINE CAPACITY GRAPH=NO

NETWORK PHYSICAL LINE INBOUND TRAFFIC GRAPH

This Option specifies whether or not the Network Physical Line Inbound Traffic Graph is to be produced.

Options:

NETWORK PHYSICAL LINE INBOUND TRAFFIC GRAPH=YES (produce the graph)

NETWORK PHYSICAL LINE INBOUND TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK PHYSICAL LINE INBOUND TRAFFIC GRAPH=NO

NETWORK PHYSICAL LINE OUTBOUND TRAFFIC GRAPH

This Option specifies whether or not the Network Physical Line Outbound Traffic Graph is to be produced.

Options:

NETWORK PHYSICAL LINE OUTBOUND TRAFFIC GRAPH=YES (produce the graph)

NETWORK PHYSICAL LINE OUTBOUND TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK PHYSICAL LINE OUTBOUND TRAFFIC GRAPH=NO

NETWORK PHYSICAL LINE RETRANSMITTED TRAFFIC GRAPH

This Option specifies whether or not the Network Physical Line Retransmitted Traffic Graph is to be produced.

Options:

NETWORK PHYSICAL LINE RETRANSMITTED TRAFFIC GRAPH=YES (produce the graph)

NETWORK PHYSICAL LINE RETRANSMITTED TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK PHYSICAL LINE RETRANSMITTED TRAFFIC GRAPH=NO

NETWORK PHYSICAL LINE TOTAL TRAFFIC GRAPH

This Option specifies whether or not the Network Physical Line Total Traffic Graph is to be produced.

Options:

NETWORK PHYSICAL LINE TOTAL TRAFFIC GRAPH=YES (produce the graph)

NETWORK PHYSICAL LINE TOTAL TRAFFIC GRAPH=NO (don't produce the graph)

Default Option:

NETWORK PHYSICAL LINE TOTAL TRAFFIC GRAPH=NO

NCP SLOWDOWN TIME COUNT GRAPH

This Option specifies whether or not the NCP Slowdown Time Count Graph is to be produced.

Options:

NCP SLOWDOWN TIME COUNT GRAPH=YES (produce the graph)

NCP SLOWDOWN TIME COUNT GRAPH=NO (don't produce the graph)

Default Option:

NCP SLOWDOWN TIME COUNT GRAPH=NO

NCP ACTIVITY GRAPH

This Option specifies whether or not the NCP Activity Graph is to be produced.

Options:

NCP ACTIVITY GRAPH=YES (produce the graph)

NCP ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

NCP ACTIVITY GRAPH=NO

NCP FREE BUFFER HIGH WATERMARK GRAPH

This Option specifies whether or not the NCP Free Buffer High Watermark Graph is to be produced.

Options:

NCP FREE BUFFER HIGH WATERMARK GRAPH=YES (produce the graph)

NCP FREE BUFFER HIGH WATERMARK GRAPH=NO (don't produce the graph)

Default Option:

NCP FREE BUFFER HIGH WATERMARK GRAPH=NO

NCP FREE BUFFER LOW WATERMARK GRAPH

This Option specifies whether or not the NCP Free Buffer Low Watermark Graph is to be produced.

Options:

NCP FREE BUFFER LOW WATERMARK GRAPH=YES (produce the graph)

NCP FREE BUFFER LOW WATERMARK GRAPH=NO (don't produce the graph)

Default Option:

NCP FREE BUFFER LOW WATERMARK GRAPH=NO

NCP CHANNEL HOLD QUEUE LENGTH GRAPH

This Option specifies whether or not the NCP Channel Hold Queue Length Graph is to be produced.

Options:

NCP CHANNEL HOLD QUEUE LENGTH GRAPH=YES (produce the graph)

NCP CHANNEL HOLD QUEUE LENGTH GRAPH=NO (don't produce the graph)

Default Option:

NCP CHANNEL HOLD QUEUE LENGTH GRAPH=NO

VTAM FILTERED VALUES

This Option is used to specify whether INCLUDES/EXCLUDES are to be applied to the CIMS Capacity Planner table data when compiling the NETWORK Trends reports and graphs.

Options:

VTAM FILTERED VALUES=YES (use INCLUDE/EXCLUDE)

VTAM FILTERED VALUES=NO (don't use INCLUDE/EXCLUDE)

Default Option:

VTAM FILTERED VALUES=NO

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems. A new statement

TOP LABEL=YES

causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time:End as specified
- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

Another statement

FOOT1=UP TO 64 CHAR FOOTNOTE

is used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that are included on all reports where there is room.

FOOT2

A third statement

FOOT2=UP TO 64 CHAR 2ND FOOTNOTE

is used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that is included on all reports where there is room.

Additional Report Options

In addition to the INCLUDE/EXCLUDE options described in *DD Statements on page 5-40*, it is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it is useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report is prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting Options:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****"). Finally, the statement :

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK Option causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK Option does not require the use of the INCLUDE OR EXCLUDE facilities.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a GENERIC ELEMENT MASK=00001100. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ***LL**. If GENERIC ELEMENT MASK=11111111 were to be specified, all records would be summarized into a single entry.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the Network report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK Option differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *
      xxxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "???23?". The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the Network report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK Option differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *  
  
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "???23??. The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

NETWORK ISPF/PDF INTERFACE

The CIMS Capacity Planner NETWORK Subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM graphs, generate the full range of batch reports to be printed on the system printer, and to generate input to the PC Presentation Graphics Subsystem (Harvard Graphics application).

The ISPF/PDF interface is menu driven with separate sets of menus for each CIMS Capacity Planner subsystem (WORKLOAD, NETWORK, etc.).

On-line Graphs

The following NETWORK GDDM graphs are available through the ISPF/PDF Interface:

Transaction Response Graph

This graph depicts the average response time for all transactions in a VTAM region (specified by APPLID), divided into CPU, WAIT, and Suspend time. Within the range of dates specified by you, the graph can show Prime Shift averages by hour, day, week, or month.

Transaction Volume Graph

This graph depicts the average volume for all transactions in a selected VTAM region (specified by APPLID). Within the specified date range, the graph can show Prime Shift averages by hour, day, week, or month.

Terminal Activity Graph

This graph depicts the average number of terminals active in a selected VTAM region (specified by APPLID). Within the specified date range, the graph can show the number of terminals active during the Prime Shift by hour, day, week, or month.

On-line Reports

Terminal Session Statistics Report

The VTAM Terminal Session Statistics Report shows, for a specified VTAM region, which terminals on the network were among the top 40 in the following categories:

- Number of terminal messages processed
- Terminal active time
- Terminal response time
- Number of output messages
- Output traffic in KBytes
- Input traffic in KBytes

Physical Line Statistics Report

The VTAM Physical Statistics Report shows the top 40 lines (within a group of lines specified by the INCLUDE Option) in terms of the following:

- Line utilization
- Baud rates
- Total traffic in KBytes
- Inbound traffic in KBytes
- Outbound traffic in KBytes
- Retransmitted traffic in KBytes

Network Summary Report

The Network Summary Report provides an analytical narrative summarizing the data encountered for the measured period. The following areas are covered:

- Peak Terminal Activity
- Average response times, broken down by Host and Network
- Total transactions processed, including a breakdown by terminal name and shift
- Total outbound messages sent, including a breakdown by terminal name and shift

VTAM Trends Analysis Report

The VTAM Trends Analysis Report shows the trends in VTAM usage in terms of the following:

- The average number of transactions processed
- The average response time during Prime Shift
- The average input message traffic during prime shift (characters per second)
- The average output message traffic during Prime Shift (characters per second)
- The ratio of Early Shift to Prime Shift transactions
- The ratio of Late Shift to Prime Shift transactions
- The ratio of Early Shift to Late Shift terminals
- The ratio of Late Shift to Prime Shift terminals

Network Session Exception Report

The VTAM Session Exception Report shows, for the period specified, any 15-minute intervals during which the XCPT Parmlib Option "Network Terminal Response Time" was exceeded. Each terminal that experienced an average response time in excess of the time specified in the Parmlib member is listed on the report.

Network Physical Line Exception Report

This report identifies any physical communication lines and/or NCPs that exceeded the "Physical Line Percentage Busy" Option specified in the CPPR Option library. The Report shows the Lines and NCPs that exceed the limits within each 15-minute period during which the limits were exceeded.

Network Ad Hoc Reports

Network Ad Hoc reporting is not yet implemented under the ISPF/PDF Interface. For the time being, please use the JCL contained in the DNETADHC member of the CNTL Library.

Network Batch Report Submission

The batch reporting ISPF/PDF interface provides a convenient mechanism to submit the various CIMS Capacity Planner report requests without having to code the required JCL. The Batch interface provides the functions required to:

- Extract data to be down-loaded to the Harvard Graphics applications on the PC
- Submit the Jobs required to run the reports enumerated under *Network Workload and Performance Reports on page 5-3* of this section:
 - NCP Reports
 - Network Summary Report
 - NCP Slowdown Time Count Graphic Report
 - NCP Activity Graphic Report
 - NCP Free Buffer High Watermark Graphic Report
 - NCP Free Buffer Low Watermark Graphic Report
 - NCP Channel Hold Queue Length Graphic Report
 - Physical Line Reports
 - Network Summary Report
 - Line Capacity Graphic Report
 - Line Outbound Traffic Graph
 - Line Total Traffic Graph
 - Line Activity Report
 - Line/NCP Exception Analysis
 - Line Inbound Traffic Graph
 - Line Retransmitted Traffic Graph
 - Line Statistics Report
 - Cluster Controller Reports
 - Network Summary Report
 - Inbound PIU Graphic Report
 - Outbound PIU Graphic Report
 - Inbound Traffic Graph
 - Outbound Traffic Graph
 - Retransmitted Traffic Graph

- Total Traffic Graph
- Statistics Report
- Activity Report
- Network Session Reports
 - Network Summary Report
 - Exception Analysis
 - Performance Report
 - Terminal Statistics Report
 - Inbound Activity Graph
 - Outbound Activity Graph
 - Terminal Activity Graph
 - Transaction Response Graph
 - Inbound Traffic Graph
 - Outbound Traffic Graph
 - Total Traffic Graph
 - Terminal Profile

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch JOBS instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch JOB
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- These options are explained in detail below.

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual Harvard Graphics charts. As originally designed, the data members are created by Batch JOBS, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through Harvard Graphics. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch JOB is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of several key phrases.

```
GDDM BATCH PROCESS=YES
```

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

```
GDDM PRINTER NAME=pppppppp
```

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQ DD Statement in the JCL.

Several members have been added to the Distribution CNTL file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch JOB. The member name "GDDMBATN" shows how to produce any of the Network GDDM graphs in Batch.

GDDMBAT Parmlib Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (IPO1201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST0001) WHICH MAY THEN BE
/* SENT TO A GDDM PRINTER (PRRINT01).
/* * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
```

```

//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
//* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR,DSN=SYS1.GDDMGDF
//* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN  DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00)'
//SYSIN    DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PPRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*          | | | % CHARACTER LENGTH
*          | | % CHARACTER WIDTH
*          | HORIZONTAL OFFSET
*          VERTICAL OFFSET
//INCLUDE DD *
IPO1201
GDDMBATN Parmlib Member
//SSAGDDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
//* * * * *
//* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR ALL OF THE

```

```

/* SUPPORTED NETWORK CHARTS AND SENDING THEM TO A PRINTER.
/* * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * * *
//STO      EXEC PGM=IKJEFT01,REGION=5000K,DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING DD NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD30)' /* NETWORK RESPONSE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD31)' /* TRANSACTION VOLUME /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD32)' /* TERMINAL VOLUME /*
//SYSIN   DD *
VTAMNAME=CICSPROD
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST001
GDDM PRINTER NAME=PPRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*          | | | % CHARACTER LENGTH
*          | | % CHARACTER WIDTH
*          | HORIZONTAL OFFSET
*          VERTICAL OFFSET

```

Specifying GDDM Batch Print Controls

It is possible to specify certain Print Controls that are to take effect when printing a chart in Batch. The key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter allows you to specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

You might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. In order to invoke these local templates in a Batch JOB, a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter is added to the SYSIN, where the "nnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=dddddddd
```

where the "dddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeeee
```

where the "eeeeeeee" is the name under which the graph is saved in the GDF library.

Presentation Graphics

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the Network workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics—release 3.0, plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the Network workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, that describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and can be used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, is used to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner IDMS graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Network graphs specified in the list set forth in *Graph Descriptions on page 5-75* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters:

"&sidpiiq" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly
- "iiq" is the graph identifier that consists of two alpha-numeric characters. The "piiq" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under *Presentation Graphics Distribution Disk on page 5-74*. The qualifier "q" is or is not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpiiq.DAT where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

```
HGDpii.TPL
```

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to download the data points from HGDLIB to the PC using IND\$FILE. The BAT files can be modified as required to fit the naming conventions of your installation.

The names of the BAT files all conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
 - "WKLD"—Workload graph data
 - "DASM"—DASM graph data
 - "NETWORK"—NETWORK graph data
 - "CICE"—NETWORK Extension graph data
 - "IDMS"—IDMS graph data
 - "IDME"—IDMS Extension graph data
 - "IMSG"—IMS graph data
 - "NETG"—Network graph data
 - "NETE"—Network Extension graph data
 - "NETL"—Network Physical Lines graph data
 - "NETN"—Network Physical NCP graph data
 - "DB2S"—DB2 Subsystem graph data
 - DB2U"—DB2 Connect ID graph data
 - M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTH- LY	MONTH- LY	MONTH- LY	MONTHLY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB data point members from the Host processor (assumes the use of IND\$FILE for the file transfer). The naming conventions are as follows:

.BAT File Naming Conventions

The .BAT filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1= Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a subsystem identifier:

- WKLD=Workload
- DASM=DASM
- NETWORK=NETWORK
- CICE=NETWORK Extension
- IDMS=IDMS
- IDME=IDMS Extension
- IMS=IMS

- NETG=Network
- NETE=Network Extension
- NETL=Network Physical Lines
- NETN=Network Physical NCP

For example, the .BAT file used to download the HGDLIB members required to produce the Workload Daily graphs would be named HGD2WKLD.BAT.

.TPL File Naming Conventions

The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1=Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a 2 character graph identifier. In some cases, a 1 character qualifier is added extending the graph identifier to 3 characters. Please refer to the Graphic Report Tables for the specific template names.

Graph Periods

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for a description of the hourly, daily, weekly and monthly graph periods.

Graph Descriptions

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for descriptions of each of the graphs produced by the Network Subsystem.

Graphic Report Tables

The following tables summarize the relationships of members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-based graph program.

Network Subsystem Hourly Graphs

These graphs represent specific VTAM APPLIDs executing within a system specified by a specific SMF SID. Therefore, specific APPLID suffixes are required to uniquely identify a particular VTAM APPLID.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETG	HGD1N0	&sid1N0*	(none)	VTAM Terminal Response
	HGD1N1	&sid1N1*	(none)	VTAM Transaction Activity
	HGD1N2	&sid1N2*	(none)	VTAM Outbound Activity
	HGD1N3	&sid1N3*	(none)	VTAM Inbound Traffic
	HGD1N4	&sid1N4*	(none)	VTAM Outbound Traffic
	HGD1N5	&sid1N5*	(none)	VTAM Terminals Active
	HGD1N6	&sid1N6*	(none)	VTAM Transactions/User
	HGD1N7	&sid1N7*	(none)	VTAM Throughput
	HGD1NA	&sid1NA*	(none)	VTAM Traffic Ratio

* plus a one character suffix representing the VTAM APPLID

** &sid is the SMF System ID (e.g, IPO1)

Network Subsystem Extension Hourly Graphs

These graphs represent ALL VTAM APPLIDs executing within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETE	HGD1NE	&sid1NE	&sidNETL	VTAM Xactns by Line
	HGD1NF	&sid1NF	&sidNETL	VTAM Traffic by Line
	HGD1NP	&sid1NE	&sidNETN	VTAM Activity by APPLID
	HGD1NQ	&sid1NF	&sidNETN	VTAM Traffic by APPLID

** &sid is the SMF System ID (e.g, IPO1)

Network Subsystem Physical Configuration Hourly Graphs

These graphs represent the entire VTAM physical configuration within the system specified by the &sid. Therefore no VTAM APPLIDs are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETN	HGD1NI1	&sid1NI1	&sidNCPL &sidNCPV	1st VTAM NCP Percent Busy
	HGD1NI2	&sid1NI2	&sidNCPL &sid NCPV	2nd VTAM NCP Percent Busy
	HGD1NI3	&sid1NI3	&sidNCPL &sidNCPV	3rd VTAM NCP Percent Busy
	HGD1NI4	&sid1NI4	&sidNCPL &sidNCPV	4th VTAM NCP Percent Busy
	HGD1NI5	&sid1NI5	&sidNCPL &sidNCPV	5th VTAM NCP Percent Busy
	HGD1NI6	&sid1NI5	&sidNCPL	6th VTAM NCP Percent Busy
	HGD1NT	&sid1NT	&sidLGRP	All VTAM Lines % Busy
SSA1NETL	HGD1NT1	&sid1NT1	&sidLGRP	1st VTAM Line Group % Busy
	HGD1NT2	&sid1NT2	&sidLGRP	2rd VTAM Line Group % Busy
	HGD1NT3	&sid1NT3	&sidLGRP	3rd VTAM Line Group % Busy
	HGD1NT4	&sid1NT4	&sidLGRP	4th VTAM Line Group % Busy
	HGD1NT5	&sid1NT5	&sidLGRP	5th VTAM Line Group % Busy
	HGD1NT6	&sid1NT6	&sidLGRP	6th VTAM Line Group % Busy
	HGD1NU	&sid1NU	&sidLGRP	All VTAMLines Traffic
	HGD1NU1	&sid1NU1	&sidLGRP	1st VTAM Line Group Traffic
	HGD1NU2	&sid1NU2	&sidLGRP	2nd VTAM Line Group Traffic
	HGD1NU3	&sid1NU3	&sidLGRP	3rd VTAM Line Group Traffic
	HGD1NU4	&sid1NU4	&sidNETL	4th VTAM Line Group Traffic
	HGD1NU5	&sid1NU5	&sidNETL	5th VTAM Line Group Traffic
	HGD1NU6	&sid1NU6	&sidNETN	6th VTAM Line Group Traffic

** &sid is the SMF System ID (e.g, IPO1)

Network Subsystem Daily Graphs

These graphs represent specific VTAM APPLIDs executing within the &sid system. Therefore, VTAM APPLID suffixes are required to uniquely identify a particular VTAM APPLID.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETG	HGD2N0	&sid2N0*	(none)	VTAM Terminal Response
	HGD2N1	&sid2N1*	(none)	VTAM Transaction Activity
	HGD2N2	&sid2N2*	(none)	VTAM Outbound Activity
	HGD2N3	&sid2N3*	(none)	VTAM Inbound Traffic
	HGD2N4	&sid2N4*	(none)	VTAM Outbound Traffic
	HGD2N5	&sid2N5*	(none)	VTAM Terminals Active
	HGD2N6	&sid2N6*	(none)	VTAM Transactions/User
	HGD2N7	&sid2N7*	(none)	VTAM Throughput
	HGD2N8	&sid2N8*	(none)	VTAM Transactions/Shift
	HGD2N9	&sid2N9*	(none)	VTAM Terminals/Shift
	HGD2NA	&sid2NA*	(none)	VTAM Traffic Ratio

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Extension Daily Graphs

These graphs represent all VTAM APPLIDs executing within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETE	HGD2NE	&sid2NE	&sid NETL	VTAM Transaction by Line
	HGD2NF	&sid2NF	&sidNETL	VTAM Traffic by Line
	HGD2NG	&sid2NG	&sidNETR &sid NETL	VTAM Response by Line
	HGD2NP	&sid2N4*	&sid NETN	VTAM Activity by APPLID
	HGD2NQ	&sid2N5*	&sid NETN	VTAM Traffic by APPLID
	HGD2NR	&sid2N6*	&sid NETR &sid NETN	VTAM Response by APPLID

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Subsystem Physical Configuration Daily Graphs

These graphs represent the entire VTAM physical configuration within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETN	HGD2NI1	&sid2NI1	&sidNCPL &sidNCPV	1st VTAM NCP % Busy
	HGD2NI2	&sid2NI2	&sidNCPL &sidNCPV	2nd VTAM NCP % Busy
	HGD2NI3	&sid2NI3	&sidNCPL &sidNCPV	3rd VTAM NCP % Busy
	HGD2NI4	&sid2NI4	&sidNCPL &sidNCPV	4th VTAM NCP % Busy
	HGD2NI5	&sid2NI5	&sidNCPL &sidNCPV	5th VTAM NCP % Busy
	HGD2NI6	&sid2NI6	&sidNCPL	6th VTAM NCP % Busy
	HGD2NT	&sid2NT	&sidLGRP	All VTAM Lines % Busy
SSA1NETL	HGD2NT1	&sid2NT1	&sidLGRP	1st VTAM Line Group % Busy
	HGD2NT2	&sid2NT2	&sidLGRP	2nd VTAM Line Group % Busy
	HGD2NT3	&sid2NT3	&sidLGRP	3rd VTAM Line Group % Busy
	HGD2NT4	&sid2NT4	&sidLGRP	4th VTAM Line Group % Busy
	HGD2NT5	&sid2NT5	&sidLGRP	5th VTAM Line Group % Busy
	HGD2NT6	&sid2NT6	&sidLGRP	6th VTAM Line Group % Busy
	HGD2NU	&sid2NU	&sidLGRP	All VTAMLines Traffic
	HGD2NU1	&sid2NU1	&sidLGRP	1st VTAM Line Group Traffic
	HGD2NU2	&sid2NU2	&sidLGRP	2nd VTAM Line Group Traffic
	HGD2NU3	&sid2NU3	&sidLGRP	3rd VTAM Line Group Traffic
	HGD2NU4	&sid2NU4	&sidNETL	4th VTAM Line Group Traffic
	HGD2NU5	&sid2NU5	&sidNETL	5th VTAM Line Group Traffic
	HGD2NU6	&sid2NU6	&sidNETN	6th VTAM Line Group Traffic

&sid is the SMF System ID (e.g. IPO1)

Network Subsystem Weekly Graphs

These graphs represent specific VTAM APPLIDs executing within the &sid system. Therefore, VTAM APPLID suffixes are required to uniquely identify a particular VTAM APPLID.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETG	HGD3N0	&sid3N0*	(none)	VTAM Terminal Response
	HGD3N1	&sid3N1*	(none)	VTAM Transaction Activity
	HGD3N2	&sid3N2*	(none)	VTAM Outbound Activity
	HGD3N3	&sid3N3*	(none)	VTAM Inbound Traffic
	HGD3N4	&sid3N4*	(none)	VTAM Outbound Traffic
	HGD3N5	&sid3N5*	(none)	VTAM Terminals Active
	HGD3N6	&sid3N6*	(none)	VTAM Transactions/User
	HGD3N7	&sid3N7*	(none)	VTAM Throughput
	HGD3N8	&sid3N8*	(none)	VTAM Transactions/Shift
	HGD3N9	&sid3N9*	(none)	VTAM Terminals/Shift
	HGD3NA	&sid3NA*	(none)	VTAM Traffic Ratio

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Extension Weekly Graphs

These graphs represent all VTAM APPLIDs executing within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETE	HGD3NE	&sid3NE	&sid NETL	VTAM Transaction by Line
	HGD3NF	&sid3NF	&sidNETL	VTAM Traffic by Line
	HGD3NG	&sid3NG	&sidNETR	VTAM Response by Line
			&sid NETL	
	HGD3NP	&sid3N4*	&sid NETN	VTAM Activity by APPLID
	HGD3NQ	&sid3N5*	&sid NETN	VTAM Traffic by APPLID
	HGD3NR	&sid3N6*	&sid NETR	VTAM Response by APPLID
			&sid NETN	

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Subsystem Physical Configuration Weekly Graphs

These graphs represent the entire VTAM physical configuration within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETN	HGD3NI1	&sid3NI1	&sidNCPL &sidNCPV	1st VTAM NCP % Busy
	HGD3NI2	&sid3NI2	&sidNCPL &sidNCPV	2nd VTAM NCP % Busy
	HGD3NI3	&sid3NI3	&sidNCPL &sidNCPV	3rd VTAM NCP % Busy
	HGD3NI4	&sid3NI4	&sidNCPL &sidNCPV	4th VTAM NCP % Busy
	HGD3NI5	&sid3NI5	&sidNCPL &sidNCPV	5th VTAM NCP % Busy
	HGD3NI6	&sid3NI6	&sidNCPL	6th VTAM NCP % Busy
SSA1NETL	HGD3NT	&sid3NT	&sidLGRP	All VTAM Lines % Busy
	HGD3NT1	&sid3NT1	&sidLGRP	1st VTAM Line Group % Busy
	HGD3NT2	&sid3NT2	&sidLGRP	2nd VTAM Line Group % Busy
	HGD3NT3	&sid3NT3	&sidLGRP	3rd VTAM Line Group % Busy
	HGD3NT4	&sid3NT4	&sidLGRP	4th VTAM Line Group % Busy
	HGD3NT5	&sid3NT5	&sidLGRP	5th VTAM Line Group % Busy
	HGD3NT6	&sid3NT6	&sidLGRP	6th VTAM Line Group % Busy
	HGD3NU	&sid3NU	&sidLGRP	All VTAMLines Traffic
	HGD3NU1	&sid3NU1	&sidLGRP	1st VTAM Line Group Traffic
	HGD3NU2	&sid3NU2	&sidLGRP	2nd VTAM Line Group Traffic
	HGD3NU3	&sid3NU3	&sidLGRP	3rd VTAM Line Group Traffic
	HGD3NU4	&sid3NU4	&sidNETL	4th VTAM Line Group Traffic
HGD3NU5	&sid3NU5	&sidNETL	5th VTAM Line Group Traffic	
HGD3NU6	&sid3NU6	&sidNETN	6th VTAM Line Group Traffic	

&sid is the SMF System ID (e.g. IPO1)

Network Subsystem Monthly Graphs

These graphs represent specific VTAM APPLIDs executing within the &sid system. Therefore, VTAM APPLID suffixes are required to uniquely identify a particular VTAM APPLID.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETG	HGD4N0	&sid4N0*	(none)	VTAM Terminal Response
	HGD4N1	&sid4N1*	(none)	VTAM Transaction Activity
	HGD4N2	&sid4N2*	(none)	VTAM Outbound Activity
	HGD4N3	&sid4N3*	(none)	VTAM Inbound Traffic
	HGD4N4	&sid4N4*	(none)	VTAM Outbound Traffic
	HGD4N5	&sid4N5*	(none)	VTAM Terminals Active
	HGD4N6	&sid4N6*	(none)	VTAM Transactions/User
	HGD4N7	&sid4N7*	(none)	VTAM Throughput
	HGD4N8	&sid4N8*	(none)	VTAM Transactions/Shift
	HGD4N9	&sid4N9*	(none)	VTAM Terminals/Shift
	HGD4NA	&sid4NA*	(none)	VTAM Traffic Ratio

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Extension Monthly Graphs

These graphs represent all VTAM APPLIDs executing within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETE	HGD4NE	&sid4NE	&sid NETL	VTAM Transaction by Line
	HGD4NF	&sid4NF	&sidNETL	VTAM Traffic by Line
	HGD4NP	&sid4NE	&sidNETN	VTAM Activity by APPLID
	HGD4NQ	&sid4NF	&sid NETN	VTAM Traffic by APPLID

* plus a one character suffix representing the VTAM APPLID

&sid is the SMF System ID (e.g, IPO1)

Network Subsystem Physical Configuration Monthly Graphs

These graphs represent the entire VTAM physical configuration within the &sid system. Therefore, no VTAM APPLID suffixes are required.

CPPR PROGRAM NAME	HARVARD GRAPHICS Template Name	HGDLIB Member **	CPPRPARM Library Member	Report Description
SSA1NETN	HGD4NI1	&sid4NI1	&sidNCPL &sidNCPV	1st VTAM NCP % Busy
	HGD4NI2	&sid4NI2	&sidNCPL &sidNCPV	2nd VTAM NCP % Busy
	HGD4NI3	&sid4NI3	&sidNCPL &sidNCPV	3rd VTAM NCP % Busy
	HGD4NI4	&sid4NI4	&sidNCPL &sidNCPV	4th VTAM NCP % Busy
	HGD4NI5	&sid4NI5	&sidNCPL &sidNCPV	5th VTAM NCP % Busy
	HGD4NI6	&sid4NI6	&sidNCPL	6th VTAM NCP % Busy
	HGD4NT	&sid4NT	&sidLGRP	All VTAM Lines % Busy
SSA1NETL	HGD4NT1	&sid4NT1	&sidLGRP	1st VTAM Line Group % Busy
	HGD4NT2	&sid4NT2	&sidLGRP	2rd VTAM Line Group % Busy
	HGD4NT3	&sid4NT3	&sidLGRP	3rd VTAM Line Group % Busy
	HGD4NT4	&sid4NT4	&sidLGRP	4th VTAM Line Group % Busy
	HGD4NT5	&sid4NT5	&sidLGRP	5th VTAM Line Group % Busy
	HGD4NT6	&sid4NT6	&sidLGRP	6th VTAM Line Group % Busy
	HGD4NU	&sid4NU	&sidLGRP	All VTAMLines Traffic
	HGD4NU1	&sid4NU1	&sidLGRP	1st VTAM Line Group Traffic
	HGD4NU2	&sid4NU2	&sidLGRP	2nd VTAM Line Group Traffic
	HGD4NU3	&sid4NU3	&sidLGRP	3rd VTAM Line Group Traffic
	HGD4NU4	&sid4NU4	&sidNETL	4th VTAM Line Group Traffic
	HGD4NU5	&sid4NU5	&sidNETL	5th VTAM Line Group Traffic
	HGD4NU6	&sid4NU6	&sidNETN	6th VTAM Line Group Traffic

&sid is the SMF System ID (e.g. IPO1)

IMS SUBSYSTEM

Subsystem Overview	6-2
IMS Workload and Performance Reports	6-2
IMS Summary Analysis Report	6-2
IMS Response Performance Report	6-3
IMS Transaction Statistics Report	6-3
IMS Terminal Statistics Report	6-4
IMS Exception Analysis	6-4
IMS Graphs	6-5
IMS Transaction Response Graph	6-5
IMS Transaction Activity Graph	6-5
IMS Terminal Activity Graph	6-5
IMS CPU Activity Graph	6-5
IMS I/O Activity Graph	6-6
IMS Trends Analysis Reports	6-6
Operating JCL	6-7
IMS Subsystem Installation JCL	6-7
Data Reduction	6-7
Report Production	6-16
Additional Report Options	6-28
ISPF/PDF IMS Subsystem Interface	6-30
On-line Graphs	6-31
On-line Reports	6-31
Batch GDDM Graphs	6-33
Presentation Graphics	6-37
Graph Production Procedures	6-37
Naming Conventions	6-38
Presentation Graphics Distribution Disk	6-40
Graph Periods	6-41
Graph Descriptions	6-41
Graphic Report Tables	6-41

Subsystem Overview

The CIMS Capacity Planner IMS Subsystem provides a variety of reports and graphs related to the workload and the performance of your IMS system and applications. The IMS reports provide detail and summary level information designed to meet the needs of both the technical staff and technical management. A variety of graphs are provided through both GDDM and the PC based Harvard Graphics system. GDDM provides the capability to instantly display a number of IMS graphs on the Host through the use of the CIMS Capacity Planner ISPF interface. The Harvard graphics system is used primarily to produce professional quality colored graphs that are used for presentations to users and management.

IMS Workload and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner IMS reports and describes the contents of each report.

IMS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the IMS workload and how well the IMS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time
The Summary Report is produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the IMS system.
- The average number of active terminals (high watermark)
A terminal is considered active during each 15-minute period throughout the day only if at least one transaction is submitted. If a terminal is logged on throughout three shifts, but transactions are only submitted during two 15 minute periods, the terminal is considered to be active for only 30 minutes.
- The average number of IMS transactions processed per minute during the prime shift
- The average number of data base calls made per second during the prime shift
- The average number of IMS-related EXCPs per second during the prime shift
- The average transaction response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average percentage of the CPU time consumed by IMS and the IMS applications during each shift within the measured period.

- The average number of transactions processed by shift for the measured period.

The 25 most frequently executed transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions are reported by shift.

IMS Response Performance Report

This report shows, by shift, for a single day or a range of days, how the IMS transactions flowed through the system for a specified IMS region. Each of the 25 most frequently executed IMS transactions are listed individually, by shift, with an indication of what percentage of the response times fell within certain predefined time-frames. An Overall totals line is accumulated and written to indicate how all transactions performed as a group.

This is a key report in measuring how well your performance goals are being met within the CPU by the IMS system and the application programs.

IMS Transaction Statistics Report

This report shows, for a specified IMS region, which IMS transactions were among the top 40 in the following categories:

- Most frequently executed
- The transaction codes and the transaction counts of the 40 most frequently executed transactions are listed along with their respective percentages of the total IMS transaction workload. The transactions are listed in descending order by frequency of execution.

- Most cumulative elapsed time

The transaction codes and the total amount of elapsed (residence) time for each of 40 transactions experiencing the highest residency time are listed along with their respective percentages of the total IMS transaction residence time. The transactions are listed in descending order by residency time.

- Most cumulative CPU time

The transaction codes and the total amount of CPU time consumed for each of the 40 transactions that consumed the most CPU time along with their respective percentages of the total IMS CPU time. The transactions are listed in descending order by CPU time.

- Most File Accesses

The transaction codes and the number of file access calls are listed for each of the 40 transactions issuing the most file access calls along with their respective percentages of all the file access activity. The transactions are listed in descending order by the number of file accesses.

Transactions are excluded from the Transaction Statistics Report through the use of the EXCLUDE facility (see JCL). Similarly, transactions that would not normally be included are included through the use of the INCLUDE facility (JCL).

IMS Terminal Statistics Report

This report shows which terminals, for a specified IMS Region, were among the top 40 in the following categories:

- Most Active

The Terminal ID and the amount of time active (hours:minutes) are reported along with the percentage of the combined active time for all active terminals for the 40 most active terminals. A terminal is considered active within any given 15 minute period only if transactions are received from the terminal. A terminal that is signed-on, but does not submit any transactions is not considered to be active.

- Most Transactions Executed

The Terminal ID and the number of transactions submitted are reported for the 40 terminals submitting the most transactions along with their respective percentages of the total number of transactions submitted during the measured period. The listing is produced in descending order by the number of transactions submitted.

Selected terminals/transactions are excluded from the report through the use of the EXCLUDE facility (see JCL). Similarly, terminals that can not qualify as among the top forty are included in the report through the use of the INCLUDE facility (see JCL).

IMS Exception Analysis

The IMS Exception Analysis Report shows, for each fifteen minute period, for a specified IMS region, all transactions and terminals whose response times exceeded the pre-defined thresholds specified in the CIMS Capacity Planner PARMLIB member &sidXCPT. The PARMLIB member &sidXCPT contains a number of threshold values for the installation. This member is used to identify the performance thresholds for the various components of the overall environment.

IMS Graphs

The IMS Subsystem produces a number of graphs on the mainframe that depict the magnitude of the workload and the level of performance of a specified IMS Region.

IMS Transaction Response Graph

The Transaction Response Graph shows, by day of the week, the average and peak average transaction response times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the response time (both average and peak average). In the middle of each line, both the average and peak average values are printed. The lines, in bar graph format, form the graphic representation of the values. The peak average values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak values and the average values are either equal or very nearly so. The unit of measure represented by the values along the horizontal axis is specified in the third line of the report title.

The Transaction Response Graph is produced either for a single day's activity or several days' activities. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the `BEGIN DATE` and the `END DATE` (see parameters).

IMS Transaction Activity Graph

The Transaction Activity Graph is organized in the same format as the Transaction Response Graph described above. It shows, by day of the week, the average and peak average number of transactions submitted from all the IMS terminals during each 15-minute period throughout the day. The horizontal values represent the number of transactions submitted per minute.

IMS Terminal Activity Graph

The Terminal Activity Graph shows, by day of the week, the number of terminals during each 15-minute period throughout the day. Both the average and the peak average values are reported. The values along the horizontal axis represent the number of active terminals.

IMS CPU Activity Graph

The CPU Activity Graph shows, by day of the week, the percentage of the CPU processing capacity that was expended in processing IMS transactions for a specified IMS region. The average and peak average percentage values are reported in 15 minute intervals throughout the day. As is the case with the other graphs, the measured period is specified to include anywhere from a day to several months.

IMS I/O Activity Graph

This graph shows, for a specified IMS region, the amount of EXCP activity related to processing transactions. Both the average and the peak average I/O rates are reported. The format of the I/O Activity Graph is identical to the graphs described above, except that the values along the horizontal axis represent the number of I/Os per second related to processing the IMS workload.

IMS Trends Analysis Reports

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. The IMS elements that are selected for comparison are:

- CPU Utilization Statistics, including a breakdown in terms of IMS overhead and Message Region usage
- The number of transactions executed per minute computed as an average during Prime shift
- The average number of EXCPs per second during Prime shift
- The average response time during Prime Shift
- The ratio of Early shift to Prime shift transactions
- The ratio of Late shift to Prime shift transactions
- The ratio of Early shift to Prime shift terminals
- The ratio of Late shift to Prime shift terminals

This information is gathered and presented to you in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right hand side of the report.
- In data suitable for graphing with the Harvard Graphics Program once it has been down loaded to a Personal Computer. The manner in that the data is down loaded is left to the user. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Operating JCL

The IMS Subsystem functions in much the same manner as the Workload Analysis Subsystem, but there are a few differences that stem from the nature of IMS systems:

- Each IMS region is treated as a separate entity.

When you want to examine the performance of IMS, you normally look at a specific region executing on a specific MVS system. Thus each of the IMS regions in the data center must be registered for each of the MVS systems under which it can execute.

- Separate ONLINE table files.

You might want to allocate separate ONLINE table files for the IMS tables related to a given MVS system or you might even want to allocate separate ONLINE table files for each separate IMS region.

IMS Subsystem Installation JCL

The JCL required to install the CIMS Capacity Planner IMS Subsystem is described in the *CIMS Desktop Installation and Getting Started Guide*.

Data Reduction

The data reduction process for IMS data differs substantially from the data reduction processes for the other CIMS Capacity Planner subsystems because the data to be included in the Consolidated Tables is taken from both the IMS or IMF Logs and selected Type 30 SMF records from the SMF Log. The CIMS Capacity Planner includes a Log Preprocessor for use during the IMS Data Reduction step. For more information, refer to the CNTL Library member named DIMSPROD.

The IMS Subsystem data reduction program accepts input from the following types of records:

- The IMS Log data set or the IMF Log data set
- SMF Type 30 data that was extracted from the SMF/RMF input to the Workload Analysis Data Reduction program by way of the SYSUT2 data set and the FILTER keyword option.

The IMS data is reduced into a set of Consolidated tables that are stored in the CIMS Capacity Planner On-line Performance Data Base. These Consolidated Tables are later used by the report generation programs to produce the reports and graphs and to generate the data required to support the PC-based systems.

Regardless of the format of the input, all data must be presented to the IMS data reduction program in chronological order. The IMS Subsystem maintains a table showing the date and time of the last 30 transaction logs processed. Should a duplicate log or a transaction be read into the data reduction program that is time stamped earlier than the latest transaction in the file, the transaction file is considered a duplicate of a previously processed file and discarded, unless the checking feature is overridden by the FORCE IMS LOG=YES parameter.

Both the IMS Log and the IMF Log files contain much more data than is required by the data reduction programs. Rather than sorting the entire file, three programs are provided that extract the relevant records required by the data reduction programs, and pass them on to be sorted and processed. Sample JCL is provided to illustrate the use of the extract and data reduction programs. In addition, model JCL DIMSPROD is provided in the &PREFIX.CPPR.Vnnn.CNTL dataset. The following programs perform the data extraction:

- CIMSLP61 and CIMSLP51 performs data extraction from IMS Logs
- IMFCOPY5 performs data extraction from IMF Logs

Processing SMF Type 30 Records

Because IMS operates in several address spaces, the Data Reduction programs must gather resource usage data from the SMF Type 30 records related to all the Address Spaces that are related to IMS. This is accomplished by providing the Data Reduction programs with a list of up to five job names under which the programs related to IMS operate. The data reduction programs process both the step termination and interval accounting records. If interval accounting records are detected, the step termination records are ignored as the interval accounting records provide a much more accurate resource consumption data over time than the step termination records. The recommended time period for generating the interval accounting records is specified in *Chapter 1, Installing CIMS Capacity Planner* found in the *CIMS Desktop Installation and Getting Started Guide*.

Processing IMS Log Records

IMS Logs must be preprocessed by the IMSCOPY program that consolidates data from several types of IMS log records prior to passing the data to the IMS Data Reduction module (SSA1IMSW). The CIMSLP51 and CIMSLP61 programs are used to process IMS 5.1 and IMS 6.1 and later logs respectively. As each log is processed, a record is kept of the last 30 logs processed to avoid processing any duplicate log data sets. The key used to detect duplicate logs consists of the IMS ID and the DATE/TIME contained in the IMS START log record. Only one IMS log is processed during a single execution of the data reduction program. Partial logs are processed if they are a continuation of the last log in the list maintained by the data reduction program or if the log is being forced.

Processing IMF Log Records

The IMF Log written by the BMC IMF Monitor must be preprocessed by the IMFCOPY5 program to produce an IMF Extract File prior to being sorted and processed by the SSA1IMSW. A record of the last 30 IMF Logs processed is maintained to avoid processing duplicate IMF Logs. Only one IMF log is processed during a single execution of the data reduction program.

Operating JCL for Data Reduction

Following are sample job streams to operate both the IMS and IMF data reduction programs.

IMS Data Reduction JCL (DIMSPROD)

Following is a set of JCL that is used to run the IMS Data Reduction program using input from an IMS 6.1 Log and the corresponding SMF data.

```
//SSACPPR JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST01 EXEC PGM=SSAIWKLD, REGION=5000K, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.V530.LOADLIB, DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUT1 DD DISP=SHR, DSN=SMF.INPUT.FILE
/*SYSMANO DD DISP=SHR, DSN=SYS1.MANO
/*SYSMAN1 DD DISP=SHR, DSN=SYS1.MAN1
/*SYSMAN2 DD DISP=SHR, DSN=SYS1.MAN2
/*SYSMAN3 DD DISP=SHR, DSN=SYS1.MAN3
//ONLINE DD DUMMY
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CIMSPASS DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CNTL(CIMSNUM)
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
/*****
/* THE FOLLOWING PASSES SMF TYPE 30 RECORDS TO ST03
/*****
//SYSUT2 DD SPACE=(CYL,(100,50)), UNIT=SYSDA, DISP=(,PASS)
//SYSUT3 DD SPACE=(TRK,(1,1)), UNIT=SYSDA, DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSMSGG DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=INCLUDE(5)
* * * DONT PUT ANYTHING INTO THE ONLINE PERFORMANCE DATABASE
SCANONLY
* * * PASS SMF TYPE 30 RECORDS INTO THE SYSUT2 FILE * * *
SYSUT2=YES
FILTER=30
/*****
/*
/* PREPROCESSES THE IMS LOG RECORDS
/*
/* USE STEPS ST0261A AND ST0261B FOR IMS 6.1
/*
/*****
/* STEPS ST0261A AND ST0261B
/*
/* CIMSLP61 TO PREPROCESS THE IMS LOG FOR BOTH THE CIMS CHARGEBACK
/* AND THE CAPACITY PLANNER SYSTEMS IN A SINGLE PASS OF
/* THE IMS LOG DATASET. IMS RELEASE 6.1
/*****
/*-----
/* IF CHARGEBACK IS NOT BEING USED, THEN THIS STEP MAY BE OMITTED.
/*-----
//ST0261A EXEC PGM=IEFBR14
//DELETE1 DD DSN=&PREFIX.CIMSIMS.CIMSIMS1, DISP=(MOD,DELETE),
// SPACE=(1,1)
//DELETE7 DD DSN=&PREFIX.CIMSIMS.CIMSIMS7, DISP=(MOD,DELETE),
// SPACE=(1,1)
/*-----
//ST0261B EXEC PGM=CIMSLP61, REGION=OM, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
```

```

//SYSUT1 DD DSN=IMS.LOGTAPE,DISP=(OLD,KEEP,KEEP)
//SYSUT2 DD DSN=&&IMSLOG,DISP=(,PASS)
//CPRSTAT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN001 DD *
* PROCESS=CHARGEBACK,CAPACITY PLANNER
* PROCESS=CAPACITY PLANNER,CHARGEBACK
* PROCESS=CHARGEBACK
  PROCESS=CAPACITY PLANNER
* LAST RUN=YES
  IMS SYSTEM=IIII
//*-----
//* THE FOLLOWING DD STATEMENTS ARE NECESSARY ONLY IF THE
//* IF THE CIMS CHARGEBACK LOG PROCESSING IS BEING PERFORMED.
//*-----
//CIMSPRNT DD SYSOUT=*
//CIMSIMS1 DD DSN=&PREFIX.CIMSIMS.CIMSIMS1,DISP=(,CATLG,DELETE),
//          LRECL=80,DSORG=PS,RECFM=FB,BUFNO=10,BLKSIZE=3120,
//          SPACE=(CYL,(25,25))
//CIMSIMS7 DD DSN=&PREFIX.CIMSIMS.CIMSIMS7,DISP=(,CATLG,DELETE),
//          LRECL=27994,BLKSIZE=27998,DSORG=PS,RECFM=VB,
//          BUFNO=10,SPACE=(CYL,(50,50))
//CIMSPASS DD DISP=SHR,DSN=&PREFIX.CPPR.V530.CNTL(CIMSNUM)
//*-----
//* END OF SAMPLE JCL FOR CIMSLP61 IMS LOG PREPROCESSOR *
//*****
//*****
//* THIS STEP SORTS THE SELECTED IMS DATA - ALL IMS RELEASES
//*****
//SORT EXEC PGM=SORT,REGION=4096K,TIME=10
//SYSOUT DD SYSOUT=*
//SORTIN DD DSN=&&IMSLOG,DISP=(OLD,DELETE)
//SORTOUT DD DSN=&&LOGSRT,UNIT=3380,DISP=(,PASS),
//          SPACE=(CYL,(10,3)),
//          DCB=(DSORG=PS,BLKSIZE=19069,LRECL=3120,RECFM=VB)
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SYSIN DD *
  SORT FIELDS=(5,4,PD,A,9,4,PD,A),SIZE=E60000
  END
//*****
//* THIS STEP PERFORMS THE DATA REDUCTION - ALL IMS RELEASES
//*****
//ST03 EXEC PGM=SSA1IMSW,REGION=5000K,TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ABNLIGNR DD DUMMY
//*-----
//* THE FOLLOWING INPUT COMES FROM THE SMF TYPE 30 RECORDS
//*-----
//SYSUT1 DD DISP=(OLD,DELETE),DSN=*.ST01.SYSUT2
//*-----
//* THE FOLLOWING INPUT COMES FROM THE IMS LOGTAPE
//*-----
//SYSUT2 DD DISP=(OLD,DELETE),DSN=&&LOGSRT
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS

```



```
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CIMSPASS DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CNTL(CIMSNUM)
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT3 DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
IMS SYSTEM=IIII
IMS CONTROL=IMSCTL
DBRC REGION=IMSDBRC
DLI REGION=IMSDLI
DSNMSTR REGION=DB2MSTR
DSNDBM1 REGION=DB2DBM1
IMS DUMP=YES
```

For IMS 5.1 Log processing see steps ST0251A and ST0251B in DIMSPROD.

IMF Data Reduction JCL (DIMSPROD)

Following is a set of JCL that is used to run the IMS Data Reduction program using input from an IMF Log and the corresponding SMF data.

```
//SSACPPR JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST01 EXEC PGM=SSA1WKLD,REGION=5000K,TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUT1 DD DISP=SHR,DSN=SMF.INPUT.FILE
/*SYSMAN0 DD DISP=SHR,DSN=SYS1.MAN0
/*SYSMAN1 DD DISP=SHR,DSN=SYS1.MAN1
/*SYSMAN2 DD DISP=SHR,DSN=SYS1.MAN2
/*SYSMAN3 DD DISP=SHR,DSN=SYS1.MAN3
//ONLINE DD DUMMY
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CIMSPASS DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CNTL(CIMSNUM)
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//*****
/* THE FOLLOWING PASSES SMF TYPE 30 RECORDS TO ST03
//*****
//SYSUT2 DD SPACE=(CYL,(100,50)),UNIT=SYSDA,DISP=(,PASS)
//SYSUT3 DD SPACE=(TRK,(1,1)),UNIT=SYSDA,DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=INCLUDE(5)
* * * DONT PUT ANYTHING INTO THE ONLINE PERFORMANCE DATABASE
SCANONLY
* * * PASS SMF TYPE 30 RECORDS INTO THE SYSUT2 FILE * * *
SYSUT2=YES
FILTER=30
//*****
/*
/* STEP ST02 FOR IMF
/*
//*****
```

```

//ST02      EXEC PGM=IMFCOPY5,REGION=2000K
//STEPLIB  DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//SYSUT1   DD DSN=IMS.LOGTAPE,DISP=(OLD,KEEP)
//SYSUT2   DD DSN=&&IMSLOG,DISP=(,PASS),UNIT=SYSDA,
//          DCB=(DSORG=PS,BLKSIZE=19069,RECFM=VB),
//          SPACE=(CYL,(10,10))
//CIMSPASS DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CNTL(CIMSNUM)
//SYSUDUMP DD SYSOUT=*
//*-----
//*
//* THE FOUR-CHARACTER IMS SYSTEM ID MUST BE SUBSTITUTED FOR IIII
//*
//*-----
//SYSIN     DD *
IMS SYSTEM=IIII
/*
//*****
//* THIS STEP SORTS THE SELECTED IMS DATA - ALL IMS RELEASES
//*****
//SORT      EXEC PGM=SORT,REGION=4096K,TIME=10
//SYSOUT    DD SYSOUT=*
//SORTIN   DD DSN=&&IMSLOG,DISP=(OLD,DELETE)
//SORTOUT   DD DSN=&&LOGSRT,UNIT=3380,DISP=(,PASS),
//          SPACE=(CYL,(10,3)),
//          DCB=(DSORG=PS,BLKSIZE=19069,LRECL=3120,RECFM=VB)
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(10,3))
//SYSIN     DD *
SORT FIELDS=(5,4,PD,A,9,4,PD,A),SIZE=E60000
END
//*****
//* THIS STEP PERFORMS THE DATA REDUCTION - ALL IMS RELEASES
//*****
//ST03      EXEC PGM=SSA1IMSW,REGION=5000K,TIME=60
//STEPLIB  DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//ABNLIGNR DD DUMMY
//*-----
//* THE FOLLOWING INPUT COMES FROM THE SMF TYPE 30 RECORDS
//*-----
//SYSUT1   DD DISP=(OLD,DELETE),DSN=*.ST01.SYSUT2
//*-----
//* THE FOLLOWING INPUT COMES FROM THE IMS LOGTAPE
//*-----
//SYSUT2   DD DISP=(OLD,DELETE),DSN=&&LOGSRT
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CIMSPASS DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CNTL(CIMSNUM)
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT3   DD SPACE=(CYL,(10,8)),UNIT=SYSDA,DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN     DD *
PARAMETER1

```

.
. .
PARAMETERn

Data Reduction DD Statements

■ CIMSPASS

CIMS product passwords.

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all IMS Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of the CIMS Capacity Planner report generation programs.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ SYSUT1

The SYSUT1 DD statement specifies the input to the CIMS Capacity Planner Data Extraction Programs (IMSCOPY and IMFCOPY) and the SMF data input to the IMS and IMF Data Reduction programs (SSA1IMSW and SSA1IMWF).

■ SYSUT2

The SYSUT2 DD statement is used to specify the data set containing the IMS/IMF Log data.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the IMS Ad Hoc report. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA, LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to the IMF SORT step and the CIMS Capacity Planner data reduction programs. The parameters that relate to the IMS data reduction phase are described in the next paragraph.

CIMS Capacity Planner IMS Data Reduction Parameters

The CIMS Capacity Planner Parameters permit you to specify certain run-time options and to provide selected data required by the data reduction program in processing the IMS history files and producing the Ad Hoc report.

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMFSID of the MVS system for that data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID is specified by this parameter in running the IMS data reduction program.

Parameters

The SELECTED SYSTEM= parameter is specified in several ways:

- The SID itself (e.g. SELECTED SYSTEM=SYS1)
- An "*", indicating the SID for the system on which the program is being executed (e.g. SELECTED SYSTEM=*)

Default Parameter

SELECTED SYSTEM=*

IMS SYSTEM

This keyword is used to specify the four-character IMS System ID that identifies the particular IMS system being monitored. There is no default value for this parameter.

Specify IMS System ID

IMS SYSTEM=imsid (e.g. IMS SYSTEM=IMS1)

Default IMS System ID

No default value is defined

IMS CONTROL

This keyword is used to specify the job name of the Address space in which the IMS Control Region operates. The Control Region job name is used to identify the SMF Type 30 records that apply to the IMS Control Region.

Specify Job Name

IMS CONTROL=jobname (e.g, IMS CONTROL=IMS1CNTL)

Default Control Region Name

No default value is defined.

DBRC REGION

This keyword is used to specify the job name of the DBRC Address space. The DBRC Region job name is used to identify the SMF Type 30 records that apply to the IMS DBRC Region.

Specify Job Name

DBRC REGION=jobname (e.g, IMS CONTROL=IMS1DBRC)

Default Control Region Name

No default value is defined.

DLI REGION

This keyword is used to specify the job name of the DLI Address space. The DLI Region job name is used to identify the SMF Type 30 records that apply to the IMS DLI Region.

Specify Job Name

DLI REGION=jobname (e.g, DLI REGION=IMS1DLI)

Default DLI Region Name

No default value is defined

DSNMSTR REGION

This keyword is used to specify the job name of the DB2 DSNMSTR Address space. The DSNMSTR Region job name is used to identify the SMF Type 30 records that apply to the IMS DSNMSTR Region. This keyword parameter is required only if you are running DB2 in conjunction with IMS.

Specify Job Name

DSNMSTR REGION=jobname (e.g, DSNMSTR REGION=DSNMSTR)

Default DSNMSTR Region Name

No default value is defined.

DSNDBM1 REGION

This keyword is used to specify the job name of the DB2 DSNDBM1 Address space. The DSNDBM1 Region job name is used to identify the SMF Type 30 records that apply to the DSNDBM1 Region. This keyword parameter is required only if you are running DB2 in conjunction with IMS.

Specify Job Name

DSNDBM1 REGION=jobname (e.g. DSNDBM1 REGION=DSNMSTR)

Default DSNDBM1 Region Name

No default value is defined.

FORCE IMS LOG=YES

The FORCE IMS LOG parameter governs whether or not the dates of the IMS history file are checked to avoid duplicate updates. It is useful to process logs that have been skipped for one reason or another.

Parameters

FORCE IMS LOG=YES

FORCE IMS LOG=NO

Default Parameter

FORCE IMS LOG=NO

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should only be used when working with the CIMS Lab to isolate problems.

Parameters

None

Default Parameter

None (equivalent to DEBUGOFF)

Report Production

All the IMS batch reports and graphs are produced by one of the following report generation programs:

- SSA1IMSR—the IMS Report Generation Program
- SSA1IMSP—the IMS Transaction Profile
- SSA1IMST—the IMS Trends Analysis/Summary Program
- SSA1IMSG—the IMS Subsystem Graph Program
- SSA1IMSE—the IMS Subsystem Extension Graph Program

DD Statements

The DD Statements required to execute the above programs are described below:

- STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

- CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

- CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all IMS Job streams.

- INDEX

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (the Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

- ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

- HGDLIB

The HGDLIB DD statement specifies the Partitioned Data Set used to save the data points required to produce the various CIMS Capacity Planner graphs using the Harvard Graphics subsystem. The data points contained in the members of the HGDLIB PDS are down-loaded to a PC and matched up with a set of CIMS Capacity Planner graph templates that specify how the graphic data is to be presented. The parameters required are the Disposition and the Data Set Name.

- SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the IMS printed reports. Unless otherwise specified, the DCB characteristics are:

```
RECFM=FBA, BLKSIZE=133, LRECL=133
```

- SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

- SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ EXCLUDE/INCLUDE

This is the (optional) Element Exclusion file. It is used by the IMS Report Program to force certain elements to be excluded from the reports or to force the inclusion of certain elements in the reports that would not normally be included. When using the INCLUDE, only the specified elements are included in the reports.

The EXCLUDE/INCLUDE file must consist of 80 character records, 72 of which are interpreted. Columns 73 through 80 are used for sequence numbers. Multiple elements are entered in a single record, but they must be separated by commas. The format of the input statement follows:

- An element is up to eight characters in length.
- Multiple elements are included in a single record provided that they are separated by commas or blanks.
- Elements are terminated by the wild card character "*", indicating that all elements that match up to the "wild card" are to be excluded/included.

Example

The following statement in the EXCLUDE file:

```
TSOP,CS*
```

would cause the report program SSAIMSR to exclude any data related to the transaction named TSOP. It would also cause the exclusion of all transactions beginning with the characters "CS".

The use of the INCLUDE causes the report program to limit the contents of the report to only those elements specified in the INCLUDE data set.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the IMS data reduction phase are described in *Additional Report Options on page 6-28* below.

Report Job Streams

CIMS Capacity Planner Batch Report JCL

The standard CIMS Capacity Planner IMS batch reports are all produced by the CIMS Capacity Planner IMS Report Program—SSA1IMSR.

Following is an example of the JCL required to run the report program:

```
//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IMSR, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
```



```
//SYSUDUMP DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYMSMGS DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=*
IMS SYSTEM=IMSA
IMS SUMMARY REPORT=YES
IMS EXCEPTION ANALYSIS=YES
IMS PERFORMANCE REPORT=YES
IMS TRANSACTION STATISTICS REPORT=YES
IMS TERMINAL STATISTICS REPORT=YES
IMS TRANSACTION RESPONSE GRAPH=YES
IMS TERMINAL RESPONSE GRAPH=YES
IMS TRANSACTION ACTIVITY GRAPH=YES
IMS TERMINAL ACTIVITY GRAPH=YES
IMS CPU ACTIVITY GRAPH=YES
IMS I/O ACTIVITY GRAPH=YES
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
BEGIN DATE=06/02/2005
END DATE=07/22/2005
//
```

The SYSIN parameters used to request reports and to specify reporting options are described below under *Additional Report Options on page 6-28*.

Member DIMSREPT in the PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that is customized for your installation.

IMS Transaction Profile Report JCL

The following JCL is used to run the IMS Transaction Profile Report.

```
//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1     EXEC PGM=SSA1IMSP,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYMSMGS DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=*
IMS SYSTEM=IMSA
BEGIN DATE=*-7
END DATE=*
//INCLUDE DD *
ACCRT*
//
```

This example would result in the production of a Transaction Profile report that would summarize the data for transaction codes beginning with ACCRT and reporting the results as though they were one composite transaction.

IMS Trends Analysis Summary Program

Following is an example of the JCL required to run the IMS Trends Analysis Report Program—SSA1IMST.

```
//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IMST,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
IMS SYSTEM=IMS1
1ST PERIOD BEGIN DATE=02/16/2005
1ST PERIOD END DATE=04/15/2005
2ND PERIOD BEGIN DATE=05/16/2005
2ND PERIOD END DATE=07/155/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
LUNCH BREAK BEGIN HOUR=11
LUNCH BREAK END HOUR=12
```

IMS Presentation Graphics Program

The CIMS Capacity Planner IMS Subsystem provides an extensive number of presentation-quality graphs that depict the IMS workload and performance. The graphs are produced for individual IMS regions operating under a specific `SELECTED SYSTEM`. The data to be included in the graphs are produced by the CIMS Capacity Planner program (SSA1IMSG) and written to selected members of the HGDLIB data set to be down-loaded to a PC and subsequently plotted by the Harvard Graphics presentation graphics system.

IMS graphs can also be produced for All IMS regions that are registered in the Element Registration Table for a specific `SELECTED SYSTEM`, consolidated together, by executing the CIMS Capacity Planner IMS Extension Graphs Program (SSA1IMSE)

IMS Graph Program—(SSA1IMSG) JCL

Following is an example of the JCL required to operate the SSA1IMSG program to generate the data points for the IMS graphs.

```
//SSAGRPH JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1IMSG,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
```

```
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYMSMSG  DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=CPPR
IMS SYSTEM=IMS1
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of sample Job streams for producing the CIMS Capacity Planner IMS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

- DIMSHGRF—Produce Hourly Graphs
- DIMSDGRF—Produce Daily Graphs
- DIMSWGFR—Produce Weekly Graphs
- DIMSMGRF—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program - SSA1IMSG. To tailor the Job streams, merely modify the selection criteria as required.

IMS Consolidated Graph Program—(SSA1IMSE) JCL

Following is an example of the JCL required to operate the SSA1IMSE program to generate the consolidated IMS graph data points.

```
//SSAGRPH JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1     EXEC PGM=SSA1IMSE,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.IMS
//ONLINE  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.IMS
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP  DD SYSOUT=*
//SYMSMSG DD SYSOUT=*
//SYSIN   DD *
SELECTED SYSTEM=CPPR*
IMS SYSTEM=IMS1
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of sample Job streams for producing the CIMS Capacity Planner consolidated IMS graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

- DIMSHGRE—Produce Hourly Graphs
- DIMSDGRE—Produce Daily Graphs
- DIMSWGRE—Produce Weekly Graphs
- DIMSMGRE—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program SSA1IDME. To tailor the Job streams, merely modify the data set names and the selection criteria as required.

IMS Reporting Parameters

SELECTED SYSTEM

The `SELECTED SYSTEM=` parameter is used to specify the SID of the system for which reports are to be produced or trends to be analyzed. This is the SMF identifier that is specified in the `SMFPRMxx` member of `SYS1.PARMLIB`.

The `SELECTED SYSTEM` parameter has the following options:

- The SID itself (e.g., `SELECTED SYSTEM=MVSA`)
 - An "*", indicating the SID for the system on which the report program is being executed (e.g., `SELECTED SYSTEM=*`)
- Default SID
 - `SELECTED SYSTEM=*`

BEGIN DATE

The `BEGIN DATE=` parameter is used to specify the beginning of the period for which data in the Performance Data Base is to be collected for analysis and reporting.

Specify a beginning date in one of the following formats:

- YYDDD (Standard Julian format)
- MM/DD/YY (USA standard Gregorian format)
- DD.MM.YY (Standard European Gregorian format)
- *, *-n (relative day format where *=today)

Default Beginning Date:

`BEGIN DATE=00001`

END DATE

The `END DATE=` parameter is used to specify the last date (inclusive) for which an IMS record is to be included for analysis.

Specify End Date

The format of the parameters for the `END DATE` are identical to the `BEGIN DATE` parameters specified above in *BEGIN DATE on page 6-22*

Default End Date

`END DATE=2099365`

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Parameters

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Parameter

`REPORT LANGUAGE=ENGLISH`

DEBUGON

This parameter is used in combination with the `SYSNAP DD` statement to request SNAP dumps be taken at various points in the data reduction program. It should only be used when working with the CIMS Lab to isolate problems.

Parameters

`DEBUGON` (Turn on debug feature)

Default Parameter

None (equivalent to `DEBUGOFF`)

PRIME SHIFT FIRST HOUR

The `PRIME SHIFT FIRST HOUR=` parameter specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Specify the beginning of the Prime Shift

`PRIME SHIFT FIRST HOUR=7`

Default begin time

`PRIME SHIFT FIRST HOUR=6`

LATE SHIFT FIRST HOUR

The `LATE SHIFT FIRST HOUR=` parameter specifies the hour of the day on which the LATE Shift begins. It is expressed as an hour using the 24 hour convention.

Specify the beginning of the Late Shift

`LATE SHIFT FIRST HOUR=19`

Default begin time

`LATE SHIFT FIRST HOUR=18`

IMS SYSTEM

The IMS SYSTEM= parameter designates the IMS region for which data is to be collected and reported.

Specify IMS ID

IMS SYSTEM= any IMS Region Name (VTAM APPLID) that has been registered using the SSA1REGI program (see *Chapter 1, Installing CIMS Capacity Planner* in the *CIMS Capacity Planner Installation and Getting Started Guide*).

Default IMS ID

None

IMS SUMMARY REPORT

This parameter is used to specify whether or not the IMS Summary Report is to be produced.

Parameters

IMS SUMMARY REPORT=YES (produce the report)

IMS SUMMARY REPORT=NO (don't produce the report)

Default Parameter

SUMMARY REPORT=NO

IMS TRANSACTION STATISTICS REPORT

This parameter specifies whether or not the IMS Transaction Statistics Report is to be produced.

Parameters

IMS TRANSACTION STATISTICS REPORT=YES (produce the report)

IMS TRANSACTION STATISTICS REPORT=NO (don't produce the report)

Default Parameter

IMS TRANSACTION STATISTICS REPORT=NO

IMS TERMINAL STATISTICS REPORT

This parameter specifies whether or not the Terminal Statistics Report is to be produced.

Parameters

IMS TERMINAL STATISTICS REPORT=YES (produce the report)

IMS TERMINAL STATISTICS REPORT=NO (don't produce the report)

Default Parameter

TERMINAL STATISTICS REPORT=NO

IMS PERFORMANCE REPORT

This parameter specifies whether or not the IMS Performance Report is to be produced.

Parameters

IMS PERFORMANCE REPORT=YES (produce the report)

IMS PERFORMANCE REPORT=NO (don't produce the report)

Default Parameter

IMS PERFORMANCE REPORT=NO

IMS CPU ACTIVITY GRAPH

This parameter specifies whether or not the IMS CPU Activity Graph is to be produced.

Parameters

IMS CPU ACTIVITY GRAPH=YES (produce the Graph)

IMS CPU ACTIVITY GRAPH=NO (don't produce the report)

Default Parameter

IMS CPU ACTIVITY GRAPH=NO

IMS I/O ACTIVITY GRAPH

This parameter specifies whether not the CPU I/O Activity Graph is to be produced.

Parameters

IMS I/O ACTIVITY GRAPH=YES (produce the graph)

IMS I/O ACTIVITY GRAPH=NO (don't produce the graph)

Default Parameter

IMS I/O ACTIVITY GRAPH=NO

IMS TERMINAL ACTIVITY GRAPH

This parameter specifies whether or not the IMS Terminal Activity Graph is to be produced.

Parameters

IMS TERMINAL ACTIVITY GRAPH=YES (produce the graph)

IMS TERMINAL ACTIVITY GRAPH=NO (don't produce the graph)

Default Parameter

IMS TERMINAL ACTIVITY GRAPH=NO

IMS TRANSACTION RESPONSE GRAPH

This parameter specifies whether or not the IMS Transaction Response Graph is to be produced.

Parameters

IMS TRANSACTION RESPONSE GRAPH=YES (produce the graph)

IMS TRANSACTION RESPONSE GRAPH=NO (don't produce the graph)

Default Parameter

IMS TRANSACTION RESPONSE GRAPH=NO

IMS TERMINAL RESPONSE GRAPH

This parameter specifies whether or not the IMS Terminal Response Graph is to be produced.

Parameters

IMS TERMINAL RESPONSE GRAPH=YES (produce the graph)

IMS TERMINAL RESPONSE GRAPH=NO (don't produce the graph)

Default Parameter

IMS TERMINAL RESPONSE GRAPH=NO

IMS EXCEPTION ANALYSIS

This parameter specifies whether or not the IMS Exception Analysis Report is to be produced.

Parameters

IMS EXCEPTION ANALYSIS=YES (produce the report)

IMS EXCEPTION ANALYSIS=NO (don't produce the report)

Default Parameter

IMS EXCEPTION ANALYSIS=NO

1ST PERIOD BEGIN DATE

This parameter is used to specify the beginning of the first of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1IMST) or the beginning date for the entire period to be used for developing the Trends Analysis Graph (SSA1IMSG) for which consolidated table information is to be selected for analysis.

Specify a beginning date in one of the following formats:

- YYDDD (Standard Julian format)
- MM/DD/YY (USA standard Gregorian format)
- DD.MM.YY (Standard European Gregorian format)
- *, *-n (relative day format where *=today)

Default date

1ST PERIOD BEGIN DATE=00001

1ST PERIOD END DATE

This parameter is used to specify the ending date of the first period for the Trends Analysis Summary Program (SSA1IMSTT) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1IMSG).

Specify date

Specify an ending date as in 1ST PERIOD BEGIN DATE.

Default date

END DATE=2099365

2ND PERIOD BEGIN DATE

This parameter is used to specify the beginning of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1IMST) or the Trends Analysis Graph (SSA1IMSG).

Specify date

Specify a beginning date in the same format as the 1ST PERIOD BEGIN DATE.

Default date

2ND PERIOD BEGIN DATE=00001

2ND PERIOD END DATE

This parameter is used to specify the ending date of the second period for the Trends Analysis Summary Program (SSA1IMST) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1IMSG).

Specify date

Specify an ending date as in 1ST PERIOD BEGIN DATE.

Default date

2ND PERIOD END DATE=2099365

GENERAL REPORT STATEMENTS

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems.

Example

TOP LABEL=YES

This statement causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time: End as specified
- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

FOOT1=UP TO 64 CHAR FOOTNOTE

This statement is used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that is included on all reports where there is room.

FOOT2

FOOT2=UP TO 64 CHAR 2ND FOOTNOTE

This statement is used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that is included on all reports where there is room.

Additional Report Options

In addition to the INCLUDE/EXCLUDE options described in *DD Statements on page 6-17*, it is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allows you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center can have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it is useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report is prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

GENERIC ELEMENT MASK=xxxxxxx

INCLUSIVE ELEMENT MASK=xxxxxxx

EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is `type ELEMENT MASK=xxxxxxx`.

Where "type" is GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxx is any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved.

Example

GENERIC ELEMENT MASK=0000000

This statement specifies that the key is to remain untouched, whereas the statement:

```
GENERIC ELEMENT MASK=11111111
```

Specifies that the entire record key should be masked out (set to "*****"). Finally, the statement:

```
GENERIC ELEMENT MASK=10101010
```

Specifies that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The `GENERIC ELEMENT MASK` parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the `GENERIC ELEMENT MASK` parameter does not require the use of the `INCLUDE OR EXCLUDE` facilities.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a `GENERIC ELEMENT MASK=00001100`. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form `****LL**`. If `GENERIC ELEMENT MASK=11111111` were to be specified, all records would be summarized into a single entry.

INCLUSIVE ELEMENT MASK

The `INCLUSIVE ELEMENT MASK` keyword phrase causes the IMS report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the `INCLUDE` selection module. In this case, the key is restored to its original value once the `INCLUDE` selection processing is completed.

The `INCLUSIVE ELEMENT MASK` parameter differs from the `GENERIC ELEMENT MASK` in that the `INCLUSIVE ELEMENT MASK` is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the `GENERIC ELEMENT MASK` is used.

The mask specification procedure is the same as for the `GENERIC ELEMENT MASK` described above.

The `INCLUSIVE ELEMENT MASK` is used with an `INCLUDE` data set specified in the JCL by the `//INCLUDE DD` statement. The format of the `INCLUDE` entries are:

```
//INCLUDE DD *
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

Where `xxxxxxx` is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the `INCLUDE` data set would contain "????23??. The "?" characters in the `INCLUDE` selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The `EXCLUSIVE ELEMENT MASK` keyword phrase causes the IMS report processing programs to intercept the key of each element (in this case, the VTAM terminal name), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the `EXCLUDE` selection module. In this case, the key is restored to its original value once the `EXCLUDE` selection processing is completed.

The `EXCLUSIVE ELEMENT MASK` parameter differs from the `GENERIC ELEMENT MASK` in that the `EXCLUSIVE ELEMENT MASK` is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the `GENERIC ELEMENT MASK` is used.

The mask specification procedure is the same as for the `GENERIC ELEMENT MASK` described above.

The `EXCLUSIVE ELEMENT MASK` is used with an `EXCLUDE` data set specified in the JCL by the `//EXCLUDE DD` statement. The format of the `EXCLUDE` entries are:

```
//EXCLUDE DD *
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

Where `xxxxxxx` is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the `EXCLUDE` data set would contain "????23??. The "?" characters in the `EXCLUDE` selection statements correspond to the "*" positions in the modified keys.

ISPF/PDF IMS Subsystem Interface

The CIMS Capacity Planner IMS subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM Graphs, generate the full range of batch reports to be printed on the system printer, and generate input to the PC Presentation Graphics subsystem (Harvard Graphics).

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, IMS, etc.).

On-line Graphs

The following IMS GDDM graphs are available on-line through the ISPF Interface:

IMS Transaction Response Graph

The Transaction Response graph depicts the average response time for all IMS transactions. The graph shows the prime, early, and late shift averages by hour, day, week, or month within the range of dates specified in the report request.

IMS Transaction Volume Graph

The IMS Transaction Volume Graph depicts the average volume for all transactions executed. Within the range of dates and the period specified in the report request, the graph shows the prime, early, and late shift averages by hour, day, week, or month.

IMS Terminal Activity Graph

The IMS Terminal Activity Graph depicts the average number of IMS terminals active within the range of dates specified in the report request. The graph shows the average number of terminals active during the prime, early, and late shifts by hour, day, week, or month, depending upon which is requested.

On-line Reports

The following reports are produced on-line through the CIMS Capacity Planner IMS ISPF Interface.

Statistical Reports

IMS Transaction Statistics Report

The Transactions Statistics Report provides a list of the IMS transactions that rank among the top 40 in the following categories:

- Most frequently executed
- Required the most cumulative transaction residence time
- Consumed the most CPU time
- Required the most file accesses

IMS Terminal Statistics Report

The IMS Terminal Statistics Report shows which IMS terminals were among the top 40 in the following categories:

- Submitted the most transactions
- Accumulated the most active time (a terminal is active during each 15 period throughout the day if it is actively doing work. Terminals that are logged on, but not submitting transactions are not considered to be active).

Analysis Reports

IMS Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the IMS workload and how well the IMS system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time
- The peak number of concurrent active terminals (high watermark)
- The average number of IMS transactions processed per minute during the prime shift
- The average number of data base calls per second during the prime shift
- The average number of IMS-related EXCPs per second during the prime shift
- The average transaction elapsed time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average percent of the CPU consumed by IMS for each shift by the IMS Message Processing Regions and the IMS System Regions
- The average number of transactions processed by shift for the measured period

The 25 most frequently executed transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions are reported by shift.

IMS Trends Analysis Report

The on-line Trend Analysis report shows the trends in selected characteristics of the IMS usage and performance that are indicators of the future IMS workload. Trends are reported for:

- CPU utilization, including a breakdown in terms of IMS overhead and task related Usage.
- The average number of transactions executed per minute during prime shift
- The average number of EXCPs per second during prime shift
- The average response time during prime shift
- The ratio of early shift to prime shift transactions
- The ratio of late shift to prime shift transactions
- The average ratio of active terminals on early shift compared to prime shift
- The average ratio of active terminals on late shift compared to prime shift

IMS Exception Analysis Report

The IMS on-line Exception Analysis report shows, for the period specified in the report request, any 15 minute intervals during which the following predefined limits were exceeded. The limits are specified in the CPPR.PARMLIB member &sidXCPT.

- Average transaction processing time
- Average response time

IMS Ad Hoc Report

The IMS Ad Hoc report is not yet implemented under the ISPF Interface. Please use the DIMSADHC member of the CPPR.CNTL library for the JCL required to run the Ad Hoc report.

Batch Reports

A Batch Job submission panel is provided so that the full range of CICS batch reports and graphs described earlier under *IMS Workload and Performance Reports on page 6-2*, are produced through the ISPF Interface. This considerably simplifies the processes of operating CIMS Capacity Planner.

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch Jobs instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch Job
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- Saving the Chart as a GDF File (Batch)

These options are explained in detail below

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual Harvard Graphics charts. As originally designed, the data members are created by Batch Jobs, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through Harvard Graphics. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch Job is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of several key phrases.

GDDM BATCH PROCESS=YES

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

GDDM PRINTER NAME=pppppppp

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQ DD Statement in the JCL.

Several members have been added to the Distribution CNTL file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch Job. The member name "GDDMBATI" shows how to produce any of the IMS GDDM graphs in Batch.

GDDMBAT PARMLIB Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (CPPR201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST0001) WHICH CAN THEN BE
/* SENT TO A GDDM PRINTER (PRRINT01).
/* * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//         DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR, DSN=SYS1.GDDMSYM
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR, DSN=SYS1.GDDMGDF
/* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR, DSN=SYS1.GDDM.REQUEST.QUEUE
```



```

//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00)'
//SYSIN DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*
* | | | % CHARACTER LENGTH
* | | % CHARACTER WIDTH
* | HORIZONTAL OFFSET
* VERTICAL OFFSET
//INCLUDE DD *
CPPR201

```

GDDMBATI PARMLIB Member

```

//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR all
/* THE SUPPORTED IMS CHARTS AND SENDING THEM TO A PRINTER.
/* * * * * *
//STO EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
// DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD DISP=SHR, DSN=SYS1.GDDMSYM
//ADMCDATA DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF DD DISP=SHR, DSN=SYS1.GDDMGDF
/* THE FOLLOWING DD NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR, DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD40)' /* IMS RESPONSE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD41)' /* TRANSACTION VOLUME /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD42)' /* TERMINAL VOLUME /*
//SYSIN DD *
IMS SYSTEM=IMSA
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST0001
GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */

```

```
*GDDM PRINT CONTROLS=00,00,100,100
*           | | | % CHARACTER LENGTH
*           | | % CHARACTER WIDTH
*           | HORIZONTAL OFFSET
*           VERTICAL OFFSET
```

Specifying GDDM Batch Print Controls

It is possible to specify certain Print Controls that are to take effect when printing a chart in Batch. The key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter allows you to specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

You might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. In order to invoke these local templates in a Batch Job, a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter is added to the SYSIN, where the "nnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=dddddddd
```

Where the "dddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeeee
```

Where the "eeeeeeee" is the name under which the graph is saved in the GDF library.

Presentation Graphics

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the IMS workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required IMS workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics—release 3.0, plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, that describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and is used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, are used to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner IMS graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in *Graph Descriptions on page 6-41* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters in the form:

"&sidpiiq" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly
- "iiq" is the graph identifier that consists of two alpha-numeric characters. The "piiq" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under *Presentation Graphics Distribution Disk on page 6-40*. The qualifier "q" can or can not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpiiq.DAT

Where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As is seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to download the data points from HGDLIB to the PC using IND\$FILE. The BAT files are modified as required to fit the naming conventions of your installation.

The names of the BAT files all conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
 - "WKLD"—Workload graph data
 - "DASM"—DASM graph data
 - "CICS"—CICS graph data
 - "CICE"—CICS Extension graph data
 - "IDMS"—IDMS graph data
 - "IDME"—IDMS Extension graph data
 - "IMSG"—IMS graph data
 - "NETG"—Network graph data
 - "NETE"—Network Extension graph data
 - "NETL"—Network Physical Lines graph data
 - "NETN"—Network Physical NCP graph data
 - "DB2S"—DB2 Subsystem graph data
 - "DB2U"—DB2 Connect ID graph data
 - "M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M 204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGD LIB data point members from the Host processor (assumes the use of IND\$FILE for the file transfer). The naming conventions follow.

.BAT File Naming Conventions

The .BAT filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1= Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a subsystem identifier:

- WKLD—Workload
- DASM—DASM
- CICS—CICS
- CICE—CICS Extension
- IDMS—IDMS
- IDME—IDMS Extension
- IMS—IMS
- NETG—Network
- NETE—Network Extension
- NETL—Network Physical Lines
- NETN—Network Physical NCP

For example, the .BAT file used to down-load the HGDLIB members required to produce the Workload Daily graphs would be named HGD2WKLD.BAT.

.TPL File Naming Conventions

The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1=Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by 2 character graph identifier. In some cases, a 1 character qualifier is added extending the graph identifier a 3 characters. Please refer to the Graphic Report Tables for the specific template names.

Graph Periods

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide*, for a description of the hourly, daily, weekly, and monthly graph periods.

Graph Descriptions

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide*, for descriptions of each of the graphs produced by the CIMS Capacity Planner IMS Subsystem.

Graphic Report Tables

The following tables summarize the relationship of members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-based program.

IMS Subsystem Hourly Graphs

These graphs represent specific IMS Regions executing within the &sid system. Therefore, IMS Region suffixes are required to uniquely identify a particular IMS Region.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1IMSG	&sid190*	(none)	IMS Transaction Response
	&sid191*	(none)	IMS Transaction Volume
	&sid192*	(none)	IMS CPU Utilization
	&sid193*	(none)	IMS I/O Activity
	&sid194*	(none)	IMS Terminal Activity
	&sid195*	(none)	IMS Transactions/Terminal
	&sid196*	(none)	IMS Transactions/Day

* plus a one character suffix representing the IMS Region

&sid is the SMF System ID (e.g, CPPR)

IMS Subsystem Daily Graphs

These graphs represent specific IMS Regions executing within the &sid system. Therefore, IMS Region suffixes are required to uniquely identify a particular IMS Region.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1IMSG	&sid290*	(none)	IMS Transaction Response
	&sid291*	(none)	IMS Transaction Volume
	&sid292*	(none)	IMS CPU Utilization
	&sid293*	(none)	IMS I/O Activity
	&sid294*	(none)	IMS Terminal Activity
	&sid295*	(none)	IMS Transactions/Terminal
	&sid296*	(none)	IMS Transactions/Day
	&sid297*	(none)	IMS Transactions/Shift
	&sid298*	(none)	IMS Terminals/Shift

* plus a one character suffix representing the IMS Region

&sid is the SMF System ID (e.g, CPPR)

IMS Subsystem Weekly Graphs

These graphs represent specific IMS Regions executing within the &sid system. Therefore, IMS Region suffixes are required to uniquely identify a particular IMS Region.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1IMSG	&sid390*	(none)	IMS Transaction Response
	&sid391*	(none)	IMS Transaction Volume
	&sid392*	(none)	IMS CPU Utilization
	&sid393*	(none)	IMS I/O Activity
	&sid394*	(none)	IMS Terminal Activity
	&sid395*	(none)	IMS Transactions/Terminal
	&sid396*	(none)	IMS Transactions/Day
	&sid397*	(none)	IMS Transactions/Shift
	&sid398*	(none)	IMS Terminals/Shift

* plus a one character suffix representing the IMS Region

&sid is the SMF System ID (e.g, CPPR)


IMS Subsystem Monthly Graphs

These graphs represent specific IMS Regions executing within the &sid system. Therefore, IMS Region suffixes are required to uniquely identify a particular IMS Region.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1IMSG	&sid490*	(none)	IMS Transaction Response
	&sid491*	(none)	IMS Transaction Volume
	&sid492*	(none)	IMS CPU Utilization
	&sid493*	(none)	IMS I/O Activity
	&sid494*	(none)	IMS Terminal Activity
	&sid495*	(none)	IMS Transactions/Terminal
	&sid496*	(none)	IMS Transactions/Day
	&sid497*	(none)	IMS Transactions/Shift
	&sid498*	(none)	IMS Terminals/Shift
			(none)

* plus a one character suffix representing the IMS Region

&sid is the SMF System ID (e.g, CPPR)



DB2 SUBSYSTEM

Subsystem Overview	7-3
DB2 Workload and Performance Reports	7-3
DB2 Subsystem-Wide Reports	7-3
DB2 Connect Name Reports–by Plan	7-4
AUTH-ID Reports	7-9
DB2 Ad Hoc Reports	7-13
DB2 Trends Reports	7-16
Data Reduction	7-17
Processing SMF Records	7-18
Ad Hoc Request	7-19
Data Reduction JCL	7-20
CIMS Capacity Planner DB2 Data Reduction Parameters	7-21
DB2 Reporting	7-24
DB2 Reporting DD Statements	7-24
DB2 Reporting Parameters	7-26
Report Job Streams	7-36
DB2 ISPF/PDF INTERFACE	7-38
On-line Reports	7-38
Additional Report Options	7-39
GENERIC ELEMENT MASK	7-40
INCLUSIVE ELEMENT MASK	7-40
EXCLUSIVE ELEMENT MASK	7-41
Presentation Graphics	7-41
Graph Production Procedures	7-42
Naming Conventions	7-42
Presentation Graphics Distribution Disk	7-44
Graph Periods	7-45

Graph Descriptions	7-45
DB2 Presentation Graphics JCL	7-45
Graphic Report Tables	7-46

Subsystem Overview

The CIMS Capacity Planner DB2 Subsystem provides a variety of reports and graphs related to the workload and the performance of your DB2 system and applications. The DB2 reports provide both detail and summary level information designed to meet the needs of both the technical staff and technical management. A variety of graphs are provided through both GDDM and the PC based Harvard Graphics system. GDDM provides the capability to instantly display a number of DB2 graphs on the Host through the use of the CIMS Capacity Planner ISPF interface. The Harvard Graphics system is used primarily to produce professional quality colored graphs that can be used for presentations to users and management.

DB2 Workload and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner DB2 reports and describes the contents of each report.

DB2 Subsystem-Wide Reports

DB2 Subsystem Summary Analysis Report

The DB2 Subsystem Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 workload and how well the DB2 system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the DB2 system.
- The average number of threads processed during the prime shift
- The average number of successful commits that took place during the prime shift
- The average percent of threads that required a create thread during the prime shift
- The average percentage of create threads that were delayed by queuing during the prime shift
- The average number of times that a short-on-storage condition arose during the prime shift
- The average percentage of LOAD CT SECTION requests that required DASD access during the prime shift
- The average percentage of LOAD DBD requests that required DASD access during the prime shift
- The low watermark percentage of EDM POOL PAGES that were free during the prime shift

- The average number of BUFFER POOL expansions that took place during the prime shift
- The number of deferred write conditions that arose during the prime shift
- The number of times that Data Manager Critical Threshold Conditions arose during the prime shift
- The percentage of Sequential Prefetch operations that resulted in Prefetch Disabled during the prime shift
- The percent of the time that Get Page requests resulted in physical reads during the prime shift
- The percent of the time that System Page Writes resulted in physical writes during the prime shift
- The average number of EXCPs per second during the prime shift
- The number of Write Active Log Buffers that took place during prime shift
- The number of SQL calls that were processed by shift for: INSERT, SELECT, UPDATE, DELETE, FETCH, and OTHER call types.

DB2 Connect Name Reports—by Plan

DB2 Connect Name Summary Analysis Report

The DB2 Connect Name Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 connected applications workload and how well the DB2 system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the DB2 system.
- The average number of threads per minute that were processed during the prime shift
- The average number of commits per thread during the prime shift
- The average percent of threads that required a create thread during the prime shift
- The average number of SQL commands per thread during prime shift
- The average transit time for DB2 threads during the prime shift
- The average number of SQL INSERTS per minute during the prime shift
- The average number of SQL SELECTS per minute during the prime shift
- The average number of SQL UPDATES per minute during the prime shift

- The average number of SQL DELETES per minute during the prime shift
- The average number of SQL FETCHES per minute during the prime shift
- The average number of OTHER SQL calls per minute during the prime shift
- The average percent of the CPU utilized by each DB2 Connect Name by shift throughout the measured period
- The number of DB2 threads processed by shift for each Connect Name throughout the measured period

DB2 Transit Time Performance Report—by Plan

The DB2 Transit Time Performance Report shows, by shift, for a selected (INCLUDED/EXCLUDED) set of Plans within a specified DB2 connect name, how the DB2 threads flowed through the system. The 25 most commonly executed threads are listed individually with the percentages of the time that the threads were processed within a set of time bounds defined by the entries contained within the CPPR.PARMLIB member &sidDB2R. The threads that are not among the 25 most common are summarized under "OTHER" and reported as a single line on the report. The counts of each of the threads are also reported. The transit times are defined in seconds and tenths of seconds. Specific threads that might be of interest, but not reported individually (part of "OTHER") can be reported using the INCLUDE facility, which allows you to specify which threads are to be reported. For information related to specifying the times in the CPPR.PARMLIB member &sidDB2R, please refer to the Installation Instructions in the *CIMS Capacity Planner Installation Guide*.

DB2 Plan Statistics Report

The DB2 Plan Statistics Report shows, for a specified DB2 region, which plans were among the top 40 plans in terms of volume and resource consumption based upon the following categories:

- Most frequently executed

The plan name and the number of threads executed are shown along with the percentage of all activity for each of the 40 most commonly executed plans.
- Transit Time

The plan name and the total transit time is shown along with the percentage of all transit time for the plans requiring the most total transit time.
- CPU time

The plan name and the total CPU time is shown along with the percentage of all CPU time for the plans requiring the most total CPU time.
- Selects

The plan name and the total number of Selects issued is shown along with the percentage of all Selects for the plans issuing the most total Selects.

■ Inserts

The plan name and the total number of Inserts is shown along with the percentage of all Inserts for the plans requiring the most total Inserts.

■ Updates

The plan name and the total number of Updates is shown along with the percentage of all Updates for the plans requiring the most total Updates.

■ Deletes

The plan name and the total number of Deletes is shown along with the percentage of all Deletes for the plans requiring the most total Deletes.

■ Fetches

The plan name and the total number of Fetches is shown along with the percentage of all Fetches for the plans requiring the most total Fetches.

■ Other

The Plan name and the total number of OTHER DB2 commands is shown along with the percentage of all other commands for the Plans issuing the most commands not itemized above.

In the event the activity for any given Plan is not reported among the top 40, the data is available through the use of the INCLUDE facility. Similarly, you can use the EXCLUDE facility to exclude selected Plans from the report.

DB2 Exception Analysis

The DB2 Exception Analysis Report shows, for each 15-minute period, for a specified DB2 Connect Name, all threads whose transit times exceeded the thresholds specified in the CIMS Capacity Planner Parmlib member &sidXCPT.

DB2 Thread Transit Time Graph

The DB2 Thread Transit Time Graph shows, by day of the week, the average and peak average thread processing times for each 15-minute period throughout the day. The report page is broken into four sections - each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the average transit time (both average and peak average). In the middle of each line, both the average and peak average values are printed. The lines, in bar graph format, form the graphic representation of the transit time values. The peak average values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak values and *the average values are either equal or very nearly so. The unit of measure represented by the values along horizontal axis is specified in the third line of the report title.

You can produce the DB2 Transit Time Graph for a single day's activity or for several days. When more than one day's activity is graphed, separate graphs are produced—one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

You can control the data used to construct the DB2 Transit Time Graph using the INCLUDE/EXCLUDE facility. You can exclude data related to specific plan names from the graph using EXCLUDE facility. (All plan data except the plans that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of plans .

DB2 Commit Activity Graph

The DB2 Commit Activity Graph is produced in the same format as the Thread Transit Time Graph described above, except that the horizontal axis represents the number of commits per minute (both average and peak average) processed during each 15-minute period throughout the day.

You can control the data used to construct the DB2 Commit Activity Graph using INCLUDE/EXCLUDE facility. You can use the EXCLUDE facility to exclude data related to specific plan names from the graph. (All plan data except the plans that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of plans.

DB2 SQL Command Activity Graph

The DB2 SQL Command Activity Graph is produced in the same format as the Thread Transit Time Graph described above, except that the horizontal axis represents the number SQL commands processed per minute (both average and peak average) processed during each 15-minute period throughout the day.

You can control the data used to construct the DB2 SQL Command Activity Graph through the INCLUDE/EXCLUDE facility. You can use the EXCLUDE facility to specify that data related to specific plan names can be excluded from the graph. (All plan data except the plans that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of plans.

DB2 Plan Profile

The DB2 Plan Profile graphically depicts the daily average values of selected characteristics of a specified plan operating under a specified SYSID. The characteristics that are reported are:

■ Threads processed per day

The number of threads processed per day represents the number of threads processed during the period specified by the BEGIN DATE and the END DATE divided by the number of days contained within the specified period.

■ Transit time

The average transit time for all threads processed during the specified time period.

■ CPU time

The average number of CPU seconds required to process the threads during the specified period.

■ Commits per thread

The average number of COMMITs per thread for all threads processed during the specified period.

■ SQL Inserts, Selects, Updates, Deletes, and Fetches

The average number of each type of principal SQL command issued per thread processed throughout the specified period.

■ SQL Other

The average number of Other SQL commands (not specified above) issued during the processing of each thread during the specified period.

■ Maximum Page Locks

The maximum Page Locks held during the processing of each thread.

AUTH-ID Reports

DB2 AUTH-ID Summary Analysis Report

The DB2 AUTH-ID Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 connected applications workload and how well the DB2 system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the DB2 system.
- The average number of threads per minute that were processed during the prime shift
- The average number of commits per thread during the prime shift
- The average percent of threads that required a create thread during the prime shift
- The average number of SQL commands per thread during prime shift
- The average transit time for DB2 threads during the prime shift
- The average number of SQL INSERTS per minute during the prime shift
- The average number of SQL SELECTS per minute during the prime shift
- The average number of SQL UPDATES per minute during the prime shift
- The average number of SQL DELETES per minute during the prime shift
- The average number of SQL FETCHES per minute during the prime shift
- The average number of OTHER SQL calls per minute during the prime shift
- The average percent of the CPU utilized by each DB2 AUTH-ID by shift throughout the measured period
- The number of DB2 threads processed by shift for each AUTH-ID throughout the measured period

DB2 Transit Time Performance Report—by AUTH-ID

The DB2 Transit Time Performance Report shows, by shift, for a selected (INCLUDED/EXCLUDED) set of AUTH-IDs within a specified DB2 connect name, how the DB2 threads flowed through the system. The 25 most Active AUTH-IDs are listed individually with the percentages of the time that the threads were processed within a set of time bounds defined by the entries contained within the `CPPR.PARMLIB` member `&sidDB2R`. The AUTH-IDs that are not among the 25 most active are summarized under "OTHER" and reported as a single line on the report. The counts of each of the AUTH-IDs are also reported. The transit times are defined in seconds and tenths of seconds. Specific AUTH-IDs that might be of interest, but not reported individually (part of "OTHER") can be reported using INCLUDE facility, which allows you to specify which AUTH-IDs are to be reported. For information related to specifying the times in the `CPPR.PARMLIB` member `&sidDB2R`, please refer to the Installation Instructions in the *CIMS Capacity Planner Installation Guide*.

DB2 AUTH-ID Statistics Report

The DB2 AUTH-ID Statistics Report shows, for a specified DB2 region, which AUTH-IDs were among the top 40 AUTH-IDs in terms of volume and resource consumption based upon the following categories:

- Most frequently executed

The AUTH-ID and the number of threads executed are shown along with the percentage of all activity for each of the 40 most active AUTH-IDs.

- Transit Time

The AUTH-ID and the total transit time is shown along with the percentage of all transit time for the AUTH-IDs requiring the most total transit time.

- CPU time

The AUTH-ID and the total CPU time is shown along with the percentage of all CPU time for the AUTH-IDs requiring the most total CPU time.

- Selects

The AUTH-ID and the total number of Selects issued is shown along with the percentage of all Selects for the AUTH-IDs issuing the most total Selects.

- Inserts

The AUTH-ID and the total number of Inserts is shown along with the percentage of all Inserts for the AUTH-IDs requiring the most total Inserts.

- Updates

The AUTH-ID and the total number of Updates is shown along with the percentage of all Updates for the AUTH-IDs requiring the most total Updates.

- Deletes

The AUTH-ID and the total number of Deletes is shown along with the percentage of all Deletes for the AUTH-IDs requiring the most total Deletes.

- Fetches

The AUTH-ID and the total number of Fetches is shown along with the percentage of all Fetches for the AUTH-IDs requiring the most total Fetches.

- Other

The AUTH-ID and the total number of OTHER DB2 commands is shown along with the percentage of all OTHER commands for the AUTH-IDs issuing the most commands not itemized above.

In the event the activity for any given AUTH-ID is not reported among the top 40, the data is available through the use of the INCLUDE facility. Similarly, you can use the EXCLUDE facility to exclude selected AUTH-IDs from the report.

DB2 Exception Report

The Exception Analysis shows, for each 15-minute period, for a specified DB2 Connect Name, all threads whose transit times exceeded the thresholds specified in the CPPR.PARMLIB member &sidXCPT.

DB2 Thread Transit Time Graph

The DB2 Thread Transit Time Graph shows, by day of the week, the average and peak average thread processing times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the average transit time (both average and peak average). In the middle of each line, both the average and peak average values are printed. The lines, in bar graph format, form the graphic representation of the transit time values. The peak average values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak values and the average values are either equal or very nearly so. The unit of measure represented by the values along horizontal axis is specified in the third line of the report title.

You can produce the DB2 Transit Time Graph for a single day's activity or for several days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

You can control the data used to construct the DB2 Transit Time Graph through the INCLUDE/EXCLUDE facility. You can use the EXCLUDE facility to specify that data related to specific plan Auth-ids can be excluded from the graph. (All Auth-id data except the Auth-ids that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of Auth-ids.

DB2 Commit Activity Graph

The DB2 Commit Activity Graph is produced in the same format as the as the Thread Transit Time Graph described above, except that the horizontal axis represents the number of commits per minute (both average and peak average) processed during each 15-minute period throughout the day.

You can control the data used to construct the DB2 Commit Activity Graph through the INCLUDE/EXCLUDE facility. You can use the EXCLUDE facility to specify that data related to specific plan Auth-ids can be excluded from the graph. (All Auth-id data except the Auth-ids that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of Auth-ids.

DB2 SQL Command Activity Graph

The DB2 SQL Command Activity Graph is produced in the same format as the as the Thread Transit Time Graph described above, except that the horizontal axis represents the number SQL commands processed per minute (both average and peak average) processed during each 15-minute period throughout the day.

You can control the data used to construct the DB2 SQL Command Activity Graph through the INCLUDE/EXCLUDE facility. You can use the EXCLUDE facility to specify that data related to specific plan Auth-ids is excluded from the graph. (All Auth-id data except the Auth-ids that are excluded will be reflected in the graph). Also, you can use the INCLUDE facility to specify that the graph reflect data limited to a specific set of Auth-ids.

DB2 AUTH-ID Profile

The DB2 AUTH-ID Profile graphically depicts the daily average values of selected characteristics of a specified AUTH-ID operating under a specified SYSID. The characteristics that are reported are:

- **Threads processed per day**

The number of threads processed per day represents the number of threads processed during the period specified by the BEGIN DATE and the END DATE divided by the number of days contained within the specified period.

- **Transit time**

The average transit time for all threads processed during the specified time period.

- **CPU time**

The average number of CPU seconds required to process the threads during the specified period.

- **Commits per thread**

The average number of COMMITs per thread for all threads processed during the specified period.

- SQL Inserts, Selects, Updates, Deletes, and Fetches

The average number of each type of principal SQL command issued per thread processed throughout the specified period.

- SQL Other

The average number of Other SQL commands (not specified above) issued during the processing of each thread during the specified period.

- Maximum Page Locks

The maximum number of Page Locks held by this AUTH-ID during thread processing.

DB2 Ad Hoc Reports

The DB2 Subsystem gives you the option of producing three separate Ad Hoc reports during the data reduction process. The Ad Hoc reports contain detail data that you can summarize or discard as you maintain the CIMS Capacity Planner Performance Data Base.

The Ad Hoc reports are generated by the data reduction module of the DB2 subsystem (SSA1DB2W). The Ad Hoc reports should be generated on an as-needed basis and it is therefore likely that by the time the need is determined, the DB2 data has already been reduced and stored in the Performance Data Base. In such a case, it is imperative that the data not be stored into the Performance Data Base a second time. The possibility of updating the Performance Data Base a second time can be eliminated by specifying DD DUMMY for both the INDEX and the ONLINE DD statements.

Three separate DB2 Ad Hoc reports are available:

DB2 System Wide Ad Hoc Report

The DB2 System Wide Ad Hoc Report is selected by including a report request in the SYSIN Parameters in the DB2 Data Reduction JCL. For example, by specifying the following in the SYSIN file:

```
DB2 SYSTEM WIDE AD HOC REPORT=YES
```

```
DB2 SUBSYSTEM NAME=DB2
```

```
BEGIN DATE=06/02/2005
```

```
BEGIN TIME=08.00.00
```

```
END DATE=06/07/2005
```

```
END TIME=10.00.00
```

A system Wide Ad Hoc report would be produced for the DB2 Subsystem named DB2 for the period beginning on June 2nd at 8:00 AM and ending on June 7th at 10:00 AM.

For each record selected, the following data will be included in the report for each input record selected for inclusion in the report:

- The DB2 Subsystem name
- The Date from the time-stamp
- The time from the time-stamp
- The number of threads processed
- The number of commits processed
- The number of threads created
- The number of SQL commands processed
- The number of GET PAGEs
- The number of READ I/Os
- The number of UPDATE PAGEs
- The number of UPDATE I/Os

The JCL required to invoke the Ad Hoc reports is illustrated under the DB2 Data Reduction discussion.

DB2 Plan Name Ad Hoc Report

The report is selected by specifying a DB2 Connect Name through the use of the DB2NAME parameter and a DB2 AD HOC Plan Name through the use of the DB2 AD HOC PLAN NAME parameter. Both of these parameters, along with the specified date and time period are specified in the SYSIN file. For example,

```
//SYSIN DD *  
DB2NAME=CICSPROD  
DB2 AD HOC PLAN NAME=XYZ*  
BEGIN DATE=06/02/2005  
BEGIN TIME=08.00.00  
END DATE=06/07/2005  
BEGIN TIME=10.00.00
```

would result in a DB2 PLAN Ad Hoc Report being produced for the DB2 Connect Name CICSPROD for all Plans beginning with the characters XYZ and would include data for all records cut within the period beginning on June 2nd at 8:00 AM and ending on June 7th at 10.00.00 AM.

The Data contained in both the Plan Name and AUTH-ID Ad Hoc reports are:

- The SMF System ID
- The DB2 Connect Name
- The DB2 Plan Name
- The DB2 AUTH-ID
- The date the record was written
- The time the record was written
- The number of commits processed (this thread)
- The number of SQL commands processed (this thread)
- The CPU time required (this thread)
- The transit time (this thread)

DB2 AUTH-ID Ad Hoc Report

The contents of the AUTH-ID Ad Hoc Report are identical to the Plan Name AD Hoc report except that the data are reported by AUTH-ID. Please see [DB2 Plan Name Ad Hoc Report on page 7-14](#) above for a description of the report. The report is selected by specifying a DB2 Connect Name through the use of the DB2NAME parameter and a DB2 AUTH-ID through the use of the DB2 AD HOC AUTH-ID parameter. Both of these parameters, along with the specified date and time period are specified in the SYSIN file. For example,

```
//SYSIN DD *  
DB2NAME=CICSPROD  
DB2 AD HOC AUTH-ID=Q10*  
BEGIN DATE=06/02/2005  
BEGIN TIME=08.00.00  
END DATE=06/07/2005  
BEGIN TIME=10.00.00
```

would result in a DB2 Plan Ad Hoc Report being produced for the DB2 Connect Name CICSPROD for all plans whose name begins with the characters Q10 and would include data for all records cut within the period beginning on June 2nd at 8:00 AM and ending on June 7th at 10.00.00 AM.

DB2 Trends Reports

The CIMS Capacity Planner approach to DB2 Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. Two separate Trends Analysis Summary programs are available, one for DB2 Subsystems (SSA1DB2T) and the other for DB2 Connect Names (SSA1DB2U). Two Trends Analysis Graph Interface programs are also available, one for DB2 Subsystems (SSA1DB2G) and one for DB2 Connect Names (SSA1DB2H).

The elements that are compared are as follows:

- CPU Utilization Statistics
- The number of threads executed per minute computed as an average during Prime shift
- The number of commits executed per minute computed as an average during Prime shift
- The number of SQL commands executed per minute computed as an average during Prime shift
- The average number of EXCPs per second executed by DB2 during Prime shift (Subsystem only)
- The average thread transit time during prime shift (Connect name only)

The above information is gathered and presented to you in two ways:

- In summary report format, showing the values of each of the elements cited above both during a baseline period and a secondary period with a comparison of the baseline and secondary periods, with the slope of each comparison indicated at the right side of the report.
- In data suitable for graphing using the Harvard Graphics Program once the data points have been down-loaded to an IBM compatible Personal Computer capable of running Harvard Graphics. The manner in which the data is down-loaded is left up to you. The data points are written into a partitioned data set (HGDLIB) to await down-loading.

Data Reduction

The DB2 data reduction program (SSA1DB2W) accepts input from the following:

- SMF Type 100 data that is extracted from the SMF/RMF input to the Workload Analysis Data Reduction program (SSA1WKLD) by way of the SYSUT2 and FILTER Parameter or that comes directly from the SMF/RMF cluster or a historical repository.
- SMF Type 101 data that is extracted from the SMF/RMF input to the Workload Analysis Data Reduction program (SSA1WKLD) by way of the SYSUT2 and FILTER Parameter or that comes directly from the SMF/RMF cluster or a historical repository.

The data is reduced into a set of Consolidated Tables that reside in the CIMS Capacity Planner Performance Data Base. The Performance Data Base is stored on DASD in the data set specified by the //ONLINE data set. These Consolidated Tables are used later by the various report generation programs in preparing the CIMS Capacity Planner reports, graphs, and data extractions.

All data must be presented to the CIMS Capacity Planner DB2 Data Reduction program in chronological order. CIMS Capacity Planner maintains a table showing the date and time of the last transaction processed. Should an input record be time stamped with a date and time earlier than the last transaction processed, it is considered a duplicate and rejected.

During the first pass through the SMF data, the type 100 records (statistics) for each registered DB2 subsystem are processed and the type 101 records (accounting) for each of the DB2 Connect Names are written to a spill file for later processing in Connect Name order.

Processing SMF Records

If DB2 monitoring records are written to the SMF/RMF data set, the CIMS Capacity Planner Workload Data Reduction program (SSA1WKLD) pass the SMF Type 100 and 101 records through to the DB2 data reduction program (SSA1DB2W) via the //SYSUT2 data set. Alternatively, the entire SMF file can be processed by the DB2 data reduction program. An example of a job stream that can be used to pass the Type 100 and 101 records through the Workload Subsystem data reduction program to the DB2 data reduction program via the //SYSUT2 data set is shown below.

Note • If you are using DB2 6 or later, you no longer need to supply the GMT OFFSET. The adjustment for GMT will be received dynamically from the DB2 100 and 101 SMF records.

```
//jobname JOB (.....)
//*****
//*   THIS STEP WRITES THE TYPE 100,101 DATA TO A SEQUENTIAL
//*   DATA SET AND PASSES IT ON TO THE DB2 DATA REDUCTION
//*   PROGRAM WHILE ALSO PERFORMING THE PROCESSING FOR THE
//*   WORKLOAD DATA REDUCTION.
//*****
//ST1      EXEC PGM=SSA1WKLD,REGION=4096K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOAD
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE(WKLD)
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX(WKLD)
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1   DD DISP=OLD,DSN=SMF.INPUT.FILE
//SYSUT2   DD DISP=(,PASS),DSN=CPPR.DB2.INPUT
//SYSUT3   DD SPACE=(CYL,(10,10)),UNIT=SYSDA,DISP=(,PASS)
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=CPPR
SYSUT2=YES
FILTER=100,101
//*****
//*   THIS STEP PERFORMS THE DB2 DATA REDUCTION PROCESSING
//*****
//ST2      EXEC PGM=SSA1DB2W,REGION=4096K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOAD
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE(DB2)
```

```
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX(DB2)
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1 DD DISP=PASS,DSN=CPPR.DB2.INPUT
//SYSUT3 DD SPACE=(CYL,(10,10)),UNIT=SYSDA,DISP=(,PASS)
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR
```

* The following presumes Local time is 6 hours west of GMT

```
GMT OFFSET=W,6
```

Ad Hoc Request

The following JCL is suggested to request a DB2 Ad Hoc report:

```
//jobname JOB (.....)
//*****
//* THIS STEP PERFORMS THE DB2 AD HOC PROCESSING
//*****
//ST2 EXEC PGM=SSA1DB2W,REGION=4096K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOAD
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE DD DUMMY
//INDEX DD DUMMY
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSUT1 DD DISP=PASS,DSN=CPPR.DB2.INPUT
//SYSUT3 DD SPACE=(CYL,(10,10)),UNIT=SYSDA,DISP=(,PASS)
//SSASPILL DD DISP=(,PASS),SPACE=(CYL,(30,8))
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR
DB2 SYSTEM NAME=DB2
BEGIN DATE=06/02/2005
BEGIN TIME=08.00.00
END DATE=06/07/2005
END TIME=10.00.00
DB2 SYSTEM WIDE AD HOC REPORT=YES
```

The above sample JCL would produce a DB2 System Wide Ad Hoc report.

Data Reduction JCL

The JCL has been designed to conform to IBM system utility conventions to the extent possible. The DD statements required to execute the DB2 data reduction are described below.

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all DB2 Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ SYSUT1

The SYSUT1 DD statement specifies the input to the CIMS Capacity Planner Data Reduction Program—SSA1DB2W.

■ SYSUT3

The SYSUT3 DD statement is used to specify a temporary data set used to hold a working copy of the Performance Data Base during the data reduction process. It is no longer needed after the performance data is copied back to the ONLINE file at the conclusion of the data reduction phase. The parameters that must be specified are: Disposition, Unit, and Space.

■ SSASPILL

The SSASPILL DD statement is used to specify a temporary data set used to hold type 101 SMF records during the processing of the SMF type 100 records and during the repetitive passing of the type 101 data during the data reduction processing. The parameters that must be specified are the Disposition, Unit, and Space.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the DB2 Ad Hoc report. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input parameter data to CIMS Capacity Planner. The parameters that relate to the DB2 data reduction phase are described in the next paragraph.

CIMS Capacity Planner DB2 Data Reduction Parameters

The CIMS Capacity Planner Parameters permit you to specify certain run-time options and to provide selected data required by the data reduction program in processing the DB2 history files and producing the Ad Hoc report.

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMFSID of the MVS system for which data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID can be specified by this parameter in running the DB2 data reduction program.

Options:

The SELECTED SYSTEM= parameter can be specified in several ways:

- The SID itself (e.g. SELECTED SYSTEM=SYS1)
- An "*", indicating the SID for the system on which the program is being executed (e.g. SELECTED SYSTEM=*)

Default Option:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the earliest date (inclusive) for which data is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which a DB2 record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in *BEGIN DATE on page 7-22*

Default Date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

FORCE DB2 INPUT

The FORCE DB2 INPUT parameter governs whether the dates of the DB2 history file are checked to avoid detecting duplicate updates. It is useful to process logs that have been skipped for one reason or another.

Options:

FORCE DB2 INPUT=YES

FORCE DB2 INPUT=NO

Default Option:

FORCE DB2 INPUT=NO

SMFILE

The SMFILE parameter is used to specify the type of input that is to be processed by the DB2 data reduction program. Input can come from one of the following three sources:

- 1- HISTORY (used for SMF records from a QSAM)
- 2- LIVE (used for SMF data from VSAM SMF Clusters)
- 3- ACTIVE (used for SMF data from Active SMF Cluster)

Options:

SMFILE=HISTORY

SMFILE=LIVE

SMFILE=ACTIVE

Default Option:

SMFILE=HISTORY

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

DB2 Reporting

With the exception of the DB2 Ad Hoc Report, all the DB2 batch reports are produced by one of the following report generation programs:

SSA1DB2R—the DB2 Report Generation Program

SSA1DB2T—the DB2 Trends Analysis Subsystem Summary Program

SSA1DB2U—the DB2 Trends Analysis Connect Name Summary Program

SSA1DB2G—the DB2 Trends Analysis Graph Program

SSA1DB2H—the DB2 Connect Name Regular Hourly Graph Program

DB2 Reporting DD Statements

The DD Statements required to execute the above programs are described below:

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.⁵⁶

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all DB2 Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the Index to the ONLINE (the Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ HGDLIB

The HGDLIB DD statement specifies the Partitioned Data Set used to save the data points required to produce the various graphs using the Harvard Graphics subsystem. The data points contained in the members of the HGDLIB PDS are downloaded to a PC and matched up with a set of CIMS Capacity Planner graph templates that specify how the graphic data is to be presented. The parameters required are the Disposition and the Data Set Name.

■ **SYSPRINT**

The SYSPRINT DD statement specifies the data set that will contain the DB2 printed reports. Unless otherwise specified, the DCB characteristics are:

```
RECFM=FBA
```

```
BLKSIZE=133
```

```
LRECL=133
```

■ **SYSNAP**

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ **SYSUDUMP**

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ **EXCLUDE/INCLUDE**

This is the (optional) Element Exclusion file. It is used by the DB2 Report Program to force certain elements to be excluded from the reports or to force the inclusion of certain elements in the reports that would not normally be included. When you use the INCLUDE, only the specified elements are included in the reports.

The EXCLUDE/INCLUDE file must consist of 80 character records, 72 of which are interpreted. Columns 73 through 80 can be used for sequence numbers. Multiple elements can be entered in a single record, but they must be separated by commas. The format of the input statement follows:

- An element can be up to eight characters in length.
- Multiple elements can be included in a single record provided that they are separated by commas or blanks.
- Elements can be terminated by the wild card character "*", indicating that all elements that match up to the "wild card" are to be excluded/included.

For example, the following statement in the EXCLUDE file:

```
TSOP,CS*
```

would cause the report program SSA1DB2R to exclude any data related to the Plan named TSOP. It would also cause the exclusion of all Plans beginning with the characters "CS"

The use of the INCLUDE causes the report program to limit the contents of the report to only those elements specified in the INCLUDE data set.

■ **SYSIN**

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the DB2 data reduction phase are described in *DB2 Reporting Parameters on page 7-26* below.

DB2 Reporting Parameters

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SID of the system for which reports are to be produced or trends to be analyzed. This is the SMF identifier that is specified in the SMFPRMxx member of SYS1.PARMLIB.

Options:

The SELECTED SYSTEM parameter has the following options:

The SID itself (e.g., SELECTED SYSTEM=MVSA)

An "*", indicating the SID for the system on which the report program is being executed (e.g., SELECTED SYSTEM=*)

Default Option:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the beginning of the period for which data in the Performance Data Base is to be collected for analysis and reporting.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which a DB2 record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in *BEGIN DATE on page 7-26*

Default Date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Options:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Option:

REPORT LANGUAGE=ENGLISH

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should be used only when working with the CIMS Lab to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

PRIME SHIFT FIRST HOUR

The PRIME SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Prime Shift

Default Hour:

PRIME SHIFT FIRST HOUR=8

LATE SHIFT FIRST HOUR

The LATE SHIFT FIRST HOUR= parameter specifies the hour of the day on which the LATE Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Late Shift

Default Hour:

LATE SHIFT FIRST HOUR=16

LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter *does not cause* any data to be removed from the Performance Data Base.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

DB2NAME

The DB2NAME= parameter designates the DB2 region for which data is to be collected and reported.

Specify Name:

DB2NAME=any DB2 Connect Name (VTAM APPLID) that has been registered using the SSA1REGD program (see DB2 Installation Procedures in the *CIMS Capacity Planner Installation Guide*).

Default Name:

None

DB2 SYSTEM WIDE SUMMARY REPORT

This parameter is used to specify whether the DB2 System Wide Summary Report is to be produced.

Options:

DB2 SYSTEM WIDE SUMMARY REPORT=YES (produce the report)

DB2 SYSTEM WIDE SUMMARY REPORT=NO (don't produce the report)

Default Option:

DB2 SYSTEM WIDE SUMMARY REPORT=NO

DB2 CONNECTION:PLAN SUMMARY REPORT

This parameter specifies whether the DB2 Summary Report—by Plan is to be produced.

Options:

DB2 CONNECTION:PLAN SUMMARY REPORT=YES (produce the report)

DB2 CONNECTION:PLAN SUMMARY REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:PLAN SUMMARY REPORT=NO

DB2 CONNECTION:PLAN PERFORMANCE REPORT

This parameter specifies whether the Thread Transit Time Performance Report is to be produced.

Options:

DB2 CONNECTION:PLAN PERFORMANCE REPORT=YES (produce the report)

DB2 CONNECTION:PLAN PERFORMANCE REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:PLAN PERFORMANCE REPORT=NO

DB2 CONNECTION:PLAN STATISTICS REPORT

This parameter specifies whether the DB2 Plan Statistics Report is to be produced.

Options:

DB2 CONNECTION:PLAN STATISTICS REPORT=YES (produce the report)

DB2 CONNECTION:PLAN STATISTICS REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:PLAN STATISTICS REPORT=NO

DB2 CONNECTION:PLAN EXCEPTION REPORT

This parameter specifies whether the DB2 Exception Analysis Report - by Plan is to be produced.

Options:

DB2 CONNECTION:PLAN EXCEPTION ANALYSIS=YES (produce the report)

DB2 CONNECTION:PLAN EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

DB2 CONNECTION:PLAN EXCEPTION ANALYSIS=NO

DB2 CONNECTION:PLAN THREAD TRANSIT TIME GRAPH

This parameter specifies whether the Thread Transit Time Graph—by Plan is to be produced.

Options:

DB2 CONNECTION:PLAN THREAD TRANSIT TIME GRAPH=YES (produce the graph)

DB2 CONNECTION:PLAN THREAD TRANSIT TIME GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:PLAN THREAD TRANSIT TIME GRAPH=NO

DB2 CONNECTION:PLAN SQL ACTIVITY GRAPH

This parameter specifies whether the DB2 SQL Command Activity Graph is to be produced.

Options:

DB2 CONNECTION:PLAN SQL ACTIVITY GRAPH=YES (produce the graph)

DB2 CONNECTION:PLAN SQL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:PLAN SQL ACTIVITY GRAPH=NO

DB2 CONNECTION:PLAN COMMIT ACTIVITY GRAPH

This parameter specifies whether the DB2 Commit Activity Graph is to be produced.

Options:

DB2 CONNECTION:PLAN COMMIT ACTIVITYGRAPH=YES (produce the graph)

DB2 CONNECTION:PLAN COMMIT ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:PLAN COMMIT ACTIVITY GRAPH=NO (don't produce the graph)

DB2 PLAN PROFILE

This parameter specifies whether the DB2 Plan Profile is to be produced.

Options:

DB2 PLAN PROFILE=YES (produce the graph)

DB2 PLAN PROFILE=NO (don't produce the graph)

Default Option:

DB2 PLAN PROFILE=NO

DB2 CONNECTION:AUTH-ID SUMMARY REPORT

This parameter specifies whether the DB2 Summary Report - by AUTH-ID is to be produced.

Options:

DB2 CONNECTION:AUTH-ID SUMMARY REPORT=YES (produce the report)

DB2 CONNECTION:AUTH-ID SUMMARY REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:AUTH-ID SUMMARY REPORT=NO

DB2 CONNECTION:AUTH-ID PERFORMANCE REPORT

This parameter specifies whether the Thread Transit Time Performance Report is to be produced.

Options:

DB2 CONNECTION:AUTH-ID PERFORMANCE REPORT=YES (produce the report)

DB2 CONNECTION:AUTH-ID PERFORMANCE REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:AUTH-ID PERFORMANCE REPORT=NO

DB2 CONNECTION:AUTH-ID STATISTICS REPORT

This parameter specifies whether the DB2 AUTH-ID Statistics Report is to be produced.

Options:

DB2 CONNECTION:AUTH-ID STATISTICS REPORT=YES (produce the report)

DB2 CONNECTION:AUTH-ID STATISTICS REPORT=NO (don't produce the report)

Default Option:

DB2 CONNECTION:AUTH-ID STATISTICS REPORT=NO

DB2 CONNECTION:AUTH-ID EXCEPTION REPORT

This parameter specifies whether the DB2 Exception Analysis Report - by AUTH-ID is to be produced.

Options:

DB2 CONNECTION:AUTH-ID EXCEPTION ANALYSIS=YES (produce the report)

DB2 CONNECTION:AUTH-ID EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

DB2 CONNECTION:AUTH-ID EXCEPTION ANALYSIS=NO

DB2 CONNECTION:AUTH-ID THREAD TRANSIT TIME GRAPH

This parameter specifies whether the Thread Transit Time Graph - by AUTH-ID is to be produced.

Options:

DB2 CONNECTION:AUTH-ID THREAD TRANSIT TIME GRAPH=YES (produce the graph)

DB2 CONNECTION:AUTH-ID THREAD TRANSIT TIME GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:AUTH-ID THREAD TRANSIT TIME GRAPH=NO

DB2 CONNECTION:AUTH-ID SQL ACTIVITY GRAPH

This parameter specifies whether the DB2 SQL Command Activity Graph is to be produced.

Options:

DB2 CONNECTION:AUTH-ID SQL ACTIVITY GRAPH=YES (produce the graph)

DB2 CONNECTION:AUTH-ID SQL ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:AUTH-ID SQL ACTIVITY GRAPH=NO

DB2 CONNECTION:AUTH-ID COMMIT ACTIVITY GRAPH

This parameter specifies whether the DB2 Commit Activity Graph is to be produced.

Options:

DB2 CONNECTION:AUTH-ID COMMIT ACTIVITY GRAPH=YES (produce the graph)

DB2 CONNECTION:AUTH-ID COMMIT ACTIVITY GRAPH=NO (don't produce the graph)

Default Option:

DB2 CONNECTION:AUTH-ID COMMIT ACTIVITY GRAPH=NO (don't produce the graph)

DB2 AUTH-ID PROFILE

This parameter specifies whether the DB2 AUTH-ID Profile is to be produced.

Options:

DB2 AUTH-ID PROFILE=YES (produce the graph)

DB2 AUTH-ID PROFILE=NO (don't produce the graph)

Default Option:

DB2 AUTH-ID PROFILE=NO

1ST PERIOD BEGIN DATE

This parameter is used to specify the beginning of the first of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1DB2T or SSA1DB2U) or the beginning date for the entire period to be used for developing the Trends Analysis Graph (SSA1DB2G or SSA1DB2H) for which consolidated table information is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

1ST PERIOD END DATE

This parameter is used to specify the ending date of the first period for the Trends Analysis Summary Program (SSA1DB2T or SSA1DB2U) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1DB2G or SSA1DB2H).

Specify Date:

Specify a beginning date as in 1ST PERIOD BEGIN DATE.

Default Date:

1ST PERIOD END DATE=00001

2ND PERIOD BEGIN DATE

This parameter is used to specify the beginning of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1DB2T or SSA1DB2U)

Specify Date:

Specify a beginning date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD BEGIN DATE=00001

2ND PERIOD END DATE

This parameter is used to specify the end date of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1DB2T or SSA1DB2U)

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD END DATE=00001

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems. A new statement

TOP LABEL=YES

causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time:End as specified
- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

Another statement

FOOT1=up to 64 char footnote

can be used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that will be included on all reports where there is room.

FOOT2

A third statement

FOOT2=up to 64 char 2nd footnote

can be used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that will be included on all reports where there is room.

Report Job Streams

CIMS Capacity Planner DB2 Batch Reports

The standard CIMS Capacity Planner DB2 batch reports are all produced by the CIMS Capacity Planner DB2 Report Program - SSA1DB2R.

Following is an example of the JCL required to run the report program:

```
//SSAREPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1DB2R,REGION=4096K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.DB2
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.DB2
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
BEGIN DATE=01/25/2005
END DATE=01/25/2005
*
DB2 SYSTEM WIDE SUMMARY REPORT=YES
DB2 SUBSYSTEM NAME=DB2P
*
DB2NAME=**OTHER
DB2 CONNECTION:PLAN SUMMARY REPORT=YES
DB2 CONNECTION:PLAN PERFORMANCE REPORT=YES
DB2 CONNECTION:PLAN STATISTICS REPORT=YES
DB2 CONNECTION:PLAN EXCEPTION ANALYSIS=YES
```

```

DB2 CONNECTION:PLAN THREAD TRANSIT TIME GRAPH=YES
DB2 CONNECTION:PLAN THREAD THROUGHPUT GRAPH=YES
DB2 CONNECTION:PLAN COMMIT ACTIVITY GRAPH=YES
DB2 CONNECTION:PLAN SQL ACTIVITY GRAPH
DB2 PLAN PROFILE=YES
DB2 PLAN NAME=RTPPU001
*
DB2NAME=**OTHER
DB2 CONNECTION:AUTH-ID SUMMARY REPORT=YES
DB2 CONNECTION:AUTH-ID PERFORMANCE REPORT=YES
DB2 CONNECTION:AUTH-ID STATISTICS REPORT=YES
DB2 CONNECTION:AUTH-ID EXCEPTION ANALYSIS=YES
DB2 CONNECTION:AUTH-ID THREAD TRANSIT TIME GRAPH=YES
DB2 CONNECTION:AUTH-ID THREAD THROUGHPUT GRAPH=YES
DB2 CONNECTION:AUTH-ID COMMIT ACTIVITY GRAPH=YES
DB2 CONNECTION:AUTH-ID SQL ACTIVITY GRAPH=YES
DB2 AUTH-ID PROFILE=YES
DB2 AUTH-ID NAME=DPA6

```

The SYSIN parameters are used to request reports and to specify reporting options that are described above under *DB2 Reporting Parameters on page 7-26*.

Member DDB2REPT in the PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that can be customized for your installation.

DB2 Subsystem Trends Analysis

Following is an example of the JCL required to run the DB2 Subsystem Trends Report Program—SSA1DB2T.

```

//SSAREPT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1DB2T, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.DB2
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.DB2
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *

```

```
SELECTED SYSTEM=*  
IDMSNAME=IDMSPROD  
1ST PERIOD BEGIN DATE=02/16/2005  
1ST PERIOD END DATE=04/15/2005  
2ND PERIOD BEGIN DATE=05/16/2005  
2ND PERIOD END DATE=07/155/2005  
PRIME SHIFT FIRST HOUR=7  
LATE SHIFT FIRST HOUR=19  
LUNCH BREAK BEGIN HOUR=11  
LUNCH BREAK END HOUR=12
```

The SYSIN parameters used to request the trends reports and to specify the reporting options are described above under [DB2 Reporting Parameters on page 7-26](#).

Member DDB2TRND in the &PREFIX.CPPR.Vnnn.CNTL library contains a set of model JCL that you can customize for use in your installation.

DB2 ISPF/PDF INTERFACE

The CIMS Capacity Planner DB2 subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing several of the DB2 reports on-line.

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, DB2 etc.).

On-line Reports

The following reports can be produced on-line through the CPPR DB2 ISPF Interface.

DB2 Subsystem Summary Analysis Report

The DB2 Subsystem Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 workload and how well the DB2 system is performing.

A description of the contents of this report can be found under [DB2 Subsystem Summary Analysis Report on page 7-3](#) above.

DB2 Connect Name Summary Analysis Report

The DB2 Connect Name Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 connected applications workload and how well the DB2 system is performing.

A description of the contents of this report can be found under [DB2 Subsystem Summary Analysis Report on page 7-3](#) above.

DB2 AUTH-ID Summary Analysis Report

The DB2 AUTH-ID Summary Analysis Report contains key summary information required to determine the magnitude of the DB2 connected applications workload and how well the DB2 system is performing.

A description of the contents of this report can be found under [DB2 AUTH-ID Summary Analysis Report on page 7-9](#) above.

Additional Report Options

In addition to the INCLUDE/EXCLUDE options it is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it can be useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report can be prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is: type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an asterisk and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

GENERIC ELEMENT MASK=00000000 specifies that the key is to remain untouched, whereas the statement:

GENERIC ELEMENT MASK=11111111 specifies that the entire record key should be masked out (set to "*****"). Finally, the statement :

GENERIC ELEMENT MASK=10101010 would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE OR EXCLUDE facilities.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, we would merely specify a GENERIC ELEMENT MASK=1110011. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL*. If GENERIC ELEMENT MASK=11111111 were specified, all records would be summarized into a single entry.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the DB2 report processing programs to intercept the key of each element (in this case, the Plan Name or the Auth-ID), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "???23??. The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the DB2 report processing programs to intercept the key of each element (in this case, the Plan Name or Auth-ID), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *
      xxxxxxxx,xxxxxxxx,xxxxxxxx,xxxxxxxx,....
```

where xxxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "???23??. The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

Presentation Graphics

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the MVS system workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics—Release 3.0, plots each set of data points in accordance with the characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the workload, the performance, and the trends that are not highly technical are supported by the presentation graphics subsystem.

The Harvard Graphics templates, which describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and can be used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, can be used to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner Workload graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in *Graph Descriptions on page 7-45* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs.
- Down-load the data points from the host to the PC.
- Invoke the Harvard Graphics system on the PC.
- Select the proper graph template.
- Import the data points into the template.

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters:

&sidp*iiq*" where:

- "sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly

- "iiq" is the graph identifier and consists of two alpha-numeric characters. The "pii" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated under *EXCLUSIVE ELEMENT MASK on page 7-41*. The qualifier "q" might or might not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpii.q.DAT where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to down-load the data points from HGDLIB to the PC using IND\$FILE. The BAT files can be modified as required to fit the naming conventions of your installation.

The names of the BAT files all conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above
- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
 - "WKLD"—Workload graph data
 - "DASM"—DASM graph data

- "CICS"—CICS graph data
- "CICE"—CICS Extension graph data
- "IDMS"—IDMS graph data
- "IDME"—IDMS Extension graph data
- "IMSG"—IMS graph data
- "NETG"—Network graph data
- "NETE"—Network Extension graph data
- "NETL"—Network Physical Lines graph data
- "NETN"—Network Physical NCP graph data
- "DB2S"—DB2 Subsystem graph data
- "DB2U"—DB2 Connect ID graph data
- "M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M 204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB members from the Host processor. The naming conventions are as follows:

- The .BAT Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator (1=Hourly, 2=Daily, 3=Weekly, 4=Monthly) followed by a subsystem identifier (WKLD=Workload, DASM=DASM, CICS=CICS, CICE=CICS

Extension, IDMS=IDMS, IDME=IDMS Extension, NETG=Network, NETE=Network Extension, NETL=Network Physical Lines, NETN=Network Physical NCP). So, for example, the .BAT file to be used to download the HGDLIB members that would be used to produce the Workload Daily graphs would be named HGD2WKLD.BAT.

- The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator (1=Hourly, 2=Daily, 3=Weekly, 4=Monthly) followed by a 2 character graph identifier (see the tables below) followed by a 1 character qualifier in some cases (see the table below). Please refer to a specific subsystem document for Template names.

Graph Periods

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for a description of the hourly, daily, weekly, and monthly graph periods.

Graph Descriptions

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for descriptions of each of the graphs produced by the CIMS Capacity Planner DB2 Subsystem.

DB2 Presentation Graphics JCL

Following is an example of the JCL required to operate the SSA1DB2G program to generate the DB2 graph data points.

```
//SSAGRPH JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
//ST1 EXEC PGM=SSA1DB2G, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.INDEX.DB2
//ONLINE DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ONLINE.DB2
//CPPRERT DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//HGDLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD SYSOUT=*
SELECTED SYSTEM=CPPR*
DB2NAME=CICSPROD
```

```
BEGIN DATE=02/16/2005
END DATE=04/15/2005
PRIME SHIFT FIRST HOUR=7
LATE SHIFT FIRST HOUR=19
GRAPH PERIOD=WEEKLY
//
```

Several sets of sample Job streams for producing the CIMS Capacity Planner DB2 graphs are provided in the CIMS Capacity Planner CNTL library under the following member names:

- DDB2HGRF—Produce Hourly Graphs
- DDB2DGRF—Produce Daily Graphs
- DDB2WGRF—Produce Weekly Graphs
- DDB2MGRF—Produce Monthly Graphs

All four of the above Job streams contain the JCL required to execute the program - SSA1DB2G or SSA1DB2H To tailor the Job streams, merely modify the selection criteria as required.

Graphic Report Tables

The following tables summarize the relationship of members in the HGDLIB library, the Harvard Graphics templates, and supporting members of the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar graph program.

DB2 Subsystem Regular Hourly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2G	&sid1B0*	(none)	DB2 Threads /Hour
	&sid1B1*	(none)	DB2 Commits/Hour
	&sid1B2*	(none)	DB2 SQL Commands/Hour
	&sid1B3*	(none)	DB2 % CPU Busy
	&sid1B4*	(none)	DB2 EXCPs/Second/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Connect Name Regular Hourly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2H	&sid1BA*	(none)	DB2 Threads /Hour
	&sid1BB*	(none)	DB2 Commits/Hour
	&sid1BC*	(none)	DB2 SQL Commands/Hour
	&sid1BD*	(none)	DB2 Thread Transit Time
	&sid1BE*	(none)	DB2 % CPU Busy/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Subsystem Regular Daily Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2G	&sid2B0*	(none)	DB2 Threads /Hour
	&sid2B1*	(none)	DB2 Commits/Hour
	&sid2B2*	(none)	DB2 SQL Commands/Hour
	&sid2B3*	(none)	DB2 % CPU Busy
	&sid2B4*	(none)	DB2 EXCPs/Second/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Connect Name Regular Daily Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2H	&sid2BA*	(none)	DB2 Threads /Hour
	&sid2BB*	(none)	DB2 Commits/Hour
	&sid2BC*	(none)	DB2 SQL Commands/Hour
	&sid2BD*	(none)	DB2 Thread Transit Time
	&sid2BE*	(none)	DB2 % CPU Busy/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Subsystem Regular Weekly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2G	&sid3B0*	(none)	DB2 Threads /Hour
	&sid3B1*	(none)	DB2 Commits/Hour
	&sid3B2*	(none)	DB2 SQL Commands/Hour
	&sid3B3*	(none)	DB2 % CPU Busy
	&sid3B4*	(none)	DB2 EXCPs/Second/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Connect Name Regular Weekly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2H	&sid3BA*	(none)	DB2 Threads /Hour
	&sid3BB*	(none)	DB2 Commits/Hour
	&sid3BC*	(none)	DB2 SQL Commands/Hour
	&sid3BD*	(none)	DB2 Thread Transit Time
	&sid3BE*	(none)	DB2 % CPU Busy/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Subsystem Regular Monthly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2G	&sid4B0*	(none)	DB2 Threads /Hour
	&sid4B1*	(none)	DB2 Commits/Hour
	&sid4B2*	(none)	DB2 SQL Commands/Hour
	&sid4B3*	(none)	DB2 % CPU Busy
	&sid4B4*	(none)	DB2 EXCPs/Second/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

DB2 Connect Name Regular Monthly Graphs

These graphs represent specific DB2 subsystems executing within the &sid system. Therefore, DB2 subsystem suffixes are required to uniquely identify a particular DB2 subsystem.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1DB2H	&sid4BA*	(none)	DB2 Threads /Hour
	&sid4BB*	(none)	DB2 Commits/Hour
	&sid4BC*	(none)	DB2 SQL Commands/Hour
	&sid4BD*	(none)	DB2 Thread Transit Time
	&sid4BE*	(none)	DB2 % CPU Busy/Hour

* plus a one character suffix representing the DB2 subsystem

&sid is the SMF System ID (e.g, CPPR)

GENERIC SUBSYSTEM

Introduction	8-4
Components of the Generic Subsystem	8-5
Functional Description	8-8
Data Reduction—SSA1GENW	8-8
Report Generation	8-15
Generic Subsystem Operation	8-18
Data Reduction	8-18
User Table Specification	8-34
Ad Hoc Reports	8-38
Batch Reports	8-40
Generating Reports From CIMS Capacity Planner Maintained Tables	8-45
Generic Subsystem JCL (Data Reduction)	8-45
SYSPRINT	8-45
SYSUT1	8-45
CPPRERT	8-45
CPPRPARAM	8-46
INDEX	8-46
ONLINE	8-46
SYSNAP	8-46
SYSUT3	8-46
SYSIN	8-46
INPUT PARAMETER STATEMENTS FOR DATA REDUCTION	8-47
SELECTED SYSTEM	8-47
SMFILE	8-47
BEGIN DATE	8-48
END DATE	8-48
BEGIN TIME	8-48
END TIME	8-49

GENERIC ELEMENT KEY	8-49
GEOGRAPHIC LOCATION	8-49
LOCAL HOLIDAYS	8-50
DEBUGON	8-50
SAMPLE JOB STREAMS FOR DATA REDUCTION	8-50
OPERATING JCL FOR BATCH REPORTS	8-54
SYSPRINT	8-54
CPPRERT	8-55
CPPRPARM	8-55
INDEX	8-55
ONLINE	8-55
EXCLUDE/INCLUDE	8-55
SYSNAP	8-56
SYSIN	8-56
INPUT PARAMETER STATEMENTS FOR BATCH REPORTS	8-56
SELECTED SYSTEM	8-56
GENERIC STATISTICS REPORT	8-57
GENERIC DAILY OVERVIEW	8-57
GENERIC ELEMENT ACTIVITY LIST	8-57
BEGIN DATE	8-58
END DATE	8-58
BEGIN TIME	8-58
END TIME	8-59
PRIME SHIFT FIRST HOUR	8-59
LATE SHIFT FIRST HOUR	8-59
SELECTED DAY	8-59
GEOGRAPHIC LOCATION	8-60
DEBUGON	8-60
LOCAL HOLIDAYS	8-60
TABLEDEF STATEMENTS FOR BATCH REPORTS	8-61
USER TABLE NUMBER	8-61
OUTPUT ELEMENT CELL	8-61
SAMPLE JOB STREAMS FOR BATCH REPORTS	8-62
PRODUCING BASIC PC GRAPHICS DATA POINT MEMBERS	8-63
OPERATING JCL FOR BASIC GRAPH DATA POINTS	8-63
SYSPRINT	8-63
CPPRERT	8-64
CPPRPARM	8-64
INDEX	8-64

ONLINE8-64

HGDLIB8-64

EXCLUDE/INCLUDE8-64

SYSNAP8-65

SYSIN8-65

INPUT PARAMETER STATEMENTS FOR BASIC GRAPH DATA POINTS8-66

FILTER8-66

SELECTED SYSTEM8-66

BEGIN DATE8-67

END DATE8-67

PRIME SHIFT FIRST HOUR8-67

LATE SHIFT FIRST HOUR8-68

SELECTED DAY8-68

GEOGRAPHIC LOCATION8-68

DEBUGON8-68

LOCAL HOLIDAYS8-69

GRAPH PERIOD8-69

TABLEDEF Statements for Basic Graph Data Points8-69

USER TABLE NUMBER8-69

OUTPUT ELEMENT CELL8-70

SAMPLE JOBSTREAMS FOR BASIC CELL GRAPH DATA POINTS8-71

Introduction

Although the pre-programmed reports and graphs provided by CIMS Capacity Planner generally satisfy the majority of an installation's reporting requirements, there are cases where the pre-programmed reports must be supplemented with specialized reports that are unique to each installation. CIMS Capacity Planner addresses the preparation of specialized reports through the **Generic Subsystem**.

The Generic Subsystem supplements the other CIMS Capacity Planner Reporting Subsystems by providing the capability to:

- Reduce data that is otherwise unsupported into CIMS Capacity Planner User Tables and report on it

Software products such as ADABAS, COMPLETE, other program products, or in-house developed systems, which are not supported by CIMS Capacity Planner but which produce logs containing performance and system utilization data, can be supported through the use of the Generic Subsystem by reducing the log data and storing them in user-defined tables that reside in the CIMS Capacity Planner Performance Data Base. Reporting capabilities are provided so that you can produce both reports and graphs from the data contained in the tables.

The SMF and RMF records have many fields that are not used by CIMS Capacity Planner. The Generic Subsystem allows you to report on those fields as well. For example, the RMF Type 71 record (Paging Activity) has several fields that show how storage is allocated in terms of CSA, LPA, SQA, LSQA and Nucleus both in Central and Expanded storage. You can produce your own specialized reports based upon the paging data contained within the Type 71 records.

- Build reports from User Tables or standard CIMS Capacity Planner Tables, using values not normally combined in one report

The Generic Subsystem lets you extend the reporting capability of CIMS Capacity Planner by producing additional reports using the data contained in the standard CIMS Capacity Planner tables, by screening the data as desired and combining data in ways not supported by CIMS Capacity Planner.

For example, let's say you want to examine the activity of a specific channel during a specific period of time, but you can't find a report in the WKLD subsystem that allows you to see that level of detail. The Generic subsystem provides this capability.

- Build PC Graphs from standard CIMS Capacity Planner Tables, using values not normally combined in one graph

The Generic Subsystem provides the capability to produce User-defined graphs by producing a set of data points that you can down-load to Harvard Graphics or a similar package operating on a PC and subsequently graph using graph templates that you define.

- Create Ad Hoc reports from virtually any input source

Ad Hoc reports are produced at the time the input enters the data reduction phase. Data used in the preparation of the Ad Hoc reports need not be saved in the CIMS Capacity Planner User Tables.

For example, there is a great deal of valuable information in the SMF type 30 record, not all of which is captured by CIMS Capacity Planner. You might be interested in studying the number of specific and non-specific tape mounts resulting from the execution of a set of Jobs beginning with the three characters "XYZ". By specifying the Jobname field in the type 30 record as the key and selecting the specific and non-specific tape mount cells in the type 30 records, you can process the SMF history tapes and put together a very nice Ad Hoc report.

The Generic Subsystem, in addition to extending the CIMS Capacity Planner reporting capabilities, provides a mechanism to set up short term analytical projects focusing on very specific objectives.

Components of the Generic Subsystem

The Generic Subsystem consists of several components that let you specify the input, the data to be saved in the user-defined tables, the screens (filters) to be applied to the data, and the reports and graphs to be produced. The various components are briefly discussed below. They are covered in considerable detail in subsequent sections.

Data Reduction

The Data Reduction program reads the input records, determines whether the record meets the screening criteria you specified (filters) and, assuming the filters are satisfied, extracts the relevant data from the record for processing. The Generic Subsystem processes records from a variety of sources. This requires it to support a number of different input record formats. To accommodate the varied record formats, an input record description language is provided to describe the record types, the record organization, and the various fields within the records.

The Data Reduction program also constructs any Ad Hoc reports you specified. They are produced from selected data elements that come from the original input stream. An Ad Hoc report can consist of a record key, a date and time stamp, and up to eight columns of data, each with column headings of up to twelve characters.

The principal objective of the Generic Subsystem is to build CIMS Capacity Planner User Tables, preserving specific elements of data in a highly compressed and summarized form so that they can subsequently be used in producing reports and graphs without having to repeatedly process the original input.

Data in CIMS Capacity Planner User Tables are summarized at the hourly level within specified eight-character keys. For example, if you wanted to keep track of the number of DASD (data cell value) and TAPE (data cell value) EXCPs by JOBNAME (element key value) over a period of time, this data could be reduced and summarized by hour in a CIMS Capacity Planner User Table.

Note • Table numbers 128 through 159 have been reserved for CIMS Capacity Planner User Tables. It is important that this number range be used, as CIMS Capacity Planner User Tables have several special characteristics. Also, it prevents the user-specified tables from overlaying standard CIMS Capacity Planner Tables, should the user inadvertently store the User Tables in a WKLD database.

Values stored in a User table have the same general characteristics as those found in any of the standard CIMS Capacity Planner tables. That is, the table itself is identified by:

- A Date
- An SMF SID
- A Table Number (128-159 are reserved for User Tables).

The values stored in a table are associated with:

- An Element Key (for example, JOBNAME)
- A Time period during which the event took place
- A Cell number (to separate contiguous values)

In selecting the values to be entered into the User Tables or displayed in the Ad Hoc reports, the Generic Subsystem applies the tests of all Filters that have been specified in the SYSIN and TABLEDEF input streams. A Filter is simply a test of the form:

IF an element value is EQUAL to a specified value

THEN process the record;

ELSE discard the record.

There are two types of Filters:

- Record Filters

A Record Filter determines whether a Record is eligible for processing.

- Segment Filters

A Segment Filter is examined after all Record Filters have been satisfied. It is used to determine whether a specified Segment within a record is to be processed.

- //UEXALIAS

In the Generic Subsystem Data Reduction Step it is now possible to specify a unique name for a user exit through the use of the //UEXALIAS DD statement. This allows the user to have many different User Exits (e.g, SSA1UX00=MYEXIT01) in the same LOADLIB. For more information, refer to the CNTL Library member named DGENR200.

Report Generation

The Generic Subsystem produces reports from the original input data (Ad Hoc), Batch Reports from either the User Tables or the standard CIMS Capacity Planner tables or trends analysis data points that can be used to build graphs or can be imported into a spread sheet for further analysis.

Batch Reports Driven from the User Tables

The Batch reports produced by the SSA1GENR program use the CIMS Capacity Planner User Tables as a source of input. Date and Time selection as well as specific element key selection can be used at reporting time to control the scope of the report.

■ Batch Reports Driven from Standard CIMS Capacity Planner Tables

The Batch report processor, SSA1GENR, also generates user-specified reports from the standard CIMS Capacity Planner Tables. Date and Time selection as well as specific element key selection can be used at reporting time to limit the scope of the report.

Graph Generation

The Generic Subsystem Graphics production program produces sets of points that, when processed by a graphics program such as Harvard Graphics, produces presentation quality graphs.

■ Harvard Graphics Points Driven from User Tables

Data points used to produce user-defined graphs are produced using the Generic graph production program, SSA1GENG, which extract data from the User Tables. You can use Date and Time selection as well as specific element key selection at reporting time to limit the points included in the graphs.

■ Harvard Graphics Points Driven from Standard CIMS Capacity Planner Tables

Data points used to produce user-defined graphs are produced using the Generic graph production program, SSA1GENG, which also extract data from the standard CIMS Capacity Planner Tables. You can use Date and Time selection as well as specific element key selection at reporting time to limit the points included in the graphs.

Functional Description

This section contains a description of the components within the Generic Subsystem and the functions they perform.

Data Reduction—SSA1GENW

The Generic Data reduction Program (SSA1GENW) reads the input records, determines whether each record meets the screening criteria you specified (filters) and, assuming the filters are satisfied, maintains the User Tables and produces the Ad Hoc Reports. Input to the Generic Subsystem can come from virtually any machine readable source, so long as the input records are self contained (that is, do not rely on subsequent records to supplement the information contained in the record itself) and they are date and time stamped. Also, one of the fields in the input record (8 characters or less) must be selected as the key to the user table.

Following are several examples of input that users routinely process using the Generic Subsystem.

■ Local Monitors

Any locally developed or commercially available monitor that is not specifically supported by CIMS Capacity Planner is a candidate for the Generic Subsystem, so long as the records produced by the monitor conform to the specifications noted above [Data Reduction—SSA1GENW on page 8-8](#).

■ SMF/RMF Input Records

All forms of SMF/RMF records are supported, although some of the more esoteric formats (Type 64, for example) are a bit complex to define.

■ Time and Date Stamped Logs and Journals

Logs and Journals (CICS, IDMS) are also candidates for the Generic Subsystem. They typically contain both time and date stamp information and, although the stated purpose for many of these Audit Trails is to provide a methodology for backing out transactions, they frequently contain many fields that can be used to track performance.

Input Description Language

In order to process the input records, their format must be defined using an Input File Description Language. In describing an input record, it is very important to tell CIMS Capacity Planner the shape of the record. A record can be simply a string of fields (Header) or it can be a string of fields followed by several sets of related fields (Segments). Or it could be even more complex. The CIMS Capacity Planner accepts Floating Point Values in the input record.

■ Header Only

The Header, the simplest case, is a fixed length input record that contains a string of values. Each and every record contains the same set of values located in the same relative location within each record. This type of record can be considered to be a Header Only. That is, there are no subsequent dependent segments following the set of fields.

■ Header plus one or more Segments

The next case, in terms of complexity, is a header followed by one or more segments, each of which must be described somewhere within the header (length, count, offset). If the values you want to save in the user table fall within a segment and the segment must be located through the offset in the header, this is considered a segmented record. Of course, if all the values you want to process are in the Header portion of the record, there is no need to describe the segment.

■ RMF Type Input Records (Triplets)

The classic example of a Segmented record is the RMF record. The Header of the RMF record describes each of the segments within the record through a 6 byte Triplet (Offset, Count, Length) located at a fixed offset in the Header.

For example, the RMF Type 72 record contains a Header plus several types of Segments, each of which is described in the Header. The data pertinent to the Performance Group Numbers is contained in one of these segment types. The segment is repeated within a single RMF Type 72 record, once for each PGN in the system.

In order to capture the data in an RMF Type 72 record and relate it to a specific PGN, it is necessary to be able to locate the beginning of the segments and to process each segment in turn, using the PGN as the key to the User Table.

■ CICS (110) Input Records (Dictionary Records)

A third type of record is the CICS 110 SMF record. This type of record is described by a Dictionary record that precedes the actual records themselves and that identifies each of the fields contained in each logical record (length, type, identification). The problem is that the Dictionary record can change even within a single pass of the data. When the Dictionary record changes, the field offsets within subsequent records also change.

Rather than implement complex logic that takes all of this into account, the CIMS Capacity Planner Generic Subsystem lets you specify that the input consists of CICS 110 records (SMFILE=C110), and to key on the identifier in the Dictionary record to identify fields that are to be preserved during data reduction.

Input Filters

An input filter is a user-specified (screen) value against which a field within an input record is compared in order to determine whether the record is to be processed or skipped. If all records in an input file were to be processed, there would be no need for filters. This, however, is normally not the case. For example, consider the SMF records. The layout of the SMF records varies widely from the Type 0 (IPL) to the Type 74 (RMF IODEVICE), even though the Header is consistent for the first 16 or so bytes. That is, the length field (greenword), the time and date stamp, the SMF SID, and the Record Type are always in the same relative positions within the Header.

Unfortunately, before the Generic Subsystem can process a record, the layout of the record must be communicated to CIMS Capacity Planner. This is only possible by describing a specific record based on the value contained in the Record Number field as specified in the description of the Header. The Record Number, in this case, becomes a Filter.

Selecting Records Based On Header Values

The Filters in the Header are used to determine whether the record conforms to the layout description that follows, and whether the Generic Subsystem is to proceed with processing the record and/or its segment values.

- **Selecting Records based on SMF SID (SMF/RMF)**

SMF records, by definition, always have the four character SMF SID in the Header. This is the most common type of Filter. If this value does not match the SMF SID value as specified in the SYSIN parameter `SELECTED SYSTEM=`, the record is discarded.

- **Selecting Records by Record Number (SMF/RMF)**

This is the most important Filter for SMF/RMF records, as it is used to specify the record type to which the layout pertains. If this value does not match the value specified in the TABLEDEF statement `RECORD NUMBER=`, the record is discarded.

- **Selecting Records by a Header Type Qualifier**

Occasionally, it might be useful to specify a field in the Header that is to be used in addition to (or in place of) the record number. This Filter performs that task.

- **Selecting Records based on Subsystem Name**

Some types of input records have an embedded name (for example, VTAM APPLID) in the header that can be useful as a Filter. For example, you might want to examine Netview Type 39 records that only apply to the CICSTOR1 VTAM APPLID. In this case, the Subsystem Filter is used to make that distinction.

Selecting Records Based On Segment Values

Filters can also be applied to data contained within a record Segment. They are used to determine whether to process a specific segment. That is, the Segment Filter value is always checked before processing the current segment.

■ Segment Subtype

IF the value located through the RECORD SUBTYPE pointer does not match the value in the INPUT ELEMENT CELL then discard the segment (a Segment Subtype can be anywhere)

When processing SMF Type 30 records, for example, there are several Subtypes from which to choose. The format of the data contained within any given Subtype differs from the data contained within other Subtypes. The subtype is located in a segment apart from the header. In this case, CIMS Capacity Planner must know how to find and validate the Segment Subtype before processing the data contained within the segment.

■ Segment Subtype Qualifier

IF the value located through the RECORD SUBTYPE pointer does not match the value in the INPUT ELEMENT CELL then discard the segment (a Segment Subtype can be anywhere)

In the Subtype segment itself (the one selected for reduction) there can be a value field embedded within the segment that can be used as a Filter. For example, in the EXCP section of the SMF Type 30 record, Subtype 4, there is a Device Type field that, if x'20', indicates a DASD Device and, if x'80', indicates a Tape Device. This field can be used as a Segment Subtype Qualifier Filter.

Selecting Records Through a User Exit

The Generic Subsystem gives you four separate User Exits that can be used to implement Filters. If the User Exit routines pass back a non-zero return code in Register 15, the record is discarded. The CIMS Capacity Planner convention for a non-successful return is to place the value twelve in Register 15. Several sample User Exits are contained in the Distribution SOURCE Library. The four User Exits are described below.

■ When the Record is First Read (SSA1UX00)

When the record is first read, the record is passed to the 1st User Exit. Register 0 points to the beginning of the record. You can do anything to the contents of the record, including moving the fields around to different locations or comparing for a variety of values. At the conclusion of the exit processing, if the record is to be processed, set R15 to zero and return. If the record is to be discarded, set R15 to twelve and return.

■ When the Date Stamp is Processed (SSA1UX01)

The Generic Subsystem supports many forms of dates in an input record. Date and Time stamps are two essential fields in the input records. Should the input record contain a date that is not in one of the supported formats, you can write an exit to pass a valid date back to the Generic Subsystem. If you decide not to process the record at this point, set the return code in Register 15 to twelve before returning control to the Generic Subsystem. If the record is to be processed, set Register 15 to zero.

■ When the Time Stamp is Processed (SSA1UX02)

The Generic Subsystem supports many forms of time values in an input record. Date and Time stamps are two essential fields in the input records. Should the input record contain a time value that is not in one of the supported formats, the user can write an exit to pass a valid time value back to the Generic Subsystem. If you decides not to process the record at this point, set the return code in Register 15 to twelve prior to returning control to the Generic Subsystem. If the record is to be processed, set Register 15 to zero.

■ Just before the Cells are Updated (SSA1UX03)

You have one last chance to interrupt processing just before the values are stored in a user table. The final User Exit is entered just before the tables are updated. If you decide not to update the table, set Register 15 to twelve; otherwise, set Register 15 to zero to indicate that the table is to be updated.

Examining the Results of Filtering

Including the statement:

```
DUMP SMF STATISTICS=YES
```

in the //SYSIN data stream during the data reduction process lets you determine the number of records coming into the data reduction program and to see whether they passed the various filters. This is extremely useful when trying to narrow or expand the set of filters in place for a given run.

Electing to print the SMF statistics causes a matrix to be printed at the conclusion of the data reduction process that shows how many records were read and how many records passed each type of filter. The top row of the matrix reflects input filters, the second row reflects output filters. These elements of the matrix are explained below.

■ Row 1: Input Filter Elements

The 1st element in row 1 of the matrix is a count of the number of records read by the Input Processor.

The 2nd element in row 1 is a count of the number of records that successfully passed through User Exit 00 (SSA1UX00)

The 3rd element in row 1 is a count of the number of records that also matched the RECORD NUMBER as specified in the TABLEDEF input stream.

The 4th element in row 1 is a count of the number of records that also matched the RECORD HEADER TYPE QUALIFIER value as specified in the TABLEDEF input stream.

The 5th element in row 1 is a count of the number of records that also matched the SMF SID value as specified in the SELECTED SYSTEM key phrase in the SYSIN input stream.

The 6th element in row 1 is a count of the number of records that also matched the SUBSYSTEM ID as specified by the CICSNAME key phrase in the SYSIN input stream.

The 7th element in row 1 is a count of the number of record segments that also had at least one non-zero input cell value as specified in the INPUT ELEMENT CELL key phrases in the TABLEDEF input stream.

At this point, the input record is written to SYSUT2 if it has been so indicated in the SYSIN input stream.

Once a record segment has passed all Input filters, it is passed to the output processor where it is subjected to additional tests.

■ Row 2: Output Filter Elements

The 1st element in row 2 is a count of the number of record segments that produced one or more non-zero output cell values as specified in the OUTPUT ELEMENT CELL key phrases in the TABLEDEF input stream.

The 2nd element in row 2 is a count of the number of record segments that also had a valid date.

The 3rd element in row 2 is a count of the number of record segments that also passed successfully through User Exit 03 (SSA1UX03).

User Table Organization

The User Tables are stored in much the same format as the standard CIMS Capacity Planner tables. Each table record contains a header segment that contains the length of the table, a date, an SMF SID, and a table number. Following the header is the user-specified portion of the table that consists of a number of table element sets, each of which contains a key and 24 Cell Sets, one Cell Set corresponding to each hourly period throughout the day. A Cell Set consists of up to 99 full word values, depending upon the number of cells specified in the table definition. The tables and their contents are defined using a table definition language supplied through an input data stream defined by a //TABLEDEF DD statement.

The User Tables, unlike the standard tables, are self defining. Once the tables are built, you need not know the exact format of the table. You only need to refer to the data by key and field number. The Generic Subsystem retains the field names (column headings) as part of the User Tables.

Each record processed by the Generic Subsystem must contain a key field, referred to as an ELEMENT KEY. An element key is any descriptive piece of information you can use to report on later. For example, a JOBNAME can be the key to a report describing resources that are used by department, presuming that the JOB naming conventions allow JOBS to be associated with departments. In that case, the JOBNAME would be used as the key to the table that contains the resource usage information.

On the other hand, say you wanted to report on I/O activity by device address. The device address becomes the key to that particular table. Or perhaps it is important to analyze service unit usage by Performance Group Number. The PGN becomes the key to the service unit table. And so on.

When generating reports or graphs from the tables, the values within the tables are accessed through reference to the element keys.

A special type of element key, the GENERIC ELEMENT KEY, is used to determine whether the data contained within each input record is to be included in the Ad Hoc report. The GENERIC ELEMENT KEY is discussed in more detail in the section devoted to the generation of the Ad Hoc Reports.

The ELEMENT KEYS and other data fields within the input records, the User Tables, and the Generic Subsystem reports are defined by a set of table definition statements that are supplied to the Generic Subsystem through an input data set defined by a //TABLEDEF DD statement. These input statements are referred to as TABLEDEF statements.

The Record Level filters are specified through a set of parameters supplied through an input data set defined by a //SYSIN DD statement.

When it comes to selecting an element of data in an input record for use as the key to a CIMS Capacity Planner User Table, that data element must possess the following three characteristics:

- It can be displayed
- It is in a fixed position within a record or segment
- Keys may be up to 64 bytes in length and may be built from up to 4 non-contiguous segments

In the case of a JOBNAME, the data element itself normally possesses the three characteristics. A JOBNAME by its nature is displayable (must be Alphanumeric characters) and is less than eight characters long.

A Device Address, on the other hand, is normally kept in binary form in a record and must be converted to its displayable equivalent (for example, x'050E') before using it as a key.

The Performance Group Number is also usually maintained in binary form in an input record, although it is customary to refer to a PGN by its decimal value (for example, PGN00010) when using it as a key.

The data reduction portion of the Generic Subsystem lets you refer to a key in any of the three manners listed above:

- As an 8-character (or less) displayable field, padded on the right with blanks as necessary
- As a 4-byte binary value converted to its hexadecimal equivalent (8 hexadecimal digits, left filled with 00)
- As a 4-byte binary value converted to its decimal equivalent (8 decimal digits, left filled with 00)
- As an up to 64 byte key which may be built from up to 4 non-contiguous segments

A User Table must have one and only one key field.

Report Generation

Ad Hoc Reports

As mentioned earlier, the Ad Hoc reports are generated during the data reduction phase. The records containing data to be included in the Ad Hoc reports are selected based upon the element key contained in the record matching a GENERIC ELEMENT KEY you specified through one or more GENERIC ELEMENT KEY= parameters contained in the /SYSIN data stream. The Ad Hoc report contains a row of column headings, followed by subsequent rows of data containing the key field, the date and time stamps from the input record, and up to eight additional fields defined by OUTPUT ELEMENT CELLS. While more than eight OUTPUT ELEMENT CELLS can be specified, only the first eight are included in the Ad Hoc report.

The GENERIC ELEMENT KEY can contain up to 8 characters. The Wildcard format of the GENERIC ELEMENT KEY is supported in which all records whose element keys begin with a specified set of characters are to be selected. The Wildcard format lets you specify a common Root followed immediately by an asterisk. This format of the GENERIC ELEMENT KEY is interpreted to mean that any record containing a element key beginning with the characters preceding the asterisk will be selected. For example, if the keys in the records to be selected contain a job name and you only want to select those records where the job name begins with the characters ACCT, then GENERIC ELEMENT KEY=ACCT* would be specified.

Reports Generated From Tables

The Generic Subsystem provides several generic report formats that require little specification on your part. Their formats parallel standard CIMS Capacity Planner pre-programmed reports. The tables themselves contain self-describing dictionary data, so it is not necessary to define the makeup of either the User Tables or the standard CIMS Capacity Planner Tables to the batch reporting program (SSA1GENR). The available report types defined to the batch reporting program are:

Statistics Report

The Generic Subsystem Statistics report shows, for a keyed file, the top 40 elements in each selected category. The category can consist of a single user table cell or a value derived by applying an algorithm against a set of user table cells and/or literals. For example, presume that the user table consists of several elements whose key is a CICS USERID. Presume further that each USERID is associated with the following values in a cell set (one cell set for each hour in a 24 hour period):

- A cumulative volume of transactions
- The cumulative response time of the transactions
- A cumulative volume of EXCPs
- A cumulative value of Page operations
- The high water mark in storage for any given element

The statistics report could be used to determine the top 40 USERIDs in one or more of the above categories.

Daily Overview

The Generic Daily Overview shows, for each hour of the day for each day specified in the date range, the cumulative values as specified in the OUTPUT ELEMENT CELL key phrase(s) in the input stream. Because INCLUDE or EXCLUDE processing applies, it is possible to see the actual values for a specified element for a particular hour of a particular day in this report.

Activity Report (All Elements)

The Activity Report contains all specified elements, normally in alphabetical order by key. You can use the INCLUDE/EXCLUDE facility to limit the population of keys in this particular report.

Although the report is normally produced in Element Key ascending sort order (for example, JOBNAME), you can alter this order through the use of one of the following two key phrases:

ASCENDING SORT COLUMN=nn

or

DESCENDING SORT COLUMN=nn

where the "nn" is the OUTPUT ELEMENT CELL number as specified in the TABLEDEF input stream. This causes the sort to be ordered, either in ascending or in descending sequence, based on the values in column "nn".

Substitution (&sid.GENF)

In any report or graph that uses the element key as part of the report (Statistics, Activity List), the GENF member of CPPR.PARMLIB can be used to actually replace keys with synonyms. In this manner it is possible to combine sets of keys into common applications (for example) and to report on them as a group.

In order to do this, specify in the SYSIN data stream:

```
GENERIC ELEMENT MASK=11111111
```

and set the &sid.GENF member of PARMLIB to handle up to 64 sets of synonyms. See the model CPPRGENF member in the Distribution PARMLIB.

Using the ISPF/PDF Interface

The Generic portion of the ISPF/PDF interface lets you produce several reports and GDDM graphs from either CIMS Capacity Planner Standard or User Tables. Because the reports and graphs are generated from the tables already in the data base, it is not necessary to supply the details regarding the form of the original input data. It is only necessary to provide the following:

- The Table Number
- The Table Suffix (CIMS Capacity Planner Standard Suffixed Table: for example, CICS Transaction Statistics Table)
- The Element Key (INCLUDE/EXCLUDE processing)
- The order in which values are stored in the table (Cell Number)
- The values that you want to see in the final report or graph and how they are calculated from the values in the table along with the ordinary CIMS Capacity Planner input data that describes the date range for the period and the SMF SID.

Generic Subsystem Operation

The information presented above is intended to serve as an introduction to the Generic Subsystem and focuses on its structure and general capabilities. This section provides the detailed information required to operate the Generic Subsystem.

Data Reduction

Acceptable Forms of Input

The Generic Subsystem supports both fixed and variable record formats (blocked or unblocked). The Undefined format is not supported. In order to be processed by the CIMS Capacity Planner Generic Subsystem, each input record must have the following minimum characteristics:

- A Date Stamp for each record
- A Time Stamp for each record
- Fixed or Variable length records (Not Undefined)
- Binary Cell Values no longer than four bytes wide
- Keys of 64 bytes or less in length which may be built from up to 4 non-contiguous segments

The date and time stamps are mandatory because CIMS Capacity Planner summarizes and compresses the input into hourly sets of data cells that are kept in user tables that, themselves, are maintained by date. In this manner, it is possible to produce reports on an hourly basis for selected dates, regardless of the original form of the input.

Fixed Length Records

The simplest form of an input record is one where the entire record is contained in a fixed length array. For example, presume that a data center has a homegrown online system supporting fifty users. Presume further that the transactions processed by the online system have eight character identifiers, and that the online system includes an integrated monitor that is used to track I/O operations by transaction.

If the local monitor produces 20 byte records, and each record contains:

- A Julian Date in the first four positions,
 - A Time stamp in the next four positions,
 - A Transaction Identifier in the next eight positions
 - The number of I/O operations used by that transaction in the next four positions,
- the Generic Subsystem could easily reduce that information into a user-defined table.

VLR Records with a Fixed Length Header

A slightly more complex form of input is comprised of Variable length records with a standard VLR length word at the beginning of the record. In this case, the record typically consists of a fixed length header followed by one or more variable length sections, each of which has its own length defined by a field that resides in the section itself. The VLR sections might or might not contain a "type" field that can be used to select the section for further processing.

Using the example above (fixed length records), presume that the homegrown monitor has become more complex, and that now several data sections can follow the header. Presume further that the sections that follow the header can be Transaction sections (type X) or Terminal sections (type T). In this case, the record consists of a 12 byte header:

- VLR Record length word
 - Julian Date in the next four positions,
 - Time stamp in the next four positions, plus one or more Transaction sections containing:
 - A 1 byte length of the Transaction section (18),
 - The character "X",
 - A Transaction Identifier in the next eight positions,
 - The number of I/O operations used by that transaction in the next four positions
 - The amount of CPU time used in the last 4 positions
- plus one or more Terminal sections containing:
- A 1 byte length of the Terminal section (14),
 - The character "T",
 - A Terminal Identifier in the next eight positions,
 - The number of terminal Write operations for that terminal in the next four positions

The Generic Subsystem could easily parse the record and, using the "type" field in each of the sections following the header, could select either the Terminal (T) sections or the Transaction (X) sections for processing.

VLR Records with Segments

Many commercially available monitors produce more complex record structures than those described above. In order to process these records, the Generic Subsystem requires more information. For example, the record that is produced can have a Record number in the header (records that are written to the SMF file MUST have a record number in the header in order to distinguish them from other SMF records) plus a record number qualifier (where many record subtypes are produced) plus other subqualifiers that further qualify the format and contents of the record being processed.

The Generic Subsystem lets you to specify the following selectivity criteria for selecting specific records and/or specific record subsections:

- A Record number
- A Record number qualifier
- A Segment Identifier
- A Segment identifier subqualifier

so that the proper subset of records can be selected for processing. This record and segment selectivity information is necessary when the layout of the record depends on the type and subtype of the record.

Using the example above, presume that the homegrown monitor has become even more sophisticated, and that now the records produced are being written as SMF User records (record number 233) to the SMF Cluster. Presume further that the Transaction sections (type "X") that follow the header have been expanded to include subtypes: a subtype "01" contains I/O operation counts and a subtype "02" contains CPU times. In this case, the record consists of a standard SMF 18 byte header:

- A VLR Record length word
- A System Identifier (MVS/370, MVS/XA, MVS/ESA)
- The record number (233)
- A Time stamp in the next four positions,
- A Julian Date in the next four positions
- A four byte SMF System ID (for example, CPPR)

plus one or more Transaction I/O sections containing:

- A 1 byte length of the Transaction section (15),
- The character "X",
- A section subtype ("01")
- A Transaction Identifier in the next eight positions,
- The number of I/O operations used by that transaction in the next four positions

plus one or more Transaction CPU sections containing:

- A 1 byte length of the Transaction section (15),
- The character "X",
- A section subtype ("02")
- A Transaction Identifier in the next eight positions,
- The amount of CPU time used in the last 4 positions

plus one or more Terminal sections containing:

- A 1 byte length of the Terminal section (14),
- The character "T",
- A Terminal Identifier in the next eight positions,
- The number of terminal Write operations for that terminal in the next four positions

The Generic Subsystem could easily parse the record and, using the SMF record number plus the "type" field in each of the sections following the header, plus the subtype field for the Transaction sections, could select either of the Terminal (T) sections or the Transaction (X) sections for processing.

Segment Descriptor Triplets (RMF)

RMF records are a special case of SMF records where the sections following the header are described by three values at a fixed location in the header itself. These three values are:

- The offset from the beginning of the record to the 1st occurrence of a contiguous set consisting of one or more of the sections being described
- The length of this type of section
- The actual count of this type of section in the record

Indeed, several other SMF records have adopted this scheme of imbedding self-describing information in the header. The Generic Subsystem respects the presence of this self-describing information and lets you indicate whether each of the three fields are 1, 2, 3 or 4 bytes wide themselves.

VLR Segments with Fixed Length Prefix

Some monitors further complicate the situation by providing an imbedded length at a fixed offset into a variable length segment. For example, Candle® CL/SUPERSESSSION® has implemented this scheme of encapsulating information into a highly compact SMF User record. The Generic Subsystem allows for this scheme as well.

Date and Time Stamp Information

As noted above, the Generic Subsystem requires that each record have both a date and a time stamp. The varieties of formats for each of these stamps is almost infinite. The Generic Subsystem provides for perhaps 95% of the standard forms, and lets you supply a User Exit for the other 5% of the cases.

Forms of Time Stamp Information

Supported Time Stamp information includes at least the following:

- A four byte field containing the count of hundredths of a second since midnight
- A four character field in the form HHMM (hours, minutes since midnight)
- A two character field (HH)

- A four byte packed field in the form HHMMSSsf (Hours, Minutes, Seconds, Tenths of a second)
- An eight byte field containing the STCK value (CICS)
- A four byte field containing the top word of the STCK value
- A four byte field containing the time in 64 microsecond units

See the discussion on parameters for further information.

Forms of Date Stamp Information

Supported Date Stamp information includes at least the following:

- A four byte packed field in the form 0YYDDDC that represents the Julian Date
- A six character field in the form MMDDYY (Month, Day, Year)
- A six character field in the form DDMMYY
- A six character field in the form YYMMDD
- A five character field in the form YYDDD (Julian Date in character form)
- An eight character field in the form MM/DD/YY
- An eight character field in the form DD.MM.YY

See the discussion on parameters for further information.

Element Keys

The user table that is built by the Generic Subsystem must be keyed. That is, data must be gathered into cell sets associated with a particular key (for example, Transaction ID or Terminal Name) and the reports thus generated reflect the data associated with a particular key. The advantage of this approach is that specific keys can be eliminated or included for a report. For example, if you wanted to see all I/O operations for all Transactions beginning with the characters "ZQ", the INCLUDE function would allow this type of selectivity if the User Table had been built using the transaction ID as the key.

Forms of Element Keys

Keys must be 8 characters or less in length. Keys can be in one of three forms:

- From 1 to 8 Characters, in which case the stored key are padded on the right with blanks
- A four byte hexadecimal field that is to be stored as the EBCDIC equivalent (for example, 01C4), in which case the stored key are padded on the right with blanks
- A four byte hexadecimal field that is to be stored as the Decimal equivalent (for example, 15), in which case the stored key are padded on the right with "00" for a total of seven characters (for example, 0000015) ending with a blank

For more information, see the section on parameters.

Input Data Cells

An input data cell refers to the location where the selected information from the user-specified input resides. An input data cell has two characteristics:

- The offset into the segment to the beginning of the cell containing the information
- The size of the cell itself (1, 2, 3 or 4 bytes; or in the case of a key, a time or a date, up to 8 bytes)

All input data cells except keys, times and dates are presumed to be in binary format. Dates are converted to Julian form based on the format of the original in the input as specified. Times are converted to hundredths of a second since midnight based on the format of the original in the input as specified.

User Defined Tables

The table built by the Generic Subsystem is self-defining, using the parameters you specify. The table has the following characteristics:

- It must have a Table Number ranging from 128 (x'80') through 159 (x'9F') as specified by the user
- It must be keyed
- It contains 24 hourly cell sets per key entry
- Each cell set consists of up to 99 four-byte value fields that represent the unique values gathered by the Generic Subsystem as specified.

Describing the Input File (TABLEDEF)

The key phrases used to describe the characteristics of the input file are read from the TABLEDEF input file. CIMS Capacity Planner input parameter conventions are used (* in column one is a comment, double space terminates a scan, etc.).

Identifying the SMF System ID

If the input file has an SMF System ID associated with it, the location of the SMFSID must be specified. Otherwise, the SMFSID is set from the "SELECTED SYSTEM" parameter that comes from the SYSIN stream.

SMFSID OFFSET

This key phrase tells the Generic Subsystem where to find the SMF System Identifier in the record header. If this field is omitted, the SMF SID is set according to the "SELECTED SYSTEM" parameter in the SYSIN stream.

Providing Selectivity Criteria

If each of the logical records in the input file has a record type that is to be used to identify incoming records to be reduced, the location and size of the field containing the record type must be specified. In addition to a record number, it is possible to specify that a record type qualifier be matched in order to select the record. The record type qualifier value must reside in the record header.

If the incoming records do not have a record type or if the format of all incoming records is identical and no specific record types are to be selected, the following parameters can be omitted.

RECORD NUMBER OFFSET

This key phrase tells the Generic Subsystem where to find the record number in the record header. If the record does not contain a record number field, this parameter is omitted.

RECORD NUMBER LENGTH

This key phrase tells the Generic Subsystem the length of the record number field in bytes. If the record does not contain a record number field, this parameter is omitted.

RECORD NUMBER

This key phrase tells the Generic Subsystem the decimal value of the record number field in determining whether to select a specific record for processing. If the record does not contain a record number field, this parameter is omitted. You can select only a single record type for processing during a single pass of the input file.

RECORD HEADER TYPE QUALIFIER OFFSET

This key phrase tells the Generic Subsystem where to find the record type qualifier in the record header. If the record does not contain a record type qualifier field, this parameter is omitted.

RECORD HEADER TYPE QUALIFIER LENGTH

This key phrase tells the Generic Subsystem the length of the record type qualifier field in bytes. If the record does not contain a record type qualifier field, this parameter is omitted.

RECORD HEADER TYPE QUALIFIER VALUE

This key phrase tells the Generic Subsystem the decimal value of the record type qualifier field in determining whether to select a specific record for processing. If the record does not contain a record type qualifier field, this parameter is omitted.

Logical Records With Segment Pointer Triplets (SMF)

Some logical records consist of a header portion plus one or more Segment Pointer Triplets (4 byte Segment Offset, 2 byte Segment length, 2 byte Segment count). For example, many of the standard SMF and RMF records have this format. The fields shown below are used to process this type of logical record.

RECORD SUBTYPE SEGMENT POINTER OFFSET

This key phrase tells the Generic Subsystem where to find the pointer (in the record header) to the segment that contains the record subtype. This parameter applies to SMF and SMF-compatible record formats. If record subtype checking is not performed, this parameter is omitted.

RECORD SUBTYPE SEGMENT POINTER LENGTH

This key phrase tells the Generic Subsystem the length of the segment pointer (in the record header) for the segment that contains the record subtype. This parameter applies to SMF and SMF-compatible record formats. If record subtype checking is not performed, this parameter is omitted.

SUBSYSTEM ID SEGMENT POINTER OFFSET

This key phrase tells the Generic Subsystem where to find the pointer (in the record header) to the segment that contains the Subsystem ID. This parameter applies to SMF and SMF-compatible record formats. If no Subsystem IDs are to be processed, this parameter is omitted.

SUBSYSTEM ID SEGMENT POINTER LENGTH

This key phrase tells the generic subsystem the length of the segment pointer (in the record header) for the segment that contains the Subsystem ID. This parameter applies to SMF and SMF-compatible record formats. If Subsystem IDs are not to be processed, this parameter is omitted.

RECORD SEGMENT LENGTH FIELD LENGTH

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the length of the length field.

RECORD SEGMENT LENGTH FIELD OFFSET

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the offset from the beginning of the record to the field that contains the length of the segment that is to be processed.

RECORD SEGMENT OFFSET FIELD LENGTH

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the length of the offset field.

RECORD SEGMENT OFFSET FIELD OFFSET

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the offset from the beginning of the record to the field that contains the offset of the segment that is to be processed.

RECORD SEGMENT COUNT FIELD LENGTH

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the length of the count field.

RECORD SEGMENT COUNT FIELD OFFSET

This field is used for RMF records whose contents are described by triplets (length, offset, count). It describes the offset from the beginning of the record to the field that contains the count of the segments that are to be processed.

Segmented Logical Records

Some logical records consist of a header portion plus one or more variable length segments, each of which has its own length value defined as a fixed portion plus a length field. In this case, the header portion must have a length field that contains the entire length of the logical record (standard VLR format). The segment lengths are derived by adding a fixed length for each segment (the length of the fixed portion may, in fact, be 0 if the segment is entirely variable in length) to the length field contained in the segment (the variable length portion may, in fact, be 0 if the segment is fixed in size).

RECORD TOTAL LENGTH FIELD OFFSET

This key phrase tells the Generic Subsystem where to find the logical record length in the record header. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the record does not contain a collection of variable length segments, this parameter is omitted.

RECORD TOTAL LENGTH FIELD LENGTH

This key phrase tells the Generic Subsystem the length of the logical record length in the record header. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the record does not contain a collection of variable length segments, this parameter is omitted.

RECORD HEADER LENGTH

This key phrase tells the Generic Subsystem the length of the record header. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the record does not contain a collection of variable length segments, this parameter is omitted.

VARIABLE LENGTH HEADER LENGTH FIELD LENGTH

This field is used for records (such as the SMF VSAM CLOSE Type 64 record) that have variable length headers followed by a segment that is to be processed. This field and the field immediately following are used in conjunction with the "RECORD HEADER LENGTH" key phrase to describe the manner in which to calculate the length of the header. This field describes the length of the length field.

VARIABLE LENGTH HEADER LENGTH FIELD OFFSET

This field is used to describe the offset from the beginning of the record to the field that contains the length of the remainder of the variable length header. This value is added to the value provided by the "RECORD HEADER LENGTH" value to determine the actual length of the header.

RECORD SEGMENT FIXED LENGTH

This key phrase tells the Generic Subsystem the length of the fixed portion of the variable length segment. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the segment is entirely variable in length (the length field in the variable length portion includes the entire segment length), this parameter is omitted.

RECORD SEGMENT VARIABLE LENGTH FIELD OFFSET

This key phrase tells the Generic Subsystem the offset into the segment of the field that contains the length of the variable portion of that segment. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the segment is entirely fixed in length, this parameter is omitted.

RECORD SEGMENT VARIABLE LENGTH FIELD LENGTH

This key phrase tells the Generic Subsystem the length of the field that contains the length of the variable portion of that segment. This parameter applies to logical records that contain a fixed length header plus a variable number of variable length segments. If the segment is entirely fixed in length, this parameter is omitted.

Segment Descriptor Key Phrases

The record segments themselves are described by the following key phrases.

Phrases Identifying the Input Record Sub-Type

If each of the logical records in the input file have a record sub-type that is to be used to identify incoming records to be reduced, the location and size of the field containing the record sub-type must be specified. If the incoming records do not have a record sub-type or if the format of all incoming records is identical and no specific record types are to be selected, the following parameters can be omitted.

RECORD SUBTYPE OFFSET

This key phrase tells the Generic Subsystem where to find the record subtype in the segment described above. The record subtype is used to filter selected record segments when scanning for a value defined through the INPUT ELEMENT CELL statement (see below). If record subtype checking is not performed, this parameter is omitted.

RECORD SUBTYPE LENGTH

This key phrase tells the Generic Subsystem the length of the record subtype field in bytes. If the record does not contain a subtype field, this parameter is omitted.

RECORD SUBTYPE QUALIFIER OFFSET

This key phrase tells the Generic Subsystem where to find the record subtype qualifier in the segment described above. The record subtype qualifier is used to further filter selected record segments when scanning for a value defined through the INPUT ELEMENT CELL statement (see below). If record subtype qualifier checking is not performed, this parameter is omitted.

RECORD SUBTYPE QUALIFIER LENGTH

This key phrase tells the Generic Subsystem the length of the record subtype qualifier field in bytes. If the record does not contain a subtype qualifier field, this parameter is omitted.

Identifying the Input Record Element Key

Each of the logical records in the input file has a field that uniquely identifies an element key that is to be used to further qualify cell values. The location and size of the field containing the element key must be specified.

ELEMENT KEY OFFSET

This key phrase tells the Generic Subsystem where to find the element key in the segment described above. If element keys are not to be collected, or if the element key is identified through the INPUT ELEMENT CELL statement (see below), this parameter is omitted.

ELEMENT KEY LENGTH

This key phrase tells the Generic Subsystem the length of the element key field in bytes. Only the first eight bytes of the key are saved. If the key is less than eight bytes long, the key is padded on the right with blanks. If element keys are not to be collected, or if the element key is identified through the INPUT ELEMENT CELL statement (see below), this parameter is omitted.

Identifying the Input Record Subsystem ID

If each of the logical records in the input file have a field that uniquely identifies the Subsystem that is to be used to further filter out unwanted records, the location and size of the field containing the Subsystem ID must be specified. The Subsystem ID must be in the record header. If the Subsystem ID is eight bytes or less and is in character form (for example, a VTAM APPLID), the actual Subsystem name to be matched is input through the SYSIN key phrase "VTAMNAME". If the incoming records do not have Subsystem IDs, or if Subsystem IDs are not to be taken into consideration, the following parameters can be omitted.

SUBSYSTEM ID OFFSET

This key phrase tells the Generic Subsystem where to find the Subsystem ID in the segment described above. If Subsystem IDs are not to be processed, this parameter is omitted.

SUBSYSTEM ID LENGTH

This key phrase tells the Generic Subsystem the length of the Subsystem ID field in bytes. If Subsystem IDs are not to be processed, this parameter is omitted.

INPUT ELEMENT CELL

This key phrase describes a specific Input Element Cell from which data is to be collected. The format of the argument for this key phrase is positional. The parameters are specified as follows:

- A field identifying the Cell Number (01-99)
- A field identifying the input record subtype, if any. If no input record subtype is used, this field must be 00.
- A field identifying the input record subtype qualifier, if any. If no input record subtype is used, this field must be 00.

- A field identifying the location of the offset field in the record header, if any, that points to the segment that contains the value being described. If the record does not contain segments, or if the value is in the record header, this field is 00.
- A field identifying the length of the offset field in the record header, if any, that points to the segment that contains the value being described. If the record does not contain segments, or if the value is in the record header, this field is 00.
- A field specifying the offset into the segment for the field that contains the value being described. This value must be specified.
- A field specifying the length of the field that contains the value being described. This value must be specified.
- A field that describes the format of the value being described.

The following table itemizes the various formats of the data elements that are supported by the Generic Subsystem. All the value formats are categorized as either **Normal**, **Time of Record**, or **Date of Record**. Values that are categorized as **Normal** are values that are used for calculation and reporting purposes that are **not** related to the time and date of an event. All values not related to dates and times are categorized as **Normal**. Date and Time fields can be classified as either **Normal**, **Time of Record**, or **Date of Record** depending upon their meaning. A date or a time field is classified as **Normal** if it is **not** used to determine into which of the 24 hourly cell sets the data within the record is summarized. If, on the other hand, a date or time field is used in determining into which of the 24 sets of hourly counters (cells) the data contained within a record is to be summarized, it is classified as **Time of Record** or **Date of Record**. Only one date field and one time field can be specified as **Time of Record** or **Date of Record**.

Table 8-1 • Table of input cell value type definitions

Value Formats	Normal
Binary Value	00
Packed Decimal Value	01
Character Value	02
Element Key (character)	03
Element Key (hex-decimal)	04
Element Key (hex-EBCDIC)	05
Numeric (EBCDIC + punctuation)	06
Floating Point Value (8 bytes)	07
Floating Point Value (/1000)	08
Cumulative Value, Derive Delta	0A
Display-only Value (Binary)	0B

Table 8-1 • Table of input cell value type definitions

Value Formats		Normal
Display-only Value (Character)		0C

Time Value Formats	Normal	Time of Record
SMF HDR (.XX Seconds)	10	50
PL4'HHMMSSSC'	11	51
XL4'.X SECONDS'	12	52
8 BYTE STCK FORMAT	13	53
4 BYTE STCK FORMAT	14	54
CICS 16 MS Units (8 Byte)	15	55
PL2'OHHC' (Actual Hour)	16	56
10,000th SECOND (IDMS)	17	57
64 MICROSEC UNITS (TMON)	18	58
PL4'MMSSTTTC' (RMF Interval)	19	59
PL4'OHHMMSSC' (RMF Start Time)	1A	5A
CL4'HHMM' (VM Monitor Format)	1B	5B
16 MS Units (4 Byte)	1C	5C
Reserved	1D	5D
Reserved	1E	5E
User Exit 02	1F	5F

Date Value Formats	Normal	Date of Record
SMF HDR (00YYDDDC)	20	60
CL8'MM/DD/YY' USA	21	61
CL8'DD.MM.YY' EUROPE	22	62
CL6'MMDDYY' USA	23	63
CL6'DDMMYY' EUROPE	24	64
CL5'YYDDD'	25	65
CL6'YYMMDDD'	26	66

Time Value Formats	Normal	Time of Record
STCK Top Word	27	67
User Exit 01	2F	6F

Sample Input File Definitions

Example 1

```

*****
* THIS EXAMPLE DESCRIBES VARIOUS FIELDS IN THE SMF TYPE 30
* RECORD, SUBTYPE 5 (JOB TERMINATION).
*****
* FIRST WE IDENTIFY THE SYSTEM
SMFSID OFFSET=14
* NEXT WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=144 /* X'90' */
* NEXT WE NOTE THE LOCATION OF THE RECORD NUMBER
RECORD NUMBER OFFSET=5
RECORD NUMBER LENGTH=1
RECORD NUMBER=30
* NEXT WE NOTE THE LOCATION OF THE RECORD SUB-TYPE
RECORD SUBTYPE SEGMENT POINTER OFFSET=24
RECORD SUBTYPE SEGMENT POINTER LENGTH=4
RECORD SUBTYPE OFFSET=0
RECORD SUBTYPE LENGTH=2
* NEXT WE NOTE THE LOCATION OF THE CIMS Capacity Planner TABLE KEY (JOBNAME)
ELEMENT KEY SEGMENT POINTER OFFSET=32
ELEMENT KEY SEGMENT POINTER LENGTH=4
ELEMENT KEY OFFSET=0
ELEMENT KEY LENGTH=8
* NEXT WE DESCRIBE THE INPUT CELLS OF INTEREST
* NOTE THAT ALL CELLS ARE CONTAINED IN THE SUBTYPE 5 RECORD ONLY
* INPUT CELL NUMBER (MUST BE UNIQUE)
* | INCOMING RECORD SUBTYPE
* | | INCOMING RECORD SUBTYPE QUALIFIER
* | | | SEGMENT OFFSET LOCATION IN RCD HDR
* | | | | SEGMENT OFFSET LENGTH
* | | | | | VALUE FIELD OFFSET INTO SEGMENT
* | | | | | | VALUE FIELD LENGTH
* | | | | | | | VALUE FIELD TYPE
* | | | | | | | |
INPUT ELEMENT CELL=01,05,00,0,0,6,4,50 /* TIME FROM SMF HEADER
INPUT ELEMENT CELL=02,05,00,0,0,10,4,60 /* DATE FROM SMF HEADER
INPUT ELEMENT CELL=03,05,00,0,0,0,0,0 /* USED FOR FREQUENCY

```

```

INPUT ELEMENT CELL=04,05,00,32,4,56,4,10 /* PROGRAM INIT TIME
INPUT ELEMENT CELL=05,05,00,32,4,64,4,10 /* JES READER TIME
INPUT ELEMENT CELL=06,05,00,40,4,4,4,0 /* TOTAL EXCP COUNT
INPUT ELEMENT CELL=07,05,00,56,4,4,4,0 /* TCB CPU TIME
INPUT ELEMENT CELL=08,05,00,56,4,8,4,0 /* SRB CPU TIME
INPUT ELEMENT CELL=09,05,00,88,4,8,4,0 /* TAPE MOUNTS
INPUT ELEMENT CELL=10,05,00,88,4,12,4,0 /* TAPE MOUNTS
INPUT ELEMENT CELL=11,05,00,56,4,12,4,0 /* INIT TCB CPU
INPUT ELEMENT CELL=12,05,00,56,4,16,4,0 /* INIT SRB CPU

```

Example 2

```

***** THIS
EXAMPLE ILLUSTRATES THE DEFINITION OF A CL/SUPERSESSION RECORD

```

```

* FROM CANDLE CORPORATION.

```

```

*****

```

```

//TABLEDEF DD *
* FIRST WE IDENTIFY THE SYSTEM
SMFSID OFFSET=14
* NEXT WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=145 /* X'91' */
* THESE ARE LINKED LENGTHS. DEFINE THE LENGTH VECTORS.
RECORD TOTAL LENGTH FIELD OFFSET=0
RECORD TOTAL LENGTH FIELD LENGTH=2
RECORD HEADER LENGTH=28
RECORD SEGMENT FIXED LENGTH=10
RECORD SEGMENT VARIABLE LENGTH FIELD OFFSET=9
RECORD SEGMENT VARIABLE LENGTH FIELD LENGTH=1
* NEXT WE NOTE THE LOCATION OF THE RECORD NUMBER IN THE HEADER
RECORD NUMBER OFFSET=5
RECORD NUMBER LENGTH=1
RECORD NUMBER=254
* NEXT WE NOTE THE LOCATION OF THE RECORD SUB-TYPE
RECORD SUBTYPE SEGMENT POINTER OFFSET=0
RECORD SUBTYPE SEGMENT POINTER LENGTH=0
RECORD SUBTYPE OFFSET=0
RECORD SUBTYPE LENGTH=1
* NEXT WE NOTE THE LOCATION OF THE CIMS Capacity Planner TABLE KEY
ELEMENT KEY SEGMENT POINTER OFFSET=0

```

```
ELEMENT KEY SEGMENT POINTER LENGTH=0
ELEMENT KEY OFFSET=10
ELEMENT KEY LENGTH=8
*   NEXT WE DESCRIBE THE INPUT CELLS OF INTEREST (CL/SUPERSESSION)
*   NOTE THAT ALL CELLS ARE CONTAINED IN THE SUBTYPE 227 RECORD ONLY
INPUT ELEMENT CELL=01,227,0,0,0,1,4,50 /* TIME FROM SEGMENT HEADER
INPUT ELEMENT CELL=02,227,0,0,0,5,4,60 /* DATE FROM SEGMENT HEADER
INPUT ELEMENT CELL=03,227,0,0,0,0,0,0 /* USED FOR FREQUENCY
INPUT ELEMENT CELL=04,227,0,0,0,48,4,0 /* INBOUND MESSAGE COUNT
INPUT ELEMENT CELL=05,227,0,0,0,52,4,0 /* INBOUND BYTE COUNT
INPUT ELEMENT CELL=06,227,0,0,0,40,4,0 /* OUTBOUND MESSAGE COUNT
INPUT ELEMENT CELL=07,227,0,0,0,44,4,0 /* OUTBOUND BYTE COUNT
INPUT ELEMENT CELL=08,227,0,0,0,56,4,0 /* REFRESH MESSAGE COUNT
INPUT ELEMENT CELL=09,227,0,0,0,60,4,0 /* REFRESH BYTE COUNT
INPUT ELEMENT CELL=10,227,0,0,0,64,4,0 /* UPDATE MESSAGE COUNT
INPUT ELEMENT CELL=11,227,0,0,0,68,4,0 /* UPDATE BYTE COUNT
INPUT ELEMENT CELL=12,227,0,0,0,72,4,0 /* TERMINAL INPUT MESSAGE COUNT
INPUT ELEMENT CELL=13,227,0,0,0,76,4,0 /* TERMINAL INPUT BYTE COUNT
```

User Table Specification

User Tables are stored in much the same format as the standard CIMS Capacity Planner tables. Each table contains a header segment that contains the length of the table, a date, an SMF SID, and a table number. Following the header is the user-defined portion of the table that consists of a number of table element sets, each of which contains a key and 24 Cell Sets, one Cell Set corresponding to each hourly period throughout the day. A Cell Set consists of up to 99 full word values. The tables and their contents are defined using a table definition language supplied through an input data stream defined by a // TABLEDEF DD statement.

The User Tables, unlike the standard tables, are self defining. Once the tables are built, you need not know the exact format of the table. You only need to refer to the data by key and cell number when specifying reports or graphs.

Describing the User Table

In general, the following points describe the user tables built by the Generic Subsystem:

- User tables are numbered 128 (x'80') through 159 (x'9F')
- A user table must have a single element key (such as a JOBNAME)
- Each cell set for a user table consists of one hour's worth of data for up to 99 separate fields
- Input for any given user table must come from a specific input record type. That is, multiple input record types cannot be used to build a single user table in a single data reduction execution. It is possible, however, to make several passes, collecting data from different input record types on each pass of the input file into the same user table.
- User Tables are self defining

Once a CIMS Capacity Planner User Table has been built, it is no longer necessary to provide details of the shape of the table to the reporting or graphing programs. The tables describe themselves. This means that the report program or graph program understands the layout of the table and can select values that have been defined in the data reduction phase through the OUTPUT ELEMENT CELL key phrases in the TABLEDEF input stream.

This also means that you must be very careful when assigning CIMS Capacity Planner User Table numbers. If the same Table number is used for different tables within the same PDB, strange results occur, especially if the tables themselves are different shapes!

- Table Extract can also be used with User Tables

Because the CIMS Capacity Planner User Tables are self describing, you can use the Table Extract Utility (SSA1TBLX) to extract elements from the table and pass them to a locally defined report writer or other analytical tools.

TABLE NUMBER

This key phrase tells the Generic Subsystem that CIMS Capacity Planner User Table number to assign to this table. Table numbers 128 (x'80') through 159 (x'9F') are reserved for CIMS Capacity Planner User Tables.

OUTPUT ELEMENT CELL

This key phrase describes a specific Output Element Cell into which data is to be collected. The format of the arguments for this key phrase is positional. The parameters are specified as follows:

- A field identifying the Output Cell Number (01-99)
- A reserved field
- A reserved field

- A field that describes the format of the value being described. The following value formats are supported by the Generic Subsystem:

Table 8-2 • Table of output cell value type definitions

Value Formats	Normal
Binary Value	00
Packed Decimal Value	01
Character Value	02
Time Value (.xx Sec)	03
Frequency Counter	04
Percentage (xx.xx%)	06
.xx Binary Value	07
Save Highest Value (.xx value)	08 (e.g., high water mark)
Average (.xx value)	09
Save Lowest Value (Not 0)	0A (e.g., low water mark)

- A formula describing the algorithm for deriving the output cell value. The formula consists of input cell numbers and/or literals plus mathematical operators. The following rules apply to the formula:

- 1 Up to 12 levels of parentheses are supported
- 2 Numbers refer to Input Cell Numbers
- 3 Literals are enclosed in single quotes ("")
- 4 Operators supported are as follows:
 - (+) A + B Means add A to B and save the sum
 - (-) A - B Means subtract B from A and save the difference
 - (/) A / B Means divide A by B and return the integer quotient
 - (*) A * B Means multiply A by B and save the product
 - (:) A : B Means divide A by B and return the remainder (Modulo)

- A twelve byte label for the resultant value

Sample User Table Definition

Example 1

This table describes a User Table containing several values extracted from SMF Type 30 records. The table is keyed on JOBNAME. This table is the result of processing the input described in the example illustrated under *Example 1 on page 8-32*.

```
*****
*   NEXT WE DESCRIBE THE OUTPUT CELLS TO BE CREATED IN THE CIMS Capacity Planner TABLE
*
*           OUTPUT CELL NUMBER (MUST BE UNIQUE)
*           | OUTPUT VALUE FORMAT
*           | |FORMULA FOR DERIVING OUTPUT CELL VALUE
*           | || 12 CHAR VALUE LABEL
*           |   |   |   |
OUTPUT ELEMENT CELL=01,0,0,03,((04-05)),IN JES QUEUE
OUTPUT ELEMENT CELL=02,0,0,03,((01-05)), THROUGHPUT
OUTPUT ELEMENT CELL=03,0,0,04,((03)), FREQUENCY
OUTPUT ELEMENT CELL=04,0,0,00,((06)), TOTAL EXCPS
OUTPUT ELEMENT CELL=05,0,0,07,((07+08+11+12)), CPU SECONDS
OUTPUT ELEMENT CELL=06,0,0,00,((09+10)), TAPE MOUNTS
OUTPUT ELEMENT CELL=07,0,0,00,((06/07)),RANDOM VALUE
OUTPUT ELEMENT CELL=08,0,0,06,(((04-05)*'10000')/(01-05)), INQ:THRUPUT
```

Example 2

Following is a sample table used to store selected information from the Candle CL/SUPERSESSION record described in *Example 2 on page 8-33*. The record key has been described under the input specifications.

```
*   NEXT WE DESCRIBE THE OUTPUT CELLS TO BE CREATED IN THE CIMS Capacity Planner TABLE
OUTPUT ELEMENT CELL=01,0,0,04,((03)),LOGON COUNT
OUTPUT ELEMENT CELL=02,0,0,00,((04)),INPUT MSGS
OUTPUT ELEMENT CELL=03,0,0,00,((06)),OUTPUT MSGS
OUTPUT ELEMENT CELL=04,0,0,00,((11/07)),CMPRS RATIO
OUTPUT ELEMENT CELL=05,0,0,00,((05+07)),APPL<=>SS
OUTPUT ELEMENT CELL=06,0,0,00,((09+11)),SS<=>TERM
OUTPUT ELEMENT CELL=07,0,0,00,((07/06)),OUTMSGLEN
OUTPUT ELEMENT CELL=08,0,0,06,(((11)*'10000')/(07)),CMPRS PCT
```

Ad Hoc Reports

The Generic Subsystem gives you the option of producing an Ad Hoc report from the original input data. Generic Ad Hoc reports are generated by the data reduction program (SSA1GENW). Ad Hoc reports are particularly useful in preparing reports based upon detailed input data that can be summarized or discarded by CIMS Capacity Planner. Ad Hoc reports can be generated during the generation of user tables or by themselves. Ad Hoc reports, by their general nature are generated on an as-needed basis and are not generally associated with the creation of user tables, although an Ad Hoc can be produced coincident with the generation of a user table. It is important that multiple updates are not made to the user tables contained in the Performance Database as the result of generating Ad Hoc reports. **Unless an Ad Hoc report is associated with the generation of a user table, then the ONLINE and the INDEX DD statements should be specified as DD DUMMY.**

Record Selection Parameters

Records are selected for inclusion in the Ad Hoc based upon several parameters that are specified through parameters entered through the //SYSIN data set:

GENERIC ELEMENT KEY=

The Generic Element Key keyword specifies the element key or the common root of multiple element keys related to records that are to be selected for inclusion into the Ad Hoc report. The Generic Element Key parameter is required if an Ad Hoc report is to be produced. To include all records, the element key is specified as an asterisk (GENERIC ELEMENT KEY=*)

BEGIN DATE=

END DATE=

The beginning and ending dates (inclusive) are used to constrain the amount of data to be selected for inclusion into the Ad Hoc report. The date range also constrains the records to be included in maintaining the user tables.

BEGIN TIME=

END TIME=

The beginning and ending times (inclusive) are used to constrain the amount of data to be selected for inclusion into the Ad Hoc report.

User Exits

The inclusion or exclusion of an input record is also influenced by the user exit routines. If a non-zero return code is returned from one of the user exits, the record currently being processed is excluded from the Ad Hoc report.

IMBED=

This command, placed in the TABLEDEF input stream, allows you to imbed an input record definition from the Generic library into the TABLEDEF input stream. For more information, refer to the CNTL Library member named DGENR200.

//UEXALIAS

In the Generic Subsystem Data Reduction Step it is now possible to specify a unique name for a user exit through the use of the //UEXALIAS DD statement. This allows the user to have many different User Exits (e.g. SSA1UX00=MYEXIT01) in the same LOADLIB. For more information, refer to the CNTL Library member named DGENR200.

//PARMEMBER

The use of this DD statement allows you to specify a unique alias for the PARMLIB member being used during a specific job. For example, if the IPO1PGNS member is to be used in multiple Job Steps, but with modifications, it is now possible to specify, through the PARMEMBR input stream, IPO1PGNS=MYPGNS01 in the first step, IPO1PGNS=MYPGNS02 in the second step and so forth. For more information, refer to the CNTL Library member named DWKLTGRT.

Output Specification

Once a record satisfies the inclusion tests, based upon the //SYSIN parameters, the data to be included in the output is determined automatically. The data elements printed are:

- Output Element Key
Taken from the User Table definition
- Date
Taken from the input record date stamp
- Time
Taken from the input record time stamp
- First eight output element cells
Taken from the User Table definition

Fields taken from the User Table Definition **are not** taken from the contents of the User Table cells. The values presented to the Ad Hoc processor are taken from the input record and operated on according to the formulas specified in the table output cell definitions.

When generating an Ad Hoc report during a table update run, consideration of the types of data to be included in the report must be made and, depending upon the specification of the first eight output cells, steps can be required to limit the use of the first eight cells to fields to be included in the Ad Hoc report. Any of the first eight cells that are not to be included in the Ad Hoc report should be set to null values (zero) with no column titles specified by the OUTPUT ELEMENT CELL definitions.

Batch Reports

Once the user tables are built, you can produce several types of batch reports based on the values in the tables. The tables themselves contain self-describing dictionary data, so it is not necessary to define the makeup of the tables to the batch reporting program (SSA1GENR). The reports thus far defined to the batch reporting program are as follows:

Statistics Report

The Generic Subsystem Statistics report shows the top 40 elements in each selected category. The category can consist of a single user table cell or a value derived by applying an algorithm against a set of user table cells and/or literals. For example, presume that the user table consists of several elements whose key is a CICS USERID. Presume further that each USERID is associated with the following values in a cell set (one cell set for each hour in a 24 hour period):

- A cumulative volume of transactions
- The cumulative response time of the transactions
- A cumulative volume of EXCPs
- A cumulative value of Page operations
- The high water mark in storage for any given element

The statistics report could be used to determine the top 40 USERIDs in one or more of the above categories.

Executing the Generic Statistics Report

The Generic Statistics Report is executed by specifying **GENERIC STATISTICS REPORT=YES** in the **//SYSIN** dataset and supplying the appropriate **OUTPUT ELEMENT CELL** definitions under the **//TABLEDEF DD** statement. (see following sample job stream for details)

```
//SSA1GENR JOB (...),CLASS=A,MSGCLASS=X
//STEP01 EXEC PGM=SSA1GENR,REGION=5000K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE.GENERIC,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX.GENERIC,DISP=SHR
//CPPRERT DD DSN=&REFIX.CPPR.Vnnn.CPPRERT,DISP=SHR
//CPPRPARAM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
BEGIN DATE=07/01/2005
END DATE=07/30/2005
```

```

GENERIC STATISTICS REPORT=YES
/*
//TABLEDEF DD *
*   Specify the user table number containing the report data
USER TABLE NUMBER=144   /* X'90' */
*
* User tables are self describing.  It is not necessary to describe
* the incoming table.  Cell numbers refer to the number assigned to the
* cells at the time that the table was built.  (See dgenwkld in the
* CIMS Capacity Planner cntl library).
*****
*   The following statements define the output cells to be contained *
*   on the report.                                               *
*
*           output cell number (must be unique)                 *
*           |           output value format (see table)         *
*           |           |   formula for deriving value reported *
*           |           |   | on the statistics report          *
*           |           |   |           title (up to 12 characters *
*           |           |   |           | to be used as column heading *
OUTPUT ELEMENT CELL=01,000,00,07,((03)),FREQ_CNT
OUTPUT ELEMENT CELL=02,000,00,03,((01)),JES_QUEUE

```

Daily Overview

The Generic Daily Overview shows, for each hour of the day for each day specified in the date range, the cumulative values as specified in the OUTPUT ELEMENT CELL key phrase(s) in the input stream. Because INCLUDE or EXCLUDE processing applies, it is possible to see the actual values for a specified element for a particular hour of a particular day in this report. You can combine fields contained within the specified CIMS Capacity Planner or User Tables through the use of formulas specified as part of the OUTPUT CELL definitions. Only those fields specified by an OUTPUT CELL definitions are included in the report. Up to eight output cells can be specified for inclusion in the Daily Overview Report, each of which can have a column heading of up to 12 characters in length.

Executing the Daily Overview Report

The Daily Overview Report is executed by specifying GENERIC DAILY OVERVIEW=YES in the //SYSIN dataset and supplying the appropriate OUTPUT ELEMENT CELL definitions under the //TABLEDEF DD statement. (see following sample job stream for details)

```

//SSA1GENR JOB (...),CLASS=A,MSGCLASS=X
//STEP01 EXEC PGM=SSA1GENR,REGION=5000K,TIME=10
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR

```

```

//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE.GENERIC,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX.GENERIC,DISP=SHR
//CPPRERT DD DSN=&REFIX.CPPR.Vnnn.CPPRERT,DISP=SHR
//CPPRPARAM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
BEGIN DATE=07/01/2005
END DATE=07/01/2005
GENERIC DAILY OVERVIEW=YES
/*
//TABLEDEF DD *
* specify the user table number containing the report data
USER TABLE NUMBER=144 /* X'90' */
*
* user tables are self describing. it is not necessary to describe
* the incoming table. cell numbers refer to the number assigned to the
* cells at the time that the table was built. (see dgenwkld in the
* CIMS Capacity Planner cntl library).
*****
* The following statements define the output cells to be contained *
* on the report. *
* output cell number (must be unique) *
* | output value format (see table) *
* | | formula for deriving value reported *
* | | | on the statistics report *
* | | | title (up to 12 characters *
* | | | | to be used as column heading *
OUTPUT ELEMENT CELL=01,000,00,07,((03)),FREQ_CNT
OUTPUT ELEMENT CELL=02,000,00,03,((01)),JES_QUEUE
/*
*****

```

Generic Element Activity List

The Generic Element Activity List processes the specified user table and makes each key and all data stored under each element key available for inclusion into the Generic Element Activity List. As in the other types of reports, the various cells stored under each key are selected for inclusion in the report based upon the specification of the OUTPUT ELEMENT CELLS. All element keys and the specified cells are selected for reporting unless the element keys are excluded through the use of the INCLUDE/EXCLUDE facility.

The values are displayed in columns, one value set per column. Up to eight value sets can be displayed in a single report. The report is normally produced in ascending Element Key sequence (for example, JOBNAME). You can, however, alter this order using one of the following two key phrases:

ASCENDING SORT COLUMN=nn

or

DESCENDING SORT COLUMN=nn

where the "nn" is the OUTPUT ELEMENT CELL number as specified in the TABLEDEF input stream. This causes the sort to be ordered, either in ascending or in descending sequence, based on the values in column "nn".

The titles of all the report columns except the one containing the Element Key are defined through the OUTPUT ELEMENT CELL definitions. A keyword parameter (GENERIC USER KEY LABEL=llllllll), supplied through the //SYSIN data set, is provided to label the column containing the element keys with a user-specified label of up to eight characters.

Because the OUTPUT ELEMENT CELL definitions let you define formulas, the Generic Element Activity List provides a powerful tool for producing a variety of reports based upon the data contained within the user tables or the various CIMS Capacity Planner maintained tables. The report columns can contain the results of calculations performed upon the data contained within the tables (within the same element key) as well as the contents of the various cells.

The Generic Element Activity List report is selected by including the keyword parameter

GENERIC ELEMENT ACTIVITY LIST=YES

in the job stream.

Executing the Generic Element Activity List

The Generic Element Activity List is executed by specifying GENERIC ELEMENT ACTIVITY LIST=YES in the //SYSIN dataset and supplying the appropriate OUTPUT ELEMENT CELL definitions under the //TABLEDEF DD statement. (see following sample job stream for details)

```
//SSA1GENR JOB (...),CLASS=A,MSGCLASS=X  
//STEP01 EXEC PGM=SSA1GENR,REGION=5000K,TIME=10  
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
```

```

//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE.GENERIC,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX.GENERIC,DISP=SHR
//CPPRERT DD DSN=&REFIX.CPPR.Vnnn.CPPRERT,DISP=SHR
//CPPRPARAM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
BEGIN DATE=07/01/2005
END DATE=07/01/2005
GENERIC ELEMENT ACTIVITY LIST=YES
/*
//TABLEDEF DD *
* specify the user table number containing the report data
USER TABLE NUMBER=149 /* X'95' */
*
* user tables are self describing. it is not necessary to describe
* the incoming table. cell numbers refer to the number assigned to the
* cells at the time that the table was built. (see dgenwkld in the
* CIMS Capacity Planner cntl library).
*****
* The following statements define the output cells to be contained *
* on the report. *
* output cell number (must be unique) *
* | output value format (see table) *
* | | formula for deriving value reported *
* | | | on the statistics report *
* | | | title (up to 12 characters *
* | | | to be used as column heading *
OUTPUT ELEMENT CELL=01,000,00,00,((01)),VR_COUNT
OUTPUT ELEMENT CELL=02,000,00,00,((02)),VR_TOTAL
OUTPUT ELEMENT CELL=03,000,00,00,((03)),VR_T1
OUTPUT ELEMENT CELL=04,000,00,00,((04)),VR_T2
OUTPUT ELEMENT CELL=05,000,00,00,((05)),VR_T3
OUTPUT ELEMENT CELL=06,000,00,00,((06)),VR_T4
/*
*****

```


Generating Reports From CIMS Capacity Planner Maintained Tables

You can use the Generic Subsystem to generate user-defined reports from the CIMS Capacity Planner tables in the Performance Database. In using the CIMS Capacity Planner tables, the CIMS Capacity Planner table number is specified using the USER TABLE NUMBER= parameter, just as in the case where a userdefined table is processed. When the Generic report program (SSA1GENR) senses that a CIMS Capacity Planner table number, as opposed to a user table number has been specified, it assigns the standard CIMS Capacity Planner cell numbers to each of the cells in the table. In defining the OUTPUT ELEMENT CELLS, you must specify the cells to be extracted from the CIMS Capacity Planner table using the standard CIMS Capacity Planner cell numbers in the Formula portion of the OUTPUT ELEMENT CELL definitions. You can find a list of the CIMS Capacity Planner tables, the cell numbers, and the data formats in the UTILITIES Chapter of the *CIMS Capacity Planner Reference Guide* under [Extracting Data from the Performance Database \(SSA1TBLX\)](#).

Generic Subsystem JCL (Data Reduction)

The operating JCL has been designed to conform to system utility JCL, following IBM conventions wherever possible. Input comes from the Consolidated Table File (ON-LINE).

The EXEC statement references:

PGM=SSA1GENW for the Data Reduction program

The DD Statements required for program execution are described below.

SYSPRINT

This describes the report output file. Unless otherwise specified, the DCB characteristics are as follows:

RECFM=FBA

BLKSIZE=133

LRECL=133

SYSUT1

This DD Statement describes the input file unless the input comes from the live SMF clusters (SMFILE=LIVE) or from the Active SMF cluster (SMFILE=ACTIVE).

CPPRERT

This DD Statement describes the CIMS Capacity Planner Element Registration Table data set.

CPPRPARM

This DD Statement describes the CIMS Capacity Planner Parameter Library.

INDEX

This DD Statement describes the CIMS Capacity Planner ON-LINE Table Index data set. The INDEX data set greatly reduces the amount of time required by the Trends Analysis Programs.

ONLINE

This DD Statement describes the permanent file used to save activity records.

SYSNAP

This DD Statement describes an optional print file for diagnostic hexadecimal snap dumps

SYSUT3

This DD Statement describes the scratch file that is used to process all changes to the ONLINE file before the changes are committed back into the ONLINE file. It must be at least as large as the ONLINE file.

SYSIN

This is the parameter input file. For detailed information regarding the input parameters, see the section following.

INPUT PARAMETER STATEMENTS FOR DATA REDUCTION

SELECTED SYSTEM

This parameter is used to identify the SID of the System for which SMF/RMF data is to be analyzed. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Parameters:

This parameter has two options.

- The SID itself (e.g., SELECTED SYSTEM=MVSA)
- An '*', indicating the SID for the system on which the program is being executed (e.g., SELECTED SYSTEM=*)

Default Parameter:

SELECTED SYSTEM=*

SMFILE

This parameter is used to identify the general stored format and source of the input records that are to be processed by SSA1GENW.

Parameters:

This parameter has several options.

- SMFILE=VLR (standard VLR format)
- SMFILE=FIXED (fixed blocked input)
- SMFILE=HISTORY (input from the SMF history file)
- SMFILE=LIVE (input from the Live SMF clusters)
- SMFILE=ACTIVE (input from the Active SMF cluster)

Default Parameter:

SMFILE=HISTORY

BEGIN DATE

This parameter is used to specify the earliest date for which an input record is to be selected.

Parameters:

Specify a beginning date in one of the following forms:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (European Standard Gregorian format)

* or *-n (relative date where * is today)

Note • All dates can be specified in one of the above formats.

Default Parameter:

BEGIN DATE=00001

END DATE

This parameter is used to specify the last date for which an input record is to be selected.

Parameters:

Specify an ending date.

Default Parameter:

END DATE=2099365

BEGIN TIME

This parameter is used to specify a time of the day before which no data is to be selected for analysis. For example, if specific shift-related activity were being analyzed, the beginning of the shift would be specified by BEGIN TIME and the end of the shift would be specified by END TIME.

Parameters:

The beginning time, in the form HH.MM.SS

Default Parameter:

BEGIN TIME=00.00.00

END TIME

This parameter is used to specify a time of the day after which no data is to be selected for analysis. See the discussion of BEGIN TIME above.

Parameters:

The ending time, in the form HH.MM.SS

Default Parameter:

END TIME=24.00.00

GENERIC ELEMENT KEY

This parameter is used to select specific input records for the Ad Hoc report based on Element Key.

Parameters:

The element key or a fragment of the element key followed by an asterisk (*).

Default Parameter:

none

GEOGRAPHIC LOCATION

This parameter is used to determine the format of printed numerics (commas and decimal points) as well as the format of the date that is printed within the reports.

Parameters:

GEOGRAPHIC LOCATION=USA (date=mm/dd/yy)

GEOGRAPHIC LOCATION=EUROPE (date=dd.mm.yy)

GEOGRAPHIC LOCATION=GREAT BRITAIN (date=dd/mm/yy)

Default Parameter:

GEOGRAPHIC LOCATION=USA

LOCAL HOLIDAYS

This parameter tells the program whether to exclude holidays when preparing a set of reports. It is used in conjunction with the CSECT named "SSA1520", which contains a list of all local holidays.

Parameters:

LOCAL HOLIDAYS=EXCLUDE (don't include holidays) or

LOCAL HOLIDAYS=INCLUDE

Default Parameter:

LOCAL HOLIDAYS=INCLUDE

DEBUGON

This parameter is used to request SNAP dumps at various points in the program. It should be used only if problem determination is required.

Parameters:

None.

Default Parameter:

none

SAMPLE JOB STREAMS FOR DATA REDUCTION

The Job stream shown below could be used to save JOB-related data from the SMF Type 30 record, Subtype 5.

```
//SSADGEN JOB (...), 'SSA', CLASS=A, MSGCLASS=X
//STA EXEC PGM=SSA1GENW, REGION=5000K, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE, DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.Vnnn.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB, DISP=SHR
//HGDLIB DD DSN=&PREFIX.CPPR.Vnnn.HGDLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
SMFILE=ACTIVE
```

```

* THE FOLLOWING TWO STATEMENTS CAUSE AN AD HOC TO BE PRODUCED FOR
* ALL ELEMENTS (IN THIS CASE, JOB NAMES) BEGINNING WITH "A"
GENERIC ELEMENT KEY=A*
TITLE=JOB THROUGHPUT ANALYSIS AD HOC REPORT
//TABLEDEF DD *
*****
* THIS INPUT FILE IS USED TO DESCRIBE THE FORMAT OF THE INPUT
* AS WELL AS THE LAYOUT AND CONTENTS OF THE USER TABLE
* THE EXAMPLE SHOWN DESCRIBES VARIOUS FIELDS IN THE SMF TYPE 30
* RECORD, SUBTYPE 5 (JOB TERMINATION).
*****
* FIRST WE IDENTIFY THE SYSTEM
SMFSID OFFSET=14
* NEXT WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=144 /* X'90' */
* NEXT WE NOTE THE LOCATION OF THE RECORD NUMBER
RECORD NUMBER OFFSET=5
RECORD NUMBER LENGTH=1
RECORD NUMBER=30
* NEXT WE NOTE THE LOCATION OF THE RECORD SUB-TYPE
RECORD SUBTYPE SEGMENT POINTER OFFSET=24
RECORD SUBTYPE SEGMENT POINTER LENGTH=4
RECORD SUBTYPE OFFSET=0
RECORD SUBTYPE LENGTH=2
* NEXT WE NOTE THE LOCATION OF THE CIMS Capacity Planner TABLE KEY (JOBNAME)
ELEMENT KEY SEGMENT POINTER OFFSET=32
ELEMENT KEY SEGMENT POINTER LENGTH=4
ELEMENT KEY OFFSET=0
ELEMENT KEY LENGTH=8
* NEXT WE DESCRIBE THE INPUT CELLS OF INTEREST
* NOTE THAT ALL CELLS ARE CONTAINED IN THE SUBTYPE 5 RECORD ONLY
* INPUT CELL NUMBER (MUST BE UNIQUE)
* | INCOMING RECORD SUBTYPE
* | | INCOMING RECORD SUBTYPE QUALIFIER
* | | | SEGMENT OFFSET LOCATION IN RCD HDR
* | | | | SEGMENT OFFSET LENGTH
* | | | | | VALUE FIELD OFFSET INTO SEGMENT
* | | | | | | VALUE FIELD LENGTH

```

```

*           | | | | | VALUE FIELD TYPE
*           | | | | | | |
INPUT ELEMENT CELL=01,05,00,0,0,6,4,50 /* TIME FROM SMF HEADER
INPUT ELEMENT CELL=02,05,00,0,0,10,4,60 /* DATE FROM SMF HEADER
INPUT ELEMENT CELL=03,05,00,0,0,0,0,0 /* USED FOR FREQUENCY
INPUT ELEMENT CELL=04,05,00,32,4,56,4,10 /* PROGRAM INIT TIME
INPUT ELEMENT CELL=05,05,00,32,4,64,4,10 /* JES READER TIME
INPUT ELEMENT CELL=06,05,00,40,4,4,4,0 /* TOTAL EXCP COUNT
INPUT ELEMENT CELL=07,05,00,56,4,4,4,0 /* TCB CPU TIME
INPUT ELEMENT CELL=08,05,00,56,4,8,4,0 /* SRB CPU TIME
INPUT ELEMENT CELL=09,05,00,88,4,8,4,0 /* TAPE MOUNTS
INPUT ELEMENT CELL=10,05,00,88,4,12,4,0 /* TAPE MOUNTS
INPUT ELEMENT CELL=11,05,00,56,4,12,4,0 /* INIT TCB CPU
INPUT ELEMENT CELL=12,05,00,56,4,16,4,0 /* INIT SRB CPU
*****
* NEXT WE DESCRIBE THE OUTPUT CELLS TO BE CREATED IN THE CIMS Capacity Planner TABLE
*           OUTPUT CELL NUMBER (MUST BE UNIQUE)
*           | OUTPUT VALUE FORMAT
*           | |FORMULA FOR DERIVING OUTPUT CELL VALUE
*           | | 12 CHAR VALUE LABEL
*           | | | |
OUTPUT ELEMENT CELL=01,0,0,03,((04-05)),IN JES QUEUE
OUTPUT ELEMENT CELL=02,0,0,03,((01-05)), THROUGHPUT
OUTPUT ELEMENT CELL=03,0,0,04,((03)), FREQUENCY
OUTPUT ELEMENT CELL=04,0,0,00,((06)), TOTAL EXCPS
OUTPUT ELEMENT CELL=05,0,0,07,((07+08+11+12)), CPU SECONDS
OUTPUT ELEMENT CELL=06,0,0,00,((09+10)), TAPE MOUNTS
OUTPUT ELEMENT CELL=07,0,0,00,((06/07)),RANDOM VALUE
OUTPUT ELEMENT CELL=08,0,0,06,(((04-05)*'10000')/(01-05)), INQ:THRUPUT

```

The following **JOBSTREAM** shows a segmented record being processed into a CIMS Capacity Planner user table. The record is a **CL/SUPERSESSION** record from Candle Corporation.

```

//SSADGEN JOB (...),'SSA',CLASS=A,MSGCLASS=X
//STA EXEC PGM=SSAIGENW,REGION=5000K,TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.Vnnn.CPPRERT,DISP=SHR

```



```
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB,DISP=SHR
//HGDLIB DD DSN=&PREFIX.CPPR.Vnnn.HGDLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
SMFILE=ACTIVE
GENERIC ELEMENT KEY=A*
TITLE=CL/SUPERSESSION SESSION LOGOFF REPORT
//TABLEDEF DD *
* FIRST WE IDENTIFY THE SYSTEM
SMFSID OFFSET=14
* NEXT WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=145 /* X'91' */
* THESE ARE LINKED LENGTHS. DEFINE THE LENGTH VECTORS.
RECORD TOTAL LENGTH FIELD OFFSET=0
RECORD TOTAL LENGTH FIELD LENGTH=2
RECORD HEADER LENGTH=28
RECORD SEGMENT FIXED LENGTH=10
RECORD SEGMENT VARIABLE LENGTH FIELD OFFSET=9
RECORD SEGMENT VARIABLE LENGTH FIELD LENGTH=1
* NEXT WE NOTE THE LOCATION OF THE RECORD NUMBER IN THE HEADER
RECORD NUMBER OFFSET=5
RECORD NUMBER LENGTH=1
RECORD NUMBER=254
* NEXT WE NOTE THE LOCATION OF THE RECORD SUB-TYPE
RECORD SUBTYPE SEGMENT POINTER OFFSET=0
RECORD SUBTYPE SEGMENT POINTER LENGTH=0
RECORD SUBTYPE OFFSET=0
RECORD SUBTYPE LENGTH=1
* NEXT WE NOTE THE LOCATION OF THE CIMS Capacity Planner TABLE KEY
ELEMENT KEY SEGMENT POINTER OFFSET=0
ELEMENT KEY SEGMENT POINTER LENGTH=0
ELEMENT KEY OFFSET=10
ELEMENT KEY LENGTH=8
* NEXT WE DESCRIBE THE INPUT CELLS OF INTEREST (CL/SUPERSESSION)
* NOTE THAT ALL CELLS ARE CONTAINED IN THE SUBTYPE 227 RECORD ONLY
```

```
INPUT ELEMENT CELL=01,227,0,0,0,1,4,50 /* TIME FROM SEGMENT HEADER
INPUT ELEMENT CELL=02,227,0,0,0,5,4,60 /* DATE FROM SEGMENT HEADER
INPUT ELEMENT CELL=03,227,0,0,0,0,0,0 /* USED FOR FREQUENCY
INPUT ELEMENT CELL=04,227,0,0,0,48,4,0 /* INBOUND MESSAGE COUNT
INPUT ELEMENT CELL=05,227,0,0,0,52,4,0 /* INBOUND BYTE COUNT
INPUT ELEMENT CELL=06,227,0,0,0,40,4,0 /* OUTBOUND MESSAGE COUNT
INPUT ELEMENT CELL=07,227,0,0,0,44,4,0 /* OUTBOUND BYTE COUNT
INPUT ELEMENT CELL=08,227,0,0,0,56,4,0 /* REFRESH MESSAGE COUNT
INPUT ELEMENT CELL=09,227,0,0,0,60,4,0 /* REFRESH BYTE COUNT
INPUT ELEMENT CELL=10,227,0,0,0,64,4,0 /* UPDATE MESSAGE COUNT
INPUT ELEMENT CELL=11,227,0,0,0,68,4,0 /* UPDATE BYTE COUNT
INPUT ELEMENT CELL=12,227,0,0,0,72,4,0 /* TERMINAL INPUT MESSAGE COUNT
INPUT ELEMENT CELL=13,227,0,0,0,76,4,0 /* TERMINAL INPUT BYTE COUNT
* NEXT WE DESCRIBE THE OUTPUT CELLS TO BE CREATED IN THE CIMS Capacity Planner TABLE
OUTPUT ELEMENT CELL=01,0,0,04,((03)),LOGON COUNT
OUTPUT ELEMENT CELL=02,0,0,00,((04)),INPUT MSGS
OUTPUT ELEMENT CELL=03,0,0,00,((06)),OUTPUT MSGS
OUTPUT ELEMENT CELL=04,0,0,00,((11/07)),CMPRS RATIO
OUTPUT ELEMENT CELL=05,0,0,00,((05+07)),APPL<=>SS
OUTPUT ELEMENT CELL=06,0,0,00,((09+11)),SS<=>TERM
OUTPUT ELEMENT CELL=07,0,0,00,((07/06)),OUTMSGLEN
OUTPUT ELEMENT CELL=08,0,0,06,(((11)*'10000')/(07)),CMPRS PCT
```

OPERATING JCL FOR BATCH REPORTS

The operating JCL has been designed to conform to system utility JCL, following IBM conventions wherever possible. Input comes from the Consolidated Table File (ONLINE).

The EXEC statement references:

PGM=SSA1GENR for the Batch Report program

The DD Statements required for program execution are described below.

SYSPRINT

This describes the report output file. Unless otherwise specified, the DCB characteristics are as follows:

RECFM=FBA

BLKSIZE=133

LRECL=133

CPPRERT

This DD Statement describes the CIMS Capacity Planner Element Registration Table data set.

CPPRPARM

This DD Statement describes the CIMS Capacity Planner Parameter Library.

INDEX

This DD Statement describes the CIMS Capacity Planner ONLINE Table Index data set. The INDEX data set greatly reduces the amount of time required by the Trends Analysis Programs.

ONLINE

This DD Statement describes the permanent file used to save activity records.

EXCLUDE/INCLUDE

■ EXCLUDE

This is the (optional) Element Exclusion File. It is used by the Batch Report Program (SSA1GENR) to eliminate elements from the Reports that you do not want to have included. The format of the input statements is as follows:

- An element can be up to eight characters in length
- Multiple elements can be included in a single record
- Elements included in a single record must be separated by commas
- Elements can be terminated by a "wild card" (an *), indicating that all elements that match the characters up to the "wild card" are to be excluded from the reports.

For example, the following input statement to the Exclusion file:

```
TSOP,CS*
```

would indicate that the Reports would not include data with respect to the element named "TSOP" nor would they include any data reflecting the use of any element beginning with the characters "CS".

The exclusion file must have an LRECL of 80. Only the first 72 columns of the statement are interpreted, leaving columns 73-80 for sequence numbers.

■ INCLUDE

This is the (optional) Element Inclusion File. It is used to limit the elements included into the graphs to the elements specified in the Include list. The format of the input statements is as follows:

- An element can be up to eight characters in length
- Multiple elements can be included in a single record
- Elements included in a single record must be separated by commas
- Elements can be terminated by a "wildcard" (an *), indicating that all elements that match the characters up to the "wildcard" are to be included (if present) into the reports.

For example, the following input statement to the Inclusion file:

```
TSOP,CS*
```

would indicate that the graphs would include only data with respect to the element named "TSOP" and any data associated with any element beginning with the characters "CS".

The inclusion file must have an LRECL of 80. Only the first 72 columns of the statement are interpreted.

SYSNAP

This DD Statement describes an optional print file for diagnostic hexadecimal snap dumps

SYSIN

This is the parameter input file. For detailed information regarding the input parameters, see the section following.

INPUT PARAMETER STATEMENTS FOR BATCH REPORTS

SELECTED SYSTEM

This parameter is used to identify the SID of the System for which SMF/RMF data is to be analyzed. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Parameters:

This parameter has two options.

- The SID itself (e.g., SELECTED SYSTEM=MVSA)
- An '*', indicating the SID for the system on which the program is being executed (e.g., SELECTED SYSTEM=*)

Default Parameter:

```
SELECTED SYSTEM=*
```

GENERIC STATISTICS REPORT

This parameter is used to request the Generic Statistics Report.

Parameters:

GENERIC STATISTICS REPORT=YES (produce the report)

or

GENERIC STATISTICS REPORT=NO (Don't produce the report)

Default Parameter:

GENERIC STATISTICS REPORT=NO

GENERIC DAILY OVERVIEW

This parameter is used to request the Generic Daily Overview report.

Parameters:

GENERIC DAILY OVERVIEW=YES (produce the report)

or

GENERIC DAILY OVERVIEW=NO (Don't produce the report)

Default Parameter:

GENERIC DAILY OVERVIEW=NO

GENERIC ELEMENT ACTIVITY LIST

This parameter is used to request the Generic Element Activity List.

Parameters:

GENERIC ELEMENT ACTIVITY LIST=YES (produce the report)

or

GENERIC ELEMENT ACTIVITY LIST=NO (Don't produce the report)

Default Parameter:

GENERIC ELEMENT ACTIVITY LIST=NO

BEGIN DATE

This parameter is used to specify the earliest date for which an input record is to be selected.

Parameters:

Specify a beginning date in one of the following forms:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (European Standard Gregorian format)

* or *-n (relative date where * is today)

Note • All dates can be specified in one of the above formats.

Default Parameter:

BEGIN DATE=00001

END DATE

This parameter is used to specify the last date for which an input record is to be selected.

Parameters:

Specify an ending date.

Default Parameter:

END DATE=2099365

BEGIN TIME

This parameter is used to specify a time of the day before which no data is to be selected for analysis. For example, if specific shift-related activity were being analyzed, the beginning of the shift would be specified by BEGIN TIME and the end of the shift would be specified by END TIME.

Parameters:

The beginning time, in the form HH.MM.SS

Default Parameter:

BEGIN TIME=00.00.00

END TIME

This parameter is used to specify a time of the day after which no data is to be selected for analysis. See the discussion of BEGIN TIME above.

Parameters:

The ending time, in the form HH.MM.SS

Default Parameter:

END TIME=24.00.00

PRIME SHIFT FIRST HOUR

This parameter is used to specify the first hour of Prime Shift for reporting purposes.

Parameters:

Specify an hour.

Default Parameter:

PRIME SHIFT FIRST HOUR=6

LATE SHIFT FIRST HOUR

This parameter is used to specify the first hour of Late Shift for reporting purposes.

Parameters:

Specify an hour.

Default Parameter:

LATE SHIFT FIRST HOUR=18

SELECTED DAY

This parameter is used to specify a specific day of the week for which reports are to be generated. The day is specified by name (e.g., MONDAY). Multiple days can be specified, separated by commas (e.g., MONDAY,WEDNESDAY).

Parameters:

Specify one or more names of days, or ALL DAYS.

Default Parameter:

SELECTED DAY=ALL DAYS

GEOGRAPHIC LOCATION

This parameter is used to determine the format of printed numerics (commas and decimal points) as well as the format of the date that is printed within the reports.

Parameters:

GEOGRAPHIC LOCATION=USA (date=mm/dd/yy)

GEOGRAPHIC LOCATION=EUROPE (date=dd.mm.yy)

GEOGRAPHIC LOCATION=GREAT BRITAIN (date=dd/mm/yy)

Default Parameter:

GEOGRAPHIC LOCATION=USA

DEBUGON

This parameter is used to request SNAP dumps at various points in the program. It should be used only if problem determination is required.

Parameters:

None.

Default Parameter:

none

LOCAL HOLIDAYS

This parameter tells the program whether to exclude holidays when preparing a set of reports. It is used in conjunction with the CSECT named "SSA1520", which contains a list of all local holidays.

Parameters:

LOCAL HOLIDAYS=EXCLUDE (don't include holidays) or

LOCAL HOLIDAYS=INCLUDE

Default Parameter:

LOCAL HOLIDAYS=INCLUDE

TABLEDEF STATEMENTS FOR BATCH REPORTS

The output elements in the batch reports are described through the TABLEDEF input file. The table itself is identified by the table number assigned through SSA1GENW. Each output element to be included in the report is described by the "OUTPUT ELEMENT CELL" key phrases.

USER TABLE NUMBER

This key phrase tells the Generic Subsystem which CIMS Capacity Planner User Table number is assigned to this table. Table numbers 128 (x'80') through 159 (x'9F') are reserved for CIMS Capacity Planner User Tables.

OUTPUT ELEMENT CELL

This key phrase describes a specific Output Element Cell that is to be included in the report. The format of the argument for this key phrase is positional. The parameters are specified as follows:

- A field identifying the Output Cell Number (01-16)
- A reserved field
- A reserved field
- A field that describes the format of the value being described. The following value formats are supported by the Generic Subsystem:

Table 8-3 • Table of output cell value type definitions

Value Formats	Normal
Binary Value	00
Packed Decimal Value	01
Character Value	02
Time Value (.xx Secs)	03
Frequency Counter	04
Percentage (xx.xx%)	06
.xx Binary Value	07
Save Highest Value (.xx value)	08 (e.g., high water mark)
Average (.xx value)	09
Save Lowest Value (Not 0)	0A (e.g., low water mark)

- A formula describing the algorithm for deriving the output cell value. The formula consists of input cell numbers and/or literals plus mathematical operators. The following rules apply to the formula:
 - 1 Up to 12 levels of parentheses are supported
 - 2 Numbers refer to Input Cell Numbers
 - 3 Literals are enclosed in single quotes ("")
 - 4 Operators supported are as follows:
 - (+) A + B Means add A to B and save the sum
 - (-) A - B Means subtract B from A and save the difference
 - (/) A / B Means divide A by B and return the integer quotient
 - (*) A * B Means multiply A by B and save the product
 - (:) A : B Means divide A by B and return the remainder (Modulo)
- A twelve byte label for the resultant value

Note that, for a Performance-Against-Objectives report, the first "OUTPUT ELEMENT CELL" statement must describe a volume value and the second "OUTPUT ELEMENT CELL" statement must describe a cumulative elapsed time value unless the OUTPUT ELEMENT CELL statement describes a Comparative Cell Set.

SAMPLE JOB STREAMS FOR BATCH REPORTS

```
//SSAINIT JOB (...), 'SSA', CLASS=A, MSGCLASS=X
//STA EXEC PGM=SSA1GENR, REGION=5000K, TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.Vnnn.ONLINE, DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.Vnnn.INDEX, DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.Vnnn.CPPRERT, DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.Vnnn.PARMLIB, DISP=SHR
//HGDLIB DD DSN=&PREFIX.CPPR.Vnnn.HGDLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=*
BEGIN DATE=01/01/2005
END DATE=01/30/2005
GENERIC STATISTICS REPORT=YES
GENERIC PERFORMANCE REPORT=YES
```

```
//TABLEDEF DD *
*****
*   FIRST WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=144      /* X'90' */
*****
*   FOR THE PERFORMANCE REPORT, THE FIRST OUTPUT CELL MUST BE A   *
*   FREQUENCY VALUE AND THE 2ND CELL MUST BE A CUMULATIVE TIME VALUE
*****
*   NEXT WE DESCRIBE THE OUTPUT CELLS TO BE CREATED IN THE REPORT
*
*           OUTPUT CELL NUMBER (MUST BE UNIQUE)
*           | OUTPUT VALUE FORMAT
*           | |FORMULA FOR DERIVING OUTPUT CELL VALUE
*           | || 12 CHAR VALUE LABEL
*           | | | |
OUTPUT ELEMENT CELL=01,0,0,07,((03)),FREQ_CNT
OUTPUT ELEMENT CELL=02,0,0,03,((01)),JES_QUEUE
```

PRODUCING BASIC PC GRAPHICS DATA POINT MEMBERS

Basic graph data point members reflect the values associated with a single selected user table. They can be constructed for Hourly, Daily, Weekly or Monthly graphs. The values thus constructed can include all elements in the table or be limited to specific elements that can be Included or Excluded (see the JCL section for details).

OPERATING JCL FOR BASIC GRAPH DATA POINTS

The operating JCL has been designed to conform to system utility JCL, following IBM conventions wherever possible. Input comes from the Consolidated Table File (ONLINE).

The EXEC statement references:

PGM=SSA1GENG for the basic graphing program

The DD Statements required for program execution are described below.

SYSPRINT

This describes the report output file. Unless otherwise specified, the DCB characteristics are as follows:

RECFM=FBA

BLKSIZE=133

LRECL=133

CPPRERT

This DD Statement describes the CPPR Element Registration Table dataset.

CPPRPARM

This DD Statement describes the CPPR Parameter Library.

INDEX

This DD Statement describes the CPPR ONLINE Table Index dataset. The INDEX data set greatly reduces the amount of time required by the Trends Analysis Programs.

ONLINE

This DD Statement describes the permanent file used to save activity records.

HGDLIB

This DD Statement describes the Partitioned Data set used to save the members produced by the Graphics programs (SSA1GENG and SSA1GENE). These members are later down loaded to a Personal Computer and used as input to the Harvard Graphics (or other graphing) program to build the actual graphs.

EXCLUDE/INCLUDE

■ EXCLUDE

This is the (optional) Element Exclusion File. It is used to eliminate elements from the graphs that you do not want to have included. The format of the input statements is as follows:

- An element can be up to eight characters in length
- Multiple elements can be included in a single record
- Elements included in a single record must be separated by commas
- Elements can be terminated by a "wildcard" (an *), indicating that all elements that match the characters up to the "wildcard" are to be excluded from the reports.

For example, the following input statement to the Exclusion file:

```
TSOP,CS*
```

would indicate that the graphs would not include data with respect to the element named "TSOP" nor would they include any data reflecting the use of any element beginning with the characters "CS".

The exclusion file must have an LRECL of 80. Only the first 72 columns of the statement are interpreted.

■ INCLUDE

This is the (optional) Element Inclusion File. It is used to limit the elements included into the graphs to the elements specified in the Include list. The format of the input statements is as follows:

- An element can be up to eight characters in length
- Multiple elements can be included in a single record
- Elements included in a single record must be separated by commas
- Elements can be terminated by a "wildcard" (an *), indicating that all elements that match the characters up to the "wildcard" are to be included (if present) into the reports.

For example, the following input statement to the Inclusion file:

```
TSOP,CS*
```

would indicate that the graphs would include only data with respect to the element named "TSOP" and any data associated with any element beginning with the characters "CS".

The inclusion file must have an LRECL of 80. Only the first 72 columns of the statement are interpreted

SYSNAP

This DD Statement describes an optional print file for diagnostic hexadecimal snap dumps

SYSIN

This is the parameter input file. For detailed information regarding the input parameters, see the section following.

INPUT PARAMETER STATEMENTS FOR BASIC GRAPH DATA POINTS

FILTER

This parameter is used to determine the suffix of the data point member produced for basic data point members. If no FILTER value is specified, the HGDLIB member produced has a default suffix (&sid.pG0, where &sid is the SMF SID of the Selected System, and "p", the Graph Period, is 1 for HOURLY graphs, 2 for DAILY graphs, 3 for WEEKLY graphs and 4 for MONTHLY graphs.

Parameters:

The FILTER keyword has a single parameter.

- The parameter specifies the suffix for the basic graph data point member

Default Parameter:

FILTER=0

SELECTED SYSTEM

This parameter is used to identify the SID of the System for which SMF/RMF data is to be analyzed. This is the SMF System Identifier as specified in the SMFPRMxx member of SYS1.PARMLIB.

Parameters:

This parameter has two options.

- The SID itself (for example, SELECTED SYSTEM=MVSA)
- An '*', indicating the SID for the system on which the program is being executed (for example, SELECTED SYSTEM=*)

Default Parameter:

SELECTED SYSTEM=*

BEGIN DATE

This parameter is used to specify the earliest date for which an input record is to be selected.

Parameters:

Specify a beginning date in one of the following forms:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (European Standard Gregorian format)

* or *-n (relative date where * is today)

Note • All dates can be specified in one of the above formats.

Default Parameter:

BEGIN DATE=00001

END DATE

This parameter is used to specify the last date for which an input record is to be selected.

Parameters:

Specify an ending date.

Default Parameter:

END DATE=99365

PRIME SHIFT FIRST HOUR

This parameter is used to specify the first hour of Prime Shift for reporting purposes.

Parameters:

Specify an hour.

Default Parameter:

PRIME SHIFT FIRST HOUR=6

LATE SHIFT FIRST HOUR

This parameter is used to specify the first hour of Late Shift for reporting purposes.

Parameters:

Specify an hour.

Default Parameter:

LATE SHIFT FIRST HOUR=18

SELECTED DAY

This parameter is used to specify a specific day of the week for which reports are to be generated. The day is specified by name (for example, MONDAY). Multiple days can be specified, separated by commas (for example, MONDAY,WEDNESDAY).

Parameters:

Specify one or more names of days, or ALL DAYS.

Default Parameter:

SELECTED DAY=ALL DAYS

GEOGRAPHIC LOCATION

This parameter is used to determine the format of printed numerics (commas and decimal points) as well as the format of the date that is printed within the reports.

Parameters:

GEOGRAPHIC LOCATION=USA (date=mm/dd/yy)

GEOGRAPHIC LOCATION=EUROPE (date=dd.mm.yy)

GEOGRAPHIC LOCATION=GREAT BRITAIN (date=dd/mm/yy)

Default Parameter:

GEOGRAPHIC LOCATION=USA

DEBUGON

This parameter is used to request SNAP dumps at various points in the program. It should be used only if problem determination is required.

Parameters:

None.

Default Parameter:

none

LOCAL HOLIDAYS

This parameter tells the program whether to exclude holidays when preparing a set of reports. It is used in conjunction with the CSECT named "SSA1520", which contains a list of all local holidays.

Parameters:

LOCAL HOLIDAYS=EXCLUDE (don't include holidays) or

LOCAL HOLIDAYS=INCLUDE

Default Parameter:

LOCAL HOLIDAYS=INCLUDE

GRAPH PERIOD

This parameter is used to specify the period for the graph data points. The parameter has four options.

- GRAPH PERIOD=HOURLY (produce hourly graph data points)
- GRAPH PERIOD=DAILY (produce daily graph data points)
- GRAPH PERIOD=WEEKLY (produce weekly graph data points)
- GRAPH PERIOD=MONTHLY (produce monthly graph data points)

Parameters:

Specify a Graph Period.

Default Parameter:

none

TABLEDEF Statements for Basic Graph Data Points

The output elements in the graph data point members are described through the TABLEDEF input file. The table itself is identified by the table number assigned through SSA1GENW. Each output element to be included in the graph data point member is described by the "OUTPUT ELEMENT CELL" key phrases.

USER TABLE NUMBER

This key phrase tells the generic subsystem which CIMS Capacity Planner User Table number is assigned to this table. Table numbers 128 (x'80') through 159 (x'9F') are reserved for CIMS Capacity Planner User Tables.

OUTPUT ELEMENT CELL

This key phrase describes a specific Output Element Cell that is to be included in the graph data point member. The format of the argument for this key phrase is positional. The parameters are specified as follows:

- A field identifying the Output Cell Number (01-16)
- A reserved field
- A reserved field
- A field that describes the format of the value being described. The following value formats are supported by the generic subsystem:

Table 8-4 • Table of output cell value type definitions

Value Formats	Normal
Binary Value	00
Packed Decimal Value	01
Character Value	02
Time Value (.xx Secs)	03
Frequency Counter	04
Percentage (xx.xx%)	06
.xx Binary Value	07
Save Highest Value (.xx value)	08 (for example, high water mark)
Average (.xx value)	09
Save Lowest Value (Not 0)	0A (for example, low water mark)

- A formula describing the algorithm for deriving the output cell value. The formula consists of input cell numbers and/or literals plus mathematical operators. The following rules apply to the formula:
 - 1 Up to 12 levels of parentheses are supported
 - 2 Numbers refer to Input Cell Numbers
 - 3 Literals are enclosed in single quotes ("")
 - 4 Operators supported are as follows:
 - + Means add A to B and save the sum
 - Means subtract B from A and save the difference
 - / Means divide A by B and return the integer quotient
 - * Means multiply A by B and save the product
 - : Means divide A by B and return the remainder (Modulo)
- A twelve byte label for the resultant value

SAMPLE JOBSTREAMS FOR BASIC CELL GRAPH DATA POINTS

```

//SSAGENG JOB (...),'SSA',CLASS=A,MSGCLASS=X
//STA EXEC PGM=SSA1GENG,REGION=5000K,TIME=60
//STEPLIB DD DSN=&PREFIX.CPPR.V106.LOADLIB,DISP=SHR
//ONLINE DD DSN=&PREFIX.CPPR.V106.ONLINE,DISP=SHR
//INDEX DD DSN=&PREFIX.CPPR.V106.INDEX,DISP=SHR
//CPPRERT DD DSN=&PREFIX.CPPR.V106.CPPRERT,DISP=SHR
//CPPRPARM DD DSN=&PREFIX.CPPR.V106.PARMLIB,DISP=SHR
//HGDLIB DD DSN=&PREFIX.CPPR.V106.HGDLIB,DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=CPPR
GRAPH PERIOD=HOURLY
BEGIN DATE=01/01/2005
END DATE=01/30/2005
* ORDINARILY, THIS PROGRAM WILL PRODUCE A GRAPH WITH THE SUFFIX 1G0.
* IF YOU WISH TO CHANGE THE SUFFIX TO 1GN, SPECIFY THE "N" WITH
* THE "FILTER" KEYWORD THUS:
FILTER=1 /* THIS PRODUCES HGDLIB MEMBER CPPRIG1 */
//TABLEDEF DD *
*****
* FIRST WE PROVIDE THE CIMS Capacity Planner USER TABLE NUMBER
USER TABLE NUMBER=144 /* X'90' */
*****
* USER TABLES ARE SELF DESCRIBING. THUS IT IS NOT NECESSARY TO
* DESCRIBE THE INCOMING TABLE. CELL NUMBERS REFER TO THE NUMBER
* ASSIGNED WHEN THE TABLE WAS BUILT (SEE DGENWKLD)
*****
* NEXT WE DESCRIBE THE FORMULA TO BE USED TO DERIVE THE VALUES
* WHICH ARE TO BE PRESENT IN THE DATA POINT MEMBER.
* OUTPUT CELL NUMBER (MUST BE UNIQUE)
* | OUTPUT VALUE FORMAT
* | |FORMULA FOR DERIVING OUTPUT CELL VALUE
* | || 8 CHAR VALUE LABEL

```

■ GENERIC SUBSYSTEM

SAMPLE JOBSTREAMS FOR BASIC CELL GRAPH DATA POINTS

```
*           | |   |   |  
OUTPUT ELEMENT CELL=01,0,0,07,((01)/'60'),JES_QUEUE  
OUTPUT ELEMENT CELL=02,0,0,07,((02-01)/'60'),THRUPUT
```

Operations Analysis Subsystem

Subsystem Overview	9-2
Shift Turnover Report	9-2
Batch Window Analysis	9-4
What Is a Batch Window	9-4
Events Inside a Batch Window	9-5
Data Collection	9-6
Reports	9-6
Key Capacity Factors Analysis	9-9
Which Factors are Analyzed	9-10
Data Collection	9-10
Reports	9-11
Application Profile Analysis	9-13
Which Values are Maintained	9-13
Data Collection	9-14
Application Profile Report	9-14

Subsystem Overview

The Operations Analysis subsystem brings together several reports and graphs that have been developed jointly between CIMS Lab and CIMS Capacity Planner users. These reports and graphs are useful in managing the day-to-day challenges that arise during normal operation of a data center.

Shift Turnover Report

The shift turnover report is produced from a new table, Table 065, which is created during Workload data reduction (SSA1WKLD). This table contains information regarding the resources that were used by shift during a specific 24 hour day, and compares those values to the average daily work day values for the previous week, and to the average daily work day values for the previous month.

The items listed in the Shift Turnover Report are:

Available Time. This report line shows the amount of time (HH:MM) in the shift

Processor Active. This report line shows the amount of time (HH:MM) during the shift that the processor was active

Processor Idle.

Execution Time. This report line shows the cumulative Elapsed time (HH:MM) from Step Start to Step End for all Batch Job Steps during a specific shift. This value does not include Initiation time.

Throughput. This report line shows the cumulative Elapsed time (HH:MM) from Batch Job initiation through Batch Job termination for all Batch Jobs during a specific shift.

Initiation Ratio. This report line shows the percentage of time spent processing Jobs as opposed to the time spent in Initiation. It is derived by dividing Throughput into Execution Time. An initiation ratio of .98 would imply that Jobs spent 2% of their elapsed time in initiation.

Multiprogram Factor. This indicates the average number of Jobs active during a specific shift. It is derived by dividing Throughput by Processor Active time.

Number of IPLs. This report line shows the total number of times the system was IPLed during a specific shift.

Number of Runs. This line shows the number of Batch Jobs that were started during a specific shift.

Prob Prog CPU Time.	This report line shows the total amount of CPU time (HH:MM) used by all Batch Jobs during a specific shift.
System CPU Time.	This report line shows the total amount of CPU time (HH:MM) used by the System (Started Tasks, MVS overhead, TSO). It represents the difference between Processor Active and System Wait time (CPU Active Time) minus Prob Prog CPU time.
System Wait Time.	This report line shows the cumulative amount of time (HH:MM) during a specific shift that any processor was waiting for work.
PCT CPU Time.	This report line shows the percentage of time during a specific shift that the CPU was actually processing work. It is derived by dividing the CPU Active Time by the Processor Active Time.
System Pages / Min.	This report line shows the auxiliary paging activity (pages out plus pages in) per minute for all Batch Jobs processed during a specific shift.
TSO LOGONs.	This report line shows, for a specific shift, the total number of actual Logons that occurred.
TSO Connect Time.	This report line shows the total Elapsed Time (HH:MM) for all TSO sessions during a specific shift.
PCT of Throughput.	This report line shows, for a specific shift, the percentage of Batch work that was processed. It is derived by dividing the sum of the Elapsed Times (Throughput) for all shifts into the Elapsed Times (Throughput) for each shift.
ABENDs.	This report line shows, for a specific shift, the total number of Batch Jobs that terminated abnormally.
ABEND Time.	This report line shows, for a specific shift, the total amount of Elapsed time that was lost due to ABENDs .
PCT Execution Time.	This report line shows, for a specific shift, the percentage of Elapsed Time that was lost due to ABENDs.
Specific Tape Mounts.	This report line shows, for a specific shift, the total number of tape mounts issued that were referred to by Volume Serial Number.
Non-SPC Tape Mounts.	This report line shows, for a specific shift, the total number of tape mounts issued for Scratch Tapes.
Printed Pages.	This report line shows, for a specific shift, the total number of lines printed by all JES printers attached to the processor.

The report is produced by the SSA1SHAQ program. Sample JCL to produce the report is contained in the distribution CNTL library member named DACTSHAQ.

Batch Window Analysis

Processing in a mainframe environment is traditionally divided into Batch and Online workloads. In the case where users are connected directly to the machine via remote or local terminals (for example, CICS or TSO or IMS applications) the workload is generally classified as Online. Work that is performed without the direct interaction of a person is generally classified as Batch. This might include tasks such as long running mathematical calculations, DASD backups and Payroll.

Online users have specific expectations of levels of performance from the data center. These expectations are usually formalized in Service Level Agreements that spell out the parameters of performance promised to users during a specified period of time, such as Prime shift. Frequently Service Level Agreements include constraints as well as promises. That is, the Online users are promised a specified level of performance so long as they don't overload the capacity of the system.

Online performance is concerned primarily with response time. If response times are unnecessarily prolonged, production suffers. Thus, during the period when Online work is scheduled, Batch tasks normally are assigned a lower priority than Online tasks.

Most regularly scheduled Batch tasks are therefore run overnight, or off shift. The definition of off shift varies considerably by data center. Some data centers have Online applications that are available twenty four hours a day, seven days a week. Nonetheless, Batch work must get done, frequently within a relatively short period of time. Daily reports need to be available in a timely fashion, DASD backups cannot be postponed, journals must be reconciled, payments must be posted.

Whereas Online performance deals primarily with response time, Batch performance deals primarily with meeting schedules. This gives rise to the concept of a Batch Window.

What Is a Batch Window

Briefly, a Batch Window is any period of time that has the following characteristics:

- It begins when a specific JOB goes into execution. That JOB can have several different names, one for daily, another for weekly and so forth. But the execution of that JOB always signals the opening of this logical Batch Window.
- It ends when a specific JOB terminates. That JOB can have several different names, one for daily, another for weekly and so forth. But the termination of that JOB always signals the closing of this logical Batch Window.
- It includes a specified set of Jobs

Any given Batch Window can occur one or more times in a day, but a Batch Window can not be iterative. That is, a Batch Window can not include the execution of the same Batch Window inside itself.

The daily or weekly schedule of an MVS system can include several distinct logical Batch Windows, each with the above three criteria. One or more of the logical Batch Windows can be candidates for analysis.

Events Inside a Batch Window

Because the duration of the Batch Window is of paramount importance in measuring the Performance of the batch jobs inside the Batch Window, the events that affect the duration of that window should be isolated and analyzed. The following events affect the duration of the Batch Window:

- Start/Stop (elapsed) time for each individual task in the Batch Window
- Allocation of DASD data sets and Tape Drives
- Releasing of Tape Drives
- CPU time within an individual task
- DASD EXCPs within an individual task
- JOB execution overlap (M/P factor)
- Task ABEND occurrences

When the duration of the Batch Window exceeds a specific level, the Batch work is in danger of not being completed on time. An analysis of the events inside the Batch Window allows you to determine why the duration of the Batch Window has changed.

For example, if an ABEND event causes the jobs to be re-run, the elapsed time of the entire window is necessarily elongated. On the other hand, if a DASD data set is not available (is owned exclusively by another task), a specific task might be delayed, causing the duration of the Batch Window to increase. The lack of availability of Tape drives or other serially reusable resources can cause tasks inside the window to be temporarily suspended. Or an operator can put all the jobs on Hold for a period of time.

Not all events that occur while the Batch Window is open are related to the duration of the window. For example, a TSO user whose session was canceled might have nothing to do with the duration of a Batch Window that was open at the time. It is therefore important to isolate the events that relate directly to the duration of the Batch Window.

Data Collection

Several SMF records contain data that is essential to the analysis of the Batch Window. Type 30 records contain information relating to the tasks contained in the Batch Window. Type 14 and 15 records contain information relating to tape device allocation and usage. Type 74 records contain device performance statistics by device for a specified period of time. Type 72 records contain information regarding the Service Units consumed by different workloads that can affect the availability of those resources to the tasks inside the Batch Window. Type 71 records contain information about the overall use of memory resources inside the CPU. Type 70 records contain information regarding the saturation of the CPU.

The tables used during Batch Window analysis are created by SSA1WKLD when the SMF/RMF data is reduced into the Workload database. One set of tables is created for each logical Batch Window so that each window can be analyzed individually.

The CIMS Capacity Planner PARMLIB is used as a trigger during data reduction. Member names beginning with the SMF SID, followed by the characters "BW" and a two position numeric (01-99) tell SSA1WKLD to capture data for this particular Batch Window. For example, if the CIMS Capacity Planner PARMLIB contains a member named "CPPRBW01" and the SMF SID of the system for which data is being reduced is CPPR, the member is read and data is reduced for logical Batch Window number 01.

Reports

Several reports are available once the tables have been created. These reports are produced by the SSA1OPA3 program and are described below.

Batch Window Summary Analysis

This report shows, for a specific Batch Window as indicated in the //INCLUDE stream, the following values:

Average Window Size	The time, in Hours, Minutes and seconds, from Batch Window open to close
Earliest Window Open	The earliest wall clock time (hh:mm:ss) the Batch Window opened
Latest Window Close	The latest time (hh:mm:ss) the Batch Window closed
Window Elapsed Time	The cumulative elapsed times (hh:mm:ss) for all jobs in the Batch Window
Window INIT Gap	The average amount of time (hh:mm:ss) the tasks in the Batch Window spent in Initiation
Total Idle Gaps all Windows	The total number of times <i>NO</i> jobs included in this Batch Window were executing while the Batch Window was open

Average Window Idle Gap Time	The time (hh:mm:ss) during which no jobs were executing
Cumulative Idle Gap Time	For all Batch Windows that occurred during the measured period, the total time (hh:mm:ss) during which NO jobs were executing
Average Window CPU Time	The total amount of CPU time (hh:mm:ss) for all tasks in an average Batch Window during the period
Average Window DASD EXCPs	The total number of DASD I/O operations for all tasks included in the Batch Window during an average Batch Window for the period
Maximum Concurrent Tasks	For all Batch Windows in the period, the highest number of concurrently executing tasks included in the Batch Window
Average Overall M/P Factor	The average number of tasks included in the Batch Window that were executing concurrently
Maximum Concurrent Tape Units	For all Batch Windows in the period, the highest number of concurrently allocated Tape Drives for all tasks included in the Batch Window
Average Tape Data set Accesses	The average number of Tape data sets that were opened by Tasks included in the Batch Window during the entire measured period
Tape Drive Early Idle Time	The time (hh:mm:ss) during which a Tape Drive was allocated, but had not yet been used to read or write to a data set
Tape Drive Late Idle Time	The time (hh:mm:ss) after the last close was issued for a data set on a Tape Drive until the task terminated
Tape Drive Window Active Time	The cumulative elapsed time (hh:mm:ss) during which a Tape Drive was allocated to a task included in the Batch Window and a data set on the Tape was open but had not yet been closed for the last time. Each Tape Drive counts for its full elapsed time.
Tape Drive Window Allocated Time	The cumulative elapsed time (hh:mm:ss) during which a Tape Drive was allocated to a task included in the Batch Window. Each Tape Drive counts for its full elapsed time.

Batch Window Event Chronology

This is a listing of the events that took place during the Batch Window. Each event is associated with one of the JOBS named in the &SID.BWnn member of CIMS Capacity Planner PARMLIB. The list can be filtered through the use of the JOBNAME parameter in the SYSIN stream. If a JOBNAME is specified, only events associated with jobs whose names match the JOBNAME parameter are shown. Wildcards (*) are allowed.

The following list identifies the events that are tracked during the period a Batch Window is open:

Batch Window Open	The time (hh:mm:ss) when the Opening JOB as identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB started.
Batch Window Close	The time (hh:mm:ss) when the Closing JOB as identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB terminated
JOB Step Start	The time (hh:mm:ss) when the a task associated with a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB went into initiation
Program Start	The time (hh:mm:ss) when a program associated with a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB started.
Tape Drive Allocate	The time (hh:mm:ss) when a tape drive was allocated by a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB
Open Tape Data Set	The time (hh:mm:ss) when a data set on a tape associated with a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB was opened
Close Tape Data Set	The time (hh:mm:ss) when a data set on a tape associated with a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB was closed
Unallocate Tape Drive	The time (hh:mm:ss) when a Tape Drive owned by a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB was released
JOB Step End	The time (hh:mm:ss) when a task associated with a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB terminated.
JOB Step ABEND	The time (hh:mm:ss) when a task associated with a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB terminated abnormally

Batch Window Tape Drive Chronology

This is a listing of events associated with Tape Drives that took place during the Batch Window. A separate listing is provided for each Tape Drive that was allocated at any time to one of the JOBS named in the &SID.BWnn member of the CIMS Capacity Planner PARMLIB.

The list can be filtered through the use of the JOBNAME parameter in the SYSIN stream. If a JOBNAME is specified, only Tape Drive events associated with jobs whose names match the JOBNAME parameter are shown. Wildcards (*) are allowed.

The following list identifies the events that are tracked during the period a Batch Window is open:

Batch Window Open	The time (hh:mm:ss) when the Opening JOB as identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB started.
Batch Window Close	The time (hh:mm:ss) when the Closing JOB as identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB terminated
Tape Drive Allocate	The time (hh:mm:ss) when a tape drive was allocated by a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB
Open Tape Data Set	The time (hh:mm:ss) when a data set on a tape associated with a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB was opened
Close Tape Data Set	The time (hh:mm:ss) when a data set on a tape associated with a JOB identified by the &SID.BWnn member of CIMS Capacity Planner PARMLIB was closed
Unallocate Tape Drive	The time (hh:mm:ss) when a Tape Drive owned by a JOB identified by the &sid.BWnn member of CIMS Capacity Planner PARMLIB was released

Key Capacity Factors Analysis

In terms of operations analysis, capacity planning involves analyzing the trends of several key capacity factors. While an investigation into the anticipated business needs of an organization might be based on several different business factors, so the future capacity requirements of the organization can be anticipated through the analysis of a number of specific data processing metrics. At a very high level, these key factors provide a valuable picture of the direction in which the utilization of the resources in the data center is headed.

Which Factors are Analyzed

CIMS Capacity Planner maintains a daily table (i.e., one table for each day) in the Performance Database that contains the following specific key capacity factors:

- The hourly percentage busy for the CPU.
- The total number of DASD I/O Operations per hour
- The maximum number of Batch Jobs executing concurrently during each hour
- The average number of Batch Jobs executing concurrently during each hour
- The maximum number of TSO users logged on concurrently during each hour
- The average number of TSO users logged on concurrently during each hour
- The maximum number of Started Tasks active concurrently during each hour
- The average number of Started Tasks active concurrently during each hour
- The number of Swaps that took place during each hour
- The number of Demand Paging operations that occurred during each hour
- The number of Service Units that were accumulated by all active Address Spaces during each hour
- The number of ended transactions during each hour for all Address Spaces
- The number of Tape I/O operations per hour
- The Serial number of the CPU the workload used for that hour
- The model and engine count for the CPU
- The CPU Adjustment factor (SRM) of the CPU

These factors are used in the Key Capacity Factors reports. They are also available to other programs in the CIMS Capacity Planner family of subsystems.

Data Collection

Several RMF records contain data that is essential to the analysis of the Key Capacity Factors. Type 70 records contain information relating to the percent busy of the processor, as well as the CPU Serial number, model and configuration information. Type 71 records contain information about the overall use of memory resources inside the CPU, including Demand Paging and Swap Sequences. Type 72 records contain information regarding the Service Units consumed by different workloads. Type 74 records contain DASD and Tape device statistics. The table used during Key Capacity Factor reporting is created by SSA1WKLD when the SMF/RMF data is reduced into the Workload database.

Reports

Two reports are available once the table has been created. These reports are produced by the SSA1OPA1 and SSA1OPA2 programs and are described below.

Key Capacity Factors Analysis (SSA1OPA1)

This report shows, for a specific Time period as indicated in the //SYSIN stream, the following values:

CPU Percent Busy	the average Prime Shift percentage of utilization of the CPU.
DASD Rate	the average Prime Shift value of the total number of DASD I/O Operations per second
JOB MAX	the maximum number of Batch Jobs executing concurrently during Prime Shift
JOB AVG	the average number of Batch Jobs executing concurrently during Prime Shift
TSO MAX	the maximum number of TSO users logged on concurrently during Prime Shift
TSO AVG	the average number of TSO users logged on concurrently during Prime Shift
STC MAX	the maximum number of Started Tasks active concurrently during Prime Shift
STC AVG	the average number of Started Tasks active concurrently during Prime Shift
Swap Rate	the average number of Swaps per second that took place during Prime Shift
Demand Paging	the average number of Demand Paging operations per second that occurred during Prime Shift
Service Rate	the average number of Service Units per second that were accumulated by all active Address Spaces (in millions)
Transaction Rate	the average number of ended transactions per second during Prime Shift for all Address Spaces

The Prime Shift average value for each capacity factor is shown for each day of the week included in the report. The values are presented in columns, beginning with Monday's data.

Selected Key Capacity Factors Analysis (SSA1OPA2)

Any two of the thirteen Key Capacity Factors can be selected for this report. The report shows detailed hourly values for the two factors for each day included in the date range. Selection is done by means of the FILTER keyword in the SYSIN for the report. The values that can be selected are:

FILTER=nn,nn

where

01=CPU % Busy

02=DASD I/O Operations per second

03=Maximum Jobs Executing Concurrently

04=Average Jobs Executing Concurrently

05=Maximum TSO Users Active Concurrently

06=Average TSO Users Active Concurrently

07=Maximum Started Tasks Executing Concurrently

08=Average Started Tasks Executing Concurrently

09=Swaps per Second

10=Demand Paging Rate per Second

11=Service Units (Million) per Second

12=Ended Transactions per Second

13=Tape I/O Operations per Second

The default is

FILTER=01,10 (CPU/Demand Paging)

Application Profile Analysis

CIMS Capacity Planner provides an Operations Analysis facility that allows you to track the behavior of a set of specific Applications during a particular date range. An application group consists of a specified set of Jobs, based on Jobname. For shops that enforce Job naming conventions, this approach to analyzing workloads by application is useful. Applications can be defined at report time, as all information in the table is kept by Jobname.

Which Values are Maintained

CIMS Capacity Planner maintains a daily table (i.e., one table for each day) in the Performance Database that contains the following application profile values for each unique Jobname:

- The daily CPU time in .xx seconds
- The daily elapsed time in .xx seconds
- The number of times that particular Job was executed
- The number of times that particular Job ABENDED
- The total DASD I/O operations for that Job
- The total Tape I/O operations for that Job
- The total Non-specific Tape Mounts for that Job
- The total Specific Tape Mounts for that Job
- The SYSOUT class and Page count for the 1st encountered print class
- The SYSOUT class and Page count for the 2nd encountered print class
- The SYSOUT class and Page count for the 3rd encountered print class
- The SYSOUT class and Page count for the 4th encountered print class
- The SYSOUT class and Page count for the 5th encountered print class
- The SYSOUT class and Page count for the 6th encountered print class
- The SYSOUT class and Page count for the 7th encountered print class
- The SYSOUT class and Page count for the 8th encountered print class
- The SYSOUT class and Page count for the 9th encountered print class

These values are used in the Application Profile report. They are also available to other programs in the CIMS Capacity Planner family of subsystems.

Data Collection

Two SMF records contain data that is essential to the construction of the Application Profile. Type 30 records contain information relating to the amount of CPU and Elapsed time by Job, as well as DASD and Tape I/O Operations. The SMF Type 30 record also contains ABEND information along with Tape Mounts by type. The SMF Type 6 record contains Job-related print information by SYSOUT class. The table used during Application Profile reporting is created by SSA1WKLD when the SMF/RMF data is reduced into the Workload database.

Application Profile Report

The Application Profile report is driven by two external tables along with the parameters in the SYSIN stream. The SYSIN stream specifies the date range for the report and the name of the system (SELECTED SYSTEM). INCLUDE/EXCLUDE filtering can optionally be used to limit the names of the Jobs included in the report

Identifying the Application

The 1st table, pointed to by the //ACCTABLE DD statement, contains the association between a Jobname and a specific Application. The table consists of two columns. The 1st column is 8 bytes wide, starting in position 1, and contains the Jobname or fragment of a Jobname. Wild cards are allowed.

The 2nd column begins in position 23 and runs for 32 positions. It contains the name of the Application itself. This name is displayed on the Application Profile Report. It can contain upper and lower case characters.

Grouping SYSOUT Classes into Categories

The 2nd table is a CPPR PARMLIB member named &sid.SOUT and contains an association between a category of Print work and one or more SYSOUT classes, separated by commas. Only the 1st four categories in the PARMLIB member are used in the report.

The name of the SYSOUT category is printed on the report. It can be up to 7 characters long and can contain upper and lower case characters. A sample member is distributed in the Distribution PARMLIB named CPPRSOUT.

Application Profile Report (SSA1ACSR)

This report shows, for each listed Application in the table pointed to by the //ACCTABLE DD statement, for a specific Time period as indicated in the //SYSIN stream, the following values:

JOBNAME	The name of the Job
Number of Runs	The number of times during the period the Job was executed
Number of ABENDs	The number of times during the period the Job terminated abnormally

Average Start Time	The average start time (Hour : Minute based on a 24 hour clock) the Job started
Average CPU Time	The average amount of CPU time used by the Job (Minutes : Seconds)
Average Elapsed Time	The average amount of elapsed time for the Job (Hours : Minutes)
Total CPU Time	The total CPU time used by all executions of the Job during the period (Minutes : Seconds)
Total Elapsed Time	The total elapsed time for all executions of the Job during the period (Hours : Minutes)
Non-specific Tape Mounts	The total number of non-specific tape mounts associated with this Job
Specific Tape Mounts	The total number of specific tape mounts associated with this Job
DASD I/O	The total number of DASD I/O Operations associated with this Job
Tape I/O	The total number of Tape I/O Operations associated with this Job
Pages Printed Category 1	The total number of pages printed in the SYSOUT classes associated with category 1
Pages Printed Category 2	The total number of pages printed in the SYSOUT classes associated with category 2
Pages Printed Category 3	The total number of pages printed in the SYSOUT classes associated with category 3
Pages Printed Category 4	The total number of pages printed in the SYSOUT classes associated with category 4

Model 204 Subsystem

Subsystem Overview	10-3
Model 204 Workload Analysis and Performance Reports	10-3
Model 204 Summary Analysis Report	10-3
Model 204 Response Performance Report	10-4
Model 204 Transaction Statistics Report	10-4
Model 204 Terminal Statistics Report	10-7
Model 204 Exception Analysis	10-7
Model 204 Transaction Activity List	10-8
Model 204 Ad Hoc Report	10-8
Model 204 Workload and Performance Graphs	10-9
Model 204 Transaction Response Graph	10-9
Model 204 Terminal Response Graph	10-9
Model 204 Transaction Activity Graph	10-10
Model 204 Terminal Activity Graph	10-10
Model 204 CPU Activity Graph	10-10
Model 204 I/O Activity Graph	10-10
Model 204 User ID Profile	10-11
Model 204 Trends Analysis Report	10-11
Operating JCL and Parameters	10-12
Model 204 Subsystem Installation JCL	10-12
Data Reduction JCL	10-12
Additional Report Options	10-32
ISPF/PDF Model 204 Subsystem Interface	10-34
On-line Graphs	10-35
On-line Reports	10-35
Model 204 Summary Analysis Report	10-36
Model 204 Trends Analysis Report	10-37
Model 204 Exception Analysis Report	10-37

Model 204 Ad Hoc Report	10-37
Batch Reports	10-38
Batch GDDM Graphs	10-38
Presentation Graphics Interface	10-43
Graph Production Procedures	10-44
Naming Conventions	10-44
Presentation Graphics Distribution Disk	10-46
Graph Periods	10-48
Graph Descriptions	10-48
Graphic Report Tables	10-48

Subsystem Overview

The CIMS Capacity Planner Model 204 Subsystem provides a variety of reports and graphs related to the Model 204 workload and the performance of your Model 204 system and applications. The Model 204 reports provide both detail and summary level information designed to meet the needs of both the technical staff and technical management. A variety of graphs are provided through both GDDM and the PC based Harvard Graphics system. GDDM provides the capability to instantly display a number of Model 204 graphs on the Host through the use of the CIMS Capacity Planner ISPF interface. The Harvard Graphics system is used primarily to produce professional quality colored graphs that are used for presentations to Users and management.

Model 204 Workload Analysis and Performance Reports

The following discussion enumerates the various CIMS Capacity Planner Model 204 reports and describes the contents of each report.

Model 204 Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the Model 204 workload and how well the Model 204 system is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time

The Summary Report can be produced for a single day or portion of a day or it can encompass any number of days, thereby allowing an evaluation of the overall long term and short term performance of the Model 204 system.

- The peak number of concurrent active terminals (high watermark)

A terminal is considered active during each 15-minute period throughout the day only if at least one transaction is submitted. If a terminal is logged on throughout three shifts, but transactions are only submitted during two 15 minute periods, the terminal is considered to be active for only 30 minutes.

- The average number of Model 204 transactions processed per minute during the prime shift
- The average number of Model 204-related EXCPs per second during the prime shift
- The average transaction elapsed time (in seconds) during the prime shift
- The average terminal response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift

- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift
- The average percentage of the CPU time consumed by Model 204 and the Model 204 applications during each shift within the measured period
- The average number of transactions processed by shift for the measured period

The 25 most frequently executed transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions are reported by shift.

Model 204 Response Performance Report

This report shows, by shift, for a single day or a range of days, how the Model 204 transactions flowed through the system for a specified Model 204 region. The User IDs for the 25 users who most frequently executed Model 204 transactions are listed individually, by shift, with an indication of what percentage of the response times fell within certain pre-defined time-frames. Users that are not among the 25 most active are grouped together and reported under "OTHER". A Totals line is accumulated and written to indicate how all users performed as a group.

This is a key report in measuring how well your performance goals are being met within the CPU by the Model 204 system and the application programs.

Model 204 Transaction Statistics Report

This report shows, for a specified Model 204 region, which Users were among the top 40 in the following categories:

- Most transactions executed

The User IDs and the transaction counts of the 40 most active Users are listed along with their respective percentages of the total Model 204 transaction workload. The Users are listed in descending order by transaction volume.

- Most cumulative elapsed time

The User IDs and the total amount of elapsed (residence) time for each of 40 Users experiencing the highest residency time are listed along with their respective percentages of the total Model 204 transaction residence time. The Users are listed in descending residency time order.

- Most cumulative CPU time

The User IDs and the total amount of CPU time consumed for each of the 40 Users that consumed the most CPU time are listed along with their respective percentages of the total Model 204 CPU time. The Users are listed in descending CPU time order.

■ Most records added

The User IDs and the number of records added are listed for each of the 40 Users issuing the most record addition calls along with their respective percentages of all Model 204 record addition activity. The Users are listed in descending order by the number of records added.

■ Most records deleted

The User IDs and the number of records deleted are listed for each of the 40 Users issuing the most record deletion calls along with their respective percentages of all Model 204 record deletion activity. The Users are listed in descending order by the number of records deleted.

■ Most records sorted

The User IDs and the number of records sorted are listed for each of the 40 Users causing the most records to be sorted along with their respective percentages of all Model 204 record sorting activity. The Users are listed in descending order by the number of records sorted.

■ Server size

The User IDs and the size of the server are listed for each of the 40 Users requiring the most storage. The transactions are listed in descending order by the amount of storage required.

■ Server reads

The User IDs, the number of server reads, and their relative percentage of all server reads are listed for the 40 Users requiring the highest number of server reads. The Users are listed in descending order by the number of server reads.

■ Server writes

The User IDs, the number of server writes, and their relative percentage of all server writes are listed for the 40 Users requiring the highest number of server writes. The Users are listed in descending order by the number of server writes required.

■ Model 204 EXCPs

The User IDs and the number of EXCPs are listed for the 40 Users that require the highest number of EXCPs along with their respective percentages of all EXCPs related to Model 204 processing. The listing is produced in descending order by the number of EXCPs.

■ Directed lines

The User IDs and the number of directed lines are listed for the 40 users requiring the highest number of directed lines along with their respective percentages of all directed lines. The listing is produced in descending order by the number of directed lines.

■ Terminal lines

The User IDs and the number of terminal lines are listed for the 40 users requiring the highest number of terminal lines along with their respective percentages of all terminal lines. The listing is produced in descending order by the number of terminal lines.

■ Find requests

The User IDs and the number of find requests are listed for the 40 users issuing the highest number of find requests along with their respective percentages of all find requests. The listing is produced in descending order by the number of find requests.

■ Direct searches

The User IDs and the number of direct searches are listed for the 40 users requiring the highest number of direct searches along with their respective percentages of all direct searches. The listing is produced in descending order by the number of direct searches.

■ Record reads

The User IDs and the number of record reads are listed for the 40 users requiring the highest number of record reads along with their respective percentages of all direct searches. The listing is produced in descending order by the number of record reads.

USERIDs are excluded from the Transaction Statistics Report through the use of the EXCLUDE facility (see JCL). Similarly, users that would not normally be included are included through the use of the INCLUDE facility (JCL).

Whenever any of the reported statistics for a transaction exceed six digits (1 million or more), the numbers are expressed in thousands and suffixed with the character "K".

Model 204 Terminal Statistics Report

This report shows which Model 204 terminals were among the top 40 in the following categories:

- Most Active

The User ID and the amount of time active (hours:minutes) are reported along with their respective percentages of the combined active time for all Model 204 terminals for the 40 most active terminals. A terminal is considered active within any given 15 minute period only if transactions are received from the terminal. A terminal that is signed-on, but does not submit any transactions is not considered to be active. The list is produced in descending order by the amount of time active.

- Most Transactions Executed

The User ID and the number of transactions submitted are reported for the 40 terminals submitting the most transactions along with their respective percentages of the total number of transactions submitted during the measured period. The listing is produced in descending order by the number of transactions submitted.

Any counts that exceed 6 digits are expressed in thousands by suffixing the number with the character "K".

Data pertaining to selected Users are excluded from the report through the use of the EXCLUDE facility (see JCL). Similarly, terminals that do not qualify as among the top 40 can be included in the report through the use of the INCLUDE facility (see JCL).

Model 204 Exception Analysis

The Model 204 Exception Analysis Report shows, for each 15 minute period, for a specified Model 204 region, all User IDs whose response times exceeded the pre-defined thresholds specified in the CIMS Capacity Planner Parmlib member &sidXCPT. The Parmlib member &sidXCPT contains a number of threshold values for the installation. This member is used to identify the performance thresholds for the various components of the overall environment.

Model 204 Transaction Activity List

The Model 204 Transaction Activity List shows daily average values for each individual M204 transaction in terms of the following:

- Transaction Count (average number of transactions per day)
- Average Response per transaction
- Total daily CPU time per transaction per day
- Total M204 I/O calls per transaction per day
- Total records added per transaction per day
- Total records sorted per transaction per day
- Total records deleted per transaction per day
- Total daily average I/O to a terminal per transaction

Model 204 Ad Hoc Report

The CIMS Capacity Planner Model 204 Subsystem supports the production of an Ad Hoc report from the original data before portions of the detail data are summarized or discarded. The Ad Hoc report is produced by the Data Reduction program - SSA1M20W.

To specify that the report is to be produced, you merely specify the report parameters in the Data Reduction Job stream as part of the SYSIN data set. The parameters that apply to the Ad Hoc report are:

M204 ACCOUNT NAME=

M204 USERID=

The Model 204 Ad Hoc report contains the following information for each record selected for inclusion in the report.

- The Model 204 Last ID(Transaction Identifier)
- The time that the transaction completed execution
- The User ID
- The Account Name
- The Model 204 Procedure Name
- The elapsed time for the transaction
- The CPU time used
- The Model 204 I/O count
- The counts of records added and deleted

- The server size
- The number server reads and writes
- The number terminal lines read and written

In addition to selecting records for inclusion into the Ad Hoc report by specific Account and User IDs, groups of both accounts and users are selected through the use of the wild card character "*".

For example, all accounts with the characters "TR" are selected by specifying MODEL 204 ACCOUNT NAME=TR*. Similarly, all User IDs beginning with "L422" are selected by specifying MODEL 204 USERID=L422*. This convention is extended to include all records by specifying MODEL 204 ACCOUNT NAME=*. M204 USERID=* would result in the selection of all User IDs. If both the account and User names of "*" were specified, all records would be selected.

The selection of records for inclusion into the Ad Hoc report can be further restricted by specifying date and/or time ranges by using the BEGIN DATE=, END DATE=, BEGIN TIME=, and END TIME= parameters.

Model 204 Workload and Performance Graphs

The Model 204 Subsystem produces a variety of bar graphs that depict the workload and the performance of the Model 204 system.

Model 204 Transaction Response Graph

The Transaction Response Graph shows, by day of the week, the average and peak average Model 204 transaction response times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the horizontal axis represents the response time (both average and peak average). In the middle of each line, both the average and peak values are printed. The lines, in bar graph format, form the graphic representation of the values. The peak values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak values and the average values are either equal or very nearly so.

The Transaction Response Graph is produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

Model 204 Terminal Response Graph

The Terminal Response Graph shows, by day of the week, the average and peak average Model 204 terminal response times for each 15-minute period throughout the day. The report page is broken into four sections—each covering a period of 6 hours. Within each section of the page, the time-of-day is represented by the vertical axis while the

horizontal axis represents the response time (both peak and average). In the middle of each line, both the average and peak values are printed. The lines, in bar graph format, form the graphic representation of the values. The peak values are represented by dashes while the average values are represented by asterisks. In those cases where no dashes are present, they have been overlaid by asterisks because the peak and average values are either equal or very nearly so.

The Terminal Response Graph is produced for a single day's activity or several days. When more than one day's activity is graphed, separate graphs are produced - one for each day of the week (e.g. Monday, Tuesday, etc.). Multiple days' activity is included on a single graph only when multiple occurrences of a given weekday fall within the measured period defined by the BEGIN DATE and the END DATE (see parameters).

The Terminal Response Graph differs from the Transaction Response Graph in two ways:

- Not all transaction responses are directed at terminals
- Through the judicious use of the EXCLUDE facility, it is possible to exclude all printers, monitors, and responses to long running transactions from the graph.

Model 204 Transaction Activity Graph

The Transaction Activity Graph is organized in the same format as the Transaction and Terminal Response Graphs described above. It shows, by day of the week, the average and peak average number of transactions submitted from all the Model 204 terminals during each 15-minute period throughout the day. The horizontal value represents the number of transactions submitted per minute.

Model 204 Terminal Activity Graph

The Model 204 Terminal Activity Graph shows, by day of the week, the number of terminals during each 15-minute period throughout the day. Both the peak and the average values are reported. The values along the horizontal axis represent the number of active terminals.

Model 204 CPU Activity Graph

The Model 204 CPU Activity Graph shows, by day of the week, the percentage of the CPU processing capacity that was expended in processing Model 204 transactions for a specified Model 204 region. The average and peak average percentage values are reported in 15 minute intervals throughout the day. As is the case with the other graphs, the measured period is specified to include anywhere from a day to several months.

Model 204 I/O Activity Graph

This graph shows, for a specified Model 204 Region, the average and peak average I/O rates for each 15 minute period throughout the day. The format of the I/O Activity Graph is identical to the graphs described above, except that the values along the horizontal axis represent the number of I/Os per second related to processing the Model 204 workload.

Model 204 User ID Profile

The User ID Profile report shows how a specified Model 204 User ID behaved for the following categories:

- How many times per day, on average, it was invoked
- The average response time per transaction
- The average CPU time per transaction
- The average number of adds, deletes, and sorts per transaction
- The average number of terminal I/Os per transaction
- The average number of Model 204 I/Os per transaction

Model 204 Trends Analysis Report

The CIMS Capacity Planner approach to Trends Analysis centers on isolating a number of capacity and performance related elements and providing either a summary or graphic comparison of the values of those elements over time. The Model 204 elements that are selected for comparison are:

- CPU Utilization Statistics, including a breakdown in terms of Model 204 overhead and Task related CPU usage
- The number of transactions executed per minute computed as an average during Prime shift
- The average number of EXCPs per second during Prime shift
- The average response time during Prime shift
- The ratio of Early shift to Prime shift transactions
- The ratio of Late shift to Prime shift transactions
- The ratio of Early shift to Prime shift terminals
- The ratio of Late shift to Prime shift terminals

This information is gathered and presented to you in one of two ways:

- In summary report format, showing the values of each of the elements listed above during a baseline period and comparing them to a secondary period, with the slope of each comparison indicated at the right hand side of the report
- In data suitable for graphing with the Harvard Graphics Program once it has been down loaded to a Personal Computer. The manner in which the data is down loaded is left to you. It is stored at the mainframe, however, in a PDS with the DDNAME of HGDLIB.

Operating JCL and Parameters

The Model 204 Subsystem functions in much the same manner as the Workload Analysis Subsystem, but there are a few differences that stem from the nature of Model 204 systems:

- You might want to allocate a separate ONLINE Performance Data Base file for the Model 204 tables
- Model 204 input data is normally in the form of a Model 204 journal record. The input can, however, come from several sources:
 - The SMF MANx Clusters, included with the other SMF records, if the Model 204 SMF option is selected
 - An SMF history file (QSAM)
 - A specific Model 204 journal file (QSAM)

The manner in which the input is processed varies, depending upon the source of the input. The Model 204 input is processed by a single Model 204 data reduction program, SSA1M20W, regardless of the source of the input. The source of the input is identified to the data reduction program through the use of the SMFILE parameter as part of the Job stream.

Model 204 Subsystem Installation JCL

The JCL required to install the CIMS Capacity Planner Model 204 Subsystem is described in the *CIMS Capacity Planner Installation Guide*

Data Reduction JCL

Model 204 Input data is normally in the form of an SMF record, however it can come from a variety of sources:

- The SMF MANx clusters, included with other types of SMF records
- A collection of off-loaded SMF data in the form of a SMF historical data set (QSAM)
- A specific Model 204 Journal file produced by the Model 204 Monitoring Facility

Regardless of the format of the input, all data must be presented to the Model 204 data reduction program in chronological order. CIMS Capacity Planner maintains a table showing the date and time of the last transaction processed. Should a transaction be read into the data reduction program that is time stamped earlier than the latest transaction in the file, the transaction file is considered a duplicate of a previously processed file and discarded unless the checking feature is overridden by the FORCE MODEL 204 INPUT=YES parameter.

Processing Model 204 SMF Records

Following is a sample set of the JCL required to run the Model 204 Data Reduction program using Model 204 Journal QSAM history records as input.


```
//SSADM20 JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//*
//*      RUN CIMS Capacity Planner DATA REDUCTION STEP
//*
//STA      EXEC PGM=SSA1M20W,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.INDEX.M204
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.ONLINE.M204
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.CPPRERT
//SYSUT1   DD DISP=(OLD,PASS),DSN=M204.JOURNAL
//SYSUT3   DD DISP=(,PASS),SPACE=(CYL,(10,8)),
//
//          UNIT=SYSDA
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSG   DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=&sid
//*      ONE OF THE FOLLOWING STATEMENTS IS USED TO DESCRIBE THE FORMAT
//*      OF THE INPUT.  UNCOMMENT THE SELECTED STATEMENT.
SMFILE=JRNL /* CCA JOURNAL INPUT
*SMFILE=SMFA /* SMF RECORDS FROM CURRENTLY ACTIVE CLUSTER
*SMFILE=SMFL /* SMF RECORDS FROM A SMF CLUSTER
*SMFILE=SMFH /* SMF RECORDS FROM QSAM HISTORY FILE
DUMP SMF STATISTICS=YES
```

Other forms of the input are readily accommodated by respecifying the SMFILE parameter.

Member DM20PROD in the &PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL which you might want to customize for your installation.

CIMS Capacity Planner Model 204 Ad Hoc Report

CIMS Capacity Planner provides you the option of producing one or more reports from the original input data. These reports can contain much of the original detail information that might be discarded or lost during summarization during the data reduction process.

The Model 204 Ad Hoc report is generated by the Model 204 data reduction program (SSA1M20W), when all the detail data is available. The Ad Hoc reports are generated on an as-needed basis and it is likely, therefore, that the need for the Ad Hoc reports is determined after the original data are already processed and stored into the Performance Data Base. In such a case, it is important that the tables created by the data reduction program not be stored into the Performance Data Base twice. This is accomplished by specifying DD DUMMY for both the ONLINE and the INDEX data sets.

In general, the records selected for inclusion in the Ad Hoc reports are limited by including the following parameters in the SYSIN data set in the Job stream:

- A specific Account Name
MODEL 204 ACCOUNT NAME=LO*
- A specific USERID
MODEL 204 USERID=CS*
- A date range
 - BEGIN DATE=02/04/2005
 - END DATE=02/07/2005
- A time period
 - BEGIN TIME=08.00.00
 - END TIME=10.00.00

Using the data provided above, the Ad Hoc report produced would contain only data generated between 08:00 A.M. on February 4, 2005 and 10:00 A.M. on February 7, 2005. The inclusion of the data would be further limited to all accounts beginning with the characters "LO" that executed transactions for Users whose IDs begin with the characters "CS".

To request a Model 204 Ad Hoc report, merely specify the accounts and User IDs that you want to include in the report. To request that an Ad Hoc report not be produced, merely remove any selection criteria from the SSA1M20W Job Step.

Data Reduction DD Statements

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all M204 Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the INDEX to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ SYSUT1

The SYSUT1 DD statement specifies the input to the CIMS Capacity Planner Data Reduction Program—SSA1CICW.

■ SYSUT3

The SYSUT3 DD statement is used to specify a temporary data set used to hold a working copy of the Performance Data Base during the data reduction process. It is no longer needed after the performance data is copied back to the ONLINE file at the conclusion of the data reduction phase. The parameters that must be specified are: Disposition, Unit, and Space.

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the M204 Ad Hoc report. Unless otherwise specified, the DCB characteristics are:

RECFM=FBA

BLKSIZE=133

LRECL=133

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*

■ SYSIN

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the Model 204 data reduction phase are described in the next paragraph.

CIMS Capacity Planner Model 204 Data Reduction Parameters

The CIMS Capacity Planner Parameters permits you to specify certain run-time options and to provide selected data required by the data reduction program in processing the M204 history files and producing the Ad Hoc report.

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMF SID of the MVS system for which data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID can be specified by this parameter in running the Model 204 data reduction program.

Parameters:

The SELECTED SYSTEM= parameter is specified in several ways:

- The SID itself (e.g. SELECTED SYSTEM=SYS1)
- An "*", indicating the SMF SID of the system on which the program is being executed (e.g. selected system=*)

Default Parameter:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the earliest date (inclusive) for which data is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which a Model 204 record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in *BEGIN DATE on page 10-17*

Default date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Parameters:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Parameter:

REPORT LANGUAGE=ENGLISH

SMFILE

The SMFILE parameter is used to specify the type of input that is to be processed by the M204 data reduction program.

Specify Input Source:

SMFILE=JRNL (CCA Journal)

SMFILE=SMFH (used for SMF history records from a QSAM file)

SMFILE=SMFL (used for records from a VSAM SMF cluster)

SMFILE=SMFA (used for the currently active SMF cluster)

Default Input Source:

SMFILE= JRNL (CCA Journal)

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should only be used when working with the CIMS Lab to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

MODEL 204 ACCOUNT NAME

This parameter is used to specify a particular Model 204 account or group of accounts to be selected for inclusion into the Model 204 Ad Hoc Report. It is used in conjunction with the Model 204 USERID parameter in selecting a single User account or group of Users accounts. The wild card character "*" is supported so that a group of accounts, all beginning with a common character string, can be conveniently and easily specified. Please refer to the discussion of the Ad Hoc report specified above in *CIMS Capacity Planner Model 204 Ad Hoc Report on page 10-14*.

Specify a Model 204 Account Name:

MODEL 204 ACCOUNT NAME=acct*

Default Account Name:

None

MODEL 204 USERID

This parameter is used to specify a particular Model 204 User or group of Users to be selected for inclusion into the Model 204 Ad Hoc Report. It is used in conjunction with the Model 204 ACCOUNT NAME parameter in selecting a single User account or group of Users accounts. The wild card character "*" is supported so that a group of User IDs, all beginning with a common character string, can be conveniently and easily specified. Please refer to the discussion of the Ad Hoc report specified above in *CIMS Capacity Planner Model 204 Ad Hoc Report on page 10-14*.

Specify a Model 204 User ID:

```
MODEL 204 USERID=user*
```

Default Parameter:

None

Model 204 Report JCL and Parameters

A variety of reports and graphs are produced by the Model 204 Report Program SSA1M20R. All reports are selectable and all the report generation modules take their input from the ONLINE Performance Data Base (Consolidated Table File) that is maintained by the Model 204 Data Reduction Program.

An optional Element Exclusion/Inclusion file can be utilized to limit the User IDs that are analyzed in preparing the various reports and graphs. For example, all data related to users whose User IDs begin with the characters "CS" might be excluded from the reports by specifying an EXCLUDE DD statement and entering "CS*" as an excluded name. Please refer to the discussion related to the EXCLUDE entry below.

The JCL and parameters required to run the Model 204 reports are described below.

Model 204 Report Generation JCL

```
//D204REPT JOB (...),'SSA',CLASS=A,MSGCLASS=X
/*JOBPARM S=*
//STA EXEC PGM=SSA1M20R,REGION=5000K,TIME=60
//STEPLIB DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX DD DISP=SHR,DSN=&PREFIX.CPPR.INDEX.M204
//ONLINE DD DISP=SHR,DSN=&PREFIX.CPPR.ONLINE.M204
//CPPRERT DD DISP=SHR,DSN=&PREFIX.CPPR.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSNAP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN DD *
SELECTED SYSTEM=&sid
M204 EXCEPTION ANALYSIS=YES
```

```
M204 PERFORMANCE REPORT=YES
M204 TRANSACTION STATISTICS REPORT=YES
M204 TERMINAL STATISTICS REPORT=YES
M204 SUMMARY REPORT=YES
M204 TRANSACTION RESPONSE GRAPH=YES
M204 TERMINAL RESPONSE GRAPH=YES
M204 TRANSACTION ACTIVITY GRAPH=YES
M204 TERMINAL ACTIVITY GRAPH=YES
M204 CPU ACTIVITY GRAPH=YES
M204 I/O ACTIVITY GRAPH=YES
M204 LINEAR LIST=YES
M204 USERID PROFILE=YES
M204 USERID=XXXX
PRIME SHIFT FIRST HOUR=07
```

Member D204REPT in the &PREFIX.CPPR.Vnnn.CNTL library contains a model set of JCL that you might want to customize for your installation.

Report Production DD Statements

■ STEPLIB

The STEPLIB DD statement specifies the data set name and the disposition of the CIMS Capacity Planner load module library.

■ CPPRPARM

The CPPRPARM DD statement specifies the data set name and the disposition of the CIMS Capacity Planner parameter library.

■ CPPRERT

The CPPRERT DD statement specifies the data set name and the disposition of the CIMS Capacity Planner Element Registration Table data set. The CPPRERT data set must be specified in all M204 Job streams.

■ INDEX

The INDEX DD statement specifies the data set name and the disposition of the INDEX to the ONLINE (Performance Data Base) data set. The INDEX data set greatly improves the performance of CIMS Capacity Planner.

■ ONLINE

The ONLINE DD statement specifies the name and the disposition of the CIMS Capacity Planner Performance Data Base.

■ EXCLUDE/INCLUDE

The EXCLUDE or INCLUDE DD statements specify the User IDs that are to be used to limit specified elements of the Consolidated Table data from the reports (EXCLUDE) or to precisely define those elements of the data to be included in the reports. The format of the input statements is as follows:

- An element can be up to eight characters long
- Multiple elements can be specified in a single record
- Multiple elements specified in a single record must be separated by commas
- Element definitions must not go past column 72
- Elements can be terminated by a wild card character "*" to specify that all User IDs beginning with the character string preceding the "*" are to be EXCLUDED or INCLUDED.

For example, the following input statement in the Exclusion file:

```
TSOP,CS*
```

would indicate that all data pertaining to the User ID TSOP and any users whose IDs begin the string CS are to be excluded from the reports.

If these same input statements were to come through the INCLUDE file, then the report would contain data related to the User whose ID is TSOP and all users whose IDs begin with the characters "CS".

■ SYSPRINT

The SYSPRINT DD statement specifies the data set that will contain the Model 204 reports. Unless otherwise specified, the DCB characteristics are:

```
RECFM=FBA
```

```
BLKSIZE=133
```

```
LRECL=133
```

■ SYSNAP

The SYSNAP DD statement is used only in conjunction with the CIMS Capacity Planner Parameter DEBUGON. It specifies a SYSOUT data set (usually SYSOUT=*) used to print snap dumps.

■ SYSUDUMP

The SYSUDUMP DD statement is generally included in the sample JCL, but is not mandatory. It is usually specified as SYSOUT=*.

■ **SYSIN**

The SYSIN data set is normally specified as: SYSIN DD *, but can also reference a data set on disk. The SYSIN data set is used to input Parameter data to CIMS Capacity Planner. The parameters that relate to the M204 data reduction phase are described in the next paragraph.

CIMS Capacity Planner Model 204 Report Production Parameters

The CIMS Capacity Planner Parameters permits you to specify certain run-time options and to provide selected data required by the data reduction program in processing the M204 history files and producing the Ad Hoc report.

SELECTED SYSTEM

The SELECTED SYSTEM= parameter is used to specify the SMF SID of the MVS system for which data is to be analyzed. This is the parameter specified in the SMFPRMxx member of SYS1.PARMLIB. Only a single SMF SID can be specified by this parameter in running the Model 204 data reduction program.

Parameters:

The SELECTED SYSTEM= parameter can be specified in several ways:

- The SID itself (e.g. SELECTED SYSTEM=SYS1)
- An "*", indicating the SID if the system on which the program is being executed (e.g. selected system=*)

Default Parameter:

SELECTED SYSTEM=*

BEGIN DATE

The BEGIN DATE= parameter is used to specify the earliest date (inclusive) for which data is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

BEGIN DATE=00001

END DATE

The END DATE= parameter is used to specify the last date (inclusive) for which a Model 204 record is to be included for analysis.

Specify Date:

The format of the parameters for the END DATE are identical to the BEGIN DATE parameters specified above in **BEGIN DATE**.

Default Date:

END DATE=2099365

REPORT LANGUAGE

This parameter specifies the language or languages to be used in producing the narrative sections of the report.

Specify Language:

ENGLISH, DEUTSCH or any combination of languages separated by commas.

Default Language:

REPORT LANGUAGE=ENGLISH

LUNCH BREAK BEGIN HOUR

This parameter, in combination with the LUNCH BREAK END HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports. Use of this parameter does not cause any data to be removed from the Performance Data Base.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK BEGIN HOUR=11)

LUNCH BREAK END HOUR

This parameter, in combination with the LUNCH BREAK BEGIN HOUR, defines a period during the day for which the data resident in the Performance Data Base is to be excluded from the reports.

Specify Hour :

Specify hour in the 24 hour format (e.g, LUNCH BREAK END HOUR=13)

DEBUGON

This parameter is used in combination with the SYSNAP DD statement to request SNAP dumps be taken at various points in the data reduction program. It should only be used when working with the CIMS Lab to isolate problems.

Options:

None

Default Option:

None (equivalent to DEBUGOFF)

PRIME SHIFT FIRST HOUR

The PRIME SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Prime Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Prime Shift

PRIME SHIFT FIRST HOUR=hh

Default Hour:

PRIME SHIFT FIRST HOUR=08

LATE SHIFT FIRST HOUR

The LATE SHIFT FIRST HOUR= parameter specifies the hour of the day on which the Late Shift begins. It is expressed as an hour using the 24 hour convention.

Specify Hour:

Specify the beginning of the Late Shift

LATE SHIFT FIRST HOUR=hh (24 hour clock)

Default Hour:

LATE SHIFT FIRST HOUR=16

MODEL 204 SUMMARY REPORT

This parameter specifies whether or not the Model 204 Summary Analysis Report is to be produced

Options:

MODEL 204 SUMMARY REPORT=YES (produce the report)

MODEL 204 SUMMARY REPORT=NO (don't produce the report)

Default Option:

MODEL 204 SUMMARY REPORT=NO

MODEL 204 TRANSACTION STATISTICS REPORT

This parameter specifies whether or not the Model 204 Transaction Statistics Report is to be produced

Options:

MODEL 204 TRANSACTION STATISTICS=YES (produce the report)

MODEL 204 TRANSACTION STATISTICS=NO (don't produce the report)

Default Option:

MODEL 204 TRANSACTION STATISTICS=NO

MODEL 204 TERMINAL STATISTICS REPORT

This parameter specifies whether or not the Model 204 Terminal Statistics Report is to be produced

Options:

MODEL 204 TERMINAL STATISTICS REPORT=YES (produce the report)

MODEL 204 TERMINAL STATISTICS REPORT=NO (don't produce the report)

Default Option:

MODEL 204 TERMINAL STATISTICS REPORT=NO

MODEL 204 PERFORMANCE REPORT

This parameter specifies whether or not the Model 204 Performance Report is to be produced

Options:

MODEL 204 PERFORMANCE REPORT=YES (produce the report)

MODEL 204 PERFORMANCE REPORT=NO(don't produce the report)

Default Option:

MODEL 204 PERFORMANCE REPORT=NO

MODEL 204 CPU ACTIVITY GRAPH

This parameter specifies whether or not the Model 204 CPU Activity Graphic Report is to be produced

Options:

MODEL 204 CPU ACTIVITY GRAPH=YES (produce the report)

MODEL 204 CPU ACTIVITY GRAPH=NO (don't produce the report)

Default Option:

MODEL 204 CPU ACTIVITY GRAPH=NO

MODEL 204 I/O ACTIVITY GRAPH

This parameter specifies whether or not the Model 204 I/O Activity Graphic Report is to be produced

Options:

MODEL 204 I/O ACTIVITY GRAPH=YES (produce the report)

MODEL 204 I/O ACTIVITY GRAPH=NO (don't produce the report)

Default Option:

MODEL 204 I/O ACTIVITY GRAPH=NO

MODEL 204 TERMINAL ACTIVITY GRAPH

This parameter specifies whether or not the Model 204 Terminal Activity Graphic Report is to be produced

Options:

MODEL 204 TERMINAL ACTIVITY GRAPH=YES (produce the report)

MODEL 204 TERMINAL ACTIVITY GRAPH=NO (don't produce the report)

Default Option:

MODEL 204 TERMINAL ACTIVITY GRAPH=NO

MODEL 204 TRANSACTION RESPONSE GRAPH

This parameter specifies whether or not the Model 204 Transaction Response Graphic Report is to be produced

Options:

MODEL 204 TRANSACTION RESPONSE GRAPH=YES (produce the report)

MODEL 204 TRANSACTION RESPONSE GRAPH=NO (don't produce the report)

Default Option:

MODEL 204 TRANSACTION RESPONSE GRAPH=NO

MODEL 204 TERMINAL RESPONSE GRAPH

This parameter specifies whether or not the Model 204 Terminal Response Graphic Report is to be produced

Options:

MODEL 204 TERMINAL RESPONSE GRAPH=YES (produce the report)

MODEL 204 TERMINAL RESPONSE GRAPH=NO (don't produce the report)

Default Option:

MODEL 204 TERMINAL RESPONSE GRAPH=NO

MODEL 204 EXCEPTION ANALYSIS

This parameter specifies whether or not the Model 204 Exception Analysis Report is to be produced

Options:

MODEL 204 EXCEPTION ANALYSIS=YES (produce the report)

MODEL 204 EXCEPTION ANALYSIS=NO (don't produce the report)

Default Option:

MODEL 204 EXCEPTION ANALYSIS=NO

MODEL 204 USERID PROFILE

This parameter specifies whether or not the Model 204 User ID Profile Graphic Report is to be produced

Options:

MODEL 204 USERID PROFILE=YES (produce the report)

MODEL 204 USERID PROFILE=NO (don't produce the report)

Default Option:

MODEL 204 USERID PROFILE=NO

MODEL 204 USERID

This parameter is used to specify the User ID for which the User ID Profile is to be produced.

User ID Specification:

MODEL 204 USERID=userid (any legitimate M204 User ID)

Default User ID:

None

1ST PERIOD BEGIN DATE

This parameter is used to specify the beginning of the first of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1M20T) or the beginning date for the entire period to be used for developing the Trends Analysis Graph (SSA1M20G) for which consolidated table information is to be selected for analysis.

Specify Date:

Specify a beginning date in one of the following formats:

YYDDD (Standard Julian format)

MM/DD/YY (USA Standard Gregorian format)

DD.MM.YY (Standard European Gregorian format)

*, *-n (relative day format where *=today)

Default Date:

1ST PERIOD BEGIN DATE=00001

1ST PERIOD END DATE

This parameter is used to specify the ending date of the first period for the Trends Analysis Summary Program (SSA1M20T) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1M20G).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

DefaultDate:

1ST PERIOD END DATE=99365

2ND PERIOD BEGIN DATE

This parameter is used to specify the beginning of the second of two time periods that are to be compared in developing trends for the Trends Analysis Summary Report (SSA1M20T)

Specify Date:

Specify a beginning date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD BEGIN DATE=00001

2ND PERIOD END DATE

This parameter is used to specify the ending date of the second period for the Trends Analysis Summary Program (SSA1M20T) or the ending date for the entire period for the Trend Analysis Graph Program (SSA1M20G).

Specify Date:

Specify an ending date in the same format as the 1ST PERIOD BEGIN DATE.

Default Date:

2ND PERIOD END DATE=2099365

General Report Statements

Several general capabilities have been added to the reports produced by all the CIMS Capacity Planner Subsystems. A new statement

TOP LABEL=YES

causes the report to include a Top Line that contains the following specific information pertaining to the creation of the report:

- Date Range as specified
- Begin Time:End as specified

- An Hour Map that shows Early Shift, Prime Shift, Ignore Period (Lunch Break) and Late Shift as specified
- A Day map that shows Selected Days
- A running page number for the entire set of reports

FOOT1

Another statement

```
FOOT1=UP TO 64 CHAR FOOTNOTE
```

is used to specify an optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that is included on all reports where there is room.

FOOT2

A third statement

```
FOOT2=UP TO 64 CHAR 2ND FOOTNOTE
```

is used to specify another optional Foot Note Line (up to 64 characters with the end being denoted by a pair of blanks) that is included on all reports where there is room.

Trends Analysis JCL

The following is an example of a Jobstream that produces the Trends Analysis Report, comparing the period beginning on February 13, 2005 and ending on April 13, 2005 to the period beginning on April 14, 2005 and ending on June 14, 2005. The prime shift is defined as beginning at 07:00 and ending at 19:00. The SMF SID is IPO1.

```
//SSADM20 JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/*JOBPARM S=*
/*
/*      RUN CIMS Capacity Planner TRENDS ANALYSIS REPORT
/*
//STA      EXEC PGM=SSA1M20T, REGION=5000K, TIME=60
//STEPLIB DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.LOADLIB
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//INDEX    DD DISP=SHR, DSN=&PREFIX.CPPR.INDEX.M204
//ONLINE   DD DISP=SHR, DSN=&PREFIX.CPPR.ONLINE.M204
//CPPRERT  DD DISP=SHR, DSN=&PREFIX.CPPR.CPPRERT
//SYSPRINT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSMSGGS DD SYSOUT=*
//SYSIN    DD *
SELECTED SYSTEM=IPO1
```

```
1ST PERIOD BEGIN DATE=02/13/2005
1ST PERIOD END DATE=04/13/2005
2ND PERIOD BEGIN DATE=04/14/2005
2ND PERIOD END DATE=06/14/2005
PRIME SHIFT FIRST HOUR=07
LATE SHIFT FIRST HOUR=19
LUNCH BREAK BEGIN HOUR=12
LUNCH BREAK END HOUR=13
```

Member D204TRPT located in &PREFIX.CPPR.Vnnn.CNTL contains a set of model JCL nearly identical to the above Jobstream that can be customized for your installation.

Additional Report Options

In addition to the INCLUDE/EXCLUDE options described in *Report Production DD Statements on page 10-20*, it is possible to limit the various data elements that are contained in reports or graph data point members through the use of Element Masks. The element masks allow you to mask out selected portions of a key field such as the transaction name or a terminal name and screen and/or report on the basis of the masked name.

For example, a data center might have a naming convention for VTAM terminals in which the first and second characters designate the type of terminal, the third and fourth characters identify the department in which the terminals are located, the fifth and sixth characters designate the terminal locations, and the seventh character designates the functional characteristics of the terminal (display or printer).

In this case, it is useful to develop a report showing the response times for a given location, regardless of the departments involved, but excluding printers. This type of a report is prepared by masking out selected portions of the terminal name so that only the terminal location and the terminal type show through and then applying selection criteria to determine which locations and terminal types are to be included or excluded.

The masking and selection required are supported through the provision of three additional reporting parameters:

- GENERIC ELEMENT MASK=xxxxxxx
- INCLUSIVE ELEMENT MASK=xxxxxxx
- EXCLUSIVE ELEMENT MASK=xxxxxxx

The form of the keyword phrase is:

type ELEMENT MASK=xxxxxxx

where "type" can be GENERIC, INCLUSIVE, or EXCLUSIVE and xxxxxxxx can be any combination of "1"s and "0"s, where a "1" specifies that the corresponding character in the record key should be replaced with an "*" and a "0" specifies that the corresponding character in the record should be preserved. Thus the statement:

```
GENERIC ELEMENT MASK=00000000
```

specifies that the key is to remain untouched, whereas the statement:

```
GENERIC ELEMENT MASK=11111111
```

specifies that the entire record key should be masked out (set to "*****"). Finally, the statement :

```
GENERIC ELEMENT MASK=10101010
```

would specify that key positions 1,3,5 and 7 are to be set to "*" and that the remaining positions (2,4,6, and 8) are to remain unchanged.

GENERIC ELEMENT MASK

The GENERIC ELEMENT MASK parameter causes the data selection routines in CIMS Capacity Planner to mask out up to eight characters of the data element key by replacing it with an "*". The masking is specified by placing a "1" in each position to be replaced by an "*" and a "0" in each position that is to show through (not be masked). This has the effect of creating records (after masking) with duplicate keys where the keys were different prior to the masking process. The reporting program groups the records together prior to reporting so that the statistics or other data are summarized and reported under the modified keys.

The use of the GENERIC ELEMENT MASK parameter does not require the use of the INCLUDE or EXCLUDE facilities.

Using the example of the VTAM terminal names cited above, to report the average terminal response times for all terminals at the various locations, you merely specify a GENERIC ELEMENT MASK=11110011. All records in the terminal table with the same characters in positions 5 and 6 of the terminal name would be grouped together for reporting. There would be as many summarized entries as there are unique locations and the terminal names in the report would be in the form ****LL*. If GENERIC ELEMENT MASK=11111111 were to be specified, all records would be summarized into a single entry.

INCLUSIVE ELEMENT MASK

The INCLUSIVE ELEMENT MASK keyword phrase causes the M204 report processing programs to intercept the key of each element (in this case, the Model 204 USERID), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the INCLUDE selection module. In this case, the key is restored to its original value once the INCLUDE selection processing is completed.

The INCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the INCLUSIVE ELEMENT MASK is used only in selecting records for inclusion in the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The INCLUSIVE ELEMENT MASK is used with an INCLUDE data set specified in the JCL by the //INCLUDE DD statement. The format of the INCLUDE entries are:

```
//INCLUDE DD *  
  
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be selected. In the example of the VTAM terminal name masking cited above, to select all terminals in location 23, the selection mask specified in the INCLUDE data set would contain "???23??. The "?" characters in the INCLUDE selection statements correspond to the "*" positions in the modified keys.

EXCLUSIVE ELEMENT MASK

The EXCLUSIVE ELEMENT MASK keyword phrase causes the M204 report processing programs to intercept the key of each element (in this case, the Model 204 USERID), to change the key as specified by the mask, for comparison purposes only, and pass the changed key to the EXCLUDE selection module. In this case, the key is restored to its original value once the EXCLUDE selection processing is completed.

The EXCLUSIVE ELEMENT MASK parameter differs from the GENERIC ELEMENT MASK in that the EXCLUSIVE ELEMENT MASK is used only in selecting records for exclusion from the report. The keys that appear in the reports are not modified by the masking as they are when the GENERIC ELEMENT MASK is used.

The mask specification procedure is the same as for the GENERIC ELEMENT MASK described above.

The EXCLUSIVE ELEMENT MASK is used with an EXCLUDE data set specified in the JCL by the //EXCLUDE DD statement. The format of the EXCLUDE entries are:

```
//EXCLUDE DD *  
  
xxxxxxx,xxxxxxx,xxxxxxx,xxxxxxx,....
```

where xxxxxxx is a combination of "?" characters and the unmasked key characters to be excluded. In the example of the VTAM terminal name masking cited above, to exclude all terminals in location 23, the selection mask specified in the EXCLUDE data set would contain "???23??. The "?" characters in the EXCLUDE selection statements correspond to the "*" positions in the modified keys.

ISPF/PDF Model 204 Subsystem Interface

The CIMS Capacity Planner Model 204 Subsystem is supported by an ISPF/PDF interface that greatly simplifies the task of invoking the reporting facilities of the CIMS Capacity Planner system. The ISPF/PDF interface supports running and viewing the full range of CIMS Capacity Planner reports on-line and, in addition, provides the option to generate a series of on-line GDDM Graphs, generate the full range of batch reports to be printed on the system printer, and generate input to the PC Presentation Graphics subsystem (Harvard Graphics).

The ISPF/PDF interface is menu driven, with separate sets of menus for each CIMS Capacity Planner subsystem (Workload, DASD, M204, etc.).

On-line Graphs

The following Model 204 GDDM graphs are available on-line through the ISPF Interface:

Model 204 Transaction Response Graph

The Transaction Response graph depicts the average response time for all Model 204 transactions, divided into CPU, and WAIT time. The graph shows the prime, early, and late shift averages by hour, day, week, or month within the range of dates specified in the report request.

Model 204 Transaction Volume Graph

The Model 204 Transaction Volume Graph depicts the average volume for all transactions executed. Within the range of dates and the period specified in the report request, the graph shows the prime, early, and late shift averages by hour, day, week, or month.

Model 204 Terminal Activity Graph

The Model 204 Terminal Activity Graph depicts the average number of Model 204 terminals active within the range of dates specified in the report request. The graph shows the average number of terminals active during the prime, early, and late shifts by hour, day, week, or month, depending upon which is requested.

On-line Reports

The following reports are produced on-line through the CIMS Capacity Planner Model 204 ISPF Interface.

Statistical Reports

Model 204 Transaction Statistics Report

The Transactions Statistics Report provides a list of the Model 204 users who rank among the top in the following categories:

- Most transactions executed
- Required the most cumulative transaction residence time
- Consumed the most CPU time
- Added the most records to the data base
- Deleted the most records from the data base
- Sorted the most data base records
- Issued the most server reads
- Issued the most server writes
- Incurred the most EXCPs

- Handled the most directed lines
- Handled the most terminal lines
- Issued the most find requests
- Issued the most direct searches
- Issued the most record reads

Model 204 Terminal Statistics Report

The Model 204 Terminal Statistics Report shows which user's terminals were among the top 40 in the following categories:

- Submitted the most transactions
- Accumulated the most active time (a terminal is active during each 15 period throughout the day if it is actively doing work. Terminals that are logged on, but not submitting transactions are not considered to be active).

Model 204 Summary Analysis Report

The Summary Analysis Report contains key summary information required to determine the magnitude of the Model 204 workload and how well the Model 204 subsystem is performing. The report provides the following information for the time period specified in the report request:

- The period measured by date and time
- The peak number of concurrent active terminals (high watermark)
- The average number of Model 204 transactions processed per minute during the prime shift
- The average number of Model 204-related EXCPs per second during the prime shift
- The average transaction elapsed time (in seconds) during the prime shift
- The average terminal response time (in seconds) during the prime shift
- The average ratio of transactions submitted during the early shift compared to the prime shift
- The average ratio of transactions submitted during the late shift compared to the prime shift
- The average ratio of terminals active during the early shift compared to the prime shift
- The average ratio of terminals active on the late shift compared to the prime shift
- The average percentage of the CPU time consumed by Model 204 and the Model 204 applications during each shift within the measured period
- The average number of transactions processed by shift for the measured period

The 25 most frequently executed transactions are listed individually. All remaining transactions are reported collectively under the title of "OTHER". The total number of transactions are reported by shift.

Model 204 Trends Analysis Report

The on-line Trends Analysis Report shows the trends in selected characteristics of the M204 usage and performance that are indicators of the future Model 204 workload. Trends are reported for:

- CPU utilization, including a breakdown in terms of Model 204 overhead and task related usage
- The average number of transactions executed per minute during prime shift
- The average number of EXCP's per second during prime shift
- The average response time during prime shift
- The ratio of early shift to prime shift transactions
- The ratio of late shift to prime shift transactions
- The average ratio of active terminals on early shift compared to prime shift
- The average ratio of active terminals on late shift compared to prime shift

Model 204 Exception Analysis Report

The Model 204 on-line Exception Analysis Report shows, for the period specified in the report request, any 15 minute intervals during which the following predefined limits were exceeded. The limits are specified in the CPPR.PARMLIB member sidXCPT.

- Average transaction processing time
- Average terminal response time

Model 204 Ad Hoc Report

The Model 204 Ad Hoc Report contains the following information for each record selected for inclusion in the report.

- The Model 204 Last ID (Transaction Identifier)
- The time that the transaction completed execution
- The User ID
- The Account Name
- The Model 204 Procedure Name
- The elapsed time for the transaction

- The CPU time used
- The Model 204 I/O count
- The counts of records added and deleted
- The server size
- The number server reads and writes
- The number terminal lines read and written

Batch Reports

A Batch Job submission panel is provided so that the full range of M204 batch reports and graphs described earlier under *Model 204 Workload Analysis and Performance Reports on page 10-3*, are produced through the ISPF Interface. This considerably simplifies the processes of operating CIMS Capacity Planner.

Batch GDDM Graphs

Several options for GDDM graphs have been added to CIMS Capacity Planner in order to expand its capabilities to include all currently available CIMS Capacity Planner graphs created as input to Harvard Graphics and to make it easy to process GDDM Charts through Batch JOBS instead of limiting the process to the ISPF interface. The options are:

- Charting an HGDLIB data point member through GDDM
- Sending a chart to a GDDM Printer in Batch
- Specifying several Print controls in the Batch JOB
- Selecting a GDDM Template for the Chart (Batch)
- Saving the Data portion of the Chart (Batch)
- Saving the Chart as a GDF File (Batch)

These options are explained in detail below.

Charting an HGDLIB member

Members in the HGDLIB represent the data for individual Harvard Graphics charts. As originally designed, the data members are created by Batch JOBS, downloaded to a PC and combined with Harvard Graphics Templates to form a graph. Although there are several dozen GDDM graphs available through CIMS Capacity Planner, there are perhaps ten times as many available through Harvard Graphics. It is now possible to chart any of the HGDLIB members through GDDM to form a GDDM graph.

Through the ISPF interface to CIMS Capacity Planner, go to the Utilities section (Option "U") and select Option 6 (CPPRGDDM). At that panel, enter the name of the HGDLIB member you want to use as the data portion of a GDDM graph, along with a TITLE for the graph and a Template name (Optional) and press <ENTER>. The HGDLIB member is imported into a GDDM graph.

Printing a GDDM Chart in Batch

It is now possible to produce a GDDM chart and send it directly to a GDDM printer without going through the ISPF interface. For example, you might want to set up a procedure whereby a Batch JOB is automatically submitted to produce several GDDM graphs on a monthly basis and to have them printed overnight and available for presentation the next morning. CIMS Capacity Planner allows this through the addition of several key phrases.

```
GDDM BATCH PROCESS=YES
```

This key phrase is added to the SYSIN to inform CIMS Capacity Planner that the operation is taking place in Batch and no terminal is present for interactive processing.

```
GDDM PRINTER NAME=pppppppp
```

This key phrase is added to the SYSIN to supply the name of the GDDM printer ("pppppppp") to which the output is to be directed for printing or plotting. Please ensure that the GDDM Printer Queue is correctly identified through the ADMPRNTQDD Statement in the JCL.

Several members have been added to the Distribution CNTL file that demonstrate the manner in which these new key phrases are to be used. The member named "GDDMBAT" shows how to chart an HGDLIB member through GDDM via a Batch JOB. The member named "GDDMBATM" shows how to produce any of the Model 204 GDDM Graphs in Batch.

GDDMBAT Parmlib Member

```
//SSAGDDM JOB (...), 'SSA', CLASS=A, MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH, SENDING
/* A MEMBER FROM THE HGDLIB (IPO1201) THROUGH GDDM TO CREATE A
/* MEMBER IN THE ICUDATA LIBRARY (TEST001) WHICH MAY THEN BE
/* SENT TO A GDDM PRINTER (PPRINT01).
/* * * * * *
//STO      EXEC PGM=IKJEFT01, REGION=5000K, DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB, DISP=SHR
//          DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT  DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//SYSUT1   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//HGDLIB   DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.HGDLIB
//ADMSYMBL DD DISP=SHR, DSN=SYS1.GDDMSYM
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR, DSN=SYS1.GDDMGDF
/* THE FOLLOWING NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR, DSN=SYS1.GDDM.REQUEST.QUEUE
//ADMCDATA DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR, DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
//SYSTSIN  DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD00)'
//SYSIN    DD *
GRAPH PERIOD=DAILY
TITLE=CPU STUFF
GDDM BATCH PROCESS=YES
GDDM FORM NAME=PIECHART
GDDM CHART NAME=TEST001
*GDDM GDF NAME=GDFCHT01      /* TO SAVE THIS CHART AS A GDF FILE */
*GDDM PRINTER NAME=PPRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*                               | | | % CHARACTER LENGTH
```

```

*           | | % CHARACTER WIDTH
*           | HORIZONTAL OFFSET
*           VERTICAL OFFSET

//INCLUDE DD *
IPO1201

GDDMBATM Parmlib Member

//SSAGDDM JOB (...),'SSA',CLASS=A,MSGCLASS=X
/* * * * * *
/* THIS JOBSTREAM IS USED TO INVOKE THE GDDM ICU IN BATCH FOR ALL OF
/* THE SUPPORTED M204 CHARTS AND SENDING THEM TO A PRINTER.
/* * * * * *
//ST1      EXEC PGM=IKJEFT01,REGION=5000K,DYNAMNBR=64
//STEPLIB DD DSN=&PREFIX.CPPR.Vnnn.LOADLIB,DISP=SHR
//         DD DSN=WHATEVER YOUR GDDM LOADLIB NAME IS (SYS1.GDDMLOAD)
//SYSPRINT DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//SYSNAP   DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//CPPRERT  DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.CPPRERT
//CPPRPARM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.PARMLIB
//ONLINE   DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ONLINE.WKLD
//INDEX    DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.INDEX.WKLD
//ADMSYMBL DD DISP=SHR,DSN=SYS1.GDDMSYM
//ADMCDATA DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUDATA
//ADMCFORM DD DISP=SHR,DSN=&PREFIX.CPPR.Vnnn.ICUFORMS
/* THE FOLLOWING DD NAME REFERS TO THE GDDM GDF MEMBER LIBRARY
//ADMGDF   DD DISP=SHR,DSN=SYS1.GDDMGDF
/* THE FOLLOWING DD NAME REFERS TO THE GDDM PRINTER QUEUE
//ADMPRNTQ DD DISP=SHR,DSN=SYS1.GDDM.REQUEST.QUEUE
//SYSTSIN  DD *
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD60)' /* M204 RESPONSE /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD61)' /* TRANSACTION VOLUME /*
CALL '&PREFIX.CPPR.Vnnn.LOADLIB(SSA1GD62)' /* TERMINAL VOLUME /*
//SYSIN    DD *
SELECTED SYSTEM=*
BEGIN DATE=03/01/2005
END DATE=03/02/2005
GRAPH PERIOD=DAILY

```

```
GDDM BATCH PROCESS=YES
*GDDM FORM NAME=PIECHART
*GDDM CHART NAME=TEST0001
*GDDM GDF NAME=GDFCHT01 /* TO SAVE THIS CHART AS A GDF FILE */
GDDM PRINTER NAME=PRRINT01 /* TO SEND THIS CHART TO THE PRINTER */
*GDDM PRINT CONTROLS=00,00,100,100
*           | | | % CHARACTER LENGTH
*           | | % CHARACTER WIDTH
*           | HORIZONTAL OFFSET
*           VERTICAL OFFSET
```

Specifying GDDM Batch Print Controls

It is possible to specify certain Print Controls that are to take effect when printing a chart in Batch. The key phrase:

```
GDDM PRINT CONTROLS=xx,yy,wid,len
```

This parameter allows you to specify the Vertical offset ("xx"), the Horizontal offset ("yy"), the % character width ("wid") and the % character length ("len") in order to accommodate different printers and plotters.

Selecting a GDDM Template (ADMCFORM)

A user might want to tailor the GDDM charts to local specifications by creating a series of unique GDDM Templates in the ICUFORMS library. In order to invoke these local templates in a Batch JOB, a new key phrase:

```
GDDM FORM NAME=nnnnnnnn
```

The GDDM FORM NAME parameter can be added to the SYSIN, where the "nnnnnnnn" is the name under which the **template** was saved in the ICUFORMS library.

Saving a GDDM Chart (ADMCDATA)

It is also possible to save the data in GDDM format for later retrieval. This is done via the key phrase:

```
GDDM CHART NAME=dddddddd
```

where the "dddddddd" is the name under which the **data** is saved in the ICUDATA library.

Saving a GDF File (ADMGDF)

It is also possible to save the entire chart in GDF format for later retrieval. This is done via the key phrase:

```
GDDM GDF NAME=eeeeeeee
```

where the "eeeeeeee" is the name under which the graph is saved in the GDF library.

Presentation Graphics Interface

CIMS Capacity Planner provides the ability to produce a wide variety of high quality graphs that depict virtually all aspects of the M204 workload, the level of performance being experienced by the user community, and the trends required to support your capacity planning efforts.

The graphs are produced using a combination of programs operating on both the MVS host computer and an IBM-compatible PC. The data points required to produce the various graphs are generated through the use of a series of host-resident programs that extract and summarize the required Model 204 workload, performance, capacity, and trends data from the On-line Performance Data Base. The points required to produce each graph are then stored in a separate and distinct member of a MVS partitioned data set designated as HGDLIB. The members of the HGDLIB are down-loaded to the PC where the graphics support software, Harvard Graphics—release 3.0, plots each set of data points in accordance with a set of characteristics specified by a graph template.

The resulting graphs are of high quality and are easy to read and understand. Graphs are provided for all the CIMS Capacity Planner subsystems. Virtually all aspects of the workload, the performance, and the trends are supported by the presentation graphics subsystem.

The Harvard Graphics templates, which describe the characteristics of the various CIMS Capacity Planner graphs, are distributed on a floppy disk and can be used directly from the floppy disk or loaded onto a hard disk on the PC. The floppy disk also contains a series of PC scripts that, when customized to conform to your data set naming conventions, are used to down-load the data points to the PC.

Graph Production Procedures

The process of generating the CIMS Capacity Planner M204 graphs through the Harvard Graphics System is as follows:

- Determine the graphs to be produced from the list of Workload graphs specified in the list set forth in *Graph Descriptions on page 10-48* below.
- Compile a composite list of the data point generation programs that must be run to produce the required data points.
- Determine the period(s) to be graphed (hourly, daily, weekly, or monthly) and set up a job step to run each required program for each graph period to be produced.
- Run the required data point generation programs
- Down-load the data points from the host to the PC
- Invoke the Harvard Graphics system on the PC
- Select the proper graph template
- Import the data points into the template

The Harvard Graphics system then produces the graph.

Naming Conventions

The process of combining the graph characteristics specified through the templates with the data points requires that each set of points be associated with its corresponding template. This is accomplished through the naming of the data points and the templates.

Naming the HGDLIB Members

The HGDLIB members are given a name consisting of up to eight characters in the form:

"&sidpiq" where:

- "&sid" is the SMF SID for the MVS system to which the data pertains
- "p" is the time period specified for the graph. The possible values are:
 - "1"—Hourly
 - "2"—Daily
 - "3"—Weekly
 - "4"—Monthly

- "iiq" is the graph identifier that consists of two alpha-numeric characters. The "pii" portion of the HGDLIB member name is used to identify the data points on both the host and the PC. It is also used in naming the corresponding templates as indicated below under *Presentation Graphics Distribution Disk on page 10-46*. The qualifier "q" might or might not be required, depending upon the nature of the graph. Some series of graphs consist of a summary graph, along with a separate set of individual graphs depicting the data that makes up the summary graph. The qualifier is generally used to specify the individual sets of data.

Naming the PC Data Point Files

The files into which the HGDLIB members are down-loaded are given the same names as the HGDLIB members, except that a file type of "DAT" is appended to conform to the following format:

&sidpii.DAT, where the "p" and the "iiq" components of the name are identical the HGDLIB naming convention.

Naming the Graph Templates

The graph templates are named using the same convention as the HGDLIB members except that a file type of TPL is appended to the name:

HGDpii.TPL

As can be seen from the above descriptions of the naming of the HGDLIB members, the data point files on the PC, and the Harvard Graphics templates, the "pii" portion of the names are common to all components of each graph. By knowing which member of HGDLIB contains the data for a specific graph, the graph template can easily be determined. The User Manual contains a separate section pertaining to the Presentation Graphics within each subsystem that specifies, by graph, which files and templates are required.

Naming the BAT Files

A set of BAT files are included with the graph templates pertaining to each subsystem. The BAT files contain the basic commands required to down-load the data points from HGDLIB to the PC using IND\$FILE. The BAT files are modified as required to fit the naming conventions of your installation.

The names of the BAT files all conform to a common format:

HGDpsubg where:

- "HGD" is common to all the BAT files
- "p" specifies hourly, daily, weekly, or monthly as specified above

- "subg" specifies the subsystem to which the graphs apply and the series of graphs supported by the data to be down-loaded:
 - "WKLD"—Workload graph data
 - "DASM"—DASM graph data
 - "CICS"—CICS graph data
 - "CICE"—CICS Extension graph data
 - "IDMS"—IDMS graph data
 - "IDME"—IDMS Extension graph data
 - "IMSG"—IMS graph data
 - "NETG"—Network graph data
 - "NETE"—Network Extension graph data
 - "NETL"—Network Physical Lines graph data
 - "NETN"—Network Physical NCP graph data
 - "DB2S"—DB2 Subsystem graph data
 - "DB2U"—DB2 Connect ID graph data
 - "M204"—Model 204 graph data

Presentation Graphics Distribution Disk

The graph templates and the BAT files are distributed by the CIMS Lab on CD-Rom. The structure of the diskette is described below:

ROOT DIRECTORY							
WORKLOAD	DASM	CICS	IDMS	IMS	NETWORK	DB2	M 204
HOURLY		HOURLY	HOURLY	HOURLY	HOURLY	HOURLY	HOURLY
DAILY		DAILY	DAILY	DAILY	DAILY	DAILY	DAILY
WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY	WEEKLY
MONTHLY	MONTH- LY	MONTH- LY	MONTH- LY	MONTHLY	MONTHLY	MONTHLY	MONTHLY

Each of the sub-directories contains the templates for the graphs produced for that subsystem for the graph period specified, as well as the .BAT file(s) necessary to download the HGDLIB data point members from the Host processor (assumes the use of IND\$FILE for the file transfer). The naming conventions are as follows:

.BAT File Naming Conventions

The .BAT filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1= Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a subsystem identifier:

- WKLD=Workload
- DASM=DASM
- CICS=CICS
- CICE=CICS Extension
- IDMS=IDMS
- IDME=IDMS Extension
- IMS=IMS
- NETG=Network
- NETE=Network Extension
- NETL=Network Physical Lines
- NETN=Network Physical NCP

For example, the .BAT file used to download the HGDLIB members required to produce the Workload Daily graphs

would be named HGD2WKLD.BAT.

.TPL File Naming Conventions

The Harvard Graphics Template (.TPL files) Filenames all begin with the 3 characters "HGD" followed by a 1 character graph period indicator:

- 1=Hourly
- 2=Daily
- 3=Weekly
- 4=Monthly

The graph period identifier is followed by a 2 character graph identifier. In some cases, a 1 character qualifier is added extending the graph identifier to 3 characters. Please refer to the Graphic Report Tables for the specific template names.

Graph Periods

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for a description of the hourly, daily, weekly, and monthly graph periods.

Graph Descriptions

Please reference Chapter 3, CIMS Capacity Graphics, in the *CIMS Capacity Planner Reference Guide*, for descriptions of each of the graphs produced by the CIMS Capacity Planner Model 204 Subsystem.

Graphic Report Tables

The following tables summarize the relationship of the members in the HGDLIB library, the Harvard Graphics templates, and supporting members in the CPPR PARMLIB to specific graphs produced by the Harvard Graphics program or a similar PC-based graph program.

Model 204 Subsystem Regular Hourly Graphs

These graphs represent all Model 204 Regions executing within the &sid system.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1M20G	&sid1M0	(none)	Model 204 Transaction Response
	&sid1M1	(none)	Model 204 Transaction Volume
	&sid1M2	(none)	Model 204 CPU Utilization
	&sid1M3	(none)	Model 204 I/O Activity
	&sid1M4	(none)	Model 204 Terminal Activity
	&sid1M5	(none)	Model 204 Transactions/Terminal
	&sid1M6	(none)	Model 204 Transaction/Day
	&sid1MA	(none)	Model 204 Terminal Response
SSA1M20E	&sid1MB	2040,204R	Model 204 Response by Organization
	&sid1MC	204T	Model 204 Volume by Organization
	&sid1MD	204R	Model 204 Volume by Response

&sid is the SMF ID (e.g, IPO1)

Model 204 Subsystem Regular Daily Graphs

These graphs represent all Model 204 Regions executing within the &sid system.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1M20G	&sid2M0	(none)	Model 204 Transaction Response
	&sid2M1	(none)	Model 204 Transaction Volume
	&sid2M2	(none)	Model 204 CPU Utilization
	&sid2M3	(none)	Model 204 I/O Activity
	&sid2M4	(none)	Model 204 Terminal Activity
	&sid2M5	(none)	Model 204 Transactions/Terminal
	&sid2M6	(none)	Model 204 Transaction/Day
	&sid2M7	(none)	Model 204 Transactions/Shift
	&sid2M8	(none)	Model 204 Terminals/Shift
	&sid2M9	(none)	Model 204 Terminal Response
SSA1M20E	&sid2MB	2040,204R	Model 204 Response by Organization
	&sid2MC	204T	Model 204 Volume by Organization
	&sid2MD	204R	Model 204 Volume by Response

&sid is the SMF ID (e.g, IPO1)

Model 204 Subsystem Regular Weekly Graphs

These graphs represent all Model 204 Regions executing within the &sid system.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1M20G	&sid3M0		Model 204 Transaction Response
	&sid3M1	(none)	Model 204 Transaction Volume
	&sid3M2	(none)	Model 204 CPU Utilization
	&sid3M3	(none)	Model 204 I/O Activity
	&sid3M4	(none)	Model 204 Terminal Activity
	&sid3M5	(none)	Model 204 Transactions/Terminal
	&sid3M6	(none)	Model 204 Transaction/Day
	&sid3M7	(none)	Model 204 Transactions/Shift
	&sid3M8	(none)	Model 204 Terminals/Shift
SSA1M20E	&sid3M9	(none)	Model 204 Terminal Response
	&sid3MB	(none)	Model 204 Response by Organization
	&sid3MC	204T	Model 204 Volume by Organization
	&sid3MD	204R	Model 204 Volume by Response

&sid is the SMF ID (e.g, IPO1)

Model 204 Subsystem Regular Monthly Graphs

These graphs represent all Model 204 Regions executing within the &sid system.

CPPR Program Name	HGDLIB Member	CPPRPARM Library Member	Report Description
SSA1M20G	&sid4M0	(none)	Model 204 Transaction Response
	&sid4M1	(none)	Model 204 Transaction Volume
	&sid4M2	(none)	Model 204 CPU Utilization
	&sid4M3	(none)	Model 204 I/O Activity
	&sid4M4	(none)	Model 204 Terminal Activity
	&sid4M5	(none)	Model 204 Transactions/Terminal
	&sid4M6	(none)	Model 204 Transaction/Day
	&sid4M7	(none)	Model 204 Transactions/Shift
	&sid4M8	(none)	Model 204 Terminals/Shift
	&sid4M9	(none)	Model 204 Terminal Response
SSA1M20E	&sid4MB	2040,204R	Model 204 Response by Organization
	&sid4MC	204T	Model 204 Volume by Organization
	&sid4MD	204R	Model 204 Volume by Response

&sid is the SMF ID (e.g, IPO1)

Tape Subsystem

Tape Subsystem Overview	11-2
Tape Volume Activity Reports	11-2
Tape Volume Activity Graphs	11-3
Tape Drive Activity Reports	11-3
Tape Drive Activity Graphs	11-4
Input Data Supported	11-5
SMF Data (SSA1WKLD)	11-5
Tape Support for VTS	11-6
Tape Management System Control Information (SSA1TAPW)	11-6
The Reporting Component	11-7
Creating Tape Volume Ad Hoc Reports	11-7
Creating Tape Volume Activity Reports	11-8
Creating Tape Drive Activity Reports	11-8
The PC Graphing Component	11-9
Graph Periods	11-9
Graph Descriptions	11-9
Tape Volume Activity Graphs	11-9
Using Batch JCL from the Distribution CNTL Library	11-9
Using the Batch Panel in the ISPF/PDF Interface	11-9
Tape Drive Activity Graphs	11-10
Using Batch JCL from the Distribution CNTL Library	11-10
Using the Batch Panel in the ISPF/PDF Interface	11-10
Tape Ownership Graphs	11-10
Using Batch JCL from the Distribution CNTL Library	11-10
Using the Batch Panel in the ISPF/PDF Interface	11-10

Tape Subsystem Overview

The Tape Subsystem deals with magnetic tape volumes, both reels and cartridges, as well as physical tape drives. It keeps track of the manner in which the tape volumes are used (the percentage of space used by volume) and how many of each category of tape volume (active, inactive, scratch) reside in the Tape Library. It further keeps track of the manner in which the tape drives are used, whether or not the drives are over utilized, whether mounts are satisfied in a reasonable time and the rate at which the tape drives are transferring data while they are allocated.

Input data for the Tape subsystem comes from a variety of sources, including:

- SMF Type 21 records
- RMF Type 74 records
- The control file from a Tape Management System, such as:
 - TMS (Computer Associates)
 - RMM (IBM)
 - Control/T

The data is loaded into a series of tables in the Online Performance Database where it is available to the various reporting and graphing programs in the Tape subsystem, as well as to the ISPF/PDF Interface portion of the Tape subsystem.

Tape Volume Activity Reports

Tape volume activity reports are produced either by the Tape subsystem data reduction program (SSA1TAPW) when data is reduced into the Online Performance Database or through the Tape Volume reporting program (SSA1TAVR).

Tape Storage Allocation by Volume Type Report

This report shows, based on OWNERID, how many volumes of each type (Reel, 3480 Cartridge, 3490 Cartridge, 3590 Cartridge and VTS) are assigned for use and how much cumulative storage is owned by you. It also shows the percentage of storage that is unused and how much has been unreferenced for 60, 120 and 240 days.

Tape Volume Summary Report

This report shows, for any given Tape Library, how many volumes of each type (Reel, 3480 Cartridge, 3490 Cartridge, 3590 Cartridge and VTS) are available, how many are designated as Scratch, Active and Inactive, how many volumes are less than 1% used, less than 5% used, less than 10% used, less than 25% used, less than 50% used, less than 75% used and over 75% used. And finally, how many volumes contain 1 dataset, 2 datasets, 3 datasets and more than 3 datasets.

Tape Volume Activity Graphs

The presentation quality tape volume activity graphs are produced on a PC using Harvard Graphics or any graphing package that supports Comma Separated Value (delimited ASCII) files. Alternatively the graphs can be produced by GDDM on the mainframe if desired. Data point members are built in the HGDLIB library by batch programs that execute on the host machine. Each data point member represents the data for a specific graph. The data point members are subsequently down loaded to the PC and are combined with distributed templates to form high quality graphs.

Tape volume activity graphs allow you to compare the characteristics of the inventory in the Tape Library over time.

Refer to Chapter 3 of the *CIMS Capacity Planner Reference Guide* for detailed descriptions and samples of the Tape Volume Activity Graphs.

Tape Drive Activity Reports

Tape drive activity reports deal with individual tape drives by address (CCUU). INCLUDE/EXCLUDE processing is supported for all reports. In order to Include or Exclude a specific tape drive, please specify the four characters "TAPE" plus the hexadecimal tape drive address (CCUU).

Example

```
//INCLUDE DD *  
TAPE0580
```

The statements above would include only the tape drive at address 0580 in any selected report. Normal CIMS Capacity Planner INCLUDE/EXCLUDE rules apply. That is, you can specify multiple entries on one or more lines, separated by commas. Wild cards can be used. That is, to exclude all tape drives from 580-58F, specify:

```
//EXCLUDE DD *  
TAPE058*
```

Tape Drive Exception Analysis

This report shows, for each fifteen minute period, the tape drives whose Mount Pending, Active or Allocated percentages exceeded the thresholds specified in the &sid.XCPT member of CPPR.PARMLIB.

Tape Drive Statistics Report

This report shows, for all included tape drives, the top 40 in terms of Device Allocated time, Device Active time and Mount Pending time. A tape drive is considered to be active when data transfer is actually taking place.

Tape Drive Activity List

This report shows, for all included tape drives, in ascending sort order by address, the average % of time the drive was allocated, the average % of time the device was active, the average % of time a mount was pending for the drive, the total mount count for the drive and the cumulative number of read or write errors that occurred while the drive was active.

Tape Drive Device Allocation Graph

This graph shows, by individual tape drive (TAPECCUU), for each fifteen minute period, for each day of the week (Mon, Tue, Wed) the % of the period the tape drive was allocated. Since a separate graph is produced for each selected tape drive, it is strongly recommended that the INCLUDE/EXCLUDE facility be used for these graphs.

Tape Drive Device Active Graph

This graph shows, by individual tape drive (TAPECCUU), for each fifteen minute period, for each day of the week (Mon, Tue, Wed) the % of the period the tape drive was active. Since a separate graph is produced for each selected tape drive, it is strongly recommended that the INCLUDE/EXCLUDE facility be used for these graphs.

Tape Drive Mount Pending Graph

This graph shows, by individual tape drive (TAPECCUU), for each fifteen minute period, for each day of the week (Mon, Tue, Wed) the % of the period the tape drive had a mount pending. Since a separate graph is produced for each selected tape drive, it is strongly recommended that the INCLUDE/EXCLUDE facility be used for these graphs.

Tape Drive Activity Graphs

The presentation quality tape drive activity graphs are produced on a PC using Harvard Graphics or any graphing package that supports Comma Separated Value (delimited ASCII) files. Alternatively the graphs can be produced by GDDM on the mainframe if desired. Data point members are built in the HGDLIB library by batch programs that execute on the host machine. Each data point member represents the data for a specific graph. The data point members are subsequently down loaded to the PC and are combined with distributed templates to form high quality graphs.

Tape drive activity graphs allow you to compare the behavior and usage of tape drives or strings of tape drives over time.

Refer to Chapter 3 of the *CIMS Capacity Planner Reference Guide* for detailed descriptions and samples of the Tape Drive Activity Graphs.

Input Data Supported

There are actually two separate programs involved in building the tables for the Tape subsystem: SSA1WKLD (Workload Data Reduction) and SSA1TAPW (Tape Library Control File input). Because the SSA1WKLD program processes the SMF data during its normal operation, it has been enhanced to support the Tape subsystem as well. The tables thus built are processed by the various tape drive reporting and graphing programs. These tables typically reside in the ONLINE.WKLD Performance Database alongside the workload tables.

A second program, SSA1TAPW, reduces data from one of the Tape Management Systems into Volume Inventory tables that are used by the various tape volume reporting and graphing programs. These tables typically reside in the ONLINE.TAPE Performance Database.

SMF Data (SSA1WKLD)

Two SMF record types are processed by the Workload data reduction program, SSA1WKLD, into Table 67, the Tape Drive Statistics table. Table 67 is subsequently used by the Tape subsystem programs to build reports and graphs. For a detailed description of the specific fields in Table 67, refer to the section in Chapter 1 of the *CIMS Capacity Planner Reference Guide* that deals with the Table Extract utility, SSA1TBLX.

RMF/CMF Type 74 Records

The I/O Device Detail records (RMF/CMF Type 74) provide the Tape subsystem with information regarding Mount Pending conditions, tape drive allocation and tape drive active states for each online tape drive in the system. Granularity for the data depends on the RMF reporting interval, although the data is kept in fifteen minute buckets in the CIMS Capacity Planner Tape Drive Statistics table.

SMF Type 21 Records (ESV)

The Error Statistics by Volume records (SMF Type 21) provide the Tape subsystem with information regarding Mounts and Error counts by individual tape drive. The ESV record is written to SMF each time a tape volume is dismounted.

Tape Support for VTS

VTS (Virtual Tape System) support is implemented for CA1 (TMS), RMM and Control/T through the use of a new PARMLIB member named `sid.VTSX`, where `sid` is the SMF SID of the system. The PARMLIB member has a single entry in the form:

01. VTS

On the first line and one or more Volume Serial Numbers that represent VTS volumes on the second and following lines. A Volume Serial Number may be all six characters or it may be 5 or less terminated by an '*'. For example, if you wanted to include all VOLSERs in the range 900000-939999, you would enter:

```
90*,91*,92*,93*
```

In this case, the program SSA1TAPW would first examine the VOLSER of the tape being processed and, if it is a match with the VTSX-specified criteria, the Volume would be classified as a VTS volume. VTS volumes are always presumed to be 100% full.

Tape Management System Control Information (SSA1TAPW)

Three Tape Management Systems are supported by the Tape subsystem of CIMS Capacity Planner:

- TMS (Computer Associates)
- RMM (IBM)
- Control/T (BMC)

Each of the various input sources has its own unique characteristics. They all, however, provide the Tape subsystem with information regarding the types of volumes in the tape library and their usage.

The input is reduced by the Tape data reduction program, SSA1TAPW, into Table 66, the Tape Volume OWNERID Statistics table and Table 68, the Tape Volume Summary table. Tables 66 and 68 are subsequently used by the Tape subsystem programs to build reports and graphs. For a detailed description of the specific fields in Tables 66 and 68, refer to the section in Chapter 1 of the *CIMS Capacity Planner Reference Guide* that deals with the Table Extract utility, SSA1TBLX.

TMS (CA1) from Computer Associates

A sequential copy of the TMS TMC (Tape Management Control) file is used as input to SSA1TAPW. Please specify:

```
SMFILE=TMS
```

in the SYSIN and identify the TMC input file through the SYSUT1 DD Statement.

RMM from IBM

In order to produce an input data file from RMM, please execute the IBM utility EDGHSKP with `PARM='RPTTEXT,DATEFORM(J)'`. This will create a sequential file that is used as input to SSA1TAPW. Please specify:

```
SMFILE=RMM
```

in the SYSIN and identify the RMM Extract file through the SYSUT1 DD Statement.

Control/T from BMC

A sequential copy of the Control/T Master File (CTLCD5) file is used as input to SSA1TAPW. Please specify:

```
SMFILE=CTLT
```

in the SYSIN and identify the CTLCD5 input file through the SYSUT1 DD Statement.

The Reporting Component

Reports from the Tape subsystem can be divided into two categories:

- Volume oriented reports
- Tape drive oriented reports

Tables for the volume oriented reports reside in the ONLINE.TAPE Performance Database, whereas the tables for the tape drive oriented reports reside in the ONLINE.WKLD Performance Database.

Creating Tape Volume Ad Hoc Reports

Depending on the format of the input, the Tape subsystem supports Ad Hoc reports by Dataset name or, for both EPIC and RMM, Ad Hoc Reports by Volume Serial Number.

The DSNAME Ad Hoc reports show when the dataset was last referenced, how many blocks it occupies on the volume, the DCB characteristics and the approximate amount of space used.

The Volume Ad Hoc reports show when the volume was last referenced, how many datasets reside on the volume, the high level qualifier of the first dataset and the approximate amount of space used on the volume.

Batch JCL from the Distribution CNTL Library

The DTAPADHC member can be used to produce Ad Hoc reports. Simply follow the instructions in the SYSIN stream.

Using the Batch Panel in the ISPF/PDF Interface

There is a separate ISPF/PDF Panel for each of the options:

- DSN Ad Hoc report for TMS
- DSN Ad Hoc report for RMM
- Volume Ad Hoc report for RMM
- DSN Ad Hoc report for EPIC
- Volume Ad Hoc report for EPIC

Creating Tape Volume Activity Reports

Tables for the volume oriented reports reside in the `ONLINE.TAPE` Performance Database. The various reports are selected through the `SYSIN` stream.

Online Reports Available Through the ISPF/PDF Interface

The Tape Volume Summary report is available through the ISPF/PDF Interface, as is the Tape Volume Trends Analysis report. Both reports are created dynamically and displayed through the Browse facility of ISPF/PDF.

Using the Batch Panel in the ISPF/PDF Interface

The Batch Panel of the ISPF/PDF interface allows you to select either the Tape Volume Summary Report or the Tape Storage By OWNERID Report.

Using Batch JCL from the Distribution CNTL Library

The CNTL Member named `DTAPVRPT` contains the JCL necessary to produce either of the Tape Volume reports. The CNTL member named `DTAPTRND` is used to produce the Trends Analysis Report.

Creating Tape Drive Activity Reports

Tables for the tape drive oriented reports reside in the `ONLINE.WKLD` Performance Database. The various reports are selected through the `SYSIN` stream.

Online Reports Available Through the ISPF/PDF Interface

The Tape Drive Activity List is available through the ISPF/PDF Interface, as are the Tape Drive Statistics report and the Tape Drive Exception Analysis. All three reports are created dynamically and displayed through the Browse facility of ISPF/PDF.

Using the Batch Panel in the ISPF/PDF Interface

The Batch Panel of the ISPF/PDF interface allows you to select any of the Tape Drive Activity Reports.

Using Batch JCL from the Distribution CNTL Library

The CNTL Member named DTAPURPT contains the JCL necessary to produce any of the Tape Drive Activity reports.

The PC Graphing Component

General Guidelines

The process of building a graph with Harvard Graphics is greatly simplified with CIMS Capacity Planner. The templates for each graph are distributed on a floppy disk and are copied onto the hard disk of the PC that is used to create the graphs. The template describes the "body" of the graph. That is, the template contains information regarding the type of graph that is to be built (Pie, Bar, Mixed), the colors that represent the various objects, some minimal labeling information (X-Axis, Y-Axis), line sizing and footnotes.

The actual data for each graph is contained in the HGDLIB, a Partitioned Dataset that resides with the Host Processor. The HGDLIB is built by batch jobs that reduce the data from the Performance Database into separate members of the HGDLIB, each representing the data for a specific graph.

These Data Point members are then down-loaded from the Host to a PC using some sort of file transfer protocol (IND\$FILE is one of the more popular methods). Finally Harvard Graphics is started, the template is fetched from the directory where it was stored, and the Data Point Member for the specific graph is imported into the Template to form a graph.

Graph Periods

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide* for a description of the hourly, daily, weekly and monthly graph periods.

Graph Descriptions

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide* for descriptions of each of the graphs produced by the CIMS Capacity Planner Tape Subsystem.

Tape Volume Activity Graphs

Tables for the volume activity graphs reside in the ONLINE.TAPE Performance Database.

Using Batch JCL from the Distribution CNTL Library

The CNTL Member named DTAPVGRF contains the JCL necessary to produce the data point members in the HGDLIB for the Tape Volume Activity Graphs.

Using the Batch Panel in the ISPF/PDF Interface

The Batch Panel of the ISPF/PDF interface allows you to create the JCL that will produce the data point members in the HGDLIB for the Tape Volume Activity Graphs.

Tape Drive Activity Graphs

Tables for the tape drive activity graphs reside in the `ONLINE.WKLD` Performance Database.

Using Batch JCL from the Distribution CNTL Library

The CNTL Member named `DTAPUGRF` contains the JCL necessary to produce the data point members in the `HGDLIB` for the Tape Drive Activity Graphs.

Using the Batch Panel in the ISPF/PDF Interface

The Batch Panel of the ISPF/PDF interface allows you to create the JCL that is used to produce the data point members in the `HGDLIB` for the Tape Drive Activity Graphs.

Tape Ownership Graphs

Tables for the ownership graphs reside in the `ONLINE.TAPE` Performance Database.

Using Batch JCL from the Distribution CNTL Library

The CNTL Member named `DTAPOGRF` contains the JCL necessary to produce the data point members in the `HGDLIB` for the Tape Ownership Graphs.

Using the Batch Panel in the ISPF/PDF Interface

The Batch Panel of the ISPF/PDF interface allows you to create the JCL that is used to produce the data point members in the `HGDLIB` for the Tape Ownership Graphs.

UNIX Reporting Subsystem

Overview	12-2
Input Data Supported	12-2
Using Standard UNIX Reports as Input	12-3
Identifying the UNIX Server	12-4
Consolidating Multiple SAR Files	12-5
Processing the Consolidated Output File	12-5
Using PARMLIB Member UNIXSRVR	12-6
Format of the UNIXSRVR Member Record	12-6
Creating or Editing a UNIXSRVR Member Record	12-7
The Reporting Component	12-8
Generating Reports and Graphs	12-8
UNIX Summary Report	12-9
UNIX Statistics Reports	12-10
UNIX Technical Graph Reports	12-12
UNIX Server Availability Reports	12-13
Tracking UNIX File System Capacity	12-14

Overview

The CIMS Capacity Planner UNIX Reporting Subsystem deals with activity within a specific UNIX server or group of servers, examining the use of UNIX server resources such as Memory and Disk Storage. Processor use is also shown, divided into User, Local, Remote, Waiting for I/O and Idle states. Remote File Sharing statistics are tracked, including Page Invalidation incidents and specific RFS calls. Finally, Disk access statistics are maintained for each arm attached to the server. The various reports and graphs can be selected to address a specific audience: executive, management or technical.

Performance data is kept for each individual UNIX server, although data for several servers can be combined (by owning department or location, for example) into a single composite set of reports or graphs. A Summary report is produced for each server when the data is reduced into the UNIX Performance Database. This summary is used to assess the immediate status of the server. Other reports such as the Disk Statistics report might include more than one server and can be used as the basis for a comparative analysis across servers.

Input Data Supported

The UNIX Reporting Subsystem contains data reduction support for SAR (System Activity Reporter) reports that are produced by UNIX.

Although these reports might be treated differently by different versions of UNIX (Release, platform) the information contained in the various reports is essentially similar. The reports contain information relating to the utilization of the CPU on a specific UNIX server (portion of time running in User mode, System mode, Idle mode, servicing remote requests and servicing local requests). In addition, system call activity is logged, along with the average run queue length and the run queue size. The status of several internal tables is monitored along with the number of process switches that took place during the interval.

CIMS Capacity Planner scans the ASCII version of the reports and extracts specific information that is then summarized into hourly buckets and saved in several tables in the UNIX Performance Database.

Using Standard UNIX Reports as Input

The standard UNIX reports that are available on all UNIX platforms are produced in an ASCII report format suitable for browsing. The reporting interval for each of the reports is specified by you. This interval value can be changed during the generation of the reports (it might be hourly up until 7 in the morning, then every ten minutes until 5 in the afternoon, then back to hourly), but the time the interval ends is always shown on the far left-hand column of each report.

All the reports are produced as a single flat file that is uploaded to the mainframe (ensure that it is converted to EBCDIC). This flat file is used as input to the data reduction module of the UNIX Subsystem.

CIMS Capacity Planner identifies the various report types by scanning the header lines and finding keywords that are standard across platforms. There are, however, several platform-specific variables that must be communicated to CIMS Capacity Planner via the `FILTER` keyword in the `SYSIN` so that CIMS Capacity Planner is familiar with the architecture of the report.

Remote Versus Local Data

Some platforms provide both Remote and Local statistics for the CPU % Busy Report, others report only Local activity. This information is communicated to CIMS Capacity Planner through the first positional parameter of the Filter statement:

- 0 = both Remote and Local statistics
- 1 = Local statistics only

The Count of Date Values in the Header Line

CIMS Capacity Planner must have a date to key on in order to store the data in the proper table. The date of the report is stored in the header line that precedes the entire set of reports passed to CIMS Capacity Planner. Unfortunately, the precise format of the header line differs across platforms. In the first place, the manner in which the header line is identified requires you to inform CIMS Capacity Planner of the first six characters that begin the header line itself. This is done through the CIMS Capacity Planner input statement:

```
DATE LINE TAG=.....
```

where the is the actual first six characters on the header line, including blanks. Upper/lower case is important!

Next, CIMS Capacity Planner must be informed as to whether or not the date in question is the first date on the Header line or the second date on the Header line. This is done by specifying the second positional parameter in the `FILTER` statement:

- 0 = the second date on the header line
- 1 = the first date on the header line

Identifying the UNIX Server

The tables in the Performance Database built by the UNIX data reduction process are keyed by date as noted previously. They are also keyed by UNIX server name in much the same manner as the rest of the CIMS Capacity Planner tables are keyed by SMF SID. The UNIX server name must be represented by a pseudo-SID.

To enable the system to handle individual UNIX servers, you can use the PARMLIB member UNIXSRVR to assign a pseudo-SID to each server. Within the UNIXSRVR member, you can manually assign pseudo-SIDs or you can allow the data reduction program, SSA1UNXW, to assign pseudo-SIDs automatically for the incoming UNIX server names. For more information about the UNIXSRVR member, see [Using PARMLIB Member UNIXSRVR on page 12-6](#).

The pseudo-SID must be 4 characters long and must begin with the letter "U". If the pseudo-SID is assigned automatically, the three characters that follow the U are a number value that ranges from 000–999 (i.e., U000, U001, U002, etc.).

If the pseudo-SID is assigned manually, the remaining characters are user-determined. The second position of the name might represent a location, the third a department, and so forth. For example, the pseudo-SID UCIT might represent a server in the Chicago IT department.

The alternative to using the UNIXSRVR member to assign pseudo-SIDs is to use the following SYSIN statement followed by a pseudo-SID as shown in the following example:

```
SELECTED SYSTEM=UXDE
```

However, this method does not allow you to assign separate pseudo-SIDs for each UNIX server. The SID specified by this statement is used for all servers.

Consolidating Multiple SAR Files

An alternative to processing one SAR file at a time is to combine all of the SAR files into a single stream using the UCCLSARX program. This program concatenates the various input streams of the original SAR data provided the SYSUT1 file. The consolidated output file is written to the SYSUT2 file.

The UCCLSARX program recognizes several versions of SAR files (e.g., SunOS, HP-UX, AIX) and dynamically appends a prefix to the beginning of each SAR record. This prefix is extracted from the SAR file header and contains the file date, originating system, and server name.

To support the various date formats that can appear in the SAR file header, the EXEC statement includes a PARM= field that defines the format of the date. The date format may be specified in one of the following formats:

- PARM=MM/DD/YY (this is the default)
- PARM=MM/DD/CCYY
- PARM=DD/MM/YY
- PARM=DD/MM/CCYY
- PARM=DD.MM.YY
- PARM=DD.MM.CCYY
- PARM=YY/MM/DD
- PARM=CCYY/MM/DD
- PARM=YY.MM.DD
- PARM=CCYY.MM.DD

Processing the Consolidated Output File

The SYSUT2 file containing the consolidated SAR data is processed by the SSAIUNXW program using the SYSIN statement:

```
SMFILE=SARX
```

Using PARMLIB Member UNIXSRVR

PARMLIB member UNIXSRVR enables the SSA1UNXW program to associate a unique pseudo-SID with one or more UNIX servers based on the incoming UNIX server name in the SAR data. Each time the SSA1UNXW program encounters a new UNIX server name, a new record is created in the UNIXSRVR member. This record contains the pseudo-SID, the UNIX server name, and the date of the last SAR file processed from the server (if applicable). (See [Format of the UNIXSRVR Member Record](#)).

The Unnn value derived for each new pseudo-SID is equal to the pseudo-SID in the preceding record plus 1 (i.e., U000, U001, U002, etc.) The Unnn number values range from 000–999.

You can also use the UNIXSRVR member to manually assign a pseudo-SID to a UNIX server. To manually assign a pseudo-SID, you need to create a record in the UNIXSRVR member for the pseudo-SID/server name pair as shown in [Creating or Editing a UNIXSRVR Member Record on page 12-7](#).

For users who have a UNIX Performance Database containing existing pseudo-SIDs, the UNIXSRVR member can be created using the ISPF/PDF editor. The records containing the existing pseudo-SIDs can then be edited as needed (see [Creating or Editing a UNIXSRVR Member Record on page 12-7](#)).

The UNIXSRVR member can also be used to prevent duplicate data from being loaded into the UNIX Performance Database. In the UNIXSRVR member, the date of the most recent SAR file that was processed is included in the line containing the Unnn and server name. The next time a SAR file for the UNIX server is encountered, the date in the file is checked. If the date is not chronologically later than the date of the last SAR file processed, a message is issued and the data is skipped.

If you want to override the duplicate check feature, use the following SYSIN statement:

```
SKIP DUPLICATE CHECKING=YES
```

Format of the UNIXSRVR Member Record

The UNIXSRVR member contains a separate record for each pseudo-SID/server name pair. The format for each record is as follows:

- Columns 1–4 contain the pseudo-SID value (Unnn).
- Columns 5–20 contain the UNIX server name. *Note that UNIX server names are case-sensitive.*
- Columns 21–26 contain the date in YYMMDD format for the last SAR file processed or are blank (000000) if a file has not been processed.

An example UNIXSRVR member is shown on [page 12-7](#). In this example, the Unnn values are UXDE, U000, and U001. The corresponding UNIX server names are p01250e, tinux01, and p01170e, respectively.

Example UNIXSRVR Member

In the following example, the Unnn pseudo-SID in the first record was manually assigned. The Unnn value in the second and third records were dynamically assigned. SAR records for each of these servers were last processed on the dates shown in the last columns.

```
UXDEp01250e      041204
U000tinux01     041203
U001p01170e     041201
```

Creating or Editing a UNIXSRVR Member Record

When you create or edit a UNIXSRVR member, the records must be in the format shown in *Format of the UNIXSRVR Member Record on page 12-6*.

For example, if you want to create a new record to assign the pseudo-SID UXDE to the server p01330e, the record would appear as follows:

```
UXDEp01330e      000000
```

Note that the last file processed date 000000 will be updated when a SAR file for this server is processed.

Creating Alias UNIX Server Names

If you want a name other than the actual UNIX server name to appear in reports, you can assign an alias for the server by replicating a record, retaining the pseudo-SID, and changing the server name as shown in the following example:

```
UXDEp01250e      041204
UXDEoracle server 5 000000
U000tinux01     041203
U000oracle server 6 000000
U001p0117e0     041201
U001oracle server 7 000000
```

In this example, the UNIX server names have been assigned the aliases oracle server 5, oracle server 6, and oracle server 7, respectively. These are the server names that will appear in reports. Note that the last file processed date is not applicable for alias entries and the date is 000000.

The Reporting Component

The following sections describe the reports and graphs produced by the CIMS Capacity Planner UNIX subsystem. For more information about these reports and graphs, including examples, refer to the [CIMS Capacity Planner Reference Guide](#).

Generating Reports and Graphs

If the PARMLIB member UNIXSRVR has been created, you can generate reports for one or more UNIX servers in a single step using the SERVERS file as shown in the following example:

```
//SERVERS DD *  
oracle server 5  
oracle server 6  
oracle server 7
```

If the UNIXSRVR member has been created, all UNIX reports include the server name in the report header. The server name that appears reflects the server name in the SERVERS file. This can be the actual UNIX server name or an alias (see [Creating Alias UNIX Server Names on page 12-7](#)).

If the UNIXSRVR member has not been created, use the `SELECTED SYSTEM=` statement to generate reports for all UNIX servers using one pseudo-SID. That is, only one pseudo-SID can be specified for each set of reports in a single run. This pseudo-SID appears in reports.

UNIX Summary Report

The UNIX Summary Report is produced for all UNIX servers that were specified for the report. The summary report pertains to an "average" UNIX server that is the accumulation of all data for each server included in the population of servers.

Disk statistics are reported separately, however, with the top 10 arm/server combinations being referenced in the following categories:

- Disk percentage busy
- Disk queue time in milliseconds
- Disk response time in milliseconds

The summary report is organized into the following line items:

- The period reported upon
- The average run queue size during prime shift
- The average run queue percent occupied during prime shift
- Swap in operations per second (prime shift)
- Swap out operations per second (prime shift)
- Page in/out operations per second (prime shift)
- Process switches per second (prime shift)
- Local system calls per second (prime shift)
- Local read/write cache hit ratio (prime shift)
- Memory allocation failures (prime shift)
- % CPU Busy by shift divided into:
 - User time
 - Local time
 - Remote time
 - Wait I/O
 - Idle time
- The top 10 arm/server disk % busy by shift
- The top 10 arm/server disk queue time by shift
- The top 10 arm/server disk response time by shift

UNIX Statistics Reports

The UNIX Statistics Reports are produced for all UNIX servers that were specified for the report. The statistics reports pertain to the top 40 UNIX servers included in the population of servers.

Disk statistics are reported differently, however, with the top 40 arm/server combinations being referenced in the following categories:

- Disk percentage busy
- Disk response time in milliseconds
- Disk I/O count

All UNIX Statistics Reports recognize the `BEGIN TIME/END TIME` statements to delimit the time period for the report. That is, all UNIX statistics are kept in hourly increments, and you specify:

```
BEGIN TIME=hh.mm.ss
```

```
END TIME=hh.mm.ss
```

in order to select a specific time period for the report. The default time period is:

```
BEGIN TIME=00.00.00
```

```
END TIME=24.00.00
```

UNIX Disk Statistics Report

This report shows, for the top 40 arm/server combinations as noted above, the following statistics:

- Highest average response time in milliseconds
- Highest percent busy
- Highest I/O count

UNIX Process Table Statistics Report

This report shows, for the top 40 servers in each category, the following process table statistics:

- Maximum entries in the process table
- Maximum percent used
- Process table overflows

UNIX Lock Table Statistics Report

This report shows, for the top 40 servers in each category, the following lock table statistics:

- Maximum entries in the lock table
- Maximum percent used
- Lock table overflows

UNIX INODE Table Statistics Report

This report shows, for the top 40 servers in each category, the following INODE table statistics:

- Maximum entries in the INODE table
- Maximum percent used
- INODE table overflows

UNIX File Table Statistics Report

This report shows, for the top 40 servers in each category, the following file table statistics:

- Maximum entries in the file table
- Maximum percent used
- File table overflows

UNIX Technical Graph Reports

The UNIX technical graph reports are produced for each day of the week (Monday, Tuesday, Wednesday) in the period being reported upon. For example, if a full month's worth of data were being examined, seven separate graphs would be produced; one for all Mondays in the month, another for all Tuesdays and so forth.

Each graph is in the form of a horizontal bar chart, with one bar representing each fifteen minute period of the day. The bars themselves are comprised of dashes (peak values) and asterisks (average values). For example, on a Monday chart, the asterisks represent the average value for that fifteen minute slot for all Mondays over the entire period whereas the dashes represent the peak value for that time slot for one of the Mondays.

Each graph is divided into four quadrants: one quadrant representing a 6 hour period. The time period from midnight to 6 a.m. is positioned at the top left hand side of the page. Six a.m. until noon is positioned at the bottom left hand side of the page. Noon until 6 p.m. is at the top right hand side of the page and 6 p.m. until midnight is at the bottom right hand side of the page.

A column down the middle of each quadrant of the graph presents the actual values represented by the asterisks (averages: left hand side of the column) and the dashes (peaks: right hand side). This column can be turned off by specifying the following in the `SYSIN`.

Example

```
DETAIL=NONE
```

UNIX Buffer Write Activity Graph

This graph presents the number of UNIX Buffer Writes per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Physical Write Activity Graph

This graph presents the number of UNIX Physical Writes per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Logical Write Activity Graph

This graph presents the number of UNIX Logical Writes per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Paging Activity Graph

This graph presents the number of UNIX Page In and Page Out operations in terms of pages per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Swapping Activity Graph

This graph presents the number of UNIX Swap In and Swap Out operations in terms of swaps per second that occurred on a specific server for each fifteen minute period of the day.

UNIX System Calls Graph

This graph presents the number of UNIX System Calls per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Buffer Read Activity Graph

This graph presents the number of UNIX Buffer Reads per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Logical Read Activity Graph

This graph presents the number of UNIX Logical Reads per second that occurred on a specific server for each fifteen minute period of the day.

UNIX Physical Read Activity Graph

This graph presents the number of UNIX Physical Reads per second that occurred on a specific server for each fifteen minute period of the day.

Graph Periods

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide*, for a description of the hourly, daily, weekly and monthly graph periods.

Graph Descriptions

Please reference Chapter 3 of the *CIMS Capacity Planner Reference Guide*, for a description of each of the graphs produced by the CIMS Capacity Planner UNIX Subsystem.

UNIX Server Availability Reports

The SSA1UNGA program creates graphs that show a UNIX server's availability over a specified period of time including the number of server outages and the amount of time lost due to outages.

The SSA1UNRA program creates a report that shows a UNIX server's availability over a specified period of time. The granularity of the report is in fifteen minute increments. The report is in matrix format where an asterisk indicates that the server was up and a blank indicates that the server was down.

Tracking UNIX File System Capacity

The CIMS Data Collector for UNIX can be used to provide information for the CIMS Capacity Planner UNIX subsystem. To process this information in the UNIX Performance Database, the SSA1UNXW program is used with the following SYSIN statement:

```
SMFILE=DASD
```

A table is built that contains the following information for each file system tracked:

- File system mount point.
- Fully qualified device name.
- Amount of space allocated and used in gigabytes.

The SSA1UNXR program is used to provide a report for each selected UNIX server. To generate reports, use the following SYSIN statement:

```
UNIX FILE SYTEM REPORT=YES
```

Example File System Capacity Report

Server Name=underdog		CIMS LAB TESTING ENVIRONMENT		
		12/16/2005		
unix File System Mount Point	Fully Qualified Device Name	G_Bytes Alloc	G_Bytes Used	
/	/dev/hd4	.00	.00	
/	/dev/vg00/lvo13	.13	.03	
/home	/dev/hd1	.27	.10	
/home	/dev/vg00/lvo15	.01	.01	
/opt	/dev/vg00/lvo16	.85	.62	
/ora	/dev/vg00/lvo19	4.19	3.27	
/stand	/dev/vg00/lvo11	.08	.03	
/tmp	/dev/hd3	.03		
/tmp	/dev/vg00/lvo14	.05		
/usr	/dev/hd2	.53	.32	
/usr	/dev/vg00/lvo17	1.15	.59	
/var	/dev/hd9var	.00	.00	
/var	/dev/vg00/lvo18	.83	.48	
Totals:		8.18	5.48	



Index

A

- activity graphs
 - channel 1-30, 1-49
 - CPU 1-47
 - DASD 1-48
 - device 1-32
 - IDMS
 - terminal 4-40
 - transaction 4-8
 - Model 204
 - CPU 10-10, 10-26
 - I/O 10-10, 10-26
 - NCP 5-58
 - Network
 - inbound 5-51
 - inbound message 5-12
 - outbound 5-51
 - outbound message 5-14
 - processor 5-17
 - terminal 5-12, 5-52, 5-63
 - transaction 5-12
 - paging 1-47
 - tape
 - volume 11-3
 - Tape drive 11-4, 11-10
 - TSO 1-48
 - TSO transaction 1-47
 - UNIX
 - buffer read 12-13
 - buffer write 12-12
 - logical read 12-13
 - logical write 12-12
 - paging 12-12
 - physical read 12-13
 - physical write 12-12
 - swapping 12-12
 - Workload
 - CPU 1-29
 - paging 1-30
 - TSO 1-29
- activity lists
 - Tape
 - drive 11-4
- activity reports
 - cluster controller 5-10, 5-55
 - CPU 1-45
 - DASD 1-45
 - DASM 2-42
 - Generic 8-16
 - Network
 - physical 5-55
 - physical line 5-8
 - Tape
 - drive 11-3
 - volume 11-2
 - Tape drive
 - creating 11-8
 - Tape volume
 - creating 11-8
 - Workload
 - JOBNAME CPU 1-21
 - printer 1-26
- Ad Hoc reports
 - CICS 3-7
 - DASM 2-6 to 2-10
 - DB2 7-13 to 7-15
 - Generic 8-15, 8-38
 - IDMS 4-7, 4-15, 4-42

- IMS 6-33
- Model 204 10-8, 10-14, 10-37
- Network 5-6, 5-64
- Tape
 - creation of 11-7
- Workload 1-18 to 1-19, 1-67
- Ad Hoc requests
 - DB2 7-19
 - Workload 1-12
- alias for UNIX server, creating 12-7
- analysis reports
 - IMS 6-32
- Application Profile
 - data collection for 9-14
 - report 9-14
 - values 9-13
- AUTH-ID reports
 - DB2 7-9 to 7-13

B

- batch GDDM graphs
 - IDMS 4-43
 - IMS 6-33
 - Model 204 10-38
 - Network 5-66
 - Workload 1-69
- batch reports
 - DB2 7-36
 - Generic 8-40 to 8-44
 - IDMS 4-23, 4-43
 - IMS 6-33
 - Model 204 10-38
 - Network 5-42
 - submission of 5-65
 - Workload 1-69
 - input parameter statements for 1-38
- Batch Window
 - data collection for 9-6
 - defined 9-4
 - events inside 9-5
 - reports 9-6 to 9-9
- BEGIN DATE
 - parameters 8-58, 8-67
- BEGIN DATE parameters
 - Generic 8-48
- buffer read activity graphs
 - UNIX 12-13
- buffer write activity graphs

- UNIX 12-12

C

- capacity factors
 - contents 9-10
 - data collection for 9-10
 - reports 9-11 to 9-12
- capacity graphs
 - Network
 - physical line 5-14, 5-56
- channel
 - graphs
 - activity 1-49
- CICS
 - graphs 3-11 to 3-12
 - CPU activity 3-12
 - I/O activity 3-12
 - terminal activity 3-12
 - terminal response 3-11
 - transaction activity 3-12
 - transaction response 3-11
 - JCL 3-14 to 3-46
 - data reduction 3-14 to 3-26
 - report production 3-27
 - subsystem installation 3-14
 - lists
 - transaction activity 3-8
 - report production, DD Statements 3-27
 - reports
 - Ad Hoc 3-7
 - exception analysis 3-6
 - response performance 3-4
 - summary analysis 3-3
 - system overview 3-10
 - terminal activity 3-9
 - terminal statistics 3-6
 - transaction profile 3-12
 - transaction statistics 3-4 to 3-6
 - trends analysis 3-13
 - subsystem 3-3
- cluster controller
 - graphs
 - inbound PIU 5-53
 - inbound traffic 5-54
 - outbound PIU 5-53
 - outbound traffic 5-54
 - retransmitted traffic 5-54
 - total traffic 5-54

- reports
 - activity 5-55
 - statistics 5-55
- cluster controller graphs
 - inbound PIU 5-19
 - inbound traffic 5-20
 - Network
 - outbound PIU 5-20
 - outbound traffic 5-21
 - retransmitted traffic 5-22
 - total traffic 5-22
- cluster controller reports
 - activity 5-10
 - statistics 5-9
- commit activity graphs
 - DB2 7-7
- connect name
 - DB2 7-4 to 7-8
- contacting CIMS Lab xvi
- CPPRERT
 - DD Statements 8-45, 8-64
 - JCL 8-55
- CPPRPARM
 - DD Statements 8-46, 8-64
 - JCL 8-55
- CPU
 - graphs
 - activity 1-47
 - ratio 1-31
 - reports
 - activity 1-45
- CPU activity graph
 - CICS 3-12
- CPU activity graphs
 - IDMS 4-9
 - IMS 6-5
 - Model 204 10-10, 10-26

D

- D/B activity graphs
 - IDMS 4-9
- daily graphs
 - IDMS 4-53
 - Model 204
 - regular 10-50
 - Network 5-78
 - extension 5-79
 - physical configuration 5-80

- DAILY OVERVIEW
 - parameters 8-57
- daily overview reports
 - Generic 8-16, 8-41
- DASD
 - graphs
 - activity 1-48
 - device activity 1-32
 - device busy 1-33
 - I/O service time 1-31, 1-49
 - queue delay 1-32, 1-48
 - reports
 - activity 1-45
 - detail 1-28, 1-44, 1-45
 - volume mapping 2-42
 - space allocation
 - by device type 2-12
 - by device volume 2-12
 - volume map
 - by CCHH 2-13
 - by Dsname 2-13
 - workload reports 1-28
- DASM
 - data
 - DCOLLECT 2-4
 - PARMLIB 2-11
 - SMF 2-3
 - VTOC 2-3
 - data reduction 2-3 to 2-10
 - DD Statements 2-29 to 2-31
 - ISPF/PDF 2-42
 - naming conventions 2-44
 - operating JCL 2-29
 - parameters 2-31 to 2-35
 - presentation graphics 2-43
 - reports 2-12 to 2-29
 - activity 2-42
 - Ad Hoc 2-6 to 2-10
 - data reduction 2-6 to 2-10
 - data set activity 2-13
 - data set detail 2-13
 - selected VSAM summary 2-15
 - space allocation 2-42
 - volume activity 2-13
 - volume summary report 2-14
 - sample jobstreams 2-39
 - subsystem 2-3

- data
 - DASM
 - DCOLLECT 2-4
 - PARMLIB 2-11
 - SMF 2-3
 - VTOC 2-3
 - data collection
 - Application Profile 9-14
 - Batch Window 9-6
 - capacity factors 9-10
 - data reduction
 - DASM 2-3 to 2-10
 - reports 2-6 to 2-10
 - DB2 7-17
 - for NPM and NET/MASTER physical
 - configuration data 5-37
 - Generic 8-5, 8-18
 - input parameter statements for 8-47 to 8-50
 - job streams for 8-50
 - IMS 6-7
 - Model 204
 - DD Statements 10-15
 - parameters 10-16 to 10-19
 - Network
 - input parameters 5-27 to 5-31
 - JCL 5-25
 - Workload 1-3
 - input parameters 1-7
 - JCL 1-5
 - reports 1-5
 - sample jobstreams for 1-14
 - data reduction DD statements
 - IDMS 4-16
 - data reduction JCL
 - CICS 3-14 to 3-26
 - DB2 7-20
 - Generic 8-45 to 8-46
 - IDMS 4-11
 - IMS 6-13
 - Model 204 10-12
 - data reduction parameters
 - DB2 7-20, 7-21
 - IDMS 4-17
 - IMS 6-14
 - data set activity report
 - DASM 2-13
 - data set detail report
 - DASM 2-13
- date and time stamp
 - Generic 8-21
- DB2
 - Ad Hoc requests 7-19
 - data reduction 7-17
 - data reduction JCL 7-20
 - data reduction parameters 7-20, 7-21
 - DD Statements 7-24
 - graphs
 - commit activity 7-7
 - plan profile 7-8
 - regular daily 7-47
 - regular hourly 7-46
 - regular monthly 7-49
 - regular weekly 7-48
 - SQL command activity 7-7
 - thread transit time 7-7
 - ISPF/PDF 7-38
 - parameters
 - EXCLUSIVE ELEMENT MASK 7-41
 - GENERIC ELEMENT MASK 7-40
 - INCLUSIVE ELEMENT MASK 7-40
 - presentation graphics
 - distribution disk 7-44
 - presentation graphics JCL 7-45
 - reporting parameters 7-26 to 7-36
 - reports
 - Ad Hoc 7-13 to 7-15
 - AUTH-ID 7-9 to 7-13
 - batch 7-36
 - connect name 7-4 to 7-8
 - connect name summary analysis 7-4
 - exception analysis 7-6
 - job streams 7-36
 - on-line 7-38
 - plan statistics 7-5
 - subsystem trends 7-37
 - subsystem-wide 7-3
 - summary analysis 7-3, 7-38
 - transit time performance 7-5
 - trends 7-16
 - workload and performance 7-3 to 7-16
 - SMF records 7-18
 - subsystem 7-3
- DB2 naming conventions 7-42
- DB2 presentation graphics 7-41
- DCOLLECT data
 - DASM 2-4

- DD Statements
 - CICS
 - report production 3-27
 - DASM 2-29 to 2-31
 - DB2 7-24
 - Generic
 - CCPRPARM 8-46
 - CPPRERT 8-45, 8-64
 - CPPRPARM 8-64
 - EXCLUDE/INCLUDE 8-64
 - HGDLIB 8-64
 - INDEX 8-46, 8-64
 - ONLINE 8-46, 8-64
 - SYSIN 8-46, 8-65
 - SYSNAP 8-46, 8-65
 - SYSPRINT 8-45, 8-63
 - SYSUT1 8-45
 - SYSUT3 8-46
 - IDMS 4-21
 - data reduction 4-16
 - IMS 6-13
 - Model 204
 - data reduction 10-15
 - report production 10-20
 - Network 5-25 to 5-27, 5-40
 - Workload 1-5 to 1-6, 1-34 to 1-37
 - DEBUGON
 - parameters 8-60, 8-68
 - DEBUGON parameters
 - Generic 8-50
 - detail reports
 - DASD 1-28, 1-44, 1-45
 - device busy graphs
 - DASD 1-33
 - disk statistics reports, UNIX 12-10
 - drive activity
 - Tape
 - graphs 11-4, 11-10
 - lists 11-4
 - reports 11-3, 11-8
 - drive device
 - Tape
 - active graphs 11-4
 - allocation graphs 11-4
 - drive exception analysis
 - Tape
 - reports 11-3
 - drive mount pending graphs
 - Tape 11-4
 - drive statistics
 - Tape
 - reports 11-3
- E**
- ELEMENT ACTIVITY LIST
 - parameters 8-57
 - element activity lists
 - Generic 8-43
 - ELEMENT KEY parameters
 - Generic 8-49
 - element keys
 - Generic 8-22
 - END DATE
 - parameters 8-58, 8-59, 8-67
 - END DATE parameters
 - Generic 8-48, 8-49
 - EPIC 11-6
 - error statistics by volume records
 - Tape 11-5
 - ESV records
 - Tape 11-5
 - events
 - Batch Window 9-5
 - exception
 - workload reports 1-27
 - exception analysis reports
 - CICS 3-6
 - DB2 7-6
 - IDMS 4-6, 4-42
 - IMS 6-4, 6-33
 - Model 204 10-7, 10-27, 10-37
 - Network 5-6, 5-50
 - physical 5-55
 - Tape
 - drive 11-3
 - Workload 1-66
 - exception reports
 - Network
 - physical line 5-64
 - session 5-64
 - EXCLUDE/INCLUDE
 - DD Statements 8-64
 - JCL 8-55
 - EXCLUSIVE ELEMENT MASK parameter
 - DB2 7-41

- extension daily graphs
 - IDMS 4-54
- extension hourly graphs
 - IDMS 4-52
- extension monthly graphs
 - IDMS 4-58
- extension weekly graphs
 - IDMS 4-56

F

- file table statistics reports
 - UNIX 12-11
- FILTER
 - parameters 8-66
- Floating Point Values 8-9

G

- general report statements
 - Workload 1-51
- generation options
 - Network
 - batch report 5-46
- Generic
 - data reduction 8-5, 8-18
 - input parameter statements for 8-47 to 8-50
 - JCL 8-45 to 8-46
 - job streams for 8-50
 - date and time stamp 8-21
 - DD Statements
 - CPPRERT 8-45, 8-64
 - CPPRPARAM 8-46, 8-64
 - EXCLUDE/INCLUDE 8-64
 - HGDLIB 8-64
 - INDEX 8-46, 8-64
 - ONLINE 8-46, 8-64
 - SYSIN 8-46, 8-65
 - SYSNAP 8-46, 8-65
 - SYSPRINT 8-45, 8-63
 - SYSUT1 8-45
 - SYSUT3 8-46
 - element activity lists 8-43
 - element keys 8-22
 - input file definitions 8-32
 - ISPF/PDF 8-17
 - JCL
 - batch reports 8-54
 - CPPRERT 8-55
 - CPPRPARAM 8-55

- data reduction 8-45 to 8-46
- EXCLUDE/INCLUDE 8-55
- for basic graph data points 8-63
- INDEX 8-55
- ONLINE 8-55
- SYSIN 8-56
- SYSNAP 8-56
- SYSPRINT 8-54
- job streams 8-62
 - for basic cell graph data points 8-71
- parameters
 - BEGIN DATE 8-48, 8-58, 8-67
 - DAILY OVERVIEW 8-57
 - DEBUGON 8-50, 8-60, 8-68
 - ELEMENT ACTIVITY LIST 8-57
 - ELEMENT KEY 8-49
 - END DATE 8-48, 8-49, 8-58, 8-59, 8-67
 - FILTER 8-66
 - for basic graph data points 8-66
 - for batch reports 8-56
 - GEOGRAPHIC LOCATION 8-49, 8-60, 8-68
 - GRAPH PERIOD 8-69
 - LATE SHIFT FIRST HOUR 8-59, 8-68
 - LOCAL HOLIDAYS 8-50, 8-60, 8-69
 - PRIME SHIFT FIRST HOUR 8-59, 8-67
 - SELECTED DAY 8-59, 8-68
 - SELECTED SYSTEM 8-47, 8-56, 8-66
 - SMFILE 8-47
 - STATISTICS REPORT 8-57
- report generation 8-15
- reports
 - activity 8-16
 - Ad Hoc 8-15, 8-38
 - batch 8-40 to 8-44
 - daily 8-41
 - daily overview 8-16
 - generated from tables 8-15
 - statistics 8-16, 8-40
 - substitution 8-17
- tables
 - user 8-34
 - user defined 8-18, 8-23
 - user table definitions 8-37
- GENERIC ELEMENT MASK parameter
 - DB2 7-40
- GEOGRAPHIC LOCATION
 - parameters 8-60, 8-68

- GEOGRAPHIC LOCATION parameters
 - Generic 8-49
- graph descriptions
 - IDMS 4-52
- GRAPH PERIOD
 - parameters 8-69
- graph periods
 - IDMS 4-51
- graph production procedure
 - Model 204 10-44
- graph production procedures
 - IMS 6-37
- graph production procedures, IDMS 4-48, 5-71
- graphic report tables
 - IDMS 4-52
- graphics
 - IDMS
 - presentation 4-25
 - report tables 4-52
- graphs
 - channel
 - activity 1-49
 - CICS 3-11 to 3-12
 - CPU activity 3-12
 - I/O activity 3-12
 - terminal activity 3-12
 - terminal response 3-11
 - transaction activity 3-12
 - transaction response 3-11
 - cluster controller
 - inbound PIU 5-19, 5-53
 - inbound traffic 5-20, 5-54
 - outbound PIU 5-20, 5-53
 - outbound traffic 5-21, 5-54
 - retransmitted traffic 5-22, 5-54
 - total traffic 5-22, 5-54
 - CPU
 - activity 1-47
 - DASD
 - activity 1-48
 - device activity 1-32
 - device busy 1-33
 - I/O service time 1-31, 1-49
 - queue delay 1-32, 1-48
 - DASM production procedures 2-43
 - DB2
 - commit activity 7-7
 - plan profile 7-8
 - production procedures 7-42
 - regular daily 7-47
 - regular hourly 7-46
 - regular monthly 7-49
 - regular weekly 7-48
 - SQL command activity 7-7
 - thread transit time 7-7
 - IDMS 4-8 to 4-9
 - batch GDDM 4-43
 - CPU activity 4-9
 - D/B activity 4-9
 - daily 4-53
 - descriptions 4-52
 - extension daily 4-54
 - extension hourly 4-52
 - extension monthly 4-58
 - extension weekly 4-56
 - hourly 4-52
 - I/O activity 4-9
 - monthly 4-57
 - on-line 4-40
 - periods 4-51
 - production procedures 4-48, 5-71
 - terminal activity 4-8, 4-40
 - transaction activity 4-8
 - transaction response 4-8, 4-40
 - transaction volume 4-40
 - weekly 4-55
 - IMS 6-5 to 6-6
 - batch GDDM 6-33
 - CPU activity 6-5
 - I/O activity 6-6
 - on-line 6-31
 - subsystem daily 6-42
 - subsystem hourly 6-42
 - subsystem monthly 6-44
 - subsystem weekly 6-43
 - terminal activity 6-5, 6-31
 - transaction activity 6-5
 - transaction response 6-5, 6-31
 - transaction volume 6-31
 - Model 204
 - batch GDDM 10-38
 - CPU activity 10-10, 10-26
 - I/O activity 10-10, 10-26
 - on-line 10-35
 - production procedures 10-44
 - regular daily 10-50

- regular hourly 10-49
 - regular monthly 10-52
 - regular weekly 10-51
 - terminal activity 10-10, 10-26, 10-35
 - terminal response 10-9, 10-27
 - transaction activity 10-10
 - transaction response 10-9, 10-27, 10-35
 - transaction volume 10-35
 - workload and performance 10-9 to 10-10
- NCP
- activity 5-58
 - channel hold queue length 5-19, 5-58
 - free buffer high watermark 5-18, 5-58
 - free buffer low watermark 5-19, 5-58
 - slowdown time count 5-57
- Network 5-11 to 5-23
- cluster controller outbound PIU 5-20
 - daily 5-78, 5-79
 - extension hourly 5-76
 - extension monthly 5-84
 - extension weekly 5-82
 - hourly 5-76
 - inbound activity 5-51
 - inbound message activity 5-12
 - inbound message traffic 5-13
 - inbound traffic 5-52
 - monthly 5-84
 - on-line 5-63
 - outbound activity 5-51
 - outbound message activity 5-14
 - outbound message traffic 5-13
 - outbound traffic 5-52
 - physical configuration daily 5-80
 - physical configuration hourly 5-77
 - physical configuration monthly 5-85
 - physical configuration weekly 5-83
 - physical line capacity 5-14, 5-56
 - physical line inbound traffic 5-15, 5-56
 - physical line outbound traffic 5-15, 5-56
 - physical line retransmitted traffic 5-16, 5-57
 - physical line total traffic 5-17, 5-57
 - processor activity 5-17
 - processor slowdown count 5-18
 - programs 5-45
 - terminal activity 5-12, 5-52, 5-63
 - through ISPF/PDF 5-63
 - total traffic 5-13, 5-53
 - transaction activity 5-12
 - transaction response 5-11, 5-52, 5-63
 - transaction volume 5-63
 - weekly 5-81
- storage activity
- Auxiliary to Central 1-34, 1-51
 - Central to Auxiliary 1-34, 1-51
 - Central to Expanded 1-33
 - Expanded to Auxiliary 1-34, 1-50
 - Expanded to Central 1-33, 1-50
- Tape
- drive activity 11-4, 11-10
 - drive device active 11-4
 - drive device allocation 11-4
 - drive mount pending 11-4
 - ownership 11-10
 - PC components 11-9
 - volume activity 11-3
- TSO
- activity 1-48
 - transaction activity 1-47
 - transaction response 1-48
- UNIX
- buffer read activity 12-13
 - buffer write activity 12-12
 - logical read activity 12-13
 - logical write activity 12-12
 - paging activity 12-12
 - physical read activity 12-13
 - physical write activity 12-12
 - swapping activity 12-12
 - system calls 12-13
- Workload 1-29
- batch GDDM 1-69
 - channel activity 1-30
 - CPU activity 1-29
 - CPU ratio 1-31
 - on-line 1-61 to 1-63
 - paging activity 1-30
 - paging ratio 1-31
 - production procedures 1-74
 - TSO activity 1-29

H

HGDLIB

DD Statements 8-64

hourly graphs

IDMS 4-52

Network 5-76

extension 5-76

physical configuration 5-77

I

I/O activity graphs

CICS 3-12

IDMS 4-9

IMS 6-6

Model 204 10-10, 10-26

I/O device detail records

Tape 11-5

IDMS

data reduction parameters 4-17

DD Statements 4-21

data reduction 4-16

graphs 4-8 to 4-9

batch GDDM 4-43

CPU activity 4-9

D/B activity 4-9

daily 4-53

descriptions 4-52

extension daily 4-54

extension hourly 4-52

extension monthly 4-58

extension weekly 4-56

hourly 4-52

I/O activity 4-9

monthly 4-57

on-line 4-40

periods 4-51

production procedures 4-48, 5-71

report tables 4-52

terminal activity 4-8, 4-40

transaction activity 4-8

transaction response 4-8, 4-40

transaction volume 4-40

weekly 4-55

ISPF/PDF 4-39

JCL

data reduction 4-11

operating 4-11 to 4-39

report production 4-21

subsystem installation 4-11

trends analysis reports 4-24

naming conventions 4-49

presentation graphics 4-25, 4-48 to 4-58

distribution disk 4-51

reports

Ad Hoc 4-7, 4-15, 4-42

additional options 4-36

batch 4-23, 4-43

exception analysis 4-6, 4-42

job streams 4-23

on-line 4-40

parameters 4-27

response performance 4-4

statistical 4-40

summary analysis 4-3, 4-41

terminal statistics 4-6, 4-41

transaction profile 4-9

transaction statistics 4-4, 4-40

trends analysis 4-10, 4-42

workload and performance 4-3 to 4-8

subsystem 4-3

trends analysis 4-34

IMS

control 6-15

data reduction 6-7

data reduction parameters 6-14

graph production procedures 6-37

graphs 6-5 to 6-6

batch GDDM 6-33

CPU activity 6-5

I/O activity 6-6

on-line 6-31

subsystem daily 6-42

subsystem hourly 6-42

subsystem monthly 6-44

subsystem weekly 6-43

terminal activity 6-5, 6-31

transaction activity 6-5

transaction response 6-5, 6-31

transaction volume 6-31

ISPF/PDF 6-30

JCL 6-7

batch report 6-18

data reduction 6-13

subsystem installation 6-7

transaction profile report 6-19

trends analysis summary program 6-20

- naming conventions 6-38
- presentation graphics 6-20, 6-37, 6-40
- reports
 - Ad Hoc 6-33
 - additional options 6-28
 - analysis 6-32
 - batch 6-33
 - exception analysis 6-4, 6-33
 - on-line 6-31
 - parameters 6-22
 - production 6-16
 - response performance 6-3
 - statistical 6-31
 - summary analysis 6-2, 6-32
 - terminal statistics 6-4, 6-31
 - transaction statistics 6-3, 6-31
 - trends analysis 6-6, 6-32
 - workload and performance 6-2
- subsystem 6-2
- system 6-14
- INCLUSIVE ELEMENT MASK parameter
 - DB2 7-40
- INDEX
 - DD Statements 8-46, 8-64
 - JCL 8-55
- input data
 - UNIX
 - sar 12-2
 - sar, consolidating 12-5
 - standard reports used as 12-3
- input data supported
 - Tape 11-5
- input file definitions
 - Generic 8-32
- ISPF/PDF
 - DASM 2-42
 - DB2 7-38
 - Generic 8-17
 - IDMS 4-39
 - IMS 6-30
 - Model 204 10-34
 - Network 5-62
 - Tape 11-9
 - Workload 1-61

J

- JCL
 - CICS 3-14 to 3-46
 - data reduction 3-14 to 3-26
 - report production 3-27
 - subsystem installation 3-14
 - DB2 presentation graphics 7-45
 - Generic
 - CPPRERT 8-55
 - CPPRPARM 8-55
 - data reduction 8-45 to 8-46
 - EXCLUDE/INCLUDE 8-55
 - for basic graphs data points 8-63
 - for batch reports 8-54
 - INDEX 8-55
 - ONLINE 8-55
 - SYSIN 8-56
 - SYSNAP 8-56
 - SYSPRINT 8-54
- IDMS
 - data reduction 4-11
 - operating 4-11 to 4-39
 - report production 4-21
 - subsystem installation 4-11
 - trends analysis reports 4-24
- IMS 6-7
 - batch report 6-18
 - subsystem installation 6-7
 - transaction profile report 6-19
 - trends analysis summary program 6-20
- Model 204
 - data reduction 10-12
 - installation 10-12
 - operating 10-12
 - report generation 10-19
 - trends analysis 10-30
- Network
 - data reduction 5-25
 - operating 5-24
 - report production 5-40
 - SSA1NETW 5-45
 - subsystem installation 5-25
- Tape
 - batch 11-9
- trends reports, for 1-57

- Workload
 - batch reports 1-34
 - data reduction 1-5
 - sample 1-36
- job stream reports
 - IDMS 4-23
- job streams
 - DB2 7-36
 - Generic 8-62
 - data reduction 8-50
 - for basic cell graph data points 8-71
- JOBNAME reports
 - Workload 1-18

L

- LATE SHIFT FIRST HOUR
 - parameters 8-59, 8-68
- lists
 - CICS
 - transaction activity 3-8
 - Generic element activity 8-43
 - Model 204
 - transaction activity 10-8
 - Tape
 - drive activity 11-4
- LOCAL HOLIDAYS
 - parameters 8-60, 8-69
- LOCAL HOLIDAYS parameters
 - Generic 8-50
- lock table statistics reports
 - UNIX 12-11
- logical read activity graphs
 - UNIX 12-13
- logical write activity graphs
 - UNIX 12-12

M

- Model 204
 - DD Statements
 - data reduction 10-15
 - report production 10-20
 - graphs
 - batch GDDM 10-38
 - CPU activity 10-10, 10-26
 - GDDM 10-35
 - I/O activity 10-10, 10-26
 - on-line 10-35
 - production procedures 10-44

- regular daily 10-50
- regular hourly 10-49
- regular monthly 10-52
- regular weekly 10-51
- terminal activity 10-10, 10-26
- terminal response 10-9, 10-27
- through ISPF 10-35
- transaction activity 10-10
- transaction response 10-9, 10-27, 10-35
- transaction volume 10-35
- workload and performance 10-9 to 10-10

ISPF/PDF 10-34

JCL

- data reduction 10-12
- installation 10-12
- operating 10-12
- report generation 10-19
- trends analysis 10-30

lists

- transaction activity 10-8

naming conventions 10-44

parameters 10-12

- data reduction 10-16 to 10-19
- report production 10-22 to 10-31

presentation graphics 10-43

presentation graphics distribution disk 10-46

profiles

- user id 10-11

reports

- Ad Hoc 10-8, 10-14, 10-37
- additional options 10-32
- batch 10-38
- exception analysis 10-7, 10-27, 10-37
- JCL and parameters 10-19
- on-line 10-35
- performance 10-25
- response performance 10-4
- statistical 10-35
- summary 10-24
- summary analysis 10-3, 10-36
- terminal statistics 10-7, 10-25, 10-36
- through ISPF 10-35
- transaction statistics 10-4, 10-25, 10-35
- trends analysis 10-11, 10-37
- workload analysis and performance 10-3

SMF records

- processing of 10-12

terminal activity 10-35

monthly graphs

IDMS 4-57

Model 204

regular 10-52

Network 5-84

extension 5-84

physical configuration 5-85

MVS Release 5.1, special considerations for 1-17

N

naming convention

Model 204 10-44

naming conventions

DASM 2-44

DB2 7-42

IDMS 4-49

IMS 6-38

Network 5-72

Workload 1-75

NCP

graphs

channel hold queue length 5-19, 5-58

free buffer high watermark 5-18, 5-58

free buffer low watermark 5-19, 5-58

slowdown time count 5-57

NCP free buffer watermark graphs

high 5-18

low 5-19

NET/MASTER

physical data 5-39

Network

DD Statements 5-25 to 5-27, 5-40

generation options

batch report 5-46

graphs 5-11 to 5-23

cluster controller outbound PIU 5-20

daily 5-78, 5-79

extension hourly 5-76

extension monthly 5-84

extension weekly 5-82

GDDM 5-63

hourly 5-76

inbound activity 5-51

inbound message activity 5-12

inbound message traffic 5-13

inbound traffic 5-52

monthly 5-84

on-line 5-63

outbound activity 5-51

outbound message activity 5-14

outbound message traffic 5-13

outbound traffic 5-52

physical configuration daily 5-80

physical configuration hourly 5-77

physical configuration monthly 5-85

physical configuration weekly 5-83

physical line capacity 5-14, 5-56

physical line inbound traffic 5-15, 5-56

physical line outbound traffic 5-15, 5-56

physical line retransmitted traffic 5-16, 5-57

physical line total traffic 5-17, 5-57

processor activity 5-17

processor slowdown count 5-18

terminal activity 5-12, 5-52, 5-63

through ISPF/PDF 5-63

total traffic 5-13, 5-53

transaction activity 5-12

transaction response 5-11, 5-52, 5-63

transaction volume 5-63

weekly 5-81

ISPF/PDF 5-62

JCL

data reduction 5-25

operating 5-24

report production 5-40

SSA1NETW 5-45

subsystem installation 5-25

naming conventions 5-72

parameters

data reduction input 5-27 to 5-31

presentation graphics 5-71

distribution disk 5-74

programs

presentation graphics 5-45

trends analysis summary 5-44

records

processing of 5-31 to 5-37

reports

Ad Hoc 5-64

additional options 5-60

batch 5-42

batch submission 5-65

exception analysis 5-6, 5-50, 5-55

on-line 5-63

performance 5-51

physical activity 5-55

- physical exception analysis 5-6
- physical line activity 5-8
- physical line exception 5-64
- physical line statistics 5-7
- physical statistics 5-56
- response performance 5-4
- session exception 5-64
- summary 5-50, 5-64
- summary analysis 5-3
- terminal profile 5-53
- terminal statistics 5-5, 5-51
- trends analysis 5-24
- workload and performance 5-3 to 5-10
- subsystem 5-3
- terminal profiles 5-23
- traffic
 - physical line inbound traffic 5-15
 - physical line outbound traffic 5-15
- NPM
 - physical data 5-38

O

- ONLINE
 - DD Statements 8-46, 8-64
 - JCL 8-55
- on-line graphs
 - IDMS 4-40
 - IMS 6-31
 - Model 204 10-35
 - Network 5-63
 - Workload 1-61 to 1-63
- on-line reports
 - DB2 7-38
 - IDMS 4-40
 - IMS 6-31
 - Model 204 10-35
 - Network 5-63
 - Workload 1-63 to 1-65
- operating JCL
 - DASM 2-29
- operating JCL for batch reports 8-54
- Operations Analysis. *See* Application Profile
- ownership graphs
 - Tape 11-10

P

- paging activity graphs
 - UNIX 12-12
- paging graphs
 - activity 1-30, 1-47
 - ratio 1-50
 - Workload
 - ratio 1-31
- parameters
 - DASM 2-31 to 2-35
 - DB2
 - EXCLUSIVE ELEMENT MASK 7-41
 - GENERIC ELEMENT MASK 7-40
 - INCLUSIVE ELEMENT MASK 7-40
 - Generic
 - BEGIN DATE 8-48, 8-58, 8-67
 - DAILY OVERVIEW 8-57
 - DEBUGON 8-50, 8-60, 8-68
 - ELEMENT ACTIVITY LIST 8-57
 - ELEMENT KEY 8-49
 - END DATE 8-48, 8-49, 8-58, 8-59, 8-67
 - FILTER 8-66
 - for basic graph data points 8-66
 - for batch reports 8-56
 - GEOGRAPHIC LOCATION 8-49, 8-60, 8-68
 - GRAPH PERIOD 8-69
 - LATE SHIFT FIRST HOUR 8-59, 8-68
 - LOCAL HOLIDAYS 8-50, 8-60, 8-69
 - PRIME SHIFT FIRST HOUR 8-59, 8-67
 - SELECTED DAY 8-59, 8-68
 - SELECTED SYSTEM 8-47, 8-56, 8-66
 - SMFILE 8-47
 - STATISTICS REPORT 8-57
 - IDMS
 - data reduction 4-17
 - reports 4-27
 - Model 204 10-12
 - data reduction 10-16 to 10-19
 - report production 10-22 to 10-31
 - Network
 - data reduction input 5-27 to 5-31
 - Workload 1-38 to 1-55
 - data reduction input 1-7
- PARMLIB data
 - DASM 2-11
- PC graphing components 11-9

- performance reports
 - Model 204 10-25
 - Network 5-51
 - TSO 1-44
 - Workload 1-23
 - PGMNAME reports
 - Workload 1-19
 - physical configuration data
 - data reduction for NPM and NET/MASTER 5-37
 - physical data
 - NET/MASTER 5-39
 - NPM 5-38
 - physical exception analysis reports
 - Network 5-6
 - physical read activity graphs
 - UNIX 12-13
 - physical write activity graphs
 - UNIX 12-12
 - plan profile graphs
 - DB2 7-8
 - plan statistics reports
 - DB2 7-5
 - presentation graphics 1-74
 - DASM 2-43
 - DASM distribution disk 2-46
 - DB2 7-41
 - distribution disk 7-44
 - JCL 7-45
 - IDMS 4-25, 4-48 to 4-58
 - distribution disk 4-51
 - IMS 6-20, 6-37, 6-40
 - Model 204 10-43
 - distribution disk 10-46
 - Network 5-71
 - distribution disk 5-74
 - programs 5-45
 - Workload 1-74
 - PRIME SHIFT FIRST HOUR
 - parameters 8-59, 8-67
 - process table statistics reports
 - UNIX 12-10
 - production procedures
 - DASM 2-43
 - DB2 7-42
 - profile reports
 - CICS
 - transaction 3-12
 - profiles
 - Model 204 user ID 10-11
 - Network terminal 5-23
 - programs
 - Network
 - graph 5-45
 - presentation graphics 5-45
- Q**
- queue delay graphs
 - DASD 1-32, 1-48
- R**
- ratio graphs
 - paging 1-50
 - regular daily graphs
 - DB2 7-47
 - regular hourly graphs
 - DB2 7-46
 - Model 204 10-49
 - regular monthly graphs
 - DB2 7-49
 - regular weekly graphs
 - DB2 7-48
 - report generation
 - Generic 8-15
 - report production
 - Network
 - JCL 5-40
 - report production JCL
 - CICS 3-27
 - IDMS 4-21
 - reporting component
 - UNIX 12-8
 - reporting parameters
 - DB2 7-26 to 7-36
 - IMS 6-22
 - reports
 - Application Profile 9-14
 - Batch Window 9-6 to 9-9
 - capacity factors 9-11 to 9-12
 - CICS
 - Ad Hoc 3-7
 - exception analysis 3-6
 - response performance 3-4
 - summary analysis 3-3
 - system overview 3-10
 - terminal activity 3-9

- terminal statistics 3-6
- transaction profile 3-12
- transaction statistics 3-4 to 3-6
- trends analysis 3-13
- cluster controller
 - activity 5-10, 5-55
 - statistics 5-9, 5-55
- CPU
 - activity 1-45
- DASD
 - activity 1-45
 - detail 1-28, 1-44, 1-45
 - volume mapping 2-42
- DASM 2-12 to 2-29
 - activity 2-42
 - Ad Hoc 2-6 to 2-10
 - data reduction 2-6 to 2-10
 - data set activity 2-13
 - data set detail 2-13
 - selected VSAM summary 2-15
 - space allocation 2-42
 - volume activity 2-13
 - volume summary 2-14
- DB2
 - Ad Hoc 7-13 to 7-15
 - AUTH-ID 7-9 to 7-13
 - batch 7-36
 - connect name 7-4 to 7-8
 - exception analysis 7-6
 - job streams 7-36
 - on-line 7-38
 - plan statistics 7-5
 - subsystem trends 7-37
 - subsystem-wide 7-3
 - summary analysis 7-3, 7-38
 - transit time performance 7-5
 - trends 7-16
 - workload and performance 7-3 to 7-16
- Generic
 - activity 8-16
 - Ad Hoc 8-15, 8-38
 - batch 8-40 to 8-44
 - daily 8-41
 - daily overview 8-16
 - generated from tables 8-15
 - statistics 8-16, 8-40
 - substitution 8-17
- IDMS
 - Ad Hoc 4-7, 4-15, 4-42
 - additional options 4-36
 - batch 4-23, 4-43
 - exception analysis 4-6, 4-42
 - job streams 4-23
 - on-line 4-40
 - parameters 4-27
 - response performance 4-4
 - statistical 4-40
 - summary analysis 4-3, 4-41
 - terminal statistics 4-6, 4-41
 - transaction profile 4-9
 - transaction statistics 4-4, 4-40
 - trends analysis 4-10, 4-42
 - workload and performance 4-3 to 4-8
- IMS
 - Ad Hoc 6-33
 - additional options 6-28
 - analysis 6-32
 - batch 6-33
 - exception analysis 6-4, 6-33
 - on-line 6-31
 - parameters 6-22
 - production 6-16
 - response performance 6-3
 - statistical 6-31
 - summary analysis 6-2, 6-32
 - terminal statistics 6-4, 6-31
 - transaction statistics 6-3, 6-31
 - trends analysis 6-6, 6-32
 - workload and performance 6-2
- Model 204
 - Ad Hoc 10-8, 10-14, 10-37
 - additional options 10-32
 - batch 10-38
 - exception analysis 10-7, 10-27, 10-37
 - JCL and parameters 10-19
 - on-line 10-35
 - performance 10-25
 - response performance 10-4
 - statistical 10-35
 - summary 10-24
 - summary analysis 10-3, 10-36
 - terminal statistics 10-7, 10-25, 10-36
 - transaction statistics 10-4, 10-25, 10-35
 - trends analysis 10-11, 10-37
 - workload analysis and performance 10-3

Network

- Ad Hoc 5-6, 5-64
- additional options 5-60
- batch 5-42
- batch submission 5-65
- exception analysis 5-50, 5-55
- on-line 5-63
- performance 5-51
- physical activity 5-55
- physical line activity 5-8
- physical line exception 5-64
- physical line statistics 5-7
- physical statistics 5-56
- response performance 5-4
- session exception 5-64
- summary 5-50, 5-64
- summary analysis 5-3
- terminal profile 5-53
- terminal statistics 5-5, 5-51
- trends analysis 5-24
- workload and performance 5-3 to 5-10

shift turnover 9-2**Tape**

- Ad Hoc 11-7
- component 11-7
- drive activity 11-3, 11-8
- drive exception analysis 11-3
- drive statistics 11-3
- storage by Ownerid 11-2
- storage by volume type 11-2
- volume activity 11-2, 11-8
- volume summary 11-2

TSO

- performance 1-44
- USERID 1-19

UNIX

- component 12-8
- disk statistics 12-10
- file system capacity 12-14
- file table statistics 12-11
- inode table statistics 12-11
- lock table statistics 12-11
- process table statistics 12-10
- server availability 12-13
- standard as input 12-3
- statistics 12-10
- summary 12-9
- technical graph 12-12

VTAM

- physical line statistics 5-63
- terminal session statistics 5-63
- trends analysis 5-64

Workload 1-18

- Ad Hoc 1-18 to 1-19, 1-67
 - Ad Hoc request parameters 1-12
 - batch 1-69
 - data reduction 1-5
 - exception analysis 1-66
 - general statements 1-51
 - job statistics 1-24
 - JOBNAME 1-18
 - JOBNAME CPU activity 1-21
 - on-line 1-63 to 1-65
 - performance 1-23
 - PGMNAME 1-19
 - printer activity 1-26
 - printer statistics 1-26
 - program statistics 1-24
 - statistical 1-63
 - statistics 1-24
 - summary analysis 1-20, 1-66
 - trends analysis 1-66
 - TSO command statistics 1-25
 - TSO user statistics 1-25
 - workload and performance
 - DB2 7-3 to 7-16
 - IDMS 4-3 to 4-8
 - Network 5-3 to 5-10
 - workload DASD 1-28
 - workload exception 1-27
- reports generated from tables
- Generic 8-15
- response graphs
- CICS
 - transaction 3-11
 - IDMS
 - transaction 4-8, 4-40
 - Network
 - transaction 5-63
 - Network transaction 5-11, 5-52
 - TSO transaction 1-48
- response performance reports
- CICS 3-4
 - IDMS 4-4
 - IMS 6-3
 - Model 204 10-4

- Network 5-4
- RMF/CMF type 74 records 11-5
- RMM 11-6
- S**
- sample jobstreams
 - DASM 2-39
 - Workload 1-14
- sar input data (UNIX) 12-2
 - consolidating 12-5
- SELECTED DAY
 - parameters 8-59, 8-68
- SELECTED SYSTEM
 - parameters 8-56, 8-66
- SELECTED SYSTEM parameters
 - Generic 8-47
- SMF data
 - DASM 2-3
 - Tape 11-5
- SMF files
 - Workload
 - processing history 1-3
 - processing live 1-3
- SMF logical records 8-24
- SMF records
 - DB2 7-18
 - Model 204
 - processing of 10-12
- SMF type 21 records
 - Tape 11-5
- SMF/RMF records
 - Workload 1-4
- SMFILE parameters
 - Generic 8-47
- space allocation reports
 - DASM 2-42
- SQL command activity
 - DB2 7-7
- standard UNIX reports as input 12-3
- statistical reports
 - IDMS 4-40
 - IMS 6-31
 - Model 204 10-35
 - Workload 1-63
- STATISTICS REPORT
 - parameters 8-57
- statistics reports
 - cluster controller 5-9
 - Generic 8-16, 8-40
- IDMS
 - terminal 4-41
 - transaction 4-4, 4-40
- Network
 - physical line 5-7
- Network physical 5-56
- Network terminal 5-51
- PGN service unit 1-26
- Tape
 - drive 11-3
- UNIX 12-10
 - disk 12-10
 - file table 12-11
 - inode table 12-11
 - lock table 12-11
 - process 12-10
- VTAM
 - physical line 5-63
 - terminal session 5-63
- Workload 1-24
 - job 1-24
 - PGN service unit 1-26
 - printer 1-26
 - program 1-24
 - TSO command 1-25
 - TSO user 1-25
- storage activity graphs
 - Auxiliary to Central 1-34, 1-51
 - Central to Auxiliary 1-34, 1-51
 - Central to Expanded 1-33
 - Expanded to Auxiliary 1-34, 1-50
 - Expanded to Central 1-33, 1-50
- storage by Ownerid reports
 - Tape 11-2
- storage by volume type reports
 - Tape 11-2
- substitution reports
 - Generic 8-17
- subsystem
 - CICS 3-3
 - DASM 2-3
 - DB2 7-3
 - IDMS 4-3
 - IMS 6-2
 - Network 5-3
 - Tape 11-2

- subsystem daily graphs
 - IMS 6-42
- subsystem hourly graphs
 - IMS 6-42
- subsystem installation JCL
 - CICS 3-14
 - IDMS 4-11
 - IMS 6-7
- subsystem monthly graphs
 - IMS 6-44
- subsystem trends reports
 - DB2 7-37
- subsystem weekly graphs
 - IMS 6-43
- subsystem-wide reports
 - DB2 7-3
- summary analysis reports 10-36
 - CICS 3-3
 - DB2 7-3, 7-38
 - IDMS 4-3, 4-41
 - IMS 6-2, 6-32
 - Model 204 10-3
 - Network 5-3
 - Workload 1-20, 1-66
- summary programs
 - Network
 - trends analysis 5-44
- summary reports
 - Model 204 10-24
 - Network 5-64
 - Tape
 - volume 11-2
 - UNIX 12-9
- swapping activity graphs
 - UNIX 12-12
- SYSIN
 - DD Statements 8-46, 8-65
 - JCL 8-56
- SYSNAP
 - DD Statements 8-46, 8-65
 - JCL 8-56
- SYSPRINT
 - DD Statements 8-45, 8-63
 - JCL 8-54
- system calls graphs
 - UNIX 12-13
- system overview reports
 - CICS 3-10

- SYSUT1
 - DD Statements 8-45
- SYSUT3
 - DD Statements 8-46

T

- TABLEDEF 8-69
- TABLEDEF input file 8-23
- TABLEDEF statements 8-61
- tables
 - Generic
 - user 8-34
 - user defined 8-18, 8-23
- Tape 11-8
 - error statistics by volume 11-5
 - ESV records 11-5
 - graphs
 - drive activity 11-4, 11-10
 - drive device active 11-4
 - drive device allocation 11-4
 - drive mount pending 11-4
 - ownership 11-10
 - PC components 11-9
 - volume activity 11-3
 - I/O device detail records 11-5
 - ISPF/PDF 11-9
 - JCL
 - batch 11-9
 - lists
 - drive activity 11-4
 - reports
 - Ad Hoc 11-7
 - component 11-7
 - drive activity 11-3, 11-8
 - drive exception analysis 11-3
 - drive statistics 11-3
 - storage by Ownerid 11-2
 - storage by volume type 11-2
 - volume activity 11-2, 11-8
 - volume summary 11-2
 - SMF data 11-5
 - SMF type 21 records 11-5
 - subsystem 11-2
- Tape Management Systems supported 11-6
- technical graph reports
 - UNIX 12-12
- technical support, contacting the CIMS Lab xvi

- terminal activity graph
 - CICS 3-12
- terminal activity graphs
 - IDMS 4-8, 4-40
 - IMS 6-5, 6-31
 - Model 204 10-10, 10-26, 10-35
- terminal activity reports
 - CICS 3-9
- terminal profile reports
 - Network 5-53
- terminal profiles
 - Network 5-23
- terminal response graph
 - CICS 3-11
- terminal response graphs
 - Model 204 10-9, 10-27
- terminal statistics reports
 - CICS 3-6
 - IDMS 4-6, 4-41
 - IMS 6-4, 6-31
 - Model 204 10-7, 10-25, 10-36
 - Network 5-5
- thread transit time graphs
 - DB2 7-7
- TMS 11-6
- traffic graphs
 - cluster controller
 - inbound 5-20
 - outbound 5-21
 - retransmitted 5-22
 - total 5-22
 - cluster controller inbound 5-54
 - cluster controller outbound 5-54
 - cluster controller retransmitted 5-54
 - cluster controller total 5-54
 - Network
 - inbound message 5-13
 - outbound message 5-13
 - physical line inbound 5-15, 5-56
 - physical line outbound 5-15, 5-56
 - physical line retransmitted 5-16, 5-57
 - physical line total 5-17, 5-57
 - total 5-13
 - Network inbound 5-52
 - Network outbound 5-52
 - Network total 5-53
- traffic reports
 - cluster controller statistics 5-55
- transaction activity graphs
 - CICS 3-12
 - IDMS 4-8
 - IMS 6-5
 - Model 204 10-10
- transaction activity lists
 - CICS 3-8
 - Model 204 10-8
- transaction profile report JCL
 - IMS 6-19
- transaction profile reports
 - CICS 3-12
 - IDMS 4-9
- transaction response graphs
 - CICS 3-11
 - IDMS 4-8, 4-40
 - IMS 6-5, 6-31
 - Model 204 10-9, 10-27, 10-35
- transaction statistics reports
 - CICS 3-4 to 3-6
 - IDMS 4-4, 4-40
 - IMS 6-3, 6-31
 - Model 204 10-4, 10-25, 10-35
- transaction volume graphs
 - IDMS 4-40
 - IMS 6-31
 - Model 204 10-35
- transit time performance reports
 - DB2 7-5
- trends analysis
 - IDMS 4-34
 - Network
 - summary programs 5-44
 - Workload 1-56
- trends analysis reports
 - CICS 3-13
 - IDMS 4-10, 4-24, 4-42
 - IMS 6-6, 6-32
 - Model 204 10-11, 10-37
 - VTAM 5-64
 - Workload 1-66
- trends analysis summary program JCL
 - IMS 6-20
- trends reports
 - DB2 7-16
 - operating JCL for 1-57
 - sample jobstreams for 1-60

TSO

- graphs
 - activity 1-48
 - transaction activity 1-47
 - transaction response 1-48
- reports
 - performance 1-44
 - USERID 1-19

U

UNIX 12-11

- graphs
 - buffer read activity 12-13
 - buffer write activity 12-12
 - logical read activity 12-13
 - logical write activity 12-12
 - paging activity 12-12
 - physical read activity 12-13
 - physical write activity 12-12
 - swapping activity 12-12
 - system calls 12-13
- input data
 - sar 12-2
 - sar, consolidating 12-5
 - standard reports used as 12-3
- reports
 - component 12-8
 - disk statistics 12-10
 - file system capacity 12-14
 - file table statistics 12-11
 - inode table statistics 12-11
 - lock table statistics 12-11
 - process table statistics 12-10
 - server availability 12-13
 - standard as input 12-3
 - statistics 12-10
 - summary 12-9
 - technical graph 12-12
- server
 - alias for, creating 12-7
 - identifying 12-4, 12-6 to 12-7
- user defined tables
 - Generic 8-18, 8-23
- user ID profiles
 - Model 204 10-11
- user table definitions
 - Generic 8-37

user tables

- Generic 8-34

V

values

- VTAM
 - filtered 5-59
- volume activity graphs
 - Tape 11-3
- volume activity report
 - DASM 2-13
- volume activity reports 11-8
 - Tape 11-2
- volume graphs
 - IDMS
 - transaction 4-40
 - Network
 - transaction 5-63
- volume map
 - DASD
 - by CCHH 2-13
 - by Dsname 2-13
- volume mapping reports
 - DASD 2-42
- volume summary report
 - DASM 2-14
- volume summary reports
 - Tape 11-2
- VTAM
 - reports
 - physical line statistics 5-63
 - terminal session statistics 5-63
 - trends analysis 5-64
 - values
 - filtered 5-59
- VTOC data
 - DASM 2-3

W

weekly graphs

- IDMS 4-55
- Model 204
 - regular 10-51
- Network 5-81
 - extension 5-82
 - physical configuration 5-83
- Workload 1-74
 - activity reports

- JOBNAME CPU 1-21
- data reduction 1-3
 - input parameters 1-7
 - JCL 1-5
 - sample jobstreams for 1-14
- DD Statements 1-5 to 1-6, 1-34 to 1-37
- graphs 1-29
 - batch GDDM 1-69
 - channel activity 1-30
 - CPU activity 1-29
 - CPU ratio 1-31
 - on-line 1-61 to 1-63
 - paging activity 1-30
 - paging ratio 1-31
 - production procedures 1-74
 - TSO activity 1-29
- ISPF/PDF 1-61
- JCL
 - batch reports 1-34
 - data reduction 1-5
 - sample 1-36
- naming conventions 1-75
- parameters 1-38 to 1-55
- reports 1-18
 - Ad Hoc 1-18 to 1-19, 1-67
 - Ad Hoc request parameters 1-12
 - batch 1-69
 - data reduction 1-5
 - general statements 1-51
 - input parameter statements for 1-38
 - job statistics 1-24
 - JOBNAME 1-18
 - on-line 1-63 to 1-65
 - performance 1-23
 - PGMNAME 1-19
 - printer activity 1-26
 - printer statistics 1-26
 - program statistics 1-24
 - statistical 1-63
 - statistics 1-24
 - summary analysis 1-20, 1-66
 - trends analysis 1-66
 - TSO command statistics 1-25
 - TSO user statistics 1-25
- SMF files
 - processing history 1-3
 - processing live 1-3
- SMF/RMF records 1-4
 - trends analysis 1-56
- workload and performance graphs
 - Model 204 10-9 to 10-10
- workload and performance reports
 - DB2 7-3 to 7-16
 - IDMS 4-3 to 4-8
 - IMS 6-2
 - Model 204 10-3
 - Network 5-3 to 5-10
- workload DASD
 - reports 1-28
- workload exception reports 1-27

